PHYSICAL ACOUSTICS SUMMER SCHOOL



DTIC
ELECTE
MAR 2 8 1995
B

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

ASILOMAR CONFERENCE CENTER

19950327 072

DTIC QUALITY INSPECTED 1

1994 Physical Acoustics Summer School

Volume I - Transcript

This work relates to Department of Navy Grant N00014-94-1-0468 issued by the Office of Naval Research. The United states Government has a royalty-free license throughout the world in all copyrightable material contained herein.

Volume I
CONTENTS

Page

FORWARD
Dr. Logan Hargrove 1

25 June

MOLECULAR ACOUSTICS
H.E. Bass 2

DROPS AND BUBBLES: EXPERIMENTAL ASPECTS
P.L. Marston 48

27 June

RESONANCE ULTRASOUND SPECTROSCOPY
O.L. Anderson 105

NONLINEAR ACOUSTICS
S.L. Garrett 146

28 June

SONOLUMINESCENCE
S. Putterman 210

PERIODIC, RANDOM, & QUASIPERIODIC PHENOMENA
J.D. Maynard 263

29 June

FAR-OFF-EQUILIBRIUM ACOUSTIC PHENOMENA
B.C. Denardo 306

30 June

THERMOACOUSTIC ENGINES AND REFRIGERATORS
G.W. Swift 358

PHYSICAL ACOUSTICS IN MATERIALS SCIENCE
R.E. Green 401

REPORT DOCUMENTATION PAGE 433

per form 50
Codes
and/or
Dist Special
A-1

Forward

These are the Proceedings of the 1994 Physical Acoustics Summer School, The lectures were recorded and the verbatim transcripts were subsequently edited by the authors for publication here.

The Office of Naval Research (ONR), in cooperation with the Acoustical Society of America (ASA) and the National Center for Physical Acoustics (NCPA) sponsored the 1994 Physical Acoustics Summer School (PASS 94), 24 June - 1 July 1994, at the Asilomar Conference Center in Pacific Grove, California, the site of the first and similar 1992 Summer School (PASS 92). Participation in both Schools was limited to a total of 50 which included Students, Lecturers, and Discussion Leaders.

The purpose of these Summer Schools is to bring graduate students, distinguished lecturers, and discussion leaders together to discuss a wide variety of subjects in physical acoustics. This gives the students the opportunity to meet experts and talk about topics most students ordinarily wouldn't encounter at their own colleges and universities. The focus was on graduate students and academic participants. Approximately half of the participants have been advanced graduate students in physical acoustics.

The Physical Acoustics Summer Schools had their beginning in an ONR Principal Investigators meeting in 1988 where it was decided that the best investment of this kind would be in a summer school focused on graduate students. In 1990 an informal summer school for students of ONR Principal Investigators and some invited guests was held as part of a Principal Investigators meeting at the Naval Postgraduate School, Monterey, California, and this controlled experiment set the pattern and influenced the site selection for the subsequent PASS 92 and PASS 94. We hope that the Summer Schools will continue as biennial events with the high standards and wonderful success we have thus far enjoyed. This has happened and will happen again because everyone involved does their best, and for this I say, "Thank You!".

Logan E. Hargrove

MOLECULAR ACOUSTICS

DR. BASS: We are going to start this summer school with a lecture on molecular acoustics. Molecular acoustics and can actually be traced back into the last century, at the time when scientists were only first beginning to appreciate that as molecules and atoms actually existed.

(Transparency 1)

Throughout the history of the evolution of molecular acoustics there has been a synergism between people who are interested in the physics at the fundamental scale, at the molecular scale or at the atomic scale, and like; Kirchhoff and Stokes, who were more interested in what happened to the acoustic wave as it propagated than about the microscopic processes that took place that allowed the propagation. As time progressed, quantum theory was introduced. Acoustics measurements were used to test the new quantum theory. These macroscopic measurements tested microscopic models of how nature behaves at the molecular and atomic levels. This process has continued throughout the history of molecular acoustics, and I will talk about some of those examples today.

I will develop this topic in more or less a historical framework, not because I am especially interested in history but because we start out with the simpler problems and work up to the more complex.

Recalling your summer school applications, most of you have had at least one course in acoustics. During that one course in acoustics you were probably introduced to some simple concepts about wave propagation in a gas. I am going to be talking today mostly about wave propagation in gases.

That choice is made for several reasons. There are two types of physical systems that we can really treat reasonably well. One is a gas, because it is completely disordered, and one is a regular solid when it is completely ordered. Those things in-between, like liquids, we avoid; they are too complicated. You do not have either simplification working for you.

So we are going to be concerned primarily with gases, and we are going to start out, as you would in your first acoustics course, talking about the propagation of simple plane waves through an ideal gas, a gas that has no means for dissipation.

Next we will put in dissipation mechanisms, starting out with viscosity (which I think most of you have probably seen), and throw in some heat conduction (which you probably also have

encountered). These two, combined, we call classical loss mechanisms, classical loss mechanisms because they, generally speaking, do not require us to invoke quantum concepts to describe their source and, also, because they were discovered before the quantum theory.

Then we will go into what I consider to be the more interesting part of the lecture; we will begin to talk about those types of energy-transfer processes that lead to absorption of the acoustic wave that rely on a quantum description for their understanding. That will be the primary subject of the second part of the lecture, relaxation absorption.

Then, finally, I will talk about some things I know nothing about, just to toss out some ideas about areas that need further exploration, some areas that you might be interested in pursuing in the future.

Before I start this, I will point out Doug Shields in the audience. Doug and I have worked together in molecular acoustics for longer than either one of us care to admit. Much of the work that I present here is his since I learned much of what I am presenting today from him. If you have questions, especially questions that I cannot answer, ask Doug. He has all the answers.

(Transparency 2)

Let us start out with a simple case: The wave equation with no dissipation.

We start out with conservation of mass, conservation of momentum, and conservation of energy, and then we try to come up with some type of equation of state. The equation of state that I will use most for this lecture is the ideal gas law; we take the most simple equation of state.

If I do this, and I assume that I am propagating a plane wave in the +x direction, one solution is a wave oscillating with angular frequency ω. I make those substitutions into my four equations -- I come up with a solution for the propagation speed C_0. This derivation is outlined in your notes (I am not going to go through it, I suspect most of you have seen it at one time or the other in some of detail).

One possible solution to the four conservation equations is a propagating plane wave in the x direction, provided that the speed of propagation is given by C_0, which can typically be written for a gas, as the $\sqrt{\frac{\gamma RT}{M}}$. This is the adiabatic speed of sound and we are familiar with that as the speed of sound for an ideal gas.

An important point to note is that propagation velocity is purely real. There are no attenuation mechanisms introduced up to this point.

(Transparency 3)

Now we are going to start complicating matters. The first complication we will introduce is viscosity. Introducing viscosity is not that tough; basically what we do is to introduce an additional term in our equation of motion -- an additional term here (momentum equation) -- that includes a viscosity. Actually, there are two types of viscosity; a shear viscosity and a bulk viscosity (I will talk about those in a minute).

So what I have added in is an additional term. Which allows me to include viscosity, or losses due to viscosity, in the wave equation.

If you think about this for a moment. An acoustic wave squeezes the gas. Why should just squeezing the gas lead to viscous loss, especially when the quantity that is being entered into the equation is a shear viscosity? Why do you have shear if you are just compressing the gas? This is not all that clear to me, and I have gone through this argument a hundred times.

(Transparency 4)

Basically, you can think about it this way, if you like. Assume you have a wave going in the x direction. If you uniformly squeeze a volume of gas or let a volume of gas expand, then you have displacements in the x, y, and z directions. But if you are propagating in only the x direction, you expect displacement in only the x direction. What that means is that you do not have a uniform displacement in the x, y, and z; you have it in only the x direction.

In order to describe the process, you can start out with a bulk displacement or bulk compression of a test volume. Next remove the displacement in the y direction and then remove the displacement in the z direction. Removing those displacements requires that you distort the volume. You are distort the volume in a direction that is orthogonal to the direction of propagation.

You might think about this for a little while (it is written up in the notes). After a few minutes, you will understand it better than I do, but the effect is that you introduce a shear. You have to introduce a shear if you are going to disturb this volume in only the x direction and maintain a constant volume.

(Transparency 3)

Now, let us look at the magnitudes of different terms. Assume that the density of gas is approximately the static density. You will notice that there are two viscosity terms here, an η and an η' (Stokes introduced these way back in the 1800s). One of them is referred to as a bulk viscosity; the other is the shear viscosity. For now, I will set the bulk viscosity term equal to zero. I now can eliminate this η' term from the equation of motion. Next, introduce a sinusoidal time dependence, as I did in the other case of a plane wave propagating in the x direction, and then I can come up with an equation for the propagation constant.

The propagation constant now has both real and imaginary parts. The imaginary part is an absorption α and the real part is k_r. The real part is almost the same as the propagation constant in the absence of viscous losses since the propagation velocity is approximately the same.

The term α, though, appears in the exponential. As you can see, the amplitude decays with time according to the absorption term α.

I take this solution and substitute it back into the wave equation (which results from equation 15). I get both an imaginary part and a real part of the propagation constant. The imaginary part of the propagation constant depends upon the viscosity (the larger the viscosity, the more absorption you get). It also depends on ω^2.

This ω^2 dependence is characteristic of many absorption mechanisms both in gases and in liquids and also, quite often, in solids.

I give you two examples to illustrate the magnitude of these terms: oxygen and water. The transparency shows the ratio of alpha or the imaginary part of the propagation constant compared to the real part of the propagation constant. It turns out that in order to make a clean separation between the real and imaginary parts of the propagation constant it is necessary that I assume that this ratio is small.

As you can see, for most materials that we might want to work with, gases or liquid, at reasonable frequencies this ratio is very small. Attenuation is a much smaller quantity than is the speed of sound or the real part of the propagation constant.

This ratio also increases linearly with the frequency.

In the case of water, the attenuation is down by a factor of 10^4 from that in a gas. Water attenuates sound much less than does a gas. Furthermore, if you go to a crystalline solid, you will find that the attenuation is even smaller. We will see some fundamental reasons for this trend as we progress, but it is something that I would like for you to keep in mind.

(Transparency 5)

In addition to viscous losses, we also have to take into account heat conduction. The source of loss due to heat conduction you can envision reasonably well through hand-waving arguments. If a plane wave is propagating in the x direction, there are alternate compressions and rarefactions. As the gas is compressed, it heats. So in the region of pressure maxima there will be an increase in temperature; in a pressure minima, a decrease in temperature.

If you could freeze the wave in time and space, there would be a flow of heat from the high-temperature region to the low-temperature region and, given time, these would average out. There is conduction of energy between these two regions.

Conduction of energy can be described through Fourier's law. We simply write down Fourier's law in the energy equations.

Then we again make the substitution for plane wave propagation. Doing that allows us to solve for a new propagation constant, k, which will again have real and imaginary parts. The real part is approximately the same as it was in the absence of dissipation -- ω/C_0 -- but α is now a different attenuation constant, an attenuation constant that is due to thermal conduction. As you can see, it depends upon the coefficient of thermal conductivity -- κ --.

It also depends upon ω^2.

The attenuation also depends upon γ, the ratio of specific heats, which is C_p/C_v.

(Transparency 6)

Now we will get to the interesting stuff, or at least things that are not quite as simple and straightforward to envision and to compute: the contribution to absorption and dispersion due to relaxation processes. We now must consider the gas at the microscopic level; otherwise we have no hope of understanding these relaxation processes.

At the microscopic level we visualize atoms and molecules running around, bouncing into each other. In a typical

container these molecules and atoms have several degrees of freedom available to them. They can move around in 3-D translationally; there is translational energy associated with that motion. They can rotate around some axis. If they are diatomic molecules, they can rotate in 2-D, or they can spin about the axis joining the atoms. A large molecule will have three degrees of rotational freedom.

An atom has no degrees of rotational freedom; a diatomic molecule has two degrees of rotational freedom.

Most interesting are the vibrational degrees of freedom. You can think of molecules as atoms connected together by little springs. Atoms connected to each other with springs, of course, vibrate. You know what happens when you put in three atoms and two springs: The number of modes goes up. If you put in four atoms and three springs, you get even more vibrational modes. The number of vibrational modes available to the molecule increases as the complexity of the molecule increases.

I am a pretty simple physicist and it is difficult for me to think in terms of more than two atoms and one spring, so we will spend most of our time today talking about diatomic molecules. I would feel better if we could just treat atoms, but in order to get vibrational modes you have to have at least one spring, so I will allow myself that much complexity. Once or twice we will even talk about a three-atom system.

If one wants to examine all the vibrational modes that are possible and the relaxation processes associated with those vibrational modes you would need a large enough mind to be able to envision four or five atoms in a single big molecule (or at least big by a physicist's standard, probably small by most chemists' standards). In any case, in general, there are three degrees of freedom: translation, rotation, and vibration. All of them have associated with them some energy.

These three energy modes at any one time would like to be in equilibrium with each other. There is a fundamental theorem of statistical mechanics called the equipartition theorem, which says, basically, that each of the energy modes should have associated with it the same energy; they should be at equilibrium with the translational temperature in order to be in equilibrium.

If I were to disturb the number of molecules in a vibrationally excited state, the system would try to return to equilibrium with the other modes, rotation and translation. In most cases equilibrium can only be established through collisions.

It is possible in a gas for two molecules to exchange energy by releasing a quantum of electromagnetic radiation, a

photon, which is absorbed by another molecule. Those processes occur but are typically much less effective in establishing equilibrium. I am going to assume that if you take a gas away from equilibrium, the only way that it can return to equilibrium is through collisional processes. One atom or molecule, collides with another and, through multiple collisions, equilibrium is established.

The term "collision" here means that two gas molecules interact. One of the nice things about gases is that the molecules are far enough apart so that they do not see each other most of the time; they see each other for only very small instants of time; when they come close they interact; we call that a collisions.

If we were talking about billiard balls, it would be pretty easy to envision when a collision takes place. That is when they strike each other and go "bang." Different types of molecules interact over different ranges, and so molecular collision are not quite as simple as billiard balls.

Collisions in a gas at atmospheric pressure take place at a rate of about 10^{11}/s. There are a lot of collisions taking place between the gas molecules in this room. There are 10^{11} opportunities per second for the gas to reestablish equilibrium. It is this process of reestablishing equilibrium that will consume the remainder of this first lecture and most of the second lecture.

Let me discuss this process in more detail. There are three degrees of freedom: translation, rotation, vibration. In your quantum mechanics course, you learned that the translational energy states for a particle in a reasonable sized container -- this room -- are very finely spaced; so finely spaced that for all practical purposes we can treat the manifold of energies available as a continuum. We will treat translational energies as a continuum.

That means that if two billiard balls collide, there are no quantum mechanical that must be applied when transferring energy from one to the other. You have to satisfy only the classical equations of motion for energy-momentum transfer and everything works.

Since there are no quantum rules necessary for these types of collisions, it is fairly simple to transfer energy between translational modes.

Rotation invokes a little more complication. We know, from our study of quantum mechanics, that rotational motion is, in fact, quantized. The simplest case is a diatomic molecule.

A diatomic molecule has rotational energy states that are unequally spaced. When you have a collision that is going to change the rotational state of the collision partners, the collision must be such that the energy transferred equals a quantum of rotational energy.

This represents an additional restriction placed upon the collision. For most molecules, in the case of rotation, that is not a severe constraint. If you put in some numbers, it turns out that the rotational energy levels for most molecules that we will deal with, oxygen, nitrogen, CO_2, things like that, are very close together. If you look at the most probable quantum number for rotation and look at the energy-level spacing, which is the energy required to cause a transition from one level to the next, a molecule at 1K has sufficient energy. In other words, the energy necessary to cause a change of quantum number of one is possessed by molecules with a translational energy characteristic of a gas with a temperature of 1K. Since we are typically interested in room-temperature or about-room-temperature gases, almost every colliding pair has more than enough energy to cause a quantum of rotational energy to be transferred.

We would expect, then, that rotational energy transfer is fairly efficient and, indeed, it is. The only thing that prohibits the transfer of rotational energy in every collision is the fact that there are geometric parameters that come into play. Not all collisions are conducive to the transfer of rotational energy. Imagine a collision in which the partners are coming in directly along the chemical bond of each. During the collision, there will be compression of the bonds (springs) and then a rebound; there is not much chance to transfer rotational energy in those types of collisions. A few collisions will not transfer rotational energy, but most will.

On the average, for nitrogen, oxygen, molecules like that, about one in five collisions are effective in transferring a quantum of rotational energy. In the case of translation, about one in one. Rotation takes more time to equilibrate, but if you have 10^{11} collisions per second, who cares? So it takes five of those 10^{11} collisions to cause re-equilibration of rotational energy.

There are a couple of exceptions to this general rule. The most notable one is hydrogen. Hydrogen atoms do not have much mass and they are on this big spring, so that when they rotate the rotational energy states are fairly widely spread. As a result, there are some very strong quantum effects that take place when you try to transfer energy from translation to rotation of hydrogen. As a result, the number of collisions required to re-establish rotational equilibrium in hydrogen at room temperature is about 350. About one in 350 collisions can cause a quantum of rotational energy to be transferred to

translation, or vice versa.

Because of the large rotational energy-level spacings, many of the collisions occur between collision partners that do not have the energy required to undergo a transition from one rotational state to the other. But still, if there are 10^{11} collisions per second who cares? Instead of re-equilibrating with a relaxation time of 10^{-11}, it is now 10^{-9}. Still, 10^{-9} seconds is pretty fast.

Q: We are talking about a large number of molecules. You mentioned that per second there would be 10^{11} interactions, but generally we have -- 10^{23} molecules. Considering this number, each molecule only suffers a few collisions.

A: Each one of the molecules undergoes 10^{11} collisons per second. That is 10^{34} collisions that take place per mole. There are 10^{11} for each molecule.

Going back to rotational relaxation and who cares that it takes 10^{-9} seconds to re-equilibrate things, how important is that? Well, it turns out that this is a time regime that really can be significant for ultrasonic studies. If you think about it for a moment, the number of collisions that take place per second depends upon how many molecules or atoms are in the volume; it depends on the density.

If the density is decreased, the collision frequency decreases. At a pressure of one 10th of an atmosphere, there are 10^{10} collisions per second and, for a hundreth of an atmosphere, 10^{9}. As the pressure decreases further the number of collisions that take place per second gets smaller and smaller. If you get into a good laboratory vacuum, where you are down to 1000ths or ten thousands of an atmosphere, the time frame for rotational relaxation shifts down to 10^{-4}, 10^{-5}, 10^{-6} seconds.

We will find out later that the acoustic frequency or period should be about the same as the time it takes to transfer a quantum of energy to observe maximum acoustic effects. This means that in a laboratory vacuum an acoustic period of 10^{-4},10^{-5}, 10^{-6} seconds will show a large effect. The problem is to propagate sound through a very low pressure, maybe a 1000th or ten thousands of an atmosphere. If we turn the gain up high enough and really zap an ultrasonic transducer we can get a little sound through the gas, so these are experimentally explorable regimes.

If you get one out of 350 collisions contributing to the energy transfer, that process can be experimentally observed. Rotational relaxation of hydrogen has, in fact, been carefully mapped out. It turns out that this is important for

astronomical reasons. The lifetime of rotational hydrogen can be a useful piece of information in trying to study the lifetimes or ages of certain intergalactic molecular clouds.

If you look at the number of rotationally excited molecules in an intergalactic hydrogen cloud, know the lifetime of rotational states, and know the collision rate (from the approximate density of the cloud), you can come up with an estimate of how long ago that cloud was at a very high temperature which would have left it in a rotationally excited state. This is one way of determining the age of far-away hydrogen clouds.

For most molecules rotational relaxation is fast and we will not investigate rotational relaxation in greater detail.

(Transparency 7)

Vibrational relaxation is much more important for most conditions we normally encounter. Quantum mechanics tells us that the energy associated with vibration of a molecule is typically much, much greater than that associated with rotation. In general, in order to excite a vibrational oscillation, more energy is required than that necessary to change the rotational quantum number by one. But this depends upon the vibrational energy level spacing for a particular mode. Molecules like propane, ethane, methane, have many vibrational energy states. The lowest vibrational energy state has an energy gap that is very small. In fact, it is a small fraction of the thermal energy that molecules have in this room.

Most molecules have sufficient translational energy to excite a quantum of vibrational energy during a collision with propane or ethane, something like that. Simple diatomic nitrogen, which is nice because lots of nice perturbation treatments can be applied, has a vibrational energy which requires for one quantum exchange about 10 or 20 times the typical thermal energy of a molecule in this room. A large amount of vibrational energy must be given to a single nitrogen molecule in order to cause energy transfer to take place. Very few collisions have sufficient energy to cause such a transition. It is a very rare event.

As a result, molecules like propane can re-equilibrate vibration in as few as 10 collisions. Nitrogen, on the other hand, might take 10^{12} collisions.

Q: What is the energy gap for nitrogen? Is it just a few electron volts?

A: I think I can convert for you in a second. When, dealing with molecular energy transfer as we are doing here,

even for vibrational states, the energy-level spacings are very small compared to, let us say, electronic states. These are comparatively small quanta.

A typical vibrational quantum might be on the order of a 10th of an eV or smaller -- I cannot do the conversion immediately in my head. Typically, we give the energy for vibrational states, or for rotational states, in terms of a multiple of the typical thermal energy that a molecule would have at that temperature.

We could say that the vibrational energy level is 10 kT, or 2 kT, or a 10th of a kT. This puts us into energy units that are consistent with the problem that we are treating. Although these vibrational states in nitrogen are fairly widely spaced by thermal standards, they still are pretty small compared to electronic states in hydrogen.

Q: Can't you argue that very few are going to be excited, maybe 10 in a whole box of atoms, so if you just change the state of one of those, you have changed 10% of the entire number in excited states? You have to have only one collision to really significantly reduce the state. So it is not only the collision frequency but the number of excited states, that is important.

A: There are very few molecules in excited states, that is correct. If I cause one molecule to become vibrationally excited that box is going to remain out of thermal equilibrium for a long period of time. Once it returns to thermal equilibrium that full quantum of energy will then be distributed over all the translational and rotational modes so that it will equal out. But in the meantime, just a few molecules have the dominant effect as far as carrying the energy for that mode. Since there are 10^{23} molecules/mole, even a very energetic vibrational state will still have a 10^{10} to 10^{15} molecules, not one or two.

Q: Hank, we are trying to remember back here, room temperature is about 1/40th of an eV?

A: Yes, so that would put nitrogen vibration about a third, I think. (Note added: At T=300K, $kT=4 \times 10^{-14}$.) That is about right.

Let us go through some of the equations. These are equations that you might or might not have seen. I am not going to deal a lot with equations; I would like to talk more about the physics. But, I would like to describe a simple relaxation equation.

Next, I want to establish a mathematical link between the phenomenon that I just described, energy transfer, and the

quantity that we can measure, which is velocity, or absorption of sound. In order to do that, we set up a simple relaxation equation. We then describe the change in the vibrational energy as a function of time, which we assume depends upon the difference between the present vibrational energy of a state and what that vibrational energy would be if it were in equilibrium with the translational motion of the gas. The time that it takes to return to equilibrium is measured by a relaxation time, τ.

The only other thing necessary to invoke in order to get a solution is detailed balance. In equilibrium, there will be a certain number of excitations up the quantum ladder, there will be a certain number of excitations down; in equilibrium the energy carried both ways will be equal. The rate constant for transitions from the first state to the ground state times the concentration of molecules in the first state minus the reverse rate times the number of molecules in the ground state, that has to be equal to zero (Eq. 35). Detailed balance is just a chemical equation.

Next, I define an effective specific heat. An effective specific heat is just the specific heat in the absence of the relaxing mode plus the specific heat associated with the relaxing degree of freedom which has a heat capacity, C'. We will talk about how we come up with that relaxing specific heat in a few moments. We will have the specific heat satisfy a relaxation equation.

Doing that, I come up with equations for the absorption per wavelength and the velocity in terms of relaxation times and frequency (Eqs. 42, 43). The ε represents an amplitude.

If I have multiple relaxation, then there will be several of these terms to sum over.

The relaxation time is directly related through an equation referred to as the Landau-Teller relation to the rate of energy transfer which comes from the rate equation. This rate, the energy-transfer rate, is proportional to one over the probability (equal in the proper units) of a quantum exchange during a single collision.

So within a few multiplicative factors we can look upon the relaxation time as being inversely proportional to the probability of energy transfer in a single collision. The probability of energy transfer and the associated rate constant can be derived from quantum theory. If you do a quantum scattering calculation and sum over all the scattering angles, you come up with a scattering cross-section. By summing over all impact parameters, you come up with the probability of a transition; the probability of a quantum exchange.

The probability of a quantum exchange can be used to compute k. Once you have k, (Eq. 37) allows you to relate the k computed from quantum theory to a relaxation time. The relaxation time which enters into these expressions is what we will deduce from our experimental observations. So this is the equation that allows us to relate experiment to theory.

This equation is valid only when there is a single relaxing degree of freedom and the modes are harmonic. For a complex molecule with multiple modes, a more complicated form is required. I am not sure it contributes much here to go into the more complicated form. We refer to the process as multiple relaxation. For multiple relaxation processes, the math gets a little tougher; the concepts do not change.

Q: Did you say you measure k?

A: I can compute k from fundamental quantum theory, or deduce it from experiment.

Q: Then in your measurement you are just verifying that computation or --

A: You compare the two. What we would like to have is a measurement tool that will allow us to determine if our calculation is correct. A theory which allows me to compute k, the energy transfer probability, depends upon such things as the interaction potential, the orientation during collisions, and other factors.

Since we cannot do a full quantum calculation, we develop an approximation such as a quasi-classical approximation. I take these approximations plus the best quantum I can do, put them into a calculation and compute k. The quantity I compare is τ which I measure.

If I get a good comparison, I go away happy that this microscopic model represents the way that nature works. The Nobel Prize in chemistry was awarded a few years ago to a person who did this.

Q; What frequency is represented by ν in equation 37?

A: It is the frequency of the vibrational mode; ν is what you would get if you solved the quantum mechanical problem for the lowest vibrational frequency of this harmonic oscillator. At times I use ν as the acoustic frequency. I have used them interchangeably without telling you. If you cannot keep up with my changes, force me to explain. $h\nu$ represents a quantum of vibrational energy.

DR. MARSTON: Could you say something about the classical moment of equation 37? Maybe the question makes no sense, but if you take h going to zero, you have a singular condition.

A: I will have to think about that. I cannot answer right off. Can you, Doug?

DR. SHIELDS: I think you ought to point out that equation 37 holds only on the assumption that you have equally spaced vibrational modes. There are only two places where you can relate k and τ. One is where you have a two-state gas. In that case, you have a plus sign instead of a minus sign in 37. In the case where you have a harmonic oscillator where the wells are equally spaced, you invoke the assumption that the transitional probability is proportional to the spacing then you can get a single relaxation time for the whole mode that is represented by τ and k_{10} is just the transitional probability for the first excited state.

A: I do not think that answers Phil's question. Why can't you allow $h\nu$ to go to zero?

Q: What he said was that there should be a plus sign instead of a minus sign.

A: For a two-state gas it is a plus. For a multi-state gas it is a minus.

Q: As $h\nu \to 0$, $\tau \to \infty$, because you do not have any transitions because you do not have that series of levels in classical harmonic oscillators.

Q: If $h\nu \to 0$, that is similar to $T \to \infty$. When $T \to 0$, $\tau \to \infty$. When $T \to 0$, there are almost no transitions. The relaxation time is infinity.

A: If $\nu \to 0$, $k_{10} \to \infty$. The probability for a transition gets bigger as the energy spacing gets less.

Q: If you go in the notes, you have the equation above, 34, you just get, in that case, that $\frac{dE\nu}{dt}$ is going to go to zero. You are not going to get any change in your energy state. If you go from 34 to 37, then you are dividing by zero

A: That still goes back to the argument that τ then becomes infinite.

DR. ATCHLEY: Shouldn't you go back in the calculation of k and see what happens when you let that go classical as well? I mean, k up there may not be constant. You cannot let $h\nu$ go to zero on one side and not worry about what happens to k.

A: Let us take a break now and we will pick up here. I have just one more overhead to summarize.

(Transparency 8)

What I have tried to do in this first part of the lecture is to give you an idea of the three absorption mechanisms that we are going to be discussing in more detail as the lecture progresses.

Viscosity and thermal conduction are pretty well understood. When I get toward the end of the lecture, though, I will come back and try to convince you that these two mechanisms are probably less understood than relaxation. We understand viscosity and thermal conduction -- we think we do -- because we typically operate in a frequency range, or frequency-over-pressure range, where the materials behave like they are a continuum. For a continuous medium, the viscosity and thermal conduction are reasonably well understood.

At very low pressures, for example, intergalactic space, or even in the laboratory in a good vacuum, one can observe a breakdown of that assumption. In those cases, thermal conduction and viscosity become much more complex.

In the next period we are going to talk more about relaxation processes and we will talk more about the type of information that one can learn from the study of absorption measurements.

DR. BASS: Let us resume. Amit Lal has an answer to Professor Marston's question concerning equation 37. Let me put that back up so that we know which one we are dealing with.

(Transparency 7)

Basically, the question is how about the classical limit of equation 37, which we would arrive at by setting h=0 or, alternatively, $h\nu = 0$.

DR. LAL: I think k is going to be proportional to $1/h\nu^2$. A perturbation calculation in quantum mechanics gives that h. So, $1/h\nu^2 (h\nu)$ will give you -- the whole thing will become $h\nu/kT$ which goes to zero. So you do get the relaxation time

approaching zero as $h\nu \pounds 0$.

A: I think that is quite reasonable. That is what we would expect τ to do. Thank you.

(Transparency 9)

Now, the fun part of this lecture: relaxation absorption. It is the fun part, because you never finish. You can make a career of doing this, and many of us have. It turns out that in order to get support you have to find the right niche, the right application for the results.

Professor Amme, at the University of Denver, does research in this area and he still gets good funding for it. He investigates relaxation processes at very high altitudes. He is interested in molecules important to such phenomenon as ozone depletion.

In those problems, the chemical reactions depend to a large extent upon the degree of vibrational or rotational excitation of reactants. Vibrational and rotational excitations occur because of photons being absorbed in the upper atmosphere. As a result, if you want to try to predict the long-term effects of these reactions, you have to know the vibrational lifetimes of the reactants.

So although Doug and I have not done real well at making a living out of this type research in the past five or 10 years, other people do and they continue studies. In fact, as you might know, Hertzberg in Canada, has studied the lifetimes of many of these molecules for astronomical applications. He made much of his living that way for a number of years and ended up getting a Nobel Prize for that work.

(Transparency 10)

I am going to start out by talking about experimental techniques. I did this because I knew Steve would be here, and he defines physics in terms of experiments.

The difficulties involved in measuring vibrational lifetimes of molecules are not insignificant. (Transparency 10b) Absorption of sound as a function of frequency, looks like the curve shown here. This curve represents the equation for α as a function of frequency for a single vibrational relaxation.

This curve illustrates velocity dispersion. (I think that goes back to Figure 2, so you have to go back in your notes.)

In order to measure the relaxation time, we must determine

either the location of the peak of the absorption curve, or the inflection point of the velocity dispersion curve, one of the two.

(Transparency 10c)

First consider the measurement of absorption. We measure absorption in some type of a chamber in which we have a gas. If we put the gas in a chamber, there will be viscous and thermal losses due to the walls. There will also be viscous and thermal walls due to the gas which are not relaxation phenomena.

We might also need to consider radiation or leakage. For certain geometries there might be spreading losses. There is also absorption due to rotational relaxation or modes other than those modes that are of interest. All of these contribute to the absorption coefficient we measure.

We must measure the change in the acoustic level very accurately if we are going to hope to get meaningful results. Let me give you an example. If I were to try to measure the attenuation of sound in nitrogen at a frequency of 10 or 20 Hz, where it is measurable -- it is near maximum at 10 or 20 Hz -- then I would find that the absorption due to relaxation is about 1dB/10km.

Now, I defy you to create a sound tube that allows you to measure attenuation that has a value as small as 1dB/10km. That is pretty tough, it is pretty tough to build a 10km-long tube, much less measure 1 dB.

On the other hand, for uranium hexafluoride you have to be able to measure attenuation on the order of a few hundred dB per meter.

You obviously do not allow a signal to attenuate a few hundred dB before you measure it. You use a much shorter propagation distance, but then you can get into all types of near-field effects, etc., if you get the transmitter and the receiver too close together. There is a huge range in the absorption that might be encountered, many, many orders of magnitude. That is one of the experimental challenges.

Another challenge concerns liquids. Liquids, it turns out, relax faster than do gases, which means that measurements must be made at very high frequency. We saw earlier that viscous losses in fluids are three or four orders of magnitude less than for gases. Other absorption losses are also three or four orders of magnitude less. Instead of having an absorption of a 1dB/km, you might have 1dB/100km, associated with absorption in liquids. That is a tough experiment.

In many cases it is very difficult to make the

measurements. Some very innovative ways have been devised to make the measurements. I will talk about the ones that those of us in this room have developed and ignore those used by others.

When I came to the University of Mississippi a number of years ago -- Doug Shields was one of the reasons I was interested. He had a unique way of measuring absorption of sound in gases.

Before Doug, most of the measurements of absorption of sound in gases relied on an interferometer. You had a quartz crystal transmitter and a reflector or receiver. Whichever the case, as you brought the transmitter and receiver together you went through resonances, in a column. The load on the crystal would change so the current through your transmitting crystal would also change. By measuring the distance change required to go through those peaks and valleys in current you could accurately measure the wavelength and velocity. By repeating the measurement at different frequencies, a dispersion curve results.

Not only that, but the current through that quartz crystal provides a measure of the amplitude of the acoustic pressure. As you move the transmitter and receiver closer and closer together, the amplitude of the acoustic pressure can be plotted as a function of distance, and from that one could determine the attenuation.

In order to perform this experiment, the quartz crystal had to be very pure, and the two surfaces had to be very well-aligned. Quartz crystals work well only at fairly high frequencies so there was not much possibility of using this method at low frequencies.

(Transparency 11)

Doug, at the University of Mississippi, had come up with a different method of making these measurements. He was a brave soul and studied some very exotic gases, things like iodine, chlorine, bromine, etc. Once I joined his research group, we never studied those again. I do not think he trusted me with the vacuum system with all those exotic gases. In any case, what he had was a transmitter and a receiver in a cylindrical tube. He varied the distance between the transmitter and the receiver and looked for the amplitude and phase of the signal as the distance between the transmitter and receiver was changed.

DR. SHIELDS: We should say this system was designed first by Anthony Angona at UCLA.

A: At the two ends of the tube there were little fiberglass cones which attenuated the sound to eliminate

standing waves, reflections, inside the tube. This tube was used for a large number of sound absorption measurements. Much of the information that we have about the halogen gases with impurities, carbon dioxide, carbon dioxide with water, etc., came out of this very apparatus.

I guess people should be reminded there was actually a time when as you changed the distance between the transmitter and the receiver, you actually wrote down the result on a piece of paper. That shows our age. How many of you have actually collected data and written the results down by hand? But it can be done. Tough to analyze, but it can be done.

Simple measurement: Amplitude as a function of distance; $A^{-\alpha x}$. You get the slope of the amplitude as a function of distance; we get α. We get the phase as a function of distance to determine the velocity.

(Transparency 12)

Time progressed and equipment became more sophisticated. More recently instruments have become available that allow you to transmit a burst of sound waves. If you are allowed to have a burst of sound waves at a given frequency (those really are not consistent statements, as some of you recognize, but let me get by with it anyway) and let it propagate down the tube, hit the opposite end, measure its amplitude, comes back, hit the transmitter, reflect off of it, go back, hit the opposite end again, and measure its amplitude again, etc., you have amplitude verses distance. As it bounces back and forth in the tube, you measure the amplitude every time the tone burst hits the receiving end.

That tone burst eventually dies out and you introduce another one, so you introduce tone bursts and allow the tone bursts to decay. This method has a lot of good things going for it. One of the things it allows you to do right off is to look at the amplitude for successive bounces, and different transmitter receiver distances and subtract out the effects of losses at the end walls. So you are left with a more accurate determination of the attenuation within the tube.

This method also allows you to average over many tone bursts, so you can improve your results just through statistics. Now we have these things hooked up to computers so this process can be automated. You can send out as many bursts as you want, you can change the nominal frequency of the tone burst, etc.

You can measure the time of flight from one end of the tube to the other. You do this over and over again for the same tone burst to get a very accurate determination of velocity.

DR. GREEN: How do you measure the distance?

A: It is the distance between --

DR. GREEN: No, I know where it is. How do you measure it?

A: There are various ways. In the tube shown here, we had a rod attached to our transmitter. The rod had a scale on it and we pushed the rod through the end plate. As we pushed the rod through the end plate we watched the scale that went through.

The whole tube here is about 4m long. So accuracies on the order of a millimeter were plenty good.

DR. GREEN: The reason I asked that, in solids it is not good enough. The inaccuracy is always the distance measurement.

A: But you are dealing with very small things.

Q: For the tone bursts, what you are trying to look at for best results is the steady-state region. The backside will have transient effects. Is that where you want to look?

A: We typically operate with tone bursts that have many cycles with a very broad band transducer, so by the time you have your first half-cycle, you have typically reached the steady state amplitude. When you go to higher frequencies, transducers tend not to be as broad band, so you do not have a burst envelope that is square. Then you have more difficulty making that decision.

DR. HAMILTON: How do you eliminate the tube wall effects from your measurements to get at what you are looking for? The reason I ask is because in my recollection you need a theory for the tube wall effects and, for example, the Kirchhoff or other classic theories in this area, as I recall, tend to be off by 10 or 20% or so from measurements I have seen. Don't you need to know that specifically to be able to extract it to get at your other results?

A: Yes, we go at that in two ways. We do correct for the tube walls. We use a little better solution than Kirchhoff's solution. If we compare our corrections to measurements made in rare gases where wall losses dominate, we typically do much better than 10%, nearer 1%.

The other thing we do is to make our tube as large in diameter as is practical. In that case the fractional contribution from the tube walls is reduced.

DR. SHIELDS: I want to say that the unique part about this experiment was that we did make these waves in a large-diameter (10in) tube. In order to do this at these frequencies you had to use frequencies far above cutoff. Ordinarily when you measure above cutoff you have to worry about transverse modes. We had a special solid dielectric transducer that was able to generate plane waves far above cutoff.

A: For the benefit of the students. When you make these measurements in a tube, you try to keep the frequency low enough so that the only mode that will propagate down the tube is a plane wave mode. I forget the exact equation, but there is a cutoff frequency, below which only the plane wave will propagate. You do not have radial modes or other types of modes propagating in the tube.

The alternative to keeping the frequency low is to generate a very plane wave initially so that other modes are not excited. That is tougher, but it can be done by making the transducers flat and very well-aligned.

Corrections for the tube wall will always be significant, but I think that for the gases we have worked with, we have had pretty good success doing tube-wall corrections to within a couple percent.

Q: The point you were making about correcting for the reflectance, how do you get rid of that?

A: If you are loosing energy out the end, you are going to lose a constant factor, let us say 0.1, for each bounce. The other absorption mechanisms are going to depend upon the distance of propagation, but the losses through the end depend only upon the number of times the burst reflects. The signal strength as a function of distance will have a slope associated with all those absorption terms that depend upon how far the wave has propagated.

The fraction that depends upon losses out the endcaps will not depend upon the propagation distance. If we make the measurement at two transmitter receiver distances this allows us, by solving two equations simultaneously, to factor out all losses at the end of the tube. At medium and high frequencies those are usually pretty small, but at low frequencies they are not necessarily small.

Q: You have to make measurements at various distances between the transmitter and the receiver?

A: Yes, we make the measurements at various distances.

(Transparency 13)

There are other techniques for measuring relaxation times.

I will discuss one, the spectrophone. Pat Arnott, in the back of the room, is now doing some spectrophone measurements for a somewhat different reason, but basically what this involves is finding a molecule that will absorb energy at a particular wavelength.

You hit the molecule with a short burst of electromagnetic radiation at that wavelength. The molecule becomes vibrationally excited. Now you turn the light off. That vibrational energy is going to re-equilibrate itself with translation. This means that the translational temperature is going to rise.

If the translational temperature rises, the pressure will increase. So when the gas is zapped with a laser pulse, a pressure rise is observed. The pressure does not increase instantaneously. There are a number of complexities. If energy goes into a high-lying vibrational state, the energy cascades down and during each intermediate transition on its way to translation it gives up a little energy. This process requires collisions, so it takes a little time. So you see a little release of energy, a little more, and a little more release of energy. Each one of these causes a pressure rise in the gas, and corresponds to an energy-transfer step.

The spectrophone can be very rich in the information that it can provide about energy transferred through collisions from high-lying energy states to translation.

Q: That would be done in a very small test cell?

A: It could be done in a very small test cell. The test cells we use are typically on the order of a few cm in length and a cm in diameter. I think that Pat is working on one that is somewhat larger.

This is the way the time signal might look. A laser pulse is absorbed in a test cell, and a hydrophone or microphone picks up the pressure. The output of the receiver looks something like this.

The first thing that happens in this example is a decrease in temperature. You add energy to the gas and the temperature decreases, the translational temperature decreases.

With that, we will have coffee.

(A brief recess was taken.)

DR. BASS: Let me talk about some experimental results.

(Transparency 14)

This solid curve represents attenuation and this curve

represents dispersion measured in the propagation tube that I described in an earlier slide. Someone asked me to point out, and it is an important point to make, that when doing the absorption calculation from measured levels in these propagation tubes, it is necessary to measure the level not only for the tone bursts as they bounce back and forth in the tube and decay, but also for different separations of the transmitter and receiver.

This allows me to solve simultaneous equations from which I can determine end losses. It is necessary that I do the measurement for different separations between transmitter and receiver as well as observe a number of bounces back and forth inside the tube.

Q: Are there any adjustable parameters which can be used to eliminate tube wall effects?

A: That is certainly possible; measurements have been done in different size tubes. We feel, though, that we can predict those losses fairly precisely and do not consider them to be of much concern. The end losses are a little tougher. Losses might vary with frequency because of the resonance of the end plate, or something like that.

(Transparency 14)

Let us look at some results. This curve right here is a curve computed from the relaxation equation that I presented earlier. These are experimental data points. Not only are they experimental data points, but some of those experimental data points are taken at different pressures.

Think about it for a minute. The relaxation absorption, like viscous absorption and thermal conduction, depends upon collisions. If you decrease the number of collisions, you increase the relaxation time. So the effect of pressure and frequency are opposite.

When I measure the relaxation absorption, f/p is the important quantity; not the frequency or pressure, but the ratio of those two. Some of these points represent different pressures and different frequencies to give the same f/p value. As you can see, they nicely overlay each other. This is further evidence that the absorption does depend upon the ratio of f/p.

That observation has an underlying meaning. If we had very many collisions that were three-body collisions as opposed to just two-body collisions, you would not see that simple f/p dependence. If you measured the attenuation at different frequencies and different pressures and plotted them versus f/p, the points would not overlay.

This transparency is for fluorine -- nasty stuff to work with. The adjustable parameter is relaxation frequency. The amplitude, the maximum absorption, can be determined from the relaxing specific heat so that is really not adjustable.

This transparency shows results for chlorine -- again, a nasty gas (Transparency 15). In fact, Doug has done not only chlorine and fluorine but also bromine and iodine. All of it is pretty nasty stuff, but he did come up with some interesting results from measuring those absorption curves.

From the absorption curves you deduce relaxation time. From the relaxation time you deduce rate constant. From the rate constant you deduce the probability of energy transfer during a single collision. It is a fairly long train -- from this you get this, from this you get this -- but each step is based on good physics.

(Transparency 16)

We finally arrive at a probability of energy transfer -- P -- versus temperature plotted on a rather unusual scale. Sometimes I will plot collision number, which is 1/P, so it is the number of collisions required for energy transfer, but here we have calculated the probability.

Note the peculiar way temperature is plotted, higher temperature is to the left -- we do anything to confuse you. There is actually a reason for this plot. If one goes through a quasi-quantum calculation of the probability of energy transfer during a diatom-diatom collision, the probability of energy transfer depends upon an exponential $T^{-1/3}$. If you take the log, you should get that the log of the probability varies as $T^{-1/3}$.

You would expect that if this simple calculation was correct the experimental points would lie along a straight line. It is not bad.

There are a few trends that are obvious. One is that the probability of energy transfer increases as the temperature goes up. You would expect that, because at higher temperature the colliding molecules carry in more kinetic energy. If they have more kinetic energy, they are more likely to have sufficient energy to cause a quantum of energy to be exchanged. The more energy, the better chance of having a transition.

There are some interesting issues involving the separation between vibrational levels and the amount of vibrational energy required for single transition. If there is a large separation between vibrational levels, you would expect it to be difficult for energy transfer to occur. For small vibrational energy-level spacings, you would expect the probabilities to be

higher. That is observed.

There is another factor that is not so obvious. That is the effect of the mass of the colliding molecules. These effects are summarized in your handouts, but let us look at them for a moment.

Q: In the figure description it says to multiply by 10^{-5}. I wanted to ask about this, too. You are just going to discuss mass. Is the probability increasing just because of the mass?

A: You have the vibrational energy-level spacings and the mass effect mixed together. It is not clear on this chart, but, typically, less massive collisions tend to transfer energy more rapidly.

(Transparency 17)

Let us look at some experimental results. First, the probability is very sensitive to the energy gap. The probability of energy transfer is very sensitive to the duration of the collision. That is not so obvious.

If you have a very sudden collision, the probability of energy transfer is higher than if you have a collision that takes place over a longer period of time.

The duration of the collision depends to some extent on the shape of the interaction potential. P^{10} also depends upon the interaction potential, because the interaction potential, to a large extent, determines how long range and how sharp the interaction is.

(Transparency 18)

In collisions you have to conserve both energy and momentum. For a transition, you have to have enough energy to cause a quantum of energy to be transferred and also satisfy conservation of momentum. Conservation of momentum is most easily satisfied if you have a very hard collision. If the collision is such that the molecules interact at the steep part of the potential barrier, the collision will be very efficient as far as energy transfer.

Consider this interaction between collision partners. At first there is an attraction which causes the two molecules to accelerate as they come together, and then they run up against this steep wall, which is like running into a barrier, and then there will be a rebound.

The probability of energy transfer depends upon the shape

of this curve. Molecules that have very large electric dipole moments, can interact over relatively long ranges, and their probability of energy transfer falls in a class different from nonpolar molecules. They have to be treated in a separate class because the nature of the interaction is different.

The probability of energy transfer also depends upon the geometry of the mode. Some vibrational modes are more easily excited and de-excited than others. For example, consider CO_2. The OCO can vibrate in a bending mode, or asymmetric or symmetric stretching modes. There are three different vibrational modes available to CO_2.

If you were going to use a hammer to drive the molecule into an excited state, which mode would be the easiest to drive? In order to drive the stretches, you have to hit the molecule on the end, so you have to position yourself just right, like when striking a nail, and hit it head on.

But if you wanted to excite the bend, you hit it anywhere, so the bending mode is going to be pretty easy to excite. The geometry of the mode is also important in dictating the efficiency of energy transfer.

Much of what we know about the dynamics of molecules and energy transfer has been learned from acoustic and ultrasonic measurements. We have begun to look at even more sophisticated systems. An example system that I have chosen to discuss is SO_2, sulfur dioxide. Sulfur dioxide is a molecule that Doug studied for six or seven years trying to map out the energy transfer pathway. It is a very complex system, but it serves as a useful example of what one can learn about the microscopic details of energy transfer in a complex molecule.

SO_2 is not like CO_2; it is not linear. But it does have those same three modes. It has an asymmetric mode, a symmetric stretch, and a bending mode. When one SO_2 molecule collides with another SO_2 molecule, you can have transfer of energy from one vibrational mode in one molecule to a different vibrational or same vibrational mode in the collision partner. But when SO_2 collides with a rare gas atom, argon for example, the only possibility is to take a quantum into or out of the SO_2.

This opens up new possibilities. If one studies SO_2 in self-collisions, then SO_2 in collisions with rare atoms, etc., you get additional information.

More insight is gained when studying SO_2 colliding with O_2. SO_2 has an asymmetric stretch that has an energy level close to the fundamental mode of oxygen. That means if I put energy into SO_2 vibration, it can transfer to oxygen and back

much more efficiently than it can transfer to translation. Oxygen and SO_2 share their vibrational energy.

This is very important, because the relaxation time for the SO_2 asymmetric stretch is much faster than for oxygen. If I put vibrational energy into oxygen, it is very slow to decay. When, I put a little SO_2 in as a mixture, the vibrational energy can transfer over to the SO_2 stretch and then go to translation. A different energy pathway has been created.

Q: Would you go through that, again?

A: Okay. Imagine putting energy into oxygen. Oxygen decays vibrationally very slowly in self collisions. Now put in a little SO_2. Energy transfers from oxygen vibration to SO_2 vibration very efficiently, because there is a very small ΔE. SO_2 vibration relaxes rapidly, because it has a cascade of energy levels to pass energy through. No step has a large ΔE, so it serves as a shunt to the normal vibrational relaxation process of O_2. The relaxation rate for oxygen, then, is much more rapid than it would be in the absence of the SO_2.

Q: I do not know if you are going to talk about it, but I wanted you to talk about air and water.

A: I am going to. The same thing, oxygen and H_20, in that case, as opposed to SO_2 and oxygen. But it is the same phenomenon.

To study this phenomenon we make measurements at different concentrations of SO_2 and oxygen. Even in pure SO_2, the absorption cannot be described with a single relaxation time -- the width of the curve is a little bit too broad -- so SO_2 is one of the few molecules with observable multiple relaxation.

(Transparecy 19)

Q: What is the horizontal axis? Is this inverse cm?

A: That is f/p. Do not worry too much about that, because we are going to shift that around as we progress. The vertical axis is attenuation per wavelength.

(Transparency 20)

If I put in more oxygen, two things happen. One is the emergence of another peak -- here -- at lower frequencies, and a decrease in the amplitude of the high frequency peak. It is tempting to associate the low frequency peak with oxygen relaxation, and the high frequency peak with SO_2 relaxation. But if you use the Einstein relation to compute the vibrational

specific heat of oxygen and use that vibrational specific heat to compute this low frequency absorption amplitude, the measured absorption is much too great. The relaxation absorption for oxygen should be much less.

So what has happened is that SO_2 is sharing its vibrational energy with oxygen. The oxygen peak is effectively larger because of the presence of the SO_2 and because these two vibrational levels are strongly coupled through collisions.

To continue this, put in more oxygen, you get even more sharing. (Transparency 21) When you finally go to pure oxygen, the low frequency peak goes way down in frequency and magnitude. So what? Who cares?

If we examine an energy level diagram for SO_2 and study the shifts in relaxation frequencies amplitudes, we can find out how energy is transferred between the modes of SO_2.

(Transparency 22)

There are a lot of possibilities. It could be that oxygen shares energy with any of several modes. There are a lot of other ways that the energy could be transferred during a collision. But, in fact, from measurements we find that the way the energy flows is through the asymmetric stretch of SO_2.

We are able to assign a rate constant, and a probability of energy transfer for each energy-transfer step. So we are able to map out how the energy flows from one molecule to the other and between energy levels by using these ultrasonic absorption measurements. It is a very difficult process, and it turns out that when creating such a map we might need optical measurements to complete the picture.

There are certain laser techniques that allow you to isolate some of the v-v processes, for example, and identify rate constants for those processes. The spectrophone also provides additional information. It is a matter of using different experimental techniques to build up the overall picture. But this picture was pretty much mapped out on the basis of ultrasonic absorption and dispersion measurements.

Q: Why did you choose SO_2?

A: Doug chose it. Why did you choose SO_2, Doug?

DR. SHIELDS: That was the only gas at that time that showed multiple relaxations. Everybody said there ought to be multiple relaxations but SO_2 was the only one around that demonstrated it. Since the modes there are already separated, then we had another chance of mapping the coupling into a

particular mode.

Q: What is the wavelength of the $518cm^{-1}$ line?

A: This turns out to have a wavelength of about $15\mu m$. You are talking about the electromagnetic radiation.

Q: This can make like a laser kind of structure if you can just excite the SO_2 molecules and transfer energy to O_2. It is the same principle as the laser.

A: It turns out that for SO_2, the bending mode at $518cm^{-1}$ is only weakly infrared active. You do not have a strong infrared dipole to work with. The asymmetric stretch at $1361cm^{-1}$ is strongly active in the infrared so you would have a better chance of getting population inversion on that mode.

Q: The reason that SO_2 is more efficient in converting vibrational energy into translational is because you have all these steps in-between?

A: Transitions take place, during collisions. During a given collision there has to be sufficient energy available to cause a quantum of energy to exchange. There are a lot of molecules and a lot of collisions, but most of them have fairly low energies. If you could have a number of low-energy events add up to a large energy event, then there are a sufficient number of the small energy events to do so.

But quantum mechanics does not allow us to just add one event to the next; either it happens or it does not during every event. In SO_2 there are a number of discrete steps that allow you, one collision at a time, to go through a series process. There are only five steps. If each step had unit efficiency, only five collisions would lead to a de-excitation. There is only one molecule in 10^5 or so that has enough energy to cause one big transition to take place.

(Transparency 23)

There are other interesting examples, but I am running a little long. Relaxation of CS_2 is an interesting example, because it is easily handled in both the liquid and gas states. The liquid behaves as though it is a very dense gas. There are some liquids that behave that way --Kneser liquids.

Essentially, treating the CS_2 liquid as though it were just a dense gas, means increasing the collision frequency according to the increase in density. We find out that the relaxation time is predicted well just by considering this increase in density.

Most fluids are not that easy. That is why we have spent a lot of time with CS_2; we try to deal with the easy fluids. Unfortunately, one of the most interesting fluids, water, is tough. In water, if you put in a little thermal energy, the water changes. You do not have the same internal structures that you had before adding energy. Every time you put in a little thermal energy, there is shifting of chemical bonds --we refer to this as configuration relaxation. We try to avoid dealing with water when at all possible. Unfortunately, the Navy is interested in water.

(Transparency 24)

Let me show you an example where our understanding of energy transfer has been successful in making predictions. It has to do with the absorption of sound in air. I am going to go through this rapidly, although it is a very interesting example. It is outlined reasonably well in your handouts.

Air is of obvious interest to most people. The study of absorption of sound in air goes back well beyond this century and perhaps multiple centuries. It is a natural thing to wonder what happens to sound once you speak it into the air. Where does it go? What happens to it?

One of the things that happens, of course, is that viscous losses and thermal conduction losses cause the energy to be dissipated. Those losses we talked about an hour or so ago. We know how to compute them. All we have to do is put in the proper numbers, and we can predict the classical absorption. People have measured the viscosity and the thermal conduction of air to 189 significant digits for the past 400 years, so you do not have to guess about this stuff. It is pretty well-documented.

The kicker is the fact that sometimes air has water in it -- in this part of the country you do not have so much as back where I live -- but a small amount of water can change the thermodynamic properties of air very dramatically. So the problem is not as clean as I might have led you to believe; maybe we know the absorption to only two significant digits as opposed to three, but compared to other contributions it is still pretty good.

Thermal conduction and viscosity come about microscopically for the same reason; collisions transfer translational energy and momentum. From the transfer of energy, we derive conduction; from the transfer of momentum, we derive viscosity. On the microscopic level, these occur during the same collisional processes.

You would expect, then, that there should be, at some fundamental level, a relationship between viscosity and thermal

conduction, and there is. There is quite a nice relationship between thermal conduction and viscosity. It is called the Eucken relation between thermal conduction and viscosity. Both viscous and thermal loss increase as the f^2. So I have some constant times frequency^2 for the classical absorption of sound in air, probably good within a couple percent.

Another contribution is due to rotational relaxation. The rotation of nitrogen and oxygen contributes some attenuation because of relaxation. The relaxation frequency of those modes is very high. In air it is about 100 MHz at atmospheric pressure. Even very high in the atmosphere, the relaxation frequency is much above that of interest for acoustic propagation studies.

In the low-frequency limit, relaxation absorption behaves very much like viscosity and thermal conduction. In other words, it is proportional to f^2. Measurements in oxygen and oxygen-nitrogen mixture measurements, show that it takes about five collisions, on the average, to de-excite rotationally a nitrogen molecule or an oxygen molecule. Using this result, I can come up with a relationship for combined classical and rotational relation losses in air. This is probably accurate to within a few percent -- you can never test it, because it is so small compared to relaxation absorption.

The "biggie" is vibrational relaxation of oxygen, nitrogen, and perhaps, in some cases, CO_2. It turns out that these take patience to predict. Consider all the vibrational energy-transfer mechanisms which must be taken into account. There is the transfer between oxygen and nitrogen. Oxygen can de-excite nitrogen in a VT transition. Nitrogen can de-excite oxygen in a VT transition. Oxygen and nitrogen can exchange energy through a v-v interaction. We have to include the probabilities for each of those events.

There is also energy transfer between oxygen and water vapor. That is an interesting one, because water vapor has a symmetric stretch which is almost resonant with the oxygen vibration so there is a very rapid exchange between oxygen and H_20. Now, H_20 self-relaxes very rapidly, so you put just a little bit of H_20 in with oxygen to provide a shunt, an alternative pathway, for relaxation of oxygen that is very efficient. Because this process depends upon the product of H_20 times H_20 you get a quadratic dependence of relaxation frequency on H_20 concentration. You also get this in the case of nitrogen.

H_20 and nitrogen do not have vibrational energy levels that are close, so they do not transfer energy as rapidly, but, still, water has a major effect on how rapidly nitrogen de-excites. CO_2 effects, especially on oxygen, are not

unimportant. For very low levels of H_20, the 335 ppm CO_2 in the atmosphere provides a more efficient way for oxygen to de-excite than through oxygen-oxygen collisions. These interactions must be included.

If you write down all of the processes, all of the different types of collisions that can have a significant impact on the relaxation times of nitrogen and oxygen, in the atmosphere, there are 28 of them. It has been a very daunting task to measure these 28 rates. It has not been easy, it took a while, but there is nothing like job security.

Many of the rates required were measured not using acoustic techniques. The lifetime of the CO_2 ν 3 mode is most easily measured using a laser fluorescence technique. There are other rates that were measured using other than acoustic techniques.

The drill is to measure the rates one at a time. You measure the vibrational relaxation times in different mixtures of oxygen and nitrogen, and from that you determine the rate that nitrogen de-excites oxygen, oxygen de-excites oxygen, oxygen de-excites nitrogen, nitrogen de-excites nitrogen, and the rate at which they exchange energy between them. This took Bauer and some colleagues of his seven or eight years.

Then you study CO_2 and oxygen and go through the same process. It took Doug and me almost that long with CO_2 and oxygen. Then you do H_20 and O_2. Doug worked on that a long time as did others. Then H_20 and nitrogen. Allen Zuckerwar and Doug worked on that rate.

Once all the rate constants have been determined, you put them into the equations that govern absorption of sound in air.

(Transparency 25)

You measure the absorption of sound in air and you compare measurement to theory. It works.

(Laughter)

The only measurements that were missing when I entered this process the CO_2-O_2 and the nitrogen-H_20 rates. Most of the others were in place.

Q: How far different were the measurements from the theory as far as the relaxation effects?

(Transparency 26)

Here, this shows you the magnitude. If there are no

relaxation effects, you get the lower curve, so relaxation absorption dominates. Without relaxation, you basically would not have absorption of sound in air; it would not be a measurable quantity.

DR. SHIELDS: This curve shown spans a region where the rates were not known. At lower humidities, we were able to establish a better rate from the relaxation in oxygen from these measurements. That has no affect over the range where this theory was applied. But there was some deviation and we were able to correct the deviations by using a proper rate constant.

A: This is a typical comparison of measurements and predictions (Transparency 25) for absorption of sound in air based upon 28 rates, all of which, in certain regimes, can cause a large difference between measurement and experiment. That means you have to know every one of them to cause this type agreement.

(Transparency 26)

Q: Just from that curve it seems that there are two linear regions, one in the high frequency and one in the low frequency, and there is some connecting link. Does that explain it?

(Transparency 25)

A: Let me go back to this curve. In the low-frequency limit, well below nitrogen relaxation, absorption increases as f^2. In the frequency where relaxation is occurring there is a rolloff. Well below the relaxation frequency for oxygen the f^2 is observed again. In the region of relaxation of oxygen, it begins to fall off. So what you get is up, over, up, over, up, until all the relaxation processes have been worked through.

Q: Is there a master equation somewhere that you could distribute for computing attenuation?

A: There are two answers to that question. There has been an American National Institute standard based on this calculation. It is a standard for absorption in the atmosphere. The ISO standard uses the same equations. These give empirical equations that come from the calculations.

We also have a computer program that takes the fundamental rate constants and the vibrational energies associated with each of the reactions and computes these same curves from a microscopic point of view. It is a little bigger, as you might expect. It is very useful if you make a better measurement of a rate constant or if you want to incorporate spectrophone results in a gas mixture which is different from the standard.

We have such a program and you are welcome to a copy.

Q: How much does the speed of sound change as you keep increasing the frequency?

A: The dispersion in air is fairly small. There is some humidity dependence of the velocity of sound in air, because of the molecular weight term, but if you look at the dispersion in this frequency range, it is essentially negligible. At higher frequencies, where rotational relaxation begins to become important, dispersion becomes measurable.

DR. HAMILTON: This curve here and those on the previous slide include many relaxation frequencies, some of which cluster together and give you what appears to be separate peaks in the spectrum, is that correct?

For the curves you have up there right now you have two distinct relaxation frequencies and you just add those two components. How well does a simple theory like that, which is often used for sonic boom propagation, work?

A: I have always felt it was pretty good. The reason for my optimism is that oxygen and nitrogen coupling is fairly slow, so you do not get exchange of energy, specific heat, between those two modes very much, not like you do with SO_2 and O_2.

Rich Raspet argues that there is, perhaps, an exception to this at very high altitudes where the relaxation frequencies of all these processes shift to very, very low frequencies. In that case, the role of CO_2 apparently cannot be ignored. CO_2 vibration is very strongly coupled to O_2, and it shares some specific heat with O_2. He thinks he can argue for absorption associated with CO_2 relaxation. That makes things much more complicated. If you had asked me that a month ago, I would have said they are separate, you can ignore coupling, but I am less certain right now.

DR. GARRETT: I would just like to make a comment that applying this kind of theory to sonic booms is going to be misleading, no matter how good the theory is, simply because for larger amplitude waves the attenuation process is unaffected by the particular relaxation or absorption mechanisms.

A: That is wrong.

DR. GARRETT: It is wrong?

A: Yes.

DR. GARRETT: It works for sonic booms?

A: It works very well. What happens is that absorption, attenuation of the boom, is constantly at work. You do, of course, have nonlinear effects -- you are going to tell us about this. (Dr. Garrett is trying to get some publicity for his talk on my time.)

(Laughter)

You can look at nonlinear effects in terms of harmonic generation or wave steepening. Relaxation absorption constantly serves to counterbalance those. Molecular energy-transfer is always present, and will limit how sharp a wavefront can get. The limit to how sharp a wavefront can get is dictated by these same relaxation processes whether you are operating in a linear regime or in a nonlinear regime.

Many of these rate constants are measured in shock tubes where you have a very high-amplitude wave and look at the decay of that amplitude. The relaxation times of the microscopic processes and the link to the acoustics is in no way limited by the linear acoustics approximation.

I have added them to the linear acoustics approximation in my derivation but all of the physics is the same.

DR. ATCHLEY: Whatever waveform you have is a time-domain waveform and harmonic generation by nonlinear processes does not matter, because you know the waveform in time and you let that be operated on by this formula.

A: The same physical processes are always occurring. If you have a sine wave going through the atmosphere, or a sudden disturbance going through the atmosphere, these same energy distribution processes are governed by the same rates at work.

A very high-amplitude shock wave can cause a significant temperature change or change in the chemical composition of the atmosphere. I am going to treat that, if I have enough time, in just a few minutes.

(Transparency 26b)

I will show this transparency for Steve. We used the same rate constants to determine the rise time of shocks emanating from nuclear explosions, from small blasting caps; from small amounts of C4, 3006 rifles, and tank canons firing off. The predicted waveforms and decay rates compare well with experiments.

I have 35 minutes left and I want to cover some new topics.

(Transparency 27)

I want to mention some topics that I have waved my hands over. There are absorption mechanisms that I have ignored. One of those absorption mechanisms is that due to diffusion. Consider diffusion a moment.

(Transparency 28)

Consider a mixture of two gases with different molecular weights. If you cause the temperature to change, the lighter molecules are going to respond more rapidly to that temperature change than the heavier ones because they have a higher velocity. There is going to be a natural separation between heavier and lighter molecules, diffusion.

The light molecules are all going to hurry over and be trapped in the low-temperature trough and the sound wave is dampened. There is going to be absorption associated with this diffusion. The calculation is not tough. It is not given in the typical text, but it is given in the reference in the back of your handout.

In most cases, diffusion is not a very important mechanism for absorption for two reasons. First, in air, oxygen and nitrogen have almost the same mass so the mass difference is small and so is diffusion. The coefficient of thermal diffusion for air is pretty small. So, in air at least you can ignore diffusion. It is down a couple orders of magnitude compared to thermal conduction.

But that is not always the case. If you take a mixture of hydrogen and UF_6, this could become the dominant source of absorption. But there are other cases where diffusion absorption can, in fact, be dominant. Mixtures of helium and xenon, for example.

Q: Does γ change if you have a mixture of diatomic and polyatomic molecules?

A: Earlier, the equation we derived was for two gases that had the same γ. To answer your question, in these equations γ is a concentration averaged value. At least, I think that is the case.

There are other potential applications of this equation. One of these is propagation in a dusty gas. Consider a mixture of gases like xenon and helium; very massive particles and very light particles. But if you make the massive particles even more massive, you might visualize very large molecules embedded in a gas which has very small mass. There has been some work done through the years by the group at Yale on such mixtures that are referred to as dusty gases. In these mixtures, dust

particles with a mass 10 or 100 times the molecular mass are mixed with a gas. The waves that might be supported individually by those two types of media are strongly coupled.

Going to the extreme, we can go over to a Biot solid. A two-minute digression. Think for a moment about sand. Sand has air pores in it. If you irradiate sand with a speaker, you can set up waves in the sand matrix or in the pores. The two are coupled, because if the matrix shakes, it causes pore motion; if the pore fluid moves, it causes the sand to move. We call that a Biot solid; it supports two types of waves, one in the matrix, one in the pore fluid.

Now imagine that you fill the pores with water. There is very little density contrast between the matrix and the fluid. There is not a nice distinction between the two. Let us go one step further. Make the matrix of very small particles, so that they essentially float like a suspension; they are no longer in contact.

There is now a transition from a matrix filled with a fluid, a gas, to one where you have a dusty gas, dirty water -- I guess dusty gas and dirty water are about the same thing. To my knowledge, no one has either applied this approach to saturated sediments or have they done a measurement to show the similarity.

Q: That might have some applicability to the composites problem, particularly.

(Transparency 29)

A: Another case that I want to discuss, but I do not understand is acoustics at very low pressures. I talked earlier about making the continuum approximation. I assumed that the gases that we dealt with could be treated as though they have a viscosity and thermal conductivity, as though they were continuous.

At very low pressure you really have to worry about whether or not you can treat the gas as a continuum. In that case, we invoke the Boltzmann equation. The Boltzmann equation takes the place of our continuum model. The Boltzmann equation describes the media in terms of a distribution function for the velocity, which is a function of position and time. Collisions cause changes in that distribution function.

An acoustic wave in a media causes a time dependent change in the distribution function. I can no longer define such things as pressure which are continuous media properties, but I can speak in terms of the average values of position, velocity, etc.

Collisions between molecules change the distribution

towards the equilibrium case. In order to find the time evolution of this distribution function it is necessary to know the collision cross-section for energy and momentum transfer between individual atoms and molecules. These are the same quantities required for the relaxation calculation in air.

One must invoke Boltzmann mechanics to describe this case, a regime where the translation is not keeping up with external perturbations. There is now translational relaxation.

(Transparency 30)

An excellent series of measurements made by Moe Greenspan illustrates this case. For those of you who did not know Moe, he was at the National Bureau of Standards almost forever and made a large number of measurements in a variety of gases.

Typically, in a gas attenuation due to classical mechanisms increases as the square of the frequency. Decrease the pressure, same effect until some value is reached where the collision frequency is approximately equal to the acoustic frequency. On this transparency that corresponds to $\omega\tau = 1$. Above $\omega\tau = 1$ there is a rollover and the attenuation actually begins to decrease. This is very interesting. Acoustic waves propogate through collisions. Some information is lost as the wave propogates through the gas because collisions do not happen instantaneously.

At very low pressure, collisions do not take place very often. If you have a transducer and hit a molecule, the molecule might go all the way to the receiver without hitting any other molecules. Any information imparted to this one molecule gets all the way over to the receiver without being lost. Collisions are necessary to transfer energy in the normal way.

But in the absence of collisional processes there is still a possibility of transferring information through individual molecules. That is what happens at very low pressure or high frequency. The description of this propagation is more complex. Since Moe Greenspan left this field there has not been much work devoted to propagation of sound through very low-pressure gases. This lack of interest is surprising. Evolution of galaxies and very large explosive events on a galactic scale involve energy transfer and very low pressures.

In this regime you quite often want to deal with enough particles to be statistically significant and track them over large distances, so the computational time is not cheap.

(Transparency 32)

This transparency addresses free molecule propagation. I

wanted to point out that for free molecule propagation both the absorption and phase velocity depend upon distance. A propagation velocity which depends upon distance is troublesome. We like to think of the propagation constant as being a constant. In the free propagation regime, the molecules that go fastest are going to get to the receiver prior to being scattered. Up close, the receiver will encounter all the molecules hitting. Those molecules will have a range of velocities from low to high. Far away only the fastest reach the receiver. So the velocity you measure depends upon how far the receiver is from the source. That is a prediction. To my knowledge, it has not been observed.

(Transparency 33)

Another interesting effect that Doug and I worked on for some time, is the amplification of sound as it propagates through the same energy-transfer processes.

If we can maintain an inversion of some type in the populations of the various states -- what that means acoustically is different from optically -- then as the sound propagates through the gas, one can observe net gain instead of absorption. Doug has demonstrated gain in a few gases.

Another good example of this effect, which we will hear about from Greg Swift, is a thermoacoustic engine.

(Transparency 34)

One more topic: thunder. Everyone knows about thunder. Steve Garrett was asking about nonlinear waves. Is there any acoustic event more nonlinear than thunder? Overpressures reach 100 Atm.

DR. ATCHLEY: You know, Seth will argue with that when he comes here.

A: I know. He would argue with anything, though.

(Laughter)

Let us think about thunder a moment. In the atmosphere electric fields build up. Then some type of disturbance causes electrical conduction to begin to take place from one part of the atmosphere to another. The path is rather tortuous since it depends on the position of free electrons in the atmosphere at that particular instant. When the conduction path is opened, there is a lowered resistance along this path where current can travel. The current travels very rapidly leading to a rapid heating of the channel and an immediate expansion, supersonic expansion. As the expansion slows down, there is an acoustic wave or a shock wave, thunder, propagating outward, which you hear at great distances.

Everyone likes to sit around and listen to thunder and watch the lightning and think about what is taking place. Let me tell you what is taking place.

(Transparency 34)

In order to describe thunder you must account for the energy, deposited in a discharge channel by the flow of current, i^2R. No big thing. But what is R? The conductivity of the channel is challenging to predict. The conductivity of the channel depends upon the number of free electrons available along the channel at any instant of time. The number of free electrons available along the channel at any instant of time depends upon the temperature of the channel. The temperature of the channel changes very rapidly as the current increases.

So you have a resistance or a conductivity that is a strong function of time. You must solve for this conductivity as a function of time by assuming some type of temperature rise and then going back and checking for self-consistency.

This calculation requires knowledge of the temperatures at which the various atoms ionize and the lifetime of those ionized states. It is the ions and the electrons from those ions that conduct current. You must know a great deal about the microscopic rates that take place within the lightning channel. This really goes back to detailed modeling of the energy-transfer processes, in this case including electron transfer.

Myron Plooster, at NOAA in Denver, set about trying to solve this problem beginning in 1962. In 1980, he finally collected a sufficient number of rate constants to predict, from a microscopic level, how the shock wave would evolve as a function of time. It took him a long time, but he did it.

(Transparency 35)

This transparency show the computed conductivity, as a function of temperature.

(Transparency 36)

This transparency shows the input to the discharge channel, the current for i^2R heating.

(Transparency 37)

And from the energy dissipation, Plooster computed short waveforms. Shown here are such quantities as pressure density as a function of time. You can see the buildup is very rapid leading to large overpressure, density and temperature as time

progresses.

(Transparency 38)

Now, we will examine a typical shock overpressure as a function of the radius of the lightning stroke. This is how the experimental observations appear.

Plooster does a full hydrodynamic calculation, including the time dependence of the electrical conductivity, etc., out to a time where these terms have become ignorable, when the shock separates from the lightning channel. At this point the overpressure of the shock is typically on the order of only a few atmospheres. From this point on out to the listener we apply standard acoustic techniques to compute the amplitude and shape as a function of distance. Very far away from the lightning stroke you see the effects of sound from different branches of the stroke coming together. These shocks form caustics and all kinds of complicated stuff, which we have a tough time dealing with. Those effects make it difficult to compare to the theory.

Up close to the stroke, where you see a direct lightning channel we predict the shape, energy, and time duration of that thunder reasonably well. That is the crack you hear. When you are close, thunder goes "crack." Far away where you hear "Boom, boom, boom, rumble, rumble, rumble," we lose our ability to predict quantitatively.

Q: The Maxwell stresses are another way the magnetic field pressures enter into the calculation. Is that a significant term?

A: In these calculations it is not. Whether or not that is ever significant is subject to debate. There is a group at Rice that has argued for a number of years that these are important factors. They have not received a lot of attention, because the calculations that are done based on the model I described, agree so well with the data.

But the people at Rice would argue that the electromagnetic contribution will lead to a strong angular dependence. You should observe the largest effect at some small angle beneath the charged region.

Most of the measurements have been made in the very far field where the angle is large. So you cannot prove them wrong.

Q: What is the distance from the lightning bolt typically where you can use simple acoustical theory without worrying about ionization problems and so forth?

A: On the order of 50 channel diameters. A channel

diameter depends precisely on the energy, but it is on the order of half a meter.

Q: Could you put the graph back up with the channel diameter on the horizontal axis where you plotted a bunch of data points and had one straight line?

(Transparency 38)

What is the vertical axis?

A: That is shock overpressure as a function of the radius of the channel.

Q: So the shock wave pressure increases as the radius gets smaller? Why?

A: Small radius means closer to the source.

Q: I thought you were talking about the radius of the bolt itself. I am sorry.

Q: You mentioned there is very little dispersion in air. When you hear a thunder clap very far away, you definitely hear high frequencies and then low frequencies.

A: Good question, and it is very fundamental. Recall that absorption in many cases is proportional to ω^2.

Very far away from a broad-band source we do not hear the high frequencies. The farther away the source, the lower the frequencies we hear.

For such an impulse, if you get it very close, you hear all the frequency components, a sharp crack. Move away a little and the sound is not so sharp a crack. Even farther away, it is a boom. As you go farther and farther out, you get only the lower and lower frequencies.

The evolution of species, of insects, of birds, of people takes this effect into account. Our range of speech communication is dictated by absorption which is proportional to f^2. Bats use high frequency only when they are at close range because absorption is proportional to f^2. Little bitty bugs that communicate over short distances do so at high frequency. If an elephant wants to call another elephant 15 km away, what frequency does he use? Very low frequency because it is only that very low frequency sound that will propagate 15 km. It is logical that they would use low frequencies; they could not communicate over large distances at high frequencies.

DR. ATCHLEY: I think the point Bob was trying to make, though, is that thunder goes on for a long time. I think part

of the answer to your question is that lightning bolts can be very long and very tortuous, so you have a lot of distributed sources, a km every three seconds or so, if you have a 10km long channel.

Q: It seems to me that you hear the high frequency from a lightning clap before you hear the low frequency.

DR. ATCHLEY: It could be if you have a path that goes, for instance, like this, you will hear the high frequency from this, and then propagation delay; later you will hear the low frequency from that part, because high frequency from that part has been attenuated. So you are not hearing sound from the same source location.

Q: My question is related. Why does it go crack, rumble, rumble, instead of rumble, rumble, crack?

DR. HARGROVE: If it goes crack, rumble, rumble, crack, rumble, rumble, but you think it goes rumble, rumble, crack, that depends on when you tuned in.

(Laughter)

Q: I think the sound could follow different paths. The high frequency follows the direct path and the low frequency goes some other.

A: Low frequency sound follows a more direct path due to diffraction. High frequency sound follows paths dictated by ray acoustics. They are not necessarily the same.

I have, in the past three hours, tried to get you to think about acoustics in terms of processes that take place on a scale that is very small. Hopefully, you now understand that there is a relationship between microscopic processes in physics and acoustic propagation.

These microscopic processes can be used to explain acoustic observations and acoustic studies can be used to study these processes. Throughout the history of the science it has worked both ways. It is important to understand that there are acoustic waves in outer space. There are long wavelength waves propagating on the galactic scale. They are not the same frequency we are used to, but they can be used to study processes at the galactic scale.

If we understand the basic phenomena that govern acoustics, the transmission of energy by a mechanical wave, we have the ability to study all those processes.

Q: There seems to be a hotbed of activity measuring the propagation and damping characteristics of sound in gases at high pressures because they want to build acoustic microscopes

using air coupling and so on. Anything important there?

A: I think that there will be some interesting things emerging. If you increase the pressure some, everything scales pretty well and we can predict changes. But you are going to reach a point where the pressure is great enough that three-body interactions cannot be ignored.

If you start seeing three-body interactions, there are possibilities for some interesting physics. I am not sure what the effects will be. That is not a regime that has been explored up to this time.

DR. SACHSE: A second comment is brought to mind

-- Steve's shirt has a picture of the Challenger or something -- I was in Beijing in 1986 and visited an acoustic institute. They were very proud of the fact that they heard the Challenger blow up. Is that possible?

A: Yes. In fact, as some of you might be aware, there was a program for several years in this country to monitor nuclear explosions. It was not uncommon for the low-frequency sound to propagate around the earth one way, and propagate around the earth the other way. By measuring the arrival from both directions you could determine approximately where the explosion occurred. Sometimes you see multiple passes.

Q: How low a frequency?

A: I do not know. I have seen the Beijing data, also. I think the Beijing data were about a 0.1 Hz. I think that is not atypical.

For such events you get a ducted wave trapped by the upper atmosphere. The energy cannot escape because of refraction. There is almost no attenuation, so it can just go and go and go.

DR. GARRETT: Cook, who is at the Bureau of Standards, was really an expert in this, what is called infrasonic propagation, and he claimed that in Gaithersburg he could hear the surf cracking in Malibu.

(Laughter)

I heard him make that claim, that low-frequency line sources generated due to large storms in the Pacific could be picked up on the Atlantic coast.

A: Ted Brown at NOAA in Boulder and the University of Colorado also did a lot of work in that area. There is a reasonable body of literature to suggest that you can detect low frequency sound at very long distances.

The wave is trapped in a mode created by the upper atmosphere and the ground. There is leakage on both boundaries, so it is a leaky waveguide. The leakage was not all that well-defined when these experiments were done.

Q: How do you measure these waves? I mean, it is just weird to me to measure a 0.1 Hz signal when you do not even know when it is coming?

A: If you had a pressure that was varying on a time scale of a day you would use a barograph, a barometer. A microbarograph operates under a very similar principle. You make the time constant of the microbarograph be 0.01s and you use that to observe the acoustic wave. Steve is going to tell me a better way.

DR. GARRETT: Cook used a capacitance microphone, but because of the problem of the impedance going to infinity he would run it as an oscillator, so it was an FM microphone. It would generate a MHz or something. As you changed the spacing of the capacitor plates, that would change the C in an LC circuit, so you could get responses down in the middle Hz area that bridged the gap between barometry and microphones.

A: There is actually a company, GUS, Incorporated, that built those for a while. I think they are out of business. Cook died and they did not have anyone else to buy them.

Q: Could you focus an explosion on the other side of the world, I mean if you had your source spread out cylindrically? Would it then focus? Is it isotropic enough that you could actually recreate, almost, the same explosion on the exact opposite side of the world?

A: The focusing element has to be on the same order of the wavelengths, and the wavelengths we are talking about now are 10s or 100s of kms, so it makes it pretty challenging.

Q: How do you separate the effects of turbulence in the atmosphere from the signal?

A: It is not clear which is signal and which is turbulence. At that scale, turbulence is signal. Turbulence on the mean gives you more or less zero. If you look at an explosion, it is a variation about that zero.

On the other hand, if you look at things like diurnal heating, is that turbulence or is that an event? I guess it depends on how you make the interpretation. Also, infrasound arrays are sometimes used to look for storm fronts. Is a storm front a large turbule or is it an event? Well, it is both, isn't it? It is a disturbance in the atmosphere about the mean, so in that sense it is an event.

I do not think there is a clean answer to your question. We average over an appropriate period of time and average out all those turbules that are smaller than what we want to observe.

DROPS AND BUBBLES: EXPERIMENTAL ASPECTS

DR. MARSTON: I was asked to speak on experimental aspects of drops and bubbles, and I took that as a message to avoid talking about equations as much as possible, and that is what I will do, avoid talking about equations.

(Transparency 1)

I refer you to discussions that are either in the material you have been given or in the references listed there. This talk would not be possible without contributions from several present or former students, many of whom are listed here, and only a few of the students have been concerned with drops or bubbles in their Ph.D. work, and I have underlined those. One of the reasons why there are other students listed is that sometimes I have suggested a drop or bubble problem in the first couple of summers of their work in my program to learn something about experimental methods and then they might go on to something else.

(Transparency 2)

Since I was asked to speak about experiments, I am going to take a brief liberty and say something about my favorite experimentalist. He is Michael Faraday. There is something about dedication in his life. He basically educated himself by becoming a bookbinder's apprentice so he could read the books he was binding. Then he applied for the position at the Royal Society as an aide and they were fortunate enough to hire him. In the early 1820s he was asked to give a report on the relationship between electricity and magnetism, and at that time he was able to repeat all of the previous experiments and discuss them in his report. Then, of course, later he discovered electromagnetic induction, which is probably his most important discovery.

But if you have any doubts about the importance of experiments in the development of theory, I would encourage you sometime to look up Maxwell's discussion of Faraday in Maxwell's collected papers. After Faraday's death, there are a couple of things that Maxwell wrote about Faraday, and it is quite clear that Faraday really had a major impact on Maxwell's thinking, and others, and he was under a great deal of attack. He proposed the idea of field theory and a lot of people did not like it.

Next I wanted to say something about modes of thought. This is from Maxwell's collected papers, the preface thereto. Niven wrote: "Whenever the subject admitted of it he had recourse to diagrams, though his fellow students might solve the question more easily by a train of analysis." In any event, Maxwell was clearly at ease with many ways of thinking

about problems. I would encourage you to study that quote in more detail when you get the copies of the transparencies.

(Transparency 3)

We will show some computer calculations, but we will also look at the interpretations thereof. Drops and bubbles is a big field. Let us take bubbles alone. After I prepared these transparencies I received this in the mail. It is Leighton's book on "The Acoustic Bubble". It has roughly 1400 references, most of which I have not read.

(Laughter)

Anyway, therefore, this follows, a quote from my favorite character from C. S. Lewis: "If you had felt your self sufficient, it would have been proof that you were not."

Just a few other things about experimental physics. History and experiment in physics are interwoven. I think in an audience of this size someone would enjoy these books. I just bring them to your attention: this remarkable book on the history of experimental mechanics, that is, the mechanics of solids, and Everitt's book on Maxwell discusses a great deal of the history of kinetic theory.

I actually brought copies of these (Braginsky and Welford), because I thought somebody might be interested in them. Welford's book, recently published, has a lot of interesting optical tricks in it. Then, if you are interested in surface tension phenomena, which is where we are going, Boys' lectures from basically the 1890s are remarkable.

Then, concerning introductions to bubble dynamics, I would refer you to, though somewhat dated, these articles. (Clay and Medwin's book and Prosperetti, 1982 and 1984).

(Transparency 4)

Why bubbles? There are many reasons. Some people think that bubbles play an important role in the climate equation, because the bursting of bubbles at the surface of the sea produced naturally by wind, that is, depending on the sea state, is thought to affect the production of salt nuclei and thereby affect the rain rate. So bubbles have really global significance, it is asserted.

In this picture you can see persistent bubbles in the wake of a boat and they show up white, and we will talk a little bit about that later. This is from a Bakerian Lecture. A previous Bakerian Lecture dealt with the subject at the top of the overhead. One aspect shown is scattering. Does anybody recognize whose diagram that is? It is Maxwell's. This was, actually, from a series of papers that began with his Bakerian

Lecture on kinetic theory. So we will talk about scattering, but not in the sense of molecular transport.

(Transparency 5)

This is an outline of where we are going to go. I suggested some natural breaks. Also, after each of these sections I will stop and entertain questions, because there is sort of a logical succession. If you think I have made a misstatement, please interrupt me and get that resolved quickly.

(Transparency 6)

We are going to start out talking about levitation. I would like to just look at the general principles of levitation. Oil drops and small gas bubbles in water are attracted to pressure antinodes. One can write down equations to describe this, and they are given in some of the references, but I would like to talk about one of the basic mechanisms that is important in the long wavelength limit, because that is principally what we will be interested in, in acoustic levitation devices.

We are all familiar with calculation of the force on a fluid mass, the negative of the volume times the pressure gradient. In the long wavelength limit you can consider that the time average of this force is the time average of the product of the oscillating volume and the local pressure gradient and, in the very long wavelength limit, neglect the scattering.

There are a number of corrections to this, and you can see papers, say, by Yosioko and Kawasima [Acoustica 5, 167 (1955)] and others that describe these corrections, in particular, in the case where the compressibilities of the inner material, that is, what you are levitating -- call that β_i --is comparable to the compressibility of the outer fluid -- β_0.

A way to think about this average force is diagrammed here, where I show at the pressure as a function of vertical position for two different extreme times, one where the pressure is positive relative to the hydrostatic pressure. In that case, for the indicated problem, you have compression of the fluid mass, so it has a smaller volume at that time. So if you consider the instantaneous value of this $-\langle V(t)\nabla p\rangle$, you would have a relatively short arrow here.

On the other hand, 180° later, you have negative pressure and you have a larger volume, whereas the pressure gradient is the same and, as a result, you have a larger force

for this half the cycle, and if you time average that, you end up with a net attraction to the pressure antinode.

As a consequence, if you have, say, a low compressibility drop in water, you can levitate this. It will reach its equilibrium position and typically be then levitated above the pressure antinode, as shown here. Or if you have a small bubble in water, likewise, that is, one that is smaller than resonant size.

Throughout the discussion we are going to limit ourselves to a linear description of the standing wave. In fact, there are certain cases where there are very important corrections to that, and I will mention those at least briefly

-- Taylor Wang has done some work on that.

(Transparency 7)

The system of a bubble that is smaller than resonant size is the one that is used in sonoluminescence. This is an example of such a signature -- the short pulse probably not so short in our apparatus, it is only a few nanoseconds wide. But what might be remarkable is that same levitator was able to suspend a very large bubble -- shown here -- that is several millimeters in diameter. I would like to explain now how that happens.

(Transparency 8)

This is a diagram of the levitator (we will talk some more about the technology a little bit later). There is a piezoelectric cylinder, which is radially poled, and because of the Poisson coupling there is a mode where this becomes longer and shorter and then you can set up the standing wave.

(Transparency 9)

What you need to appreciate is the problem of the natural frequency of the bubble and what this does to the levitation condition. This is the monopole resonance frequency of a bubble, and we will be interested in fairly large bubbles, where the surface tension correction to the frequency is weak and the corresponding frequency is given here in terms of the hydrostatic pressure, P_0, and the density of the outer fluid, ρ_0, and the diameter, D.

You can derive this expression, for example, by writing down the Lagrangian terms of the kinetic energy of the fluid for radial motion and the potential energy associated with compression of the gas (something I usually assign as a homework problem).

For the previously indicated levitator, 22 kHz, this natural frequency is about 300 μ. So that means, for bubbles larger than 300 μ, there is a phase reversal of the relationship between pressure and the response of the bubble. I should also mention that there is an important damping of this mode. This is from a rather dated treatment, but I chose it because it is very simple, from Clay and Medwin's text.

One can divide the principal damping mechanisms into viscous, radiation damping, and thermal -- the capital R here means it is evaluated at the resonant condition and in this notation, then, Q is one over this delta_R. For some of the larger sized bubbles of interest these Q's might range between, say, 10 to 20.

(Transparency 10)

We have already mentioned this, that you have this phase reversal in the response. That means we have to change our previous diagram. So when the pressure is positive, now the bubble is larger than when the pressure is negative. So that says there is a reversal in the sense of these arrows and the larger arrow now becomes attracted to, that is pointed to, the pressure node.

In order to reach equilibrium, actually such a bubble would be levitated above the pressure node and have a net force downward that balances out the buoyancy of the bubble.

There is a problem, however, and that is because of buoyancy the bubble will be levitated slightly above the pressure node, where typically the magnitude of P decreases away from the vertical axis. The same argument then suggests that the bubble would tend to be pushed horizontally off the vertical axis and have an instability. There is a way around that and I will mention that later.

(Transparency 11)

The result of the analysis is the bubble tends to be levitated close to and slightly above the pressure node, and this is a scan through the vertical axis of the levitator described in Asaki's paper.

(Transparency 12)

In addition to the radiation force, the bubble will be oblate, and that was readily evident from the previous picture. You can understand a principal component of that since the bubble is levitated near a velocity antinode and you have high-speed flow past the equator. It turns out the quantitative part of this is a little tricky because of the

relative velocities of the bubble and the surrounding fluid -- actually, at the acoustic frequency the bubble is moving relative to the surrounding fluid, and that is addressed in one of our previous publications. But the basic issue is that you end up with a low-pressure equator.

This is a basic diagram in the upper part of the overhead. This transducer is typically three inches in diameter.

(Transparency 13)

What about the horizontal stability problem? It turns out that with this levitator, when it was discovered we could do this, it certainly was not obvious to me why it was stable. So Asaki took a number of transverse scans found and the coupling into this system is such that the pressure varies horizontally, it actually increases as you move off the axis horizontally, and that is the right sign to give you stability. In fact, you can enhance that effect, as discussed in the publication you have been given, by attaching the ring to the levitator. So, actually, as you move away here, the pressure ends up going up, and that is just what you need to give you horizontal stability.

In fact, if you scan all the way out to the walls, you find that the pressure goes down near the walls. I would like to comment on that.

(Transparency 14)

It turns out you really need to do a full analysis of the coupled system of the elastic rectangular chamber and the water. However, there is one limiting case that is fairly simple to analyze, and it is actually based on some work that Kargl did on scattering for a study of longitudinal resonances in a shell.

The idea is what does the reflection coefficient look like from a plate at very low frequencies? If you have water here and you consider the case of a vacuum-backed plate, the initial reflection is positive, because this is a higher impedance material. But you have this infinite series of reflections. This must be unimodular if you get all the energy back.

But when you add on this infinite series and take the low-frequency limit, it turns out you can sum this analytically, you end up with minus one, which looks like a free surface.

So that is at least one aspect of the issue as to why, when you scan horizontally, you end up with a pressure decrease near the edge of the levitator.

(Transparency 15)

What about the problem of the radiation force on large bubbles? There has been a lot of work done on radiation force on small bubbles, smaller than resonant size. Larry Crum and some of his colleagues have made a number of careful measurements of this. For reasons that may become evident later on in our presentation, we wanted to verify (that is, Tom Asaki wanted to verify) whether or not the radiation force was just a smooth function of size for a shrinking bubble.

What he did is set up the following experiment. He took two CCD cameras, one which was some distance away that has a wide field of view and he could accurately measure where the bubble was located in the chamber. Then he had another CCD camera where he could focus on the bubble and get an accurate measure of the size of the bubble. Then he also did hydrophone scans through the sound field, as indicated here, and this is the pressure profile relative to the pressure node. That is where the bubble would be attracted to.

Essentially what he did, then, was inject the bubble at a location above the pressure node and, as the bubble dissolves, keeping the amplitude of the drive constant, the bubble would seek its equilibrium position and migrate as it dissolves, so you get a relationship for a given bubble between position and bubble radius.

There is enough information in the quantitative radiation pressure equations, and you can see these discussed

-- a classic one is Eller in 1968, Crum and Prosperetti in 1983. There is also a recent paper by Taylor Wang and Lee, which, it turns out, contains the result that is in the Crum and Prosperetti paper. This is the resulting comparison and the measurements.

The calculation is based on a spherical bubble. As I previously mentioned, particularly for a large bubble, it becomes nonspherical, and that is this region, and maybe that is part of the reason for the deviation for large bubbles.

(Transparency 16)

We have already made reference to the deformation of bubbles. I would like to now look at the levitation of drops in air. If you take a drop in air and place it near a velocity antinode, you are going to have high-speed flow past the equator; it is much denser than the surrounding air, so you have stagnation points at the poles and it becomes oblate.

It would be at the velocity antinode in zero gravity, but that is a little expensive, to go to zero gravity,

so what normally happens is this situation, where you have an equilibrium between the weight of the drop and the radiation force and the drop hangs below the velocity antinode. Again, there are equations by King and others in the literature for the radiation force, but one way to think about this is that, look, this top part of the drop is closer to the velocity antinode than the bottom and there is a higher velocity there and an associated lower pressure on the surface.

We will talk some more about the technology of such levitators a little bit later.

(Transparency 17)

To give a detailed analysis of this -- that simple argument works only in the long wavelength limit -- to give a detailed analysis of this, some years ago I considered the radiation stress equations, how the radiation stress was distributed over the surface, and you can write, in some sense, a radiation traction involving a radiation stress tensor, and there is a simple discussion of this in a recent paper by Lee and Wang.

The net result is, by considering this net radiation force on the drop, you can do two things. You can integrate that over the surface of the drop and find the net force projected along an axis and find the radiation force equation (actually, you can recover the results of Yoshoko and Kalaseem that way). The other thing you can do is, from the stress distribution, calculate the equilibrium shape in a weak perturbation approximation. The point is, you are starting with a spherical calculation, you are finding the radiation stress distribution for a sphere, and that was discussed in Marston et al. (JASA [1981]).

The interesting result is that when you write down the aspect ratio here, this is D/H for the drop, as a function of the acoustic pressure amplitude of a standing wave -- a is the drop radius, and B is this coefficient that involves compressibility of the outer medium and the surface tension -- in the long wavelength limit this $F \rightarrow 0$ and this 1 here is what you get from this very simple incompressible potential flow calculation. All the work ends up getting this frequency-dependent correction, where $F = ka$. There is a very similar expression for bubbles.

(Transparency 18)

This works when it should and it fails when it should. Eugene Trinh tested this expression for weak deformations, that is, D/H, sufficiently close to unity. He ended up with adequate agreement with the theory that I just gave you, though I should say he left out the correction in the

theory for allowing a drop to move off of the velocity antinode.

For large deformation this breaks down, as indicated here, because the calculation is simplified for two reasons. One, the calculation is done for scattering by sphere, and that is not the case. This is basically a partial wave series calculation for calculating the full pressures in and outside the drop.

Secondly, the equilibrium condition was based on a weak deformation from a spherical shape.

(Transparency 19)

If you want to go to high frequencies, you also have to add in more terms of the partial wave series. This was done in a nice paper by Jackson and Barmatz, but if we were going to go through the equations we would be here all afternoon. This is a comparison between this long wavelength approximation and normally you do not run levitators up here in this high-frequency condition. But, again, they found that drops become oblate in air and bubbles become oblate in water, which is what we have seen, large bubbles become oblate in water.

(Transparency 20)

Recently, in Apfel's group, they have improved upon this calculation by describing the scattering and calculating the scattering from the spheroid and then doing a self-consistent calculation so the basic scattering is done for whatever the shape that is at the equilibrium condition.

Sure enough, for weak-deformation-aspect ratio this goes over to the old result. But in addition, as previously mentioned, there is going to be an offset from the velocity antinode, depending on the magnitude of the pressure. This is an experimental aspect. On the left of the middle figure is for high pressure and then, as you decrease in intensity, you lower pressure, you move off the antinode here at H = 0, and the drop becomes more spherical. But you have to actually take that shift of position into account, which they have done in their theory, comparison of the theory. It is an improvement over the previous approximations.

These improvements become important if you want to properly interpret experiments first on the earth in order to go to low gravity, as they have done.

(Transparency 21)

For the bubble case, Asaki has recently looked at this. This is the same kind of data discussed before for the equilibrium position, but here is the measured aspect ratio as

a function of bubble radius for different bubbles in a given field and comparison with the kind of approximate theory previously discussed, namely, the theory that is based on a spherical bubble and the calculation (not the Apfel, et al. theory).

(Transparency 22)

Just a couple more comments on levitators. For objects in air it has become customary now to use acoustic horns. This is actually a design that we borrowed from Apfel's group. It is also related to one that Eugene Trinh has been using. The way this is set up is, you have PZT disks that are longitudinally poled along the axis, so they run their thickness mode.

You support this by some thin metal plate that is relatively lossless; there is not much mechanical energy carried away from this. And then you have a resonant stub down below (in our case, aluminum), and if you figure out the signs you get a reflection coefficient here of minus one. So the appropriate thing is to take this, $\frac{\lambda Bar}{4}$ that is to say, the bar wave that you get by considering Young's modulus for the rod.

You are at a situation where the wavelength of the mechanical wave in this, considering this axis to be a rod, to be enormous compared to the size of the rod. That may not be evident from the picture. So that means you have a velocity antinode here, velocity node, velocity antinode here, velocity node, and velocity antinode up here, but you have this horn that ups the velocity at this tip and then you drive the standing wave in this chamber in order to levitate the drop.

A related design is discussed in Trinh's article, where here what he has done is to add a plate in some sense to better impedance match to air.

For a compressible drop in water, this diagram at the bottom is probably one of the simplest designs, though I do not necessarily recommend it, but if you just want to go home and make one of these things without having to go to the shop, you just basically take the 2-inch size microscope plates and cement them together and put a styrofoam piece in the bottom, coat a PZT disk, stick it in the water, and tune the system, and you can levitate oil drops in water.

(Transparency 23)

There are other ways of levitating and trapping. One of the things that we explored a number of years ago is the optical trapping of bubbles in water. This was done with a "doughnut" mode laser beam. The reason you had to use a

doughnut mode is for lateral stability. Basically, you focus this doughnut mode beam down and you have radiation pressure of light on the bubble because of the large reflection when you go beyond the critical angle of incidence.

With the doughnut mode, if the bubble moves off the axis, there is a restoring force. Similarly, for high-refractive-index objects, say, drops, in air, Art Ashkin at Bell Labs did a whole series of experiments discussing that case. We have also worked out the theory of what happens if you make this light circularly polarized. It turns out, if you have absorption, you can put torques on an object.

(Transparency 24)

But this is the realization of that levitator -- this is some work done by Bruce Unger -- where there is a bubble optically trapped here and this is acoustically quiet. You can see sort of a scattering pattern due to the total reflection region off the bubble. Unfortunately, because the radiation pressure varies as the intensity divided by velocity of the wave, and the velocity of light is enormous compared with that of sound in water, this means that the required intensities are very large, so you typically need a few-watt laser to do this.

That is really the introduction to levitation. Do you have any questions on this block of material?

Q: You talked about small bubbles being trapped at the pressure antinodes. Can you get the same thing with drops, small drops in water, being at pressure antinodes as well?

A: (Note added: The upper figure of Transparancy 16 explains why drops in air are ordinarily attracted to pressure nodes.) If you take a highly compressible oil drop, put it in water, that tends to go to pressure antinodes, correct.

Q: But you do not see the same droplets in air?

A: You have to consider the resonance responses. Normally that would not happen, because the resonances for the droplet are extremely high frequency. In other words, the thing about a bubble is that the monopole resonance is at a very low frequency compared with its characteristic size. That is not the case for the monopole resonance of a drop.

Q: I did not quite understand completely as far as horizontal instability. The reason that the levitation point is off the node or antinode is because of the body force in the vertical direction, right? But why should it be horizontally unstable?

A: Because if the pressure decreases as you move away from the axis --

Q: Oh, I'm sorry, so it is not a plane-wave resonator?

A: No, it cannot be fully, I mean, we have finite dimensions to our system, so you really have a variation of pressure with horizontal position. Many of the typical levitators that you try, if you do pressure scans through, as you move away from the axis -- say you use a spherical resonator, something like this, as you move away from the axis, the pressure is going to fall off.

If you try to levitate a large bubble in that case, it gets kicked sideways.

Q: I understand. I thought this was a plane-wave construction.

A: It is not truly a plane wave. That is a good point.

Q: If it were a plane wave --

A: It would have no restoring force, so it would in some sense be marginally stable horizontally.

Q: You were mentioning the length below PZT disks (in transparancy 22). Have you taken into account the difference between the PZT and the rest of the rod?

A: Yes. Actually, that is a good point. If you really want to calculate it, you need to put in the fact that there is a little different mechanical property with the PZT, but if you just want to estimate how long do I cut my rod off to make such a system, I would use this first guess. Thank you.

Q: Why is it quiet in transparancy 24?

A: This does not make any noise. As far as I know, the only way to levitate a bubble where you can actually go in there and do something with it without making any noise, the light is levitating the bubble.

Q: I was not understanding why the smaller bubbles are going to be attracted to the pressure antinodes and the larger bubbles are going to be attracted to the nodes.

A: If we go back and reason through the phase relationship, the phase relationship between the response of the bubble and the sound field will vary, depending upon the size of the bubble relative to the resonance. So when the

bubble is larger than resonant size, when the pressure is large, the bubble is large. When the bubble is smaller than resonant size, when the pressure is large, the bubble is small. That is the key to understanding the sign.

Q: You mentioned the long wavelength limit. How good is it in these ultrasonic levitators? Can you give me an idea of the a-k parameter?

A: Most of the k-a parameters we use are well below a half. It depends upon which property you are calculating. You can find discussions of it in the literature for the radiation force. Transparancy 19 has to do with that part of the radiation force that deforms the drop or bubble and, in this case, you can see, when k-a gets larger than about a half, you start to deviate from the long wavelength limit.

Q: Can you give us ballpark numbers for the acoustic pressure required to levitate a droplet in air?

A: The best way would be to put up Eugene's figure here.

(Transparency 18)

I think there are RMS pressures on here, so you can read this scale, RMS pressure x 10^{-4} $(dynes/cm^2)$. You can stare at that later.

Q: The viewgraph you just had up there with the bubble shapes, that is just a comparison of the oblate shape under levitation as opposed to the natural shape of a bubble, say, when it is rising to the surface?

A: That is a good point, and we will get to this later. Bubbles naturally become deformed and drops also become deformed when they fall for the same reasons, you have high-speed flow past the equator. So in some sense you can use levitation to simulate what nature does.

I would like to go on to the next part, having to do with shape oscillations.

(Transparency 25)

These are the lower modes, shape oscillation modes. This is Lamb's equation for the frequency based on an inviscid calculation; that is, you neglect viscosity. You can find this equation, a later derivation, summarized in Landau and Lifshitz' book ("Fluid Mechanics"), where you can write down the Lagrangian by considering the kinetic energy associated with spherical harmonic deformations and the potential energy associated also with such deformations.

Rayleigh considered the case of a drop in air, where he neglected the outer fluid. Again, this contains the density of the outer fluid, ρ_0, and ρ_i for the inner fluid and surface tension.

The point of putting up this equation is that these frequencies are typically fairly low, around 100 Hz for a millimeter-sized object, because surface tension is a weak restoring force compared with forces involving compressibility.

Q: Are you looking down from the top?

A: This is for a particular choice of modes where it is symmetric around the vertical axis. You could take other azimuthal symmetries and for deformation from a spherical shape you have a degeneracy in the frequency for those other shapes.

Q: So we are looking in from the side?

A: You are looking in from the side, that is correct.

(Transparency 26)

A number of years ago I had the good fortune to take a postdoctoral with Bob Apfel and he said, basically, find some way to drive these shape oscillations. What is the problem? The problem is that these frequencies are so low that if you put in sound at those frequencies, 100 Hz, the wavelength at that frequency is going to be enormous compared with the drop, and you have very weak coupling and weak deformation.

On the other hand, if you want to choose something that strongly interacts with the drop, you want to have something where the wavelength is comparable to the size of the drop for the sound field. But that is at the wrong frequency to drive the resonance.

(Transparency 27)

So what suggests itself is to use a nonlinear coupling mechanism, modulate the sound to drive the shape oscillations. The way that this came about is, I think, what people should do when they have ignorance in a field; they just go read about related fields and see what they can come up with. At least for me, it was seeing these diagrams. This is from actually an old paper by Hertz and Mende.

What they did was take an ultrasonic -- in those days, quartz -- transducer, put it in a fluid, say, carbon tetrachloride, and there is water above, and they turned on high-frequency sound, probably megahertz frequency, and they

saw a deformation due to the radiation stresses.

In fact, there are arguments about the details of those equations. Some of them are reviewed in a nice article by Apfel and B.T. Chu in JASA (1982). But let us just look at the observations. What else did they find? They found, for a different choice of liquids, water and aniline, sound, once again, still going up, the liquid goes down. So the deformation of the liquid is not related to the intrinsic direction of propagation.

In fact, you can see that even better in this lower experiment, where they took an acoustic mirror here, sound down, sound up, and here is what they saw in the two cases. In one case the deformation is down in both cases, beam coming down, beam going up, and here the deformation is up for a beam going down, beam going up.

If that is the case, then why not do the following experiment? Take, say, an ultrasonic beam and shoot it through the drop and if, as asserted, in this case, the carbon tetrachloride is attracted to the region of the intense beam, then you should have such a deformation. If you turn the beam on and off, then you can drive such shape oscillations. So that is basically the origin of these experiments.

(Transparency 28)

There are a number of ways you can design them. The initial experiments involved a separate levitation standing wave and the standing wave that is used to drive the shape oscillations, and that actually gives you an extra degree of freedom, because you can control where the drop is relative to the standing wave for the high-frequency carrier that is used to drive the shape oscillations.

In fact, what we used in this is what you ham radio operators will recognize as a double-side-band-suppressed carrier. You basically take this carrier signal, multiply it by the modulation -- this is what you get. Then if you take the time average of the square of this, you end up with a constant plus a low-frequency modulation, and that you can use to drive the shape oscillations.

In fact, you can also consider the low frequencies to be the nonlinear response to the side bands in this signal that arise because of the nonlinearities in the full hydrodynamic equations.

(Transparency 29)

If you do this and you tune the modulation frequency, you should be able to sweep through the shape oscillation resonances of a drop, and this is what we did for

the quadrupole resonances of it -- I think this is a drop of benzene in water. If you calculate what the response is in terms of the modulation frequency, you would have this typical resonance response peak, where you have a quasi-static deformation here, and then there are oscillations.

It turns out benzene is slightly soluble in water, so you can inject a drop of benzene and watch it dissolve and then monitor, track, the frequency as a function of size. As size goes down, the frequency goes up. Compared there, also, is a Lamb's equation and then a first viscous correction. It turns out there are some subtleties in this system, which we will see in the next section, when you want to include the effects of the viscous boundary layer.

The way that these experiments were done, this is relatively crude technology, were at very low-amplitude oscillations, sub-micron. How could one detect those? You can use some optical tricks.

(Transparency 30)

It turns out, if you pass light through the axis of such a drop and look near the "rainbow" angle, you see this kind of pattern here. We will get to discussing the rainbow conditions way at the end of this lecture, I hope. You have essentially an Airy maximum. But what is important for this first part of the experiment is these very fine fringes you see here, and those are associated reflection here, and this transmitted ray, which dominate the scattering.

Then if you look at infinity, that is, you look with a telescope focused to infinity, you see this pattern. If you undergo oscillations, then the optical path difference between these paths as a function of time, and your eye cannot track the motion of these fringes, so the fringes blur when you are on resonance and they do not blur when you are not on resonance. So it is a very simple, small-amplitude oscillation detector.

To make this more quantitative, if you stare at this diagram for a while for this particular orientation of the axis of oscillation, you realize that, as a function of shape, the rainbow angle will shift, so this whole pattern will go back and forth. What we did was took a photometer and put it at a given position here in the light scattering and monitored the optical intensity at that angle.

(Transparency 31)

This is the apparatus. With this rotatable mirror you could adjust it so you are right near the rainbow angle. Then, by using synchronous detection, you can monitor the amplitude of oscillations, and there are some ways of

calibrating this.

(Transparency 32)

As you sweep through a resonance, you see resonance curves. This is for the in-phase part of the response and this is for the quadrature part or the amplitude here, and the scale has been cut down by two so this did not swamp the upper one.

These are the old days when you had to use a storage scope.

Q: Question. You were saying that when you shine the light there and it is going through shape oscillations, your eyes begin to blur. Are you currently using laser beams?

A: Yes.

Q: I was wondering if you might try a deflector so that you would phase the laser light to come onto the drive when it was at a certain phase and then reflect --

A: You could do that. But we actually want to see the blur and we want to see the shift in the Airy pattern, because that had the information content in it.

This was when we were looking at relatively weak oscillations. Even with the early apparatus we discovered that you could drive larger oscillations and you could break up drops, and then we went back to this and took high-speed pictures.

(Transparency 33)

This is the breakup phase for a drop of silicon oil (it is dyed) in water undergoing large-amplitude oscillations.

(Transparency 34)

How about drops in air? This is some work that I did when I was on a sabbatical at JPL in 1984. In this case, it is a little bit simpler technology. We used basically the kind of levitator, though not necessarily very well-drawn, shown here, and now have a very weak modulation just of the levitation signal and there is modulation of these radiation stresses that cause the drop to become oblate. You can couple into the modes, the modes are very high Q, much higher Q than for the drop of oil in water, and thereby scan through the resonances.

I would like to say something, though, about how the resonances are detected, because it is a useful experimental method. We will return to it later.

(Transparency 35)

It is called pseudo-extinction. Basically, what you do is you illuminate the drop with a horizontal propagating laser beam (this is a view from above). This looks like the profile of the oscillating drop. It turns out, if you work out the area of that profile, that will be a function of time due to the shape oscillations, and that is what these equations say.

If you can measure that cross-sectional area, you can get a handle on the shape oscillations in a way that is space invariant; that is to say, you can translate this drop around but you can still do the measurement, and that, unfortunately, is a complication in certain experiments.

The way that we did this is to take an expanded beam of the laser, focus that down, put a hole in the focal plane of this lens, but you choose the hole to be much larger than the forward diffraction peak of the drop. That is because if you make it small, so you do a true extinction measurement, then you are also sensitive to other problems, any thermal gradients in the system. And you can even detect the acoustic sound wave that way, you can use that to tune up a levitator.

Q: Excuse me, would you explain what an extinction measurement is?

A: A true extinction measurement is if you have a light beam and you shrink this pinhole down to a very small size, so it is basically equal to the diffraction-limited spot size for that light beam, and then you monitor the power that gets through that hole, if there is extinction of light, light is removed from that beam and you measure the change of that power.

What we are doing is not a true extinction measurement, it is a pseudo-extinction measurement. We are losing the light not because of the forward diffraction peak (of the drop) but because of the light refracted and reflected from the drop, and that goes to larger scattering angles. We will return to this later when we talk about pseudo-extinction for bubbles. But the net result is you can then use synchronous detection and you get tuning curves.

(Transparency 36)

Because I was at JPL for only something like three months or four months, we did not really get the system working any better than what is shown here. The point is, you can have very sharp resonance peaks. An interesting complication is a satellite peak -- shown here. This is a case where a drop is evaporating, so as it evaporates, the frequency shifts up. What we did is basically scanned back and forth through the

resonance region and monitored the output of the lock-in amplifier.

It turns out this satellite peak is undoubtedly related to the equilibrium oblateness of the drop, so that breaks the symmetry of the different modes, and there is some weak coupling into a non-axisymmetric mode.

(Transparency 37)

This is a demonstration of this class of coupling. In some cases you can see there are specific drops, like drop 7. We monitored this as it evaporated over a significant range and then tracked the frequency.

So that is the end of the section on drop dynamics. Are there any questions on this?

Q: Pardon me for being a stick in the mud, but would you motivate for me why one is interested in exciting these modes? I mean, what do you intend to do with them?

A: That is actually a very good question and it has to do with how people got this funded. Apfel's goal was to measure the interfacial tension between superheated liquids and water. At that time, there was significant interest in liquified natural gas, and one of the problems is, if you spill it, there could be some disastrous consequences, particularly in the Boston Harbor.

So he got funds from the Petroleum Research Fund to explore this. You could not measure the interfacial tension by conventional methods. Since then a number of people have pursued this -- not myself -- for other reasons, including Apfel's group, for measuring interfacial dynamics of contaminated systems, say, a drop with a surfactant on it (we are going to see a little bit later, when we get to bubbles, that surfactants can affect the dynamics), and those kinds of dynamics are very hard to measure by some conventional methods, because you always have some other boundary present.

Apfel, for example, they have arranged to have experiments done on both USML-1 and the next relevant Space Shuttle flight, USML-2, for these kinds of quantitative measurements. Thank you.

(Transparency 38)

I would like to go on to bubbles now. I really need to acknowledge Tom Asaki at the outset and then, also, the work with Eugene Trinh. It is perhaps not entirely fair that I talk about all of this, because Tom, who was at the previous summer school, is still working on his thesis project, but some results I am able to show you.

If you calculate the shape oscillations of an air bubble in water, this shows the frequency versus size. At characteristic low-frequency modes you can see 100 Hz up here and 10 Hz down there. Eventually we are going to be interested in some applications or problems in low gravity, where we would like to slow the whole process down, and then we want to be at the larger size bubble, which will be basically below 10 Hz.

Most of Asaki's experiments are really done up here, where your frequencies are up around 100 Hz or above.

(Transparency 39)

Suppose you have Newtonian fluids, that is, Newtonian water and Newtonian air, and an ideal interface. We will talk later about what that means. But you can calculate the Q of the bubble and, say, the lifetime of such an oscillation. We see that the Q of the bubble is in that range.

(Transparency 40)

We can anticipate how to do this kind of experiment based on what we saw previously. If you trap a bubble close to the velocity antinode, modulate the radiation pressure as a function of time, or modulate the carrier -- actually, we use a very small modulation percentage -- then you can drive the shape oscillations. That way you can see, also, the higher modes, most of our work has been done on the quadrupole mode.

(Transparency 41)

There are several ways to detect this. The simplest way is to just look in here and estimate how large an amplitude oscillation we have. Most of these data were taken in that way. These are the same data that are in the publication you were given, except there is a shift in the comparison with theory because of calculation of a certain offset that I will show you in a minute.

The key results are identical, namely, that when you get to larger bubbles, there is a shift off of Lamb's equation. Some of the other data on here, namely, those with the triangles, was taken by a light-scattering method I will show you in a minute that was done at very low-amplitude oscillations, so you can see whether the amplitude of oscillation was an issue.

In addition, we see data from USML-1 in these crosses, where there is good agreement both for the n equal to two or quadrupole mode, and n equal to three mode, even though the bubble is very large. That is because there is not a large background static deformation of the bubble, whereas in this case there is significant background just to be able to

levitate the bubble when you calculate the equilibrium position. We will get to the USML-1 data later on.

(Transparency 42)

We will also talk about light scattering a little more later, because it is a useful experimental method, but let me just comment on the method Asaki used in his measurements. That is, if you consider the rays through a bubble and go into the region where you have a reflected ray and a transmitted ray, these are the dominant contributors in this scattering angle, more forward in the critical angle, which, again, we will return to, but these are displayed in these Mie scattering calculations, which is the exact partial wave series for a spherical object.

But it turns out, if you were to change the radius of this bubble, the optical path length between here and here shifts, and that means this coarse structure shifts, so what Asaki did is focus near the bubble, but in the region where these rays overlap you can see an interference pattern, and then if you get a little oscillation, very weak-amplitude oscillation, this interference pattern would blur.

He did this for light through the plane of the equator of the bubble, so this diagram is applicable there.

(Transparency 43)

I mentioned there is a kind of first-order correction you can do for shape, and that is to realize that if you do a linearized theory for the second-order response, that is, that flow that is induced by the modulated radiation pressure, you really have a sum of a static deformation -- shown here -- and an oscillating deformation -- shown here. R_E is the equilibrium radius for the associated sphere.

What you measure is an equilibrium-aspect ratio D/H and then a resonance condition. When you work backwards to can correct what R_E is, which goes into the Lamb's equation, you find out to order Δ, where Δ is this measure of oblateness. R_E is the radius of the spherical bubble having the same volume as the static oblate bubble, which is the correction indicated.

(Transparency 44)

For very large bubbles, Asaki discussed this before, we had these instabilities (some people think it is a Faraday instability on the surface of a very large bubble), this is without any modulation. It is associated with an instability from the accelerations of the acoustic field that levitates this bubble and is the topic of current interest.

(Transparency 45)

There is another complication, as most people who have worked with levitation will tell you, and that is, what you are studying can have its properties vary as a function of time. That is true here. If you take a bubble and levitate it, and these are done for roughly the same amount of gas dissolved into the water, the bubble will dissolve and, with no modulation, this variation of diameter with time is along that curve.

If you drive shape oscillations with different modulation percentages but always keep the frequency of the modulation to track the frequency of the shape oscillations, so you are always driving resonance, what you find is, when you up the modulation, you have a more rapid dissolution, and that is because as the bubble oscillates there is going to be oscillating flow past the bubble and thereby you can enhance the rate at which gas goes into solution.

(Transparency 46)

In order to make any of these ideas more quantitative, we need to consider a couple of boundary layers that show up in the problem. One is the acoustic, our first-order boundary layer, where you can estimate the thickness from this expression, where ν is the kinematic viscosity for water and f is the acoustic frequency. You end up with a boundary layer thickness that is typically very much smaller than the size of the bubble.

Because there is negligible change in normal pressure across that boundary layer, it turns out to be a pretty good approximation to use in inviscid theory for the scattering of a sound wave by that bubble. However, if you consider the low-frequency oscillations, that is, the second-order flow induced by the radiation pressure, then you can have boundary layer thicknesses whose frequencies are much lower and boundary layer thicknesses become comparable to the size of the drop or bubble, in this case, and that ratio is an important parameter in the theory, which I will not attempt to write down all the equations for.

(Transparency 47)

You can see one reason why it is an important parameter by considering Lamb's equation. In particular, we are going to see that when you put in viscous effects there is an oscillating vorticity near the surface of the drop or bubble. Lamb solved this problem for inviscid flow, that is, potential flow.

If I look at the quadrupole mode and I take a cut

here at 45°, and I look at the θ component of velocity, μ_θ, which is the tangential component, as a function of radius, from Lamb's equation you get a discontinuity at the surface, and that cannot happen for a real system, because you do not have this slip condition at the surface. You really cannot have slip. So the associated flow is as diagrammed in Lamb's equation away from the interface, but you really generate vorticity at this interface.

It turns out the nature of that vorticity will be affected by the nature of the interface. If you have a clean interface, that is to say, you have Newtonian fluids and you put in viscosity, you have one kind of viscous boundary layer.

If you have a dirty interface where you, say, put in elasticity or interfacial viscosity, then you have a different kind of boundary layer and it strongly affects the damping of the shape oscillations, and that is why Asaki is doing this work.

It is actually funded by the Navy ARI on dynamics of bubbly flows, where they would like to understand the evolution of certain bubble properties as a function of time. Actually, he is doing some of this in sea water.

(Transparency 48)

Again, I do not want to write down all the equations, but I should write down at least what the boundary

conditions are, again. In order to model the effect of viscosity and what these boundary conditions do to the oscillation modes, let us look at the conditions. For a "clean" interface you would have the radial for small-amplitude oscillations, the radial components become equal and the tangential components become equal at the interface.

Then you also have a radial stress balance. There are these viscous terms and then there is a curvature term. In fact, the p^{-r} here is the radiation stress, which is distributed over the surface of the drop or bubble, and this is the equation that you solve to find, say, the equilibrium shape. Or if you look at the time modulation of this radiation stress, that is what drives the second-order flow.

However, there is also a tangential stress balance that contributes to the dissipation. The simplest way to model this is to neglect the tangential radiation stress, and that is a pretty good approximation in certain frequency ranges. I discussed this in this Drop and Bubble Symposium article from 1981 about some conditions where this tangential radiation stress assumption will break down, but for the conditions we are interested in we will make that assumption.

Then you end up with a condition, as indicated here, where $e_{\tau r}$ is the appropriate rate of strain tensor in the inner or outer fluid. Then, when you solve these equations, you end up with the equations for the damped modes. For the case of free oscillations, freely decaying oscillations, the leading order terms were given by Miller and Scriven, and then Prosperetti and I discussed a correction to that.

There are some modifications with surfactants and this is really an ongoing problem, but Lu and Apfel discuss that.

(Transparency 49)

What is the nature of the equations qualitatively? If you look at the characteristic equation for the free modes, you can write down the oscillation amplitude is $e^{i\Omega t}$, where there is a real part, which is the free-decay frequency and a damping, which is $i(\tau^{-1})$.

The characteristic equation for Ω looks a little bit different than you might be used to from elementary physics. If you take away this term, this boundary layer term, then it looks like a freshman physics equation; that is to say, there is the mass-related part here, the inertia part from, actually, the outer fluid in the bubble case, there is the spring part, which actually shows up as the square of the Lamb frequency in this normalization, and then there is the sort of ordinary viscous damping part, i.e., dash-pot damping.

However, because of this tangential boundary condition, you end up with a boundary layer damping, and that affects both the damping and the real part of the frequency, and that is because of the inertia of the boundary layer. There is a very nice discussion of this kind of damping in the fluid mechanics book by Landau and Lifshitz.

We simplify the notation if we replace the ω^{Lamb} by just ω. Then the free-decay frequency has this $\omega - \gamma^2$ part. That is the ordinary shift in frequency for a dash-pot-damped oscillator. But it also has a reduction in frequency due to the inertia of the boundary layer, so measuring that frequency is a probe of the boundary.

The damping, which I will write here as S, which is the reciprocal of this decay time, also has an important boundary layer damping term. The details of this coefficient α and, to some extent, γ, would be affected -- for example, if you do not use Newtonian fluids you get something else, generally larger damping (and we worked out such as the dirty

bubble limit).

(Transparency 50)

That was free decay. If you look at the forced oscillator equation, then the characteristic form is -- this gives you the deviation from the equilibrium shape. You have a sum of static terms, which I write in this way, and oscillating terms multiplied by spherical harmonic functions -- this is for the M = 0 case. The associated dynamical equation, this looks like the forced part of the harmonic oscillator equation, where the forcing term, the inhomogeneous term, involves the projection of the radiation stress onto a given mode. Again, you see the boundary layer term shows up.

(Transparency 51)

How can we make these kinds of measurements more quantitative than the way previously described? We go back to our pseudo-extinction measurement. In this case, you make a levitator with a quartz or glass bottom and drive this plate, drive this up and down and levitate a bubble. We are primarily interested in the quadrupole mode, though we looked at some higher modes.

What happens is, as you undergo oscillation, if you view this from above, the profile -- actually, it is the apparent spherical or circular radius as viewed from above -- will oscillate with time. We would like to be able to measure that, measure those oscillations.

We pass laser light along the axis from an expanded beam and one way to do the experiment is to monitor the oscillations with a lock-in amplifier.

(Transparency 52)

Let me return to the problem of the pseudo-extinction. It has long been known, for objects like a drop, that if you calculate the total scattering cross-section at short wavelengths, that is approximately twice the area of the object, the profile or cross-sectional area of the object. The way people in light scattering define this is they describe the scattering efficiency, which is the ratio of the scattering cross-section to this cross-sectional area, for a spherical object.

If you have a lossless object, that is just in the extinction, ratio of extinction cross-section to that area. You see, for the drop, which is the solid curve, that goes over to two at high frequencies, short wavelengths, and that is the so-called extinction paradox. We calculated this some years ago for bubbles, and that is the somewhat smoother curve, and that turns out to be because bubbles do not have strong optical

resonances, the resonances are strongly damped. But you also still go over to two.

(Transparency 53)

One way to understand that extinction paradox is to look at the features of the scattering diagram. This is a Mie scattering diagram for bubbles (and, again, we are going to get back to this later, but it is the solution of the partial wave series for scattering of light by a spherical bubble), plotting the logarithm of that intensity, and it turns out -- and it is also true for objects other than bubbles -- you end up with a forward diffraction peak, which, if you integrate through this peak, you end up with a scatter power roughly equal to the incident radius times the cross-sectional area of the object.

Then you also, if you integrated everything else, get something that is comparable to that number. Hence, the total scattering is twice the cross-sectional area. The point is that this forward diffraction peak is narrow. We designed the system to accept this forward diffraction peak and reject all of the rest of this, and that gives some stability to the measurement system.

(Transparency 54)

So we have something where the variation of power is proportional to the beam intensity times this variation of the pseudo-extinction cross-section and that, in turn, is just a variation of the cross-section area. You will get a copy of this transparency later, but that is the way to make all this quantitative.

(Transparency 55)

These were actually some very early results from Asaki's work, where he has scans through a resonance. He tries to do this fast enough that the bubble does not dissolve very much, and it is pretty hard to do. But you get a phase change of about 180°, which is what you expect for scanning through resonance. You get a resonance peak in the amplitude and then this is, I think, based on essentially a two- or three-parameter fit, where you adjust the parameters that go into the boundary layer-damped oscillator equation.

You end up, in this case, as sort of a dirty bubble; that is to say, the damping is large compared with the theoretical value (for a clean bubble). It turns out to be much better for the kinds of exploration he is doing to look at free-decay signals (the lower curve), and that is actually a very early result. We will show you some more of this in a minute.

(Transparency 56)

The way this is done, it is the same apparatus now, but you capture the free decay on a digital oscilloscope, where you band-pass this in the frequency range of interest to get rid of any low-frequency fluctuations. You monitor the bubble size with a CCD camera.

(Transparency 57)

When he does this, he sees a number of interesting features. I am going to only introduce you to them. The bubble dissolves as a function of time, and this is the measurement of a radius over time, where he starts off with something that is about 1.5 mm and slowly dissolves.

What he does is, he adjusts the oscillator frequency so he tracks this natural frequency of the bubble and then he cuts off the drive and monitors the free decay, and then he fits the decay rate. Then he compares the experimental decay rate to the theoretical value for that particular size and monitors the ratio.

For reasons that we honestly do not understand at this point, but it is seen in a number of cases but not all cases, you can actually have a dramatic drop in the decay rate. It might start off large, then drop. It goes down to a theoretical value of near unity and then starts drifting back up. It would appear to drift up because surfactants or other contaminants in the water, despite his best effort, would collect on the bubble. It turns out that anybody who works with pure water systems knows this always happens.

So this part of the curve makes sense. This part we do not understand yet. You might say, well, this did not really happen. It really did happen, and you can see that here.

(Transparency 58)

This is the decay rate for the last trace right before the drop, the free decay, and this is for the very next record, and it decays more slowly. The decay rate really does fall off in many of the cases.

(Transparency 59)

In closing this section, perhaps surprising, this can be done in sea water. This is a record for a sample from Puget Sound water. It turns out there are a number of things that happen, but these are still under investigation. One always finds, for sea water, excess damping, but this excess damping can evolve in time, so you can inject a small bubble and watch this damping evolve.

That is basically the end of the section on bubble dynamics on the earth. The next section has to do with a very brief discussion of low-gravity experiments.

Q: May I ask a quick question, please? I have a question on your interface conditions there. You are defining a clean bubble as a Newtonian fluid, and you have a kinematic condition, and then you have the dynamic condition. The dynamic condition changes for the non-Newtonian case. You are modeling the surfactant as a non-Newtonian?

A: That is right, surfactant would be a non-Newtonian example.

Q: Okay. Well, how do you -- I missed -- all of a sudden you must have linearized something when you got some harmonic oscillator equation.

A: Yes, all of this is for infinitesimal amplitude oscillations about the equilibrium shape.

Q: So are you doing polar and -- I did not understand what your equation was for --

A: Actually, I did not write out the non-Newtonian equations. Conveniently, if you want to see them, this is not in the transparencies that will be handed out, but --

(Transparency 145)

-- following the method of Miller and Scriven you can write down the dirty bubble limit. It is either here or, actually, on the next one. Sorry.

(Transparency 146)

The ratio of these parameters in these two cases is about 130 for water. I refer you to Miller's and Scriven's paper for a discussion.

DR. MARSTON: Before I begin the next section, I have a couple of brief comments.

(Transparency 57)

This decrease in damping may not be unique to this experiment. Bob Apfel's student, Lu, in her thesis, did damping measurements of some kind of hydrocarbon drop in water and also found situations where there would be a sudden decrease in damping, and the reason why is not known.

(Transparency 15)

It is that drop in damping that is really the reason

why we did these kinds of measurements of, say, equilibrium, position versus radius, to see if there is some big change in the radiation pressure, and it is always a smooth curve. You never see a big jump.

(Transparency 60)

The next section is on low gravity. I guess I could introduce this really briefly if I had the slide projector on. Here the collaborator, Eugene Trinh, went to great heights to do this experiment which was launched during this summer school two years ago.

(Slide)

This is a levitated drop with a reflector above it.

(Slide)

This is one of Eugene's pictures from his platform.

(Slide)

This is operation of USML-1. You will notice there is a facility here -- that is a glove-box facility -- and that provided the electrical power to do the experiment that I will now describe.

(Transparency 60)

This is probably the lowest-budget experiment to go in space. NASA budgeted $20,000 a year for three years. Why would you want to investigate dynamics of bubbles to radiation pressure? It turns out control of bubbles in low gravity is important; you no longer have buoyancy. Bubbles can be present in melts, molten materials, and other liquid systems. In fact, there was a problem in USML-1. One of them was that there were bubbles that people did not want in everything. So you would like to have some way of controlling these, and radiation pressure turns out to be one of the ways.

The other issue has to do with basic bubble dynamics, which we have alluded to earlier, this sort of breakdown of Lamb's equation at very large sizes. In anticipating the results, we found, for the cases where we could make measurements, yes, Lamb's equation did work for the large sizes in the absence of the complications of buoyancy.

So our primary goals were to investigate the oscillatory dynamics of large centimeter diameter air bubbles; look for the agglomeration of bubbles in an ultrasonic standing wave in micrography as a possible way of controlling bubbles; investigate the coalescence of two bubbles -- that is, fluid

dynamics-type people are interested in what happens when two bubbles coalesce, it is a very difficult problem to discuss in detail, but one can do some approximations you would like to test; and, finally, to investigate surfactant effects.

I will not have time to discuss those but, in addition, we became interested in the response of a bubble in water coated with an oil layer on the inside to see whether we could find core centering, and that was partly because Tom Asaki found that you could produce coated bubbles in the laboratory with one of these levitators.

(Transparency 61)

This is the flight hardware. I should also give some credit here. John Depew at JPL actually built the hardware and Curtis Boswell developed the electronics box. The chamber, of course, had to be sealed and filled in advance. Unfortunately, it had to be filled many months in advance, which led to some complications. There is a driver down here, and this is the electronics that drive this levitator.

This (lower photograph) is the ground-support equipment, but in the flight these are plugged into the power supplies from the glove box. In this little black box there is all of the following that enables one to do this experiment.

(Transparency 62)

There is a dual-frequency drive system, dual frequency because we found that we could not guarantee stability with a single drive frequency, say, 22 kHz, in low-gravity simulations, that is, KC135 flights. The easiest other frequency to drive was the third harmonic, and so we have a third harmonic oscillator.

We can adjust the frequency of the carrier and can adjust the amplitude modulation. We can monitor the frequency of the drive with an LED display or we can switch this display over to monitoring the power -- actually, it is the current -- into this output amplifier, and that is the easiest way to tune up the system, because when you are driving the resonance you can drive more power into this chamber and thereby be sure you are on the right acoustic frequency.

(Transparency 63)

We have already shown you the data on shape oscillations, very low-amplitude oscillations of large bubbles, and particularly the ratio of the modes is in good agreement with Lamb theory and the actual frequencies are as good as the size measurements.

(Transparency 64)

What about the problem of agglomeration? Sometimes the term "Bjerknes force" is used when you talk about radiation pressure on bubbles. Larry Crum and others like to use this distinction between the primary Bjerknes force, which is just this radiation pressure from the incident wave, and the secondary Bjerknes force, which involves a mutual attraction or repulsion of a bubble in a wave.

The way to see this is actually this expression here. If you have two bubbles near each other, if bubble one scatters whatever the incident sound is and bubble two is subjected to the scattered sound from bubble one as well as the incident amplitude, and the dominant interaction term comes from the volume variations of bubble two due to the incident sound wave multiplied by the negative gradient of the pressure of the sound scattered from bubble one to bubble two, that is, that is the dominant interaction term, and that gives rise to the approximation of the secondary Bjerknes force, which you can show. It varies as $1/r^2$ at low frequencies and is attractive if the bubbles are both higher, that is, larger than resonant size, or both smaller than resonant size (in our experiment they are always larger than resonant size).

(Transparency 65)

And sure enough, you can get agglomeration due to this attraction. This is a Trinh's comment "Turning on the field gathers them up in about the same place". One of the paradoxes is we did not get very much coalescence. We do not know why, except after we injected a surfactant some time after we did get a few coalescence events.

(Transparency 66)

This illustrates one of them. This is right before coalescence. There is a small bubble here attracted to a large bubble, and this breaks into the large bubble.

(Transparency 67)

And you can see a wave that runs around the large bubble and then reappears back at this initial point of coalescence. It turns out, if you count the video frames, it is about a sixth of a second later.

(Transparency 68)

It is kind of remarkable, the volume of this bubble is very close to one that Longuet-Higgins had previously calculated for the case of a Gaussian bulge on a spherical bubble where he included only inertial effects (there are no viscous effects in this calculation), and the dynamics is very

complicated. You are driving many modes in this system, as you can see.

What happens is that after a certain time, and this is a scan through his analytical calculation, you have a reappearance of this bulge.

(Transparency 69)

That is, in fact, at the same time that we calculated for this experiment, that we found in the experiment. So certainly some aspects of the inertia-dominated case seem to be observable.

(Transparency 70)

The final aspect of this I would like to comment on is the problem of core centering. This was studied in Taylor Wang's group in the early 1980s for various fluid systems but not for the case where you have an air-bubble/liquid-layer/outer-liquid, which is the one we will look at. But for their systems you could get these kinds of modes where, during the course of the oscillation, you either have a situation where, during the oscillation, the liquid layer is relatively constant in thickness, or one where the phase of the relative oscillations is such that the liquid layer varies greatly in thickness during the oscillations.

If you calculate these (and, actually, Tom Asaki worked these out for this system), what is called the "slosh mode" here is at these fairly low frequencies, and we were interested in driving the slosh mode. We could not quite get our driver down to that frequency, because it was a little bit below the low-frequency limit of the modulation, but even though you could not quite get on the resonance frequency, Eugene Trinh found that you could see this core centering.

(Transparency 71, Right Hand Side)

So this is a case of air bubble, silicon oil, water on the outside, and you could see evidence and you could push this bubble away from the wall and be centered. One thing that I should add, it is clear, when you are doing the experiment, because you can look at it at various angles -- we have a video record of only one dimension because of the kinds of constraints you have in such experiments.

(Transparency 71, Left Hand Side)

However, there are other cases, such as this one, where the bubble is much smaller than the drop, and here there was an enormous amplitude oscillation for the drop driven and the bubble would not break away from the pole.

(Transparency 72)

These are Eugene's comments about this experiment and this may have some applications. Taylor Wang has done a lot of work on various core-centering problems, because he thinks it is important for certain fabrication problems.

(Transparency 73)

I would like to switch gears now to just a number of experimental methods on bubble measurements, and I am going to give you some ancient history. It is part of my thesis work. It is in liquid helium. I really will not spend much time on this, but I thought that something here might be of general interest.

One of the problems was to improve upon previous measurements of the tensile strength and get some idea of the bubble dynamics in superfluid liquid helium. The difficulty was that the early tensile strength measurements, using ultrasound, gave numbers on the order of millibars, where the predicted values of homogeneous nucleation theory were typically a few bars.

So this apparatus was designed, where, in order to calibrate the sound field, we used the diffraction of light, so the calibration of ultrasonic amplitude is based on essentially Raman-Nath diffraction of beam passed through the acoustic standing wave.

Then, by looking at the diffraction pattern in the focal plane or optical-transform plane of its lens, and making a number of measurements as a function of amplitude of the drive, you could get an estimate of the actual standing-wave amplitude in the system. It turns out the tensile strengths we measured were orders of magnitude larger than the previous values and closer to theoretical results. Now there are a number of people who have used this kind of method for such experiments.

Some comments on the bubble dynamics. The bubbles are going to scatter light out of the beam and hence appear dark. In addition, the pressure antinodes appear relatively bright, whereas at the pressure nodes you can have large transverse gradients in the refractive index and deflect light out of the aperture, out of the received light, and that is helpful in interpreting the results.

The first set of experiments I will show you are below the lambda-transition temperature.

(Transparency 74)

This is a blowup of a bubble in liquid helium. This

is successive frames at 6000 frames a second. In previous frames there is no apparent bubble. The bubble appears to be nucleated, for whatever reason, and blows up. Notice the bubble is at a pressure antinode because it is bright here, and that is what you would expect, because you have pressure extrema in that region. Then the bubble becomes stable and then all these guys get shoved off, because they are unstable at that position because they are large bubble fragments.

(Transparency 75, Right Hand Side)

It turns out, if you keep track of what the fragments are doing low enough, they are acted on by Bjerknes forces. This is a bubble that is larger-than-resonance size and we monitored it for several ms and it is apparently trapped -- larger-than-resonance size-- near a pressure node, as observed. Then it becomes eventually unstable. We do not know if that is from condensation or what, but it starts oscillating, apparently, back and forth.

We also did some measurements above the λ transition (Left Hand Side)-- there were some complications -- but one of the interesting things was you could produce clouds of bubbles near this resonant diameter, 500 μ. But in addition, you could produce stable microbubbles, and these are visible along here. Here is a little row of bubbles, and you know that by observing the experiment, monitoring the history, and you could see these bubbles get trapped in the sound field. That is something of interest to people who are trying to understand dynamics of bubbles where the dynamics is dominated by vapor.

(Transparency 76)

I would like to now go on to other kinds of measurement methods and in this very nice experiment of Pumphrey and Elmore they discovered an apparent source of noise, a dominant source of noise, due to rain in certain spectral regions, where they monitored here the impact of a drop on a free surface at the same time they monitored -- with a split screen -- the output of a hydrophone located near the impact point.

What happens is, from momentum conservation, when this drop impacts, there is a crater created -- here -- and then this dimple, there is a wave that runs down and pinches off a bubble and that bubble is not in equilibrium when it pinches off, and rings, because it has a differential in pressure from equilibrium. That ringing gives rise to this decay of oscillation shown here.

That is nice. It is good to have done these kinds of measurements and there has been theoretical work, but with this kind of measurement you cannot resolve the actual volume

oscillations themselves, because they are very high frequency, typically around 10 kHz.

(Transparency 77)

So I asked John Stroud for his masters project to see whether he could resolve these oscillations. He used the pseudo-extinction method we talked about before. What he did was, he skimmed the laser beam, an expanded beam, right below the surface of this water and he adjusted the height of this drop source -- the impact velocity of the drop is right in the stable range of so-called regular entrainment.

Let us look at the gross picture down here with a long time scale. In the hydrophone output there is a little bit of impact noise. Then, later on, there is a ringing signal that would appear associated with the pinch-off of a bubble at the end of this dimple.

With the photodiode output this dimple comes down and starts to block the laser beam. This comes down, you can see this. Then, as it pulls back up, the bubble is pinched off, which is right in here. What John did is, he took this photodiode output and high-pass filtered it and then amplified it by a large value, which I do not remember exactly, but it is on the order of 1000 to 10,000, then band-pass filtered it with a wide-band filter, but it is one whereby you would include this frequency range of oscillation.

This is the amplified photodiode signal and this is the hydrophone record. By measuring the ratio of this voltage to this voltage here, which is the steady-state voltage in the beam, and then estimating the diameter of the bubble that is being pinched off from the frequency, he could get the fractional volume changes (actually, the fractional radius change from such oscillations) and he got a number that is about half a percent.

It turns out nobody has done a calculation at this nonterminal velocity case, but Oguz and Prosperetti have done calculations at terminal velocity and they get numbers, I think, that would run between 1 and 2 % or thereabouts. This generally supports the idea that there are real volume oscillations. In this case it is from the oscillations of the pinched-off-bubble profile that can be measured dynamically with this optical technique.

Q: Could you estimate that size just simply from the amount of acoustic amplitude that was irradiated? You know the distance and --

A: Yes, that actually would have been a good thing to do. This was done with the leftover pieces that were not used everywhere else and, therefore, unfortunately, that

calibration was not done. But we did estimate the size from the hydrophone frequency.

(Transparency 78)

This compares FFT for the two cases.

Q: Is the hydrophone signal upside down?

A: That is a good question. I think it is addressed in the paper. I do not recall. I think, in fact, it is. Thank you.

Q: (Inaudible.)

A: What we are doing is estimating the oscillation amplitude by measuring the variation of the cross-sectional area of the bubble as it oscillates after it is pinched off. So that gives rise to this high-frequency component of the signal coming out of the photodiode. This is a way to get such information with good time resolution.

Q: (Paul Elmore) When I was working with this, the bubbles seemed to get smaller and larger due to the way the surface oscillations were going and the way it was oriented. How would you be able to account for that?

A: What we did was to calculate the shape oscillation frequencies of these bubbles -- these are typically around 2 kHz and below are the dominant modes. So the band-pass filter is set around, say, 5 kHz, to reject that shape oscillation, to avoid confusing it with the volume part.

However, we believe that the shape oscillations, which are superimposed, may be one of the reasons why this signal is not a nice exponential decay, because if you look at the equations for extracting the radial part of the oscillation, there is a mixture term that is affected by the shape.

So you can, really, only get a pretty good order-of-magnitude estimate of the volume part of the oscillations (in this application), but you might have difficulty with this system in doing better than that, as you know from your photographs, because there can be superimposed large-amplitude shape oscillations.

Q: Are you saying that the bubble oscillations you thought would were driven by a pressure gradient?

A: No, when the bubble is pinched off, it is not at equilibrium with the surrounding pressure. That is all I am trying to say.

Q: What Longuet-Higgins was suggesting, and I think I am in agreement with the photographs, is that there is a jet that goes into the bubble after it is pinched off --

A: And that would contribute to this momentum.

(Transparency 79)

Some elementary methods: Greg Kaduchak did this experiment recently, and a lot of people have done this kind of thing. One way to get an estimate of bubble size -- this is done at the bottom of an 8 ft. tank -- is to have air flow into a needle in a very careful way to produce a bubble. With a hydrophone nearby you get this record and you can then work backwards to find the bubble size.

One of the things we do not understand, and in discussion we can talk about it, if somebody knows, we do not understand why this decay is as rapid as it is. Whether or not it is because of the bubble rising, I do not know.

(Transparency 80)

That is enough of these transient methods. I would like to shift gears to scattering. This is probably the first major quantitative paper on light scattering from bubbles calculated geometrically by Davis. This is something that I recall looking at as a graduate student and realizing there are a number of interesting problems.

One of them is, this is scattered intensity as a function of angle, that is, the deviation from the forward direction. This is a critical angle edge, which we will talk about in a minute, but the important thing is there is a divergent derivative of this, and that cannot be preserved. You cannot go to the far field and have such a divergence.

In addition, if you take the equations geometrically and calculate it and go to the backscattering, you get a divergent amplitude along this axis. So that poses an interesting problem.

The way that Davis did his calculation was this kind of ray diagram and adding up many such amplitudes. It turns out, I should add, there are some other things that are wrong with that paper. He does not deal with unpolarized light properly. If you want unpolarized light, you really have to do the calculation and keep track of the two orthogonal cases, $E_{\perp}$ to this plane, and $E_{\parallel}$, trace those all the way through the bubble, because those polarizations are preserved as they go through the bubble, and then combine whatever the total irradiance is afterwards, whereas Davis calculates an average reflection coefficient at each vertex based on averaging the

polarization. That is a relatively minor correction.

(Transparency 81)

Let us talk about some elementary ideas in conventions concerning light scattering. We are looking at the far-field scattering. One of the first questions you can think about is the reflected light for different rays, for different impact parameters, S, and that corresponds to different angles of incidence, i, and different scattering angles, θ. If the angle of incidence exceeds the critical angle, then you would expect to have total reflection from the surface, because you have water out here and gas on the inside. That corresponds to scattering angles less than 82.8° for water.

One way to think about this, one way to make this more quantitative, is to use a normalized irradiance. Remember, irradiance is watts per, say, square meter at some distant observer. This is going to be proportional to the incident irradiance, a geometric factor, which we choose to be the radius of the bubble divided by two times the distance to the observer, and then you square this quantity.

The reason we use this factor is that, then, if you had a perfectly reflecting sphere and used flux conservation to calculate the far-field irradiance, then this i_j, this normalized irradiance here, would be unity. So we plot things compared with that. Some people call that the gain.

(Transparency 53)

There are two polarization cases that are of interest, $j = 1$ for $\mathbf{E}_\perp$ to the scattering plane and $j = 2$ for the $E_\parallel$ polarized case. That is what is basically plotted here, is $\log i_j$. This is calculated from Mie theory, which is the exact partial wave series for a sphere for the two different polarizations. This we see, gives a rounded transition to total reflection. We will talk some more about that.

In addition, for $E_\parallel$ there is a dip here and a loss of fringes. This is associated with the suppression of the reflection, that is, a Brewster scattering angle. Stefan Baumer observed this in his masters work, but I will not have time to talk about that.

(Transparency 82)

How are these kinds of experiments done? This is before readily affordable CCD cameras, so we did this with film. Illuminate a bubble in this chamber and then view this with a camera focused on infinity. In this case, this chamber is cut near the critical angle to get the largest possible aperture in that range.

The early experiments were somewhat simpler, not as good an aperture, and are discussed here [JOSA 69, 1205 (1979)]. If you try to read this paper, I do not mind telling you a brief story. I sent the proofs in for this paper and it came out with the key inequality reversed. Why is this? I called up the publisher and one of the publication people got interested in the paper and read it and decided that it looked better the other way, so he reversed it after the proofs --

Q: Was he right?

A: No, he was wrong. Otherwise, I would not have told you the story.

(Laughter)

(Transparency 83)

Anyway, Dean Langley cleaned up the experiment and this is the kind of photographic records he has for scans near the critical angle. You can see these kinds of structures. He has these broad bright, bright, bright, dark, dark broad peaks called coarse structure. You get a broad angular period as opposed to these fine superimposed fringes. The fine fringes are due to the interference of these widely spaced rays -- here and over here. (See upper figure on Transparency 42).

(Transparency 84)

One can use this information to measure bubble sizes and so on and learn something about the condition and the surface of the bubble.

This is a comparison between Mie theory, which is the partial wave series, with geometrical optics, like Davis, but corrected for the polarization, the divergent derivative, and then a physical optics approximation, which I will summarize briefly. These are data for a bubble the same size, and the size was 123-μ radius for a freely rising bubble.

(Transparency 85)

Physical optics considers both the reflection and this transmitted ray and makes some approximations in the early version of this, and we will talk something about improvement in a minute. The basic idea is that you would model the outgoing wavefront near the bubble using geometrical methods, then plug that into a propagation integral -- from the acoustics standpoint, a Rayleigh-Summerfield propagation integral -- to go to the far field.

That must remove this divergent derivative. This is

actually the reflection coefficient, which is a divergent derivative, versus some horizontal coordinate -- μ -- which is an effective coordinate for cutting across the wavefront.

Initially, because we could not figure out how to do the integrals otherwise, we did the simplest approximation, which is to take a step and throw away this energy, and that gave you Fresnel integrals. Then, later on, summarized later, we put this in and that gave you much more complicated integrals. Anyway, that is the origin of the physical optics approximation.

(Transparency 86)

This (upper figure) shows some more of Dean Langley's data where he measured the fine structure both for freely rising bubbles and pendant bubbles, compared them with independent measures of bubble radius. This (lower figure) is somewhat of a remarkable case where he had high-quality data on this bubble and then he did a Mie scattering calculation and adjusted the size, such that he found that for one particular size (within a 10th of a μm) these modulations would line up in both the Mie theory and the observations -- the modulations in the fine structure -- and for no other size would they line up.

(Transparency 87)

This transition to total reflection can readily be observed and it should give rise to colors. We photographed these in an aquarium in Hawaii and I could not find my slide, I am sorry, but we have the original publication if someone wants to look at this color carefully.

If you think about the configuration, what they had is an outdoor aquarium that was simulating the surf zones, and they had water crashing down into this aquarium, creating a cloud of bubbles. The free surface of the water would decay away, become nearly flat, before the bubbles had time to rise out. So then the sunlight would come down and scatter off this cloud and you can see, if you look down here, you have large scattering angles and small scattering angles up here, so it should be bright up here and darker down here, and a transition region near the critical angle, which you can argue can give you colors, and that is evident in this colored band. We have also seen this in the laboratory, say, for a bubble in glass.

(Transparency 88)

It is natural to wonder what happens with bubbles in sea water, what are their real surface properties? One of the reasons for wondering this is the number of measurements that suggest they are highly stabilized. One of these measurements is the measurement of Johnson and Cooke, where they captured

some bubbles from sea water and they pressurized them and the bubbles would be stabilized with regard to response to pressure and then there was actually a little residue left, in this case, suggesting that the bubbles coated. There were also a number of other experiments, some Russian experiments, where they actually had a chamber, they plunged the chamber to large depths in sea water and monitored with a camera what the bubbles were doing.

So it is natural to wonder, with this coating, how would this affect the bubble in sea water? We just saw how with sea water we could see critical-angle transition.

(Transparency 89)

Stuart Billette, for his masters thesis, made a program to evaluate what is called the Aden-Kerker series. Aden and Kerker calculated the light scattering, wrote down the partial wave series for calculating the light scattering of a coated object. What we did is, we took the coating to be something that was, say, an oily material (refractive index of about 1.5), and thin, and what happens is, as you increase the coating, there is a shift of this feature, and you can see it with this arrow.

But the coating does not quench the critical-angle scattering regime. This is a transition from partial reflection to total reflection and the core structure that is associated with -- I should emphasize that core structure is associated with the diffractive effects near the critical angle and the interference with this P equal to one ray, this once refractive ray.

(Transparency 90)

You can do a perturbation calculation of what does the coating do to the critical angle by asking -- if you look at the critical angle at this interface, the inner interface, the inside of the coating, and ask at what angle is it scattered to, you find a shift in the critical angle with δ, and that, in fact, describes the shift found in those previous calculations.

(Transparency 91)

You are going to hear some more about this in Seth Putterman's talk, I hope, that is, the measurements of bubble radius as a function of time, but measurement of bubble radius with light-scattering methods. But I would like to say something about the theory behind this.

We had pointed out that because of the simplifications of the ray optics near the critical angle you

should have, apart from the fine structure, a monotonic relationship between the scattered irradiance and size, and that could be useful for working backwards and finding out the bubble size as a function of time, for example.

Larry Crum assigned this to a student, Gary Hansen, to check out this relationship. I do not think Hansen believed me, because he did not pick the critical angle, he picked 80°, and I think that is why he has wiggles up here. The point is that from his smooth Mie scattering calculation at 80° he gets this monotone relationship (lower figure), and these are the data for the scattered intensity versus radius versus the theory, whereas you pick some other scattering angle, you get something that is much more complicated.

It turns out, if you want to measure only very small-amplitude oscillations, then it could be advantageous to tune your system, say, to right here (upper figure), because you have a higher slope here than here, but then you have to know everything about tuning your system. Normally you are interested in dynamic range, and it is much better to do the experiment down here.

Can we understand this curve any better? Let me take just a moment to at least underline the understanding.

(Transparency 92)

I asked Cleon Dean to look at this, to work through the Sommerfeld-Watson transformation for a bubble and then figure out the detailed asymptotics at short wavelengths for the scattering at the critical angle, and he did that. The results are summarized here. At the critical angle, the geometric optics limit, the normalized irradiance, I, goes to 1.

There is another way to plot it, I will call it dimensionless irradiance, i, that scales with size properly, and the way to make the actual irradiance dimension is to put in all these constants, but; in any event, this is basically a R^2 curve in terms of, here, its size parameter, wavenumber in water, k, times bubble radius.

What is shown here is also the physical optics, which does not agree with Mie theory, because the physical optics, as originally posed, was not an asymptotic theory, as we pointed out, and the new asymptotic series, which is the dashed curve. Let me comment on the form of the series, because it is some interesting scattering physics.

(Transparency 93)

The form of the series is that the electric field

correction term, the leading electric field term, goes as $\frac{1}{(ka)^{\frac{1}{4}}}$. That makes the series very slowly converging. Most of you who know something about scattering realize that the leading correction ordinarily is $\frac{1}{(ka)}$. This, in terms of the normalized irradiance, shows there is the Mie theory, which is the curve, and then, on top of it, there is the two-term asymptotic approximation, and if you take only the leading order correction, the $\frac{1}{(ka)}$ correction, it is this.

It turns out the correction is enormous, even when you go to a boulder-size bubble, that is, one that is several millimeters in radius, unlike normal optics. If you go to such a high size relative to the wavelength you are usually over to geometric optics.

(Transparency 85)

Why is that, why do you get this $\frac{1}{(ka)^{\frac{1}{4}}}$ correction? That can be understood in terms of this diagram and this concept of Fresnel width. What is the Fresnel width? If I reflect light from the bubble, I will get this ray, as indicated, at the critical angle. However, physical optics tells us that the scattering contribution at that angle is not simply a ray contribution; it includes information in the nearest Fresnel zone. So the outgoing wavefront is this curved wavefront, and you can consider paths that contribute to this direction but deviate in phase from the Fermat path, which is the actual ray, by π rad, and that gives the associated Fresnel width of this reflection.

If you are interested in this, my forthcoming paper in JASA on scattering from shells uses this concept.

The point here is that if you average this reflection coefficient with this divergent derivative, over the Fresnel width you find that the leading order correction is $\frac{1}{(ka)^{\frac{1}{4}}}$, which is what all the mathematics and the Watson transformation give you. So there is a simple explanation.

(Transparency 24)

To close this section on total reflection from bubbles, we saw that early on it was important in the levitation of the bubbles -- I might just flip back quickly -- this cone is from the total reflection region.

(Transparency 94)

It is also important in the photoacoustic response of the bubble, and this is from work done by Bruce Unger, where he took a bubble, illuminated it with a doughnut-mode beam -- not very high power, maximum power around 2 W -- and with a hydrophone (in order to get some gain he used an elliptical reflector here) he could monitor the response due to the light impulse incident on this bubble. This is a ringing due to this pulse. He could then Fourier transform that, calculate the natural frequency of the bubble, and then drive the resonance; that is to say, turn this light on and off at that resonant frequency, and you have resonant photoacoustic excitation. This is from his work in around 1985.

That is the end of critical-angle scattering, and we will talk about some other things in scattering after the break. Any questions?

Q: I assume you answered, but I did not quite understand the answer. Why is it that that the bubbles look so light behind the ship?

A: Because they have such high reflectivity. There is this wide range of angles where a bubble looks totally reflecting. It is basically there is this wide range between zero and 83° where, to go back to the Mie scattering calculation, which is the intensity versus scattering angle, you exceed that of a perfect reflector, on the average.

(Transparency 53)

Notice here, $\log_{10}$ of zero would be the perfect reflector region, so you have this transition to total reflection in a big angular region where you have high reflectivity.

Q: Related to that, why do those bubbles last so long?

A: That is a good question. I suspect that is one of the reasons why the Navy is putting money into this accelerated research initiative. Probably they are coated. Undoubtedly, in any sea water situation, the surface property of the bubble, once you create it, will become dirty and that will affect the rise of the bubble and so on. But the evolution of that coating is not at all well understood.

DR. MARSTON: Let me just comment, if you are interested in critical-angle scattering for measuring, say, radius time curves, you should also be aware that Nussenzveig and his colleagues have also obtained this leading order correction -- I may have forgotten to say that before -- in

agreement with Clon's results, and they also found that the diffraction integral could be expressed in terms of parabolic cylinder functions. So those of you who are so inclined can look up Nussenzveig's papers.

(Transparency 95)

The next section on scattering has to do with backscattering and also caustics. We will start off with a discussion of "glory scattering," first from drops and then from bubbles. If you stand on a mountain and look at your shadow below, as in this picture here, the light that comes back from your shadow is nearly backscattered to your eye.

What you see is this -- more often seen, of course, now, from an airplane. To study this phenomenon is the reason why Wilson built the cloud chamber, but he discovered cosmic rays, so he went off and did something else.

(Laughter)

But this remained a puzzle and there were a number of incorrect explanations in the literature until van de Hulst pointed out that there is a focusing mechanism of a different type than previously considered, and it is called axial focusing, that is, focusing along the backscattered axis.

(Transparency 96)

And that is associated with an outgoing toroidal wavefront, where you have an off-axis ray that goes in the backward direction. Then, as you deviate away from this condition, you would have this angular wavefront, and it is toroidal because you can rotate this around the axis.

(Transparency 97)

I would like to comment on the mechanism for focusing, because it shows up for bubbles and first remind you of a few things about the description of wavefronts. Consider at some point on a wavefront here P. It is characterized by two principal curvatures and two principal directions. You can take a slice through this point and rotate that surface around and you have a curvature that varies with angle of that slice.

The extreme values of curvature are the principal curvatures and the associated directions are the principal directions. In this notation those principal curvatures are ρ_1 and ρ_2.

(Transparency 98)

If you know the principal curvatures on a wavefront,

you can predict the different kinds of focusing you could have. This is a sketch of the associated ray tubes for that focusing for the case where both principal curvatures are converging but not the same; so you have one focus here and another focus here; that is, one caustic here and another caustic here.

If you try to calculate the amplitudes geometrically, you get a divergence at this caustic, because this cross-sectional area of this tube shrinks to zero. You can find a nice discussion of that in Allen Pierce's book, for example.

(Transparency 99)

What is the focusing that occurs for the glory scattering? This (upper left side) looks like the wavefront when you rotate the diagram around the axis of the drop. That is a toroidal wavefront. It turns out, if you stare at this, you realize --out here -- you have a positive Gaussian curvature, that is, the product of the principal curvatures are positive, and in here it is negative, so it has to go through zero, and it goes through zero right along here.

If you look at this previous expression for this ray tube calculation, if you shove the caustic to infinity, that corresponds to the Gaussian curvature going to zero. That means that the rays from this circle here are, in effect, focused at infinity. If you move in from infinity, it is the rays from inside here, a loop along here, due to the negative -- actually the converging part of that surface or wavefront will focus along the axis, so, in fact, the axial focusing remains as you move in from infinity.

Because we will not have time to go back to it later, let me just comment on, as a very natural question, what happens to this caustic, this bright region along the axis, when you deform this wavefront? The simplest deformation you can consider is this harmonic perturbation shown here. If you use a lower harmonic it just corresponds to tilting this, and then just rotate the wavefront to another direction.

Michael Berry made a conjecture that this would give rise to an astroid caustic -- the one shown here -- and Pat Arnott, we will see later, has observed this caustic and actually verified it with analytical methods.

(Transparency 100)

Let us go back to the simple circular torus again. This shows up for bubbles and the relevant ray diagram is shown here, where you have off-axis rays that are backscattered, and if you try to use Davis' equations, that is, geometric flux conservation, you get the backscattering amplitude, you get an infinity due to such rays.

So you have to put in a diffractive correction. The way this is done is to approximate the near-field amplitude geometrically and then plug this into a propagation integral to go to far field and, when you do that, you have an amplitude factor here of $ka^{1/2}$, which is characteristic of axial focusing. It is weaker than the perfect focusing from a lens, that is, aperture-limited focusing, where that would be, in effect, ka scaling. This is the axial-focusing factor.

This is for cross-polarized scattering, and we will define that in just a minute, but there is a J_2 Bessel function dependance and a dependance on azimuthal angle, as indicated. This is for a spherical scatter.

(Transparency 101)

We did a study of that (Dean Langley did this work) and the first measurements were done for bubbles in a highly viscous oil, because we did not need a very high-power laser to study them. The way the experiment is done is to illuminate with a laser beam -- let us say polarized out of the plane of the projection screen, so the E field comes out of the plane -- and the light that goes by the bubble goes into beam dump.

Then you look at the backscattering with a camera focused on infinity so you can map scattering intensity versus angle that way near the backward direction, but you look through a polarizer. You can either polarize the polarizer so that it passes light with E field parallel to the screen (and that would be cross-polarized scattering) or you can rotate the polarizer and pass light out of the screen (and that is co-polarized scattering).

(Transparency 102)

When you do that, this is what you see. You get this symmetric pattern for the cross-polarized scattering for a spherical bubble. The symmetry, the sin 2ϕ dependance of the amplitude of the corresponding square for the intensity, is actually true of any spherical object. But the reason this is so strong near the backward direction is because of this axial focusing.

This is the co-polarized, where now you actually have a peak right in the backward direction and, if you pull the polarizer out, you see the total. Pat Arnott returned to this problem and studied this for bubbles rising in water, where he had to use an argon-ion laser, and this is the associated cross-polarized scattering for a bubble that is sufficiently small to be spherical on the scale of a wavelength.

(Transparency 103)

This is the calculation of the kind of theory I just summarized, physical optics approximation, with Mie theory for this bubble and you get very good agreement.

(Transparency 104)

However, as the bubble rises, it does not remain spherical if it is large enough, and that is because of the flow past the equator and the low-pressure region becomes oblate. There are various approximations in the literature for this oblateness, and this is one of the approximations for this factor, which is $\left(\frac{D}{H}\right)^2 - 1 = \Gamma$ versus bubble diameter, which is this curve.

What Pat did was to measure the opening of this so-called unfolding of the glory into this asteroid pattern as a function of size and then work backwards to find the oblateness, and this is what he found.

(Transparency 105)

The way that this wavefront distorts can be calculated by considering two different cuts through the bubble, one through the horizontal plane, where you have a circular cross-section, one through the vertical plane, where it is no longer circular and gives rise to this disordered wavefront. He looked at the cross-polarized scattering.

(Transparency 106)

But you actually have some different choices now on how you orient the cross-polarized scattering relative to the vertical axis, but for one of the choices this is one of the simple results. This is the observed pattern where the bright regions correspond to the cusp. You do not have just one cusp, because you have many such rays. This is a calculated pattern from a physical optics approximation. We do not claim that every dot lines up, but we have the essential features of the observed scattering. This can be used to infer information about backscattering and bubble shape.

(Transparency 107)

I would like to comment before we close on diffraction catastrophes. Just briefly, these ideas are applicable to elastic objects in water with an axis of symmetry. This was demonstrated by Kevin Williams for Rayleigh waves on a spherical scatter and by Steve Kargl for thick shells and by Greg Kaduchak for thin shells, where the relevant scattering diagram is shown here. You get a very strong enhancement near the coincidence frequency of this fluid-loaded

shell.

We have specular reflection off this point here and a guided-wave contribution, which, according to theory, is about three times the specular reflection for any class of shell where this is a distinct effect from other contributions and, in fact, there is good agreement with theory. It is a highly dispersive system, so if you chirp the input signal in the right sense, because of the dispersion of this waveguide the signals pile up on each other and you get an even greater enhancement of the scattering.

(Transparency 144)

You might think this is of limited application. Let me go to my very last transparency in the series. Tim Stanton, at Woods Hole, has been applying these ray models now to critters, and it is particularly important for something near and dear to us, namely, food. These are beginnings of certain food chains in the ocean, and there are attempts to model these acoustically, and he wants to understand the scattering cross-sections of such natural objects and he has found that these kinds of ray models are useful in understanding those scattering cross-sections.

(Transparency 108)

So can we find some other ways of getting this elastic information from scatter? This is work done by Kaduchak and Kwiatkowski, who is here. It was presented in Denver and submitted to JASA. It is a very broad-band source, which, we will see in a minute, has application to bubble studies.

Consider an ideal transducer which is impedance matched to the water and very thin -- in fact, we use a large sheet of 52μ thickness. One can show, if you have an impedance-matched system, that if you were to apply a δ function voltage to this system you would get two pulses coming out.

You can show this by writing down the full piezoelectric equations for a 1-D system with a source term and, as 1-D, you know how to integrate the domain, domain of influence, and you end up with this result.

This looks something, from the point of view of a long time scale -- if I am not interested in time scales on the order of the propagation of sound across this sheet -- it looks like the derivative of a δ function. So if you put in a step function in voltage, what should be radiated is approximately a δ function. In fact, the detailed theory for an impedance-

matched system gives this pulse.

There are a number of corrections you have to put in, the fact that you have reverberations in the sheet, and so on, but if you look at it on a time scale long compared with this width, what you see is something that is approximately a pressure impulse that is affected by your hydrophone response.

(Transparency 109)

This is from their manuscript. We have the amplified hydrophone output -- not a very good hydrophone by medical ultrasonic conditions, but good for underwater acoustics, it has a resonance around 350 kHz. There is this pressure pulse and then there are some hydrophone wiggles. If you wait a long time, you see an edge signal comes in. You make the sheet as big as possible, so your experiment is all over before the edge signal comes in.

From this theory I just summarized the prediction is that the relatively slow-scale pressure components are proportional to the current going into the sheet. So you monitor that, and for this system the peak current is about 7.5 A for a step-function voltage applied. Now, you cannot actually apply a step-volt function; you have to pay a lot of money for a voltage source to do that, so there are some limitations there.

But in fact, this is on the same time scale as this one and if you overlay them you see, sure enough, they have the same kind of profile, and this is the FFT of the hydrophone record.

(Transparency 110)

What can you do with this? First of all, you can look at resonances of objects, exciting the modes, and this is an experiment that Greg did in this 8 ft deep by 12 ft diameter tank. You can see the elastic features of the structure -- this is for this a_0-, that is, coincidence frequency guided wave that has a frequency of about 300- kHz. There is a very low frequency response present that you might think is the breathing mode. It is not. It is due to another quadrupole mode of the system.

It turns out the specular reflection has information. You can extract the areal mass density of the shell from details of that specular, and we worked out the theory for that. This is based on the partial wave series that are filtered in a way that approximates the hydrophone, or has the same kind of features as shown here.

(Transparency 111)

We were interested in whether or not this can be used for bubbles, because it can be hard to measure the impulse response of a bubble. Greg has done this experiment, and it is a hard experiment to qualify the results. He took a polymer mesh and attached some bubbles on this, put it next to the PVDF sheet. He recorded a record and then he went down with some water and squirted the bubbles off, recorded a record, and subtracted the two.

(Transparency 112)

You get a lot of stuff here, partly because there is probably a little bit of motion to the PVDF sheet in the process, but there is some low-frequency stuff here. He could also estimate the size of the bubbles. If you window this in time, you see this. If you Fourier transform it, you see a peak here, which is remarkably close to the estimated frequency for the size of bubble. Whether that is bubbles, well, we want to do some more experiments, but it is promising.

(Transparency 113)

I would like to give you really condensed discussion of diffraction catastrophes and explain at the end why you would do these sorts of experiments. Acoustic levitation affords you the ability to take a scatter like a drop in air and to adjust its shape for the reasons we have seen earlier.

This is a view from above, a top view, and a side view. You illuminate this drop, vary the acoustic amplitude -- you do two things. We focus a camera on the drop to measure its shape, we focus a camera at infinity to measure the far-field scattering.

(Transparency 114)

You all know what you should see if the drop is spherical. For spherical objects you have this transition from no rays of the once reflected, twice refracted class to two rays over here. That curve is at the rainbow angle transition. Airy modeled this with a cubic wavefront and so on.

(Transparency 115)

And sure enough, if you make this system nearly spherical, this drop nearly spherical, you see this pattern. This is the Airy pattern for transition from no rays to two rays. However, the surprise was that if you take this now and you flatten the drop, you see this sequence of patterns as increasing flatness. The transverse cusp comes into view. There is a special focal section here where the scattering is even stronger and then it unfolds.

Q: What are the different parameters of the panels?

A: What you are doing is making the drop flatter as you go up in letter. A is nearly round, a little bit flatter, a little bit flatter, and so on. What we are looking at is the angular scattering pattern, vertical scattering up, horizontal scattering horizontal. And you are looking at the scattering at infinity by taking this camera and focusing it on infinity.

(Transparency 116)

It turns out there are a lot of reasons for understanding these kinds of patterns and we will get to these right at the end, but let us look at this problem, because it is a very well-defined problem in the sense that you can measure the curvatures of the drop and the light goes in and rattles around and you can calculate a number of things about it.

The way to think about this problem, how to classify these patterns, was discussed in Berry's Citation Classic paper called "Waves and Thom's Theorem."

(Transparency 117)

It is in terms of what is commonly called catastrophe theory. What catastrophe theory does is look at the far-field scattering, where we are interested in the ways in which the Gaussian curvature vanishes. Catastrophe theory allows you to classify the different ways the Gaussian curvature can vanish.

(Transparency 118)

That is all summarized in this transparency and we can talk about it afterwards if you want to know about that.

(Transparency 119)

What are the qualitative features? It has different parameters. The parameters can be scattering angle (these different horizontal and vertical angles) or they can be things like shape of the wavefront, the drop shape. Those are known as control parameters.

This shows the singularity of the differential mapping that describes these caustics, that is to say, the differential mapping describes where the rays are, the mapping becomes singular at caustics. The ordinary Airy caustic is the transition from no ray to two rays. The cusp caustic is a transition from one to three, and then the hyperbolic umbilic we will be interested in is the one down here, where you can have zero, two, or four rays.

(Transparency 120)

What Michael Berry and others claimed is that you can express such wavefields in terms of this canonical list of diffraction integrals, where this so-called state variable can be a transverse coordinate. What is the difficulty? The difficulty is that in order to realize such a description you have to be able to transform the variables in your problem into such a canonical form. That is, you have to carry out a smooth transformation that preserves the class of catastrophe and puts the integral in the right form.

(Transparency 121)

We worked this out for a number of cases and we will see later on some of the practical applications of this. The way we have done that is what is sometimes known as generalized ray tracing. It is something that is used in design of imaging systems, where one, instead of just tracing a ray, you trace the principal curvatures along a given ray for a wavefront as it goes through a system.

What that tells you is what all the rays are doing in an infinitesimal neighborhood of the ray of interest. By tracking those principal curvatures you can then construct the outgoing wavefront and carry out the required smooth transformation.

(Transparency 122)

A common such caustic is this tea cup caustic. If you have a glazed cup and reflect light into your tea, you can see this pattern. It turns out this is a cusp-like caustic, it is a stable caustic. If you took the cup and deformed it, it would still have a transition from a one- to a three-ray region, though the details of the catastrophe are not changed.

(Transparency 123)

A long time ago -- in the 1940s -- Pearcey described that wavefield. It is one of the canonical diffraction integrals, the so-called Pearcey function. It is a 1-D integral and has a number of interesting features. There is a three-ray region, three real rays, a one-ray region out here, one real ray, but you actually have a complex ray here, where you can have interference between the complex ray and the real ray (a complex ray being one that has complex coordinates for this S).

Then you have one real ray out here, where you no longer have such interference. The problem is how can we describe this transverse cusp, because the transverse cusp is transverse to the orientation of the outgoing light.

(Transparency 124)

Here is the drop. You put a screen a long way away. You look at the screen and you see this pattern. The caustic is oriented orthogonal to the way the light is going. It is an intrinsically three-dimensional wavefield described by a two-dimensional diffraction integral.

It turns out that by tracing rays through the drop and you think about it, you can figure out what the relevant wavefront is.

(Transparency 125)

Which is shown here. By writing down the diffraction integral, it turns out you can convert it into the Pearcey form.

(Transparency 126)

To close with a couple of examples, we verified this was the right shape by bending the reflecting surface in that form. You put in ground-glass screen so you can image the light, illuminate it from a divergent source, you move the screen around, and you see this caustic that opens up transverse to the reflected light.

(Transparency 127)

Carl Frederickson took this kind of reflector, put it in a tank of water and illuminated it by ultrasound and found

-- this is the measured pattern. It is bright where the sound is intense. This is the calculated pattern without any adjustable scaling parameters. He could also vary the wavelength by a factor of about 10^3 by switching over to light with the same reflector and he could measure out where the caustic was.

(Transparency 128)

You can do this in a time domain and you have this so-called triplication in the different arrivals.

(Transparency 129)

We worked through this analysis for the spheroidal drop. It is a very well-defined problem. You calculate the parameters of the outgoing wavefront. The way that you would think about it is the vertical curvature has to vanish for the light going to the caustic. By symmetry, the vertical

curvature will vanish if this light is focused on the back of the drop.

(Transparency 130)

If you calculate at what scattering angle that is, that should be the scattering angle of this cusp point -- here.

(Transparency 131)

And sure enough, scattering angle versus aspect ratio, and it is in agreement with measurements. It turns out you can also understand the transition, numbers of rays, and so on.

(Transparency 132)

You can do this for white light, and there is some interest in what would you see from nature if you had sunlight-illuminated oblate drops, which is what you have in rain. It turns out this cusp is remarkably white. John Nye in Bristol and also Michael Berry have written papers trying to explain why that is so white.

At the hyperbolic umbilic focal section (plate 15) you have these rather bright patterns. This (plate 16) is part of an E_6 catastrophe that we will not have time to go into.

(Transparency 133)

Also, doing the wavefront tracing, you can calculate the rate at which the cusp spreads. Cleon Dean did that and compared that with theory.

(Transparency 134)

You can also do this for this pattern near the focal section. It turns out the wavefront in that case is shown here, and then you work through the diffraction theory.

(Transparency 135)

Without any adjustable scaling parameters, this is a comparison between measurements and the analysis.

(Transparency 136)

Nye predicted that these cusps should come together in the backscattered axis for a certain aspect ratio. We call that a lips event.

(A translucent plastic sheet is pleated between two clear transparencies.)

You can actually see that in the following demonstration. You see there is a cusp there and, as you pull this sideways, you can lose those cusps, and it turns out you can lose them by two cusps merging.

(Transparency 139)

Harry Simpson did a study of this and analyzed the wavefront and the caustic shape and he observed this by looking at the backscattering. And, sure enough, if you look at the backscattering you find that these cusps come in at the right aspect ratio.

(Transparency 140)

Let me just comment, there is a higher number of catastrophes that show up for the E_6 six rays merge.

John Nye carried out an analysis of this case so we could understand Greg Kaduchak's data and, if you are interested in that case, I would refer you to these forthcoming publications.

(Transparency 142)

In closing, does this have broader applications? The answer is yes. The same kinds of diffraction integrals, the same programs that we used in some of this other work have been applied in collaboration with Kevin Williams at APL, University of Washington, and a graduate student at Washington State University, John Stroud, in this forthcoming paper.

To model the caustic wavefield produced by a randomly curved surface -- this shows up if you reflect sound off the surface of the ocean or off the sea bottom. Actually, the original motivation was that Kevin was interested in modeling the wavefield under the ice canopy in the polar Arctic, and he actually sees evidence for such caustics.

But we have now done this also in the laboratory and understand a number of the parameters. These same kinds of transformations are used to analyze the problem, but they show up in other problems.

(Transparency 143)

This is a photograph of a sea floor on the Sinai coast and you see these caustic networks caused by the ripples of water on the free surface, and the same kind of mathematics is applicable. Thank you for your attention.

Q: What were you doing with the plastic sheet?

A: What I was trying to show -- and it normally works when I try this, but I am not sure if I can get it to work.

(A translucent plastic sheet is pleated between two clear transparencies.)

If I deform this in the right way, I will get lips. I will get a cusp that comes together. What happens is that as I vary a parameter -- there is a cusp right here, so there are three sheets in this region and one sheet out here and out here. Then as I move this around, this cusp moves. If I did it the right way, you would see these lips. You vary the parameter and these cusps come together at this lips event.

It is analogous to the transition from three rays to one ray. The intersection of those dark lines is analogous to the caustic. If I did it the right way, I could actually get a cusp on this end and, as I stretch it, these cusps merge and disappear at a critical event, this lips event, which is what Harry observed. That was his masters work, I should add.

RESONANCE ULTRASOUND SPECTROSCOPY

DR. ANDERSON: I have put my references over here on the board, and the last sheet of the handouts has got a list of references. I am sorry I do not have the references written out so that they come in, in a natural way, but this might do.

This field is called resonant ultrasound spectroscopy, and I will explain why those three words came about. They came because we in solid state physics are interested more and more in measuring smaller and smaller things. This is where resonance comes into its forte. It is better than ultrasonics for small samples in a number of ways [1], and the following explains why.

If we take a simple calculation for the frequency and put in the numbers, we find that to get a specimen 1 mm in size with a sound velocity of 6 km/sec, we have to have a frequency of about 150 MHz. We are getting into the gigahertz range, and it is tough doing gigahertz experiments; the electronics are very difficult. On the other hand, if we take the resonant-type equation, we find that for a 1 mm specimen, our frequencies go down to about 10 MHz. Megahertz frequencies are easy to handle (it is radio frequency), and so the electronics are greatly simplified. But there are some other reasons why this is a good technique.

First, in resonance, in one run you have all the modes: from those modes you may find all the elastic constants. Elastic constants are (in some complicated way) linear functions of a subset of the modes. Once you have got the spectrum, you have got all the experimental information.

In ultrasonics, on the other hand, you are propagating a plane wave -- at least, that is my definition of ultrasonics -- and, therefore, you get only the information in that direction. So therefore, you must propagate in several different directions, and this requires different cuts of the crystal, or perhaps cuts of different crystals. If you start with a small sample, of, say, an orthorhombic crystal, the crystal is very hard to cut in such a way that you will get all the propagation directions you need for ultrasonics. A number of cuts is required. A tiny crystal is hard to cut into a number of subcrystals, each carefully oriented.

The third advantage, and perhaps the most important, is that with resonance you almost completely eliminate the Q of the measuring system. In ultrasonics you have to worry about the glue and the transducer and how they are affected by frequency and temperature. This requires a big data reduction. If you use ultrasonics on a small sample, your transducer takes a lot of the mass of the system, so your corrections get to be

very large and, therefore, your Q problem is difficult. So the Q of the system obscures the Q of the sample.

But in resonance you have hardly any Q at all coming from the system in the data measurement. Look at this. I am going to hold this parallelepiped at two corners between the palms of my hands. I resonate this by my lower palm (which is my primary transducer), and let the normal modes pick up on my upper palm (which is my secondary transducer). Where is the glue? There is not any glue. This specimen is being held very lightly by gravity, so the force on the ends of the sample is small. It is possible to reduce the force gradually until it is vanishingly small. You can easily find the frequency of the mode at zero load by extrapolation (as Don Isaak will show you later), and in that case the parallelepiped is essentially vibrating in free space.

So because there is no glue, the Q of the modes you measure is of the order of 5,000 to 50,000, and the Q of the system arising from the electronics does not affect the measurement. We are finding out, because of this measurement, that the Q's we measure in ultrasonics perhaps have no meaning. While we thought we had been measuring the Q of the sample, we had really been measuring the Q of the system. So the fact that you have eliminated the system's effect on the modal frequency is an advantage that leads to high accuracy and sensitivity.

The other thing is, because there is not any glue, you can take this to high and low temperatures without the problems associated with glues and, therefore, the temperature range goes from a few degrees Kelvin up to 2000 K (we are not quite at 2000 K yet; we are up to 1800 K), but there is no limit, in principle, that I can see. I am sure, if we could solve the solid state creep problems in our holding devices, which is a materials science problem, we could probably get up to 2000 K or 2200 K (or in that range) with these measurements, all because no glues are used.

The experiment for finding the spectra of modes is not difficult. The analysis tends to be difficult, because you have to do the spectroscopy. The measurement gives a sequence of vibrational frequencies called modes. But just as in a Raman spectroscopy, you have to identify those modes; and you have to be sure they are identified correctly. Then you have to calculate an equivalent spectrum, and that is tough for low-symmetry crystals.

It is not nearly as tough, I suppose, in this day of computers, as it was for Rayleigh, Lamb, and Love. Nevertheless, the analysis takes you into the domain of mathematical physics quite deeply: into potential theory, where you have to deal with long strings of LeGendre polynomials or Bessel functions or spherical harmonic equivalent potential

functions. You have seen books on library shelves about potential theory. These books you get acquainted with.

There are two ways to handle this problem. One way is to study potential theory and find out how you handle Bessel functions and LeGendre polynomials. You see, you have to have those, because essentially you have got to be sure that you can define a plane surface mathematically, and you have got to guarantee that in your analysis there is zero pressure at all the surfaces. Therefore, you must have an analytical method that not only describes the size of the specimen, but also accounts for the corners, the edges, and so forth. Love did it with spherical harmonics because he analyzed a sphere. But for non-spherical specimens, a whole string of Bessel functions or LeGendre polynomials or a power series is needed. You must be able to cope with that, so either you have to learn it or go over to the mathematical physics department and get somebody to write your program for you. In our lab, we have probably steered more to the latter than to the former.

(Transparency, Figure 1)

I have put a schematic diagram up here to show you where our names come from. This is not intended to be inclusive, but in the discipline of physical acoustics I list two fields (out of many). One of the fields is ultrasonics, and one of them is resonant ultrasound spectroscopy (RUS).

In both fields there are techniques. In ultrasonics there is the phase-comparison technique and the pulse superposition technique. I was told last night that the latter should be called pulse echo, and that is an equivalent name. The list of acoustic techniques is not exhaustive; there are many other ultrasonic techniques.

Corresponding to the techniques of ultrasonics there are the techniques of RUS, or resonant ultrasonic spectroscopy, and they have to do with shape: the resonant sphere technique and the rectangular parallelpiped resonance technique are the most important. So you often hear the letters RUS, which is a field of physical acoustics, or you might hear RPR, which is the technique in RUS.

The reason that the techniques of RUS have different names arises from the mathematics of doing the inversion, in which we have to use all those Bessel or Legendre polynomials. The mathematics of each technique is quite different, depending on the shape of the specimen, so it is easier to identify the inversion scheme (the mathematics) by name.

(Transparency, Figure 2)

Now I would like to introduce you to the resonance of

spheres. I am taking this from Love. Love's book, A Treatise on Mathematical Theory of Elasticity, (Dover, 4th Edition, N.Y., 1944)[2]. On Page 284 of my edition, it says, "Frequency Equations for Vibrating Spheres," and the next four pages are filled with the math. You can read that at your leisure and see how this is done, but I will give you a brief sketch of Love's derivation here.

First, you have equations of motion, and he defines the harmonic function (Equation (1) in the transparency) and, of course, the Laplacian. You have displacements, u, v, and w, in the x, y, z coordinates, and the equations of motion are given by Equation (2), where this is ρ (density), and this is the potential energy of strain. You put in a small vibration (Equation (3)), and you introduce an ω (for frequency) (called p by Love) squared into the equation of motion (Equation (4)). Then you transfer the whole thing into spherical coordinates, and you have Equation (5) to handle.

It takes two pages of Love's mathematics to get the equations of motion for a sphere. By these equations of motion, you have the displacements u,v, and w vibrating back and forth everywhere. But now you put in the boundary condition to find the normal modes. In other words, we ask, what is the frequency at which the surface has displacements, u,v, and w, that are nonzero, but at the same time, the traction at the surface is zero. In Love's words, "the traction disappears at r equals a".

So you use a bunch of other mathematics and end up with two equations, Equations (6) and (7), which can be solved analytically. How do the original equations of motion reduce to the situation in which the stress is zero across the surface? I won't go into details, but Love introduces spherical harmonics and requires that the pressure be zero at r equals a.

(Transparency, Figure 3)

And you have a complicated equation that has two main roots, and these are the equations right here. You have to satisfy both of these -- Equation (8) and Equation (9) of this transparency -- both of equations leading to transcendental numbers for the roots. This is an analytical solution. Only a sphere or an infinitely long right circular cylinder leads to analytical solutions. The sphere must be homogeneous (that is, isotropic), and it must be elastic.

This transcendental equation has roots, and they are spaced in an infinite series, and, referring to the transparency, these are the roots, and they are spaced in an infinite series, and they are the vibrational modes. Whenever the frequency satisfies one of these equations, you have a

normal mode of vibration -- it is an eigenvalue problem, you can think of it that way -- arising from the fact that the stress is zero, or the pressure is zero, on all the surfaces.

This is Love's solution. A few years before that, Lamb solved this with Bessel functions, but he did not get nice simple equations for roots. The reason he did not was because he did not use spherical harmonics. Because Love used spherical harmonics, he easily got a solution at the surface, r = a. So in Lamb's solution, it is done numerically, and in Love's solution it is done analytically. But I think this points out the idea that unless you use spherical harmonics on a homogeneous spherical object, you must solve for the roots by numerical methods.

The first type of mode was named by Love as a rotary solution, but this was later changed to torsion. It is labeled T. This second solution is called -- Love did not name it, but we call it -- spheroidal, or S.

Now Sato and Usame [3] solved the problem again as a basis for finding the normal modes of the Earth, and they prepared tables of the modal solution. Frazer and LeCrew [4] did the first resonant experiment for a small sphere, and they prepared graphical solutions for the modal frequencies.

It turns out that the transverse solution, the torsional solution, is independent of Poisson's ratio. In this equation, the $\kappa a/\pi$ (under Equation (8)), is a function only of the elastic constant, μ, the shear constant of an isotropic material. There is no Poisson's ratio in the solution, and so the torsional modes are independent of Poisson's ratio. For that reason, Poisson's ratio is a ratio of frequencies. That is to say, you can measure Poisson's ratio to the accuracy of a measured frequency, which can be accurate to five significant figures. The second solution (under Equation (9)) has two elastic constants, λ and μ. Because two elastic constants are involved, you cannot determine the shear velocity without evaluating the radius of the sphere, a. As a consequence, the determination of v_S is limited by the ability to measure a diameter, rather than a frequency.

The frequencies you can get to one part in 10^5, but it is very hard to measure the diameter of a small sphere to that accuracy, so usually it is one or maybe two orders less. Nevertheless, we have really got some accuracy with this method.

(Transparency, Figure 4, from reference [3])

Here are some of the modes (plotted in tabular form) calculated from the equations presented above. They are

dimensionless modes, because the factor a makes them dimensionless. Below are the torsional solutions; they are independent of Poisson's ratio. Above is a table of the modes versus Poisson's ratio.

In solids (materials of lower symmetry) one mode is always independent of all but the shear constant, C_{44}, just as the homogeneous sphere is independent of all constants but μ.

Now for the nomenclature of these modes. You see ${}_1S_2$ or ${}_1T_2$. This was Love's notation. Consider ${}_nT_p$. The n means the number of nodal radial surfaces, and the p is the number of wavelengths around the polar circumference. So ${}_1T_7$ means there are seven wavelengths in the sphere, and there is one radial nodal surface. You can measure 40 or 50 modes with this type of experiment, because the Q of the system does not affect the experiment, and the resonance arises only from the vibrating sphere. It is not uncommon to measure 50 modes. Since a large number of modes can be measured, another advantage crops up.

Suppose your crystal has eight independent elastic constants, and you measure 48 modes. Then you have overdetermined the elastic constants six times, so there is a great possibility for statistical minimization in the data analysis. Some important cross checking is available. If you find that you are not getting the same elastic constants with all these different combinations of modes, then you have made a mistake some place. So we always do overdetermine the elastic constants, to make sure that we have not slipped up some place.

Q: Could you remind us what Poisson's ratio is, please?

A: In an isotropic solid, it is the ratio of the extension of a rod to the lateral contraction. In a cubic crystal it corresponds to C_{44} over C_{12}. It is the ratio of a longitudinal constant to a shear constant, or something like that. It is a dimensionless constant, ratio of elastic constants, which has to do with the ratio of the shear motion to the longitudinal motion. In crystals of symmetry lower than cubic, you have more than one Poisson's ratio, up to three.

Q: If you pull something, it is really how much it shrinks versus how much it stretches?

A: That is right.

Q: And the ratio of it, if it is something that is incompressible, can be a half. If you pull it this way, it is going to have to go down.

A: Well, we are going to dwell on spheres for a while, because I think there is a lot of useful pedagogy here about

principles that comes out in analyzing resonating spheres, although I do not think there are many applications today in vibrating spheres.

(Transparency, Figure 5, from reference [4])

Here is a plot of normal mode solutions for elastic spheres, showing Poisson's ratio on the bottom, and the various modes. Here is $_1T_2$. This is the fundamental mode in torsion. Note that it is independent of Poisson's ratio. This is the third harmonic, fourth harmonic. This is the second harmonic of the first node, and so on, and you see that these lines go straight across, independent of Poisson's ratio. Spherical modes, identified as dashed lines, go up as Poisson's ratio increases. A vertical line on this plot shows the sequence of modes. The sequence is different for each value of Poisson's ratio. Therefore, this plot can be used to identify all the modes.

You can find the solutions to first order by matching the spectrum of the sphere to the chart shown in Figure 5. Don Isaak measured the spectrum of a ball bearing last night in this room. Here it is. Note how sharp the peaks are.

(Transparency, Figure 6)

Now, as we look at the spectrum of the ball, we observe a double mode, followed by a single mode, followed by three modes together, and so on. Referring back to Figure 5, this sequence falls at a Poisson's ratio of about 0.26 and between the modes $_1T_3$ and $_1T_4$. This is the approximate solution, and it suffices to identify all the modes. These modes are not the same size, of course, but there is one and only one Poisson's ratio where you can get a matched spectrum. Once the match is made, the mode identification follows. That is the easy way to start the analysis, but actually, in the final data reduction, the Poisson's ratio is ratios of certain frequencies.

Q: What kind of ball bearing is that, a stainless steel ball bearing?

A: This was one of the new things ceramicists make. It is nitrogen silicon, isn't it, Don?

A (from Don): Silicon nitride. We were talking about it last night.

It is very hard. It doesn't seem as though it has got nitrogen in it, but NASA likes these. It is very hard and used as a ball bearing in NASA some place, I am not sure where. Now let us look at what we mean by the nomenclature in mode identification.

(Transparency, Figure 7)

In this figure, note the mode $_1T_1$, meaning there is only one nodal radius, one nodal sphere, and then we are looking one wavelength in the circumference. Now for $_2T_1$. There are two nodal surfaces. If you start getting any higher harmonics, it gets very complicated to draw.

Q: Does this mean there is displacement inside the sphere?

A: Yes, it means that the displacement of the atoms is going this way in this nodal zone, and this way in the other nodal zone. The radial component is guaranteed to be zero at the contact.

Q: These are torsional modes, right?

A: That is a torsional mode.

Q: And the harmonic refers to what?

A: The overtone. I cannot draw the overtone. It is too complicated geometrically.

Q: Stress is zero at the surface, or the displacement is zero at the surface?

A: The radial traction is zero; the stress is zero in the radial direction of the sphere. It does not have to be zero in the tangential direction when we consider the transducer. Q: Is this the cross-sectional view of a transducer?

A: Well, you see the transducer there and the transducer is oriented this way.

Q: So you have to use different transducers for driving the spheroidal modes?

A: It does not really matter how you drive the transducer, as it turns out. What happens if you drive it differently is that the spectra will have different amplitudes but not different frequencies. If you drive that with a different transducer, oriented another way, you may not get these peaks this high, and maybe these peaks will be higher, and so on. But the frequency spectrum itself is only affected by the height as we change transducers.

So this is what you do in this business. If you wonder about a peak that is just above noise, you might put in a different orientation on top of the transducer or change the transducer. If your peak is a true mode, then it will become quite apparent.

Q: It is the coupling of the modes that is different for different driving configurations, but not the frequencies?

A: You can think of it this way. If you look at the slow-motion camera shot of a kicked soccer ball, it goes through these modes. It does not matter how you kick it, it still goes through the same modes. Its translation depends on how you kick it, but not these motions here. You could kick the soccer ball underneath, like this, and it would rise up, but it will undertake these same internal modes.

Q: In theory, coupling may be zero, but in practice it is rarely zero, but it may be small enough to be an annoyance, so you may miss some of the modes.

A: The advantage of this beautiful system is that you are going to get the spectra no matter how you load it, although the amplitudes may be different, depending on how you do it. If you go to low-symmetry crystals, then, of course, the direction of the transducer is going to identify certain modes differently from corner to corner. The spectrum will look different when it is oriented like that than it will like this in that some modes will be higher and some will be lower, but the position will not change. This is a great advantage to you. If you are uncertain about a mode, you just rotate the sample until that mode comes up.

You see, you could have a situation in which the transducer is pushing this way and the normal motion at that particular time due to the structure is that way, almost perpendicular to this motion, so the mode is barely perceptible. But you turn the sample over, and now that motion will be different; that same mode will be reinforcing the motion of the transducer.

Q: I think the point is just that a transducer does not couple very strongly to a breathing mode.

A: I know, but it does, a little. Then, of course, you have already got the answer: I had better get a different transducer. But there is not any transducer that is good for the whole spectrum, that is optimum for the whole spectrum.

We got into the spherical resonance business 25 years ago, when I was at Columbia University. The reason we did was because of the Apollo lunar science program. NASA was bringing back lunar samples, and we wanted to get into that program to measure some of the lunar samples. Clearly there were some things that could be done. It was obvious that the p velocity and shear velocity of rocks could be measured ultrasonically.

But we decided to go after something a little different. We wanted the glass spherical grains in the lunar soil, and it

was a good thing we did, because the rocks turned out to be a mess. But consider the moon in a vacuum, a cold vacuum. A meteorite comes in -- pow! Around the edge of the meteorite crater there are ejecta going out. Part of the ejecta is melted. Splashes, being liquid, form as spherules to reduce surface tension and are cooled rapidly by radiation. There is no conduction of heat away from the spherule, because there is no atmosphere.

So we figured that small droplets would be beautifully quenched glass and they would be in the soil at a certain distance from the crater. We had no idea how big the droplets would be, but it seemed as though, if we wanted to get close to a spherical shape, we would have to be content with quite small samples, because the larger samples would almost surely be elongates, not spherical enough to be of use. A heat radiation calculation convinced us that spherical droplets would be close to 1 mm in diameter, certainly not as much as 5 mm.

We wrote a proposal to NASA and it was accepted, and resonating lunar spheres was our main project. But in order to have our proposal accepted, we had to do other things, such as make an ultrasonic determination of wave speed of the lunar rocks. Nobody else had thought of resonating glass spheres, but we still had to do other things NASA wanted done. To prepare ourselves for the measurement of small spheres, we measured the elastic constants of small spherical tektites before the lunar landing.

(Transparency, Figure 8)

Now we go to the Soga and Anderson reference, [5]. At that time, John O'Keefe was king of this theory of tektites, and he thought they came from the moon (nobody believes that anymore, except O'Keefe). We now are sure tektites are splashes from craters made by meteorite impacts on the Earth. But at that time they were thought to be splashes that came from the moon.

So we thought, fine, a tektite is lunar glass. We will get some tektites and make a small sample about the size of what we expect to find in the lunar soil and measure it to show NASA that we could really get good results. Those are our results up there on top of the transparency, Figure 8. You see we have torsional and spheroidal modes identified in the spectra. There are two spectra, one for Moldavite and one for Indochinite tektites. The middle picture shows the tektite we got, and the sphere shows the sample we resonated. The sample is about 4 mm in diameter.

We were very careful to make it spherical. We had Walter Bond at Bell Labs help us make good spheres. So they were real idealized spheres.

The bottom of the transparency shows the results. The accuracy shows what it is possible to measure. Look at Poisson's ratio here, four significant figures. I have since wondered whether we were justified in publishing four figures for the shear velocity, but there they are. Good results. We had two elastic constants to measure, two independent constants, and we had about 10 modes there to consider, and those modes are sharp, no question about what the frequency is.

We submitted that to NASA, and on the basis of it, we passed the stringent requirements to obtain samples of the lunar soil. We got the soil, and we searched and searched. It was a terrible job. We finally found some approximate spheres and were able to measure the spectra. Here is the result (Reference 3).

(Transparency, Figure 9, from references [6] and [7])

That designation no. in the upper left means the lunar geology program with the number given by the curators of the lunar receiving lab. We found the samples and then we had to label them according to NASA's instructions. You can see the diameters out here in millimeters, less than 1 mm. The samples were not quite spherical and they were rough. They were not smooth like the Moldavites tektites we had made.

As a consequence, we got at most five modes, and they were not peaked, but rounded like this, you see. That is because the measurements were not done on an ideal sphere. The quenching process put little microscopic pits on the surface, and they were not real spheres, but ellipsoids.

So all in all, we were lucky to get a number of these samples that had five modes. But five modes overdetermine the two elastic constants. You do not see five significant figures for Poisson's ratio, as in the tektite case, but you see two (under heading called s). The resonant frequencies were in the seven to thirteen Mhz range. In some of these samples we could not get any peaks at all; in some, unresolved peaks; and in others, no signal. But out of that we got a number out here that gave us the answers.

In the lower part of the figure, we see that the compressional velocities were around 6.4 km/sec, and comparing those results to those for the tektites, they have this particular property, that the compressional velocity is roughly proportional to the density, but the shear velocity is independent of density. That was considered a proof that we were truly measuring the sound velocity of the lunar material. We were measuring the sound velocity of nonporous glasses that had been formed by melting of the rocks at that crater. This convinced lunar seismologists that the lunar crust was made of normal rocks (where $v_P \sim 6$ km/sec, but that the low velocity

they were measuring was entirely due to a texture effect arising from meteoritic impact.

An interesting thing was that the rocks we got out of the same crater did not have normal compressional velocity (6 km/sec). They had a compressional velocity of 2 km a second, and that was really very strange [8], and it took a long time to find out why. I think for people in physical acoustics the answer is very interesting. It came out after many experiments in a big hard vacuum, that cracks made by impact in a hard vacuum, in the absence of atmosphere, are mathematically perfect cracks. They do not have the water that we find everywhere on the Earth to fill in between crack surfaces.

Because of the absence of the water film, then, the acoustic impedance is quite different across the crack and the sound wave will not go across the crack. So what we were measuring with our ultrasonics on the lunar rocks (at 2 km/sec) was the travel time as the wave found its way around all these cracks, giving a large travel time and a false reading of 2 km/sec instead of 6 km/sec, as seen in the glasses.

I do not want to go deeply into the acoustic measurements of the rocks, because this lecture is not about ultrasonics, but this experiment proved to the community interested in lunar science that there was something different about the meteorite-impacted rocks on the surface of the moon from what we have on Earth. The answer is simply that the complete absence of water impedes the propagation of sound waves in the surface rocks of the moon.

Q: I have a question about a statement you made earlier. You said because the glass is not round, that broadens the peak and somehow increases the dissipation?

A: Don Isaak brought some marbles to show you. You can compare the peak of the marbles to the peak of the ball bearing and you can see what we are talking about. What is different about them? The Q of the glass is big, and the marble is not perfectly round and it has roughness on the surface, so you do not measure an ideal mode.

Q: Is there any reason why you did not make them smooth perfect spheres?

A: We did not know how to make them -- these lunar specimens were rather small (0.4 mm). We know how to make spheres of larger radii spheroids, but we were not able to do that to the lunar spheroids. Besides, we could not get permission, after we tried twice unsuccessfully. NASA would not let us try any more. They thought these samples were too precious.

Q: How many spheres did you find?

A: We found 20, I think -- the list is on the top of the transparency.

Ultimately NASA took our spherical glass samples back, and they are still registered in the Lunar Receiving Laboratory in Houston. We didn't get to keep them. I think that this an example of showing that there still is room for work in spherical resonance.

In the next transparency, I show the resonance spectrum of the Earth.

(Transparency, Figure 10)

If you take the equations derived by Love and ask, what would be the period of the lowest mode of the Earth?, it is about one hour (we actually find 1.1 hour). In the spectroscopy of the Earth we find a spectrum that at first looks like the one we found for Moldavite, tektite, but it is not. There are extra modes, and there is some shift in frequency. Why are they shifted? Because the Earth is not homogeneous. It has got layers in it, and there are jumps in density between the layers. It is highly compressed due to the pressure, and it rotates. So seismologists have amongst their group people who specialize in finding out about the structure of the Earth in terms of the shift of these modes from the homogeneous sphere.

Two important effects that geophysicists have been seeking are the effect of the Coriolis force, and the centrifugal force on the internal structure of the Earth. These two forces should cause a splitting of certain modes. UCLA geophysicists measured the vibrational spectrum of the Earth at the South Pole, where both effects must vanish. In the South Pole spectrum, the two splits are closed: that was due to the Coriolis force of the Earth and its rotation. Once the particular modes are absolutely identified, then the magnitude of the shift is correlated to internal structure.

They calculated the mode splitting and supported UCLA's experiment on Earth's normal modes at the South Pole to see if it was true, and it was true. So, using this kind of idea, the science of seismology, has developed to the point where they can actually measure the Coriolis force and its effect on the vibrational spectrum of the Earth.

You get this kind of spectrum whenever you have a large earthquake, because that is equivalent to kicking the soccer ball. The thing starts to vibrate and it goes on for several years when they get a spectrum like that. Seismologists are always waiting for the Richter scale 8 earthquake, 8 and above, because that is enough to activate these modes of the Earth.

Q: Was there was one of those centered under South

America?

A: It does not matter where you have it on the Earth, you can have it any place, but it must be large.

Q: This was a once-in-two-decades kind of event.

A: Yes, it was. I believe you are referring to the famous Chilean earthquake. Seismologists argue about how deep it has to be, and how large, to kick the Earth hard enough to make it start to ring and goes through these long period vibrations. Every time there is an earthquake like that, they can get some new information by these specially located long-period-instrument seismometers. The first mode was finally discovered. It is about an hour for the period.

Q: What is the Q for the Earth?

A: Two thousand. It is as though the Earth were made of steel. The crust does not matter because the mass of the crust is relatively small. The crust has got a low Q, but the rest of the Earth, consisting of most of the mass, has got a high Q. Lord Rayleigh made the calculation of the first few fundamental modes of the Earth, assuming that it was steel and he got close to the correct answer. It is as though the Earth were made of steel, but without self-compression (due to high pressure). In other words, suppose you had a homogeneous sphere as big as the Earth, but did not have a gravity effect, and it did not self-compress, and was made of steel, you would get about a one-hour period for the fundamental mode.

Q: If it has a period of an hour at a Q of 1000, it has got to ring for something like 60 days.

A: It rings for years. Seismologists can see the modes for two or three years afterwards.

Before I go into the theoretical aspects of RUS, I would like to present a slide of the typical RPR setup in our laboratory.

(Transparency, Figure 11)

On the lower left is a representation of the parallelepiped held lightly between two transducers, with a small load (several grams) by means of a chemical balance. In many of our experiments, we actually use a chemical balance. On the right side is shown the dispersion of several modes as the load is reduced, and our extrapolation of the resonance frequency to zero load. In general, modes that are strongly related to shear or torsion will exhibit dispersion, while modes that are strongly related to dilation will not. In some studies, we carefully account for this dispersion; in other studies, where accuracy is not so crucial, or where the

frequency data are degraded by other, less accurate, data (such as specific heat), we simply use the frequency measured at one load (typically four grams). The upper part of the diagram shows the electronic components and their relation to the resonating solid.

Now let us proceed to discuss the inverse calculation of the spectrum. In the theory, you have to choose the functional relations that describe surfaces and edges. These are called basis functions. They may be polynomials, or a long series of LeGendre polynomials, or a long series of spherical harmonics, or a long series of Bessel functions. They could even be the product of power functions.

Then you have to set up your differential equations of motion involving displacement. The displacement, u,v,w, has to be expanded as a sum (out to some large number), where each term will involve the appropriate degree of basis functions. If the x dimension of the solid is a_1, a_2 for y, and a_3 for z, then the displacement has a large number of terms. For examples, each term is of the form $P\ell(x/a_1)\ Pm(y/a_2)\ Pz(z/a_3)$, where ℓ, m, and n are the degree of order of the Legendre polynomial. How many terms are required? Well, that depends on how many are needed to define the surface of the parallelepiped under measurement (including edge and corners). The lower the crystallographic symmetry, the more terms are required. So when the equations of motion are defined this way, you find long equations with many terms.

These will be set up as a long vector equal to a large square matrix multiplied by another long vector. The eigenvalues of the matrix give us the normal modes. This problem proved intractable until a student of mine, Harry Demarest, in a flash of genius, decided that the proper approach was not to use equations of motion, but to use the Lagrangian for deformation. The normal modes are solutions of the Lagrangian stationary points.

Now the Lagrangian is made of terms involving the difference between the kinetic energy and the potential energy. The kinetic energy will involve displacement and the appropriate elastic constants. So all the displacements will involve a long series involving the basis functions. Many of these terms will be zero for symmetry reasons. This is where group theory comes in. Given a certain crystallographic symmetry, we can reduce by group theory the number of terms, and further come up with the number of independent modes of first order (2 for isotropy, 3 for cubic, 9 for orthorhombic, etc.). All other modes will be overtones of the first order modes.

Once we have solved the symmetry problem, we can identify each and every term in the matrix relating the Lagrangian and

the displacement. It is a large square matrix, and for very simple symmetries the degree of the matrix is closely approximated by 3xn, where n is the number of terms in the basic expansion (and the 3 comes from three dimensions). For an expansion out to 20 terms of the basis function (e.g., out to the 20th order Legendre term), the matrix is 60x60.

Having defined the matrix, the next problem is to find the eigenvalues where the derivative of the Lagrangian is zero. This is a formidable problem, but Harry Demarest (1969) [9] found the best computational solution, which has since been used by most experimentalists and theorists in RUS. He found that with big computing power, you can use the Rayleigh-Ritz method of approximation to find the zeros of the matrix; that is, the normal modes.

The bigger the matrix, the more computing time it takes to solve for the roots. So up until recently, we have had to be very cost conscious and have fought to hold down the size of the matrix. That situation has changed due to a new supercomputer, so the size of the matrix is no longer a big cost issue, but it is of interest to see how big a matrix has to be in order to get a stable solution. We see this in the next transparency for orthorhombic symmetry.

(Transparency, Figure 12)

This is the solution for the first 15 modes of an orthorhombic crystal (Mg_2SiO_4). To get stability of the high frequency modes, we need a 60x60 matrix. But to get stability of the lower modes, we need a 120x120 matrix.

Now this theory is a simplified sketch of complicated mathematical physics, so I am going to repeat what I said again, using the very slides that John Sarrao used, 2 years ago, in his final presentation to his Ph.D. advisory committee at UCLA [10] (He had no trouble passing his thesis). Transparency 13 is what John called his theoretical underpinnings.

(Transparency, Figure 13)

(Note: the instructor went over each entry on the transparency).

The question of how the displacements are expanded in terms of the basis functions is given in Transparency 14:

(Transparencies, Figure 14)

(Note: the instructor went over each entry on the transparency).

In the next transparency, we see how the normal modes deform a solid (greatly exaggerated):

(Transparency, Figure 15).

In the last of Sarrao's slides,

(Transparency, Figure 16),

we see his summary.

In John Sarrao's thesis [10], he used RUS to find the normal modes on a high-T superconductor, lanthium copper oxide ($La_{1.86}Sr_{0.14}CuO_4$), which is an orthorhombic crystal. Using the RPR technique, he measured all nine elastic constants of this crystal down to the superconducting temperature, which was 53°, and found the behavior of all C_{ij} approaching the superconducting transition.

He found that C_{66} went through a strong plunge near 53 K, indicating a probable vanishing of this elastic constant and a soft phonon mode, but the other elastic constants were largely unaffected. He verified the answer that the theoreticians were looking for. Indeed, C_{66} is closely associated with the high T transition.

I would like to add a dimension of human interest here concerning my graduate student, Harry Demarest, who today is acknowledged by theorists and experimentalists alike as the one who cracked open the solution of resonance modes of crystalline solids by his insight into defining the normal mode as the stationary points of the Lagrangian, and by using the Rayleigh-Ritz method to solve for the stationary points of the Lagrangian in a large matrix. You remember when I spoke of Soga and my vibrating spheres of the lunar soil. That was at UCLA in 1966, and Harry Demarest needed a thesis, so we said, "Why don't you try to do this for a cube of an isotropic material, Harry?" And, to my amazement, he did, and he also did it for a parallelepiped of a cubic solid. He had the insight for a fantastic advance in theory and the ability to compute it out. Everybody uses his methods today, and his paper [9] is considered a landmark.

But there is a sad part about this. Harry is a very bright guy, and he got his bachelor's degree from Reed College before he came to my laboratory at Columbia University. Now Reed requires each senior to submit a thesis for graduation. Two years had gone by at my lab, and just as he was finishing this great work, Reed College wrote the registrar at Columbia saying, "Demarest has not put in his bachelor's thesis, which is required for graduation. We are going to give him six weeks to submit a thesis, and if he does not, we will withdraw his credits." To have his credits withdrawn would have washed him

out as a student at Columbia.

So here he had this beautiful research result, and everybody was going ape over it, but he had to give it to Reed College as his senior's thesis to stay at Columbia. (So it is not a Ph.D. thesis; it is a senior's thesis at Reed).

Well, we had to find something different for Harry to do, but it never reached the same pinnacle. His brilliant research work was published in The Journal of the Acoustical Society of America [9], and everybody refers to it. While subsequent scientists improved his math, they have not replaced his method or his insight. Where Harry learned how to find the stationary points of a Lagrangian, I do not know, but he certainly did not learn it from me.

While Harry Demarest was doing his theory, and Nahiro Soga was resonating glass spheres from the moon, there was a postdoc from Nagoya University who was working on ultrasonic determination of elastic constants [11, 12]. His name is Mineo Kumazawa, and he later became a professor at Tokyo University. He was very impressed with Harry's work. When he went back to Nagoya, he and his student at Nagoya, I. Ohno, advanced Harry's method up to orthorhombic symmetry [13]. Ohno later came to my UCLA lab to extend Harry's method to triclinic symmetry [14]. Ohno's work is used in our RPR work today. Recently, Ohno has extended his work to include piezoelectric crystals [15] and is reported to be working on monoclinic symmetry.

In summary, now, using the theoretical method to find a set of modes, you ask the computer to find a set of modes by choosing a set of elastic constants, and test the results against the experimental modes. Of course, there is a difference. The program changes the input of elastic constants systematically until the calculated frequency spectrum is close to the experimental spectrum already measured. When this has been done, you have determined the measured set of elastic constants.

I want to comment that the advance in experimental technique using Ohno's theory was first done in Japan by Sumino [16,17] and then by Sumino at UCLA [18], and then by Suzuki [19] at UCLA, and also by Goto [20,21] at UCLA. Don Isaak [22] learned the experimental technique of RPR by working with these scientists from Nagoya.

About five years ago, Al Migliori at Los Alamos independently reinvented the RPR technique of measuring elastic constants of small parallelepipeds. He perfected this technique further to use as a measurement of the elastic constants of high T superconductors [23]. He has put a tremendous spurt into the measuring technique ability and is responsible for making the Quatro device available to us all.

This is the end of the first lecture, and now it is time for a break. Then I am going to turn the time afterwards over to my friend, Don Isaak, and he is going to talk more about the experiments in our laboratory at UCLA.

NOTE: The first 15 minutes or so of Dr. Isaak's presentation were not recorded because of crossed signals as to whether he would be speaking or demonstrating.

DR. ISAAK: I would like to show a few figures about our high temperature experiments. The first is a schematic of our basic experiments at UCLA, which we have demonstrated here last night.

(Transparency, Figure 17)

The next figure concerns the high T experiment mechanism using buffer rods, which separate the transducer from the heated zone where the specimen is located. Buffer rods are also used in ultrasonics for the same purpose. They are made of Al_2O_3.

(Transparency, Figure 18)

(Note: the lecturer went over the details of the drawing in the figure).

If the specimen has any iron in it, we have to protect it from oxidation. That is, we control the partial pressure of oxygen at the surface. In chemical terms, we control the fugacity. The next figure shows the diagram of this device. The solid lines show the buffer rods holding the specimen in place.

(Transparency, Figure 19)

(Note: the lecturer explained the parts of the drawing in detail).

We control the fugacity by keeping a flow of gas, which is a mixture of CO_2 and CO, moving past the specimen. This is automatically controlled, and here is the diagram of this system.

(Transparency, Figure 20)

(Note: the lecturer pointed out in detail how oxygen is prevented from entering the sample at high T).

In the following diagram, I show our complete run at high T for Mg_2SiO_4, forsterite. The error bars are less than the height of the size of the dots.

(Transparency, Figure 21)

At about 800 K, you can see modes crossing. In some cases they cross in a very strange way, and I do not fully understand the reason why. Maybe somebody in this group understands that and could offer some suggestions.

(Transparency, Figure 22)

Now we see the crossing is a kind of repelling. The open dots, B_3U_1, are curved upward on the right, and the closed dots are curved downward. We regard this as mysterious because for most of the crossings, there is no repelling, as in MgO (see next transparency).

(Transparency, Figure 23)

The group in Los Alamos who are doing this work at low temperatures see similar cases in which they get some mode crossing with repelling. There is some kind of a coupling between these special modes that is interesting.

Q: All it takes is the slightest amount of coupling to cause that behavior.

A: Yes, some modes do that and some do not.

DR. GREEN: Do you have some fair amount of evidence for that or is the computer telling you that?

A: This is an experiment.

DR. GREEN: But the data are coming from a computer, right? I have cases where the computer makes me believe the modes are changing across, and it is not true.

A: I will tell you that these data were taken about five or six years ago; this was done without a computer, just manually scanning and looking at those modes.

Q: How do you distinguish modes other than just by a frequency calculation?

A: How do I know they crossed? In this case here, as Orson said, it is a problem where you overdetermine the elastic constants, so we are measuring something like 40 modes, so I take the other 38 and they tell me that B_1U_1 is up here at that temperature. The mode crossing for forsterite involves shear

modes, by the way.

I would now like to change the topic to the problem of mode splitting.

(Transparency, Figure 24)

Here is a graph showing the spectra of a polycrystalline sphere of MgO [24], and it is a good example of mode splitting. This torsional mode here -- marked as $_0T_3$ -- is supposed to split into seven peaks. It is seven fold degenerate if the specimen is perfect. This $_0T_2$ mode below is splitting into five peaks, because it is five fold degenerate, and so on. This is polycrystalline periclase (MgO), and our interpretation is that this kind of mode splitting is due to the fact that the specimen is not exactly spherical. This work was done by our colleague, Dr. Hitoshi Oda, from Okayama University in Japan, using spherical resonance. The mode splitting looks really gross, but you have got to look at the scale here -- the range in frequency is only a few kHz on the plot.

The next graph shows you the kinds of problems that we had last year when trying to measure the elastic properties of a polycrystalline sphere of iron to very high temperatures. This is another case of mode splitting.

(Transparency, Figure 25)

We had trouble even at room temperature making sense out of the measured spectra. In a pure iron polycrystalline sphere, we could not make the theoretical spectrum agree with the measured spectrum. We calculated these modes, as shown by the arrows at the bottom. Instead of a clean, sharp spectrum, we measured a spectrum with bunches of modes that tended to overlap (see top of graph). This is something we cannot quite understand. I just want to show you the kinds of problems that mode splitting can produce.

What Oda found was, when they do split by a reasonable amount, and he takes the average value, he gets very good agreement in terms of the theory and also in terms of cross-checks with his own data. But that splitting I showed you on the previous slide for polycrytalline MgO is a lot less than we found for polycrystalline iron.

To make a long story short, what we finally deduced was that for polycrystalline iron we were having problems with crystalline grain sizes in the specimen.

You can have mode splitting because you do not have a perfect sphere, but you can also have problems if it is not isotropic as well. So most likely our grain size distribution was causing the specimen to appear non-isotropic, with the

complications in the mode spectra that was introduced.

Q: Instead of just averaging frequencies, could you somehow weight them according to their strength?

A: Yes, maybe. I, myself, have done only a little bit of work on the spheres -- most of my work is with parallelepipeds. But I could probably respond to the question: it is hard to weight them according to strengths. Remember, for example, that the amplitudes of the modes vary with their orientation to the transducers. In fact, some of those modes I just showed you of Oda's work I reproduced in the lab last week and the spectra look different in every orientation, but the modal frequencies do not shift, so identifying according to intensity is more trouble than it is worth.

(Transparency, Figure 26)

I now want to show you our results on a single crystal of pure iron [25] at high temperature. The sample is not spherical; it is a single-crystal of iron that has been cut into a rectangular parallelpiped. This is showing the bulk elastic modulus as a function of temperature. One of the problems with this iron experiment is that we kept losing the modes. As we heated it up, the modes became very difficult to see. I had one run out to a little bit beyond 900°, where we had seven modes left that we are able to see. We usually like to see 20 or 30, just to make sure of a good redundancy check.

Here at 690 K is a 17 mode result, but at higher T, I have only seven modes, in terms of the range of overlap of the 7- and 17-mode data reduction, there is pretty good agreement, so I do not have to worry about my results down in here due to the fact that I have a sparsity of modes. The modulus-temperature plot is pretty well defined.

(Transparency, Figure 27)

Here is the equivalent data for the shear modulus of iron.

Q: How do the values that you recovered compare to published values?

A: Great. I will show you right now the data for the C_{11} and shear modulus in comparison with other ultrasonic data.

(Transparency, Figure 28; Transparency, Figure 29)

I did not mean that to be too glib. In some of the silicates and oxides that we have done, in all seriousness, we get very good agreement. But iron is a more problematic type of material. Here I show the Curie temperature and the bcc-fcc phase transition. I think we get pretty good agreement up in

here. These are old ultrasonic data by Leese and Lord here, with these crosses -- it is a little bit low.

Back in 1968, Leese and Lord did not really publish error bars in their data, or else they published real small error bars.

(Laughter)

The error bars that they published include only some of the errors of their experiment. But anyway, you can see the kind of agreement we get. We consider it to be very good. This would be kind of typical in terms of the agreement.

(Transparency, Figure 30)

Let me show you the bulk modulus here in comparison with other experiments. What you just saw with the C_{11} modulus is very similar in terms of how it deviates. Even though we get fairly good agreement in the bulk modulus, my data actually tend to systematically deviate from some earlier data of Dever. That is a problem I do not understand. These are relatively new data and I do not quite understand the reason why.

(Transparency, Figure 31)

When you look at the shear modulus, C_{44}, in comparison with Dever's data, you do not see that systematic deviation. It is only in the compressional-type moduli. Here is the slide for the isotropic shear modulus, showing all the data, including some in the 1 Hz region.

(Transparency, Figure 32)

I think that is one significant result we are getting; we do not understand this difference for the compressional moduli in our work and that of others. I did this data run six or seven times, actually, up to 850 K and 900 K, and it is reproducible and all that kind of stuff. I am not quite sure of the reason for the difference between ours and Dever's. If anybody knows anything about it -- I just have not had time to look at it, but I am wondering out loud if there is some kind of problem with magnetic-acoustic interactions. If anybody has any kind of expertise in this, I would like to talk with you afterwards or even now.

Q: It is just a coil that you are pushing current through, right?

A: Yes.

Q: Is it two counterwound coils, so that you are canceling the magnetic field in the center of the two coils?

A: No, it is one coil.

Q: You could use a counterwound coil. You wind one coil this way and one coil this way, hopefully trying to cancel the field in the center. Can you do something like that?

A: Sure can -- with much effort, but -- I have to admit I never thought of that.

Q: Is there any dependence on the applied field?

A: You are thinking it might be coming from a magnetic coupling. That is something to check out, yes.

Q: Are you using A.C. or D.C?

A: It is A.C. It is line frequency.

Q: Sixty Hz, right?

A: Yes, it is just coming out of the line into transformers and you vary it that way.

Q: Put up the C_{11} data that you just had, please.

A: There is C_S; that is not a compressional mode.

Q: The steps actually do not quite appear here.

A: What do you mean, "steps?"

Q: Put the other one on.

(Transparency, same as Figure 28)

There you can kind of see little steps as you increase the temperature.

Q: Are those the places where you turned up the heater current?

A: No, I turn up the heater current at every point. Basically, what we do is, we come to a temperature and sit there, let it equilibrate, and then scan.

Q: (Inaudible.)

A: There may be. I do not have it, but if you plot the C_{12} modulus, which is really the difference between K_S and C_{12}, it scatters around pretty much. That is usually the case.

Q: Could you turn off the heater while you do the

measurement and let it just sort of cool down a little bit?

A: That is a good idea. It will drop pretty quick here at these temperatures, but it might be worth a try.

Q: Do you scan both directions in temperature, up and down?

A: Usually just going up.

Q: I was just wondering if you see any hysteresis if you go back through those regions?

A: No, because I do go back. I do not systematically, but a point here and there, as it is cooling down. Not every run, but usually you do not see enough hysteresis to make a difference there.

ISAAK: I was wanting to know if anybody knows whether or not the speed of sound is going to be affected by the magnetic domains in the samples? Some of the professors in the back might know. I do not.

DR LEVY: They certainly are. Domains moving around will change the speed of sound and the attenuation. Now, I do not know by how much, but they certainly will change it.

DR. ANDERSON: I think what is important to note, also, here is rate of change of the bulk modulus with temperature is very high. This is unusually high compared to the transition elements and for anything else we have ever measured. It is really dropping fast at the higher temperatures. This is another reason I think maybe there is some magnetic coupling of the elastic constants, but we do not know how to couple them with phases.

DR. LEVY: That is why you need pulse ultrasonics.

(Laughter)

Q: Are you planning on seeing how these velocities change with pressure?

A: Not with the iron, I do not think so.

Q: That would give you some information about the higher order elastic constants, seeing how that Lossy changes with pressure.

A: Actually, I might just want to move over to that area now.

My last little section of this presentation is that I want to emphasize that we are working on the pressure dependence of

these elastic moduli. They have been measured by different methods already, some of them have, but never by resonance method using single crystals and so on.

For us to embark on this type of work we had to ask ourselves are we going to be able to measure changes in frequency as a function of pressure, is it a measurable quantity? We did a quick calculation, which I want to show you, basically just a calculus with a little bit of thermodynamics included. This here as an identity is that the change in the frequency as a function of temperature at constant volume can be expressed in terms of the pressure and temperature derivatives at constant temperature and constant pressure, with this coefficient of one over αK_T, where α is the thermal expansion and K_T is the isothermal bulk modulus.

(Transparency, Figure 33)

To make a long story short, in this expression here, in a quasiharmonic-type approximation you would say this quantity, $(\partial f/\partial T)_V$, is zero. Oftentimes the effect of temperature is simply just to change the volume and that changes the physical properties, but a change of f with T a constant volume is negligible.

But if you think of this derivative here as something very small, because you are allowing temperature change but holding it at constant volume, then this right-hand side here would equal zero, in which case we would have

(Transparency, Figure 34)

-- in which case we could guess at rearranging that last equation and moving this term over here to the left side. I might be missing a minus sign -- yes, this would be a minus sign here.

Anyway, we have an estimate of how the pressure affects frequencies here if we know αK_T, and the temperature derivative of the frequencies. That gives us a chance at least to estimate what kind of effect the pressures would have on the frequencies. Again, we have measured this well for a lot of materials, αK_T is basically known for these types of materials.

(Transparency, Figure 35)

Then just doing a quick calculation, again, to get an order-of-magnitude idea of the effect, we are able to calculate there the pressure effect on the frequencies in the last column of the figure, knowing the temperature effect on the frequencies and αK_T. This example here is for sodium chloride.

We might add that these predicted $(\partial f/\partial P)_T$ lead to a value of $(\partial M/\partial P)_T$, where M is a modulus, that agree well with actual measurements made by acoustics for NaCl under pressure.

This number here depends not only on this, but also on αK_T, so in setting up our first experiment we are starting with materials that have a low value of αK_T, because the lower αK_T is, the bigger the pressure effects will be on the frequency, and that is what we use to guide us in constructing our apparatus and deciding how to proceed.

But that does not obviate the fact that in this experiment we are going to need to control temperature very, very precisely, because if we have effects that are relatively insensitive to pressure, we have to worry considerably about the kinds of temperatures -- I will not get into the numbers there, but I show you how we did that and the fact that we think we can measure those effects. Here is an overall description of our experiment.

(Transparency, Figure 36)

What I am going to do over here is demonstrate a spectrum or two on the adjoining equipment and show you the old and the new way of doing experiments with Quatro.

DR. ANDERSON: There have been some questions during the recess about the vibrational mode structure of the Earth. So I would like to return to Figure 10, and open up the floor to questions.

The theorists said that mode splitting would be due to rotation and Coriolis force, mostly Coriolis force, and that turned out to be the case.

It is not apparent to the eye here (you have to look at it very closely), but that was a very important experiment, and we had to keep that long-term instrument alive -- this was a UCLA project -- at the South Pole for several years, waiting for our earthquake.

Q: What are you actually measuring? It is not a compressional wave?

A: It is a gravimeter. You are measuring the motion of the surface of the Earth this way or that way at the point where the gravimeter lay. We were measuring the x,y, and z components of surface acceleration.

Q: What sort of displacement do you actually get in absolute numbers?

A: Very small, so we need a superconducting gravimeter.

Q: What is the order of the displacements, meters, centimeters, kilometers?

A: No, microns with maximum displacement a fraction of a millimeter. It is not as big as a tide.

Q: Why is Coriolis effect strongest?

A: I do not know. The theoreticians say that is, so we had to go to a place where the Coriolis effect is vanishingly small. If you were trying to determine the absence of a Coriolis effect, you would go to the South Pole. You could measure it at the North Pole, except there is not a continent at the North Pole -- there is moving ice, so that is not a good place to measure small components of gravity.

That is an ongoing experiment even now, an example of spheres, spherical resonance.

Q: I could not tell what you were pointing at.

A: This is the spectrum taken by the instrument at UCLA. This is its sister instrument at the South Pole for the same earthquake.

Q: What is the difference?

A: The difference is mode splitting due to Coriolis force.

Q: We cannot see the mode splitting, so will you point to it for us?

A: I will be glad to lend you the slide and you can see it with a reading glass of some sort. But I wanted to show you that as an example of how important spherical resonance is to seismology today.

Now I would like to return to the main theme. I will make a comment on this demonstration over on the side. The reason that we have gone to the automatic system from our old system is that at high temperature we have temperature drift in our lab; we cannot escape it. We cannot keep things to an absolute temperature level at 1800° for two hours, because there is thermal drift, there is drift in the line voltage, for one thing, in the laboratory, and then there is air-conditioning and all kinds of effects that make it very hard to keep a measurement to a fraction of a degree at high temperatures.

The answer is -- and you could see, as Don turned that crank, you would probably be spending at least an hour, and maybe two, getting the whole spectrum, if you measured each

mode manually. But here, with this Quatro machine, we get the spectrum in less than a minute, the whole spectrum. During that time we know the temperature very accurately and we know that it has not drifted very much, so that enables us to define the temperature with more accuracy.

For numbers, a tenth of a degree -- we have to have things within a fraction of a degree in order to protect the fourth significant figure in frequency. We can report five significant figures, but if the temperature has drifted a degree, then that sensitivity is lost. So that is the reason why for these applications we need the automatic system. We have got to get in and out fast, get the spectrum and be done with it.

This machine was designed by a commercial outfit with something different in mind, and I will show you what that is. Suppose this is the spectrum of a ball bearing, over here on top. You put that on top of your scope -- this is from a good ball bearing, a perfect one. If you put another ball bearing in that you manufactured and it has got a flaw, some of those modes will be missing. It will be the orange line and you immediately can reject it. So within a fraction of a minute you can find the bad ball bearing.

It was for that purpose that this Quatro device was made: to find immediately some flawed spectrum.

Q: In the case where we were looking at the ball bearings, I guess the bad ball bearing, would that be due to a crack?

A: Typically a crack inside.

Q: It would not be due to an oblate shape.

A: We could see that, too, because then we would have mode splitting due to lack of sphericity. There are two modes, one for each axis. This is not a problem in materials, but I am just pointing out that the Quatro machine was made for that purpose, and it does very well. It gives a spectrum very fast, but it does not do other things we would like.

One of the reasons it does well is because all the electronics and the solid state are on the inside and the losses due to the system are largely gone. Especially, I think, the wires running back and forth were giving us a background noise. No they are gone in that solid state system, so we are getting, we think, the background noise down to a low level. But it is not completely right for us, because it does not have the software we would like. We would like it to tell us more about the individual modes.

Right now it is printing on top of each mode what the mode

number is and that is great help to us. We do not have to peer at the frequency synthesizer and write it down, which is the reason it takes two hours when we do it by hand, so it is a big help, but it does not give us all the information we would like about the mode. It is not made as a functional and versatile research system, and we have got to find a system that is better. But its performance is telling us that the fundamental idea is correct.

At low temperature, in superconducting work, you would want to have the capability of taking a complete spectrum very fast. You do not want to mess around and have to take several parts of the spectrum independently and allow your specimen to get out of equilibrium with its temperature environment. You do not want to be measuring the beginning of a reaction for the first mode and the end of the reaction at the second mode. You want to get in and out fast. So that is the state of the art.

Q: I have a suggestion. I notice, even when you are computing, it takes at least a few seconds to do this thing. One thing you could do, if you have the right equipment, you can sweep through in a fraction of a second.

A: Well, there is a menu there. You can make it sweep fast. You have got lots of choice.

Q: Another thing in the built-in computers, like this computer or even a faster one, it still takes some time to do that, but you can use a digital scope to record the response of that, because you usually have several channels. You can do the FTT or whatever or you can just record it on a disk and later on analyze it.

A: Yes, we could also put in an XY recorder or something else. There are all kinds of things you would like. So this is not as flexible, and so far it is restricted to one purpose only. But it costs $30,000 to do what is done. We are looking for improvements on this. Al Migliori says Quatro is learning how to make the improvements UCLA needs.

DR. GARRETT: I do not think you can just arbitrarily increase the speed, but one thing that you can do and one thing that we do for monitoring processes is to pick a single resonance and then create a phase-locked loop, so you track that one resonance. That gives you only one elastic constant per run, but if you know, as in the case of the C_{66} constant for the superconductors, if that is one you are really interested in, then you can track that one resonance and have that as a continuous time record of your elastic behavior.

DR. HARGROVE: One resonance does not give you one elastic constant.

A: Yes, unfortunately, the elastic constants are functions of all the modes, linear functions, hopefully.

I am going to go on now with some other applications. I am going to give some more slides from John Sarrao's thesis, which was on superconductivity of several materials.

(Transparency, Figure 37)

Here is Sarrao's cryostat, and there is his experiment. It is a classical doer, but you can see it is basically the same problem we have in high temperature as far as temperature control is concerned. We have got to have some extension from the place where the temperature is controlled up to where we measure.

(Transparency, Figure 38)

The next question is the sample itself. I forgot to mention this morning the champion of us in small size. That belongs to Jay, I think [26]. What was the thickness of your material that you measured, Jay?

DR. MAYNARD: Nine microns.

DR. ANDERSON: And the sample itself?

DR. MAYNARD: Fifty or 100 μm.

DR. ANDERSON: So a 100-μg sample in a few-micron transducer has got to take the prize here in smallness.

Essentially what Sarrao did is something like you did, Jay, which is to use piezoelectric plastic films, PVDF. You put a strip across there of this transducer and that gives you a signal, and Sarrao's sample is, in this case, the superconductor, below.

In this particular design, he had a lithium niobate disk that was very thin, and he had this moving back and forth with the magnetic field. But this was a very small sample, too (I have forgotten the size), very, very small.

(Transparency, Figure 39)

The question that John was trying to find was the mode, the modal frequencies to determine, the elastic mode at the superconducting temperature of his superconductor. These are two of his modes down here near the superconducting temperature, which diverged from each other in the same way we were talking about before. They do not cross, the repel each other.

The next transparency shows the three shear elastic

constants versus T. One of them, C_{66}, is plunging down, at the superconducting temperature.

(Transparency, Figure 40, from [14])

We can get down to the high T superconducting temperatures, and I am sure we could get down very much farther with further applications. The technology for piezoelectrics is here; the standard for small size is here; there is no question but what this is opening up a great field. In this particular case, the individual elastic constants near the superconducting temperature were measured, and the behavior of each elastic constant was found. This got John his Ph.D. thesis.

I would like to attend for a few minutes to high-temperature work that we do at our lab.

(Transparency, Figure 41)

I will repeat what Don has already shown you. We have the transducers down here on buffer rods, which go up and touch the specimen there, and we control the weight down to a few grams, usually 4 gm, after we get it seated (if this took fewer than 4 gm this might fall down). So we are sure that we are not getting very much effect due to the force on the transducer itself.

We put the oven around the sample, and then, if the sample has iron we sheath it with a fugacity control device. If we do not have any iron in our specimen, we do not need that control, although I suspect, if we get above 1800°, we will probably need some atmospheric control for any oxide or silicate.

(Transparency, Figure 42)

This is the complete spectrum of Al_2O_3 versus T, 40 modes (I think there are less than 40 there -- there might be some higher ones) just as Don showed you for Mg_2SiO_4, and the dots are larger than the size of the frequency error. We need that. You can see that because of the mode crossing we have got to really take a lot of points. We have got to take in many, many temperature points; otherwise, we will get lost in the crossings, so that is another reason for automation. We have got to get in and get out fast.

(Transparency, Figure 43)

Here is the spectrum of Al_2O_3, showing relative amplitudes of the modes at high T [14]. We measure the modes, then we calculate them. We assume elastic constants, and the machine program goes over and over, again, and finds the improved elastic constants, which finally yields a match to the

frequency spectrum.

We have to impose a statistical criterion to know when the match is good enough, so the computer can stop. Here is an example of what we find for a match at 3 temperatures.

(Transparency, Figure 44)

Here is the observed and the calculated; the match is fairly good. This is for 296, 1000, and 1825 K. We follow those modes very well. We see the decrease of the mode frequencies at three temperatures, and see the match between theory and experiment. The mode splitting at high T is a little bit more than at ambient. Here these two modes are together, and here they are split apart. It is not perfect but it is good enough.

This technique is not good enough to start from zero. If you do not know what the elastic constants are, you cannot use this technique. The reason is that in this phase space that we are dealing with there are many minima. If the machine finds a minimum that is not the right minimum, it will stay there and will never get out, even though the experimental and theoretical spectra are far apart. So you have got to know the numbers approximately. That is what we consider to be the purpose of ultrasonics, Dr. Levy. It is to give us the first approximation.

(Laughter)

How close you have to be, to begin with, we do not really know, but we have got into trouble a couple of times because we did not guess it right very close and then the program just cannot cope with that, it has got to be reasonably close.

For instance, supposing one of the elastic constants is wrong out of nine for an orthorhombic crystal? The program could probably cope with that. But if two were wrong out of nine, I do not know whether it would cope or not. That is the kind of thing we are up against. This is a process for finding the refined constants, in my opinion. However, Al Migliori is optimistic that he can identify when we are in a false minimum.

You can start in a region where the ultrasonics is good and get the elastic constants there and have the computing method verify them -- this is really what happens. You take Levy's elastic constants, determined from ultrasonics, and you put them in, and, from the results, you say, "Yes, he's good, he's right."

So you may say, "Why use RUS when you get good results with ultrasonics?" The answer is that there is a limit to the temperature at which good ultrasonic work may be done. RUS gets us into a temperature range that he cannot get into with

ultrasonics. That is really how the whole system works. The two techniques are coupled. Every physical acoustics lab should have both techniques. We could not get to high T if Levy or ourselves or someone good in ultrasonics did not do the ultrasonic experiments, or if there were not some good experimental results around.

(Transparency, Figure 45, from [14])

Don pointed out the importance to geophysics, and I would just like to amplify what he said with the notion of isotropic shear velocities. We get the elastic constants in geophysics on single crystals, which would yield C_{ij}, but we then turn around and average them to get polycrystalline constants. Why? Because the Earth gives us isotropic values of elastic constants, through seismology. So to compare our type of data with the data from the interior of the Earth we have got to go to isotropy again.

We have a little program that converts single-crystal data to polycrystalline data. Then we can calculate some results. Here, for example, are the calculated v_P and v_S for corundum from 100 K to 1800 K. That is where we really wanted the data. The Debye temperature is around 1000 K, and you see some curvature below the Debye temperature -- it is curving down -- and it will, before it goes to absolute zero, go in with a zero slope.

But in the upper range, the velocity is virtually linear with T, and this was a very important result. Actually, the frequencies are curved with T; the single crystal constants are virtually straight with T; and the isotropic velocities are almost straight with T [27]. We wanted to know if sound velocities were linear in the temperature range of the deep Earth. It is. Over and over, again, for all solids, it is, so that is an important and significant result. We could tell the theorists, "yes, you can assume it is linear, and this is accurate to the limit given by our measured second derivative with respect to T" -- we can see the change with temperature of the first derivative of the velocities with temperature, so we are getting second-derivative information here. This allows us to extrapolate down further. That is vital information for geophysics.

Q: Why should it be linear?

A: Because it turns out that elastic properties are linear in the quasiharmonic approximation. That is the solid state physics question I want to get to. I was not certain the experimental results would be linear. Even though you can prove that specific heat has anharmonic effects, at the same time the same data are showing that the velocity properties are linear with T. This is something for theorists to chew on. We

might have expected that the v_P vs. T curve was nonlinear because the quasiharmonic approximation would not be good to that high temperature.

(Transparency, Figure 46, from [28])

We now switch from geophysics to physics. Here is the quasiharmonic approximation at the top of the page, in the high-temperature limit. For anharmonicity, the AT^2 term over here, on the right. If you think some anharmonic corrections need to be made. That is the Helmholtz energy. The internal energy is the second equation, and the specific heat is the third equation. It is the Dulong and Petit limit if there is no anharmonicity. Or if there is anharmonic correction, you would add this 2AT term as a first correction.

Our data at the bottom of the transparency reveal this. They reveal that Al_2O_3 is quasiharmonic and did not need any correction up to more than twice the Debye temperature. Mg_2SiO_4, on the other hand, seemed to require a little bit. MgO did not, so MgO is quasiharmonic; Al_2O_3 is quasiharmonic; CaO is quasiharmonic; but the garnets and Mg_2SiO_4 are not quasiharmonic. They required the anharmonic term.

Q: What is the Debye temperature of Mg_2SiO_4?

A: 890 K. All these geophysical solids are 800 to 1100 in the Debye temperature.

Surely at 2 theta you ought to be able to see the anharmonic effects. What this is saying is that some of the specific heat properties are quasiharmonically good up to 2 theta and some are not.

Then we used the same data to calculate the thermal pressure versus T.

(Transparency, Figure 47, from [28])

Here is the quasiharmonic high-temperature limit there on top, as before. And here is U, the internal energy from that. And here is P, pressure, and it has this term, called thermal pressure, and it should be linear in temperature if you have only quasiharmonic effects. You should see a T2 term if there is any anharmonicity. Those same six solids seen in the last slide all turn out to be linear up to 2 theta, even though three of them show anharmonic effects in C_V.

For geophysics the answer was great. It says that it is not necessary to worry about that T^2 term in your equation of state down deep in the Earth. That is the message -- we are sure that is true (The same result turns out for all kinds of

solids: noble gas solids, alkali halides, and alkali metals). But for physics this result leaves us in the lurch. In particular, why is it we need an anharmonic correction for garnets in specific heat but not for the thermal pressure?

I think I know the answer to that, but I just want to pose that problem to you. This is one thing that high-temperature work does, it bears on anharmonic theory. It has given us data for the first time and on lots and lots of materials.

The people in anharmonic theory were a little lazy, I am afraid. They always chose to prove their results by specific heat, but they never went to the other thermodynamic properties to check it out. I am talking about Cottrell; I am talking about all the great names who write books on anharmonic theory.

There is this to be said for it. If you take the Helmholtz energy, F, and you take a derivative with respect to T to get U and then another derivative to get specific heat, you are taking a derivative of that quantity there (in the graph), and you are going to get an AT or an AT^2 or something down below. But if you take the derivative with respect to V in order to get pressure, then the T^2 is carried down, and you are interested in the derivative of A with respect to V. So if the T^2 doesn't appear in the thermal pressure, something is wrong.

So the simple answer is, I expect, those coefficients for the T^2, etc., anharmonic terms in the Helmholtz energy are insensitive to V. If that were the case, then a derivative of the anharmonic term with respect to V would vanish, agreeing with our experimental results. So it may turn out that the coefficients in the anharmonic expansion might be insensitive to V, and any time you take a V derivative, that anharmonic term vanishes. That would explain our results.

However, this is a guess, not theory; we propose an empirical observation, and somebody has to go back in the theory of anharmonicity and straighten the matter out. We are finding this over and over. When we first ran into this, we said maybe it is just true for oxides. But it is not true for just oxides; it is true for alkali halides; it is true for noble gas solids --all Swenson's data take noble gas solids up to high Debye temperatures -- and it is true for alkali metals. We went back and looked at all the data, and it was the same story all over again.

All the metals we examined, except gold [29], seemed to have this particular property -- we could see anharmonicity in thermal pressure in gold but not iron, or copper, or nickel. We could not see it in alkali-metals, lithium (there are a lot of data on that, and we went and looked at that). Apparently there are no anharmonic effects in pressure or any of its

derivatives.

Q: Is there anything in the shock-wave data that has a bearing on anharmonicity with respect to thermal pressure?

A: I spent some months with Bob McQueen at Los Alamos trying to probe into this, and the accuracy of measurement just does not permit us to look for this. What usually happens is people assume that the anharmonicity is there and then derive it as though it is true, but nobody has done a careful experimental investigation to find out if it is true, and so there are problems in the interpretation. I am not a shock-wave physicist. It is something they can cope with, I am sure, but it has yet to be done.

Here is how the thermal pressure varies with T/Θ for a number of solids according to our experiments. The slope of the line depends on the initial density.

(Transparency, Figure 48, from [28])

I would like to now discuss something important to chemical physics. It turns out that αK_T always pops up (where K_T is isothermal bulk), as calculus will show you. (Writes formula on blackboard). From a Maxwell's relationship, $(\partial S/\partial V)_T = (\partial P/\partial T)_V = \alpha K_T$.

So entropy can be calculated as the derivative of αK_T dV, between V_1 and V_2, and that will be entropy of V_1 minus the entropy of V_2. Our experiments are just delightfully tuned to that kind of a calculation, because we measure both α and K_T as a function of T, and do it quite well. So we can calculate entropy as a function of V and T from our measurements [30].

(Transparency, Figure 49, from [30] and [31])

This is an example of what we found. You can calculate the internal energy, too, from the pressure, so you can get, then, internal energy minus TS and you have free energy, or Helmholtz energy. This is now calculated from our results.

(Transparency, Figure 50)

We are looking at the minimum here, which is the condition of crystal stability. The interesting thing about this is that the theorists now are calculating a lot of things from first-principles methods [32,33,34]. What they do is assume the potential and then calculate the free energy, and then one is free to take derivatives of the free energy to get all the physical properties they want, such as entropy, specific heat, enthalpy, pressure, bulk modulus, and thermal expansivity.

This is our experimental check on the accuracy of their assumed potentials. In their calculation of Helmholtz energy, it had better be like we find from experiments, or they are in conflict with experiment. So we are bearing in on the experimental tests for calculations from first-principle calculations. The theorists are not paying any attention to us now, but later on, they will, I am sure. (They currently feel that theory is true, and if experiment doesn't agree, then experiment is wrong).

You see, what they do is they calculate the Helmholtz energy and then they calculate the elastic constants as a function of volume and temperature, which is what we measure. Often they do not agree with our elastic constants. What is happening is they agree on the derivatives of the elastic constants with respect to P and T, but not very well on the elastic constants themselves. That is the state of the art in *ab initio* work.

But with our experiment we have a way to enter chemical physics and condensed matter physics. This all results from our accurate RUS experiments at high T.

Before I summarize, I would like to comment on the RUS experiments and theory now going on in Japan. I Suzuki and H. Oda are now at Okayama University, and the focus of RUS has shifted from Nagoya to Okayama. The Rayleigh-Ritz method found by Demarest [9] has been improved by Mochizuki [35]. The Okayama group has concentrated on developing further the RST (Resonant Spherical Technique) in contrast to the RPR technique, discussed at length here, which is prominent in the United States. Oda et al. [36] have developed a theory for the free oscillations of elastically anisotropic spheres and ellipsoids (with orthorhombic, tetragonal, cubic, or isotropic symmetries). They have calculated how the degenerate eigenfrequencies split due to anisotropy. This was based, in part, on the theory of Visscher et al. [37]. Oda [38] developed the RST theory to the point that it is now possible to compute eigenfrequencies of an elastic anisotropic ellipsoid with arbitrary crystal orientation, so that the directions of the crystal axes as well as the elastic constants may be simultaneously determined from only the observed resonant frequencies. All problems have now been solved in order to make RST competitive with RPR.

I have tried to give you some examples of where we are today in the general field. Don has given you some examples of where we are today with the experiment itself, what our problems are, and we tried to give you historical background, with some references, and we tried to give you the beginning of the ideas as a pedagogical teaching approach, and I hope that we have succeeded.

Let me review the problems of where we are today. We do not know why some polycrystals give us no signal and some give us beautiful signals. We do not know why iron is giving us a lot of trouble. We do not know what things are affecting the modes in these cases. We need to have help on that. We need to get spectra faster and better so we can move the experiment further up in T, and we need to learn to work with smaller and smaller samples.

If we can do that, then I think this field will move on, and I hope it can. One other thing we have to do, we have to get into the pressure measurement, and for that reason we are in Hargrove's program, and we are attempting to proceed with pressure measurements.

Let me repeat the difficulties of that. You can see that if you put pressure on the outside of a body you are violating the conditions that we have to derive the equations, aren't you? We derive these equations by saying the traction is zero across that surface (e.g., P=0 at the surface). That is no longer zero if we apply pressure. But we can put a little pressure on and without violating things too much, a few bars, maybe a hundred bars, and still make the equations work by accounting for a small correction.

The question is, can we get pressure derivatives of the elastic constants with those few bars, 80 to 100 bars, of pressure, and the answer is: yes, we can, if we can take extraordinary efforts to control the temperature. That seems to be the answer and that is where we are and we are starting our experimental program now on that.

The thing you have got to remember is that we are able to get great accuracy in frequency, and if we can measure the pressure with great accuracy, which I believe we can -- we probably can measure the pressure to one part in 10^4 in this experiment -- then we will get the derivatives to one part in 10^3, and that is quite high accuracy. That is probably higher than you can get by having a gross measurement of pressure at a high pressure.

Since these elastic constants (and frequencies) at low pressure are going to obey Hooke's law, elastic constants are going to be proportional to pressure up to quite high pressures. So measuring at 100 bars and getting a precise value is enough. We will get a lot of information out of that. That is a challenge that we have today. I think that is all the challenges I want to talk about. If I put any more challenges on you, you will all go to sleep.

(Laughter)

Thank you very much.

REFERENCES

1. O.L. Anderson, J. Acoust. Soc. Am. **91**, 2245 (1992)
2. A.H. Love, *A Treatise on Mathematical Theory of Elasticity*, Dover, 4th Ed. (1944)
3. Y. Sato and T. Usame, Geophys. Mag. (Japan. Meteorological Agency, Tokyo) **31**, 15 (1962)
4. D.B. Frazer and R.C. LeCraw, Rev. Sci. Instrum, **35**, 1113 (1964)
5. N. Soga and O.L. Anderson, J. Geophys. Res. **72**, 1733 (1967)
6. E. Schreiber, O.L. Anderson, N. Soga, N. Warren, and C. Sholtz, Science **167**, 732 (1970)
7. N. Warren, O.L. Anderson, N. Soga, Geochem. Cosmochem. Acta, **Suppl. 33**, 1877 (1972)
8. E. Schreiber and O.L. Anderson, Science **168**, 1579 (1970)
9. H.H. Demarest, J. Acoust. Soc. Am. **49**, 768 (1969)
10. J. Sarrao, Ph.D. thesis, UCLA (1990)
11. M. Kumazawa and O.L. Anderson, J. Geophys. Res. **74**, 5961 (1969)
12. M. Kumazawa, J. Geophys. Res. **74**, 5973 (1969)
13. I. Ohno, J. Phys. Earth **24**, 355 (1976)
14. I. Ohno, S. Yamamoto, O.L. Anderson, and J. Noda, J. Phys. Chem. Solids **47**, 1103 (1986)
15. I. Ohno, Phys. Chem. Miner. **17**, 371 (1990).
16. Y. Sumino, J. Phys. Earth **27**, 209 (1979)
17. Y. Sumino, O. Nishizawa, T. Goto, J. Ohno, and M. Ozima, J. Phys. Earth **25**, 377 (1977)
18. Y. Sumino, O.L. Anderson, T. Suzuki, Phys. Chem. Miner. **10**, 38 (1983)
19. I. Suzuki, O.L. Anderson, and Y. Sumino, Phys. Chem. Miner. **10**, 125 (1983)
20. T. Goto and O.L. Anderson, J. Sci. Instr. **59**, 1405 (1988)

21. Y. Sumino, O.L. Anderson, and I. Suzuki, Phys. Chem. Miner. **9**, 38 (1983)

22. D.G. Isaak, O.L. Anderson, and T. Goto, Phys. Chem. Min. **16**, 704 (1989)

23. A. Migliori, W.M. Visscher, S. Wong, S.E. Brown, J. Tanaka, H. Kojima, and P.B. Allen, Phys. Rev. Lett. **64**, 2458 (1990)

24. H. Oda, S. Isoda, Y. Inouye, and I. Suzuki, J. Geophys. Res. **99**, 15,517 (1994)

25. D.G. Isaak, and K. Masuda, submitted, J. Geophys. Res. (1994)

26. J.D. Maynard, J. Acoust. Soc. Am. **91**, 1754 (1992)

27. O.L. Anderson and T. Goto, Phys. Earth Planet. Inter. **55**, 241 (1989)

28. O.L. Anderson, D.G. Isaak, and H. Oda, Rev. Geophys. **30**, 57 (1992)

29. O.L. Anderson, D.G. Isaak, and S. Yamamoto, J. Appl. Phys. **65**, 1534 (1989)

30. O.L Anderson and K. Zou, Phys. Chem. Miner. **16**, 463 (1989)

31. O.L. Anderson and K. Zou, J. Phys. Chem. Ref. Data **19**, 69 (1990)

32. R. Hemley and R.G. Gordon, J. Geophys. Res. **70**, 7803 (1985)

33. R.E. Cohen, Geophys. Res. Lett. **14**, 37 (1987)

34. L.L. Boyer, M.J. Mehl, J.L. Feldman, J.R. Hardy, J.W. Flocken, and C.Y. Fong, Phys. Rev. Lett. **54**, 1940 (1985)

35. E. Mochizuki, J. Phys. Earth **35**, 154 (1987)

36. H. Oda, J. Hiro, I. Suzuki, W.M. Visscher and O.L. Anderson, Geophys. J. Int. **118**, 555 (1994)

37. W.M. Visscher, A. Migliori, T.M. Bell, and R.A. Reinert, J. Acoust. Soc. Am. **90**, 2154 (1991)

38. H. Oda (letter), J. Phys. Earth **42**, 97 (1994)

NONLINEAR ACOUSTICS

DR. GARRETT: I would like to thank the organizing committee for running this experiment, for allowing me to teach a course in a subject in which I have not worked in 15 years. The idea is that since I am at a fairly low level, I should be able to communicate well with the students, who are at about the same level.

The way I would like to start this is with a quote from Andrès Larraza, who has been a friend and colleague of mine for some time.

(Transparency) (1)

He says, "An acoustician is merely a timid hydrodynamicist." What I am going to try to do is, in some sense, what Professor Bass did; try to show the importance of having a macroscopic view -- Professor Bass looked at the microscopic connection, but I am going to look at the sort of the macroscopic understanding of acoustical systems. That view must be based on a hydrodynamic picture.

Professor Bass said that you have to have a good balance between the microscopic and the macroscopic. Actually, there is somebody who said it much better than you -- but do not feel bad, it was Albert Einstein. He said the true testing ground of physical theory was thermodynamics, because it is model-independent. So what we are going to do here is look at a model-independent hydrodynamic analysis of nonlinear acoustic effects.

My purpose for wanting to teach this course is that when I was a graduate student I found this stuff to be fascinating and ultimately did my Ph.D. thesis in the area. I was just so intrigued by what happens when you look beyond the linear level that I felt that it would be something that would really be wonderful to expose some of you to who have not seen this.

So one of my motivations is to introduce the students to the simplest concepts and techniques; that is automatically guaranteed, because I can do only the simplest concepts and techniques. If you heard Hank last night, he made a crack, he sent me up to the board but he said beware, because I am a "computation cripple". The reason he said that is when I applied for summer school, you had to list your food preferences and your roommate preferences, and he also said to

list your handicaps, so I put down "computational cripple."

But like any other handicapped person, I have developed techniques to compensate for that, and that is what I mean by simple concepts and techniques. There are many different ways to look at a problem and by looking at it in many different ways I think we get some insight. Although you may not believe me, Feynman has said that.

I think there is also another reason for studying nonlinear acoustics. As this week progresses, there will be several other talks on nonlinear acoustics. This is the introduction to nonlinear acoustics. In some sense, I tend to call this, at least when I am around Denardo and Larraza, the history of nonlinear acoustics. Bruce Denardo and Seth Putterman and others will be lecturing on really recent developments, but I think having some of these basic concepts under your belt will make appreciation of what is going on in these more advanced topics, more contemporary topics, more penetrable.

Finally, I would like to present the strange and the bizarre. The reason for that is, again -- this is like my second point -- if you take a look at a standard result in nonlinear acoustics, if you can go look at another system that behaves differently from that, it gives you a better insight into that nonlinearity, just as nonlinearity gives you a better insight into the linear picture.

My guidance, at least my spiritual guidance, in putting these lectures together has been a quote that I have always enjoyed from Professor Feynman, who said, "One measure of our understanding is the number of different ways in which we can arrive at the same result." So what I plan to do is do the same problem, or the same three problems, over and over, again.

There are two reasons for that. One is that, of course, it will presumably generate a deeper understanding of those problems. The other is I probably know how to do only three problems, so I had to stretch it out some way. This is, you know, a post hoc justification.

(Laughter)

There is a problem with having various ways to attack a problem and that is that, a man with one watch knows the time and a man with two is never sure.

(Transparency) (2)

The lecture is divided into two basic areas. One is what I like to call wave-wave interactions. There we will be talking about the effects that convection and material nonlinearity have on the propagation of a sound wave. We will look at shock-wave formations, which is self-interaction, a wave interacting with itself. We will look at nonlinear wave mixing, where you have two different waves, possibly even moving in different directions, interacting with each other to produce some other effect through a nonlinear reaction.

And I had planned on showing a video on superfluids that ran 17 minutes, but I do not think I have the 17 minutes. I will be talking about some of these effects in superfluids, because they defy the standard classical shock-wave behavior. As I said, it gives you a better insight into what happens in ordinary systems.

The second area I will be talking about is non-zero time-averaged effects. In acoustics you have sound waves which, at the linear level, vary with a sinusoidal time dependence, so the time average of the pressure is zero, the time average of the velocity is zero, the time average of the density change is zero.

When you go to the nonlinear level, those time-averaged properties do not disappear and, therefore, they can lead to acoustic forces and acoustic torques that are, to a large extent, explained simply by Bernoulli's law. I will be deriving a particular acoustic form, a strong form of Bernoulli's law, and applying it to spheres and disks and levitated objects.

Nonlinear acoustics is a very rich field and, to be fair, I have excluded topics that will offend almost everyone. I am not going to talk about nonlinear bubble dynamics. Seth and Anthony will tell you that is the most fundamental thing in nonlinear acoustics, how can you possibly give a lecture and not include that? I am not.

Pierce and Rodgers, no N-waves, no sonic booms. Mark, I am not going to talk about nonlinear aperture theory or parametric arrays. I am not going to talk about wave turbulence. I am going to mention Burgers' equation only twice. Burgers' equation -- that was the second time.

(Laughter)

If Dave Blackstock were here, he would say, "How can you do that?"

I am not going to talk about streaming or solitons. There will be others who will do that.

(Transparency) (3)

This says "genealogy," but what it really means is acknowledgments. Almost all the material I am going to be presenting here was passed to me by my advisor, Isadore Rudnick, and the insights here are not mine, they are his. I just took them. And the stuff that I did not take from Professor Rudnick I took from Professor Putterman. I just wanted to make two points with this viewgraph.

One is that I am representing the material of my teachers and, as a message to the students, one person can make a tremendous amount of difference. What I have done here is put somewhat of a genealogical chart up that shows Professor Rudnick, who hired Seth Putterman and Jay Maynard around the time I was there and educated me and Robert Keolian. Greg and I work together and Greg was at UCLA for some amount of time and was strongly influenced by Professor Rudnick and these gentlemen over here, and Bruce worked for Seth. Anyway, if we were doing sex instead of science, that would be called incest.

The point of the slide is that Professor Rudnick was able to create an environment and present a style of science that really brought a tremendous number of other people to doing good work. One person can make quite a large difference and that person could very well be, and probably is, in this room.

(Transparency) (4)

I did not think that 20 years ago, when I was at the Physical Acoustics Summer School. This is the group picture. I regret to say that at this meeting 20 years ago the only person giving a talk who was at that meeting 20 years ago is me, so 20 years from now, when somebody is going to be showing our picture, which has not been taken yet, one of you out there is going to carry this culture along, or it is going to be lost

The other comment I had was that the old guys who were there -- see, I am almost an old guy now -- seemed much older to me then. This is Warren Mason, D. Sette, and Bruce Lindsey and Professor Rudnick. The guy in the back is Charles Manson.

(Laughter)

With that said, let us do some nonlinear acoustics.

(Transparency) (5)

The place to start, of course, is what is linear acoustics? I believe the two fundamental levers that you have in linear acoustics are superposition and stability. If you have a linear wave theory, the assumption is that the presence of the sound wave in the medium does not perturb the properties of the medium. That is what linearity is. That is why Einsteinian gravitation is a nonlinear theory, because the field is created by the mass and the mass is in the field. It is intrinsically nonlinear theory. It can be linearized.

That is what we have done with acoustics, with hydrodynamics. It is an intrinsically nonlinear theory. We have linearized it. So for enduring the limitations of linearity, the fact that we cannot deal with some of the topics that I am going to show you in nonlinear, we are rewarded by having the principle of superposition.

Therefore, if we take two wave packets of whatever arbitrary ugliness -- this one, say, traveling in this direction and this one traveling in this direction -- when they occupy the same space, they superimpose, they look like hell, but lo and behold, they pass through each other unchanged.

Secondly, there is no effect of that interaction outside the interaction region. None of the disturbance propagates outside, the original waves just pass through one another.

The other thing that you have in linear acoustics is stability. If you put in a sinusoidal waveform of a certain frequency and wavelength, then that is all you get. It may get smaller due to dissipative processes that Professor Bass described, but you do not get new frequencies and you do not get new wave shapes. The waveforms are stable. They have a very boring life cycle. They fade away.

Nonlinear waves have a much more interesting life cycle.

To start discussing the nonlinear problem, what I would like to do is take advantage of the fact that we are in Asilomar and only a few hundred yards from the beach and start talking about surf. That is the deep philosophical question in California: Why is there surf?

(Transparency) (6)

What I am going to do is take a look at the propagation of very idealized surf, shallow water gravity waves, because I think it is a very simple system in which most of the interesting phenomena of shock waves will be exhibited and then we will go through that system and extend it to ideal gases, liquids, quantum fluids, solids, et cetera. But I think it is reasonable that when we are in Rome, we should just do as the Romans.

What I will do is assume a uniform beach, that is, that the water is of uniform depth and we have a small wave propagating on the surface, and the restoring force for that wave is gravity.

This "assume a spherical cow" is the punch line from a joke that some of you actually probably will not find funny. You know that recently, with the employment problems in physics, there has been a lot of articles in Physics Today and AIP about what great lives people have when they leave physics. It is kind of an interesting thing to read.

One guy actually took a job in the dairy industry, he was a theorist, he was being paid by Knudsen's. After a couple of months, the guys who were paying his salary said, "Well, what did you learn?" He said, "Well, I'll tell you, let's set up a blackboard and I'll give you a seminar." He walks to the board and says, "Okay, now, I'd like to tell you about the dairy industry. We'll start by assuming a spherical cow."

(Laughter)

So I am assuming a uniform beach and uniform depth. We will make it very simple, but it will exhibit the proper physics that we need to understand almost everything in nonlinear shock-wave formation.

To describe this system, we will impose a wave that deviates from equilibrium by an amount h_1. (I will get much more fanatic about what the subscripts mean as we proceed). The total height of the fluid column squiggle, in element dx, is h_0, the equilibrium height of the water plus whatever the deviation is due to the wave, h_1.

We can write a continuity equation. We can say that the change in height of that slab at point dx is given by the divergence of the flow through the slab; that is, if the flow is uniform through here, the height stays uniform. But if more liquid flows in than flows out, if there is a gradient in that x velocity, since we are assuming incompressible flow, it has to go up or go down, depending on the sign of the gradient. So this is a continuity equation; a mass conservation equation.

We can also apply Newton's law. The acceleration of a fluid particle, mass times acceleration, has to be due to the forces on that particle due to the gradient of the pressure. Since this is a one-dimensional system, we need only the x component of this equation, so dv in the x direction, dt, is one over $\frac{1}{\rho}\frac{\partial}{\partial x}\rho g\varsigma$. The pressure is $\rho g\, \varsigma$. Since the time-dependent or spatially dependent part is the h_1, if we take that derivative we come up with that answer.

So we have two equations that describe the motion of the fluid at the linear level, linear meaning that I do not have v^2's in here -- all of my variables are only to first power.

(Transparency) (7)

If we take the time derivative of the continuity equation and subtract it from h_0 times the space derivative of Euler's equation, then we get something that looks very familiar, probably, to everyone in the room. We get a wave equation where the constant that plays the role of the propagation of speed squared is just gh_0.

The speed of a gravitational wave, which is the ratio of the frequency to the wavenumber, is $\sqrt{gh_0}$. I have made simplifications in saying that the wavelength is very long compared to the depth, and if you do it right, if you go look it up in Landau, you get a result that involves hyperbolic trigonometry, but in the limit of long wavelength it reduces to that answer. So it is a nondispersive wave, as I have

calculated it, whose speed depends on g and h_0.

For this application, g is considered a constant. With the exception of earthquakes, that is usually true in California. For Denardo, that is not true. Denardo modulates local gravity; that is the way he drives his resonators. But for this particular application, we have a constant and the wave speed depends on the height.

So now we can ask ourselves what happens if we let the height get larger and larger and the troughs get deeper and deeper? Since the sound speed depends on the height, then the local velocity of propagation at any given point on the wave will depend on the change in sound speed with height. In addition, the wave is moving to the right (that was on the previous graph) and, therefore, since the fluid motion, v_x, is in the same direction as motion, the local velocity is shifted by two effects, and this is important, so I will stress it again.

When you allow the wave to affect the propagation speed, it affects it in two ways. It affects it because the fluid is moving itself and, therefore, it is Doppler-shifted. If the fluid is moving at a velocity v, the sound speed is c + v, and if it is on the troughs, it is moving in the other direction, the local velocity is c - v.

The other thing is that the wave itself affects the propagation speed, what I like to call the thermodynamic sound speed or the equilibrium sound speed. In this case, the way it is affected is that the speed is a function of height, height is a function of velocity. Then we have a term that we can add in to the convective velocity. The reason for expressing it in terms of this velocity, is that we know how to correct the sound speed for the Doppler shift; we just add the velocity. So now we are going to change the thermodynamic dependence on the sound speed into something that is also proportional velocity.

We can do that because we know from the continuity equation that h and v are connected, that is, we take the time derivative of h and we multiply that by the gradient in v, $\frac{\omega}{k}$ just gives you c and you end up with this quantity: the amplitude over the equilibrium height is equal to the flow velocity over the sound speed and that ratio has a name. It is the acoustic Mach number. That is why I introduced the big M.

The sound speed, on the other hand, c^2 is gh, and I will introduce

the only advanced mathematical technique I plan to use. The integral of dx over x is log x plus c. Those of you who have never seen that, I suggest that you visit a tatoo parlor. It is useful.

If we take the log of both sides, we get, because of the square here, two-log gravity is log of the depth plus a log of g. We have already said that g is not changing. So if we take the derivative of these logarithms, we regenerate by the inverse -- you know, the first law of calculus. Delta c over c is a half delta h over h, which is a half h_1 over h_0. Therefore, dc by dh is equal to one-half c over h. You see, that is a technique for doing calculus if you are a computational cripple and you have to do differentiations.

If we plug those results back in here, what we find out is that the local sound speed depends on the particle velocity, which I showed depends on the height, and c as a function of x is your equilibrium sound speed plus three halves times the x velocity component. One-half is due to this, the equilibrium sound speed being a function of depth. Two halves is just due to the convection.

Q: I have a question on your last slide. You have air above it, and water, air and water, and then you have a kinematic condition on the bottom and you have some dynamic condition on the top. It is going to be nonlinear as hell, it looks like to me. What assumptions do you make to get the dynamic -- you have linear equations, what assumptions do you make to --

A: I put some of those assumptions down, that the wavelength is much greater than the depth.

Q: Which ones eliminate the dynamic condition at the top?

A: I let the top move. Let me show you the structure of what I am getting at. Basically, as I say, I am going to do the same problem over and over, again, with more detail and completeness included in each iteration. I want to get the basic phenomenon out.

It is a perfectly valid question, and the answer is, first of all, I made the depth so shallow that all those hyperbolic tangents go away. So the bottom boundary condition went away. I have just said it is impenetrable and if the velocity is moving this way and there is a gradient, the surface has to move up.

Q: I understand that, but then what about on the free boundary there?

A: That was the free boundary. I let the height go up. I could put oil on there, I could put a membrane on there, but that is homework, right? You do that.

What is the consequence of the fact that the high parts are moving faster than the low parts? The answer is, the consequence is, "surf's up."

(Transparency) (8)

If we start with an initially sinusoidal wave propagating in this direction, the zero crossings will travel at the thermodynamic equilibrium sound speed and the crests will advance and the troughs will retard. After a certain amount of time you can see the waveform has become distorted; it is no longer sinusoidal, it has that shape. You saw that in Professor Keolian's demonstrations.

If we allow this to continue, of course, this will get over this, this will get over that, and for a free surface this is allowed; it can become multi-valued. If you do not have a free surface, the gradients will get so large as to stabilize that (we will get to that shortly). Basically, since the crests are traveling faster than the troughs, they finally start to approach the zero crossings and you get the wave forming a shock.

We can ask ourselves how far the wave has to go or how large the wave has to be so that that shock formation actually takes place. We can ask ourselves how long does it take for the crest to get right above the trough and the zero crossing? If we ask that question, it is very easy to solve. The distance between this point and this point is well-known, it is a quarter wavelength, because I have got a whole wavelength there.

The excess speed is the part 3v over 2, so I can ask myself how long does it take for the crest to overtake the zero crossing, and I can define that as what I call the shock-inception distance, although that is not exactly right (I will refine the definition shortly).

But you can see that after a time T, shock the wave will have traveled this distance, the zero crossing will have traveled this distance, speed times time, and I know what the distance is and I

know what the speed is and, therefore, the shock-inception distance is simply the wavelength over six times the Mach number.

DR. HARGROVE: The crest has to go forward only one radian to get infinite slope at the zero crossing.

A: Thank you. That is what I said, this is approximate. If you read on the next line, exactly. D_S is usually defined as the distance at which some part of the wave becomes vertical. The slope goes to infinity at some part of the wave, which, of course, is the zero crossing. If you do that, then my simple hand-waving argument is wrong by a quantity known as the Rudnick (since Izzy is not here, I can say that). Professor Rudnick, as I mentioned to some of you, used to be able to get any result at all in one or two lines starting with Newton's law, but you can just bet it was off by pi over two.

(Laughter)

So I am trying to carry that tradition on. The numeric actually goes from 6 to 3 pi, if you ask where it actually gets vertical, but the details of that derivation are not going to stay with you. This model with stay with you and you will be within a Rudnick, because you just ask how long does it take? That defines the shock-formation distance.

If I had done it right, the way Logan pointed out, this number would not be 6 but it would be 3 pi.

(Transparency) (9)

What we have done so far is to explain how a wave distorts. It happens to be a shallow water gravity wave, but I chose that because it is clear that the speed depends on the height and the convection. We have other media that we are much more interested in, certainly, for acoustics. For that, I like to define a constant in the way that Professor Putterman taught me, in terms of the Grüneisen coefficient. This is a term borrowed from solid state physics.

As Professor Anderson was saying, thermal expansion is something that affects materials and the reason it affects materials is due to nonlinearity of the elasticity. Robert Keolian said that for a linear system you have a potential well that is parabolic, so no matter how large an amplitude oscillation you have, the mean position is at the center of that parabola, because it is symmetric.

But if we add a nonlinearity, say, third order, if you have a spring that gets a lot stiffer as you close it up, as most springs do, then if you have thermal motion in that well, as the motion gets larger and larger, the equilibrium position will switch around, and that is what thermal expansion is. So the solid state people look at the nonlinear coefficients and the elastic tensor for solids and define this deviation from perfect linear behavior as the Grüneisen coefficient. We have adopted that for acoustics, because it is just a convenient thing to do. I label it here Gamma.

What we are saying is that, in general, if you have a wave system, that wave will affect the propagation speed in two ways: through the convective and the constitutive contributions. If we have a wave that has some velocity, c_0, then it will have a local sound speed that depends on the local velocity. However, c depends on whatever parameter that wave is changing times however that wave changes with respect to velocity times velocity. We now can have this convenient form of writing the first correction to the velocity as u times this Gamma, and the Gamma is the Grüneisen coefficient, as I refer to it, that includes both constitutive and convective contributions.

If we take the example, say, of a wave whose speed depends on density, then the Grüneisen coefficient will be dc d ρ times d ρ dv. If we go back to the continuity equation we can evaluate the d ρ dv contribution, which is ρ_1 over v_1. This is just ρ_0 over c, and there is a simple expression for the distortion parameter or the Grüneisen coefficient.

What I would like to do now is apply this to ideal gases. Sound waves in air will be the first example. In air, propagation is adiabatic -- I assume you guys worked it out in that evening session, is that right, you are all satisfied with adiabatic propagation, that Newton was wrong and Laplace was right?

Under those circumstances, when you compress a gas its temperature goes up, because the heat has no place to go, so the change in temperature over the ambient temperature is related to the change in pressure over the ambient pressure by this constant, which is gamma minus one over gamma, where gamma is the exponent polytropic coefficient for ratio-specific heats.

The speed of sound is temperature-dependent, c^2 in an ideal gas is gamma RT over m. This made Mike Moldover rather famous, because he used sound speed measurement to determine the gas constant. But using differentiation, we can see that the change in sound speed over the ambient sound speed is one-half the change in temperature over the ambient temperature.

(Transparency) (10)

If we stick this back into my generalized expression for the Grüneisen coefficient, we can see that here the chain -- obviously the convective part is straightforward -- but the chain of events that leads to the change in sound speed with pressure is the sound speed changes with temperature, the temperature changes with pressure, and then the pressure is related to the velocity through the Euler equation.

So when we stick in all those derivatives, as I have motivated them on the previous slide, the answer is that Gamma for an ideal gas is gamma plus one over two. That is the distortion parameter. Gamma for shallow water gravity waves was three halves.

If we ask ourselves what is the typical sound wave that would be generated by a top rock and roll group, say, my former group, Clebsch Gordon and the Coefficients, when we were playing our greatest hits, it is roughly 120 dB. That gives us a Mach number of about 2 parts in 10^{-4}. If we plug that back into our expression and now use the correction, 2 pi, in here, we find that in roughly a third of a kilometer, the zero crossing has gone vertical and we have started shock formation.

If we were to increase the frequency by a factor of 10, this distance would decrease by a factor of 10, so the distortion scales with the wavelength. There is a natural length scale in there that tells you really how many wavelengths before it distorts and it scales with the inverse of the Mach number. If I increase the amplitude to 140 dB, 10 times the acoustic amplitude, then it would shock in about 23 feet, which is roughly how far I would like to be ahead of the loudspeakers.

What I have left out of this picture, of course, is dissipation. If there is a large amount of dissipation, this shock formation will never happen. In a medium that has no dissipation, shocks will always form; it is just a question of how far out it is going to go before it shocks up. Some of the work that Bruce will talk about for surface waves on

water, the stuff that you see when the wind is blowing on the water, the Goldberg numbers -- that is what this is -- are in the millions. The nonlinear effects will always dominate.

Here it is not as obvious, but if you go to the famous work of Professor Bass and colleagues you can look up what the attenuation rate is for sound in air and, therefore, you have a one over E distance for dissipation of roughly 4.3 km at a kilohertz, which I chose arbitrarily. Therefore, the Goldberg number, which is the ratio of the distance it takes to attenuate to one over its E value to the distance it takes to shock, is about 13, which means it shocks. There is plenty of signal there. It will shock before it peters out. That is what the Goldberg number is about.

Now we have introduced shock-inception distance, Goldberg number, constitutive and convective nonlinearity.

DR. SWIFT: That is for a plane wave, so if you have got sound radiating from a point source or something you have also got the one over R kind of falloff.

A: That is correct, Dr. Swift. Now, let us continue.

(Laughter)

I loaded the audience with guys who still work in this field, so I can punt. Here, do you want to talk about spherical divergence? Yes, a different quantity. When you are falling off by one over R it is a different quantity.

Q: So it is harder to get shocks to form?

A: Yes. The opposite is horns. Professor Rudnick did some excellent work in the early days for Altec Lansing, because horn loudspeakers sounded distorted, and there you have higher intensity of air near the source. It shocks up before it gets out the horn, which is an effect that I paid about $500.00 a side for to get into my loudspeakers.

DR. HAMILTON: Just to answer that question about spherical spreading, if the plane wave distortion increases in proportion to distance, the spherical wave increases only in proportion to longer distance, so you have a much slower distortion scale, so that Goldberg

number can kick in a lot earlier and beat out the shock. It is much harder to get shocks in spherical waves, logarithmically so, than in plane waves.

A: And vice versa, it is much easier to get it in horns. In fact, that is why the beach is such a good example: It is a horn. Not only are the crests catching up with the troughs, but the amplitude is catching up just by energy conservation. You are going into a horn.

(Transparency) (11)

What is the ultimate fate? The ultimate fate is that the crests overtake the troughs and you end up with a waveform which is stable. You end up with a sawtooth. I have just written up here the Fourier expansion of a sawtooth, because I am going to use it later. But ultimately the crests will take over the troughs.

It is possible to calculate the attenuation rate of this sawtooth waveform. One way to do that is, again, through a phenomenological argument which is known as the Rankine-Hugoniot relations. You have a shock front and if you are moving with that shock front you have basically got a mass flow through the shock front because the velocities ahead and behind the front are different, and the densities in front and behind the front are different and, therefore, there is a flow of mass through that front. The stuff that comes into the front must be coming out of the front, so you can write a mass conservation equation.

Similarly, the momentum is conserved. This is ρ u times u. Momentum in the fluid is the density times the velocity and we are convecting that with velocity u, so you get the ρ u squared terms ahead and behind the front. Similarly, the energy is conserved through the front. Therefore, one can ("one" being Landau) use those three conservation equations coupled with the ideal gas laws and show directly that the entropy discontinuity across the front is cubic; it is proportional to the third power, the relative pressure difference to the third power.

(Transparency) (12)

Given that we know what the entropy discontinuity is across the front, it is therefore possible to calculate the rate of change of energy

as this thing propagates along, because the energy is just simply TdS, the temperature times the change in entropy. So if we know the rate of energy dissipation in the front, which depends now on the amplitude to the third power, we can calculate what the attenuation is.

But that requires the ability to do algebra. Geometry is much simpler. If we take a look at the fully developed shock now -- we have gone many shock-inception distances and the stuff is just piled up -- what is it that finally stabilizes that wavefront? It cannot do this, it cannot climb over the edge.

What keeps it from climbing over the edge is that the gradients get so large that the dissipative terms that Professor Bass was talking about, the thermal conductivity, the relaxation, and the shear viscosity, will get larger and larger, because the energy dissipation there, is positive definite, depends on the square of the gradients of the velocities, and the square of the gradients of the temperatures, and if those gradients go to infinity (which is what a vertical line here means), the shock fron can consume as much energy as it needs to.

What we can do to calculate the attenuation is actually quite simple, because we know that this shock front is remaining nearly vertical, it cannot do anything else. If it tries to get close to vertical the gradients will go up and it will stabilize.

But back here each portion, since this is a straight line in a sawtooth, moves ahead in proportion to how high it is above the axis, so you basically get the equivalent of solid body rotation about that point. We can, therefore, draw where the wavefront would be and this would be the amount that that would have to have to be dissipated by the shock.

By similar triangles we know that du, the change in the amplitude, over gamma u dt, which is the distance that that moved ahead in a time dt, must be the same as the ratio of u, if this were an infinitesimal amount (I have shown this as being one-quarter of the effect but I meant it to be infinitesimal), to half the wavelength. So we can write that the similar triangles give us this relationship and if we let dt equal dx over c_0, then we can integrate the expression to obtain our attenuation.

You can see that this has nothing to do with exponential attenuation. In fact, I believe this is what Seth calls long-time tail.

This is the initial position and the initial Mach number, so we can set that to whatever we want and we see that this decays -- this is one over the amplitude -- so the amplitude decays like one over x or one over t for a shock wave.

(Transparency) (13)

There are two interesting facts, therefore, about the attenuation of shock waves. I alluded to that during Professor Bass' talk. The decay of a shock is not exponential. It does not fall off like an exponential; it falls off much slower, number one. Number two, that expression, which has been experimentally verified, as you will notice, has no transport coefficients in it. It does not depend on the relaxation, it does not depend on the viscosity, it does not depend on thermal conductivity.

The reason is that, of course, the amount of energy that has to be eaten gets eaten. If you have something with high thermal conductivity and no viscosity, the thermal conductivity will eat it up. But that result, I think, is an amazing result for the simple fact that it, one, does not fall off exponentially and, two, it does not depend explicitly on the transport coefficients.

Q: With your shock fully developed, neglecting any loss mechanisms at that point, you would picture just the fact that the wave speed is different along that random shock that the shock should just keep stretching out?

A: It would rotate. It cannot stretch out, because it has got a shock ahead of it and a shock behind it. When Professor Keolian talked about the N-waves, the N-wave was a single pulse, so the fact that the crests were going faster and the troughs were going slower stretched that guy out. This is a repeated sawtooth, so it cannot stretch, it has no option of stretching.

DR. BASS: I think you covered it, but I think it is important to note here the physics is the same. Basically it is still a transport phenomenon. It is a question of how you observe a transport phenomenon. You have not done anything to the physics by introducing shock, but you now have a different way of including that same physics.

A: Again, I am going to do this two ways. I am going to do it the

right way and I am going to do it the cheating way. I do want to refer to this, as Professor Bass points out. This is the equation that he put up, that the attenuation in a classical medium depends on the frequency squared, the shear viscosity, the thermal conductivity, and squiggle, which are all the relaxation phenomena.

One can show that bulk viscosity is merely a sixth relaxing variable in the hydrodynamics. But I can then write this whole thing, for reasons you will see later, in terms of a mean collision time. Whether that collision time really represents relaxation processes, viscosity, or thermal conductivity, I do not care. I am just going to define a time and that time times the speed will be the mean free path corresponding to that effective collision time.

One can calculate the thickness of the shock using these transport equations. What you do when you look at the attenuation of the shock is you do not get the attenuation length, you get the thickness of the shock required to give you the attenuation length, and that can be written in the following form, where this derivative is a measure of the nonlinearity, the change in the sound speed with pressure and, for an ideal gas, again, we can work that out. It is dc by dt at constant entropy times dt by dp at constant entropy, which gives you this quantity.

(Transparency) (14)

If we now substitute this expression for the junk in there, we can write Landau's result as delta, the shock thickness. It is the simplified expression including all the other thermodynamics and nonlinearity that I threw in. The answer is that for a shock wave whose amplitude is roughly atmospheric pressure, the shock wave thickness is only about 30 mean free paths. That gradient is steep, for a large-amplitude sound wave it is a very steep gradient. I would not put too much stock in that number 30. It is just telling you that the shock thickness is roughly on the order of a few mean free paths.

Again, there is another way to understand this that I think has a certain amount of intuitive appeal. Because the attenuation is proportional to ω^2 and because the Fourier composition of a sawtooth wave has components that fall off only like one over ω, the term in the linear attenuation theory that dominates the amount of attenuation you get is the term whose half-wavelength is the shock thickness.

That is the highest frequency component that can determine what the steepness of this Fourier series is; that is, we are going to truncate the Fourier series at some term and the half-wavelength corresponding to that term would be the steepness of the shock front.

If we make that assumption and say that the energy contained in the wave is the kinetic energy, so it is some constant times u squared times the wavelength, because we want the energy contained in the volume defined by one cycle of the sawtooth wave. k is simply a constant that depends on the density of the fluid and the integration. We can say that the change in energy over the energy for the shock front is proportional to du over u and can go back to that nice geometrical construction which gave me an expression that produced du over u. I can plug all that stuff in and get this very nice expression for the attenuation in a shock.

The linear attenuation, because it is an exponential process, is just 2 alpha dx, because if the attenuation rate is alpha, then the energy goes like the linear variable squared, so you get a 2 here -- I do not care about factors of two. If we take and equate the shock losses over a distance equal to that wavelength to the attenuation of a wave whose wavelength is half that of the shock thickness, we put those things together, it should generate for us what that half-wavelength is, and I have done that for you here.

Sure enough, if you plug in the numbers you get, the shock thickness, using the expression, is 20 mean free paths. For a calculation like this, 20 is 30 for large values of 20.

What we have done now is, we have taken a look at this distortion process. We have said that a waveform will distort until it can distort no longer. Then that waveform will be stabilized by dissipation, irreversible processes that occur within the shock front, and that shock front thickness adjusts itself to absorb the requisite amount of energy.

(Transparency) (15)

Let us proceed from gases to liquids.

DR. HAMILTON: I just wanted to bring up kind of a theoretical paradox that your slides bring up, and I am not sure I really know the answer. A couple of slides ago you showed that the entropy jumped

over a shock as a cubic of the acoustical variables, let us say, yet all the theory on the slides that followed, your equal area rule and your shock-thickness formulas, all of those came out of quadratic nonlinear theory.

What I am saying is, we began the analysis by assuming the process is isentropic and then used that to derive the relations for the shock waves, which owe their very existence to the fact that there is viscosity and losses at those shocks. So you have kind of an interesting circular argument here. You could not have gotten to the answer without assuming it is isentropic, but the answer you find is that you have got these losses that are correctly described by this weak shock theory. It is kind of unusual.

A: How do you feel about vorticity?

DR. PUTTERMAN: It is quite right, what Mark says, that the jump in entropy is a cubic shock, but in the interior of the shock it is quadratic. So the theory is correct to quadratic order, so I do not see what the paradox is. If it is quadratic in the middle, then you have done everything right, keeping corrections to the entropy everywhere but your quadratic in the fundamental variables.

If the entropy, for instance, inside the shock had not been quadratic, then there would have been a contradiction in how everything was formulated, or if it had been first order, for instance.

DR. HAMILTON: You could get the answer from two limits. One is the second way Steve showed, you could get it from the continuous wave-shock solution, take a limit of no viscosity to get your step. But he was coming at it the first time through the Rankine-Hugoniot relations, and there you are throwing out your cubic terms, you are assuming no loss, you assume a perfect gas and you will still get the proper decay --

DR. PUTTERMAN: But the thickness of the shock is adjusting itself so that the gradients always yield the same dissipation, so the viscosity drops out, the thickness times the gradient --

DR. HAMILTON: There is no thickness in the weak shock theory. Anyway, it is kind of a red herring that comes up from time to time.

A: I think I have an answer to that, but maybe we will save it

until the discussion. If you believe in vorticity, then you have no problem with ruling out the singularity and still having a good fluid dynamic theory. I would approach it in a different way.

In any case, what we have done so far is we have looked at shock wave formation on the surface of water, we have looked at shock wave formation in air, which is sort of the common problem. I would now like to move to fluids, because it is going to take us into more bizarre types of effects. As long as everybody is comfortable with the first two examples, it is time to raise the stakes.

When you get into fluids, it is common practice to expand the equation of state in terms of a power series. Frequently the power series is expanded in the condensation. This is problematic for me -- I do this all the time in lectures and I do not realize it until the end -- but now I am using S to imply condensation, when I have been using it previously all the time to mean entropy, so here I am, stuck, again. My record is three different values for K.

We can expand the equation of state, the relationship between pressure and density in a power series -- this is called the virial expansion -- and we can then identify these terms, the various coefficients in the expansion, from the Taylor series terms. So the first term is dp by d ρ, which we know is the sound speed squared. The next one is d squared by d ρ squared, and the next term is d cubed p by d ρ cubed, et cetera. All these things can be expanded in various ways, but they all come down to the sound speed or derivatives of the sound speed with pressure, higher derivatives of the sound speed with pressure.

Therefore we can write the speed of sound in terms of these coefficients. What we have done up to this point is we have truncated the series at this term and defined the Grüneisen coefficient -- in fact, I have rewritten it here. You can see, since the Grüneisen coefficient depends on one plus half this term, B over A.

Under those circumstances, when the pressure was high and the density was high, and the velocity worked forward, this term caused everything to speed up, and when the velocity was in the opposite direction, the pressure was low, and the density was low, this was negative and they worked together.

If we include the next term in the series, this term is positive definite, it is quadratic, so if this term here, say, is negative, if C is negative, then high pressure, high amplitude, goes slower, and low amplitude goes slower. If we go to the higher order nonlinearities, instead of getting the cooperation between convective and constitutive nonlinearities, constitutive nonlinearity can introduce a term which does not depend on the sign of the flow velocity.

We can ask ourselves, one, what happens if that is the case and, two, is it ever the case? Perfectly valid questions.

(Transparency) (16)

What happens is, of course, with nonlinearity there is always a scale, there is always a threshold or an amplitude at which the effect is important and below that the effect is not important. That is what the Goldberg number was telling you. So we can define a velocity that I call v_d for reasons that I do not remember, where the term that is positive definite in the velocity correction exactly cancels the convective and constitutive terms, the Gamma Grüneisen coefficient.

That can be expressed in terms of the virial expansion, as I have done up here. I suspect it is correct, because I had Dave Blackstock check it. If we go to an amplitude where the velocity is, say, half of this critical velocity, half of the velocity where the terms take over, you can see what happens.

On the negative part, the trough slows down for three reasons. It slows down because the velocity is in that direction, it slows down because the Grüneisen coefficient slows it due to the constitutive first-order contribution, and it slows down because of the velocity squared term. So the trough shocks up faster. But at the crests, the second-order term in the virial expansion almost cancels the Grüneisen term and the crests are much less distorted.

In fact, when this velocity equals that velocity, uds the wave is pretty much symmetric here and pretty well shocked down here. As we go to higher amplitudes, eventually, instead of forming a single shock, we form a double shock. Rather interesting.

Does that make sense? That is, a higher order term can now overwhelm a lower-order term, because that is what nonlinear is about. Higher order terms always win if you can go high enough.

(Transparency) (17)

We can ask ourselves, is this observed in an ideal gas? The ideal gas is particularly easy, because we can express the adiabatic gas law, as I have shown here, as one plus the condensation raised to the polytropic exponent, and we can do a binomial expansion on that and pick the terms off one by one. B over A is this term, C over A is this term -- in the A part the gamma gets divided out.

Lo and behold, if you take those results and plug them into the previous slide, the denominator is identically zero. There is no amplitude whatsoever in an ideal gas that will create double shock formations. Interesting result. You will not see it.

You can ask yourself whether they are observable in classical liquids. The answer there appears to be no, also. If you take the measured B over A and C over A coefficients, which I do not necessarily trust, you find that the velocity at which the second-order term is equal to the Grüneisen term corresponds to a velocity that is 20 percent higher than the speed of sound. Therefore, the pressure amplitudes have to be about 26,000 Atm, and that is a good two orders of magnitude higher than the highest cavitation threshold that has ever been measured (that was measured, by the way, by Moe Greenspan, whose name seems to be popping up a lot around here). That is very good.

Q: How does it mean that that is not seen, because that would mean Mach number is infinity?

A: Let me get back to the equation.

(Transparency) (16)

For an ideal gas, the C's and B's are calculated. You plug them in here, and this goes to zero. That means that v_d, the velocity amplitude you have to go to is infinity before that term takes over. It is not that this equation goes to zero; it is the denominator of the equation that goes to zero.

(Transparency) (17)

In classical fluids, this would not be seen, but there is a

gentleman who has been doing calculations on classical fluids near their critical points and he claims that this effect should be seen in fluids near the critical points. Just yesterday Mark (Hamilton) and I were having a discussion and Mark claims that for rubber the third-order nonlinearity is dominant. Your student just did a calculation on this, correct, and got basically the same result for double shock formation?

DR. HAMILTON: For shear waves.

A: But it has not been observed. It has been calculated?

DR. HAMILTON: It has been observed.

A: Okay, great, you will have to give me the reference on it. How long ago?

DR. HAMILTON: Some experiments, I think, done in Moscow by _____.

A: I would love to see that. I actually started to worry about this for a completely different reason. I was interested, when I was working with Dick Packard and Greg Swift at Berkeley on critical velocity effects in superfluid helium-3 -- I will be saying more about superfluids, so I want to give a little bit of background.

Below a temperature of about 2° above absolute zero, helium, which is the only thing that remains liquid down to absolute zero because of its zero-point motion and its light mass and its high symmetry, (there are not great van der Waals forces that tend to condense the stuff), it stays liquid all the way down to absolute zero. At about 2.2° above absolute zero it undergoes a transition. A fraction of the fluid behaves as though it has absolutely no viscosity. By no viscosity I mean that you can start up a flow of this fluid and the time it takes to decay by 10 percent is 10 times the age of the universe. It is persistent current; it has no viscosity.

But because the fluid behaves as though it is made out of two components, one without viscosity and one with viscosity, it is possible to generate a whole variety, a menagerie, of sound modes. One sound mode that can be generated is a sound mode in which you take the fluid and put it into a porous medium, so the porous medium immobilizes the component that still has some viscosity and the

inviscid part is free to move.

This particular sound mode is called fourth sound, so you must know there are zero, one, two, three, four, five --there are six sound modes if you count the viscoelastic mode. One of them is this mode where only the superfluid moves and the normal fluid is locked by the viscous interaction with the porous medium. It has a sound speed, Cu, that is determined by the compressibility of the liquid as represented by the ordinary first-sound speed and the density of the superfluid component. This is the compressibility and this is the inertia, and because only the superfluid is moving, the inertial term may be temperature-dependent.

(Transparency) (18)

The total density of the fluid, of course, is just the sum of the viscous and inviscid parts, but because the superfluid fraction is a thermodynamic term, it can depend on only other thermodynamic terms. That is, the dependence of ρ_s over ρ must be on temperature or pressure or some other Galilean invariant.

When it comes to relative velocity, v_n and minus v_s squared is an invariant. It does not matter whether you are going right or left, it is positive definite. So the superfluid density depends on the density with no flow minus v_s squared, the superfluid velocity, over some critical velocity. That is the lowest order that satisfies thermodynamics.

Therefore, when you have flow, you have a suppression in the superfluid density and that suppression shows up as a change in sound speed. If you look at superfluid helium that I just mentioned, that becomes superfluid at 2°, that critical velocity, v_c is dependent on what is called the roton minimum, the smallest amplitude excitation one can put into the superfluid and thereby dissipate energy and make it nonsuperfluid.

It turns out that that is a fairly large number, a couple of meters per second, and that it is way above the measured cavitation threshold. But in about 1974 superfluidity was discovered in the rare isotope of helium-3 at a much, much lower temperature, a couple of millikelvin and, therefore, the excitations necessary to generate a reduction in the superfluid density have much lower energy.

For superfluid helium-3, the critical velocity is on the order of only a few hundred microns per second. When I did these calculations, I believed that it was possible to take a look at the depression in the superfluid state by looking at this bizarre shock formation in helium-3B near a temperature of a couple of millikelvin above absolute zero.

There was an experiment that was done that, actually, when you plugged in the numbers, they were a factor of 20 above this, so it should have been shocking like crazy with these double shocks. Because of the nature of the resonator (it was rather unconventional, it did not use sound transducers), nobody ever saw the waveform. So I am glad to hear what Mark says, that actually in some system like shear waves and rubber this behavior has been observed.

(Transparency) (19)

We can go on, again, and look at another anomalous situation. We have shown, in the Grüneisen coefficient that both the thermodynamic, or the constitutive contribution, and the convective contribution are in the same direction. But that is not necessarily true. I would like to show you a couple of systems where, in fact, the constitutive correction can be exactly canceled by the convective correction and, therefore, you get propagation without distortion at arbitrarily large amplitudes. This has been observed in two systems. In fact, you can force the Grüneisen constant to go negative.

The most dramatic of these systems is what is called third sound in superfluid helium. If you have a flat substrate and you adsorb a film of helium on there -- you can adsorb films that are only 1, 2, 3, or 4 atomic layers; you make a very, very thin film. But since superfluids have an inviscid component, again, the normal fluid is locked by its viscous interaction with the plate, but the superfluid is free to move around without dissipation.

Professor Rudnick made this statement, which I think is one of the great statements in quantum fluids, which is that "Superfluids obey the laws of quantum mechanics on the macroscopic scale and the laws of classical mechanics on the microscopic scale."

When you think about it, how does one define hydrodynamics for a system that is only one atomic layer thick? Really, in your mind,

you think sort of statistically, you must have more helium atoms in one place than in the other, on the average. But, in fact, superfluids do obey the laws of hydrodynamics on the microscopic scale, but because these films are so thin, the restoring force is not gravity but the van der Waals interaction.

So at one atomic layer it is all solid, because the van der Waals pressure is over 25 Atm and helium becomes solid at 25 Atm. But the next layer is fluid, and the next layer is fluid, and the next layer is fluid, in fact, until you get a nice thick film. So you have got basically a gravity wave, just what I calculated for you before. Its inertial term is modified by ρ s over ρ, and here it is an average ρ s, because you have got to take into account the solid layer thickness (that is not superfluid).

The restoring force here in the shallow water surf out there on the Asilomar beach was due to gravity, but the force on the helium film here is the van der Waals attraction, which is proportional to one over the thickness to the fourth power. Therefore, the sound speed does not get faster as you go deeper, it gets slower as you go deeper.

(Transparency) (20)

Given the weighted superfluid density, we can write the excess third sound speed in terms of the velocity. The third sound speed is a function of height times the change in height, and the third sound speed is a function of velocity times the velocity. Again, we can write the equivalent of a Grüneisen coefficient, where all the convective and the nonlinear terms are contained here, and you can see that when h is equal to 3D, that is, when you have about a 10-Å-deep ocean, this term vanishes. Therefore, you can propagate arbitrarily large (if such a thing exists in atomic-layer stuff) waves that will not distort.

For thin films, shocking is normal, because the quantity is positive, but for very thick films that quantity is negative and the waves break backward. Perfect surf that no human will ever enjoy. But even a small-amplitude wave, a fractional angstrom amplitude is tsunami to an electron. Electrons are very, very tiny. In fact, people have surfed electrons on waves in superfluid, but not humans.

This is a good time to take a break and then we will come back and look at some more deviant cases.

DR. GARRETT: I received some very interesting comments over the break. One, I am ashamed to say, I did not know. Apparently it is well-known to everyone except me that fused silica also exhibits this backward-breaking shock -- not breaking, but the distortion in the reverse direction -- and that is used regularly by the guys at Washington State University, where Phil Marston is. WSU was the center of shock wave physics for three decades.

When they create a shock wave and want to see how it propagates through the medium, they want to be able to have various steepnesses of shocks, but the way they generate the shock is they pump up some gas behind a big diaphragm, just as in the demonstration, they pop the balloon. The thing goes roaring down the tube, it is a big shock wave, and what they do to tailor the shock wave is they put fused silica in front of different thicknesses to bend it back.

Anyway, I am surprised to know that it has more than just amusement for electrons.

DR. MARSTON: What goes roaring down the tube is a projectile.

A: That is right. I am sorry.

(Transparency) (21)

I want to show you one other example of, again, distortion being zero or being, in fact, negative. This again comes out of superfluidity. I learned more about propagation in classical systems because I studied quantum systems and saw what the difference was. I would like to do the same thing one more time with a sound wave that won Lev Landau and Pyotr Kapitsa the Nobel Prize in 1965 and 1978, respectively.

In superfluid helium, when you have bulk liquid, you do not force the normal fluid to be immobilized. If the fluid is in a pore or trapped because it is in a thin film, the normal fluid is viscously locked. In the bulk liquid, you can have counterflow of the superfluid and normal fluid components. Under those circumstances, you have an additional hydrodynamic velocity field. I will get to this eight variable case in this section of the talk, but in ordinary hydrodynamics five variables completely describe the system. Three variables for the velocity field v_x, v_y, v_z, and two thermodynamic variables like P and T

or ρ and s.

In superfluid, Landau received the Nobel Prize for showing that a second velocity field -- so you have now eight variables -- forms a complete description of the superfluid. That second velocity field is the superfluid velocity field. But in order to close the description, you need now, instead of five equations, eight equations, and he supplemented the Euler equation and the continuity equation with this equation. It is the superfluid acceleration equation, which looks very much like the Euler equation, except that the driving force for the superfluid is not pressure gradients but gradients in the chemical potential.

The chemical potential, as you see, is defined in this way. It is one over the density times the change in pressure, entropy times change in temperature, and then this rather insignificant term, which, again, is a Galilean invariant, since it is quadratic in the velocity difference. Its sole purpose was to get me a Ph.D., that term there, because it is an intrinsically nonlinear thermodynamic variable. You will hear a little more about that later.

For ordinary sound, we can argue by analogy, the square of the sound speed is the derivative of pressure with respect to density at constant entropy. So, if we look at the fact that the restoring potential is no longer the enthalpy but the chemical potential, we can see that there should be an analogous sound wave, where the sound speed squared is dT/d one over s, dT/d(1/s), and, sure enough, there is. It is called second sound and it travels at about one-tenth the speed of the compressional or ordinary first sound waves.

The second sound speed can be expressed in terms of the superfluid densities, the entropy, the temperature, and the specific heat.

But that second sound wave distorts in a different way. It has a distortion parameter, the Grüneisen coefficient, which happens to be the temperature derivative of the natural log of some combination of sound speeds and thermodynamics. At a temperature of 1.884° this vanishes. That is, the convective part is exactly compensated by the constitutive or thermodynamic parts.

(Transparency) (22)

You can see that easily here with a plot of first sound versus temperature, fourth sound versus temperature, and second sound versus temperature -- you see second sound is rather slow. We will get back to some of this when we talk about mode conversion and parametric arrays, so I ask you to give this a little bit more attention than it may deserve for just the shock formation problem.

You can see that the sound speed for second sound speed is pretty much temperature-independent at low temperatures and, therefore, the convective part must be the dominant part. If the convective part is dominant and the constitutive part is unimportant, then it shocks normally, and the Grüneisen coefficient is basically essentially unity.

But here you can see that as the temperature goes up the second sound speed goes down, just the opposite of an ideal gas. In an ideal gas, when the temperature increases, the propagation speed goes up. So it is possible that the curvature in the sound speed vs. temperature near the lambda transition here is sufficient that the constitutive part exactly cancels the convective part. That is, the hot crest travels slower, unlike an ideal gas, where the hot crest traveled faster. The real reason that I introduce second sound and third sound, is that there are situations that are completely different from what we see in shallow water gravity waves and what we see in propagation in gases and normal liquids.

Up to this point, I have used a fairly simple model that involves the constitutive and convective corrections to affect or to modify the local sound speed. I showed that under those circumstances the ultimate fate of a large-amplitude wave, a wave whose Goldberg number is large, is a repeated sawtooth shock wave.

What I would like to do now is to investigate the initial growth of that instability.

(Transparency) (23)

I am going to do it several ways, because, again, my claim is that the number of ways that we can see how this theory develops will be proportional to the depth of our understanding.

I would like to do it first in the way that Earnshaw did it about the time Americans were cutting each other apart during the Civil

War. It is a very old calculation. What we are going to do in this calculation is to take a look at the early events leading to shock formation. We take a sine wave which stops being sinusoidal as other Fourier components start to become important. The amplitude of the wave represented by the velocity or the pressure or the density, is treated as a series of terms. The linear term v_1 is the ordinary first-order solution to the wave equation, and v_2 will now be the first correction to that, and v_3 may be the next correction to those two.

The assumption will be that we are in this weak shock limit: that v_1 is much less than the sound speed, and that v_2 is much less than v_1. You will see that v_2 will be the second harmonic contribution and v_3 will be the third harmonic contribution, and so forth. It is very easy to see this in terms of the mechanism that we have developed of a wave dependent local sound speed by simply writing the first-order sound field as some amplitude, cos $\omega t - kx$, propagating wave.

(Transparency) (24)

If we regroup the terms, I can write this as ω times t minus x over c. There is your standard linear solution for propagating waves in some medium with the equilibrium thermodynamic sound speed indicated by the subscript "zero".

We can now go the next step by substituting into this equation the fact that the sound speed is proportional to the particle velocity amplitude. If we make that substitution for c_0, c_0 plus gamma v_1, and we assume that we are starting out and the amplitudes are not too large, then we can expand this argument in a binomial series. We get rid of the plus sign and we get this minus sign.

If we regroup the terms, we have got cosine of ω this plus cosine of ω this, and we know from our trigonometry identities how to deal with something that is cosine of something times A plus B. It is cos A cos B minus sine A, sine B, where I suppress the ω.

If I substitute my first-order solution, v_1, into this equation, and I assume that we are far from the shock-inception distance, then what we end up with is the original sound field minus a term that depends on distance, the distortion parameter, the amplitude squared, and these two trigonometric terms.

Again, by a trigonometric identity, we know that sine times cos is sine of twice the arguments and, therefore, the v_2 term can be expressed in this way. I have chosen to reexpress it as equation 61 in this form, because I think it is a little clearer as to what is going on. The second harmonic term occurs at a frequency that is twice the frequency of the disturbance and half the wavelength. Its amplitude is proportional to how hard you started and the distortion parameter. The growth, if we scale it in terms of the wavelength, depends on the Mach number.

Long before we form the shock, the distortion starts out by adding a second harmonic; that is, waveforms are not stable. If you start out with something in ω, you go a little way, you have got something that is ω, and 2ω and, if you go farther, that 2ω terms grows and several terms I will show you grow, also.

The assumption that we made in making this small was that v_2 was a small fraction of v_1, so now that we have a result for v_2 we can plug it back in and ask what the regime of applicability. We can set v_2 equal to v_1. The approximation will certainly fail by that point since we have chosen v_2 to be smaller than v_1.

Sure enough, the distance beyond which this approximation does not work is related to the shock-inception distance. So as long as you are much less than the shock inception distance, this Earnshaw approximation gives you the initial growth of the second harmonic component.

Q: Is there any significance as to why v_2 is a sine and your initial waveform is a cosine when you are showing your shocks?

A: There are two possible explanations. One is that I screwed up. The other is that there is a phase relationship between them that I have properly produced. I do not know which is the proper answer.

DR. HAMILTON: If you start with a sine, you will keep getting sines for all the harmonics, but if you start with a cosine, it will go cosine, cosine, for the higher harmonics.

A: So I did do it right? Thank you. I would never know, but I knew there would be somebody out in the audience who would worry

about it.

But, again, the point is that we have a self-consistent description at small amplitude, because when it falls apart, we said it was going to fall apart, and it does fall apart.

(Transparency) (25)

We can do this successively for higher and higher iterations; that is, we can take the result and plug that result back in as our initial result and generate a new result, plug the new result in, but you can see it was no fun. It was four trigonometry identities and eight lines of algebra, which exceeds my pain threshold by a factor of two.

There is actually an easier way to do this. The iterative resubstitution was done back in 1935, but there is a clever trick that a gentleman came up with a good 35 years ago

-- it is not clear if the guy is still around. One can take this idea of distorting the wave, that is, taking a sine wave, taking an amplitude y, which is sine of theta, and then letting theta be some value plus an amount that is shifted forward by that amplitude, again. So, again, it is an iterative scheme. You can draw the sine theta (this happens to be shown for a value of k equals 1, which is the shock inception distance).

Then you can just do Fourier analysis. Fourier analysis would not work, because superposition is not valid in nonlinear systems, but we have already taken care of the nonlinearity and now we just have a picture. We want to take a Fourier transform of a waveform, and that is the waveform. So this is a legal trick.

(Transparency) (26)

When one does that, you take your answer for that distorted waveform, and you Fourier analyze it. You know that integrals of trigonometric functions that have trigonometric functions within their arguments generate Bessel functions, so you end up with the Fourier coefficients, the B_n's, for the sine wave components that are as a series of Bessel functions, and that is quite satisfying.

Remember when I did the Earnshaw approximation the first-order solution did not change even though we generated more energy in the second harmonic solution, which, of course, has to be wrong. It

violates energy conservation. If you do it properly, you can see that the coefficient of the fundamental is dropping, it is a J_1 over its argument.

It is dropping where the second harmonic contribution is growing linearly with distance, but it eventually saturates out. The third harmonic is growing quadratically with distance. The fourth harmonic is growing cubically with distance. So rather than going through this iterative scheme as Fubini-Ghiron -- this is known as the Fay-Fubini solution if you do it the hard way -- you can see that you generate this series of higher harmonics, which, in fact, does, as you accumulate, become a sawtooth.

What we have seen is, in weak shock limit, when we start out -- we can see the evolution from a pure sinusoidal wave to the accumulation of a whole series of harmonics. Those harmonics then stabilize and form a shock wave.

(Transparency) (27)

We are at a good point here to review and summarize. The reason we are at a good point is because what we have been doing here is piecing together a whole bunch of concepts: convective and constitutive changes in the local sound speed; shock-wave formation; growth of higher harmonics; stabilization of wavefronts. It is all sort of a hodge-podge of putting these things together.

We will go on and do this right. We will go right back to the hydrodynamics and then generate all of this stuff and some new phenomena for you. But I wanted to go through it this way so that you get a feel for the physical phenomena before we get into the mathematics.

To summarize, we have looked at self-interactions of waves. We have looked at what happens when a wave modified a medium, and two things happen. The sound speed depends on the sound wave itself in two ways: The particle velocity adds a convective contribution to the local sound speed, that is, it Doppler-shifts itself. The thermodynamic sound speed is affected by the thermodynamic changes that the wave imposes on the medium -- there is a pressure wave and a temperature wave, it changes the pressure, changes the temperature and, therefore, it changes the local thermodynamic sound speed.

We have also shown now that nonlinear effects accumulate with distance. In the weak shock limit the harmonics grow like x over lambda to a power, n minus one, so the second harmonic grows linearly. The third harmonic grows quadratically, et cetera. That after a distance, lambda over 2 pi gamma M, a shock wave will form if nonlinear effects dominate dissipative effects (remember, we introduced the Goldberg number).

That waveform is stable and attenuation takes place due to the enormous gradients that are required to absorb the shock energy. That is, the dissipation within the shock front will control the attenuation of the sawtooth wave, not the magnitude of the transport coefficients, say, viscosity and thermal conductivity, or the internal relaxation processes. But the thickness of the shock front will adjust itself to accept the required amount of energy. For large amplitudes, I remind you that that thickness is on the order of a few tens of mean free paths.

(Transparency) (28)

We will move to the next phase of the presentation, which is what I call El Camino Real.

Q: I just wondered, how do you pick the parametric representation Eq 63 and 64? Where did that come from?

A: That is exactly what I did back when I was shocking up the water wave. If I start out with a wave whose amplitude is sine of theta, where theta is this nondimensional axis. I used a half sine wave in that picture. If I now say that the amplitude, say, 6, will depend on its location, it will not be where it used to be at theta, it will be at some point farther along. The distance farther along will depend on its amplitude. That is how I generated the parametric representation.

Therefore, I had a sine wave function whose argument had a sine within it and I generated Bessel functions beaucoup Bessel functions.

Q: In the first wave speed, second wave speed, and fourth wave speed, the first one is that it is like a kind of phase speed due to the pressure, the second one due to entropy. Where is the one due to the energy flow? And what does the fourth speed mean?

A: This is a designation that maybe I should not have introduced, but I could not control myself. The first sound is the ordinary sound in a fluid. It is a pressure-density wave where the motion of the fluid particles is restored by compressions and rarefactions. So its sound speed depends on the compressibility and the density as it does in a classical fluid.

Because superfluids respond not only to gradients in pressure but gradients in temperature, the second sound is a wave of temperature and entropy, so there is no net mass flow, there is no center of mass motion, there is no mass flux.

For second sound the superfluid and the normal fluid are moving an anti-phase. As superfluid piles up at some location, it is getting colder, because a larger superfluid density corresponds to a lower temperature. At the point where the superfluid is moving away and the normal fluid is coming in, you have a higher normal fluid density and that greater normal fluid density corresponds to a higher temperature. This is because the density of superfluid goes from zero above the transition and grows monotonically as you get colder and colder. You get more and more superfluid and less and less normal fluid.

So if you jam a lot of normal fluid together, you raise the temperature, you are on that equilibrium curve. So first sound is a pressure wave, the second sound is a temperature wave, the third sound is surf in shallow films of helium, the fourth sound is the Biot mode. Remember when Professor Bass talked about fluid in a solid matrix, like air in sand or sea water in sludge? You have two media there. One of them is the matrix and one of them is the fluid that penetrates the matrix.

It is possible to have two modes there. There is a matrix mode, where the fluid moves with the matrix, and there is a mode where the fluid and the matrix are in anti-phase. The fourth sound is much like that. In fourth sound, the matrix is fixed and the superfluid component is moving. What you have done is you have artificially immobilized one of the components and now you are just looking at the flow of the other component, which has no viscosity so you cannot immobilize it.

Q: What is the speed of energy from?

A: The fourth sound speed is nondispersive, so the group speed and the phase speed are identical. The energy flow for compressional waves of first sound is 240 m/sec, the sound speed in the liquid. The heat in superfluid does not diffuse. In ordinary fluids heat diffuses. It does not have a velocity, it has a diffusion rate. In superfluids it turns out that heat does travel in waves. I was pushing it. I had hoped to show a video on superfluids to get you a basic understanding, and I chose not to do that.

What I would like to do now is go back and redo this same problem of shock formation, weak shock theory, but do it from the basic hydrodynamic approach. Hydrodynamics is a description of a fluid which is macroscopic. The magic of that description is, I guess, easiest to exhibit through the ideal gas laws. In the ideal gas situation you have 10^{23} particles or some multiple of that. Yet two variables will completely describe the system, pressure and temperature, or density and entropy, or something else. You can have a complete description of a nonmoving fluid, its thermodynamics, with only two variables, even though you are really describing an ensemble of zillions of particles.

For fluid flow you have to have a description that involves five variables. There are people who claim they can derive that from spontaneously broken symmetries. I do not understand it, but I think people can actually derive the number of variables it takes to describe a system from the symmetry principles. I am not one of those people.

But for hydrodynamics in a single-component, homogeneous fluid, five variables completely describe the system, the two thermodynamic variables that you would need for a gas anyway plus the three-component vector velocity field. Therefore, one can produce a closed description of a fluid by having five independent equations and, therefore, the five variables are completely determined.

Those five equations are: conservation of mass, which I have written here. The rate of change of density depends on the divergence of the mass flow; conservation of entropy. That is, the time rate-of-change of entropy plus the divergence of the entropy flux is zero. If the fluid has a non-zero thermal conductivity and viscosity and internal relaxation, then, of course the entropy is not conserved. Those are all dissipative processes. If you have a fluid that is inviscid, all these terms are exactly zero (I just wrote them up there because I

mentioned the dissipation being positive definite in the gradients); and, finally, the momentum conservation equation, where the momentum, in this case, Newton's law, is changed by pressure gradients, body forces and viscous shear.

Q: (Inaudible.)

A: Yes, that is why I wrote Euler-Navier-Stokes, because this part is Navier-Stokes. If I hide that part, then it is the Euler equation. It is basically Newton's law. It is just a question of how many phenomena you want to tag onto the source side. I can put surface tension, I can put all kinds of forces on there, if I like.

You will notice, though, that I have not solved the problem, because I have not got five variables, I have got six variables. I have s, ρ, and P.

(Transparency) (29)

But I can define an equation of state, and if I supplement those conservation equations with an equation of state, that is, that pressure is a function of density and entropy, which is true, then I have a complete description, because I can expand the dependence of pressure on density and entropy in a Taylor series, if I like.

DR. SWIFT: Why can't pressure be a function of velocity as well?

A: Pressure can be a function of velocity, but only to second order, because it is a thermodynamic variable and it has to be dependent on other Galilean invariants.

DR. SWIFT: So why wouldn't it be a function of velocity there like it was for your superfluid?

A: Because this is being evaluated at equilibrium.

DR. SWIFT: But equilibrium fluid can rotate.

A: So can I. Seth, you may have a better answer. I bet you do.

DR. PUTTERMAN: Pressure cannot, absolutely, positively cannot depend on velocity because of the Galilean transformation.

A: But it can depend on velocity squared.

DR. PUTTERMAN: It cannot depend on velocity squared or any function of velocity because of the requirements of the Galilean transformation. Equations of motion must be the same in all reference frames. In liquid helium the pressure can depend on the difference of the superfluid and normal fluid velocities, because the difference is invariant under Galilean transformation.

But I would like to say something nasty.

(Laughter)

A: Let us say you have earned it, but I am armed, Seth, be careful.

DR. PUTTERMAN: Absolutely, positively no one has ever shown that five variables are a complete description of a classical fluid.

A: Thank you.

DR. PUTTERMAN: _____ could resurrect Boltzmann.

A: Okay, good, because I never understood those arguments, but I thought people did make those arguments, and I am glad they are invalid. Divine intervention is the only way to get those five variables, is that correct?

DR. PUTTERMAN: Yes, that is Landau's method.

(Laughter)

A: Ask God.

(Transparency) (29)

Given that we have a complete description, we can then solve the equations if we can solve the equations. One way to solve the equations is to expand the solutions in a series, what is called a perturbation expansion. This now gets us to the heart and notation of nonlinear effects and why we call certain things second order.

The pressure, we can say, is some ambient value plus some value

that depends on linear deviations from equilibrium, and then some value that depends on deviation from equilibrium squared, et cetera. That is what this ordering means, the 0, 1, 2 subscripts. For our problems here, we will let v_0 = zero, so we are not going to talk about flowing fluids at this point.

(Transparency) (30)

So given that expansion, we can linearize the continuity equation. Since ρ_0 is a constant, its time derivative will vanish. The first derivative would be $d\rho_1$ dt, and now the divergence of ρv to first order has a term that is ρ_0, which is non-zero, and v_1, and that is the lowest order term that is in there.

There is another term, which is $\rho_1 v_1$, but that is the product of two first-order terms. There is $\rho_0 v_2$, but that is a second-order term. We truncate all those terms and in the next viewgraphs I will show you what it was that allowed me to throw those terms away for the linear solution -- that is what I paid for making that simplification.

Similarly, I have linearized the Euler equation. I have kept only the first-order contributions, v_0 = zero, insignificant, and I have taken p_1/ρ_0.

I can also then truncate the equation of state (if we assume it is adiabatic) to include only terms that are linear in deviations from equilibrium density. This is what you have probably done in most of your undergraduate courses in acoustics. If you take the time derivative -- again, we do the same trick we did with the shallow water case -- we take the time derivative of the continuity equation and ρ_0 times the divergence of the Euler equation. Put them together, supplemented with the equation of state, and we have a wave equation for density deviations in terms of the sound speed, and it is homogeneous. It is a second-order differential equation in the first-order variations in the fluid density.

Solutions: We could have substituted for ρ instead of p, and p instead of ρ or velocity. We would end up with any of these field variables at the linear level being traveling waves, or superpositions of traveling waves.

(Transparency) (31)

What we did to get to that point is we threw away a whole variety of higher order terms. When we linearized the equation for continuity, we kept d ρ, dt, but we threw away v_1. grad ρ_1. We can substitute for that, because we know gradients generate a jk dependence and time derivatives a jω dependence.

We said, in throwing away the numerator in favor of the denominator, what we had to do was assume that v_1 was much smaller than c_0 -- low Mach number.

If we compare all of the terms that we threw away, here in the Euler equation we threw away the v.dv term. That is a $v_1 2$ k and this ρdt is an $\omega\rho$ -- again, v/c. The same thing with the second-order term in the pressure gradients. The same thing compared to both the pressure gradient term we kept and to the acceleration term we kept. All these generate the fact that we were making an assumption that v_1/c_0 is small, the linear approximation.

Just for convenience, probably the convenience of Greg Swift, at the linear level we can also multiply v_1 by the Euler equation and multiply p_1 into the continuity equation, add them together, and you can see you get a continuity equation. The divergence of a flux is equal to the time derivative of a density. We get this term, which we recognize as kinetic energy density, this term, which we recognize as potential energy density, so time derivative of the energy is the flux of the intensity, the divergence of the intensity. That came straight out of linear theory, but it generates second-order terms as products of linear terms.

(Transparency) (32)

Now these things are starting to look like an eye chart. Let us keep all the terms. You notice the equations are getting smaller, because you cannot get all the terms on one line. There you see both first- and second-order terms kept in the continuity equation, all of them. In the Euler equation, both first- and second-order terms, the first-order velocity contribution, second-order velocity contribution, density-velocity products, velocity-velocity products, the works.

And we have expanded the equation of state out to second order, so there is dp/d ρ times p_1, ρ_2, s_2, all the possible combinations, except we are going to again make the adiabatic approximation and let $s_1 = 0$, $s_2 = 0$, all $s = 0$.

If we do that, we can collect the terms together and the first-order wave equation is in there and it cancels out, because we know that d2 ρ_1 dt2 -$c_0$2 times the Laplacian of the p_1 is 0, so we get rid of all the first-order stuff and what do we end up with?

We end up with a wave equation for strictly the second-order contributions, but it is no longer homogeneous. It has, on the right-hand side, a source term whose source strength depends on the product, the square of the first-order fields. A lot of people like to refer to this as the Lighthill equation, or the Westervelt equation.

I can take, again, the thermodynamic derivatives and reexpress this result in a term that is now familiar to you, because you remember this guy, this is just the old Grüneisen coefficient -- convective term, constitutive term. So bang, we find that the second-order fields are driven by quadratic combinations of the first-order fields.

(Transparency) (34)

This allows us to take a completely different view of distortion as we are revisiting this weak shock theory, because now we can take the first-order field and square it. If we assume that the first-order field is ρ', some amplitude, cos ωt, and we square it, we end up with the 2f component and a dc part. The dc part is the last section of the lecture.

What we see is that when we double this we have got twice the frequency and twice the wavenumber and, therefore, its phase speed is just the old thermodynamic sound speed. Therefore, this source term is in resonance with this wave equation.

The way that we like to look at this, the way that it is convenient to look at that, is that source term is a virtual array. That is, if I have a loudspeaker generating p cos ω t minus kx, out here the square of that is generating a term that is driving that equation all through this medium. There is an infinite linear array of these virtual sources. I

am in one dimension, but we will soon get out of one dimension.

The phase speed of this contribution is the same, so when this pressure is generated, it propagates in this direction and this direction (we will worry about that in a moment), and when it gets to this point, this source is in phase with the wave when it gets here, so it adds coherently, and then this wave plus this wave is twice the amplitude, it travels to this point and it picks up another factor, so now it is four times the amplitude. You have got a linear endfire array antenna, except you do not have any real sources out there. What you have is a bunch of quadratic combinations of your first-order field.

So what we see here is the Earnshaw solution repeated from the hydrodynamic view. We get the linear growth of the second harmonic because the source terms add coherently to become an antenna for that second harmonic generation. But it also gives us a convenient way to deal with another problem, which is not self-distortion but wave mixing.

That is, there is no reason why I cannot drive that loudspeaker with two frequencies. Therefore, the ρ will be some ρ" times cos ω_1 t minus kx 1t plus cos ω_2 t minus k2x. Again, if we just take a look at what happens when you square that trigonometric function, you get a constant term plus two times the first frequency, which is shocking of one of the waves, two times the second frequency, which is the shocking of the other wave, but you also get the sum and the difference frequencies as you generate two new frequencies which are the sum and the difference in addition to the second harmonic. If both of these beams being sent from the same speaker and are colinear, all the phase speeds match.

If we have reasonably large amplitudes, Goldberg numbers on the order of one, so we have got a fairly high degree of nonlinearity, but it is not really shocking up, and we have a high enough amplitude that the wave mixing will generate virtual sources of nonnegligible strength, then the virtual array allows us to generate a difference frequency that has the directionality that is determined roughly the attenuation length of the "pump" waves. This is a wonderful thing, because, as you know from linear acoustics, to get a directional beam, the directionality depends on the ratio of the wavelength to the diameter of the transducer.

If we take two very high frequencies of slightly different value, you can generate two long beams that attenuate over a fairly reasonable distance. In that distance, the difference frequency will grow. It will grow as though it were generated not from an aperture this size but an aperture that size. So it allows you a means of using nonlinearity to generate a very long directional low-frequency source without having to carry, on the hull of your boat, a transducer that is 40 meters in diameter.

The ability to take two sound sources of slightly different frequency, higher frequency but a small difference, allows you to generate highly directional low-frequency sound waves from a very small aperture, but not with particularly efficiently.

(Transparency) (35)

What you basically have are two waves that attenuate over some length, as I have shown here schematically. The self-distortion, the shocking of those waves, will occur fairly rapidly, but because it is at twice the frequency and attenuation goes like frequency squared. They will decay four times faster after you get out of the region of interaction, since the driving goes like the amplitude of the linear fields squared. Once this gets down to half its value, it is generating only a quarter of the amount of virtual array, so it falls off rather quickly -- like so -- for the sum frequency.

(Transparency) (36)

But the difference frequency will grow linearly until those things peter out. It will asymptotically reach some value, but because it is at a frequency 10 or 50 times lower, its attenuation is 100 or 2500 times more slowly and it propagates far out of that virtual array region. This is an explicit violation of what we have in linear theory, which says that when waves interact, they interact only in the region of interaction and when they leave the region of interaction it is all over. Nothing happened and if you did not see them interacting, you would not have any evidence that the interaction occurred.

Here you explicitly have something that goes far beyond the region of interaction. It was not there before the interaction. It is due to nonlinearity.

Q: Does this get used?

A: I cannot answer that. Who would like to answer that?

Q: Yes. Depth sounding. It is commonly known --

A: By Tom Clancy, or who are you talking about here?

(Laughter)

Q: What about dispersion?

A: Exactly. That is just where we are headed now, phase matching.

(Transparency) (37)

If we follow the method that I am supposedly following here, the way to understand this is to take a look at sources that do not phase match. So if we take two beams and we interact them at some angle other than zero, then we get what I call the scissors effect.

You have two waves moving like this at some speed but their intersection is moving at a much higher speed. The way to show that mathematically is to show that the phase speed, say, for two waves, k_1 and k_2, of equal frequency and equal wavenumber, but at some angle other angle theta, the phase speed is going to be the ratio of ω_1 plus ω_2 of the sum term, and k1 plus k2 of the sum term. But k is a vector, so this k is going to be shorter. It is going to be the modulus of k times cos of theta over 2, the half angle between the pump beams, and since cosine is always less than one, that speed is always going to be greater than c.

Therefore, you do not have phase matching. What you have is the virtual sources creating sound in phase and out of phase and you just get a mush. Nothing propagates out of the region of interaction, except for the very last part that did not get canceled by the part right before it.

But there are circumstances where a non-zero angle will lead to resonant nonlinear generation. One of the obvious cases is when you have dispersion.

(Transparency) (38)

In a dispersive medium the speed of sound can be a function of frequency. You can have dispersion where the speed of sound increases with frequency and you can have dispersion where the speed of sound decreases with frequency. If you have a speed of sound that increases with frequency, then the sum frequency will require you to have a non-zero interaction angle in order to satisfy the resonance condition. So the mixing is allowed for the sum frequency.

If, on the other hand, dispersion bends down instead of bending up, then that particular process is not allowed. Bruce is going to talk about that, because these are, in some sense, selection rules for whether or not three-wave mixing will occur or not occur.

DR. HAMILTON: "Allow" might be a little strong. It is just not as efficient. You will still generate that sound but slightly less than linear growth.

A: Yes, it will oscillate. It will grow and then it will eventually be out of phase with the sources ahead of it and it will get cancelled. There will be constructive and destructive interference.

DR. HARGROVE: In fact, they suppress second harmonic generation in some very high-frequency acoustic microscopes by pressurizing the working fluid into the range where it bends the wrong way to feed energy into the second harmonic.

A: Right, because you cannot satisfy the phase-matching conditions.

If I were a solid state physicist, I would have taken that equation that you saw there and put h-bar in front of absolutely everything, since it will all cancel out anyway. I will tell you that the momentum of the wave was conserved as well as the energy, h ω_1 plus h ω_2 equals h ω_1 plus ω_2. Then you think, God, this guy really understands his solid state physics, because I have conserved phonon momentum and phonon energy, but it is garbage. It is just phase matching.

You have to watch like a hawk where Planck's constant enters. Is that correct, Professor Putterman?

DR. PUTTERMAN: That is what daddy says.

(Transparency) (39)

A: There are some media that have two sound speeds, or more than two sound speeds, four sound speeds in the case of superfluid helium. We know from the talks that took place earlier today that solids support both shear waves and longitudinal waves. The square of the shear wave speed is the shear modulus over the density, and the shear modulus is related to the Young's modulus -- E. I have written it here -- by Poisson's ratio.

For most materials Poisson's ratio is between about a quarter and a half and, therefore, from stability requirements, the shear wave speed is less than the longitudinal sound speed. If one were to take a pair of transducers in, say, aluminum, which has a Poisson's ratio of almost exactly one-third, and chose to take those two shear wave transducers and aim them at one another at a half-angle of 52°, which puts the phase velocity of the two interacting shear waves up to the longitudinal wave speed, you will get conversion of shear waves to longitudinal waves. This has been shown by Rollins, Taylor, and Todd on a block of aluminum that was machined to the right dimensions and the shear transducers were put here and longitudinal waves were detected here.

So it is possible to have mode conversion, that is, coupling between two different modes of a system that are mediated by a nonlinear interaction.

Q: (Inaudible.)

A: Yes, you then end up with a wave equation for longitudinal waves that has the product of shear waves in it, ρ_1 and ρ_2, except shear waves do not have any density variation, so there will be a different equation. You would have a nonlinear source term. That is called nonlinear resonant mode conversion.

DR. SACHSE: This, in fact, is also a way of measuring the nonlinear elastic properties of the material. You can measure residual stresses, for example, by measuring the amount of this mixing.

A: It has a lot of practical applications for, as you say, finding

nonlinear elastic constants -- possible, not necessarily practical.

DR. ANDERSON: If you refract a shock wave off the surface of a solid, it goes into another shock wave at a different angle plus the longitudinal wave. Is that because there is some nonlinear conversion?

A: No, that is true at the linear level. When you take, say, a compressional wave, like a longitudinal wave, and you whack it up against a free surface, in order to satisfy the pressure release boundary condition on that surface you have to superimpose both longitudinal waves and shear waves. That is the reason you get mode conversion at a surface. It is not necessary to impose nonlinear effects; linear boundary conditions will do that. It is a good way to launch Rayleigh waves and other suface disturbances. This is the way a scorpion on the surface of the sand detects a bug burrowing below the surface.

(Transparency) (40)

Here we have second sound and first sound for two different modes that differed by an order of magnitude in their sound speed and, therefore, it is, in principle, possible to couple two second sound waves intersecting at the proper angle to a first sound wave. Generating two nice collimated beams of second sound in two different directions, as they did in the aluminum case is difficult in a superfluid. You want thos beams to be very well collimated, so you need the transducers to be big compared to a wavelength. Put it into a medium where attenuation is virtually non-existant, what do you get? Basically you are singing in the shower -- there is just reverberation all over the place.

Fortunately, to observe this effect -- I had Seth Putterman as my advisor and Isadore Rudnick as my advisor, and Professor Rudnick came up to me one day and said, "Well, you know, if you take the higher order mode of a waveguide, it corresponds to plane waves bouncing back and forth down the waveguide and, therefore, if you operate a waveguide just above cutoff, by controlling the frequency, you control the angle. Therefore, you can control the angle of the intersection very, very carefully, because frequency is determined by a synthesizer and, in principle, you should be able to interact two second sound waves and generate a first sound wave."

But one needs to do it in a waveguide, and that is one of the reasons I put this up here. You can change this from a temperature-dependent angle to a frequency-dependent angle, because the frequency controls that interaction angle and the boundary conditions in the wave quide guarantee that you have perfect plane waves.

(Transparency) (40A)

Now comes the obligatory Chinese slide for a nonlinear talk showing the waveguide in which this type of measurement was made, because you need a long interaction path, and since the liquid helium dewars are typically 6 inches in diameter, you sort of just coil the whole thing up.

(Transparency) (41)

Let me give you the English translation for simplicity. There is a photograph of the waveguide and a blowup of the thermal sound transducer. You can see this thing is just a wire heater, but the heater has been woven in a sinusoidal profile -- each half was driven 90° out-of-phase -- which preferentially excites the higher order mode of second sound -- that is, this one. It goes up and down.

By controlling the frequency, you can then cause the mode conversion to take place and this plot is part of the proof that this, in fact, works and is quite efficient. These are the mode conversion frequencies from about 1 K to about 2 K of superfluid helium at the various temperatures of the dots. The line is the theory with no adjustable parameters.

The reason the fit is so good is that Jay Maynard knows how to measure thermodynamic quantities to very high precision, so you could actually plug in the values that are really well-known in helium and compare the theory. Of course, if you measure the growth in amplitude, it is quadratic, et cetera.

It is possible, then, by using these nonlinear source terms, to generate other modes.

Q: Who has done that work?

A: It was done by a rather obscure Chinese physicist.

That was my thesis experiment, and the net effect of that, besides demonstrating that it was possible, was it gave you the value of d ρ n dv_n minus V_S squared. It measured the weirdest thermodynamic derivative in existence. It had absolutely no importance. It was read by a total of three people. And I am glad you brought it up, because it is important.

(Laughter)

This is important for the students. The advisers, you guys can laugh. The quality of a thesis project really depends on its educational value to a much larger degree than its scientific value. If you have a thesis project and it teaches you theoretical and experimental and mathematical and analytical techniques and all different kinds of data analysis and it turns out that it does not necessarily solve the energy crisis or replace CFC's, it could very well be an important part of your career.

In fact, this experiment did not work at all for six months. You can imagine how depressed I was during those six months. In retrospect, those were the six most important months of my graduate career. I would run the damned experiment, I would go into Professor Rudnick's office and say, "This didn't work, and this didn't work." And he would say, "Try this, try this, try this." I would go back and do that. That did not work.

So I got him to expose almost every trick he knew for six consecutive months, at least two per week, and he would talk to me and he would come up with these spectacularly interesting things and I would rush back into my lab, lock the door, and write down everything I could remember of what he said. Pull out the phone and close the door, I have this lovely notebook with every trick.

It turns out it did not require a trick at all. Throughout all of this, Professor Rudnick said, "Don't worry, it's got to work, helium is more nonlinear than water. Parametric arrays work in water. This thing is going to work. Helium is really nonlinear." So we kept working on it and working on it.

It turns out, of course, it was just the machine shop that had screwed up. The first version they made of the spiral had an enormous beach in it, a full 6 mil.

(Laughter)

I discovered that only by accident. As I was miking the thing out and putting in these wave-vector filters and crap, which were great tricks that I had to learn, but were irrelevant, so I am glad it did not work, because I picked up all these tricks.

Anyway, I "stumbled" across "the beach" and he found out that Bud Knox and the Hackers -- B.K. and the Hackers -- the guys who run the shop, they are another rock-and-roll group, they should never have been involved in research. They just blew the machining and I did not have the brains at that time to check stuff when it came out of the shop to see if it was to spec.

I mention that not to show that I was a complete idiot, but because during your graduate career, although it looks like the darkest days, in retrospect, they are the best days, when things are not working, and you have an adviser and can plant yourself in his office and spend time talking about how to solve the problems, even if the solutions were not the relevant solutions for that problem. So keep that in mind. It may make you feel better on those dark days.

(Transparency) (42)

Let me summarize at this point, give you one more break, then we will come back and do DC effects for half-an-hour. The motivation at this point is we will review and summarize the results derived since we introduced the formal perturbation approach. Then in the next section we will look at the non-zero time-averaged effects.

First of all, products of first-order fields act as sources for the second-order wave equation. As the Mach number increases, the assumptions used to linearize the hydrodynamics are no longer valid -- I showed you that.

The source strength depends on the square of the first-order fields and on the distortion parameter, the thing that I call the Grüneisen coefficient.

The cumulative effects of these virtual sources depend on the phase matching between the source and the wave propagation speed. In the absence of significant dispersion, the first-order sources have the same phase speed as the second-order contributions. If they are

colinear, things grow.

For colinear waves of different frequencies, the difference frequency has a directionality which is determined by the attenuation length of the pump waves. Very important.

In systems which support more than one mode of propagation, like solids, superfluids, or wave turbulence systems that Bruce will tell you about, the slower mode can convert to the faster mode through nonlinear processes, and this gives a way for energy to cascade. That is why this is called wave turbulence. The big vorticities create little vorticities and long wavelengths convert to short wavelengths.

DR. GARRETT: We will talk about the non-zero time-averaged effects that come out nonlinear acoustics. And I wanted to mention, before I begin, that Craig brought up an interesting comparison, which is, what we have talked about here in acoustics is actually very important in optics.

In the area of optics, exactly the same thing happens, because optics is less nonlinear than hydrodynamics, you do not get shock waves forming in light, but you do get second harmonic generation, and that is a technique that is used consistently in optics for taking red light and generating blue light, drive it very hard into a crystal where the nonlinearity causes second harmonic generation, and then you can get clever by stacking different materials together to correct for the dispersion.

If it grows and starts to cancel, you can put in a crystal with the opposite nonlinearity so that it will add to the stuff that is being brought in the other way. Anyway, there are lots of things in optics, both that I am aware of and was not aware of. The fact is that the nonlinear conversion process is by no means exclusively an acoustic phenomenon. It is, in fact, a property of nonlinearity. It is just that acoustics tends to be about five orders of magnitude more nonlinear than, say, optics.

(Transparency) (43)

What I would like to do is go back -- when we squared the first-order term, I did not discuss the DC portion. I ignored that, I was looking at only the wave effects. I would like to come back now and show you a couple of demonstrations and describe -- again, using the

same technique I have been using thus far -- it in a couple of different ways.

We know that all the first-order fields are time harmonic and, therefore, their time average over sufficiently long times is zero. Sine and cosine have equal positive and negative contributions.

But the second-order terms contain a constant which, when time-averaged, does not vanish. This was demonstrated by a gentleman I never met who did some spectacular work in the 1940s. His name was Hillary St. Clair. He was the inventor of the St. Clair generator and also did a lot of siren work. I saw a wonderful picture of a St. Clair siren in a reflector levitating pennies. That is, the time-averaged effect of high-amplitude sound field was sufficient to take something with a density nearly nine times that of water and suspend it in space against earth's gravity!

So these time-averaged nonlinear effects can be substantial, and this has been known for 50 years, although it has really reached its peak in the space program in acoustic levitation for materials processing, containerless processing. I will get to that.

The basic idea behind these non-zero time-averaged effects is the Bernoulli pressure, which you can see as a second-order effect. If you have a flowing fluid, the faster it flows, the lower the pressure. This is important for those of you who came here by means of commercial air transportation. That is why the wing is curved on the top and flat on the bottom, so if you get that v2 up large enough, you can sustain enough pressure differential to lift something the size of a 747. Always amazing to me you can lift something like that on a pair of wings just because air is flowing over them.

This is the same effect, I claim, that takes your shower curtain and splats it against your body. Now, Jay disagrees with me. Jay believes that it is not the entrainment of the air, the fact that the air is moving faster inside the shower than outside the shower. If you are a grad student and single, that is a pretty ugly shower curtain. Things from other universes live on shower curtains. (The women who take nice care of their apartments do not know what I am talking about.)

(Laughter)

Jay claims that it is the hot water inside, the air temperature is

higher and, therefore, the rises in the DC flow is not entrainment but it is the air coming up this way.

Q: Does it work with cold water?

A: Jay says that he is so old now that he cannot remember the last time he had to take a cold shower.

(Laughter)

Thank you. And it does work with cold water.

Q: Why do you have to be in the shower to observe it?

A: The reason he cannot watch from the outside has to do with an old Hitchcock movie. It has nothing to do with Bernoulli effect. Mother is not feeling well today.

(Laughter)

The fact that it is quadratic in velocity means that it does not matter whether the fluid is flowing this way or that way, the pressure goes down. Therefore, when you time average, there will be a net effect. These effects have been understood and observed for a very long time.

I was planning on bringing a Kundt tube, but I did not, and if people have not seen this, later on I will bring it and put in there. The old experiment was observed by Kundt and then explained by Rayleigh, they would normally take a tube, and at the end of the tube they take a bar, and at the end of the bar they take a violin bow. They get large-amplitude sound in this tube at a tube resonance.

The tube resonance would be demonstrated by putting cork dust in, and you see the cork dust agglomerate at the velocity antinodes, but you would see striations. In fact, there is a chapter in Rayleigh's book called "Striations."

(Transparency) (44)

These are taken right out of Theory of Sound. If you have fluid flow with two particles that are aligned in the direction of flow, then the stream lines have to go around them and there is stagnation in

between. Therefore, there is no velocity here and the particles repel, so in the directions the striations form, they block the acoustic velocity and, therefore, repel each other.

But they stand up, that is, the particles lift. The reason they lift up is if the particles are perpendicular to the flow, then the gas has to go between the particles, it has to speed up to do that in order to conserve mass and, therefore, the pressure is lower. This is known as the "Maidenform Effect," because it takes the particles and lifts them in this direction and separates them in this direction.

Only old guys who remember the commercials will know what that was about.

(Laughter)

I have not put arrows on these stream lines. The direction of that flow is irrelevant.

We can look at the same thing with a Helmholtz resonator. I live in Carmel and, as you know, Clint Eastwood used to be the mayor. So if you live in Carmel, you could be stopped on the street and if you could not do this -- (Clint's movie sound) -- you were arrested. That is a Helmholtz resonator. I voted for Clint.

A Helmholtz resonator is basically a mass-spring system, where the mass is the plug of the fluid oscillating back and forth and the restoring force is the compressibility of the gas within the volume. A Helmholtz resonator has a neck which is open to the atmosphere and, therefore, the pressure at this point must be atmospheric, because there is a big sink at p_0.

This pressure has to be lower than the pressure inside, because there is no motion in here. So if this is at p_0, then this must be a p_0 plus one-half ρ v2. The only way one can sustain an excess pressure in here is that fluid must be accelerated out. Very difficult to do, since this is closed.

In fact, I will show you, in the demonstrations, that if you put a Helmholz resonator in a large sound field there will be sufficient flow out in this direction to blow out a match and, in fact, sufficient flow to make a radiometer, one of these things you buy at the observatory

that will go around.

Q: Wait a minute, if the stuff blows out, then you deplete the pressure on the inside.

A: No, it does not blow out. What happens is that the gas at this point is entrained, so none of the gas from the outside comes in. You are always accelerating fluid ahead of this out to maintain that pressure.

You are right, if you had just this and there was no acoustic drive on this, it is an unstable situation, the stuff just leaks exponentially. But if you are putting in energy to sustain this Bernoulli effect in the neck, that gas will be entrained by the radiation pressure and accelerated out in sufficient amount to keep the pressure in here above atmospheric.

Q: Could you take an excursion back to the previous viewgraph?

(Transparency) (43)

Wasn't it just an hour and a half ago that Seth Putterman said that under no circumstances can you have a pressure that depends on velocity?

Q: No, Greg asked why it could not and Seth said under no circumstances can it.

DR. PUTTERMAN: That pressure is a function of density, that pressure that he wrote down, which is p_0 -_____. The pressure is a function of density but that function of density can be changed by the velocity.

(Transparency) (45)

In addition to forces due to the Bernoulli effect which agglomerate and separate and accelerate, one can also apply torques to extended objects. If you take a look at a disk in a sound field -- again, I have the stream lines drawn roughly here, but there is a stagnation point here, where the pressure is high, and a stagnation point here, where the fluid either goes up or the fluid goes down. At this point it must be at rest.

There is high velocity flow at this point and high velocity flow at this point. Therefore, since the pressure here is lower due to the Bernoulli effect than the pressure here, and the pressure here is lower due to the Bernoulli effect than the pressure here, then there will be a net torque that will tend to take that disk and rotate it perpendicular to the sound flow.

One can calculate, by integrating the moment of that force over the disk, that the torque is four-thirds density times the radius squared, which gives you the area, times the radius, which gives you the moment, so it has to be ρ aq times the velocity squared, times the sine of the angle, because when it is this way the torque is canceled. If it is this way, it will get twisted that way.

In fact, at this aspect it is also zero torque, but it is unstable, so if it flips a little bit this way, it will go that way, and if it flips this way, it will go that way. That is why people in Hartford, Connecticut, are not dead, because when the windows fall out of the Hartford Insurance Building, they do not come down and slice people like a deli sandwich. They stay perpendicular to the flow and you can get out of the way.

(Laughter)

I put this up here because this happens to be a Rayleigh disk apparatus. The Rayleigh disk was, prior to the advent of electroacoustics, <u>the</u> way you determined the absolute amplitude of an acoustic field. You did not have calibrated microphones and reciprocity and all these other wonderful techniques. In fact, you used to take a disk and mount it on a torsion fiber. You could measure the torsion very easily, because you would calculate the moment on the disk. You could oscillate it, and thereby determine what the torsion fiber constant was. That is how you would measure the absolute amplitude of sound using a Rayleigh disk.

This, I believe, although I cannot prove it, was the last Rayleigh disk ever built. It was built in England to investigate the counterflow of superfluids. I wanted to point this Rayleigh torque out, because when I show you the levitation demonstration, you will see that the disk rights itself and is supported in the sound field.

(Transparency) (46)

Given the hand-waving Bernoulli expression for forces and

torques, we can now go through and do it carefully. We can take the Euler equation correct to second order. We can express the pressure gradient in terms of the enthalpy, since the enthalpy is defined as the internal energy plus pV. When we take the differential, we can take this and, for adiabatic purposes, we end up with the equation looking like that.

That is convenient, because if we can get this equation into the form of the gradient, as I am headed for down here, then we can say that this is a constant. Having done that, what we are left with is a gradient in the enthalpy, a gradient in the velocity squared, but no gradient here. We can get rid of that, because in potential flow we can define a velocity potential whose gradient is the velocity.

Therefore, the gradient of this entire quantity is zero and, therefore, this has to be a constant. So this is the strong form of the Bernoulli equation. The Bernoulli equation in most of your fluid dynamics books is given for steady flow and they give you this result without the velocity potential part and say that it is constant along a stream line. You can remove that restriction by doing the trick that I just illustrated there.

(Transparency) (47)

Then we can do just what we said before. We can expand the thermodynamic dependence, we can expand the enthalpy or the heat function in terms of the pressure and the pressure squared, in a Taylor series, and then write the expression that I gave you before correct to terms of second order. I have included second-order contribution from the enthalpy, the enthalpy, the first-order pressure, second-order pressure, second-order velocity, and the velocity potential.

Since we are interested in time-averaged quantities, we can let the total deviation in the pressure be the first-order part plus the second-order part, as we tend to do. If we can look at only the time-averaged portion. Since this is a linear velocity with a time derivative, this will average to zero. This is just a constant, so I will lump it into that constant. This will time average to zero, because it is a linear contribution. The second-order contribution will time average to that quantity.

In a traveling wave, P is equal to ρ cv. When we substitute for v we find that there is no time-averaged pressure. The convective term

and the compressive terms cancel.

DR. DENARDO: The second-order pressure here, which could be considered, I think, the radiation pressure in this case. It is just the potential energy density minus the kinetic energy is a nice way to remember that.

A: Yes, I meant to do that, because I wrote those two densities down on the previous viewgraph. Thanks, Bruce.

(Transparency) (48)

On the other hand, for a standing wave, where pressure is a trigometric function and the time dependence is sinusoidal, if we plug that into the previous equation we find out that there is a time-averaged non-zero component to that pressure and it varies with twice the periodicity, half the wavelength, of the linear field.

If we have a static pressure and we place a body in that field, say, a sphere, since it is the easiest one to deal with at first, and we integrate that force over the sphere, the force will be four-thirds times the surface area of the sphere times this nonlinear contribution to the pressure times ka, which is effectively the gradient -- that is, the bigger the sphere, the larger the pressure difference between the average disk here and the average disk here.

Therefore, we get a non-zero force on this object in the field. If that non-zero force is equal to or greater than the gravitational force, which depends on the mass of the sphere and the acceleration due to gravity, as written here, then it is possible to levitate a solid object -- in this case, a sphere or a disk or something else -- in an acoustic field against the weight of gravity, which I claimed is what Hillary St. Clair was capable of doing.

With that, let me show you some video demonstrations of just the effects that I have described.

(Video)

What you will see first is a Helmholtz resonator. There is a loudspeaker here at the resonant frequency and as we crank it up you see that DC flow blowing a match out. But a more dramatic demonstration is the acoustic radiometer, which shows two Helmholtz

resonators front to back. You can see now that that thrust is sufficient to make a lawn sprinkler.

It is just Christmas tree ornaments without the little hangers on them. And the bearing is nothing special, it is just a point.

Now here we have a Rayleigh disk, so I am in this sound field, again. I am going to put a piece of cardboard hanging from the roof and we are going to excite these two loudspeakers and you can see that the disk orients perpendicularly to the sound field. We are driving a standing wave. We are at a velocity antinode here. We are in the middle of the box.

There we go. There is just a disk now, and you notice it is being made vertical by the Rayleigh effect. These are the pressure antinodes and it is being driven to the pressure node; the velocity antinode.

What I will do at this point is erect the tube so that gravity will tend to bring the disk down. It will levitate a little bit low in the tube, but it will maintain itself there. You can see it is sitting there just below a velocity antinode. Then I am going to reduce the sound wave and trap it at successively different levels.

Notice it is oscillating here, and that the oscillation is not dying out. There we go, entrapped it there. See it waving and jumping up and down? You hear the modulation in the amplitude and the waving around there? That is a very important effect.

I backed off on the sound amplitude until it lost lock and then I jammed it back up and caught it. This looks like about the 5, 0, 0 mode.

I will put up here the last demonstration, without explanation. You can worry about both that oscillating business and about this little demonstration in whatever discussion session happens a week from Thursday. I will tell you that two modes are being excited in here. What we have just done is changed the relative phase of those two modes.

We have got a half-wavelength in this direction, a half-wavelength in this direction and these two modes are at the same frequency, because these lengths are the same but the phases are shifted by 90°.

We can do the same thing at the second harmonic with four cups. And change the phase, it will rotate the other direction. There we go, I switch the phase, again.

(Transparency) (49)

Final review and summary: the waveform life cycle. In the linear approximation superposition is valid and waveforms are stable. A monochromatic wave remains monochromatic. You put in a frequency, it propagates, it attenuates.

In the nonlinear case, if nonlinear terms are included, a large amplitude-sine wave or any other wave will distort into a stable repeated sawtooth shock wave.

Energy is dissipated in the shock front and as that shock wave attenuates -- it starts out, it becomes a shock by gaining high-frequency components, then it dies by losing high-frequency components, and, finally, all the high frequencies are dissipated and you are back to your monochromatic wave and it attenuates.

These phenomena are due to the effects of the sound wave on the medium. The way the sound wave affects a medium is that it changes its local sound speed by convection and the constitutive relations.

When we make an expansion in the weak shock limit we see that the wave equation for the second-order variables is driven by quadratic combinations of the first-order variables. If these first-order variables phase match to the second order wave equation, you get linear second harmonic growth and, in fact, if there are two different modes that can possibly exist, then by changing the angle of intersection we can force that phase matching.

Second-order effects lead to sum and difference frequencies. One way to think of this is if you take a single sound wave, the sum frequency is 2f, which is the start of your distortion, but the difference frequency is DC, and that DC term will lead to forces and torques that are not time harmonic.

I hope that this prepares you for the next series of Putterman, Denardo, and Maynard, some of the other UCLA alumni who will be

addressing you over the next two days. Thank you for your attention.

DR. MARSTON: On your last transparency with the levitation equations, lest there is any confusion, this equation 112 is the same equation or result that you get when you take what I called, in my lecture, normal traction, the radiation stresses, project that on the x-axis in your notation and integrate over the surface.

A: Exactly, it is another way to look at what Phil was talking about, where he was taking the second-order pressure, is taking the time-averaged value of that over the volume.

DR. MARSTON: Yes, but the specific issue at hand is to take a rigid sphere if you have no volume changes, but if you take this radiation stress tensor and take the normal part, the projection of force on the x-axis and integrate over the surface and get that equation.

A: We assume it is an incompressible solid levitated in a compressible liquid. Phil is doing more interesting varieties which allow you to have a compressible solid in a fairly incompressible liquid. That is a more general result.

DR. ATCHLEY: I just want to point out, for those of you who have not read Lord Rayleigh, that that term "graver" might need some explanation.

A: Rayleigh always called the lower modes the graver modes. He also uses a for the speed of sound. Greg also observes that convention.

DR. HAMILTON: I wanted to follow up on your comment about nonlinear optics earlier. You made a comment, I think, that the nonlinearity in nonlinear optics is much weaker than that in acoustics. In general, it is the opposite when you --well, Maxwell's equations are linear, but once you introduce a medium that will give you the nonlinearity, typically it is very strong nonlinearity. The reason you do not get the shocks is because your dispersion is so high, which you were mentioning.

For example, as you said, when you do your wave-vector matching, you get waves at different angles. You can resonate and that is how you can get efficient second harmonic generation and that

can be done over fairly short path lengths. But in acoustics, even though you have very weak nonlinearities, since you have no dispersion you have a lot of time for those cumulative effects to come up and give you the shock wave.

I think crystals for nonlinear optics can have very high nonlinearities associated with them and are far more efficient than we --

A: You mean the Grüneisen coefficients are bigger than three or four or five.

DR. HAMILTON: They can be much higher, orders of magnitude, in optics than you can get with liquids and gases, at least, under standard conditions.

DR. MARSTON: It depends in some sense what you are comparing. What is large in optics is energy densities are enormous.

A: Yes, exactly, and that is why I made the claim that compared to the energy density in acoustics and the nonlinear effects as opposed to the energy density in optics and the nonlinear effects that the nonlinear effects in acoustics were stronger. But that is a prejudice.

What you are saying is, okay, throw away the amplitude part and just look at the Grüneisen coefficient. The energy densities in laser are kilowatts per square meter and megawatts per square meter. In acoustics you do not need those kinds of levels to generate those kinds of effects. So it is a somewhat prejudicial choice. I just happened to pick another.

Q: _____streaming, so none of that Helmholtz resonator is blowing out candles?

A: No. There is an important distinction. Streaming is a circulation that is generated by the acoustic fields. By the Kelvin circulation theorem -- the stuff I have done here nonlinearly has been inviscid and, therefore, I cannot create streaming flows, because I cannot break the Kelvin circulation theorem, unless I happen to have a viscous interaction. So there is no level, you cannot take this to any order. At arbitrarily high order you will not get streaming, because you do not have viscosity. In order to get streaming you must have viscosity.

Greg will be talking about streaming in some strange sense. He will be talking about thermal streaming, which, again, takes place only in a dissipative system. But streaming is not included in what I showed here, and the easiest weasel to ride out of that argument is that I would have to violate the Kelvin circulation theorem because it is potential flow.

Q: You were saying there is a vorticity argument, you will only get vorticity at boundary, if you have the gas equations there --

A: Did I say vorticity?

Q: Yes, you said there was a vorticity argument.

A: No, no, when we were talking about the shock waves I said how do you handle vorticity? Vorticity has a singularity in it just like shock has a singularity.

Q: Oh, you are talking analogy?

A: Yes. Thank you.

SONOLUMINESCENCE

DR. PUTTERMAN: May I show you the first slide, please.

(Slide)

Let me show you something amazing that we have been having fun with for the last few years. You have here a flask. This is a quartz flask. You can buy a spherical flask about yay big from the chemistry supply house wherever you are. On the walls of the flask are mounted piezoelectric transducers. When you put an oscillating voltage through the PZT it oscillates. Typical frequencies of oscillation are in the range of 30 kHz. That is just so you can do this without wearing ear muffs.

This is filled with water and in the center of the water is a trapped bubble of air. Even though this is illuminated with an incandescent light, the trapped bubble of air is making enough light by transducing sound into light so that you can actually see it in a real-time photograph. This is the phenomenon of sonoluminescence from a single bubble and this is what I would like to talk about today.

I would like to tell you why I find the phenomenon so amazing, what interested us in looking into it, and what are the experimental properties of this effect? We, of course, have looked at it and it has gone through many stages of technological change, but this is the basic effect. You can all set it up in your labs and one of the themes of today's discussion is please go out and set it up, do anything, whatever you do will find something interesting and amazing.

We have the upper hand in acoustics. This is a great day for acoustics, because in order to measure the properties of the flashes of light which are coming out you need to use photodetectors, which are faster than the high-energy physicists would have used at the SSC. All this to study the properties of an audible sound field.

One small thing about this is, of course, if you can change sound into light, if you can turn up the sound field you can get even more light out, right? So every day I go in the labs pestering the students and the postdocs, "Please turn up the sound field, get more light. Get more light." So one day they turned it up and the flask started vibrating, I got these Love waves running around the flask and everything was nonlinear, then the Love waves started running down the table and into the floor. I show you in the next slide what happened when they got the sound wave really up.

(Slide)

That is nonlinear physics.

(Laughter)

(Transparency 1)

I heard about sonoluminescence about six years ago from a good friend and colleague of mine, Tom Erber, at the Illinois Institute of Technology, where he is a professor of math and physics. He likes to do science by taunt and tease. If you are his friend, he will find the weak points in what you know and press on them, especially if it is an area where you think you know something.

One day at the coffee house at UCLA Tom Erber said to me, "You're interested in fluid mechanics, right?" I said, "Yes, I've been working on nonlinear fluids my whole life." He says to me, "Well, if you know so much about fluids, then please explain to me how you can turn sound into light?" I said, "That can't be done, you can't turn sound into light." I mean, if you think in terms of condensed matter physics, as Steve Garrett told you yesterday about two waves making a third wave, like the en fire array, in the language of condensed matter physics, when you multiply everything by bar, it is called a three-phonon process. For acousticians it is a three-wave process.

If you go to Landau's and Kolotnokov's books on phonons and liquid helium -- when the dispersion goes the wrong way, as Steve mentioned, you got four-phonon processes, two phonons hit and make two coming out. Then if you go to some really complex dispersion where it sort of curves up here and then curves down over there, you need a five-phonon process. Landau and Kolotnokov did a five-phonon process.

To turn sound into light you need a 10^{12} phonon process, and I said at the time, "This just can't go." It was, to the largest extent, my lack of scholarship that I was so amazed on hearing about this effect. I was convinced that to see sonoluminescence they must have been putting sound through nitroglycerin or something like that.

(Laughter)

He said, "No, no, do your homework. Two Germans did it 50 years ago." So with Ritva Löfstedt, then a senior at UCLA, we went back through the old literature to study the discovery of sonoluminescence and see if the phenomenon was for real or whether it was some artifact.

(Transparency 2)

I would like to show you in the first viewgraph my

motivation, why I was so excited about sonoluminescence and why I was so skeptical about it, and then to come back a little bit to the history and then to proceed headlong into the science.

In my mind I put sonoluminescence into parallel with the photoelectric effect strictly in terms of input and output. In the photoelectric effect you have some piece of metal, a photo cathode, light hits the metal and ejects an electron. The energy of the electrons is on the order of electron volts. If you think of the light as being a classical wave spread out over space, then you take a typical 1 watt light source at 3 meters from the photo cathode and this light delivers 10^{-4} eV per atom per second.

You can see that in order to eject an electron, the light would have to be on 10^4 seconds in order to deliver enough energy to one atom (we presume the electron comes from an atom) to create the free electron.

In a typical photoelectric experiment, when you turn on the light you immediately get some electrons coming out. Something is concentrating the energy. Sonoluminescence is quite similar in terms of input-output. A sound wave which you can imagine as being uniformly distributed in space, say, the pressure amplitude is about 1 Atm, is in a fluid and photons come out.

Knowing our quantum mechanics, we know the photons have energy in the range of electron volts and if we calculate the energy, which I will show you in a moment, the sound field with 1 Atm of amplitude, that is about 10^{-11} eV per atom. Knowing that photons come from an atom, you have here a situation where somehow the fluid must have concentrated the acoustic energy by 11 orders of magnitude in order to make light. In fact, this situation was ruled out by the people who won the Nobel Prize for quantum electrodynamics.

(Transparency 3)

If we go to Tomonaga's book on quantum mechanics, one of the best books to use as a textbook for elementary quantum theory, he shared the Nobel Prize with Schwinger and Feynman for quantum electrodynamics, he remarks about the photoelectric effect.

He says there is no classic way to explain the photoelectric effect. He says right here, in "the wave theory of light, energy is considered to be distributed continuously in space. Hence it takes quite a long time for the atom to absorb" enough energy to kick out an electron, "unless the atom has a sort of supernatural method of collecting energies instantly from an extended region of space."

Sonoluminescence is exactly the supernatural process which was ruled out in order to say the only thing left is quantum theory, we have got to bring in h-bar and quantize the light into photons, which are bundles of energy.

Key to my prejudice about why sonoluminescence would be so amazing is the energy density of a standing wave.

(Transparency 4)

Permit me to show you something that is really very elementary, but it was this thinking that compelled us to look into sonoluminescence. If we have a standing sound wave with velocity amplitude V' -- and as you all know, V' is related to the pressure amplitude of the sound field through the acoustic impedance, U is the speed of sound -- then the energy density is ρ V2, which is the sum of the kinetic and the potentials, that is why there is no half. And this is one-quarter the acoustic and pressure amplitude squared over ρ times the speed of sound squared.

If we now put in the known values and convert from ergs to electron volts and look at the amount of sound energy within a cubic Angstrom, which is the region of a molecule, you find that the energy of vibration of a uniform sound field is about 10^{-11} eV per $Å^3$.

Therefore, if light comes out we have a prejudice that the light came from an atom or a molecule, it came from a region characterized by cubic Angstroms. If the sound energy is only 10^{-11} eV per $Å^3$ and we are getting light out -- I am using some current knowledge -- in the ultraviolet, 7 eV, then something has concentrated the ambient or average sound level by 12 orders of magnitude in order to make the ultraviolet photons.

The problem of sonoluminescence is to understand and figure out how does the sound energy concentrate by this many orders of magnitude so as to deliver a trillion phonons to make a single photon coming out of the fluid. While the body of experimental insight is building, I think one of the themes of today's discussion is that we do not understand what the mechanism is, we do not understand what are the key aspects of the process, and I hope through this discussion to lead us to at least a definition of what are the phenomenological issues that we do not know and why.

Looking ahead a little bit, the lowest pressure amplitude at which sonoluminescence has been achieved was done recently by Brad Barber. He got acoustic pressures of about a third of an atmosphere, which means this is now 10^{-13} eV, so you are up to 13 orders of magnitude amplification of the sound field.

(Transparency 5)

Speaking of that, let me show you my cohorts --actually, the perpetrators -- Brad, Bob, and Keith are here and Ritva is not here, so I thought I would show you a couple more slides.

(Slide)

This is Ritva Löfstedt, who went back through the old literature and is now working on the theory of sonoluminescence.

(Slide)

That is Tom Erber. Can't you tell how he likes to do the physics?

(Slide)

This is an unsuccessful experiment. Hamamatsu was in the process of shipping a streak camera from Japan to Iraq in order to make up for last year's deficit. They found it very useful to have it stop at UCLA on the way {joke}, so we used the streak camera here to try to resolve the width of the flashes of light coming out of the sonoluminescence, which, as I will show you, are so short that we really do not have a reliable measurement of the flash widths yet. So this is an example of an unsuccessful experiment that we had set up.

(Transparency 6)

Here is an outline of the talk: sonoluminescence from a single bubble. The experimental phenomenology and motivation is what we will go through probably before and perhaps a little bit after the first break. Then we turn to the theory, but this will still be most of the talk. The theory is attempts to lay the groundwork for what we do not understand by doing the theoretical aspects that apply to the problem for all the different obvious cases. Then showing that nothing works and then we have our marching orders.

First we have to go through Navier-Stokes hydrodynamics plus mass diffusion and we will consider the processes which store and concentrate acoustic energy. At low Mach number there is the Rayleigh-Plesset equation, which shows how the bubble of air oscillates in a sound field at low Mach number.

At high Mach number we have an imploding shock wave and we will discuss imploding shock wave theory a little bit. Then there is trapping a bubble. The trapping is linear in the sound field for sonoluminescence, not quadratic. The DC radiation pressure is linear in the acoustic drive, that is the key point there.

We have a virtuoso sound field in sonoluminescence. First,

the sound field traps the bubble against gravity. Then the sound field causes huge amplitude oscillations in the bubble which turn sound into light. Thirdly, the sound field causes the amplitudes of motion of the bubble to be just right to keep the bubble from dissolving into the liquid, because normally a bubble would want to dissolve (that is directified diffusion).

The sound field does all of this for free. It wants to turn sound into light.

The effect of impurities: You saw a picture there of the purified water giving off the light. It appears that if you add a few drops of ethyl alcohol -- this is the big difference between the research in Mississippi and UCLA, in Mississippi they study the effects of glycerine and at UCLA we study the effects of ethyl alcohol -- add five drops of ethyl alcohol to the bubble and that cuts the light out. That is the coming attraction there.

Maintaining a spherical bubble: Why do we have such a nice spherical bubble for sonoluminescence to turn the sound into light? At Livermore they desperately tried to get spherical collapses of bubbles containing deuterium and tritium in order to see inertial confinement fusion and they want to know why we get nice spherical collapse for free.

The light-emitting mechanism is unknown. I can chat with you about black body and thermal Bremsstrahlung, that is, light which comes off of accelerating charges.

Then we have the unsolved problems. Why water? Until today the only system in which a single-bubble sonoluminescence can be seen is water.

Then bubble technologies: photonics, medical, chemical, inertial confinement fusion.

(Transparency 7)

Let me show you the setup, which is a takeoff on the important discovery of Felipe Gaitan that a single bubble could be trapped in a liquid and made to give off a continuous stream of light. Tom Erber, through his taunting, teased me to get interested in sonoluminescence. Brad Barbor set up an apparatus at UCLA and then I went traveling around to speak to the keepers of cavitation theory of the day and I went up to visit Anthony at the P.G. School to learn what he could tell me about cavitation then I went to Mississippi to see Larry Crum's group.

There I learned that they had actually achieved to get away from the original many-bubble, transient, irreproducible sonoluminescence, in fact, the kind of stuff that goes into all the sonochemistry, and get to what a physicist revel in,

getting the reductionized state, the hydrogen atom, the single bubble trapped in the sound field to give off light so that it could be studied.

We at UCLA got very excited about that and, with Felipe's advice and help, we got a single bubble going, and I would like to show you the experimental results that came out of that. Here is the typical arrangement, as you saw in the photo, a flask, a trapped bubble of air, the bubble of air is giving off flashes of light, one flash of light with each cycle of the sound field.

It is going to be important to have a certain straw man picture in mind as we go through the sonoluminescence. The straw man picture is the Santa Ana condition down south. It is the following.

If you have air and the air goes from a higher elevation to a lower elevation, it gets compressed adiabatically. An adiabatic compression heats the air. It is hot in Los Angeles because the wind blows from the desert down to the basin, and it is not hot because the desert is hot, it is hot because the desert is one mile above the basin, and as the air drops one mile it compresses, and $T(V)^{\gamma-1}$ is the constant, so the temperature goes up 30° F per mile of drop. (V is the specific volume of gas).

In sonoluminescence there are lots of miles of drop, because you have a small bubble of gas in a sound field. During the rarefaction part of the sound field this trapped bubble of gas on the order of 50 μ in radius expands and then during the compressional part of the sound field the bubble experiences a runaway collapse, first calculated by Rayleigh in 1917 and the bubble crashes in, going supersonic.

This collapse takes it from about 50 μ maximum radius down to about 0.5 μ collapsed radius -- that is a 1 million change in volume -- and the heating which accompanies this, as a straw man, just for a picture of going through the viewgraphs as we discuss the experiments, is what leads to the emission of light, a kind of incandescence.

This conversion is remarkably efficient, as you will see. You get a lot of light out. If you keep that in mind, a bubble expanding, storing sound, and then concentrating as it collapses, and then it does in sync, this explosion take place in sync with the sound field, every cycle of sound makes another flash. If you are running at 30 kHz, that is 30,000 flashes per second.

(Transparency 8)

Let me get to the first set of experiments that were

carried out at UCLA. Thinking as physicists, our first concern was what is the temporal width of the flashes of light? Physicists are very anxious to know about the time scales in a problem. It became very clear after a couple of experiments that we would need the fastest photomultiplier tube in the world and we need the fastest oscilloscope in the world.

I called up Hewlett-Packard and told them that we were thinking of buying their $50,000 oscilloscope, but we wanted to see if it worked.

(Laughter)

And I told Hamamatsu we were thinking of buying their photomultiplier tube and we wanted to see if it could handle our particular application, and these things showed up.

Then when the salesman showed up to reclaim it, somehow Brad was never there. I think that is when Brad got started on the midnight to 8:00 a.m. shift.

Let me try to tell you how to read the response of a photomultiplier tube. By convention, the response is negative. Light is coming in here and there is a signal -- the zero, of course, is no light. The key quantity of a photomultiplier tube is the rise time.

The rise time is the waveform you get, the time it takes to go to the maximum negative value for a delta function source of light. If this photomultiplier tube is hit with a delta function of light, zero width, then it would take 172 psec -- right now a picosecond is a millionth of a millionth of a second, 10^{-12} sec. So it would take 172 psec for the photomultiplier tube to roughly come to its minimum value here.

The width of the actual pulse that you are looking at will then cause, if it is much wider than a delta function, this to become longer and perhaps come over here. Here what is being compared is a picosecond laser pulser, which is actually 35 psec wide, which is the dashed line, and the sonoluminescence.

This is the time here for the tube to recover and come back and be ready for the next helping of light. Everything is happening right here. Notice that the rise time of the two is indistinguishable. This means that the response of the tube to a 35-psec flash of light at the sonoluminescence is basically indistinguishable, in that sonoluminescence is really a very short flash, indeed.

Secondly, you would have thought that the pulser which you buy for $30,000 would be the good Greens function of the system, but it looks as though sonoluminescence is actually the good characteristic, because you notice the sonoluminescence

returns faster than the pulser, which to us indicated that the pulser had an after-pulsing, after the main pulse there was another bit of light leaking out. The $30,000 laser pulser has a dirty pulse, so we used the sonoluminescence to tell Hamamatsu that their pulser had some after-ringing characterizing it. This indeed is the case.

It is also interesting to point out that there is a technological application here. The peak power of this picosecond laser pulser is measured in about 1 mW. The peak power of the sonoluminescence, this particular one, is approaching 30 mW. So we have here a light source which is putting out light flashes at a fraction of the cost of an equal power to a $30,000 picosecond laser pulse.

Q: Does the laser pulser have all this light directed in one direction?

A: Exactly. For some people that is an advantage. We actually heard from some high energy physicists at MIT who wanted to have short pulses of light to calibrate an array, calibrate the time of rise of an array of photomultiplier tubes. In that case, as you pointed out, having a 4-pi light source lets them synchronize over a large solid angle, so it has different applications. It would be nice, though, to figure out a way to get this all in one direction.

Q: How does the energy intensity compare?

A: The intensity for the laser, of course, is huge by the ratio of R2 to square millimeters. Let us say, for argument's sake, there is a factor of a thousand there.

(Transparency 9)

The next quantity to measure is if you are getting one flash of light with each cycle of the sound field, what is the jitter in the time between flashes? How good is the synchronicity of the successive flashes of light coming out?

With a time-interval meter we measured the histogram of the time between flashes. As you can see here from the scale, the histogram is very tight; in fact, it is limited by the electronics of the stopwatch which is used to measure the jitter in the time between flashes.

You are getting one flash of light with each cycle of the sound field and the flashes are coming out with a synchronicity on the order of 50 psec, a jitter of about 50 psec.

(Transparency 10)

As I said, anything you do causes something interesting to happen. One day the students were not in the lab, they were

off seeing some great movie, and they came back with a great idea for an experiment -- I think the movie was "Wierd Science", no it was actually "Real Genius" -- and they said we should cool the water and see what happens.

After that the water was cooled and the number of photons in a burst went up from about 35° down to about -6° supercooled. It goes up about a factor of 1000. To this day we do not understand why cooling the water in which the sound is being transduced into light, why that should have such a spectacular effect on the emission of light.

(Transparency 11)

The next quantity, you would say, well, what is the spectrum of light coming out of the bubble? Of course, this was done for the cooled water. You have here the standard way spectroscopists usually plot the data. Four hundred nanometers is violet, 700 nm is red, and 200 is the far ultraviolet. Beyond here you are into what is called the vacuum ultraviolet region, and that is the region where water and the quartz flask cut off the light so the light cannot get out of the system.

As you notice from the data, on a log-log plot there is no sign of any structure. There is just a straight line on a log-log plot. There are no peaks, no lines. It would appear from this that the sonoluminescence is very, very hot, indeed, and it is on the tail of some very hot distribution. If we would compare it to the curvature of a typical cutoff from various theories of light emission which would go as minus Planck's constant times the speed of light over wavelength times kT, you see that this line is clearly straight compared to 100,000°, which has that much curvature, as you can see there.

Q: Are there finite size effects that would affect the spectrum that you get out of the --

A: There certainly should be. There have been a number of papers looking at emission from drops of laser dye and the size of the drop changes the density of state, so you would take something that might be a broadband emission and it becomes peaky. Our drop is smaller than those, so there should be finite size effects all over the place, but there are not.

Q: What sort of resolution?

A: This was scanned at about 5 nm resolution. This is calibrated for various processes like the transmission of light through water and quartz. For this particular case, at 1° C, not at room temperature, Bob Hiller did a scan with about 1-nm resolution and found no lines. So that is the accuracy, but we do not have a spectrum to calibrate that. He just verified there were no lines in that particular case.

(Transparency 12)

Let me summarize the time scales.

Q: Sorry to interrupt. On one of these previous evenings I was sitting in a room with Bob and one of the students mentioned that he thought in some cases there were lines for H+ and OH- that were seen. Who was the person who said that?

Q: I think Ron Roy mentioned that some people in Europe think there might be.

A: I think it is very important at this point to realize that there is a huge baggage from the past which we call transient or multi-bubble sonoluminescence. That is a very complicated arrangement, and I think there could be a long discussion about that whole field, which has gone off into sonochemistry and other applications.

The ability to study a single bubble is probably -- in fact, I will show you in a few minutes -- going to show that there might be more than one kind of sonoluminescence.

Q: Now that you mention it, I recall that it was four transient bubbles, the person said.

A: As I told you, the breathing motion of the bubble, which absorbs sound energy as it expands and then creates a flash of light as it collapses, the breathing motion of the bubble is the key to at least getting some picture of the theoretical interpretation of sonoluminescence.

(Transparency 13)

I would like you to trust me on a bit of theory that I am going to derive for you later. This is the key equation which is at least a guiding light in interpreting the bubble motion. I would like to write down the equation, describe it for you for a moment, and then use it and come back to it later when we discuss the theory of sonoluminescence.

The key parameters are the bubble radius, the velocity of sound, the ambient radius (that is the radius the bubble would have if you instantly turned off the sound field and just asked what is the radius of the bubble before it dissolves -- it would be R_0), there is the sound field amplitude P acoustic, there is the pressure inside the gas in the bubble, for which we use the adiabatic equation of state of a van der Waals gas. This is your PV^{γ} is constant, except we correct V by the van der Waals hard core of the bubble, which is the ambient radius over 8.5.

Then the equation of motion of the bubble is that the

acceleration of the wall of the bubble is given by the pressure difference across the wall, which is the driving acoustic pressure, the ambient 1 Atm, and the back pressure in the gas.

Q: In the van der Waals equation of state you have an A and a B. You have only A in that equation. Why do you write it like that, for computational simplicity?

A: We are neglecting the van der Waals attractive part. It is using only the hard core. If you put in the attractive part, it is absolutely negligible considering the phenomena we are looking at.

Fluid mechanics at low Mach number, where the time derivative of the radius -- that should be U_0 -- over the speed of sound is small and long wavelength, where the wavelength of the sound field is large compared to the radius of the bubble, yields, then, the Rayleigh-Plesset equation, which is your basic equation of bubble motion.

(Transparency 14)

In fact, I would like to show you a typical solution of the Rayleigh-Plesset equation. Here is the sound field, a sinusoidal sound field. Then, when the sound field goes negative, the radius of the bubble expands to a maximum. The sound field turns positive, the bubble suddenly crashes, this is the point at which we expect the light to be emitted, so this is the collapse of the bubble according to the Rayleigh-Plesset equation.

(Transparency 15)

At small amplitude, this equation describes the well-known formula for the adiabatic breathing of a bubble. The bubble will expand and contract sinusoidally in time for a small amplitude disturbance in the radius. Omega $_0 2$ is that frequency, and it goes to three times the density of the gas over the density of the liquid, the speed of sound in the gas squared over the ambient radius squared. This frequency will characterize the period of these afterbounces when the bubble collapses. These are relatively small amplitude and these characterize those oscillations.

Q: (Inaudible.)

A: The light? Okay. Notice the parameters here. We are working with an ambient radius of about 5 μ. The maximum radius is 38, so there is a collapse here of about nine to one, or an expansion ratio of about 9:1, to go from one here up to here.

What causes that expansion? The expansion is caused by

the amplitude of the sound field, which in this case is 1.3 Atm. If you go to lower sound field amplitude, the expansion goes down. When the expansion ratio hits about 7, you do not have the sonoluminescence any more. When you push the expansion ratio, in certain cases you are permitted to get an expansion ratio up to about 14. In that case, you get the largest signals that we have been able to see.

Q: But the sonoluminescence starts in this jittering area?

A: Here there is no light. All the light comes out here. Efforts to see light down here have seen nothing.

Q: What was your light?

A: That is the spectrum of the light. That is in the ultraviolet.

Q: Sorry, the frequency of the sound?

A: You can see this is one period, so this is 35 μsec long.

Q: With a change of frequency the light will still come out?

A: Yes, if you change the frequency, you are just scaling the graph this way and you will see changes. The light will still come out but the intensity will change a little bit.

DR. ATCHLEY: Seth, you might point out that because this is in a resonant system you do not have continuous control of the frequency. You have to hit a resonant mode, right?

A: Yes, you have to hit a resonant mode, but you can choose different resonators.

Q: What is the ambient pressure?

A: One atmosphere.

Q: What is the 1.3 Atm?

A: That is the amplitude of the sound field driving the system. That is the peak to zero. It is going negative. It is essential.

(Transparency 16)

Let us go back to the Rayleigh-Plesset equation. It is important that you go negative here, because you need P acoustic to be more negative than P_0; otherwise, you will not

get the expansion if you push the sines through.

Q: Put that graph up you had just a minute ago.

(Transparency 17)

The point where it shocks is almost completely vertical. How sensitive is that to the position in your driving sound field, because at that point your sound field would seem almost constant on the time scale where it drops precipitously.

A: Right. I can give you the answer. There will be a viewgraph to answer exactly your question later on, when we get to the theory. The answer is, if you got the bubble up to here and you had some miraculous way of turning off the sound field, you would see exactly the same drop.

Q: Wouldn't you see the light modulated?

A: The light all comes out at the minimum.

Q: It seems to modulate.

A: This is the radius of the bubble and the light comes out at the bottom of the crash when the inside of a bubble is so hot that it is, say, 100,000°, all the light comes out. Then when the bubble comes back up, everything is cooled down, again, and it never gets as hot until it comes around for another helping. So the whole spectrum and all the light, the milliwatts, are all down there.

(Transparency 18)

Obviously, determining the bubble motion is essential. Brad set up a light-scattering measurement to find the radius of the bubble as a function of time. Let me give you the overall picture of how this works.

The idea is that light scattering off of a spherical dielectric in a certain direction, a certain angle from the forward, so here is light coming in, hitting a bubble, and then scattering in some angle, the scattering in a particular angle goes as proportional to the square of the radius of the bubble in the classical limit. It depends on the surface area of the bubble.

If you could measure the light scattering as a function of time, then you get radius squared as a function of time. Take the square root of the voltage and calibrate it and you then have the radius of the bubble as a function of time.

There are problems that Phil Marston is very happy to study, and that is that the classical limit is achieved very, very slowly. Even when the wavelength of the light is a 30th

of the radius of the bubble you have huge variations on a log scale in the intensity scattered as a function of angle.

Brad overcame this problem. The classical limit is the solid line. So you have diffraction and interference on top of diffraction and interference. So much for the WKB approximation. Brad overcame that problem by collecting light from a large number of angles through a lens and averaging over that graph.

First, there is the overall picture, where things were tuned so that you could see -- this is just the response on the photomultiplier tube. This is the radius of the bubble increasing. This is the radius of the bubble decreasing. Here is the flash of light. There is a spike at the decrease. So each cycle repeats like that: maximum radius; minimum radius: flash of light.

These are the afterbounces, which will be resolved by improving the signal-to-noise measurement.

(Transparency 19)

Here is the radius of the bubble versus time compared to the solution to the Rayleigh-Plesset hydrodynamic equation that I showed you before. The solid line is the solution to Rayleigh-Plesset and the boxes are the experimental measurements. As you can see, the Rayleigh-Plesset equation is an excellent description of how the bubble is primed, how it collapses, and even the period of the afterbounces. Of course, it would be fair to say that one of the ways of calibrating the motion is to choose the period of the afterbounces to let you determine R_0 through the equation for the breathing frequency of a free bubble.

Q: Is there any significance to the little blip on the right?

A: That is the acousto-optic modulator turning off. There is a lot of significance. That is why we do not blow a tube every time we throw the switch, because when you try to pick out certain regions of the system you have got to have the light on for only a certain length of time, because the dynamic range of the photomultiplier tube cannot deal with that.

Q: It looks like you assume you can infer that that adiabatic approximation is valid?

A: I would say you can infer that good old-fashioned low-Mach-number hydrodynamics as first solved by Rayleigh in 1917 describes 99.99 percent of the cycle.

Now we are going to look in there and try to get at the

0.01 percent of time that we do not know, where everything is happening.

(Transparency 20)

We are going to take this box and blow this up, so that box is full scale, to look at the after-bouncing. This entails using a very beautiful acoustic device, which we came to appreciate only as a result of our interest in sonoluminescence. You see, when you want to look at the portion of a cycle when the radius is small, you need a lot of gain in order to get a signal on the photomultiplier tube.

But when you turn up the gain, you run into the problem that when the bubble gets to be large, you are going to get a huge signal and instantly blow your photomultiplier tube. Brad used an acousto-optic modulator, which is indicated here, which is just a small crystal with a traveling wave. The sound wave in the crystal sets up a variation in density in the crystal, as we all know, due to the sound wave, and that variation in density has a certain wavelength which acts as a diffraction grading.

When it is on, the laser going through the crystal is deflected by so many milliradiants and, therefore, it is deflected off the bubble. When it is off, it is back on the bubble. You run it so that when it is on, it is on the bubble and, when it is off, it is off the bubble. A sound field is used to set up the grading the deflect the laser so that you can run the experiment to look at the sonoluminescence.

(Transparency 21)

When that is blown up, you get the afterbounces in detail here. This is the crashing bubble and here are the afterbounces. Finally, if you take this little region and blow that up in time, you then get the final collapse and here is the bubble radius coming in from 10 μ crashing down to under 1 μ. You are following the bubble to sizes which are smaller than the wavelength of light.

Right down here is where the flash of light is emitted. The maximum slope on this graph is about Mach 1.2. It appears that the bubble is crashing at a supersonic velocity just at the moment when the light is emitted. The emission of the light is shown as a bump on this graph here. In this case, a filter was in to cut out the sonoluminescence and look only at the laser light, which was a helium-neon laser, a red laser, scattering off the bubble, so a filter was used to block ultraviolet light, and therefore block the sonoluminescence, to get the best possible radius curve.

Here that filter is removed and you see a bump due to the sonoluminescence right near the minimum.

Q: How well time-resolved is that?

A: As good as we could?

(Laughter)

Q: You were using a different detector for that. I was wondering --

A: The micro-channel-plate photomultiplier that can resolve on the 100-psec scale is built for pulses and cannot handle anything near the DC levels that come from a light-scattering experiment. In fact, I have the viewgraph that answers your question.

(Transparency 22)

Shown here is the radius time curve. Here is the radius time curve with the bump and here is the response curve of the photomultiplier tube to a delta function. As you can see here, the time resolution for this experiment is not nearly as good, because we needed the dynamic range or something more clever. The resolution is not nearly as good, because that is 10 nsec. As you can see, we are mixing together signals which are 10 nsec wide now.

I will give you my best opinion on what is happening here. The actual collapse is even steeper than we can resolve, because this response time of the photomultiplier tube is mixing in the slow part with the fast part, and we would like to find a high dynamic range, high bandwidth photomultiplier tube to look in more closely. This is a guess at what the convolution looks like, but I am not going to sell that now.

DR. SACHSE: When you said you reached 1.2 Mach number --

A: That line was just to show you what Mach 1.2 is. That the maximum slope on the curve if you choose the two pixels that do that.

DR. SACHSE: But is that the Mach number with respect to the speed of the sound in water?

A: No, that is with respect to the speed of sound in the ambient gas.

DR. SACHSE: But the speed of sound depends on the temperature, so you do not know what the temperature is in there, so I do not understand how you can plot a Mach number.

A: No, we can define one. The Mach number we always talk about relative to the ambient gas.

Q: It really does not mean what we would think it would mean, because you are talking about ambient room-temperature gas --

A: The Mach number, that is when the bubble is right here. So at that point the speed of sound is 330 m/sec.

Q: But not for the gas in that bubble it is not.

A: At that point it is. At that point it is going through ambient radius and at the ambient radius you are below that point by some amount.

(Transparency 23)

Brad said, well, he could measure the radius versus time of the bubble, why not measure the radius versus time of the complete parameter space of a trapped bubble of air in water? What is shown here is increasing acoustic drive level. This is the radius of the bubble versus time. As you increase the sound field from, say, 1 Atm to 1.3 Atm, there were two regions apparent.

The first region, unshaded, is the region of a trapped stable bubble which does not give off any light. This ramp scale shows the relative amount of light being given off as you increase the sound field level driving the system.

Here we have what we call the bouncing bubble. This bubble is entirely described by fluid mechanics and no light comes out. As you increase the amplitude of the drive, there is a sudden drop to a smaller bubble which then emits light. That is, you increase it further, you go into the strong sonoluminescence regime where you reach collapse ratios of, say, 10:1 and the small afterbounces that I was showing you before.

If you were to set up this experiment in your lab, you would trap a bubble and then, at low sound field amplitudes you would be in this regime. As you increase the amplitude, you will notice the bubble will suddenly shrink at a critical value of the sound drive and when it shrinks there you are in the sonoluminescence regime. A further increase in amplitude will then bring you to a bright bubble.

(Transparency 24)

I would like to talk for a moment about this transition. That is absolutely the key. This is a bubble at the top of the bouncing regime matched to the Rayleigh-Plesset equations. As you can see, the fit is relatively good. The drive level is 1.05 Atm. When the drive level is slightly increased, the radius of the bubble shrinks dramatically; it goes from am ambient radius of 7 μ to an ambient radius of under 2 μ. The

smaller bubble is the bubble which emits the light.

This is the bifurcation that takes you from ordinary fluid mechanics to sonoluminescence. We do not understand the physical mechanism that is responsible for this transition. If we could, then we could probably know how to get it in other liquids or why we cannot get it in other liquids. We have been unable to see this phenomenon in other liquids, because we have been unable to see this transition in other liquids. This is one of the key unsolved questions, as far as we are concerned.

Q: This looks like something that might depend on the rep rate, so have you tried changing the frequency to see what happens to this drop?

A: Largely, it appears to be independent of the frequency, but in the experiments I am going to show you later with different gases rather than air -- we cannot change the liquid but we have been able to change the gas -- there are some indications that have to be followed up that indicate that perhaps we might be able to get at this dynamic process through changing the rep rate. It is a perfectly solid point you are raising and I am sorry I cannot give you an answer.

Q: You have reported the radius equivalent to about 5, as I recall, and the maximum was about 40. What is different now?

A: I told you I pestered these guys to get the most light possible.

(Transparency 25)

The transition is right here and those experiments are taken up here.

DR. ATCHLEY: On that graph right there, are those phase-shifted just to make them look nice or is that actual --is the bubble collapsing earlier in the cycle?

A: I think this was strictly for display. This can be answered. What you are asking is can we put the sine wave for the sound field on each of these so we know which phase is doing it. That can be done.

Q: What liquids have you tried?

A: In addition to water? Alcohol.

(Laughter)

Q: The light is related to the temperature, so if you lower the alcohol temperature so it is relatively into water temperature close to the breathing point --

A: Excellent point. You could be motivating experiments with your thinking. That is an excellent point. We lowered the temperature of the alcohol to see whether we could get the vapor pressure down to the point where it is comparable to water. No luck.

DR. ATCHLEY: We tried silicon oil, which the sonochemists evidently got a lot of success from, and we did not see anything. Sonochemists say if you use silicon oil you will blind yourself with transient cavitation sonoluminescence, and we could not see anything.

A: That is because they walk around with image intensifiers all day.

(Laughter)

Q: (Inaudible.)

A: We tried to list all kinds of parameters that I could depend on, and there are just a lot of parameters.

Q: If you used alcohol, did you try to get rid of all the water? Was it pure?

A: Yes, this is a problem, but it was a reasonable effort. It was not done -- I have to explain the current state of affairs. There are millions of things to try and so many of them yield something interesting that we just go from this to that. If it looks like it is going to take a lot of work, we will go to something else and find something interesting there, rather than try to really press down hard on why can't we get it from alcohol.

Q: If you have a little alcohol in the water, you could not see any light, so if you have a little water in the alcohol you still could not see light.

A: Right, I agree. We have also tried propylene carbonate, which has a very, very high dielectric constant. Chemists recommended that. The people in electrical engineering wanted us to do heptane. The reason they wanted us to do heptane is that microwaves can propagate through heptane and they would have really loved to study microwaves coming out of a bubble. Water is a great insulator against microwaves.

(A brief recess was taken.)

Q: I have a question. You can calculate some total energy contained in that system at a given cycle. When it makes that drop, the energy has decreased. Is that related to the amount of light that is being emitted?

A: There are various efficiencies that you can talk about

for this process. Let me define, just for purposes of characterizing the phenomenon, not for the purpose of setting up a company, first there is the efficiency, which you can say, if I had a Bob Leonard resonator, where the attenuation was due to the shear and bulk viscosities, and I compared at this amplitude of sound how much energy went into heat due to Kirchhoff's law, [such as in Landau and Lifshitz in the chapter on sound]. If I compared the amount of sound energy going into heat in such a resonator to the amount of light that comes out when a bubble is present at the same sound amplitude, the amount of light which comes out is equal to the amount of energy that would go into heat.

So from the point of view of Boltzmann's equations, and so on, this is not a small effect. This is a hydrodynamically significant effect.

Coming to the efficiency, which determines how much it is going to cost from the outlet to run the experiment, that is, of course, much less, because, as you pointed out, the key quantity is the amount of acoustic energy stored in the bubble as it goes up to the maximum compared to the amount of energy that comes out as light at the minimum. That is about 0.1% efficient. It is a little better than an argon-ion laser.

The other point is that Brad asked me to emphasize that these are all steady-state graphs. You sit at this pressure, take a steady-state picture. It repeats like this as far as you go. Then you suddenly increase the parameters of the phase plane of the system to this point, and then this is the steady picture. You have an abrupt transition in parameter space, but this is a steady bubble and that is a steady bubble. The experiment is run by very slowing increasing the acoustic so that you can stabilize it at each state.

Q: Are these the shock curves or averages?

A: Those are averages. Each curve is averaged.

DR. HARGROVE: If all you could monitor is the bubble-bounce trajectory from just after the flash on, if that was the only way that you could see, could you tell whether or not there had been a flash? Does that structure change enough?

A; Yes, but the thing is, two things are happening at once. As I mentioned, when you get to the sonoluminescence, the collapse ratio or the expansion ratio increases dramatically relative to here, and you can see that these afterbounces are tiny compared to those afterbounces.

The reason these afterbounces are tiny compared to those is probably more due to acoustic radiation due to the higher acceleration. But you need a higher acceleration to get sonoluminescence, so you get more acoustic radiation, so if you

look at the graph, you can immediately diagnose that this is a sonoluminescence graph, because there is that much more acoustic radiation that caused the afterbounce to damp out.

DR. HARGROVE: I guess what I am getting at is did the light emission take enough energy out that you see a direct effect on the ringing?

A: This would be very exciting and we do not want to rule out that possibility, that the light could then bite back on the hydrodynamics. A shortcoming of the theory that I will try to develop later is that I am always going to be separating the hydrodynamic concentration of energy from the light-emitting mechanism, whatever that is.

Q: Surely you have got less bubble radius, and you are still not seeing the oscillations.

(Transparency 26)

A: Back to this graph. Yes, but, in fact, that is just the scale of the graph, because the expansion ratio here is up to around -- it is getting closer to 10:1.

DR. ATCHLEY: I think it is important to point out that hydrodynamics models both of those curves pretty well, yet there is nothing about light emission in the hydrodynamic code. It is more the parameters that model it.

DR. MARSTON: It is interesting, if, in fact, hydrodynamics models both of those curves and your bubble is, in fact, sitting in the center of a spherical resonator, you never see, it would appear, any effect from that radiation that is going out from the bubble being focused back on the bubble.

A: There is such an effect. The bubble changes the sound field. In fact, the first thing you will notice when you sit down to play with this is that obviously the first thing you have to do is take the water, as Felipe noticed, you have to de-gas it somewhat so that you do not get all this transient cavitation all over the place. Then you have a spherical resonator with water and you drive it with the piezoelectric transducer.

Then you say, okay, it is time to introduce a bubble. If you find the resonance inside that resonator and you then introduce the bubble, you will find that before that bubble gets to the center sometimes you are not quite at the right parameter space to trap the bubble, it walks around, and the microphone looking at the system will see the sound field fluctuate all over the place.

One tiny bubble in a sound field in a resonator completely messes up the sound field, so the bubble changes it, and the P

acoustic that is listed here is the net P acoustic due to the resonator and the bubble.

DR. MARSTON: No, I agree with that. What I am trying to ask about is the transient signature. In other words, as you have indicated in this collapse phase, the mechanics, if I understood you right, are modified significantly by the radiation of the bubble, and the question is whether you ever see an event because of the reflection, the high speed of that? In other words, before that drop, it has a low-frequency component --

A: But the collapse of the bubble sends out some shock wave which bounces off the wall and comes back and might cause something to happen later on in the cycle.

DR. MARSTON: Yes, have you ever seen any evidence of that?

A: We have seen various anomalies but nothing that was reproducible. Sometimes we see humps on these curves and we think it is perhaps due to the bubble drifting out of the laser beam or into it. We look for those kinds of explanations. The answer is no, but this is something that would be very interesting to look for in a tweaked-up arrangement. It has not stared at us.

Q: Do you know what the temperature of the bubble is as it passes through R_0?

A: No.

DR. LEVY: You have got two curves that look very similar, one that cavitates and one that does not, and yet you can reproduce both of them, theoretically. What do you do to do that? What parameters are you using where the maximum bubble diameter is about the same?

A: There are various ways that the curves are calibrated. I would prefer to put that off to when we do the theory, but I could make a couple of comments just to hold you for the time being. The slope here -- notice this long extended linear slope -- that slope calibrates P acoustic (as I will show later) and then the period of the bounce calibrates the ambient radius.

(Transparency 27)

As I showed you here, we do not know why we cannot see this transition in other liquids. But if you cannot see sonoluminescence in other liquids and changing parameters gives you such a surprise all the time, then perhaps you could change the gas inside the bubble. I would like to show you a couple

of slides taken of a system which was designed by Bob Hiller in order to change the gas inside the bubble.

(Slide)

We are now to cylindrical resonators. This is a quartz cylinder. Quartz is chosen because it is friendly to passing ultraviolet light so we can make the best spectral measurements. This is a thin brass cap. On top you can see a piezoelectric transducer which drives the resonance. This is the light being given off by a xenon bubble. This bubble is roughly 2% xenon gas and 98% nitrogen gas.

The wire sticking in here is a thermocouple to measure the temperature and, as you have seen, the system is very sensitive to temperature variations, so since we are going to cool the xenon bubble eventually, why not measure the temperature to track that?

There is another piezoelectric on the other side. Bob aimed for an extremely symmetric arrangement so as to have good acoustics. When you start weighting things down, the acoustics can go haywire, so the PZT down here, which is used as a pickup to measure the sound field, a microphone PZT, through this small tube comes, from the mixing chamber, water with the appropriate dissolved gas which we want to turn into the bubble.

The trick here is as follows. The container is filled with water with, for instance, the dissolved xenon mixture. First the air is removed from the water by de-gassing under a vacuum. After the air is removed from the water, then the xenon mixture is mixed into the water.

Now, you say, but we want a bubble of gas inside the water and how do we get that particular bubble? To do that, Bob mounted here some toaster wire, a nichrome wire, and a pulse through the nichrome wire locally boils the water near the nichrome wire and forms a vaporous cavity.

But before that vaporous cavity can collapse as much as a vaporous cavity would collapse in a boiling pot of water, before it can collapse, gas that was dissolved in the water diffuses into the vaporous cavity and cushions the collapse, keeps it from collapsing, so that more gas can come in, and you then turn from a vaporous cavity to a gas-filled cavity, and that gas-filled cavity then gets trapped at the center to give off the light which you can see.

There are a couple of very, very important things here, like, for instance, you have this tube coming here to bring the water with the gas dissolved in it. Then you have another tube on the other side going here. This is the "politically correct" way to relieve the gas pressure at 1 Atm. The vice-

chancellor for political correctitude came by -- we had this unspeakable way of bringing things to atmospheric pressure which we were told had to go, so instead of going to the pharmacy we had to go to Fisher Scientific.

(Laughter)

Bob has these nice supports here to get the maximum Q for the resonator. Bob is here and I am sure he can tell you more about this gem.

(Slide)

This is a close-up of the bubble. This bubble has laser scattering from it, in addition -- maybe, perhaps, afterward you can come up and see a little red ring around the bubble. This is also taken in incandescent light. Xenon is a very bright bubble.

Q: Do you know how much residual water is in the bubble?

A: No, we just assume we have got the usual 20 mm during expansion. We have been unable to assay the bubble properties. When I show you on these graphs pressure head, this is the partial pressure at which various mixtures have been put into the system.

Punch line first: The first thing Bob did was, he said, look, air is 80 percent nitrogen, let us look at a nitrogen bubble. Take the small step first. So he got a nitrogen bubble, turned up the amplitude, and it gave off very, very little light, roughly 1/30th to 1/50th the light you get from an air bubble. Nitrogen did not work.

We were all very excited, oxygen must really be amazing stuff, so we put in an oxygen bubble and oxygen bubble did not work. Then he put in an 80/20 mixture, 80% nitrogen,

20% oxygen, and that did not work. We were convinced that there was some grease in the lines bringing the gas into the system and we were corrupting the system every time an experiment was run.

We could not find the stupid mistake. Then we realized air is 1% argon. When the 1% argon was mixed with the 80 and the 19, or take the 1% argon and just the rest nitrogen, you then got a complete sonoluminescence signal. The argon is the key to sonoluminescence. We do not know why. In fact, any noble gas is the key to sonoluminescence and we have no explanation for what is going on there. We call that "Hiller's thriller."

I want to backtrack to recent literature. There was a paper by Anthony and Felipe and Glynn Holt and a fellow from

Germany, Holtzfutz, and they said in the paper that the discovery of single-bubble sonoluminescence by Felipe was fortuitous, and that is just not true. Felipe worked really hard.

(Laughter)

But the statement that the discovery of single-bubble sonoluminescence was fortuitous is fortuitous, because if it was not for the 1% argon he never would have found it.

(Transparency 28)

If we look at sonoluminescence normalized to an air bubble for different pressure heads, these are the pressure heads at which the gas is mixed into the water. The water is de-gassed to some extent and the extent to which it is de-gassed is shown here. For instance, here a saturated solution would be 760 mm and here you have nitrogen mixed in at roughly a third of saturation, or something like that, and this is a log scale, so you can see the sonoluminescence intensity of a pure nitrogen bubble is extremely low.

You can see that a pure argon bubble is lower than the 1%. So not only does 1% argon in nitrogen give you the sonoluminescence but it is larger than either the pure argon or the pure nitrogen.

Another point to notice on this graph -- of course, there were these little peaks here, which we do not understand and are reproducible, this has been taken more than once -- is that as you take the 1% argon mixture and go closer and closer to saturation, transient multi-bubble sonoluminescence is out here, there is a cross between the curves.

In transient sonoluminescence the air bubble of the 1% is going to be very, very far below the pure argon, at least by extrapolating this curve. I think these phenomena here near the peak have different physical aspects from the phenomena out here. The difference is the de-gassing.

DR. GREEN: Do you know what gases they put in fluorescent light bulbs? Are they 100% argon, neon, stuff like that, or are they some mixture?

A: There are all kinds of mixtures.

DR. GREEN: There is some reason they are using this mixture, right, in these fluorescent lights?

A: Yes, but the gas discharges are affected by impurities.

DR. GREEN: So what if you used the same mixture as that,

would that be more efficient?

A: I do not know. We have called around to neon sign makers to ask whether they can give us any advice on what is going on here. As a rule, they did not, but we are very open minded that there is something going on of an atomic-discharge nature being run by the sound field.

Q: Have these mixed gases ever been tried in other host liquids?

A: That is what is happening, because obviously if we cannot get it in other liquids, maybe nitrogen is the key to another liquid, absolutely.

(Transparency 29)

Here is a graph of the sonoluminescence as a function of xenon, argon, and helium. Basically, in each case the pure nobel gas is here and the mixture at the 1% level is there. As you can see, it is a very quick turn-on as you increase the percentage mixing to about the 1% level, which is roughly where the peak is.

(Transparency 30)

Then Bob measured the spectrum of these different cases and the 1% argon and nitrogen -- these are log-log scales, again, here is red, violet, ultraviolet -- the 1% argon and nitrogen is roughly the air spectrum at room temperature here. You can see a number of interesting aspects to these spectra.

For instance, the helium spectrum is enormously ultraviolet, like the helium is going straight off into the far ultraviolet until the properties of water prevent us from taking further data. The xenon spectrum is particularly interesting, because for the first time we have got the sonoluminescence to have a spectral peak, and the spectral peak for xenon is at about 300 nm.

You know, with diodes, people have red diodes, and they are desperately looking for these blue-green diodes and trying to get to higher energy. To learn about sonoluminescence we are starting off with something in UV acoustics and we are desperately pushing, trying to get the peak into the red so we can learn something about it.

This is 100% xenon and that is the 2% xenon mixture. This is the enhancement effect, as we mentioned. But they both have a spectral peak more or less in the same range.

Q: What sort of time frame here?

A: Let me go to the next graph.

(Transparency 31)

As we have said, cooling things down can have a dramatic effect. This is 2% xenon cold. This is air cold. This is the helium at room temperature to show this is not the cold helium. The first thing to notice is that these graphs have the same tail, which goes as one over wavelength to the 2.5 power.

The procedure for taking these data has been improved over various previous spectral measurements and exponents of slope are very sensitive, especially since you have such a limited range in which you can look at light. Your ears are much better than your best spectrometer or your eyes. Sound goes a factor of 1000 and light is a factor of two or three.

Bob calibrated these data by forming a real image of a calibrated deuterium lamp where the bubble would have been, and letting that thing go through the apparatus. This makes an improvement over previous calibrations and gives a slope which is more in the range of one over wavelength to the 2.5 power. This is the best that we have been able to do at this point. This is corrected for the transmission of the water and the quartz flask, which has its own lines around 240 nm, which has been corrected for as best as possible.

So it looks as though there is a chance things are heading toward some kind of universal slope in the tail, one over wavelength to the 2.5. There are a number of comments that can be made about this. The average power of the cold xenon is about half a microwatt. If these flashes, which have not been measured, are still determined by, say, 100-psec time scales, then the peak powers are pushing about .2 W.

Helium, if you would say this is the tail of a curve which has the same energy as this curve, except it is cut off, to the eye the helium is dim, but that is because all the energy is in the ultraviolet, then this curve would turn over beyond 100 eV. That is, of course, wishful thinking that that curve would go to 100 eV, because that would be soft x-rays.

But if this curve turned over at 100 eV, there would be no violation of energy conservation, because this curve has enough energy so that if it were put here, this would turn over at 100 eV. So it cannot be ruled out that this is the case. We do not understand this variation, but we do maintain it is curious.

X-rays do not go through water until you hit 20-30 KeV. Water cuts off from 70 to about 30 KeV. You are going to see neutrons before you are going to see x-rays.

(Laughter)

And before you do that, you are going to see hallucinations.

(Laughter)

(Transparency 32)

This is the phase-plane curve of an argon bubble and there is a remarkable difference in the pure argon bubble from the 1% argon, namely, the air bubble. In this curve there is no abrupt transition to sonoluminescence. The sonoluminescence just turns on smoothly. So the phase transition which characterized an air bubble in water for a pure argon bubble is gone. It is a first-order phase transition and for the argon bubble it is second order. This is not unusual when you look at, for instance, superconductivity and external magnetic fields: limit of zero magnetic field, second order; finite magnetic field, first order.

(Transparency 33)

Everything begs to have the spectrum of a nitrogen bubble. Sorry. The intensity of light from the nitrogen bubble just oscillates with time. There is some regularity here, but the dimness plus the nonsteady behavior of the light emission from a nitrogen bubble just precludes more careful measurement.

(Transparency 34)

I forget who had the brilliant idea, but this is helium-4. Obviously, we should do helium-3 and see if there is some kind of isotrope effect, and cooling, right? Cooling helps sonoluminescence. We are going to be doing low-temperature helium-3 research.

Some people said what about the vapor pressure of water and thermal conduction? Greg Swift said, "Seth, what you should do is take heed from this work on similitude in thermo-acoustics. Why don't you set down the completely general theory that has all the parameters and then you can scale it around and make predictions?"

(Transparency 35)

So I made this graph for Greg. I wanted to show all the parameters I can guarantee are important in this phenomenon: There is the temperature of the water, the latent heat of the water; the ambient pressure; the density of the water; the resonator volume; shear viscosity; thermal conductivity; speed of sound; van der Waals hard core; the hard core for ionizing collisions (hey, at the temperatures we are at it is a plasma there); the ionization potentials to kick electrons out (this nobel gas stuff might have something to do with that); the molecular weight; the surface tension; amplitude of sound;

concentration of gas in a liquid; concentration of impurities; and basic constants (there is no evidence that Planck's constant or polarizability of the water should play a role, so these I just list, but there is no evidence we need those, I am happy to drop those).

Here is an example of dimensionless parameters you can form in your grand unified similitude theory, GUSTY: ionization potential over shear viscosity; speed of sound; van der Waals hard core resonator volume to the one-third power -- that is dimensionless.

(Laughter)

Q: What is it called?

A: The Swift number, right?

What is the theory that has all those parameters that we can start scaling around? The only thing I know is Schröedinger's equation or something like that. I think the lesson from this is that we do not know the underlying theory that includes all these parameters which can be tweaked for sonoluminescence. The amazing thing is that the sonoluminescence works in spite of our inability to approach it from the point of view of similitude. So we have to embrace that perversity here.

(Transparency 36)

Having said that, the only thing I know how to do is what I have been studying, what Tom Erber taunted me with, hydrodynamics. Certainly the hydrodynamics will give us some insight into the sonoluminescence; it will at least give us the Rayleigh-Plesset equation which describes so much of the cycle that precedes the flash of light.

Let me review here for you the equations of fluid mechanics. Fluid mechanics of a simple fluid such as water is characterized by five variables. The variables are: the mass density as a function of position in time; the velocity with which the fluid is moving in centimeters a second (that is the local velocity of convection of the fluid); and the entropy per gram, S.

I am selecting from the equations of fluid mechanics the paired-down equations which have the key phenomena that at least we can connect to the sonoluminescence that we have done so far. There is, first, of course, conservation of mass, the amount of water that is conserved or the amount of gas that is conserved (the same equation applies).

Then there is Newton's law for a liquid: mass per unit volume times acceleration is the gradient and pressure that

accelerates the liquid, the shear viscosity times the diffusion of the velocity, shear stresses, plus gravity (this will be important only for discussion of how the bubble is trapped in the sound field).

Then there is the law of conservation of entropy. The change in entropy as you follow the motion of the liquid is zero. We are going to assume adiabatic flow. D/Dt means the convective derivative, the partial with respect to time, plus dr/dt times the gradient of V. So these are the five basic equations of fluid mechanics. You do not even need this one to get Rayleigh-Plesset. These two are sufficient to get the Rayleigh-Plesset equation.

Since we have been talking about gas dissolved in the liquid, we also need the equation of mass diffusion, so the time change of the concentration of gas in the liquid plus the rate at which it is convecting along, which can be very sizable because this V will be pushed by the bubble, which is moving at high velocity, is the diffusion coefficient for mass diffusion times the Laplacian of C.

I do not know if I will get to barrodiffusion today. If we do, there will be one extra term on that equation which will depend on the Laplacian of the pressure. Let us not worry about that right now.

The key equations are: mass conservation; Newton's law; mass conservation of impurity; and adiabatic flow. Obviously, fluid mechanics can be more general. I have taken that subset of fluid mechanics which will match to our best attempt to get some insight into a theory of sonoluminescence.

(Transparency 37)

First, I would like to derive the Rayleigh-Plesset equation. If you start from Newton's law for the center of mass and write the velocity as the gradient of a velocity potential, you can then remove gradients everywhere. We will ditch the shear viscosity, since it is a small correction. We will ditch gravity, since it is small. So we will work with just those two terms.

Doing that we will get Bernoulli's law, the time change of the velocity potential plus velocity squared over two plus pressure is a constant everywhere. Therefore, this quantity evaluated at the radius of the bubble equals this quantity evaluated at infinity, which is P_0 over rho, P_0 being 1 Atm, the ambient pressure in the system.

Q: (Inaudible.)

A: No problem. The viscosity, even if it is irrotational

and viscid, the viscosity would still come into the boundary condition.

Now what we are going to do is look for a solution to Bernoulli's law, where phi is the incident phi of the sound field plus the scattered velocity potential due to the motion of the bubble. The sound comes in, causes the bubble to oscillate, but the oscillating bubble radiates sound.

If we want to include this consistently, we expand the velocity potential as the incident plus the scattered. For the incident for our purposes we will just take a plane standing wave and then for the scattered field we will expand in the Mach number. As we mentioned before, the Rayleigh-Plesset equation is an expansion in the velocity of the surface of the bubble divided by the speed of sound. It is a low-Mach-number expansion.

Doing that, we get the first term which comes for incompressible flow. This solves del squared phi equals zero plus the correction due to the compressibility of the water. I should have used a different U. Even for our purposes, the only reason to carry this term around is to get the radiation damping of the sound field.

The scattered velocity potential is obtained by using the boundary condition, that the velocity of the surface of the bubble is R-dot, so D phi DR, which is the velocity in the radial direction, must equal R-dot, which is the velocity of the surface of the bubble. Then we use the continuity of pressure at the surface of the bubble. The pressure in the gas inside the bubble equals the pressure in the liquid.

So we can even come very far if we just took incompressible flow and took del squared phi equals zero. This term will track through only to the radiation damping. So if we substitute this expression for phi into Bernoulli's law, pressure plus velocity squared over two is a constant, we then obtain the Rayleigh-Plesset hydrodynamics, which I wrote down before.

(Transparency 38)

Which is the acceleration of the radius of the bubble depends upon the acoustic pressure, the gas pressure, the ambient pressure, and all those retarded terms with the velocity of sound all combine to give the acoustic radiation damping of the bubble motion.

As we said, that equation has in it already the seeds to describe the expansion and collapse of the bubble.

(Transparency 39)

But before I do that, we have to pay homage. Rayleigh, in 1917, studied this equation. He is the first to have considered the properties of a bubble in a sound field and the collapse of a bubble. He worked on this because he wanted to study the sounds emitted by water in a kettle as it comes to a boil, vaporous cavities formed and the vaporous cavities collapsed. That was his interest, he connected it to the collapse of bubbles.

But even Lord Rayleigh had grant monitors: the Royal Navy. The way he got them to support this was by saying this was connected to cavitation behind screw propellers and everyone knew that that was what caused damage to propellers of the boats, and so "help me study this and I will help you know what is going wrong with all the Royal Navy boats."

(Transparency 40)

He then proceeded to set down the equation which I showed you and to calculate the pressure developed when a bubble collapses. He found, from that equation, that when a bubble collapses, the pressure reached 10,000 Atm, which he says is pretty sizable.

But he made a spectacular error in calculating this result. That error is why I am talking to you today. This field would have been studied to death by now. Right here. He says the collapsing bubble can be assumed to obey Boyle's law. Boyle's law means temperature constant. He assumed temperature was constant in the collapsing bubble. He ruled out the Santa Ana condition.

If he had taken it adiabatic he would have found here, in addition to 10,000 Atm for his parameters, he would have found 10,000° and he would have said, hey, some light should come out.

(Laughter)

(Transparency 41)

In fact, he can really be faulted, because he was not up on the literature. Right there, in the Bible it says, Exodus 20:18, "All the people saw the sounds." I tried to find out why people did not take that seriously and they said that God's proposal to study sonoluminescence was turned down, not only wasn't the work refereed but no one could reproduce it.

(Laughter)

(Transparency 42)

The actual discovery of the sonoluminescence came about, and this is what Tom Erber was taunting me with, in a 1934

paper by Frenzel and Shultes. They discovered the transient sonoluminescence based on some colleagues who were irritating them. They had this colleague, Meckie, who never wrote a paper on the topic and Meckie said the American chemists found that when they blast a fluid with sound the sound affects the rate of chemical reactions and produces peroxide.

He said if you can make peroxide with sound, you can make light with sound, so he instigated these guys to do it. They thanked him in the paper. However, in the good German tradition, these guys say right here, at the end of their paper announcing the discovery of multi-bubble sonoluminescence, they say it is our duty as good German scientists to report this interesting work to you, but as regards the future, we have more important things to do, so this will be our last paper on the topic.

(Transparency 43)

Now I would like to go back to the Rayleigh-Plesset equation and try to give you a theoretical feel for what I think is the key insight this equation can give you regarding the bubble motion.

What we should understand is the asymmetry of the expansion and collapse of the bubble. When the sound field goes negative, the bubble expands. Then, as the sound field goes positive, there is this catastrophic collapse that even interested Rayleigh in 1917, or was derived by him in 1917.

Why does it say symmetric?

(Transparency 44)

If we go back to the Rayleigh-Plesset equation, we are going to take it in its simplest form. We have the acceleration is the pressure difference. We will drop the acoustic radiation damping, and we are going to write for P acoustic something proportional to sine omega t. This is the acoustic pressure at the location of the bubble.

(Transparency 45)

If we expand, calling this point t equals zero, if we expand around that point, the right-hand side goes as the difference between the amplitude of the sound field and one atmosphere and then the term in t squared, which is the expansion of the sine or cosine around this point.

The solution to this equation has a term, a constant, a term linear in time, and a term which goes as time to the fourth power. In other words, the solution to the Rayleigh-Plesset equation around the minimum of the sound field, the antinode in time, is accurately linear, because the deviations

from linearity comes not at quadratic or cubic but at fourth order.

This quantity alpha, which is then determined by the Rayleigh-Plesset equation, is the speed with which the bubble is expanding in this region. That speed is roughly 1% of the speed of sound in the bubble and it is given by the difference between the sound field amplitude, one atmosphere divided by density to the one-half power. So there is a relatively slow expansion of the bubble, Mach 0.01, linear in time and extended linear regime until it hits a maximum and then collapses.

(Transparency 46)

Let us now look at the catastrophic collapse. For the collapse everything is happening fast. You put the right-hand side to a constant in the Rayleigh-Plesset equation, so now the right-hand side of the Rayleigh-Plesset equation is just a constant, because this radius is changing so quickly.

(Transparency 47)

That is the approximation we will make in order to get an insight, as people say. Of course, once you have the equation you can just get it out of the computer, but I guess permit me some old-fashioned manipulations here.

The collapse is catastrophic, so we put the right-hand side to a constant. Then we get what looks like an equation of energy conservation. It is the time change of the kinetic energy of the bubble plus the potential energy of the bubble is zero. This is an ordinary differential equation that can be solved time-independent, can be solved straightaway, to get the velocity of the bubble and, as the radius of the bubble gets smaller, the velocity gets faster. It runs away and if you integrate the equation, the time to go from the maximum to zero radius is the maximum radius divided by the ambient pressure over the density, which all turns out to be about 5 μsec. This is the catastrophic collapse of the bubble in 5 μsec.

You have a linear rise and then the catastrophic collapse. If you then say what is the speed of the bubble as it is moving through the ambient radius divided by the velocity of sound in the gas, which is the true Mach number, then this is the square root of the gas density over the liquid density times the maximum radius over the ambient radius cubed. This is a very nice combination of large and small numbers.

The gas density over the liquid density is 1/1000. It is the spacing between gas molecules divided by the van der Waals hard core. In order to get sonoluminescence, which we think requires this Mach number to be order one, you then need an expansion ratio, maximum radius over ambient radius, which is

of order of 10, so that 10^3 cancels the 1000 to give you a quantity of order one. That is the key criterion for sonoluminescence; namely that this quantity calculated from the Rayleigh collapse be of order one.

(Transparency 48)

That is the picture in a nutshell of the slow rise, the catastrophic collapse. The real question is what happens when it goes past Mach 1 and what is the nature of the shock emission or what is going on there?

Before the next break, I would like to show you what the theory is for the trapping of a sonoluminescence bubble. I guess I should go back here.

(Transparency 49)

The trapping force is the volume of the bubble times dp/dz. For simplicity, I am imaging a cylindrical resonator where z is along the axis of the resonator, dp/dz is the gradient in pressure. The gradient in pressure times the volume of the bubble gives, then, the net force trapping the bubble.

Because of the wildly nonlinear motion of the bubble, the trapping force will be dominated when the bubble in this region where the radius is maximum, because since the trapping force goes as the volume, the volume goes as R^3, that is dominated up here, so you figure out what the pressure gradient is up here, multiply that by the volume, and you have the trapping force. The trapping force is, therefore, linear in the acoustic driving amplitude. So you go to high-amplitude motion and you end up with a linear effect.

(Transparency 50)

The buoyancy force goes as maximum radius cubed times gravity. There is a one-quarter because it is at the maximum radius for one-quarter of the acoustic cycle, so that is your duty cycle. Acoustic trapping force is volume of the bubble times the derivative of pressure with respect to z. Here is our expression for delta P for a standing wave and a resonator of height H.

So if we balance the trapping force against the buoyancy force and scale everything to the wavelength of the sound field, then the deviation of the bubble from the velocity node divided by the wavelength goes as rho g times the wavelength over the acoustic drive level. Typically, this turns out to be on the order of 1 mm, if lambda is 10 cm and P acoustic is 1 Atm.

The bubble will be located roughly 1 mm from where the velocity node would have been in the absence of the bubble. As I mentioned, the trapping force is linear in the acoustic drive level. This is in contrast to the long-standing Bjerknes force for trapping bubbles of small amplitude, which, of course, is quadratic in the sound field. You go to higher amplitude and it switches from being quadratic to being linear.

Q: Can you go over that, again, Seth?

A: It happens because the bubble motion is so nonlinear that the force up here dominates the force anywhere else in the cycle. For the Bjerknes effect the maximum radius differs from the minimum radius by an amount epsilon, and you have to subtract the two.

Thanks. We will have another break.

(A brief recess was taken.)

DR. PUTTERMAN: Let us summarize where we are with the theory.

(Transparency 51)

We can see how to go from classical hydrodynamics to the Rayleigh-Plesset equation, which applies a low Mach number, to derive the solid line which fits the radius time curves of the bubble with remarkable accuracy through over 99% of the acoustic cycle.

Also, we see how energy is stored and how the collapse initiates. Obviously we cannot trust the theory to those speeds of bubble wall collapse which are approaching the speed of sound, since the theory was, as we said, derived at low Mach number. So there is a second process going on, where you are going beyond the Rayleigh-Plesset equations.

I think there is another step to the energy concentration, which is the handover from a collapsing bubble and the adiabatic heating inside the collapsing bubble to an imploding shock wave.

(Transparency 52)

I could best show you that on this viewgraph here. It is because of this possibility that, for instance, the bomb code people at Livermore have gotten interested and, in addition, the inertial confinement fusion Mafia has developed a lot of interest in the problem.

Here is the water. At this interface is the interface between the water and the air. Then inside the air bubble, (or the xenon, or whatever you have), there can be another strong

gradient in the density, a shock front, which is your second radius. We will call this radius R_s, and this is radius R, so from here to here is the bubble.

If all you could have was a bubble collapse, the concentration of energy would be limited, because the bubble could collapse down roughly to the van der Waals hard core and the collapse would stop. But a shock wave can run through a compressed medium. It is just a very high-amplitude sound wave. So if you hand over from the collapsing radius of the bubble to an imploding shock wave, you can achieve still greater degrees of energy concentration.

(Transparency 53)

So it is back to Landau, second edition, where he has the section on strong spherical shocks. By the way, there is also some old stuff in this literature in some journals we do not have in the UCLA library. It is about the dazzler in Dazzler Comics right here. Right here it says, "Not too odd, considering my mutant ability to channel sound into light."

(Laughter)

(Transparency 54)

Now we will go back to the fluid mechanic equations. We are now going to take the equations of fluid mechanics without any damping whatsoever, so we have conservation of mass, Newton's law for a continuum, mass times acceleration as the gradient in pressure, and conservation of entropy.

(Transparency 55)

The key point here, which was figured out by a German physicist, Guderley, in the 1940s -- and I still have not been able to ascertain why he did it, but he studied imploding shock waves and he showed that those three equations which I have reproduced here in spherical coordinates, this is conservation in mass in spherical coordinates, the one over R comes from your spherical divergence, that is what distinguishes 3D from 1D, this is what dramatically curtails shock formation in three dimensions, is this term.

This is conservation of entropy in an ideal gas, log PV^{γ} is your entropy, and then here is acceleration, Newton's law, gradient P as ρ times the convective derivative.

The amazing property of fluid mechanics is that when you go to very, very high amplitude, that means Mach numbers much bigger than one, you would then get an exact solution. You have an exact solution in the limit of small Mach number -- that is acoustics and perturbative nonlinearities. But in the

limit of Mach No. 1, it is brutal. High Mach number, you again get an exact solution.

I have the exact solution written down, but what is more important to you is the fact that there exists such an exact solution and the summary of the exact solution is this term right here. The radius of the imploding shock wave is a constant times time to the alpha power. You say what is alpha? Alpha is 0.72. So this is the sum total of the exact solution. How you derive that, of course, is fairly hard, but let me indicate to you the basic ideas behind the theoretical derivation and leave out the algebra.

The key insight is, first, to introduce a dimensionless variable, which is R divided by the radius of the shock, and then to look for a solution where the radius of the shock goes as the power of time. So everything is scaling as a power law of time as the shock comes in.

In that limit you look for density being a constant times the function of the dimensionless parameter. The velocity of sound squared is r2/t2 -- that is the only quantity with the dimensions of speed of sound squared times the function of the dimensionless parameter, and velocity of the fluid being r/t times the function of the dimensionless parameter.

If you assume that everything depends on this dimensionless parameter and that the radius of the shock goes as time to a power, which was known by G. I. Taylor -- he did the opposite problem of the exploding shock and he knew hydrodynamics becomes exact to high amplitude, even high amplitudes generated by a hydrogen bomb, so when he saw on film the explosion of a hydrogen bomb, he immediately could match it to the solution to these scaled equations and tell them exactly what the yield was, to the consternation of a lot of people, because he published it in the open literature right away.

(Laughter)

So we are heading back into the hydrodynamic regime, perhaps, as the bubble collapses. If we substitute these expressions, these dimensionless forms, into fluid mechanics for an imploding shock, the key fact is that r and t drop out and you end up with ordinary equations involving only dimensionless variable, which is the radius coordinate over the size of shock.

(Transparency 56)

I will write them down just to show you that these dimensionless equations exist. The key fact here is you have now got three ordinary differential equations for Z, ∇, and G. And everywhere you have a derivative it is only with respect to

the dimensionless variable. So these three ordinary equations now possess a solution and there is also an eigenvalue problem. The solution is finite for only one value of alpha and that is the value that determines the form of the implosion of the shock wave.

(Transparency 57)

The key fact from all this is that in the limit of very high amplitude fluid mechanics possesses an imploding shock where the radius goes as time to an alpha power, where alpha, for our purposes, is 0.7.

You can take from these expressions, you can find the asymptotic value, the value for large dimensionless variable. You find, for instance, that velocity goes as r to the minus one over alpha minus one. That is one over the square root of r, roughly speaking. Pressure goes as one over r, roughly speaking.

That means at the moment of focusing everything is asymptotic, because the radius of the shock has gone to zero, so at the moment of focusing of the shock wave, when the shock wave has collapsed to the origin, pressure diverges to infinity and velocity diverges to infinity. As the shock comes in, it goes faster and faster until you hit some incredible singularity when it focuses on the origin. So that is the first key fact and how it comes about.

(Transparency 58)

The second key fact is shock fronts. I am glad Steve spoke about this yesterday. Here is your shock front moving with velocity, U_s, continuity of mass flow across the shock front, namely, the mass to flow through is equal to mass to flow that way, continuity of momentum flow, continuity of energy across the shock front. If we use an ideal gas and use continuity of mass flow, momentum, stresses, and energy, and we introduce a Mach number, which is the velocity minus the speed of the shock front over the velocity of sound on the left or the right, then we find the key result, which is the next idea we need to work with.

We need two ideas. One is the radius of the shock as a function of time. The other is the temperature jump across the shock wave goes as the cp over cv -- that is your gamma times gamma minus one, gamma plus one squared times the square of the Mach number. That is the coldest one. The temperature jump across a shock wave goes as Mach number squared.

To be precise, we are looking at the situation where the gas here is at rest, the shock is moving into a gas at rest -- think of the bubble, the gas is at rest, the shock is moving

into the gas at rest. In that case, this Mach number is the speed of the shock wave and the temperature jump goes as Mach number squared.

Here is the ticket. We are trying to construct a scenario. This is our straw man scenario that we keep in mind when we are trying to motivate new experiments in sonoluminescence, and the goal of any new experiment, of course, is to disprove it.

The scenario is that the bubble is collapsing, it hits a supersonic velocity, it hands over to an imploding shock wave. As the shock wave implodes, it speeds up, the Mach number goes up, the temperature jump goes up as Mach number squared.

But worse happens. Then the shock wave goes through focusing. Then the gas which was behind the shock wave as it focuses is now in front of the shock wave as it expands and you get an extra factor of Mach number squared. The temperature jump goes as Mach number to the fourth. So right after the moment of focusing there is a catastrophic moment when all the gas which thought it had already made one pass through the shock wave and gotten heated is suddenly heated from that high temperature up by another factor of the Mach number instantaneously.

(Transparency 59)

Here is the picture. We have the radius of the shock which goes as time to the alpha power. The problem is to calculate A. We call that the Hubble constant (that is not contained in the theory, that is your launch condition, that tells you how it started).

The temperature jump goes as Mach number to the fourth. The Mach number of the shock wave is simply the time derivative of this, so it goes one over time to the one minus alpha, just take the time derivative of the imploding radius. From that you can then get the time spent at each temperature, because you can now, from this expression, find the Mach number for a given temperature. All temperatures up to infinity are reached, but as you get to higher temperatures, they are far shorter times. From the time spent at each temperature you can then characterize the spectrum of the system for different assumptions. The time spent at 100,000° is roughly 100 psec.

(Transparency 60)

So this is consistent with what we know about sonoluminescence. The shock wave is in the region of 100,000° or greater, which is more or less consistent with our spectral measurements. That would give us our ultraviolet light for about 100 psec and the radius of the shock at that time is about 1000 Å, so according to this picture the light comes out

as the shock wave passes through 0.1 μ and it is at that region or smaller for about 100 psec.

If you now get courageous and follow it in, at 5 psec the shock radius is 200 Å and you get soft x-rays, 100 eV, and at 0.1 psec, if the shock should exist down to such ridiculously short scales, is 60 Å. You are hot enough for fusion, but the question is how much, because that is a pretty small bubble at that point. There might not be anything left.

(Transparency 61)

But it turns out it is remarkably sizable. If you use the known formula for neutron emission, which goes as the radius of the bubble at that time times the cross-section for neutron emission, you get on the order of one neutron a second if the shock should exist down at such ridiculously small scales.

The interesting parallel here is that, of course, that has not been looked for, that is a risky area to go into, but what is curious is that the issue of how much energy focusing can be achieved with sonoluminescence has become, from this perspective, identical to the issue of how do you do inertial confinement fusion? Those guys are constantly suffering with how do they make a better imploding shock wave.

Here the acoustic field gives us a beautiful spherical implosion, because the acoustic field positions the bubble at the velocity node and the velocity node is the point of maximum sphericity of the sound field to help keep everything in tow. In their case, they build these very complicated gold devices in order to position this little tiny 100-μ pellet in just the right location and, of course, they get two or three implosions on a good day. We get 30,000 per second, and that helped those numbers along there.

(Transparency 62)

Back to the SL graph.

Q: Have you looked at what kinetic theory is doing during all this?

A: No.

We think that a key aspect of the sonoluminescence is the diffusion of mass. If there are two solid theoretical perspectives I would like to leave you with, one would be the Rayleigh-Plesset equation, because it describes so much of the bubble cycle and how, from Bernoulli's law and the solution to del squared, the velocity potential is zero, you can get such a versatile equation.

The other would be the issue of mass diffusion into and

out of the bubble in the sonoluminescent state. We think there is a curious difficulty here, and I would like to try to show you a bit about it.

If we look at the radius of the bubble versus time, it goes from the ambient to the maximum, back down. There is gas flowing into and out of the bubble. In the steady state this mass flow into and out of the bubble had better balance. Let me describe where it is coming from.

When the bubble expands, the pressure inside the bubble goes down. It goes below the partial pressure at which the gas was dissolved into the liquid, so gas dissolved in the liquid flows into the bubble. When the bubble comes up here, gas is flowing into the bubble. Down here the pressure is 1 Atm inside the bubble. If the pressure in the bubble goes to 1 Atm, that is larger than the pressure at which gas is dissolved in the water, because the system is always run de-gassed, so gas flows out of the bubble into the water.

The balance between these two should be achieved in the steady state. Otherwise, the bubble has to dissolve or expand and sonoluminescence has to run down. An amazing aspect is that we have steady sonoluminescence. Through use of a mode-locked arrangement it can run for hours or days, overnight, so the mass flow is balancing.

But there is a contradiction. The amount of mass to flow in here, using diffusion theory, does not balance the amount of mass to flow out there using diffusion theory. So there is some process, which I would like to motivate now, which goes beyond the diffusion of mass, as I have just described it, which accounts for and makes it possible to have sonoluminescence.

(Transparency 63)

I would like to revisit the old equations of Eller and Flynn on rectified diffusion. The key quantity here is c, which is the grams per cubic centimeter of gas dissolved in the water. The diffusion equation is that DC dt is del squared c, that is your standard diffusion, and then we have the convective term, v. gradient c, but v is the velocity due to the motion of the bubble. It is the gradient of the velocity potential and goes as $1/r^2$.

The goal is to solve that equation and see what the consequences are of the solution. The key parameter for the solution is the diffusive penetration depth, which is the square root of the diffusivity coefficient of mass of the gas molecules in water divided by the acoustic frequency.

If this is small compared to the radius of the bubble,

then we can make a transformation, following Eller and Flynn, which converts this equation with the convective derivative into a new diffusion equation and different variables. The transformation is to introduce a stretched coordinate and a stretched time and we get a diffusion equation in terms of these new variables, h and tau.

The solution to that equation is that the concentration at position and time can now be expanded in a Fourier series, as usual, and the key point is in the steady state the a_0 must vanish. Otherwise, there is a DC flow in one direction or the other. So if the a_0 must vanish in the steady state, that means the time average of this must be zero.

(Transparency 64)

Therefore, if you permit me to use Henry's law -- I am sorry, I did not mark that down -- the concentration is the saturated concentration times the pressure of the gas divided by the saturated pressure. So the concentration is c_0, the saturated concentration of gas in water times the ambient radius cubed over the radius cubed of the bubble at that moment.

The gas inside the bubble is at a certain pressure, depending on where you are in the cycle. That pressure of the gas in the bubble then determines what the concentration is in the water just right near the bubble through Henry's law. That becomes the boundary condition.

Then we want to know, with that boundary condition, as that concentration goes up and down with time with the motion of the bubble, how the gas flows into and out of the bubble.

(Transparency 65)

Requiring that a_0 vanish, that you have a steady state, tells you that the concentration far from the bubble, which is the concentration that you impose on the system when you prepare the gas and the water, divided by saturated concentration is roughly three times the ambient radius over the maximum radius cubed.

DR. GARRETT: Aren't you making the same mistake that Rayleigh made in assuming in that concentration that it is isothermal? You are talking about the pressure being only a function of radius.

A: You mean thermodiffusion in addition? Sure, there is thermodiffusion, barrodiffusion, and diffusion. We have looked at those quantities and we think for the case of the gas in the water the thermodiffusion is really very, very small, because the temperature is really getting hot for only a very short

time.

(Transparency 66)

If we go back to Brad's waterfall data, in this region the ambient radius over the maximum radius cubed is, in fact, a good match to the concentration of gas in the water. As you move into the sonoluminescence regime, the ambient radius over the maximum radius cubed becomes much smaller, because the expansion ratio goes up and you get a contradiction with the diffusion.

An interesting sidelight is that you can, from the radius time curves of the bouncing bubble, now use that as a meter to determine what the concentration of gas is in the liquid by inverting this procedure.

(Transparency 67)

For the bouncing bubble, R_0 over R_m is 2.5. You cube it, multiply by 3 and you get your fraction of an atmosphere at which the gas was dissolved in the liquid. For the sonoluminescent bubble R_0 over R_m is about 10. Cube it, multiply by 3, and you get an incredibly small number, which means that there is some other kind of mass flow going on, there is a contradiction there, and there must be some steady flow of mass from the bubble to infinity.

(Transparency 68)

You can calculate what that flow must be. The conclusion is that in order to maintain the steady state for sonoluminescence there must be a nondiffusive ejection of mass on each cycle and the amount of mass to be kicked out on each cycle is roughly a part in 10^4 of the mass in the bubble. But that adds up, right, that is less than one second. I mean, that is more than 10^4 cycles per second.

So if there is some other process which kicks mass out in the sonoluminescent state, and this process is what accounts for the abrupt transition to the sonoluminescence, where the radius shrinks, whatever is causing this is also telling us why we can have sonoluminescence only in water or is accounting for the transition to the sonoluminescence. There is some nondiffusive sudden mass flow that characterizes the sonoluminescent state.

DR. ATCHLEY: Seth, in going from the nongray area to the gray area on the parameter space plot, is there any obvious change in the bubble, if you change the pressure amplitude by just 10 percent and the bubble looks the same if it is flowing? When Felipe did his initial work, there was a lot of jumping around and jittering and apparent shedding of mass in that

transition. Did you see that?

A: Yes, you can see --

(Transparency 69)

-- up in this region, right around here, not so much down here, but right around here you can see some jittering around of the bubble.

Q: Did that happen with the argon, too?

A: Yes, argon tends to jitter everywhere, and that jittering has not been quantified. I can try to give you a picture of it. We cannot quantify it. We can guess, but it would be just building one scenario on top of another scenario.

Q: Is that motion chaotic?

A: It is hard to predict.

(Transparency 70)

Another reason the people working on inertial confinement fusion are very interested in the sonoluminescence is that we apparently get a spherical collapse for free. If the bubble developed convolutions, such as shown in brown here, then you would not get such a nice collapse -- it might splatter --and the bubble might not come back for another helping for the next cycle of the sound field a few microseconds later.

In our attempt to understand why we are getting this for free, where they struggle with hundreds of millions of dollars in order to fight Rayleigh-Taylor convoluting instabilities, we just thought that this effect would somehow, if it existed, help us understand the lower threshold for the sonoluminescence or the upper threshold for sonoluminescence where it disappears.

Not only don't we understand how, as you increase the sound field level, you get this abrupt transition to the sonoluminescent state with a small bubble, we also do not understand why, when you increase the sound field level to the largest bubbles, which give off the most light, there is an upper threshold when the phenomenon disappears. We do not know what the physical process is that takes it away.

We thought perhaps the convolution of the bubble would have time to build up when you had a bigger cycle with a larger radius and you might get some wiggles on the surface and so on. On the side I would like to say probably this is the best candidate for understanding the jittering that Anthony was bringing up a few moments ago in the bouncing bubble.

To make the long story short, our opinion is that the Rayleigh-Taylor or convolution instability is not a problem for sonoluminescence. The reason we believe that is the case is that we have a situation where, initially, the heavy liquid is on the outside and the lighter stuff is on the inside, liquid pushing gas.

For the inertial confinement fusion they have exactly the opposite situation. We think that is the reason, but here is the theory, and I would like to show you roughly how the theory goes. First, we have an average radius, R-bar, and then we have a radius as a function of theta, which is your convoluted radius whose equations we want, so we expand R theta in terms of the Rayleigh-Plesset variable, R-bar, plus now the convolutions expanded in terms of your spherical harmonics and these coefficients depend on time. The goal is to find how these coefficients vary with time to see if there is any instability that builds up.

The velocity potential on the outside of the bubble we will take to satisfy $\nabla^2 \varnothing = 0$ and so the velocity potential on the outside has the term from the Rayleigh-Plesset equation -- we are neglecting acoustic radiation in this problem -- so we then expand this in terms of solutions to Laplacian phi equals zero in terms of the YLMs and the appropriate powers of your distance from the center of the bubble.

The boundary condition, again, is that the velocity at the surface is equal to $\dot{R}$. This boundary condition, combined, again, with Bernoulli's law, but now we are equating the angular components of Bernoulli's law, gives us, then -- let us look at this equation here.

(Transparency 71)

It gives us the Rayleigh-Taylor equation appropriate to being driven by the Rayleigh-Plesset equation. The acceleration of the convolution times the radius is balanced by the acceleration of the radius times the convolution.

Here is the key point. There is a Hamiltonian which underlies that equation. The Hamiltonian is the momentum of the surface convolutions divided by the mass, ρR^3 cubed minus the acceleration of the bubble times A squared. Notice, when the acceleration is negative, which means you have the collapse, this term is a restoring term, since L is equal to two or more. If this is a restoring term, it is stable.

So during the collapse, which is the key time, $\ddot{R}$ is negative (being directed in) and so this term is also negative, therefore the effective potential is positive, so it is a harmonic oscillator, P squared plus Q squared, as we would say.

Now, however, the mass is a function of R, so there can be a slight change of scale increase in A, and there is. A does increase, not exponentially, because this is stable in that regard, but it increases as a weak algebraic function of R. It goes as one over R to the one-quarter times the initial value of the convolution. But this is very small. This one-quarter power knocks out any effect. If this had been the fourth power, that is another story. So the change in bubble size does introduce a kind of change of scale which does permit some kind of amplification, but there is no exponential amplification.

I cannot see how, within a cycle, the Rayleigh-Taylor instability will lead to an effect. And here we have not even included shear viscosity.

(Transparency 72)

If we put in shear viscosity, there is an enormous damping term which damps things as e^{-15} even during half a cycle. So it would appear that the Rayleigh-Taylor instability can be -- it is going to have to be something very tricky and clever to invoke this instability. The more we look at it, the more it appears as though it should have been obvious that there was sonoluminescence.

(Transparency 73)

You notice I have not said much about the light-emitting mechanism. I think there are three mechanisms that exist in the literature. One is black body, the other is Brehmsstrahlung, and finally there is accelerated zero-point motion that Julian Schwinger has been writing about.

Let me talk to you about black body and Brehmsstrahlung. I believe Brehmsstrahlung is a better picture than black body. This is theoretically speaking; the data stand for themselves. To achieve a black-body spectrum it is necessary that the photons have a chance to interact with the matter, that the photon matter mean free path be small compared to the size of the black body. That is the only way you can have a situation where the radiation for the surface dominates.

In plasma physics they never worry about black-body radiation. A tokomak in the small diameter would have to be 600 miles for them to have -- that is their photon matter mean free path. I cannot imagine how, in a 10th-of-a-micron bubble, we are getting photon matter mean free paths of less than a 10th of a micron, which is even smaller than the wavelength of the light, but we have to keep it in mind there can be surprises.

Brehmsstrahlung is the light emitted from charges which accelerate as they collide at high temperature. If you increase the temperature of a system, the matter starts to ionize, the electrons separate from the ions. These charges then move with the thermal velocity of the square root of kT over m. As they move with the thermal velocity and undergo collisions, the acceleration radiates light.

In the Brehmsstrahlung radiation the power of an accelerating charge goes as the acceleration squared. If we have an impact parameter, we can then calculate the radiated power in a collision in terms of the minimum impact parameter. Then the books do an amazing thing for the Brehmsstrahlung.

These plasma physics books say that the minimum impact parameter is the determined by the thermal deBroglie wavelength -- h-bar comes in. I am just exposing my naivete here. I cannot understand why, in these very high-temperature systems

[and if you look at all the plasma physics books, they have got an H-bar right there for the Bremsstrahlung power, and Bremsstrahlung power radiation is a problem for the tokomaks, they do not want the energy to come out] the formulas have h-bars. I just cannot see why something else does not cut it off rather than h-bar, but that is the accepted formula for radiated Bremsstrahlung power, h-bar and all.

If you would take this formula and puts 1's or 2's here and temperatures of 100,000°, you get within factors of three the integrated power of the sonoluminescent flashes. So there could be some truth to the Bremsstrahlung picture.

Q: What is the n_e?

A: I am sorry, n_e is the number density of free electrons, number density of free ions, which I took as being

-- one or two per atom. I took here 10^{22} or 10^{23}, something like that.

DR. GARRETT: I would think V min would be a property of the interacting ions and there would not be a universal --

A: But they do not want to choose the size of an electron.

DR. GARRETT: You have some effect of scattering cross-section and that is being hidden in the me, the effective mass that they throw into --

A: Yes, I would sometime like to go back there and understand all these. I can give you a million references for that one over h-bar in the plasma literature, and this is for

their 100 million° plasmas.

(Transparency 74)

I promised to make some comments about what it is good for. I think we do not have to say much more about the fact that we have an ultraviolet light source with picosecond-pulse widths; and it's cheap.

The high quench rates of about 10^{14} K per second during the collapse can lead to unusual chemical reactions, such as the synthesis of amorphous iron that Ken Suslick has patented. I am very much intrigued by the use of cavitation and aspects of the bubbles for medical imaging. As I mentioned before, the one dramatic thing that you first notice when you try to set up sonoluminescence in a flask is that if you inject one bubble of air into a liter flask with a sound field, the sound field goes really bad very fast.

So if you could find ways of injecting tiny bubbles of air into arteries, it would be an incredible contrast agent for ultrasound but, of course, injecting bubbles into arteries is questionable. When we first saw the sonoluminescence, we said, gee, this is a great excuse to ask the DoE for $200,000 for a streak camera to study the flash widths. But the biologists, when they see the tiny bubbles there, they say, "Oh, we should pour some egg white in."

So they pour some egg white in and the egg white coats the bubbles and cooks. Then you have a 5-μ bubble of air inside an albumin shell, which you can inject into an artery. They go with this idea to MBI, Molecular Biotechnologies, Incorporated, and they get $5 million for their pockets. The idea is the following.

If they get these 5-μ bubbles encapsulated in albumin, they can make a whole bunch of them --it is called Albunex -- and inject that into the arteries and then there is no problem of emboli. Five microns is smaller than the size of a red-blood cell, so these 5-μ bubbles can go right through membranes and escape and there is no problem at all. Unless other techniques come on board, this is going to revolutionize acoustic imaging.

The other advance, it was found that when these Albunex bubbles are coated with a monolayer of cholesterol esters then they go preferentially to tumor sites. They can, for instance, then be used as little capsules delivering medicine to tumor sites or, in the case of various rat tumors, they just insulated the bubbles and they blasted the tumor sites, whatever good that was, but there is an enormous amount of work on this by D'Arrigo and Simon.

Let me turn now, in the last couple of minutes, to the unsolved problems and perhaps looking ahead.

(Transparency 75)

First, I think the sonoluminescence is amazing because you are transducing sound at the kilohertz level into picosecond flashes of ultraviolet light. The energy scales which are mixing are amazing. The ratio of the acceleration at the top of the curve to the acceleration at the minimum is 10 orders of magnitude. This is nature's most nonlinear oscillator. If we could understand it, then it would be a great paradigm for energy concentration in other systems.

(Transparency 76)

So what are the problems that stand in the way of understanding it? First, why water? I think that is a good practical key. Why have all our investigations been limited to water? What is the theory of the ambient radius? The key parameter that is measured and characterizes all the graphs and simulations is the ambient radius R_0. We do not know a theory for R_0. The reason we do not have a theory for R_0 is that we do not understand the anomalous mass flow which characterizes the sonoluminescent state.

For the bouncing bubble, which does not give off light, we know R_0. We can calculate and predict R_0 in that case from the diffusion equation. In the sonoluminescent state we do not know R_0. I would say that is the key unknown. It may be connected with why water. Everything is wrapped together.

What is the spectrum beyond 6 eV? We have not figured out the experimental trick which enables us to look beyond the cutoff of water. There is light out there but we cannot calibrate in that region and by the time you get to 7 eV water has completely cut off.

One way around this is to get sonoluminescence in other liquids which transmit light to higher energies, but we know what the problems are there.

We do not understand the upper threshold of sonoluminescence, why it turns off above a certain sound field amplitude. Most of all, we do not understand why such an incredible energy-concentrating mechanism should be so robust in spite of all the incredible number of physical properties which contribute to it.

Thank you very much.

Q: I am thinking about a collision. The percentage of the air you have to have to have the light and the temperature

dependence, all are pointing to the collision of all the gas. In objects, if you want to build a gas laser, you have to have a certain percentage of various kinds of gas. You have to have the gas percentage to really have the laser going, and this is because of the collision of different molecules.

A: So you are saying that the grand unified similitude theory should have also had the excited states of the various atoms. With the laser that is what you are playing with.

Q: Like in H_20 and N_2, the molecular weights are compatible. You transfer energy real efficiently. It probably is not because of the energy level; it is more like similar weight molecules, so they collide. I think the collision should consider the various gases, not just consider them as a whole just a gas.

A: So what experiment would you say we should do to check that?

Q: I think a theoretical calculation, that with a certain percentage of gas the collision --

A: I would not know how to do the theory, because everything is so dense and so hot. Researchers working on the bomb codes at Livermore got involved in this, because they thought they could contribute something with regard to the equation of state of gas inside the bubble. And even that is not known. They have been through three iterations and each time the equation of state changes. I would say we have a problem knowing what is the precise equation of state that characterizes the system.

Everything is within a van der Waals hard core, so you cannot think in terms of a dilute gas where things are colliding. A collision implies a two-body problem and this is going to be an n-body problem.

Q: The other night I was asking you what analogy there was between what was going on during the collapse of the bubble and a supernova explosion. I was wanting to know, I guess you were saying the supernova people might be trying to use this to try and understand supernovas better, type II supernovas?

A: Some of the researchers in that area at Caltech have asked for various details. I think their interests go in that direction, but I do not know what they are pursuing.

Q: Their theory about what goes on in those events is not very good, either?

A: Right. They certainly felt that they -- they were contacting us because they thought some of the details here

might help them. I asked them to recommend an experiment based on what they are interested in and they did not have anything in particular to request.

Q: Have you ever had any other kind of _____ that has higher efficiency in transmitting sound into light?

A: The xenon, the cold xenon bubble, I think, perhaps has the highest efficiency.

Q: When you normalize that, it is still below _____.

A: I did not show any graphs of efficiency. The efficiency in general depends on how you define it. For the best case, if you take the definition in terms of energy stored over energy released, it is about 0.1%. That is about the best we have gotten.

DR. GARRETT: I was curious about how the sonoluminescence stops at the high amplitude. Apparently sonoluminescence stops before you get to amplitudes sufficient to blow the bubble apart into several tinier bubbles?

A: I do not think that is the case. Let me see if I can find the graph.

(Transparency 77)

I can take 10 seconds and show you a graph of the death of a bubble. If you suddenly boost the sound field at time zero, first the radius increases -- here -- and then the bubble crashes. So first, the intensity of light emission goes up and then the bubble crashes. It even has a little glurp here where it tries to come back, but then it dies.

So these are measurements of the radius versus time and the intensity of light coming out is a function of time when you burst the sound above the upper threshold. It is the transient death of a bubble.

DR. MAYNARD: First, let me tell you that the viewgraphs and the handouts are the not the same, so as you watch the viewgraphs, do not try to follow in the handouts. The handouts are different for a reason, and that is: on the viewgraphs I am going to be describing some concepts and some experiments. The experiments are pretty easy, you can remember those, and they are in the reprints that are in the notebook. What is not on the viewgraphs are the grubby details. Usually you just give references for the grubby details, but the concepts that I am going to be talking about were sufficiently exciting that physicists were running around trying to get prizes for this stuff, and nobody took the time to sit down and write down the grubby details. So since I did not have any such constraints, I sat down and wrote down the grubby details, and they are in the handout and they are not available in any store. (Laughter)

(Transparency 1)

A fundamental problem in solid state physics is to solve the Schrodinger wave equation for an electron in the potential field of some 10^{23} ions or scattering sites in a solid. Now, 10^{23} is a big number. To solve this problem probably would take Mathematica a little while.

But fortunately, the solid could be a crystal, in which case the ions or the scattering sites are arranged periodically. In that case, you have to solve the Schrodinger wave equation for only one unit cell, and that is something you can do. Electrical engineers have computer programs that they can use for calculating the band structure of a semiconductor and so on. The bad news is that a lot of materials like alloys or amorphous materials are highly disordered. If they are highly disordered, then you are back to having to solve the Schrodinger wave equation for some 10^{23} nonperiodic scattering sites, and that is a big problem.

Fortunately, in the real world most experiments are done at finite temperatures. At finite temperatures, the ions or the scattering sites are moving. In other words, in a solid you have phonons running around. The electron can inelastically scatter off of that phonon and that destroys its phase coherence. You can do statistics. Because of this inelastic scattering you can assume that certain things will average to zero, and so by doing the statistics you can convert the wave equation with some 10^{23} parameters into a Boltzmann equation or diffusion equation with only a couple of parameters, maybe just a mean free path for the inelastic scattering. That is an equation that you can do in Mathematica. Because of this inelastic scattering, you get an

equation you can solve and, because of this, solid state physics has been able to continue on its merry way for these disordered materials.

Recently a couple of things have happened that have changed all this. One thing is: the devious experimentalists made their samples so small and so cold that the size of the sample was smaller than the mean free path for the inelastic scattering. Without the inelastic scattering the statistics are not valid any more, and you are back to having to solve the wave equation for some 10^{23} nonperiodic scatterers.

That is actually a problem that dates back to Lord Rayleigh. It is basically wave propagation in a disordered array of scatterers. A lot of people have worked on that problem since Lord Rayleigh and it is recently that, motivated by these experiments, other theoreticians have worked on the problem and have actually made some progress.

These contemporary theoreticians do not want to admit that they have been working on a problem that Lord Rayleigh worked on, namely, the wave equation, so what they say, instead, is that they have been working on solving a problem that has long-range phase coherence. They do not actually solve the wave equation; they go back to the equation they can solve, the diffusion equation, and do perturbation theory back toward the original wave equation. They introduce quasi-particles called diffusons and so on. They try to include as much phase coherence as they can.

There has been a lot of excitement and activity in this field. The buzz words to look for are "Anderson localization," which I will be talking about, "universal conductance fluctuations," the "Ahronov-Bohm effect" and, most recently, "normal electron persistent currents."

I said that a couple of things happened that changed this. Another thing that happened is that quasicrystals were discovered. These quasicrystals are made with an aluminum alloy that is rapidly quenched, and originally you got very small samples, maybe 100 μm in size. Quasicrystals have unusual properties and I will tell you more about quasicrystals later.

Let me go back to these small cold samples. The experiments that people did involved making very small wires, about 100 Å across and only a μm in length. They would measure electrical conductivity in these small wires. For that kind of scale of a wire, they had a name for that physics, and that was mesoscopic physics.

(Transparency 2)

And that is the study of phase coherence on the scale of mm. In our research, we also coined a phrase, and that is megascopic physics. That is looking at phase coherence on the scale of millions of μm. A million μm is a meter, so we like to claim that we can see phase coherence in a 1-D wire that is not a mm in length, but more than 10 m in length.

When we look at the density of states and the properties of a quasicrystal, that is not 100 μm in size but more than a meter in diameter, easily the world's largest quasicrystal.

Both of these statements are worded so as to shock and give heart failure to people who actually work in these fields. Our experiments are classical analog systems. Of course, they are acoustic analog systems. Our 1-D wire is not one of these 100Å wires. Our 1-D wire is actually a piece of steel wire that we bought at the music store. Our quasicrystal is made out of aluminum alloy, but it is the aluminum alloy you find in the machine shop.

You can think of these acoustic analog systems as analog computers, and they have some very definite advantages over the experiments that are done on the quantum mechanical systems. First of all, they are very precise analogs of the quantum mechanical systems. Remember, without the inelastic scattering, you are dealing with a time-independent Schrodinger equation.

You can rewrite the time-independent Schrodinger equation like this. This is just an eigenvalue problem. In quantum mechanics the q, this eigenvalue parameter, is the $\sqrt{\frac{2mE}{\hbar^2}}$, and E would be the energy eigenvalue in a quantum mechanical system. In acoustics q would be $\frac{\omega}{C}$. When you solve this equation, you get eigenfunctions and eigenvalues; in quantum mechanics you get energy eigenvalues, and in acoustics you get eigenfrequencies or resonant frequencies or natural frequencies.

Anyhow, the point is, you can rig up acoustic systems that are mathematically identical to the quantum mechanical systems. All of the physics is in this potential field, and the potential field could be periodic, disordered, or quasiperiodic.

In the acoustic analog experiments, all the conditions and the parameters in the experiment can be very precisely controlled, or at least measured, which is hard to do in a 100Å wire. In a measurement you can measure the eigenvalues, eigenfunctions, densities of states, and all those things directly. In a quantum mechanics experiment,

essentially it is impossible to measure eigenfunctions directly.

All of these things here you can do also with digital computer simulations. But in the acoustic experiments you can do some additional things. You can study time-dependent effects and nonlinear effects. That is, this potential field can be made time-dependent. It can also depend on the amplitude of the wave function, so you get a nonlinear equation. These things, time-dependent fields and nonlinear effects, are very difficult to simulate with digital computer simulations. So with these acoustic analog systems you can do serious contemporary physics literally using balls and springs.

Before I tell you about the experiments that we did and the physics involved, I want to show you how we got involved in this business. There is lots of acoustics at Penn State and one of the things that people do is: they study noise control.

(Transparency 3)

Here is a problem in noise control. Suppose you have a plate, like the floorboard in the little aircraft that fly into Monterey. At one end of the plate you have a source of vibration that puts transverse waves in the plate. They travel down the plate and, as the plate vibrates, it makes noise, and makes an annoyance.

Usually, for structural reasons, a plate will have a rib on it, and that rib reflects the vibration, so there is less vibration transmitted, and less noise at the other end of the plate. So the engineers up north at Boeing said, well, if one rib reflects the vibration, why not a whole bunch of ribs?

(Transparency 4)

Possibly for ease of manufacture or aesthetic reasons they put down a nice periodic array of identical ribs. What they found was that the first couple of ribs reflect the vibration, but all the rest transmit without any further reduction. So this is analogous to something that you learn in solid state physics for an electron in a metallic crystal.

In the metallic crystal you have an electron and positive ions arranged periodically. An electron is very strongly scattered by a positive ion; it is the strongest force. If you use this strong scattering to calculate the electrical conductivity in the metal, you would have to conclude that the metal should have very low electrical conductivity, because the electron just keeps getting deflected out of its path.

Of course, metals have very high electrical conductivity. There are two reasons. One is, the electron behaves like a wave, and the other reason is that the scatterers are arranged

periodically, the same as up here: these ribs were arranged periodically.

(Transparency 5)

Both of these effects can be explained with a theorem that you can derive with group theory. In mathematics it is called Floquet's theorem. Of course, in physics we had to rename it, and it is called Bloch's theorem.

What Bloch's theorem says is that if you have a system with a periodic potential or, in acoustics, a periodic impedance, then the eigenfunctions are extended. That is, they have this particular form. There is a phase factor, e^{ikx}, that looks like a plane wave, times a function which is periodic with the period of the potential field.

So if you take the modulus of this eigenfunction, then the phase part just goes away and you are left with just this periodic part. Because it is periodic, it has the same nominal amplitude everywhere throughout space.

In quantum mechanics, for the electron in the metal, if the electron eigenfunction, its wave function, has the same nominal amplitude everywhere throughout the system, then the electron has equal probability of being anywhere in the metal. In other words, the electron can travel freely through the metal, and that gives it its high conductivity. So the fact that this thing has the same nominal amplitude everywhere is the extended eigenfunction.

I have taken a textbook (this is actually Tinkum's group theory textbook) and copied some Bloch eigenfunctions. You can see that if you took the modulus of these you would get the same nominal amplitude throughout the infinite system.

Actually, Bloch's theorem and Bloch functions are taken for granted in solid state these days. They are a lot more complicated than the textbook would tend to show, and some of the grubby details are in the handout so you can see what it is really like.

Q: I wanted to ask if the vibration in the front of an airplane, if you decrease the length between the ribs, do you prohibit vibration? It would stop the vibrations from traveling through to the passengers? On your previous slide.

A: The only point I want to make here is that they used a periodic array that had special properties. That is all I want to get to here. The actual spacing I do not know about.

(Transparency 6)

This just restates the situation. If you have a periodic potential field, then the eigenfunctions are extended. So now you can ask what happens if you add a random part to the potential field? I would have bet money that putting the random part in, you would just get random variations in the amplitude of the eigenfunction, so maybe this one goes up, this one down; just random variations.

But it turns out that that is not what happens. What happens is the eigenfunctions become exponentially localized, that is, there is some site in the system where the eigenfunction has a maximum value and then as you go away from that site, the eigenfunction amplitude exponentially decays, and there is some characteristic length here which is the localization length.

Phil Anderson and Sir Nevill Mott used this to explain the metal-to-insulator transition for disordered metals. This is called Anderson localization. It is a very important effect. Phil Anderson got the Nobel Prize for this idea.

The fundamental idea involves wave propagation in disordered potential fields, and that dates back to Lord Rayleigh. Lord Rayleigh worked on the problem and he almost got Anderson localization, but he did not have Mathematica, so I guess he did not completely get the problem.

Since Lord Rayleigh a lot of famous people worked on the problem: Wigner, von Neumann, Freeman Dyson. They made a lot of progress. The first rigorous theorem for Anderson localization in 1-D was by Furstenberg in 1963. That is Furstenberg's theorem and that is a fairly famous theorem. Furstenberg was a mathematician.

Anyhow, 1963 is quite recent on the time scale of Lord Rayleigh, so this has been a long-outstanding problem and it has been solved in only 1-D. In two or three-dimensions there are no rigorous theorems, but there are some compelling arguments.

When I first heard about this, I wanted to try to understand Anderson localization. As I said, I would have bet money that putting in disorder would just give you random variations in the eigenfunction amplitude. So I looked in the literature. There are thousands of papers written about Anderson localization, but most of them use Anderson localization as a way of explaining some observation in solid state physics. I tried to read Furstenberg's paper. Furstenberg's paper was 25 pages of theorems and lemmas. There were references to other papers by mathematicians full of theorems and lemmas. I am an experimentalist. I have a lot of trouble reading a paper that starts off with "let M be a manifold." (Laughter)

So after quite a while I finally found a paper that gives a beautiful explanation of the Anderson localization. Actually, Furstenberg's theorem, the statement of the theorem, is quite easy. You can understand it. What the theorem says is that in a disordered system, as you go off to $\pm\infty$, the eigenfunctions decay exponentially to zero with probability one. I can understand those words.

The problem is that, as an experimentalist, I do not have an infinite-size system. All my systems are finite in size. Even if I had an infinite system, I would have to march off to infinity and look for something happening with probability one. Well, how would I know if it happened or not?

Anyhow, the problem is to understand what all this means for a finite-sized system and what the probability means and so on.

That is all in this paper by Marshall Luban. I have taken Marshall Luban's explanation and sort of summarized it in the handout, so you have that explanation. It is a very satisfying explanation.

This Anderson localization effect has to do with the statistics of this disordered potential field. In statistics, one of the most important parameters is the dimensionality of the system. It turns out that in 1-D any amount of disorder will give you some degree of Anderson localization. If the disorder is small, then the localization length may be large. If you have a finite-sized system, the localization may be bigger than your system, in which case you do not really see the localization. As you increase the amount of disorder, the localization length gets smaller and then you do see the Anderson localization.

In any case, for an infinite system in 1-D any disorder will give you localized eigenstates. In 3-D you have to exceed a certain critical amount of disorder before you get localization. So for a small amount of disorder you still get extended eigenfunctions. You have to exceed a critical disorder before you start getting the exponentially localized eigenfunctions.

Q: With a periodic medium and with random distributions, are you going to have more than one localized spot, one spot where the eigenfunctions are strong?

A: There is going to be only one place where you have a maximum amplitude. That is where you have to start worrying about infinite-sized systems and so on.

Q: If you have an envelope there with a peak and it dying off, do you have another envelope elsewhere?

A: What happens is, you can have places where it goes up, again, above this envelope, but you would just count that as a fluctuation. It will still go down to zero with probability one, but it can fluctuate above it. That is the sort of thing you have to deal with in an actual experiment, and that is what is treated in Luban's paper.

That is 1-D and 3-D. 2-D is the so-called critical dimension. At first, the theoreticians were not sure whether there was a finite amount of critical disorder in 2-Ds or not. They think they understand it now, but at least when we first learned about this, that was an open question. So we thought we would do an experiment on 2-D Anderson localization.

(Transparency 7)

We thought we would just take a large plate, an aluminum plate, maybe 1/8" thick, maybe a meter or so in size, and put a shaker on the plate and drive it into its normal modes of vibration, excite transverse waves in the plate, so that is a 2-D wave medium.

Then for the potential field we were just going to put a rib structure on the plate. On one plate we could put a periodic hexagonal rib structure, and we could drive that plate and look at its normal modes of vibration, and we should see extended Bloch eigenstates.

Then you can take the centers of these hexagons and displace them by random amounts and do a Wigner-site cell construction and you can come up with a plate that has random polygons in the rib structure.

So now we could drive this plate and look at its normal modes of vibration and we should see 2-D Anderson localized states. Actually, the different eigenfunctions are localized at different positions, so one idea was to put noise into this plate and the Anderson localization would filter different noise components, or bands, into different locations on the plate where you could use active cancelation or something. So it might be a clever way of doing noise control. That is still a viable idea.

Before we did this experiment, we thought we ought to try to understand how this Anderson localization worked. This was before we found Marshall Luban's paper, so it was a mystery as to why disorder does not just give you disordered amplitude. So before we did this 2-D experiment, we thought we ought to do a 1-D experiment.

(Transparency 8)

We went down to the local music store and bought a spool of music wire, just steel wire used for guitar strings. We took about 60' of this wire and hung it down the stairwell in the physics building and put a weight at the bottom, so you have a string with a tension in it. Then we put a shaker next to the wire so that we could excite transverse waves in the wire. So that is our 1-D wave medium.

For the potential field we took little lead masses (in fact, these are split shot used for fishing lines; we bought them at a sporting goods store). We just put these lead masses along the wire and those were our scattering sites. In fact, this wire with the masses on it is a very accurate model of a Kronig-Penny potential field that you study about in solid state and quantum mechanics courses.

We could put the masses on the wire periodically and look for extended Bloch eigenstates, or we could put them on there randomly and look for the Anderson localization.

DR. ATCHLEY: The transparency you showed with the periodic structure that had randomized potential wells, in the randomized it looked like the wells were periodically spaced and had different heights. Now you are talking about having things of the same mass but at different spacings, is that right?

A: Yes, in fact, there is a thing called diagonal disorder and off-diagonal disorder. That is getting into the grubby details (I think maybe that is covered here in the paper). But there are two ways you can introduce the disorder.

Running parallel to our wire was an aluminum beam with a little trolley on it and a C-shaped magnet which put a magnetic field locally across the wire. So when the wire vibrated in that magnetic field, there was a current induced in the wire, and the amplitude and phase of the current was proportional to the amplitude and phase of the vibration at the site of the magnet. We could translate the trolley along the wire, so we could make a plot of the amplitude and phase of vibration as a function of position, which is exactly what the eigenfunction is, so in this experiment we could measure eigenfunctions directly.

Before you measure eigenfunctions, you need to measure eigenvalues. To do that, we just put the magnet in one position, and then we would drive it and sweep the frequency. When the drive frequency coincided with an eigenfrequency you get a resonance in the response to the system.

Before I show you the experimental results, I want to show you what you would expect to see in the periodic system, where you can use Bloch's theorem.

(Transparency 9)

There are two approaches. The most convenient one is the tight-binding approach. I have drawn the wire horizontally instead of vertically, so here is the wire. You first imagine that the masses are infinite. If the masses are infinite, it is as though the wire were clamped at the positions of the masses. But then it is easy to understand what the eigenfunctions and eigenvalues are; they just correspond to fitting an integral number of half-wavelengths between the two clamps, just like the modes of a guitar string.

The frequencies that correspond to that are just a harmonic sequence of frequencies. When the masses are infinite, the little sections of wire are just local oscillators with sharp eigenfrequencies. But the masses are finite. What that means, with a finite mass, is that when this section of wire is vibrating it gets coupled to its neighbors, and so when you take these local oscillators and couple them together, these sharp frequencies broaden out into bands. This is a very nice demonstration of how you get band structure in solid state physics.

So in solid state physics there are energy bands and gaps, and in microwave and acoustic transmission lines, there are pass bands and stop bands. Anyhow, you get this band structure.

Our system was finite in size; we had about 50 masses on the wire. With 50 masses, each one of these bands contains 50 discrete eigenvalues. Now you can see what happens when you sweep the frequency. You can start off with the frequency in a gap. As you go into a band you should see 50 resonances and then into the next gap, then into the next band, where there are 50 frequencies, and so on.

(Transparency 10)

Here, now, are the experimental results. This is response as a function of frequency. At these frequencies there is no response, you are in a gap, but as you go into a band, you get 50 resonances and then into the next gap. As you go up higher in frequency you will go into the next band, and so on.

This band, in fact, is the one that corresponds to fitting approximately one-half wavelength between the two masses. This, again, is for the periodic system.

Once you find the eigenfrequencies, you can just sit at any one of these frequencies, and translate that magnet and make a plot of the eigenfunction.

(Transparency 11)

These are plots of the amplitude of the eigenfunction, so this is amplitude as a function of position along the wire. These are two different eigenfunctions that correspond to two different eigenvalues.

If you look in solid state books, you will not find any that look quite like this. Usually in the solid state books they plot eigenfunctions at some symmetry point in the Brillouin zone; however, an arbitrary Bloch eigenstate can look quite different. But it is certainly clear that these eigenfunctions have pretty much extended amplitude.

There are some variations and that might be due to the fact that the masses, commercial fishing split shot, vary by about 13% in their mass. That corresponds to off-diagonal disorder, which is a weaker effect; it does not give you much Anderson localization.

The important thing is the spacing between the masses. We very carefully measured that spacing with calipers to make sure it was the same for each section of wire. So we could get pretty good extended Bloch eigenstates.

This is actually pretty amazing, because if you take the wire and you put one mass on it, you get almost 100% reflection of a wave on the wire. In other words, those masses are fairly large, and it is really a big perturbation on the wire. You get almost 100% reflection. But if you put the masses down periodically, the wave just goes right through the system. So Bloch's theorem is pretty amazing. So that is the periodic system.

What we did for the random system: we took the masses and changed them from their periodic positions by only 2%. There was just a random amount, but within 2%, no more than 2%, so it was a fairly small change. We changed all the masses by just some random amount.

(Transparency 10)

When you do that, there is a dramatic change in the eigenvalue spectrum. The eigenfrequencies kind of clump up. One thing that happens is that you get a frequency that would have been in the gap for the periodic system, so you can say you get states in the gap, which is the sort of thing people in solid state physics say all the time.

There are some theories about the clumping of these eigenvalues and so on. One theory is that if you have a state that is sort of all by itself, that will be highly localized. If you have a state that still has some neighbors around, that will be less localized, and so on.

For the disordered system we could sit at one of the eigenfrequencies, and look at the eigenfunctions.

(Transparency 11)

This is what we got. These eigenfunctions here are for five different eigenvalues for the disordered system. This one right here, this is just textbook Anderson localization. Actually, this is the first experiment where an Anderson localized state was observed directly.

These other states are localized to more or less different degrees. The fact that the localization length varies for different eigenstates is an example of the universal conductance fluctuations that people are talking about in the mesoscopic electronic systems. Also, you will notice that the localization occurs in different sites.

This was nice, but we had no intention of publishing this, because we did this experiment just to educate ourselves as to how this Anderson localization works. We found out that it is actually quite easy to do digital computer simulations that will give you the same sort of eigenfunctions.

But when we got this result, we realized we could do something that was not so trivial. What we thought of doing is the following. We said: suppose this eigenfunction right here is at 800 Hz, but if you drive it to 810 Hz you get the Anderson localization over here. These are just two different localized eigenstates at different eigenfrequencies.

We said: well, suppose we drive it at 800 Hz but then modulate the tension in the wire at 10 Hz, the difference frequency. When you modulate the tension in the wire, the little masses are moving back and forth, and that looks just like a longitudinal phonon going through the system.

The transverse waves in the wire are analogous to electron Schrodinger waves. Modulating the tension is like sending a longitudinal phonon through the system. So we could simulate electron-phonon scattering and the effect of that electron-phonon scattering on Anderson localization. This was a serious problem in Anderson localization at the time: How does inelastic phonon scattering affect Anderson localization?

In the experiment, you start off with the amplitude localized at one site. You turn on the modulation and what you see is the amplitude going back and forth between the two different sites. What happens is, the phonon couples the two localized states and they just beat against one another at the difference frequency.

In the solid state system, at some point in time the phonon goes running off to the thermal bath, and depending on

when it does that there is some probability that the energy gets left in the second site, so this is phonon-assisted hopping between two Anderson localization sites. This was a hot topic in physics at the time, so we were able to publish this result.

In this experiment we could measure the individual eigenfunctions, we could measure the results of the modulation. This plot down here (h) is a time average of the amplitude going back and forth, which shows you the relative amplitudes of the two eigenfunctions. That relative amplitude gives you the probability of hopping from one localized site to the other. In this analog experiment we could measure this hopping probability. You could do it as a function of the strength of modulation or the strength of the phonon field, and you could measure everything in just gory detail, and that was what we were able to publish.

The original idea was to study 2-D Anderson localization, but we got derailed from doing this experiment, because quasicrystals were discovered.

(Transparency 12)

Not too long ago it was thought that solids could exist in two basic forms, crystalline and amorphous. In an amorphous solid the atoms are more or less randomly placed and it has isotropic and homogeneous properties, but a crystal is quite different. To make a crystal, you take a unit cell and repeat it periodically to fill all space. In order to fill all space without leaving any gaps, only certain shapes are allowed, and there are 14 different shapes and they form the basis for the 14 Bravais lattices in classical crystallography.

Because only a finite number of shapes are allowed, there are only certain rotational symmetries that are allowed. In particular, fivefold rotational symmetry is not allowed. In fact, if you look in Kittel's solid state book (this is in an earlier edition, he has fixed it in later editions), it says, "A fivefold axis of symmetry cannot exist in a lattice because it is not possible to fill all space with a connected array of pentagons."

Some time ago, Schechtman, at what was then the National Bureau of Standards, made an aluminum alloy and he put it in his x-ray diffraction machine and found 10 spots in a circle, which means he had a mirror plane and fivefold rotational symmetry. He showed this to people and, of course, they just laughed at him: "You cannot have fivefold rotational symmetry."

He and other people repeated the measurement and as they got better and better measurements, better and better aluminum quasicrystals, these spots got sharper and sharper. A sharp

spot in a diffraction pattern means you have long-range order. So he had long-range order and fivefold rotational symmetry, which people thought was impossible in a crystal.

It turns out that mathematicians and physicists knew what was happening here. It is impossible to have fivefold rotational symmetry and long-range periodic order, but you can have fivefold rotational symmetry and another type of long-range order, which is called quasiperiodic.

Actually, mathematicians get mad when you call it quasiperiodic. It is actually quasicrystalline or Penrose tile symmetry but, anyhow, I will just call it quasiperiodic.

(Transparency 13)

Here are some electron micrographs of actual aluminum alloy quasicrystals. You can see there are five facets here. It is macroscopic; this is on the order of 100 μm. You can see that it has macroscopic fivefold symmetry.

(Transparency 14)

Here is another example. This is aluminum alloy that Alcoa made. Apparently it is something they make especially for the Chrysler Corporation.

(Transparency 15)

In 1-D it is not too difficult to understand quasiperiodicity. Suppose I take a line and put points on it periodically with a lattice constant a. If I do a diffraction experiment or just take the Fourier transform, then I get a sharp line at $\frac{\pi}{a}$, and it is a sharp line indicating that I have long-range order, in this case periodic order.

On the other hand, if I take a line and put dots on it randomly and take a Fourier transform, then I get a broad spectrum, indicating that I do not have long-range order. There is no length that characterizes the system.

On the other hand, I could take another line and put dots on it with a lattice constant b (you put these on periodically) and then superimpose these two lines. I could get a line with dots on it that might not look all that different from the random array of dots. But if these two lattice constants, a and b, are commensurate, that is, if their ratio is equal to a rational number (equal to a ratio of two integers) then this sequence of dots will have some pattern that will repeat periodically.

If this ratio is equal to the ratio of two large numbers, then this thing may have a large unit cell, but at least it is guaranteed to repeat periodically and you can use Bloch's theorem and other powerful theorems for periodic systems.

On the other hand, if these two lattice constants are incommensurate, that is, if the ratio is equal to an irrational number, then the points on this line will have no pattern that will repeat periodically. But on the other hand, if you take their Fourier transform, you just get two sharp lines, because it is just a linear superposition of two periodic structures.

The two sharp lines tell you that there is some long-range order, but because these spatial frequencies are incommensurate, it is not periodic, so it has quasiperiodic long-range order.

Q: Do you have a third sharp line that corresponds to the really long wavelength that is the repeat period?

A: This one does not have a repeat period, if they are incommensurate. If they are commensurate, then you will just get two commensurate spatial frequencies. It is just a Fourier transform of a superposition of two linear systems.

This is actually kind of an oversimplification of the quasicrystalline system. In higher dimensions (this is in 1-D) it is a much more sophisticated concept. To explain quasicrystals in higher dimensions, I am going to show you a second way of getting a 1-D quasicrystalline sequence of dots.

In this method you start off with a periodic lattice in a higher dimension. In this case we will start off with a periodic square lattice in 2-D. Then you take that higher dimensional lattice and you intersect it -- question?

Q: I am having trouble understanding how you can take two periodic lines of dots and try and get something that is commensurate.

A: You get something that is not periodic when you superimpose, because the spatial frequencies are irrational. It is an irrational ratio. There is nothing that repeats.

Q: There is no least common denominator is another way of looking at it.

A: Yes.

So we start off with a 2-D square lattice -- that is the higher dimension. We intercept it with a lower dimensional surface. Then you define a window and you project points inside the window onto the lower dimensional surface.

In this case here you intersect the 2-D periodic array of dots with a line at some angle, you define a window, and then project points down onto that line. For that lower dimensional surface, if the directional cosines of that surface are irrational, then the dots on that lower dimensional surface will form a quasicrystalline pattern.

The dots will not form a periodic pattern because of the irrational direction cosines. In this case you can look at the tangent of this angle. That tangent is an irrational number. Then the points on this line will not be periodic. But because they are projected from a periodic lattice, then they will have long-range order, so that gives you another example of a quasiperiodic sequence in 1-D.

For higher dimensions: to get a 3-D quasicrystal you start off with a 6-D periodic lattice and you intersect that with a 3-D surface. I am sure you will have not trouble at all imagining that.

(Laughter)

Anyhow, you intercept it with a 3-D surface, have a window and project points onto that surface and get a 3-D quasicrystal. For a 2-D quasicrystal you start off with a periodic lattice in 5-D and intersect it with a plane, project the points, and you will get a quasicrystal that is more popularly known as a Penrose tile.

(Transparency 16)

These 2-D Penrose tiles have been fairly popular in the literature. There were a number of articles in *Scientific American*, in Martin Garner's section, and there was a cover article on these in *American Scientist* some time ago, and that is a good article, if you want to look for that.

These Penrose tile quasicrystals have some very interesting properties. To make an ordinary crystal, you take one unit cell and repeat it periodically to fill all space. To make a quasicrystal, you are allowed to use more than one unit cell and, in fact, in this sort of standard Penrose tile there are two unit cells. There is a fat rhombus and a skinny rhombus. The ratio of the areas of those two rhombuses is the Golden Mean, $\sqrt{\frac{5+1}{2}}$. That is an important number.

(Transparency 15)

Let me back up here a little bit, back in this 1-D example. If the tangent of this angle, we want that to be an irrational number, is equal to the Golden Mean, $\sqrt{\frac{5+1}{2}}$, then you

get special properties. This Golden Mean is also known as the Divine Ratio. It is also known as the most irrational number. I think the reason is if you do a continued fraction to get the Golden Mean, you get a continued fraction that converges more slowly than any other one.

If the tangent of this angle is equal to the Golden Mean, then it turns out that the points on this line can correspond to a Fibanacci sequence of numbers. The Fibanacci numbers have all kinds of miraculous properties and, as a consequence of that, these 1-D Fibanacci sequences of points on a line also have incredible properties.

(Transparency 16)

This Golden Mean also shows up in this 2-D Penrose tile as the ratio of the areas of those two unit cells. In order to put these two-unit cells down to fill all space, you have to obey some very strict rules. If you do not follow those rules, then somewhere you will get a gap where you cannot fit a tile in.

Because of those rules, when you tile the plane with these tiles, you will have long-range order, because you have to follow those rules. But because the unit cells involve this irrational number, there will be no pattern in here which repeats periodically, so it is has long-range order but it is not periodic order.

These Penrose tiles have some other amazing properties. If you take these tiles and draw some lines on them in some way, then put the tiles down according to the rules, then the lines that you drew on the tiles will also form a Penrose tile pattern, but it will be scaled down, and this is called inflation or deflation symmetry.

There is another theorem, and this theorem says that if you take some pattern with some nominal diameter, then within two diameters and usually within one diameter, you will find that pattern again. This thing sort of has global fivefold symmetry, so it is not too hard to see that here, but you can go off this symmetry point and that theorem is still valid.

Take a pattern of some diameter, within two diameters you will find that pattern again, and that is pretty amazing. I do not know how they discovered that or proved that theorem, but it is pretty amazing.

The thing is that even though this pattern has to occur according to this theorem, it cannot recur periodically. This theorem is known as Conway's theorem and I think it is going to be very important in the properties of this Penrose tile.

(Transparency 17)

There are three fundamental types of solids. There are crystalline solids, or periodic solids, random or amorphous solids, and now there are quasicrystalline solids. So now there is a fundamental question. Suppose you have a wave equation with a potential field that is quasicrystalline. How does that quasicrystalline symmetry show up in the eigenvalues and the eigenfunctions? It is a very basic question about this new fundamental symmetry.

(Transparency 18)

In 1-D, thanks to the Fibanacci numbers, you can actually prove rigorous theorems. For 1-D quasiperiodic sequences you can show that the eigenvalues form a Cantor set (this is related to "let M be a manifold"). The eigenfunctions can be either extended, having more or less constant amplitude, they can be localized, or they can be so-called critical eigenfunctions that just fluctuate a lot. That is in 1-D; people have proved rigorous theorems.

(Transparency 17)

But in two and 3-Ds no one has been able to prove any rigorous theorems. There is no quasi-Bloch's theorem for 2 and 3-D quasicrystals. The problem is they are not periodic, so you cannot use Bloch's theorem; they are not random, so you cannot use statistical techniques, either. They are right in-between and there are no theorems.

In the absence of a theorem, we thought we would do an experiment and measure the properties to find out how a quasicrystal affects (or how that symmetry affects) eigenvalues and eigenfunctions in a wave equation.

Q: How would you make a Fibanacci 1-D crystal with piano wire and split shot?

A: The spacing between the dots, there are only two spaces, and they are related by the Golden Mean and they occur in a sequence that corresponds to Fibanacci numbers. The way you do that is given by a formula in the handout. It is kind of an interesting looking formula in the handout here.

If you look at that article in *American Scientist*, it gives you the explanation of how that occurs. The way the Fibanacci numbers came up, you start off with two rabbits, a big rabbit and a small rabbit, and they breed and they have certain offspring, and then they breed, and so on, and that process generates the Fibanacci numbers, or generates a ratio of big and small rabbits.

So now we want to measure the effect of a 2-D quasicrystal symmetry on eigenvalues and eigenfunctions of a wave equation.

What we want to do is to use the tight-binding approach again. You want to have a local oscillator at the centers of these unit cells, so the local oscillator has some sharp frequency, but then you want to couple these together. That sharp frequency broadens out into a whole spectrum of frequencies and then you want to look for structure in that spectrum of frequencies.

Another thing you want is not very much damping, because all this stuff is based on long-range phase coherence, interference of the waves as they go around this quasiperiodic pattern, so you do not want much damping.

If you want a local oscillator that is a high-Q oscillator, it is hard to beat the old tuning fork. I can get this tuning fork going here, and it will ring for a couple of minutes. It has a Q of about 10,000 or so. Not only that, I can hold it by the stem, so this is a lossy support, but that does not affect the Q very much, because the two tines will work against one another and they do not lose much energy out the support.

Also, you cannot hear it very well, and that is good, because it is not losing energy to radiating sound out into the room. So after buying a couple hundred feet of guitar wire, we went back to the music store and bought 300 tuning forks.

(Transparency 19)

We took an aluminum plate that was about 1-1/2 m on a side, about 3/4 in thick, and drilled holes in the plate in a Penrose tile pattern, and then epoxyed the tuning forks into the plate. So there is our local oscillator, each one at 440 Hz, by golly, and then to couple them together we took arcs of steel wire and spot-welded the steel wire between the tines of the tuning forks.

We tried other coupling schemes but you still need a fairly high-Q coupling mechanism, so the arc of steel wire that was spot-welded turned out to work very well.

For the coupling scheme, I will show you the details here. Here is one of the rhombuses. The tuning fork was in there so that the two tines were across the short diagonal. Then the four sides of the rhombus identified the four nearest neighbors. The local oscillator is actually kind of complicated. It is like a diatomic molecule here. But the fact that it has Penrose tile symmetry is inescapable. There is our balls-and-springs Penrose tile.

Then, to measure the eigenfrequencies, we just put an electromagnet next to the side and put AC current through it to drive one of the tines. Then, to measure the response of the system, we went back to the music store and bought half a dozen

electric guitar pick-ups. We put the electric guitar pick-up next to the tine of the tuning fork, so when you vibrated it, you get a signal in the electric guitar pick-up. We put them at 20 random places throughout the quasicrystal, and swept the frequency and looked for the resonances.

(Transparency 20)

Here now are the resonant frequencies. This is just what the frequencies are. This is the spectrum of eigenfrequencies for that quasicrystal. Here is 440 Hz; it actually showed up in a gap. The span here is on the order of 100 Hz. We actually used 150 tuning forks, so there were 300 tines in the array, so in this spectrum there are 300 eigenfrequencies within 100 Hz. In order to identify the individual lines, you had to have a very high-Q resonance. Using the tuning forks was a nontrivial part of this project. In fact, Paul Chakin at Princeton University had a postdoc who was working on an experiment like this for a year. He took a rubber membrane as a 2-D wave system and was putting epoxy dots on it in a Penrose tile pattern, but that experiment never worked, and it is probably because there is too much damping in the rubber membrane. So using the tuning forks was a serious matter.

We could identify all the lines and we found that there was a definite structure in the spectrum. There are these gaps and these gaps were very dramatic. It was not just that we missed a line or something. We would be sweeping along in frequency and we would get a resonance and a resonance and a resonance and then nothing. You turn up the drive, you turn up the gain of all the amplifiers, there is just nothing happening in those frequencies, so these were real gaps in the spectrum.

These bands and gaps have widths, and if you take the ratio of these widths, you get the Golden Mean. That was something that was not predicted by any theory. It happens in 1-D; that Cantor set has a scale that is the Golden Mean, but this is something that was not predicted by theory for 2-D.

Also, people had done computer simulations of 2-D Penrose tiles and they did not see this effect. There is a lesson there about computer simulations. They made a simplifying assumption that, on the face, looked okay, but it turns out that the real world had a better idea. Essentially, in the computer simulations they wound up throwing away the entire wave nature of the problem. In solid state theory, instead of solving the Schrodinger wave equation, they often study what is called a hopping Hamiltonian, and often you lose the actual wave properties of the original problem. The equation you should be solving is the Schrodinger wave equation, which is analogous to our acoustic wave equation here.

Q: Can you identify any finite side effects because of the finite crystal?

(Transparency 16)

A: I can sort of explain why you see this. In this Penrose tile pattern there are some patterns that recur. In fact, Conway's theorem guarantees that patterns recur. In the computer simulations they essentially assumed that they had things coupled with massless springs, and the speed of sound on a massless spring is infinite, so in their system their wavelength is always basically infinite.

In our system, these coupling wires have a finite mass; they have a stiffness and also a finite mass and a finite speed of sound in the coupling wires. As you vary the frequency, you are varying the wavelength in the system, so that varying wavelength is probing the length scales in the system. You can go to a frequency where you can fit an integral number of wavelengths around these patterns and it is going to like that; it sort of locks up there and it gives you a gap in the spectrum.

Another way you can think of it is as this wavelength is varying, it is probing all the length scales in this pattern, and the Golden Mean is all over the place in this thing; the wavelength is proportional to the frequency, so you ought to see the Golden Mean in the frequency. All that sort of wave nature was missing from the computer simulation.

Your question was the finite size. You can see that this is going to work even if this thing has a finite size, because those patterns still recur and you are still going to lock up on those frequencies.

Q: Would it matter a lot whether you chose, say, the points of the intersection to put the forks in?

A: If you put them at the intersection, then the distance to the nearest neighbor is the same in each case, but the number of nearest neighbors is different. If the number of nearest neighbors varies, I do not know how to hook up the coupling wires for the tuning forks.

The way I did it, this way, is an obvious way of hooking up the coupling wires. In fact, if you work it out, the lengths of coupling wires, I think there are only four different lengths of coupling wires that you need to hook up all possible combinations. Of course, those lengths will have the Golden Mean in them.

Q: Were the coupling wires a loop or were they tight strings?

A: They were an arc. So you have a couple of tines like this and then it was an arc, so the arc just bends like this.

We got the eigenvalues. Now you have to get the eigenfunction. To get the eigenfunction, you need to measure the amplitude of vibration of all the tines of the tuning forks. You cannot really see the tuning fork move -- you can kind of feel it. So to measure the motion of the tuning forks we took the whole tuning fork array and put it up on one wall of the lab and then, on the end of a tine, we put a little mirror.

Then you shine a laser on it and the laser gets reflected to the opposite wall in the lab, so as the tine oscillates, the laser beam gets deflected, and even though this is a small angle, you have this optical lever arm, so even though you cannot see the motion here, on the wall you could see the streak of the laser beam that shows you how the tine is vibrating.

(Transparency 21)

That is shown here. We put a little mirror on the end of each tine and then, when the laser hits it, you get a streak over here which shows the motion of that tine. The laser went to a scanning system so it would scan over the whole tuning fork array and put streaks over here. Then, for the measurement, we just turned the lights off in the room, opened up a camera to look at this, and then scanned the array.

In order for this to work, each mirror has to be adjusted so that the streak appears at the corresponding position on the opposite wall. The poor graduate student was up there for months aligning 300 little mirrors to get the eigenfunctions, but he succeeded.

(Transparency 22)

Here are the eigenfunctions. The original photographs were a lot better, but you can see this sort of uniform distribution here. This one has some highly localized states, and so on. The problem with eigenfunctions is that the little spot welds for the arcs of wire are not perfect. There are going to be variations throughout the thing. If you have little perturbations, there is an adiabatic theorem that says those little perturbations will not affect the eigenvalues very much, so we have got this nice spectrum with the Golden Mean in it. But the perturbations do affect the eigenfunctions strongly, so that is something you should remember. So eigenfunctions were not great. They were fairly good but they did suffer from the perturbations.

Q: The figure you just showed, Figure 3 of this paper that is in the notes, what are those white lines that go through there?

A: It is an aid to the eye in seeing the symmetry. There is some structure in the streaks here that you cannot really see in this reproduction. In the original photograph you can see it a lot better. This is sort of uniform streak amplitude. This one has higher streak amplitude in the middle. This one is sort of a fivefold pattern of streak amplitude, and so on. Some of them we could not understand, like that one. Then there were some that looked like they were quite localized. So that is what the white lines are for.

(Transparency 23)

Anyway, this experiment was a lot of fun and people liked it. In fact, it showed up in that reputable science journal, the Science Section of *The New York Times*, which was kind of thrilling.

I think this is a good point to stop. Next time I will talk about disordered systems where you add nonlinearity.

DR. MAYNARD: What I will start with now is disordered systems, where you have Anderson localization, but now add nonlinearity.

(Transparency 24)

In condensed matter physics there has been a lot of research with waves in disordered fields. There is Anderson localization, universal conductance fluctuations, the normal electron persistent currents, thousands of papers published, books have been published, there has been one Nobel Prize and two Buckley Prizes for properties of waves in disordered fields, so there has been a lot of activity for disordered systems.

At the same time, but independently, there has been a lot of research on nonlinear physics. There have been studies of solitons in chaos, thousands of papers published, lots of books, popular books, even jokes about chaos.

(Transparency 25)

(Laughter) Here is one that I found. It says, "What led you to the mathematics of chaos, Dr. Maynard?" They got my name but I do not know where they got such a good likeness? I do not know who did this cartoon. Somebody sent me this in the mail.

(Transparency 24)

With so many interesting things going on in these two fields separately, let us take a look at all the activity that is going on in the overlap of the two fields. Actually, there is relatively little activity in systems that have combined disorder and nonlinearity. It is not because there is nothing interesting happening there. The reason there is so little activity is this is a very new field in physics. The first meeting on disorder and nonlinearity was in 1989 and since then there has been a second meeting, so those were not all that long ago. All the papers at those meetings were theory papers. I think we are the only people who are doing a lot of research on disorder and nonlinearity. The only other research I know about is stuff that Bruce Denardo and Seth did and that was just a small amount, right? So there is a big opportunity here to study systems that are both disordered and nonlinear.

Theoretically, it is a difficult field. A lot of the people who had papers at those meetings were hard-core mathematicians.

(Transparency 26)

In the field of disorder and nonlinearity, a fundamental question is: Does nonlinearity weaken Anderson localization? Remember, if you have disorder, you are going to have Anderson localized eigenfunctions. But if you turn on nonlinearity, does that destroy the Anderson localization?

If you survey the literature, there are about eight theory groups who have addressed this problem. About half of them predict that, yes, the nonlinearity weakens localization, and about half predict that, no, it does not weaken the localization. Some groups hedged their bets by voting both ways. (Laughter)

So it looks like a lot of fudging and perhaps a big controversy here. But it turns out there is not a controversy. It is nonlinear physics, and one thing you could say is that some nonlinear problems do not have unique answers. Actually, in this case, it turns out that the answer that you get depends on how you ask the question precisely, or how you do the experiment.

(Transparency 27)

For linear wave systems, you can think of two ways of doing an analysis, a normal mode analysis or a pulse analysis. Or, if you like, this is the Orson Anderson way of doing it, and this is the Moises Levy way of doing it.

For normal mode analysis, you have a wave system and you drive it at one end, $\cos(\omega t)$, and at the other end you look at

transmission, and then you vary the frequency and look at transmission as a function of frequency. If the frequency corresponds to one of the natural frequencies of the system, you excite a normal mode, and you will get a peak in the transmission. In the normal mode analysis you look at a transmission spectrum.

In the pulse analysis you send in a real sharp pulse. That pulse rattles around in the system and at some point at the exit you get a bunch of pulses coming out. In this case, what you measure is the displacement as a function of time at some point.

In a linear system, these two things are just Fourier transforms of one another because these two things are Fourier transforms of one another, and it is a linear system. But if it is nonlinear, these two ways of analyzing the system are no longer equivalent.

It turns out that if you do a normal mode analysis, then the theory predicts that the nonlinearity will not destroy the Anderson localization. If you have disorder and nonlinearity, you will still get localized normal modes.

We wanted to show that there was a part of the theory where there was some question about what would happen and we wanted to study that. It turned out that the string with the masses on it was a good system for studying that, because the nonlinearity of the string was kind of interesting.

(Transparency 28)

If you have a string without the masses now, you take the string and you stretch it some tension, T_0, and the wave speed is the $v = \sqrt{\frac{T_0}{\frac{m}{L}}}$. But if you give the string a finite transverse displacement, then the arc length of the string is changed and so the tension will have changed.

If the tension changes, then the speed of sound changes (that is the arc length right there) since the tension depends on the displacement, you get a nonlinear equation. But the interesting thing is that if I give it a finite displacement here, so that the tension changes here, that change in tension will propagate along the wire with the speed of longitudinal waves in the wire, and that speed is much higher than the speed of the transverse waves in the wire. What happens, essentially, is that if you give it a finite displacement, the tension changes throughout the entire length of wire. It gets pretty much distributed along the entire wire uniformly. That is an approximation but it is a reasonable approximation.

Here is what we thought would happen for tbe disordered system. You have a disordered system, so the linear eigenfunctions might look that and another one might look like that, but if you drive it at finite amplitude, then this finite amplitude here would modulate the tension in the wire.

If you take a string and modulate the tension, you can excite the transverse eigenmodes of the string, i.e. you can parametrically excite it. What we thought might happen is that this Anderson localized state at finite amplitude would modulate the tension in the wire and that might parametrically excite a different Anderson localized state, so it would have nonlinear-assisted hopping between Anderson localization sites. When that happens, if you can excite more of this, you will get more transmission. So as we drive our system at finite amplitudes, we should see the transmission increase if we get this nonlinear-assisted hopping. That is actually a point that the theory paper could not really address.

(Transparency 29)

Here are the results of the experiment. The experiment was the same as for linear Anderson localization. We just literally turned up the volume to see what would happen. Here is the transmission spectrum. This is transmission as a function of frequency. These are different drive amplitudes. These curves are normalized by the drive amplitude.

If the system were strictly linear, all of these plots would be the same, but they are not. It is a nonlinear system. If you look at some of these peaks here, this peak is increasing, this peak down here is getting bigger and bigger. But then they eventually drop off and, in fact, if you integrate under these curves to get a total transmission as a function of amplitude and normalize it, say, to this first one, then you find that the total transmission actually decreases, so apparently we are not seeing this nonlinear-assisted hopping that would increase the transmission. The transmission is actually decreasing.

So instead of one localized state parametrically exciting another state, the localized state is actually parametrically exciting itself or, as Bob Keolian would say, it digs a deeper hole for itself, so the states actually get more localized, if anything. So that was consistent with the theory.

(Transparency 30)

Here is one of these normal modes. You can see that as you go to higher drives some of the sections of wire actually get a larger amplitude, but the Anderson localization does not really change that much.

(Transparency 27)

For normal modes the nonlinearity does not destroy the Anderson localization, so that was good and we got that published. But remember, there are two ways of analyzing the thing, the normal mode analysis and the pulse analysis. So what happens if you have a disordered system and a finite amplitude pulse?

(Transparency 31)

If you want to find out what a pulse does in a system, some people actually try to keep track of the pulse and its reflections and transmissions and so on. In fact, doing this problem that way has actually led to some incorrect results.

If you want to find out what a pulse does in a linear system, you should actually use the normal mode analysis. You should find out what the transmission as a function of frequency is and then do a Fourier transform. In fact, we did some computer simulations doing that and showed that a lot of the results in the literature are wrong.

If you want to calculate this theoretically, what you do is you start at one end of your system, where you have some boundary condition, and then as you go past each scatterer, you will find that you will be multiplying by a 2 by 2 matrix. The reason it is 2 by 2 is because the wave equation is second order, you have two linearly independent solutions, so to go from one section to the next, you multiply by a 2 by 2 matrix.

You multiply by 2 by 2 matrices as you go past each scatterer until you get to the other end, where you have to satisfy some condition there. Satisfying both conditions gives you some equation for the eigenfrequencies.

Q: Is this just equivalent to looking at the reflection and transmission off of each of these?

A: Yes, exactly.

The point is that to study a pulse, you have to multiply by 2 by 2 matrices from one end of the system to the other end of the system, over the whole system. That is exactly what Furstenberg's paper was about, the product of random matrices.

If the system is large enough and you make this product of random matrices over a large enough system, you are going to get Anderson localization. That is Furstenberg's rigorous theorem. A pulse in this linear, disordered system, will exponentially decay as it goes through the system.

The crucial part is that you have to satisfy conditions over the entire disordered system and so you have to get Anderson localization.

A nonlinear pulse has an extra degree of freedom. In fact, for a soliton, the way you can mathematically treat solitons is with inverse scattering theory. The extra degree of freedom for the soliton is the number of eigenvalues that go into making that soliton.

So a nonlinear pulse has an extra degree of freedom. In satisfying conditions, you can adjust that extra degree of freedom to satisfy conditions only locally. You do not have to satisfy conditions over the whole system for a nonlinear pulse. Or in other words, when you have a nonlinear pulse, there is a second length, which is a nonlinear length. It is essentially the width of the nonlinear pulse.

Now, when you have this nonlinear pulse in disorder, there are two lengths: this nonlinear length and the Anderson localization length. This is what the theory predicts. This is transmission on a log scale as a function of distance that the pulse has traveled. If you have a strong soliton or a strong nonlinear pulse, so that the nonlinear length is less than the Anderson localization length, then the nonlinear pulse does not sample enough of the disorder for Furstenberg's theorem to be valid and so it does not know about Anderson localization, so a strong soliton will just propagate right through the system.

On the other hand, if you have a weak soliton, where the nonlinear length is much greater than the Anderson localization length, then you do sample enough of the disorder to learn about Anderson localization, and so a weak soliton (essentially just a linear pulse, I guess) will exponentially decay as it goes through the system.

The interesting case is when the nonlinear length is the same order as the Anderson localization length. What the theorists predict (guess is what I think they did) is that the pulse will go along, decaying some, and then it will start to exponentially decay.

We wanted to do an experiment to look for that. We wanted a good nonlinear pulse and, as Steve pointed out, a good nonlinear wave can be formed with a surface wave on a fluid.

(Transparency 32)

The speed of a surface wave involves the depth of the fluid, but if you have a finite amplitude wave, that modifies the depth, and so if you put this in the wave equation you will get a nonlinear wave equation. Steve already did that. And you

can get strong nonlinearity with large-amplitude surface waves.

We also want low attenuation, we want low damping, because when we have a linear wave we have to get that long-range phase coherence so we can see the Anderson localization in the linear system. If you want low-attenuation surface waves, you can use superfluid helium and, as Steve pointed out, the surface waves are called third sound. The restoring force, instead of being gravity, is van der Waals force; it goes like this.

The experiment uses a superfluid helium film on a substrate. The substrate we use is just a microscope slide. At one end there is a third-sound transmitter, which is simply just a heater. The third sound has a temperature component, so you can drive it with a heater and you can detect it with a thermometer.

For the thermometer we use small narrow aluminum strips. The experiment is at a temperature where the aluminum is just going through its superconducting transition. So if I make a plot of resistance versus temperature for an aluminum strip: there is an abrupt drop when it goes superconducting. The temperature where we are operating is right here, so if we have any small excursions in temperature we get large excursions in the resistance of that strip, so that is a sensitive (thermometer) receiver for third-sound.

For the scattering field we took a diamond wire saw and put cuts across the microscope slide. First we put cuts across periodically, because we always want to check and make sure we can understand the periodic system, the Bloch eigenstates; then we made a slide with a random array of scatterers.

Before we do any scatterers at all we want to make sure we have a nonlinear pulse to deal with. By the way, I should mention that this research was started by Doug Meegan, who had great success. He said he did not get results in his talk, but in this experiment he got lots of great results.

(Transparency 33)

First, we want to take a look at what third sound does, finite amplitude third sound, when there are no scatterers. In this system, pulses are launched here, there is a time of flight down the substrate, and then there is the received pulse at the superconducting strip. As you increase the amplitude, the thing just scales up, so this is still in a linear regime.

(Transparency 34)

As you increase the amplitude further, you find that this pulse right here saturates. Apparently what happens is the

superfluid moving back and forth eventually is moving back so fast that it reaches the critical velocity, and there it does not go any faster. So that saturates. This, then, does not change as you increase the power.

What you find is that a little blip appears on the side of that. As you increase the power further, that little blip moves along and in order to see that, what we have done here is taken one drive level and plotted it, and then a slightly higher drive level and plotted that. The difference between those two is this little blip that moves out as you increase the power level. You can make that even more visible by actually subtracting the two traces.

(Transparency 35)

That is shown here. This, now, is the time of flight of this little nonlinear thing, whatever it is, and you can see that as you increase the drive amplitude, the time of flight of this nonlinear pulse changes, so this thing has a velocity of propagation that depends on the drive amplitude, so it is certainly some kind of nonlinear thing. Unfortunately, no one has really worked out a theory for what this thing is.

One thing to notice here, these are normalized -- I forget how -- if you normalize them with the drive power, you see that as you increase the drive, this normalized amplitude does not increase. It changes its velocity but the amplitude does not change.

Q: There is a small bump, just like in the first one, but the other one is just --

A: There are a couple of little bumps. This bump right here is sound propagation in the vapor above the superfluid film. When I subtract these two, sometimes the subtraction is not perfect and you will get a little extra blip there, just some little noise or something in the data.

(Transparency 36)

Now we put the scatterers in. The first thing we did is we put the scatterers in periodically. For the linear system you get a pulse like that. With periodic scatterers you actually get a coherent pulse out, because they are periodic. There is lots of scattering back and forth, but at the far end of the thing you actually get a coherent pulse out, and then you get something that looks like echos but, remember, there are about 40 scatterers, but they are periodic, so you get coherent echos out here.

If you take the Fourier transform of this, you get bands and gaps. There is a gap here, there is a gap here, and so on. Doug actually got some much better data but with this run

he went to finite amplitudes, so I can show that. Now you increase the amplitude and, as you increase the amplitude, you find that these sort of coherent echoes go away. We'll, the band structure goes away.

What is happening here, as you increase the amplitude, that nonlinear length is getting shorter and shorter and it is sampling less and less of the periodic array of scatterers. Pretty soon it is maybe sampling only a couple of scatterers, and so that is not really enough to get a decent band structure, so the disappearance of this band structure and these coherent reflections indicates that we are getting a smaller and smaller nonlinear length, it is consistent with that.

(Transparency 37)

Now we go to the disordered system of scatterers. This looks a lot messier, because, remember, now the pulse is not going through nice coherent periodic scatterers, it has gone through a disordered array of scatterers, but we still see that little extra nonlinear thing that appears at finite amplitudes.

If you subtract the two curves to bring out the nonlinear part, if you look at the pulse in the nonlinear periodic system, you get that the pulse shifts, with increasing amplitude, indicating it is a nonlinear something -- but, again, its amplitude does not really change that much as you go through the periodic system.

On the other hand, if you look at the disordered system, the thing shifts over, its velocity changes, but, also, it gets larger in amplitude.

(Transparency 38)

I can show this for different distances. The nonlinear pulses are normalized by what the linear pulse is doing. In the disordered system, unfortunately, the linear pulses are undergoing conductance fluctuations, so they do not really decay exponentially; they fluctuate. But on the whole they decay exponentially.

This now shows you the ratio of what the nonlinear pulse is doing normalized to what the linear pulse is doing. You can see that as you go farther and farther, the nonlinear pulse is relatively larger. What is happening is: because of Furstenberg's theorem, the linear pulse has to be exponentially decaying.

If I turn this this way, so that that is exponentially decaying, then you see that the nonlinear pulse is not decaying like the linear pulse was.

This dashed line, by the way, is on the periodic system. There is a little scatter there, but this scatter we think is just due to the fact that there are conductance fluctuations in the disordered system.

(Transparency 31)

If I take this turned sideways and actually plot it on the original theory graph -- here -- then our data points; there is still scatter in the points, but the points pretty much follow the line that was predicted for that intermediate behavior.

I think we have actually seen this effect. In fact, if you do a calculation of what the Anderson localization length is for the linear system, the pulse width is about the same as the Anderson localization length, so I think we are really seeing this intermediate behavior. What we are doing now is we are measuring some more slides with disorder and we will average them and try to take out the conductance fluctuations.

So that is the nonlinear pulse in the disordered system and it looks like it is agreeing with the theory. As I say, we are the only people, I think, who are doing experiments on disordered nonlinearity.

I told you about the experiment where we used the wire with the masses on it to study nonlinearity and disorder. We also did an experiment with the periodic sequence of masses and drove it nonlinear to see what would happen. We were kind of looking for something interesting that is based on what I am sure Bob Keolian told you about, what he calls the bent tuning curve.

(Transparency 39)

If you have a mass on a spring and you look at response as a function of frequency, you get a resonance curve; but if the spring is nonlinear (for example, if the spring stiffens at larger amplitudes) then you get, as you go to high amplitudes, a bent tuning curve. So you go up in frequency, it goes over the tuning curve like this, and then drops off. If you sweep down in frequency, it goes along and it jumps up to here and comes back. So that is if you have one resonance curve.

(Transparency 40)

Remember, for the string with the masses when they were periodic, you get a band of resonances, and these can be fairly closely spaced. We did a wire, in this case it had 24 masses on it, so we have 24 peaks and one band, so now the question is, what if you drive this nonlinear; are these individual curves going to curve over and drop down; and what happens if

one of these things curves beyond several of the others; what can happen?

We just turn up the amplitude in the experiment and you can see that there is a little noise here (we were able to eliminate this later, but not in this viewgraph). The curve goes over and drops off, just like the individual resonant curve did. You keep doing that. It gets a little more complicated as some of these curves bend over their neighbors.

But then, right at the beginning, we got this really noisy looking part, where you change the frequency and it just goes to a completely different amplitude. We thought maybe this was just noise or temperature variations in the room, and we improved the experiment, but this was still there.

We took finer and finer frequency sweeps. This was at about 1 kHz and we were sweeping with, I think, 0.05 or maybe 5 mHz -- we were sweeping in 5 mHz steps, and this busy stuff just got busier. It sort of looks like chaos, but it is not chaos, because the wire is still oscillating as $\cos(\omega T)$. So what is this? This is a new effect.

One problem that we had in the experiment was that we are driving the string in transverse oscillations like this and we have these mass scatterers, if you drive a string at finite amplitude, it modulates the tension. If you take a wire and modulate the tension, then you can excite the transverse waves, but you can excite them up and down and you can also excite them back and forth.

What happens is, at finite amplitudes you start getting torsional effects showing up. The masses were not really point masses, so at some amplitudes you could see the masses start doing this. It was at real low frequency but it is a nonlinear system and one of the things that nonlinear systems do is they will feed a little bit of energy into low frequencies.

We thought that might be a problem and we tried to get away from it with some clever experimental things, but it was getting too messy. We finally said, oh, my gosh, we are going to have to do a computer simulation.

(Transparency 41)

We actually took the nonlinear equations for the string and put them in the computer and did a Runge-Kutta integration of the nonlinear equations. Here is what we found. This is essentially the same sort of spectrum. It looks a little different from the experiment but that has to do with the finite output impedance of our driver.

There is the spectrum, one of the bands for the periodic system. As you increase the amplitude in the computer simulation, you find that this first thing goes over and drops off; it gets busy right in here. So we took finer frequency steps in the computer simulation and it just got busier.

(Transparency 42)

It is a real effect. It was not due to just torsion on the masses. We tried to find the right buzz words to assign to this and came up with ``fractal nonchaotic behavior," but that just gives a name to the effect but does not really say why it happens.

I think we now understand what is happening. It has to do with the fact that more than one tuning curve is bending over, and they are bending over top of one another. We should be able to calculate that for the periodic system.

Q: Is that a fractal dimension?

A: This structure is fractal. If we did finer steps, we would see even more structure in there, and so on.

Q: In your simulation, were you sweeping frequency or would you just sit on a frequency, integrate until you came to some steady state solution and then you called that the amplitude?

A: Yes. Then, with that system we let it run until it came to steady state. Then at some point we changed to the next frequency, let it run, come to steady state, and so on, pretty much the same way we did the experiment. The experiment was done with a computer-controlled frequency synthesizer, so it would run, come to steady state, and we would change the frequency by 5 mHz, and so on.

Q: Are those forward sweeps?

A: We also did backward. For the regular part of the spectrum you do see the usual hysteresis and so on, but when you get back to the front of the spectrum where you had the busy stuff, you get the busy stuff again.

When you are doing the bent tuning curve, you can take the equation for the regular tuning curve, do some algebra, and you will get the bent tuning curve. We wanted to do the same thing for the system where you have many tuning curves, the band structure for a periodic system.

In order to do the bent tuning curve for that, you have to have an analytic solution for the periodic system. We sat down and looked at the textbooks. Here is the problem.

(Transparency 43)

We have a 1-D wave system with some boundary condition at the end. For periodically spaced scatters what is the driving point impedance? It is periodic, so it ought to be easy to do, because we know that in a periodic system -- we already know what the eigenfunctions are. They are Bloch eigenfunctions.

Well, you go and try to do this problem and you find out that Bloch periodic problems are not as easy as the textbooks would have you believe. We actually did the analytic theory for this thing and found out the driving point impedance.

(Transparency 43)

(Laughter) That is not it. That is just the real part of the numerator. There is an imaginary part and a denominator. It is a lot more complicated than the textbooks will have you believe, but it is all in your handouts, all the gory details. It is not quite this bad -- I mean, I did this on purpose. You find that some things recur and so you can call this S+. You can have a number of smaller formulas. But it is quite complicated.

We got this analytic solution, so we could just plug in the parameters and we will get the driving point impedance for the linear system and then we can do the bent tuning curve analysis to see what happens when they bend over.

It is a messy formula and we wanted to test it. Another way of calculating this thing is to multiply by the 2 by 2 matrices. That is something that is easy to do numerically. You would get an even bigger mess if you did that just analytically for n masses. But we could at least test this formula numerically, so we could plug in values for the parameters, get a number, multiply the matrices and get a number, and compare.

We did that and they did not agree. It turned out there was just a sign error, one sign error which we fixed and so they agreed. So we now have the analytic formula for the periodic system, and that is what it looks like.

(Transparency 44)

You will notice that that agreed with the low-amplitude Rungacutta simulation. We also have data that show this. It is bigger at one end than the other. If you think about it, you can see why that is. It is an impedance thing.

This gives you the $A = F\left(\frac{\omega}{c(A)}\right)$ for the linear system. In the nonlinear system, a finite amplitude changes the tension and

that changes the speed of sound. For the nonlinear system c is now a function of A and you need to find solutions where these A's are the same.

You pick a frequency, start stepping through A and you plug in here and you look at this A. Is it the same? No. You keep going. And when they are the same, you plot that point.

(Transparency 45)

And this is what you get. You get the bent tuning curves, but you see that right in this region here, at this particular frequency you can have as many as 10 different modes running all at the same frequency, just with different amplitudes. This is right where we got all that busy stuff.

That is a step in the right direction, but there is still a puzzle, because the usual theory for the bent tuning curve says that as you change the frequency you should go up the bent tuning curve and then `fall off, up this next one, fall off, or you might fall off to a different one. But the thing is, after you fall off you should follow that stable line, but it does not do that, so there is more going on.

We think the explanation has to do with what is called basin crowding.

(Transparency 46)

This is a phase plot, displacement versus velocity, a simple harmonic oscillator just goes around in a circle. Here the thing is running in one of its modes. At some point -- say here -- we changed the frequency, and when you change the frequency, what happens after that?

What it was doing right before you changed the frequency becomes the initial conditions for what it does after you change the frequency. When you give something initial conditions in a nonlinear system, where it goes depends on what basin of attraction you are in. The thing is, for a complicated system like this, apparently the basins are crowded. You can change things by just a small amount and you might go into a basin where the amplitude drops way down or you might go into a basin where it just increases slightly. But if the basins are crowded, there is sort of no telling where it is going to go.

It may be that just a little bit of noise can send you one way or the other.

Q: This sounds very similar to what happens in chaos.

A: Right, yes, it all comes from the same thing, except that nothing has gone chaotic yet.

Q: It is hypersensitivity to initial conditions.

A: It is all part of that. The only difference here is that usually the chaos theories are done for essentially a single-mass system, very small degrees of freedom, and here we have lots of degrees of freedom and yet we are seeing something systematic in it, so I think that is kind of exciting.

Q: Is that buzz at the beginning of the trace repeatable or do you get a different structure each time?

A: I do not think we ever looked. We were worried about this business, which certainly put noise into the thing. The computer simulation, of course, repeats. It does have digital noise in it. That is the whole problem with these nonlinear things. You know your experiment is going to have some noise, cosmic rays, whatever, but computer simulations also have noise. Numerical roundoff is the same as noise.

Anyhow, that is something that one of these days we will have to write this up, but not yet. I think it is time for another break.

DR. MAYNARD: I mentioned that one of the hot topics for disordered systems is the normal electron persistent current.

(Transparency 47)

This is an experiment that was done by Richard Webb at IBM Research. He got the Buckley Prize for this. What he did is, he made a loop of gold wire. The gold was about 100 Å wide and the loop was a couple of μm in diameter. He cooled it down to 40 mK. At that temperature and this size the electron could go around the ring a couple of times before it underwent inelastic scattering.

So, for at least once around, the electron can maintain its phase coherence and, because of that, it could take part in a persistent current. This would be a persistent current for a normal electron. If you have a superconductor, in a superconductor the electrons pair up through a phonon interaction and that uses up the phonon scattering. For the pairs of electrons, the Cooper pairs, there is no more phonon scattering, so those Cooper pairs can move without inelastic scattering, and so in a superconducting wire you can have a persistent current of these Cooper pairs.

But here, in Richard Webb's experiment, he got rid of the electron-phonon scattering by brute force, just getting rid of the phonons, and so he could get a normal electron persistent current. If he has a persistent current or if the electron maintains its phase coherence around the ring, then it turns

out that the magnetic flux (Φ_0) through this ring has to be quantized, and it is quantized in units of the normal electron flux quantum, which is h/e. In a superconductor that would be h/2e, because the electrons are paired.

Also, if the flux is quantized, then the current around the ring is also quantized, so in Richard's experiment he made a plot of the current in the ring as a function of the magnetic flux through the ring. He found that in the experiment the current oscillated and the period of the oscillation was h/e, the normal electron magnetic flux quantum.

But he found that the amplitude, which is how the current should have been quantized, was more than an order of magnitude, almost two orders of magnitude, larger than what had been predicted by the theory, so there was a big discrepancy. All the theoreticians were running around crazily trying to explain this.

I took a look at this and found that you could do an acoustic analog of this normal electron persistent current experiment.

(Transparency 48)

First of all, if you go to a real simple level, the electron obeys the Schrodinger equation, which is a wave equation, so if you have an ideal ring with a circumference, L, then you can just have traveling waves going around the ring.

The quantization condition is that you have to have an integral number of wavelengths around the ring. If you take this wave function and calculate the electrical current around the ring, this is the $i=\frac{Nneh}{mL}$ (if you have more than one electron). That just boils down to this. The wavelength is just the L/N, and so you get that the current -- this n has to be an integer -- is quantized in units that are given by this.

You can rewrite this like this, and this quantity here just turns out to be the Fermi velocity for a collection of electrons. The bottom line L´is that the current is $\frac{ev_F}{L}$ of thering. This is a real hand-waving derivation; you can do a more careful job and get the same result.

You notice that if you take the circumference and divide it by the Fermi velocity, the velocity of the electron, you get the time it takes the electron to go around the ring. If you take the charge of the electron and divide it by the time it takes to go around the ring, you get a current. That is a real simple way of remembering that result.

(Transparency 49)

Richard Webb's gold ring was not a perfect waveguide for the electrons. It had disorder in it, so it had a disordered array of scatterers. So now we have a wave, the electron Schrodinger wave, going around the ring, but it is undergoing elastic scattering off these defects.

An easy way to look at this is that the electron, instead of going around a distance L actually follows a tortuous path, and so there is a tortuosity correction. Instead of using L we should use L′, which is the actual path length that the electron follows.

If you put that into the equation we had before, you can write it this way. There is a correction for the current. Since L′ may be much larger than this circumference L, that makes this current a lot smaller. The theory predicts that you should get this small current that has this correction factor.

If you look at this combination here, because the electron is following this tortuous path, it seems to take it longer to get around or it looks like it is moving slower. This is something we encounter in acoustics -- it is called the tortuosity correction -- for waves propagating in a porous medium, for example. This is in the case where the wavelength is much larger than the size of the scatterer, so you actually do not get Anderson localization, you get sort of an effective medium. That effective medium makes the sound wave look like it is moving slower. It is a tortuosity correction.

In acoustics, when Steve and I and other people did fourth sound we had superfluid waves propagating in a powder, so we had to worry about this tortuosity correction.

(Transparency 50)

There is a model that you can use to treat this tortuosity correction. It goes like this. Suppose you have a black box of length L and inside this black box you have just a straight waveguide. You send in a wave, $e^{i\frac{\omega}{c_0}x}e^{i\omega t}$, you send that in and you measure the phase shift across the black box. You may have to count two p's or something in some mysterious way, but you measure the phase shift across the box, and that will be given by $\frac{\omega L}{c_0}$. You can turn that around and solve for c_0: $\frac{\omega L}{\Delta\phi_0}$.

Now suppose you have a black box that inside has a meandering waveguide. A model you can actually do calculations with is a square wave meandering waveguide. Now the length of

the waveguide inside is now L′. If you send in a wave, you measure the phase shift across that. Phase shift is going to be given by L′. But this is inside a black box, so you do not know that, so when you do the calculation of the speed of the wave, you use the length L.

So if you plug this in, you get that the speed that you measure is $tc = c_0 \frac{L}{L'}$. So this is a simple model of how you get a tortuosity correction for the speed.

(Transparency 51)

Now suppose you take this meandering waveguide and bend it around to make an annular waveguide. Now the meandering waveguide models the tortuosity correction that you would have if you had the disordered scatterers.

If you take a source and send a wave around this meandering annular waveguide and measure the phase shift, or measure, say, half of the phase shift, and calculate the effective speed of the wave, then you get the tortuosity correction.

But sending in a wave like this and having it go around would coincide with a transport measurement in a linear system, but that is not what Richard Webb did. He did not make a transport measurement. What he did was, he had a magnetic field through his system, and to actually make his measurement he oscillated that magnetic field.

Here is the question. If you have a magnetic field in Richard Webb's experiment with the electron, what is analogous to that in the acoustic system with a fluid? Rotation. In fact, if you have a magnetic field, there is a vector potential that gives the electron wave a phase shift up in one direction and phase shift down in the other direction.

In the acoustic system, if the fluid is flowing in here, then one wave is Doppler-shifted up and the other wave is Doppler-shifted down. So the analog of having a magnetic field is to have the fluid flowing in the waveguide.

In Richard Webb's experiment he oscillated the magnetic field and that would correspond to taking this waveguide and oscillating it. So that is a different way of driving this thing than having a source of waves here and a receiver over here. You are driving it like this.

(Transparency 52)

Here, now, is what would happen if you drove this system like this. When you drive it like that, then you can induce a

phase shift in the azimuthal parts of the waveguide but not in the radial parts. In the radial parts, when you are doing this, you are taking that waveguide and just sloshing it, or you are taking the fluid and just sloshing it back and forth like this. That excites a mode in the waveguide like this, and that has infinite phase velocity.

So by rocking the waveguide like this, you induce phase shift here but not along the radial parts, and these parts, when you add them up, where there is a phase shift, you just get a length L and not L′. So there is a hint that maybe by driving the thing like this, the L that appears in the equation should be the actual circumference L and not the L′.

If you put that into the equation for the current -- remember, the current, $\Delta i = \frac{ev_F}{L}\left(\frac{L}{L'}\right)$ but if that is not right, then this becomes much larger and it is just the right size now to fit the experiment. The scattering can change the transport in the system, but by driving it this way somehow you avoid that scattering effect.

There is a problem in that if you excite phase change this way, that wave can just turn the corner and undergo phase change that way. In other words, you can excite modes that do have phase change along the radial parts. But if there is damping in the system -- in fact, the whole reason you have to drive it is to overcome the damping of the wave -- if you drive it like this, you are certainly going to favor driving the modes that have phase shift only along this way and not ones that have phase shift along the radial parts.

It is a question now of how well the drive couples to the different modes and overcomes the attenuation, so that is a more complicated problem. To study this, we thought we would do an experiment to see what different modes of driving would do. You want to see if you can see the difference between L, which is going around just these parts, and L′, which goes around the whole thing.

If you want to see the difference between L and L′, of course, you use a waveguide shaped like the letter L. You can drive this thing in such a way as you can drive modes that would go around like this and then that total distance there would be L′, or you could just drive the thing back and forth so that you are just fitting waves along this edge, and that would be the length L. So we just use an L-shaped waveguide.

(Transparency 53)

You want to look at wave modes that have frequencies that are less than the resonances in the material you are making the waveguide out of. We decided to use slow surface waves.

First, we tried water and discovered that water is a terrible fluid to use to study surface waves. The surface gets messed up very quickly and it is a pain to clean and so on.

We switched to mercury. Mercury has much better surface effects. In fact, it has a fairly high-Q surface wave mode, but you do not want to use it too long. For our L-shaped waveguide with the surface waves we took a block of teflon and machined the L out of it (Laughter)

We put in a capacitor-receiver that would sense the change in the height of the mercury. In one experiment we put just a point plunger to drive it, so we could drive modes that had waves going around the corner and could also drive modes that had waves parallel to this edge, and so on, so we could look at all the modes with that driver.

Then we took the whole thing and just shook it back and forth like that to see if it would favor modes that had no phase change along this section of the thing. This one would see modes with an L´ in it. This would see modes with just the length L in it.

(Transparency 54)

Here are the results. This is with the point source drive and all the different modes that you can get. I will not go into what the nomenclature is, except to point out that there is an L´ in these. But if you drive it back and forth, which is equivalent to taking a meandering waveguide and rocking it, then you find that the biggest peaks are the ones that involve just the length of L.

It is really kind of obvious, but it is something that has to be demonstrated experimentally.

Q: Are you below the cutoff for the width mode?

A: There is a width mode in here.

So it works for the acoustic analog. The problem now is to translate this into the quantum mechanical system. That is a problem I am not going to tackle. We are driving this thing and we are doing that so we have to feed energy back into the system to account for the attenuation.

In Richard Webb's experiment with the electron, the quantum mechanical electron, there is inelastic scattering that takes energy out of the system, and so he puts on an AC magnetic field to -- well, his inelastic scattering would kill

his persistent current in $10^{=12}$s or something, so to see the effect at all he has to have this AC magnetic field.

That AC magnetic field does give you a drive force that is only azimuthal, it has no radial part, so there is some parallel but I just do not know how to do the quantum mechanics problem to really prove to the theoreticians that this has anything to do with that discrepancy.

It is still a discrepancy. The theoreticians are trying to blame it on many-electron effects and in these wires there definitely are many-electron effects. But the point of this experiment was to show that maybe you do not really understand the single-electron system, either, and how you are driving the quantum mechanical system and so on. Someone who knows a lot more about quantum mechanics is going to have to take over from here.

NONLINEAR WAVE PHENOMENA

DR. DENARDO: What we are going to be looking at this morning are three phenomena in nonlinear waves: self-localized waves, absorption of sound by noise, and wave turbulence. There will be a break between each of these. The second and third parts of this lecture were prepared with Andrés Larraza.It is actually good that this lecture is later in the summer school, because a lot of this material will complement many of the lectures you have seen before, particularly some of Bob Keolian's demonstrations and Steve Garrett's talk.

The focus here is going to be on nonlinear waves in a variety of systems, not just nonlinear acoustics. There will be many similarities but there will also be some contrasts.

(Transparency 1)

We will begin with what we call self-localized waves. A self-localized wave is a wave that is exponentially spatially localized due to its finite amplitude, and not due to nonuniformity. Because it is due to its own amplitude it is called self-localized.

A solitary wave is a self-localized wave of permanent form in the limit of no dissipation. Note that we do not include shock waves here. We also do not include a wave on a string, which moves along undistorted, so we add that self-localized waves should be dispersive waves.

A soliton is a solitary wave that collides elastically with all other waves in the system. The first discovery of a soliton, a solitary wave which turned out to be a soliton, was as a hump in a canal in Scotland by John Scott Russell in the 1830s. This was generated, I believe, by mules dragging barges. When the mules stopped, a pulse would be sent out. Scott Russell noticed that this hump could propagate for long distances, on the order of a mile, without essentially any distortion. It would just slowly dissipate away.

This is remarkable, because the surface waves here are dispersive. Steve Garrett derived that they are nondispersive in the limit of very long wavelength compared to depth. This is not the case here, and the dispersion turns out to be essential.

On the linear level we can think of this hump as being as being a superposition of different frequencies. Those frequencies travel at different speeds. So this thing should spread, but it does not spread. What is going on here, and it is not obvious, is that the nonlinearity tends to make the waves shock, or break in this case, and there is a stable balance between that tendency and the tendency to spread due to

dispersion. It is not obvious how that could come about and also how it could be stable.

What we are going to do later is explicitly look at such a case in a certain class of solitons. For all solitons there is some kind of balance to achieve a permanent waveform. This is always the case.

The Scott Russell soliton is referred to as a KdV soliton due to the authors (Korteweg and de Vries) of the theory that came about around the turn of the century. The solution is a hyperbolic secant squared for the surface height. It is exponentially localized in the wings, and it moves along undistorted. There is a free parameter here -- I have expressed it in terms of the amplitude here. Higher amplitude solitons are shorter. This is always the case, as far as I know. And they travel faster.

Here is a collision of two solitons shown in space and time. They are both traveling in the same direction and the higher amplitude one overtakes the smaller amplitude one. These two humps emerge from the collision with no change in shape. That is what we mean by an elastic collision.

I want to make sure everyone understands here that the disturbance is not the sum of the two in the interaction region here. Superposition breaks down, because the system is governed by nonlinear differential equations. You can see that there is some information that a collision has occurred. What is that?

Q: The tall wave looks like it is ahead of where it would have been.

A: Exactly. There is a phase shift. So there is information that the collision has occurred. You see that, the higher amplitude soliton is farther along, and the smaller amplitude soliton is not as far along. The center of mass moves with constant velocity.

(Transparency 2)

There is a very important, very broad class of solitons that are called nonlinear Schrödinger solitons. Such a soliton can be thought of as a monochromatic wave that has an amplitude modulation that is exponentially localized. It moves along -- the whole packet moves along -- with approximately the group velocity. Actually, the bigger the amplitude is, the faster it will go, and the lower amplitude ones move at approximately the group velocity. These have been observed as surface waves on liquids, in flexing shells and, perhaps most importantly, in optical fibers. According to AT&T Bell Labs, these are going to be employed in trans-Atlantic fiberoptic communications before the end of this decade. There is an interesting story

there. When you send a fiberoptic wave over large distances, it turns out to be unstable. This is a nonlinear instability. You will see, in a videotape, the analogous instability in a lattice. People fought this for some time. What happens in this instability is that the energy wants to clump into solitons. It was finally realized that this should not be fought. Why not go with the natural modes of the system over large distances, which are solitons? You have the advantage of these also being digital. Each is like a bit: it is like zero; or one, it is either not there or there.

There are other types. Here is a kink-type soliton. If you imagine a pendulum lattice, where the pendulums are coupled by a torsional wire, it should be clear that there is a state in which all the pendulums are hanging straight down on one side, and all of them are hanging straight down on the other side, but in-between we can have a 360° kink. That can just sit there, or it can move. Here is the general solution for this moving kink. This is called a sine-Gordon kink soliton. It is a genuine soliton.

DR. HARGROVE: Bruce, it may not be clear that the displacement is in the plane normal to the line -- some of them look like they are swinging the other way.

A: Yes, you are right. I should have stolen one of Bob Keolian's drawings; he does a better job. This is a transverse lattice. Displacement is perpendicular to the lattice.

(Transparency 3)

There have been other observations. Analytical work became a hot topic in the early 1970s. There were not that many observations in the beginning and there is still an absence of observations, but they are now coming in more frequently. There are an incredible number of predictions of solitons -- every month you will see another one -- in all kinds of systems. We are going to focus on what we call cutoff solitons and related localized states..

(Transparency 5)

Here is an example of an observation. It is now settled that this is a soliton. This is a satellite picture of the Andaman Sea. What you see here are KdV-like solitons. The humps are about a kilometer apart. These are soliton-like internal waves in the ocean. An internal wave is a temperature-type wave that can exist in a fluid that is stratified due to a gravitational field.

Q: Bruce, if somebody sees a picture like this, how does one know that this is a soliton wave or that it is not just a traveling wave or a standing wave or whatever? How does somebody know that?

A: One thing here is that these humps are very far apart compared to the width of a hump. But if you think of a normal traveling wave, the length of the hump is comparable to the distance between the humps.

But this was actually questioned for a long time, I believe. Ultimately, I think it takes a comparison to theory before anybody really believes it.

DR. PUTTERMAN: Isn't what they did was they followed it for so many hours and then, based on theory, as you said, they calculated how much the width of each individual one would have spread had it been a linear dispersive wave running on the thermocline of the ocean? They claimed that it did not spread that much and, therefore, they concluded it is a soliton.

A: Right. It is like what John Scott Russell established, the fact that the hump did not spread. That is another strong indicator. If you have a dispersive wave that is not spreading, that is a very strong indicator that you have a soliton. Technically speaking, they also have to collide elastically, although people a lot of times do not insist on that.

Q: Did you say this was a temperature wave?

A: They are temperature internal waves. It is a temperature-type wave.

Q: What was the temperature fluctuation?

A: I do not know. I just found out a couple of weeks ago in an Acoustical Society meeting there have been other solitons like this in thermoclines where the temperature in the ocean changes rather abruptly. It is now pretty well-settled, evidently, that these are solitons (I want to get that paper).

Q: Can the waves in this sea, the solitons in this sea, a satellite would have the capability of taking periodic pictures, say, every 20 minutes, and I suppose -- is it true that if these persist for a long time it is a soliton? There are certainly observations there in the sea that persist.

A: But normal waves can persist sometimes, too, right? I am not sure what you are getting at.

Q: I guess I am asking a question. Does a soliton persist for a long time and is that proof that it is a soliton?

A: No. I will show you an example of solitons that do not last very long at all because of the dissipation in a particular system -- you will see that in a videotape -- but they are solitons.

(Transparency 6)

To understand cutoff-type solitons, we are going to go back where nonlinear physics begins, and that is with a pendulum. It is a planar pendulum. As I am sure all of you know, the exact equation of motion is given by this. We are going to look at the weakly nonlinear motion here. You always model nonlinear waves;
you always make approximations.

There is an exact solution in this particular case of an oscillator but it really does not buy us much, so we will look at the weakly nonlinear level, expand the sine to the cubic or to the first nonlinear order. This is the weakly nonlinear equation of motion.

For the pendulum, this quantity here is equal to $\frac{\omega_0^2}{6}$. This is called the nonlinear coefficient. For the pendulum it is a positive number. We want to generalize this and consider nonlinear oscillators where this nonlinear coefficient can be positive or negative. That is going to be particularly important, as you will see.

Bob Keolian talked about this, so I will move through it quickly. A simple way to look at what is going on here is that for the linear motion, if we look at the potential energy as a function of angle, we have a parabola. This is the Hooke case, where we have an exactly linear storing force.

In the case of the real pendulum with a positive nonlinear coefficient the force does not rise sufficiently fast with distance and we have what is called a softening nonlinearity. So if a particle is oscillating in this potential well, it takes a longer time to complete one cycle.

As you go to small-amplitude motion, the period becomes amplitude-dependent. That is the linear level. In the hardening case, the frequency goes up with amplitude. In the softening case, the frequency goes down with amplitude.

(Transparency 7)

It is easy to show this with a pendulum. You could time it, but there is an easier and more interesting way of demonstrating the fact that the period depends on the amplitude. Here is an actual pendulum. The claim is that the time it takes to execute a large amplitude cycle is a little longer than the time it takes for a small amplitude cycle. Rather than measure it, we will utilize the fact that this pendulum can go in two directions (this is a spherical pendulum). There is an approximately isotropic support up

here.

On the linear level, what is the general motion of a spherical pendulum? Two harmonic oscillators with the same frequency, x and y motion? The general solution is ellipse. If you look at that, it looks like an ellipse to me. The motion closes on itself. Why does the motion close on itself? In this case the frequencies are the same, so the motion closes on itself. There is a degeneracy here.

But what if one of the amplitudes is made large? Now watch what happens. This can be described as a precessing ellipse. It is an approximate ellipse over one cycle, but it does not close on itself. We can describe this as a forward-precessing ellipse.

This can be simply understood if we restrict ourselves to the case where we have an infinitesimal amplitude in one dimension and a weakly nonlinear amplitude in the other. Then we can use superposition. We can imagine that this is the superposition of a motion in the y direction and a motion in the x direction. Because the pendulum softens, the time it takes to complete the y motion is longer --that is the claim -- than it takes to complete the x motion.

What happens is this: when it has gone one cycle of the x motion, it has not yet completed the y motion. It will then go out here and that is why the softening nonlinearity will create a forward-precessing ellipse. A hardening nonlinearity will create a backward-precessing ellipse.

I am sure most people know that there is a famous case of this kind of precession: the planet Mercury. For the Kepler problem there is a degeneracy between the radial frequency and the angular frequency. The fact that those two are the same gives you closed orbits, but that degeneracy can be lifted by either an oblateness of the sun or general relativity. Mercury's orbit does not precess as fast as the one you just saw there.

(Transparency 8)

How do we handle this nonlinearity analytically? I am going to give you the simplest case and then comment on more complicated cases. Here is our weakly nonlinear equation of motion. We know that the motion will be approximately sinusoidal with some frequency which will be close to but not equal to ω_0. This is due to the fact that the motion is weakly nonlinear.

If we imagine substituting the sinusoid into the nonlinear term higher harmonics will be generated, and those in effect drive the linear equation. Steve Garrett pointed this out and

I am going to do it again and again today. That is how we usually approach nonlinear wave equations. Because it is being driven at these harmonics there will be a response at those frequencies. There will be higher harmonics here at $3\omega t$ and $5\omega t$, etc. But because we are on the weakly nonlinear level, those will be smaller and smaller.

All you do here, and this is very simple, is just substitute this into the equation of motion and equate like-frequency terms. When we cube the $\cos(\omega t)$ we get a $\cos(\omega t)$ term back. The Fourier decomposition of the cube of the $\cos(\omega t)$ is a $\cos(\omega t)$ term and a $\cos(3\omega t)$ term. What is relevant for us is $\cos(\omega t)$. We then set the like frequency terms equal and we get to lowest order how the frequency depends on amplitude, as a quadratic correction to amplitude. I will show you a curve of this in a moment.

You might wonder, what if there is a quadratic nonlinearity in addition to the cubic nonlinearity? This is typically the case in wave systems.

If you square the $\cos(\omega t)$, you get a dc term and you get a $\cos(2\omega t)$ term. You have nothing in $\cos(\omega t)$, so does that mean there is no frequency shift? No, it means you have got to go to the next order. You have got to generate the second-order solution, which, in this case, will have a dc plus a $\cos(2\omega t)$ terms, and then substitute that back into the quadratic term, and then you will get back to the original frequency. When you do this, you will find that this shift is also proportional to the square of the amplitude. There is a very important point here in dealing with nonlinear oscillators and nonlinear standing waves: You cannot neglect the cubic nonlinearity compared to the quadratic nonlinearity, at least with regard to the frequency shift. They both enter at the same level, proportional to the square of the amplitude. The first mistake everybody makes with nonlinear oscillations is to neglect the cubic term compared to the quadratic. You cannot do that.

This formula happens to be consistent with the exact frequency of a pendulum given in terms of elliptic functions. If you take the exact solution and look near the linear limit, you get that.

Q: Why can you neglect the fifth order term?

A: That is a good question and I think I know what you are getting at. I was wondering if someone was going to ask that. This is a sharp group, better than two years ago, much better. I think what you are saying is if you take the $\cos(3\omega t)$

term, so if I imagine a coefficient B here, I will get contributions at $\cos(\omega t)$. That is a higher-order correction, however. To really believe that, you have just got to sit down and go through the calculation.

(Transparency 9)

What happens when you put in drive and dissipation? Bob Keolian talked about this. Here is the linear level. Here is what we just solved for. This is sometimes called the backbone curve. If we look at steady-state response due to a drive, we get tuning curves. Here are linear tuning curves. These will conform to the amplitude-dependent frequency here. They will be bent over like this. In fact, as Bob pointed out, a curve can bend so much that it becomes multi-valued here and this gives rise to hysteresis. Did everyone understand when Bob talked about the hysteresis? It is one of the simplest cases of hysteresis. If you start here and slowly increase the frequency, when you get to this point the oscillator will jump up and ring down above here and then will go along the upper curve. If you decrease the frequency, it will jump down here. Jay Maynard also talked about this.

This is an example of a direct drive. You are all familiar with this. Another way to drive a system is with a parametric drive. Moving the support of a pendulum offers two ways of appreciating this. Suppose the support is moved back and forth horizontally. In the frame of reference of the support, we have an additional force, an inertial force that comes in, because we are in a noninertial frame of reference, so there will be an oscillating force back and forth on the mass, and we will get resonance if we drive this at the natural frequency of the pendulum. That is an example of a direct drive.

Parametric drive is where you oscillate the support vertically. In this case, as Bob pointed out, if we go into that frame of reference, we now see there will be an additional noninertial force. It is oscillating in time and it acts just like an additional contribution to gravity. It is as if we were modulating acceleration due to gravity. Because we are modulating that parameter, this method is called parametric excitation.

Here is what the tuning curves look like. In this case, $\theta = 0$ is always a solution, by symmetry. The issue is whether or not it is stable. If you are driving at near twice the natural frequency you can feed energy in. My hand is going up and down at twice the rate of the pendulum; for every one cycle of my hand the pendulum is making half a cycle.

There is an infinity of discrete frequencies at which you

can feed energy in here, but the easiest way to get energy in is when the drive frequency is twice the natural frequency.

(Transparency 10)

Let us now couple these pendulums together. This is a transverse lattice, which is a particularly simple and very useful model. You can imagine that there are springs that couple the pendulums. Here is the weakly nonlinear equation of motion; this is f = ma for the lattice. Here is the acceleration. Here is the coupling, the usual coupling that occurs for the mass-and-spring lattice that I am sure you are all familiar with. Here is the linear restoring force and then here is the nonlinearity due to finite amplitude motion. We are looking on the weakly nonlinear level and, again, we will consider positive or negative values of the nonlinear coefficient.

Always in any wave system one of the first things you have to do is understand the dispersion law, how monochromatic waves of different frequency propagate. We get that by plugging in a monochromatic wave. When we speak of dispersion law, we almost always mean on the linear level. The dispersion law is given by this. It gives the frequency as a function of wavenumber.

There are two cutoff modes here. The lower cutoff mode, where all the pendulums are going together, is the lowest frequency excitation the lattice can have on the linear level. The upper cutoff motion, where they are out of phase, is the highest frequency. The way you can appreciate this is to imagine you are shaking one end of the lattice back and forth. If you are driving below the lower cutoff frequency, you will not get propagation. You will get an evanescent wave. If you are driving above the upper cutoff frequency, you also do not get propagation. You get an evanescent wave. Propagation only occurs in this band here.

(Transparency 11)

What happens on the nonlinear level? Entirely new phenomena that have no linear analog arise. We can demonstrate this in a simple case. Here are two coupled pendulums. This is a particularly simple coupling; you do not need springs or anything. With this demonstration, we can demonstrate beats on the linear level . If we start the energy off in one oscillator, because these are nearly identical pendulums, we will get a complete energy transfer to the other one.

What you see is linear beating, because we started the system off in a superposition of normal modes, which are the symmetric and antisymmetric modes.

On the nonlinear level we can ask what happens to these modes. Nothing happens to the antisymmetric mode. But watch

what happens to the symmetric mode. There it goes. It was very close to the instability threshold. There is an amplitude threshold hidden here. The nonlinear world is full of instabilities just lurking and waiting to happen. And they can have disastrous results.

There is an instability here. We put the system on the computer, we solved it analytically, and we had this demonstration to believe what was going on here. It turns out that the symmetric mode parametrically drives the asymmetric mode. This is a parametric instability. It is not obvious why there is an amplitude threshold and what is going on there, but I encourage any of you to just sit down and, if you are interested, write out the equations of motion. Do not try to handle this coupling on the nonlinear level. It is a beautiful way of coupling pendulums experimentally, but it is a mess analytically. You do not want to deal with this coupling on a nonlinear level. The reason is that the motion is no longer planar.

Q: What were we supposed to see on this last demonstration when you had both of them?

A: Let me do it, again. This mode is not stable on the nonlinear level.

You can put this on the computer and, with perfect pendulums, start it off and give a slight fluctuation and sometimes one will soak up the energy and sometimes the other one will. It is a genuine instability. And it is very general. It should occur in any system of coupled nonlinear oscillators. Absolutely general.

Q: Does the energy stay on one side?

A: No, it will beat back and forth in a kind of complicated beating, similar to the linear case.

Q: So the only way that you can tell the difference between this being beating and some nonlinear thing is that the energy does not go back and forth like it does in a linear system?

A: No, once the instabilities come in --

Q: But it does not go back and forth with the next beat frequency?

A: Oh, it will go back and forth with some frequency but it will be amplitude-dependent and that frequency will change in time. The key thing here is that on the linear level this mode is stable.

(Transparency 12)

In the lattice, what arises? New phenomena that arise are localized states, spatially localized states, and we are going to look at the simplest case here first. These are sometimes called standing solitons. Imagine looking down on the pendulum lattice. These arrows represent the motion of the pendulums.

There are breather localized states and kink localized states. The kind you get depends upon whether you are looking at an amplitude modulation of the lower cutoff mode or an amplitude modulation of the upper cutoff mode. It also depends on what kind of nonlinearity you have. These are easily explained and let us do that right now.

Let us look at the lower cutoff softening breather. Start off with this profile. Imagine we have these pendulums going like this. I will argue that this is a mode of the system, that everything is moving with the same frequency. That will give us an appreciation for the existence of this state.

What is happening here is we can think of the body of the breather as driving -- out here in the wings, where we have linear theory -- as driving the rest of the lattice. It is very important here to note that this state exists at a frequency below the linear cutoff frequency. Let me just state that for now and we will soon see why.

Because the body is driving the rest of the lattice below cutoff, there will be an evanescent wave. This is an exponentially evanescent tail here.

Inside the breather we have to understand how we can have the same frequency all along here below the cutoff frequency. Here the frequency will be shifted from the linear frequency by two effects. First of all, there is a finite amplitude effect. Because we have a softening system, that will lower the frequency.

However, there is curvature here, there is coupling to the neighbors, and because it is negative curvature, this will cause the frequency to rise. So there are competing effects here between the curvature, the coupling, which wants to raise the frequency, and the nonlinearity, which wants to lower it. In order to get a frequency below cutoff, the nonlinearity has to beat out that coupling.

We need to consider a state that is sufficiently spread out compared to its amplitude, so that we get a frequency below cutoff. This is an argument that indicates how you can have a state here where everything is moving at the same frequency. Everything is consistent.

The picture that emerges here is that this is a self-trapped state. The energy in the body of the breather cannot

get out, because in the wings here the system is driving the lattice below cutoff. In other words, it digs its own hole. That is how Professor Isadore Rudnick described it.

The kink is very different. Let us imagine we have a lattice with free boundary conditions, and let us start off with a linear mode. The first mode will have a uniform profile, and the second mode up will have a sinusoidal profile. And everything will be moving at the same frequency. It is a mode.

Now we ask what happens as we crank up the amplitude of the response, but we still want to have a mode, we want everything moving at the same frequency. What does the system do, if it can, to have everything going at the same frequency?

Let us just first scale up the cosine profile here. The frequency in the wings will increase due to the nonlinearity, but it will stay the same in the nodal region, because this is linear. So we do not have a monofrequency state. The response has to deform; it has to do something so that everything is going at the same frequency.

What it does is everything it can do. It flattens the profile in the wings to reduce the coupling and thereby lower the frequency and it increases the curvature in the nodal region thereby increasing the coupling, to bring these frequencies up. So it brings these frequencies up and these frequencies down such that everything is going at the same frequency.

It does not sound like a big deal, but look what happens. As you will see in the video, for all intents and purposes the wings become flat, so that the nodal structure does not know that the boundary is there any more. This becomes like a particle and, in fact, it is free to move around and you will see that it can be very sensitive to nonuniformities. This is a kink.

Let me show the video. Before I do, let me mention you will also see some surface waves. The pendulums going back and forth near the lower cutoff mode are a lot like water sloshing back and forth in a trough. Because surface waves with sufficiently deep liquid soften, we expect to get breathers here and, in fact, we do.

I want to make the point that historically the surface wave breather was the first of these cutoff states to be observed. It was observed without any theory. People were working in the undergraduate acoustics lab at UCLA, in Isadore Rudnick's lab, and found it. This led to all this other work on cutoff-type solitons.

DR. MAYNARD: You can obviously get all the pendulums to

swing together at a large amplitude at a frequency that is below cutoff without the breather, but is this going to be unstable to the formation of the breather?

A: Yes, and you will see this in the video. This is a manifestation of the Benjamin-Feir instability, -- and that is the same thing that causes the light to want to clump into bundles of energy in a fiber. It is kind of the many-body version of the instability you saw in the two-pendulum demo. It is similar to that, but that is the a few-degrees-of-freedom case.

Surface waves in water have the remarkable property that if you go to shallow depth the mode hardens. So now we expect to see kinks if we go to shallow depth, and you will see that, too.

(The video is shown)

The word "soliton" is now firmly in the public domain. As you can see, it made it onto Star Trek.

Here is a very simple pendulum lattice. This was the idea of Bob Keolian. It is made of kite string and machine nuts. There is a support bar on top. There are 35 pendulums with that Keolian-type coupling, and the whole thing rests on a loudspeaker that acts as a shaker table to parametrically drive the system.

Here is the instability that Jay asked about. There is Andrés in his early days as an experimentalist. The pendulums have been released as carefully as possible such that their amplitude is all the same. You can see how the energy clumps there. You can put this on a computer even with no nonuniformities, start the motion, put in slight fluctuations, and the energy will clump like this. In our real system, the motion is influenced by nonuniformities, but this is a genuine instability.

What you see here is the normal type of extended standing wave you expect to see in a system like this, but if you drive such that you get response below the lower cutoff frequency, you can get a breather soliton.

There it is. It is a self-localized state, and that can occur anywhere in the lattice. You can move it along. Bob moved the one along in the trough by blowing on it. Here you can use your fingers to move it along.

Here is a highly localized state. You can see that the nut is rotating sporadically. The breather is a very robust state because it can still exist with all these other things going on.

Here is the upper cutoff mode. Usually when you fire this thing up at a frequency that favors this mode, you do not get a pure upper cutoff mode. We were upset that we got these dislocations there. It took us awhile to appreciate that that was where the interesting physics was.

Here is a dislocation. This is a kink, as I described. It is an upper-cutoff softening kink. The description I gave for the lower-cutoff hardening kink is essentially the same as this. There is a domain here of upper-cutoff motion, and a domain on the other side, but the domains do not match up. Due to the nonlinearity of the system, it can accommodate that mismatch with a stable state.

Now you are looking at an angled mirror below the lattice. This is a different pendulum lattice. This is a non-cutoff mode; it is a wavelength-four mode. We discovered that there can be kinks in this mode.

In one region there is a wavelength-four mode, and in the other there is a wavelength-four mode. But they do not match up, there is this spatial phase mismatch. And there is a kink structure there.

You can even have different wavelengths going at the same frequency, a stable juxtaposed state. We call this a domain wall.

Here is a trough. He missed creating a breather the first time; it was not in the basin of attraction. There is the surface wave breather, the first of these states to be discovered. You saw that live in the demonstration. Here is a different perspective of it. The drive amplitude is typically roughly a couple of millimeters. This was about a 10 Hz drive frequency.

For shallow liquid, now, because we have changed the sign of the nonlinear coefficient, this mode is now stable. It is not subject to the Benjamin-Feir instability any more. Here is what looks to be the next mode up, but it is not. It is better described as a kink.

The node is to the left here. This looks like it is a standing wave, a normal standing wave, where the location of the node may be due to nonuniformity. I think this kink-dragging technique was developed by Brad Barber in his early days. You can drag the node from one side of the trough to the other. That is not your normal extended standing wave. It is better to characterize this as a particle. As it is a localized kink, it is a particle that you can move around.

Here is a longitudinal magnet lattice. The coupling is supplied by magnets with like poles facing each other. This is a hardening system. These magnets, as they get closer, cause

the upper-cutoff mode will harden. Its frequency will go up with amplitude.

Here we are now driving parametrically above the upper-cutoff frequency. This is the upper-cutoff hardening breather. It is a localized state.

Now what you are going to see is something really remarkable. We were motivated to do this by the observation of a similar effect in gravity surface waves. We are going to drive the end -- there is the driver. We are now driving in the band. Those traveling waves are the kind of response you would expect to see. We are driving in the band of propagation.

Now we are driving above the upper cutoff frequency. At low drive amplitudes, you get an evanescent wave that exponentially attenuates into the system. Now, as you turn up the amplitude, but do not change the frequency -- we are still driving outside the band -- now look what happens. You see these pulses that move along here? Those are solitons. You can see as each goes by. They do not last very long, because of the dissipation in the system. They are being spontaneously shed from the system. As the energy builds up, there is some instability that comes in that makes these things propagate, then the energy builds up, again.

The next system our student called a working model of the Bohr atom. It is an annular resonator. We are looking at surface waves here. This is the kind of standing wave you expect to see. There are an integral number of wavelengths around the annulus.

You are now looking down on a conical mirror so you can see 360° of the motion. Here is a kink. It is a very violent thing. There is turbulent motion, there are bubbles (you cannot see them right now) being created. As you go around the annulus, there is not an integral number of wavelengths. There is a mismatch, and the system can accommodate that indefinitely. In this case it eventually disappears, as you see. We now have results in which if you wait long enough in certain regimes the kink will spontaneously come back, and then die out sometime later.

Now we are going to look at mode hopping. There are seven wavelengths around the annulus. As the frequency is slowly lowered, the system will down-hop to six wavelengths. There, it did it. That happened quickly. We know there is a connection between this mode hopping and these kink-type structures. I will comment on that in a minute.

Let us now look at the up-hopping. This is really bizarre. Note that there is a traveling wave component in addition to the standing wave. Do you see it converging, a

traveling wave from both directions? Now watch what happens. It has moved a little bit. It is getting bigger. There-an extra wavelength was just added.

We believe this kink is taking part in this mode hopping. The reason we think that is that as we drive harder and harder, this intermediate structure that is involved in the hopping from one mode to the next persists longer and longer and it gets more and more violent. Eventually it persists indefinitely, so it is playing some kind of role in the mode hopping.

Q: Are you going to try to reproduce that and use a high-speed videocamera?

A: Exactly. Yes, we need to take more data on that.

Q: If you have a lattice with 36 degrees of freedom on it, you should get 36 normal modes?

A: There can be many more. This is a good point. There are all kinds of modes. Sorry, everything you learned is just only true in this little tiny region of parameter space we call linear physics.

Q: Is it just because the nonlinearities are so strong that -- I always thought that even if things were nonlinear you would still get 36 normal modes.

A: It is not true.

Q: Does the breather change when you change modes? You showed that you had more than one mode at the same frequency.

A: You can have all kinds of stuff going on. You saw two different wavelengths at the same frequency, because the frequency depends on amplitude. Jay Maynard talked about this. If you have too many curves that bend over like this, at some frequency you can have two waves, standing waves, in our case, at the same frequency. That is not surprising.

What is remarkable here is that they can be juxtaposed. They can be right next to each other and that mode not get destroyed. I would call that a mode of the system, that domain wall. Everything is going at the same frequency, and it can occur anywhere along the lattice.

I just think the bottom line here is that there are many more modes than numbers of degrees of freedom. It is a good question.

DR. ATCHLEY: When you showed the mode hopping in the annular fluid from seven down to six, you were slowly decreasing frequency?

A: Yes, we were very slowly decreasing the frequency. It is quasi-static. Eventually you are going to reach a point where it hops.

DR. ATCHLEY: Suppose you had a linear system and you had two modes and you have a high Q, so what you would expect if you swept frequency was that you would see one mode come and grow and die out, and then another one come and grow, right? But if the Q was low enough, is it possible that those Q curves would overlap? Then, as you changed frequency, what would you see on the linear level?

A: Yes, I think that can happen. Are you imagining like an acoustic case with a direct drive?

DR. ATCHLEY: How would that change of mode occur?

A: An extra node. If you are up-hopping, the node comes in from a boundary in that case, if you think of a tube. It will just continuously come in, there will be no localized structure on the linear level involved there.

If you have a linear system that is periodic and you are now driving it and you ask -- I guess what you are asking is how would it mode hop or how does it go to the next mode on the linear level?

DR. ATCHLEY: Yes, if the modes overlap.

A: And on the linear level, right? I do not know. Good question.

DR. MARSTON: Basically, it would depend on how you are coupling to the modes of the system. I guess the way you would pose the problem is if you could force the linear system with what would amount to a spatial white noise, so you would couple into all spatial modes at this particular frequency, and then you vary the frequency. In other words, you have all spatial wavenumbers excited at this frequency, then you vary the frequency. That would be a way to think about -- actually, it is how to pose the question as to what would happen.

Otherwise, what you would see would be dominated by what the excitation mechanism is, how you are coupling into the drive.

A: If you imagine a linear case, annular, how is it going to go from six wavelengths up to seven? It almost has to be done locally; it has to introduce an extra wavelength somehow. Perhaps this kink is sort of a linear effect, but we find that at higher levels it continuously goes into this nonlinear localized state. There is some kind of connection that we do not know.

Q: On the linear level does it hop or do you have to make it hop?

A: This is parametric drive, which, as Bob points out, is fundamentally nonlinear, right? When you have parametric excitation, nonlinearity saturates it. I do not know if I can actually say that we can go to the linear level.

Q: (Inaudible.)

A: In the steady state on the linear level the system will respond to the frequency of the drive.

Q: There are two modes, but they are just off resonance.

A: But as you continuously go from one mode to the next you have got to add a wavelength. How do you do that on the linear level? That is the question Anthony is asking. How does that happen? I know how it happens when you have boundaries.

Q: Just superimpose two different modes, one with two wavelengths and one with three.

DR. ATCHLEY: And they each respond to their own separate Q curve?

A: Well, we can discuss that. I have trouble seeing that. Let us discuss that.

(Transparency 13)

Let me try to wrap up this lecture. I promised you that we would look analytically at solitons. This is one of the simplest cases to look at. Here is the equation of motion of our model pendulum lattice at the weakly nonlinear level.

If we are after a lower cutoff softening breather or a lower cutoff hardening kink, what is appropriate to do here is to consider an amplitude modulation. On the linear level, imagine this to just be a constant. Then we are describing this mode, the lower cutoff mode if it is a constant.

To get the weakly nonlinear motion out, which is the simplest case, we will allow this function to be slowly varying in space and time. The proper way to do this is with multiple scale analysis, but a lot of times you do not have to do that if you do not want to. I invite all of you to take this, plug it into this equation, make these approximations, and it is an easy matter to arrive at the following equation.

This is an example of a nonlinear Schrödinger equation. It is called that because it has got the imaginary unit times

the first time derivative and there is a term here like the kinetic energy term here in quantum mechanics. This other term is playing the role of a potential. There is a nonlinearity in it and that is the only reason it is called a nonlinear Schrödinger equation.

Q: What are your variables?

A: For convenience, a is the lattice spacing and c is a characteristic speed in the lattice. Just think of this as a coupling term, coupling strength.

The standard nonlinear Schrödinger equation does not have this term, but there is a very close connection between this nonlinear Schrödinger equation and others. In fact, in other, more complicated, systems with a variety of modes you can get coupled nonlinear Schrödinger's equations. When you amplitude-modulate waves on the weakly nonlinear level, you almost always end up with some kind of nonlinear Schrödinger description. This is a very important class of nonlinear wave equations.

If we look at the static case where nothing is changing in time and want to solve for the frequency detuning here, the shift of the square of the frequency, we now see analytically how everything can be going at the same frequency. This is a mathematical representation of what we argued before physically. Everything is going at the same frequency. We are specifying a mode here. The reason everything can go at the same frequency is that there is a balance between the coupling, which is due to the curvature of the profile, and the nonlinearity. This will be negative. This will not, in general, be zero; there will be some frequency detuning. You can see now to get a mode of constant frequency there is this balance between nonlinearity and coupling. In all solitons there is always some kind of balance like that.

(Transparency 14)

You can also look at a slowly varying amplitude modulation of the upper cutoff mode and you get a similar nonlinear Schrödinger equation. Solutions? The simplest solution is static in time, and is a hyperbolic secant. So this is the breather. This delta here is this detuning parameter. That is treated as a free parameter here. This is a one-parameter family.

For the lower cutoff hardening kink, there is a hyperbolic tangent solution. In other words, the cosine wave on the linear level, as the amplitude goes up, continuously transforms into a hyperbolic tangent.

(Transparency 15)

These solutions can be generalized. You can have a

periodic standing wave. The way that is nice to look at this is to imagine you are looking on top of the pendulum lattice. Here is a breather. Here is an anti-breather 180° out of phase with this one. These things are just oscillating back and forth. This is just a nonlinear standing wave.

The function that describes this is this elliptic cn function, so these are called "conoidal" waves, and they have been discovered before. I used to be afraid of elliptic functions, as I think most people are, but they are obvious. What they do is continuously take you -- it is a one-parameter family in these cn functions -- from the linear level of a cosine-type wave through this, which we can think of as a breather-anti-breather chain, to a single breather as that parameter goes from zero to one, or something like that. One limit of the cn function will be a single hyperbolic secant. The cn function goes from cosine to hyperbolic secant.

Here is a kink. There is an anti-kink, a kink 180° out of phase. This turns out to be described by the sn function. We did not find this in the literature, so we were very happy to name these "snoidal" waves. And we have observed these, we have observed snoidal waves.

There is another way you can generalize these. The single states have propagating solutions. Here they are. You can see that this hyperbolic secant is going to move along at some speed v. I am treating that now as the free parameter. There is a range over which you can specify v.

There is some phase modulation inside this breather. After all, if you have a breather like this, I do not know if it really can move one way or the other. It has to have some information on which way to go. What happens here is it looks like this. There is a phase difference along there, given by that, and this thing kind of worms its way along.

This is what you saw in the magnet lattice in the longitudinal case, although the phase variation was not apparent, but you saw those upper-cutoff solitons kicked out, spontaneously kicked out and then move along. That was this upper solution.
And it has been shown that these are solitons.

(Transparency 16)

These self-localized waves are a property of most nonlinear wave systems. When you look at nonlinear waves, there are all kinds of regimes out there that you can go to. There are all kinds of phenomena that you will stumble across, and there is almost always one where you can have some kind of soliton or soliton-like object. It is very typical. Many of these have been discovered now.

But there is still a need for a lot more experimental work. The ones that we looked at, the cutoff states, the non-cutoff kink, and the domain wall, these states like to exist. They are robust, and they are rich in behavior.

Let me give you an example of the richness of behavior here. Consider the soliton shedding that you saw when the end of the system was driven and a soliton was quasi-periodically kicked out. We have now seen that when you take a single soliton and drive it too hard in a certain part of the parameter plane there, it is going to become unstable eventually. One instability that this thing can have is it can go quasi-periodic. What we did was, we then looked off to the side, and guess what we saw coming off this soliton? Other solitons. The soliton was spontaneously shedding other solitons. This kind of thing happens all the time; you cannot keep up with publishing it, especially if you publish at my rate. This is very common with nonlinear phenomena; you are always encountering things that are very remarkable, very rich in behavior.

These things are easily accessible experimentally and they are observable in a very wide variety of systems. We are now looking, in fact, for an acoustic breather in a tube of some medium that acoustically softens. We want to generate these things in, let us say, a radial mode. If it softens, a breather should be spontaneously kicked out and move down the tube. We may look for this in solids, too; I do not know. We are also working on quantifying that mode hopping.

(A brief recess was taken.)

DR. DENARDO: We looked at coherent waves earlier. On the linear level waves travel right through each other. On the nonlinear level, as Steve Garrett has explained, you can get scattering, sum and difference waves, and we will look at this explicitly.

(Transparency 17)

One thing that can happen when you have finite amplitude noise is that a signal will attenuate due to its interaction with the noise. The noise will, through nonlinearities, take energy out of the signal. That is the present subject. This lecture was prepared in collaboration with Andrés Larraza, who is here.

(Transparency 18)

Believe it or not, we are going to start off with a pendulum again. We do not care here about the gravitational nonlinearity. Here what is relevant is the nonlinearity in the dissipation. If you pull this pendulum aside and let it go at large amplitude, in the beginning there will be turbulent flow.

If the dissipation is weak, it is not difficult to show that the amplitude falls off algebraically. If you want to know how to do that, see me later.

In the beginning the amplitude falls off algebraically, but then eventually, at small amplitude, the flow will become laminar. In that case we now have a linear system. The dissipation is proportional to the velocity. In that case, as we all know, there will be exponential relaxation. There is, of course, no abrupt boundary here, this is continuously happening. This part of the amplitude curve is an exponential.

This is near-equilibrium. We are going to call the first part far-off equilibrium. Finally, in equilibrium the pendulum is not motionless (a real pendulum will be nearly motionless), but due to thermal fluctuations there will be $\frac{kT}{2}$ for each degree of freedom in the pendulum, so it will have these relatively small fluctuations.

(Transparency 19)

This exponential relaxation is supposed to be a property of all systems near thermal equilibrium. That is what I am told. Is it true? It is true, certainly, in a lot of cases. Is it true for a shock wave? I was wondering. What happens to an old shock? Steve Garrett showed that a shock wave in the beginning does not attenuate exponentially. What happens eventually?

DR. HAMILTON: Eventually the viscosity offsets the nonlinearity. The crest nonlinearity is going down as about 1/x and then eventually your rise time increases and attenuation kicks in and then it becomes more rounded and exponential.

A: Good. Okay, so this is still true.

An example of exponential relaxation is the so-called Landau-Rumer process. This involves phonons in a dielectric at low temperature. We can imagine that the phonons are in equilibrium due to nonlinear interactions, if you like.

An impressed wave will attenuate due to its interaction with these phonons. Landau and Rumer calculated that. It was also calculated, purely classically, for all frequencies of noise by Westervelt. He arrived at the following attenuation constant. This is the expression.

I will give an indication of how this comes about and we will actually solve for something similar to this in one case. The quantity β here is the nonlinear coefficient. It is dimensionless and of order unity. We will see it explicitly in

a moment. The quantity ρ is the density, c is the speed of sound. This u is the energy density per unit magnitude of the wavenumber. In this expression we are assuming the noise is isotropic.

(Transparency 20)

This Westervelt formula was confirmed for low-frequency noise by Stanton and Beyer in a series of experiments. They took an aquarium and banged on it and put an ultrasonic wave in there and looked at its attenuation due to the noise and confirmed the low-frequency limit of the Westervelt formula.

To be accurate here, I should say that the noise is not necessarily in equilibrium in the Westervelt result. So you might wonder, if we have a signal that is going through some noise and interacting with the noise, then certainly the noise will interact with itself. What will happen, what kind of equilibrium, if there is an equilibrium, will come about due to the nonlinear action of the noise with itself?

In this case, you can show that what happens is equipartition. There is an equal amount of energy in each volume of k space. I called it q here, the wavenumber. The energy-per-unit magnitude of the wavenumber will then go as the square of the wavenumber. That is equipartition: an equal amount of energy in each little cell in wavenumber space. That has an ultraviolet catastrophe as you go to high frequency. It diverges; there is an infinite amount of energy.

We all know how Planck overcame that. It is interesting, and this relates to a discussion last night with Jay Maynard on how you can obtain quantum mechanical results without doing quantum mechanics, or how you do not necessarily have to quantize in order to get quantum mechanical results.

Cabot and Putterman derived the Planck distribution for sound, for phonons, essentially from a classical argument. They postulated two spectra. One was a zero-point spectrum that exists independent of temperature (that is why it is called zero point). Through nonlinear interactions the energy can redistribute itself, and they deleted one channel. They specified that you cannot remove energy from the zero-point spectrum. If you could do that, you could violate the second law of thermodynamics. That is like taking energy at zero degrees and you can get work out of it. With that postulate, they were able to actually derive the Planck spectrum from an essentially classical point of view.

Paul Roberts applied this to light and derived the Planck spectrum. I should comment, when Orson Anderson talks about the mathematical physics department at UCLA, there is no mathematical physics department at UCLA. There is Paul

Roberts.

(Laughter)

I mean, there is a mathematical physics department, and it is Paul Roberts. The Cabot-Putterman calculation for light is horrendous because of the vector-field nature of light.

Before we look into the physics of what is going on, let us ask where might you see absorption of sound by noise.

(Transparency 21)

There is noise in the ocean, so a sound signal will be attenuated due to its interaction with the noise. But this is evidently not a big problem, because people have sent sound halfway around the world. Is there anyone who has not seen this picture? Good, a few people. The purpose here is long-term monitoring of the temperature of the ocean by looking at the speed of sound, sending sound waves through the ocean and observing how their speed changes, in order to monitor a possible rise in temperature of the ocean.

DR. ATCHLEY: You might point out that right-hand line comes from Monterey.

A: That is right.

DR. SACHSE: That is very controversial now, isn't it?

A: You mean the whales?

DR. SACHSE: Yes, you are going to turn the ocean into a discotheque or something.

A: Maybe we should have a discussion about this later, and Bob can lead it. There were some fundamental acoustics mistakes made regarding the sound intensity level of this experiment in the ocean. People thought it would kill whales or make them deaf, and there is just no way.

Dr. Keolian: What is used as the reference level in the ocean is different than in air. Whales make sounds that are louder than the ones that we are talking about.

DR. ATCHLEY: I should point out that even though there is a fundamental error in the calculation of the sound pressure, even though that mistake has been acknowledged, it is still a very politically hot problem. Protesters have been out in front of Woods Hole and all sorts of stuff. It is serious.

A: Yes, we learned the story at the last Acoustical Society meeting, those of us who went to the talk, and maybe we should discuss this later.

There is a natural two-dimensional sound waveguide in the ocean, because the pressure increases with depth and the temperature decreases, and the waves refract. This is called the SOFAR channel. These sounds were sent through the SOFAR channel, but there is noise --

Q: You were making a point about noise. I was just wondering how that changes anything.

A: That is a good question. It becomes essentially a two-dimensional problem in the absorption of noise by sound. I will talk a little bit about that. It is still absorbed.

(Transparency 22)

Let us look just briefly at the ocean. Most of the noise is in the low frequencies. Here is the spectrum from 0.1 Hz to 20 Hz. We can model this with approximate power law. If we are optimistic here and extend that power law down to some cutoff wavelength the size of the earth (it is a cutoff frequency of about a third of a millihertz) we get the total energy density of noise and now if we plug this into the low-frequency limit of Westervelt, we get an attenuation that is proportional to the square of the frequency of the signal. The bulk losses also go, as Dr. Bass mentioned, approximately like the square of the frequency.

At 100 Hz the absorption of sound due to noise is about a tenth of the normal attenuation due to thermal viscous losses, which is not negligible. We do not know if anyone has looked into this. You might be able to turn this around, if it is possible in the environment of the ocean to do any kind of controlled experiment. By measuring the absorption of sound by noise you might be able to use that to probe the low-frequency noise in the ocean. This may become important some day. Some day people may say the low-frequency noise in the ocean is killing our whales, or something, and we need to measure it. It is something to keep in mind, that this phenomenon could be going on there and might be measurable.

Let us look at the physics of what is going on. At the heart of all this, as Steve Garrett mentioned, is the fact that when you have two primary waves that are passing through each other, there are quadratic nonlinearities and these will generate sum and difference waves. Here is a schematic diagram of this. We will explicitly go through a calculation that will complement what Steve did in a moment, but for right now let me remind you that quadratic nonlinearities will give you a sum wave whose wavenumber is the sum of the two wavenumbers (vector sum) and a frequency that is equal to the sum of the frequencies. There is similarly a difference wave.

This happens no matter what the angle of the two primary

waves is. You can ask if a secondary wave, say the sum wave, has the right dispersion law. It is not, in general, true that the frequency will equal c times the magnitude of the wavenumber. When that does not happen, the waves are off resonance.

When the frequency of sum or difference waves does equal ck, we have a resonant condition. The two primary waves are creating sum and difference waves that propagate. Particle physicists say that the particle is "on the mass shell". What do they call a particle that is off the mass shell? A virtual particle.

In acoustics, as Steve pointed out, resonance occurs colinearly. The only way you can get the frequency to equal ck is with co-linear, unidirectional waves. The two waves have to be going in the same direction. You do not get resonance at an angle. What that means, for example, if you have two primary waves that fill all space and they are at some angle, you will generate sum and difference waves. If you can imagine switching off those primary waves at an instant of time throughout all space, the secondary waves should instantly disappear.

On resonance, what happens if you have these two primary waves and you switch them off? You have created propagating waves that carry energy. They are real; they continue to exist.

It is hard to imagine how you could do such an experiment. An alternative is to have one beam of sound going one way and another beam going another way. At a nonzero angle there should be no resonance. There should be sum and difference waves in the interaction region, but because these waves do not propagate, there should be no energy carried away by the sum and difference waves outside the interaction region. This is what I think most theorists think. Is that right, Andrés?

DR. LARRAZA: Some.

A: The problem is, it turns out that this is a big controversy. I just found this out a few weeks ago at the Acoustical Society meeting. People have done experiments where they do measure, let us say, the sum wave. I would like to know -- let me ask Mark, who has done such experiments -- is the frequency of the wave, if we are looking at the sum process, what is the wavenumber?

DR. HAMILTON: What happens is, when you have a nonco-linear interaction, you are forcing a wave at that wave vector, sum or difference, but to that you have got a free wave, because you have got a homogeneous solution in there, and that is what is causing your, say, sum frequency component to oscillate in space.

As you approach this resonance condition that you were talking about, that means your forcing function wave vector is becoming parallel and the same length as your free wave that would ordinarily exist at that frequency, and that is what gives you that singularity with your resonant --

A: Wait a minute. At a resonance here you get an additional wave, you are saying, a real propagating wave whose frequency is the sum?

DR. HAMILTON: That is right.

A: And whose wavenumber is dictated by the dispersion law, I guess, because it is a real wave?

DR. HAMILTON: I should mention we did an experiment that did what you said would not happen. We had a plane wave and a higher-order waveguide mode interact and eliminated one of the primary waves and launched this scattered component, so you can launch these waves from these nonco-linear --

A: Well, it is a controversy. Didn't Junru Wu do an experiment where he carefully had crossed beams and did not detect sum and difference waves? I hope that we can discuss this later, the scattering of sound by sound.

(Transparency)

How is sound absorbed by noise? We will treat this from a perturbative or weakly nonlinear point of view. On the zeroth level we have two primary waves -- consider wave vector q to be a component of the noise. Our wave of interest has wave vector k.

At the first level we get sum and difference waves. In order to calculate how the k wave is affected, we need to go to third order to get back to the original wave vector. This is a purely schematic diagram. This q wave is the same as this q wave, and we want to consider resonant interactions. All waves are co-linear. We look at the sum and difference waves created here. This wave can then scatter off the noise -- let us say this q component. The difference process will then get us back to k. Similarly, there is this other channel that leads back to k.

When you do these calculations for two and three-dimensional noise, they require a fair amount of work. (We are going to go through a one-dimensional calculation, which is much simpler.) When you go to third order, you always find, as you must by energy conservation, that the pressure swing is reduced, and that it is reduced by an amount proportional to the distance the wave travels. You can then infer that this is the leading term in an exponential and derive the exponential

attenuation. More sophisticated treatments will actually derive the fact that it is an exponential. I am just trying to give an indication here of how this comes about when you scatter sound off noise.

(Transparency 23)

We will now go through a calculation of what happens in the one-dimensional case. I call this the unidirectional gas equation. What do you call it, Mark?

DR. HAMILTON: It is a progressive wave equation that Poisson derived. Riemann, too.

A: Well, we are going to call it the unidirectional gas equation, because it is for unidirectional waves all traveling one direction and it is for a gas. Let me quickly make some comments about this. On the linear level the equation describes waves moving to the right, to positive x, at speed c.

On the nonlinear level there are two contributions that go into the nonlinear term. Note that v is the particle velocity. As you crank this up and start to see nonlinear effects, the equation says that the wave moves faster in the case where the particle velocity is in the positive direction.

That is due to two effects, as Steve Garrett explained. It is due to the fact that the temperature is greater, because when the particle velocity is moving to the right for a traveling wave that is also a compression, and the temperature goes up, so the speed of sound goes up. It is also due to the fact that the particle velocity convects the wave with it. The speed is going to be c plus the particle velocity. The two effects are combined in the nonlinear term here. The quantity β is the nonlinear coefficient, which is given by this expression, where γ is the ratio of specific heats.

We are going to rewrite the equation as a linear operator, and put the nonlinear term on the other side. We imagine that the nonlinearity is driving the system, and we are going to go through a perturbation calculation, very much like the one that Steve explained, except we now need to go to third order to see how the wave attenuates. We have our signal, which we imagine to be a small-amplitude impressed wave, and our noise which is finite amplitude. For the moment we will assume that the noise is in steady state (I will comment on that in a little while), and now we just crank through the calculation.

To first order you just find the linear dispersion law. At second order we see that the first order solution acts as a drive; the first-order solution drives the second-order solution. We are interested in only interactive terms now. We are not interested in self-interaction (self-terms) and,

indeed, our signal is small amplitude anyway, so those are negligible. So the relevant piece here is going to be $v_1 \times v_2$. That is third order.

We also must impose boundary conditions. Let us imagine we have a piston at x = 0 and it is moving sinusoidally, generating the signal and the noise.

(Transparency 24)

Note that we are using complex numbers hare. I think Steve used cosines and sines, but as you work in this field, you very quickly get out of that and start using exponentials. They are much easier to deal with, but you have to make sure you recognize the complex conjugate here when you square things. In this way, we find sum and difference waves. The second-order solution is effectively driven by these waves. Because it is driven at those wavenumbers and frequencies, it will respond at those wavenumbers and frequencies, and we get a solution. We can solve for v_2.

Because all waves are moving in the same direction, v_2 is resonantly driven. The differential equation is secular. The solution grows linearly in space, and so is limited in the range of its validity. It holds only as long as its amplitude remains small compared to v_1, after which our perturbation expansion breaks down.

Now third order. Again we pick out the relevant interaction terms. All we care about is getting back to the original signal: wavenumber k and the frequency ω. When you do that, a simplification emerges here. The effect of the noise is completely embodied in this term here, which is the dimensionless energy per unit volume of the noise.

A component of the third-order solution is proportional to x^2, because an x arises at each level of perturbation. There is also a component that is 90° out of phase. This is how our signal is altered. We can see that the dissipative part is quadratic in the distance. Multiplied by the amplitude A is one minus something times x^2.

It was shown by Rudenko and Chirkin that the exact solution is a Gaussian. The signal attenuates as $\exp(-X^2)$. They missed the part that is 90° out of phase here. They missed that, because they employed an energy approach. This serves to change the speed of the wave. The waves move slower due to the noise, by the same amount independent of the frequency.

(Transparency 25)

Let me quickly mention that our perturbation result does not necessarily have to apply to noise. You can think of it as

also applying to two co-linear waves. In that case, it is amazing that there is an exact solution for two waves of finite amplitude, pre-shock, moving in the same direction. It is called the Fenlon solution.

If you look in the case where you have a low-frequency finite-amplitude wave and a high-frequency small-amplitude wave, the pressure falls off as the absolute value of a Bessel function.

(Transparency 26)

Here is an experiment that shows that. This was taken from Mark Hamilton's handbook, and if you are interested in nonlinear acoustics and you do not have a copy of this, I strongly suggest you get a copy of it. Here is the absolute value of the Bessel function. This is the signal attenuating with distance.

At larger distances, restitution occurs. Due to these higher order terms, the net energy is going back into the wave. The initial amplitude can never be completely restored, by energy conservation. We solved for small x. There is our parabola. If you expand that Bessel function for small x, you get our result. So our perturbative analysis has given us this part of the curve. I just wanted to point that out.

(Transparency 27)

Let us now look at this Gaussian. This is a strange result. There is a characteristic relaxation distance. The signal is a fluctuation that we impose, and the system subsequently relaxes to equilibrium.

The relaxation length goes like the wavelength divided by the square root of the total energy density of the noise. The fact that it is the total energy density of the noise is an indication that our assumption regarding frozen noise was correct. Remember what I said about steady-state noise, we did not take those amplitudes to be functions of time? What we are doing there is assuming that a component of the noise lasts long enough for the signal to interact with it. That is sometimes called a frozen noise assumption. You can question this, because the noise is certainly interacting with itself. But the fact that the result depends only upon the total energy of the noise is an indication that it is okay that the noise is cascading to higher frequencies, it is okay that it is evolving. Rudenko has shown that this is exact, even though the noise does evolve. The absorption depends only on the total energy density of the noise. If you are interested in doing experiments, the square root is of the utmost importance, because the dimensionless energy density, as you will see, is typically a small quantity. The fact that it goes like the square root makes an experiment feasible.

Before I get into that, let me just mention that some of the bizarre aspects of this Gaussian. If you look at a signal launched at x = 0, it is predicted to decay as a Gaussian in x. If you launch another signal downstream with the same amplitude that your original signal has at that point, you would think that it would follow along the same curve. Well, it is not going to do that. The system somehow has memory. We do not understand it, but that is what is predicted.

We can ask, if we think of the noise as being in equilibrium, why isn't this exponential relaxation? From the mathematical point of view, that x^2 arises because we get an x for each level of perturbation theory. The x^2 is unavoidable from a mathematical point of view.

What we can infer here is that this system is not in equilibrium. One-dimensional acoustic noise cannot be in equilibrium with itself because this signal or deviation from steady state, is not relaxing exponentially. So this is evidently a far-off-equilibrium system and we are interested in it from that point of view. We are hoping that this system will shed some light on far-off-equilibrium systems. The fact that it is a far-off-equilibrium state is presumably due to the fact that every component of noise is in resonance with every other component.

Q: Why do you assume that the second-order harmonics were small compared to first harmonics compared to the fundamental but yet you are using it even when the amplitude has fallen off halfway, which means half your signal energy is one-fourth?

A: We have an impressed wave. It is going to scatter off the noise and create sum and difference waves.

Q: And then you assumed that all the harmonics that were generated were low amplitude compared with the original wave?

A: In the perturbation scheme that is correct.

Q: But once your initial wave has dropped to half its energy, that means half of the energy is traveling along in the other harmonics.

A: By that point there will be all kinds of sums and differences, right? What we looked at is just when that signal loses a little bit of energy. When you look way downstream when it has lost a lot of energy, by then there are all kinds of sum and difference waves. It is very high order in the perturbation theory.

Or you can imagine applying the perturbation theory over a little bit of time and then starting off all over, again,

applying the perturbation theory again. That is one of the ways of deriving the fact that it is exponential.

DR. MARSTON: If you could put up that Gaussian. Maybe this is what he was driving at.

(Transparency 28)

Perhaps a way to think about the distinction of launching the fundamental signal at this later time and it now has this Gaussian from that given point is that let us go back to the original experiment, where you launch from x = 0. By the time you get to this offset point, there is a lot of energy in these harmonics and you are plotting here only the amplitude at this signal frequency.

When you launch from this offset point, there is none of the energy in these other frequencies yet, so it takes a while to build up the energy in these other frequencies, i.e., there is some propagation distance compared with that buildup, which is why you now have the offset Gaussian behavior for that case.

A: Can't I make the same argument about noise in two and three dimensions, for which case you do get an exponential -- can't I make the same argument you just made for noise in two and three dimensions?

DR. MARSTON: I guess what I am trying to say is that this diagram in the one-dimensional case is not paradoxical from that viewpoint.

A: I think that you have made a very good point. It is the two and three dimensional cases that are paradoxical, not the one-dimensional case!

DR. LARRAZA: The original equation conserves energy. The basic assumption established is that the signal is weak. Linear acoustics cannot describe it any higher. From just that point of view _____. The noise conserves energy all throughout the interaction length.

A: Let me continue.

(Transparency 29)

Andrés got Anthony and me interested in doing an experiment to verify this strange result. If we consider sound in air in a 2-inch diameter tube, it will have a 4 kHz frequency cutoff, so as long as we are below 4 kHz only plane waves will propagate.

The thermoviscous attenuation is due to losses at the walls in this case. (The thermoviscous losses in the boundary layer that dominate those in the bulk.) The attenuation length

for a 2 kHz signal is 20 meters. If we have 150 dB of noise, the 1/e Gaussian length from the Rudenko solution is 5 m. So this should be a big effect. This is the kind of experiment that I like to do.

(Transparency 30)

We made the following apparatus. It has got two compression drivers at one end. It has got seven 10-foot sections of aluminum pipe. In the seventh there is steel wool that is carefully tapered here in its overall density and it absorbs sound very well. Reflections are 40 dB down, so if you hit this with a pressure wave of some amplitude, 1% of that pressure gets reflected. We were careful with construction of the junctions here. We also have reflections of about 1% off them.

This is a long apparatus, so it is attached to a wall in the hallway in the basement of Spanagle Hall. Maybe some of you saw it last Sunday. We are always worried someone is going to hang from it, and try to do a chin-up.

A standard problem with nonlinear waves is that you want to make sure that the nonlinearity is coming from your medium and not from your drivers. Your drivers are oscillators, too, and nonlinear effects can come from there. In the notes you will see that we did a check on what we call cross-modulation. We do not want sum and difference frequencies to come from nonlinear effects in the drivers.

Those sum and difference waves will fall off with distance, but the ones that we care about will grow with distance, because they are being driven by this extended source, these primary waves that fill the tube. We checked that and found good agreement.

(Transparency 31)

Let me get now to the experiment. Here are noise spectra at three different intensity levels. We have put white noise through a brick-wall filter set from 500 Hz to 1 kHz. This is at a microphone near the drivers. Here are the spectra at a microphone in the fifth tube, compared to the microphone in the first tube.

You can see how the noise has scattered into higher frequencies. You actually see, in the beginning, the sum waves of the noise being created. We are getting something above the cutoff frequency there, but it is way down.

When we put the signal in you can see that it is attenuated with distance. This is, again, at a location in the fifth tube. It gets attenuated with distance. This line right there is the amplitude for the greater level of noise; the

signal attenuates more, the greater the noise.

(Transparency 32)

Here are the reduced data. This is the amplitude of the signal as a function of distance. With no noise it is exponential. Here are the data for two levels of noise. The third level was buried in the noise and we have not yet been able to dig it out.

Q: How did you turn on the noise?

A: One of the drivers generates the noise, the other driver generates the signal.

Q: What is the noise source?

A: It is white noise out of a synthesized function generator; the usual digital white noise.

Q: Could you go back to the previous viewgraph?

(Transparency 33)

For any one of them, what are the three aspects?

A: This is a superposition of three sets of data with three noise levels, three different runs with three different noise levels. We just cranked up the noise -- this is low-level noise. What specified this, incidentally, was that noise shocks were starting to occur toward the end of the tube. We could have driven it higher, but in this case that is what specified this high level.

Q: What is that first _____? One position is a function of frequency?

A: Right, this is the microphone near the drivers. This is the microphone far from the drivers. You see how the noise has evolved down the tubes. It has gone into higher frequencies.

DR. HAMILTON: This was of a lot of interest, these types of spectra, to the Air Force for airplane noise from the point of view that as that band of frequencies is increased in intensity, it is pumping all that energy to higher frequencies, but the rate of dissipation has gone up much higher, so your energy in the wave attenuates far faster than one might expect based on linear theory, so you get much more enhanced dissipation even without the shock formation.

A: It is enhanced dissipation due to the fact you are cascading to higher frequencies and those higher frequencies attenuate faster.

DR. HAMILTON: There is also an interesting detail on your lower right figure there.

A: You are right, I should have pointed this out. Here is the sum wave being created by the signal and the noise, and it is reproducing the noise profile remarkably well, as you can see. Thank you.

(Transparency 34)

We are going to conclude this part of the talk with this last graph. Here is the greatest amount of noise so far in which we have been able to observe the attenuation of the signal. Here are the data on a linear-linear plot. This dashed curve is our perturbation theory. That is the thing we calculated earlier. It is the parabolic approximation. The solid curve is Rudenko's Gaussian result. This is the signal with the normal losses subtracted off.

(Transparency 35)

(Transparency 36)

You can see that the theory predicts greater attenuation, because the theory does not yet include the fact that the noise is being attenuated. Andrés is now working on that. That is an interesting problem involving fractional calculus, because the losses here go as the square root of the frequency. These are wall losses. To incorporate those into the theory, you can use something called a fractional one-half derivative in time, because if you operate that on a frequency, you get the square root of the frequency. That is fractional calculus.

(A brief recess was taken.)

DR. DENARDO: I failed to emphasize something. With this Gaussian attenuation, if you remotely detect that wave, in principle you can specify the location and strength to the source, which you cannot do with exponentially attenuating waves. You can not remotely determine where they came from.

(Transparency 38)

For the final part of this talk, we are again going to look at random waves, specifically, how energy redistributes itself among nonlinear interacting random waves. This is the subject of wave turbulence. Again, this was prepared in collaboration with Andrés Larraza.

(Transparency 39)

First some preliminary remarks about normal turbulence. Consider steady flow, where the velocity is not changing in

time at every point. For sufficiently low velocities, laminar flow occurs and you will have smooth stream lines that look like this. This flow has the properties that if you follow two neighboring fluid particles that start off close, they will always remain close or at most slowly diverge. In the flow field the variation of the velocity has one length scale. It is the characteristic distance of the object that is interrupting the flow.

Suppose now that you crank up the velocity. How high you have to go is dictated by a quantity called the Reynolds number, which is a characteristic velocity times a characteristic distance divided by the kinematic viscosity, which is the shear viscosity divided by the density. When that gets to be typically 100 or so instabilities develop in the flow and eventually you will get turbulence.

Turbulence is characterized by strong divergence of neighboring fluid particles. Randomness occurs and the flow field has variations on all length scales. If you zoom in, you will still find variation of the velocity, unlike the laminar case.

(Transparency 40)

This turbulence in its simplest form would be in an incompressible liquid and it is associated with swirling or vortex motion, so we will call it vortex turbulence. Its importance lies in the fact that it is ubiquitous in nature. In fact, if you can take your finger and move at speeds greater than roughly 15 cm/sec, you will generate turbulent flow behind your finger. This can be checked with smoke from, say, a cigar or whatever you want.

There is another type of turbulence, however, and this is turbulence among random waves that are interacting. That is the subject of this part of the talk. Here there is random redistribution of wave energy in a continuum band of wavelengths. It is not as common as vortex turbulence but it can be handled analytically.

Regarding the difficulty in handling vortex turbulence, there is a quote from the famous British hydrodynamicist, Lamb: "I am an old man now, and when I die and go to Heaven there are two matters on which I hope for enlightenment. One is quantum electrodynamics, and the other is the turbulent motion of fluids. And about the former I am really rather optimistic." He was right.

(Transparency 41)

I want to quickly move through the vortex turbulence. It is important to set some ideas here. Vortex turbulence is described by the Navier-Stokes equation, which is f = ma for a

fluid. You have seen this equation several times before during the summer school.

If we take the curl, the equation can be put in the following form. The curl of velocity is called the vorticity. This equation describes how the vorticity changes in time. The changes are driven by two terms here: a nonlinear term and a dissipation term. This characteristic ratio of nonlinearity to dissipation is the Reynolds number. When we have high Reynolds number, we get turbulence.

(Transparency 42)

Here are some data. This is tidal flow in a channel near Vancouver. What they call the wavenumber here is 2π divided by the characteristic length scale. What is plotted along here is energy density multiplied by the wavenumber to the five-thirds power. That the data is approximately constant means that the energy density is approximately proportional to the wavenumber to the minus five-thirds power.

I have written on the graph the Reynolds number at various length scales. Here is a Reynolds number of 10^8, which is extremely large compared to many other systems. This is why they did this experiment, to get to these very high Reynolds numbers.

There is some kind of scaling law here. As you can see, there is a simple power law relationship. As you go to smaller wavelenghts, dissipation eventually comes in, and the spectrum falls off more quickly. Down here the dissipation and the nonlinearity are roughly comparable.

Kolmogorov and Obukhov were able to derive, based on very simple arguments, the scaling law here. We are going to go through that now. I want to point out that other people have worked on other parts of the spectrum. Here, dissipation is playing a fundamental role in the motion. Heisenberg worked on this part of the spectrum. That has got to tell you something. He has a theory and there are a number of other theories of what is going on there.

(Transparency 43)

Here is the Kolmogorov idea. It is very simple. We have a system here that is being driven. This is an off-equilibrium system. There is some external agent dumping energy into the system. The idea is that the energy is being dumped in at some very large length scale. In the case of vortex turbulence, you are stirring a fluid only on large length scales.

There is then a regime of length scales in which you have high Reynolds number where the system is neither being forced

nor is dissipation playing any essential role. It is dominated by the nonlinearities. This is called the inertial regime. As you look at smaller and smaller length scales, you will eventually get to a point where dissipation plays a role. This is called the dissipation regime.

The effect is that energy is being dumped into the system at large length scales, it is being taken out at small length scales. In between, there is a cascading of energy to smaller and smaller length scales.

Q: Bruce, when you talk about flow of fluid or waves of fluid --

A: This is a swirling vortex-type motion. Imagine the liquid to be incompressible.

Q: How do you get a wave?

A: There are no waves here: You look at the energy at different length scales. I looked into this in this experiment. They had effectively a hot wire anemometer and they looked at the frequency spectrum of the velocity. From that, at each frequency they got a length scale, and they got an amplitude. That is all they did.

Q: They must have made a measurement at more than one point.

A: They did not. It sounds like there are some assumptions there, right?

DR. LARRAZA: If you know the _____ velocity and you know the power spectrum of the turbulent flow, you can get a length scale.

A: But here they were establishing the power law, right? They are not assuming the power law.

DR. LARRAZA: No, they measured the frequency.

A: That is actually all they did. I was surprised, too.

Q: (Inaudible.)

(Transparency 44)

A: They looked at the RMS average of the velocity at each frequency. That is the spectrum of the velocity of the flow. From the frequency and velocity, they got an amplitude. From the velocity you can get a kinetic energy. The average in time is presumably the same as a spatial average. There is a homogeneity assumption here. So they can get the energy density at that length scale. I think that is all they did.

Q: Which length scale?

A: The length scale corresponding to the velocity. They got a velocity of a certain frequency. That corresponds to a certain length scale which is the frequency divided by the velocity.

Q: Has this experiment been repeated?

A: That is a good question. I do not know the answer. Does anybody know the answer? This was done in 1962.

Q: The island does not exist any more. They did it in the Bay of Fundy at the wake of an island where the tide comes in and since the study was done, the island --

DR. GARRETT: But this is quoted a lot in the atmosphere and turbulence studies and the same Kolmogorov spectrum is --

A: The problem is that you are limited by how high a Reynolds number you can get. That is why they were particularly attracted to this location. It happens in the atmosphere, certainly.

DR. MARSTON: Put that one, again.

(Transparency 45)

In the atmosphere, I do not know if they have the range of wavenumbers, but that has been measured by various methods, but there also is a peak at high wavenumbers, i.e., small spatial scales, which is due to another effect, and that poses problems, like propagation of light to the atmosphere.

A: It is important to set the Kolmogorov idea here. We will do the same thing in wave turbulence.

(Transparency 46)

There is an assumption that there is an inertial regime where all that is important are nonlinearities. The system is not being forced and the dissipation is negligible. There is also an assumption of stationarity; you have to average over some time scale, but things are not changing over a longer time scale. The spectrum is not slowly moving around. It is also assumed to be homogeneous.

Most importantly, and this is an assumption in the Kolmogorov case, there is a local cascading of energy. This is local in k space. This means that energy of vortex on one length scale drives vortices on the next length scale down, and so on. A vortex like this does not drive tiny little vortices. It is local, the energy transfer is from one length scale to

the next length scale. I know that is very hand-waving, but we will later see mathematically how this happens.

(Transparency 47)

Here is the Kolmogorov argument on one transparency. We let E be the energy per unit volume at some length scale, L_n. We will scale the lengths geometrically, so the next one down is, let us say, half the length. The next is then half of that.

The wavenumber, as I said, is 2π divided by the characteristic length. The velocity at the length scale must go as the square root of the energy divided by the density. This is the only quantity that has dimensions of velocity. Similarly, the rollover time, which is the time required for energy to cascade from one length scale to the next one down, has to go as the length scale divided by the velocity. There are no other quantities around with dimensions that will give the dimensions of time.

What remains is simply to note that by the stationarity and the homogeneity assumptions, the rate at which energy goes from one length scale to the next must be the same rate at which energy goes from that length scale to the next one and so on. We will call that Q. It must be independent of the length scale.
Now all you do is make use of the expression for Q and solve for the energy density per unit wavenumber as a function of wavenumber. The scaling law is found to be that energy density is proportional to the minus 5/3 power of the wavenumber. This is a remarkably simple argument. It is based essentially on dimensional analysis.

I will not be covering some transparencies, because we are very far behind.

(Transparency 48)

Now we are going to look at wave turbulence. Instead of vortex-type motion, we are going to look at waves in a continuum band of frequencies or wavenumbers. I should take the word "nonlinear" out of there, because all wave systems are nonlinear. Light in vacuum is nonlinear, if you have energy densities of 10's of kiloelectron volts per cubic fermi.

(Laughter)

You laugh, but I am looking into this possibility. There are people who believe light may scatter off light in some astrophysical processes where the energy density can be very high.

Wave turbulence is not as common as vortex turbulence, but it apparently appears in nature.

(Transparency 49)

Here is a famous graph for surface waves on the ocean in a stormy sea. This is a log-log plot. The ordinate is proportional to the energy density and the abscissa is proportional to the frequency. A simple analysis shows that in a stormy sea the energy input is predominantly at these low frequencies here. If you have wind of high enough speed, you can show this, and that is commented on in the notes.

Similarly, it can be shown that for amplitudes all along this higher frequency range inertial motion will occur. In other words, the dissipation is not important; it only comes in at much higher frequencies There is thus an inertial regime, where the forcing is negligible and the dissipation is negligible. The fact that power laws are observed there suggests that there may be some kind of turbulence going on, and that some kind of scaling argument may apply.

(Transparency 50)

At the root of what is happening in this transfer of energy from these low frequencies to these high frequencies is resonant wave-wave interactions, which we discussed earlier. In acoustics I remind you that this is a co-linear process. The dispersion law for gravity waves is $\omega^2 = gk$, where g is the acceleration due to gravity. These are dispersive waves. Furthermore, there cannot be a resonant three-wave process in this case.

The fact that gravity waves have this dispersion law can be arrived at by a simple Rayleigh energy argument, where you estimate the potential and kinetic energies and set the two equal for a standing wave. This approach is in the notes (I will not
go over it here).

Why is there no three-wave resonant process? Why can't you have two gravity waves and find an angle for which you get resonance, where energy is dumped into a third wave. There is an interesting geometry argument that proves this. If you are interested please seek out me or Andrés for the explanation.

For a dispersion law like that for gravity waves, where ω as a function of k bends over, you cannot have a three-wave resonant process. You cannot have two waves resonantly creating sum and difference waves. You must have an additional wave.
Capillary waves have a dispersion law that bends upward, and you can have three-wave process for that case. Acoustics (no

dispersion) is on the boundary, and that has caused some interesting problems and calculations.

(Transparency 51)

Let us now do the Kolmogorov argument for gravity waves. The ideas are the same as for vortex turbulence, except now we have an actual wavenumber. We again have successive length scales, and we now have a dispersion law. The characteristic surface height at the nth scale we will call ζ. To see how the energy per unit area goes like, think of potential energy. This is mgh, where h is the amplitude of the wave, but m, if you think of a wave, is proportional to h, so the energy per unit area goes like $\rho g h^2$.

We now require the rate at which energy cascades from one length scale to the next. However, unlike the Kolmogorov case, we have to put some physics in because now we have a dimensionless number. The wavenumber times the amplitude is dimensionless. That is essentially the steepness of the wave, and is the Mach number for gravity waves.

We can insert this to any power and still preserve our dimensions. The point here is that the number of interacting waves specifies this power. Because we have a four-wave process, because three waves are resonantly dumping energy into one, the cascading rate will be proportional to the product of three energies. We therefore take m = 2 in this expression.

You now carry through as in the Kolmogorov case. The result is a power law spectrum that is ω^{-4}, and that is exactly the power law for the ocean wave data that I showed you. This is suggestive that wave turbulence is going on there. But the spectrum rolls over to ω^{-5}, and that it is still not in the dissipation regime. So what is happening here?

(Transparency 52)

This is interesting. If you go along the ω^{-4} spectrum to higher frequencies and look at how the Mach number is changing, you find that the Mach number is actually increasing with frequency. So as energy cascades, as it goes to higher and higher frequency, the amplitude is falling off, the energy is falling off, but the frequency is going up. The frequency goes up faster than the energy falls off. The waves thus become steeper, they become more nonlinear. This is counterintuitive. You would think that you would be approaching linear waves, but the waves get steeper for this spectrum.

What is going to eventually happen is that the weakly nonlinear description, in our case a four-wave process, is not going to describe what is going on. In fact, higher and higher

order wave processes will be brought in, in this resonant transfer of energy.

In some spectra it goes the other way; the waves become less nonlinear. In the gravity wave case, they happen to become more nonlinear. The scaling law will break down, because it is only weakly nonlinear. Andrés and Seth were the first ones to handle this, and they did it in an amazingly simple way. It is evidently now believed by the Russians. Andrés tells me one of them has just published this and now the rest are all referring to that paper.

(Transparency 53)

You go back to the exponent m and you argue that as the system becomes more and more nonlinear, and it is not in the weakly nonlinear limit anymore, higher and higher order wave processes will come in, so we really should solve the spectrum for an arbitrary m and let m go to infinity.

(Transparency 54)

Believe it or not, when you do that, you find an ω^{-5} law. Moreover, the spectrum does not depend upon the input energy. It is saturated, because the Mach number is 1. You find a spectrum which is in agreement with what is observed. A remarkable fact.

So what happens here in the gravity wave case is that energy is pumped in at low frequency, and as long as the energy is not too high, the cascading of energy will give a spectrum that falls off as ω^{-4}. At higher frequencies, it will roll over to ω^{-5}, where it becomes saturated. This is universal in the sense that all ω^{-4} curves will eventually hit the single ω^{-5} curve.

This is beautifully simple physics and you can apply it to any wave system. You derive the scaling laws, basically by dimensional analysis with some essential physics, the number of waves involved. You can apply it to acoustics, in which case you get ω^{-3} for the weak spectrum. The saturated spectrum turns out to be ω^{-1}, or 1/f noise. It has been pointed out it may be that some forms of 1/f noise might be the result of wave turbulence. That has not been established yet, but it is certainly provocative.

(Transparency 55)

How is the theory done rigorously? First let me say that I do not do it rigorously. This is beyond me. What Seth and Andrés and other people do is to develop a kinetic equation for how the energy in each wavenumber is changed due to interaction

with other waves. What is found is a Boltzmann-type equation for how the energy at each wavenumber changes with time.

In this final part of the talk I want to look into that. There are some very interesting predictions here. Let me point out that from this weakly nonlinear approach of generating this Boltzmann equation, they are able to derive that the energy cascading is local. This follows from the analysis. It is an assumption in the Kolmogorov case. I think that people working in vortex turbulence now believe that there are cases where it is not true that the cascading of energy is not local. It is one of the complications. But it can be shown to be true in wave turbulence.

(Transparency 56)

How do we describe what is going on with this sea of waves? We begin by considering a wave in a slowly varying background. The slowly varying background can arise in two ways. It could be due to some nonuniformity of the medium or it could be due to other waves.

Because the background is slowly varying, the quantity that should be considered here is that which is invariant. As the wave moves along, what is invariant as it goes through these slow changes is the wave action, which is the energy density divided by the frequency. This is the energy per unit k per unit volume at a time divided by the frequency, the frequency at that k and at that point in space. This is the wave action.

(Transparency 57)

I want to say a few simple words about that. Let me talk a little bit about adiabatic invariance. It plays a fundamental role in a lot of physics. It has to do with the fact, as I am sure all of you know, that when you make a slow change in the parameter of a system, there will be a quantity that is constant. It is called the adiabatic invariant.

A simple system to exhibit this is a pendulum whose length you slowly change. On the linear level -- this is all linear here -- you can show that the oscillation energy divided by the frequency is a constant. In fact, you can show it is true for any system that undergoes simple harmonic motion. The energy divided by the frequency is a constant. This also happens to be true for some nonharmonic oscillators, for example, a ball banging back and forth elastically between two walls.

(Transparency 58)

To carry on this further, let us apply it to a resonator. Consider a closed-open resonator (you will see why in a minute) and imagine we have an undamped mode in this resonator and has

uniform cross section initially. Suppose that we, as an external agent, slowly constrict the open end of the tube slowly compared to the oscillation frequency. The flow is approximately incompressible here, due to the pressure-release boundary condition. So there will be a dc Bernoulli pressure drop. When we constrict the end, because of this pressure drop, the system does work on us; the system gives up energy to us. That means that the energy of the wave must go down. If its energy goes down, what must its frequency do? E over ω is a constant, so the frequency must drop. There are other ways to argue this, but this is one way of looking at it. You can look at constrictions or enlargements anywhere along the resonator, but now, because flow is in general compressible, you have to realize that there is another contribution to the dc pressure. Whereas the Bernoulli term is the negative of the kinetic energy density, you have to add the potential energy density to get the general radiation pressure in this case.

Then you can quantitatively calculate how the frequency shifts, and that turns out to be equivalent to what Rayleigh did long ago, of course.

Here is a resonator. It has got a step in it; it is necking down toward the open end. This is not a slapper, this is a clapper, so let me not try to slap it. If I did slap here, the frequency would be lower than the frequency if I slapped it on the other end. To show you that, instead of slapping it, I am going to set it on a smooth surface. We thus have a closed-open resonator that necks down toward the open end. Now I will compare the frequencies. This resonator has an intermediate configuration where there is essentially no change in cross-sectional area. It just makes a jog halfway. That will give us a third frequency. This was carefully designed and machined to give musical intervals.

(Demonstration)

There is a simple way to appreciate what is going on here. In the case that necks down toward the open end, the system approaches a Helmholtz resonator, and so the frequency is going to drop. In the case that necks up, the quarter-wavelength fundamental of an initially straight resonator approaches a quarter-wavelength mode of a closed-open resonator of half the length, so the frequency increases.

This is a way of appreciating the effect, but calculating is not so obvious. Adiabatic invariance offers a way of appreciating it and being able to calculate it. To consider a constriction (or enlargement) at any location of the resonator, the full dc pressure, or radiation pressure, must be used. This turns out to be the potential energy density minus the kinetic energy density.

(Transparency 59)

Let us now go back to wave turbulence. As we follow a wave in a slowly varying background, the energy divided by the frequency, called the action, will be constant. The rate-of-change is therefore zero. Because we have to follow along in real space, there will be a convective derivative in addition to the partial time derivative at a point. Specifically, we have got to move with the group velocity to see how the action is changing in space. But the wave can also evolve with respect to wavenumber, so we must include a convective derivative in k space. This turns out to be the expression for that. In summary, if we move with the wave in phase space, the change in action will be zero.

The next step is to recognize that there can be collisions or resonant wave-wave interactions. These are incorporated into the term on the right side. Due to the resonant transfer of energy among the waves, energy can be dumped into this wavenumber. The collision integral looks like this for the three-wave case. It always has this form regardless of the system. The particular system will dictate what the coefficients are. You can see the generation of sum and difference waves here. The four-wave case is similar, but the collision integral is cubic in the wave action n.

(Transparency 60)

We end up with a Boltzmann-type equation. We can ask for homogeneous and stationary solutions. In this case, there is no net energy being dumped by collisions into any wavenumber. There are two types of solutions. There is the closed system, which is the equilibrium state. And there is the open system, in which the system is driven far off equilibrium by some input of energy away from frequencies described by the Boltzmann equation.

The first solution is equipartition of energy, and thus contains the ultraviolet catastrophe. The second solution typically turns out to be power law spectra. When you go through this analysis for gravity waves, you end up with a ω^{-4} spectrum. When you go through this for acoustics, you find $\omega^{-3/2}$. These are precisely the spectra found by our Kolmogorov argument in the weakly nonlinear case.

(Transparency 61)

Next, we consider small departures from the far off equilibrium steady state. Here is the steady state value of the wave action as a function of wavenumber. We look at a small departure from this in k and in space and time. Something else has to be done here, however. We must include the fact that a wave can be altered not only by resonant wave interactions but also by moving in the background of another

wave. Let me give an example of that.

(Transparency 62)

This is from Bascom's Waves and Beaches and he describes how to do this demonstration. You have a wave tank, and put some tracer particles in and generate a wave. For a low-amplitude deep gravity wave, the motion of the tracer particles is circular, as I am sure a lot of you know. At higher amplitudes you find that the paths do not close on themselves. Superimposed on the circular motion is some drift velocity in the direction of the wave. On the nonlinear level of this wave there is thus a dc drift, first calculated, I think, by Stokes. Rayleigh has an explanation for this that we do not yet understand.

(Transparency 63)

If we now generate a wave of frequency ω moving against the original wave, what is its frequency? Its frequency will not be ω, due to the presence of the dc motion of the original wave; it will be Doppler-shifted down.

So waves can interact, and not necessarily resonantly. The frequency of the wave can change due to the presence of other waves. This is not a collisional process; there is no sum-and-difference resonant generation of waves. We can include this effect by writing the frequency of the wave as the bare frequency plus some integral. In the case of gravity waves, the bare frequency is $(gk)^{1/2}$. The integral depends upon the energy of all the other waves in the system, and you have to work this out for your system.

(Transparency 64)

You put all this together to get a linearized Boltzmann equation, linearized about the steady state. Here is the frequency shift term. This equation describes how fluctuations of the steady state relax back to the steady state. Remarkably, there are collective solutions. By collective we mean that the solution has a wavenumber and a frequency that is the same for all k's in the sea of waves here, although the amplitude out front here will be a function of k. These are called collective modes because they involve motion at all the length scales, at all the frequencies. To me, the existence of these solutions is a miracle, because, on the linear level, a disturbance travels at the group velocity, and the group velocity is not the same from one frequency to the next. Somehow the combination of the convective term and the collisionless frequency shift term makes everything go at the same speed, for a family of wavelengths and frequencies with some amplitude function. It is a collective mode.

Let me conclude by saying there are two regimes here, because there are now two time scales of the motion. There is the frequency of the collective mode and there is this collision time, the time it takes δn to relax to zero due to collisions. When the frequency of the wave is fast compared to the collision time, we have the collisionless limit, and the collective mode is called zero sound. The other extreme, when the frequency of the collective mode is small compared to the collision time, involves compressions and rarefactions of the turbulent sea of waves and, due to the collisions of the waves, a local steady state is reached. This is analogous to in superfluid helium. In that case, the phonons are responsible for the temperature. Second sound in liquid helium is a temperature wave corresponding to compressions and rarefactions of the gas of phonons. A compression represents a hotter temperature, and a rarefaction a cooler temperature.

(Transparency 65)

It turns out that the zero sound for three-wave systems is overdamped. For gravity waves, for example, these collective modes should be underdamped. Andrés has worked out predictions. They remain to be observed, however. This is in zero sound.

There is a problem with the second sound, as nice as the picture is, about the local steady state arising due to collisions and its analogy with the normal type of second sound. The problem is that on top of this there is the cascading energy in all length scales, and that causes the wave to damp. Its degree of damping is unknown at present.

(Transparency 66)

I want to conclude by talking briefly about two experiments that are in progress. One is at UCLA experiment, and another is here at NPS. There are currently no published controlled experiments on wave turbulence. This is exciting, because, as I said, any wave system is a candidate for wave turbulence.

(Transparency 67)

The UCLA group consists of Bill Wright, Raffi Budak, and Seth Putterman. They are investigating capillary waves. They shake a dish of liquid up and down. They have a wave-height probe here, and they measure the spectrum. Here are two spectra on a log-log plot.

(Transparency 68)

They get a big response at half the frequency because it is a parametric drive. This peaky spectrum of harmonics is due

to nonlinearities. As they drive harder, what happens is that the valleys start to fill in, and the eventual result is a broad-band spectrum that goes like ω^{-4}. This may sound great, because that is the right spectrum for gravity waves, but these are not gravity waves, these are capillary waves. That is the wrong spectrum. Some other investigators have trouble understanding this.

DR. PUTTERMAN: Sorry, it is the right spectrum. Nature chose it.

(Laughter)

A: Right, you are right.

DR. PUTTERMAN: It is the right spectrum. It is the wrong theory.

A: You are right.

This is the predicted spectrum based on the inertial regime. It should go like this, and can be derived with our simple scaling argument. One might think here, because the observed spectrum is steeper, that dissipation is coming into play. I do not know what progress they have made on that kind of understanding. One would think that perhaps this is not an inertial regime; maybe dissipation is dragging down the spectrum.

(Transparency 69)

It has been pointed out that just because you have a broad band spectrum, just because energy is in a band of frequencies in a wave system, does not necessarily mean that you have energy in a band of wavelengths? This is an interesting idea of, I believe, Andrés'. It is conceivable that energy is in just a few modes but that the motion is chaotic. The energy may not be in a band of wavelengths. This potential problem prompted the UCLA group to look at the system spatially. They have come up with, as I understand it, a beautiful light extinction method, and perhaps Seth or someone else can talk about this later. The higher the amplitude is, the less light will pass through the liquid. They are able to generate pictures and not have to worry about caustics when obtaining amplitudes.

(Transparency 70)

Here are some images. This one is at a lower drive level. What you see is the actual image of the surface, and here is the two-dimensional Fourier transform. The bright spots represent the k vectors being excited. You can see there are a number of rings.

(Transparency 71)

When they drive harder, here is what the surface looks like. It appears that there is motion on smaller length scales. The rings are fainter now; they are starting to get washed out. This is an indication that energy is on different length scales and it is not just chaotic motion of one or several length scales.

Q: Did they experimentally measure length scales from the _____, taking FFT and getting --

DR. PUTTERMAN: There is a largest length scale which comes out, which is a few millimeters, it comes out at the FFT.

(Transparency 72)

A: Here is Bob Keolian's experiment, which is in the basement of Spanagle Hall at the Naval Postgraduate School. The tank is 20 m long, about 1 m wide, and 1.5 m deep. The sea of waves is generated by wind. It is pumped in by these fans and the top is covered to keep the wind confined. This generates a broad-band spectrum of waves. A paddle at the same end as the wind source puts in a burst of energy.

(Transparency 73)

Here is the spectrum as a function of time for a probe down some distance in the channel here, for a burst with no wind. The colors represent the intensity at certain frequencies. Here is the burst going by the probe. Here is the reflection of that burst off the end of the tank. The wind is turned on and here is what is seen.

(Transparency 74)

This is the wind spectrum represented here. The burst is launched, and here it is moving by the probe. Some time later you see that there is a notch in the data. Furthermore, although it does not show up well on here at all, this notch is over a broad band of frequencies. It is a candidate for a collective mode.

Bob is busy collecting reduced data on the speed of the notch. There is proof that it does not travel at the same speed as the burst. It is traveling at approximately the group velocity of the spectral peak, right around where the theory predicts it ought to travel if it is a collective mode. This is very exciting.

(Transparency 75)

In the meantime, there has been a recent publication by

Chu, Long, and Phillips, and this is the famous Phillips who has the ocean wave spectrum named after him. They see a similar effect. Here is their burst, and they see a notch. Actually, their notch is followed by greater energy and it does move at a different speed. It is in the neighborhood of the right speed for being a collective mood. These are oceanographers, and it is my understanding that they refuse to even sit down and think about the wave turbulence theory. They would just rather not worry about it.

(Transparency 76)

To conclude, we have looked at three aspects of nonlinear waves. Let me remind you that this behavior is completely absent on the linear level; it only exists at only the nonlinear level. There are a lot of predictions here that need to be explored experimentally in all of these areas, and it is certainly true that people will experimentally stumble onto entirely new phenomena when it comes to nonlinear waves. This is a very exciting area in which to work. Thank you.

THIS PAGE HAS BEEN PURPOSELY LEFT BLANK.

THERMOACOUSTIC ENGINES AND REFRIGERATORS

DR. SWIFT: Logan warned me that everybody was up real late last night, so I am going to go only, maximum, five equations at a time between demonstrations and photographs, or something like that. The first hour or so is going to be to try to give you a broad, qualitative, intuitive understanding of thermoacoustics, and I hope that will carry us through until coffee.

Let us start by reviewing a couple of things just to give us a common vocabulary here.

(Transparency) (1)

In your thermodynamics classes you saw heat engines and you probably remember there were two kinds of heat engines, the prime mover and the refrigerator. The prime mover is a heat engine where heat Q_H, comes out of a high-temperature reservoir, temperature T_H, and heat Q_C goes into a low-temperature heat sink T_C, and some work is produced in the process.

The first and second laws of thermodynamics put a bound on the efficiency of such a machine. The efficiency is what you want, work production, divided by what you have to spend to get it, the heat flow. The first law is just conservation of energy: What goes in, Q_H, must equal what comes out, the sum of W and Q_C.

The second law of thermodynamics says that the entropy of the universe always increases or, at best, stays the same. This universe is made of three pieces: this reservoir, the machine here, and this reservoir. The machine is in steady state, so its entropy is not changing in time. So for the second law you just write down that the entropy increase of this reservoir has to be greater than or equal to the entropy decrease of this reservoir.

If you work on these two with a little bit of algebra, eliminating Q_C, you can form this ratio, with efficiency being bounded by this ratio of temperatures, and that, you remember, is the Carnot efficiency.

For a refrigerator you just reverse the directions of all the energy flows, so you are putting work into a refrigerator to lift heat from a low temperature to a high temperature, and if you do the same kind of an analysis here you end up with its efficiency, which is usually called coefficient of performance, the heat of refrigeration divided by the work you have to put in to get it being bounded by another temperature ratio.

(Transparency) (2)

Next part of the review: wave equations. A simple wave equation in pressure. Here is time and here is space. If you let your oscillating pressure P be a complex amplitude times e to the i omega t, then you can turn it into just a spatial wave equation in the complex amplitude P_1 of x. So this is an ordinary kind of wave equation without any extra frills at all.

It was not until about 20 years ago that anybody figured out how to do a wave equation in a duct with a temperature gradient along the direction of acoustic propagation with the duct being small enough that viscous and thermal wall effects were important. That wave equation is Rott's wave equation.

We are not going to look at the details of this right now but just so you can see it, it looks like a wave equation. This 1 times P_1 is that thing. This a squared over omega squared d by dx dp_1/dx, that is this term. All the rest of this stuff is corrections for viscous effects at the walls and thermal effects at the walls and then the big correction here for the temperature gradient along the wall, which can actually give you either gain or loss in this wave equation.

(Transparency) (3)

If you put these two pictures together, the heat engine picture and Rott's wave equation, you get what we call thermoacoustic engines. There are two types, just like there are two types of heat engines. There is the prime mover and the refrigerator, the prime mover taking heat from a high-temperature source, converting some of it to work; the refrigerator accepting work to lift heat out of a low-temperature heat sink.

Then the acoustics part is mostly just some kind of a resonator. You need heat exchangers to take the heat in from the outside and distribute it across the resonator -- here and here. Then there is a thing we call a stack -- you will see a lot about stacks in the next hour -- that is kind of a passive heat exchanger that temporarily stores heat for half the cycle of the sound wave and then returns it to the other half cycle.

(Demonstrations)

So we have two demonstrations, a prime mover demonstration and a refrigerator demonstration. We will start with the prime mover demonstration, which is over here. Between my two hands is a quarter of a wavelength in this little pipe. It is open on this end and closed on this end, oriented just the way the figure is.

Here is our heat source. Here is our heat sink. As soon

as the temperature gradient gets high enough, this thing will start to oscillate at a frequency that makes that a quarter of a wavelength. There it goes, right on cue. If you cork it up, it will stop, and then it will start again.

Here is something from Bruce's talk. Remember, Bruce talked about taking a resonator like this and changing the cross section near the mouth and he predicted which way the frequency would change and he showed us that nice little clapping demonstration? If I stick this in here and effectively neck down the area, which way is it going to go? There we go.

Another thing that is kind of interesting, when I take my finger off, it is not that there is sound in there now and you cannot hear it because my finger is in the way; it is not oscillating right now. It will start up. It will take a tiny fraction of a second to start up. Do you hear the little bit of a delay?

This is one of these things where it starts with some little fluctuation, picks up a fluctuation from the room and amplifies it and it gets louder and louder over the course of a 10th of a second or so. Okay, that is the prime mover demonstration.

(Transparency) (4)

Tom Hofler brought a refrigerator demonstration and he is going to tell us about that. Essentially, his refrigerator demonstration is oriented like this. It has got a piston driving it here, and there is a volume here, and so it is about a quarter of a wavelength from there to there.

(Demonstration)

DR. HOFLER: This is a demonstration refrigerator that we got running a couple of days ago, so there are no measured data to speak of, except for some temperature measurements. The driver was designed and built by Jay Adeff and Lt. Kevin Mode, a student at NPS, and the demonstration refrigerator resonator itself was designed by Lt. Brent Brooks. We got this thing assembled, as I said, just a few days ago.

Basically, it is the same kind of design that Greg was describing. It was used in my old thesis refrigerator. However, a lot of these thermoacoustic refrigerators end up being primarily laboratory curiosities and connected to water-recirculating chillers and so on. Jay was just saying, in fact, this might be the first thermoacoustic refrigerator that is truly CFC-free, because it does not use a water recirculator with a CFC chiller inside it.

The object here was to produce frost in a small portable

demonstration unit and to have it air cooled so it was totally self-contained and more or less autonomous and could run indefinitely without external chillers and apparatus.

As far as the basic design, the acoustic resonator is up here in this region. It is really roughly a quarter-wavelength design, where the vibrating piston -- down here -- is the pressure antinode, which is like a rigid termination even though, in fact, it is vibrating. The velocity antinode or pressure node is about in the middle of the small tube, about this position here somewhere.

If you want to, you can think of this dead volume, flow volume, here as simulating an open end to the tube, so it is like Greg's demonstration, where you have a rigid termination and an open end. Of course, there is a diameter change. You go from the vibrating piston to the hot heat exchanger here and there is a dead volume in-between, then the hot heat exchanger, then the thermoacoustic stack, made out of Mylar film, then another copper heat exchanger, which is the cold heat exchanger, and then we reduced the diameter down to a smaller diameter tube, again, like Bruce Denardo's resonators and what Greg was just describing with his pointer inside the tube.

One of the original motivations for doing that is because this particular choice of diameters minimizes the acoustic losses on the cold end of the refrigerator. With this design, the diameter is actually slightly smaller than optimum, because the smaller diameter you make that small-diameter tube, the shorter the resonator gets.

An actual quarter-wavelength resonator for this -- a full quarter wavelength would probably be out to here, but this design is quite a bit shorter because of the diameter change. So that helps make the thing more compact and exposes less cold surface area to external heat loads.

Down here we have a magnet moving-coil electrodynamic driver. The coil is connected to a reducer cone and a vibrating piston with a flexible bellows. Up here, next to the hot heat exchanger -- you cannot really see it very well in this drawing -- are about 20 or 30 aluminum fins. This outer structure is an air duct and there is a muffin fan down here which sucks air down through the fins, around the outside, and then back down out through the bottom. So the hot end and the driver are all air-cooled and will sit here doing what it is doing right now indefinitely.

In terms of performance, as I said, we do not really have any measured stuff, but the computer model that we used to optimize the performance we optimized for a 40° temperature span. We measured about a 38° temperature span, with the hot end maintained at 23° C. with the water recirculator, and the cold end went down to about -14° C.

With the air flow in here, the hot end is now up to about 27° C. The cold end is about -- actually, it is -10° C., which is 5° colder than we have ever gotten before. I was going to say it is a 32° temperature span, but I guess it is 36° or 37° now. So it is pretty much right on our ballpark of a 42° temperature span.

The COPs we calculate from the model are on the order of one. The acoustic losses on the cold end, depending on whether they all end up as a load on the cold end of a stack or whether some of them go out into the room, you can get COPs that are either a little less than one or a little greater than one. That gives you a COP relative to Carnot on the order of 20%, which is not any higher than what we have done before, although the COP is a bit higher.

The driver was really designed for our cryocooler experiments. With the low acoustic impedances for cryocoolers, we got driver efficiencies, electric to acoustic, that are in the range of 50%; although with this acoustic load the impedance is higher and the efficiency is probably about half that, 25% electric to acoustic driver efficiency.

That is about all we know about it right now. There is another thesis student who is going to quantify the performance of this thing and package the electronics a little bit better for demonstration purposes.

Q: Could you describe the resonator in a little more detail?

(Transparency) (4)

A: The drawing here is a little busy, but basically you have a constant bore tube here with your vibrating driver piston. As I said, it is the lowest velocity point in the resonator. It is not a true rigid termination, because it is vibrating, but it is the velocity minimum in the system -- right down here.

So you have an empty tube region here and then you have a stack, which to first order is more or less transparent to the sound wave as far as the velocities and so on. Then here you reduce the diameter. Normally, you would think that that would increase the acoustic viscous losses because the velocity goes up here but, in fact, it decreases the total acoustic losses, because the surface area of the inside of the tube is going down rapidly, even though the velocity is going up.

As you decrease that diameter, two things are happening. One is, the diameter is decreasing, but the other thing that is happening is the length gets much shorter for a fixed-frequency operation. So the two combine to reduce the surface area more

quickly than the velocity squared increases.

Q: Is the shape of this curved portion important?

A: I think it is, but I do not have any real measurements to prove it. In other words, I have not really tried to make bad ones to see how bad they are compared to what I think are good ones. I mean, that is something a good scientist should do, but when you get in a mode where you are trying to produce working things as quickly as possible, you tend not to do things that you think might work poorly.

The answer is I do not really know. But if you buy standard plumbing fittings or any steady-flow hardware for high-velocity steady flows, they all have radii. So I think if you have a hard corner, you are going to shed vortices off of any internal hard corners.

Q: Are the details of the shape here important in determining the acoustic losses and resonant frequency?

A: Sure, every detail affects, in some sense, the reactive sound field, but in a low-amplitude calculation, whether it is a hard corner or a curve does not really matter; it is only when you are talking about the high amplitudes in the nonlinear regime where you might be shedding vortices or generating turbulence or who knows what.

Let me say one other thing I forgot to say. In terms of the reactive resonant frequency, you can also consider this thing to be a Helmholtz resonator, where you have one gas volume or gas spring down here and another gas volume and gas spring up here and a small-diameter tube as your gas mass for the Helmholtz resonator. So it is really one mass and two gas springs on either end of the mass.

That is a little bit more crude than thinking of it as an open end there. Actually, both of those calculations will get you correct ballpark frequencies.

Q: What gas are you using?

A: It is pure helium. It is 90 psi or about 6.2 bar. The pressure amplitude is about 5% of the mean peak pressure swings.

Q: The fan is not the sound coming out?

A: The swooshing sound is the fan, but the [whistle] is the driver, the acoustic frequency.

Q: What is the maximum displacement in your piston?

A: I think it is about 15,000th of an inch, which is

about a third of a millimeter.

Q: Can you tell us more details about the driver?

A: This is a hi-fi driver and supposedly a fairly powerful 2-inch version. It has a Ferrofluid in the magnet gap, which is there in order to cool it --at least one function of it. In the early days I left the Ferrofluid in there, figuring that would give me more power and more current handling. More recently we have been thinking, well, hey, we have helium in there and helium has five times the thermal conductivity of air, maybe we can just get rid of the Ferrofluid, so, in fact, on this one we stripped everything apart, and removed the Ferrofluid. It has our own custom suspension in there, although we are using the same original coil. Really, we have just as good a current-handling capability without the Ferrofluid and with helium in there as we did with the Ferrofluid.

The other problem is, at the higher amplitudes, the Ferrofluid will spray out of the gap. It will handle about 4 amps continuous.

Q: Are there any precursors to failure?

A: We have an accelerometer on the driver and often you can see distortion from various mechanisms. One of the common problems we have is we used to use the original fabric suspension, which was not robust enough and it would go out of alignment and the coil would start dragging on the magnet and then, after awhile, you would start getting intermittent shorts through the insulation of the coil, and you can see that if you monitor the current waveform with the accelerometer, so you can see that kind of nasty stuff.

But what we have done most recently is we can measure the coil temperature <u>in situ</u> while it is running using a combined DC resistance and AC drive method. So we think at the 4-amp level we are about 80° C, which is about as high as we want to take the coil in temperature.

Q: What is the cooling capacity that you get out of the calculation and, also, how might your relative position and performance compare to a pretty good normal refrigerator?

A: Basically we were calculating for a cooling heat load of about 12 W and the acoustic power going in is about 10 W, so that gives you a COP of about 1.2, although it might be a little less than that, depending on how you deal with your acoustic losses.

In terms of conventional refrigerators, again, as I said in a previous talk, I think if you would compare the kinds of

COPs we are getting now, acoustic to thermal efficiencies, and you compare that to CFC refrigerators that were being built in the 1950s or something, you might find that they are pretty competitive or comparable, but if you compare them to CFC refrigerators being built today, we are probably a factor of two worse in COP compared to a modern refrigerator.

DR. SWIFT: Continuing getting oriented here in thermoacoustic things, what are some of the orders of magnitude involved?

(Transparency) (3)

The lengths tend to be a quarter or a half of the wavelength of the frequency you choose to operate at. That is of the order of a meter. In these things it is more like 10 cm and we have some that are more like 10 m, but, ballpark, a meter for the length of a thermoacoustic apparatus.

The other key dimension is the separation between the little fins in the heat exchangers. That is a few thermal penetration depths at the frequency you have chosen to operate at, and that is typically a millimeter -- it might be a tenth of a millimeter or a few millimeters, but, ballpark, think of millimeters here and much bigger distances, like meters, here.

Pressure amplitudes, as Tom mentioned, are typically 5% of mean pressure or, at Los Alamos, usually 10%, so that means that the acoustic Mach number out here, the velocity divided by the sound speed is about 5% or 10%. The displacements here are of the order of 10% of the distance from here to here. The Reynolds numbers out here can be 10,000 or even higher quite easily, so this is not acoustics in the usual sense that we think about it; this is high-pressure amplitude, big displacement stuff.

(Transparency) (5)

I want to give you a very quick survey of some of the applications that are being attempted. The green dots here are places where kind of academic-style thermoacoustics research is going on, stuff that is supported by the government, for the most part, exploratory kinds of things.

The red dots are places where usually corporate money but sometimes very directed government money is being spent to try to develop real applications of thermoacoustics. Ten years or so ago there would have been just two or three green dots and no red dots, so this field is really taking off.

I am going to spend about two minutes on each of maybe four of the red dots, just to give you an idea of the breadth of applications that people are contemplating and to show you some more real hardware, because that is fun.

(Transparency) not available for proceedings

This is a refrigerator at Ford Motor Company. It is essentially a quarter wavelength from here to here and there is a driver down here in the bottom with some kind of a big stereo woofer, and the two heat exchangers are in this plane and in this plane, with the stack in-between. This thing has, ballpark, 100 W of cooling power and 20° C temperature difference. And I do not know why Ford is interested in cooling.

(Transparency) (6)

Here is another refrigerator, built here at Monterey, essentially half a wavelength from here to here, and I say essentially because these cross-sectional area changes shift the resonant frequency a little bit, as Tom explained a few minutes ago.

Two stacks, two sets of heat exchangers, hot here and here, cold here and here, and two loudspeakers to drive them. This one was 250 W and maybe 30° C. This one is also notable because it is the first time that a new temperature-sensing technique, called the "lost bud method," was used.

(Laughter)

Those of us who grew up in cold climates learned as children that things like porch railings and door handles -- well, never mind.

Q: What is the phase relation between the two speakers?

A: Out of phase, because it is a half wavelength between the two ends, so the pressure amplitudes on the two ends are out of phase.

(Transparency) Not available for proceedings

These guys are also doing something that we learned in winter climates not to do. They are doing an experiment in January in the middle of New York state. Starting here and extending off into the distance somewhere there is the acoustic resonator of this engine. This is intended as a sonar projector, so it is an engine, a heat-driven sound source, and they are getting it ready to sink into a very cold lake in upstate New York. I think they learned that you cannot get a soldering iron to even melt solder in a 30 mph wind in the winter.

(Laughter)

(Transparency) Not available for proceedings

This is the biggest thermoacoustic engine I know of. It is in Scotland. It is about 5.5 m tall and a meter and a half in diameter. It is a prime mover. The stack is up here. It uses air at atmospheric pressure as working substance and burns natural gas inside the resonator as a heat source. And I honestly do not know why they are doing this.

(Transparency) Not available for proceedings

In a Los Alamos-NIST collaboration we made a heat-driven cryocooler comprised of a thermoacoustic engine and a refrigerator called a pulse-tube refrigerator, which is not quite like a thermoacoustic refrigerator, but it also has no moving parts.

The thermoacoustic engine here is under this fuzzy insulation here. The resonator starts behind Bob, about there, and goes along here and up to the ceiling. It would not fit in the room, it was so long, so we had to bend it into a hoop here, and it ends kind of over here. So that thermoacoustic engine turned heat into acoustic power, which was then coupled directly to this pulse-tube refrigerator, which is this little U-shaped thing down here. This part of the U would get to 90° K and had about 5 W of cooling power at 120° K.

At Los Alamos we are spending most of our time now working on follow-ons to this project. In collaboration with Tektronix and especially with Charles Jin we are trying to make this very tiny for cooling electronic components, and in collaboration with a company called Cryenco in Denver we are trying to turn it into a combustion-powered natural gas liquefier for use in the natural gas industry.

(Transparency) (7)

This is a cartoon that would give you an idea of the size of thermoacoustic engine, resonator extending from here to here, about a meter in diameter, with thermoacoustic engines in here and in here, drive- and pulse-tube refrigerators to liquefy natural gas.

(Transparency) Not available for proceedings

One more application that is under way, this guy is in Dhaka, Bangladesh. He is bending a copper bar into a circle here to make the hoop around a heat exchanger of a thermoacoustic engine. They want to make a kerosene-fired refrigerator for keeping medicines cold in rural medical centers where there is not any electricity. I guess the standard way that is done now is to haul ice from the cities, and that is a little bit trickier than hauling kerosene. Anyway, this will be an air-working-substance device, just bronze castings and welded steel, pretty simple geometry.

I hope that one or two of these things pays off in the long run. There is a lot of activity right now and if none of them works we are going to be really embarrassed.

(Transparency) (8)

Advantages and disadvantages of thermoacoustics relative to other things in widespread use: They are not as efficient as steam turbines, internal-combustion engines, Rankine-cycle CFC refrigerators, and they tend to be a little big bigger for a given power-handling capacity, but they are easy to make, they are cheap, they should be really reliable, because there are no moving parts. These are the kinds of tradeoffs that are going to be played in the real-world applications.

Why are they not so efficient? That is the worst disadvantage of thermoacoustic engines. About a third of Carnot is typical of current design and, as Tom mentioned, about a fifth of Carnot may be typical of things that have been built so far.

The losses come basically from four places. No. 3 here is the easiest one to describe. It is just stupid conduction of heat from hot to cold. You have got material in there, gas and solid, and it conducts heat. That is a loss mechanism. You have viscous dissipation on all the surface area in that stack. You have what we call intrinsic losses, again, on all the surface area of the stack within a thermal penetration depth, and we will talk about that a lot more in about 10 minutes. You have the same kind of losses, these first two, happening on the heat exchangers and the resonator.

For instance, if you have a pressure vessel here, you have conduction of heat from hot to cold through the pressure vessel losses.

(Transparency) (9)

I want to move on to a little bit more quantitative look at some of the processes that go on in thermoacoustics. I think you will find that a lot of this connects back to earlier lectures in the week.

First, we want to look at viscosity, not because it is the most important thing in thermoacoustics but because it is the easiest thing to understand. Suppose you have oscillatory flow in a gas next to a wall like this and you want to understand what is going on.

The equation of motion -- we have seen this about 10 times this week -- Newton's law, mass times acceleration is pressure gradient force and viscous forces. Boundary condition, there is no velocity on the surface. In acoustics we often make the

acoustic approximation that each variable has a mean value and then a small oscillating part. If you substitute this into this equation and keep only the lowest order terms, you get that the velocity has a pressure gradient piece and a viscous piece.

And now this is a differential equation for this velocity u_1 in the x-direction as a function of y, the perpendicular direction. You solve it and you get this thing. You can stare at this all day and it is hard to see what is going on. This is a complex function of y and somehow the imaginary part tells you something about which part of the velocity moves out of phase with the pressure gradient, and so on, and I can never keep track of those things.

(Transparencies of computer program "viscous")

I hope this will help give you a better feeling for just how this works. This is supposed to be a resonator, half a wavelength from here to here. There is the standing wave in that resonator. When the pressure peaks, you will notice the velocity passes through zero. It is like water sloshing in a bathtub, right?

I mean, we all know what standing waves are, but it still helps to stare at this thing for about a minute to really appreciate the phasing and so on. When the velocity is maximum or minimum is the instant of time when the pressure is going through zero.

Mark some places in the gas, put some little smoke lines in there or something, and watch how they go back and forth. Here, when the pressure goes up, the density goes up, they get closer together. The pressure goes down, the density goes down. When it is moving to the left the fastest, that is when the velocity is the most negative.

Are you with me so far? Now, we are supposed to be understanding viscous effects here. I have got this gas moving along the wall here. We know that cannot happen. We know that it is locked to the wall right at the surface of the wall.

The next frame is going to be a magnified view of what is going on right here, where the velocity is the highest and where the pressure gradient is swinging back and forth like that. This is the solution to that differential equation that I wrote down on the bottom of that last viewgraph.

That is what the motion of the gas is like in the viscous boundary layer in acoustics. It does not really matter whether it is a traveling wave or a standing wave. That is what that functional form looks like. Each one of these tic marks is one viscous penetration depth and out here, several penetration

depths away, it is just flat motion in response to the pressure gradient. The boundary condition pulls it to zero here and in-between that is what it does.

(Transparency) (9)

The characteristic distance over which these viscous effects take place is the square root of twice the viscosity over the frequency and the density. Typical distance is a millimeter or a fraction of a millimeter, stuff like that.

(Transparency) (10)

Power. This is left over from the last viewgraph. There is the picture. There is the equation. Viscosity and gradients in velocity dissipate acoustic power. Per unit volume that dissipation in this simple flow in one direction x, with dependence on space in the other direction y, is just this.

If you work that out, you will find that that is largest on the surface and it falls off exponentially. I like to think of these resonators and thermoacoustic engines as having these little shaded surfaces. I imagine a viewgraph where the hard metal wall is drawn black and then you shade in red, about a viscous penetration depth's worth of stuff on each surface everywhere, and that is lossy. That is why you shade it in red, because red is a bad color. So much for viscosity.

(Transparency) (11)

It is a little harder to think about the temperature, although that is more important for thermoacoustics. Suppose now we have a solid wall here and we have oscillating flow kind of toward the wall, so we have the pressure going up and down here a lot, and we want to ask how does it work out when the heat capacity of the wall imposes an isothermal boundary condition on this gas motion that is trying to be adiabatic?

We have seen this before this week, an equation of heat transfer. The change in entropy in unit volume is just heat transfer, thermal conductivity. There is our boundary condition and we need a little thermodynamic identity relating changes in entropy to changes in temperature and pressure. We play the usual acoustics trick of linearizing to get this differential equation. Does that look like a differential equation we saw about 10 minutes ago? Yes. Same solution, the old $e^{-(1 + i)}$ and so on.

This time the characteristic length is the thermal penetration depth, which, instead of having viscosity, has thermal conductivity and specific heat.

(Transparencies of computer program "thermal")

Let us go back to our standing wave, again. There it is, velocity and pressure. I have pulled a fast one on you now, I have switched from velocity to temperature, so there are the adiabatic temperature oscillations in the standing wave moving along in phase with the pressure oscillations.

Now we are going to magnify that region to see the solution to the equation that we had on the last viewgraph. Here is the pressure moving up and down as a function of time. So this is just the left end of that cosine wave that we looked at in the last frame. There is the temperature moving up and down.

Out here the temperature goes up and down adiabatically and within a few thermal penetration depths of the wall -- that tic mark right there is one thermal penetration depth away from the wall -- the wall exerts its thermal influence. That function is just the same function as we had for viscous losses.

I did not do this for viscosity, because I think dissipation in viscosity is kind of obvious. It is like friction. Its rubbing turns mechanical power into heat. But the source of the dissipation in the thermal boundary layer is maybe not quite so obvious.

What I want you to do here is imagine one of these lines of gas -- think of it as a piston. We can ask how much work does that piston do on the gas in front of it? That is the power that this penetration depth is soaking up, absorbing from the wave.

If you are going to compute the power that a piston does on the gas in front of it, what do you do? You compute the integral of pdV. The distance between the front of this piston and the wall is the volume and the pressure is this other line, so if we look at the intersection between those two lines, that is p and V at any instant of time, so the area of that loop is the integral of pdV.

The gas over here is not a resilient spring, it is a hysteretic spring, and the area of that loop is a measure of the hysteresis. Why is that area there? Let me pause it at some extreme of its motion here. I will pause it up at the top.

At that instant of time the pressure is high, the temperature is like this, the temperature of the gas here is hotter than the wall, so heat is flowing into the wall from the gas. The gas there is thermally contracting a little bit, so its volume gets a little smaller. That is why the little point there that is tracing out the ellipse moves a little bit to the

left right there. It is because of the thermal contraction that this gas is experiencing. That is the source of the dissipation in the thermal boundary layer.

One more thing you can ask is what would the area of this ellipse look like if you made the piston here or here or here or out here. It turns out that that is what the loop area as a function of position looks like. So by the time you are about three, four, or five thermal penetration depths out from the wall you have gotten to the infinite limit, and the gas out here, farther to the right, is resilient, it does not dissipate any more power.

Q: If all that energy loss is going into the wall, shouldn't the work done be the same right up to the wall? I am talking the gap between that piston and the next one along.

If you draw a piston right at the last line you have drawn, there is no energy being absorbed in the gap between that piston and the piston before, in the piston you have got drawn in right now.

A: Oh, yes, the gas in here is not a perfectly resilient spring, so it absorbs a little bit of work. If it were really thermally isolated from everything else, it would be a resilient spring, but a little bit of heat oozes out of it when it is at high pressure and oozes into it at low pressure.

DR. PUTTERMAN: Is the non-zero power dissipation a manifestation of nonlinear hydrodynamics?

A: I do not think so, any more so than are any of the powers in acoustics which are quadratic in the amplitude.

DR. PUTTERMAN: They are quadratic quantities.

A: Yes, but intensity is a quadratic quantity.

DR. PUTTERMAN: Right, but in this case how do you know that there are not intrinsic nonlinear properties which, treated in a linear fashion, would give, also, contributions to the area?

A: Well, maybe there are, but this kind of an analysis seems to account for what happens pretty well. As when Bruce was talking yesterday about absorption of sound by sound, you know, there are processes like that, sure, and then they add to this one. I think this is just, then, one process that absorbs sound. Fair enough?

Q: Over here you have not included any nonlinear terms, right?

A: No. Now, you remember the equations I had on the

viewgraph? It was just linearized stuff.

DR. LEVY: Is the loop area the amount of work done? So isn't there work done at large distances, according to your diagram?

A: I guess I am not following you.

DR. LEVY: Is there any energy loss at large distances? Wouldn't the area enclosed in one of those loops equal the amount of work done, or energy loss, or something?

A: Yes, so at large distances, like if you were out here, if you drew the loop out here, it would be leaning more steeply than the one that is drawn is, but it would not have any more area. That piston does not do any more work on what is ahead of it than that piston does.

DR. LEVY: I agree with that, but the amount of work done is equal to the energy lost, is it not? I would not say the spring is resilient far away from the wall.

A: Yes, if you look all the way to the wall from here, it is not resilient, but if you look at just the gas between this white line and this white line, then you would be looking at the difference in the areas of little ellipses for those two pistons.

DR. LEVY: Oh, so this is the net work. I did not understand that. I thought it was work at each point.

A: Good point. That is a little confusing there. Thank you.

(Transparency) (12)

This is kind of a connection, actually, back to what Hank talked about on the first day. He gave us a review of bulk attenuation of sound with the attenuation constant having a viscous term and a thermal conduction term. You can rewrite this in terms of these penetration depths just doing a little bit of algebra like this. Bulk attenuation goes like one over the wavelength, viscous and thermal penetration depth squared over wavelength squared.

The wall losses that we have just talked about, the viscous and thermal wall losses, they do not go like penetration depth over wavelength but instead penetration depth over the radius of the duct that the sound is propagating down. It is really the same physics. It is whether heat can diffuse in an appreciable fraction of the wavelength cause the gas to and become a nonresilient spring, or whether heat can diffuse in an appreciable fraction of the distance to the wall in your pipe and cause the gas to become a non-resilient spring.

If you look here, you can decide in a given problem whether bulk attenuation is going to be important or wall attenuation is going to be important. It is going to be a ratio of these two. I guess, if the penetration depth times the radius of the duct divided by wavelength squared is big, then bulk attenuation is more important than wall attenuation.

DR. HAMILTON: Why is your tubal attenuation proportional to omega and not root omega?

A: There are more frequencies buried in here, remember? Frequencies buried in here, too.

Q: That equation for the attenuation of pipes is only below cutoff, right, for a plane wave? You are not including any transverse?

A: No.

(Transparency) (13)

Time to get back to thermoacoustics instead of these simple little problems. Back to our simplest problem. We are going to neglect viscosity for a minute now, because it is too complicated to do everything at once. Here is an oscillating velocity in the gas next to a plate. Now we have a temperature gradient along the plate that might be non-zero, and y and x are in those directions.

If we look at that equation of heat transfer, as we did just a couple of viewgraphs ago, and linearize it, we get that the temperature oscillates at any point out here because of the sum of three reasons. One is, if the pressure is oscillating, the temperature is oscillating; that is just adiabatic temperature oscillation.

A second reason that the temperature oscillates is because there is a temperature gradient and the gas is moving past, and if you have a temperature gradient, say, hot over here and cold over here and something that is moving past you, what you see at that point is something that moves up and then down and then up and then down, and that is an oscillating temperature.

The third reason that the temperature oscillates is thermal contact to the wall. So if you try to solve this differential equation for T_1 as a function of y, you get this same factor here, again, and you get a coefficient of magnitude out in front of it that has P_1 and u_1 and it has got the square root of minus one and it has the temperature gradient.

The only point I want to make here is that you can get anything you want out of a function like this, because you have

so much complexity over here, with a function that is doing all kinds of things. By picking the phasing between u_1 and P_1 to be standing wave or traveling wave you can get different phases in this factor, and if you happen to pick standing waves, so that iu_1 has the same phase as P_1, by picking the value of the temperature gradient you can get this to be positive or negative. So you can get anything you want.

(Transparencies of computer program "fancy")

Let us look at what some of those things are, some of those different options that you can get in thermoacoustics. Now we are going to put in a stack and in thermoacoustics you want the stack at a place where there is both oscillating pressure and oscillating velocity in the wave. You do not want it all the way at a velocity node and you do not want it all the way to pressure node, either, so you put it in here.

Now we are going to take a look at what happens in a little -- this is excessive, I know.

(Laughter)

Down here, this is supposed to be a magnified view of what is inside this little yellow space. These are two of those plates in the stack.

We are going to look at this little blob of gas right here that is moving back and forth. There is the magnified view of that blob of gas going back and forth. You can see it is changing size, because the density changes. When the pressure is high, the density is high, and when the pressure is low, the density is low.

Now let us play our pressure-volume game, again, look at the pressure and volume of that blob of gas as a function of time. Not quite a resilient spring. Again, there is an area to this loop that represents work that the sound wave is doing on that blob of gas and you can understand the origin of that work by looking at all these things at once.

(Laughter)

You really do have to think about all these things at once to understand it. This blue line represents the temperature gradient in the plate, so I have no temperature gradient here, zero temperature gradient. The position of this little speck is the position of the blob of gas at any instant of time, so you see they move together back and forth like this.

The vertical axis, the white, is the temperature of that blob of gas, and it is mostly just experiencing adiabatic temperature oscillations, making it go up and then down in

temperature as the pressure goes up and then down. See, when the pressure is high, the temperature is high?

But when the gas is hotter than the plate, like it is right there, there is enough thermal contact between that blob of gas and the plate for heat to flow out of the gas into the plate. So the gas cools a little bit, thermal contraction lets it shrink, so the volume can move a little bit to the left here, and that is what opens up the area of the loop here.

With respect to heat, you can see that what this blob of gas does. It lifts some heat out of the plate in this region and then dumps that heat onto the plate over at this region. That is kind of like a refrigeration action. This blob is typical of all the blobs of gas in there and they are all moving heat from right to left.

The net effect is to lift heat from this end and carry it over to this end and that is still the case if you let there be a little bit of a temperature gradient along the plate, with colder temperatures being over here on the right and hotter temperatures being over here on the left.

The signs of the heat flow and the sign of the work is still the same and so now we have a thermoacoustic refrigerator with this being the cold heat exchanger, where heat is allowed to come in from an external load, and all the little blobs of gas pick up heat, move it over a little bit and dump it, so it is like a bucket brigade carrying heat from here -- dup, dup, dup -- and dumping it up here, and each blob of gas absorbs a little bit of work from the sound wave.

Q: When you put the temperature gradient in -- well, before, without the temperature gradient, your equation of state assumed constant temperatures as expanding or contracting. Now, with the gradient, as you trace out each point on that curve, you are putting the temperature as given by that gradient as the temperature in your equation of state, is that correct?

A: Yes and no. The mean density tracks the temperature gradient, of course: $\rho_m = P_m/RT_m$. The oscillating part of the density goes like $\rho_1 = \rho_m (P_1/P_m - T_1/T_m)$, with ρ_m and T_m tracking the temperature gradient. But T_1 has three parts, as we saw a few minutes ago.

Now let us let this temperature gradient be a little bit steeper. Here is a steeper temperature gradient that basically lines up with the adiabatic oscillations that the gas experiences. So you see how we have matched them? We have matched the adiabatic temperature swings that the gas experiences and the displacement that the standing wave carries the gas along. We have matched that with the temperature

gradient of the wall. Well, nothing happens. It is pretty boring. That is what we call the critical temperature gradient in thermoacoustics jargon.

At steeper temperature gradients, now the sign of everything has changed. You see now, when the gas is, let us say, at that extreme of its motion at low pressure and large volume, it has been adiabatically cooled from its average temperature, but not as cold as the wall next to it is. So it is hotter than the wall and heat flows from the gas to the wall, so that it cools and it shrinks a little bit there.

Now the sign of the loop here is opposite of what it was in the previous frames. The line integral of pdv now is positive, so that gas is doing work on its surroundings, pumping energy into the sound wave, and this is a thermoacoustic engine. The critical temperature gradient separates steeper temperature gradients, which give you engines or prime movers, from shallower temperature gradients, which give you refrigerators.

The rest of the lecture is going to be a complete drag, so let us get coffee.

(A brief recess was taken.)

DR. SWIFT: I would like the newcomers we have to introduce themselves.

(Introductions)

Question here?

Q: In the sample where you were showing compression and rarefaction and you were showing net work done and it was far away from the boundary, the net work going into that gas parcel, what was the loss mechanism? It would seem, if this were just a spring oscillator with no damping, you should not be doing any net work.

A: The spring, though, is hysteretic if it is a gas that has thermal contact to its surroundings.

Q: So it was thermal -- I do not want to say radiation.

A: Relaxation. It is a hard thing to get your mind around. It really is.

Now you kind of understand that all the action in thermoacoustics takes place within a thermal penetration depth of the nearest solid surface and that is why we put these things in called stacks. You want a lot of power, you want a lot of area in your sound wave, but area with big distance to the wall does not do you any good, so you need to fill that

whole area up with metal plates or something, spaced a couple or three or four thermal penetration depths.

(Transparency) Not available for proceedings

Here is an example of such a thing the way we are building them these days. This is a stack in our collaboration with Tektronix. It is a spiral of metal with little ribs on the ends to maintain the spacing. The intersection of each rib with each turn of the spiral is actually welded to hold it all together. This is thin sheet metal. The spacing is five times the thickness of the sheet metal. The gaps are a few thermal penetration depths and the length is about a 20th or a 30th of a wavelength.

Q: You mentioned thermal conduction along x was a problem in these things. Why does it have to be continuous sheet metal all the way along the length of it?

A: It does not have to be. That is something to improve. But actually, the conduction of the gas is pretty bad, too. If you pick a poor thermal conductivity solid and it is a small fraction of the cross-section, then the helium gas starts to eat you alive, too. But usually the pressure vessel is the worst conductor of heat of all.

Q: Was the center of that stack blocked?

A: Just blocked, yes, a little spool to start the spiral.

(Transparency) (14)

In Steve Garrett's talk this would have been called "The Royal Road." This is the hard-core way to do it. Now that we have some intuitive grasp, let us look at how bad it is to really calculate thermoacoustic effects. This viewgraph illustrates the total of the starting point of Rott's calculations.

You need some gas properties so that you can change differentials of one variable into others. This is entropy, temperature, and pressure. You need to know the heat capacity of your gas, its density, its ratio of specific heats, and its sound speed.

You need the continuity equation, equation of motion, the heat-transfer equation, and then you -- "you" being Rott -- linearize all that by substituting in each variable with a mean value and a small oscillating part. You can integrate with respect to y and z perpendicular to the sound wave analytically in a couple of geometries if you are Rott, and a few more if you are Pat Arnott.

We now have expressions for circular pores, parallel plate

geometry, triangles, squares, rectangles, pins, a lot of useful geometries. In the x direction, the direction of sound propagation, you are usually stuck with a numerical integration. We will get back to that later.

(Transparency) (15)

If you do all those things on the last viewgraph, you end up with Rott's wave equation that I showed you real early this morning. The wave equation part, there is p and here are two derivatives of p. Then there are these viscous f_v and thermal f_k corrections that have to do with those boundary layer processes that we looked at in the animations.

In a large duct, the f_k is just thermal penetration times the perimeter of the duct divided by the area, with a

one minus i, just to make it confusing. Phasing in these thermoacoustic things is always the hardest part. Parallel plate geometry, with a spacing of $2y_0$, you get a hyperbolic tangent. You know, it is just a function and it can be computed.

The temperature gradient term here is the one that can give you gain in the wave equation for a prime mover and it gives you loss in a refrigerator. You can sort of see how dp_1/dx is like velocity, that is the way it goes in acoustics -- so this is temperature gradient times velocity. This is not a bad equation to integrate numerically as a function of x. That is the wave equation.

(Transparency) (16)

There is another important piece of this in thermoacoustics. It is how the energy flows. If you have had a hydrodynamics class like out of Landau's and Lifshitz's book, you know that enthalpy is the "right" energy to use in fluid mechanics. That is because if you take the velocity times this quantity, that is the energy-flux density, which is made up of a kinetic energy piece, a half rho squared, plus an enthalpy piece, which you can break up further into internal energy and pressure.

With critical people like Seth in the back of the room, I am sort of nervous about saying this, maybe it is not right, but I always think about energy flow in fluids, if there is a blob of fluid moving along, changing shape, et cetera, crossing some plane, I always think about it like this. The kinetic energy density and internal energy are just convected along and then the pressure term has something to do with the work that is done on the fluid ahead of that surface by the fluid behind the surface.

In acoustics we always look at energy flows to second order, that is, the product of first-order things, and if you take this to second order -- well, you throw away the kinetic, because that is v x v^2, that is v^3, so you can throw that away. You end up with just the time-averaged product of enthalpy-density times velocity.

Then if you use the definition of differentials of enthalpy, you get an energy-flux density in acoustic problems that has two pieces. One is just the intensity that we always work with in acoustics, but that is the piece we always work with in acoustics, because we are usually thinking of adiabatic sound waves where the oscillating entropy of the sound wave is zero. In thermoacoustics that is not true any more, because entropy can flow in and out of pieces of gas by thermal contact to the walls. So we need this extra piece.

You can gain a little bit of intuition by thinking of those two pieces separately sometimes.

(Transparency) (17)

Here is a schematic of a thermoacoustic refrigerator. I want to plot here the energy flows as a function of position and break it up into these different pieces. First, let us look at the work, the pv intensity of the sound wave. You put in work at the beginning here. You do not get any out at the end here of this refrigerator.

In-between you know that it has got to look something like this, because most of the work is going to get sucked up where all the surface area in the device is, so that is where the work is going to drop off the most quickly. There are not any discontinuities in this curve, because pressure and volume-velocity are things that are continuous in acoustics. It is kind of a smooth curve like that.

Total energy, the enthalpy-flux density, that has to be piecewise constant like that. There is a constant piece, and there is another constant piece. If everything is insulated, insulated from its surroundings, except where the heat exchangers are -- because conservation of energy tells you that at a place like this, for instance, the energy flux this way here has to equal the energy flux this way here, it cannot escape to the outside, so that makes the energy flux piecewise constant.

At the beginning it has to equal the work, because it is clear that is the energy flux at the beginning. At the end it has to be zero if we have insulation around our refrigerator. The discontinuities in energy flux at the ends of the stack have to be the amount of heat we are putting in or taking out at the heat exchangers.

The difference between here and here is Q_H and the difference between here and here is Q_C.

(Transparency) (18)

That is the picture to have in mind when you look at Rott's energy equation, which is just the hard-core integration and multiplication of $h_1 \times u_1$ spatially averaged and then time-averaged. Again, it has the acoustic pressure all over the place and the temperature gradient and all these f viscous and f thermals, and Prandtl number shows up in a few places here, too. You can kind of see the pieces of it here.

Pressure times velocity -- right here -- and actually it includes this f viscous here, is just the acoustic intensity, so that is what I labeled "work" a few viewgraphs ago.

This pressure times velocity times this stuff over here is the contribution to the heat flow that comes from pressure times velocity. This whole term has velocity squared in the temperature gradient, so that is another kind of heat-flow term. Then this last thing is just conduction of heat from hot to cold.

DR. ATCHLEY: May I ask you to go back to the previous viewgraph to say a couple of words about how you know which sign the heat flows and things are, what direction things are? Is that a reasonable question?

(Transparency) (17)

A: Yes, it is a good question, and I am smiling because it is a hard question.

(Laughter)

DR. ATCHLEY: This is the most confusing thing, I think, about these diagrams, is how do you know when the discontinuity is up or down and all that stuff?

A: In part, it is just convention. I tend to always think of positive energy flows as being things to the right and negative energy flows as being to the left.

Say you are going from here to here. You bring in extra energy. There is more energy flow to the right here than there was here, so the discontinuity -- it has got to be a step in the positive direction. So here is a step in the positive direction.

Here you are taking energy out, so the energy flow here to the right has to be less than it is here. So that is a step

down. I always have to think it through each time, though, and make sure I am getting it right.

(Transparency) (18)

So Rott's wave equation and his energy-flow equation, that is the whole hard-core thermoacoustics right there.

(Transparency) (19)

You can use those in a number of ways. There are quite a few approximate methods out these days. You can use a spread sheet or a hand calculator or a piece of paper or just about anything you want, and it is really useful for getting started and getting something designed to within an order of magnitude or so.

The hard-core way to do it is to do a numerical integration of those equations in the x-direction. Those equations you rewrite in terms of the pressure and the volume velocity and the temperature all as functions of x and this is really five variables, because these two are complex functions, so they each count for two.

You just numerically integrate from one end of the apparatus to the other, integrate those through whatever geometry you encounter. If you come to a thing like this, you might use Webster's horn equation. If you come to a thing like this, you might just use propagation in a duct. When you are in here, you use Rott's equations.

I am going to try a computer experiment on you here that might not be as successful as the previous ones. We have a computer program called DELTAE, Design Environment for Low-Amplitude Thermoacoustic Engines, that is available. I thought we would just go through kind of a simple example, and a fun one, also, to see how DELTAE does some of these things, and hopefully we can get a few physics lessons out of the deal.

(Transparency) (20)

This is the example I wanted to do for you. Remember the guy from Bangladesh I showed? This is what his refrigerator looks like. Overall it is a half-wavelength resonator with thermoacoustic stuff symmetrically on each end, engine stack here and here, so H is hot temperature and R is room temperature. Then refrigerator stacks here and here. Down here is a bigger diagram of one end of that. Here is the scale.

There will be a hot heat exchanger here. This is going to be a kerosene burner out here with flames licking through these fins. Inside there is going to be air at pressure of 5 Atm or something like that, something you can get with a tire pump.

That is Third World technology.

The engine stack will be between here and here -- it is not shown, because that would really clutter up the drawing too much. There will be a room-temperature heat exchanger here with cooling water just free convecting through this to cool it. Another stack here and then the cold heat exchanger will be inside the little refrigerator box.

Q: Is there room for gravity feed?

A: Gravity feed. He has in mind, once you get away from the apparatus, pipes like that big and a tower up on the roof of the building or something. If you can evaporate anything in that humidity, fine, let it evaporate to cool, but mostly I think it is just going to be conduction off the surface of the tank.

This is what would be sticking inside the refrigerator box and then some stuff to finish out the resonator.

(Transparencies of computer program) (21)

This computer program, DELTAE, accepts an input file that has mostly the geometry of the apparatus you have in mind and some other stuff as well. It marches through the apparatus in sequence. On the previous viewgraph it is from left to right. The rules for this program are that the only thing that matters is what it encounters first on each line, either one of these segment names that tells you what kind of geometry to expect or a number that tells you a value. All the rest of the stuff out here is just comments for us humans to keep track of everything.

So here we are. We have got 6 Atm mean pressure, 60 Hz, maybe we will start out at 600 K on the hot end, 5×10^4 pascals pressure amplitude, so that is a little bit less than 10% of mean pressure, and a bunch of zeros, air, and solid type ideal. That is kind of a nuisance, we always have to have a solid type even for segments that do not need them.

You remember 10 cm is the typical dimension, on the cartoon. So we have an endcap, we have a duct, 6-inch diameter. It has some perimeter. There is a hot heat exchanger that comes next. I have guessed that we might have to put 400 W of heat in from the burner at a temperature like that.

The stack. Let us see, we have maybe 5-mm gaps in this stack. That is going to be a cinch to build. It is only 3 inches long. We will try 0.002 inch sheet metal to make it. There is a heat exchanger in the middle that takes the waste heat out to room temperature, whatever your room is at in

Dhaka.

Another stack, the refrigerator stack. It is a little bit longer and a little bit thinner. Where did these numbers come from? I cheated. I knew the answers that I wanted. There is the cone that adapts you from the end of the heat exchanger to the final duct, and that is the end.

So if we want to do DELTAE on that, this computer program now has all that geometry in it and there is a bunch of different things we can do, like we can plot variables as functions of other variables and we can even ask it what it thinks the thermal conductivity there is. What we are going to do now is just run this thing. It is going to integrate those equations of Rott from one end to the other using that sequence of geometry. That is all it is going to do.

Now we can display what it has done, but the display is so much smaller on the computer screen than it is on the overhead projector that instead of displaying it over here I am going to display it over here.

(Transparency) (22)

Here is this sequence, again, of all the segments, heat exchanger, stack, heat exchanger, stack. The data that are in columns here are the real and imaginary parts of the complex pressure amplitude as functions of position all the way through the apparatus, real and imaginary parts of the volumetric velocity, the temperature, and the energy flux and intensity flux, or work flux, actually.

This is the kind of thing that it will compute for you. It took that input file and it computed along and -- well, look, it did some things that I am not very happy with it. I marked them in red here.

When it integrated along starting at the beginning, it found a temperature here in the middle that is supposed to be room temperature of only 293 K, that is 20 C. I think that is too cold for rooms in Bangladesh. So this thing is not going to be able to dump its waste heat unless we change something. We will have to try and make this temperature come out hotter.

Another thing it did when it integrated along, look at this, it has 20 W of work going out the end of the apparatus. That was the mirror-image plane in the center of the resonator. You cannot have 20 W of work going out a mirror-image plane. What is going on here?

What is going on is that often, in solving differential equations, you have mixed-up boundary conditions. You do not know all the boundary conditions at the beginning in this problem. You know some of them at the other end. You know the

impedance at the far end. It is a mirror-image plane. The pressure amplitude is supposed to be zero.

So you have to use a shooting method kind of technique, just like integrating some other differential equations. You have to fiddle with something up in the beginning here to try to get these things to come out right.

(Transparency of computer program)

DELTAE will do that fiddling for us. The way that it does that fiddling is it lets you select some target boundary conditions near the end of the apparatus and some fiddle variables near the beginning of the apparatus. I am going to introduce two targets and the two fiddle variables right now.

I am going to use ob, which I know to be the frequency, as one of the fiddle variables to try to get to the end of the apparatus in a quarter of a wavelength instead of some other fraction of a wavelength. I am going to use the heat that is taken out at the middle heat exchanger as another fiddle variable. Then as targets at the far end I am going to use the impedance at that mirror-image plane.

So now I am going to run it, again.

If we display the results, again, remember these were some of the things we were upset with. Now it has gotten the impedance at the end correct. I am going to add a third, because we do not have the temperature correct yet. It is even colder now. This would be a good thermoacoustic refrigerator for Antarctica.

Now let us target that waste heat exchanger.

We are going to try to make that temperature come out 320 instead of 200-and-whatever-it-was. We are going to use, as a guess, another one of the heats. We are going to run it and when it is done we are going to display our results.

(Transparency) (24)

Now we have 320 K for our dump temperature and we have got zero pressure amplitude, real and imaginary part at the mirror-image plane. This is just to give you an idea of what this computer program can do. Now we can tell the guy in Bangladesh that he should expect to have to put in 300-and-something watts in his burner at a temperature like about that. If he does that, he can expect to get 35 W of cooling power in his refrigerator at -10 C. That ought to be cold enough to keep medicine cold.

This is the end of the hard-core part of my talk, so it would be a good time to break for questions and more coffee.

(A brief recess was taken.)

DR. SWIFT: Between the first break and the second break, that was the extent of hard-core thermoacoustics for us today. The one remaining thing I want to do for you is show you the application of similitude to thermoacoustics, mostly because I think similitude is a tool that every acoustician ought to have in his tool box, because it might turn out to be useful for whatever you are doing. It was really a surprise to us that it has been as fun and useful as it turned out to be.

(Transparency) (25)

We had an experimental engine in the lab at Los Alamos that looked like this. It was a half-wavelength resonator. This is about 4 m long. The stack was about that big around and 1-mm-sized gaps in there. Hot heat exchanger that was electrically heated, and a cold heat exchanger that was water-cooled. This did not do anything useful; this was just for experiments to see whether we knew what we were doing.

You could measure the heater power that you put in with electric heaters and you could measure the temperatures and you could measure the oscillating pressure at the end of the resonator. One thing that we did, we thought, well, let us see what the harmonic content of the wave is like.

This plot shows the amplitude of the harmonic P_2, so the second mode, divided by P mean plotted as a function of the square of P_1, the amplitude of the fundamental of the sound wave. So it is not too surprising that this should be a straight line, that the second-order things should be the square of the first-order thing. That is the normal sort of deal.

But I was just amazed that it did not matter whether you used a monatomic gas or a diatomic gas or a gas mixture with a different Prandtl number, and it did not matter what mean pressure you selected, you got the same slope for all of these things. That actually is what got us thinking about what tools you could bring to bear on nonlinear effects in thermoacoustics that we do not understand yet.

(Transparency) (26)

Similitude is one such tool. This is one of the simplest examples of the application of similitude in fluid mechanics. Suppose you have a sphere in a fluid and either the fluid is moving past the sphere or the sphere is falling through the fluid, I do not care. If you think about all the variables that the force between the sphere and the fluid might depend on, they are the fluid velocity, the density of the fluid, the

viscosity, and the size of the sphere.

If you think about how you can group those variables together in dimensionless groups, you quickly find out that rather than having a function of four variables, you really have a function of only one variable. The force is a function of the Reynolds number and nothing else. So this is the kind of argument that has led people to realize that you can make scale models of airplanes and stick them in wind tunnels and make measurements and know exactly what you want to know about the full-sized airplane.

Or a little more complicated example is drag on ships. If you are going to build a new kind of oil tanker and you want to know how many horsepower it is going to take to shove the thing along at a certain speed, people will make a scale model of the ship and put it in a water tank and drag it along. Actually, it is a bit more complicated, because you cannot get the viscous drag and the wave drag to both scale in the same way.

People can compute the viscous drag but they cannot compute the wake of a complicated hull, so they will use the scale model to predict the wave drag on the full-sized thing and they calculate the viscous drag on the full-sized thing.

Q: Isn't that because they are still using water? If you use a different fluid with a different viscosity or different properties, couldn't you get that to work, then?

A: Get both of them to work? That sounds right. I am not sure. Maybe it is a matter of just not using such exotic or nasty fluids. I am not sure about that.

Anyway, similitude is the science of picking your dimensionless variable groups right. You can look in textbooks to see really how to do that.

(Transparency) (27)

Jeff Olson thought about that problem with respect to thermoacoustics. If you think about some generic kind of thermoacoustic apparatus, there it is. What are all the variables it takes to describe the problem? You have all the geometry, the geometric variables. Some of those we selected out for special attention, like the overall length of the apparatus we call L, the cross-sectional area we call A, the gap spacing in the stack we call h, and then all the other dimensions we will just call x_j.

If you have an ideal gas working substance, that is uniquely specified by the ratio of specific heats, the sound speed, and the transport properties; and these b's are the exponents and power-law fits to the temperature dependence of

the transport properties. That is one way you can do it, a pretty simple way.

If you have solid properties, the thermal conductivity of the solid is something that matters. If you have some miscellaneous variables like temperature, mean pressure, and heat flow, I have this subscript $_{ref}$ running all over the place here. You always want to pick one location that you call your reference location. That will be a reference temperature and that temperature will define a reference sound speed and so on.

These might be all the input variables or variables that you control in an experiment, and you want to find out some values for some dependent variables like the frequency that this thing chooses to oscillate at, and the temperature distribution throughout the whole apparatus, the velocity and pressure distribution, maybe the power that it puts out.

If you go through this similitude analysis, it will tell you some of the ways you can group variables together to make them dimensionless. Of those ways, this is the one we kind of settled on as making the most sense to us.

Some independent variables over here -- geometric ratios, those are kind of boring. Some dimensionless properties of the gas, like the Prandtl number and the ratio-specific heats. The heater power scales like mean pressure, sound speed, and area. Of course, seeing thermal penetration depth divided by pore size in the stack, that is no surprise. Then here is the dimensionless version of all of our dependent variables.

What does this buy for us? If you stare at this long enough, you can actually figure out how to make scale models of thermoacoustics apparatus. It turns out, for instance, if you want to make a half-sized scale model of an apparatus that you designed for helium gas, if you cut all the dimensions in half, these ratios of dimensions do not change, that is okay. If you switch to another monatomic gas, these do not change much -- close enough. If you switch from helium to argon and change the mean pressure just a little bit, you can get the thermal penetration depth to drop just exactly by half. So this dimensionless variable stays the same.

Then you just turn down the heat and you have it, you have got all the independent variables the same as they were for the full-sized helium engine but in a half-sized argon engine, and similitude says that all of these variables have to be the same for that scaled argon apparatus, including all the nonlinear stuff that we do not understand.

You can do size scaling and switch gas like that. You can also do temperature scaling, like if somebody would ever pay us to work on a power supply for the Voyager-13 spacecraft that

would run off some red-hot plutonium heat source and dump its waste heat into space off a radiator at 500 K, we could do a temperature-scaled model, build it with reasonable materials and not have red-hot parts in the lab, and learn a lot before going after the hard materials that you have to use at higher temperatures.

(Transparency) (28)

I will show you an example of how this similitude works out. In that apparatus I showed you a couple of viewgraphs before -- and here is kind of a reminder of crudely of what it looks like -- just a simple engine that does not do anything useful, it just sits there and oscillates, the plots here have the square of the acoustic pressure amplitude of the fundamental divided by mean pressure squared, as horizontal axis.

So out here at point 0.010, this is 10% pressure oscillations at the end of the resonator there. Vertical axes, T_H/T_C, 2.0, 2.2, 2.4, and heater power required to make the thing run normalized by its group of other variables. The way to think about this is what heater power does it take to make it run at a given amplitude and, when it is doing that, what hot temperature has it settled down at?

The three different sets of points are for three different gases with their mean pressures chosen so that they all have the same value of this ratio of thermal penetration depth to stack pore size, because that is one of the variables you want to keep constant in similitude.

What this shows, the fact that these data all line up together just shows that, hey, similitude works. It should not be a surprise. You can switch from helium to argon and drop the pressure by about a factor of three and if you plot it this way, you get the same answer.

A couple of other things that this viewgraph illustrates: The lines are calculations from that computer program I was showing you half-an-hour ago, so in the low-amplitude limit it does a good job. That is the way an acoustic approximation calculation ought to be, right? In a low-amplitude limit it ought to do a good job.

At higher and higher amplitudes it deviates more and more. Order of magnitude of the deviations? At 10% pressure amplitude, 2.0, 2.2 -- well, 10% errors. That is reasonable for an acoustic approximation theory.

But it would really be great to be able to do better, because when you design an apparatus for somebody who wants to make a prime mover and he wants to squeeze as much efficiency

out of it as possible, make it run at the hot end as hot as possible, if you tell him to expect that it will run at this hot temperature and it turns out to run at this hot temperature, the stainless steel has just melted, or the Inconel [phonetic] has just melted and you should have used tungsten. From a practical point of view it can be pretty important to try to go after these effects that are a little beyond the acoustic approximation.

Q: Did you go back and nondimensionalize your field equations? You have got some nondimensional work variables.

A: Nondimensional versions of Rott's equations?

Q: Yes, the temperature, your field equations, your Navier-Stokes?

A: Yes, Navier-Stokes, right.

Q: In that way I was thinking if you had problems, you would like to do better than the acoustic approximation, you could back in there and see if you get something.

(Transparency) (29)

A: This an example of switching from one set of mean pressures to another, one value of delta kappa over h to another value of delta kappa over h. At that value, for some reason the nonlinear effects do not affect the hot temperature too much, but they lead to more error in the powers. I do not know why.

(Transparency) (30)

As another example of how you can get a pleasant surprise out of similitude, that apparatus had a strange behavior in some regions of parameter space; instead of humming nicely and steadily like the demonstration did, it would shut off for a minute or two and then start up screaming really loud, scream for a minute or two and then shut off, again. It would do that all day. We do not know what is going on there, we just do not know.

But if you plot the stability region of that problem in the dimensionless variables that similitude tells us to choose, you get the same stability map for two different gases. So I think that maybe should give you a hint of how powerful a tool occasionally this similitude can be. It can at least help you correlate things that you do not understand.

(Transparency) (31)

I have some summary viewgraphs and then we can break up into probably multiple small groups for 20 minutes or so to

kick around some thermoacoustics things, because there are so many thermoacoustics people here.

To summarize, you remember this from the beginning of the morning. Thermoacoustic engines are nothing but the marriage of a heat engine picture with acoustics a la Rott's wave equation. You can get both prime movers and refrigerators. Whether these things ever get used in a commercial sense will depend on the details of the tradeoffs between these advantages and shortcomings. Many of us in this room are hopeful about it, but it is not a done deal yet.

(Transparency) (32)

Efficiency, I think, is the most serious shortcoming of thermoacoustics and so I want to come back to this viewgraph, again. Why don't thermoacoustic engines have Carnot's efficiency? There are four reasons, and I hope those computer animations helped you to appreciate some of them.

There is viscous dissipation on all the surface area in the apparatus wherever there is velocity. Think of a little highlighted line along each surface here. Pick a color -- red. Make a little highlighted line. All that red volume is dissipative. Red should be brighter where the velocity is the highest and you can make it zero where the velocity is zero.

Intrinsic losses in delta kappa. What do I mean by that? The whole effect of thermoacoustics depends on heat transfer in that thermal penetration depth. You remember the red arrows in the computer animation? That heat flow from the blob of gas to the walls was across a distance of about a thermal penetration depth, and diffusion of heat is a lossy process.

It turns out that that process going on in all the surface area of the stack region not only gives you all the effects you want but it is also lossy and we are stuck with that in thermoacoustics.

The third loss mechanism is just dumb conduction of heat from hot to cold by the gas, by the stack, by the resonator walls.

The fourth I guess I have already talked about. All those processes take place not only in the stack but also along the resonator walls.

About a third of Carnot is typical of present designs. That is good compared to Peltier coolers and thermoelectric-effect electric generators and stuff like that, but it is not good compared to steam turbines and car engines.

Q: Are the sizes meant to represent the level of the

problem?

A: It seems to me a lot of the designs end up giving you the first three being equal if you fiddle with your geometry properly to minimize the total. That is the way it seems to come out. I do not know if there is a reason why that should be.

For large apparatus like the natural gas liquefier, the resonator losses are negligible, because there is very little surface area on the walls of the resonator compared to everything else. So how No. 4 balances in with the other three is a matter of size.

Q: How many of those do you think are sort of surmountable in the near future? Which ones do you think attention should be applied to trying to reduce their effect?

A: We need more ideas. The ratio of viscous to other stuff is tied to the Prandtl number of the gas, viscosity times C_p over the thermal conductivity. That is two-thirds, for monatomic gases. Gas mixtures, which is an idea that just came into our community five years ago, or maybe a little more than that, that helps there. If we keep coming up with ideas like that every few years, we will improve the efficiency 10% every few years.

It really needs more invention to improve on these things.

Q: Suppose you could pick a working fluid where you could tune the relaxation, that is to say, you have a relaxing fluid a la Hank Bass' talk, where you could tune the relaxation time to whatever you wanted. Is there ever a case you could win by having a relaxing fluid?

A: There is an idea, and that is the kind of new ideas we need. I do not know, but it surely sounds like something worth pursuing.

Q: I guess there is a version of Rott's equation that has that in it.

DR. ATCHLEY: It is not only viscosity, because I think gamma minus one shows up everywhere. If you put something complicated in and get relaxation, then gamma drops. Hank can tell us everything we want to know about that.

DR. BASS: But what if you make gamma complex? You can get some phasing there as well. I am not sure where it goes. I am saying what if?

Q: Those are all losses that are occurring in the bulk, right, and those, typically, are small compared with the losses

that occur at the walls.

A: Maybe you can beat that into shape so it is not -- I think you would still need the contact to the walls, but if you could mess up the adiabatic temperature oscillations that go along with the pressure oscillations and get some phase out of that, maybe you could win.

DR. BASS: It is not true that the bulk losses are always small compared to the wall losses. It is for the type of gas that we talked about, but you could certainly come up with a gas where bulk losses dominate. I can think of a few, UF_6, for example.

A: Well, SF_6, then.

DR. SHIELDS: There is evidence that the thermal conductivity for vibrational energy is different from thermal conductivity translation energy for gas, and also the combination of vibrational energy for these wall collisions is different from translational combination, so that is something else to be measured before you would really know how the vibrational energy is being conducted through the gas.

Q: Can you get to that answer by introducing a relaxing variable in the similitude arguments?

A: It will not tell you whether it helps or hurts. It will just tell you how to group variables, or how to switch from sulfur hexafluoride to uranium hexafluoride once you already know what sulfur hexafluoride does.

Q: Does it help to use different kinds of gases and put a membrane between them, one kind of gas on one side and one on the other?

A: I think that could probably help. That will help in issues like power-per-unit weight of apparatus. In fact, Tom has been thinking about putting a membrane in and getting rid of the gas on the other side just to save size. Use a compliant mechanical element to take away some of the resonator.

Q: In the Carnot cycle the entropy is the same, the entropy did not waste. In other words, the temperature between high temperature and low temperature, the difference is extremely small. In the ideal gas it is almost close to zero, so it just a very minor step by step to reach the Carnot efficiency.

In your case, would that help, if you made lots of stacks, real small stacks one after another, so the temperature difference is small? You can increase your efficiency.

A: I think it will not help as much as you think, because we really have that already. Typically, the gas displacement amplitude might be 10% of the length of the stack. You can really think of it as 10 little engines in series, each engine being the length of a gas displacement amplitude. So we really already kind of have that.

Another way to answer your question is that there is a research group in Switzerland that is doing just exactly what you said, breaking up the stack into well-separated segments like that just to try to understand what goes on there.

Q: When you think about it, it seems that thermoacoustics works very similarly to a regular refrigerator on a micro-scale, you are compressing the gas and getting rid of the heat. Then you put it another place and you expand it and get cooling. The difference is, in a regular refrigerator you have a genuine phase change, which increases your efficiency a lot. Is there any way that can be used here?

A: That is another new idea, phase-change working fluid in thermoacoustics. There may be some people who have looked at that in the context of Stirling machines, but I do not think they have looked very hard.

Q: In your first scaling argument, where you cut the size in half and you went from helium gas to argon, I did not quite follow why you used the argon. It seems to me you have to have a balance between your conduction along x and the power-density of the system.

(Transparency) (27)

A: What you want to do is change things in such a way that you keep these variables the same. Suppose you decide that a factor of two in size is about right, because if you could find a quarter-inch bolt you can find a half-inch bolt.

Then what you do is you look at the thermophysical properties of the monatomic gases and you see which one comes the closest to keeping this constant for you when you just change gas and change h by a factor of two. It is a little more complicated than you think, because delta kappa depends on frequency, and when you change size by two and change gas, the resonance frequency is way different, because the length and the sound speed are different.

Q: That is if you change the length and the channel spacing together. You could let those scale independently.

A: Not to do hard-core similitude, no. That would be like a scale model of an airplane where you decided you would make the wings half as long but the same width or something.

You just cannot count on similitude under such circumstances.

Q: If you take the original Rott equation with some of the other stuff deleted and you scale the length of the whole tube down by a factor of two and the channel spacing down by root two, then you have similarity in those equations.

A: You do, but the Rott equation is within the acoustic approximation, so if you want to be hard core and go beyond that, you have got to be a little more careful.

Q: By introducing argon don't you trash the relationship between the conduction along x and the cooling power and heat-flow density?

A: Oh, yes, you do, actually, a little bit, but of all the things that happen in there, that conduction along x is the one that I know how to calculate the best. I am willing to calculate that out and then add it back in, like the viscous drag on scale models of ships.

Q: It is not similar, then?

A: Not exactly. Actually, if you look at the ratio of thermal conductivity of helium to argon, you can go looking for solids whose thermal conductivity is the same fraction of other solids as argon is to helium. It turns out stainless and copper have about the same ratio of thermal conductivities as argon and helium. So in the one scale model we are building we actually have stainless steel fins in our heat exchangers to make that part also scale.

(Laughter)

DR. PUTTERMAN: I would like to ask a question about the derivation of Rott's equation from basic Navier-Stokes equations. Since you linearized, obviously the Mach number had to be taken small in order to get Rott's equation. I was wondering, what are the other quantities that had to be taken small in addition to the Mach number?

Obviously, there is the viscous penetration depth over the wavelength of sound. That had to be taken small to get even your simple relations with dp_1/dx times one minus e to the i one plus i y over delta. What other quantities have to be small in addition to the penetration depths over the wavelength and the Mach number?

A: There are a couple of other assumptions built in. I think the biggest one is the one you just mentioned, the penetration depth over wavelength.

DR. PUTTERMAN: I was concerned as you were flashing

through that that there might be an assumption on the length scale which determines the critical gradient in temperature in the stack. There is some characteristic length scale there, and that length scale somehow must be large or small compared to the viscous penetration depth in order, also, to apply Rott's equation.

The reason I am asking is, has someone taken Navier-Stokes equations and carried out a systematic asymptotic expansion in terms of which quantities are small and shown that gives Rott at some order of epsilon where epsilon is these very small quantities?

A: It probably has not been done at the level of rigor that you would like. Of course, displacement amplitude compared to lengths along x, you have to neglect that, but displacement amplitude compared to penetration depths I do not think you have to neglect when you go through it, which is lucky, because the displacement amplitudes can be centimeters, with penetration depths millimeters.

DR. PUTTERMAN: Even that small Mach number guaranteed linearization, because you have now got diffusion mixed in, and then you get nonlinear terms that do not go away with Mach numbers small necessarily.

A: I would like to see that. I do not understand that.

Oh, yes, because there is one -- let me see, when you start out with Navier-Stokes, the viscous term, is it like there is a derivative and then the quantity, viscosity, and more derivatives of velocity? It may be that we slipped that derivative in past the viscosity there, for instance.

DR. PUTTERMAN: I guess I was asking for a reference or for a list of the small quantities. I was curious if that has been done.

Q: That is what I was asking before. You have gone through some basic stuff on similitude here and that is why I asked earlier did you go back to the field equations, or he likes Navier-Stokes equations, and go through? Now you have got all this stuff right here, you worked it out, and then you can tell within the parameters you chose here what is small and big, and then you can carry out what Dr. Putterman is suggesting.

A: With the dimensionless versions of the equations.

Q: Yes, that is the way they do it in engineering a lot.

Q: We actually wrote out all the dimensionless and the Navier-Stokes heat transfer, all that stuff. I mean, it is out there, but we have not done the analysis of this thing is this

small, this other thing is this small, and put it all together and how big is the agglomeration.

Q: But that might be a very good thing to do.

Q: I may not be picky enough, but I really liked the discussion in your appendix to your review article as you go through the Navier-Stokes development there. That is a good place to look.

A: I may not be picky enough, either, but that was as picky as I could think of how to do it.

DR. ATCHLEY: Wheatley's paper at the last Varenna conference goes through a derivation similar to what you have done in JASA but uses different words, and some of the approximations may be more to Seth's liking in there, the way it is spelled out. That is the only other place I know of that has a detailed derivation.

Q: The stack has so much surface area that there is lots of dissipation in it. Now the stack acts kind of like a bucket brigade where each little parcel of gas is going back and forth and shuttles some heat down so that it is ultimately shuttling heat from the hot exchanger to the cold exchanger.

What if you set the separation between the exchangers such that a displacement length is the separation between them so that instead of a bucket brigade you are taking the bucket and tossing it directly on the fire instead of going bit by bit? Can you get rid of all that lossy surface area?

A: I am not going to answer that. That is a topic for discussion.

DR. ATCHLEY: They had a sound source that had just heat exchangers and no stack. They put the heat exchangers close enough together so you would create that gap.

Q: And is it five times better?

A: I do not see any of them around today.

(Laughter)

Q: When we were going through patent processes somebody dug up Bell Lab's patent, and it was an electricity generator, I think. It had a heat source and a heat sink, no stack in-between, they were close together, and then something like a loudspeaker to generate electricity.

DR. GARRETT: I think you have to realize the stack, in some sense, is like the regenerator in a Stirling engine. You can do that, but you cannot stand large temperatures if you do

not have an intermediate step, I do not see how you stand the larger temperature range. There may be applications where you do not need a larger temperature range or that you get the acoustic amplitude large enough that in one jump you make it all the way to what you need. That would be wonderful if you did not have to dissipate --

Q: Or if you had a design where your penetration depth was a 100th of a millimeter. Trying to build a stack to a 100th of a millimeter size would be terrible, but you might be able to do screens for the heat exchangers and not stack.

Q: I was going to mention that Anthony and I did a prime mover experiment to study heat exchangers under very large amplitudes. It was room temperature to liquid nitrogen temperature, a prime mover. The highest amplitudes we achieved were a peak pressure over mean pressure of 24%.

If you calculate what the displacements in the stack region were, the peak-to-peak displacement amplitudes at our best values were about equal to the hot-to-cold heat exchanger distance. We thought about what would happen if we pulled the stack out, but I think you could never get it started. If you could hit it with an incredibly big hammer, you might be able to get it started.

Q: How did you tune your speaker frequency to reach the resonance?

A: That is a question for Tom, actually, having to do with tuning the frequency in a loudspeaker-driven apparatus. In a prime mover, like the demonstration, it picks its own frequency, and that is that, but in the refrigerators you pick the resonance frequency.

DR. HOFFLER: We played around with that a little bit. Normally, what we do is we have instrumentation where we can determine the phase between pressure and velocity at the driver piston, an accelerometer on the piston and a microphone right next to it. Usually we tune it so there is zero phase between pressure and velocity, so that is sort of a gas resonance condition.

We have also tried shifting that phase a lot, and while generally your piston stroke goes up, it really does not seem to have a lot of effect. It depends on the Q of your system. If you have a high-Q system, it takes a small frequency shift to make a big phase shift. If the Q is quite low, then you have a much bigger frequency shift for a given phase shift.

But even in some of our lower Q systems, if you shift the frequency and the phase quite a lot, it generally does not have a lot of effect on the performance. But you have to remember that what you are doing is shifting the position of the

pressure antinode relative to the position of the stack.

In fact, we have gotten slight increases in temperature span and efficiency by, in some cases, pushing the pressure antinode closer to the stack by raising the frequency. It also has a pretty drastic effect on the electric-to-acoustic efficiency of the driver, because you can either tune to the gas resonance or you can tune to a system resonance, where the whole mass-spring of the piston and the gas are all self-resonant and you get rid of some of the reactions in the mechanical part of the driver; that improves the efficiency.

DR. GARRETT: There is another comment I would like to make, since I am a big fan of gas mixtures, which is it is possible to configure a refrigerator that does not require frequency control by using gas mixtures and having some medium in the cold end that can absorb or desorb the heavier species. If you have, say, helium and xenon, xenon freezes out at 131, the helium does not freeze out until 4°.

Therefore, if you have, say, carbon granules on the cold side of your fridge, as the temperature of that fridge goes down, you selectively desorb from the gas mixture the heavier species and, therefore, C^2 being gamma RT over m, you can pretty much lock the resonant frequency to be independent of temperature and then drive off your 60- Hz utility and e frequency or some harmonic of that.

There are a lot of ideas in frequency control. Not all of them involve having to track the resonant frequency. In Tom's fridge it is a lot harder, because he is getting it real cold, and you have to span a fairly large frequency range because you are spanning a large temperature range.

But for the question of frequency control there may be passive solutions that do not involve electronics at all. If you are trying to build something that is cheap and efficient, those solutions are also interesting.

DR. HOFFLER: With this demo, though, first of all, it is a lower temperature span, so the frequency does not want to change much there, but also because most of the apparatus is warm, when only a small part of it gets cold there is almost no frequency variation on that at all. It is only half a percent or less, so we ellipsis?

DR. HARGROVE: Why don't we let Greg wind down and we can go off line and you can ask those questions.

A: Summary of the summary here.

(Transparency) (33)

Here is what we know how to do the best. Start with your basic fluid mechanics stuff. Make the acoustic approximation. Integrate analytically in the planes like this to get Rott's equations and then integrate numerically along x to get everything else you want.

I think this is mostly understood in the acoustic approximation, although I think we do not really understand a lot of details at the heat exchangers yet properly.

Other things for the future besides the heat exchanger details include all the high-amplitude effects. You want to run these things at the highest amplitude you can, because that puts more power per unit volume of apparatus and applications are always going to be sensitive to power-per-unit volume or unit weight, so we are going to have to understand the high-amplitude effects.

Then I think that where the action is really going to be in the near future should be in the real applications, both in picking applications that make sense from an economic point of view and also in terms of all the invention that is required to bring these physics ideas into reality and make them work economically. That is it. Thanks.

PHYSICAL ACOUSTICS IN MATERIALS SCIENCE

DR. GREEN: Before I start, I want to tell you two stories of how I got to be a professor doing teaching and experimental research. I went to Brown University for graduate school when Bruce Lindsey was chairman of the department. Just as now, since they did not know what great talents I might have, they gave me a teaching assistantship and, also, I got the favor of teaching an undergraduate lab for nonphysics majors, like physics for poets.

There were two labs each day of the week and I was assigned the first lab on Monday. That meant that there were nine other T.A.'s watching me conduct the lab. As you can tell I am not from Boston, and my accent was more deep South than it is now. The first experiment was to measure the acceleration due to gravity. The Physics Department was in an old building called Wilson Hall that had a spiral staircase. The deal was to put some students in the basement, some on the first floor, second floor, and third floor, and go up into the attic with the remainder. A student in the attic was supposed to drop a red rubber ball down the spiral staircase. The acoustic part was to measure the time which it took the ball to pass each floor by listening to the tick of a metronome. I had some students stand at the various levels and I went up to the attic with the student with the ball and the student with the metronome. I said, "Are you all ready in the basement," and everybody laughed because I said "you all." I said, "Are you all ready on the first floor, second floor, and third floor" getting more and more nervous, and more and more laughs. So finally I told the student holding the ball in the attic, "Okay, let it go," and this guy threw over the metronome.

(Laughter)

The other guy said, "What should I do with the ball?" So I told him.

(Laughter)

They did not do that experiment the same way any more afterwards. They never gave anybody the new metronome. They always played the sound of the metronome over the P.A. system. So when they did that experiment afterward, you could hear tick, tick, over the whole building.

Now I will tell you how I got to Johns Hopkins and then I will quit telling stories about my past. All my degrees are in physics. I made a mistake going to Brown, because I wanted to be a high energy physicist and they had a Conkcroft-Walton accelerator there, where you had about 10 or 15 guys -- that

was in the old days, now they have got 30 or 40 or 50, as you probably know -- running the experiments and you had to get in line to get onto the experiment. So Lindsey suggested that I measure the elastic constants of ice crystals. Someone had done it with an optical technique and he suggested I should do it with an ultrasonic technique. I got about two years into my masters research when a hurricane came into Rhode Island, the power went out, and my research work melted. They were tough in those days, despite my begging and pleading, and made me start over, again. So it took me three years for my masters degree and I was ready to quit when a Professor came from Germany and wanted to know if I wanted to work with aluminum crystals. First, he asked did anybody know anything about crystals and they said, well, Green has been messing with ice crystals for three years, so he ought to know something about crystals. He then asked me did I want to work with aluminum crystals and I asked him what was the melting temperature. (Laughter)

So I became a materials scientist. I spent a postdoctoral year in Germany, thanks to the German Professor, and then I got this letter from Johns Hopkins asking did I want a job as Assistant Professor in Mechanical Engineering? I said the same thing I am sure all of you would say, "I am a physicist. I don't want to be an engineer, I don't even like engineers." I got this letter back saying, "Wonderful. You're just the kind of guy we want."

(Laughter)

So I went, and I have been there 34 years now. The only reason I am telling all this stuff is I know that some physicists are having a hard time getting jobs nowadays, so since I feel that my time at Hopkins has been worthwhile, I suggest that there might be other options for your young physicists. You might be able to help engineers by teaching them some fundamental physics or even using your knowledge of physics for practical application.

Most of the talk I am going to give is quite similar to that which the other speakers have done, except I am going to be discussing physical acoustics of solids rather than fluids. I guess the main difference is to talk a little bit about some anisotropic behavior rather than isotropic.

(Transparency 1)

I apologize if some of this is too elementary, but I thought that more of the people here would be second-year graduate students and maybe never even had a course in elasticity. When I was in the physics department I learned a simple version of Hooke's law and that is about as far as I

went, so after I got to Hopkins I had to learn a lot of elasticity the hard way. The normal way people start off with solids is to consider a volume element that has different stress components applied to it. There are normal stress components, normal to the faces of the volume element, and then there are shear stress components. To be in stable equilibrium under translation, you have to have a normal stress σ_{11} in the +1 direction balanced by a normal stress σ_{11} in the -1 direction. In order for the volume element to be in stable rotational equilibrium you have got to have shear stress component σ_{12} equal to shear stress component sigma 21. So eventually you put all these things together and you have the normal stresses, shear stresses and, along with them, the normal strain components and the shear strain components.

Then, for linear elasticity, you employ a generalized form of Hooke's law such that each of the strain components is a function of the stress components and vice versa. This is the way you do it experimentally, because normally you apply a force and you measure what happens, i.e. the strain which is the change in length per unit length.

(Transparency 2)

We call them all the time elastic constants, but they are not very constant. Nevertheless, I am in the habit of doing that, too, although it is better to refer to them as elastic moduil.

(Transparency 3)

If it is an isotropic solid, Hooke's law simplifies considerably. If you use the C's -- I will not show the inverse (It is similar for the S's.) But if you use the notation most often used by physicists, you would used the Lame constants, λ and μ.

(Transparency 4)

However, most engineers use Young's modulus, Poisson ratio, and shear modulus. With isotropic materials there are a lot of names given to these different moduli. For isotropic materials in the linear approximation you need only two.

(Transparency 5)

Everything I am going to do is like this. It is a very simple relationship. I can show, by considering the case of a long thin rod, where I have a little volume element here, with area A and length dimension dx. On one side of this volume element I am going to have a stress, σ_{xx} in the -x direction,

and in the +x direction I am going to unbalance it with a stress $\sigma_{xx} + \frac{\partial \sigma_{xx}}{\partial x} dx$. Thus, the little volume element experiences an acceleration toward the right in the +x direction. If we let, u, represent the displacement of the volume element in the +x direction and designate the density by, ρ, and apply Newton's Second Law, we find that the net unbalanced force is $\left(\sigma_{xx} + \frac{\partial \sigma_{xx}}{\partial x} dx - \sigma_{xx}\right) A = ma = m\ddot{u}$. Next I use the generalized Hooke's law and assume that the linear expansion is large compared with any lateral contractions, so Poisson's ratio is small; I obtain a simple 1-D wave equation with $v = \sqrt{\frac{E}{\rho}}$ (E is Young's modulus)

It is kind of interesting, if you do not deal much with practical engineers, to learn that in the engineering world many really believe that any material they pick up in the shop and pull in a uniaxial tensile test, that the slope of the linear portion of the stress-strain curve is the Young's modulus. Since I am a professor in a materials science and engineering department, I have to tell you there are very few solids that are actually isotropic, or homogeneous for that matter. If you solidify a metal, the crystallization process follows the temperature gradients which usually proceed in from the outside of a cold mold. If you deform a metal plastically by rolling it into a sheet you will see long grains along the rolling direction similar to a piece of wood. If you recrystallize it, which means you take a piece like that has been heavily plastically deformed and then heat it, solid state diffusion will take place and, if the grains are deformed enough, new grains will grow at the expense of the old in a solid-state transformation. So all of these things result in textures, which sometimes are single crystals, sometimes are isotropic, and most often something in-between.

(Transparency 6)

In 3-D, I will first consider the simpler isotropic case. Since all directions are the same for an isotropic material, it is sufficient to again consider wave propagation in the x-direction only. I have to consider all of the stress components in the x-direction. I go through this very similarly to what I did in the 1-D one.

(Transparency 7)

The unbalanced stresses look like this now in the x-direction.

(Transparency 8)

And I get an equation like this. Stress is a second-rank tensor, so writing it in tensor notation, we can do these equations very fast. The generalized Hooke's law is like this, where the elastic moduli are a fourth-rank tensor. Perhaps this is a good place to point out that a lot of people draw analogies between elastic waves and light waves, i.e. electromagnetic radiation. It is very dangerous to do this, because in the linear approximation the electrical conductivity is a second-rank tensor, the dielectric constant is a second-rank tensor, and the magnetic permeability is a second-rank tensor. So elastic wave propagation is much more complicated than in electromagnetic wave propagation, even in the linear approximation. It is not a one-to-one analogy. Sometimes you can find an analogy and sometimes you cannot.

(Transparency 9)

The strains look like this in terms of the displacement gradients. This one-half in front of the shear strains gives people problems sometimes. That means when you do the matrix of the elastic moduli you need to also put in a two if you define the shear strains with the one-half included. If you do not put it in here, then you do not have to do that, but you have got to be careful to be consistent. I could substitute all this in, just as I did in the 1-D case, and assume some solution, but now I want to proceed a little bit differently. These are the components of the displacement. I will call them u, v, and w for the three directions, u along the x-direction, v along the y-direction, and w along the z-direction.

Next I assume a solution as plane waves propagating along some direction given by this combination of wave numbers, and in a minute I am going to put in some direction cosines down here for that. This would be like a general three-dimensional case. In fact, I am going to use *l* to be the direction cosine of the wave with respect to the x-axis, m to be the direction cosine of the wave with respect to the y-axis, and n to be the direction cosine of the wave with respect to the z-axis. But I have to worry about something else, how the particles move, because I can have longitudinal and transverse waves. The direction cosine of the particle displacement with respect to the x-axis is α, the direction cosine with respect to the y-axis β, and the direction cosine with respect to the z-axis γ. If $\alpha = l$, $\beta = m$, and $\gamma = n$, we have a pure mode longitudinal wave. However, if the sum of the products $\alpha l + \beta m = \gamma n = 1$, then the wave is a pure mode transverse wave. In general for anisotropic media, the waves will neither be pure mode longitudinal nor pure mode transverse.

(Transparency 10)

I am going to jump way ahead, because I want to show you some other things. When I substitute all of this in, I can come up with an equation to solve that looks like this and includes the elastic moduli, the direction cosines, the density of the material, and the wave speeds.

These λ's, in this expression you can write out for different crystal systems. Then, in matrix notation, you can reduce the four indices on the elastic moduli down to two. If I remember right, this is the Einstein convention. It is contracted notation.

(Transparency 11)

These λ's, in general, are functions of the direction cosines *l*, m, and n, and the elastic constants.

(Transparency 12)

To solve this equation, notice, too, the way people do in the forward direction and the way I saw everybody doing it this time, I think, something that always bothered me when I was in physics, is that if you know the solution, you can put it in. I never could figure how they knew the solution to start with. But nevertheless, that was a good way to solve differential equations; if you know the solution, you put it in. If it works and makes sense, do it, even though there are other ways.

But this is the other way. It is a general formulation that if you do not know what type of wave it is in advance, this tells you what kind it is. This is an eigenvalue equation. You remember that λ is a function of direction cosines, the elastic moduli of the material, and the density. This is a sextic equation, but it turns out to be solvable in terms of v^2.

(Transparency 13)

The other thing that you have is the eigenvector equation, where these are the direction cosines of the particle displacement, the same λ's, so on and so forth.

(Transparency 14)

If α, β, γ equal *l*,m,n in turn, it is a pure-mode longitudinal wave. If the sum of $\alpha l + \beta m + \gamma n = 1$ it is a pure-mode transverse wave. This obviously means that the wave is propagating in one direction while the particles are vibrating perpendicular to that direction.

In summary, the general way of finding out what is going

on without making any assumptions in a linear regime is to solve this determinant, you get three v's, take each v in turn and plug it into this eigenvector equation and solve for the αs, βs, and γs, and then you know exactly what type of wave it is.

(Transparency 15)

There is something else that goes on. If you have an isotropic material, the energy flow goes in the same direction in which you are trying to send the wave, but in anisotropic material it does not do so except in some special symmetry directions.

(Transparency 16)

So you calculate the energy-flux vector by using the kinetic energy and potential energy and working that all through you obtain the formulas for the energy flux vector components.

(Transparency 17)

The energy flux vector is analogous to the Poynting vector of electromagnetic wave progation. Sometimes analogies work. The energy flux vector will tell you in which direction energy flows. In order to use it practically we need to take some time averages. What I usually do is I do not try to calculate the absolute magnitude, but just the direction in which the energy flows. The question is: does it coincide with the wave propagation direction or does it exhibit refraction?

(Transparency 18)

Let me give you an example. This is not the simplest way to do an isotropic solid, but if we treat the case of an isotropic solid and we get the same answer obtained by an easier method, then we can assume that we can use this method for more complicated crystal systems which exhibit anisotropy. These are the moduli, the λ's, and the eigenvalue matrix. Solve it and you obtain three real roots. For this isotropic case, one is unique and the other two are different from the first, but equal to each other.

(Transparency 19)

We know already from past experience the first real root corresponds to the wave speed of a longitudinal wave, while the other two correspond to transverse waves. But right here I do not know; I just know I have got three roots.

(Transparency 20)

It we now substitute each of these roots in turn back into the eigenvector equations we will be able to find the direction cosines of the waves associated with each of the roots.

(Transparency 21)

To make a long story short, I find that for the first root $\alpha = l$, $\beta = m$, and $\gamma = n$. This means that the particle directions of the wave move in the direction of wave propagation. Therefore, the wave speed for the first root corresponds to a longitudinal wave. For either of the other two I find that the direction cosines indicate that the particle direction is perpendicular to the propagation direction, indicating transverse waves. That is $\alpha l + \beta m + \gamma n = 1$.

Q: When you talked about examining the α, β, and γ, you could determine whether it was a longitudinal or transverse wave, is it possible to get linear combinations of the two?

A: In general, that is what you get. They are called quasi-longitudinal or quasi-transverse waves and that is what I am going to go into. You usually get quasi-waves in most cases in anisotropic media, except along symmetry directions or directions where there are particular values for the density and elastic moduli for a given material.

(Transparency 22)

If you calculate the energy-flux vector direction cosines to determine the direction of energy flow. I picked a, b, and c for them. In isotropic solids and in certain symmetry and special directions in anisotropic solids, the energy will flow along the propagation direction. However, in most directions in anisotropic solids the energy flows in directions different from the "propagation directior" for all three waves, i.e. refraction occurs.

(Transparency 23)

A summary of this technique for determining the basic features of elastic wave propagation in solid materials is as follows: You select the material, the density, and the elastic moduli. You select the direction of interest with direction cosines l,m,n. When you do this be careful. My first student, who did these type of calculations using a computer, made a mistake. He picked all three l,m,n. This was a big mistake, because no computer is exact and will crank out crazy answers if you do this. You have got to remember the sum of the squares of the direction cosines for any vector must equal to 1. The proper procedure is to pick any two of the

direction cosines and let the computer use the above relation to calculate the other one or else you get strange answers. You think you have discovered some new phenomenon and all you discovered is you did not do it right.

(Laughter)

So you substitute this in the eigenvalue determinative equation, solve for three real wave speeds, substitute each wave speed in turn in the eigenvector equations, solve for the particle displacement direction cosines, and finally substitute into the energy-flux vector component equations and solve for the energy-flux vector direction cosines.

(Transparency 24)

Since all of this will be in your handout I am not going to go through all of it here but tell you what the answers are. For linear elastic waves in a solid, if the solid is isotropic, you do get pure-mode longitudinal and pure-mode transverse waves for all directions -- isotropic means the same in all directions.

If you put a longitudinal wave transducer on top of a block of isotropic materials the energy flows straight down underneath the transducer. If you put a shear-wave transducer on top of a block of isotropic material and rotate it, you will get the same wave speed for any angle of rotation and in all cases the energy flows straight down underneath the transducer.

If you go to anisotropic materials, it is much more complicated. What happens, in general, is the waves are not pure modes, meaning they are neither pure longitudinal nor pure transverse, and the energy flux always deviates from underneath the transducer.

(Slide)

In the handout I do some calculations of these things. Since I told you I am more or less experimentally inclined, I wanted to make sure this was true, despite the classical calculations that were done -- I guess Lord Kelvin may have been the first one -- so there is nothing new in what I have told you.

I looked through the literature and found this Schlieren picture of a transducer up on the top of a block of quartz. This happens to be along an axis of six-fold rotation symmetry. With the Schlieren technique you take a point source of light followed by a lens where you expand the point source to get a parallel beam going through the transparent material under investigation; in this case, the block of quartz. You use a very similar lens on the opposite side of the specimen to bring

the light back to a point and then block that light off exactly as you can with what is called a spatial filter. One way of doing it is to put a photographic plate there, make an exposure, go develop it so it gets black, and put it back in the exact place where the exposure was made. Another way is just to take some India ink and make a bunch of ink spatters and move the plate around until you block of the exit light spot. Then when you send the ultrasonic wave through the specimen, the strain field of the wave changes the optical index of refraction locally, causing the light which was blocked to form a bright field image on a screen placed behind the opaque spatial filter.

The other thing I should say, since some of you talked about the Schlieren technique previously. For a given amplitude of the ultrasonic wave, it is very easy to perform the experiment in a gas. It is more difficult in a liquid and it is very difficult in a solid. But some Russians did this, Merkulov and Yakovlev, and they showed that, sure enough, the ultrasonic wave goes straight down if you are propagating in essentially an isotropic-type direction.

(Slide)

If you propagate the ultrasonic wave along a nonsymmetry axis, this is what you get. This photo shows the results for the case of the same transducer along a different axis in the quartz crystal and you see you get refraction of the ultrasonic beam even for normal incidence. That never happens in optics; for normal incidence the light beam always goes straight down. If you go off of axis you can get refraction in optics. So here is one place right where the optical analogy does not work.

Q: How big is the quartz crystal?

A: I do not remember, I will have to go look it up, but we did some similar Schlieren photos ourselves and we used specimens about one inch cubed or two inches cubed. Bigger is a little bit better, because you do not want to hit the edges of the specimen with the ultrasonic beam.

Q: What happens to your reflected or transmitted wave?

A: It goes back on itself.

Q: The wave fronts are in the same direction as they were in the original ones, so it is not that the wave is bending and traveling in this direction, the wavefronts are still horizontal.

A: Not in this case. Let us say this is almost a longitudinal wave, actually a quasi-longitudinal one. I have got some drawings to show you. I should have done that first.

(Transparency 25)

This is a general case. If you have a longitudinal wave transducer on top of an anisotropic material propagating in a non-symmetry direction, the ultrasonic beam will deviate off axis. In addition the particle displacements will have a small transverse component. In the case of a shear wave transducer, you will have energy flux deviation, and a small longitudinal particle displacement component. That is what you get in general. They are called quasi-longitudinal and quasi-transverse waves.

Q: Why is there so little spreading of the beams?

A: That is why I used the Russian photographs, because when we did it, we got diffraction beam spreading. In fact, in our pictures it looks like the ultrasonic beam spreads and after hitting the bottom keeps on spreading when reflects back up. I do not know just how the Russians did it.

Q: The frequency of your transducer may be very high compared to its diameter -- therefore diffraction from the transducer may be very low because the frequency might be quite high. That would also explain why the beam is so bright. The higher the frequency, the more scattering you get and the brighter your beam would be.

A: In the paper that is in your notes there is another picture, which I do not have a slide of, where they moved the transducer over a little bit, turned up the power, and they generated two waves, a quasi-longitudinal and a quasi-transverse, from the same transducer. When they turned up the power even more and they got a multitude of ultrasonic beams in the specimen. If you do not watch out when you increase the poser, you will go into resonance and then you get all kinds of complicated waves in the specimen.

DR. HARGROVE: You can also adjust Schlieren's spatial filter beam stop to get almost anything that you want. You can be in a range where it will be dark field for those side lobes -- you can adjust things to get the picture you want.

(Slide)

A: This shows the sort of experiments people did to measure elastic moduli before they started doing pulse-echo or travel time ultrasonics. This shows a photo of the result of shooting a helium-neon laser beam through the three-fold rotation axis of a quartz crystal while resonating the whole crystal with ultrasound. This was taken with the specimen between crossed polarizers, because the strength of the interaction of the light with transverse waves is different than it is with longitudinal waves. By adjusting the analyzing (output)

polarizer you can make both images have equal intensity on the recording film. Measurement of the diameters of the rings as a function of orientation yields the elastic wave speeds, which combined with knowledge of the density, yields the elastic moduli. These are essentially inverse velocity surfaces (slowness surfaces) in that particular plane. The inner one is for the longitudinal wave, while the outer one is for the shear wave.

(Slide)

This shows an experiment that we did where we put a transducer on the top of a quartz crystal. First we did it on top of isotropic fused quartz cube and used a laser probe system to scan across one face. The ultrasonic wave would diffract the light and we set up a photomultiplier detector on the first-order diffraction on the far side. If the interaction is linear, you can do it on either side. If it is nonlinear, there will be more light on one side than on the other, so you can measure nonlinearities that way. Over a distance of about 8 cm from the transducer you can see the beam is pretty well confined. This was half-width of the beam measurements as a function of distance from the transducer. The z-axis in quartz is also isotropic and so the difference in beam profile patterns between the isotropic results and the z-axis results is caused by material differences between quality of the fused quartz and the crystal. But when we examine a nonsymmetry direction as shown in the next slide we get a lot of diffraction beam spreading. However, depending on the particular material and directions you can also obtain focusing rather than spreading.

(slide)

You probably know, that people who are doing thermal pulses, i.e. phonon arguments rather than just wavespeeds, that the main thing they observed was focusing or hot spots in the material, and that was because of energy-flux deviation and diffraction focusing. You can get either hot spots or expansion of the beam. I am trying to tell you where all this works practically. People who are doing work in the field I work in, most often called nondestructive evaluation, use ultrasound a lot for looking for defects and flaws in metals, ceramics, and composites. They even try to measure their size so they can use fracture mechanics analysis to determine if the flaw size is critical or not. But they calibrate all their transducers in a tank of water. All it means when they do that is it works in a tank of water. It has no bearing on the beam profile inside an optically opaque material, particularly one that is anisotropic. No one has ever found any way of doing that except in transparent objects, because you cannot see the beam inside. That is a challenge. Maybe now that we are making microsensors and optical fibers, you could put a bunch of fibers inside an opaque material at different levels and

sell them to people who are doing practical work so they could really calibrate their transducers. It would not be the same material they are using but at least it would better than doing it in a liquid.

Q: Could you put a radioactive element in your medium and track, say, a Doppler shift or something?

A: I guess you could, but no one has ever tried it.

(Slide)

That is why I want the values for the elastic moduli and why I am interested in a resonance technique, hopefully to determine the elastic moduli for different materials fast. Normally, the way we do it is we take the material of interest and we cut parallel faces and perform pulse echo or through transmission ultrasonic methods to measure the travel time of the pulse, then measure the length to compute the wave speeds, which combined with density measurements yield the elastic moduli.

The maximum number of elastic moduli that I can determine in any one direction in an anisotropic material is three. This is because one quasi-longitudinal wave and two quasi-transverse waves can propagate in any given direction. As I stated before, in any direction in an isotropic solid there can be only two different waves, but since there are only two unique linear elastic moduli, measurements in any direction is sufficient.

If I propagate ultrasonic waves along an axis of twofold rotational symmetry in a cubic crystal I can obtain three different wave speeds. One from a pure mode longitudinal wave. The second is from a pure mode transverse wave with one particle displacement direction, while the other is from a second pure mode transverse wave with a particle displacement direction ninety degrees out of phase with the other transverse wave. For all three waves there is no energy flux deviation and therefore the waves go straight down. So in cubic crystals it is pretty easy. With two transducers you can measure the three linear (second order) elastic moduli.

With other directions in a cubic crystal and with all other crystal systems it is not so simple and often you have to go along a nonsymmetry axis where you get quasi-waves as well as energy-flux deviation. I was going to say earlier, when I first discussed energy flux deviation, that the predominance of ultrasonic wave energy flux deviation brings into question exactly what is meant by Snell's law in anisotropic materials? I think it means that one must consider where the energy flows goes. That is where the beam goes, that is where the energy is. In any case, we can now routinely calculate the energy flux vectors -- we did these iron crystals as shown on the next

slide, because steel is used in a lot of practical applications. These are not to the same scale because the results for the quasi-transverse waves are twice the size they should be, so you get quasi-longitudinal and two quasi-transverse. If the material was isotropic, the figures would be spheres, all three of them.

If I showed you the results for aluminum, which is very nearly isotropic, all three figures would also be very nearly spherical. The metal that is most closely elastically isotropic is tungsten. It would give spheres all the time. Molybdenum is next, but the most practical one that people use a lot is aluminum, so aluminum can be assumed to be pretty much isotropic most of the time. As you have seen iron is not quite so nice.

(Slide)

It becomes very difficult as to how to plot the energy-flux deviations. In the handout I show aluminum equal velocity contours for the quasi-longitudinal and quasi-shear waves and they have different types of contours. If you look where any two contours cross, you can determine the orientation of the crystal, so you do not need to use x-ray diffraction. In fact, the third wave will let you check the measurement you got with the first two. This shows the magnitude of the energy-flux deviation in different directions. Since it is cubic, there is no deviation in the cube face directions, [100]. But as you come out in the cube corners, [111], it is pretty big, so you get these funny-looking contour patterns, and we always debate with the students how to best portray them. You see it is worse for the shear waves than for the longitudinal waves.

(Slide)

If you do the calculations for a uniaxial composite, which simulates a hexagonal crystal, you get these type of figures. Again, the scale for the quasi-transverse waves is not the same as for the quasi-longitudinal.

Q: Tell us, again, is this the energy?

A: This is the energy, the magnitude of the average energy in different directions. This would be the magnitude of the energy along the x, y, and z directions.

Q: And you are picking that up with --

A: No, it is just calculation, just the computer computations.

Q: What do those words say? I cannot read them back here.

A: This says "energy-flux deviations for quasi-longitudinal waves, energy-flux deviations for fast quasi-shear waves, and

energy-flux deviations for slow quasi-shear waves."

Q: This is graphite/epoxy, this one?

A: This is graphite/epoxy computations. The fibers are all in one direction.

Q: I thought the one before that was what that was.

A: The one before that was iron. I have them for a lot of materials. It is a standard computer program, so we can compute them easily if you give me the linear (second order) elastic moduli and the density.

(Slide)

What are the things that we do with it? Let us assume now that we understand the fundamentals and also assume that we are working in a linear elastic regime. One thing is to look at the attenuation of the wave. What is it inside some material that might cause energy loss? There are two basic mechanisms, which is true also in fluids: absorption and scattering.

(Slide)

We have seen this solution for the wave equation before, but now because of attenuation we have to add an exponential decay, which is the standard procedure.

(Slide)

In solids you have got all these complications, and when you work with composites you also have fibers and interfaces between the fibers and layers of prepregs, and even viscoelasticity, which I am not going to talk about today. So ultrasonic examination of practical solid materials can become very complicated and extremely difficult to portray theoretically. So you try to conduct an experiment where you hold everything constant, except the property you want to measure.

(Slide)

Normally, the way we look at attenuation is to look at a pulse-echo pattern, where we bounce the pulse back and forth in the material a large number of times. This is the rectified envelope of radio-frequency pulses showing the exponential decay pattern.

(Slide)

This shows something that people asked me to do once with ultrasonic velocity measurements. One time the Navy either thought or did -- Logan knows better than I do -- whether they

did it -- to build a submarine pressure hull made of titanium. They were really worried about the weld region because melted titanium picks up oxygen, which is detrimental to welds. So we did some work where we performed a velocity scan, moved the transducer right along the weld region and looked at how the wave speed changed. As you can see it changed drastically in the weld region as well as in the heat-affected zones on both sides of the weld. So we could give them a measurement of how some aspect of the mechanical properties changed because of that.

(slide)

Q: Is this a time-of-flight test, do you have the transducers on both sides?

A: No this is pulse-echo from a single side. The reason is, if they are going to build a submarine, they are going to look at it only from the outside of the hull. They are not going to go inside and put transducers there, at least with submarines of present design.

The energy-flux deviation from the wave normal presents a problem when you are endeavoring to measure attenuation. Suppose I had limited the size of the specimen? With the previous example I had a big specimen. This is a schematic. If I had a limited size specimen and I thought the beam was coming straight down and the beam hits the side, then I get mode conversion. Longitudinal waves can convert into two shears if it is an anisotropic material or a shear can convert into a longitudinal or whatever. So I would not know whether I am having mode-conversion losses -- not really losses, it is in there, but not returning back to the transducer while in the receiver more -- or whether I have attenuation caused by some material alterations. So you have to be careful with the size of the specimen relative to the size and frequency of the transducer and the beam spread.

(Slide)

This shows some work done by Bethlehem Steel. When they are rolling sheets in a rolling mill; they would like to monitor the grain size. What they do still is to get some steel off the mill, roll it up in big rolls; cut some pieces off periodically, go in the metallurgical lab; put the pieces in some acid to attack the grain boundaries so they can see the grain size; and look at it with a microscope. Now they have a computer that scans the thing and gives them the average grain size, but it is still energy intensive and time consuming. If they did not make it right, they are in trouble.

By measuring the attenuation of ultrasound while it is being rolled, the surface condition does not make much

difference unless it is terrible. They can look at the scattering from the grains, and they found, that 5 MHz did not do much but at 10 MHz it was quite sensitive. This is not the world's best accuracy, but I am now going to show you the one when they went into the factory. It was terrible. That is probably why they do not use it. Anyway, that is the sort of thing people would like to be able to do, to control processes. I am talking just about acoustical techniques today, by using what we know about wave propagation or attenuation.

(Slide)

In my opinion, no one has a good theory of plastic deformation. We have already heard good talks on linear elastic wave propagation and nonlinear elastic wave propagation. However, no one spoke about plastic waves -- you did do shock, but that is still nonlinear, because I guess you do not have plastic deformation of fluids. But in solids you do have permanent plastic deformation.

In my opinion, the reason that no good theory for plastic deformation of solids exists is that most structural solid materials are filled with so many defects, whose distribution is not known, that you cannot calculate [a priori] the plastic behavior. You can make statistical assumptions on the one hand and you can do individual single-crystal measurements on the other side, but in the middle is where all the real stuff is, so it is a problem that still has not been solved.

When Professor Sachse was my graduate student, we thought, somewhat naively, that we would try using different polarizations and different directions of ultrasonic waves and map out the dislocation field, which is the main defect that causes plastic deformation in metals. We failed, but we did see some interesting things. The thought was if we had a single crystal of aluminum oriented for single slip, and you can determine the slip systems which are well-known for face-centered-crystals like aluminum. Using x-ray diffraction, you can determine what the orientations are. If they were dislocations moving on a slip system in a specific direction and if we had shear waves, transverse waves, with a large displacement component in the same direction, we would get a lot of interaction between the moving dislocations and the shear waves.

(Slide)

So we did this. This is the stress-strain curve for a crystal and for a second one possessing the same orientation relative to the tensile axis. You can see that the reproducibility of the stress-strain curve from one crystal to another of the same orientation is pretty much the same. When we sent the shear waves in a direction which did not get much interaction with the moving dislocations, i.e. not much

attenuation, we did not see anything much. But when we sent the shear waves in a direction so we got a lot of interaction with the dislocations, we saw a yield point. No one had ever seen evidence of a yield point in aluminum before.

Q: You got a 1% strain? That is amazing. How did you do that?

A: It is an aluminum crystal. You can extend it to 66% strain before it fractures. It is very soft.

Q: What was the vertical axis in those top ones?

A: The change in attenuation. It is hard to measure absolute attenuation. That is something else I should point out. The reason that is not often measured practically, but primarily in research labs, and then people usually do not report absolute numbers, is if you put a transducer on a metal with some coupling grease, or whatever you want to use, the inability to reproduce that coupling is usually greater than any attenuation a material will see during some change. So there has been a great desire to get away from the coupling, and I will talk about that in a moment.

Q: What is usually the cause of the increase in your attenuation?

A: In a metal under plastic deformation? It is the dislocations, vibrating, unpinning and moving.

(Slide)

I was trying to do fundamental research, but the funding situation changed and my sponsor said, "I don't want you doing single crystals, I want you to do aircraft aluminum alloys, fatigue problems in aircraft." So we set up an eccentric cam fatigue machine, put an ultrasonic transducer on the end of the aluminum alloy specimen in such a manner to generated a guided wave. The diameter of the transducer was large compared to the cross-section of the specimen, because most fatigue cracks start at the surface and I wanted to cover the surface as well as the volume.

(Slide)

We measured attenuation that way and we found, during the cycling, a spread of attenuation during the fatigue cycles in reverse bending. This technique originated with sonar. You send a pulse out and get a reflection from the far side of a specimen. If there is a crack in the middle, you will get another reflection halfway down, so you can tell where the crack is and get some idea of the size of it based on the amplitude of the reflected pulse. So this is where they would

see an indication of a crack, but prior to this in time the attenuation started to increase drastically. By cutting the specimen open and looking at it metallographically we found a lot of microcracks were forming before they grew to form a macroscopic crack.

(Slide)

People are getting very concerned now because they are making composite materials which are highly anisotropic. Most of them have been thinned for the aerospace industry and they do not know anything about energy-flux deviation. They never had any need to, because even if it deviated in a thin specimen, it did not make much difference. But the problem arises as follows. Suppose there is a crack over here inside the composite and the transducer is on the surface but not directly over top of the crack. Now when the energy-flux deviates so that it reflects from the crack you would think the crack was located beneath the transducer when it is not. When you look for it where you think it is, in order to repair the composite, and cannot find it the composite remains damaged.

One of my colleagues gave a talk not too long ago to a group of people who are doing research for ARPA on a proposed submarine pressure hull made of graphite/epoxy composites, and he came back all upset and said, "They don't believe me." I said it is not a religion, they will believe you, because it is a true physical fact. We are trying to educate the practitioners in the field about what can happen and what does happen.

(Slide)

The technique most often used in practice is what is called ultrasonic C-scan. That is just because they have three techniques, A, B, and C. The results are often plotted as the pulse amplitude versus time, called an A-scan. This one is a C-scan; the image is similar to radiography. The transducer and structure under investigation are either immersed in a tank of water, or in air and water squirters are used. In either case the transducer moves across the surface of the structure in a raster scan resulting in a two-dimensional image of the internal features of the structure.

(Slide)

This is a C-scan image we made on a metal-matrix composite recently. The reinforcing fibers were fairly large in this case. You can see all the fibers and inhomogeneities in the composite. These are some bubbles in the piece up here. This is pretty standard procedure. C-scans are now made of almost all aircraft structures, nuclear reactor pressure vessels, bridges, all kinds of large structures like this -- and people, too, in hospitals.

(Slide)

We started doing some work with ultrasonic CAT scanning (I am not going to go through the details of it). This was for a neoprene rubber specimen and there were holes in it and we could do pretty well imaging the holes sending and receiving the ultrasound through the sides.

(Slide)

This is an approximately isotropic composite. An x-ray CAT-scan image and an ultrasonic C-scan image of the same specimen is shown. Note that the ultrasonic image does not show a circle like the x-ray image does. What happens is, the fibers are not all aligned in the right direction in the material used to make the composite and the fibers tend to guide the waves. Such results gives you some idea of how bad the material was made in the first place.

(Slide)

We were doing some work on the submarine program and making some thick composite specimens. It turns out that we could not find anyone who had ever done compression-compression fatigue on composites, usually they do tension. If you are going to have a cylindrical shell that is like the fuselage of an airplane, when it is pressurized it is loaded in tension, and when it lands it comes back, again, so it is tension-relaxation, tension-relaxation. I cannot think of anything but a submarine that gets loaded from the outside and back out, again. So we managed to convince ARPA to give us some money to do some fundamental tests in the lab, fatigue testing. We put notches in the specimens so we would know where the damage occurred and then we tried a lot of nondestructive techniques.

(Slide)

Initially, we calibrated our specimen on a McDonnell Douglas test sample that had holes drilled to different depths in it. This shows the false color C-scan image of the piece, showing the depth of the holes determined with ultrasound time-of-flight measurements (B-scans).

(Slide)

Next we put an ultrasonic transducer and an acoustic emission transducer in the region of the stress concentration notch. An acoustic emission transducer is one that usually runs around 100 to 150 kHz and it picks up things like fibers breaking. It is like a stethoscope. The ultrasonic transducers usually operate in the 1 to 10 MHz range.

(Slide)

This shows that attenuation increases during fatigue testing, but does not give as early warning as it did in the aluminum before failure.

(Slide)

Acoustic emission did a little better, because things are kind of popping all the time during fatigue testing of these composites.

(Slide)

I told you we wanted to do noncontact ultrasound. This is some work one of my colleagues performed in conjunction with the National Institute of Standards and Technology. They melted the surface of a steel plate using electron beam scanning and wanted to know how far down they changed the properties and microstructure. So he used a neodymium YAG pulsed laser to generate a conglomeration of surface waves on the surface of the piece. Then an optical interferometer was used to pick up the normal displacement of the surface. Since there is no contact with the surface other than light you get absolute velocity and attenuation data, because there is no losses from anything other than the steel. Moreover, the laser systems possess a flat frequency response from DC out to about 60 MHz, so they are very broad-band sensors. They are not as sensitive as piezoelectric ones, however, about an order of magnitude less.

In this way one can capture all the ultrasonic data and put it into a spectrum analyzer to determine the wave speed and attenuation as a function of depth, because, the depth of penetration is frequency-dependent. So you obtain information about how much the surface of the steel was modified by the electron beam.

(Slide)

For NASA/Lewis we built an optical laser interferometer with fiberoptics inside. We tried to make it look like an ultrasonic system so that people familiar with ultrasonic equipment would use it. It is a box with some knobs on it and a cable coming out which normally has the transducer on the end of it, but now it is an optical fiber with a lens on the end. The other thing that happened is that we found we could get very small neodymium YAG lasers (about three inches long) from military surplus. They are range-finder lasers. They cost only about $2000 or $3000, so you can do laser ultrasonics now pretty economically, without building a big laser and turning the power down.

(Slide)

Here are a couple of applications of laser ultrasonics. People grow various types of single crystals, semiconductor crystals and others, by the Czochralski technique. It is a pulling technique, where you dip a seed into a melt of the material and you pull it out slowly and, if you do it right, the interface of the crystal with the liquid will have the atoms formed in the same crystal structure as the seed you started with. But sometimes things go wrong at the liquid-solid interface. The question is, is there some way you could monitor the interface during crystal growth?

(Slide)

Wagner and his students set up a system where essentially they used a pulse laser as an ultrasonic generator and an optical laser interferometer as a detector and generated non-contact ultrasonic waves at the liquid solid interface. Using this system they made measurements of wave speed through the crystal near the liquid solid interface as it was rotated.

(Slide)

This slide shows the resulting A-scan waveforms at 90° intervals about the circumference of the crystal. Note that three of them are the same, while one is different indicating that a differently oriented crystal was starting to grow at the location of the different waveform. Although the general shape of the crystal looks something like an overweight carrot this technique still works even though the diameter varies along its length. I can illustrate it better by just talking to you about a cubic crystal. Suppose it was a cubic crystal growing in a cubic direction. Even if the diameter varies, the crystal will possess fourfold symmetry about the growth direction, because it is a fourfold symmetry direction. If a new grain grpws, it knocks out that symmetry. When such a stray grain is detected using this technique, the crystal grower can reverse the pulling direction and remelt the stray grain, start growing again, and eliminate the stray from the final product.

(Slide)

People cannot accurately measure the temperature of materials in a furnace at high temperature. They put thermocouples on the surface but measure only the surface temperature. With optical pyrometers it is even worse, because the emissivity of the surface varies. You get oxides, sulfides, and all kinds of other films on it. So Wagner and his students have made accurate temperature measurements up to about 2000°C by using a laser pulse to generated ultrasonic waves on the back surface of a steel bar in a furnace and an optical laser interferometer to detect the ultrasonic wave transmitted through the steel bar on the front surface. Since the travel time of the wave is directly proportional to the

temperature of the steel bar, such travel time measurements yield the average temperature of the steel bar. However, no one has yet been able to measure the temperature distribution through materials parts in a furnace.

(Slide)

When you hit something with a normal pulse laser you get a broad-band frequency response. The signal-to-noise ratio of an optical detector is inversely proportional to the bandwidth, so what you would like to have is discrete frequencies, as we do with normal ultrasonic transducers. We can change transducers with different frequencies. They worked out a system where they built a laser that would generate discrete frequency ultrasound. What this looks like is this. If you use just one pulse, you get an envelope of frequencies. If you use multiple pulses in a phased array, you can get the discrete frequencies. A current student is now using 10 of these small lasers pulsed in any sequence he wants, so he can get discrete frequencies, but it can also do beam steering, as you would with a phased array of piezoelectric elements.

(Slide)

We proposed to the steel industry at one time to make a system like this using fiberoptics so we could scan a large hot casting.

(Slide)

We are collaborating with DuPont, who is making composites for various applications, to use non-contact ultrasound for process control. One way of making a composite is to take prepreg tape and use a robot to wind the tape on a mandrel to make the composite part. They would like to know, as they are making the composite part, does the top layer of the tape bond to the second layer and so forth? We are proposing to use a non-contact laser ultrasound system to monitor the tape placement process. In the event there is an unbond, there will be an ultrasonic wave reflected from the interface. So you just keep looking for a reflected wave as the tape placement process proceeds. If a reflection is found, this indicates the bond is not good. So the tape can be unwound and placed a second time.

(Slide)

An alumnus runs a company that makes stereo equipment and he asked us to do some work on his stereo speakers. Wagner used double-pulsed holography to make a hologram of the speaker, superimposed that on the image of the real speaker, and if you vibrate the speaker you can see the vibrational modes. We get written up periodically in stereo magazines for this wonderful achievement.

(Slide)

A similar technique was used to obtain information about a graphite epoxy plate. The plate was about a quarter of an inch thick and about 6 inches by 4 inches on the sides. A double-pulsed laser was used. First, a holographic image was made of the front surface of the plate, then superimposed on the real surface and hit it in the back with an infrared beam and, by thermal elastic effect, it generated the waves going out and then made the second image, so you catch it as you do with a stereo speaker, only it is a double-pulsed technique. The fibers are lying like this, so the wave speed is fast along the fibers (if it were isotropic, it would be circles), and here is retardation because there is either gas or air or something up in this slow wave speed, so you get a full-field image fast without any scanning.

(Slide)

Shortly after we did that, we were asked to do some work for Desert Storm. One had to do with land mines. If they cannot find them, they want to blow them up. They blow them up by shooting a missile down range and dropping out an explosive liquid in aerosol-like form, then exploding it, and hopefully blowing up enough path that they can run their armored vehicles down. It does not always work as well as I said it, and so now they are working with solid explosives, but they do not know what the ideal dispersion of the solid explosives is. They asked if we could record the shock waves from these explosives. It is the very same technique we did before. We put an explosive particle on a needle and made a holographic image of it. We used a double-pulsed laser, infrared, to explode it, and then, with a certain delay time, make a second picture. I do not have any pictures now, but now they are doing two needles with two particles on them to see how one sets off the other in changing the distance. By using optical delays with a graded beam splitter that has different coatings on it so that the intensity of each of these reflected pulses is the same, they get 10 images now in about 20nsec intervals, so the shock fronts develop and set off the other shock fronts.

(Slide)

This is acoustic emission, where you just listen to a crack make some sound. Most people who do this get a burst emission like this and they get a continuous emission like this and a lot of commercial vendors are selling these systems, trying to make lot of money. It turns out that Lamb developed a theoretical expression for point-source loading of the surface, making the displacement measurement on the same surface. You get a P wave, a longitudinal wave -- a shear wave is hard to find -- but a big Rayleigh wave, on that surface.

(Slide)

Using a conventional piezoelectric crystal, you get a pattern like this. If you use a laser or capacitance probe, you get a little thing like this, which is a real waveform. If you increase the time record, this is the real waveform from that displacement, and this is what people are analyzing industrially. They pick a threshold -- here -- and it is like counting all these wiggles. They have very fancy computer programs that say things like "the rise time on No. 64 divided by time between 13 and 14" has some significance, and they get pattern-recognition techniques and everything else, and it is all nonsense.

(Slide)

So what is the problem with acoustic emission? If you know what you are listening for, you can find it. Again, it is like sonar. If you know what a certain Russian submarine sounds like, you can find it -- you can tell the difference between that a whale. They know what the sounds are. But they do not know what the sounds are from materials being deformed, so we made a very quiet tensile machine, we made very small microtensile specimens, and use a laser interferometer to pick up the vibrations really close to the gauge length.

(Slide)

Since it is high fidelity and long frequency response, we got signals from zero up to about 60 MHz. Around in here, these are around 100 to 150 kHz -- these are some around 5 to 10 MHz.

(Slide)

The reason we made them so small is so we could look at them with an optical electron microscope to see what made the sound. The 100 to 150 kHz were made by dislocations moving.

(Slide)

Just before the specimen fracture we got the high-frequency ones. This is the fracture surface. Using an electron microscope, the graduate student noticed that there were ceramic particles cracked in the holes and, like all good graduate students, spent a long time counting all of them, and saw that there was a one-to-one correlation with the high-frequency sounds and the holes. It turns out there was an interesting application for this. We gave this talk at some meeting at the time the Challenger blew up because of the O-ring seals. A guy at Aerospace Corporation near Los Angeles got the idea that maybe something like this could be used for checking the quality of new seals. He got some glass microballoons and smashed them in a compression machine and

measured the acoustic emission that took place. Microballoons were different sizes, so they broke at different loads, and he got an almost linear plot of force on the microballoons versus the sounds he heard. Then he put them in some grease and told the technician, "Use my grease, not yours." The guy put them in and then he listened with his stethoscope, essentially, and said, "It's okay, they're going to work fine," and so far they have. I thought that was kind of a cute application.

(Slide)

I was going to do some stuff with nonlinear, but I am not. The only thing nonlinear used in a practical situation is to measure residual strength. They like to call it residual stress, but no one ever measured stress; they only measure the action due to a stress or force. They measure strain and then they use linear elastic equations to convert that to stress. The technique they use is called acoustical birefringence. They put a shear-wave transducer on a specimen and measure the wave speed, turn it 90° and measure the wave speed, and claim the difference between the two speeds is a monitor of residual strain and stress. But remember what I told you about energy-flux deviation. If you do that in a general material, the energy flux is going over here for one of these polarizations and over here for another wave. The wave is not going through the same part of the material. But they still use it anyway.

The next-to-last thing I am going to talk about is high-power ultrasound. The Russians did a lot of work with the influence of high-power sound on metal deformation, plastic deformation. To set up an experiment we put a 20 kHz ultrasonic horn at one end of a metal specimen and a 5 MHz ultrasonic specimen to measure attenuation on the other.

(Slide)

The reason you want to do this is to investigate the influence of the high-power ultrasound on the stress-strain curve for the material. When you cut the horn on, the load required to deform it decreases. The Russians claimed they could do things like make tungsten wires behave in a ductile fashion by superimposing high-power ultrasound during the drawing process. Tungsten is normally very brittle and often breaks into pieces when pulled through dies. Titanium, when pressed into shape in a die, tends to spring back. The Russians also claimed that if you do superimpose high-power ultrasound correctly, it would not spring back. So we wanted to look into all of this.

(Slide)

We saw that when we superimposed the high-power ultrasound, that there was a change of attenuation and a change

in the wave speed, which means change in the elastic modulus, so we checked out the Russians pretty well. But it still was not believed in the Western world, because people who did it thought that the only thing that happened was that the specimen heated up. They said if it gets hot only because of the vibration, why not just use a furnace? We know how furnaces work, so why go through this ultrasound business?

(Slide)

The student doing the work put thermocouples on the surface, too, but he found that when he took the specimens out, or touched them sometimes, that he got his fingers burned on the thermocouples and not on the specimens. So that encouraged us to set up an infrared camera to image the heat distribution in the specimens during high-power insonation.

(Slide)

And just look at this one. These are dumbbell-like specimens, like dog bones, and this is where the grips go across. The horn is always on the bottom and these are false-color infrared pictures. When the heat goes in, it looks white like that -- that is when it is pretty hot. You see, before the heat gets past the grip, the thermocouples are sources of heat, so the thermocouples do not read the right temperature, because they are vibrating vigorously and absorbing the energy. So previous workers have been misled. I cut a notch with a little saw in one of the specimens and found heat caused by ultrasound at the notch. We did more work recently looking at what caused the heat and how the ultrasound interacted with defects in metal specimens.

(Slide)

This shows the infrared image of one specimen that was bent over like a hook just due to the application of high-power ultrasound alone.

(Slide)

This is it. Ultrasound alone did that. So we are trying to figure out just what is the mechanism. There is a grain boundary where the specimen bends. Also note that the bent end is also flared out. It is a really strange shape.

(Slide)

We combined some techniques to study things and one such technique is called x-ray topography. It is something like x-ray tomography or x-ray radiography, but it is not. We do this with x-ray diffraction using relatively low voltage (soft) x-rays, around the order of 10 kV. To do something dynamic, though, we need a lot of x-ray flux, and so we are using the x-

ray ring at the National Synchrotron Light Source at Brookhaven National Laboratory. It is a 2.5 GeV ring, so it possesses a lot of high energy we do not want but, still, there is also some in the low energy regime which is useful for x-ray diffraction.

(Slide)

I have a graduate student who works for a company that cuts quartz piezoelectric crystals for resonators. This is a Laue picture of a whole resonator in the x-ray beam and, instead of getting Laue spots, we get Laue wafers. Each of these pictures is due to one set of parallel planes in a crystal. They are like Laue spots, but recorded with a large x-ray beam from a synchrotron.

(Slide)

This shows that if he resonates it while he is recording the images he can see the vibrational pattern of the resonating quartz crystal. It is similar to what would happen if you took a photograph of me and I stood still but waved my hand in the plane of the photograph. You could measure the amplitude of vibration of my hand. So this is the fundamental resonance destroying the Bragg diffraction condition. The crystals are loaded with defects, by the way, but they have very high-frequency control.

(Slide)

He can study different modes in these crystals while resonating and he can do this for all the different planes in the crystal.

That completes my formal presentation, can I answer any questions?

Q: The pictures where you were showing the spheroidal deformation, you had slow, intermediate, and fast, those were all different wave speeds?

A: Those were computational calculations based on just those elastic formulas of how the wave speeds would be in different directions. One of them is for quasi-longitudinal, the other two were fast quasi-shear and slow quasi-shear. The second set of three figures, which resembled flowers, were the magnitude of the energy flux deviation from the wave normal in different directions. It is not easy to figure out how to portray that type of data.

Q: That monochromatic laser array, what kind of a beam pattern did that produce?

A: Most of the time right now it is being used as a single

spot. But depending on the lenses used, one could get a line or a circle or different shapes. So far the only thing we have been trying to do is produce as good a set of discrete frequencies as possible. The other thing is, the first discrete frequency generating laser one was very large -- it took up two of these tables -- and the one we have now is about half this table. The idea is to try to make it smaller. We are thinking of going commercial; we never did that, but everyone is thinking that way now, at least a lot of people are. Sometimes we have to go to more practical work to get funding even at universities. But I should say, in case that bothers some of you, I have seen it cycle two or three times now in my lifetime, so it will cycle, again, in my opinion.

Q: How do you extract information from those plots?

A: If you are talking about the one where I showed that dislocation motion could be detected with shear waves of different polarizations, that is because it is pretty easy to grow good single crystals -- not perfect -- but good crystals of high-purity aluminum. The slip lines in aluminum are very nice and regular compared with some other materials. It is better for face-centered cubic materials and some hexagonal materials. In a face-centered cubic single crystal, there are 12 different slip systems, but if you orient it right, you can make only one go at a time. So that is what we did, we oriented it so only one system is preferred. Then we set up the transducer to go this way or that way. We were trying to map out the dislocation distribution inside the crystal by using different wave modes, different angles, made spheres of aluminum but, the number of dislocations and there distribution is just too complicated. So we were not successful in doing it.

Q: How were the piezoelectric ceramics resonated?

A: By application of alternating current voltage to the electrodes on both sides of the quartz crystals. The electrodes are primarily around the edges with the middle left free. You can see them if you look closely at the picture.

Q: That is why they were resonating only in the small part?
A: That is correct.

Q: What is the difference between x-ray tomography and x-ray radiography?

A: In x-ray radiography you have a white beam, a spectrum of continuous wavelength radiation, incident on the object beneath which is the film. You have the x-ray source, object, film. The image on the film is based on just the absorption of the x-rays by the material, so it is mainly a mass absorption coefficient thickness product that does it. The thing we saw

is diffraction. If you take a narrow beam of x-rays and shoot it through a crystal, some of the radiation will be diffracted by the various planes in the crystal, just like the mirrored ball on Soul Train.

(Laughter)

If you stop that ball and shine a spotlight on it, you will get light reflected, but it will not be rotating. The angle of incidence is equal to the angle of reflection. So if you know the location of the light source and where it reflected, you can find the plane that made it by bisecting the angle between the light source and the reflected spot. It is an analogous situation in x-ray diffraction. That is not the exact physics, but the analogy is true. What it means is, when you see these multiple spots, that there are parallel planes in the crystal at different angles that cause constructive interference of the x-ray beams in different directions. With the synchrotron the beam is big enough and his specimen is about the size of a dime so can orient it so that the whole crystal is illuminated by the beam. So instead of getting small spots as normally seen in Laue x-ray diffraction cases, he gets spots the same size as the whole quartz crystal. The fine structure of these spots is indicative of how good the parallel planes are that made that spot. It is as if you cracked a glass and reflected light from it. You will see fine structure in the reflected imaged caused by the crack. What we are showing is the fine structure inside individual Laue spots. Thanks to the synchrotron large beam and high flux, we can look at fairly large objects that way. We are also doing the same thing with gallium arsenide wafers, but it has nothing to do with acoustics. We can scan gallium arsenide wafers inside the plastic carriers they come in and we can see all the dislocation defects in the wafers using the same system.

Q: If you rotate it, you will see different ones?

A: If you rotate it, you will see some different spots than you saw without rotating it, but you do not want to rotate it while you are doing it, unless you have a T.V. camera recording them, because the dislocation image is going to change each time. It is really much like rotating the ball and then trying to figure out where the mirrors are. You cannot do that, you have to stop it to do that. We do rotate it, though, and then stop so we can obtain ideal pictures.

Q: Do you have to have a high-frequency limit when you are using ultrasound to detect defects in metal crystals? You could not go to a higher frequency?

A: I do not know how high a frequency you could go. I think people have gone up to about 25 or 30 MHz. Most of the time, if we go any higher, we cannot propagate a wave. Attenuation in a lot of solids is so great that you cannot go to high

frequency. You would like to go to as high a frequency as you can if you are looking for defects, because then the wavelength is small and you can find things, but there is always a tradeoff of propagating the wave through and finding small things or not being able to propagate at all. There is always a compromise.

Q: You can measure the temperature of a material in a furnace by measuring the travel time.

A: Yes, because the temperature is a linear function of the wave speed. Pressure and stress are also linear functions of the wave speed. The first nonlinear term.

Q: This is a combination of longitudinal waves and shear waves, this is a combination of both and they have different wave speeds, so how do you deal with that?

A: I should have said that. When you hit an object with a pulsed laser, you get longitudinal waves, shear waves, and Rayleigh waves all at the same time.

Q: So how do you measure it?

A: Well, we know what the speeds are, because we can measure them when the specimens are at room temperature by other more conventional means, with contact piezoelectric transducers for example. In general the speeds go from longitudinal, which is the fastest, to shear, which is roughly half the longitudinal speed, to Rayleigh which is much slower than the shear. Then we can put it in a furnace where all wave speeds will change with increasing temperature, but still remain separate relative to each other. We only have to determine the time of the first arriving wave which will be the longitudinal one.

Q: So you measure the slight difference?

A: Yes, it is not a big difference.

Q: How accurate is this?

A: I do not know precisely. I will have to get you the paper. However, it is enough that the industrial guys are giving us money to do it.

(Laughter)

Q: Is ultrasonics used commercially for weld inspection?

A: Yes, it is used very often, that is correct.

Q: Is it always a pulse echo thing or do people sometimes try to do a standing wave across the thickness of the weld?

A: Always pulse echo. There are some people who have done acoustic emission, because the solidification process sometimes makes noise. I am not enthusiastic about that. I think the best is pulse echo. However, sometimes they use a technique called pitch-catch, where they send and receive at angles to the surface.

Q: At angles?

A: Angles, because often cracks are normal to the weld surface and to see the cracks they have to use angle beams. What they always do in industry is they make up models, take real pieces of whatever material they are working with, and cut slots in it at different angles, drill flat-bottom holes at different angles, all kinds of stuff, and they always calibrate on these things. If they do not do that, they get in a lot of trouble. And they have to check out the personnel who do the inspections on these calibration blocks before they will let them work on actual structures.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 15, 1995	Final Report 01 Feb 94 - 15 March 95

4. TITLE AND SUBTITLE

1994 Physical Acoustics Summer School

5. FUNDING NUMBERS

PE 61153N
G N00014-94-1-0468
TA 3126979

6. AUTHOR(S)

H. Bass, O. Anderson, S. Putterman, S. Garrett, G. Swift, B. Denardo, J. Maynard, P. Marston, R. Green

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

National Center for Physical Acoustics
The University of Mississippi
Coliseum Drive
University, MS 38677

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
ONR 331
800 North Quincy Street
Arlington, VA 22217-5660

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

Proceedings of the 1994 Physical Acoustics Summer School, Volume 1

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release
Distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This is the verbatim transcript of the 1994 Physical Acoustics Summer School. It has been edited by the authors and includes in volume II a copy of the viewgraphs presented by the lecturers.

14. SUBJECT TERMS

Physical Acoustics

15. NUMBER OF PAGES
1080

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
unclassified	unclassified	unclassified	