

NO-A166 948


TECHNOLOGICAL PERSPECTIVES FOR AIR BASE COMMUNICATIONS
(U) RAND CORP SANTA MONICA CA W H WARE OCT 85
RAND/N-1988-AF F49620-86-C-0008


1/2

UNCLASSIFIED

F/G 17/2

NL





MICROCOPY

CHART

AD-A166 940


DTIC FILE COPY

A RAND NOTE

Prepared for

Rand

1700 MAIN STREET
P.O. BOX 2138
SANTA MONICA, CA 90406-2138



TECHNOLOGICAL PERSPECTIVES FOR AIR BASE COMMUNICATIONS

Willis H. Ware

October 1985

N-1908-AF

The United States Air Force

This document has been approved
for public release and sale; its
distribution is unlimited.

86 4 24 001

The research reported here was sponsored by the Directorate of Operational Requirements, Deputy Chief of Staff/Research, Development, and Acquisition, Hq USAF, under Contract F49620-86-C-0008.

The Rand Publication Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER N-1908-AF	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Technological Perspectives for Air Base Communications		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Willis H. Ware		8. CONTRACT OR GRANT NUMBER(s) F49620-86-C-0008
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Rand Corporation 1700 Main Street Santa Monica, CA. 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Requirements, Programs & Studies Group (AF/RDQX) Ofc, DCS/R&D and Acquisition Hq, USAF, Washington, DC 20330		12. REPORT DATE October 1985
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. NUMBER OF PAGES 133
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release; Distribution Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) No Restrictions		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Communications Networks Air Force Facilities Telephone Systems		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See reserve side		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

This Note examines the relevance of contemporary local area network (LAN) and computer-based digital telephone switch technology to the needs of CONUS airbases, in both the near term and far term. It suggests possible architectures based on such technology and concludes that in the next decade a hybrid arrangement will provide the flexibility and adaptability that differences in requirements among bases demand. The Note also considers the security aspect of handling classified information in base-level communications and concludes that security issues in a LAN-oriented base will be awkward for several years ahead. It suggests, however, that commonsense actions can be taken that will help make base communications more secure, and describes several new National Security Agency-sponsored programs that will make it possible to provide secure on-base telephone and data communications. It makes two major recommendations: (1) that a policy statement be developed outlining the context, assumptions, and guidelines for improving base communications in the next decade; and (2) that a comprehensive plan be drawn up to guide the Air Force on an evolutionary path to improved airbase communications.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

A RAND NOTE

TECHNOLOGICAL PERSPECTIVES FOR AIR BASE COMMUNICATIONS

Willis H. Ware

October 1985

N-1908-AF

Prepared for

The United States Air Force

Rand

1700 MAIN STREET
P.O. BOX 2138
SANTA MONICA, CA 90406-2138

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

PREFACE

This Note was prepared for the Air Force Communications Command as part of the study effort "Technology Planning for Future Base-Level Communications Systems" under Project AIR FORCE. Although a complete text, except for the section on security aspects, was ready in mid-1984, a number of significant developments important to security safeguards in local area networks (LANs) and other base communications matters were then under way in the communications security community; publication was intentionally delayed until this new material could be included.

The Note is divided into three parts:

- Part 1: Sections I through VII, which collectively constitute a tutorial on LAN technology, with special emphasis on technical and operational aspects pertinent to airbase usage.
- Part 2: Section VIII, which treats security aspects of LANs and other security matters relevant to airbase communications.
- Part 3: Sections IX through XI, which present the conclusions of the study; describe probable architectures for the near-term and far-term CONUS bases, plus a far-term suggestion for combat or theater bases; and offer a recommendation for Air Force action.

The reader who is already well informed on LAN technology may wish to proceed directly to Parts 2 or 3. On the other hand, some aspects of LAN technology that are significant for Air Force planners may not arise at all--or, at best, only occasionally--in the commercial world. The reader interested only in the "bottom line" may wish to proceed directly to Part 3 for projected architectures and far-reaching recommendations.

The base communications environment, as well as that of the Air Force at large, is in the process of changing from an historical world of point-to-point circuit-switched links carrying traditionally formatted teleprinter traffic to a contemporary one of multiple local area networks at Air Force installations, all-digital long-haul

COPIES
IN FILE 3

A-1	
Dr 1	or sid
Codes	

networks, electrical interfaces across the base perimeter, and high-speed data links. The new techniques collectively support an interconnected array of host computers, individual work stations or terminals, personal computers, and a broad range of information services. Many of the established ways of thinking about communications will have to change, and many new concepts will replace old ones. The discussion in this Note, quite apart from its primary purpose, is intended to help create the mindset that is consistent with a new environment of highly interconnected and integrated information and data services.

The period of this study paralleled the period during which significant changes were taking place in Air Force communications and computer organizational alignments. Among other things, the Air Force "SI community"--Information Systems--was being formed, and thousands of personnel were transferred from the major commands (MAJCOMs) to the Air Force Communications Command (AFCC). The Assistant Chief of Staff for Information Systems (USAF/SI) was consolidating his position and responsibilities; important position and guidance documents began to appear. The career fields for data processing and communications were combined into a single field of information systems.

Importantly, the position and recommendations in this Note were established independently of, and in some cases prior to, similar Air Force actions. The findings of the study, however, are generally consistent with the spirit of the actions taken by AFCC and USAF/SI, although they are not always congruent with them.

This is the third in a series of Rand Notes relating to Air Force base communications matters. The others are:

N-2162-AF, *Information Systems: The Challenge of the Future for the Air Force Communications Command*, Stephen M. Drezner and Willis H. Ware, May 1984.

N-2164-AF, *Base Communications Issues for the 1980s*, Willis H. Ware, Robert M. Paulson, and Martin M. Balaban, September 1984.

SUMMARY

This Note examines the relevance of contemporary local area network (LAN) and computer-based digital telephone switch (CBX) technology to the needs of CONUS airbases, in both the near term and far term. It suggests possible architectures based on such technology and concludes that in the next decade or so a hybrid arrangement will provide the flexibility and adaptability that differences in requirements among bases demand. "Hybrid" implies that LANs will be used where appropriate, that modern telephone technology will be used not only for voice but also for data transmission, that wideband links will be installed only where demanded by data rates, and that bases will not make a dramatic conversion to "all-coaxial wideband media." Rather, the true wideband services--video, closed-circuit television, and extremely high computer data rates--are likely to exist independently of telephony and terminal-level data rates.

Such a hybrid arrangement would

- Use a LAN for each community of interest that is relatively localized geographically and communicates heavily within itself but has limited connectivity elsewhere.
- Exploit the base twisted-pair cable plant and the base CBX (now being installed by the Air Force Communications Command (AFCC) through various SCOPE programs) to provide connectivity among LANs and with off-base long-haul circuits.
- Use broadband cable technology selectively where very high data-rate computer communications are necessary or where video services are required.

While there is a surfeit of LAN and CBX commercial equipment from which to choose, there are special Air Force needs that may make airbase installations of LANs somewhat different from those outside the military. For example, survivability can be a major issue; the vulnerabilities of commercial products and architectures must be considered. Special

services such as pagers or paging-printers may be required; and inter-LAN connectivity will be a major concern because a base subscriber should be able to reach any subscriber on any other base. Thus, technical issues such as protocols, internetwork gateways, inter-LAN bridges, and inter-LAN addressing are more immediate and important problems to the Air Force than they presently are in the commercial world. The handling of classified traffic is also a major concern, especially since crisis events or world situations can suddenly impose classification restrictions on normally unclassified data and procedures.

This Note examines the security aspect of handling classified information in base-level communications and concludes that security issues in a LAN-oriented base will be very awkward for several years ahead. It suggests, however, that commonsense actions can be taken that will contribute to greater securityworthiness of base communications; for example, a LAN should be installed according to protected-wire doctrine, which affords its transmission medium physical protection against tapping or damage; LAN and CBX equipment can be physically protected against casual access or against unmonitored maintenance by uncleared nonmilitary personnel; alternate power sources can be provided and power facilities can be protected to guard against a denial-of-service threat.

The Note describes several new National Security Agency-sponsored programs that will make it possible to provide secure on-base telephone and data communications:

- The STU-II program, which can now provide an instrument for secure communications to 2400 bits per second, but at a unit cost that will limit it to only the most urgent base applications.
- The STU-III program, which in approximately two years will be able to provide desk-top instruments for secure voice and data communications to (anticipated) 4800 bits per second and at a unit cost of approximately \$2000.

- The Commercial COMSEC Endorsement Program, which will, in cooperation with vendors of LAN equipment, lead to security controls (including encryption) embedded directly in commercially available equipment.
- The special arrangements that have been concluded to permit NSA-approved consumers to purchase STU-II, STU-III, and other key generators such as the KG-84 directly from their manufacturers.

Finally, the Note concludes that there are a number of reasons why the evolution of on-base communications can only go in the direction of the hybrid architecture proposed. These reasons include:

- The sunk investment of the existing on-base telephone cable plant.
- The requirement that the transition to any new architecture must be evolutionary in order that base communications not be degraded even temporarily.
- The fact that the Air Force will of necessity utilize commercially available equipment and not fund the development of specialized equipment for its CONUS bases.

The Note makes two major recommendations:

- The Assistant Chief of Staff for Information Systems (USAF/SI) write and widely disseminate a policy statement that sets forth the general context, assumptions, and guidelines which will be used to direct the improvement of base communications over the next decade.

The intent of this recommendation is to provide a short readable document so that all decisionmakers concerned with base communications matters can readily read and grasp its content, and so that funding and programmatic decisions can be easily fitted into an understandable cohesive framework of actions and goals.

- The Assistant Chief of Staff for Information Systems, in conjunction with Headquarters/AFCC and relevant parts of the Electronic Systems Division (ESD) (particularly the AFLANSPO¹ and ESD/OC), prepare a comprehensive plan for assuring that the Air Force, its major commands (MAJCOMs), and its base planners will be properly guided on an evolutionary path toward a significantly improved and more contemporary airbase communications environment.

The plan must address:

- Organizational issues (such as the presently uncontrolled and uncoordinated installation of LANs).
- Technical issues (such as protocol standardization and the development of any needed gateways).
- Security issues (such as TEMPEST and the smooth introduction of new security measures as they become available).

It must provide for:

- More extensive and more timely information to planners on-base and at Headquarters/AFCC.
- Acquisition of better data on base-level information flows and rates.
- Support of such R&D efforts as might prove necessary to offset shortfalls in commercial equipment or to support theater bases and their special needs for an enduring communications capability, both voice and data, in spite of damage.

In addition, AFCC should institute two specific offices:

- A "Cable Management Office" whose function is to manage the cable plant effectively as it is replaced, upgraded, extended, or replaced by new (e.g., fiber-optic) technology.

¹Air Force Local Area Network System Project Office.

- A "LAN Management Office" to bring technical cohesiveness to LAN selection and installation, and to be responsible for various technical issues related to them.

This Note also looks to the far term and observes that present developments in cable-based technology could, over time, completely supplant LAN/CBX-based arrangements. In such an environment, an enormously broader range of on-base information services could be provided--for example, delivery of specialized training and educational materials basewide including to the flight line or maintenance shops; basewide or selective alerting in emergencies; electronically based discussion groups for either recreational or job-related purposes; prompt transmission of medical records from place to place; maintaining specialized databases relevant to the continued functioning of the base; electronic distribution of specialized base-level news bulletins; plus convenient personal services such as consumer ordering, banking actions, entertainment, and access to public databases.

While the Note does not address the special circumstances of combat theater bases in general, it outlines a scheme that exploits packet-switched and ultrahigh bandwidth technology to provide extremely enduring communications.

ACKNOWLEDGMENTS

Many people have helped to bring this lengthy document on a complex subject successfully to fruition. In particular, the material for the discussion of packet technology in the context of base-level broadband communications was provided by Mr. Paul Baran, formerly of The Rand Corporation and currently Chairman of the Board of Packet Technologies, Inc., Cupertino, California.

Part 1 of the study was initially organized and drafted by David Leinweber, now of Leinweber & Company, Los Angeles, California. Subsequently, the material was revised and integrated into the final document.

Mr. P.A.T. Bibb, Jr., of TRW Defense Systems Group, Falls Church, Virginia, was very helpful in reviewing and providing material for Sec. VIII. Lieutenant Philip Mellinger (USAF), through the courtesy of the National Security Agency, provided material that is the basis of Figs. 11 and 12.

During the time period of our study, Mr. Charles J. Ludinsky and colleagues of The MITRE Corporation, Bedford, Massachusetts, were also conducting base-level communications studies. Frequent interaction with the MITRE team has been very productive in making this Note more comprehensive and complete.

The Note has been through very many drafts, with extensive revisions. It is a pleasure to acknowledge Delores Stimbert, who skillfully handled all the details of preparing each successive revision, assuring that references and footnotes tracked properly through every version, making seemingly endless changes, and preparing the final manuscript.

GLOSSARY

ADP	Automatic data processing
AFCC	Air Force Communications Command
AFLANSPO	Air Force Local Area Network System Project Office
AFLC	Air Force Logistics Command
AFSC	Air Force Systems Command
ARPANET	DARPA packet-switched data network
AUTODIN	Automatic Digital Network
AUTOVON	Automatic Voice Network
BLITS	Base-Level Information Transmission System
CBX	Computer-based exchange
CCITT	Consultative Commission on International Telephone Transmission
CCTV	Closed-circuit television
COMPUSEC	Computer-system security
COMSEC	Communications security
CONUS	Continental United States
CSMA/CD	Carrier sense multiple access with collision detection
DARPA	Defense Advanced Research Projects Agency
DDN	Defense Data Network
DoD	Department of Defense
ESD	Electronic Systems Division
Ethernet	Trademark for Xerox Corporation's local area network
IP	Internetwork protocol
IPLI	Internet private line interface
ISA/AMPE	Inter-service agency/automatic message processing environment
ISO	International Standards Organization
LAN	Local area network
MAJCOM	Major command

MEITS	Mission-Effective Information Transmission System
MILNET	Military Network (formerly part of the ARPANET and eventually to merge with the DDN)
MINET	Movement Information Network
MITRENET	MITRE Corporation Network
NATO	North Atlantic Treaty Organization
OA	Office automation
OSI	Open system interconnect
RFI	Radio frequency interference
SACDIN	Strategic Air Command Digital Information Network
SCOPE DIAL and SCOPE EXCHANGE	AFCC programs to replace old base telephone switches with modern digital ones
STU-II and STU-III	Secure telephone units developed by the National Security Agency
TAC	Terminal access controller
TCP	Transmission control protocol
TDM	Time-division multiplexed
TEMPEST	Test for control of unprotected classified electromagnetic circuits
TFS	Traffic-flow security
ULANA	Unified local area network architecture
USAF/SI	Assistant Chief of Staff for Information Systems, Hq/USAF
VHSIC	Very-high-speed integrated circuits

CONTENTS

PREFACE	iii
SUMMARY	v
ACKNOWLEDGMENTS	xi
GLOSSARY	xiii
FIGURES AND TABLE	xvii

Part 1: A TUTORIAL ON LAN TECHNOLOGY

Section

I. INTRODUCTION	3
Threat Environments	3
Changing Mix of Voice and Data	4
Integration with External Networks	4
Transition and Expansion	5
Technological Options	5
Hybrid Solutions	8
II. LOCAL AREA NETWORKS FOR BASE COMMUNICATIONS	9
Introduction	9
Broadband Versus Baseband Local Networks	12
Technical Aspects of Digital Local Area Networks	14
Network Topologies	19
Network Control	22
Reliability Considerations in Network Control	23
Other LAN Advantages	25
Radio Backup	25
Using Radio to Minimize Peacetime Delivery Times	26
Commercial LAN Products	27
III. NETWORK PROTOCOLS	28
Introduction	28
Protocol Choice	30
IV. INTERCONNECTING NETWORKS	33
Subnetworks and Bridges	34
Internetwork Addressing	37

V.	VIDEO AND VOICE APPLICATIONS ON LOCAL AREA NETWORKS	39
	Video	39
	Voice	40
VI.	INTEGRATED VOICE/DATA SWITCHES	42
	Resource Sharing	43
	Submultiplexing	44
	Compatibility with Existing Cable Plants	44
VII.	HYBRID ARCHITECTURES	46
	Introduction	46
	Hybrid Architectures in the Airbase Environment	48

Part 2. SECURITY ASPECTS

VIII.	SECURITY ASPECTS IN THE LAN ENVIRONMENT	55
	Introduction	55
	Communications Security	57
	Computer Security	58
	LAN Security	62
	Terminal Security	70
	Base-Level LANS	74
	Near-Term Security Possibilities	75
	New Security Initiatives	82
	Other Technical Options	90
	Base Communications Payoff	91
	Summary	94

Part 3. CONCLUSIONS AND RECOMMENDATIONS

IX.	FUTURE ARCHITECTURES	97
	Near Term	98
	Near-Term Architecture	106
	Cost Considerations	107
	ULANA	108
	Telecommunications Center	109
	Far-Term CONUS Bases	112
	Far-Term Combat Bases	119
X.	CONCLUSIONS	123
XI.	RECOMMENDATIONS	126
	Declaratory Policy	126
	Implementation Plan	128
	Plan Components	129

FIGURES

1.	The "M x N Problem" and a LAN as One Solution to It	15
2.	Network Topologies	20
3.	The Subnetwork Concept	35
4.	The Structure of a Bridge	36
5.	The Long-Distance Bridge	38
6.	Distributed Switch Architecture for a Military Environment	45
7.	Two LANs Internettted Through a Voice/Data Switch	47
8.	Basewide LAN and Voice Telephone Switch	49
9.	Hybrid LAN-Voice/Data Switch Architecture	51
10.	Computer Network Security Vulnerabilities	61
11.	Local Area/Office Automation Network Security Vulnerabilities	63
12.	LAN Security Aspects and Guidance	77

TABLE

1.	Representative Local Area Networks	7
----	--	---

Part 1

A TUTORIAL ON LAN TECHNOLOGY

I. INTRODUCTION

Base communications system designers must deal with a number of issues, including the anticipated growth of data communications requirements, the necessity to utilize existing facilities in the short run at least, the need to interoperate with a variety of U.S. and allied communications networks, and the emergence of a wide range of new technologies which greatly expand the options available for base communications. Much of the current base communications plant was designed and installed during a period when telephone traffic was the only on-base form of electrical communication. As the volume of data traffic has grown, in some cases displacing a portion of the existing voice load, the communications planning issues have become more complex. The proliferation of new communications technologies plus commercially available communications products further complicates the problem. In Part 1 of this Note (Secs. I through VII), we examine some aspects of USAF on-base communications and their compatibility with the new technologies available today and expected in the near future. Part 2 (Sec. VIII) treats security aspects of base communications, especially local area networks (LANs); and Part 3 (Secs. IX through XI) presents architectural proposals, conclusions, and recommendations.

THREAT ENVIRONMENTS

Two distinct environments must be considered: (1) bases in the CONUS¹ for which technological concerns center on capacity, operational reliability, and peacetime efficiency; and (2) theater bases, which additionally must be designed to sustain capability during hostilities. Bases in forward areas, such as those of USAFE, also have more extensive requirements for interoperability with allied forces and for communications security. Hardening to survive an electromagnetic-pulse threat would be desirable for obvious reasons, but hardening

¹Continental United States (a brief glossary of specialized terms is found on pp. xiii-xiv).

against direct nuclear attack will not be considered a priority in this discussion.

It is unreasonable to expect systems designed for commercial application to be robust in hostile military situations. However, it is reasonable to seek modifications or additions to the extensive available commercial products that can be made to accommodate military needs. In some instances, minor system changes can facilitate the use of technologies not specifically or originally intended for military applications. In others, wise choice of system details, for example, various redundancy arrangements, will accomplish the same purpose.

CHANGING MIX OF VOICE AND DATA

Although specific functional requirements for base communications systems are still evolving, it is generally realized that they will have to accommodate a growing need for digital communications. As computing and other information technologies become more integrated into the activities of a base, digital messages of all categories (computer-to-computer, computer-to-human, human-to-computer, human-to-human) will become the rule rather than the exception. It will be essential to provide also for continuing analog voice traffic until existing cable plants and other equipment are replaced with more contemporary or advanced facilities or until digitized voice dominates. It will be essential as well to allow for concurrent or alternate voice-data traffic for those applications that need it.

INTEGRATION WITH EXTERNAL NETWORKS

At present, the base telecommunications center is the link to the world outside the base; it will probably continue to be the principal, although not necessarily the only, gateway to the DoD long-haul networks. Therefore, most on-base systems must ultimately interface directly with the center. Moreover, they must be technically and operationally integrated with the functions of the base telecommunications center to fully modernize the handling of record traffic and to exploit contemporary technology. Such a fuller integration will be an essential aspect of the evolving base communications structure.

TRANSITION AND EXPANSION

Since on-base communications requirements are changing and evolving and cannot be predicted precisely, the system ought to be easily expandable. Much equipment in place today was procured nearly 30 years ago, and it must be presumed that a similar lifetime will be expected of the equipment purchased to replace it. Over such an extended period, things are certain to change significantly. Thus, low-cost systems that sacrifice potential for growth may prove a false economy in the long run.

While much of the existing communications equipment on bases today can be considered fully amortized or at end of life, many other facilities representing a substantial investment are in the midst of their useful lives, in particular the cable plant. Reconciling the continuing use of existing facilities with the appropriate application of new technologies will be one of the most important and challenging issues facing the base communications planner. While it would clearly be preferable to incorporate existing voice and data facilities into newly installed technology and to allow evolutionary growth to a broader scope of electrically delivered services, it is unlikely that the Air Force will completely abandon the present on-base communications plants for some wholly new approach.

TECHNOLOGICAL OPTIONS

It is possible but surely not desirable to approach the planning problem as a series of add-on or piecemeal steps from the present base scheme. The risk is that a narrowly focused decision now might not fit or might even conflict with some later action. For example, data channels can be accommodated through a conventional analog telephone plant by the use of modems at each data source. Furthermore, the existing aging telephone switches could be overhauled and expanded, as long as parts can be found. However, the combination of these two approaches is not a viable long-term strategy for a number of reasons, including cost, transmission speed, long-term reliability, and lack of channel capacity for additional expansion. It would involve adding cable and switching equipment, and possibly enlarging physical

facilities to increase capacity; there would be increasingly expensive maintenance of obsolete switches, plus extensive use of modems to handle data traffic. The need for a modem at each terminal and another at each port on each computer to which the terminal must communicate would become a major cost and operational burden as the number of terminals and computers grows. Analog transmission also limits the data rates achievable and provides no service for high-speed data users. Finally, major growth would be difficult to achieve gracefully.

A second option would be to employ integrated voice and data switching. This would allow for the use of much existing cable and termination equipment, but would eliminate the need for expensive modems at both ends of every data channel. Such an approach can accommodate a wide range of channel speeds. In addition, integrated voice/data switches are a part of a modern electronically controlled telephone system, which of itself can add a wide range of user features and conveniences to the base voice network.

Yet a third option would be bus-oriented local area networks (LANs), as typified in the commercial world by well over a hundred product lines. The details of a few commercial LAN products are summarized in Table 1. The local area network, really a conceptual elaboration of the internal bus common to most contemporary computer architectures, offers inexpensive high-speed data communications and, as now implemented, some capacity for voice as well. Other technologies, such as distributed mesh networks like the ARPANET and MILNET,² have reliability advantages over bus-oriented networks. They will become more and more feasible economically for smaller user

²The ARPANET is a data communications network established in 1969 by the DoD's (then) Advanced Research Projects Agency (ARPA) to interconnect computer resources at selected research centers at substantially lower costs than other systems available at the time. The ARPANET is a fully operational multinode network that interconnects over 200 host computers in the United States, the United Kingdom, and Norway. ARPA became the Defense Advanced Research Projects Agency (DARPA) in 1973. Recently, part of ARPANET was physically split apart to become MILNET, which services military users and will eventually be merged into the Defense Data Network that is now under construction. For a description of the Defense Data Network, see Stephen T. Walker, "Department of Defense Data Network," *SIGNAL*, October 1982, pp. 42-47.

Table 1
REPRESENTATIVE LOCAL AREA NETWORKS^a

Name	Vendor/Developer	Topology	Access Method ^b	Speed (Mbps)	Transmission Medium	Maximum Distance (km)	Maximum Stations	Packet Size (hdr/data - bytes)
Clusterbus	Nestar Systems, Inc.	bus	TDMA(CSMA)	0.2	16-wire flat cable	0.3	65	7/256
Ethernet	Xerox Corporation	bus	TDMA(CSMA/CD)	10.0	50-ohm coax	2.5	500	18/1500
Local Net 20	Sytek, Inc.	bus	TDMA(CSMA/CD)	0.128	75-ohm coax	50.0	24,000	13/64
MITRENET	MITRE Corporation	bus	FDM, TDMA (CSMA/CD)	1.2	CATV coax	--	--	10/118
Net/One	Ungermann-Bass, Inc.	bus	TDMA(CSMA)	4.0	RG-8A/U	1.21/segment	200/segment	--
Wangnet	Wang Laboratories, Inc.	bus	FDM, TDMA (polled)	12.0	50-ohm coax CATV coax	3.0	--	--
Z-net	Zilog Corporation	bus	TDMA(CSMA/CD)	0.8	RG-59 75-ohm coax	3.0	255	27/579

SOURCE: Paul Kinnucan, "Local Networks Battle for Billion-Dollar Market," High Technology, November/December 1981, p. 67.

^a This table is intended to illustrate the variety of LAN schemes, both commercial and private; it is not meant to be comprehensive, and some of the specifications are provisional.

^b FDM = frequency division multiplexed

TDMA = time division multiple access

CSMA/CD = carrier sense multiple access with collision detection

communities--even at base level--as the cost of the computing power required in the node switches drops.

HYBRID SOLUTIONS

Each approach has its own virtues and limitations, but fortunately a mixed approach is possible. Hybrid configurations--mixing network technologies with a versatile and modern cable-and-switch backbone system--allow advanced services to be provided to communities that need them without incurring the cost of forcing an excessive level of technology for applications that do not really require it.

We will next examine in more detail the technologies of LANs, voice/data switching systems, and hybrid architectures for meeting the future demand for USAF on-base communications. We then treat the security aspect of LANs, and finally return to the overall architectural issue.

II. LOCAL AREA NETWORKS FOR BASE COMMUNICATIONS

INTRODUCTION

Traditionally, an airbase has dealt separately with its telephone network, dedicated other networks for data flows or alarm systems, various radio networks, standalone video networks, and various off-base networks such as AUTODIN, AUTOVON, and the specialized data networks required, for example, by the Air Force Logistics Command (AFLC). More recently, base tenants have acquired word-processing systems which are sometimes networked together locally. Other tenants have acquired LANs; electrical interfaces have been established between on-base and off-base services; and, importantly, digital technology is gradually replacing traditional analog technology.

As the new technology has been introduced, or perhaps as a result of its presence, there has developed an increasing requirement for interconnectivity among all communicating participants on a base--computer-terminal oriented systems, voice oriented systems, computer hosts, and other data equipment.

One significant new technology is the *local area network*,³ which can be characterized as a communications network, limited in geographic scope, that provides interconnection over inexpensive media, sometimes twisted-pair wires but generally coaxial cable. As the area to be served grows, different technology will come into use, and the larger network will be called a *wide area network* (WAN). Finally, the LANs and the WANs will be interconnected among themselves and with distant places on a point-to-point basis through circuit- or packet-switched *long-haul networks*.

There is a view in the Air Force that the LAN is not only appropriate but also the only technology for some aspects of base-level data communications. LANs can provide flexible terminal access for computer users and can electrically supplant some, even much, of the

³For a brief overview of the subject, see William Stallings, "Local Network Overview," *SIGNAL*, January 1983, pp. 39ff.

current manual distribution of messages from the base telecommunications center to recipients. An on-base "electronic mail" system can also displace some traffic from the interoffice telephone and physical mail systems. Fortunately, the commercial marketplace now offers a wide range of LAN products^{4,5} and thus the question becomes, What attributes must be considered by Air Force Communications Command (AFCC) planners?

Many present designs are generally able to meet the requirements of the Air Force for CONUS bases.⁶ Constraints will be, for example, the maximum distance between nodes without having to establish multiple networks with interfaces (i.e., the maximum physical size of a LAN), the maximum number of communicating terminals attached to the LAN, and the data rates required to handle bandwidth-intensive applications such as video surveillance or high-speed data transfer. These applications will demand a disproportionate share of LAN capacity and will inconvenience users that have lighter data flows. A particular LAN design for a specific application will be influenced by the mix of low-speed data, high-speed data, voice, and video that the network must accommodate. However, unless care is taken, some choices will preclude further growth opportunities.

In contrast, the LANs for airbases in Europe or other areas of potential conflict must operate and survive under hostile conditions that are clearly not anticipated by the designers of commercial peacetime systems, and there are other technology details to be considered as well. For example, if the network is to remain useful under a conventional attack, reliability enhancements, greater connectivity among network links, and backup modes that are not required in civilian or peacetime systems will probably be essential.

⁴Paul Kinnucan, "Local Networks Battle for Billion-Dollar Market," *High Technology*, November/December 1981, pp. 64-72.

⁵Harvey Hindin and Tom Manual, "Local Networks Will Multiply Opportunities in the 1980s," *Electronics*, January 27, 1982, pp. 89ff.

⁶For a comprehensive treatment of computer-oriented communications networks, see D. W. Davies and D.L.A. Barber, *Communication Networks for Computers*, John Wiley & Sons, New York, 1983.

Within the wide range of commercially available LANs, there are many differences in system architecture, topology, and transmission media. Basic architectures include fully distributed packet-switched networks such as the ARPANET; circuit-switched networks such as those offered by Rolm, Datapoint, and Northern Telecom; baseband coaxial cable networks such as the Xerox Ethernet; and broadband coaxial networks such as the MITRENET, Wangnet, and Sytek Net. Baseband networks operate with a bandwidth of a few tens of megahertz, and the interface electronics (commonly called a Bus Interface Unit, or BIU)⁷ connect directly to the transmission cable. Broadband networks operate with a larger bandwidth, typically a few hundreds of megahertz. The latter are comparable to cable television systems in the sense that multiple channels can be present, but the cable must be interfaced by a BIU that technically resembles a cable-TV tuning device. This is required because most digital circuitry does not operate at the high speeds possible in a broadband network.

Transmission media include conventional coaxial cable like that used in radio and antenna installations, aluminum-shielded coaxial cable like that used by cable TV systems, wire pairs, fiber optics, and radio channels. Network topologies range from highly connected networks, such as the ARPANET and MILNET, to bus networks, such as MITRENET or Ethernet, in which all communicating stations are connected to a single common bus.

Ring networks, in which nodes pass messages circumferentially, are popular in Europe. Other LANs are configured as stars with a centralized control node, or as more complex hybrid topologies. Some designs are fully distributed with no central control; others require some sort of central control.

On an airbase, the architecture for a base-level network must accommodate the unique function of the base telecommunications center, which now provides the traffic gateway to the external world and to larger DoD networks. The internetting function (e.g., among bases) will

⁷Terminology is not standardized. Picking one word from each of these groups [*bus*, *network*] [*interface*] [*device*, *unit*] yields the commonly used phrases.

be of special importance during wartime, when the internal point-to-point connectivity of the on-base network will also be most critical.

Before any major USAF procurement commitments are made for use in Europe or other potential areas of conventional conflict, it is important for the communications planner to determine which combination of architectural structure, transmission medium, and network topology will provide the greatest survivability and wartime utility in degraded modes of operation. In the European threat environment, networks must be able to operate under armed assault with conventional weapons and will be potentially subject to electromagnetic pulse effects; the latter threat is a subject of some controversy but is simply noted here.

BROADBAND VERSUS BASEBAND LOCAL NETWORKS

The choice of broadband or baseband depends primarily on the amount and nature of the traffic that the network must carry. In addition to data traffic, broadband networks can also be used for surveillance video, two-way or multi-way video teleconferencing, and high-speed facsimile traffic. The price for such variety is a requirement that more costly modems be used for connecting data terminal equipment. The broadband network is in effect a multiplicity of parallel high-rate channels which are allocated as required to each data path. Thus, many simultaneous channels can be maintained. In contrast, a baseband network shares its bandwidth by allocating it to different users at different times at millisecond rates.

Since a broadband system is essentially a radio frequency system operating typically at a few hundred megahertz, it generally uses a stronger solid-aluminum-shielded coaxial cable which, as discussed later, is less vulnerable to physical damage than the plastic insulated cable used in commercial baseband networks. However, there is no reason that a baseband network could not also be implemented with the stronger cable if necessary for a military application.

Baseband networks are easier to install and implement in a small configuration, since they require only the cable and some standard interfacing equipment. A broadband network raises more difficult design issues because of the high frequencies used. In particular, more consideration must be given to the problem of radio frequency

interference (RFI) with other systems that may operate in close frequency and even physical proximity to a broadband network. This is a situation not normally encountered in the office environments where many broadband systems are installed.

The capacity of a baseband network is of the order of 10 megabits/sec, with some fraction (usually 5 to 20 percent, depending on protocols and traffic loading) consumed by overhead, i.e., system-level non-user traffic necessary for making the system function. However, the maximum data rate may be limited not by the network's inherent capacity, but by the network interface itself. Therefore, in choosing a baseband network, one must first determine the maximum data rate between any two communicating nodes and ascertain that an interface capable of supporting the data rate is available. In addition, transmission delay between subscribers can grow on a heavily loaded baseband network due to the scheme used to manage entry onto the network (e.g., contention resolution, which requires a transmitting node to wait for a clear space, or token passing, a scheme in which each node sequentially awaits its turn). Broadband networks with dedicated channels do not have this problem, although the same effects can arise within any one channel that is time-shared among subscribers.

The inherent bandwidth of a broadband cable is typically 250 MHz, which can be somewhat misleading in that the modems used to connect subscribers to the network limit the system's data-transmission efficiency. A more costly modem can utilize its segment of the channel more efficiently (e.g., a 1.5 megabits/sec modem using only about 1.5 MHz of bandwidth typically costs three times as much as one using 6 MHz); lower-speed modems are less efficient in bandwidth utilization (e.g., a 19.2 kilobits/sec modem may require 100 KHz of bandwidth). In addition, guard bands are required at the edge of each channel to minimize interference.

High-speed facsimile, surveillance video, and video teleconferencing are bandwidth-intensive applications that cannot be accommodated on a baseband network. If such services are required in a netted communications system, only a broadband network will be acceptable. Interconnection of up to a few hundred relatively low-speed data devices (e.g., computers, terminals, word processors,

printers, network gateways) is probably easier and less expensive with a baseband system.⁸

TECHNICAL ASPECTS OF DIGITAL LOCAL AREA NETWORKS

Local area networks are distinguished from long-haul networks by a number of technical and operational details. The most obvious is geographic extent, which is typically limited in a LAN to a few kilometers, substantially less than the global or continental reach of long-haul networks. A LAN is conceptually an extrapolation of the internal bus typical of computers, unlike the traditional long-haul communication network which is derived from point-to-point switched-circuit concepts. A second major distinction is the availability of cheap bandwidth--as on an internal computer bus--which eliminates the need in a LAN for many of the complex strategies used in long-haul networks to manage the limited wideband of expensive long-distance communication channels. Finally, there are multiple transmission paths between given points for some LANs. The latter two characteristics make it essential that control of routing be internal to the message structure. Hence, space for such control information must be included within each message of a LAN and is one aspect of system overhead. In exchange, though, processing at the nodes is simpler and the BIU can be less costly.^{9,10}

A fourth characteristic of the LAN is the ability to interconnect them easily by means of gateways to other networks of similar or other kinds. As shown in part (a) of Fig. 1, connecting a variety of equipment to a number of networks can require a large number of network

⁸For a more detailed discussion of the issues regarding the broadband/baseband decision, see Thomas E. Krutsch, "A User Speaks Out: Broadband or Baseband for Local Nets?" *Data Communications*, December 1981, p. 105; Ronald Gibson, "Comparing Features Aids Selecting Broadband Local Net," *Data Communications*, April 1982, pp. 127ff.

⁹For an excellent technical introduction to LANs, see David D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, pp. 1497-1516.

¹⁰A compendium of recent papers on LANs and related topics by K. Thurber and H. Freeman is published in *Local Computer Networks*, 2d. ed., IEEE Computer Society, 1981.




Fig. 1—The "M x N problem" and a LAN as one solution to it

In (a), each of three hosts at a particular site is to be connected to three long-haul networks; each host must implement the communication protocols for, and be equipped with a hardware interface to, all three networks; there may be nine different interfaces and nine protocol implementations in all. In (b), each host needs only one hardware interface and one protocol implementation; the gateway machines each handle communication between one long-haul network and the LAN.

SOURCE: David D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, p. 1499.

interfaces. However, if the heterogeneous devices are all connected to a single LAN (part (b)), which in turn has a single gateway to all external networks, only one gateway needs to be developed and maintained. Individual gateways can interface to one or more external networks, and redundant gateways can be used for extra reliability.¹¹ The use of gateways results in a substantial cost savings over the simultaneous interfacing of many different kinds of equipment to multiple networks; and while Fig. 1 refers to off-base long-haul interoperability, the gateway concept might be attractive for some on-base circumstances as well.

To evaluate the LAN and its potential contributions, it is helpful to contrast some of its functional aspects with traditional communications approaches. Much telephony and data communications, even today, are on a *circuit-switched* basis. This means that there is literally "a copper wire" between subscribers--a channel on a microwave link, a time position in a time-division-multiplexed (TDM) channel, a twisted-pair wire, a portion of the bandwidth on a satellite transponder, or some combination thereof. All are interconnected with amplifiers and other specialized electronic equipment. Importantly, for the duration of a call or a data transmission, all parts of the end-to-end connection (e.g., the time position in the TDM digital stream, the twisted pair from the nearest telephone switch to the subscriber, the microwave channel) are dedicated to that call or transmission. Because the user has full utilization of the assets, he is charged correspondingly. Thus, it is up to the subscriber to keep the circuit loaded to its maximum capacity to get his money's worth. In a way, the silences inherent in a conversation or in speech per se are costly to a communications user. The so-called statistical multiplexor alleviates this situation slightly by taking advantage of silences to combine a few voice conversations on a group of voice circuits, but the circuit dedication still exists because the group travels together over one or more circuits that are fully devoted to the purpose.

¹¹For a description of the use of gateways to connect over 20 different networks in the DARPA Internet system, see A. Sheltzner, et al., "Connecting Different Types of Networks with Gateways," *Data Communications*, August 1982, pp. 111-122.

Circuit *set-up time* is measured in seconds for even modern telephony. For calls (typically voice) that are minutes, tens of minutes, or hours long, a set-up time of seconds is inconsequential. But communications among computers, especially between a terminal and its host, are very "bursty." They may only consist of a few hundred bits, transmitted at a rate of a few kilobits per second or higher; typically such bursts are separated by relatively long intervals of time. Computer devices, especially at the terminal level, cannot tolerate long set-up times compared to the length of the messages they deal with; the circuit must be "put up" and "taken down" repeatedly in fractions of a millisecond.

Unlike circuit switching, *packet switching* assumes that all traffic will be in bursts of standard size, commonly a number of bits equal to some power of 2 such as 1024. Each packet is properly labeled with its destination, and all switches along the way simply pass it along. There no longer is an end-to-end "copper wire"; instead, the connection is said to be a *virtual circuit*. To subscribers it appears that there is a continuous end-to-end connection, whereas in fact there is a very elaborate sharing of facilities among all users but it is totally invisible to subscribers. The routing may even vary from packet to packet. Similarly, the connection is said to be a *virtual connection*.

The term *virtual* first arose in the context of mainframe computers which share their resources among a large number of users. It implies that a user perceives that his service--private network, a complete computer system, a subscriber connection--is provided by a specific and dedicated array of hardware and software facilities. In fact, the user is actually being accommodated by an elaborate and complex time-sharing scheme; but the services assigned to one or a group of users are inviolable by other users, as long as there is no hardware failure or anomalous software behavior. The virtual concept is really a very elegant and technically sophisticated extrapolation of what was at first called a time-sharing computer system.

The LAN combines the concepts of packet switching, virtual circuits and connections, and the availability of inexpensive bandwidth to achieve its functionality, its flexibility for providing connectivity,

and its fast rate of response. Since traffic on it is packetized, switching among subscribers is relegated to the address headers that accompany each packet. Thus, there is literally no set-up time, although the extra information in the header (e.g., address of recipient) has the effect of set-up time; and in some LANs, there may be a waiting period to access the network, which is also analogous to set-up. In a LAN, no subscriber ever has the network facilities dedicated 100 percent to his needs.

Since data-exchange buses within a computer configuration function in much the same sharing way, even though they may not be packetized, the LAN actually evolved as an extension of the internal computer bus rather than as an evolutionary step from traditional telephony. In a sense, a LAN can be viewed as a specialized form of a remotely accessed computer network, but one that exploits packet switching and bandwidth.

Traditional telephony and circuit-switched environments use terms such as *2-wire*, *4-wire*, *half-duplex* and *full-duplex*. They relate to the ability to communicate in both directions simultaneously. Clearly, if there are two 2-wire circuits between subscribers, each has a clear channel to the other; this is full-duplex. Conversely, in a half-duplex circuit, one end communicates at a time, using the concepts of push-to-talk and voice-operated switches, either of which turns the communication on the single 2-wire link to the opposite direction. In typical telephony, elaborate echo-cancelling line terminators and so-called hybrid coils permit voice conversations over a single 2-wire twisted pair; for digital streams, more elaborate arrangements have to be used, such as tones of different frequencies for the two directions.

None of these concepts are pertinent to LANs. On a baseband LAN or in any packet network, each packet contains its own address header; and since packets are interspersed in time and are therefore noninterfering, any subscriber can freely communicate with any other one and in either direction. In a broadband LAN, different frequencies or sometimes different cables are used for the two directions to provide the effect of full-duplex. Additional bandwidth is used to eliminate any concern for half- or full-duplexed connections.

A final observation with regard to maximum throughput of a packet network or baseband LAN: The message being transmitted is only part of a packet; the rest is header information for routing plus protocol information to facilitate the end-to-end communication. There is, so to speak, an overhead in terms of additional bits added to every packet. Thus, the effective message throughput in bits per second is somewhat less than the apparent digital bandwidth of the channel on the transmission medium. A corresponding observation holds for those broadband LANs that use packetizing in any frequency band; otherwise, the maximum digital rate in a given frequency channel depends only on its bandwidth.

NETWORK TOPOLOGIES

The topology of a network is the pattern of interconnections among its nodes. Several broad classes of network topology are illustrated in Fig. 2. The mesh, bus, ring, tree, and star topologies have all been used successfully; but the study of network topology and the allocation of channel capacity to maximize network performance is a complex subject. Topology is especially relevant to considerations of network performance under attack, when some links and nodes may be lost.

One way of evaluating susceptibility to degradation of a network that has been damaged is to use the concept of a topological cutset. In the airbase context, this involves asking the question, How many nodes and/or links must be removed from the network before any point is isolated (a) from another point or (b) from the base telecommunications center? The bus topology is the most vulnerable in this sense; any cut of the single bus partitions the network in two. A star system topology is slightly more robust; a single cut will isolate only one node. Loss of a single outer node affects only the users it services, if its failure does not otherwise affect the transmission properties of the remaining part of the network. However, the loss of the central node in the star halts all communication activity of the entire network and all useful activity stops.




Fig. 2—Network topologies (note the number of links or nodes which must be removed to partition the network into disjoint sets)

A tree-structured network is very vulnerable to the loss of a node or link; any single outage will partition the network, and each subsequent outage partitions it again.

A ring network with repeater nodes that circulate message traffic in both directions will sustain its capability if the ring is severed only once or if only one node is removed. Loss of multiple links or nodes will partition the network.

More general distributed architectures, such as that used in the ARPANET, can be made to have an arbitrarily large cutset by increasing the connectivity among the nodes. The large number of nodes and links in the network provide alternate transmission paths among users, but this technique cannot be applied to ring, star, or bus topologies. To do so would turn them into general networks and negate the concept of a simple LAN.

The "cuts" used to determine the cutsets in the discussion thus far are considered to be removals of nodes or perfect breaks. However, in a real wartime environment, electrical situations other than perfect opens may occur. The inherent hardness and failure modes for network cables have been examined for telephone cable plants and possibly for other circumstances, but such insights must be related to LAN systems. It is important to examine the dimensions of the problem; for example, there is a strong possibility of an electrical short circuit or a low-impedance fault when using either plastic dielectric or aluminum-shielded coaxial cables; also cables can be squashed by blast. Since the cable dielectric is plastic, exposure to fires and conventional weapon shocks might flatten or deform the cable and produce a complete or partial short-circuit or a serious transmission discontinuity (e.g., reflected energy). For some LAN technologies, such an event could shut down an entire network or a portion of it. For example, a ring might be completely or partially paralyzed, depending on the distribution of repeaters around it. Twisted pairs and multi-wire cables have similar damage mechanisms. While ruggedized fiber optics are not subject to short-circuits, they can be severed if not adequately protected physically. Clearly, blast damage effects must be examined when choosing a network topology and its transmission medium.

There is one damage characteristic of twisted-pair versus coaxial cable that might be very important in combat conditions. Splicing a coaxial cable is nothing, compared to the problem of splicing a severed 100-pair telephone cable. It can be done in minutes, rather than the hours required for conventional cables or even a day if personnel must be encased in chemical warfare suits. Moreover, there is no need to match proper wires together; color blindness is not a problem when fixing coaxial cable.

NETWORK CONTROL

In addition to differences of topology, networks can vary in means of control. In some networks any node may initiate transmission to one or more of its neighbors, which pass messages to their final destination nodes without a central control; but this can result in message collisions on the network or node saturation for some traffic patterns. Such congestion could become a problem in all but the most lightly loaded networks, so some form of access control is generally used. For example, a control token can be passed which is, in effect, a special control message that can never be mistaken for real traffic; it is analogous to passing a ball around a circle of people. Any node that has something to transmit must wait until the token comes; it then transmits and passes the token along. Token-passing control structures are most commonly used in ring networks, and since there is only one token in circulation at a time, collisions are avoided.

Another technique is called register insertion. In effect the ring is broken when a message is to be sent; a shift register containing the message is inserted into the ring and the message is stepped out of it. Since every node must be prepared to send, each has a shift register which can also receive the message. Switching such registers in or out adds extra complexity to the concept of a simple cable and might affect reliability adversely.

In most commercial networks, some form of access control is used on the bus. All nodes listen before transmitting, and if the channel is in use, all but the using node must refrain from transmitting until the bus becomes available. In a heavily loaded bus network, there can be many

collisions as nodes attempt to insert messages onto the bus; the new message attempts to override the in-transit message which might not have yet come along when the transmitting node first listened and decided to transmit. One way to avoid this problem is to assign a specific and different time-to-transmit (a time slot) to each node. A ring network avoids this difficulty, since a natural sequence of transmission is established as the token is passed. Decentralized unslotted networks require another approach if a low limit of network performance is to be avoided; the traffic density must be kept low enough so that collisions are rare.

The solution in many systems is to listen while transmitting and thus raise the effective utilization of the channel by early detection of garbled messages. By monitoring what is actually entered onto the cable relative to what the node wishes to transmit, a message transmission can be stopped as soon as the collision occurs, rather than waiting for complete error checking on the entire message. This is called *contention control*, and such systems are known as CSMA/CD (carrier sense multiple access with collision detection).

Unfortunately, this approach cannot work over a radio channel, because the local transmitter will overload the receiver as it listens for other signals. A different approach such as time slotting is required for control of a network using a radio channel. If there are many quiet or low-utilization nodes, use of a channel will be inefficient, but some systems are modified to allow a node with large transmission requirements to reserve oversized slots from the network controller in advance.

RELIABILITY CONSIDERATIONS IN NETWORK CONTROL

Loss of both nodes and links has been discussed previously in terms of topology. Another measure of reliability in a network is the ability to accommodate transient errors and noise. For example, if the token is lost in a ring, there must be a means for reinserting it at some restart node; some centralized oversight and control is indicated. In a contention-controlled network, noise has the same effect as a collision--retransmission of a message is required and automatically takes place. This is an attractive aspect of contention-based designs, since the

normal traffic management strategy also recovers from transmission errors. More sophisticated techniques, including well-known error-control methods, can be used to assure reliable detection of all garbled messages and adequate recovery from errors. For example, in ARPANET and MILNET, each transmitted packet contains a parity check which is checked for correctness by the receiving node. It signals the originating node for a retransmission if necessary.

In a conventional threat environment, a major technical aspect of network reliability is the repeaters which can segment the bus into electrically isolatable sections and thus localize the consequences of shorts that may occur in the cable. Commercial bus networks do not normally include segmentation features, but militarized versions may very well require automatic electrical partitioning of the bus to offset damage.

Network designers have sought to improve the reliability of repeaters by powering them from the network itself, leaving the repeaters independent of the data equipment (and its power) to which they are attached. The failure of a repeater or terminal device on a bus network could be disastrous if it were to impose a low impedance or short on the cable. In contrast, a high-impedance failure mode can leave the operation of the remainder of the network relatively intact, since it does not short out the remaining communications. Commercial products used in military combat environments must be examined for such failure modes.

Fiber-optic technology offers some reliability advantages over twisted-pair or coaxial cable; for example, it allows longer runs with high noise immunity and excellent electromagnetic pulse and ground-current isolation. This characteristic can provide a lower error rate, resulting in either better network performance or the ability to use less elaborate equipment at the nodes. Existing optical systems are intended primarily for high-speed, high-volume, computer-to-computer data transfers; commercial LANs based on fiber optics are becoming available.

OTHER LAN ADVANTAGES

There are various other advantages of LAN technology. No one of them is overwhelming, but collectively they can be important:

1. A simpler logistics stockage situation, since there is only one cable--coaxial, fiber optic--to stock and keep records on instead of many sizes of twisted-pair cables.
2. Smaller physical volume, which alleviates duct crowding or space problems.
3. Ease of connection--to splice or repair LAN cables, only two coaxial fittings need to be installed instead of tens or hundreds of soldered, wire-wrapped, or punched connections.
4. Significantly more attractive damage recovery characteristics.

RADIO BACKUP

Radio packet switching¹² can be used as a bridge or gateway between spatially separated LANs, as discussed later. Radio is also attractive as a backup channel for networks in a hostile environment. It theoretically eliminates the problems of physical severing or electrical shorting of a cable, although sufficient bandwidth cannot be made available in peacetime to replace coaxial cable or substitute for the on-base cable plant. In combat theaters, jamming will further reduce data rates during hostilities.

When a base is trying to recover from combat damage, runway and aircraft repair is likely to have higher priority than restringing cables. Thus, radio backup links may well provide better communications support to reconstitution crews and may provide interim but limited support for other communication needs as well. For operational and reliability reasons, selected systems might fall back on a radio transmission mode of operation, even if at reduced capacity. In this case, the LAN for peacetime use must be chosen to be technically compatible with a radio system. Other options are clearly also possible

¹²Robert E. Kahn, et al., "Advances in Packet Radio Technology," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, pp. 1468-1496.

(e.g., alternate or spare cables, field-strung fiber-optic cable, point-to-point microwave or other nonwire links).

A radio backup has other implications for the initial system design and must be an integral part of the system architecture rather than an add-on afterthought to a standard commercial system. For example, the communication protocols must be chosen properly. Since jamming in a combat environment is projected to be a serious threat, the system must accommodate a potentially extraordinarily high error rate.

Most packet radio techniques, e.g., the DARPA-sponsored Aloha¹³ packet radio system (a forerunner of the Ethernet), require a relatively noise-free radio channel to permit collision detection. With severe jamming or the use of spread spectrum modulation techniques, the process of detecting collisions and seizing the unused channel may be more problematic than in other channel management schemes, such as time-division multiplexing, which does not require collision detection.

A truly adaptive communications arrangement would be able to adjust its data rate in response to jamming, but it would require hardware that can operate at continually varying data rates, perhaps as low as only a few bits per second. Commercial systems have no incentive to incorporate such a feature, but a microelectronic chip under development for the DoD VHSIC program is intended for exactly this type of data communications and could be applicable in a LAN/radio context.

Use of radio channels may also require encryption of information, since the physical security of the link cannot be maintained, as can that of an on-base cable link.

USING RADIO TO MINIMIZE PEACETIME DELIVERY TIMES

A useful peacetime application of radio equipment that could also function at some or all nodes for wartime backup communications would be that of advanced portable message devices. The nature of Air Force operations often requires personnel--particularly flight-line and maintenance personnel--to be away from fixed locations. A radio link

¹³N. Abramson, "The Aloha System--Another Alternative for Computer Communications," *AFIPS Conference Proceedings*, Vol. 37, FJCC, 1970, pp. 695-702. See also "The Aloha System," N. Abramson and F. Kuo (eds.), in *Computer Communications Networks*, Prentice-Hall, Inc., New York, 1973.

interconnected to a LAN could activate a portable pocket pager-like device containing a microcomputer that drives a multi-character display or even a small printer. It could also allow long messages to be stored and scanned, thus offering a useful new service. Messages could automatically be forwarded to the portable device and the individual alerted upon receipt, irrespective of location. Public key cryptography or other techniques might be used if it is necessary to maintain privacy in these transmissions.

COMMERCIAL LAN PRODUCTS

For peacetime data communications applications in the CONUS, commercial LAN products are an appropriate technology to consider. Ethernet has emerged as a de facto standard for baseband LANs and is compatible with equipment offered by a large and growing number of manufacturers. Broadband networks, such as MITRENET or Wangnet, have a much larger capacity than baseband LANs and therefore can carry video and voice traffic as well as digital transmissions. Use of broadband network technology is appropriate when wideband channels are needed for high-speed data and video. Since there will be a substantial on-base population of voice-only users who will continue to use the base telephone system, there seems to be little reason to create a second voice network on a broadband LAN. However, there is an in-between option, namely, integrated voice/data switches for network access and internetting; they are discussed in Sec. VI.

Among the local networks in use today, the Ethernet, Cluster Bus, MITRENET, Net/One, Wangnet, and Z-net are all bus systems using some form of coaxial cable (see Table 1) and thus are topologically and electrically vulnerable to conventional attack. ARPANET-like networks such as Telenet have, to date, been limited to larger expensive nodes whose cost would be prohibitive if used in current configurations on Air Force bases. It is possible, however, to envision a smaller, microprocessor-based node supporting a highly connected distributed network that could maintain network connectivity even when severely damaged, but without the high cost associated with the present minicomputer-based distributed networks. Engineering development and possibly some R&D are necessary, but the idea is technically feasible.

III. NETWORK PROTOCOLS

INTRODUCTION

The diverse requirements and organizational structures of the various tenants on an airbase, particularly in the CONUS, are likely to result in the development of many individual internal communications systems by various communities of interest. For example, several word processors purchased for local use can evolve almost unnoticed into a small LAN. Such systems will in all likelihood meet the immediate requirements of their users, but unless a standard interface both to other groups and to the telecommunications center is provided, it may be difficult to accommodate such standalone systems into an integrated base communications system.

Protocols are used in a network to allow users and their application programs to communicate with one another and with host computers without user intervention into communication details.¹⁴ In effect, a protocol is an automated set of rules embedded (typically) in software that specifies all the details of data exchange between two intercommunicating digital devices. A protocol must function invisibly to the user and must conceal from him all the details of effecting a flow of digital information among components of a network. In computer jargon, the operation of protocols must be transparent to the user, no matter how many protocols there may be and no matter how they are nested within one another. A protocol must deal with such things as differing data rates in two networks, error control between components of a LAN, overloading of the receiving network or component, requests for special services by the network, and fragmenting or reconstructing messages.

Protocols are basically intended either for computer-computer interactions or for computer-human interactions. Many structural details can properly be concealed in a computer-computer relationship,

¹⁴ For a comprehensive discussion of protocols, see D. W. Davies, et al., *Computer Networks and Their Protocols*, John Wiley & Sons, New York, 1983; and "Special Issue on Computer Network Architectures and Protocols," *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980.

providing substantially more convenience to the end user. High-level protocols link the users' local capabilities (e.g., electronic mail, text editors, and message boxes) with the interface to the network. Lower-level protocols are those used by the system to actually pass the users' messages over the physical datalinks (e.g., cables or radio channels) to the receiving node. The receiving end, using the same datalink protocols, assembles the message for the higher-level protocols to pass it to the destination computer program, information process, or physical device.

Interconnection of multiple networks is also dependent on the incorporation of proper protocols. The DoD has standardized a set of protocols for this application; they are known as the Internet Protocol (IP)¹⁵ and the Transmission Control Protocol (TCP).¹⁶ Use of TCP/IP facilitates connection of networks with differing internal protocols. The IP is a lower-level protocol that provides *datagram* service, which is simply a message consisting of a source and destination address plus the actual message itself.¹⁷ Importantly, though, the message is delivered without error checking or acknowledgment. The appeal of a datagram is that it is so simple that it can be readily implemented in almost any network. The IP also reassembles long messages which span a number of datagrams.

The TCP, the next level of protocol, builds on the IP to provide reliable message delivery by incorporating error checking, positive acknowledgment, and flow control. This two-layer system enables reliable connections to be established between communicating computer processes. The internetting TCP need only be in the gateway, not in individual host computers and terminals; but the gateway also manages flow control and handles speed matching (i.e., differing data rates on the two sides), leaving only addressing to the individual nodes.

¹⁵"Internet Protocol. Request for Comments 791 (MILSTD 1777)," in *Internet Protocol Transition Workbook*, ARPA Internet Network Information Center (NIC), SRI International, 1981.

¹⁶"Transmission Control Protocol. Request for Comments 793 (MILSTD 1778)," in *Internet Protocol Transition Workbook*, *ibid.*

¹⁷Harold C. Folts, "X.25 Transaction-Oriented Features--Datagram and Fast Select," *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980, pp. 496-500.

In practice, protocols are defined in more than the two conceptual layers just described. The seven-layer Open System Interconnect model of communication protocols as defined by the International Standards Organization¹⁸ is a hierarchical layering of protocols that permits many functional capabilities to be implemented on communicating networks whose structures were not originally intended to support them. However, each layer of the protocol must be chosen with regard to the transmission media with which it works. For example, the Ethernet set of protocols would not be useful on a noisy radio channel for reasons discussed previously, but a hybrid protocol set with an Ethernet "front end" for the user applications and a TDM datalink control with error correction for the communications would allow the same application package to operate over a system with radio channels. Such details are invisible to the user but are essential initial design decisions.

Network protocols are not yet standardized; LAN vendors tend to choose their own. The DoD has standardized on TCP/IP, but the Open System Interconnect model is driving commercial developments. To some extent, military and civil government are on divergent paths, but these are in process of being reconciled.^{19,20}

PROTOCOL CHOICE

The following discussion suggests a number of operational features that can be provided in a military communications network through proper choice of protocols.

¹⁸Hubert Zimmermann, "OSI Reference Model--The ISO Model of Architecture for Open Systems Interconnection," *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980, pp. 425-432.

¹⁹*Transport Protocols for Department of Defense Data Networks*, Report to the Department of Defense and the National Bureau of Standards by the Committee on Computer-Computer Communication Protocols, Board on Telecommunications and Computer Applications, National Research Council, National Academy Press, Washington, D.C., February 1985.

²⁰Harvey J. Hindin, "Local-Net Standardization Gains," *Electronics*, March 24, 1983, pp. 98-99.

Speed Matching

The availability of inexpensive communications bandwidth in a LAN reduces the need for "bit conservation" and thus eliminates the necessity to implement complex flow control procedures. In contrast, long-haul networks, especially those including satellites with unavoidable long transmission delays and high data rates, require complex network control algorithms and substantial buffering to handle error control procedures. Local area networks have a much simpler task of speed matching between senders and recipients, and of providing buffer space to accommodate the delays between receipt of a message and its acknowledgment. This issue is addressed later as an aspect of internetting.

Urgent Messages

There is a requirement in military systems for an interrupt provision in protocols, a feature not normally provided in commercial offerings. This provision allows urgent messages to receive priority. One experimental method has been to signal the receiving node with an "urgent pointer," which indicates the presence of an urgent message in the data stream. An interrupt capability is particularly important for a time-shared host which may be busy with other tasks when an urgent message arrives to be sent or handled.

Locator Service

A militarily useful, high-level protocol service provided by some LAN designs is that of an address information service to facilitate message exchange. A local network interface would not be required to know the actual address of every user serviced by the network, but rather could send the message "Where is X?" and receive the address from any node that had the information. Rapid reconfiguration in wartime Air Force applications is thus possible, allowing arriving units or new nodes to be integrated readily into the base information system.

Broadcast Messages

A broadcast capability enables a node to send a message to all other nodes without having to address each one individually. This is a desirable LAN feature to announce new node interconnections and to serve as a building block for implementing higher-level protocols such as conferencing or automatic document distribution.

Message Acknowledgments

Protocols must also provide acknowledgment, which is required at many levels in a network. For some types of traffic, such as requesting an address, the response itself serves as an acknowledgment. For other types of traffic, such as digitized speech or other sampled continuous information, the acknowledgment of a message received without error would be less important than the timeliness of the message. Such details must be reflected in the protocol, but it must be noted that acknowledgment can occur at different levels. For example, receipt of an error-free message by the receiving node is automatically acknowledged by the network to the sending node, but this is invisible to the user. A "registered-mail return-receipt-requested" user-level acknowledgment would probably be implemented outside the protocol structure.

IV. INTERCONNECTING NETWORKS

A LAN offers the basic advantage of being able to connect, via one gateway, to a number of external long-haul networks, such as MILNET (and eventually the DDN), SACDIN, or other command-and-control networks. Rather than implementing and maintaining a large number of interfaces for each on-base user to each external network that he needs to access, a LAN bundles users and funnels them through a single gateway that implements the TCP/IP internet protocols of the external networks. The fundamental consideration in efficiently internetting LANs to long-haul networks or to other LANs is that of protocol compatibility. Long-haul networks may not now provide all the protocol features available on LANs. Rather, they rely primarily on the virtual-circuit concept, which means that the end-to-end circuit appears to exist to the user whether it is actually a continuous circuit or some other arrangement, such as a packet network. For the latter, all of the details of passing the traffic successively through the network nodes are completely invisible.

The format in LAN addressing will ideally provide the ability to reach destinations outside the immediate confines of the system. This is analogous to the role that an area code plays in the direct-dial telephone network. Clearly, any LAN must be able to address directly all of its subscribers; but many user actions with the LAN can be avoided if its implementation also allows him to directly address destinations outside the local user constituency. Its addressing structure should accommodate the "area code" of other LANs or of off-base networks. Including such a requirement initially will reduce the proliferation of protocols, with their associated software development and maintenance consequences.²¹

²¹For a discussion of the X.25 Standards Committee's approach to internetwork protocols, see Antony Rybczynski, "X.25 Interface and End-to-End Virtual Circuit Service Characteristics," *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980, pp. 500-510. The X.25 is an international standard protocol agreed upon by the Consultative Commission on International Telephone Transmission (CCITT) in 1976 for packet-switching networks. It has since been extended and

Another task relevant to internetting a LAN to long-haul networks is that of matching data transmission rates. Local area networks commonly have much higher data rates than most long-haul networks. For example, the fastest MILNET/ARPANET links are 50 kilobits/sec, while LANs commonly operate at 5 to 10 megabits/sec. While speed-matching to LANs is less of a problem than speed-matching to high-speed, high-delay satellite links, the internal data flow features of LANs require flow control and buffering between them and long-haul networks.

If one adapts the virtual-circuit approach for internetting, a LAN's own virtual-circuit protocol could ideally be constructed from a compatible subset of the long-haul network virtual-circuit protocol. This would avoid the necessity of implementing and maintaining two separate sets of protocols. If the external networks also support broadcast or message exchange services similar to those of the LAN, such commonalities can also be exploited. As a practical matter, the variability of internal LAN details across available commercial products very likely implies that protocol conversion will be an aspect of the LAN/long-haul interface.

SUBNETWORKS AND BRIDGES

A LAN depends on low loading of its bus for rapid acceptable access and traffic delivery, and it exploits cheap bandwidth. On the other hand, *subnetworks*²² are a particular kind of internetworking that mixes LAN technology with that of conventional transmission media, with compatible message sizes and protocols throughout. Individual subnetworks are linked by devices known as *bridges*; the overall concept is shown in Fig. 3. The function of the bridge is to pass appropriate

is now undergoing further revision and extension. CCITT is the international forum in which telephone communication practices and technical details have been resolved, and its scope has now been broadened to include communication aspects of data networks.

²²Since "subnetworking" is a technique for orderly growth, it might properly be called "supernetworking." However, the common term is "subnetworks," presumably because the ensemble, no matter how large it becomes, is considered a single LAN rather than an aggregation of many LANs. Moreover, since the initial technical design, especially of protocols and addressing, has provided for orderly interconnection of LANs, the notion of regarding the ensemble as "the LAN" is consistent.




Fig. 3—The subnetwork concept

Here, a LAN is composed of a number of subnetworks, linked in some fashion by bridges. The subnetworks, though of differing technologies, share one address space, and the same protocols are used over the entire network. Thus, the bridges can be simpler than the gateway which connects the LAN to the long-haul network. Viewed externally, from outside the dashed line in the figure, the LAN appears to be monolithic.

SOURCE: David D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, p. 1514.

messages and to provide buffering plus speed-matching features between LAN portions. A bridge can be viewed therefore as a *smart repeater*.

There are a variety of reasons why one might want to accommodate multiple technologies and data rates. In many instances, it is desirable to use existing twisted pairs if they can handle traffic requirements, rather than running a new cable system or microwave links, either of which implies construction costs.²³ Subnetting accommodates different data rates between LANs serving communities with different internal needs; e.g., one which supports video teleconferencing or high-data-rate exchanges could still exchange electronic text messages bridged to a low-speed LAN that did not include video.

²³ James Rush, "Microwave Links Add Flexibility to Local Networks," *Electronics*, January 13, 1982, pp. 164-167.

A subnetwork approach permits orderly growth. As a particular LAN grows, unless the transmission media can be upgraded to a higher speed, performance for all users must inevitably degrade. In a subnetted configuration, an additional LAN can be added or broken off from a large one, rather than overloading existing ones or upgrading communications channels. Not all commercial products allow such actions.

The common address space shared throughout a subnetwork design allows the entire collection of nodes, irrespective of subnetwork membership, to appear as a single entity to all users. Both addressing and bridges are transparent and invisible to users, just as call routing and switching in the telephone company are invisible to the individual customer. Figure 4 shows the conceptual structure of a bridge between subnetworks.

The major distinction between subnetting and internetting is that in the former, common protocols are shared by design, thereby avoiding embedding one protocol within another, as is generally required in internetting. The subnetted architecture is an a priori integrated




Fig. 4—The structure of a bridge

A bridge would most naturally be located at a point where the two subnetworks it interconnects have been made physically adjacent.

SOURCE: David D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, p. 1515.

design, as opposed to an after-the-fact interconnection. Since a protocol is in effect a necessary "overhead charge" to a message, there is economic advantage to a common set; embedding protocols within one another leads to ever-longer message lengths and therefore increasingly burdensome overhead. It is also possible to extend the subnetwork approach to longer distances by using a long-distance bridge, as shown in Fig. 5. However, if the communications link between the two halves of the long-distance bridge does not support efficient use of LAN protocols, user convenience will be reduced and response for LAN users whose traffic must cross the bridge will be degraded.

INTERNETWORK ADDRESSING

While subnetworking provides a cohesive addressing structure to all of its component parts, a different issue arises in the interconnecting of all the LANs that might exist on an airbase or between two LANs at two different Air Force installations worldwide through one or more long-haul networks. A sender must have a mechanism for addressing any subscriber on any LAN throughout the Air Force; he must have a means for fitting his message into the addressing structure of the distant LAN and be assured that the intervening long-haul networks can correctly deliver the traffic to the intended recipient. It is the analog of the telephone long-distance area-code problem, and it needs the same solution. Some central focal point must assign a unique address identifier to each LAN and/or network that is expected to interconnect.

ARPANET and its successors, and the networks with which they interconnect, have already dealt with the problem. Each network and every host on each network has a unique identifier. The central point of authority is the Defense Communications Agency which operates ARPANET, MILNET, and DDN. In principle, the AFCC, with AFSI concurrence and cooperation, could handle the problem for the Air Force, but in the long run a DoD-wide solution must be put into place.

There are implications for protocols as well. The internetwork gateways must be able to handle the overall addressing scheme and be able to route messages properly onto receiving LANS and networks, in addition to handling all the other details of interconnection, such as differing data rates and error control.




Fig. 5—The long-distance bridge

In this case, the two subnetworks cannot be made physically adjacent, so a half-bridge is attached to each, and a full-duplex communication link is employed to interconnect the two half-bridges. The control and filter functions, and the packet buffers, are replicated in each half-bridge.

SOURCE: David D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, p. 1516.

V. VIDEO AND VOICE APPLICATIONS ON LOCAL AREA NETWORKS

VIDEO

The bandwidth requirement for a video signal is much larger than that for all but the fastest data communication applications. Baseband LANs, such as Ethernet, cannot accommodate it; but broadband LANs, such as MITRENET, can readily include a number of video channels along with variable-rate data traffic.

On-base video is presently used primarily for surveillance and physical security applications, but high-speed data users will be involved in other base activities. Thus, they are likely to cross functional-area lines, whereas the security applications will not. Since data and video users may well be in separate communities, there is little reason to join them simply to justify a broadband network technology. Among other things, the use of such a network for video transmission raises its cost over that of conventional CCTV technology, because special interfaces to the network must be provided for every camera and every monitor. For most video purposes, such as surveillance, a conventional CCTV approach is more appropriate and probably more cost-effective as well.

Another issue might be the classification of different kinds of traffic. A teleconferencing scheme might involve both video and computer-generated data, in itself a functional reason to share a common network; moreover, such conferences might be classified. Even then, however, the video portion could still be a CCTV network gatewayed to whatever digital networks might have to be accessed. Issues for consideration include classification and security details, cost of additional cables versus special video interfaces to a broadband LAN, and functional interaction among classes of service.

There could even be limited, e.g., basewide, distribution of off-base video services such as videotex or commercial television.

VOICE

Local area networks can also carry digitized voice, so the question arises of whether there is any reason to expect that such networks might carry any of the voice traffic that exists on Air Force bases today and will remain for the foreseeable future.

A number of experiments have been conducted using Ethernet and other CSMA/CD networks to carry combined voice and data loads.²⁴ Analysis of their performance has included a straightforward bandwidth computation, as well as a detailed simulation that allowed for modification of the network control algorithms to facilitate simultaneous voice and data use. A straightforward calculation indicates that for a 10-megabits/sec network carrying the standard 64-kilobits/sec 8-bit PCM voice with the protocol overhead as defined in the DEC-Intel-Xerox joint specification for Ethernet, the total data rate required will be 107.2 kilobits/sec per voice channel, which divided into the network capacity of 10 megabits/sec yields a theoretical upper bound of 93 simultaneous conversations. Thus, 30 simultaneous conversations would require 32 percent of the bandwidth, 40 would require 43 percent, and 50 would require 54 percent. Such simple calculations lead one to question the feasibility of incorporating significant levels of digital voice traffic in a CSMA/CD system, since the drain on system bandwidth is substantial even for as few as 30 to 50 simultaneous conversations.²⁵

More detailed analysis (based on different methods of contention control) does not substantially change the conclusion that a baseband network, especially a contention design, is one of the least cost-effective ways of providing voice communications to Air Force base-level users.

²⁴G. J. Nutt and D. L. Bayer, "Performance of CSMA/CD Networks Under Combined Voice and Data Loads," *IEEE Transactions on Communications*, Vol. COM-30, No. 1, January 1982, pp. 6ff.

²⁵For an elaboration on the issues regarding the use of LANs for different categories of services see C. A. Niznik, "Cost-Benefit Analysis for Local Integrated Facsimile/Data/Voice Packet Communication Networks," *IEEE Transactions on Communications*, Vol. COM-30, No. 1, January 1982, pp. 19ff.

However, a CSMA/CD network implementation based on higher-bandwidth media would be able to carry more conversations; for example, a broadband coaxial or fiber-optic network can provide a 100- to 500-MHz digital path, with room for 10 to 50 times the voice traffic levels achievable on Ethernet. Advanced schemes such as linear predictive coding can also substantially reduce the bandwidth requirement for digital voice transmission at the expense of more complex equipment for each subscriber connected to the LAN. For example, if the digitized voice could be fitted into a 2400- or 4800-bps channel such as is typical for many terminal applications, the number of simultaneous conversations would be 20 to 25 times higher; and the voice loading would be correspondingly smaller. (See Sec. VIII, pp. 84-89, for further elaboration of this possibility and its costs. See also Sec. IX, pp. 114-121, for a quite different approach.)

This argument, however, could be different for a LAN using a different access control mechanism, especially one that uses a TDM scheme. For such a LAN, each voice channel would consume only 64 kilobits/sec of capacity, rather than the 107 kilobits/sec of a contention LAN.

The requirement for complex equipment to handle digital telephone service plus a total replacement of the existing analog voice plant combine to create a situation in which substantial costs would be incurred, possibly without a corresponding return in performance, reliability, or increased capability.

VI. INTEGRATED VOICE/DATA SWITCHES

One approach to networks intended to handle a large volume of both voice and data traffic is *integrated voice/data switching*.²⁶ This approach has evolved from modern telephone switching technology, unlike LANs, which are extensions of internal computer buses. In an integrated voice/data switch,²⁷ connections between terminal equipment and the switch are similar to those between ordinary telephones and their switch, but the details of the switch are based on modern digital technology. Since terminals are connected directly to the switch, using ordinary individual wire pairs, modems are not required at terminals or at computers directly attached to the switch. Typical voice/data switches can accommodate digital traffic up to 19.2 kilobits/sec on ordinary twisted-wire pairs over reasonable but limited distances, with extension to 56 kilobits/sec projected to be available soon. This would be adequate for any terminal application short of interactive graphics and for all but the most data-intensive computer-to-computer interactions. For high-speed high-volume data traffic, separate dedicated links undoubtedly will continue to be both attractive and cost-effective.

Commercial vendors offering integrated voice/data switching systems include AT&T, Intecom, Mitel, Rolm, Northern Telecom, General Telephone and Electronics, NEC, Harris, American Telecom, and Datapoint. The

²⁶Holger Opderbeck, "The PBX: Equal Rights for Voice and Data," *COMPUTERWORLD*, February 25, 1985, pp. ID/9-14.

²⁷The phrase "integrated voice/data" implies that the device is indifferent to the information content of the digital stream that it processes. An integrated telephone switch handles data streams from either voice or data sources equally well, although technical details in an integrated switch differ from those of a voice-only switch. For example, a data connection will typically be maintained much longer than a voice one; the switch must not overload under such circumstances. Alternatively, the voice and data connections might be handled in individual parts of an overall switch. Other things might be integrated as well; for example, some communications arrangements might multiplex digitized voice and data together for transmission but separate them at the switch for individual routing.

switches can incorporate full-function voice capabilities including both AUTOVON voice-priority standards and the interface to existing AUTOVON switches. The data features that are a part of modern voice-switch architectures provide the basis for a flexible yet evolutionary and expandable base communications system offering both traditional telephone services and digital data service.

RESOURCE SHARING

For communicating with computing equipment that is remote enough from the switch to require long-haul military or commercial circuits, a bank of modems can be colocated at the physical location of the switch and shared among outgoing (or incoming) data traffic. This advantage of voice/data switching is analogous to the LAN subnetting feature mentioned above with regard to standardizing internetwork protocols and procedures in the bridge and all subnet components. The voice/data switch itself is in effect a gateway to other communication networks, both intrabase and interbase. Thus, one community of interest that might require an Ethernet for access to its own word-processing hardware or dedicated computer would still have a means for connecting with any other on-base terminal, LAN, or computers via the voice/data switch. The LAN-switch interface is becoming a standard item in commercial LAN product lines.

Certain common computing resources can also be attached to the voice/data switch to provide voice or data message services to users. In addition to electronic mail,²⁸ actual digital storage of voice messages to be retrieved by phone ("voice mail") is quite feasible;²⁹ commercial products for this purpose are already available.

²⁸J. J. Garcia-Luna-Aceves and F. F. Kuo, "A Hierarchical Architecture for Computer-Based Message Systems," *IEEE Transactions on Communications*, Vol. COM-30, No. 1, January 1982, pp. 37-45.

²⁹S. Hattori, et al., "A Design Model for Real-Time Voice Storage System," *IEEE Transactions on Communications*, Vol. COM-30, No. 1, January 1982, pp. 53-57.

SUBMULTIPLEXING

Submultiplexing is a desirable feature in a voice/data switch. It exploits the fact that the digitized speech formats used in voice/digital switches require substantial data rates--between 64 and 144 kilobits/sec. Some switches allow a channel to be used for only one application, either voice or data, but submultiplexing allows a single voice channel to be subdivided into a number of lower-speed data channels. For example, a single 144-kps voice channel could be divided into 15 9600-bps data channels. This feature will become increasingly important as the volume of digital traffic increases, since it allows continuing growth in digital demand without requiring the addition of new lines to the switch. Some device, of course, will be needed to blend together the traffic from a cluster of terminals and to deliver it to a common destination (e.g., a computer or a LAN). If the combined terminal traffic must be delivered to a variety of destinations, then the situation is more complex; the switch must be able to inspect each submultiplexed channel and route it individually.

COMPATIBILITY WITH EXISTING CABLE PLANTS

One of the major advantages of an integrated switch is its ability to use the existing cable plant for connection of both voice and data terminal equipment, provided the cable itself meets the electrical requirements of the switch. Growth can be achieved by addition of cable pairs and through a bank of shared-access modems or other specialized interface equipment that can be located conveniently on a base and linked to computers attached to the switch.

Many commercial integrated voice/data switches have been designed for the military environment. Among their features, they include options for access to AUTOVON lines and the ability to distribute sections of the switch physically about the base. For example, the individual parts of a switch can serve individual communities of interest but be linked together through tie lines in a multiply connected topology. The latter feature can allow continued but limited operation even if one or more subswitches are destroyed. In

effect, contemporary switch technology can bring onto a base a multiplicity of small exchanges linked (in traditional telephone parlance) by interoffice trunking. It can go the traditional local telephone company one better by providing backup alternate routing-- a feature appearing only in the commercial direct-distance-dialing network. Figure 6 illustrates such an architecture.




Fig. 6—Distributed switch architecture for a military environment

VII. HYBRID ARCHITECTURES

INTRODUCTION

Interconnection possibilities among voice and data networks are virtually unlimited. Local area networks (either broadband or baseband) can be coupled to other LANs through direct bridges, through an integrated voice/data switch, or through an external communications network. Voice/data switches can be used to link terminals to LANs, to other terminals, to standalone computers, or to external long-haul networks.

Major issues for all interconnection arrangements will be protocol details and speed-matching. In a multiple-protocol environment, necessary conversions from one protocol to another can become a sizable processing burden to the system and may require design of specialized interface equipments. This point clearly argues for protocol standardization, especially in the context of subnetting. High-speed protocol converters, which place much of the burden on hardware, are now coming onto the market; they will facilitate the interconnection of networks using different protocols.³⁰ Speed-matching, a central issue long faced in the interconnection of all computer components, must be accommodated in any hybrid network. A high-speed source cannot transfer data to a lower-speed one without adequate buffering and flow control procedures, but this is a well-established technique and in principle presents no serious technical obstacle.

Figure 7 suggests a hybrid-based architecture using a voice/data switch serving most users, but with two LANs, perhaps corresponding to individual computing communities of interest, internettted through the switch. On one of the LANs a high-speed CPU is linked to another through a fiber-optic data channel. The example illustrates a number of properties of hybrid architectures. First, computer-to-computer communications at very high speeds (100 megabits/sec and up) can be realized without sizing the entire network correspondingly. Second,

³⁰Paul E. Green, Jr., "An Introduction to Network Architectures and Protocols," *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980, pp. 413-424.




Fig. 7—Two LANs internetted through a voice/data switch

The voice/data switch provides the duplex communications link shown in Fig. 5. If the LANs use different protocols, the half-bridges perform the conversion.

any user connected to the switch or to either of the LANs can access all resources for which he is authorized and can send messages to any other user through the message facilities associated with the switch. Third, overall reliability is improved, since many system components can fail without affecting the others. For example, one LAN can be down while the other continues normal operation; or both LANs can be down, but the switch will still be available. A switch failure does not affect the operation of the terminals on either of the LANs. Finally, failures of the high-speed data link between the two large CPUs affect only the activities immediately concerned with their mutual data stream.

HYBRID ARCHITECTURES IN THE AIRBASE ENVIRONMENT³¹

Figure 8 shows the simplest and most obvious concept for Air Force base communication modernization. A number of user groups (tenants) have been linked with a bus-oriented LAN. For a high traffic rate, a broadband network would be required; for a less communications-intensive set of tenants, a baseband LAN would be appropriate. The new base telephone switch (as procured under SCOPE DIAL or SCOPE EXCHANGE) is projected to be a modern but voice-only switch, replacing the aging electromechanical switch.

The advantage of this approach is its simplicity. There is a system for voice and a system for data, and no effort is required to deal with the complexities of internetting, subnetting or remote access. A significant disadvantage is that long and costly LAN cable runs may be required to link all data communicators on one network, even though many of them will be in different tenant organizations with minimum need to communicate among communities. Connectivity from the base computer center to off-base networks could be electrical or air-gap (e.g., hand-transferred magnetic tapes).

However, even though communications traffic between different tenant organizations may be small, those organizations will be competing for network capacity, since they are on the same single cable. In a

³¹Stuart Wecker, "Combining Voice and Data Through Local Nets and PBX," *Computerworld Focus*, Vol. 19, No. 41A, October 16, 1985, pp. 44-49.




Fig. 8—Basewide LAN and voice telephone switch

baseband network, this problem is more severe than in a broadband design, since the bandwidth pie to be divided is smaller to start with. Moreover, any network will be degraded in performance for all subscribers as the total load approaches the design limits or if one subscriber burdens the system with a very high traffic load.

Figure 9 shows an alternative and much more advantageous approach. It includes separate LANs for each tenant organization or community of interest, but all are internettted through a voice/data switch. With this architecture, nodes that do not have frequent communications do not compete with each other for network resources, a salient and certainly important aspect of system-level philosophy.

The base telephone switch has also been enlarged with a voice/digital computer-based switch--conceivably a SCOPE DIAL or SCOPE EXCHANGE switch upgraded to handle data traffic--that also allows remote terminal access to any of the LANs via twisted pairs, rather than requiring long coaxial cable runs for one or two users. Since the wire pairs are generally already in place, and distances on airbases can be quite large as far as LAN technology is concerned, this approach has an inherently important advantage, especially when new nodes must be added, as is frequently the case. Off-base data flow would be an electrical connection through an appropriate network gateway. Some data might even flow off-base through the telephone switch onto voice-grade circuits.

The separation of one all-encompassing LAN, such as that in Fig. 8, into a number of LANs for various communities of interest also provides an additional measure of fault isolation, since combat damage or a failure that brings down one network (e.g., by shorting the cable) is isolated to the one network. Clearly, a hybrid architecture, such as that in Fig. 9, has superior growth and flexibility advantages:

- A community of interest can be added, deleted, enlarged or shrunk at will.
- New tenants or tenants in new locations enter the system gracefully and easily.




Fig. 9—Hybrid LAN-voice/data switch architecture

- As much redundancy as is warranted by damage threats can be incorporated, either systemwide or for selected communities.
- Interfaces to other capabilities, such as remote paging or a radio system, can readily be established through the switch-- and traffic can readily be rerouted or exchanged as circumstances warrant.
- New technology (e.g., fiber optics) can be dropped in when ready, or as desired.
- Individual LANs can be of different rates or different vendors as long as certain interface and protocol matters are standardized.

There are other advantages as well:

- Costs are lower, since remote terminal access does not require new coaxial cable.
- With fewer users on each network, there is less competition for resources.
- The in-place cable plant can be utilized.
- The architecture provides some measure of survivability.
- The new switch technology that is available today and being installed under the SCOPE program is put to good use.
- Tenant organizations are allowed greater autonomy in the design and operation of their LANs, since they do not affect others external to their group.

While Fig. 9 shows only one "base computer center" and one "base telecommunications center," their functional capabilities could of course be replicated to enhance survivability or to provide additional capacity.

Part 2

SECURITY ASPECTS

VIII. SECURITY ASPECTS IN THE LAN ENVIRONMENT

INTRODUCTION

The LAN is a technological alternative to twisted-pair cables switched through a central, or sometimes distributed, switch for providing connectivity among computer terminals, computer hosts, file servers, printers, storage devices, and possibly other equipment. Depending upon the detailed application, the conventional telephone approach and the LAN each has its own technical and economic advantages. Although LANs are primarily associated with computer-oriented applications, they are found in most office automation systems as well. In the latter, the LAN is commonly thought of in a text-processing context, whereas in the former, it is usually regarded in a data-handling context.

Clearly, airbases are apt to have either type of installation, but the general case will involve both dimensions of information manipulation, probably through the same LAN. As vendors of office automation (OA) equipment add more general computing capability to their product lines and as ADP systems incorporate general-usage text-handling features, the distinction between "OA systems" and "ADP (or computing) systems" will blur. In a sense, that distinction concerns the kind, quantity, and diversity of equipment connected to the LAN more than the nature of the traffic on it.

Traffic on a baseband LAN is digitally represented, is packet-switched, originates on a terminal of some kind, and is delivered to a host or processor of some kind. Thus, the security issue for such LANs is largely independent of functionality in the system, although the details may vary from the OA system to the more general computing system. In the latter case, the security concerns extend beyond the LAN itself and, in particular, involve security safeguards within the hosts on the system. The same may also be true for the OA situation, although recognition of the host aspect has been slower to appear in that community.

For the broadband LAN, the situation is somewhat different. Frequency bands in this case may be assigned to nondigital services such as video or audio. In principle, the technical arrangements of a baseband LAN (e.g., time sharing of the channel, some access control mechanism) could be embedded in some frequency band on a broadband LAN, but the available bandwidth is so large that usually each subscriber is assigned to a frequency.

Thus there are potential cross-talk problems that cannot exist in a baseband LAN. In the latter, a major concern is contamination of one packet by another or contamination of a message by the misrouting of packets.

At the most general level, the security issues are the same for any LAN (e.g., missent messages), but at the detailed level, the solutions will be quite different. For example, a baseband LAN requires an encryptor for a packet environment; a broadband LAN can use an encryptor for a circuit-switched continuous traffic flow. Audio or video signals will have to be digitized for encryption, with a consequent requirement for much more bandwidth.

Perhaps because of historical timing, the vendors of computing equipment are more aware of computer security concerns than are vendors of office automation products. It is important to note that for Air Force applications, LAN security must be viewed as combining traditional communication security (COMSEC) issues with the newer computer-system security (COMPUSEC) issues, especially as the latter introduce requirements for safeguards that center around not only the terminal and hosts but also their mutual interactions.

The discussion that follows is an overview of the security aspects of LANs. It is not intended to be a manual of implementation details; rather, it suggests what might be done now for the near term (two to three years) and what new approaches will be possible five or so years from now. It will also relate the topics of COMSEC and COMPUSEC as they arise in the LAN environment.

COMMUNICATIONS SECURITY

Communications security has a long history and tradition of use by military services and governments to protect state secrets and defense information. It is generally appreciated in military circles, especially in their communication components, and it need not be extensively discussed here except to note that the classical doctrine, techniques, and specialized equipment used by the COMSEC community were developed for the point-to-point circuit-switched application, which is an entirely different operational environment from that of the LAN. Over the years a broad variety of equipments have appeared, but in general the communication safeguard is provided by an encryption component that is packaged separately from the device that originates or receives the traffic. In recent years, specialized encryption devices have appeared for application to packet-switched networks such as the Defense Data Network (DDN), notably the *bypass encryptor*, which passes the packet header around the encryption process but the body of the message through it.

LAN Differences

While the general doctrine and guidance of traditional COMSEC might prove pertinent to the LAN situation, there are important technical issues that distinguish the two cases. For example, the traditional key generator (or encryptor) could, in principle, be employed to encrypt traffic on a baseband LAN, but its usage will be inefficient, since it is intended for continuous traffic, rather than the start-stop intermittent traffic of the LAN environment. Moreover, many key generators continue to send output even when there is no input in order to conceal traffic flow characteristics. More recent units such as the KG-84A which are start-stop on the plaintext side still transmit continuously on the ciphertext side and also lack the bypass feature.³² Since baseband LANs control traffic entry onto the LAN by a scheme such

³²*Telecommunication Protection Systems Using the KG-84A*, Bendix Communications Division, Baltimore, Maryland. A successor equipment, the KG-84C, will provide a rudimentary bypass feature. See *KG-84C General Purpose Telegraphy Equipment*, National Security Agency, Attn: S711, Ft. George G. Meade, Maryland.

as detecting collisions between packets or assigning a particular time slot to a terminal, conventional key generators would excessively load the LAN and even make its proper operation impossible.

For broadband LANs that can afford to dedicate one channel to each terminal or other device, existing key generators could be used even though efficiency would be reduced (because of the intermittent nature of the traffic), and costs would probably be impractically high. Conventional encryption mechanisms are simply not a good match for many LAN environments, even though they can be employed on twisted-pair circuits through a telephone switch. A new generation of equipment that includes bypass features will be essential for proper treatment of communications security within LANs. Ideally, they will be fully integrated into terminals and hosts, rather than standalone add-ons.

COMPUTER SECURITY¹³

In contrast to the decades of use of electrical and electronic encryption devices, the tens of decades of mechanical devices, and the centuries over which various forms of communications security have been practiced, computer security is just over two decades old. It received limited attention in the 1960s, and somewhat fuller R&D attention in the 1970s; but the software aspect of computer security has been institutionalized only in recent years. Operational products with extensive software safeguards are just beginning to appear.

Software-intensive systems which incorporate safeguards to assure that system resources (e.g., computing capability, stored data) are accessed and used only by properly authorized individuals have come to be called *trusted systems*. The implication of *trust* or *trustedness* is that a system incorporates a set of security safeguards (hardware and/or software) that have been (a) designed by carefully specified methods, and (b) tested thoroughly by carefully specified means and certified by an appropriate technically qualified organization, with the result that

¹³For a technical treatment of security in contemporary computer networks, see D. W. Davies and W. L. Price, *Security for Computer Networks*, John Wiley & Sons, New York, 1984. For a different treatment, see Robert W. Shirey, *Local Area Network Security*, The MITRE Corporation, McLean, Virginia, MITRE Working Paper WP-82W00587, October 1982.

the system operator can be assured with high confidence that the system enforces appropriate security controls, including the desired rules of access to the system and to its information and computational power and storage resources. The totality of access rules enforced by a system has come to be called the system's *security policy*. While the concept of trusted systems originated in terms of computer systems, clearly trust must also be an attribute of computer-centered communication systems, either long-haul or LAN.

For example, a communications controller, which in a LAN context would be a BIU, must be trusted to (among other things) never change the security markings on a message, never disturb the association between a message (or message fragment such as a packet) and its security parameters, and guarantee that its software will be isolated from the message traffic so that the latter cannot cause anomalous behavior or security infractions in the controller.

The computer security issue was first treated comprehensively in a 1979 Defense Science Board report.³⁴ Its dimensions, as identified in the DSB report but rephrased in contemporary terminology, are:

- Physical security--locks, fire protection, entry control.
- Personnel security--trustworthiness and organizational allegiance.
- Emanation security--TEMPEST control of compromising emanations.
- Hardware safeguards--to assure hardware integrity; to guard against anomalous or unexpected modes of behavior; to assure proper execution of the software; and to enforce the system security policy in conjunction with the system software.
- Software safeguards--to enforce the system security policy, notably the rules of access to system privileges, resources, and data; to control access of individuals to the system; to provide oversight features such as audit trail information for monitoring proper performance of the system; and to assure

³⁴Ware, Willis H. (ed.), *Security Controls for Computer Systems*, Report of Defense Science Board Task Force on Computer Security, February 1970. Reissued by The Rand Corporation as *Security Controls for Computer Systems*, R-609-1, October 1979.

software integrity to avoid such software subversion as trapdoors or software bombs.

- Object control security--to control the flow of tapes, discs, paper, etc., in conformance with the access rules in the system security policy.
- Administrative security--to assure appropriate procedures for controlling certain security aspects of the system, e.g., assignment of passwords, validation of users to the system, entering classification and access privileges of users into the system, taking appropriate actions in the event of system compromise or attempted penetration; to provide appropriate system personnel (such as a system security officer) to enforce and apply the security rules and procedures.
- Management overlay--to provide overall management policy and guidance for the system; to provide general oversight of security affairs; to be ultimately responsible for the securityworthiness of the system.

Together, these dimensions enforce a security environment that determines which authorized individuals may access what classified information and for what purposes (e.g, read it, change it, create it, classify it). Each kind of protective measure is one aspect of an overall protective shield for a computer-containing system, either an ADP or a communication system. The threats against which each of them protects can be readily identified. Figure 10, reproduced from the DSB report, illustrates many of them in the context of a network configuration as it was conceived in 1970; some terminology has been changed to contemporary usage, and a few new vulnerabilities have been added.

In addition to the threats implied by the nature of the various security dimensions, there are two other threats that are not identified or treated in the Defense Science Board report. One is the so-called *HUMINT threat*, meaning penetrations or subversions that attempt to exploit the personnel in the system. From one point of view, the magnitude of the HUMINT threat is roughly proportional to the number of users who access the system--the more there are, the more likely that one or more may be subverted.




Fig. 10—Computer network security vulnerabilities

SOURCE: Willis H. Ware (ed.), *Security Controls for Computer Systems*, Report of Defense Science Board Task Force on Computer Security, February 1970. Reissued by The Rand Corporation as *Security Controls for Computer Systems*, R-609-1, October 1979.

The second is a *denial-of-service threat*, in which the penetrator attempts to capture the system, to flood it with excess traffic, or in some other way degrade or totally deny the support the user expects the system to provide. In a sophisticated form of the threat, such debilitating actions might occur only at crucial times, as far as the system user is concerned.

The severity of denial-of-service threats varies with the criticality of the mission supported by the system. For example, nuclear command-and-control systems have a very demanding requirement to be operational under all circumstances and at all times. Thus, the security safeguards must prevent an adversary from denying service to rightful users of the system, with extremely high confidence. Support systems such as logistics supply and maintenance might be seen as relatively less critical, especially in peacetime. In war, however, these systems directly affect the capability of a base to perform its mission, e.g., mount sorties. Thus, either security safeguards that are adequate for the wartime situation must be in place at all times, or they must be promptly implementable when needed. The former is obviously preferable because, as the aphorism indicates, one should practice in peacetime with the systems with which he intends to fight the war. The many aspects of security safeguards must collectively provide protection against both HUMINT and denial-of-service threats.

LAN SECURITY³⁵

Figure 11 puts the schema of Fig. 10 in the context of a modern LAN environment. All of the security concerns associated with a network that provides primarily ADP services are clearly also present in the LAN circumstance, but there are important differences of detail, especially in software aspects. For example, many of the elements in a LAN provide a well-defined but limited function, e.g., the message and protocol

³⁵For a fuller treatment of the subject that includes a bibliography of relevant documents, some discussion of LAN technology, and general guidelines for addressing LAN security with today's array of equipments and concepts, see *LAN Security Guide*, Office of Secure Data Networks, NSA/V52, National Security Agency, Ft. George G. Meade, Maryland.




Fig. 11—Local area/office automation network security vulnerabilities

details of bridging between one LAN and another, or interfacing a terminal to the transmission medium. Thus, undetected design flaws, especially those that react to particular combinations of message traffic and/or other operational details, might be of greater importance in the sense that unauthorized leakage of information could occur in larger volumes before detection of the trouble, or terminal traffic might be sporadically misrouted.

On the other hand, the operational software within various LAN components is not accessible to the traffic on the LAN. While the traffic must provide the software with certain parameters for proper handling (e.g., routing information, classification level), the traffic is isolated in the sense that it cannot change or cause anomalous behavior of the LAN software. The latter is generally not vulnerable to attacks that could be mounted through programming activities taking place on a terminal, even though such activities might in fact be an attack against a LAN-connected host that could be vulnerable to unauthorized manipulation through software actions. Thus, an operational capability that might be a normal part of LAN-delivered services (on-line programming, in this instance) and is a very important aspect of security concerns for a computer system has minimal consequence for security in a LAN installation per se, provided the software in the LAN components has the proper attributes of trust.

To put it differently, aspects of the threat which might be important for a computer host need not necessarily be important for the LAN itself, and vice versa.

Representative Security Concerns

In a broad way, the security concerns in a LAN environment center around such events as:

- Accidental misdelivery of messages so that classification or need-to-know infractions occur; or covert intentional misdelivery on a continuing basis of selected kinds/origins/destinations of messages.

- Access by unauthorized parties to classified information stored in the system for either short or long times; or similarly, access by authorized users to information for which they are not authorized, combined perhaps with covert routing of copies to other parties.
- Access by authorized or unauthorized users to use the network in an unauthorized way.
- Denial of service.

More specifically, threats can include tampering with the hardware and/or software in order to:

- Gain access to cryptologic materials.
- Circumvent cryptologic protection and thus acquire plaintext material.
- Record--perhaps in a covert fashion--classified information for later extraction from the system or for retransmission to unauthorized recipients.
- Circumvent trusted software by hardware modifications.
- Invade trusted software to modify or disable its protective features.
- Gain some advantage to support other kinds of penetration.

Another form of threat is spoofing, which can include:

- Insertion of false messages from seemingly authorized originators.
- Covertly modifying system security controls to permit other kinds of penetrations.
- Disabling cryptographic protection by sending seemingly authentic instructions to do so.
- Subverting remote keying features that might be present in the system.

Many of these concerns are of course equally pertinent to long-haul networks, both circuit-switched and packet-switched. Moreover, at this level of discussion, concerns are the same for baseband and broadband LANs.

Another risk is the failure of software protections due either to bugs that were not discovered during verification of system component trust or to anomalies of behavior induced by hardware malfunctions--either transient or permanent. If such events can be deliberately induced, then the risk, especially for denial of service, becomes a threat.

LAN Office Automation

If a LAN supports only a message function such as electronic mail, most information in the system will be in transit from one place to another. Typically, there will be short-term storage (e.g., a few days) until a message is delivered and read by its recipient; but there may also be long-term archival storage (e.g., many months) for auditing, overall system management, and so forth. The security concerns tend to be those traditionally associated with communication networks.

However, users of electronic mail tend to retain their messages for extended periods of time, especially if the messages have been organized into computer-based folders or files, as an electronic filing system. In cases where messages are stored for very long periods within the computer that hosts the mail service, security concerns transcend the predominantly COMSEC issues associated with the LAN and include COMPUSEC aspects in the host.

If the LAN supports a broad-service office automation environment, there is likely to be very-long-term storage of information (e.g., many months to years). In this case, there are many more options for misusing the system than simply eavesdropping on traffic or monitoring who-is-talking-to-whom. The security concerns now unavoidably become those of a computer-based system that provides general services to a user community. The scope of COMPUSEC attention must be much broader than that for a simple host supporting a mail service.

In the extreme, if the LAN connects many computer terminals to one or more hosts that provide the user population not only with a full range of office automation functionality but also with generalized programming and other ADP services, then the security concerns include the fullest sweep of LAN-oriented COMSEC concerns and both LAN- and host-oriented COMPUSEC ones.

Traffic Flow Security

Traffic flow security (TFS) relates to full protection of all aspects of traffic on a network so that not even message count, message frequency, message length, message format, addressees, daily periods of circuit usage, etc., are revealed. Cipher text is continuously sent over a TFS-protected circuit by the encryption device, even when message traffic is absent. Clearly, this is easy to arrange in a circuit-switched environment, since the point-to-point nature of a circuit makes clear who the recipient at the end of the link is.

In a packet-switched network, the routing information must be in the clear at the node switches in order for them to function properly. In a long-haul network such as MILNET, link encryption can be employed on the node-to-node circuits to provide traffic flow security, but within the switch the header must be in the clear. However, the body of the message can be maintained under encryption for protection.

The limited physical extent of a LAN and even its physical installation features make it unreasonable to require a corresponding arrangement. In principle, a baseband LAN could have a link encryptor between the bus and the BIU, as well as a bypass encryptor in the terminal or traffic-originating equipment; but the BIU also needs direct access to the bus in order to utilize the access control mechanism for putting packets onto the LAN. Some specialized technical arrangements are necessary.

For broadband LANs which dedicate a frequency channel to a subscriber and do not require access control mechanisms to the transmission medium, point-to-point link encryption could provide traffic flow security.

A major issue is the cost-benefit tradeoff: the cost of the link encryptors vis-a-vis the value of TFS on a LAN. There is also the technical problem of integrating encryption devices into the LAN BIUs.

Traffic flow security on a LAN is not now being advocated by COMSEC authorities, although it is on long-haul networks, where there is a high density of traffic and the exposure to interception is very much higher. The ultimate decision on LAN TFS will lie with the responsible end-using organization of the system, but it is clear, in the near term, that a broadband LAN is preferred if TFS is deemed mandatory.

LAN Physical Security

Unlike a traditional computing installation in which the equipment is centralized in one or a few locations which can be properly protected physically, the equipment in a LAN environment is much more likely to be distributed over a broader area. Thus, threats of physical damage or access to terminals by unauthorized users or of subverting terminals in unprotected areas will be more significant. Consideration must be given to physical security for the cable, amplifiers, headends, BIUs, and other LAN components. Likewise, the physical areas where terminals, hosts, servers, and other LAN subscriber devices are located must be physically protected to the same security level. Much like the weak link in a chain, the physical security of a LAN system is only as good as the physical security of its weakest component.

Many of the precepts that have been developed for the operation of secure computer-center facilities are potentially useful for the physical and operational security aspects of LANs. Handbooks and other materials exist in both the commercial and the military world; needless to say, such guidance should be considered in decisions concerning LAN installations.³⁶

³⁶For example, *Counterintelligence Support to Automatic Data Processing Security*, U.S. Army Intelligence Center and School, Fort Huachuca, Arizona, Field Circular 34-112, December 1984.

LAN Component Software

A LAN can involve a variety of electronic equipment, e.g., bus interface units for terminals, specialized couplers to other equipment, gateways to other LANs or to a switched on-base or long-haul network, line amplifiers. Unavoidably, some of the components will be software controlled, and it goes without saying that each such device of a LAN configuration that is expected to enforce security access controls must contain trusted software, since it might be a potential point of penetration. As suggested earlier, a significant aspect of the system-level and software design is determining just what aspects of trust must exist in each network component.

LAN Security Policy

As security safeguards are added to LAN environments and, for that matter, to all base-level communications, new policy issues are sure to arise. Some may require resolution through appropriate authorities, but others may be within the purview of the Air Force or the local responsible organization. Examples include the following:

- Is traffic flow security an issue?¹⁷
- Can unclassified and classified traffic be mixed on the same LAN? Only if all subscribers are cleared to the maximum classification of the traffic? Only if trusted software is utilized? What about the possibility that some--or even many--users of a LAN may not be cleared?
- What about on-base digital traffic that is routed through the on-base telephone switch? Must it be encrypted if classified? Must the software in the switch be trusted if the traffic is not encrypted? (In this regard, can the SCOPE DIAL and SCOPE EXCHANGE switches handle encrypted traffic? Are there COMSEC compatibility issues to be addressed?)

¹⁷See p. 67 for a fuller discussion of this matter.

- In times of stress when normally unclassified traffic becomes wholly or partly classified, must the telephone switch be operated in a system-high mode?³⁸ Must LANs, some or all of them on base, also switch to the same system-high mode?
- What procedures are necessary to give high confidence that terminals in a normally unclassified office environment that becomes classified have not been tampered with--the TEMPEST protection destroyed or a bug installed?
- When some base information processes become classified during crises, what arrangements must be made in transiting to off-base long-haul networks? Especially in the coming era of direct electrical connections through the telecommunications center, what policy and technical changes will be implied during classified operation?
- Given that procedures and practices may differ between unclassified and classified environments, what policies does a base need to assure that it can smoothly operate in the classified situation?

TERMINAL SECURITY

While not strictly a communications-related security matter, LAN-connected terminals do need to be discussed because they are the items most commonly connected to a LAN, either for office automation applications or for more general-purpose ADP applications.

Simple Terminals

Frequently, terminals connected to a LAN will be so-called "dumb terminals." Such devices provide a display and a keyboard capability, but no significant storage beyond that required to maintain the display. They will have no processing capability beyond that required to generate electrical signals from keystrokes, to create character displays from

³⁸"System high" is an appellation for a computer or communications system for which all subscribers are cleared to the highest level of information that might be within the system, although not all necessarily have equal needs-to-know.

electrical signals, and to manage the exchange of signals with the LAN, and through it, with the host. Specifically, they have no local diskette or hard disc storage. They may have a local printer to provide hardcopy output.

Security concerns for such installations are relatively straightforward and do not involve software issues at the terminal. There will be, of course, major security-related software issues in the host computers. The terminal must be properly protected physically for the level of classification at which it is to work, and access to its keyboard or visual observation of its display by unauthorized individuals must be prevented. Hardcopy output must be handled as any classified material would be.

If the terminal is TEMPESTed, precautions must be taken to assure that maintenance actions do not subvert TEMPEST safeguards, e.g., emplacement of a transmitting bug or bypassing line filters. Furthermore, maintenance personnel must be cleared to the appropriate level or carefully monitored while servicing a terminal to assure that inadvertent access to classified information does not occur, and to assure that electronic modifications to the terminal are not made, especially ones that might compromise security. For example, a terminal can be microprocessor-controlled even though local processing capability is not provided. The read-only memory might be replaced with a doctored one that would, unknown to the user, route all terminal traffic to a second destination at which uncleared personnel could pick it off the LAN, or it could set up some other security circumvention.

Microcomputer Terminals

Increasingly, though, LAN-connected terminals are "smart terminals". Commonly, such terminals are microcomputers--also called personal computers--of the many kinds that are available. Most of them have diskette storage; they may have hard disc storage; there is likely to be a local printer; there may be other local devices such as a mouse or a light pen. Such computers are often acquired individually and subsequently networked. The security issues for microcomputers are much more profound and much more difficult to deal with, at least for the near term.

The essence of security is that safeguards are inviolable; they must be protected against tampering. This observation is true whether the safeguard is a lock, an item of hardware, or a software protection. It is in the last instance that micros raise the most awkward security issues.

In many ways, the microcomputer of today replicates the historical situation of the 1950 era of computing. The user has the functional run of the machine, as he did with the earliest computers. There is no aspect of the machine's software that is not available to a user; there are no hardware features that deny him access to any resident software. In contemporary parlance, the typical microcomputer is a single-state machine; it does not have the "user state" and the "privileged state" commonly present in large mainframes. The privileged state is available only to the machine itself or to specially trusted system programmers; in particular, it is not accessible by any user. Thus, all the resource assignment controls and all the security and access controls run only in the privileged state and are proof against user tampering. It was the technical innovation of the 2-state machine that made the contemporary time-sharing computer system possible.

Hence, it is a vacuous exercise to insert software security safeguards in microcomputer operating systems or in other software packages resident with the user. It is a simple matter for him to copy or edit a diskette and thus to modify, disable, or circumvent security controls, or to change security-controlling parameters such as user privileges or clearance levels. Moreover, it is a simple matter for him to store classified material locally and to do with it as he wishes.

Until 2-state microcomputers exist and until the Air Force acquires them, trusted security controlling software for microcomputers is a meaningless concept. If networked, the thousands of microcomputers the Air Force has already in place represent a very awkward security situation that will be hard to deal with in other than a system-high context.

The typical microcomputer also is intentionally designed with an open architecture so that the user can readily adapt the machine's functionality to his own desires. Such a feature is an obvious

marketing advantage, but it also means that hardware cards can easily be inserted into a microcomputer with unknown, possibly unknowable, security consequences.

This means that a network of microcomputers must operate at a uniform security level; multilevel security is technically not possible with most present models.³⁹ It also means that users with smart terminals have a ready-made environment in which to construct software attacks against hosts, against communication systems, against file servers, against whatever is connected to the LAN, or against whatever can be reached through gatewayed connections to other LANs or to long-haul networks.

The issue of security control in an environment of microcomputer terminals is not yet thoroughly understood or even thought through. In the near term, and possibly into the mid term, the presence of locally available computing power to a user who is in a networked environment can only be a circumstance of extreme concern to security authorities.

End-to-End Encryption

End-to-end encryption, or "E-cubed," as it is sometimes called, simply means that a message is under full encryption from the time it leaves the originator until it reaches the recipient. Ideally, it would be encrypted as it left, say, the keyboard and would be decrypted just prior to, say, being displayed or being retrieved from a host memory for processing. However many communication networks a message might pass through en route to its destination--whether the destination is a computer or another terminal--encryption would have to be continuous. There is an implicit assumption, of course, that both the originator and the recipient handle the item with proper security attention. In the networked environment, it must be assumed that the originator or the recipient, or both, can be a person or a host computer.

³⁹Even a standalone microcomputer can be a security risk, primarily because users do not realize that security precautions are necessary, nor do they appreciate the ease with which information can be pilfered. See, for example, *Security of Personal Computer Systems: A Management Guide*, Institute for Computer Sciences and Technology, National Bureau of Standards, Gaithersburg, Maryland, NBS Special Publication 500-120, January 1985. Related companion documents are in process.

If such a technique were available, and if key distribution and key management for the encryptors were in place, some of the security concerns about LANs would vanish. It would no longer be necessary to worry about protecting the transmission media, except against a denial-of-service threat. Communities of interest, perhaps operating at different security levels, could be separated by the use of different keys. Access to particular hosts could be controlled by proper allocation of encryption keys.

There are many significant technical issues. Encryption must be an integral part of the terminal. A microcomputer terminal with an encryption capability must be proof against tampering by the terminal user. There must be an efficient means for disseminating keys to the parties who wish to exchange information. Encryption devices must also exist at the host, and there will have to be many of them if several keys are in use. There is a compatibility issue across networks, if the protection is really to be end-to-end. The encryptors must match the transmission aspects of the LAN--a packet encryptor for a baseband LAN, but possibly a continuous encryptor for a broadband one.

The issue is a complex one, but there are NSA programs under way to address it, including the STU-II/STU-III programs, the Commercial COMSEC Endorsement Program, and Trusted Systems (discussed later in this section).⁴⁰ However, until programs such as these or others come to fruition, end-to-end encryption will be limited to those very special circumstances for which the protection requirements warrant it. Even for such limited applications, the detailed implementation is not likely to be either easy or economically attractive.

BASE-LEVEL LANS

At base level, LAN technology might be applied to such functional areas as command-control and other mission applications, or to logistics and other support systems. It is likely also to be used in the office automation environment throughout the base to minimize the flow of paper and to make office procedures more efficient. In each instance, the

⁴⁰See pp. 82-90.

overall aspects of security will be the same, but the relative importance of various dimensions of security will vary, as will details.

While much normal peacetime traffic on an airbase is formally unclassified, monitoring large amounts of it can reveal such things as command-control procedures, logistics details such as inventory stocks and rates of supply, support and maintenance arrangements, organizational structures, and even relationships among individuals. Moreover, there is a new category of national defense-related information which, though unclassified, must still be protected against unauthorized access.

The threat against information in industrial installations that use LAN technology has not yet been perceived as significant, and these installations have in general not been concerned with security threats. Therefore, commercial manufacturers generally have not addressed the security concerns of the defense world, and comprehensive security safeguards do not yet exist in commercial products.

NEAR-TERM SECURITY POSSIBILITIES

Even though the full gamut of security safeguards is not yet available in LAN technology, much can be done to provide limited or even extensive protection. For example,

- Equipment should be TEMPEST-approved or located physically in conformance with accepted TEMPEST practice.
- Cable and other transmission media can be installed in such a way that they conform with doctrine for protected-wireline-distribution systems and installed in such a way that they are visually observable.
- Equipment can be located in protected areas, or at least in areas that can be protected when necessary.
- Procedures can be developed to give high confidence that only authorized personnel have access to the terminals on a LAN and to its equipment.
- Terminals can be located only in protected areas, and continuous surveillance arranged for them; personnel can be made aware of security aspects of LAN installations.

- Commercial maintenance personnel can be prevented from having unmonitored access to system components.

Aside from security in the formal sense, precautions must be taken to protect sensitive but unclassified information and to guard against acts of fraud, embezzlement, or theft committed with the aid of base-level information systems. For example, financial information is often private in nature, but exploitable for personal gain; it must be properly safeguarded. Fortunately, office automation applications require information access control, so access control and audit-trail software is available from some commercial vendors.

- Commercially available software that offers access control or other protective safeguards in LAN applications should be used to the maximum extent possible.

There exists a list of TEMPEST-approved equipment.⁴¹ While all the necessary COMSEC and trusted equipments do not yet exist for LANs, many security safeguards can nonetheless be instituted simply by using reasonable caution in choice, installation, and operation of equipment, together with monitoring and training of personnel. Figure 12 summarizes the LAN security concerns and includes, among other things, guidance to the documents for implementing what might now be done. It does not, of course, address COMPUSEC concerns in hosts.

For some situations, the LAN and all of its terminals and hosts can be installed securely, i.e., the LAN transmission medium by protected-wireline-distribution regulations and the equipment in physically secured spaces. In effect, the system in this case will run in a uniform "system-high" classification status and hence all of its users must be correspondingly cleared. It serves a community of interest at a uniform level of classification and need-to-know.⁴²

⁴¹*Preferred Products List*, Subcommittee on Compromising Emanations, Serial SCOCE-0058-83, April 1983 (or most recent issue).

⁴²Such a system has been installed in the Office of the Secretary of Defense.

	Risks	Techniques of Exploiting Risk	Effect/Result of Exploitation	Preventive Actions/Methods	Implementation Guidance	Probability of Loss with Security
Selection/ Installation	Physical	- Unauthorized physical access - Tampering - Bugging - Monitoring - Substitution of components	- Unauthorized access/modification of classified information - Denial of service	- Inspections - Access controls - No alone zones/two-man control - Intrusion detection systems - Vaulting - Passwords/encryption	DODD 5200.28 DODD 5200.28-M NSAM 90-4 NACSI 4009 NACSEM 5203	Negligible
	TEMPEST	Collection of compromising emanations	Unauthorized release of classified information	- Protected wireline distribution systems - Link encryption - Shielding - Filtering	DODD S-5200.19 NACSIM 5100A NACSEM 5112, 5200 5203 AFNAG 9A, 58 NACST 4009 National COMSEC issuances	Negligible
	Off-the-shelf Designs	Random design flaws in both hardware and software	- Unauthorized release/modification of classified information - Denial of service	- Good acquisition management practices - Trusted systems - Certified systems	NSAM 81-1,2,3 CSC-STD-001-83	Negligible (when technically possible)
Operations	Software	- Malicious entry of software - Control bypass	- Unauthorized access to classified information - Denial/disruption of service	- System-high access - Periods processing	DODD 5200.28 DODD 5200.28-M NSAM 90-4	Negligible
	HUMINT	- Blackmail - Incentives - Mole	- Release/modification of information - Service denial	- Clearance/polygraph procedures - Background investigations	Contact appropriate counter-intelligence organization	Variable
	Denial of Service	- Destruction - Deletion/delay of messages	- Disruption/denial of service - Selective denial	See physical preventive methods above	See physical implementation guidance above	Variable (varies with mission)
	Spoofing	Inserting and altering of messages	- Modification of information - Service denial	Access controls	See physical implementation guidance above	Negligible

Fig. 12—LAN security aspects and guidance

Another option depends upon what degree of trust can be developed for the interface device that permits equipment to be connected to the transmission medium. If a given physical location is authorized to receive and transmit a given classification of information, then an appropriate interface device can be trusted to properly label all material entering the network and to receive from it only its assigned level of classification. The operational concept adopted within each physical location served by the LAN should depend on personal and individual control and distribution of material from one or more terminals.⁴³

The major advantage of such an approach is that the dimensions of trust, and therefore the difficulty of implementing it in the bus interface unit, are significantly less than the corresponding aspects of trustedness in hosts that might be connected to the LAN. In fact, since interface units tend to include microprocessors to implement their other functions, trustedness might be added in an evolutionary way to installed commercial products.

Clearly, the transmission medium itself must be securely protected, and corresponding doctrine and regulations must be followed. However, it is simply standing good practice to always protect the medium physically to avoid damage and corresponding disruption of service. Thus, it may be that relatively simple trusted BIUs would be able to accommodate security concerns in many base-level installations.⁴⁴

Therefore, LANs--whether they are used to support office automation, computer systems, general data distribution, or some other application--will require a substantial effort to implement comprehensive and relevant security safeguards. There are a variety of threats to be considered, and countering them will be difficult and probably not cheap initially. Protecting against some threats may mean some degradation in performance and service, at least temporarily. In

⁴³Stephen T. Walker, "Local Area Networks: A System View," *SIGNAL*, September 1983, pp. 37-43.

⁴⁴BIUs with trusted software components are under development by The MITRE Corporation.

addition, the problem of interconnecting LANs among themselves and to other possibly multi-tiered networks (the DDN, among others) will demand standardization of interface requirements, security practices, and other details.

Encryption

Unlike a circuit-switched environment in which a circuit serves only one subscriber at each end (or possibly a few through local pony circuits or additional distribution), a LAN does not provide a uniquely identified link between subscribers. On a baseband LAN, the "link" is in essence the dynamically varying time position of packets. On a broadband, it is some assigned frequency for a given communication session, but possibly some different frequency on another occasion. For all LANs, though, the BIUs in effect give every subscriber potential access to all traffic--quite unlike the circuit environment. Anomalous behavior or failure of BIUs can lead to misrouted traffic and, in the case of protected information, to infractions.

Hence, there is an inherent requirement for source-to-destination or end-to-end encryption to provide positive security control, as opposed to depending upon the proper operation of equipment. Operationally, a LAN subscriber will communicate with a variety of others, or with several computer hosts, or with off-base networks. Therefore, convenient and automated arrangements must be made for automatic key distribution and management, and specially designed encryptors will be required. Collectively, the implications of operational convenience represent a complex technical problem that has not been solved and will require novel COMSEC approaches. Among other dimensions of complexity are the lack of standardization in interfaces and protocols and the diversity of equipment available from commercial vendors--both hardware and software.

At the moment, it is awkward to utilize traditional encryption techniques to protect traffic on LANs, partly because of cost, partly because of technical aspects, and partly because appropriate devices are not available for the broad range of commercial LAN products that might be acquired. There can also be interface problems, especially with respect to the interplay between protocol details and encryption devices.

However, new initiatives from the National Security Agency will make significant contributions to this shortfall. Importantly, they include an emphasis on COMSEC compatibility among various networks and equipments.

LAN to Long-Haul Networks

However well security safeguards might be provided in on-base LANs, other problems may arise when traffic passes through a gateway to off-base long-haul circuits. For example, encryption might be used on some LANs but not others; in the short term there could be differing COMSEC arrangements as traffic passes off base. It has already been observed that interoperability among LANs might be impeded or even made impossible because of lack of protocol standardization, but the problem might be even more extensive, depending on how individual LANs incorporate security features into the protocol structure. In fact, lack of protocol standardization may well be the major impediment to internetwork connectivity for the near future.

Message Management and Error Control

In circuit-switched military message environments, the traditional way to deal with major message errors or garbles has been to request a resend. This implies that operators are present at both ends and can have complete manual control of the traffic inserted on the link. In more elaborate cases, forward-error-correcting is used, where the supplementary information sent with the message allows the recipient to reconstruct garbles and omissions. Depending upon the nature of the errors that system designers perceive as operationally intrusive, the amount of extra error control information can become large compared to the message length and can require a significant amount of processing at each end. It is also common in military message systems to provide message receipting--an acknowledgment that a message with a given identifier has been successfully received. Among other things, receipting assures message accountability and it guards against the risk that an unauthorized person has inserted or deleted a message.

NO-A166 948


TECHNOLOGICAL PERSPECTIVES FOR AIR BASE COMMUNICATIONS
(U) RAND CORP SANTA MONICA CA W H WARE OCT 85
RAND/N-1988-AF F49628-86-C-8888


2/2

UNCLASSIFIED

F/G 17/2

NL





MICROCOPY

CHART

In a LAN environment, there are no operators; the user at his terminal, or the computer, or other data equipment insert data directly onto the LAN. Since data do sometimes get lost or damaged, automatic ways must be provided to react to the situation and handle it in such a way that operations are not disrupted.

The same technical problem had to be solved in the design of the original ARPANET, especially since the many packets of a long transmission can travel over different paths, each having different error-inducing characteristics. The typical packet-environment solution is to append a simple parity check to each packet; a receiving node checks the parity, and if it does not match, the system requests a retransmission from the sending node. Thus, there is packet-by-packet short-term receipting, but there is no long-term accountability in the network per se. This must be handled by the subscribers through whatever processes they mutually agree on.

There can be higher levels of error control as well, for example, a check at the sending end that all packets of a message have been received. Needless to say, all such automatic procedures occur invisibly to users; they are all embedded in the software of the hosts serving the communicating parties.

The same set of problems arises in computer-computer and terminal-computer communications, either through the twisted pairs of a base telephone system or over a network. If the former, or if on a broadband LAN, the terminal is in effect in a circuit-switched environment and the details of error control will have been resolved in the terminal-computer design. A parity check at the byte level is typical. On a baseband LAN, there must be additional details that resemble those developed for the ARPANET.

The details of how a message is routed through a LAN and how its end-to-end integrity is protected are embedded in the *protocols* of the network, primarily in the network software. Not all vendors of commercial equipment have included error management in their designs. For baseband networks that use CSMA/CD, error control in the protocols is essential because packets damaged by a collision have to be detected and remedied. For a broadband network in which a communications session

is essentially a circuit-switched frequency assignment, a vendor may have decided that the error performance of his equipment is sufficiently good that any residual problems will have to be dealt with by the end users.

The point of this discussion is to underscore the intricacy of the interaction among protocols, encryption, and standards. The end-to-end encryption process must be fitted into, or onto, the protocols without disturbing their normal functioning; and everything has to be done with sufficient standardization and flexibility to accommodate whatever commercial LAN products might be acquired. And it all has to be done in a way that assures the trusted attributes for security control--a guarantee that the labels and their message association have not been changed.

A LAN, or any other network, can only assure that a message has been successfully delivered without error. Additional controls, such as logging and accountability, will be the responsibility of the subscribers to the network.

NEW SECURITY INITIATIVES

At present efforts are under way to comprehensively examine security aspects of LAN networks, including in some cases cooperation with vendors to address security aspects of particular installations.

There are several major new communications security efforts by the National Security Agency that will make important contributions to the security of on-base data and voice communication traffic and to COMSEC compatibility across networks.⁴⁵ All of these efforts imply intensive interaction with private industry, and to facilitate this an Office of Industrial Relations has been established within the NSA COMSEC organization.

⁴⁵One of these is described in David Burnham, "500,000 More Spy-Proof Phones Proposed by Top Security Agency," *The New York Times*, October 7, 1984 (city and record edition: sec. 1, p. 1, col. 1; national edition: sec. 1, p. 1, col. 2).

STU-II Terminal

The STU-II fully TEMPESTed, secure-voice terminal consists of an oversized deskset (approximately the size of a multiple-key telephone instrument) with an accompanying equipment box weighing about 60 lb. It permits secure voice conversations over the switched telephone network, utilizing a normal voice-grade channel. The STU-II does, however, require access to a key distribution center to obtain a new encryption key for every call that is made; however, this is completely automated. Its use has been limited primarily to government installations that need its secure capability and for which the key center is maintained by the government.

Historically, NSA has been prevented by law from acquiring communication security equipment against future needs; it could only order equipment for which it had end-user funds in hand. Recently, however, special industrial arrangements have been made that allow producers of the STU-II units to market directly to end-users, subject only to the restriction that the purchaser meet the criteria established by NSA (i.e., it must perform a government activity or be a government contractor). In the base communications context, this would of course not be an impediment; hence, base tenants could acquire STU-II units with their discretionary funds and could directly telephone any other STU-II user. The price, however, could be too high for widespread application. Unit price varies from about \$14,500 for a STU-II with one handset to approximately \$40,000 for units with additional handsets (a maximum of six), including installation, maintenance, and warranty.

Technically, the STU-II passes analog voice through specialized processing (linear predictive coding) to produce a digitized representation that is encrypted prior to transmission. It is possible therefore to enter the STU-II electronically through a data port without voice processing; thus, the device can provide both secure voice and secure data communications at 2400 bits/sec without specially conditioned lines. It can support computer terminal traffic, with acceptable response times as long as the users are dealing primarily with textual applications or computational problems that do not involve moving large files to or from the terminal. Graphics support is of

course out of reach. In the data mode, however, synchronous input is required.

The STU-II utilizes a novel encryption concept called the *ignition key*. This is a removable item that is unique to an instrument and contains part of the key variable; without it, the instrument is inoperable in the secure mode. When unkeyed, the instrument needs only protection equivalent to high-value office equipment; when keyed, of course, the security level of the physical environment must match that of the conversation (or data transmission) taking place. In the extreme, the ignition key must be stored in a fashion consistent with the classification level of the traffic that it is intended to protect; however, the operational concept permits an individual to carry it with him. The intent is to assure that the key is not readily accessible to anyone having potential access to the instrument.

Encryption-key management is largely automatic; for each call, the user dials a 5-digit identification number in addition to the telephone number of the called party. The code is also used to arrange any special interconnect features, such as might be required in a metropolitan area that spans several area codes. From time to time, each user must take a minor action to insert new parameters. But for a secure call, the general process of dialing is essentially what one would do on any other telephone. After the calling and called instruments have each automatically interacted with the key center, the user switches to "secure mode" and carries on his conversation. There is a minor disadvantage in that a 4-wire circuit (or two 2-wire circuits) is required for full-duplex operation; with a single 2-wire circuit, the system is limited to a push-to-talk or voice-operated-switch half-duplex style of operation.

STU-III Low-Cost Secure-Voice Terminal

The *low-cost secure-voice terminal* derives from the STU-II, is fully TEMPESTed, and will make similar, but more convenient, secure voice and data communications possible. Unlike the STU-II, it will be totally self-contained in a deskset about the size of a standard multiple-key telephone instrument. Furthermore, key management will be largely self-contained, although access to a key management center (via

an area-800 number) will still be required periodically, e.g., annually. The user will dial as he normally would to a 4-digit extension, a 7-digit local number, or a 10- or 11-digit long-distance number. The price per instrument is estimated to be about \$2,000 to \$3,000. Five parallel detailed concept studies have been completed, and three concurrent engineering design efforts will lead to full-scale production.⁴⁶

The data port is expected to operate at up to 4800 bps either synchronously or asynchronously. There are other system-concept improvements as well. For example, at the time of original installation, parameters indicating the identity of the instrument (its telephone number) and its maximum permissible clearance level will be incorporated into the initial encryption set-up. Each time the instrument is used, it will automatically transmit this information to the other party, where it will be displayed electronically, i.e., the calling number and the maximum permissible clearance level for traffic from it. Thus, two parties in communication will know which handset has called and what the appropriate maximum clearance level of the transmission can be.

As with the STU-II, the sets can support either voice or data communications. Importantly, they will also be able to function full-duplex over the usual 2-wire circuit. Except for an occasional automated action from a control center that requires the user only to place the call, key management will be automated and transparent to the user. As with the STU-II, protection requirements will be minimal for both the instrument and the ignition key.

Also, as with the STU-II, the industrial arrangement will permit vendors to market directly to end-users, subject only to the restriction that the purchaser meet the criteria established by NSA (i.e., it must

⁴⁶David Burnham, "U.S. Picks 3 to Make Secure Phones," *The New York Times*, March 27, 1985 (city and record edition: p. A-21, col. 4; national edition: p. 13, col. 4). The five concept studies were performed by American Telephone & Telegraph, General Telephone & Electronics, International Telephone & Telegraph, Motorola, and Radio Corporation of America. The three companies selected to produce engineering development models are AT&T, RCA, and Motorola. GTE will function as the systems integration contractor.

be performing a government activity or be a government contractor). Similarly, each vendor will also be responsible for any specialized features that his intended market may present (e.g., special interface arrangements or cards to the many computerized telephone switches that are rapidly replacing old-style mechanical switches).

The STU-III will accommodate many, but not all, of the specialized features (e.g., call-forwarding) that contemporary computer-based telephone switches offer. For example, after call-forwarding, a user might return a call from other than his normal instrument. In such a case, the clearance level transmitted might not agree with that of the speaker and thus would not properly indicate the level of conversation (or data communication) that could take place. Voice recognition or some other ad hoc arrangement would have to be used to surmount this problem. Conference calls will also be a problem. Such shortfalls, however, are minor relative to the enormous advance provided for both secure voice and secure data transmission.

The STU-III programs in the longer run will include call-conferencing, rack-mounted versions to be used (for example) in a computing center to interface secure dial-in lines, and a cellular telephone version. The equipment will also be compatible with the wide variety of communications arrangements the post-divestiture environment affords, either for commercially provided services or for private networks. The one device that will be a problem is the statistical multiplexor. While analog voice signals have dead intervals, digitized and encrypted voice signals are steady digital streams. The STU-III modem uses the full available bandwidth so that time-division multiplexors cannot combine the outputs of several STU-IIIs.

Unlike the STU-II, there will be two versions of the STU-III low-cost secure-voice terminal. Type I will incorporate encryption appropriate to the needs of the military, defense contractors, or the civil government for protecting defense-related or other agency material at the highest security level. Type II will contain a compatible encryption scheme that is satisfactory for the protection of unclassified but sensitive government information, or for the commercial or industrial user who wishes to enforce privacy and provide communication security to his corporate affairs. In each case, NSA will

assure that the encryption details are proper for the level of protection required. Separate key management centers will be maintained by the government for the two populations of users.

Commercial COMSEC Endorsement Program

The COMSEC Endorsement Program is a cooperative effort with the commercial vendors of LAN systems to add encryption directly into the terminal. This will mean that all traffic on the LAN transmission medium will be protected against interception. The program is well beyond its initial stages, and marketable products are expected in CY 1986. Several vendors are actively cooperating on the details of this joint program, and more are expected to participate. NSA will provide guidelines and standards to the contractor who will then be responsible for the design and production. Prior to production, however, NSA will test the product design and, if it meets the COMSEC requirements, will issue a certificate of endorsement for it and add it to the master list of endorsed products. For such an approach to work properly, a supply of properly manufactured keys must be available to all users. NSA has agreed to provide keys and will specify the handling requirements for them.

Needless to say, key management questions must be addressed as well as subtle issues concerning the interface between the encryption process in each terminal and any other software that might be resident therein. It may be necessary to have some degree of trustedness⁴⁷ in the terminal operating system, especially in smart terminals with significant indigenous computational capability. Many organizations are opting for using commercially available personal computers as terminals; when these are netted by an LAN, very difficult software-related security issues can arise.

Moreover, many personal computers have an open system architecture that accepts new functional cards without any system modification. Such aspects have to be carefully thought through in the context of providing terminal-to-terminal encryption over a LAN. There may be operational issues to be addressed also. For example, some LAN-based office systems have a broadcast mode that allows one source to communicate with many

⁴⁷See p. 58 for a discussion of the concept of trust.

destinations; the encryption process in these systems will have to allow for such a feature or deny the user a possibly important capability.

As suggested earlier, it will probably be necessary for the several LANs that are likely to exist on a base to communicate with one another through gateways and mutual addressing arrangements. The end-to-end encryption from a terminal on one LAN to a terminal on another, possibly different kind of LAN, perhaps halfway around the world, is even more complex and will require the extensive COMSEC compatibility on which NSA is now focusing.

Trusted Systems

Quite aside from any terminal-embedded cryptography program, the typical LAN installation contains hosts or servers or other nonterminal components. While the cryptographic communication to such devices will have been handled properly, the security problem within these devices is a computer security problem, not a communications security problem. Where required by security regulations or by corporate requirements for controlling access to information, trusted systems will have to be employed.

The latter, too, is an aspect of NSA's responsibility through its National Computer Security Center. One of the principal reasons for creating the Center was to encourage the commercial development of software systems (and possibly mixed hardware/software systems as well) that contain appropriate security safeguards and are designed, implemented, and tested to prescribed rules. The Center will function as the certifying authority and will conduct its own tests to establish the quality of any submitted products.

Guidelines for design, specific security safeguards, and levels of certification have been published by the Center.⁴⁸ An example of applying these guidelines is given in a recent NRL report by Landwehr and Lubbes.⁴⁹

⁴⁸Department of Defense Trusted Computer System Evaluation Criteria, DoD Computer Security Center, National Security Agency, Ft. George G. Meade, Maryland, CSC-STD-001-83, August 15, 1983.

⁴⁹Carl E. Landwehr and H. O. Lubbes, *An Approach to Determining Computer Security Requirements for Navy Systems*, Naval Research Laboratory, Washington, D.C., NRL Report 8897, May 13, 1985.

Since the process of designing, testing, and certifying software so carefully and thoroughly that it

- has functional properties known with high precision,
- has high resistance to subversive attempts, and
- has integrity with high confidence,

is at most several years old, there are only a few items presently on the qualified products list.⁵⁰ However, the pace is anticipated to increase, and it is expected that a variety of products meeting the so-called B-1 level of trust will be available by the late 1980s. Even so, the emphasis has been, and probably will continue to be, on computer operating system software or major components of them (e.g., database management systems). Attention has not yet turned to the software of personal microcomputers, which may also need attention in the context of security controls in a LAN.

Depending upon many technical details of the Commercial COMSEC Endorsement Program, yet to be resolved, there may be subtle or difficult interactions with trusted software products that will have to be resolved.

Key Generators

As with the STU-II and STU-III products, NSA has concluded arrangements to allow producers of the KG-84A key generator to market directly to end-users, again subject only to the restriction that the purchaser meet the criteria established by NSA (i.e., it must be involved in a government activity or be a government contractor). This is a very versatile variable-rate device (64 kilobits/sec maximum) that can be used to secure base-level digital transmissions under present

⁵⁰*Evaluated Products List for Trusted Computer Systems*, DoD Computer Security Center, National Security Agency, Fort George G. Meade, Maryland, April 2, 1985. For a description of the SCOMP, see "SCOMP Trusted Computing Base," *IEEE/CIPHER Newsletter*, IEEE Computer Society, Silver Spring, Maryland, January 1985, p. 3.

doctrine and regulations for secure communications, e.g., for dedicated links between hosts. The unit cost is approximately \$4500 (see footnote 32 for information on product brochures).

OTHER TECHNICAL OPTIONS

To permit transmittal of classified information over the ARPANET (a part of which, MILNET, is now separate and will become part of the DDN), a new kind of encryption device had to be developed that could bypass the header needed to route packets through ARPANET nodes but would send the body of the message through an encryption process.

The first bypass key generator (or encryptor), developed under the sponsorship of DARPA, was called the Private Line Interface (PLI). A small number of PLIs were built, but they are not generally available as a production item. A follow-on item, the Internet Private Line Interface (IPLI), was also developed under DARPA sponsorship and is currently in a small production run, primarily for individual military networks, but possibly also for use by the DDN.

The IPLI is a relatively expensive device, and one is required for each connection to a network. Nevertheless, this and other encryption devices specifically designed for packet networks may be useful in certain circumstances for on-base networks, LANs and otherwise.⁵¹

Although present designs and costs might not prove generally attractive for base-level use, the concepts are sound and could be readily implemented in contemporary electronic methodology. There is no fundamental reason why cryptography, together with any necessary trusted system components, could not be integrated directly into the design of a packet switch.

Finally, the BLACKER program must be mentioned. An NSA effort, this program will provide a set of equipments exploiting the trusted system approach for providing end-to-end encryption protection through a packet-switched long-haul network. The program is intended specifically for such major efforts as the DDN, but the equipments themselves are not likely to be attractive for on-base use. However, the concepts incorporated in them will afford technical options for developing new security approaches to network security, on base and off.

⁵¹See Sec. IX, pp. 112-119, for an example of such an instance.

BASE COMMUNICATIONS PAYOFF

The case is almost self-evident for the STU-II and STU-III equipments. Either could be used to provide point-to-point dial-up voice and data security through the on-base telephone plant, and even off base through the commercial telephone system. While the price of the STU-II will probably deny its widespread use in the base-level communications plant, it could be exploited now for those circuits or places that need communications security and cannot afford to wait for a more cost-effective solution.

The STU-III is a much more attractive option for widespread application because its unit price is anticipated to be between \$2,000 and \$3,000, and production quantities are anticipated to be available in fiscal year 1987. It will provide both voice and data security, so that a single investment will fill both needs. For much on-base data traffic, the anticipated maximum data rate of 4800 bps is adequate. Furthermore, the requirements for physical protection fit readily into the typical base-level situation; special security enclaves or reconstruction of buildings would not be necessary, provided the physical location of each instrument is consistent with the security level of the traffic through it.

In this regard, it is possible to envision a situation in which the instrument would normally be used in its secure mode, but with unclassified traffic, say, in an open office environment; the ignition keys, when removed, would be properly protected for secure communications. When warranted, the locations of the instruments could be physically secured under a short-term requirement for secure operations.

Overall, the STU-III appears to be a very convenient, cost-advantageous, near-term security answer to a large part of the on-base traffic that uses--or could use--the telephone plant.

High Data-Rate Links

For computer-to-computer links that will require significantly higher data rates, point-to-point dedicated communications almost certainly will be the option of choice. For most such special links, the data rate capability of the KG-84A, and its later version KG-84C, is adequate. The KG-84s could be used to protect such links on a full-time basis, either for on-base or off-base traffic.

The KG-84A can run start-stop on the plaintext side but is synchronous only on the ciphertext side. The start-stop interval is on a character basis, and thus would be in the millisecond range. It has been used on point-to-point circuits with personal computers and word processors, but it will not work over packet-switched networks. The KG-84C has a rudimentary start-stop option on the ciphertext side; thus it might be operated in a packet-switched environment, although not efficiently. The bypass allows up to 80 characters of plaintext to be transmitted before switching to cipher under computer control. However, the KG-84C cannot be returned from ciphertext to plaintext without going through its standby mode.

Even dial-up connections which connote point-to-point or circuit-switched connections can be used for computer-computer transfers. Both the KG-84A and KG-84C work very well in these applications, since they only need recognize one another's synchronizing signals.

The industrial marketing of the KG-84 devices removes many of the problems previously associated with acquiring them.

LANs

At some point, estimated to be the very late 1980s or even the early 1990s, LAN terminals with embedded encryption will be commercially and widely available. Also, trusted products will be available for many hosts that are LAN-based, although probably not for all. It remains to be seen whether the LAN vendors who happen also to market hosts for their systems will be aggressive participants in the Computer Security Center activity. It may be possible in the time frame noted to provide full-up security safeguards for many LAN installations; on the other hand, products not now anticipated may emerge from military projects

whose developers decide to clone commercially marketable versions. The end-to-end encryption of inter-LAN or LAN/long-haul traffic might still be incompletely solved by the early 1990s.

The technical concepts in the IPLI and the BLACKER equipments are not likely to eventuate in new devices appropriate for on-base communications until the late 1990s, if even then, but their technology will be a part of NSA's Commercial COMSEC Endorsement effort. Other events and developments may well overtake the transfusion of their achievements into devices for applications other than long-haul networks.

The Near-Term LAN Outlook

In the interim, the problem of security on the LAN will remain an awkward one. In many circumstances, some, perhaps even extensive, security safeguards can be provided; but the general problem cannot readily be handled.⁵² For the most part, it will be a matter of handling each situation on a case-by-case basis.

On the other hand, there are possible architectures for LANs that rely on COMSEC encryption techniques without having to rely on COMPUSEC or trusted systems. Specifically, there are circumstances in which several groups of users at different security levels can share a common LAN for intragroup communication. No intergroup communication would be allowed if the groups or communities are at different security levels or are otherwise restricted from communicating with one another. However, for some applications, such an architecture would be very useful and more readily achievable in the near term.

Computer terminals that require high data rates (e.g., graphics terminals) also pose a problem. The KG-84 approach could handle the situation, but it might be overly expensive. Until the problems of LAN security are resolved, there may be no choice but to use dedicated links with KG-84s or protected-wireline-distribution schemes or colocation of the terminal with its host.

⁵²See pp. 75ff for possibilities.

SUMMARY

Many protective measures can be taken even now to afford protection against threats that a LAN might face. TEMPEST issues are understood and can be implemented. Much, but not all, of the experience in safeguarding computer systems is applicable to LANs as well (e.g., fire and physical protection; personnel control and training; control of physical access to the terminals in the system; administrative, management, and procedural controls). Special circumstances of one or a few communities sharing a LAN can be accommodated. A significant level of protection can be achieved for a LAN installation by exploiting what is known and available already; but to do so, there must be an awareness and concern for security from the initial planning into the operational life of the system.

Just as security safeguards must be an up-front design and installation issue for ADP computer systems, so they must be for LAN installations. Someone must "be in charge" of security for a LAN from design through installation/implementation and into operational status on an ongoing basis.

Part 3

CONCLUSIONS AND RECOMMENDATIONS

IX. FUTURE ARCHITECTURES

As the discussion in this Note has indicated, there are many technological opportunities that can be exploited for modernizing the base communications plant. In fact, there is a surfeit of choices that can be made. The problem in the near term will be to choose among them so that a smooth transition will occur from the present situation to a distinctly improved one. Moreover, the choices will have to be made to provide the flexibility that airbases need to grow and change independently of one another. Since there is no czar in charge of the base-level communications environment, the decisions also must be made to minimize the amount of centralized control and planning needed. Inevitably, there must be some cohesiveness of action across bases; hence, some things must be done for the benefit of all, whether one wishes to call it a centralized authority or not. In general, though, the commonality that must be enforced from base to base over communications improvements consists of standards (including protocols), preferred practices, interface arrangements, preferred products, and a general architecture that can tolerate the individuality of bases and their tenants, and the MAJCOMs.

At the present time, there are many LANs in place on bases, all from commercial vendors. There is continuing pressure to continue installation of even more of them. We have assumed, for this discussion, that the present situation will continue for three to five years, and we shall make some suggestions about what might be done to bring the base communications situation into functional interoperability and cohesiveness.

In spite of LAN proliferation, the Air Force has a significant program through AFLANSPO to specify and develop a standard LAN architecture and family of components; the program is known as ULANA.⁵³

⁵³*Unified Local Area Network Architecture, Functional Description*, a briefing from AFLANSPO, ESD/OCC-2, Electronic Systems Division, Hanscom AFB, Massachusetts, March 1985.

While draft documents exist for its functional reference model⁵⁴ and a prime item specification,⁵⁵ it is nonetheless still an engineering development effort of some magnitude. It remains to be seen how closely the ULANA specifications meet commercial LAN product characteristics-- or how many vendors will be willing to adapt their products as required-- before it will be possible to judge when ULANA products will overtake or replace presently installed commercial LANs.

Thus there is some uncertainty about how long the proliferation of commercial products will continue, and therefore the time frame during which the suggestions and proposals herein can be helpful. It is unlikely that ULANA products can be available in three years, but they may be in five, and should definitely be available in ten. Meanwhile, the Air Force must not temporize, lest it sacrifice important mission-effective operational capability. Hence, the general architecture proposed below⁵⁶ should hold for the near term, and probably also for the mid-term, with the assumption that ULANA products will gradually supplant commercial ones. The far term may be an extrapolation of mid-term, or it may be a radically different system and technical approach.⁵⁷

NEAR TERM

There are several factors that will strongly influence the on-base events of the next five or so years. First, an airbase has a functioning communications environment today, and it cannot be summarily junked to make way for a radically new arrangement. Quite aside from operational consequences, it would be impossible for the Air Force to justify such major expenditures.

⁵⁴*ULANA Reference Model (Draft)*, Booz, Allen & Hamilton, prepared for AFLANSPO, Electronic Systems Division, Hanscom AFB, Massachusetts, July 1985.

⁵⁵*Prime Item Specification for Standard AF Network Interface Units (Draft)*, Booz, Allen & Hamilton, prepared for AFLANSPO, Electronic Systems Division, Hanscom AFB, Massachusetts, November 1984.

⁵⁶See pp. 106-112.

⁵⁷See pp. 112-119.

- Therefore, a near-term architectural framework for base-level communications must be one that can be reached by an evolutionary path from today.

Second, it must be taken as given that the Air Force will use predominantly commercial equipment on CONUS bases. While it will develop some specialized equipment that does not exist in the commercial market,^{5*} the Air Force could never justify the R&D funds to develop a complete line of special equipment (e.g., telephone switches) for its bases.

- Thus, the architecture must be one that accommodates a variety of commercial equipment, in particular LANs and modern digital telephone switches.

Third, the base cable plant will continue to be a part of the base communications scene. Twisted-pair cables will support switched voice/data systems for a long time to come, partly because of the sunk cost in presently installed cable plants and partly because twisted-pair circuits will be the economic way to provide certain kinds of service. Undoubtedly, the Air Force will have to spend money upgrading its cable installations, partly because of increasing demand for circuits but also to replace deteriorating cable and to provide the quality of circuits needed for modern switches and digital services.

^{5*}For example, ISA/AMPE, which is a trusted message handler that did not and will not have a commercial equivalent in the time frame of need; it is the interservice version of an earlier Air Force equipment known as AF/AMPE. It will provide the electrical interface between long-haul networks and on-base communications; thus it will be able to automatically distribute incoming messages on the basis of addressee while enforcing appropriate security controls. Similarly, it can aggregate outgoing messages for insertion onto the long-haul network. In the latter case, it assumes that the security label that is attached to each outbound message is correct as received at the ISA/AMPE.

Cable Planning

Considered and reasonably optimized choices must be made in regard to size of cable, routing, etc. To the extent possible, new cabling should be installed with an eye to anticipated growth in the coming ten or so years, and to assure survivability (e.g., buried deeply, away from likely targets of dissident attack, protected when in aboveground facilities).

There is a progression of technical steps that can be taken to exploit new cables and, where quality permits, old cables as well. As load increases, such techniques as statistical multiplexors, channel banks to increase the bandwidth transmitted on a twisted pair, and terminal equipment to provide T-1 (or better) service can be employed. The life of the cable plant can also be extended for a long time, again by a progression of technical steps. Planning for new cables must properly include consideration of technical enhancements as part of an orderly phased increase in capacity. As an example, manholes for the repeaters that will be needed for T-1 service must be included in the planning, or at least provision must be made for building additional ones when needed. Moreover, they must be designed to prevent casual entrance that might lead to damage and disruption of service.

Conservation of some present cables is likely because of a historical tendency to use an individual circuit for things that underutilize the circuit, e.g., a single alarm or a simple intrusion sensor. Such extremely low data-rate devices do not need full-time circuits, but some specialized terminal equipment will be needed to multiplex several of them onto one twisted pair.

Therefore, it would seem in the best interests of maximizing the lifetimes of base cable plants for AFCC to:

- Establish a "Cable Management Office" responsible for planning the orderly extension of cable plants to accommodate increasing loads; for managing and general oversight of cable plants; for monitoring the installation of new cables to maximize survivability; and generally to assure the best return on investment of the present and to-be-installed cables, either twisted-pair or fiber-optic.

Among other things, a computer model should be developed that can give each base technical guidance for cable planning and utilization. In the far future, this office would phase out of existence when, and if, other forms of transmission (e.g., coaxial cables) completely replace current systems.⁵⁹

Switches

AFCC is presently embarked on the SCOPE DIAL and SCOPE EXCHANGE upgrade of ancient switches. On some bases, vendors may insist on cable replacement simply because the deterioration of the present cables cannot be tolerated by the new switches. The vendors will not risk the performance of the switch by connecting it to questionable circuits. If the Air Force should experience this attitude, some planning can be done to provide for future expansion rather than simply replacing present cables with ones of similar size. To the extent that the SCOPE programs also bring new cable installations with them, the cable investment--be it initial capital investment or ongoing O&M costs--will continue and will result in ongoing pressure for long-term use of the plant.

Integrated Switches

With respect to the integration of voice and data, there is an interaction between the technical characteristics and size of the switch and the distribution of the data load. If switched twisted pairs are used for data and if enough of the data load is steady-state rather than aperiodic or bursty, the switch must be larger than just for voice service or it will overload. The problem is that little is known about the distribution of data loads on an airbase, especially as Phase IV⁶⁰ comes in. Unless such projections can be developed, it will be necessary to start with the best estimate for the switched data circuits to see how the load develops, but also install a larger switch or, at least, a readily expandable one. For data loads that are beyond the

⁵⁹See, however, the discussion of Distributed Switches on p. 103.

⁶⁰The equipment replacement program for on-base functional-area computer support, and also for command-unique applications.

capacity of the usual voice-grade line, suitable dedicated wideband or T-1 service on twisted pairs or multiple twisted-pair circuits (possibly using some of the technical enhancements suggested previously) can be installed.

Should the SCOPE programs insist on integrated switches at the outset? Unfortunately, the data to support a firm decision do not seem to be available. There is no hard information to answer such questions as:

- How much dial-up data load is there?
- What are the data rates and periods of activity? During each shift? By base? By functional area? By buildings?

Certainly, a data user who is active much or all of the day should have a dedicated circuit, unless the data must be routed to a variety of recipients; in this case, the user is, in effect, dialing a sequence of addressees. An important observation is that the digital switch can act as a general-purpose gateway among technically compatible users.

Intuition suggests that there might not be very much dial-up traffic load, or at least there might be so little that it represents a minor effect on the switch. Perhaps most terminal users require, or desire, that their communications be served otherwise; or perhaps most of them will be homed on a LAN. If so, the decision concerning an integrated switch should be based on an economic analysis comparing the marginal cost of the data feature--and possibly of a larger switch--with the cost of other kinds of dedicated communication arrangements.

In this connection, the use of STU-III equipment, which provides both secure voice and secure data communications, may impact the situation because communications security will overwhelm other considerations. Its widespread use might distort projections of dial-up data traffic that had been made in an unsecured context.

Distributed Switches

Many modern switches allow for distributing parts of themselves to remote sites. So to speak, a "switch" becomes a group of switches connected (in traditional telephone parlance) by interoffice trunking. If a community of interest (say, within a building or a group of buildings) predominantly talks or transmits data among its own members, then a local or distributed switch might make sense. If a part of a distributed switch serves a community of users that form a cohesive unit that can contribute to base mission performance, even though damage has occurred, a local switch can make a contribution to functional-area survivability--which is not the usual notion of, or approach to, survivability.

The odds are that no data exist on the distribution of voice/data traffic by functional areas, so it is not possible to make a decision based on hard information. If such data can be developed, it could prove advantageous to distribute the switch to several remote locations. Again, though, an economic analysis would be needed to compare the cost of remoting the switch (not only the initial cost, but also that of ongoing maintenance) versus the savings in cable circuits or the increased survivability that might be achieved.

There is an interaction between cable-plant planning and the switch configuration. A distributed switch has the effect of concentrating the switching function in localized geographical areas and therefore acts to reduce the number of circuits and cables needed to the main central switch. In this regard, the Cable Management Office must be informed about switch arrangements and must interface with the planning for possible distribution of them.

Local Area Networks

LAN technology is of course useful for a localized community that (1) needs much connectivity within itself, (2) needs a reasonable flexibility to attach or remove devices to or from the network, and (3) has a traffic load that tends to be bursty rather than steady but may occasionally have a high data rate. Thus, it makes sense for communities of interest on an airbase to install LANs to serve their

specialized needs. But, as pointed out in the previous discussion,⁶¹ there are technical problems that must be addressed to permit connectivity among LANs.

The interface issue--standards, protocols, addressability--cannot be given too much emphasis. Unless it is handled well, LANs now in place may not be able to become part of a comprehensive base communications architecture without costly development of specialized interface gateways or equally costly LAN replacement.

Many, but not all, LAN vendors provide connectivity to a telephone network. Thus, the external traffic of some LANs--to other LANs or to other recipients--can be serviced through the switched telephone network, using all the cable options that have been discussed earlier to allow for growing loads. If the extra-LAN traffic load is too heavy for enhanced twisted-pair circuits, then other kinds of dedicated circuits must be provided. These are likely to be broadband technology of some kind (e.g., microwave, fiber optics).

Again, as already noted, there are technical issues that will have to be resolved, notably the variability of protocols and inter-LAN addressing. Interconnection, via the switched network or a dedicated link, may require specialized gateways or interface arrangements. Moreover, digital messages on a LAN must contain the internal system address of the recipient, whether he is on the same or a different LAN. Traffic originated within one LAN may require special arrangements--possibly involving protocols--to permit addressing other-LAN recipients.

Unfortunately, there has been

- An uncontrolled acquisition of LANs by base tenants which has resulted in a major operational shortfall. Connectivity among LANs--interoperability--has suffered and may be impossible in many cases without development of specialized arrangements or replacement of LAN installations.

To the extent that AFCC or USAF/SI can influence base-level users' actions, the choice between providing the required service by a LAN or

⁶¹See Secs. III and IV.

by switched circuits must include consideration of economics and the nature of the data traffic. However, from the viewpoint of a long-haul network such as MILNET or DDN, a LAN is a much more efficient means for distributing local traffic than requiring users to dial in to a network interface unit, e.g., a TAC or mini-TAC.

LAN Planning

Just as it seems wise to establish a central function for managing the cable plants, so it would also be in the best interests of the Air Force for AFCC to:

- Establish a "LAN Management Office" that will be responsible for assisting with the orderly installation of LANs, for resolving the diverse interface issues, for providing guidance to base tenants on selection of LANs, and for generally providing the ongoing oversight to assure that LAN installations fit smoothly into an integrated base communications architecture.

The AFLANSPO office at the Electronic Systems Division (ESD) is probably not well positioned to discharge such a broad scope of responsibility. It can handle technical details and has taken an important step with the ULANA⁶² program, but it cannot handle policy, guidelines, standards, etc. The job must be done by the Air Force Information Systems community which reaches onto all bases and into all MAJCOMs. The USAF/SI is obviously a participant in any such approach, but AFCC would seem to be the proper organizational location for the function per se.

Video Services

Video services are sometimes used as an argument for installation of coaxial cable systems to accommodate video, voice, and data needs. Some tenants or communities of interest may install a LAN that provides such features; but for others whose data requirements can

⁶²See p. 108 for a brief discussion of ULANA.

be met in other ways, simple CCTV approaches are adequate and much less expensive, especially with regard to interface units.

NEAR-TERM ARCHITECTURE

For the reasons already noted--proliferation of LANs, an in-place and continuing cable plant, modern switches, and because the possibilities for making changes on base are constrained by actions already taken--the general near-term architecture for at least the next three, probably five, and possibly ten years will be:

- A collection of LANs connecting to one another or to the base telecommunications center,
- Through the cable plant and the digital telephone switch;
- Supplemented by point-to-point dedicated circuits when the traffic load warrants;
- Including specialized interfaces or gateways unless the Air Force insists that all LANs be retrofitted to a standard interface that includes such things as electrical details, data rates or flow control and protocols; and
- Dealing with wideband video services separately except as some choice of a LAN opportunistically provides an option to combine them with data services on the same bus medium.

Forces presently acting will drive the base communications architecture in such a direction, but the Air Force--notably AFCC and USAF/SI--must address certain details that are not now well handled, e.g., standards, interfaces, positive management of the cable plant. Important details of the near-term architecture might not evolve properly--or evolve at all--if the Air Force does not move aggressively and assertively.

If the data loads are primarily dial-up or if the economics of the situation warrant, the switch might be an integrated voice/data switch. If data loads are predominantly steady (as they might be on an AFLC base), then the switch is likely to be voice-only or voice plus a

limited data capability, supplemented by dedicated point-to-point circuits of appropriate bandwidth.

Such an architecture can be reached by an evolutionary path from the present. Furthermore, over time, if and as bases gradually move toward comprehensive wideband communication arrangements in the more distant future, the services that are inherently wideband (video) can be blended with voice and data.⁶³ To reiterate, the emphasis is on *evolution*. Build on the SCOPE programs; exploit the cable plants; and place interfaces among LANs, the telephone plant, and off-base networks--all in a technically consistent framework.

COST CONSIDERATIONS

Quite aside from technological possibilities and preferences, there are cost considerations associated with the various approaches. For example, a scheme that can link terminals to a computer directly by a twisted pair is very likely to be the least expensive choice for a single service. Interface devices differ widely in cost; a modem to interface conventional telephone circuits will cost \$200 to \$300, but a bus interface unit for a baseband LAN will cost approximately \$700 or more, and the BIU for the computer host connection will cost several thousand dollars. If one is connecting wideband sources (e.g., a TV camera) to a broadband LAN, the BIU will be correspondingly expensive as well. At some point, these products will incorporate security features such as trusted software components or encryption capability; the prices at that point are likely to escalate beyond those suggested. If one chooses an approach such as the STU-III Secure Low-Cost Terminal, the unit cost will be about \$2000, but the system provides both voice and data security as well as the modem. The relative costs of modems to BIUs is a point in favor of using the base cable plant when this is operationally acceptable.

Transmission media vary in price as well, as do installation and maintenance costs. Splicing a coaxial cable is obviously faster and simpler than splicing a 100-twisted-pair cable.

⁶³For a commercial application of the "PBX + LAN" philosophy, see "User [Upjohn Corp] to Tie PBX, Local-Area Nets," *COMPUTERWORLD*, Vol. XVII, No. 23, June 6, 1983, p. 1.

For the foreseeable future, telephone communication will be a permanent part of base communications, so telephone switches and a cable plant will exist. The marginal cost of upgrading the switch to accommodate data and of installing new cable is likely to be small compared to the installation cost of a basewide broadband system plus the cost of all its interface units.

We are not attempting to present definitive cost data or to make a comprehensive cost argument; rather, we merely observe that cost comparisons in the absolute may very well not lead to the preferred solution based on technical considerations. The situation is so variable from base to base that one can only make the detailed cost assessment for each individual case. Attributes other than cost may dominate the choice of a preferred technical approach, e.g., flexibility, technical expandability, graceful degradation of service, graceful growth capability, or even the availability of trusted components.

Finally, some attractive attributes of a LAN may be unimportant in some instances. For example, flexibility for expansion or for moving the location of subscribers may have little value for a stable physical situation where there is already a large cable plant investment. In some circumstances a LAN may not be the cost-effective way to meet present needs.

ULANA

The Unified Local Area Network Architecture (ULANA) is a component of the Air Force Local Information Transfer Architecture, which in turn is a component of the Air Force Information System Architecture.⁶⁴ ULANA is a program to develop a standard architecture and family of

⁶⁴See the Program Management Directive noted in footnote 66. See also *Air Force Information Systems Architecture, Volume I--Overview*, Hq U.S. Air Force, Washington, D.C., May 8, 1985. This is a top-level, recent Air Force document that addresses standards (e.g., protocols), methods of acquiring information systems (e.g., Series-700 regulations), and such topics as user considerations, design guidance, and system implementation. Importantly, it also assigns responsibilities, notably to AFCC and to USAF/SI, for various actions.

components to satisfy Air Force LAN needs. It is managed through the AFLANSPO under the cognizance of USAF SITT and has the following major characteristics:

- Uses a segmented tree hierarchical structure.
- Permits any transmission rate but recommends broadband.
- Provides for a full range of network interface units and a network management system center.
- Provides bandwidth for analog grade signals which are co-resident on the cable but not functionally integrated with data services.
- Provides for the usual LAN data services among terminals, hosts, servers, and other digital devices.
- Provides for gateways to military and commercial long-haul networks.
- Has a protocol structure consistent with the ISO seven-layer model and with the DoD Protocol Model.
- Provides for full-duplex analog voice-grade products.
- Provides for the development of comprehensive security safeguards.

However the ULANA effort matures as an equipment development effort, the program is a major step by the Air Force toward bringing uniformity to the LAN situation. The ULANA documents are a very good characterization of the military LAN-type requirements as perceived by Air Force users.

TELECOMMUNICATIONS CENTER⁶⁵

At the moment, the typical base telecommunications center reflects the ways of the past. It is labor-intensive, involves manual handling of much paper, physically reproduces copies of messages for courier distribution, sometimes examines message content to assist in

⁶⁵Originally this was called the Base Communications Center but more recently it is regarded as a component of the base Information Processing Center (IPC). However, as of this writing, only two bases have colocated the functions of communications and computing. Thus, the name Telecommunications Center will be used to emphasize that the discussion concerns only communications.

determining distribution routing, and manually keyboards messages or processes hand-delivered specially typed messages through an optical character reader. Only in very limited cases is traffic delivered electrically to the end-recipient, primarily because the security problem has not yet been adequately dealt with. Messages of varying classification must be separated from one another with high confidence and delivered to users who hold the proper clearances; the trusted equipment to enforce classification rules in an all-electrical environment does not yet exist.

In general, the base telecommunications center is a "place" at which one receives or delivers hardcopy messages that travel via the off-base long-haul networks. It imposes a strict and uniform discipline over the entire base and for all users of the external networks. One observer of the situation has referred to the "tyranny of the comm center."

With the advent of AMPE--and later, ISA/AMPE--equipments, many bases can begin to electrically distribute message traffic arriving from external networks and can do so while strictly enforcing classification restrictions. This is the trend of the future.

Near Term

As the base telecommunications system for message traffic becomes more and more automated, it will acquire new equipment and new functions. In addition to direct electrical distribution of incoming traffic and aggregation of outgoing messages for insertion onto the external networks, there are several other things that may have to be taken care of, including the following:

- Speed-matching between the flow rates of the external networks and the on-base networks; buffering arrangements as well as flow control in both directions are necessary.

- Protocol conversion between on-base and off-base systems.
- Provision of modem banks for interfacing off-base telephone circuits with data traffic.
- Specialized interfaces between networks as required.
- Monitoring the technical status of the entire hierarchy of on-base systems and controlling flow onto and off base as required by external network conditions.

The base telecommunications center could become the place at which all of these functions are handled and could acquire a variety of specialized equipments such as modem banks or protocol converters. It would thereby become much less labor-intensive, more highly automated, and more like a modern-day remotely accessed computer center, rather than a post office. It could become the *Network Management Center* for on-base systems and be responsible for assuring minute-by-minute continuity of service, directing reconstitution of damaged facilities, using alternate off-base routings to avoid traffic congestion, monitoring on-base network performance, optimizing flows off/onto base, and similar duties.

These functions need not be physically concentrated at one place. For example, the interface agreement between on-base and off-base systems might stipulate that each system must handle its own speed-matching and any required protocol conversion. Furthermore, each on-base system might include the necessary modem banks to interface off-base telephone networks. The equipment might be decentralized, but for operational efficiency and probably technical cognizance, control and maintenance would still remain with the Telecommunications Center.

Far Term

As the level of automation grows and as on-base networks become increasingly self-sufficient in terms of interoperating with other networks, the Telecommunications Center will become a function rather than a place. It is likely to retain operational monitoring, technical oversight, and maintenance responsibility for on-base communications and

their interface to off-base networks, but it will not be a central repository for equipment other than that needed to monitor status. In the extreme, not every base will need the function. As the telephone companies have demonstrated, it is possible to operate, monitor, and maintain complex installations of electronic equipment remotely. Furthermore, modern digital switches already are being routinely diagnosed remotely and maintained from afar. New software is inserted remotely and the maintenance person on the spot is told explicitly what hardware replacements to make. An entirely new concept of maintenance will be used.

Thus, a cluster of bases may share a common "Communications Operating Facility," and the historic Telecommunications Center will disappear. The functions of such a facility are clear from the prior discussion; but importantly, it must be able to dispatch repair personnel on a timely basis to take necessary hardware actions or, possibly, to install new software.

FAR-TERM CONUS BASES

Proposals have been made⁶⁶ to "wire" a base with broadband cable that can accommodate all needed services--telephone, data, video, alarms, special sensors, etc. As the discussion has previously suggested, baseband systems are not efficient carriers of telephony, so one is immediately directed to a consideration of broadband LANs which have ample capability to handle all needed services.

However, the LAN approach presents two problems: geographic limitations and the costs of interface units. Traditional LAN technology may not be able to span the size of an airbase, and interface

⁶⁶For example, AFCC's Base Level Information Transmission System (BLITS) Statement of Need, which was finalized on April 25, 1983. The effort has been renamed Mission-Effective Information Transmission System (MEITS) and is supported by a Program Management Directive dated April 12, 1985. Among other things, the PMD directs the development, acquisition, and testing of the ULANA family of standard components. Under the MEITS effort, AFCC/DDR chairs a working group addressing the overall architecture of the Local Information Transfer Architecture (LITA). AFCC/EPP is also defining the Base Information Digital Distribution System (BIDDS). The overall effort is a joint AFCC/AFSC responsibility that involves the AFLANSPO at ESD and also ESD/OC. USAF/SITT is the Office of Primary Responsibility.

units to the cable will be costly whether the device is a computer terminal or a simple low-data-rate sensor. The broadband concept (e.g., coaxial cable, fiber-optic cable) is attractive, however, because of the geometric simplicity of the backbone itself (a simple coaxial or optic cable structure that is easily spliced or extended, versus a many-pair cable) and the availability of technology--that of cable television--that essentially removes any geographic limitation.

Cable television commenced as a simple distribution system for wideband television signals but has progressed to two-way systems that enable the subscriber to communicate in a modest way with some central point on the cable.⁶⁷

There is commercial interest in expanding cable systems to greater sophistication, and equipment development is actively under way that will support a wide variety of services.⁶⁸

An Approach

One company⁶⁹ has produced and is beginning to market systems that can be added to existing cable plants and will support a wide variety of services.⁷⁰ Its "PacketCable"⁷¹ system exploits packet transmission techniques to allow subscribers to address requests or messages to the cable control center, and also permits the center to return control information to service user requirements (e.g., set the local controller to select a particular channel, respond to a password entry to establish

⁶⁷For example, the QUBE interactive two-way system operated in Columbus, Ohio, and elsewhere by Warner-Amex.

⁶⁸Paul Baran, "Broad-Band Interactive Communication Services to the Home: Part I--Potential Market Demand" and "Part II--Impasse," *IEEE Transactions on Communications*, Vol. COM-23, No. 1, January 1975, p. 5 and p. 178.

⁶⁹Packet Technologies, Inc., Cupertino, California.

⁷⁰Other companies are also experimenting with telephone/cable combinations. MCI has its "Cablephone," which uses the cable to reach an entry to MCI's long-distance network. Others (for example, 3M) have experimented with frequency-division multiplexing to add telephone service to cables. The so-called "institutional network" in New York City uses cable technology to provide uptown-downtown data connectivity.

⁷¹A trademark of Packet Technologies, Inc.

some service or remove some restriction, provide videotex service). Importantly, the bulk of the intermediate electronic equipment is on the cable, not on the subscriber premises, so the subscriber is not inconvenienced with space-consuming equipment.⁷² The subscriber--in his service context--can have such features as interactive videotex, conventional television channels, special cable television channels, personal computer communications, home security and emergency services, and special television-related services, including password control of access to certain programs and user-selectable channel recall on a time/date basis. In addition, the system is planned to support an RS-232 interface for personal computers.⁷³

Such a system is of course computer-based and controlled, with the result that the cable operator can digitally address subscribers, respond selectively to any one, arrange any scheme of charges that he desires, etc. By adding additional computer capacity as required, the operator can also offer such features as electronic mail, greeting card services, bulletin boards, single- or multiple-person video games, electronic newspapers, educational services, electronic fund transfers, and bill payment.

A second system, "PacketPhone,"⁷⁴ is analogous to the traditional telephone-loop distribution technology and can stand alone or coexist with PacketCable, but it offers full voice and data services.

Using a proprietary product known as "PacketDax"⁷⁵ and exploiting both the silences inherent in voice conversation and the new CCITT adaptive differential pulse code modulation, the resultant packetized

⁷²Paul Baran, "PacketCable: A New Interactive Cable System Technology," *National Cable Television Association Conference Proceedings*, May 1982. See also the corporate brochure from Packet Technologies, Inc.

⁷³The prototype of such a system is now undergoing beta-level field testing on the United Cable system serving Cupertino, California, and the Hearst Cable system in Los Gatos, California.

⁷⁴A trademark of Packet Technologies, Inc. See *PacketPhone Tutorial*, a corporate presentation by Packet Technologies, Inc.

⁷⁵A trademark of Packet Technologies, Inc. See *PacketDax*, a corporate brochure of Packet Technologies, Inc.

voice capacity is effectively increased fourfold over the conventional T-1 transmission system. The proprietary-node equipment handles a maximum of 160,000 packets per second.⁷⁶

The costs of such installations obviously depend upon whether a cable system is in place or not. Preliminary cost analyses suggest that even if a cable system must be installed, the PacketPhone approach competes favorably with conventional telephone technology and can in many instances achieve total economic payback in a few years. No detailed cost analysis has been made for an airbase.

Since such systems will be operated by computer, they obviously can function in a common-carrier arrangement or can support independent private voice and/or data networks. The whole ensemble is software controlled, which means that *software defined components* are readily implemented. Each one, such as a private network of home computers, is in effect a *virtual network*⁷⁷ which can be dynamically modified by parameter changes in the control software. In an important way, such concepts as PacketCable and PacketPhone are the ultimate in virtual structures.

Again, since traffic on the cable (other than video) is digital, it can readily be gatewayed to off-cable destinations or sources (e.g., other cable networks, long-distance carriers, satellite circuits, specialized databases, computer systems).

On-Cable Transmission

Any cable system has to contend with signal-to-noise problems. Because economics drives commercial installations, the coaxial cable used is not of the highest quality throughout; moreover, connections will corrode in the weather environment. Thus, noise continually creeps into a cable system as it progresses outwardly toward subscribers. For systems that distribute video, the desired signal must be kept large enough to provide the desired signal-noise ratio everywhere, and the

⁷⁶This equipment is in beta-level field testing with Michigan Bell in the Detroit area.

⁷⁷See p. 17 for a discussion of the virtual concept.

best quality cable must be used on the main trunks (cable with double continuous shielding, for example) to get good quality signals out to the local distribution and individual drops.

However, the inbound signal has the opposite problem. As it travels up the cable network, it starts where the cable is most susceptible to noise and progresses into the better quality part of the network. As any one signal moves inwardly, it is joined by many others, each of which brings its own noise contribution, so the aggregated signal deteriorates rapidly. The so-called *upstream noise problem* has been dealt with in PacketCable by using a blend of two packet-switching approaches. Near the subscribers, an Aloha-type (see footnote 13) contention access system (but without carrier sensing) is used; but before the signal has degraded too seriously, bad packets are removed and everything is converted to time-division-multiplexed synchronous transmission which can be regenerated as required to maintain signal quality. This technical advance has made it possible to provide extensive upstream information flows and data rates.

Base Application

The system outlined above has been developed for commercial applications, but the concepts and technical advances embodied in it have obvious potential application to base communications. The traffic capacity of the system is in the right range for a base, and since everything is computer controlled, flexibility and adaptability are high. Expansion to new areas is simple: extend the cable, add the additional local controllers, and update the software tables in the control computer. Private voice and/or data networks to serve a community of interest or a functional area can be implemented as a *virtual LAN* of any geographical extent or configuration. Any subscriber can be hosted on many LANs or services. There can also be computer-controlled *gatewaying* to other virtual LANs, or off base, or to wherever desired.

It would seem that an easy transition could be made from a twisted-pair base to a coaxial base. At the outset, some services could be moved onto the coaxial system, with the rest following as twisted-pair cables deteriorate, are destroyed, or eventually run out of capacity.

If the government wishes, it could arrange direct access to long-distance carriers, the so-called *bypass telephone network* approach; a base could thereby become independent of the local telephone company. Interfaces to long-haul circuits would be easy, since digital traffic is already packetized and could have compatible format, protocols, and headers for such systems as the DDN. In fact, one could probably arrange *direct addressing* from any subscriber to any other subscriber on any base that is equipped with such a system anywhere in the world.

Finally, the digital representation of voice is consistent with the approach of the STU-III equipment to provide end-to-end encrypted security for voice and/or data. At present, the STU-III requires a circuit-switched transmission path because its modem has been optimized to that purpose. However, there is no fundamental technical impediment to adapting it for an all-digital packet environment, although a major engineering development program would be required. For data rates that STU-IIIs or other higher-rate devices could handle, security would appear to be rather straightforward. The principal nonelectronic problems would be physical security of the plant and adequate redundancy of routing and control facilities to avoid system collapse because of damage or outage.

The commercial system, as outlined, might or might not be directly applicable to an Air Force base; but it does offer one direction of technical approach that is consistent with an all-broadband base and that responds to the requirements as they are now understood. Not all aspects are necessarily in hand, but at least technical feasibility appears assured and economic feasibility is likely.

Network Information Services

If the Air Force decides to use a sophisticated cable-based communications arrangement for an airbase, it should think broadly about a vastly more extensive array of network information services. Many services from the commercial consumer market (see footnote 68) including high-quality commercial television and audio, cashless transactions, electronic mail, comprehensive telephone services, and access to commercial databases could also prove useful for the base. There are others unique to military requirements which are technically feasible

because the "super-cable system" is wholly computer-based and controlled:

- Training and continuing-education materials could be made available through terminals anywhere on base--in the home, at the office, even on the flight line, or while standing alert missions. Such materials could be tailored to the needs of various skill groups--including aircraft maintenance, some aspects of pilot proficiency, and simulations of combat--or even to individuals.
- Daily base news items could be distributed electrically and, if required, selectively.
- Personnel recall for emergencies could be automated through the cable-based telephone system.
- Base alerts, emergencies, or announcements could also be handed to all, or selected, telephones automatically--a "super 911" level of service. Groups of telephones could be made to ring simultaneously and each would deliver an appropriate personal or recorded message.
- Visual warnings could be displayed to all, or selected, television receivers, perhaps preempting programs that have been selected. It could even be possible to turn on sets.
- Teleconferences or telephone conferences could be convened readily and rapidly.
- On-line discussion forums of professional, Air Force, or recreational interests could be sustained.
- Interacting communities of interest could be readily established and dismantled on short notice.
- Medical records could be promptly transmitted from place to place in the event of an emergency.
- Databases essential to the continued functioning of the base in the event of natural disaster or deliberate damage could be maintained.

The extent of such services is limited only by the imagination. In the end, to use the popular phrase, "it is only a software matter."

If an airbase is to become ultimately dependent on its information infrastructure and the communications/computer complex that makes it possible, the issue of *survivability* must be addressed from the beginning, not only the problem of deliberate damage but also the usual vicissitudes of operating a very complex system. Cables would have to be run redundantly, even though this is not done in commercial systems because of cost. Computer equipment would have to be "nonstop" against all conceivable hardware and software anomalies. Terminal equipment likewise would have to be replicated or internally redundant to avoid being isolated from the network. Adequate emergency power would have to be available and within a millisecond's demand.

FAR-TERM COMBAT BASES

With the invention of packet-switched technology⁷⁸ and its implementation as the ARPANET under DARPA sponsorship, the foundation

⁷⁸A series of eleven Rand publications by Paul Baran details the Distributed Adaptive Message Block Network, a digital data communications system based on a distributed network concept. Such a system would now be called a packet network.

On Distributed Communications: I. Introduction to Distributed Communications Networks, RM-3420-PR, August 1964.

On Distributed Communications: II. Digital Simulation of Hot-Potato Routing in a Broadband Distributed Communications Network, RM-3103-PR, August 1964.

On Distributed Communications: III. Determination of Path-Lengths in a Distributed Network, RM-3578-PR, August 1964.

On Distributed Communications: IV. Priority, Precedence, and Overload, RM-3638-PR, August 1964.

On Distributed Communications: V. History, Alternative Approaches, and Comparisons, RM-3097-PR, August 1964.

On Distributed Communications: VI. Mini-cost Microwave, RM-3762-PR, August 1964.

On Distributed Communications: VII. Tentative Engineering Specifications and Preliminary Design for a High-Data-Rate Distributed Network Switching Node, RM-3763-PR, August 1964.

On Distributed Communications: VIII. The Multiplexing Station, RM-3764-PR, August 1964.

On Distributed Communications: IX. Security, Secrecy, and Tamper-Free Considerations, RM-3765-PR, August 1964.

On Distributed Communications: X. Cost Estimate, RM-3766-PR, August 1964.

On Distributed Communications: XI. Summary Overview, RM-3767-PR, August 1964.

For an article about Paul Baran, see "Packet-Switching's Unsung Hero," *New Scientist*, September 8, 1977, pp. 606-607.

was laid for the ultimate in a fully distributed, highly redundant communications network. DARPA has also sponsored variations on the theme, such as the ALOHA network (see footnote 13) and the follow-on packet-switched radio. These ideas can be extrapolated to a feasible communications architecture for bases that are subject to battle damage, and for bases that are normally bare but onto which U.S. forces might deploy for tactical operations. Without modern electronic technology, it would be impossible to envisage an approach such as the one outlined below.

We start with the basic premise that a facility that has been destroyed no longer has a need to communicate. It follows that any communications-relevant equipment which might have been in such a facility must not have any consequence for the proper operation of the remaining parts of the on-base network. The network must be proof against the destruction of one or more of its nodes, and must continue to supply the intended communications support to all surviving nodes. Such a stark view can of course be moderated. If the people from a facility that has been destroyed or severely damaged survive and require continuing communications support, some form of *park-and-drive mobile nodes* can be provided for reconstitution of the full-up network.

Imagine a communications architecture in the image of the ARPANET (or any of its successors and offshoots, such as MILNET or MINET), but in which every facility (or group of facilities) that requires communication support is a node on the network. A base might have many tens or even hundreds of nodes. If the per-node equipment were inexpensive enough and/or small enough, even buildings close to one another could be separate nodes, rather than sharing a common node and being locally linked by cable.

For the moment, assume that every node communicates with one or more others and receives from one or more others. A scheme of regular broadcasting by each node will inform all other nodes of its presence; and each can keep its own table of current membership. The network can be made self-organizing in the sense that no central network control will be required. If a node vanishes, the absence of its regularly expected announcement will permit the balance of the network to respond accordingly and update the membership. Similarly, each node can keep

track of those from which it expects to receive messages and those to which it can send. Each node, of course, would be connected to whatever traffic sources might exist--data, secure voice, or facsimile if the bandwidth were sufficient.

In the extreme, even a highly redundant, multiply connected network such as this arrangement is visualized to be can be severed if enough nodes are damaged or if enough communication links are lost.

There are many options for communications media. For the bare base onto which forces deploy, field fiber-optic cable has the advantage of minimizing interception of traffic. It is lightweight and can be restrung if damage occurs, if new links are required to bring new members into the network, or to increase the network redundancy because of traffic volume or anticipated damage.

Point-to-point laser or millimeter-wave communications are other options to facilitate reconstitution after damage, for park-and-drive terminals, or even to support mobile cadres (e.g., ground support crews that might relocate from one aircraft shelter to another for various reasons). Obviously, the bandwidth of either technology is far in excess of what on-base traffic would normally require; but it can be made to pay off in a different way.

Any communication scheme for a combat base must function in spite of dust, debris, aerosols, smoke, and other airborne consequences of wartime operations. It is likely that the line-of-sight path between, say, a laser receiver and the transmitter will not be obscured all the time; hence, the high-data-rate capability can be exploited to burst-transmit traffic when the path is open. Clearly, the observation about obscuration or absorption in the path is intuitive; experiments will be required to acquire data on the percentage availability of transmission paths under realistic battle conditions. Such data will enable one to estimate the average data rate that can be sustained over the link.⁷⁹

⁷⁹In the 1960s, an experiment was performed by the Air Force Cambridge Research Laboratories, then at Hanscom AFB. A high-frequency link was established between a site in Texas and one in California, but the frequency chosen was one with much interference. A pilot tone arrangement was used to detect when the path was clear and open, and the data were then transmitted at a high rate. It proved possible to sustain a teleprinter circuit in spite of the poor quality of the circuit. Similar techniques are reportedly in use on long-distance, high-frequency ship-to-shore circuits.

This approach means that each site wishing to communicate with another must transmit some form of pilot signal whose receipt is evidence of path availability. At that time, the transmitting site operates at its high data rate until the path again closes or the quality becomes too degraded. It is obvious that each site would need substantial buffering to hold traffic until the path opens; and it is equally clear that extensive error-correcting/detecting techniques would have to be incorporated into the traffic to help offset path uncertainties. Sophisticated error-handling is another way that the excess bandwidth of laser or millimeter-wave channels and the capability of modern electronics could be wisely exploited.

Other features could obviously be incorporated into the basic concept of a fully distributed packet-switched network combined with wideband but low average-data-rate channels--for example, embedded encryption for all traffic "on the air"; monitoring of path quality; advisories to other nodes about the maximum traffic level that can be accommodated; reporting of local status, as evidenced by the users' traffic activity, to a central location for situation evaluation and assessment; division of network capacity among non-interacting communities of interest (e.g., operations and intelligence); and allocation of network capacity on the basis of some priority scheme as service degrades with the loss of nodes.

In short, such a fully distributed network can be made *smart at the system level* as well as highly adaptive to conditions, responsive to and tolerant of damage, flexible to shift in loads and user population, and universally applicable in the sense that it can be used for small bases or large bases or--if the antennae associated with the transmission technique are made to track one another--possibly for some field deployments of forces (using relay sites on high terrain to avoid local terrain masking of transmissions).

Such are the applications that the blending of computer and communications technology as implemented in contemporary electronic systems permits one to contemplate.

X. CONCLUSIONS

There are many technologies competing for a place in future Air Force base communications activity: coaxial cable LANs, both baseband and broadband; fiber-optic and other broadband technologies; integrated voice/data switching; and digital radio. It is neither necessary nor desirable to decide that one of them is uniquely appropriate for base communications in the future. They can be used in combinations to provide the office automation facilities desired by various communities (e.g., LANs supporting networks of word processors), to use existing cable plants while still accommodating growth in data traffic (e.g., voice/data switching), to provide specialized applications requiring very-high-speed data links (e.g., fiber optics), and to provide survivable backup with paging facilities to allow personnel the freedom to roam from their terminals without loss of important communications.

By keeping LANs small and localized within a building or other community of interest, versus putting the base on one huge LAN, the risk of total service loss to an entire base is minimized; such an approach is a hedge against combat damage and another approach to survivability. Much of the commercial equipment now on the market can be used, and equipment that has already been procured can be integrated into an overall base communications system through a voice/data switch.

By applying today's LAN techniques as necessary for individual communities of interest and linking them together and to the outside world through a distributed voice/data switch, the goals of cost-efficient modernization and improved survivability can be realized without technological overkill, and without needless imposition of new information technologies on those segments of the user community that do not require it or are not ready for it. As suggested by the discussion of the many technical issues, however, inappropriate choice of a technology or approach can deter an orderly and graceful expansion of service or can directly impair some future option for new services. For example, any decision to install a LAN must include consideration of the technical points raised in this Note, lest an important hedge against presently unknowable future requirements be inadvertently foreclosed.

Base communications planners are faced with the need to accommodate growing data traffic and also the need to modernize obsolete telephone voice switches. In addition, bases facing a conventional combat threat must be modernized in a way that will allow communications to degrade slowly under damage. One of the central design decisions for all bases is how closely coupled the various communities of interest on a base must be in terms of communications. If the answer is "very closely" (i.e., broad bandwidth channels instantaneously available to all communities of interest), then a mutually acceptable overall system must be found, and incompatibilities among existing systems must be eliminated. If the various communities do not need to be closely coupled (i.e., some connection through lower-speed gateways is adequate), infrequent or short messages will dominate. No community will be completely isolated from the others, but the necessity of finding a single system meeting a diversity of requirements will vanish; in many cases, existing systems will be able to be linked through a modernized base telecommunications center, ideally including a voice/digital switch that can allow on-base internetwork access. In many cases, the existing twisted-pair cable plant will be sufficient for allowing remote access to the networks from outlying points on the base, eliminating long--and potentially very expensive--wideband cable runs for a few remotely located users.

In a conventional threat environment, failures should be isolated as much as possible. A short-circuit induced by an explosion or fire can eliminate communication on a cable-based broadband or baseband LAN. However, if a base is configured with a number of separate LANs linked through a switch, a failure will be isolated to the LAN serving its community of interest.

It is important that any future base communications network be able to access DoD and commercial networks external to the base. For DoD internetting, the gateway must support the TCP/IP protocols (see footnotes 15 and 16). While these need not be supported in the on-base networks but only by the gateway, there are nonetheless inter-LAN protocol issues that must be addressed.

Ease of transition from the present situation to an improved one and expandability of the base communications system are both important. In some instances, small LANs have already been installed by tenant organizations. Given the choice of establishing them as de facto standards for the entire base or accommodating them by means of flexible internetworking design, the latter is clearly preferable. Ease of expansion as traffic grows will be an important issue as well. Separation of the communications network into subnetworks serving the various communities of interest will allow traffic in one community to grow without interfering with the performance afforded users outside of the growing one. The problem is more serious with baseband networks which have lower total capacity, but it is a consideration in broadband networks as well, especially if extensive video or teleconferencing applications are envisioned.

A general approach to communications modernization as outlined herein allows the selection of networks for communities of interest appropriate to their traffic-mix/bandwidth requirements. Use of a redundant voice/data switch in combat areas not only provides a topology that is more resistant to failure than a single integrated bus network, but also implements network services as appropriate to meet the requirements of distinct communities of interest without imposing their individual standards on the entire base.

A hybrid approach as outlined in Sec. IX provides a cost-effective and flexible avenue to base communications modernization.⁸⁰

⁸⁰"Network Solution: Panacea or Placebo?" *COMPUTERWORLD Special Report*, February 27, 1984, pp. SR/1-SR/56.

XI. RECOMMENDATIONS

It is clear from the discussion in this Note that the base communications environment is a very complex one, not only from a technical point of view but also because many organizations throughout the Air Force are involved in the decisions to improve the situation. While it is clear that a series of recommendations could be framed that would address each of the major points in turn, the intricacy of the issue suggests that it be handled quite differently.

The USAF/SI has under way the writing of a coordinated group of architectural papers that will define many of the technical aspects and some of the policy aspects of future base communications (see footnotes 64 and 66, and the discussion of ULANA on page 110 ff). However, there is no single paper of modest length that delineates the general context or provides a general overview of what actions the Air Force intends to take or how it intends to implement them, or even of what goals it seeks for base communications. There is no one document that can provide a decisionmaker, especially one involved with advocacy and funding decisions, a comprehensive framework in which to fit various proposed actions or on which to base funding decisions. There is not, for base communications, the analog of the force structure statements that the Air Force uses to guide weapon system matters.

Rather than a number of specialized recommendations, it seems more appropriate to make just two: the first for creation of a declaratory policy that states general Air Force intentions with respect to base communications, and the second for the creation of a plan to implement such a policy.

DECLARATORY POLICY

It is recommended that

The Air Force write and widely disseminate a policy statement that sets forth the general context, assumptions, and guidelines which will be used to direct the improvement of base communications over the next decade.

Such a policy statement should be short (probably less than ten pages) so that all decisionmakers concerned with base communication matters can readily read and grasp its content and so that funding decisions can be easily fitted into a cohesive framework of actions. It should include general assumptions that have been made, indicate the general direction of improvements and progress, address relevant technological details in a generalized manner, speak to the survivability of enduring base communications, etc. In brief, such a declaratory policy should tell all parts of the service itself, and the world in which the Air Force must advocate, what the Air Force intends to do about base communications. Looked at another way, such a policy statement of intent is a concise characterization of the general communications architecture that will eventuate on base supported by the context and assumptions related to achieving it.

There is bound to be some repetition of subject matter (although the depth of treatment will vary) between a policy statement and a plan to implement it. Thus, some of these suggestions for inclusion in the policy will appear also in the discussion of plan components (see below). The discussion throughout this Note can also provide ideas for items to be included in the policy.

- Improvements must be evolutionary.
- Commercially available equipment will be used to the maximum extent.
- A combination of technologies will be used, e.g., LANs, digital telephone switches, cable plants, wideband media.
- There will be electrical connectivity from on-base to off-base communications networks.
- A maximum level of automation will be installed so that communications affairs will become labor detensive.
- Connectivity between any subscriber on one base to any subscriber on any other base worldwide will be the ultimate goal.

- The on-base communications plant will be able to support any and all kinds of information services, e.g., voice, data, secure voice and secure data, video, electronic mail, terminal-to-computer data, computer-to-computer traffic.

Since the end product is a policy declaration, USAF/SI is clearly the focal point for the action, drawing on other parts of the Air Force as appropriate.

IMPLEMENTATION PLAN

It is further recommended that

The Air Force take such action as needed to (a) create a comprehensive plan that is consistent with and implements the declaratory policy, (b) coordinate the plan across all concerned organizational elements, and (c) establish such means as necessary to enforce its guidance and features over decisions concerning base communication matters.

Such a plan will need to deal separately with CONUS and combat bases, but it should specify a minimum communications capacity and functionality for each type of base.

There are obvious parties of interest in the formulation of such a plan: AFCC; AFSC/ESD (because of AFLANSPO and other communication interests); MAJCOMs (because they will be guided by the plan in upgrading their own bases, doing their individual base planning, or laying requirements for action on AFCC and others); and the Air Staff (because of its role in Air Force-wide coordination, in advocacy of programs, budgets, and actions, and in promulgation of guidance, standards, and regulations).

Some organization must take the lead in responding to this recommendation. In principle, it could be Hq/AFCC; but given the emphasis which the Air Force has recently put on *information systems* as a functional element, it seems more appropriate for the Assistant Chief of Staff/Information Systems (USAF/SI) to accept the role. Obviously,

close coordination with AFCC is essential, since AFCC is now the owning command for all information systems personnel, base-level computing and communication assets, and specialized facilities such as those at Gunter Air Force Station and Tinker Air Force Base. Perhaps a joint effort is desirable, or even a tripartite effort with ESD/AFLANSPO (because of its technical insights into the issues, particularly for LANs), but details can be left for Air Force resolution.

The Air Force is in a position to effect such a servicewide planning effort because its recently created Information Systems community reaches into all the necessary organizations.

PLAN COMPONENTS

- The plan must start with the acknowledgment that many communication technologies will coexist on an airbase for at least the next decade. The base cable plant will be there; digital telephone switches will be present; LANs will abound, both baseband and broadband; high-speed dedicated data links will be in place; there may be video services. Each such technology has its unique contribution to make to the mission of the base and its tenants, but none will exist exclusively of the others.
- The plan must mandate evolutionary progress, to ensure that base communications improvement plus the information services that it will support move smoothly "from here to there" without disrupting base performance.
- The plan must appreciate that the bulk of on-base communication assets (at least for CONUS bases) will be commercial products, not equipment specially designed for the Air Force. Hence there is a surfeit of products from which to choose, and the plan must provide guidance to ensure that individual procurements will fit into a technically cohesive and functionally integrated overall plan.
- The plan must also provide for specially designed equipment that may be needed to meet the special requirements of combat bases or of mobility forces.

- The plan must address the present proliferation of LANs and their many technical details. Among the latter are protocols, addressability, inter-LAN interfaces, LAN/switched network interfaces, and LAN/off-base network interfaces.
- The plan must provide a mechanism for bringing the LAN situation under control and for achieving the essential interoperability among LANs themselves and with other networks; and it must state how such a goal is to be achieved. The latter could take the form of using the ULANA specification as a standard to which all LANs must be procured; or the ULANA thrust could be the basis for a competitive "fly off" such as used by the National Security Agency to select its low-cost secure-voice terminal.⁸¹ On the other hand, it may be sufficient to simply inform vendors of commercial LAN systems that the Air Force intends to acquire only products with certain specified technical and functional characteristics. Many vendors might be willing to adapt their product line to be compatible with Air Force needs; or vendors might be willing to market Air Force versions of their products, like, for example, the TEMPESTed versions of commercial products.

There are several possibilities to be explored, but the essential action is that the Air Force announce its intention to procure only LANs with the desired interoperability and other technical features. In addition, there are many LANs now in place, and the interoperability issue must be addressed for them also. Whether the application involves retrofitting extant installations with new products, the development of specialized interfaces, or some other approach, the matter must be dealt with. An AFCC "LAN Management Office" is needed.

⁸¹In a "fly-off," a group of contractors competitively conduct parallel design efforts with the understanding that some of them will subsequently be chosen for development of prototype equipments. (See pp. 84ff relative to the STU-III, especially footnote 46.)

- The plan must also provide guidance for the replacement of obsolete LANs or those that require costly maintenance support. This guidance should include standard costs for maintenance and estimates of amortization periods when replacement is considered.
- The plan must provide for the ongoing management of the base cable plant as one of the pivotal entities in the base communications environment. "Ongoing management" includes upgrading cables with present or new technology (such as fiber optics); planning for orderly expansion as required; consideration of survivability issues as may be pertinent to each base, including the survivability standards set forth in Air Force Regulation 80-38 concerning management of Air Force systems operating under conditions requiring protection from EMP and other nuclear effects; and planning for orderly exploitation of in-place and new cables (such as upgrading to higher-speed links--for example, T-1 service). In short, the plan must provide for oversight of the ongoing growth, technical upgrading, and exploitation of the "cable plant" as a communications resource and as a cohesive entity, not as simply a collection of wires that happen to go from one place to another. One approach would be an AFCC *cable-plant management office*.
- The plan should provide for the development of base-level planning aids and a planning process, so that individual bases would have a solid basis on which to forecast communications requirements and growth. In this connection, data would have to be acquired on representative bases of each MAJCOM so that information flows--voice, data, video--could be estimated and projected; and such data would have to be related to the management of the cable plant and to other communication technologies.
- The plan must also develop an awareness of communities of interest on bases so that sensible installations of LANs, decentralized digital telephone switches, and special local-

loop arrangements could be made; and it must provide access to appropriate databases (such as a comprehensive user requirements database that AFCC should maintain) and extra-AFCC documents (such as those at Hq USAF for military construction, force planning, etc.).⁸² In short, the base planning process for communications must start with firm data about present voice/data/video flows, supported by reasonable projections for each of them several years in advance; and it must be based on a rational process for specifying communication requirements many years ahead in order to accommodate the lead time of acquiring and installing the necessary assets.

It is likely that the same aids would be useful to a Hq AFCC planning office that would oversee and coordinate the base-level efforts. Such aids should be computer-based wherever possible.

- The plan should establish a planning cycle with a specified frequency. It should require the update of current plans whenever new technologies are available or whenever major new requirements are developed.
- The plan must appreciate that there are genuine differences in communications needs between CONUS bases and those in combat theaters. Such differences must be reflected in planning, particularly for survivability and mobility.
- The plan should observe that R&D initiatives on behalf of base communications may be required, and should provide a mechanism for getting such work done. These initiatives may relate to circumstances on combat bases, to new architectures for the far-


⁸²In this connection, The MITRE Corporation (Bedford), as part of its ongoing study of base communications for ESD/XPR, conducted a survey of five bases to acquire specific information on information flows. For its own internal use, a scheme was developed that related such flows to actual geographical locations of buildings on a base, to communities of interest that interact intensively with one another, and to projections for the placement of cables and other communication links. This is illustrative of a planning aid that could be sponsored and developed by AFCC and provided to base-level planners.

distant future, or to specially modified commercial equipments to meet Air Force needs (for example, TEMPESTed and trusted BIUs). Presumably, ESD would be the proper channel to accomplish such work in response to requirements originating via the USAF/SI and AFCC general planning process.

- The plan should provide a mechanism to assure awareness of new technical opportunities in on-base communication architectures, and it should provide a means for conducting appropriate demonstrations of new technology and/or applications. For example, the economics of the "all-broadband base" may prove to be such that all existing communications might be replaced with this new technology on a single base as a demonstration effort; or the demand for wideband services might expand so rapidly that particular bases would promptly need such a technological upgrade.

Suggestions for demonstration bases or situations include (1) construction of a new base, (2) use of a smaller currently operational base (to minimize the magnitude of the switchover effort), (3) use of a combat base where survivability and reconstitutability will be important, (4) use of a colocated operating base in Europe⁸³ to which deploying forces might have to bring special communication arrangements, or (5) use of bases in foreign countries onto which mobility forces might deploy in time of emergency, bringing special communication arrangements with them.

⁸³The totality of all such bases constitutes the so-called COB system of USAFE; they are NATO host-nation bases onto which forces from the CONUS are to deploy in time of emergency and which will subsequently function as NATO forces.



END
FILMED

5-86

DTIC