

Microsurveillance of the Urban Battlefield

MITRE

MITRE QUALITY MANUFACTURED

Microsurveillance of the Urban Battlefield

C. Callan (Chair)
J. Cornwall
W. Dally
A. Despain
D. Eardley
R. Garwin

P. Horowitz
W. Press
J. Sullivan
J. Vesecky
E. Williams
P. Weinberger

19960322 061

February 1996

JSR-95-125

Approved for public release; distribution unlimited

JASON
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, Virginia 22102-3481
(703) 883-6997

DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information estimated to average 1 hour per response, including the time for review instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE February 1995		3. REPORT TYPE AND DATES COVERED
4. TITLE AND SUBTITLE MICROSURVEILLANCE OF THE URBAN BATTLEFIELD			5. FUNDING NUMBERS 07-95-8534-04	
6. AUTHOR(S) C. Callan, J. Cornwall, W. Dally, A. Despain, D. Eardley, R. Garwin, P. Horowitz, W. Press, J. Sullivan, J. Vesecky, E. Williams, P. Weinberger				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The MITRE Corporation JASON Program Office, Z561 1820 Dolley Madison Blvd McLean, Virginia 22102			8. PERFORMING ORGANIZATION REPORT NUMBER JSR-95-125	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) ARPA/TIO 3701 North Fairfax Drive, Arlington, Va 22030-1714			10. SPONSORING/MONITORING AGENCY REPORT NUMBER JSR-95-125	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release -Distribution Unlimited			12b. DISTRIBUTION CODE Limiter Statement A	
13. ABSTRACT (Maximum 200 words) It is widely agreed that urban military operations demand greater "situational awareness" than now exists. Soldiers need mapping tools to tell them where they are, real time information on what's around the corner and behind walls as well as reliable data links to receive and send orders and intelligence. At the same time, commanders need accurate knowledge of "what's happening" in the city as a whole. Progress in micro-electronics and computing is so rapid that we can begin to think about a system which blankets the urban battlefield with cheap, small imaging sensors, collects all the data via a high-bandwidth communication network, converts the data flood to useful intelligence and displays information in a useful way to a wide range of users. In this study, we examine whether this vision of an urban battlefield situational awareness system survives confrontation with technological and physical reality. We conclude, with some reservations, that it does and propose some near-term demonstrations that should help in giving the vision more precise definition.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified
				20. LIMITATION OF ABSTRACT SAR

Abstract

It is widely agreed that urban military operations demand greater "situational awareness" than now exists. Soldiers need mapping tools to tell them where they are, real time information on what's around the corner and behind walls as well as reliable data links to receive and send orders and intelligence. At the same time, commanders need accurate knowledge of "what's happening" in the city as a whole. Progress in micro-electronics and computing is so rapid that we can begin to think about a system which blankets the urban battlefield with cheap, small imaging sensors, collects all the data via a high-bandwidth communication network, converts the data flood to useful intelligence and displays information in a useful way to a wide range of users. In this study, we examine whether this vision of an urban battlefield situational awareness system survives confrontation with technological and physical reality. We conclude, with some reservations, that it does and propose some near-term demonstrations that should help in giving the vision more precise definition.

1 INTRODUCTION TO THE PROBLEM

It is a commonplace that future US military operations will require greater situational awareness than is now available. Since we will have fewer warfighters and resources and will be increasingly unwilling to accept casualties and errors due to the fog of war, our forces will rely on having better and more timely information than their opponents. There is widespread agreement that ongoing advances in technology can, and should, be exploited to achieve such information dominance. The framework within which it is imagined that this can be done is the creation of a "digital battlefield" in which intelligence of all kinds is automatically collected by sensors and then processed and transported in real time to users by a battlefield version of the much-touted National Information Infrastructure. Soldiers would be given mapping tools to tell them where they are, real-time information on the location of enemy units as well as reliable data links to receive and send orders and intelligence. At the same time, commanders would be given real-time and accurate knowledge of "what's happening" on the battlefield as a whole. One of the great attractions of this concept is that it relies on an ensemble of technologies in which the US is expected to retain a sizable lead over the rest of the first world, not to speak of the second and third worlds, for many years. If the digital battlefield lives up to its military promise, its development would confer a substantial long-term military advantage on US forces over all likely adversaries. The US defense community has been thinking along these lines for some years now and the Army is beginning to construct its doctrine for twenty-first century warfighting around the existence of digital battlefield systems (FORCE 21).

Indeed, it has become a consensus view that automation and digitization of the collection and distribution of battlefield intelligence will be a transforming military technology, comparable in importance to stealth and smart weapons. In the light of this, it is surprising how little quantitative analysis exists concerning what is technologically possible, how the different types of technology might work together and what precise military benefits might accrue from their use. In the absence of such studies, it is hard to see how informed decisions can be taken concerning hardware development and the integration of hardware into systems. An important point is that, to be useful, such studies must be done in the context of fairly specific military scenarios. The quantitative implications of "information dominance" vary depending on whether the context is full-scale war against a large modern army, a relief mission in a small country in the throes of civil war, or any of the many possible cases in between. What you need to know, and at what resolution and over what spatial and temporal scales you need to know it, is very scenario-dependent and the associated technological challenges are of widely varying difficulty.

The goal of this study, then, is to provide some first-order quantitative and technical analysis of the information dominance problem in the context of Military Operations in Built-Up Areas (MOBA) in Operations-Other-Than-War (OOTW) situations. We have three reasons for choosing this context: it is operationally important (as witness recent actions in Somalia and Haiti); our forces are not ideally equipped for this mission and urgently need whatever help technology can provide (as explained in the 1994 MOBA report of the Defense Science Board); and, by virtue of the limited geographical area and small number of combatants involved, it is technologically the least demanding of the scenarios we can imagine. As we will argue in the rest of

this report, there is a good chance that useful systems can be developed for this application in the near term.

We have called this study "Micro-Surveillance of the Urban Battlefield" to reflect our notion that, in the geographically restricted case of urban operations, one can realistically consider a system which blankets the battlefield with cheap, small imaging sensors, transports the sensor output over a high-bandwidth communications network, converts the data flood to intelligence and displays useful information to a wide range of users. Our primary purpose was to examine whether this "all-seeing, all-knowing" vision of an urban battlefied situational awareness system makes physical and technological sense. We conclude, with important reservations, that it does. The secondary purpose of the study was to make suggestions about technology development. We feel very strongly that the technology alone does not define the system and that experimentation will be needed in order to decide precisely what the system should *do*. To that end, we propose a near-term demonstration that should help in giving the whole concept more precise definition.

We have organized our presentation as follows: In Section 2 we define a scenario and the mode of operation of what we feel would be a useful system. In Section 3 we proceed to a more precise definition of the technical elements of the system and the demands they must meet. In Section 4 we present detailed discussions of the individual technologies that must play together to make a working system. This section should probably be browsed on a first reading: not every topic is given a full treatment and the discussions of individual technologies are at a widely varying level of detail. Our overall point is that, while there are no obvious "show-stoppers", the needed technologies range from off-the-shelf to challenging, but doable. In Section 5 we offer some

proposals about how to organize the technology development process and in Section 6 we present conclusions and recommendations.

2 URBAN SITUATIONAL AWARENESS SYSTEMS

2.1 Mogadishu: A "Typical" MOBA Scenario?

Our first job is to define the scenario of interest. We choose to look at urban military operations in a low-intensity conflict, or peacekeeping, situation. The recent operations in Somalia and, more specifically, in its capital city, Mogadishu, are representative of what we have in mind. The setting indicated in Figure 1 is a third-world city of dimensions 10 km by 10 km and the military problem includes controlling a hinterland of dimensions 100 km by 100 km. In the city itself, and perhaps moving back and forth from the hinterland, is a small force of some few thousands of lightly-armed hostiles. A major problem is that these potential opponents move about within a civilian population of order 100,000.

There are some 10,000 US troops with perhaps hundreds of them on active patrol at any given time. There will be hundreds of US vehicles on the road with some tens of armored vehicles accompanying the active patrols. Again, these US vehicles will be moving among many thousands of civilian vehicles going about their routine business (except for the few carrying loads of ANFO!). There will be a small number of helicopters, fixed-wing aircraft and reconnaissance drones in the air and it can probably be assumed that they all belong to US forces.

In this scenario, many US troops are engaged in humanitarian, rather




Figure 1. Notional map of a Mogadishu-like scenario.

than strictly military, activity. Soldiers are often sniped at, armor-piercing rounds are sometimes fired and anti-aircraft fire is sporadic. Soldiers are more often engaged in crowd control than full-fledged firefighting.

2.2 A Notional “Mogadishu” System

What does situational awareness mean in this context? Individual soldiers and patrols want to be able to “see around corners” and “see through walls” in order to be aware of anything threatening in their immediate neighborhood. They want to know where they are, not just in the GPS sense of absolute location, but in the sense of knowing what street they are on, whether this or that building is “interesting”, and whether other friendly units are in the vicinity and so on. They also want a continuous data connection with their command center and other units, presumably by radio. This is a non-trivial problem, even for voice, because line-of-sight communication is obstructed by buildings. The command center wants to know the location of all its soldiers and vehicles, probably to sufficient accuracy to eliminate friendly fire errors. The commanders want to have a global real-time view of “what’s happening” in the city as a whole. They also want the ability to identify anomalous or threatening patterns of activity in order to be able to nip riots, for instance, in the bud.

A cartoon version of one possible realization of such a system is given in Figure 2. Video imagery is collected by a large number of autonomous fixed sensors emplaced on buildings or other vantage points. A communications network consisting of relay transmitters carried by a number of UAVs ships the imagery to a central processing location. The same links carry data back




Figure 2. Notional Mogadishu surveillance system.

out to soldiers and vehicles in the field and allow communication between units. In the situation depicted in the cartoon, soldiers on one street are using such a link to look at a crowd hidden by a building in order to make their own judgement about the presence of a threat. Depending on one's level of ambition and optimism about ATR and related subjects, the processing center can provide some level of automation in interpreting the images coming from the sensors.

All of this implies the ability to collect, transport and interpret large quantities of data. In the rest of the report, we will quantify the problem and analyse the extent to which technology can provide the means to realize the above list of functions. It is much less easy to judge the net military utility of such a system or to say what is the best way to configure it. These are important questions which can't be answered *a priori*: there is no substitute for experimentation and real-world experience. In the concluding sections of the report we will present some ideas about efficient ways of gaining such experience.

2.3 An Existence Proof: LA Freeway Status on the Web

Interestingly enough, it seems that civilian systems having many of the features that we are looking for have recently sprung into existence, courtesy of the Internet and the World Wide Web. We will digress to examine one of them, the Los Angeles Freeway Status Map, because we think it holds useful lessons for the military situational awareness problem.

If one navigates to the Web site <http://www.scubed.com/caltrans>, one is invited to click on various links to information on transportation systems in southern California. Clicking on the link to the Los Angeles and Orange County freeway map, one gets a color map display like the one shown in Figure 3. The red, yellow and green dots denote the state of traffic flow (ranges of mean traffic speed averaged over a few minutes) on a particular freeway at a particular location. The information comes from vehicle sensors embedded in the pavement at the various on- and off-ramps and is communicated (over phone lines?) to a central computer database. The file that defines the Web page is automatically updated every five minutes (the colors of some of the indicator lights change) in order to give a real-time picture of the state of traffic flow. A driver with a wireless-modem-equipped laptop and Web browser software can (at some hazard to neighboring vehicles) “see” traffic jams forming ahead of him and reroute his journey accordingly.

Although there are many hundreds of sensors reporting, the global data rate is trivial. Nonetheless, when this data trickle is displayed the right way, it can have enormous value to the informed human observer. Even staying within the Web browser conceptual framework, many useful enhancements of this facility can be imagined. Since transmission of images is a standard Web facility, one could, for instance, install video cameras and include links in the Web document to pictures of the traffic at selected locations. With some imagination, one could think of other sensors and other ways of processing or presenting the data that would have value to a variety of users (drivers, police, tow truck operators, etc.).

We think that there are some lessons for the military application to be drawn from this example. The concept of marrying sensor data to a




Figure 3. The L. A. freeway status map display on the Web.

map display is very powerful (though hardly new!). No further processing or automatic pattern recognition is needed, in this case, to generate a very powerful tool for visualizing the overall state of a complicated system. The Web browser concept of responding to user requests is likewise a powerful notion and, more importantly, a very flexible one: a simple change in an HTML file stored at the server can completely change the “look and feel” of the data display. The way the system functions can be made to evolve rapidly and cheaply in response to discoveries about what features the users find most useful. Also, there is nothing hard-wired about the data transport: anything that hooks into the TCP/IP network can see the display and the designer thinking about what to display doesn’t have to think about how his data is going to get to the user. Most of these ideas are platitudes by now, but it is instructive to see, concretely, what can be accomplished when the platitudes are incorporated in real systems!

3 PRELIMINARY QUANTITATIVE ANALYSIS

3.1 System Elements and Technology Issues

We now turn to a quantitative discussion of the components and modes of operation of such a situational awareness system. We have to deal with problems in physics (sensors, displays and platforms), electrical engineering (communications networks) and computer science (databases, processing and display). A listing of the problems that must be dealt with in each area will give a useful impression of the order of magnitude of the overall technology problem. The list is not exhaustive!

In the physics category, the primary requirement is for miniature, cheap autonomous imaging sensors. Sensors of other kinds will be wanted as well, but imaging sensors have the highest data rate and will pose the hardest problem. Since we want to blanket the battlefield with them, they have to be very cheap and very easy to emplace. Obviously, the current generation of commercial video sensors is not quite there, but it is not many generations away from the goal. We may need UAVs to provide communication links between the sensors and the users of data. The technology of UAVs *per se* is maturing rapidly, but, in this application, we need to manage, not just one flyer, but enough to support a comms network. This is a control problem that has not arisen before. Finally, there is the whole issue of getting information to the individual soldier. His equipment has to be very light and very low-

power (and preferably an integral and unobtrusive part of equipment that he *wants* to carry anyway, like his helmet). We haven't gone into quantitative detail on this problem, but it seems to be of about the same order of difficulty as the sensor problem.

In the electrical engineering category, the primary problem is that of creating an adequate data communications network. We propose to have a large number of imaging sensors, generating a large data flow, yet drawing their operating power from small batteries. This is the classic cellular telephony problem and the only solution (there's a theorem here) is to relay the data over short link distances. A further complication is that, in urban terrain, buildings can block convenient link lines-of-sight and cause multipath interference. A quantitative analysis will show that the network problem, at least for this relatively modest version of the digital battlefield, is challenging, but by no means impossible. Finally, there is the ever-present need to live by the discipline of scalable, plug-and-play architectures: in order to make these systems useful, they must be as flexible as possible with respect to the hardware that realizes the network as well as the size of the network. Ideally, when a new sensor or link node is turned on, it should integrate automatically and smoothly with the pre-existing system. This is a non-trivial but, in our opinion, manageable technology problem.

Finally, we come to the computer science issues. The dominant problem here is the sheer volume of data generated by a large number of imaging sensors. It simply will not be possible to manually convert this data flow to intelligence: some level of image/data processing will be essential. There is a wide range of information-extraction tasks that one would like to automate (pick out a particular *thing*, say Aidid's favorite Mercedes-Benz, or identify

anomalous patterns in the movement of large numbers of people or vehicles). It is probably fair to say that our current ability to automate these tasks is quite primitive, but it is probably reasonable to expect that, in ten or twenty years time, we will be able to do some amazing things. The trick will be to learn how to do something useful *now* with minimally sophisticated processing and to design a system that can easily incorporate improvements as they develop. There is a heterogeneous database problem to face as well. In the spirit of the LA Freeway Status Map, we must maintain a datafile that incorporates the current state of knowledge about the battlefield. How do we decide what part of the flood of data from our sensors (and other sources) to keep? How do we provide the different cuts through the data to meet the needs of different users? These are fairly standard database/data fusion questions, but it is our impression that powerful general strategies for solving them do not yet exist. Again, the challenge will be to find a strategy of minimal sophistication to get a working system up for experimentation and then to design things so that improvements can easily be incorporated. We will eventually present suggestions about how to manage this problem, without pretending to know how to solve it in detail. We nonetheless are convinced that it can be done.

3.2 First-Order System Overview

With these generalities in mind, we will now give a more concrete outline of the system. Implicit in everything we have said is the notion that there is great value in tying information to precise physical locations (as witness the LA Freeway Map). The foundation of our system, therefore, is a precise static 3D model of the city constructed by photographic or SAR

mapping carried out before the fact or in the initial stages of the action. Static "cultural" items, like street names, functions of various buildings and so on are posted to the map as they become known. Variable items, like the locations and identities of US military vehicles, are posted to this map database and updated as frequently as appropriate. If the map is accurate enough, it can be used for targeting precision weapons and reducing the likelihood of friendly fire casualties in close-in firefights. Mobile patrols can use the map to "look up" where they are to high accuracy using scene matching.

The eyes and ears of the surveillance system will be a large number of fixed sensors, emplaced by routine patrols, and reporting by the communications network. We are convinced that imaging sensors will be the backbone of the system, but important roles for audio and vibration sensors can be imagined. How many sensors do we need to constitute a useful system? As far as imaging sensors are concerned, it seems obvious that, at a minimum, you will need to know what is happening at all "major" road intersections. In the Mogadishu scenario that implies the emplacement of hundreds of sensors. This could easily expand to thousands of sensors for more comprehensive surveillance. Experimentation will be needed to determine how many sensors are really needed to make a useful system. These numbers set the bandwidth required of the data transport system and are therefore of crucial importance. Because large numbers of sensors are involved, it will be necessary to invent non-standard, probably wholesale, emplacement strategies (the technician clambering up a utility pole with a screwdriver is not the right model!).

We will have other suggestions concerning the data relay network, but one solution is to employ a modest number of orbiting UAVs to get re-

lay transceivers close enough to the data generators and users (how close depends mainly on the power and bandwidth we ascribe to the emplaced sensors, but it will be on the order of a kilometer or two). This architecture is one way to solve the urban obscured-line-of-sight communications problem and the associated vehicles provide ideal platforms for special mobile sensors, such as SARs. As we shall see, small (model airplane sized), cheap, low-altitude vehicles are the most appropriate vehicles for this application. Unfortunately, until recently, UAV technology development has concentrated on long-range high-altitude vehicles which are expensive and not really suited to this application.

Finally, there is the question of what to do with the data. In our view, it is important to start with some low-tech approach (from the point of view of computer science) that is nonetheless "effective" and work up from there. One possibility amounts to using the system as a switch to get the imagery from any chosen sensor to a given user (with a limit set by the network bandwidth on the number of users that can be so served). In this way, a patrol could look through cameras near its current location in order to "look round corners". One could also implement a "Guardian Angel" concept in which a human controller in the command center could follow the progress of a patrol and offer advice. The best way to implement these ideas and their real utility can only be determined by experimentation. There is also intelligence to be extracted from statistical/ATR analyses of the data flow into the central processing unit. This is surely a gold mine waiting to be exploited!

3.3 Sizing the Data Flow

In order to decide whether this general picture of a surveillance system will actually work, we need an estimate of the rate at which bits have to be transported. Let us first consider the rates at which data flows into the central database. The biggest component is the imagery from hundreds of sensors. Raw video rates are of order 10 MHz, but, as we will argue in more detail in the next section, compression and reasonable downsampling should get that rate down, on average, to 0.1 MHz. So, the bandwidth needed to transport the imagery data in this scenario is some tens of MHz. This is the equivalent of a few conventional TV channels. The data flow is large, but not daunting. Airborne EO/SAR surveillance at one foot resolution and on an hourly revisit schedule leads to a data rate of 10 MHz as well. All the other data sources we can think of (status reports from soldiers and vehicles in the field, non-imaging sensors) pale to insignificance in comparison.

Extracting information from the data will be a challenge as image data rates overwhelm current ATR capabilities. As indicated above, we think it possible, in the near term, to do intelligent things with minimal automation.

At the same time, bits flow out to the users in the field at certain rates. A few patrols at a time will want a view of "what's around the corner". They will want higher resolution and frame rate than average, so we will count this as a few users times 1 MHz per user. Commanders will want to consult the database for an overview of the situation and this will amount to a data flow of a few MHz. Orders and advice, verbal and otherwise, flow out to the field at a trivial data rate.

We will need a "cellular dataphone" system that supports such flows. The link data rates are large and the powers, at least those available to the sensors, are small, so cell sizes will necessarily be small. This is why a relay network is necessary. It should of course be built as an expandable, self-configuring digital data network that serves sensors, troops, weapons, etc. Furthermore, the data flow must be encrypted and there must be an appropriate multi-level security system that allows particular users to extract from the data stream only what they "need to know". We have not had time to study this issue in any detail so we have adopted the (plausible) working assumption that the overhead of encryption and security do not materially change our conclusions about size, cost and functionality of the various components of the system. We could be wrong about this.

4 TECHNOLOGY ISSUES AND OPPORTUNITIES

In this section we will offer more detailed remarks on various individual technology elements that will make up the overall system. In each case, we assess whether the demands of the urban battlefield surveillance problem can be met by current technology and, if not, whether the technology can be developed to do so. Because of limitations on time and the expertise of the authors, the treatment here is not completely balanced: some important topics (especially those in the automated data processing arena) are given fairly short shrift. Our overall impression is that, at least for the problem defined for this study, there is no insurmountable technological barrier to a useful system. Some of the subsections contain specific ideas for schemes to overcome current limitations of particular technologies.

4.1 Communications System Issues

4.1.1 Overview of the Communications Problem

Because of the line-of-sight problem in urban terrain, any network will necessarily have many relay nodes. A conceptually simple solution is to use UAVs to create a hierarchical network (many small flyers near the ground with a few larger flyers higher up for the final relay to the collection point). Shadowing by buildings reduces the effective ground footprint of any given

UAV, but, as we shall see, not catastrophically. Another possibility is to seed the area with cheap, low-power relays emplaced, like the sensors themselves, on buildings (they might in fact be co-located with the sensors). We will discuss aspects of both schemes in what follows.

Everything hinges on the power/bandwidth/range issue. For a given link power, as the bandwidth goes up, the range goes down and the number of cells needed to cover a given area goes up. The main problem is in serving the autonomous emplaced imaging sensors: they require large bandwidth links, yet must be powered by compact batteries. When we examine this issue quantitatively, we will see that what is required pushes the envelope of physical and information theoretic limits. Bandwidth reduction by video compression is therefore an absolutely essential technology and should be pushed to the limit for this application.

Another very important and difficult link is the “last mile” to the individual soldier, primarily because, like the sensors, the soldier has to rely on small batteries for his power. We think that the key to this problem is to realize that the soldier is never far from a vehicle (a few hundred meters) and probably doesn’t need more than audio bandwidth (a few tens of KHz). If we can connect the soldier to the net through his associated vehicle, we will see that the link power requirements to serve him are actually quite manageable.

At a system level, the first priority is to learn how to construct a self-configuring mobile data network. Since some nodes move around and multipath interference problems will vary with time, the most effective connectivity of the network will vary with time. Nodes will be added to the network when new sensors are emplaced and nodes will disappear when sensor batteries die and so on. We obviously want the management of such a dynamic

network to be as autonomous and robust as possible. Fortunately this is the subject of active research in several quarters. ARPA's Global Mobile program is asking many of the right questions and can probably be counted on to produce many of the answers needed for this application.

4.1.2 Power, Bandwidth and Batteries

For the data transport network, the most important issue is the tradeoff between power, bandwidth and battery lifetime. The power P_t needed to send data over an RF link is

$$P_t = (4\pi)^2 (kT_{eff})(BW)(SNR)L^2 / (\lambda^2 g_t g_r \eta) \quad (4-1)$$

where: T_{eff} is the receiver noise temperature, BW is the desired bandwidth, SNR is the desired signal-to-noise ratio, L is the link distance, λ is the RF wavelength, $g_{t,r}$ are the gains of the antennas used on the two ends of the link and η is transmitter efficiency. For the purposes of a robust design, we will make the very conservative choices $T_{eff} = 10^3$, $SNR = 10^2$, $\eta = .1$ and $g_t \sim g_r \sim 1$. The latter choice means no antenna directivity at all. It may be possible in some circumstances to attribute gain to the transmit or receive antenna with corresponding reductions in power requirements. There remains the choice of wavelength. The demands of miniaturization push toward small λ , while the demands of minimizing link power push in the opposite direction. We have chosen $\lambda = 10$ cm as a compromise. There is by now extensive experience (MIMIC) with microfabrication of RF generators at this frequency and, while a dipole antenna at 10 cm has one "large" linear dimension, its total volume and weight can be very small indeed. Our overall goal is to minimize the size and weight of a combined sensor/battery package

and it is our impression that this wavelength choice is about optimal. In the following table we display the required link powers for choices of link distances and bit rates that are relevant to our problem.

	10^2 m	10^3 m	10^4 m
10^6 Hz	2 mW	200 mW	20 W
10^5 Hz	.2 mW	20 mW	2 W
10^4 Hz	20 μ W	2 mW	.2 W

The three distances correspond to 1) soldier to vehicle link, 2) ground sensor to relay link 3) direct link to command center. The three bandwidths correspond to 1) high-definition compressed video, 2) low-definition compressed video and 3) audio or non-imaging sensor data rates. We will give arguments in the next subsection to the effect that 10^5 Hz should be adequate for routine imagery transmission, and we will use that as our design baseline requirement.

Now let's ask what volume or mass of battery is needed to support the link transmitter for a "reasonable" length of time. For systems carried by the individual soldier an autonomy period of a few hours, say 10^4 sec seems right (he can recharge when he eats). For emplaced sensors, the autonomy period has to be many days, preferably weeks, in order for the task of maintaining the sensor network not to be too onerous. What's needed here is a minimum lifetime of 10^6 , preferably 10^7 , sec. Battery energy densities range from 150 Whr/kg for the batteries sold in drugstores to three or four times that figure for less widely-used batteries such as lithium-thionyl-chloride or zinc-air. For purposes of illustration, we will take 400 Whr/kg and redo the table, listing the total battery mass needed to support the link, assuming a lifetime of 10^4 sec for the 10^2 meter link and 10^6 sec for the other two.

	10^2 m	10^3 m	10^4 m
10^6 Hz	1.5×10^{-2} g	1.5×10^2 g	15 kg
10^5 Hz	1.5×10^{-3} g	1.5×10^1 g	1.5 kg
10^4 Hz	1.5×10^{-4} g	1.5×10^0 g	.15 kg

It is clear from this that supporting the 10^2 meter soldier-vehicle link is energetically trivial. At the same time, supporting a long-distance (10 km) sensor link is not going to be possible with a tiny battery package. However, the intermediate distance link (1 km) does appear to be consistent with very small battery packages. We will shortly develop this observation into a preliminary design of a sensor/battery package.

4.1.3 Video Compression Techniques

As the previous discussion has made abundantly clear, it will not be possible to send raw video from autonomous battery-powered sensors. Substantial amounts of video compression will be needed. In this section, we will discuss what can be achieved by relatively standard techniques, sensibly applied to the needs of the problem at hand.

The video data stream from a wide-angle 512×512 pixel camera can easily be compressed to 100Kb/s using a simple video processor that can be included on the same chip as the active pixel image sensor (a promising technology for the sensor itself which will be discussed in the next subsection).

The raw data rate from a 512×512 camera with 8 bits/pixel grayscale operating at television rate (30 frames/sec) is 60Mb/s. The image can be compressed by discarding 2/3 of the frames, temporally delta coding the remaining frames in 8×8 patches, and then cosine transforming the remaining

patches and quantizing the resulting spatial frequencies.

For surveillance purposes a 10 frame per second sampling rate is adequate. The camera can be modified to operate at this rate, saving power, or the intervening frames can simply be discarded. Reducing the sampling rate drops the bit rate to 20Mb/s. The frames can be interpolated on playback to eliminate flicker.

With a wide angle camera, much of the scene will be static — sky, landscape, buildings, roads. The data rate can be reduced about 20:1 by sending only those 8×8 blocks of pixels that have changed by more than a threshold value since they were last transmitted. To reduce the potential of error accumulation the entire image can be transmitted every 10s without substantially increasing the data rate. Delta coding reduces the data rate to 1Mb/s (2K $8 \times 8 \times 8$ b blocks per second).

Frequency-domain image coding is performed in the final stage of the compression pipeline. A discrete cosine transform is performed on each 8×8 block to generate an 8×8 block of spatial frequencies. Frequencies below a threshold are clipped to zero and the remaining frequencies are quantized and run-length coded. It is expected that only about $1/5$ of the spatial frequencies will be above threshold and that the remainder can be coded into an average of 2 bits per pixel using Huffman coding. This sequence of steps is summarized in the following Figure 4.




Figure 4. Synopsis of video compression schemes.

4.1.4 Urban Multipath

The distributed nodes of the data network can communicate in a frequency band in the microwave (~ 3 GHz). Propagation paths will be available between most nearby pairs of nodes, but multipath propagation will be an important problem. Various modulation schemes might work, but we will adopt spread spectrum signals with code division multiplex (CDM) as the scheme. Emitted power may be a fraction of a watt.

The simplest and most effective measure to cope with the multipath problem is to equip each node with more than one antenna element. Two elements are likely to be enough in most situations, though more could be added for additional redundancy. An interelement separation of $\sim \lambda/2$ or ~ 5 cm is sufficient. If each node is a small box of ~ 10 cm longest dimension, two whip antennas could mount on opposite ends, to be deployed upon emplacement. For redundancy, the node might carry 4-6 whips and simply use the 2 or 3 that turn out to work best.

Now let us discuss several increasingly sophisticated methods of use for antennas. Assume that node 1 wants to listen to node 2, given that node 2 is transmitting. If one antenna element happens to lie in a multipath propagation minimum (Rayleigh fading) then another element is very likely to lie well out of the minimum, since nodes and antinodes of the incoming wave energy are $\sim \lambda/2$ apart. Hence the simplest method: Listen on each element and choose the one with greatest signal strength ("diversity reception"). However, there is a better method which is only a little more elaborate: Combine the whip outputs through an adaptive filter that maximizes combined out-

put of node 2's signal. This method actually takes advantage of multipath: The combined output is stronger than any single output, as all the multipath signals are added coherently.

The combining filter is probably best implemented in time domain due to our choice of CDM modulation. The signals are mixed down to baseband and processed by DSP techniques, using adaptive filter algorithms. This system also has considerable robustness against jamming: If a jammer appears in band, the multielement antennas can be adapted to null it out of the antenna pattern — even in the presence of multipath propagation from the jammer. (Multiple jammers, however, present a more serious problem, once the number of jammers is no less than the number of active antenna elements.)

Let us discuss the adaptive algorithm in a little more detail. In general, several arrivals of node 2's signal will appear at node 1, propagating along different path lengths. In the simplest algorithm, node 1 might lock onto first arrival. However, sometimes a later arrival would be stronger, so a more sophisticated algorithm would lock onto the strongest arrival, even if not first. The optimal algorithm, however, filters the incoming signal in time domain so as to combine coherently all arrivals. Again, this algorithm actually takes advantage of multipath to augment detected signal strength. Finally, several antenna elements are active in general, so that the combining filter has several inputs, each to be filtered in time domain.

Having tuned up to receive node 2, node 1 can now transmit to node 2 equally well. The adapted receiving filter can be conjugated to produce a near-optimal transmission filter. This filter will employ node 1's antenna elements to produce a signal at node 2's location that is near-optimal in strength,

and equalized for multipath arrivals. Node 2 is using some transmission filter to produce the signal that node 1 is listening to; if it conjugates this filter to create a reception filter, node 1's signal will be optimally selected.

There is still further optimization to be done, as long as either node is using more than one antenna element. In an outer adaptive loop, node 1 and node 2 can jointly adapt their filter weights into the several antenna elements, so as to jointly optimize their communication channel. Roughly speaking, each node could adjust its antenna pattern so as to put the most antenna gain in the direction of the strongest propagation path between them — although once again an antenna pattern that takes advantage of a coherent sum of paths will in general give optimal strength.

The various levels of optimization discussed here are of course done at an increasing sequence of data rates, as allowed by the current best channel. If either node moves, or if the multipath environment changes, the channel will have to be re-adapted.

A node on a moving vehicle is more stressing in a multipath environment. Here the multipath environment changes so rapidly that full adaptation is unlikely to be feasible. The simple and robust diversity strategy, switching among several whip antennas to find greatest instantaneous signal strength, is likely to be best.

In summary, multipath is an important problem for our network, but it is a solvable problem. In fact, the system can often take advantage of multiple propagation paths to increase signal strength, by coherent combination of paths.

4.1.5 Hierarchical UAV Communications Network

The analysis of the power/bandwidth constraints showed that the emplaced sensors, if they are to transmit at video rates and draw their power from small batteries, must transmit to nearby relays. One possible architecture for the overall data transport is *hierarchical*: each relay collects the data from a number of sensors and passes all their data up to a higher node, which in turn collects all the data from a number of relay nodes. Successive nodes in the chain are responsible for ever-increasing data bandwidth and use correspondingly greater transmitter power.

A natural way to organize a system like this is to put the various levels of relay on UAVs which orbit the city at appropriate altitudes. An added benefit is that, by putting the relay platforms in the air, one eliminates, partially if not entirely, the problem of shadowing by buildings of receiver-transmitter line of sight propagation paths (we will say a bit more about this problem in the next subsection).

In Figure 5, we show a cartoon of the kind of system we have in mind. For the Mogadishu scenario, where we need to service hundreds of sensors spread over an area roughly ten kilometers in diameter, we find that two levels of relay should suffice: We have of order ten low-altitude relay vehicles that each service some tens of sensors at most 3 km away and we have one high-altitude relay that collects all the data and transmits it onward to a processing center some 30 km away.

The power and bandwidth census for this scheme goes as follows: The individual sensor must send 10^5 Hz over a 3 km link. With no antenna gain




Figure 5. Notional hierarchical UAV relay scheme.

and all the other conservative assumptions we used before, it will have to emit .2W of RF power. This is small, but not microscopic: at 400Whr/kg of battery energy density, it takes 200g of battery (a 1.5in cube) to give the sensor a 400hr lifetime. The first relay sends about 10^6 Hz over a 10km link and needs 2W of RF power to do this if we can credit the receiver/transmitter pair with 10dB of antenna gain. The size of the vehicles in question is such that this seems reasonable. Finally, the top level vehicle has to send 10^7 Hz some 30km. This calls for 200W of RF power, again giving credit for 10dB of antenna gain. When we do a preliminary design of the UAVs for this application, we will find that these RF powers can be supplied by UAVs with adequate endurance and flight characteristics. This design is probably far from optimal but serves to show that the communications can, in principle, be provided by this type of architecture.

As an addendum to this discussion, we would like to make a few quantitative remarks about shadowing of direct lines-of-sight between an emitter on the ground and a UAV relay. The problem, of course, is that even if a particular relay is in range, the link path may be blocked by a building. This reduces the effective ground footprint of the relay by an amount which depends on the geometry of the city. It is possible to make a very rough estimate of this effect.

Suppose the UAV flies at altitude a and that the maximum slant range at which it can receive a certain type of transmitter is R_0 . But for shadowing, the UAV would have a ground footprint of area $A = \pi(R_0^2 - a^2)$. This number determines how many relay UAVs are needed to cover the whole city. Now suppose, for purposes of discussion, that the city buildings are all of height h , the streets of width w and that the transmitters of interest

are all located at ground level in the street. The relevant fact about the location of the transmitter is its displacement Δ from the center of the street (if the transmitter is directly up against a building, half the sky is blocked, much less if the transmitter is in the middle of the street). The geometrical relationships are indicated in Figure 6.

It is then a simple geometrical exercise to find out the size of the area in which the UAV at altitude a is in range of the transmitter. The result will be some fraction η of the unshadowed area $A = \pi(R_0^2 - a^2)$. The controlling parameters are the “aspect ratio” of the buildings and streets ($w/2h$) and the elevation angle ϕ of the UAV at max range ($\cot\phi = ((R_0/a)^2 - 1)^{1/2}$). If we average over the position Δ of the transmitter on the street we get a global measure of the effect of shadowing. We found that, with the fairly typical values $a = 1km$, $R_0 = 3km$ and $w/2h = 1$, the shadowing fraction $\eta = .4$. Concretely, this means that it would take 10 such relays to cover a city of $100km^2$. Our conclusion is that, for the UAV relay problem at least, shadowing is a significant, but not driving, effect.

4.1.6 Ground-Only, Self-Organizing Packet Radio Network

An alternative to a UAV-based radio relay network is to use a self-configuring network composed entirely of ground-based radios. Using cellular telephone technology a small radio can be built that will operate 1Mb/s data links over a 300m range with power margins large enough to tolerate the large path losses expected in an urban environment.

Each node of the network consists of a processor, a buffer memory,




Figure 6. The geometry of urban shadowing.

and a QPSK modulated 2 GHz transceiver. The nodes are deployed on a roughly 300m grid. They can be dropped from a vehicle at road intersections where they will have a good line-of-sight down several roads or air-dropped randomly over the city area. About 1K nodes are required to cover a 10km by 10km city.

Operating at 1Mb/s the nodes consume 100mW when transmitting. Because of the large numbers of nodes, duty factors of 10 per cent or less are expected giving an average power consumption of less than 10mW.

Once deployed the nodes automatically configure into a store-and-forward packet radio network. The network will automatically configure around dead nodes and obstacles that block paths between nodes (see Figure 7). This configuration is a continuing process so the network automatically reconfigures to accomodate a change to the environment such as adding or removing nodes and/or obstacles.

By operating over small distances, the network achieves high local and bisection bandwidths. The total simultaneous bandwidth available between terminals communicating through a single relay is 1Gb/s (1K relays at 1Mb/s each). The bisection bandwidth, the bandwidth between one half of the city and the other half, is 30Mb/s as there are 30 or more parallel channels across the city (one channel every 300m for 10km).

The packet radio relay nodes can provide a location service in concert with video collection nodes and/or some other means (such as GPS) for locating the nodes. Measuring round trip time between a terminal and a node locates the node on a circle. Triangulating between three relay nodes gives an unambiguous position.




Figure 7. Topology of a self-organizing network.

The nodes self-organize into a network by polling to find immediate neighbors and then iteratively exchanging routing tables to compute the transitive closure of the adjacency relation.

Each node maintains a table of neighbors that lists the serial number (SN), receive channel (RC), and power margin for each node or terminal that can be reached in one hop from that node. To generate this table, each node periodically (perhaps once per second) comes up on a channel, i , and transmits a polling request containing its SN and RC. Any node receiving the request waits a random interval (to avoid collisions) and then transmits a reply containing its SN and RC along with the power margin with which the query was received. Both nodes enter the appropriate data in their neighbor table. If a node in the table goes some number (say 5) polls without responding it is considered unreachable and removed from the table.

If many receive channels are available it is desirable (but not essential) to reduce collisions by having all adjacent nodes operate on different receive channels. This deconfliction is performed by having the node with the smaller SN switch channels if it detects a neighbor using the same RC.

Nodes compute routing tables by iteratively exchanging tables with their neighbors. Each routing table entry consists of a destination and a few (say 4) routes to reach that destination. Each route is stored as the SN and RC of the next hop of the route, and the cost of the route. Initially the neighbor table is used to create routing table entries for neighboring destinations. Neighbors then exchange and merge tables so that all nodes have entries for destinations up to two hops away. After the second exchange the routing tables extend three hops and so on. The process continues until every node has a routing table entry for every destination and a fixed point is reached.

After each routing table exchange, exchanged routes are updated by replacing the SN and RC with the SN and RC of the neighbor sending the route and incrementing the cost by the cost of sending to the neighbor (a function of power margin). For each destination the best few (say 4) routes are retained and the remainder are discarded.

4.2 Small, Cheap Sensors

4.2.1 General Considerations

Now that we have discussed data transport issues, it is time to discuss the sensors that generate the data. We will focus on imaging sensors, not because other types are uninteresting, but because the video sensor poses the hardest design problem and generates the bulk of the data.

We need a miniaturised, low-power combined sensor and communications package cheap enough to be treated as a throwaway item. The familiar CCD detector technology appears to be incompatible with the requirements of our problem for reasons of both cost and power consumption. However, a new detector technology, called CMOS active pixel arrays, appears ideally suited to our needs. The important facts about it are that: the pixel array is created by standard CMOS manufacturing steps and can be integrated on a single chip with processor and memory; individual pixels can be directly addressed just like RAM memory locations; circuit elements can be colocated with detector pixels to carry out operations, such as compression, directly. The ability to put detector, processor and memory directly on a single CMOS

chip has obvious beneficial implications for size and cost. The details of this technology and some other opportunities it presents are discussed in a companion JASON report, entitled "Unconventional Integration" (JSR-95-120 and JSR-95-121).

The primary issues in designing a sensor package for this application are size, cost and power consumption versus the bit rate needed to send useful video imagery. As we will see in the next subsection, the active pixel array technology seems to put all these parameters in an interesting range.

Although we have not had time to pursue the subject very far, it is clear that unconventional methods of emplacing the sensor packages will be needed (careful installation by trained technicians is obviously not a useful model). We like the idea of "fire and forget" emplacement in which the package is fired at, and sticks to, a convenient surface. The package, including the optics, would need a certain level of shock hardness and some means of ensuring that the emplaced sensor would have a useful field of view. Putting all the electronics, detection and RF elements on a single chip would confer a level of mechanical robustness which should make such emplacement schemes easier to develop.

4.2.2 Strawman Design of an Emplaced Video Sensor

To give a concrete sense of what can be done with the new sensor technology, we present a strawman design of a sensor/communicator package that would, supposing the optics and emplacement issues can be solved, meet the needs of our battlefield microsurveillance scheme. More details on the tech-

nology issues involved can be found in the companion JASON report entitled "Unconventional Integration". A primary goal is to make the package as small and light as possible. Given the communication needs, battery energy density is the primary determinant of size and weight. Zinc-air batteries, which store 20 Whr per cubic inch, are a very attractive choice: two cubic inches of battery volume give a 400 hour active lifetime at 100 mW average power consumption and correspondingly longer at lower power.

A conceptual miniature sensor based on this technology is shown in Figure 8. The overall size is set mainly by the battery and would be 2-3 cubic inches (the proverbial pack of cigarettes). Whether this is small enough to meet the needs of the application we have in mind is not completely obvious and must be determined by further studies of quasi-realistic scenarios. Further miniaturization is possible, but, because of the communications power requirements and the realities of battery energy density, autonomous sensors with this performance are not going to be shrunk to the size of a sugar cube.

The power budget of the sensor is as follows: A 512x512 active pixel camera draws 35 mW at video rates but only 1 mW at 1 frame/s. Whether the camera is on or off, its frame rate and the level of image compression (possibly implemented directly on the sensor chip) is controlled by network input or by trigger from other sensors. A frame rate of one per second, on average, is probably higher than is really needed (though it is important to be able to go to full video on demand). Location determination by a typical GPS receiver chip draws 150 mW for about a minute. Since this has to be done only once, upon deployment, the total energy involved is negligible and the real cost is the complication of the extra GPS hardware on board the device. Non-imaging sensors will certainly be useful and could easily be




Figure 8. Strawman layout of a miniature imaging sensor.

included in the package. Acoustic/vibration sensors would be very useful to generate an alarm signal for explosions and weapons fire. Single sensors would probably have a large false alarm rate (from vehicle backfires and the like), but correlating the output of nearby sensors might turn this into a useful tool. At a rough guess, such sensors would draw power at a rate of 1-10mW. This would be a small component of the overall power budget.

The dominant component of the power budget is the RF power needed to transmit the data. Just to recall, with conservative assumptions about efficiencies and no antenna gain, the link powers needed for typical data rates and link ranges are (at S-band): 2 mW for 10^6 bits/s at 100m, 20 mW for 10^5 bits/s at 1000m, 200 mW for 10^4 bits/s at 10 km. Depending on how the network is configured, the required communications power should be somewhere from a few tens of mW to about 100 mW.

4.3 Platforms for Sensing and Relay

As we discussed in the section on communications networks, the data path from any given sensor to the central processing unit must pass through one or more relay nodes. In the hierarchical architecture for the data relay system, the relay transceivers would most plausibly be mounted on UAVs. The required size, weight and other design characteristics extend over a surprisingly large range, depending on the precise mission assigned to the UAV.

We see three basic niches for UAVs in the urban battlefield surveillance scenario: The first is a high altitude, high power relay bird that might also carry a SAR. We'll call this the Midi UAV because it turns out to be sig-

nificantly smaller than many currently operational birds. Next, there is a low-altitude, low-power relay and optical surveillance bird which we will call the Mini UAV. Finally, there may be a need for a short endurance close-in surveillance and trailing bird which we will call the Micro UAV. High- and low-power refer to the link powers specified in our sketch of a hierarchical communications system.

Vehicle design is driven by the laws of flight and required link powers and endurance. The relevant flight laws are

$$\begin{aligned}\text{Lift : } Mg &= \rho V^2 S C_L / 2 \\ \text{Drag : } D &= \rho V^2 S C_D / 2 \\ \text{Power : } P &= \rho V^3 S C_D / 2\epsilon\end{aligned}\tag{4-2}$$

where ρ is the density of air, S is the wing area, ϵ is the efficiency of the propellor. For purposes of estimation, we take the lift and drag coefficients to be $C_L \sim 1$, $C_D \sim .05$ and also $\epsilon \sim .5$.

To design a vehicle, pick a cruise velocity and wing area. That determines the vehicle weight and the propulsion power. The energy density of the fuel then determines the endurance. If there is to be a payload, that correspondingly reduces the fuel load and the endurance. For the relay birds, the primary payload is the energy store for the relay transceiver. The following table shows the design specs for three UAVs which could fill the niches mentioned above:

	V_{cruise}	T_{endur}	Alt	Mass	S_{wing}	P_{RF}	P_{prop}
Midi	20 m/s	10 hr	10 km	15 kg	1 sqm	200W	400W
Mini	10 m/s	10 hr	1 km	1 kg	.1 sqm	1W	10W
Micro	10 m/s	1 hr	1 km	.1 kg	.01 sqm		1W

The choice of cruise velocity has to do with the distance the UAV has to

travel to get on station (a few tens of kilometers in this scenario) and the winds it has to maneuver against (not large since our vehicles fly at relatively low altitude).

A wide range of UAVs is currently available, but none quite fill the niches we are interested in. The primary military emphasis has been on comparatively big, long-range intelligence gatherers such as Raptor. For this application, we need a much more modest vehicle which should be fairly easy to develop. An important point to bear in mind is that, for the communications network application, we need a fleet of vehicles (of order ten or so) operating in concert. The fleet concept raises new control issues: the notion of one vehicle, one human controller won't work here! A single controller must manage a fleet of mostly autonomous vehicles. This is surely doable, but there is little or no experience in this area. It should be noted that bad weather will put UAVs out of business occasionally. This is a problem for a system which provides an essential service, such as communications connectivity. How serious a problem it is remains to be examined.

At this point, we would like to indulge in a small digression on the uses of antenna gain in the UAV-based communication system. The Midi UAV has enough real estate to carry an antenna with quite a lot of gain, possibly as much as 26dB at 10 cm wavelength. Of course, that will be reduced by the slant angle, etc., but it could be used to dramatically reduce the needed link power. A problem is that the antenna needs to service many links at once and seems to need to look in all directions (i.e. have $g_r = 1$ in the sense of Eqn. 4.1). The first thing to realize is that if one squirts the bits from a sensor in a small fraction of the time (with the transmitter idle the rest of the time), the peak power during transmission is higher by the squirt

ratio, but the overall energy for transmission of a given number of bits is the same. But if one combines squirt and high-gain antennas (HGA), the HGA can be scheduled to pull the various sensors, and in this way it can apply its high gain to every one of them, thus realizing the same 26 dB (or whatever) reduction in battery power for every one. This can be done in an adaptive scheduling system, and the particular transmitter can squirt when the HGA is pointing at it and sends an "out with it" tone. Furthermore, the HGA can be requested by a transmitter with something to say, by a few bits of pilot, sufficiently powerful to come into an isotropic antenna on the UAV, so that the request for HGA services need occupy only a small part of the power budget.

4.4 SAR Uses and Augmentations

SAR seems ideal for initial precise three-D mapping of the city and for real-time, all-weather vehicle surveillance (on the model of JSTARS). A useful system would provide one foot resolution, would map the city with a revisit time of a few hours and would have about 10 hour endurance. The Midi-UAV is a possible vehicle for carrying the SAR, and a bird of that size may be needed for communications relay anyway. Once the detailed city map is constructed, the main items of interest to the SAR are the locations and motions of vehicles. For that mission, continuous surveillance of *all* the major streets of the city is needed and it remains to be seen whether the single mapping SAR can serve this function as well.

To address the surveillance of moving vehicles one should consider along track, interferometric SAR. In this mode one uses two SAR antennas sepa-

rated in space along the velocity vector of the aircraft to acquire two SAR images from the same point in space, but at two different times, separated by dt . The idea can be implemented by using two separate antennas separated along the aircraft fuselage or a single displaced phase center antenna (DPCA). Co-registering the resulting complex SAR images and subtracting the phases of corresponding pixels, one acquires a phase difference $d\phi$ for each pixel and $d\phi/dt$ can be calculated, yielding an estimate for the Doppler shift of each pixel. The phase difference is related to observational parameters by $d\phi = (2\pi Bu)/(\lambda V)$ where B is the antenna spacing along the line of flight, u is the radial velocity of the target w.r.t the radar, λ is the radar wavelength and V is the aircraft speed. For an X-band SAR and a spacing B of about 3 cm between the phase centers of the two antennas (dt about 1.5 milliseconds) one can observe radial (with respect to the radar) speeds u of objects from about 0 to 40 mph without ambiguity, i.e. $d\phi$ goes from 0 to 2π . Thus, one can map moving vehicles with a single UAV SAR. This technique has been used by Goldstein, et al. (Science **246** 1282 (1989)) to observe ocean currents and the orbital velocities of ocean waves which are of the order of 0.5 to a few mph. To reduce phase noise to a useful level blocks of about 9 pixels must be averaged, but for a 1 ft. resolution SAR the resulting 3x3 pixel average resolution is still useful for resolving individual vehicles. Further details on this mode of SAR observation can be found in the reference above and in Goldstein and Zebker (Nature **328**, 707 (1987)).

Given the repeat observations of the SAR appropriate for MOBA, coherent change detection would be a very useful technique to employ. As with along track interferometric SAR, described above, two complex SAR images are collected from very nearly the same point in space, but at different times. For coherent change the time separation is of the order of from tens

of minutes to even days or weeks. As before, the phase difference between corresponding pixel blocks (3×3 to 10×10 averages to reduce phase noise) in the two images is calculated. If a given structure has not changed, then the phase difference is zero; but if the structure has been disturbed the radar echo phase will change and the structure will very likely have some phase difference other than zero. In this way one can note what has changed in a SAR scene. For example, the movement of vehicles or other transportable structures is easily noted. A related technique was used with the ERS-1 SAR satellite to note ground movement connected with the Landers earthquake in southern California. Very small changes in a structure, i.e. small fractions of the radar wavelength, cause phase changes that are very noticeable in the phase difference SAR image. Some examples would be roughening of the ground by soldiers marching over it, rotation of a tank gun barrel, use of a vehicle that is returned to the same parking place (but not exactly the same position), rotation of a radar antenna, etc. The change detection aspect makes possible very rapid identification of areas or objects that have been active at any time between the two observations. Thus, enemy activity can be detected even if the activity is stopped at the time of UAV flybys.

Let us consider in a bit more detail the operating characteristics of a UAV-borne SAR designed with the needs of the MOBA application in mind. The main novelty is that powers, altitudes and endurances are smaller than usually considered. The SAR-bearing UAV may fly at an altitude of only a kilometer or so (at least given air superiority), and have a range of only a few km on an overall target area of perhaps 10×10 km. These are smaller figures than have usually concerned SAR designers. In the preceding discussion, we assumed that the SAR actually flies over the built up area. This is good from the point of view of SAR power, mass and volume, but would expose the

UAV to Stingers, small arms fire, etc. more than one would like. There is the alternative option of having the SAR stand off from the built-up area with a range of about 10 km at an altitude of 5 to 10 km. This would require more power, mass and volume for the SAR, but would increase its survivability. This is an obvious area for tradeoffs between size and survivability.

Given a spectrum of SAR/UAV sizes for MOBA applications it should be noted that the largest SAR can be used as a public-service transmitter for other smaller SARs, operating in receive-only mode. This could prove important to avoid SAR crosstalk if there is insufficient RF channel space for all the SARs operating in the area; if RF stealth for the small UAVs is necessary; or (unlikely) if the power requirements for transmitting would be too severe for the small UAV/SARs.

We contemplate an X-band SAR (10 GHz) with a bandwidth of 500 MHz or more, yielding one foot range resolution and choose other parameters to give a matching azimuth resolution. Problems of relaying the image data are discussed elsewhere, but it is very much to the point to avoid small bandwidth in the data-relay link. The reason is that the SAR system itself can play a large part in overall MOBA communication, with its large bandwidth, good overhead lines of sight, and low duty cycle for making images.

Part of the communication needs for MOBA can be carried out by SAR in a unique way—those involving friendly location and IFF. Corner cubes a few wavelengths in size, or perhaps 10 cm at X-band and smaller at K-band, can be mounted on vehicles or even on helmets and ground sensors and seen and interrogated by the SAR to locate units in the field. The corner-cube reflectivity can be modulated to send simple messages quite covertly, if required.

A first approach to mating SARs and UAVs is simply to look at the laws of flight and the radar equation. We treat these as scaling laws only, looking for gross discrepancies in power vs size or some other figures of merit. We assume that the SAR antenna will be conformal with the underside of the wing so the wing size must scale with the antenna size. At the SAR ranges of interest it turns out that the power needed to run the SAR is a fraction of the power needed to run the UAV, and that while the UAVs fly slowly because they are small the coherent integration time needed to give an appropriate virtual antenna size is not excessive (it scales like R/v at fixed SAR frequency and ground resolution, where R is the range and v the UAV velocity).

The formulas for the lift L and drag D of a subsonic aircraft were stated at the beginning of this section. Solving the lift equation for the speed gives

$$v = (2Mg/\rho SC_L)^{1/2} \quad (4-3)$$

where M is the UAV mass. Equating the propeller thrust τ to drag yields an equation for the prime power:

$$P = v\tau/\epsilon = \rho v^3 SC_D/2\epsilon . \quad (4-4)$$

The radar equation tells us that

$$\bar{P}_t(dx)^2 A^2 T/(R^4 \lambda^2) = const. \simeq 2 \times 10^{-15} J \quad (4-5)$$

where $\bar{P}_t$ is the average radar transmit power, dx is the azimuth resolution, A is the antenna aperture area, T is the SAR integration time, R is the SAR range and λ is the radar wavelength. The constant is not universal, of course, but reflects some reasonable choices of SNR (about 100), normalized backscatter cross section (for land about 10^{-2}) and system noise temperature (about 1000K). This relation needs to be considered together with the

resolution formula

$$dx = \lambda R / (2vT) \quad (4-6)$$

We can put all these equations together (using the previous values for things like C_L , C_D and ϵ) to get a roughly consistent system design. We will also roughly equate the antenna area A with the wing area S and set the resolution at $dx=1$ foot. We are most interested in seeing where the SAR power and the UAV power lie. Taking range $R=10\text{km}$, one finds solutions like $S=A=1\text{m}^2$, $v=20\text{m/sec}$, $T=25\text{sec}$, with a UAV prime power of about 400 W and an average radar power requirement of perhaps ten percent of this (depending on SAR transmitter efficiency and other SAR power requirements). This is not the place to develop detailed designs of SAR/UAV combinations at various scales, but we believe the point has been made that these can be built to be very useful in the MOBA arena at scales rather smaller than usually considered.

4.5 Mapping and Location

4.5.1 GPS, GIS and All That

In this problem, as in real estate, the three most important things are location, location and location. All the data that is collected must have precise coordinates relative to a grid anchored in the city (and tied to the global mapping grid). This location problem has several technically distinct aspects.

First and foremost, we need a three-dimensional high resolution map (as close as one can get to a CAD description) of the fixed parts of the city. Seeing *inside* all buildings is probably asking too much (we may try to learn about the insides of selected ones), but we certainly want precise definition of their outer skin. This static mapping can be done using overhead assets, but many passes are needed to build up a three-D image. If the mapping has to be done quickly, the best option is to use UAVs carrying SAR or optical cameras.

The conversion from sensor data to a three-D model of course must be done automatically, with minimal operator intervention, and this poses an interesting technical challenge. The current state of the art is typified by the construction, from stereo overhead optical imagery, of a representation of an industrial site, for instance, so that it can be viewed, for purposes of mission rehearsal, from an arbitrary vantage point. The construction involves, not surprisingly, substantial amounts of expert intervention and tweaking. Mapping a whole city is a much more complicated process, and it will be necessary to learn how to do it in a more automated fashion. Interferometric SAR can also be used to extract three-D mapping data, but little experience exists. The SAR approach has many operational advantages (it works at night and in bad weather) and we think it should be energetically pursued.

The utility of this map comes from a layering of other information, both static and dynamic, on it in the style of what are known as Geographical Information Systems (GIS). The general idea is that the user wants most of all to know about the situation in his immediate neighborhood, so tying information of all kinds (including sensor data) to map location is helpful. An important question is the accuracy with which positions should be de-

terminated. Since munitions which guide themselves to a target using some sort of optical reference can have CEPs in the meter range, it makes sense to demand that the surveillance system produce locations with accuracies in the one meter ballpark, if possible. This level of accuracy in knowledge of where munitions are targeted and where friendly vehicles (and soldiers) are situated in principle constitutes an IFF system and should certainly go some way toward solving the friendly fire problem.

The world standard for knowing where you are these days is GPS and the whole GPS receiver apparatus can, more or less, be mounted on a chip. So GPS could supply position information to sensors as well as vehicles and this information could be reported to the database as needed. The simplest implementations of GPS don't give accuracy at the one meter level (P-code versus C-code) but differential GPS and other enhancements (viz. the Stanford/NASA GPS-based instrument landing system) can certainly do the job.

GPS and all its enhancements of course require a direct line of sight to various beacons, some on the GPS satellites themselves, others on the ground. Because of urban shadowing, some sensors or vehicles might find themselves cut off from an adequate number of beacons. What should one do in that case? It seems to us that imaging sensors can discover their location by comparing what they see with the accurate three-D map of structures which we assume has been constructed. An outline of such a scheme is presented in the next section. If it works well, it could replace GPS altogether for some purposes, with a corresponding simplification of the sensor design. Since the computation is done on the images that the sensor is transmitting in any case, the necessary processing could perfectly well be done at the central

database. No special hardware or software would have to be added to the sensor itself.

4.5.2 Location by Video Matching

The location and orientation of a camera may be determined by matching the image seen by the camera with a 3D model of the environment. Combined with video beacons or radio triangulation the camera location can be used to locate a nearby unit. This capability can be used to provide location in cases where GPS may be denied, for example by jamming or to deny location to opponents.

Matching is a search problem. The 3D database is searched to find the position (x,y,z) and orientation (t,f) for which a view of the database matches the view seen by the camera. The search can be accelerated by first matching features to find an approximate location of the camera and then matching images to fine-tune this location.

The camera transmits the image stream to a base location where a 3-D environment model is available. Key features such as buildings, roads, and other cultural artifacts are extracted from the image through a process of edge finding, clustering, segmentation and model fitting. The extracted features have an unknown perspective transformation: scale and rotation. Transform invariant parameters of the features are matched against the database. A score is generated for each feature pair and dynamic programming is used to select a group of co-located features with the best score. A rough camera position and angle is computed which aligns the image features with the

database features.

A gradient search in image space is used to fine-tune the camera position. This search is performed in a series of progressively finer scales. First, the image bandwidth is drastically reduced by low-pass filtering and gradients in camera position and angle are followed to maximize the match between image and model. The search is repeated at progressively larger image bandwidths until the best match is found on the full image. Performing the search in this scale space avoids the problem of getting stuck in a local maximum early in the search when the image is misaligned by several pixels.

4.6 Data Management, Processing and Display

4.6.1 Levels of Sophistication in Data Processing

Finally, we have to face the problem of processing the flood of data that is produced by so many sensors. How is one to extract useful intelligence from this stream, how store it and how transform it and feed it out to users? The data volume is, even in a minimal system, too great for human intervention except in the most selective sense. Ideally, one would like to automate the extraction of intelligence from this data stream in the spirit of automatic target recognition (ATR). It is our impression, however, that no systematic way of engineering such a system currently exists. Real progress is being made, of course, and the ATR-inspired approach will in the end be indispensable. In the meantime, we believe that it is important to get started with a data exploitation approach that uses minimal ATR-like sophistication, yet uses

the processing power of the computer to do something new and uniquely valuable. We have a few ideas along these lines, and experience will certainly suggest others.

We can see three fairly unsophisticated, yet clearly valuable, ways to use the system. There is a static model in which the central element is what amounts to a geographical information system (GIS) which collects data about buildings and vehicles, refers it to a map, and responds to queries. The heterogeneous nature of the information one wants to post to this database, and the problem of organizing the output of information for different classes of user, makes designing this less than completely trivial. Fortunately, the civilian GIS community is working on very similar problems and their results can no doubt be adapted to the needs of this application. There is a dynamic model in which the computer is treated as a smart switch for data streams with minimal concern for archiving or interpreting data. This is the implementation of "seeing around corners" by being able to call up, from anywhere in the field, the output of a sensor anywhere else in the field. A third model might be called "enhanced man-in-the-loop". As an example, we would offer the notion of a "Guardian Angel", an operator at the command center who follows the progress of a patrol and alerts it to activities reported by the sensor system (primarily, but not exclusively, in the immediate geographic neighborhood of the patrol's position) that could affect the mission. Experience with concrete realizations of a system will suggest more and better ideas.

Even leaving aside the classic ATR and image understanding approaches to automated information extraction, a massive data collection machine of the type we are proposing offers other opportunities. It collects what could

be a gold mine of data and thought should be given to unconventional ways of mining it. Statistical analysis and pattern matching might, for example, permit the automatic identification of suspicious patterns of movement of people and/or vehicles (preparations for a riot versus the early hours of market day). The mild innovation here is to attempt to exploit the collection of masses of information to do intelligence extraction by statistical analysis. Traffic analysis in the signals world is the model that we propose to emulate.

These more sophisticated approaches to the extraction of information from the data will presumably use vast amounts of processing power and data storage. Just how big a problem are we talking about? Supposing the global bandwidth of the collection system to be 10^7 Hz, it will take 10^2 GBytes of storage to archive one day's worth of data (a reasonable guess at the size of the permanent storage). By modern standards, this is quite manageable. As for the processing load, experience in a wide variety of contexts shows that it takes about 10^3 operations to process one image bit. At a bit rate of 10^7 Hz, this implies a processing rate of 10 Gflops. The gigaflop processor chip is almost a reality now, so an overall system processing load of this magnitude is not in itself a problem. Developing the appropriate software to do the job is the real issue.

4.6.2 Hypermedia Map Situation Display

An important aspect of the data management problem is that of displaying information to users. Again, this is a problem with many solutions at many levels of sophistication. We would like to offer a modest near-term proposal for how soldiers in the field might access information resources via

a hypermedia display.

Physically the display is a laptop computer with an RF link to the network. The primary display shows a user-centered map annotated with information and "hot links" to information feeds. The basic map depicts the local environment (streets, buildings, topography) overlaid with the positions of friendly and unfriendly units. Safe-passage and keep-out areas can also be overlaid on the basic map.

Sources of position-dependent information are graphically depicted on the map and can be selected by "clicking" an icon. For example a video feed may be depicted by shading the area of video coverage. By clicking on the camera icon, the user opens a window showing the live video feed. Position-related intelligence information can also be accessed. For example, reports on the contents of buildings, building floorplans, unit composition, etc... can be accessed by clicking on the appropriate portion of the map.

The user can add to the information database by entering a report and attaching it to a position. For example, on examining a building and classifying its purpose and occupants, the user could enter a report (possibly annotated with still pictures) into the system and associate it with the building. The user can also provide live audio and/or video feeds of his/her position. Finally, the display can serve as a communications terminal allowing the user to send or receive text and audio messages.

5 THOUGHTS ABOUT TECHNOLOGY DEVELOPMENT

5.1 More Detailed System Design

The above technology survey pretty clearly indicates that the core technologies that will be needed to build a functioning urban battlefield "micro-surveillance" system are either in hand or developing rapidly. It would be appropriate at this point to go back to the strawman system and fill in details. We did not have time in our study to do this in more than a cursory fashion, so this task will have to be deferred to a future study. In what follows, we will briefly list some observations that we think should be kept in mind in any more detailed system outline.

The most glaring omission from our study is an analysis of explicit scenarios for sensor layout. The simple sensors have a limited field of view and you certainly cannot put out enough of them to see everything everywhere. The real question is: does the number of sensors needed to do something useful exceed the number we can realistically install and service by a communications network? This is a scenario-dependent question to which a first-order answer could be obtained from an abstract simulation/gaming exercise.

There is a range of levels of performance we can aim for. To keep the system affordable, we suggest not trying to solve the general battlefield situational awareness problem. Instead, the focus should be on a simple system designed for a not-too-hostile environment. Then one can learn what

the system is good for by using it in demonstrations. One can build up from there to more demanding applications.

In our view, keeping the system cheap and simple is a primary goal: Since it won't solve all problems (at least not at first), it won't be saleable unless its cost/benefit ratio is manifestly large. The benefit is not known with certainty, and demonstrating that it is large would be a major task of any development program. As we have argued in previous sections, the hardware for such a system could be very cheap indeed: if mass market pricing were to apply (more about that shortly), one could hope for sensor/communicator packages costing in the hundreds of dollars. Of course, cost reduction plays to the strengths of ARPA's core efforts.

5.2 Civilian Security Systems: A "Killer Ap"?

At this point, we cannot help remarking that cheap imaging sensors that self-organize into a wireless local data network have obvious applications in the civilian security industry. The use of video cameras for security and surveillance is, of course, standard in certain commercial applications, but not in the home market. The reasons for their absence from the home market are cost, inconvenient installation and the problem of making use of the sensor output (how do you convert the video signal into an "alarm"?).

The sensor/communicator packages proposed above would have a very direct application to the civilian security market. The key reasons are that these systems would require no wiring, have native digital data output and would self-organize into an expandable system. The potential effects of these

properties on the cost, flexibility and effectiveness of security systems are easy to imagine.

Interestingly enough, many of the sensor, networking and data processing issues are the same for both the military and civilian applications. Indeed, the commercial market contains niches that are of similar complexity to the urban battlefield problem. The security system for a large skyscraper could have many hundreds (possibly even a few thousand) of video sensors and the problem of managing and making use of the resultant information flow is certainly conceptually and quantitatively similar to the military problem.

This civilian market is potentially huge. People are willing to spend thousands of dollars for a home security system while the cost for a skyscraper security system can be in the millions. Overall, there could be a market for $\sim 10^8$ sensor packages. Such a market, should it emerge, would turn many of the technologies needed for military systems into low-cost COTS items.

Provided the core commercial hardware and software came with the right "hooks", it could be used directly in military applications, with obvious cost benefits. It may be worth asking whether ARPA could incubate such a development by providing techno-stimulus and encouraging the adoption of appropriate data and communications standards.

5.3 A Modest Proposal

We think that it would be very valuable to do field exercises with the sort of system we have been advocating. This can be done very quickly, if

we are willing to use surrogates for the more difficult technology elements (and, in the context of controlled experiments, there is no reason not to use such surrogates). Such prototypes would serve to demonstrate the system concept, build mindshare for this type of information network, and provide a sandbox in which to experiment with new ideas.

In a demonstration, size, power, weight and some operational parameters would be compromised to get a system running quickly with existing equipment. Such a prototype will have a look-and-feel similar to a real system but with lower data and frame rates and larger, heavier, and more power-hungry components. After some experience is gained with the prototype, effort can be concentrated on improving components in the areas that give the most leverage.

Three key components would make up such a demonstration system: a self-configuring packet radio network, a wide-angle wireless video camera, and a situation display terminal.

The packet radio network can be constructed by leveraging wireless research performed at UCLA and elsewhere as well as commercial wireless network products. Existing network protocols (e.g., IP) can be layered on the network to provide compatibility with existing software. For a first prototype, data rate can be reduced, size and operating power can be increased, and security (encryption) and anti-jam considerations need not be addressed. A single module with radio, processor, and buffer can serve as a terminal as well as a relay.

A standard CCD video camera can be used in the prototype. Power constraints can be relaxed and the processor in the RF network node can

reduce data rate without compression by dropping frames to operate at a very low frame rate.

The situation display can be prototyped by using a standard Web Browser (Netscape or Mosaic) to display the situation map as a GIF image in an HTML file. A map file would provide hot links to position indexed sources of information. The problem of local information dissemination can be bypassed by feeding all available information through a standard web server.

6 CONCLUSIONS AND RECOMMENDATIONS

Our overall conclusion is that this is important: technology and culture are ripe for a major advance in military capability. There is a long way to go before a working capability is available, but we pretty much know what has to be done, so let's get going!

It is perhaps useful to point out that there is no one technological "silver bullet" here. As with stealth, one needs an intelligent integration of many "pretty good" technologies. Obviously, ARPA should keep pushing the appropriate generic technologies in the right directions. Global Mobile, MEMS and small UAVs are major examples, and they all seem to be on track. ATR, and variants thereon, will be crucial to future versions of a battlefield information system, but, at this stage, it is impossible to say which "sophisticated" information processing methods will work. We believe that useful systems can, and should, be developed without them. They will provide an experience base with which to guide development of more advanced information processing schemes.

Near-term, simple, systems have another critical function: The real problem is getting the users excited and committed to this development. We believe that the only way to do this is with field exercises and experiments. Realistic exercises will reveal what's important and what's not. In the interests of getting going quickly, off-the-shelf surrogates for the hard technologies should be used in the initial simulations. One can always proceed to more realistic demonstrations once the customer is hooked.

DISTRIBUTION LIST

Director of Space and SDI Programs
SAF/AQSC
1060 Air Force Pentagon
Washington, DC 20330-1060

CMDR & Program Executive Officer
U S Army/CSSD-ZA
Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0150

A R P A Library
3701 North Fairfax Drive
Arlington, VA 22209-2308

Dr Arthur E Bisson
Director
Technology Directorate
Office of Naval Research
Room 407
800 N. Quincy Street
Arlington, VA 20350-1000

Dr Albert Brandenstein
Chief Scientist
Office of Nat'l Drug Control Policy
Executive Office of the President
Washington, DC 20500

Mr. Edward Brown
Assistant Director
ARPA/SISTO
3701 North Fairfax Drive
Arlington, VA 22203

Dr H Lee Buchanan, I I I
Director
ARPA/DSO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Curtis G. Callan, Jr.
Princeton University
P. O. Box 708
Princeton, NJ 08450

Dr Ashton B Carter
Nuclear Security & Counter Proliferation
Office of the Secretary of Defense
The Pentagon, Room 4E821
Washington, DC 20301-2600

Dr Collier
Chief Scientist
U S Army Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0280

Dr. John M. Cornwall
Department of Physics
University of California/Los Angeles
Los Angeles, CA 90024

DTIC [2]
Cameron Station
Alexandria, VA 22314

Dr. William J. Dally
Massachusetts Institute of Technology
Dept. of EE&CS/Bldg NE43
545 Technology Square, Rm 620A
Cambridge, MA 02139

Mr John Darrah
Senior Scientist and Technical Advisor
HQAF SPACOM/CN
Peterson AFB, CO 80914-5001

Dr Victor Demarines, Jr.
President and Chief Exec Officer
The MITRE Corporation
202 Burlington Road
A210
Bedford, MA 01730-1420

DISTRIBUTION LIST

Dr. Alvin M. Despain
3273 Corinth Avenue
Los Angeles, CA 90066

Dr. Douglass M. Eardley
618 Miramonte Drive
Santa Barbara, CA 93109

Mr Dan Flynn [5]
OSWR
Washington, DC 20505

Dr Paris Genalis
Deputy Director
OUSD(A&T)/S&TS/NW
The Pentagon, Room 3D1048
Washington, DC 20301

Dr Lawrence K. Gershwin
NIC/NIO/S&T
7E47, OHB
Washington, DC 20505

Mr. Thomas H Handel
Office of Naval Intelligence
The Pentagon, Room 5D660
Washington, DC 20350-2000

Dr Robert G Henderson
Director
JASON Program Office
The MITRE Corporation
7525 Colshire Drive
Mailstop Z561
McLean, VA 22102

Dr Paul Horowitz
Harvard University
Lyman Laboratory of Physics
Cambridge, MA 02138

Dr William E Howard III [2]
Director of Advanced Concepts &
Systems Design
The Pentagon Room 3E480
Washington, DC 20301-0103

Dr Gerald J Iafrate
U S Army Research Office
PO Box 12211
4330 South Miami Boulevard
Research Triangle NC 27709-2211

J A S O N Library [5]
The MITRE Corporation
Mail Stop W002
7525 Colshire Drive
McLean, VA 22102

Dr Anita Jones
Department of Defense
DOD, DDR&E
The Pentagon, Room 3E1014
Washington, DC 20301

Mr. O' Dean P. Judd
Los Alamos National Laboratory
Mailstop F650
Los Alamos, NM 87545

Dr Bobby R Junker
Office of Naval Research
Code 111
800 North Quincy Street
Arlington, VA 22217

Dr Ken Kress
Office of Research and Development
809 Ames
Washington, DC 20505

Lt Gen, Howard W. Leaf, (Retired)
Director, Test and Evaluation
HQ USAF/TE
1650 Air Force Pentagon
Washington, DC 20330-1650

DISTRIBUTION LIST

Mr. Larry Lynn
Director
ARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. John Lyons
Director of Corporate Laboratory
US Army Laboratory Command
2800 Powder Mill Road
Adelphi, MD 20783-1145

Col Ed Mahen
ARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Arthur Manfredi
OSWR
Washington, DC 20505

Mr James J Mattice
Deputy Asst Secretary
(Research & Engineering)
SAF/AQ
Pentagon, Room 4D-977
Washington, DC 20330-1000

Dr George Mayer
Office of Director of Defense
Research and Engineering
Pentagon, Room 3D375
Washington, DC 20301-3030

Dr Bill Murphy
ORD
Washington, DC 20505

Mr Ronald Murphy
ARPA/ASTO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr Julian C Nall
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria, VA 22311

Dr Ari Patrinos
Director
Environmental Sciences Division
ER74/GTN
US Department of Energy
Washington, DC 20585

Dr Bruce Pierce
USD(A)D S
The Pentagon, Room 3D136
Washington, DC 20301-3090

Dr William H Press
Harvard College Observatory
60 Garden Street
Cambridge, MA 02138

Mr John Rausch [2]
Division Head 06 Department
NAVOPINTCEN
4301 Suitland Road
Washington, DC 20390

Records Resource
The MITRE Corporation
Mailstop W115
7525 Colshire Drive
McLean, VA 22102

Dr Victor H Reis
US Department of Energy
DP-1, Room 4A019
1000 Independence Ave, SW
Washington, DC 20585

Dr Fred E Saalfeld
Director
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

DISTRIBUTION LIST

Dr Dan Schuresko
O/DDS&T
Washington, DC 20505

Dr John Schuster
Technical Director of Submarine
and SSBN Security Program
Department of the Navy OP-02T
The Pentagon Room 4D534
Washington, DC 20350-2000

Dr Michael A Stroschio
US Army Research Office
P. O. Box 12211
Research Triangle NC 27709-2211

Dr. Jeremiah D. Sullivan
604 Burkwood Ct. E.
Urbana, IL 68101

Superintendent
Code 1424
Attn Documents Librarian
Naval Postgraduate School
Monterey, CA 93943

Ambassador James Sweeney
Chief Science Advisor
USACDA
320 21st Street NW
Washington, DC 20451

Dr George W Ullrich [3]
Deputy Director
Defense Nuclear Agency
6801 Telegraph Road
Alexandria, VA 22310

Dr. John Fenwick Vesecky
1945 Cambridge Road
Ann Arbor, MI 48104

Dr Walter N Warnick [25]
Deputy Director
Office of Planning & Analysis, ER-5.1
Office of Energy Research
U S Department of Energy
Germantown, MD 2074

Dr. Peter J. Weinberger
22 Clinton Avenue
Maplewood, NJ 07040

Dr Edward C Whitman
Dep Assistant Secretary of the Navy
C3I Electronic Warfare & Space
Department of the Navy
The Pentagon 4D745
Washington, DC 20350-5000

Dr. Ellen D. Williams
Department of Physics & Astronomy
University of Maryland
College Park, MD 20742

Capt H. A. Williams, U S N
Director Undersea Warfare Space
& Naval Warfare Sys Cmd
PD80
2451 Crystal Drive
Arlington, VA 22245-5200