

12

Semiannual Technical Summary

ADA132284

Wideband Integrated
Voice/Data Technology


31 March 1983

Prepared for the Defense Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-80-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Approved for public release; distribution unlimited.

DTIC
ELECTE
SEP 9 1983

83 09 08 007

D

DTIC FILE COPY

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Defense Advanced Research Projects Agency under Air Force Contract F19628-80-C-0002 (ARPA Order 3673).


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

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

The Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

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

FOR THE COMMANDER



Thomas J. Alpert, Major, USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

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

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

WIDEBAND INTEGRATED VOICE/DATA TECHNOLOGY

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

1 OCTOBER 1982 — 31 MARCH 1983

ISSUED 18 JULY 1983



Approved for public release; distribution unlimited.

LEXINGTON

MASSACHUSETTS

ABSTRACT

This report describes work performed on the Wideband Integrated Voice/Data Technology Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 October 1982 through 31 March 1983.

TABLE OF CONTENTS

Abstract	iii
Introduction and Summary	vii
I. COMPACT STAND-ALONE VOCODER DEVELOPMENT	1
A. Introduction and Status Summary	1
B. Speech-Processing Peripheral (SPP) Architecture	3
C. Vocoder/ Host Interface	5
D. Technology Transfer	6
II. PACKET VOICE TERMINAL AND LEXNET	9
A. Hardware	9
B. PVT Technology Transfer	10
C. Software	10
D. Traffic Emulation and Measurement	12
III. MINICONCENTRATOR GATEWAY	13
A. Group and Stream Synchronization	13
B. IP/ST Encapsulation	13
C. Histograms for Multiplexing Experiments	14
D. Speech Packet Discard Mechanism	14
E. Gateway-to-Gateway File Transfer Protocol	14
F. Terminal Control of Multiple Gateways	15
G. Other Gateway Developments	15
IV. WIDEBAND NETWORK EXPERIMENTS AND EXPERIMENT COORDINATION	17
A. System Throughput Calculations	17
B. System Operation and Experiment Coordination	24
C. Multiplexing Experiments with Emulated Voice Traffic	27
D. Packet Video Facility	33
V. VOICE CONTROL OF NETWORK VOICE CONFERENCING	35
Glossary	39

LIST OF ILLUSTRATIONS

Figure No.		Page
1	SPP Architecture	2
2	SPP Chassis	3
3	SPP Circuit Board	4
4	Two Sites, N_v Voice Slots. (Relative Sizes of Portions of Burst Not Drawn to Scale.)	18
5	N_s Sites, One Call Each. (Relative Sizes of Portions of Burst Not Drawn to Scale.)	20
6	Sample Experiment Scenarios Supported with Case 3 (772 kbps) WB SATNET Operation	21
7	Experiment Scenario Supported at Case 2 WB SATNET Bit Rates. Site 2 Carries Both LPC and PCM Traffic	22
8	Experiment Scenario Using Case 1 WB SATNET Rates	22
9	Alternative Scenario Using Case 1 Rates	23
10	Multiplexing Experiment Configuration	28
11	Delay Histogram for Experiment I	29
12	Delay Histogram for Experiment II	30
13	Delay Histogram for Experiment III	32
14	Packet Video Facility	34
15	VCOP Block Diagram	36

INTRODUCTION AND SUMMARY

An important challenge in the design of future military communications networks is to achieve overall system economy and adaptability through efficient and flexible allocation of common network resources to voice and data users. The major objective of the Wideband Integrated Voice/Data Technology Program is to address this challenge through the development of techniques for integrated voice and data communications in digital packet networks which include wideband common user satellite links. A major focus of activity in the program has been the establishment of an experimental wideband packet satellite network for realistic testing of a variety of strategies for efficient multiplexing of voice and data users. The program also serves as a focus for the development and testing of techniques for local area packet voice distribution, for speech traffic concentration, and for efficient real-time voice communication in an inter-network environment including local networks of various types connected through a wideband demand-assigned satellite network.

Through FY 82, the Packet Speech Systems Technology Program included efforts in Wideband Integrated Voice/Data Technology and Packet Speech as two subprograms. Original efforts in the Packet Speech Program led to demonstrations of conversational speech on the ARPANET and on the Atlantic Packet Satellite Network. More recent Packet Speech efforts focused on the development of compact narrowband and multirate voice processors and associated Packet Voice Terminals. The Packet Speech Program came to a close at the end of FY 82 having achieved its basic goals in packet speech techniques development, and in voice algorithm and terminal developments to support experiments in the wideband network.

This report covers work in the following areas: development of a generalized compact vocoder based on digital LSI technology; development of Packet Voice Terminal (PVT) and local access network (LEXNET) facilities for experiments in packet voice; development and experimental application of a flexible internet stream gateway (the miniconcentrator) for packet voice; coordination and execution of multi-user packet speech experiments on the experimental wideband satellite network (WB SATNET); and development and experimental testing of a Voice-Controlled Operator (VCOP).

A generalized stand-alone version of the compact Linear Predictive Coding (LPC) vocoder has been developed. This Speech Processing Peripheral (SPP) unit interfaces to a host computer through an asynchronous serial link and includes flexible software interface protocols to accommodate vocoding, speech synthesis, and speech recognition parameter extraction tasks. Two SPPs have been built and tested both at Lincoln and at Information Sciences Institute (ISI), and one unit has been used for speech synthesis in VCOP experiments. A technology transfer effort is under way for replicating the SPP for general use in the DARPA community.

A number of improvements, including capability for program control of vocoder selection, have been introduced in the PVT hardware and software. Otherwise, the PVTs and LEXNETs have remained a stable and operational experimental facility during this reporting period. The traffic emulation capability of the PVT-based measurement host (MH) has been enhanced by the addition of a stream (ST) capability to emulate traffic from multiple voice calls. This capability has been used in a set of packet speech multiplexing experiments.

The miniconcentrator gateways at Lincoln, ISI, SRI International, and Defense Communications Engineering Center (DCEC) have continued to operate reliably during this reporting period. A number of extensions have been added to the software to enhance the experimental capability of the overall system. These include: group and stream synchronization, encapsulation of ST messages in IP messages for communication with gateways not supporting ST, histograms of gateway resource utilization, a packet discard mechanism for speech multiplexing experiments, and a gateway-to-gateway file transfer program. The gateway has been used to generate test traffic for the PSAT and to log PSAT up/down status.

Experiment coordination efforts have continued to focus on the dual WB SATNET system goals of achieving regular operational status, and of increasing the operational channel bit rate to 3 Mbps. System throughput calculations have been carried out to identify experiment scenarios which can be supported at a variety of channel burst and code rates. Despite improvements in a system robustness features and channel stability and calibration, a number of system problems remain; a short-term task force effort, led by Lincoln, has been initiated to address these problems.

A series of multiplexing experiments has been initiated using emulated voice traffic generated by PVTs operating as MHs. A configuration using two gateways and two LEXNETs has been used. A sample experiment run showed that 20 full-duplex LPC calls could be accommodated, using speech activity detection, into a stream allocation sized to handle 13 full-duplex voice slots, with a total packet discard rate of about 1 percent in the gateways.

A four-participant conference has been established by voice control over the WB SATNET. A talker at ISI, after training the speech recognizer over WB SATNET, dialed up the VCOP at Lincoln and carried on a voice dialogue to provide a list of conferees and the conference type (i.e., PCM). This information was then passed from VCOP to the conference access controller which automatically dialed up the participants.

WIDEBAND INTEGRATED VOICE/DATA TECHNOLOGY

I. COMPACT STAND-ALONE VOCODER DEVELOPMENT

A. INTRODUCTION AND STATUS SUMMARY

A generalized, stand-alone version of the compact LPC vocoder has been developed. This vocoder employs the same technology (including the NEC μ PD7720 signal-processing interface chips) which was used in the Packet Speech Program to develop the Packet Voice Terminal-compatible LPC. This vocoder differs from the PVT-compatible vocoder primarily in that it interfaces to a host computer through a 9600-bps RS-232-C asynchronous serial link and is completely self-standing. A powerful, flexible I/O link management protocol has been developed and implemented in the SPP control firmware which is sufficiently general to support compatibility with a great variety of potential host operating systems. The software interface protocols have been designed so that the resulting vocoder will be flexible enough for speech recognition parameter extraction and speech-synthesis tasks as well as traditional vocoder analysis-synthesis applications. Potential applications for this unit, referred to as the Speech-Processing Peripheral, include speech bandwidth compression, speech playback, and speech-recognition front-end processing. Because of its simplicity, compactness, and potential for comparatively low-cost mass replication, it is projected to provide substantial speech-processing power as a standard, inexpensive adjunct to personal computing workstations or general computer facilities. The unit can also be used exclusive of a host computer as a stand-alone digital voice communication instrument.

Two units of the SPP have been built and tested successfully in stand-alone and host-controlled mode, both at Lincoln and at ISI. At Lincoln, a UNIX PDP-11/44 was used for testing; at ISI, the unit was interfaced successfully to a Switched Telephone Network Interface (STNI) card on a PVT. One of the units has been used for speech synthesis (see Sec. V) in the Voice-Controlled Operator (VCOP) experiments. A Project Report* has been prepared which includes an overview of

*M.L. Malpass and J.A. Feldman, "A User's Guide for the Lincoln LPC Speech Processing Peripheral," Project Report PSST-2, Lincoln Laboratory, M.I.T. (29 April 1983).




Fig. 1. SPP architecture.

the SPP architecture and detailed operational instructions of the link management software protocol and hardware programmable options for both host-based and stand-alone applications. A technology transfer effort is under way for replicating the SPP. Summary descriptions of the system architecture, the vocoder/host interface, and the technology transfer effort are included below.

B. SPEECH-PROCESSING PERIPHERAL (SPP) ARCHITECTURE

The SPP architecture, which can be viewed as a flexible LPC processor, is shown in Fig. 1. Photographs of the SPP chassis and circuit board are shown in Figs. 2 and 3. The circuitry is mounted on a single 7- $\times$ 7-in. wirewrap panel containing about 33 integrated circuits plus miscellaneous components.




Fig. 2. SPP chassis.

The LPC signal-processing computations are carried out in a triplet of dedicated NEC μ PD7720 digital-signal-processing microcomputers which have been factory programmed with Lincoln-developed LPC firmware. It is possible to conveniently vary such parameters as LPC model order, frame duration, sampling rate, pre/de-emphasis time constant, etc. The basic functions of the LPC analysis, synthesis and pitch/voicing determination are each carried out in a separate '7720.




Fig. 3. SPP circuit board.

Analog processing and A/D and D/A conversion are implemented using a μ -255 law CODEC-with-filters chip. Provisions are made for analog pre- and de-emphasis as a complement to the digital option included in the '7720s. Since the anti-aliasing filters are realized using the sampled data/analog (switched capacitor) PCM filters on the CODEC chip, audio bandwidth can be controlled by adjusting appropriate clocks in the timing chain and filter characteristics are well controlled.

Initialization and coordination of the '7720s are carried out using an 8085A-2 microcomputer complex which also has the general responsibility of managing and controlling both the serial (RS-232-C) and parallel (TTL) I/O ports. Although physical hardware for an 8-bit, TTL-compatible parallel port has been included in the current SPP design, no corresponding control firmware has yet been developed. Because of the relatively large memory capacity that has been provided (6K bytes ROM, 10.25K bytes RAM), it is possible to implement an elaborate and powerful link handler and SPP control protocol. Through this protocol the host operating system can command the SPP in a variety of ways. Voice messages, voice control, and speech-recognition front-end processing are typical applications which are implementationally straightforward using this control structure.

C. VOCODER/HOST INTERFACE

The vocoder/host interface which communicates at the rate of 9600 bps over an RS-232 asynchronous link to a host computer provides a means for the host to control the parameters of the LPC vocoder and to interact with it via a special protocol. To prevent confusion between data and control commands, data are packed 6 bits to a byte and then augmented to fall into the range 40 through 137 octal. The range 0 through 37 octal is never used except for the standard XON and XOFF characters. The range 140 through 176 octal is reserved for the special vocoder/host protocol character assignments.

As described above, the 8085 manages the initial programming of the '7720s which specifies such things as the LPC order, filter gain parameters, and the frame interrupt timer. It is the frame interrupt to the 8085 that in turn triggers the requests for the exchange of raw parameters with the '7720s. The SPP/host interface has been designed to provide the means for a host computer to initialize the '7720s and interact with them via a special protocol. The LPC analyzer and synthesizer

may be programmed independently except for the number of samples per frame, which is the same for both. The frame timing provided by the vocoder is a function of (1) the number of samples per frame specified by the host, and (2) the sampling interval determined by an external oscillator on the SPP board. Since the '7720s exchange uncoded parameters with the 8085, the 8085 is also responsible for coding and decoding and these tables may be specified by the host. Finally, the analyzer and synthesizer may be activated in independent modes. These modes are of two types: time-locked and non-time-locked. Time-locked mode requires a host computer that can maintain real time on a frame-by-frame basis; in non-time-locked mode 1K ring buffers are used by the interface to accommodate a host that can maintain real time over the long run but not on a frame-by-frame basis.

D. TECHNOLOGY TRANSFER

A solicitation-for-bid (RFI) has been prepared for replicating the Lincoln Speech-Processing Peripheral (SPP). The major element of the RFI is a detailed technical specification. The procurement package requires prospective production contractors to indicate in detail (1) the fabrication approach, (2) any alterations to the Lincoln-supplied design which would improve functionality, manufacturability, testability, or reliability (e.g., printed circuit vs wirewrap), (3) facilities available, (4) technical qualifications of engineering personnel, and (5) a proposed production schedule. A bidder list has been formulated and the RFI was released late in March.

As currently envisioned, the replication effort will proceed in two phases, the projected costs of each to be quoted separately. During the first phase a pre-production prototype will be designed, formally documented, and fabricated. Lincoln will assist in debugging and qualification of the prototype in order to validate the design and familiarize contractor personnel with details of the SPP's operation. Once the prototype design is certified by Lincoln, the second (production) phase will be initiated. The size of the production run will be determined by Lincoln at contract time, having taken into consideration available funds, per-unit costs, and supplies of critical components (e.g., NEC μ PD7720s). Vendors are required to quote unit costs in quantities of 1 to 19, 20 to 49, and over 50.

In order to maximize the potential yield of the procurement, Lincoln has taken positive steps to minimize vendor costs associated with labor, equipment capitalization, and risk. These include supplying critical components and Lincoln-fabricated SSP engineering models on a GFE basis, supplying Lincoln-developed test procedures for checkout of production units, and providing Lincoln staff-level technical assistance when required. Final acceptance testing will be conducted at Lincoln by in-house personnel using Lincoln-developed qualification procedures.

II. PACKET VOICE TERMINAL AND LEXNET

A. HARDWARE

A number of improvements were made in the PVT capabilities during the first half of this reporting period, as described below. More recently, the PVT and LEXNET hardware configuration has remained stable. All PVTs at Lincoln have had all updates installed and tested. The change descriptions and necessary parts have been sent to ISI to allow them to bring their terminals up to date.

1. Vocoder Selection

The protocol processor and PCM cards have been modified to allow program control of the vocoder selection. A two-position momentary switch on the protocol processor is tested by the program to determine the user's preference for either the external vocoder or the internal PCM capability. The actual selection is made by the program depending on the protocol negotiation during call setup. Currently, if the calling and called PVTs show different user rate preferences, and both rates are supported by both PVTs, the calling party's preference will dominate. In any case, the protocol negotiation will seek to establish communication at a compatible rate. The actual choice of encoder resulting from the negotiation is indicated by lights on the PCM card.

2. ROM/RAM Card

The ROM memory extension card has been modified so that it can accept the pin-compatible 4802 RAM chip or the 2716 EPROM chip. This allows the card to be used as a RAM for software development and then converted to ROM for operational usage. The conversion requires the change of the memory chips and one jumper package.

3. Speech Activity

The speech activity detector control switch on the PCM card has been replaced with a three-position switch. The new switch allows the user to turn off all speech transmission as well as the previous two options of transmit-during-speech and transmit continuously.

B. PVT TECHNOLOGY TRANSFER

A Request for Quotation (RFQ) specification for replication of the LEXNET PVT was written and submitted to DARPA for review. The RFQ asked for quotations on the production engineering required to produce PVTs and LEXNET/Concentrator Interfaces, and for production runs of 10 to 50 units. A detailed technical information package on the PVT and LEXNET was also included.

Based on discussions with DARPA, a decision was made not to proceed with the LEXNET PVT procurement at this time, but instead to proceed with an alternate plan based on interfacing the stand-alone LPC SPP to an off-the-shelf workstation (probably a SUN unit). The new plan includes the ongoing technology transfer of the LPC SPP, as described in Sec. I.

In addition, the following future activities are projected. The LPC SPP will be interfaced to a workstation, which in turn will include a standard ETHERNET interface. NVP/ST protocol software will be implemented to allow SPP units to communicate real-time packet voice on the ETHERNET. Lincoln will develop an ETHERNET interface which can replace the buffer control and modem cards of either the LEXNET/Concentrator Interface or the LEXNET PVT. This will allow interoperation of the workstation-based LPC units with units currently in operation on the wideband net. When this interoperation has been demonstrated, the issue of PVT procurement will be reconsidered, based on an ETHERNET-compatible PVT.

C. SOFTWARE

PVT software efforts during this period have included: (1) testing and checkout of the conferencing robustness features; (2) completion and debugging of the PVT software to support the Voice-Controlled Operator (VCOP); (3) extensions to the dial-up protocol to support the new hardware feature of automatic vocoder selection; and (4) reorganization and further modularization of the software to allow optional inclusion of features required for various applications. The first draft of a comprehensive technical report "Protocol Software for a Packet Voice Terminal" has been completed. The report is now being reviewed and edited.

Robustness features for point-to-point speech and conferencing have undergone thorough checkout locally and through gateways. In addition, a number of

conference tests have been run with ISI using the satellite, and the time-out and retransmission mechanisms for conferencing robustness have been validated in this environment. The PVTs are normally left up and running and have been quite reliable. For example, the Access Controller (AC) code was last reloaded in mid-October. We expect to continue regular testing of PVT communication over the satellite channel.

The Voice-Controlled Operation is working with the PVTs. There is one software module for the PVTs which contains all available options including VCOP. To set up a conference, a PVT places a call to VCOP (a PVT connected to the speech synthesizer and recognition system). The caller tells VCOP by voice the type of conference to set up and the names of the desired participants. The VCOP PVT then issues "PLEASE JOIN A CONFERENCE" messages to each participant. These messages are included in the time-out and retransmission mechanism so that messages lost in transit will be resent. The VCOP PVT software has been tested extensively in local configurations using one and two LEXNETs, and two conferences have been set up over the wideband satellite network (WB SATNET) from ISI.

Changes in the hardware now allow the software to control whether the PCM vocoder or the external vocoder is in use. When a PVT receives a request to talk point-to-point or a request to join a conference it will automatically change from PCM to external vocoder (or vice versa) if that will allow it to match the vocoder specified by the incoming request. Now the only requests rejected are those which specify a vocoder different from the one in the PVT's vocoder slot.

The PVT software has been completely reorganized for increased modularity and efficiency of memory usage. It now consists of six modules written in the high-level language C, one module written in the assembly (machine) language A-Natural, and sixteen library routines provided by the compiler. Two C modules handle IP messages in and out. Two C modules handle ST messages in and out. One C routine is a general control routine and the last C routine handles the interaction with the phone. This segmentation allows us to assess the space requirements of the various elements of the protocols.

The full software package occupies 31,347 bytes of the 32,000 available. The Whitesmiths C compiler allows conditional assembly. The software is set up so that any or all of a number of options can be omitted at compile time with a corresponding decrease in memory usage. Omissible options include ECVSD, LPC, debugging aids, and VCOP. ISI plans to add their Switched Telephone Network Interface (STNI) module to the PVT software. Since their module largely replaces the current phone-handling module it should fit in a PVT without omitting any optional code.

D. TRAFFIC EMULATION AND MEASUREMENT

The traffic emulation capability of the PVT-based measurement host software has been enhanced by the addition of a code to allow the use of ST as well as IP packets for carrying the emulated traffic. At present, the ST capability requires that the ST connections to be used for the traffic be set up prior to the start of the emulation itself. To accomplish the set-up process for a number of connections in a single step, we have modified the ordinary PVT program to cause it to set up a specified number of calls in rapid succession. The measurement host then makes use of these connections for sending its packets.

The new capability is being exercised by setting up a number of emulated narrowband calls through the gateway and observing their effect on other real calls. The emulated calls make use of a Markov model of talkspurt and silent intervals that gives a good approximation to the statistical behavior of a set of real talkers. The parameters of the model can be adjusted by the experimenter to observe the effects of variations in durations and duty cycles. Packet speech multiplexing experiments using this capability are described further in Sec. IV.

III. MINICONCENTRATOR GATEWAY

The miniconcentrator gateways at Lincoln, ISI, SRI, and DCEC have continued to operate reliably during this reporting period. A number of extensions have been added to the software to enhance the experimental capability of the overall system. These include: group and stream synchronization, encapsulation of ST messages in IP messages for communication with gateways not supporting ST, histograms of gateway resource utilization, a packet discard mechanism for speech multiplexing experiments, and a gateway-to-gateway file transfer program. The gateway has been used to generate test traffic for the PSAT and to log PSAT up/down status. Finally, the gateway support software has been extended to allow downline control of multiple gateways from a single user terminal.

A. GROUP AND STREAM SYNCHRONIZATION

The PSATs have been extended to provide a group and stream synchronization mechanism whereby the PSATs collectively attempt to maintain group and stream information across outages of individual PSATs. Group information is particularly resilient, surviving as long as at least one PSAT remains up at all times. The gateway was extended to take advantage of this sync mechanism by ascertaining whether the group and stream information did in fact survive across an observed outage of its PSAT. If any of the information did not survive, the gateway reestablishes the groups and/or streams. In the case of reestablishing a group, the gateway communicates with the other ST gateways telling them to join the group and asking them which groups it should join. Additionally, group information is maintained by each gateway on auxiliary disk or tape for use after a gateway outage. The result of these mechanisms is a dramatic reduction of manual intervention needed to recover across a PSAT or gateway outage.

B. IP/ST ENCAPSULATION

The setting up of ST connections for speech communication and the subsequent transmission of ST data packets may involve an internet path through gateways that do not handle ST protocol. In particular, experiments being planned in cooperation with SRI will involve packet speech connections to sites on the Atlantic SATNET

whose access to the internet is through IP gateways only. To accommodate such situations the miniconcentrator gateway's routing and dispatching software was extended to recognize such paths and consequently encapsulate the ST messages in IP messages. When receiving such an encapsulated message, an ST gateway de-encapsulates it and treats it as an ordinary ST message.

C. HISTOGRAMS FOR MULTIPLEXING EXPERIMENTS

Internal memory resources in the PDP-11 gateway were expanded in order to allow for simultaneous support of as many as 60 point-to-point speech conversations. When many conversations are open and transmitting packets, the setting of the maximum aggregation limit, i.e., the maximum number of packets that are permitted in one ST envelope, can play a major role in performance. To measure the number of packets that are actually aggregated in one envelope the gateway now maintains a histogram of this value. Additionally, the UMC-Z80 program now maintains a histogram of the number of packets it has received from the network in a given interval. These two histograms provide a reasonable measure of the adequacy of the maximum aggregation limit.

D. SPEECH PACKET DISCARD MECHANISM

During peak load situations more packets can arrive than can be dispatched, resulting in an accumulation of packets and clogging of resources. As packets accumulate, the delay before they are dispatched increases; eventually, the delay becomes so great that the packets become "stale" and are no longer worth transmitting because they would arrive too late to be usable by the speech reconstitution algorithm. For these reasons, a mechanism for discarding packets due to age was incorporated in the gateway. Using this strategy, any packets that have not been dispatched within a specifiable amount of time since their arrival are discarded. Tabulating the number of such "lost" packets in conjunction with different aggregation limits provides a measure of throughput in peak load situations.

E. GATEWAY-TO-GATEWAY FILE TRANSFER PROTOCOL

A protocol was developed between ST gateways whereby they are able to transmit files (e.g., programs) to each other via the WB SATNET for storage on auxiliary disks or tapes. A "window" allocation by the receiver together with

sumchecking and retransmission of incorrect messages provides for accurate transmission of a file. This permits the distribution of new gateway software releases to be achieved more easily than through ARPANET FTP to a host computer for subsequent downloading to the gateway computer. Furthermore, in the case of some future gateway computers there will not even be a nearby host computer connected both to the gateway and to ARPANET.

F. TERMINAL CONTROL OF MULTIPLE GATEWAYS

A major extension was produced for the downline communication program TODOWN. Previously, this program was known as TOEPOS, but its usage and flexibility have suggested the more general name, TODOWN. Heretofore, each invocation of TODOWN provided for an "attachment" to one downline processor for communication. If one wished to communicate concurrently with several downline processors, a separate invocation was needed for each processor. This made it difficult to associate outputs from several processors participating in one experiment. Moreover, it is inconvenient to tie up a terminal for each invocation, especially if one is communicating remotely via ARPANET. Now, TODOWN has been extended to provide concurrent attachment to several downline processors. One can thereby monitor and control several gateways simultaneously. The command structure in TODOWN was expanded to allow one to attach and release downline processors, specify type-through to a desired TTY line of a processor, etc.

G. OTHER GATEWAY DEVELOPMENTS

In certain experimental situations we use a loopback plug instead of the real PSAT. In more complex situations we use a "null PSAT" to connect two gateways together. In order to make such experimental environments as similar as possible to the real PSAT environment, the gateway was extended to simulate PSAT behavior when it receives a request to create or modify a stream, join a group, etc. (Such requests would be received by the gateway only when a loopback plug or null PSAT is involved.) In this way, the mainline gateway program is unaware that the real PSAT is not there and is able to proceed with its normal routing and dispatching of packets.

The time-out and retransmission of ST control messages in the gateway has now been exercised in a variety of situations, including usage through the WB SATNET. Several minor problems were found and corrected, with the result that the establishment and closing of ST connections is now quite reliable.

We are currently implementing a mechanism for logging the up/down transitions of the PSAT at a gateway site. This would provide a measure of actual availability of the network for experimentation and other uses. The mechanism is general and would therefore be usable for networks other than the WB SATNET.

IV. WIDEBAND NETWORK EXPERIMENTS AND EXPERIMENT COORDINATION

A. SYSTEM THROUGHPUT CALCULATIONS

Prior to the March 1982 Wideband Meeting, a series of calculations were performed to estimate the system throughput achievable under various assumptions, and to identify the specific factors limiting throughput. A new analysis was done prior to the March 1983 Wideband Meeting, including updated system details and a revised set of experiment scenarios. The current throughput calculations were intended to focus on the limits on experimental capability caused by operation at 772 kbps and to obtain more precise estimates as to what can be done at higher rates for some typical scenarios. To this end, more precise information had been obtained as to detailed frame and overhead structures, and increased interburst padding requirements were taken into account for mixed code rate cases. The calculations also focused on scenarios including more than four sites, since new sites are being added to the network: three by summer (RADC and Forts Monmouth and Huachuca), and three more in the near future (M.I.T., CMU, and LINKABIT), for a total of ten.


Three bit and code rate combinations were chosen for the analysis. For Case 1 the channel burst rate was 3.088 Mbps (1.544 megasymbols/second, QPSK); all the control (or header) bits were assumed to be coded at rate 1/2, and the speech (or data) bits were assumed to be uncoded. For Case 2 the burst rate was 1.544 Mbps (BPSK), with headers at rate 1/2 and speech uncoded. Case 3 assumed 772 kbps for both headers and data, with no coding. In summary, the burst/code rate combinations used in the calculations were:

Case 1: $R_c = 3.088$ Mbps (QPSK), code rate 1/2
 $R_s = 3.088$ Mbps, rate 1

Case 2: $R_c = 1.544$ Mbps (BPSK), rate 1/2
 $R_s = 1.544$ Mbps, rate 1


Case 3: $R_c = 772$ kbps (BPSK), rate 1
 $R_s = 772$ kbps, rate 1

where R_c and R_s are the bit rates for control and speech, respectively.



N_v = NUMBER OF VOICE SLOTS, REGARDLESS OF TASI

$$N_v \leq N_{conv}$$



ONE BURST PER FRAME
(the two sites alternate)

PSAT PACKET LOAD: 47 PACKETS/s (speech traffic)


Fig. 4. Two sites, N_v voice slots. (Relative sizes of portions of burst not drawn to scale.)

For each of these burst/code rate cases, and for assumed voice bit rates of 64 and 2.4 kbps, two different network scenarios were considered. The first scenario was based on the two-site configuration shown in Fig. 4, with N_V active voice slots supporting simultaneous independent conversations between the two sites. The number of conversations N_{CONV} supported in these slots will generally be greater than N_V , and will depend on the Time-Assignment-Speech-Interpolation (TASI) advantage achieved. The objective of the analysis was to find the maximum value of N_V under the various burst-rate and voice bit-rate assumptions. The second scenario used the N_S -site configuration of Fig. 5, shown with each site supporting one call, and the objective was to find the maximum number of sites N_S . The second scenario was also used to find the maximum N_S with two and three calls in progress per site.

The frame structures sketched in Figs. 4 and 5 were determined in detail by consultation with Bolt, Beranek and Newman (BBN), and were computed for each of the scenarios and cases of interest. The maximum achievable values of N_V and N_S were determined under the constraint that the time duration of a PSAT frame is 2^{15} bit durations at a bit rate of 1.544 Mbps. Figures 4 and 5 also indicate approximate PSAT packet load for handling the speech packets in each scenario.

Complete tables of results were presented at the Wideband Meeting, for all the rate and parameter combinations above. As a convenient summary, a few example scenarios were presented; these are repeated in Figs. 6 to 9. Figure 6 indicates the limitations on experimental throughput and flexibility imposed by the Case 3 channel burst rate of 772 kbps, which has been used for previous speech demonstrations. At this rate the network can support the indicated levels of either multi-site PCM traffic or two-site LPC traffic, but not both simultaneously. When the network is dedicated exclusively to one type of experimental traffic or another, the supportable traffic is either 81 LPC calls between two sites or a four-site PCM network with two calls per site.

Figure 7 shows that a moderate improvement in channel throughput, to the Case 2 rates (1.544 Mbps, coded rate 1/2 for headers and uncoded for data) makes a dramatic improvement in network capacity: for example, PCM and LPC experiments can be supported simultaneously at the levels shown. In the figure, site 2 carries both LPC and PCM traffic. Figure 7 also assumes correction of the present




PSAT PACKET LOAD AT SITE i

47 PACKETS/s TO/FROM SITE i , PLUS MODERATE
ADDITIONAL PROCESSING FOR ALL OTHER PACKETS

Fig. 5. N_s sites, one call each. (Relative sizes of portions of burst not drawn to scale.)

- $R_c = R_s = 772/1$
- EXPERIMENT SCENARIO:

4 SITES, TWO 64-kbps
CALLS EACH



OR

2 SITES, 81 2.4-kbps
SLOTS

OR




Fig. 6. Sample experiment scenarios supported with Case 3 (772 kbps) WB SATNET operation.

- $R_c = 1.544/1/2$, $R_s = 1.544/1$
- INTERBURST PADDING = $96 T_u$
- EXPERIMENT SCENARIO:
4 SITES, TWO 64-kbps CALLS EACH
PLUS
2 SITES, 57 2.4-kbps SLOTS




Fig. 7. Experiment scenario supported at Case 2 WB SATNET bit rates. Site 2 carries both LPC and PCM traffic.

- $R_c = 3.088/1/2$, $R_s = 3.088/1$
- INTERBURST PADDING = $96 T_u$
- EXPERIMENT SCENARIO:
4 SITES, THREE 64-kbps CALLS EACH
PLUS
2 SITES, 87 2.4-kbps SLOTS




Fig. 8. Experiment scenario using Case 1 WB SATNET rates.

4 SITES, SIX 64-kbps
CALLS EACH

PLUS

2 SITES, 37 2.4-kbps SLOTS

129298 N




Fig. 9. Alternative scenario using Case 1 rates.

experimentally observed need for 2048 channel symbol times of interburst padding when mixed code rates are used. The interburst padding assumed in Fig. 7 is $96 T_u$, where $T_u = (1.544 \times 10^6)^{-1}$ represents one channel symbol time at 1.544 Mbps BPSK.

Figure 8 shows an example scenario assuming the highest expected channel rates (Case 1, 3.088 Mbps, coded rate 1/2 for headers and uncoded for data). The channel supports a four-site PCM network with three calls each, as well as a two-site LPC experiment session with 87 simultaneous calls. Figure 9 illustrates an alternative arrangement under Case 3: the four-site PCM network has six calls per site, and there is still enough capacity to support 37 simultaneous LPC calls between another pair of sites.

The analysis shows the importance of operation at the Case 1 and Case 2 rates to support substantial multi-user, multi-site experiment scenarios. As an incidental output of the analysis, the calculations indicated that the packets-per-second processing requirements on the PSAT were modest under all scenarios considered, so it could be stated that this processing requirement is not the factor limiting system throughput.

B. SYSTEM OPERATION AND EXPERIMENT COORDINATION

System goals for WB SATNET continue to be the achievement of regular operational status, and of operation at system bit rates higher than 772 kbps. Important milestones have been achieved, including the 3 June 1982 packet speech demonstration at Lincoln (described in the last semiannual*) and a successful demonstration of the packet/circuit interface facilities at DCEC on 7 October. During the past six months, robustness improvements have been made at all levels; channel calibration has improved; and sources of radio frequency interference (RFI) are being identified, as noted below. The satellite channel is more stable and reliable than in the past; the bit-error rate is typically better than 5×10^{-3} .

*Packet Speech Systems Technology Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (30 September 1982), DTIC AD-A126880.

A number of problems must be resolved in order to achieve the required operational status for the system. Efforts are being made to correct these problems, and at the same time to push toward operation of the system at higher rates, as described below.

A major effort over the past six months has been to use the network for experimental purposes as much as possible, to identify and deal with causes of unreliability, and to push toward reliable operation at higher channel bit rates with appropriate code combinations. In particular, efforts were made to gather data and achieve results which could be taken into account in discussion and planning at the Wideband Meeting held at DARPA in March. One aspect of this work consisted of evaluating network operational status on a daily basis, with the cooperation of BBN and personnel at all the sites. Statistics were gathered on the availability of the channel and of the various sites and subsystems, and the results indicated problems with system availability. In particular, the time during which Lincoln and ISI (the two primary packet speech sites) were simultaneously operational has been limited. The availability problems resulted from a combination of factors, rather than being attributable to any single subsystem. The availability statistics were presented in their entirety at the Wideband Meeting, and action was undertaken to improve them, as described below.

The performance of the PSAT, ESI, and satellite channel has been tested at a variety of channel bit rates and error-correcting code rates. Most often testing has been done by BBN personnel under PSAT control, using test data both from PSAT message generators and from facilities on LEXNETs at Lincoln. In conjunction with these tests, Lincoln has conducted speech experiments with the WB SATNET running at various combinations of channel and code rates. The results indicate packet losses and other difficulties at channel bit rates higher than 772 kbps, presumably due to the higher bit-error rates. A specific problem has been observed experimentally with use of mixed code rates within a burst (e.g., rate 1/2 for headers and rate 1 for data), in that long interburst intervals appear to be required to avoid excessive packet loss in the ESI. The nominal interburst spacing is 96 channel symbols, but for mixed code rates this spacing has to be increased at present to about 2048 symbols.

A potential contributor to degraded system performance is nonuniformity as to transmitted power and frequency among the WB SATNET sites. To insure uniformity, an agreement has been reached with Western Union under which Lincoln Laboratory will serve as the power and frequency reference station for the network. A Hewlett-Packard 8566A Spectrum Analyzer has been purchased to support this function; it has a frequency stability of one part in 10^9 , and is accurate to a small fraction of a decibel in power.

Plans are in progress to automate the measurement of uplink power and frequency for every site, so that it can be carried out during low-traffic periods (probably at night) by the personnel at BBN's Network Operations Center. At present, in order to make such a measurement for a particular site it is necessary to have the NOC take all the PSATs off the air, and then to have a person at the site in question switch to the Burst Test Modem and transmit an unmodulated carrier. A set of such measurements on 14 December disclosed that Western Union's uplink power level for the other customers on our transponder had moved out of calibration, causing a substantial reduction in signal-to-noise ratio at the WB SATNET sites. This problem was communicated to Western Union; they corrected the power level problem, and subsequent tests with the spectrum analyzer showed that WB SATNET conditions were back to normal.

Some time ago, a radio frequency interference (RFI) problem was observed to exist at ISI. DARPA made arrangements with Probe Systems, Inc. to conduct RFI investigations at ISI, using a variety of techniques including wideband disk recording and correlation analysis. This work was carried out in late January, and included a number of interactions with Lincoln. Probe's investigations indicated that the RFI is due to radio altimeter transmissions from aircraft flying at about 20,000 ft over the Los Angeles airport. The altimeters operate at about 4.3 GHz, well above the 3.7-GHz WB SATNET downlink frequency, but it is thought that nonlinear effects in the front-end RF amplifier produce enough in-band energy to disrupt the desired signals. A possible solution, currently under investigation, is for Western Union to install front-end filtering to reject the interference.

One outcome of the March Wideband Meeting was recognition that difficulties have been encountered in achieving required operational status and higher bit-rate

operation for the network. In an effort to overcome these problems, a "task force" effort was formed with goals of regular operation first at 1.544 Mbps and then at 3.088 Mbps, in each case with headers coded at rate 1/2 and data uncoded. This task force is being coordinated by Lincoln, and includes representatives from BBN, ISI, and LINKABIT. A six-week program was planned, culminating in a report to DARPA in May. The essence of the program was to spend about two weeks collecting and evaluating system performance data, followed by extensive on-site testing at Lincoln by a team of BBN and LINKABIT personnel, with further iterations as necessary. Work is currently under way, and progress is being made toward the task force goals.

C. MULTIPLEXING EXPERIMENTS WITH EMULATED VOICE TRAFFIC

Work has started on a series of multiplexing experiments using emulated voice traffic generated by PVTs operating as measurement hosts (MHs). To change a PVT into an MH, we add a timer card that provides an accurate time base that is synchronized among MHs by connections to a WWVB clock. Software for the MH provides for traffic generation and the measurement of arrival-time histograms as well as counts of packets sent and received. Four traffic models are currently available. Three are data models using IP datagrams with deterministic, Poisson, or TCP file transfer traffic models. The fourth is a multiple-speaker talkspurt model that uses ST virtual circuits for its traffic. Measurements with this model have shown that a pair of MHs can emulate up to 10 conversations with 2400-bps LPC voice rates and with packet times of 40 ms. Under these conditions each MH generates an average of 125 packets per second (pps) and receives a like number from its partner in the conversation group. Measurements have also shown that the MH can sustain a steady flow of about 350 pps with packets of the same size. The total average load of 250 pps posed by 10 conversations is enough less than the peak capacity of 350 pps that packets are rarely lost due to overload of the MH under these conditions. The number lost when overloads do occur is sufficiently small that multiplexing experiments exploring packet loss rates in the order of 1 percent are not significantly affected by the MH overloads.

Experiments to date have involved traffic flowing between two gateways, mostly through a direct connection (through a "Null PSAT"), and occasionally through the




Fig. 10. Multiplexing experiment configuration.

PSAT and a loop through the wideband channel. Figure 10 shows the configuration used in our initial experiments. With two MHs on a LEXNET attached to each gateway, we get a maximum of 20 emulated LPC conversations which is a large enough number to get a TASI advantage of about 1.4 (i.e., we can handle 20 conversations with a channel capacity sufficient to handle only 14 if all transmitted continuously). The experiments involve adjusting the aggregation parameters of slot size, dispatch interval, and queue size limit in the dispatcher. The observed outputs are packet loss and dispersion of the arrival-time histograms.

Figures 11, 12, and 13 show arrival-time histograms for three experiments using the configuration of Fig. 10. The first (Fig. 11) shows the results of using a stream with a 44-ms interval and a size sufficient to carry 14 speech packets. The combination of 20 talkers and a packet time of 40 ms results in an average packet flow of 11 packets per stream interval. During periods when more than half the emulated




Fig. 12. Delay histogram for experiment II.

talkers are transmitting, queues will tend to build up in the gateway. For all three experiments, the time that a packet could wait on the queue at the stream dispatcher was limited to two dispatch intervals. Packets that could not be dispatched in that time were discarded. For each experiment the fraction of packets lost is indicated beneath the graph in the figure.

Figure 12 shows the result of using a stream interval half the size of the one used for Fig. 11. The size was reduced proportionately so that the carrying capacities of the streams were identical. Comparison of the figures shows that both the average delay and the delay dispersion are reduced by the faster dispatching, as would be expected. The differences in percentage of packets lost is not significant either from a human factors point of view or mathematically because of the short duration of the experiment (about 5 min.).

Figure 13 shows the result of decreasing the channel capacity from 14 packets per 44-ms interval to 13. Comparisons between Figs. 11 and 13 show increases in mean delay, delay dispersion, and packet loss. The conditions of Fig. 13 are about at the limit of acceptability for packet loss rate.

The goal of these experiments is to determine the optimal setting of the parameters as a function of offered traffic. This information can then be used by the gateway to adjust the PSAT stream reservation to suit the traffic flow indicated in the ST connection setup process.

The next step in the experimental program will be to mix IP data traffic with the voice traffic to explore the capability of the packet multiplexing mechanism to make effective use of channel capacity not used by voice for carrying data.

Longer range plans for the experimental program involve the use of the Funnel gateway and a Butterfly measurement host to support experiments with larger numbers of emulated talkers. The greater power of these machines should allow the emulation and multiplexing of hundreds of voice packet streams.




Fig. 13. Delay histogram for experiment III.

D. PACKET VIDEO FACILITY

A packet voice facility similar to that at ISI has been assembled and is now operating. The facility makes use of a PDP-11/45 computer and an Ikonas video processor supplied by ISI. A UMC-Z80 interface processor provides communication between the facility and either the 11/44-based IP/ST gateway or the voice funnel. Figure 14 shows a block diagram of the facility.

Steve Casner of ISI visited the Laboratory to supervise installation and checkout of the Ikonas hardware and system software. The facility is now working and has communicated successfully with the IP/ST gateway.





Fig. 14. Packet video facility.

V. VOICE CONTROL OF NETWORK VOICE CONFERENCING

A summary of results obtained with the Voice-Controlled Operator (VCOP) was presented at the DARPA Speech Research Review held at M.I.T. on 15-16 November 1982. Part of this summary described network tests of the VCOP. Four-participant conferences have been established by voice by conference initiators on a LEXNET at Lincoln Laboratory connected to the VCOP's LEXNET via a gateway. Participants in these conferences have used both LEXNETs. Conferences have also been set up by initiators on the VCOP's LEXNET.

More recently, a four-participant conference was established by voice over the Wideband Network. A talker at ISI (Steve Casner) first trained the VCOP by dialing the VCOP's number from a PVT and repeating the 20 participant names and command words needed to set up a conference five times each. The talker then hung up, redialed the VCOP, and answered prompting phrases provided by the VCOP to provide a list of conference participants, and the type of speech processing to use. This information was then passed to the conference access controller PVT which rang three PVTs at Lincoln. Participants at these PVTs joined the conference as soon as they picked up their PVT telephone handsets. Voice signals were sent using 64-kbps PCM coding and two conferences were successfully set up using voice input only.

The VOTRAX speech synthesizer that was originally used in the VCOP has been replaced with the stand-alone LPC vocoder board described in Sec. I of this report. This was desirable because of the poor intelligibility and the unnatural quality of the prompting phrases produced by the VOTRAX synthesizer. Software was developed to control this vocoder as part of the VCOP project. Independent programs were developed to store LPC parameters for prompting phrases on disk and to trim silent intervals from the beginning and end of stored phrases. A set of subroutines called from the main VCOP program was developed to play phrases back during the VCOP interaction with a conference initiator. The LPC SPP is used by first storing LPC parameters created when a talker produces a desired phrase in a named disk file. This file is stored on the host computer used with the VCOP and then trimmed to eliminate silent periods at the beginning and end of the phrase.



• PREVIOUSLY USED VOTRAX VSM-1

Fig. 15. VCOP block diagram.

Finally, when the phrase must be used by the VCOP as a prompt or response, the LPC parameters are downloaded from the host computer to the freestanding LPC vocoder. This vocoder uses these parameters to synthesize a speech waveform that is fed into the VCOP's PVT as was done with the VOTRAX synthesizer.

A block diagram of the current VCOP configuration is shown in Fig. 15. This configuration was used for the cross-net conference setup tests with ISI.

GLOSSARY

AC	Access Controller
BBN	Bolt, Beranek and Newman
CMU	Carnegie-Mellon University
CODEC	Coder/Decoder
DCEC	Defense Communications Engineering Center
EPROM	Erasable Programmable Read Only Memory
ESI	Earth Station Interface
ETHERNET	Local area network developed by Xerox Corporation
IP	Internetwork Protocol
ISI	Information Sciences Institute
LEXNET	Lincoln Experimental Packet Voice Network
LPC	Linear Predictive Coding
LSI	Large-Scale Integration
MH	Measurement Host
NOC	Network Operations Center
PCM	Pulse Code Modulation
pps	Packets per Second
PSAT	Packet Satellite Interface Message Processor
PVT	Packet Voice Terminal
RAM	Random Access Memory
RFI	Radio Frequency Interference
RFQ	Request for Quotation
ROM	Read Only Memory
SPP	Speech-Processing Peripheral
ST	Stream Protocol
STNI	Switched Telephone Network Interface
TASI	Time-Assignment-Speech Interpolation
TCP	Transmission Control Protocol

USART	Universal Synchronous/Asynchronous Receiver/Transmitter
VCOP	Voice-Controller Operator
VOTRAX	Commercially available speech synthesizer
WB SATNET	Wideband Satellite Network
WWVB	Radio signal maintained by National Bureau of Standards, used to synchronize clocks

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM									
1. REPORT NUMBER ESD-TR-83-027	2. GOVT ACCESSION NO. A132284	3. RECIPIENT'S CATALOG NUMBER									
4. TITLE (and Subtitle) Wideband Integrated Voice/Data Technology		5. TYPE OF REPORT & PERIOD COVERED Semiannual Technical Summary 1 October 1982 — 31 March 1983									
7. AUTHOR(s) Clifford J. Weinstein and Peter E. Blankenship		8. PERFORMING ORG. REPORT NUMBER									
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173-0073		10. CONTRACT OR GRANT NUMBER(s) F19628-80-C-0002									
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE 31 March 1983									
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB, MA 01731		12. NUMBER OF PAGES 52									
		15. SECURITY CLASS. (of this report) Unclassified									
		16a. DECLASSIFICATION/DOWNGRADING SCHEDULE									
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.											
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)											
18. SUPPLEMENTARY NOTES None											
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) <table border="0"> <tr> <td>packet speech</td> <td>ARPANET</td> <td>LEXNET</td> </tr> <tr> <td>network speech</td> <td>SATNET</td> <td>digital channel vocoder</td> </tr> <tr> <td>voice conferencing</td> <td></td> <td></td> </tr> </table>			packet speech	ARPANET	LEXNET	network speech	SATNET	digital channel vocoder	voice conferencing		
packet speech	ARPANET	LEXNET									
network speech	SATNET	digital channel vocoder									
voice conferencing											
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes work performed on the Wideband Integrated Voice/Data Technology Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 October 1982 through 31 March 1983.											

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)