


AD A0 67355

DDC FILE COPY



12 NW

Semiannual Technical Summary

LEVEL

Information Processing Techniques Program

30 September 1978

Volume II:


Wideband Integrated Voice/Data Technology

Prepared for the Defense Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-78-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Approved for public release; distribution unlimited.

79 04 17 052

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Defense Advanced Research Projects Agency under Air Force Contract F19628-78-C-0002 (ARPA Order 2929).

This report may be reproduced to satisfy needs of U.S. Government agencies.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Raymond L. Loiselle

Raymond L. Loiselle, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

12

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

INFORMATION PROCESSING TECHNIQUES PROGRAM
VOLUME II:
WIDEBAND INTEGRATED VOICE/DATA TECHNOLOGY

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

DDC
RECEIVED
APR 18 1979
C

1 APRIL - 30 SEPTEMBER 1978

ISSUED 1 MARCH 1979

Approved for public release; distribution unlimited.

LEXINGTON

MASSACHUSETTS

ABSTRACT

This report describes work performed on the Wideband Integrated Voice/Data Technology Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 April through 30 September 1978.

ACCESSION for		White Section ☒
NTIS		Buff Section ☐
DDC		
UNANNOUNCED		
JUSTIFICATION		
BY DISTRIBUTION/AVAILABILITY CODES		
A		

CONTENTS

Abstract	iii
Introduction and Summary	vii
I. SPEECH CONCENTRATION REQUIREMENTS STUDY	1
A. Introduction	1
B. Access-Area Design Considerations	1
C. ETHERNET Model for Local Voice Access Area	2
D. Simulation Results	6
E. Discussion	7
II. ADAPTIVE VARIABLE-RATE PACKET SPEECH NETWORKING	8
A. Introduction	8
B. Feedback Issues	9
C. Inclusion of Data Traffic	13
III. PRELIMINARY EXPERIMENT PLAN FOR THE EXPERIMENTAL INTEGRATED SWITCHED NETWORK	15
A. Introduction	15
1. Purpose and Background for the Experimental Wideband Network	15
2. Summary of Preliminary Experiment Plan	16
B. Development Plan for the Experimental Wideband Network	17
1. Pluribus Satellite Interface Message Processor (PSAT IMP)	17
2. Earth-Station Interface (ESI)	18
3. Earth Station and Channel	19
4. Integrated Local/Regional Access Node (ILRAN)	19
5. Access Facilities	19
6. Gateways	21
7. Wideband Terrestrial Links	22
8. System Control and Monitoring	22
9. Subsystem Development and Installation Schedule	22
C. System Validation	22
1. Introduction	22
2. Validation Experiments	23
D. Advanced Systems Experiments	23
1. Demand-Assignment Multiple Access (DAMA) Techniques	23
2. Multi-User Packet Speech Communications	25
3. Advanced Switching/Multiplexing Techniques for Voice/Data Integration	27
4. Rate-Adaptive Communications Techniques	28
5. Routing	30
6. Conferencing	31
7. Internetting	32
References	34
APPENDIX - Adaptive Traffic Estimation Algorithms	35



79 04 17 052

INTRODUCTION AND SUMMARY

The goal of this program is the investigation and development of techniques for integrated voice and data communication in packetized networks which include wideband common-user satellite links. Specific areas of concern are the concentration of statistically fluctuating volumes of voice traffic; the adaptation of communication strategies to conditions of jamming, fading, and traffic volume; and the eventual interconnecting of wideband satellite networks to terrestrial systems. A major focus of the current program is the establishment of an experimental wideband satellite network to serve as a unique facility for the realistic investigation of voice/data networking strategies. This facility will be jointly sponsored by DARPA and DCA and will include four ground stations sharing a leased domestic wideband satellite transponder.

The current report covers work in three areas: a study of speech concentration requirements, simulation of a technique for adaptive variable-rate packet speech networking, and the definition and planning of a program of packet speech and networks experiments to be conducted in the wideband test facility. Our speech concentration requirements study has been broadened to include the design of a local voice network suitable for providing a large community of packet voice terminals with access to a wideband switching node. In this report, we present the results of an analytical and simulation study of a proposed random-access protocol for use in that application. Work in adaptive variable-rate packet speech networks has indicated that rate instabilities can occur as a result of probing actions by individual voice terminals. A phantom probing technique has been developed to overcome this problem, and simulation results confirm that stable operation has been achieved. Our experiment definition and planning efforts have resulted in an experiment planning document that will serve as a working guide for future efforts in the wideband network program. A portion of that document is included in this report. Since the experiment planning effort is jointly sponsored by DARPA and DCA, a similar description of the preliminary experiment plan appears in the 30 September 1978 Annual Report for the DCA Network Speech Program. The experiment plan document, which has been submitted under separate cover to DARPA and DCA, is more detailed with respect to schedules and specific requirements on program participants.

REPRODUCING PAGE NOT FILMED
BLANK

INFORMATION PROCESSING TECHNIQUES PROGRAM

WIDEBAND INTEGRATED VOICE/DATA TECHNOLOGY

I. SPEECH CONCENTRATION REQUIREMENTS STUDY

A. Introduction

Consideration of several candidate topological structures for local voice access area application has led us to favor a distributed geometry similar to that used in the ETHERNET. A major question remains as to the type of communications protocol that would be best suited for a local voice network, i.e., a centrally controlled structured technique such as TDMA, or a distributed random access method similar to the ALOHA or ETHERNET systems. A TDMA-based approach was presented in our semiannual report of 31 March 1978, along with access-area design objectives and concentrator/terminal functional requirements. In this report, we analyze a random-access form of local voice network protocol, based on the ETHERNET approach. Our conclusion, based on analytical and simulation studies, is that random-access techniques can be made to perform efficiently in the voice environment. The ultimate selection between the TDMA and random-access protocols thus has to be based on relative implementation complexities, hardware implications, and reliability and robustness concerns.

B. Access-Area Design Considerations

There are several fundamental considerations in the design of a local-access network. The most important is the nature of a speech terminal's output data stream. Framed vocoder algorithms (e.g., linear predictive vocoders) analyze segments of the input speech. Typically, the duration of these segments is fixed between 20 and 25 msec and succeeding segments may or may not overlap; the exact nature of the segmentation varies according to the algorithm chosen. Each segment of speech is represented by a small number of bits; these constitute one "parcel" of speech. Parcels are assembled into packets by the network interface. In an unframed vocoder, the speech is not segmented for analysis. Rather, each speech sample is represented by a fixed number of bits (e.g., 1 bit/sample in the CVSD algorithm). The network interface forms packets in this case from a fixed number of speech samples. To reduce channel utilization in either type of vocoder, data are not transmitted during silence intervals. Parcels representing silence, either sustained (e.g., the user is listening to a conversation) or momentary (natural pauses in continuous speech), are not transmitted. Consequently, the data stream representing speech has a different character depending on the portion of speech being transmitted. During talkspurts, packets are generated periodically; during silences, no packets occur. Therefore, the data stream tends to contain bursts of activity; the duration of the bursts and the interval between them is directly related to talkspurt/silence statistics. The statistics of the duration of talkspurts and silences in conversational speech are given by Brady.^{1,2}

The transmission of speech information is seldom unidirectional; usually, a conversation is occurring between two or more users. While one user is silent (his speech terminal is not transmitting), speech is being transmitted to his terminal for acoustic synthesis. This two-way aspect of speech communication is another consideration in the design of a local-access network. One implication of this consideration is the sharing of a limited bandwidth channel by

two conversing users. The data rate required to maintain a conversation is not twice the data rate of each speech terminal. Rather, it is approximately equal to the data rate used by one terminal during the active portions of speech. This saving is termed the TASI advantage and has been incorporated into the design of other speech communication protocols.³

The characteristics of the output data stream of a speech terminal vary according to the vocoder algorithm; it is likely that the local-access network would need to cope with several vocoder types. Consequently, it must accommodate differing packet sizes, vocoder data rates, and silence detection algorithms.

Models of data transmission schemes commonly used in computer network design do not apply to the transmission of packetized speech. Often, data transmission is modeled as Poisson-distributed packet arrival times with either fixed (e.g., individual teletype characters) or exponentially distributed packet sizes. The traffic generated by one speech terminal contains bursts of constant interpacket arrival times. Packets are usually of fixed length. Furthermore, data transmission models usually describe one-way communications; speech is inherently a two-way, interactive communication method. Consequently, the transmission strategies developed for data communication may be inappropriate for speech and should be reexamined with respect to the special characteristics of speech traffic.

In this report, particular attention is given to an ETHERNET-like transmission strategy for the local-access network. In this scheme, speech terminals are connected via a single coaxial cable to a speech "concentrator." The concentrator serves as the port to the integrated speech-data network. Access to the cable by a terminal for the transmission of a packet is obtained on a demand basis. The concentrator obtains access on a similar basis to transmit packets to individual terminals. The main issue in this study is whether this strategy (details of which are provided in the next section) can be successful in sustaining speech communication. Implementation issues, particularly hardware questions, are not addressed in this study.

C. ETHERNET Model for Local Voice Access Area

The ETHERNET communication strategy was originally conceived by Metcalfe⁴ to serve as an inexpensive method of providing access for computer terminals to a central computer. The terminals are connected in parallel to a coaxial cable which serves as a two-way communication link to the speech concentrator (Fig. 1). In the original ETHERNET, access to the host computer is controlled by a contention process. Each terminal having a packet ready for transmission




Fig. 1. Access-area geometry (ETHERNET).

first monitors the cable and waits until no activity is detected. The terminal then attempts to transmit its packet to the host computer. Each terminal can monitor the activity on the cable while packets are being transmitted. If the monitored activity during a transmission by a terminal does not match the terminal's transmission, it is presumed that at least one other terminal also attempted to transmit packets at the same time; this situation is termed a "collision." If a collision occurs, the terminal attempts to retransmit the packet at a later time. A collision can only occur if two or more terminals begin to transmit within a round-trip time on the cable. If another packet is generated while a previously generated packet is awaiting transmission, the terminal queues the packets. All terminals obey these general transmission rules.

At this level of description, the design chosen for the local access of speech terminals is very similar to the ETHERNET. The contention and retransmission strategy deserves special attention; it is this portion of the communications protocol which represents pure overhead. This overhead should be minimized in order to obtain efficient communication.

The contention algorithm works in the following way. Once a speech packet is generated and no activity is sensed on the cable, the terminal attempts a transmission at that time with probability p_x . If no transmission was attempted by that terminal, it waits the equivalent of a round-trip propagation time T_r . If no other terminal attempted a transmission during this T_r seconds, the terminal again attempts a transmission with probability p_x . If the terminal senses a collision, it ceases transmission, as should the terminal with which it collided. Each terminal waits T_r seconds for the cable to clear and then each proceeds. The terminal continues with this strategy until the packet is successfully transmitted. The concept is illustrated in Fig. 2.




Fig. 2. Speech terminal access strategy.

Assume that a successful transmission has just occurred. Let M denote the number of terminals having packets to transmit. After waiting one round-trip time (insuring that packet transmission has been completed), these M terminals enter into the previously described contention algorithm to determine which of them will next transmit a packet. If the probability that exactly one terminal attempts a transmission is denoted by p_a and all terminals are assumed to use the same value of p_x , then

$$p_a = Mp_x(1 - p_x)^{M-1} \quad (1)$$

The probability that n collisions occur before the contention is resolved is given by a geometric distribution.

$$P_r [n \text{ collisions followed by success}] = (1 - p_a)^n p_a \quad n = 0, 1, \dots \quad (2)$$

To minimize the time spent in contention, the terminals need to pick a value of p_x which maximizes the probability p_a . The value which maximizes p_a is $p_x = 1/M$; the resulting value of p_a is

$$p_a = (1 - \frac{1}{M})^{M-1} \quad ; \quad (3)$$

as the number of contending terminals M becomes large, the value of p_a in Eq. (3) approaches $1/e$ rapidly.

The interpretation of Eq. (3) is clear; each terminal should adapt its transmission probability p_x to the number of terminals attempting a transmission. Note that this is not the number of terminals in the system, but is the number of terminals with packets ready to transmit at that instant. For speech traffic, the instantaneous load will have periodic components which may be exploited in the adaptation algorithm. In contrast, data traffic is generally modeled as random without periodic components. The periodicity of speech traffic occurs for several reasons. In any speech access area, vocoders of the same type are likely to be present. Consequently, these contribute periodic components. In addition, a parcel period of about 20 msec is a common number for several widely used vocoder strategies. Therefore, dissimilar vocoder types tend to produce packets at approximately the same rate thereby resulting in nearly periodic speech traffic. This periodicity is not perfect however. As users switch between speech and silence, their packet streams will start and stop. This switching does not change the period of the traffic, but alters the character of the traffic within each period. Another factor is imperfections in the terminals' clocks; the phases of packet transmissions will drift slowly with respect to each other.

As all of the terminals in the local access area can generate traffic independently, a distributed control algorithm is required to estimate the traffic load. The "true" traffic load is not known at any point in the network; each terminal must make its estimate based on its observation of the cable activity. To take advantage of the quasi-periodic nature of the traffic, the estimation algorithm should only observe the cable at the times when it has a packet to transmit. The traffic load at one transmission will be highly correlated with the load at the next transmission time; the traffic between transmissions is less so. An example transmission algorithm incorporating these ideas is given in the appendix.

An upper limit on the number of speech terminals can be derived for an ETHERNET local access network consisting of a mixture of vocoder types. The characteristics of the i^{th} terminal type can be summarized by two parameters: P_i , the size of a speech packet in bits (including overhead), and R_i , the data rate of the vocoder algorithm. N_i denotes the number of terminals of type i . The time taken to transmit a packet of size P_i is given by P_i/C where C denotes the capacity of the ETHERNET cable and the communication hardware in bits/second. The time between packet transmission on the cable is given by $(n + 1) T_r$ where n is a random variable denoting the number of collisions occurring in the contention. The probability mass function of n is given by Eq. (2). The additional round-trip time T_r is included in the expression for the inter-packet time, as terminals are required to wait one round-trip time after a packet transmission before initiating another. The average amount of time spent in contention is the

expected value of $(n + 1) T_r$. Assuming p_a to be $1/e$, the result of this computation is eT_r . Consequently, the average amount of time taken to transmit a packet by a terminal is $P_i/C + eT_r$. The rate at which the terminal generates packets is given by R_i/P_i . Hence, the total fraction of the cable capacity which is used by a terminal is given by

$$\left(\frac{P_i}{C} + eT_r \right) \frac{R_i}{P_i} . \quad (4)$$

Summing this fraction over all vocoders in the system must result in a number less than unity if the system is to work. If this sum is greater than one, speech data will be arriving faster than it can be transmitted, which results in ever-increasing queue lengths and delays; this situation will be termed as an "unstable" network. Therefore, to obtain stable transmission, we require that

$$A = \sum_i \left(\frac{N_i R_i}{C} + \frac{N_i R_i}{P_i} eT_r \right) < 1 . \quad (5)$$

This derivation assumes that each terminal has made a perfect estimate of the contending traffic. The derivation of this bound on the allowed values of N_i , P_i , and R_i did not explicitly assume that speech terminals do not transmit during silent speech intervals. It was implicitly assumed, however, that while a terminal is silent, the concentrator is transmitting speech to it. Consequently, there exists activity that can be attributed to that terminal even in silence. The left side of Eq. (5) is termed the "activity" and is denoted by A . The actual utilization of the cable for the transmission of packets corresponds to the first term in Eq. (5) and is denoted by ρ .

$$\rho = \sum_i \frac{N_i R_i}{C} . \quad (6)$$

Table I summarizes some values of vocoder parameters which are consistent with Eq. (5). Because terminals do not have precise knowledge of the number of contending terminals, the time taken to resolve a contention may be longer than the optimum value used in the derivation of Eq. (5). One might therefore expect the closer A is to 1, the more difficult it becomes in a physical (or simulated) implementation of the network to maintain stability. Note that the number of each vocoder type increases with packet size (Table I). Since the amount of time spent contending for the cable is not a function of the size of packets, longer packets mean the transmission of more data for a fixed overhead. The network thus becomes more efficient. In this case, the allowable number of users N_i can increase to maintain a constant activity level.

An expression for the average delay experienced by a typical user in transmitting a packet is easily derived. As all users are of equal importance, the expected number of trials (contentions and packet transmission) required before a terminal can transmit a packet is equal to the average number of users waiting to transmit. Therefore, letting Δ denote transmission delay (service time minus packet transmission time), the expected value of Δ is given by

$$E[\Delta] = E[M] \left(\frac{\bar{P}}{C} + eT_r \right) + \sum_i \frac{N_i}{2} \frac{P_i R_i}{C^2} \quad (7)$$

TABLE I THEORETICALLY ALLOWED VALUES OF N_i ASSUME: $C = 1$ Mbps, $T_r = 20$ μ sec (Approximately 1 km Cable Length) Three Vocoder Types: 2.4, 4.8, 16 kbps					
Packetsize Equals One Parcel (20 msec of Speech)					
N_1 (2,400)	N_2 (4,800)	N_3 (16,000)	Total Number	A	ρ
175	0	0	175	0.9	0.42
0	120	0	120	0.9	0.58
0	0	48	48	0.9	0.77
100	50	0	150	0.88	0.48
50	30	25	105	0.95	0.66
Packetsize Equals Two Parcels (40 msec of Speech)					
N_1	N_2	N_3	Total Number	A	ρ
235	0	0	235	0.88	0.56
0	140	0	140	0.86	0.67
0	0	52	52	0.9	0.83
125	65	0	190	0.87	0.61
65	45	25	135	0.95	0.77

where $E[M]$ denotes the expected number of waiting users and $\bar{P}$ denotes the average length of a packet.

$$\bar{P} = \frac{\sum N_i P_i}{\sum N_i} \quad (8)$$

The second term in Eq. (7) results from the fraction of time a newly generated packet encounters the transmission of a packet rather than a contention interval or a silent time. In this case, the terminal must wait an average of one-half the packet transmission time (P_i/C) and the probability of finding a packet of this length being transmitted is $N_i R_i/C$. An expression for the average number of queued users $E[M]$ was not derived.

D. Simulation Results

A computer simulation of the ETHERNET local access area was run to confirm the analyses and to test various adaptive load-estimation algorithms. For purposes of the simulation, a particular set of parameter values was chosen which were thought to be conservative compared with a physical network. Scaling of the results to other parameter values is straightforward. The cable was assumed to have a data rate (C) of 1 Mbps. The round-trip transit time (T_r) was 20 μ sec, which corresponds to a cable length of about 1 km. The simulation was run for various

combinations of vocoder types, and the contention algorithm described in the appendix was used. Three vocoder types were incorporated in the simulation; data rates of 2.4, 5, and 16 kbps were used. Packet sizes corresponded to either one or two parcels of speech (either 20 or 40 msec of speech data). The relative packet generation times of the terminals were selected randomly at the beginning of each simulation but remained fixed throughout the simulation. Activity measures of 0.95 or less could be tolerated in the simulations without instabilities; in those situations where activity measures greater than 1 were attempted, instability always occurred. Consequently, the restriction found in Eq. (5) is at least grossly valid.

Figure 3 depicts the average number of users waiting to transmit packets. These measurements correspond to empirical estimates of how $E[M]$ depends on A . If these empirical estimates are used in Eq. (8), the predicted values of $E[\Delta]$ correspond to the values obtained in the simulation.




Fig. 3. Average number of queued users as a function of activity level A .

Equation (7) predicts that the average waiting time $E[\Delta]$ is approximately proportional to the packet size. For the packet sizes used in the simulation, the empirical waiting time was not more than 2 msec. Consequently, the average waiting time is small compared with the interpacket time of an individual user (e.g., 20 msec for one-parcel packets). Therefore, the probability of a speech terminal being required to buffer two or more packets is low. The simulation confirmed that when A is less than 0.95 no speech terminal was forced to buffer packets longer than an interpacket time.

E. Discussion

The efficiency ϵ of the ETHERNET access strategy can be measured by the ratio of the utilization ρ to the activity A ($\epsilon = \rho/A$). Typically, ϵ ranges between 0.5 and 0.9; the larger

values occur when large packets are used. This observation corresponds to a decrease in the amount of time spent in contention when compared with packet transmission times. However, increasing the size of packets corresponds to increasing the delay in speech transmission. Packetization delays can be tolerated if they are not a significant fraction of the delay encountered by transmission and reception in the integrated speech-data network. This consideration establishes an upper bound on the size of packets.

The distinctive aspect of the access strategy described here is the impact of the characteristics of the vocoder as a data source. Because of silence detection, a sharing of bandwidth between the transmitter and receiver can be achieved. Because the vocoder generates data having an inherent periodicity, the speech terminal is able to predict with precision the time at which the next packet arrives. This observation has significant impact on the contention algorithm used to gain access to the cable: observations of the cable traffic are made only at packet generation times. If packets were generated in a random manner (e.g., Poisson), the time of arrival of the next packet would be unknown to the terminal, thereby requiring it to monitor traffic over long periods of time.

This study has indicated that an ETHERNET-like topology for a local-access network can support many speech users having different data rates and packet sizes. Furthermore, the TASI advantage is exploited to increase utilization within the context of a random-access network. This increase in utilization occurs explicitly for those situations when $\rho > 0.5$ (see Table I for some examples). In the ETHERNET, queueing (i.e., buffering of two or more packets) occurs seldom. Consequently, the control requirements placed on the interface between the vocoder and the cable are not severe. These general results imply that the ETHERNET random-access strategy is a viable alternative for a local-access network, subject to the practicalities and economies of the hardware implementation. The latter are the subjects of our continuing efforts in this area.

II. ADAPTIVE VARIABLE-RATE PACKET SPEECH NETWORKING

A. Introduction

In the previous semiannual report, initial work on a simulation of a rate-adaptive network strategy based on an embedded speech coding technique was described. Continuing work over the past six months has been directed mainly toward achieving an understanding of the effects of end-to-end feedback strategies and toward the development of a flow-control strategy which would allow the network to support a data load in conjunction with the speech load. A new strategy for end-to-end flow control has been developed which not only improves the overall efficiency of the network but also alleviates sudden dropouts which occurred in the absence of feedback, given certain traffic load conditions. In addition, a network strategy for link-sharing between data bits and speech bits was developed which gave an overall improved performance, in terms of percent utilization of the links, than was realized in the absence of data. This report overviews major developments over the past six months. A more comprehensive and detailed description of the network simulation and results is being published separately.⁵

As described in the previous report, the network configuration consists of a central node through which pass 16 paths, one from each of four "sender" nodes to each of four "receiver" nodes. The number of conversations on each path is fixed for a given run, but because users make use of TASI there are statistical variations in the actual traffic load as a function of time.

The traffic matrix, or speaker load distribution, is an important parameter of the system. For most of the discussion in this report, the assumed traffic matrix will be as follows:

$T(ij)$ = number of conversations on path (ij)

	20	40	60	80
$T =$	20	40	60	80
	20	40	60	80
	20	40	60	80

That is, sender links are all equally loaded with 200 conversations on each, whereas there is a linear increase in the number of conversations on the receiver links from 80 ($j = 1$) to 320 ($j = 4$). The assumed vocoder is one capable of synthesizing speech at nine different rates from 1,200 to 10,800 bps, in increments of 1,200 bps. This vocoder may be somewhat unrealistic compared with what may actually be feasible in terms of current vocoder technology. However, it is probably true that if the network cannot be made to stabilize with such a vocoder, it is unlikely to do any better with one with a more restricted repertoire of possible synthesizer rates, and correspondingly fewer available priority levels.

The network simulation had initially been implemented in the C language on the PDP-11/40 facility. However, computational requirements were such that even an overnight run yielded only a few seconds of simulated time. Hence, the decision was made to reimplement the simulation on the Lincoln Digital Voice Terminal (LDVT). The LDVT version transmits a number of parameters of interest to the PDP-11/40 facility for interpretation and display. A version of the LDVT program was also created which operates in real time in conjunction with a real-time embedded-coding vocoder residing in the two Lincoln Digital Signal Processors (LDSPs), with the analyzer in one, and the synthesizer in the other. The vocoder was assumed to be transmitting bits over one of the 16 network paths, and the LDVT actually performed the stripping operation according to the network results, for real-time listening.

B. Feedback Issues

In the absence of end-to-end feedback, analyzers would always transmit all priority packets, the network would strip off various lower priority packets as needed, and the synthesizer would reconstruct speech at the maximum rate the network could support. With end-to-end feedback, the synthesizer would inform the analyzer of the current received rate, and the analyzer would respond by trying to match the transmitted rate to the reported received rate. Bits would then be stripped at the source rather than at an intermediate node in the network, thus alleviating the overall network load and allowing other users of early links in the path to potentially realize a higher received rate.

In the previous semiannual, two initial end-to-end feedback strategies "continuous probe" and "periodic probe" were discussed. Both of these methods were found to be inadequate, but for different reasons. With the continuous-probe strategy, users always transmit at the next higher rate than the reported received rate, and hence can immediately respond to a decrease in the network load. However, this quick response is achieved at the cost of having to always transmit more bits than are actually being received. Unless the traffic matrix is extremely imbalanced, the method yields little or no improvement in received rate over that realized without feedback.




Fig. 4. Transmitted and received rate vs time for periodic-probe feedback strategy, for path $i = 1$, $j = 1$.




Fig. 5. Schematized plot of transmission rate vs time for two paths sharing first-sender link and percent utilization on first-sender link to illustrate instability phenomenon with periodic-probe strategy.

The periodic-probe method was an attempt to realize greater gains, but it exhibited undesirable instability problems, which were not apparent until after the simulation had been translated into LDVT code, such that longer simulation runs were possible. With this method, users transmit at the reported received rate. However, if the transmitted rate and received rate on a given path remain equal for a sufficient time period, users on that path up-probe by transmitting at the next higher rate for a short time interval. If the probe bits were not received, transmitters drop back to the original rate; otherwise they continue transmitting at the higher rate.

The results for this method for a path which includes the lightly loaded ($j = 1$) receiver link are shown in Fig. 4. In the figure, both transmitted and received rate are shown as a function of time. Wherever there are two values for rate, the higher one is the transmitted rate. Transmitted rates and received rates for this method are usually equal, except in the case of an unsuccessful probe. This means that almost no bits are being discarded at the central node. The mean rate is significantly higher than that obtained with no feedback. However, there seems to be a pattern of a staircase-like rise followed by a rather sudden collapse in the received rate. Such rapid fluctuations in the rate are not likely to be desirable for producing high-quality synthesized speech; nor would they be expected to produce the highest possible mean received rate.

This pattern arose as a consequence of several factors which may be understood by reference to Fig. 5 which is a schematized graph of received rate vs time for two paths sharing the first sender link. The staircase pattern is the rate observed for the path to the lightly loaded receiver link; the other pattern is the rate observed for the path to the heavily loaded receiver link. In the latter case, probes are always unsuccessful, due to the heavy load on the receiver link. Probe bits do get through the sender link, however, and since a large number of users are probing, they represent a heavy additional load on that link for the duration of the probe. At $t = 0$, the simulation initializes with all links accepting packets of all priority levels. There is a rapid initial collapse in the rate, as queues build up, and users respond by immediately lowering their transmission rate to the low crisis rate on the link. Once the queues spill out, the link is underutilized, and hence probes are generally successful. The result is that on the average the load on the link steadily increases. At the beginning, when the link was only being utilized to 60 percent of capacity, the short probe by users of the fourth receiver link did not pose a problem to the first sender link. However, when the load on the sender link was very near capacity, the same probe by the numerous users on the heavily loaded receiver path introduced a large excess load on the sender link which pushed its input rate far in excess of capacity. The sender link had to discard bits, but it could not discard the probe bits because they were of high priority. Hence, it had no choice but to collapse the rate of those few users directed to the lightly loaded receiver link. These users responded by setting their transmission rate to the artificially low rate imposed by the temporary probe. They would then begin their slow climb back up to 100 percent of capacity, and the cycle would repeat.

The problems, then, can be enumerated as (1) users probe simultaneously; (2) users respond too quickly and too completely to a drop in the network rate, and too slowly to an increase; and (3) fruitless probes introduce an excess load that interferes with valid transmissions. The asymmetry in user response is the main contributing factor; the simultaneity of probes is only a secondary factor. In fact, if the probe is modeled as a ramp rather than a step function, the slow rise/quick collapse pattern still occurs. The problem occurs mainly as a consequence of transfer of control to the users subsequent to each crash. Users will then continue to up-probe until the load has exceeded link capacity, and a new crash will ensue.

To solve the problem, an entirely new strategy was developed for which user response is sluggish and symmetrical, users no longer act in unison, and fruitless probes are eliminated. To eliminate fruitless probes, a new probing strategy was developed in which the network actually writes into a field in all priority one packets, information concerning the highest rate that could have gotten through. The analyzer sends the packet with the number 9 written into the field (the lowest priority possible), and each node in the path replaces the number in the field with its own stripping priority p , if p is less than the number currently in the field. By the time the packet arrives at the synthesizer, the number in the field represents the lowest priority packet that could have been successfully sent, even if that priority is lower than the lowest priority the transmitter was currently transmitting. The receiver then sends the information in this field back to the analyzer, and the analyzer adjusts his rate accordingly.

For the new strategy, users no longer set their rate to immediately match the rate reported by the receiver. Instead, users are much more sluggish in their response to network changes. Furthermore, there is no longer an asymmetry in their response to a reduction in the rate the network can accept as opposed to an increase. Each user is allowed to change his rate at periodic time intervals, where the period, equal for all users, is on the order of 5 sec. If, at the end of his period, a transmitter finds that the network can accept a higher rate than that at which he is currently transmitting, he responds by increasing his rate by only one level. Likewise, if his rate is higher than that which the network can currently support, he decreases his rate by one level. Otherwise, he makes no change until the time of his next update.

A final modification is that the users no longer act in synchrony. One could imagine, in a real network, a strategy whereby each user's first period begins at the time of dial up. Since users will dial at random time intervals, this will result in a randomization of the times of rate change for the various users on a given path. Since it is not feasible in the simulation to treat each conversation as a separate entity, a simplification was made such that it is assumed that a certain percentage of the users have probed after the same percentage of the probing interval has been exhausted.

The results for a simulation run using the new strategy are shown in Fig. 6, for the same path ($i = 1, j = 1$) with the same traffic matrix (20, 40, 60, 80) as was used for Fig. 4. Where formerly only 20 sec of simulated time were shown, this time the run was carried out for 120 sec. Only the received rate is shown, and the fractional rates are mean rate per customer rather than realizable rates. It can be seen that the system is much more sluggish, with a steady-state condition being reached only after 35 sec. However, the received rate remains steady for the remainder of the run, a much more pleasing result than was obtained formerly. Furthermore, the mean link utilization throughout the network, after discounting bits discarded at the central node, is significantly higher than that obtained with no feedback (89.1 vs 72.5 percent). Overall performance is also significantly higher than that obtained with either of the other two feedback strategies.

If the imbalance is in the sender rather than the receiver links, feedback would be expected to reduce rather than improve the overall performance. Bits will be stripped by the heavily loaded sender nodes regardless, and feedback will only result in a more sluggish response to an easing up of the network load. However, it was found that without feedback there was a tendency for sudden rate drops to occur on paths including the lightly loaded sender link, as a consequence of a sudden increase in the load on receiver links due to a rise in the stripping threshold on the heavily loaded sender link. The effect is analogous to what was observed with the




Fig. 6. Received rate vs time for the path $i = 1$, $j = 1$, using sluggish feedback strategy. Rate is mean rate per customer rather than realizable vocoder rates.

periodic-probe feedback strategy, except that recovery is more rapid. With feedback, such rate drops did not occur, because the sluggish user rate adjustment acted as a smoothing filter to prevent sudden rate increases due to changes in network stripping thresholds. The overall utilization fell, as a consequence of the feedback, from 91.4 to 88.9 percent, but the sacrifice in performance was probably more than offset by the elimination of rate dropouts.

C. Inclusion of Data Traffic

It would appear intuitively that an embedded-coding vocoder could operate quite successfully in a network environment where the link was being shared with data. The expectation would be that an increase in the data load could be dealt with by dropping the vocoder rate, and a potential sudden dropout in vocoder rate due to a temporary overload could be averted by temporarily allowing data queues to build up.

The first issue that must be resolved in combining data and voice is a strategy for portioning out the link to the two types of users, who have different requirements in terms of delays and stripping. In particular, speech users cannot afford long delays, but are only minimally disturbed by losses of all but priority one packets. With data, on the other hand, the only restriction on queue buildups is available buffer space in the node, but none of the bits should be discarded.

A possible solution would be to give all data packets a priority of one, which will assure that they are not discarded. However, the network will have no way of knowing that data packets can be delayed, and hence may insist on rushing these packets through at the expense of lower priority speech packets.

An alternative solution is to keep the data in a separate queue, for which there are less strict requirements on maximum size. Fluctuations in the data load could then be smoothed out at the point of exit on the link by allowing the queue to expand and shrink as needed.

For the experiment, the data were modeled as a stream of packets arriving at exponentially distributed time intervals. All data packets were of fixed size, 1200 bits, the same size as the

speech packets. The capacity of all eight links in the system was doubled (from 0.4 to 0.8 Mbps), and the mean data load was set to equal exactly the additional capacity on each link.

The strategy for portioning out the link to speech and data need depend only on the data traffic, since speech bit rates can be adjusted by stripping off lower priority packets as needed. Although for the simulation the mean data rate is a fixed known number, in general each node would have to predict the data load based on some past statistics. In the simulation, the mean data rate is measured for a fixed time interval and the value obtained is used in conjunction with the data queue as an estimate for deciding how often to send data packets out on the link for the next fixed interval. Hence, the model is designed to be able to deal with a data load whose mean is an unknown variable rather than a known constant, and represents a more realistic system.

The strategy developed is as follows: the incoming data rate is measured over a 1-sec interval. Data are then allocated a portion of the link such that, if that rate were sustained, the data queue would dwindle to exactly zero by the end of the next second. This portion is then quantized (either by truncation or rounding) to the nearest $1/16$. Then, for each set of 16 sequential packets sent out (see Fig. 7), speech and data are sent alternately until the one with the lower allocation has met its quota. The other type is then sent exclusively for the remainder of the set of 16. If either queue is empty, the other type is sent, instead of allowing the link to remain idle.




Fig. 7. Strategy for portioning out link to speech and data.

An important issue is whether to truncate or round the data portion. With truncation, data queues are kept on the average at a larger size, and link capacity is consequently more fully utilized. On the other hand, the system has no reserves for averting a potential crisis when an overload condition is anticipated in the speech domain. In particular, dropouts due to temporary overloads were found to be much more frequent and much more drastic than was the case in the absence of data. The reason is that data queues build up when the speech is utilizing its full capacity, because data can no longer get speech spillover. The result is that the data demand a larger portion just when the speech is consuming its entire portion, and hence a cutback in the speech portion occurs precisely when it can be least afforded. The frequency of such dropouts can be reduced by setting the threshold for when to increase the speech bit rate at a more conservative level. However, even then, the severity of such dropouts, when they occur, remains much worse than was the case in the absence of data.

If the data are allocated at a slightly larger portion than they need, data queues will not build up, even when the speech domain is overloaded. It is then possible to avert a potential catastrophe in the speech domain by temporarily giving the data a cut in allocation long enough to spill out the speech queue.

The strategy which was found to work best was to use the generous data apportionment in general, but to automatically reduce the data allocation by 50 percent for a 100-msec duration whenever the stripping priority for speech was raised. The data queue is thus intentionally allowed to build up only in those particular situations when the speech queue needs to be spilled out. With this strategy, the same thresholds could be used for stripping as were used in the absence of data, without causing any drastic dropouts in rate on any of the 16 paths.

The network with a 50-percent data load using the above strategy was found to give better performance in terms of overall link utilization than was realized in the absence of data (91.5 vs 89.1 percent). This is a somewhat surprising result, since the data introduce increased variability in the load, which ought to result in a reduced performance. This effect is apparently more than offset by the feature that speech packets can take advantage of excess capacity in the data reserve, and vice versa.

III. PRELIMINARY EXPERIMENT PLAN* FOR THE EXPERIMENTAL INTEGRATED SWITCHED NETWORK

A. Introduction

1. Purpose and Background for the Experimental Wideband Network

A wideband integrated voice/data network is currently being developed under the joint sponsorship of the Defense Advanced Research Projects Agency and the Defense Communications Agency. Organizations currently participating in the program include: Bolt Beranek and Newman (BBN), Communications Satellite Corporation (COMSAT), Information Sciences Institute (ISI), Linkabit Corporation, M.I.T. Lincoln Laboratory, and SRI International. Several additional organizations will join the program as contractors are selected for various subsystems now under competitive bid. The network will provide a unique experimental capability for the investigation of systems issues involved in a communications facility which includes wideband satellite and terrestrial links and which carries large volumes of voice and data traffic. Areas for experimental investigation include:

- (a) Demand-assignment strategies for efficient broadcast satellite communications;
- (b) Packet speech communication in a wideband, multi-user environment;
- (c) Alternate integrated switching techniques for voice and data, with particular attention to the advantages and disadvantages of packetized voice communication relative to more conventional circuit techniques;
- (d) Rate-adaptive communication techniques to cope with varying network conditions;
- (e) Routing of voice and data traffic;
- (f) Digital voice conferencing; and
- (g) Internetworking between satellite and terrestrial subnetworks.

Some of these areas are of more immediate concern to one or the other of the two sponsoring agencies. However, the problem areas are so interrelated and complementary that all are of interest to both agencies.

* A more detailed version of this plan has been submitted as a separate document.

Recent efforts related to the current wideband network program include:

- (a) The DARPA Packet Speech Program in which the ARPANET was employed for development and demonstration of packet speech communication techniques;
- (b) DCA-sponsored development of hybrid (packet and circuit) switching techniques for voice and data;
- (c) The Atlantic Packet Satellite Experiment (APSE), jointly sponsored by DARPA and DCA, along with the British Post Office and the Norwegian Telecommunications Authority, which included the development of the Priority-Oriented Demand Access (PODA) class of satellite demand-assignment algorithms.

Existing networks such as the ARPANET and the Atlantic Packet Satellite Net do not have sufficient capacity to permit experiments on a scale large enough to realistically represent the multiple and varied user environment that will be typical of future military communications networks.

The wideband integrated voice/data network is intended to provide a more realistic environment for experimental investigation and demonstration of the advanced communications techniques listed above. The complete test bed will include wideband satellite and terrestrial links and switching nodes, access facilities including concentrators and terminals, traffic emulation modules, various monitoring and support subsystems, and gateways for satellite/terrestrial internetworking experiments. The implementation of major network subsystems will proceed over the next few years. The intent is to evolve the network in such a way that experiments can begin as soon as the first major subsystems are operational.

2. Summary of Preliminary Experiment Plan

A stand-alone document representing a preliminary effort at an overall experiment plan for the wideband network has been delivered to DARPA and DCA under separate cover. The experiment plan is expected to evolve and become more detailed as the network evolves over the next few years, and the plan will be revised and expanded accordingly. This section represents a somewhat abridged version of the planning document. The intention is to include all major areas which are covered in the planning document, but to omit some of the details of schedules, equipment requirements, and issues which require coordination among the program participants. These details are of less general interest and will be subject to changes as the plan is coordinated among the sponsors and contractors.

An important purpose of the preliminary planning is to describe a development plan for the network test bed. A description of this plan, including a discussion of the functions of the various subsystems, is given in Section III-B. The other primary purpose of the planning document is to define and categorize systems experiments. The earliest experiments, which will be directed at test and validation of basic system functions and capabilities, are discussed briefly in Section III-C. Finally, advanced systems experiments are the subject of Section III-D. A broad set of experimental areas is defined, and objectives and performance measures are described for each class of experiments. Scheduling for test-bed development and for systems experiments is touched on only briefly here. More detailed scheduling information is presented in the separate planning document.

An important role of the planning document is to stimulate further interaction and coordination among the program participants. To this end, a set of unresolved issues with respect to test-bed development, system validation, and advanced systems experiments is listed in the planning document. Discussion of these issues has generally not been included here.

B. Development Plan for the Experimental Wideband Network

The experimental network will include a wideband satellite network, a wideband terrestrial network, access facilities including concentrators and terminals, and internet gateways. The satellite net will include four earth stations with planned locations at Defense Communications Engineering Center (DCEC), Reston, Virginia; ISI, Marina del Rey, California; Lincoln Laboratory, Lexington, Massachusetts; and SRI International, Palo Alto, California. At least one of the terrestrial switching nodes will be collocated with a satellite station, but the locations of all terrestrial nodes have not yet been specified. A topology for the satellite and terrestrial nodes is shown in Fig. 8. For illustration purposes, the DCEC location is depicted as the site of both a satellite and a terrestrial node, as well as a gateway interconnecting the two. Locations of




Fig. 8. Topology for experimental wideband network.

the other three terrestrial nodes are unspecified. Figure 9 shows a projected configuration for subsystems to be available at a network location which is the site of both satellite and terrestrial nodes. The functions of each subsystem are discussed below.

1. Pluribus Satellite Interface Message Processor (PSAT IMP)

The satellite network communication protocols, including the Demand Assignment Multiple Access (DAMA) protocols for the broadcast channel, will be implemented in the PSAT IMP. The PSAT IMP is a flexible device that can be programmed to serve a variety of experimental




Fig. 9. Subsystems for experimental network.

needs. Hardware and software development of the PSAT IMPs is the responsibility of BBN. The Satellite IMPs developed by BBN for APSE carried out similar functions, but were implemented as single-processor machines. The single-processor implementation was sufficient to support the 64-kbps broadcast channel employed in APSE. The satellite channel in the current experiment will be designed to accommodate approximately 3 Mbps, and a multi-processor approach referred to as the PLURIBUS technology will be employed to satisfy this significantly higher throughput requirement.

The demand-assignment algorithms to be implemented in the PSAT IMPs are Time-Division Multiple Access PODA and Contention PODA. In TPODA, the reservation subframe is shared by fixed slot allocation, while CPODA uses slotted-ALOHA type contention in the reservation subframe. The PODA algorithms provide a flexible facility for supporting a variety of DAMA-type experiments. For example, the stream capability with fixed reservations can be utilized to emulate a fixed TDMA scheme for basic satellite channel tests. As in APSE, the PSAT IMPs will have internal capabilities for network experiments and measurements, consisting of fake hosts which can be activated to serve as artificial traffic sources and sinks, and other fake hosts which can be activated to collect and transmit cumulative statistics on specified performance parameters. As was the case in the SIMP-3 version of the Atlantic Satellite Net, the PSAT network will be implemented as a stand-alone structure, capable of functioning independent of the ARPANET. However, facilities existing in host computers on the ARPANET will be utilized for experimental purposes as appropriate.

2. Earth-Station Interface (ESI)

The ESI provides an interface between the PSAT IMP and a 70-MHz Intermediate Frequency (IF) line into the satellite earth station. Linkabit Corporation has responsibility for ESI equipment development. ESI modules include a packet/burst controller which accepts packets from

the PSAT IMP and forms bursts of digital data for transmission, a burst modem, and a command/monitoring unit. The ESI will be required to operate at rates from 386 kbps to 3.088 Mbps. A similar interface developed for APSE operates in the range 16 to 64 kbps, so that a significantly scaled-up capability is required. The ESI will allow a multi-rate transmission mode within a burst. Modulation modes of Binary and Quaternary Phase Shift Keying (BPSK and QPSK) will be provided, as will optional error-correcting coding modes.

3. Earth Station and Channel

The earth station and channel is to be purchased from a carrier on the basis of competitive bid. Earth-station equipment will interface with the ESI at 70 MHz IF, and will include low-noise receiving amplifier, high-power transmitting amplifier, and antenna. A leased channel on a domestic satellite is to be provided. To minimize interference and siting problems, operation in either the 12/14- or 18/30-GHz frequency bands is preferred. However, because of the current lack of satellites operating in these bands, it may be appropriate to provide initial capability at 4/6 GHz in conjunction with a phased upgrading to the higher frequencies. The earth station and channel will be required to provide sufficient signal-to-noise ratio to support bit rates up to 3.088 MHz at specified error rates.

4. Integrated Local/Regional Access Node (ILRAN)

The ILRANs are the switching/multiplexing nodes for the wideband terrestrial subsystem. A contractor is currently being selected by the DCA to study advanced integrated systems concepts including hybrid and packet switching techniques and to design a terrestrial network (including ILRANs) to allow experimental evaluation of these concepts. Assuming successful completion of the design study, four ILRANs will be constructed. The ILRANs will be designed to be interfaced with access-area user terminals, hosts, gateways, and concentrators. In addition, each ILRAN will have an associated module for generation of emulated traffic within the terrestrial network.

5. Access Facilities

Access facilities for the wideband network will include terminals, access areas, host computers, speech concentrators, and traffic emulators. The provision of facilities for multi-user voice access is an essential item in the wideband experiment, and will receive much attention over the first few years of the program. Much of the need for data traffic can be filled by traffic emulation, and access for real data traffic can be provided by standard host computers, as in the ARPANET. Voice experiments require subjective evaluation and thus are more dependent on access for real voice users. In addition, the design of voice terminals, access areas, and speech concentrators is in itself an important topic for investigation in the wideband program. A heavy emphasis on voice problems is thus expected during the early phases of access facilities development. The associated issue of whether access for data traffic should be handled separately or integrated with voice access traffic remains open at this time.

a. Access Area

The access area provides the physical medium for collection and distribution of terminal data. A topology and a set of protocols must be specified in the access area design. Lincoln Laboratory is currently studying (under DARPA sponsorship) access-area architectures with

particular attention to distributed approaches. During FY 79, Lincoln will design and validate a pilot access area capable of supporting a few real voice terminals as well as a substantial volume of emulated traffic. This pilot access area will be interfaced to the wideband network through a minicomputer host, to allow early speech experiments on the wideband network. It is expected that this pilot access area will be the forerunner of a wideband access area capable of handling 50 to 100 terminals.

b. Speech Concentrator

The purpose of the speech concentrator is to collect the digital voice outputs from individual voice terminals and form them into one wideband data stream which is presented to a network node. It also performs the inverse process of splitting the return stream for transmission to the terminals. The concentrator has the task of transforming the local protocols of the voice terminals into the line transmission or packet protocols required by the wideband network node. It may also perform functions such as silence detection and rate control. Initially, the concentrator will be connected to the PSAT IMP, with a software gateway in the concentrator performing the necessary protocol conversion. Later, as shown by the dotted line in Fig. 9, a concentrator/ILRAN connection will also be implemented. It should be noted that the ILRAN switch/multiplexer will be designed to include its own concentration function and be able to accept inputs from individual terminals as well as from concentrators. Voice experiments in the terrestrial network therefore might not require a speech concentrator. However, a concentrator/ILRAN connection will allow the access area and concentration facilities developed primarily for the satellite network to be utilized effectively in terrestrial network-based experiments.

The requirements for speech concentration have been the subject of a study during FY 78 at Lincoln Laboratory. The study has addressed specifically the separation of functions among concentrator, access area, and terminals. No specific schedule has been set for developments of a speech concentrator capable of handling hundreds of terminals, as will eventually be required in the wideband network. However, early speech experiments will utilize only a few voice terminals, and the concentrator function for these experiments will be handled in a standard minicomputer, such as a PDP-11. This limited-capacity concentrator will be referred to here as the miniconcentrator.

c. Speech Terminals

Speech communication experiments will form an essential part of the wideband-network program, and speech terminals of various types will be utilized as the program proceeds. Speech terminals include the speech algorithm processor or vocoder function as well as an access-area interface and suitable dial-up and ringing functions. Speech terminal development is not viewed as a primary focus of the wideband-network program. However, reference is made here to current efforts sponsored by DARPA under the Packet Speech Program which will provide support for the wideband experiment in the speech terminal area.

Figure 9 shows both dedicated and programmable terminals in the access area. Dedicated terminals will run fixed vocoder algorithms and interact with the network in a fixed way. Currently available voice processors for dedicated operation include wideband encoders such as Continuously Variable Slope Deltamodulator (CVSD) devices and a more limited number of narrowband processors such as the Linear Predictive Coding (LPC) devices used in APSE. A current DARPA-sponsored program at Lincoln Laboratory and Texas Instruments (TI) is directed

at developing a small, cheap narrowband voice processor which can be deployed in large numbers in the wideband experiment. A prototype of such a unit, which is based on Charge-Coupled Device (CCD) technology and implements a channel vocoder algorithm, will become operational during FY 79.

A limited number of high-speed programmable terminals will be used for working with newly emerging voice algorithms and experimental techniques for terminal/network interaction. Of particular interest here are terminals which can communicate at varying bit rates depending on network load. Voice algorithms of the embedded coding type are of special interest in this regard because of their ability to change rates rapidly. Such algorithms, which are currently under development at Lincoln and elsewhere, will be tested on high-speed programmable devices. Currently available programmable processors such as the Lincoln Digital Signal Processor (LDSP) will be able to support such experiments in a limited way. However, the complexity of some of the proposed new algorithms (e.g., embedded coding) is such that more powerful programmable processors will be needed in the long run.

A study of the voice terminal/access-area interface issue is currently being initiated at Lincoln. The design and construction of interfaces to support experimental work in this area will be carried out as part of this effort.

d. Access-Area Traffic Emulator

The function of the access-area traffic emulator is to simulate the traffic loading effect of many voice terminals in the access area. Since voice traffic flow control is envisioned as being implemented either in the speech concentrator or at the voice terminal itself, traffic emulators resident in the PSAT IMP cannot satisfy this requirement. Rather, the emulator must act via the access area so that the concentrator's flow control mechanisms are exercised. It is expected that the required access-area traffic emulation can be effected in software either in conventional PDP-11 type machines or in currently available high-speed processors. Lincoln has responsibility for implementing a software capability for access-area voice traffic emulation. The emulation will include call initiation/termination and the use of speech activity detection for each of the simulated speakers. The emulation software will include voice flow control techniques such as bit rate selection at dial-up, dynamic modification of vocoder rates during conversations, and the embedded coding technique.

6. Gateways

Gateways allow internet communication between the wideband network and other networks. Two gateways are shown in Fig. 9, although other internet connections will probably be implemented during the course of the experimental program. An example of particular interest is a gateway to the ARPA Packet Radio Net.

An initial capability for internet communication between the PSAT IMP and the ARPANET will be implemented by BBN in the form of a software module in the PSAT IMP processor. This internal PSAT "mini-gateway" will support transfers of experimental data and control between PSAT IMP fake hosts and TENEX-based facilities residing on the ARPANET. However, experiments involving real internet (data or voice) traffic traversing the ARPANET and wideband satellite net will require a full gateway capability similar to the PDP-11 configuration used in APSE.

The PSAT IMP/ILRAN gateway will be used for wideband terrestrial/satellite internetting experiments. Design of this wideband gateway ought to begin after the initial ILRAN capability has been established.

7. Wideband Terrestrial Links

When the ILRANs are eventually delivered to network locations, wideband terrestrial links will be needed to interconnect them. The earliest need for such links is projected to occur during FY 82.

8. System Control and Monitoring

System control and monitoring support facilities will be needed throughout the experimental program. The monitoring and control technology developed for APSE will be carried over by BBN for use with the PSAT network. A monitoring and control fake host will be implemented in the PSAT IMP processor, and communication capability between this fake host and TENEX monitoring and control programs will be established. The ILRAN contractor is required to develop a Network Monitoring Center (NMC) as part of the terrestrial network implementation. Provision must be made for the control and monitoring of satellite/terrestrial internet experiments. Options include physically combining the network control centers or establishing a communications path between them.

9. Subsystem Development and Installation Schedule

A tentative schedule for development and installation of the wideband experimental facility is presented in some detail in the separate experiment planning document. A key milestone date on that schedule - 1 January 1980 - is worth calling attention to here.

At that time basic satellite subsystems, including PSAT IMPs, ESIs, and earth stations, will have been installed at the first two network locations (probably ISI and Lincoln). In addition, an access facility including a prototype access area, one or two narrowband voice terminals, a traffic emulation capability, and a miniconcentrator will be available at one of the sites (Lincoln). This date thus marks the starting point for experiments with the wideband satellite subnetwork, including emulated data and voice traffic as well as live voice.

The terrestrial network switches (ILRANs) will be developed in a parallel program which extends through FY 81. Construction and initial testing of the ILRANs is due to be completed at the end of FY 80, and systems experiments with the four ILRANs interconnected within the contractor's test facility will proceed during FY 81. Delivery of the ILRANs to specified network locations will occur early in FY 82. This represents another key milestone in the network development schedule.

C. System Validation

1. Introduction

Experiments for the wideband network can be roughly divided into two classes: (a) system validation experiments which verify or measure the performance of a network component or function and (b) system concept experiments directed at exploring some specific issue or problem in integrated voice/data networking. The concept experiments motivate the building of a test-bed system and are designed to provide answers to important systems questions. The

validation experiments have as a goal performance verification of the test-bed facility on which the concept experiments will be carried out.

2. Validation Experiments

The preparation and execution of validation test plans are generally the responsibility of the subsystem contractors. However, the separate experiment plan document gives a summary review of the minimal objectives that should be satisfied by these validation plans. This review includes discussion of validation experiments relevant to the satellite subsystem, the terrestrial subsystem, access facilities, and gateways. The review is not given here since many of the items require specific coordination with individual contractors.

D. Advanced Systems Experiments

Advanced systems experiments are defined and discussed here. Seven classes of experiments are considered in Sections III-D-1 through III-D-7. These are:

- (1) Demand-assignment multiple-access strategies,
- (2) Multi-user packet speech communications,
- (3) Advanced switching/multiplexing techniques for voice/data integration,
- (4) Rate-adaptive communications techniques,
- (5) Routing,
- (6) Conferencing, and
- (7) Internetting.

It should be recognized that there is significant overlap among the experiment areas. This overlap can be used to advantage in that particular experiments will provide results in more than one area. For each class of experiments the following items are discussed:

- (1) Introduction and background,
- (2) Experiment of objectives,
- (3) Experimental approach and requirements, and
- (4) Performance measures.

The separate experiment plan document includes more detailed discussion of each experiment area and presents a tentative schedule for each class of experiments.

1. Demand-Assignment Multiple Access (DAMA) Techniques

a. Introduction and Background

Broadcast communication satellites have a unique potential to efficiently support general-purpose communications, including both data and voice, among a large and diverse set of users. Cost analysis studies have shown that substantial savings can be achieved with systems employing hundreds of earth stations. However, a prerequisite for the achievement of these savings is the development of flexible DAMA techniques that allow one to exploit the broadcast nature of the satellite channel. The APSE program has provided an impetus for the development of such DAMA techniques and a test bed for their evaluation. Techniques investigated in APSE include fixed-TDMA, round-robin TDMA, slotted ALOHA, and several variations of PODA. A basic requirement of these algorithms is the ability of the satellite earth station equipment to transmit and receive in a flexible burst-TDMA mode. The demonstration of this capability has been an

important achievement of the APSE program. Of the DAMA schemes tested, the PODA algorithms appear to be the most promising in terms of their ability to handle the expected mix of user requirements, including block (data) and stream (voice) traffic. However, the PODA algorithms also place the most stringent requirements on the DAMA processors in terms of processing speed and memory.

DAMA-related experiments in APSE have been limited in scope because of the restricted (64 kbps) bandwidth of the satellite channel. This restriction has prohibited full testing of the ability to handle mixes of voice and data traffic, and of the reservation-handling capability of the DAMA processor. The wideband satellite network will provide an environment for DAMA experiments with traffic loads more representative of expected user communities. The PODA implementation in the PSAT IMPs will provide an effective and flexible facility for experimenting with realistic traffic volumes and mixes as well as with various DAMA schemes.

b. Experiment Objectives

DAMA experiments will be aimed at testing and evaluating the effectiveness of various multiple-access protocols in environments typical of the satellite portions of future military communications systems. Experiments shall be directed at the investigation of techniques for (1) efficiently sharing the satellite channel capacity among many earth stations and (2) taking advantage of the differing statistics and transmission requirements of voice and data traffic to achieve efficient statistical multiplexing. For example, there is strong interest in investigating and developing schemes for achieving the TASI advantage by statistical multiplexing at the satellite in cases where only a few voice users are present at each ground station. The feasibility of achieving this type of statistical multiplexing, which may require very rapid variation in satellite capacity assignments, will be investigated. The attendant advantages and costs of highly responsive DAMA schemes should be compared with simpler schemes which offer slower variation in channel allocation. Another objective is to compare the effectiveness and processing requirements of distributed and centralized control of the channel assignments in order to select an appropriate mix of these control schemes for future systems.

Of major concern in all cases is the earth station equipment requirements needed to support the algorithms (hardware size and cost, software complexity) and the robustness of the systems with respect to earth-station error or failure, transmission channel errors, etc. An important problem area is that of verifying that selected DAMA schemes can actually function efficiently with large volumes of real traffic without developing lockup problems or other limiting effects.

c. Experimental Approach and Requirements

The PODA algorithm will be used as a flexible experimental tool for carrying out the DAMA experiments. Traffic emulation in the PSAT IMPs will serve as the primary means for generating the traffic mixes required to test DAMA performance and processing requirements under a variety of conditions. Real voice and data traffic will be combined with the emulated traffic as appropriate. The measurement fake-host facility in the PSAT IMP will be used for measuring the patterns of delays, lost packets, etc., for the various experiments. These data will be delivered to a TENEX host for analysis and display.

Two important sequences of experiments in the DAMA area are identified and discussed in some detail in the planning document. The first concerns techniques for the efficient handling of full-duplex speech users, in the case where many nodes are each supporting a small number

of speakers. The second sequence is directed at assessing the effects of stressing the DAMA algorithm under different dynamic response constraints.

d. Performance Measures

Performance measures for voice include delay distributions, percentage packet loss, and the efficiency with which the burst nature of speech can be exploited. Data performance measures include average delay, buffer requirements, and probability of packet misdelivery. For the overall system, the efficiency of utilizing the satellite channel resource is of prime importance. The processing cost of achieving high channel utilization for various user mixes must be determined. The robustness of the system is an important performance measure. The system must be able to accommodate synchronization loss or hardware failure at individual stations without seriously compromising overall network performance. In addition, temporary data overloads must be accommodated without catastrophic system failure. Finally, the system should exhibit fairness in its ability to provide the same quality of service to all users of equal priority without allowing individual users to capture a disproportionate share of the channel resources.

2. Multi-User Packet Speech Communications

a. Introduction and Background

Packet speech communication has become a subject of intense interest and activity over the past several years. Factors motivating this interest include: (1) The development and demonstrated advantages of packet techniques for data communications, (2) the potential of packetized techniques for increasing channel utilization by exploiting the bursty nature of speech transmission, and (3) the possibility for the integration of voice and data in a common system based on packet switching. The basic feasibility of packet speech communication was first demonstrated in experiments on the ARPANET under the DARPA Packet Speech Program. Developments under this program included voice protocols for packet networks and techniques for maintaining speech communication quality in the face of the delay dispersion inherent in packet transmission. Later, packet speech communication techniques for a broadcast satellite environment were developed in the Atlantic Packet Satellite Experiment.

A recent economic study conducted by Network Analysis Corporation (NAC) concluded that packet-switching is potentially the most cost-effective technique for the integrated switching of voice and data. Support for this conclusion will require a demonstration of the feasibility of packet voice on a large scale. However, packet speech experiments to date have accommodated only a few voice users because of channel capacity limitations in the ARPANET and Atlantic Satellite Net.

b. Experiment Objectives

A primary objective of the experimental program is the development and demonstration of packet speech techniques for a wideband, multi-user environment. Included in this objective is the investigation of the effectiveness of packet techniques in achieving efficient statistical multiplexing of voice transmissions from a number of users. A prerequisite for multi-user packet speech experiments is the development of appropriate access facilities including concentrators, access areas, and packetized voice terminals. In fact, the investigation of advanced access and concentration techniques is in itself an important objective of the current experimental program.

Distributed architectures for the access area, where terminals are interfaced to the concentrator through a common cable or bus, are of particular interest in this regard. In addition, the design of terminals and access facilities must reflect the need for eventual compatibility with privacy/security requirements. Experiments directed at demonstrating such compatibility should be included in the program.

c. Experimental Approach and Requirements

Packet speech experiments will be phased in a manner consistent with the development of the network test bed. The initial voice experiments will make use of two voice terminals in the same access area, interfaced to the PSAT IMP by means of a miniconcentrator. Loop-back tests through the PSAT IMP will be followed by tests where packets are transmitted up to the satellite and back to the same earth station. These tests will establish the effectiveness of the interfaces and protocols between the concentrator and the PSAT IMP, and will verify the dialing and signaling protocols which allow call setup through the satellite channel. Emulated traffic will then be combined with real voice traffic in order to determine the effects of loading the satellite channel near capacity. Fake hosts in the PSAT IMP will provide a source of emulated traffic characteristic of multiple voice and/or data users.

When a second set of access facilities becomes available, these experiments will be repeated with two-way conversations between sites. Access-area traffic emulation will be used to test multiple call-handling capability and flow-control strategies in the concentrator, and to establish the throughput limitations of the complete path from access area to access area. A series of experiments will be run to test alternate voice multiplexing strategies in the concentrator. Schemes for packet aggregation, where speech packets from two or more terminals are combined into larger packets for transmission on the wideband net, will be developed and tested.

These experiments will later be scaled up to include a large number of real voice terminals interfaced to the network through a wideband concentrator. Although the discussion of packet speech experiments here has been focused on the satellite network, a similar set of experiments will be run later in the program in the terrestrial and internettted environments.

Experiments involving privacy/security will follow the initial developments of a packet speech capability. A requirement for such experiments is the availability of real or simulated packetized privacy devices. The BCR class of encryption control devices should be considered *as a primary candidate for satisfying this requirement in the satellite channel and over wideband terrestrial trunks*. Access-area privacy is an important topic for which a variety of options need to be considered.

d. Performance Measures

Performance measures for packet voice include delay distributions and percentage packet loss and their effect on speech quality. The efficiency with which speech users can be statistically multiplexed is of essential importance. It must be determined to what degree the TASI advantage can be attained, and at what cost in packet overhead, switch processing requirements, and buffer storage. The call-handling capability and achievable throughput for the access area, concentrator, and satellite network need to be assessed individually to insure that a proper match is designed and resources are not wasted. Reliability and ease of use of the packet speech system will also be important concerns.

3. Advanced Switching/Multiplexing Techniques for Voice/Data Integration

a. Introduction and Background

The trend in military communications is toward all-digital networks which will serve large volumes of voice and data traffic. Problems of fundamental concern are (1) the economics of serving voice and data applications on a common integrated communications system and (2) the comparison of alternate switching technologies for integrated voice and data networks.

Switching technologies to be considered include advanced packet and hybrid (combination of circuit and packet) approaches, as well as more traditional circuit-switched methods. Analysis, simulation, and ARPANET experience indicate that packet techniques provide considerable advantages for the handling of interactive data traffic. The previously cited NAC study concluded that packet switching is potentially the most cost-effective technique for the integrated switching of heterogeneous traffic including voice, interactive data, and bulk data. Hybrid techniques were also shown to have significant advantages over all-circuit techniques. The key factor which separated packet and hybrid systems was the potential ability of packet systems to achieve approximately a 2:1 saving in bandwidth utilization (the so-called Time-Assigned Speech Interpolation, or TASI, advantage) for voice by avoiding transmission during silent intervals. The extent to which this saving can actually be achieved in a large, distributed network remains to be determined. Another issue to be investigated is the effectiveness of fast virtual-circuit switching techniques as compared with packet approaches in achieving channel savings during speech pauses. Finally, interoperability between advanced switching techniques and existing network strategies, and the problems and costs of transition, are subjects of concern.

The efficient integration of voice and data is probably the most central issue in the wideband program and pervades all the experiment areas. This section focuses mainly on the issues of voice/data integration in the terrestrial network or ILRAN system since it is here that there are two specific switching technologies, hybrid and all packet, which have to be examined.

b. Experiment Objectives

Recent analyses have indicated that advanced switching technologies, including packet and hybrid (combination of circuit and packet) approaches, have the potential for efficient integration of heterogeneous traffic types, with significant cost savings due to the sharing of switching and transmission facilities. However, the feasibility of such techniques has not yet been demonstrated on a large scale. A key purpose of the wideband program is to achieve such a demonstration, and to carry out an implementation exercise which develops the switching concept in sufficient detail to serve as a prototype for a future integrated military network. For all-packet systems, the development, demonstration, and cost analysis of packet voice communications on a large scale is a subject of key concern. For hybrid systems, the implementation and integration of packet and circuit switching into a common system must be demonstrated. These demonstrations will provide a sound basis for comparison between packet and hybrid techniques.

The objectives of this portion of the experiment are the following:

- (1) Provide a flexible test bed for evaluating different techniques of voice/data integration both within the terrestrial network and in conjunction with the satellite network.
- (2) Perform a series of experiments which will allow a quantitative evaluation to be made of the advantages and disadvantages of both techniques.

c. Experiment Approach and Requirements

The principal vehicle for examining the voice/data integration issues will be the ILRAN system consisting of four flexible switching nodes connected by wideband trunks.

The experiments will be conducted in two parts. The first part will consist of a three-year development and test program to be conducted by the ILRAN contractor at his facility. In the second part, the equipment will be integrated with the satellite network. Experiments in this second area are described in Section III-D-7 on Internetworking.

The initial three-year effort of the ILRAN contractor is in turn divided into three phases. In the first phase, the contractor will study the promising voice/data integration techniques including at least the packet and hybrid concepts. For each promising concept, the contractor will identify the key features which should be experimentally evaluated to assess their relative advantages and disadvantages. These experiment requirements will then be used as a guide for designing a flexible experimental switching network test bed to carry out the experiments. The contractor will also develop a comprehensive, general test plan describing the concept experiments to be performed on the terrestrial network.

In the second phase, the contractor will build and test the system designed in the first phase. He will also elaborate on the general concept test plan to develop detailed acceptance and concept test procedures. In the third phase, the detailed concept test plan will be carried out at the contractor's facility.

d. Performance Measures

Performance measures for voice include delay distributions and percentage packet loss and their effect on perception, as well as the number of users that can be handled for each scheme. The effectiveness of each scheme in exploiting speech activity detection to save on channel utilization during pauses is of particular importance, since the packet versus circuit comparison hinges largely on this issue. For data traffic, delay distributions and buffer storage requirements as functions of voice traffic load and strategy for servicing the mixed traffic are of particular interest. In all cases, network throughput rates, channel utilization characteristics, and switch processing requirements will be critical performance measures.

4. Rate-Adaptive Communications Techniques

a. Introduction and Background

Conditions in an operating communications network are in a continual state of change. Traffic levels fluctuate as new users enter and leave the system and as the capacity requirements of individual users vary with time. Link capacities may vary due to jamming or fading. Portions of the network may become temporarily unavailable due to equipment failures. An essential aspect of network operation is the ability to maintain the best possible service in the face of changing conditions. Data networks such as the ARPANET include routing and flow-control algorithms designed to cope with changing conditions. In an integrated voice/data network, an important additional opportunity for effective adaptation is presented because voice can be communicated at a variety of rates with different degrees of fidelity. Assuming the availability of a sufficiently flexible speech coder, voice bit rates could be adjusted up or down to match the capacity available in the network at a given time.

One approach to voice rate control is to assign to each user entering the system a rate based on network loading at the time of dial-up and on the priority of the user. Ultimately, a user would be blocked if the network were unable to support him at his lowest bit rate. A second approach would provide more rapid response to changes in network status by allowing terminals to change rates during conversation based on control signals from the network. A third approach permits essentially instantaneous reaction to network overloads by allowing some of the speech bits in transit to be discarded by network nodes along the communication path.

This third approach requires an "embedded coding" vocoder of the type originally proposed at Naval Research Laboratory. In this technique, the voice signal is encoded into packets of different priorities. The lower priority packets may be discarded at any time without affecting the intelligibility of the speech although there will be some degradation in speech quality. In any vocoder frame, the speech synthesizer will use all the packets that arrive, and fill in for all those which are missing. This results in varying quality levels as the priority level of the arriving packets changes. To insure that intermediate nodes are not overloaded with packets that are later discarded, the embedded coding technique is combined with an end-to-end flow-control algorithm where the receiver notifies the transmitter of the current received rate, and the transmitter sends only those packets likely to be received at the destination.

During FY 78, an embedded coding vocoder based on a synthesis of channel vocoding and sub-band coding methods was developed at Lincoln Laboratory. This vocoder is capable of operating at rates from around 2 to 20 kbps with quality and robustness increasing with rate, and without perceptible transients due to rate switching. In addition, a study via simulation of the effects of end-to-end flow control in a multi-hop network with embedded coding of voice has been carried out. Control mechanisms which allow rapid adaptation and stable network behavior have been identified.

b. Experiment Objectives

The objective of experiments in the rate-adaptive communications area is to develop and evaluate techniques for maintaining network performance and stability in the face of changing conditions. Schemes which provide graceful degradation of the quality of service as network loading approaches saturation are to be developed and tested. There will be a focus of attention on voice communication because of the variable-rate nature of the voice source. However, data flow control techniques will also be tested. Of particular concern is whether the ARPANET class of flow-control techniques, which have been tested in a data-only environment, will perform effectively in an integrated voice/data network.

c. Experimental Approach and Requirements

Traffic emulation will be a primary tool in the testing of rate-adaptive techniques, since variable-rate voice coders are still in the developmental stages. Voice traffic characteristic of a large number of embedded coding vocoders will be emulated. For experiments on the satellite subnet, the priority-oriented mechanisms built into PODA will be used as a means for discarding lower priority packets as loading increases. End-to-end rate control mechanisms will be emulated and tested. To evaluate the effects on speech quality, one or two speech terminals capable of real embedded coding operation will be implemented in flexible high-speed processors. The speech from these terminals will be combined with emulated traffic, which will include other voice traffic as well as data traffic. The effect on voice performance of a

sudden large demand for data capacity (e.g., a high-priority file transfer) will be evaluated. The performance of the highly dynamic embedded coding approach will be compared with simpler approaches such as rate control at dial-up.

d. Performance Measures

The primary performance measure for voice is the speech rate and associated quality achieved for the community of users as a function of network conditions. Stability is also an important performance measure. When many users try to respond to network conditions by raising and lowering their rates, network instabilities can result if proper controls are not applied. The proper trade-off between adaptation rate and stability must be made. For data traffic, the usual measures of delay, buffering requirements, and reliability are applicable. The network-wide implications of rate-adaptation must be examined. Questions include an evaluation of the overall benefits to voice and data users as measured against the increased cost and complexity of the application of rate-adaptive techniques.

5. Routing

a. Introduction and Background

Previously, routing algorithms have been developed either primarily for voice traffic (the telephone network) or primarily for data traffic (the ARPANET). In design of routing algorithms, satellite channels have generally been utilized as a cable in the sky rather than as a demand-assigned medium for flexible broadcast connectivity. Future integrated networks will include large volumes of both voice and data traffic, with flexible satellite connectivity in addition to terrestrial links. Routing algorithms should be matched to this environment, and the wideband-network test bed provides an excellent facility for the development and testing of such algorithms.

b. Experiment Objectives

A major goal of the wideband project is the investigation of the interaction between the satellite and terrestrial components of a combined network. Of particular interest in this context are potential routing algorithms which involve choices between the utilization of satellite and terrestrial links. A mechanism must be provided for weighing the ability of the satellite to offer a path of fewer hops to a desired destination, against the disadvantage of the satellite round-trip delay. In addition, the routing algorithm must regularly operate in the context of a network with variable-capacity links, since the DAMA capability implies that capacity assignments on the satellite links will be adapted to changing traffic requirements. The design of a routing algorithm in a terrestrial/satellite network should also take into account the advantages afforded by the broadcast satellite medium for transmission of multi-addressed or conferenced data streams.

The interaction between voice traffic and data traffic will also be of key concern, since the different delay and throughput requirements of these two types of traffic may call for different routing techniques. Voice traffic might be routed by a preselected fixed path or virtual circuit which has been determined to have sufficient capacity to accommodate the required bit rate, while interactive data might be routed on an independent packet-by-packet basis. The effect of data queue sizes in the network on the choice of voice virtual circuit routes through the network must also be considered.

c. Experimental Approach and Requirements

A prerequisite for routing experiments is a study to identify potential routing algorithms for the combined satellite/terrestrial network. Such a study should consider in detail the issues described above and identify important routing experiments. The applicability of existing routing algorithms to the wideband integrated environment should be considered. The design of the PSAT/ILRAN gateway should reflect the requirements for routing experiments.

d. Performance Measures

Routing performance measures include delay, throughput, cost, and reliability. Of particular interest will be the ability to choose between satellite and terrestrial paths in such a way that performance is optimized.

6. Conferencing

a. Introduction and Background

The capability for conferencing is an important requirement for military communication systems. In addition, voice conferencing provides an excellent vehicle for exercising a network with real traffic and for direct demonstration and observation of performance. Voice conferencing experiments and demonstrations have been an important facet of the DARPA Packet Speech Program and of the APSE Program. It is expected that conferencing experiments will be of similar importance in the wideband experiment. The ability to handle greater numbers of voice users and to simultaneously communicate varying levels of real and simulated data traffic will add richness to conferencing experiments in the wideband network.

b. Experiment Objectives

Early conferencing experiments should focus on the transfer of APSE and ARPANET conferencing protocols to the wideband environment. These protocols generally implement control-signal switched conferencing, where the floor is switched from participant to participant by a chairman computer program which acts on the basis of participant request signals (push buttons, touch tones, etc.). Of particular concern in these control-switched conference experiments will be the interaction between the conference controller program and PSAT IMP and gateway software. It is expected that conferencing on the satellite net will employ the stream mechanism of PODA, and experiments determining the capability and flexibility of this mechanism for conferencing purposes should be carried out as early as possible.

Another conferencing strategy to be tested is the voice-controlled technique where a participant's speech activity is monitored and he expresses his desire to talk simply by beginning to talk. Conferencing algorithms proposed for the World Wide Military Command and Control System (WWMCCS) utilize voice-activated switching in conjunction with a distributed control algorithm. Collisions on the satellite channel and momentary speech loss result when two or more users enter talkspurt within the same satellite round-trip time. These collisions result also in temporary loss of crypto sync on the channel. Programs to simulate these effects have been developed in a local test bed at Lincoln Laboratory. Experiments using this test bed are proceeding, and future evaluations of these algorithms will be carried out in an operational WWMCCS test bed. The experimental wideband network will allow the study and evaluation of these and similar conferencing techniques in an environment which is more realistic than a laboratory test facility and more flexible than an operational conferencing system.

Consideration and simulation appropriate to the effects of security requirements will be an important aspect of these experiments. Conferencing will be carried out in the satellite subsystem. A later objective is the development, demonstration, and evaluation of an internettted conferencing capability for large numbers of conferees. This effort will serve as a follow-on to the current development of an internettted voice conferencing experimental facility involving the Atlantic Satellite Network and the ARPANET. The limited bandwidth in both these nets and the large delays in the ARPANET restrict the scope of internettted conferencing experiments in these nets. The wideband network will provide a more substantial test bed for internettted conference experiments. The location of conference control and the interaction between the broadcast satellite network and the terrestrial net are subjects of particular concern.

Future military requirements may call for conferencing capabilities which include graphical as well as voice communication among participants. The wideband network will serve as an excellent facility for the development and demonstration of such a capability. Issues to be considered include the specification of required capabilities for voice/graphics terminals, the development of protocols for conference management, and the integration of voice and graphics data in the terminal access area.

c. Experimental Approach and Requirements

Conferencing represents a user application on the network. Therefore, conferencing experiments must be preceded by the establishment and verification of network operating capability. Internettted conferencing must be preceded by the development of gateways. For conferencing on the wideband satellite network, the availability of the stream PODA capability is of particular importance. Conferencing protocols developed for APSE presuppose a stream capability for maximum efficiency. Other obvious requirements for conferencing are voice terminals with suitable signaling mechanisms (buttons, touch tones, etc.) and network access capabilities. Finally, conferencing protocols and software must be designed and implemented.

The experimental approach will be to rely on transfer of existing (e.g., APSE) conferencing protocols to the wideband environment for initial conferencing experiments and demonstration of capabilities. Planning for more advanced conferencing experiments such as voice/graphics conferencing should be a parallel effort.

d. Performance Measures

The wideband network test bed is not viewed as an appropriate facility for formal evaluation of conferencing strategies via human factors experiments. Instead, the test bed will be used for feasibility demonstrations of different conferencing systems strategies. Evaluation of the success of particular strategies will be subjective but informal. Intelligibility of the speech, recognizability of talkers, and ability to gain the floor when required are important performance measures. Also of primary importance are ease of conference initiation and reliability, including the ability to maintain the conference even when some of the equipment (e.g., one of the PSAT IMPs) temporarily becomes disabled or loses communication.

7. Internetting

a. Introduction and Background

Internetting refers to communications between networks or subnets which operate with different communications protocols. Such communication is carried out with the aid of special

processors called gateways, which reside at the internet boundaries and perform the functions necessary to transmit data from one network to another in usable form. Internetting is a subject of continuing research, as exemplified by the ongoing activities of the DARPA Internetting Program. In the packet-switching area, protocols for packet internetwork communication have been developed and implemented for experimental use in the ARPANET, Atlantic Satellite Net, and Packet Radio Net. The capability for intercommunication between different networks, whether circuit or packet switched, is a continuing requirement for operational military and civilian communications.

In the early phases of the wideband program, a limited form of internetting will serve as a convenient means for transferring central, monitoring, and measurements information between ARPANET hosts and PSAT IMP fake hosts. Later, experiments directed at internetting problems will become a focus of activity as the satellite and terrestrial networks become fully operational and development of gateways is initiated. Important problem areas to be addressed in these internetting experiments include circuit/packet interoperability, routing in combined satellite/terrestrial networks, and conferencing.

b. Experiment Objectives

The development and testing of a circuit-to-packet gateway for voice, including packetization/depacketization at the gateway, is an important experiment objective. Important issues to be investigated include requirements for gateway processing power, the effect of lost packets, and implications for privacy/security. The circuit/packet interoperability problem will arise specifically in the wideband test bed in experiments where voice is circuit switched in the ILRANs and must traverse the packet-switched satellite net. The gateway development necessary to accommodate such communication will serve as a model for the later establishment of internet connections between the experimental network and operational circuit-switched voice networks such as AUTOVON and AUTOSEVOCOM II.

c. Experimental Approach and Requirements

Initial internet experiments will utilize an ARPANET/PSAT gateway, assuming that a decision is made to implement such a gateway. Hosts on the ARPANET will be used as a source for real data traffic (file transfers, etc.) to be transmitted across the wideband satellite net. The ARPANET speech capability can be adapted for use in internetted speech experiments.

An essential requirement for wideband internetting experiments is careful design of a wideband gateway between the satellite and terrestrial network. The gateway must be flexible enough and have sufficient processing power to accommodate varied experiments with voice packetization/depaketization, satellite/terrestrial routing algorithms, and internet conferencing strategies. Study efforts directed at planning such experiments should begin as early as possible so that the required gateway functions can be specified.

d. Performance Measures

The effectiveness of the circuit/packet voice gateway will be determined on the basis of the ability to maintain voice communication with minimal added delay in passing through the gateway. This task as well as other internetting functions should be accomplished with as little cost as possible in terms of gateway processing power, memory requirements, and extra overhead added to the transmitted bit stream.

REFERENCES

1. P. T. Brady, "A Technique for Investigating On-Off Patterns in Speech," Bell Syst. Tech. J. 44, 1-22 (1965).
2. P. T. Brady, "Effects of Transmission Delay on Conversational Behavior on Echo-Free Telephone Circuits," Bell Syst. Tech. J. 55, 115-134 (1971).
3. K. Bullington and J. M. Fraser, "Engineering Aspects of TASI," Bell Syst. Tech. J. 38, 353-364 (1959).
4. R. M. Metcalfe and D. R. Boggs, "ETHERNET: Distributed Packet Switching for Local Computer Networks," Comm. ACM 19, 395-404 (1976).
5. S. Seneff, "Computer Simulation Model for a Digital Communications Network Utilizing an Embedded Speech Encoding Technique," Technical Note 1978-33, Lincoln Laboratory, M.I.T. (20 October 1978).

APPENDIX

ADAPTIVE TRAFFIC ESTIMATION ALGORITHMS

I. INTRODUCTION

A number of possible algorithms were tried for estimating the number of terminals with packets ready to send. These algorithms were tried either in the ETHERNET simulation program or in a simpler program which simulated only the contention features. The algorithm which gave the best performance and which was used in the simulation results presented is described in detail. Some observations regarding the problems with other algorithms are included.

In order to maximize the probability of some terminal successfully starting a packet transmission in a time slot, each terminal with a packet ready to transmit should try to access the channel with a probability $p = 1/q$ where q is the number of terminals with packets ready to transmit. However, since control is distributed, no terminal really knows how many other terminals are trying to gain access to the channel. Each terminal must estimate the total number by observing the activity seen at its node on the cable. Each terminal may, of course, arrive at a different estimate, but the hope is that they will be close enough to the correct value that performance will not be seriously degraded.

II. ALGORITHM DESCRIPTION

The most successful algorithm to date works in the following way:

- (a) Each terminal starts with an initial estimate $\hat{q} = 1$. Hence, initially it will transmit after the first idle slot when it has a packet ready.
- (b) If a terminal tries to transmit and collides with another user on the cable, it increments its estimate of the traffic by one. In the next slot it will transmit with a new probability $1/\hat{q}$.

If the terminal elects not to transmit, based on the probabilistic rule, and the channel is idle, then $\hat{q}$ is decremented by one. If the terminal transmits the packet successfully, the estimate is unchanged.
- (c) During the interpacket interval, the terminal ignores the activity on the cable. When the next packet is generated, the terminal starts with the estimate it used for the successful transmission of the previous packet.

III. ANALYSIS

For the first cut at the analysis, assume that all the terminals maintain exactly the same estimate of the traffic activity. In this case, we can easily obtain expressions for the probabilities of different events. If there are q terminals with packets ready and each terminal estimates $\hat{q}$ terminals and transmits with probability $P = 1/\hat{q}$, then the probability of the channel being idle during a slot is

$$P_r(I) = \left(1 - \frac{1}{\hat{q}}\right)^q .$$

The probability of a packet getting through from some terminal is

$$P_r(P) = \frac{q}{\hat{q}} \left(1 - \frac{1}{\hat{q}}\right)^{q-1},$$

and the probability of a contention is

$$P_r(C) = 1 - P_r(I) - P_r(P).$$

The probability that a particular terminal tries to transmit and encounters a contention is

$$P_r(T \cap C) = \frac{1}{\hat{q}} \left[1 - \left(1 - \frac{1}{\hat{q}}\right)^{q-1}\right].$$

These curves are shown in Fig. A-1 for the case $q = 5$.




Fig. A-1. Probability of various outcomes as a function of estimated traffic $q = 5$

The relative character of the curves is about the same for other values of q .

It can be seen that the optimum value of $\hat{q}$ is at $\hat{q} = q$ as was derived earlier. If the estimate of activity is high, it is more likely that the channel will be idle. If it is low, then the probability of contention increases.

The algorithm described above will operate at the point at which the increments to the estimate balance the decrements. This occurs when

$$P_r(T \cap C) = P_r(I).$$

This operating point is marked as A in Fig. A-1 and corresponds to an estimate of $\hat{q} = 3.2$ which is somewhat lower than the true value of $q = 5$. This results in a reduction in the probability of successful transmission from about 0.41 to 0.38, increasing the waiting time by about 8 percent over the ideal value. This value of degradation is tolerable. However, the loss becomes greater as q increases, becoming a factor of 2 when $q = 30$.

IV. OBSERVATIONS

It can be seen that the probability of successful transmission falls off very rapidly for estimates below the true value but much more gradually for estimates above the true value. This suggests using an algorithm with an operating point on the high side.

One algorithm of this type works in the same way as the previous one except that the estimate is additionally incremented when a terminal holding a packet observes the channel to be busy during a slot. This can occur either because the terminal transmits and collides with another terminal (as in the original strategy) or because it observes another packet or collision on the channel. This should produce an operating point where the probability that the channel is busy $P_r(B) = 1 - P_r(I)$ equals $P_r(I)$. This point is shown as B in Fig.A-1.

In the simulations, this point exhibited a peculiar type of instability. The activity estimates of the different terminals diverged. One or sometimes a few terminals would keep increasing their estimates to very high values thus causing them to attempt transmission very infrequently. This caused their waiting times to become large. The other terminals compensated by reducing their estimates thus balancing the busy and the idle slots.

The dynamics of this instability are not well understood. The same effect should be possible for the original algorithm but has not been observed. This suggests a hypothesis that the point A is in some sense stable while the point B is not.

Another possible algorithm would operate at the point where

$$P_r(C) = P_r(I) \quad .$$

The terminal would increment its estimate when it observes a collision on the channel whether or not it was a participant in the collision. This would result in operation at point C in Fig.A-1. This point is very close to the correct value; however, this scheme makes some added demands on the terminal hardware. It requires that the terminal be able to distinguish a packet transmitted by another terminal from a collision involving two other terminals. The hardware implications of this have not been examined. Also, the scheme has not yet been simulated so that it is not known whether it is stable.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM												
1. REPORT NUMBER 18 ESD-TR-78-366	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER 9												
4. TITLE (and Subtitle) 6 Information Processing Techniques Program * Volume II, Wideband Integrated Voice/Data Technology.	5. TYPE OF REPORT & PERIOD COVERED Semiannual Technical Summary 1 April - 30 September 1978 <i>rept.</i>													
7. AUTHOR(s) 10 Bernard Gold	8. CONTRACT OR GRANT NUMBER(s) 15 F19628-78-C-0002													
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 16 ARPA Order-2929 Program Element No. 62708E Project No. 8T10													
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209	12. REPORT DATE 11 30 September 1978													
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB Bedford, MA 01731 10 44p.	13. NUMBER OF PAGES 46													
15. SECURITY CLASS. (of this report) Unclassified														
15a. DECLASSIFICATION DOWNGRADING SCHEDULE														
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.														
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)														
18. SUPPLEMENTARY NOTES Supplement to ESD-TR-78-289 (Vol. I) and ESD-TR-78-396 (Vol. III) <i>A066247</i> <i>R NH</i>														
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) <table border="0"> <tr> <td>information processing techniques</td> <td>packet speech</td> <td>wideband satellite network</td> </tr> <tr> <td>wideband integrated voice/data</td> <td>network speech terminal</td> <td>speech concentration</td> </tr> <tr> <td>technology</td> <td>C² links</td> <td>adaptive networking</td> </tr> <tr> <td>digital voice communications</td> <td>time-varying communications</td> <td></td> </tr> </table>			information processing techniques	packet speech	wideband satellite network	wideband integrated voice/data	network speech terminal	speech concentration	technology	C ² links	adaptive networking	digital voice communications	time-varying communications	
information processing techniques	packet speech	wideband satellite network												
wideband integrated voice/data	network speech terminal	speech concentration												
technology	C ² links	adaptive networking												
digital voice communications	time-varying communications													
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes work performed on the Wideband Integrated Voice/Data Technology program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 April through 30 September 1978.														

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

207650

JTB