

# **National Security and Emergency Preparedness Communications Experiments Using the Advanced Communications Technology Satellite**

**William Kissick  
David Sutherland  
Marjorie Weibel  
Wayne McCoy**

**Mary Ruhl  
Robert Toense  
Udayan Borkar  
Moorthy Hariharan**



**U.S. DEPARTMENT OF COMMERCE  
William M. Daley, Secretary**

Larry Irving, Assistant Secretary  
for Communications and Information

August 1998



## **PRODUCT DISCLAIMER**

Certain equipment and software products are identified in the body of this report for completeness and accuracy in describing the information presented. Such identification does not imply recommendation or endorsement of the companies or the products by the National Telecommunications and Information Administration, the National Institute of Standards and Technology, the Department of Commerce, the Office of the Manager-National Communications System, COMSAT Laboratories, or the MITRE Corporation.

## FOREWORD

This report describes experimentation with an advanced telecommunications satellite—the Advanced Communications Technology Satellite (ACTS) launched and operated by the National Aeronautics and Space Administration (NASA). The experimentation included measurement and analysis of the performance of several applications as well as investigation of protocol performance over the satellite channel. The opportunity to perform the work described herein is the result of the collaboration of several organizations—Government agencies and corporations—and the support provided by a sponsor.

The ACTS Collaboration, as it has been called, developed when several ACTS Experimenters determined that they had similar goals and interests. By combining resources, a small network of ACTS Earth stations (three of them) including several kinds of terrestrial connections was assembled. Using this experimental network and complementary skills, the ACTS Collaboration was able to perform several applications performance experiments and one experiment on the performance of an Internet protocol over the satellite.

Although each collaborator was able to provide some support, the National Communications System (NCS) Office of Programs sponsored much of the experimentation. This sponsor also is an ACTS Experimenter. The NCS mission to provide communications in support of National Security and Emergency Preparedness (NS/EP) gave the ACTS Collaboration a central purpose.

The National Telecommunications and Information Administration, Institute for Telecommunication Sciences (NTIA/ITS) provided liaison with NASA and the sponsor and performed the application experiment on voice quality. The National Institute for Standards and Technology, Computer Systems Laboratory (NIST/CSL) performed an application experiment on desktop conferencing and another application experiment on local area network (LAN) bridging. COMSAT Laboratories assisted NIST with satellite access and performed an experiment using Internet communications protocols. MITRE Corporation contributed to the design of experiments useful for NS/EP and assisted all of the collaborators with the experiments.

Collaborator points-of-contact are: Dr. William A. Kissick, coordinator and principal investigator, NTIA/ITS; Mr. Wayne McCoy and Ms. Mary Ruhl, NIST/CSL; Dr. Prakash Chitre, COMSAT; and Mr. Michael Nissley, MITRE. The NCS is represented by Mr. Frank Dixon.

Section 1 of this report was prepared by NTIA/ITS with input from all of the collaborators, especially NIST/CSL. Section 2 was prepared by NTIA/ITS. Sections 3 and 4 were prepared by NIST/CSL. Section 5 was prepared by COMSAT.

## CONTENTS

	Page
PRODUCT DISCLAIMER .....	iii
FOREWORD .....	iv
FIGURES .....	vii
TABLES .....	x
ACRONYMS AND ABBREVIATIONS .....	xi
ABSTRACT .....	1-1
1. OVERVIEW AND SUMMARY .....	1-2
1.1 National Security and Emergency Preparedness Telecommunications .....	1-2
1.2 Advanced Communications Technology Satellite .....	1-3
1.3 The ACTS Collaboration and Its Goals .....	1-4
1.4 The Emergency Scenario .....	1-5
1.5 General Conclusions .....	1-8
1.6 Suggestions for Future Work .....	1-9
2. APPLICATION EXPERIMENT ON VOICE QUALITY .....	2-1
2.1 The Voice Quality Assessment System .....	2-1
2.2 Equipment and Setup .....	2-3
2.3 Procedures and Experiments .....	2-7
2.4 Results .....	2-10
2.5 Conclusions .....	2-20
3. APPLICATION EXPERIMENT ON DESKTOP CONFERENCING .....	3-1
3.1 Experiment Objectives .....	3-1
3.2 Experiment Methods and Procedures .....	3-1
3.3 Metrics .....	3-3
3.4 Experimental Procedure .....	3-4
3.5 Expected Results .....	3-6
3.6 Results .....	3-6
3.7 Analysis .....	3-14
3.8 Interpretations .....	3-15
4. APPLICATION EXPERIMENT ON LAN BRIDGE .....	4-1
4.1 Experiment Objectives .....	4-1
4.2 Experiment Methods .....	4-1

## CONTENTS (cont.)

	Page
4.3 Metrics .....	4-3
4.4 Experimental Procedure .....	4-4
4.5 Expected Results .....	4-6
4.6 Results .....	4-7
4.7 Analysis .....	4-10
4.8 Interpretations .....	4-12
5. PROTOCOL EXPERIMENT ON TCP-LFN OVER FRAME RELAY .....	5-1
5.1 Introduction .....	5-1
5.2 Experiment Methods and Procedures .....	5-2
5.3 Performance Analysis of TCP-LFN .....	5-4
5.4 Results of TCP-LFN over Frame Relay .....	5-8
5.5 Conclusions .....	5-11
6. ACKNOWLEDGMENTS .....	6-1
7. REFERENCES .....	7-1

## FIGURES

	Page
Figure 1.1 The experimental ACTS network. ....	1-5
Figure 1.2 Emergency scenario configuration. ....	1-6
Figure 2.1 The VQAS algorithm. ....	2-2
Figure 2.2 Mean opinion scores of test sentences by system from ITS to NIST .....	2-11
Figure 2.3 Mean opinion scores of test sentences by system from NIST to ITS. ....	2-13
Figure 2.4 Voice quality through the local ISDN switch (LS) using the bit-error simulator and the laptop receiver. ....	2-15
Figure 2.5 Voice quality through the local ISDN switch (LS) using the bit-error simulator and the laptop receiver, reverse direction. ....	2-15
Figure 2.6 Voice quality of the terrestrial ISDN line (TI) using the bit-error simulator from ITS to NIST. ....	2-16
Figure 2.7 Voice quality of the terrestrial ISDN line (TI) using the bit-error simulator from NIST to ITS. ....	2-16
Figure 2.8 Voice quality of the ACTS ISDN channel (SAT) using the bit-error simulator from ITS to NIST. ....	2-17
Figure 2.9 Voice quality of the ACTS ISDN channel (SAT) using the bit-error simulator from NIST to ITS. ....	2-17
Figure 2.10 Voice quality through the local ISDN switch (LS) during simulated satellite wander. ....	2-18
Figure 2.11 Voice quality through the ACTS ISDN channel (SAT) during simulated satellite wander incident from ITS to NIST. ....	2-19
Figure 2.12 Voice quality of ISDN channel through local switch during simulated thunderstorm. ....	2-20
Figure 2.13 Voice quality through the ACTS ISDN channel (SAT) during simulated thunderstorm from NIST-to-ITS path. ....	2-21

## FIGURES (cont.)

	Page
Figure 3.1 Equipment configurations for application experiment on desktop conferencing. . . . .	3-2
Figure 3.2 Terrestrial videoconference usability opinion scores. . . . .	3-7
Figure 3.3 ACTS videoconference usability opinion scores. . . . .	3-7
Figure 3.4 Terrestrial H.320 videoconference usability opinion scores. . . . .	3-8
Figure 3.5 ACTS H.320 videoconference usability opinion scores. . . . .	3-9
Figure 3.6 Terrestrial tracking response time. . . . .	3-11
Figure 3.7 ACTS tracking response time. . . . .	3-12
Figure 3.8 H.320 video frame with no errors. . . . .	3-13
Figure 3.9 H.320 video frame with block errors. . . . .	3-13
Figure 3.10 H.320 video frame with motion compensation distortion (translation). . . . .	3-13
Figure 3.11 H.320 video frame with motion compensation distortion (ghosting). . . . .	3-13
Figure 4.1 Experiment equipment configuration. . . . .	4-2
Figure 4.2 File used in FTP experiments (actual file is a color image). . . . .	4-3
Figure 4.3 LAN bridge FTP throughput results. . . . .	4-7
Figure 4.4 LAN bridge usability over ACTS. . . . .	4-9
Figure 4.5 LAN bridge usability opinion scores for the terrestrial line. . . . .	4-10
Figure 4.6 Ping packet loss results. . . . .	4-11
Figure 5.1 Experiment setup. . . . .	5-3
Figure 5.2 TCP-LFN throughput for the static bandwidth management case. . . . .	5-5



## FIGURES (cont.)

	Page
Figure 5.3 Regular TCP vs. TCP-LFN throughput. ....	5-5
Figure 5.4 TCP-LFN preallocated mode throughput. ....	5-6
Figure 5.5 Effect of slow start on throughput. ....	5-7
Figure 5.6 TCP-LFN throughput vs. BER. ....	5-9
Figure 5.7 FTP performance over TCP-LFN. ....	5-10

## TABLES

	Page
Table 2.1	VQAS Test Sentences ..... 2-2
Table 2.2	VQAS Mean Opinion Score Scale ..... 2-3
Table 2.3a	VQAS Measurement Set for Characterization ..... 2-4
Table 2.3b	VQAS Measurement Set for Simulation ..... 2-5
Table 2.4	Sequential BER's to Simulate Degradation Caused by Precipitation in a Thunderstorm ..... 2-9
Table 2.5	BER's Measured During the Satellite Wander Incident on January 21, 1995 .. 2-10
Table 3.1	Definition of Opinion Score Values ..... 3-3
Table 3.2	Steady State BER Configurations ..... 3-5
Table 3.3	Burst BER Configurations ..... 3-5
Table 3.4	Recovery Response to Noise Burst (Proprietary, ACTS) ..... 3-9
Table 3.5	Recovery Response to Noise Burst (Proprietary, Terrestrial) ..... 3-10
Table 3.6	Recovery Response to Noise Burst (H.320, ACTS) ..... 3-10
Table 3.7	Recovery Response to Noise Burst (H.320, Terrestrial) ..... 3-11
Table 4.1	Definition of Opinion Score Values ..... 4-4
Table 4.2	Steady State BER Configurations ..... 4-5
Table 4.3	Burst BER Configurations ..... 4-5
Table 4.4	Burst BER Results Over the Terrestrial Line ..... 4-8
Table 4.5	Burst BER Results Over ACTS ..... 4-9
Table 5.1	Throughput vs. Number of FTP Sessions ..... 5-10

## ACRONYMS AND ABBREVIATIONS

ACTS	Advanced Communications Technology Satellite
ATM	asynchronous transfer mode
BER	bit error ratio
BRI	basic rate interface (ISDN 2B+D)
COTS	commercial off-the-shelf (equipment)
CRC	cyclical redundancy check
DCT	discrete cosine transform
DSP	digital signal processor
$E_b/N_0$	energy per bit to noise energy ratio
ES	Earth station
ERS	Emergency Response Site
FFT	fast Fourier transform
FRACS	frame relay access switch
FTP	File Transfer Protocol
FTS2000	ISDN service in the Federal Government
GETS	Government Emergency Telephone Service
GUI	graphical user interface
H.320	ITU-T Recommendation on videoconferencing
HO	home office
ISDN	integrated services digital network
IP	Internet protocol
ITS	Institute for Telecommunication Sciences
LAN	local area network
LEO	low earth orbit
MOS	mean opinion score
NCS	National Communications System
NII	National Information Infrastructure
NIST	National Institute of Standards and Technology
NLP	National Level Program
NS/EP	National Security and Emergency Preparedness
NTIA	National Telecommunications and Information Administration
OS	opinion score
PBX	private branch exchange
PC	processing center
POTS	plain old telephone service
PRI	primary rate interface (ISDN 24B+D)
PSN	public switched network
QOS	quality of service
RTT	round-trip time
T1	digital transmission service at 1.544 Mb/s
TCP	transmission control protocol

## **ACRONYMS AND ABBREVIATIONS (cont.)**

TCP/IP	transmission control protocol/Internet protocol
TCP-LFN	transmission control protocol - long fat network
VQAS	voice quality assessment system
VSAT	very small aperture terminal
WAN	wide area network

# **NATIONAL SECURITY AND EMERGENCY PREPAREDNESS COMMUNICATIONS EXPERIMENTS USING THE ADVANCED COMMUNICATIONS TECHNOLOGY SATELLITE**

William Kissick, David Sutherland, Marjorie Weibel<sup>1</sup>  
Wayne McCoy, Mary Ruhl, Robert Toense<sup>2</sup>  
Udayan Borkar, Moorthy Hariharan<sup>3</sup>

Many government telecommunications needs, especially those that support National Security and Emergency Preparedness (NS/EP) missions, are becoming increasingly dependent on commercially available equipment and services. This is consistent with the goals and concepts of the National Information Infrastructure. This report examines the use of an advanced satellite—in this case, NASA's Advanced Communications Technology Satellite (ACTS)—with ISDN and frame relay protocols to support NS/EP communications requirements. A network using three ACTS Earth stations was established as a research facility. With this small network, several experiments were performed. Using new objective methods, voice quality was measured over the satellite and compared to other connections such as commercial, terrestrial lines. The performance of applications—desktop conferencing, file transfer, and LAN bridging—that are likely to be useful in NS/EP situations, was determined. The performance of TCP/IP running over frame relay was examined. The results indicate that advanced satellites can be very useful for emergency communications due to the rapidity that Earth stations can be deployed, the ease of reconfiguring the satellite, and the practicality of using commonly available applications running over commonly used protocols. However, there are some limitations to the performance of some applications or parts of applications due to the propagation delay of a satellite channel. Telecommunications protocols such as TCP/IP must be significantly modified to perform well over a satellite channel and to take full advantage of bandwidth-on-demand capabilities of an advanced satellite.

Key words: Advanced Communications Technology Satellite (ACTS); integrated services digital network (ISDN); frame relay; National Security and Emergency Preparedness (NS/EP).

---

<sup>1</sup>These authors are with the Institute for Telecommunication Sciences, National Telecommunications and Information Administration, U.S. Department of Commerce, Boulder, CO 80303.

<sup>2</sup>These authors are with the National Institute of Standards and Technology, U.S. Department of Commerce, Gaithersburg, MD 20899.

<sup>3</sup>These authors are with COMSAT Laboratories, Clarksburg, MD 20871.

## 1. OVERVIEW AND SUMMARY

When a crisis arises, whether natural disaster or threat to national security, there are sudden needs whose rapid fulfillment can mean the difference between life and death and loss of property. These needs include rescue operations, medical aid, law enforcement, and the distribution of food, water, and other supplies. Effective communication and coordination are essential to effectively meet these needs.

Not only the support efforts, but also the cause of the crisis itself places constraints on existing telecommunications facilities. For example, an earthquake not only creates demands far in excess of what is required under normal conditions, but also can sever terrestrial telephone lines. Preparation for and provision of the telecommunications needed by Federal and local authorities during such events is the responsibility of the National Communications System (NCS).

The NCS is responsible for ensuring that Federal Government requirements for National Security and Emergency Preparedness (NS/EP) telecommunications are met. NCS accomplishes this through a variety of programs that: (1) ensure the interoperability and availability of systems by contributing to the development of technical standards, and (2) assist the Executive Office and various Federal agencies with the management and coordination of the nation's telecommunications resources.

In order to continue to meet NS/EP communications requirements, the NCS must take advantage of the emerging high-speed, multifunctional National Information Infrastructure (NII). For the foreseeable future, the NII must include satellites in order to provide true national coverage, especially when rapid deployment is needed in an affected area or disaster site. A hybrid telecommunications system comprised of wireless and wireline services augmented with advanced satellites must be considered by NCS to best meet their mission of providing NS/EP communications for the United States.

### 1.1 National Security and Emergency Preparedness Telecommunications

The mission of the NCS was reaffirmed in a White House Memorandum, dated October 15, 1991. That memo reaffirms that the policy guidance given in Executive Order 12472 and National Security Decision Directive 97 (NSDD-97) still represents the primary mission of the NCS. The memorandum gives the functional requirements of the National Level Telecommunications Program. Specifically, these are:

Voice Band Service	The service must provide voice band service in support of Presidential communications.
Interoperability	The service must interoperate with and use the resources of selected other Government or private facilities, systems, and networks through the application of standards.

Survivability/Endurability	The service must provide for the interconnection of surviving users under a broad range of circumstances from wide-spread damage from natural or manmade disaster up to and including nuclear war.
International Interface	The service must provide access to and egress from international service.
Nationwide Coverage	The service must provide readily available nationwide coverage to support the national security leadership and intra/interagency emergency operations.
Intra/Interagency Emergency Operations	Common user service must provide NS/EP traffic with priority service.

In order to accomplish its mission, the NCS has undertaken a number of important activities and programs. The National Level Program (NLP) and evolutionary architecture have focused on the public switched network (PSN) resources to meet NS/EP telecommunications needs. The Government Emergency Telecommunications Service (GETS) program is a key element of the NLP and has been established to add NS/EP features (e.g., priority treatment) and enhancements (e.g., improved routing) into the various commercial carriers' networks to better accommodate NS/EP user requirements in emergencies.

## 1.2 Advanced Communications Technology Satellite

The Advanced Communications Technology Satellite (ACTS) provides an experimental prototype of the features and capabilities planned for future communications satellites. ACTS demonstrates two emerging technologies of great potential value to NS/EP communications: spot beams and on-board switching. Successful integration and optimization of these related technologies will enable ACTS-derived operational satellites to provide hundreds of megabits per second of switchable "bandwidth-on-demand" between Earth stations located virtually anywhere in the Western hemisphere. This is a flexibility that no practical combination of redundant terrestrial facilities can equal. Satellite-based switching facilities could be substituted for damaged terrestrial-switching facilities located anywhere in a network. Earth stations could be permanently located at key terrestrial network nodes or could be quickly deployed to other locations when required due to emergency situations. Permanent Earth stations could be used to support operational traffic under normal conditions by improving terrestrial load sharing and by providing much of the revenue needed for their installation and operation.

Satellites are an important element in providing NS/EP telecommunications and restoring of the PSN. As the PSN evolves toward new technologies and as new satellite technologies, such as ACTS, are introduced, the NCS must continue to ensure that NS/EP telecommunications requirements will be met. NCS has already undertaken experiments that demonstrate the ability of the ACTS system to accommodate required NS/EP features (e.g., access security and priority/preemption) and to support

voice communications and connectivity between NS/EP users and the PSN. Another aspect, the focus of this study, is to examine some typical, representative NS/EP communications and to determine how well the users' applications work over a satellite network.

### **1.3 The ACTS Collaboration and Its Goals**

The ACTS Collaboration is comprised of the following four Government agencies and private corporations:

1. The National Telecommunications and Information Administration (NTIA), Institute for Telecommunication Sciences (ITS), Systems and Networks Division, Boulder, Colorado.
2. The National Institute for Standards and Technology (NIST), Computer Systems Laboratory (CSL), Gaithersburg, Maryland.
3. COMSAT Laboratories, Network Technology Division, Clarksburg, Maryland.
4. MITRE Corporation, Reston, Virginia.

Each of the collaborators is an official ACTS experimenter, and three of the collaborators have an ACTS Earth station. The formation of this collaboration resulted from two opportunities. One was the ability to assemble a small network that includes the ACTS Earth stations of three of the collaborators and other telecommunications media as shown in Figure 1.1, and another was the unique combination of knowledge and skills represented in the members of the collaboration.

The general goals of the ACTS Collaboration are to demonstrate and, where possible, measure the performance of some commercial off-the-shelf (COTS) equipment and capabilities over ACTS. Each demonstration or measurement represents some aspect of NS/EP communications. Applications experiments include voice communications, desktop conferencing, and LAN bridging. A protocol experiment tested the performance of TCP operating over frame relay using variable (static and dynamic) bandwidth and TCP window size.

The objective of each experiment in the set was to determine the qualitative and/or quantitative performance of satellite communications used in tests similar to what would be needed in actual NS/EP situations. Specifically, this effort demonstrates the feasibility of using an advanced communications satellite system to enhance or replace terrestrial communications facilities, and evaluates the performance of a set of applications as viewed from the users' perspective. The performance data can support analyses and enable conclusions regarding which technology features are most effective and the quality of service levels that can be expected. The results and conclusions are useful for setting minimum performance specifications for advanced satellite communications systems and in defining other requirements for advanced satellite systems supporting NS/EP communications.



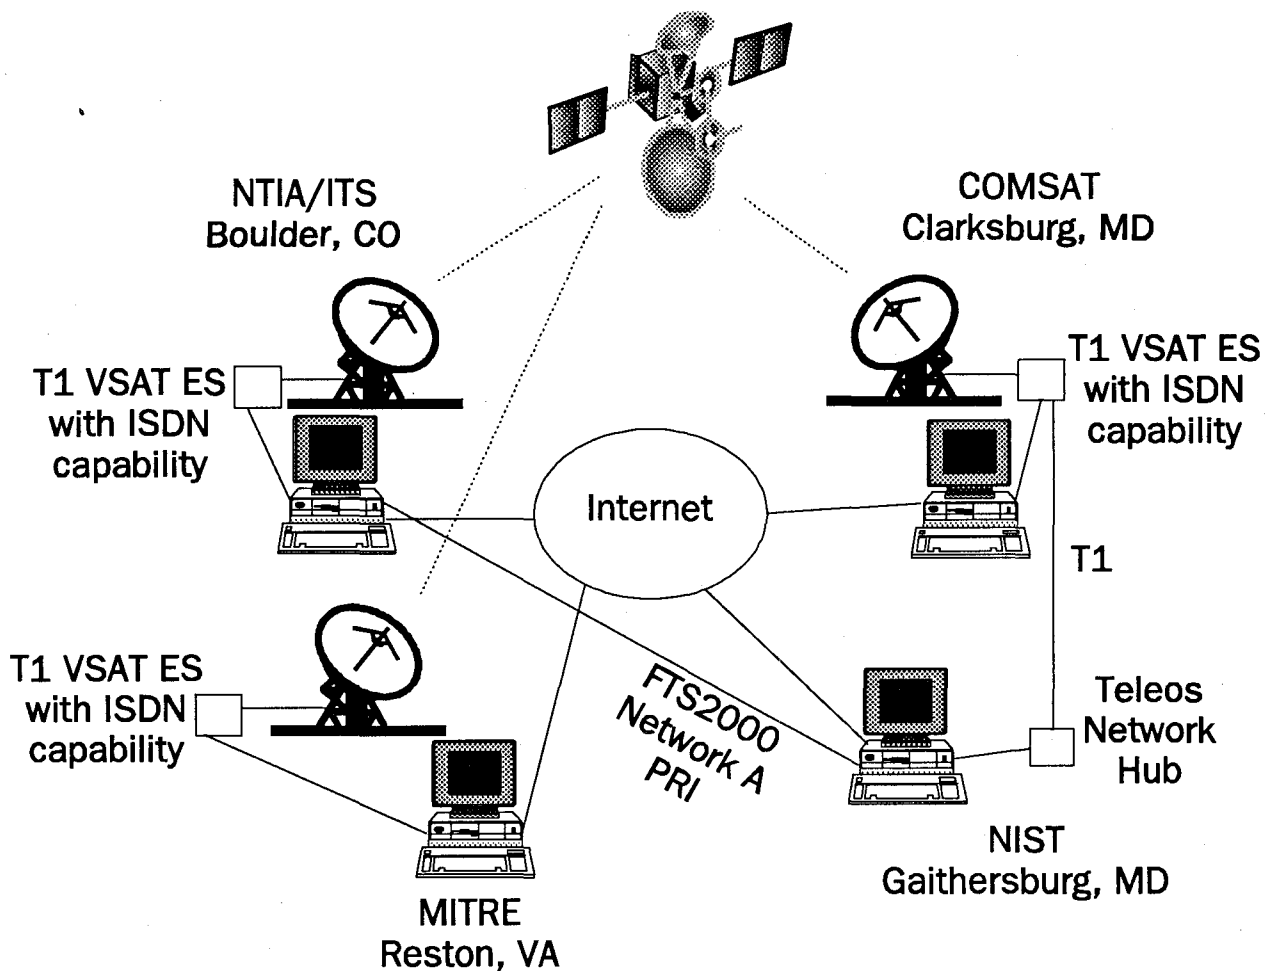


Figure 1.1. The experimental ACTS network.

#### 1.4 The Emergency Scenario

The experiments represented a hypothetical emergency scenario as shown in Figure 1.2. The purposes of the scenario were to provide a fundamental configuration of telecommunications equipment and to define the roles and activities of each site. The Emergency Response Site (ERS), located at NTIA/ITS, represents elements of emergency response organizations that require communication to locations outside the affected area. The Processing Center (PC), located at MITRE, represents some of those locations outside the affected area that can provide support for inquiries into medical records, map databases, insurance forms processing, and coordination for

supplies. The Home Office (HO), located at COMSAT and NIST represents regional or national headquarters, and can provide the ERS with additional communications connections such as a LAN or Internet bridge including applications such as E-mail and World Wide Web server access.

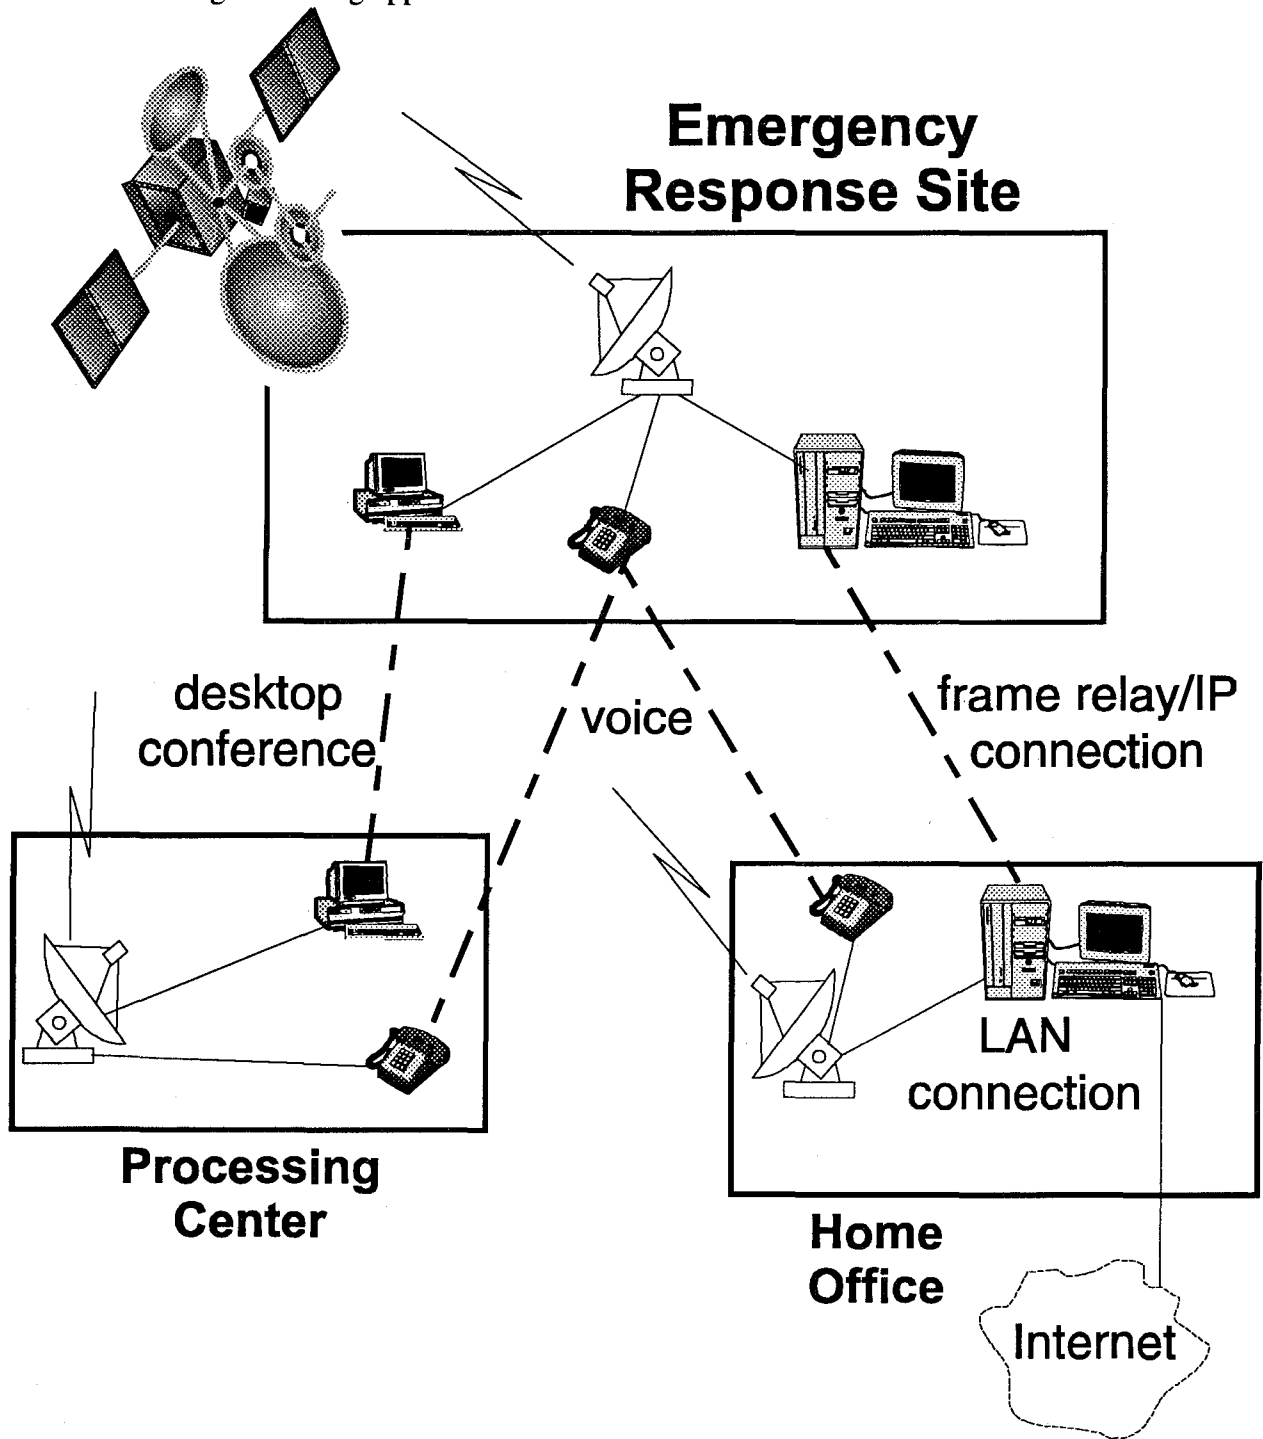


Figure 1.2. Emergency scenario configuration.

The PC and HO, in general, support the ERS with four types of communications services: voice telephone, teleconferencing, desktop conferencing, and data communications.

Voice telephone service is defined as a 64-kb/s link using a local ISDN switch. This type of communication allows for interconnections to the public switched telephone network and is the most used form of communications in emergency situations. In this scenario, a number of voice calls are made between the ERS and the PC or HO. Voice quality is measured using a voice quality assessment system (VQAS).

Desktop conferencing is supported by the PC and includes slow-speed video, voice, and data transfer over one channel. This system uses an ISDN basic rate interface (BRI) for these services. The video operates at a rate of 112 kb/s and can also provide a 640 x 480 video "snapshot" for higher quality still images. Voice communications within this system use a rate of 16 kb/s. Desktop conferencing includes the capability of document sharing.

Data communications between the ERS and the HO include a LAN connection and Internet access at the HO, as well as access to the HO's computing facilities. Through this connection, ERS personnel potentially have access to e-mail services at HO mailboxes and to the entire Internet, including file transfers, remote log-ins and World Wide Web servers. For example, the LAN connection may be used to download hourly weather updates from the Geostationary Operational Environmental Satellite (GOES) using a browser such as Mosaic® or Netscape®.

Examples of information that might flow between the ERS and the PC include requests for equipment, supplies, and other resources such as food and temporary shelter; medical support requests and medical information; requests for support for damage assessment and hazardous materials situations; damage reports; coordination of emergency response organizations; insurance form processing; and fund transfers. Communications flow through the HO includes information and data that is directly available on-line or from accessible data bases. ERS personnel access this information directly through the LAN/Internet connection provided by the HO. This information could include weather forecasts and weather satellite photographs; geological and geophysical information, activity reports and forecasts; and medical expert systems and databases.

Numerous simplifications have been assumed in order to conduct experiments in a straightforward manner. For example, only a small sample of the possible activities that occur in an emergency situation were examined. In addition, the examination was at a somewhat high level of abstraction. This was done so that the types of information movement that might occur in the applications used in an emergency situation could be determined. These interactions will be affected by the characteristics of the satellite channel in different ways and with different effects on the applications that use them.

Four experiment categories are described in detail in major sections of this report. Three of these are applications experiments that include voice quality described in Section 2, desktop conferencing in Section 3, and LAN bridging in Section 4. The fourth experiment on TCP over frame relay is described in Section 5. The subsections below present conclusions and recommendations that are drawn from the full set of experiments, and some recommendations for future work.

### 1.5 General Conclusions

For the applications tested—voice communications, desktop conferencing, and LAN bridging—the delay due to the satellite channel is not a significant contributor to reduction of the usability or quality of the applications. Delays due to processing in terminal equipment dominated. This only presents a problem on highly interactive activities, such as the use of the whiteboard function in the desktop conferencing. However, channel bit error rate (BER) is a significant contributor to reductions in usability or quality of the applications unless the bit error rates are less than  $10^{-5}$ .

For the voice communications application, there is little or no difference in voice quality, or intelligibility, across all ISDN connections, including the satellite. All ISDN connections are significantly better than long-distance telephone service.

The quality of all desktop-conferencing components—video, voice, and whiteboard—begins to degrade at a BER of  $10^{-6}$ ; they all become unusable at a  $10^{-3}$  BER. Although usable at BER's of  $10^{-5}$ , annoyance factors may be a problem for some users or in some situations. At this level, the audio signal is usable but only in good listening environments, and slow response may make the whiteboard function difficult to use. Generally, the video component degrades more quickly with increasing BER.

Each component in a hybrid network independently contributes to the quality of information an end user receives. This means that a single component can cause severe degradation of quality or even cause a loss of connection. This connection loss was confirmed during the LAN-bridging experiments and desktop conferencing, but can occur for any application. During local generation of errors, although the applications continued to operate, high BER's ( $> 10^{-3}$ ) caused the commercial, terrestrial network to terminate the connection due to administrative decision by the communication service provider. Generally, the LAN-bridging applications began to degrade at BER's of  $10^{-6}$  and became unusable for BER's of  $10^{-3}$  and higher. The degradation can be exacerbated for applications that require multiple connections.

Unlike the applications experiments, the TCP/IP over frame relay protocol experiment directly examined the supporting infrastructure; specifically, the support provided by TCP/IP and enhanced TCP (TCP-LFN) as part of the system with no applications. Over the satellite channel, TCP/IP with a default window size of 8 kB supports a peak throughput of less than 100 kb/s with any bandwidth allocation larger than this. Simply increasing the window size will increase the peak throughput to a limit of a few hundred kb/s. If the increases in window sizes are accompanied by a more careful treatment of round-trip times (RTT, needed for retransmission decisions), then the throughput rises

to fill the bandwidth—in this case 1.544 Mb/s. Enhanced TCP (TCP-LFN) has a very large window size ( $2^{31}-1$ ), time-stamps for each packet, and other features.

## **1.6 Suggestions for Future Work**

The experiments and measurements described herein can be extended or built upon to obtain more information about satellite communications. The voice quality experiments of Section 2 show that the ISDN voice channels available over ACTS are of very high quality. In fact, the satellite and terrestrial ISDN channels offered about the same, high level of voice quality. These results can provide a reference for future measurements of voice quality. For example, the quality of a voice channel that must operate with some compression and coding as would occur on a typical low Earth orbit (LEO) satellite could be compared to the voice quality measurements reported here. Indeed, measurements of voice quality can be made in the same manner over any voice communications system; for example, wireless or cellular telephone. Combinations of voice communications systems can be assessed in the same way.

The information gained from the desktop-conferencing and LAN-bridging experiments also provides a reference for any future assessments of similar applications. The protocol experiment can be extended to asynchronous transfer mode (ATM) and could include the use of Internet tools such as e-mail and web browsers to generate traffic. The ACTS will be available for experimentation for about two more years.



## **2. APPLICATION EXPERIMENT ON VOICE QUALITY**

Voice communications, used alone or in association with video communications or data transfer, represent an essential NS/EP requirement. For this reason, assessment of voice communications quality of the ISDN voice channel over ACTS is a key part of the ACTS Collaboration experiments.

The four objectives of the voice quality experiment were to:

1. Characterize the quality of voice calls over an ISDN ACTS link, a terrestrial ISDN link, and normal long-distance service (FTS2000). These link characterizations were made between Boulder, Colorado, and Gaithersburg, Maryland. For comparison purposes, a second set of link characterizations were made within the Department of Commerce Boulder Laboratories. This set was made over the local telephone switch and over an ISDN switch. A third set of measurements were made directly between the measurement equipment, in a back-to-back mode, with only a very short piece of audio cable.
2. Determine the quality of voice calls over the ISDN links under degradation provided by a simulator that injected bit errors into the digital bit stream.
3. Demonstrate the effect on quality caused by the degradation in a simulated satellite wander incident.
4. Demonstrate the effect on quality caused by rain in a simulated strong-to-severe thunderstorm occurring at an ACTS Earth station.

### **2.1 The Voice Quality Assessment System**

The Voice Quality Assessment System (VQAS) is the major tool used in this experiment. VQAS is the result of ongoing work at NTIA/ITS [1]. The ACTS Collaboration used this tool in a beta test mode. The use and evaluation of the VQAS in these experiments led to immediate upgrades of the VQAS system.

VQAS is a practical tool for assessing the quality of voice samples transmitted through a variety of channels and systems. It presently uses twelve stored voice samples that are converted to analog signals at a rate of 8 kb/s. The spectral content of the voice sample is 4 kHz, therefore this is the portion of the spectrum assessed for any channel under test. The samples are six recorded sentences spoken by a male speaker and a female speaker; prefixes M for male and F for female are used to identify the sex of the speaker, e.g., F5 or M3. The sentences are given below in Table 2.1.

Table 2.1. VQAS Test Sentences

Number	Test Sentence
1	These problems are both basic and pervasive to our society.
2	The group policy provides two types of insurance.
3	The directors raised the quarterly dividend to ninety-five cents.
4	It is difficult to pretend to be ignorant.
5	More books have resulted from somebody's need to write them than from anybody's need to read (them).
6	An employee regularly assigned to the night shift will receive night differential.

Two complete sets of test equipment are needed; there must be one set at each end of the channel under test. Each equipment set can act as transmitter or receiver. The algorithm used by VQAS to compare the reference sample with a received sample in the receiving test equipment set is outlined in Figure 2.1. The result of a measurement on a speech sample is a metric called the auditory distance. The auditory distance is mapped to the mean opinion score (MOS) scale (see Table 2.2). This MOS is highly correlated with the subjective MOS obtained from a listener panel using the same speech samples [1].

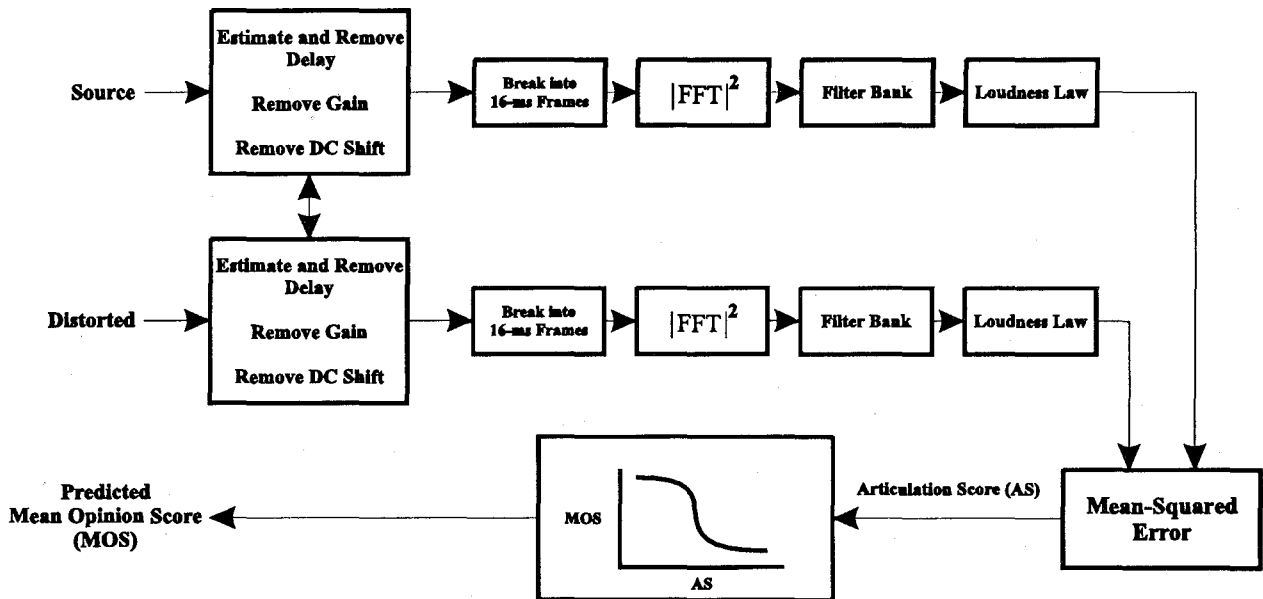


Figure 2.1. The VQAS algorithm.



Table 2.2. VQAS Mean Opinion Score Scale

Mean Opinion Score	Audio Quality
5	Excellent
4	Good
3	Fair
2	Poor
1	Unacceptable

For measuring the clean (not degraded by a bit-error simulator) channels, the MOS for each of the twelve samples is measured a sufficient number of times to achieve a 99% confidence interval at  $\pm 0.1$  about the mean for that sample set (see Tables 2.3a and 2.3b). For measurements over degraded channels, the variance was found to be sufficiently large to preclude the establishment of a standard confidence interval. The number of measurements necessary to obtain a useful confidence interval about the mean was so large over degraded channels that VQAS measurements would have used up all of the time allotted on the ACTS by NASA to the ACTS Collaboration.

## 2.2 Equipment and Setup

The test equipment needed to make the voice quality measurements includes the VQAS (comprised of hardware and software) and a bit-error simulator. The procedures for connecting and using this equipment are described below.

### 2.2.1 The VQAS System

The VQAS is comprised of hardware and software running on a computer (laptop or desktop) at each end of a telecommunications connection. On the transmitting end of the link, the software uses the hardware to convert a sample of speech, stored digitally on disk, back to its original analog form and transmit it. On the receiving end, the software uses the hardware to digitize the incoming speech and compare it to an exact copy of the original digital file. The hardware and software at each end can act as either the transmitter or the receiver. The hardware used to accomplish this overall task is the computer, a digital signal-processing (DSP) board, an expansion chassis to house the DSP board (needed if the computer is a laptop), and a hybrid adapter to change the audio (RCA phono plug) input and output of the DSP board into a signal suitable for the telephone system (RJ-11). The receiving end also needs a loudspeaker so that the human operator can hear and monitor the incoming test sentences.

Table 2.3a. VQAS Measurement Set for Characterization

Measurement Number	Description	Connection	Total VQAS Sentence Samples
1a	Reference	Back-to-back	306
1b	Reference, reverse channel	Back-to-back	90
2a	Boulder, building switch	Local voice call	95
2b	Boulder, building switch, reverse channel	Local voice call	110
3a	Boulder, ISDN switch	Local ISDN voice call	117
3b	Boulder, ISDN switch, reverse channel	Local ISDN voice call	96
4a	Terrestrial ISDN	Point-to-point 64-kb/s voice call	165
4b	Terrestrial ISDN, reverse channel	Point-to-point 64-kb/s voice call	101
5a	ACTS ISDN	Point-to-point 64-kb/s voice call	80
5b	ACTS ISDN, reverse channel	Point-to-point 64-kb/s voice call	63
6a	FTS2000	Regular long-distance voice call	73
6b	FTS2000, reverse channel	Regular long-distance voice call	55

Table 2.3b. VQAS Measurement Set for Simulation

Measurement Number	Description	Connection	Total VQAS Sentence Samples
7a	Boulder, ISDN switch with BER simulator	Local ISDN voice call	82
7b	Boulder, ISDN switch with BER simulator, reverse channel	Local ISDN voice call	80
8a	Terrestrial ISDN with BER simulator	Point-to-point 64-kb/s voice call	66
8b	Terrestrial ISDN with BER simulator, reverse channel	Point-to-point 64-kb/s voice call	54
9a	ACTS ISDN with BER simulator	Point-to-point 64-kb/s voice call	70
9b	ACTS ISDN with BER simulator, reverse channel	Point-to-point 64-kb/s voice call	66
10	ACTS ISDN, thunderstorm	Point-to-point 64-kb/s voice call	21
11	ACTS ISDN, satellite wander	Point-to-point 64-kb/s voice call	24
12	Boulder, ISDN switch, thunderstorm	Local ISDN voice call	21
13	Boulder, ISDN switch, satellite wander	Local ISDN voice call	24

A speech sample travels from the transmitting computer to the receiving computer in the following manner:

1. The VQAS program reads a speech sample from disk and outputs the speech in analog form via the D/A (digital-to-analog) functions of the DSP board in the transmitting computer.
2. The audio signal travels through a hybrid device that adapts the audio connectors and converts the impedance of the signal into the form appropriate for input to a plain old telephone service (POTS) voice channel of the system under test.
2. (Alternate) If a channel simulator is used, the POTS signal must be fed into an ISDN terminal adapter for conversion into a digital ISDN signal to be sent over ACTS via the ISDN switch. In this case, the signal is looped out of a primary rate interface (PRI) connection on the back of the ISDN switch, through the channel simulator, and back into a different PRI connection on the ISDN switch, and then into the system under test.
3. The signal is received at the receiving end of the connection where an operator hears the incoming speech via a loudspeaker. Prior to the beginning of the actual speech sample, the operator will hear the words "four, three, two, one." On the word "one" the operator hits the RETURN (or ENTER) key of the receiving computer, which begins the reception of the speech for analysis. When the ENTER key is hit, the VQAS software begins digitizing the incoming speech until it senses a period of silence, at which point it stops listening to the phone line and begins to analyze what it has captured. If the operator has approximately synchronized the beginning of the sample (by hitting the ENTER key on the "one" of the incoming speech), the software will more accurately synchronize the speech it has captured with the copy of the original it has stored on disk. The receiving computer then computes the auditory distance, maps it to an MOS, and displays the result.

### **2.2.2 Channel Simulator**

The ACTS Collaboration used a digital satellite channel simulator (AdTech SX-12) on terrestrial and satellite channels. This simulator was used to add delay, random errors, and burst errors to channels under test. For voice quality testing, various BER's were used to simulate communications in a thunderstorm and to simulate communications when the satellite had wandered out of position.

## 2.3 Procedures and Experiments

The measurements of voice quality are divided into two categories. The nondegraded category does not use the bit-error simulator to inject additional errors, and the degraded category does.

### 2.3.1 Nondegraded Tests

The test procedure was to measure the voice quality of the ACTS ISDN voice channel in both directions between ITS in Boulder, Colorado and NIST in Gaithersburg, Maryland. For the initial measurements, each test sentence was transmitted and measured for quality by the VQAS software. Each sentence was transmitted and measured five times in each direction. The MOS for each transmission was recorded automatically by VQAS. The recorded MOS's from VQAS measurements of telephone traffic (voice calls) over ACTS were analyzed. Mean and variance were computed and the 99% confidence interval was calculated. The desired confidence interval half-length is 0.1.

If the computed confidence interval width, for a particular test sentence in a particular direction, was within the established limit, there was no need for further measurement. If the interval width was greater than the desired precision, then the number of additional data points required,  $n$ , was estimated from the standard confidence interval half-width formula,

$$t_{n-1, 0.995} \sqrt{\frac{s^2}{n}},$$

where  $s^2$  is the estimated sample variance determined from the previously measured data points, and  $t$  is the upper (lower) 0.995 critical point for the t-distribution with  $n-1$  degrees of freedom [2]. The number of additional samples, when required, was usually kept to a multiple of five.

This same procedure was repeated for each test sentence. Additional MOS measurements give a more accurate estimate of the variance of the measurements for a test sentence and a more accurate estimate of the required number of data points to achieve the desired statistical precision. For comparison, other channels and systems were measured using the same equipment and procedures. These included a terrestrial ISDN link between NTIA/ITS in Boulder and NIST in Gaithersburg and a normal long-distance call between the same two sites. A reference measurement was made by connecting two VQAS systems back-to-back with only a short piece of audio cable. Other measurements were made through the local building telephone switch and through a local ISDN switch.

### 2.3.2 Degraded Tests

In emergency situations, environmental conditions may impair satellite communications. Since the ACTS operates in the Ka-band, the major impairment will be due to signal attenuation (fading) caused by moderate to heavy rain at the Earth stations. The degraded tests are based on the assumption that the signal strength fades caused by precipitation are longer than the test sentences used by the VQAS system. These sentences last 3 to 6 s. The signal-strength fades in a thunderstorm generally last at least several minutes. Background noise that affects the satellite system (galactic rf noise) also undergoes similar attenuation. Therefore, the only detriments are signal attenuation caused by absorption in the precipitation and scattering that increases noise. The BER will be constant since the signal attenuation does not vary (by assumption) and the increased noise power is directly related to the attenuation. So, bit errors are expected to have a Poisson distribution over time; this implies that the time between errors will be distributed exponentially (geometrically in the discrete sense). The log of the BER will be distributed uniformly. To reflect the rain rates, and therefore the BER's, as seen in convective rain activity (thunderstorms), it is necessary to test over a range of valid BER's.

ISDN degradation tests were conducted for both the ACTS and terrestrial voice channels and for the local ISDN channel (see Table 2.3b). The initial test determined a useful range of measurable degradation. These tests measured the quality of the single test sentence M6 at BER's of  $10^{-9}$ ,  $10^{-8}$ , . . . ,  $10^{-1}$ . At the low rates, there was no appreciable degradation while at the higher rates the degradation was enough to sever the ISDN connection. The useable BER's that affect the MOS but that do not sever the connection generally were found to be from  $10^{-6}$  to  $10^{-3}$ .

Once this useable range was determined for a specific system, then degraded-quality was tested on that system with just two test sentences: sentence 1 with the female speaker (F1) and sentence 6 with the male speaker (M6; see Table 2.1). These two sentences were chosen for stability in the measured scores and for having noticeably higher scores than other sentences measured during the nondegraded tests. Each of these sentences was transmitted five times in each direction at each different BER.

A thunderstorm model was also developed to simulate the rapidly changing voice quality that can be expected in a thunderstorm occurring at an ACTS Earth station. The test measured the changing effect on quality of test sentence M6 under rapidly changing BER. The sentence was transmitted once at each BER value and the quality was measured and recorded.

The rain degradation model is based on the expected attenuation caused by rain in a strong to severe thunderstorm that passes directly over an Earth station. Rain rates in a thunderstorm of this magnitude will vary widely as the storm passes over an Earth station; however, temporal extent is on the order of several minutes. Likewise, the BER could vary dramatically during the course of the storm but still remain constant for up to several minutes at a time. Since the test sentences used by the VQAS are less than 6 s long, it is assumed that the BER will be constant during the test sequence. The logarithm of the BER is distributed uniformly. The BER's for the thunderstorm/rain degradation experiment are given in Table 2.4.

Table 2.4. Sequential BER's to Simulate Degradation Caused by Precipitation in a Thunderstorm

STEP	BER	STEP	BER
1	$7 \times 10^{-7}$	13	$1 \times 10^{-7}$
2	$3 \times 10^{-1}$	14	$4 \times 10^{-9}$
3	$2 \times 10^{-8}$	15	$1 \times 10^{-8}$
4	$4 \times 10^{-3}$	16	$4 \times 10^{-6}$
5	$6 \times 10^{-6}$	17	$2 \times 10^{-6}$
6	$2 \times 10^{-4}$	18	$3 \times 10^{-8}$
7	$4 \times 10^{-8}$	19	$3 \times 10^{-5}$
8	$7 \times 10^{-2}$	20	$4 \times 10^{-8}$
9	$5 \times 10^{-3}$	21	$4 \times 10^{-2}$
10	$3 \times 10^{-4}$	22	$3 \times 10^{-3}$
11	$3 \times 10^{-4}$	23	$3 \times 10^{-8}$
12	$2 \times 10^{-6}$	24	$3 \times 10^{-4}$

Another impairment considered is satellite wander. This is a situation where the satellite orientation changes causing the antenna beam to no longer be accurately aimed. Wander was simulated using the data from an actual event recorded at an ACTS Earth station. The simulation used in this study is based on actual measurements made at the ACTS Master Control Station during a 15-hr period. The data from the worst 6-hr period were used to develop a "compact" version of the wander incident.

The satellite wander model is based on recorded uplink bit error rates recorded for a 15-hr period beginning January 21, 1995 from T1-VSAT 12. The downlink bit error rate for this time period was assumed constant at  $5 \times 10^{-9}$ ; that is, the downlink was not affected by the satellite wander. The BER measurements in 15-min intervals for the 6 hrs that contain the worst 4 hrs are given in Table 2.5. It is assumed that the uplink bit error rates will affect all uplink channels.

Table 2.5. BER's Measured During the Satellite Wander Incident on January 21, 1995

STEP	BER	STEP	BER
1	$2 \times 10^{-6}$	13	$2 \times 10^{-5}$
2	$2 \times 10^{-6}$	14	$1 \times 10^{-4}$
3	$3 \times 10^{-6}$	15	$5 \times 10^{-5}$
4	$9 \times 10^{-6}$	16	$2 \times 10^{-5}$
5	$9 \times 10^{-6}$	17	$1 \times 10^{-5}$
6	$9 \times 10^{-6}$	18	$3 \times 10^{-6}$
7	$8 \times 10^{-6}$	19	$1 \times 10^{-6}$
8	$1 \times 10^{-4}$	20	$9 \times 10^{-7}$
9	$3 \times 10^{-4}$	21	$7 \times 10^{-7}$
10	$1 \times 10^{-4}$	22	$5 \times 10^{-7}$
11	$8 \times 10^{-5}$	23	$3 \times 10^{-7}$
12	$2 \times 10^{-4}$	24	$1 \times 10^{-7}$

## 2.4 Results

The results of the nondegraded and degraded voice quality measurements are reported below.

### 2.4.1 Nondegraded Measurements

#### ITS-to-NIST

Figure 2.2 is a compilation of channel characterization results for comparison of voice quality as measured by the VQAS. The point-to-point measurements were made at NIST in Gaithersburg after transmission from ITS in Boulder; these are the ACTS ISDN (SAT), Terrestrial ISDN (TI), and regular long-distance telephone service (LD). The local measurements were made at ITS in Boulder; these are back-to-back measurements (BB) with no intervening communication system, measurements through the local switch (LS) of the building telephone system, and measurements over an ISDN line through an ISDN switch (TS). (This is the switch through which all ISDN connections are made.) A personal computer with the VQAS system installed transmitted to a laptop computer with the same



peripheral equipment as the personal computer. Local measurements were concluded and the laptop equipment shipped to NIST in Gaithersburg for the point-to-point measurements.




Figure 2.2. Mean opinion scores of test sentences by system from ITS to NIST.

Clearly the regular long-distance (LD) service had poorer quality over the measured VQAS sentences than the other systems. The VQAS operator at NIST reported that the perceived quality, heard over loudspeakers, was consistent with the VQAS measurements. The high-frequency components of the speaker's voice for each sentence was noticeably "clipped." The sentences spoken by the male sounded better than those spoken by the female. However, in a sentence-by-sentence comparison, the average MOS's for the male were not significantly higher than those for the female.

The three ISDN systems (SAT, TI, and TS) do not seem to be significantly different except that the ACTS satellite system had the smallest range of values (0.456) of the three. The terrestrial ISDN range was the largest of the six systems (0.5625). These ranges indicate that the test sentences with different spectral content are affected differently by the system under test. The satellite channel had a smaller range (and smaller variance) than the terrestrial ISDN channel indicating that the satellite channel has greater stability than the TI channel.

A faint echo was heard over the monitor at ITS during transmission over the satellite channel. This was apparently due to an echo cancellation feature that was not invoked in the ACTS Earth stations. The echo was heard only when the satellite channel was investigated. The echo did not appear to affect the quality measurement of the satellite channel. Variance, with the echo present, was not larger than that for other channels.

The VQAS could not record quality measurement for test sentence F5 over the ISDN systems (TS, SAT, and TI) and over LD. During testing, every measured sentence was monitored at the receive site by the test operators. The operators heard nothing discernable to indicate why sentence F5 failed to achieve a score; the sentence had the same perceived quality as the other sentences. For the BB baseline measurements, sentence F5 scored the highest average MOS (3.965) of the twelve sentences. The sentence consistently recorded an MOS of 1.00 over the ISDN systems. In rare cases, a valid measurement would be recorded, but never enough in the time allotted to obtain sufficient data points for the desired statistical precision. This effect is believed to be an artifact of the VQAS operating on the test sentence itself, such that the VQAS algorithm was precluded from properly lining up the received sentence with the recorded sentence. The test sentence F1 also had the same problem over the LD channel; however, this was due to the high-frequency clipping on the channel that prevented the VQAS from properly synchronizing the received and the stored forms of the F1 sample.

The VQAS measurements for the terrestrial ISDN (TI) channel on the ITS-to-NIST path did not meet the statistical precision criterion for locating the mean (half length of 0.1 at 99% confidence). The confidence measurements were unstable enough that the large variance precluded obtaining the desired confidence interval in a reasonable amount of time. For TI, this precision was lowered to a 0.15 interval half length at 95% confidence in order to report results for comparison with the other channels.

### **NIST-to-ITS**

Figure 2.3 is a compilation of the quality characterization of the same channels in the reverse path. That is, NIST in Gaithersburg transmitted test sentences to ITS in Boulder over the SAT, TI, and LD paths. Similarly, the BB, LS, and TS paths were tested at Boulder with the laptop system (later sent to Gaithersburg) transmitting and the PC receiving and evaluating.

Measurements of the BB, LS, and LD channels appear essentially the same; however, the quality of the LS and LD channels was slightly better in the other direction (ITS-to-NIST). As with the ITS-to-NIST path, the measured and perceived quality of the regular long-distance channel was noticeably worse than the other channels. Again this appears to be due to high-frequency clipping. This was observed by monitoring the received sentences through the loudspeakers. The high-frequency clipping was also heard on a telephone handset connected prior to the point that the transmissions entered the VQAS system.

Measured quality was better in the NIST-to-ITS direction on all three ISDN systems: TS, SAT, and TI. This may indicate a quality difference in direction in that portion of the path common to the three channels: from the ISDN terminal adaptor through the ISDN switch. Note that although the range (and variance) of quality of the TI channel was the highest of the six measured channels (higher than the long-distance line), the quality of the TI channel was stable enough to locate the mean of the MOS measures of each sentence within 0.1 at 99% confidence, which was not the case for the quality of this channel in the other direction.

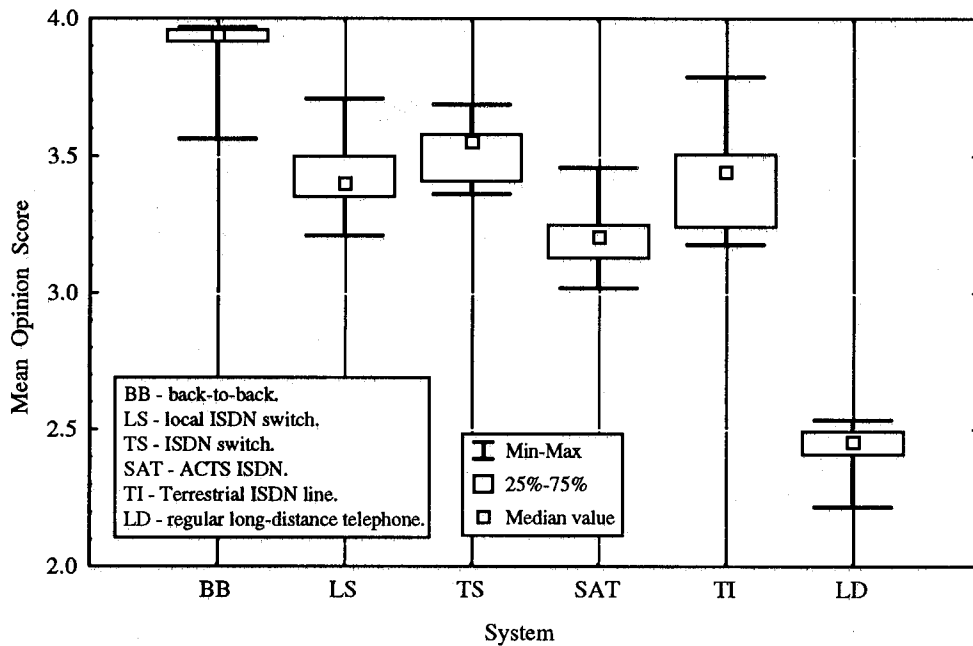


Figure 2.3. Mean opinion scores of test sentences by system from NIST to ITS.

The quality of the terrestrial channel appears to be better than the quality of the satellite channel in this direction; the opposite was true in the ITS-to-NIST direction. No reason was apparent for this discrepancy; however, it could be due to a directionality characteristic in the ACTS or in the terrestrial ISDN channel (TI). Note that the range of values (and the variance) was greater on that channel than on the ACTS ISDN (SAT) channel. This indicates that the satellite channel has more apparent stability than the terrestrial channel. The echo over the satellite was still heard by the NIST operators, but there was no apparent effect on the measurements.

As before, the VQAS could not record a quality measure for test sentence F5 over the three ISDN channels tested and the regular long-distance telephone (LD) channel. Also, the VQAS could not record a quality measure for the test sentence F2 over the long-distance telephone (LD) channel. The VQAS did record a quality measure for the sentence F1, which was not recorded at all in the ITS-to-NIST direction over the long-distance telephone (LD) channel.

#### 2.4.2 Degraded Tests

The degraded tests comprise all voice quality testing in which the ISDN channels were degraded by a BER simulator in the communications channel. The purpose was to examine effects of elevated BER on the voice quality of the end-to-end digital communications channels. The simulator was located at ITS in Boulder for all measurements.

## Degraded ISDN Measurements

Figure 2.4 is a compilation of voice quality measurements through the local ISDN switch (TS) with a bit-error simulator in the line. A data point in this graph is an average of about five MOS measurements at a particular error setting of the simulator. The receive equipment was the VQAS laptop computer. The x-axis scale is the negative logarithm of the bit error rate setting of the simulator; this scale is used for ease of display. For example, at the vertical scale line labeled 4, the bit error rate is  $10^{-4}$ ; this is the average rate (and variance) of the Poisson distribution that the simulator uses. The BER decreases to the right in the graphs. Two of the sample sentences, F1 (shown as circles) and M6 (shown as squares) are used for these tests. The curves fitted to these points are generally third-degree polynomials intended to indicate the general trend of the data points. The curves are not the result of a careful analysis of the available points. There are not enough points available for such an analysis and points due to bit error rates lower than  $10^{-6}$  were not measured. Quality at points below  $10^{-3}$  were also not measured since the communications connections were usually lost at this point.

The first observation is that quality measurements for M6 are consistently better, by about 0.5, than the measurements for F1. This is probably due to the differing responses of the VQAS due mainly to the spectral content of the two different sentences. The second observation is that the quality of the two sentences, although different, appears to follow the same trend. The curve for F6 is close to what was expected: the curve should be asymptotic to the clear channel quality on the right and approach an asymptote at the "unacceptable" reading of 1.00. However, a reading of 1.00 can also indicate a failure of the VQAS to properly synchronize.

Figure 2.5 is the quality of the reverse channel. The same general comments above apply here. Notice that this direction of the channel produced consistently better scores for each of the sentences. This indicates a directional bias in either the communications equipment or in the VQAS measurement system.

Figures 2.6 and 2.7 are similar quality measures over the terrestrial ISDN TI channel under the bit-error simulator. Figure 2.6 is the quality measure over the path ITS-to-NIST. Figure 2.7 is the quality measure over the reverse path. The measurements appear to have the same general characteristics of the measurements over the local ISDN switch (TS). Note especially the same difference in quality by direction. The curve fit for sample sentence M6 in Figure 2.7 is not included due to lack of data points.

Figures 2.8 and 2.9 show the measures of quality over the ACTS ISDN (SAT) path. These data show that there is no essential difference between the terrestrial and satellite ISDN paths using the bit-error simulator. The only noted difference is that the quality of the terrestrial ISDN (TI) channel in the NIST-to-ITS path is higher than the same direction on the SAT channel at the lower average bit error rates—lower than  $10^{-5}$ .

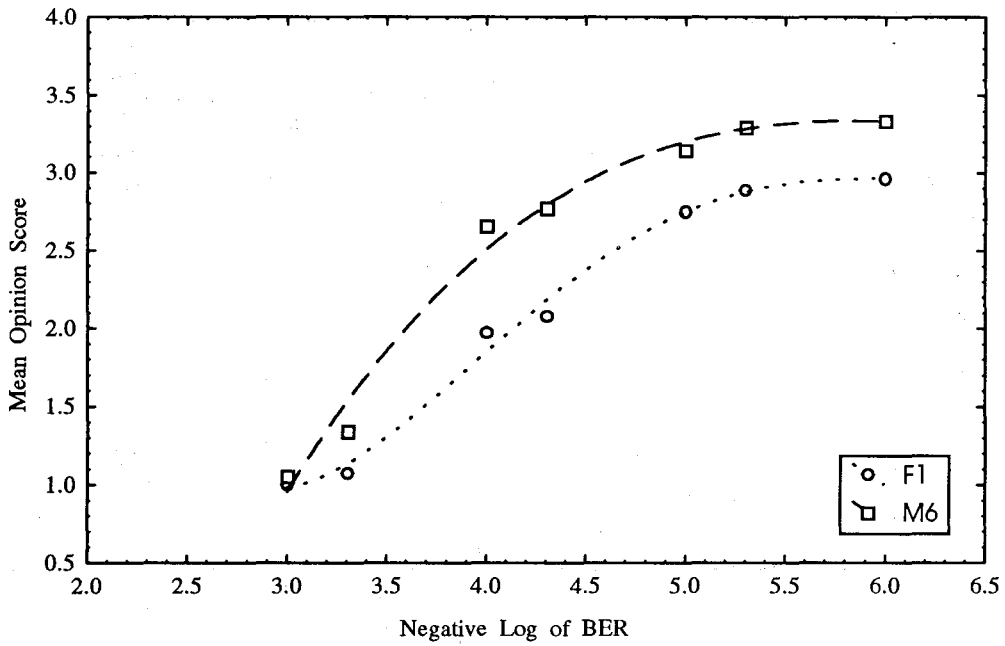


Figure 2.4. Voice quality through the local ISDN switch (LS) using the bit-error simulator and the laptop receiver.

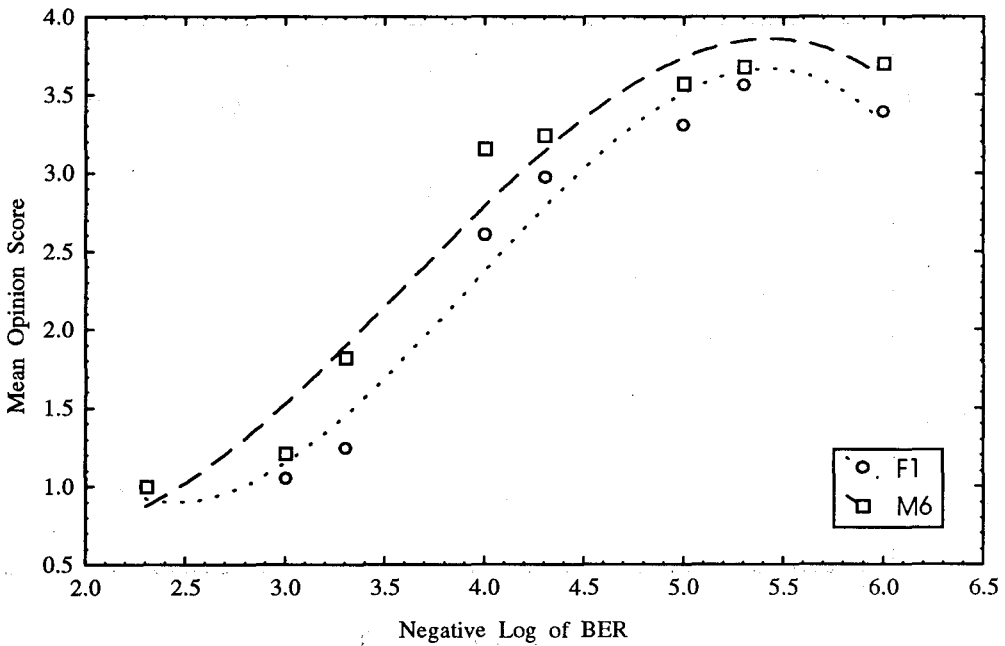


Figure 2.5. Voice quality through the local ISDN switch (LS) using the bit-error simulator and the laptop receiver, reverse direction.




Figure 2.6. Voice quality of the terrestrial ISDN line (TI) using the bit-error simulator from ITS to NIST.




Figure 2.7. Voice quality of the terrestrial ISDN line (TI) using the bit-error simulator from NIST to ITS.

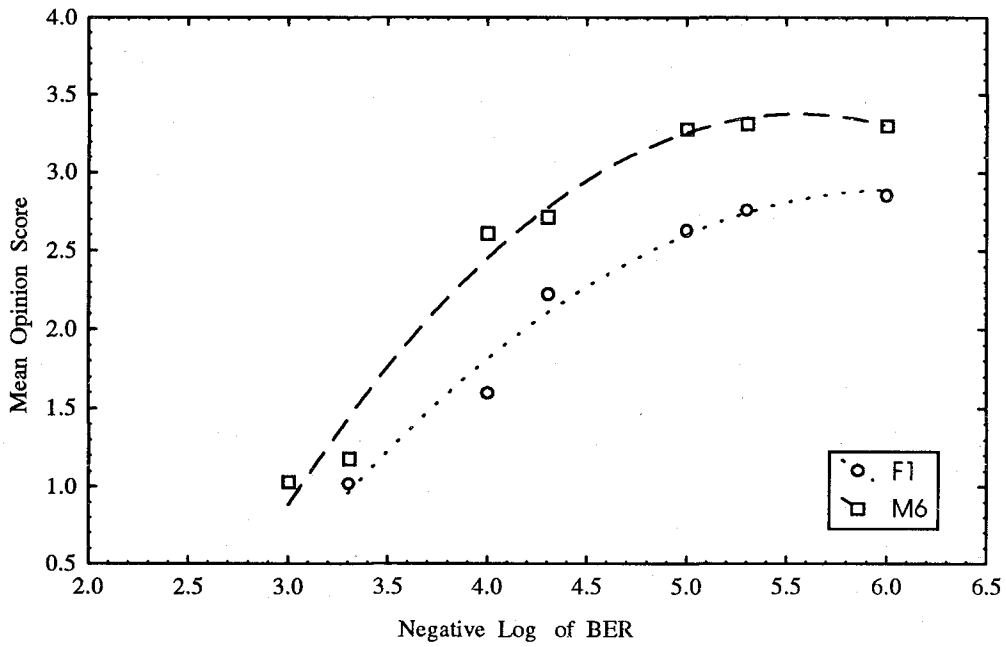


Figure 2.8. Voice quality of the ACTS ISDN channel (SAT) using the bit-error simulator from ITS to NIST.




Figure 2.9. Voice quality of the ACTS ISDN channel (SAT) using the bit-error simulator from NIST to ITS.

In general, the operators agreed with the MOS reported by the VQAS. The operators monitored the test sentences over loudspeakers as the nondegraded and degraded test sentences were received by the VQAS. With the system degraded by bit errors, the VQAS would occasionally report 1.00 (bad or not usable) since it could not properly synchronize the recorded sentence with the received sentence. If the operators' observation was in significant disagreement with VQAS, then additional measurements would be made. This helped to ensure a useable average that was indicative of the channel quality rather than of the ability of the VQAS to synchronize properly.

### Simulated Satellite Wander

Figure 2.10 shows the results of voice quality measurements over the LS channel at Boulder; the lines connecting the data points were added to make the graph easier to read. Measured BER's from a satellite wander incident (see Section 2.3.2) are used to degrade the communication channel by means of the bit-error simulator. VQAS was used to measure the quality of the channel with each change in the BER (see Table 2.5). The figure shows that there was indeed significant degradation in voice quality during the simulation. However, the operators monitoring the test sentences noted that, although there was noticeably significant degradation, the test sentences were never so garbled (from degradation) that the words could not be understood. The operators did agree that the assigned quality scores were reasonable.




Figure 2.10. Voice quality through the local ISDN switch (LS) during simulated satellite wander.



Figure 2.11 presents the same measurements for the satellite wander simulation over the SAT channel in the path ITS-to-NIST. The graph shows a similar response as the path through the LS channel. The degradation over the ACTS channel is not as pronounced and the general quality is slightly higher. A faint echo was noticed during these tests. This was probably due to an echo cancellation feature that is not invoked by the ACTS systems. No other channel exhibited an echo during the VQAS tests.

### Simulated Thunderstorm

Figure 2.12 shows the results of voice quality measurements over the local ISDN switch during a simulated thunderstorm (see Section 2.3.2). The intent was to simulate a thunderstorm at one of the Earth stations in a satellite link. The lines connecting the data points serve to indicate the changes in quality and to highlight the breaks. The breaks in the line signify that the BER was so great that the ISDN link was lost. Note that the line was re-established at a BER of  $10^{-9}$  before establishing the BER for the next degraded measurement. The three data points showing a quality of 1.0 may not be an accurate reflection of the quality of the line. The operators monitored the test sentence M6 and noted that although the line was clearly degraded with static, the sentence was clear enough to hear all the words and to decipher the meaning. The quality of 1.0 may only indicate that the VQAS failed to synchronize correctly. The severe changes in the graph indicate that voice quality may change dramatically due to rain from a thunderstorm of this magnitude. Except for the previously mentioned data points of 1.0, the operators generally report agreement with the VQAS results.




Figure 2.11. Voice quality through the ACTS ISDN channel (SAT) during simulated satellite wander incident from ITS to NIST.




Figure 2.12. Voice quality of ISDN channel through local switch during simulated thunderstorm.

Figure 2.13 shows the results of the same simulation over the ACTS ISDN (SAT) channel in the NIST-to-ITS path. The results are essentially the same as the ISDN channel through the local switch during a simulated thunderstorm and all comments above apply to these results. The only additional comment is to note the presence of the echo due to lack of echo cancellation. However, again, the echo does not appear to have any effect on the measurements.

## 2.5 Conclusions

This ACTS Collaboration experiment showed that the voice communications quality over a satellite ISDN connection was as good as or even better than the terrestrial ISDN counterpart. The delay over the satellite path was not a detriment to the one-way quality testing over the VQAS systems. Although noticeable, the delay did not hamper communications among the ACTS Collaboration test operators.

The VQAS measurements also show that the quality of the voice calls over the ACTS ISDN channel was more stable than that for the other channels. Variance of the quality measurements was small and generally consistent. The worst channel for variance was the terrestrial ISDN in the path ITS-to-NIST. The variance on this particular channel precluded collaborators from obtaining the desired statistical precision in a reasonable amount of time.




Figure 2.13. Voice quality through the ACTS ISDN channel (SAT) during simulated thunderstorm, NIST-to-ITS path.

Collaborators saw a distinct difference in quality measures between directions of the same path. It was felt that this difference was due to path differences in the communications equipment associated with the VQAS and the ISDN equipment and switch. The path difference was noted most for the TI channel. The ITS-to-NIST path was not as good as the reverse path in both quality and stability of the measurements.

VQAS is a useable tool for measuring voice quality over telecommunications channels. The inability of test sentence F5 to synchronize when measured over the ISDN systems and the regular long-distance channel is a cause for concern. This also occurred over the regular long-distance channel with sentences F1 and F2.

A significant improvement for the VQAS would be to automate the system so that:

- a single operator (computer) could control the VQAS testing at both ends of the communications channel under test (with VQAS systems at both ends),
- system synchronization could occur automatically, perhaps with signaling tones (eliminating the need for two persons at each end), and
- the addition of automated statistical processing could establish the number of measurements for a particular channel based on the previous measurements.



### **3. APPLICATION EXPERIMENT ON DESKTOP CONFERENCING**

Desktop conferencing is a highly interactive application that involves real-time voice, video, and application sharing. The real-time performance and correct transmission are both important factors for usable desktop videoconferencing. In this experiment, each of the three components (voice, video, and application sharing) were analyzed for usability and stability. Additionally, a terrestrial baseline was established for evaluating changes in application performance when a satellite link was introduced.

#### **3.1 Experiment Objectives**

The experiments were intended to identify the performance issues that should be considered when a satellite link is included in the communications system. The objectives of these experiments were not intended to evaluate the COTS implementations.

#### **3.2 Experiment Methods and Procedures**

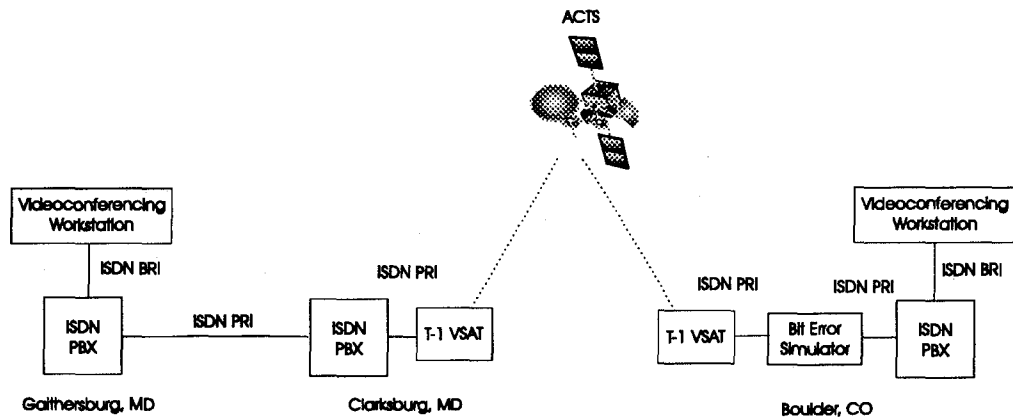
Intel Proshare<sup>®</sup> desktop-conferencing units, installed on 66-MHz/486 personal computers, were used for both proprietary and H.320 [3] experiments. In addition, one of the Proshare<sup>®</sup> units was replaced with a Picture Tel<sup>®</sup> H.320 system to test for similar behavior and interoperability. These units represent typical COTS user equipment available and already widely distributed today. The NS/EP scenario discussed in Section 1 suggests that existing user equipment that has lost telecommunications connectivity may have its connectivity restored via a satellite link without modification to the user equipment. This user equipment most likely was not designed for use with a satellite nor the latest technology. However, the function it provides may be more than adequate for daily operations, and may be critical to daily business. The user equipment selected for these experiments is intended to model such an environment.

The experiments were conducted using the two equipment configurations illustrated in Figure 3.1. For terrestrial experiments, two desktop-conferencing units were connected to a local ISDN switch (Teleos<sup>®</sup> Network Hub) with BRI ISDN. The ISDN switches were interconnected via PRI ISDN service from FTS2000. A bit-error simulator (Adtech SX-12) was connected between the FTS2000 PRI and the ISDN switch at the Boulder end of the circuit to inject errors into the data stream in both directions independently.


For the ACTS-based experiments, the same two desktop-conferencing units were connected to the same local ISDN switches with BRI ISDN service in Boulder (NTIA/ITS) and Gaithersburg (NIST). The Gaithersburg ISDN switch was connected via dedicated ISDN PRI service to an ISDN switch in Clarksburg (COMSAT). The Boulder and Clarksburg ISDN switches were interconnected via PRI ISDN service from the ACTS T1-VSATs in Boulder and Clarksburg. A bit-error simulator was

connected between the T1-VSAT PRI and the ISDN switch at the Boulder end of the circuit to inject errors into the data stream in both directions independently.

The bit-error simulator was used to simulate error scenarios that might be observed over a link in actual operations. Two scenarios were investigated. The first was based on a "steady-state" environment. For this scenario, the statistics of the random bit errors were uncorrelated and Gaussian distributed for the duration of the experiment. This simulates the errors caused by noise from a variety of background sources. Thresholds for usability and application failure were determined. The second scenario was based on a transient error environment that causes severe error bursts for some period. These tests were intended to determine the stability/recoverability of the applications to transient error rates. A moderate error background was used with a periodic severe burst of random errors of a fixed duration. Thresholds for recovery/failure were determined by varying the burst length and intensity.



**ACTS Equipment Configuration**



**Terrestrial Equipment Configuration**

Figure 3.1. Equipment configurations for application experiment on desktop conferencing.

### 3.3 Metrics

**Opinion Score:** An opinion score (OS) was used to evaluate the subjective usability for the audio, video, and whiteboard components of the videoconferencing applications. The evaluator is assumed to be a person familiar with operation of the component being evaluated. In these experiments, the experiment operator was the only evaluator. The rating scale is defined in Table 3.1 (fractional scores were not excluded from use).

Table 3.1. Definition of Opinion Score Values

Opinion Score	Description
5 - Excellent	No difference between in-person and remote perception.
4 - Good	Subtle differences between in-person and remote perception.
3 - Fair	Obviously a remote reconstruction, but not disturbing to the user in most cases.
2 - Poor	Somewhat disturbing differences in reconstruction but still usable.
1 - Unusable	Use of system is not practical.

The following examples from everyday experience serve as guidelines for OS ratings:

- OS = 3 is the audio of a toll quality telephone connection;
- OS = 3.5 is the video quality from NTSC video (U.S. broadcast television); and
- OS = 5 is equivalent to two users sitting side-by-side editing a document on a personal computer.

**Tracking response time:** This is the time for the initiator of an event to observe the response of the recipient. This metric is an indicator of the ability of the system to support real-time interaction between two users.

**Video error characterization:** Descriptions of video errors were based on ANSI T1.801.01-1995. [4].

### 3.4 Experimental Procedure

The following procedure was followed:

1. Observe link signal quality via  $E_b/N_0$  at Earth stations to verify low link BER. (A terrestrial link is assumed to have a BER less than  $10^{-9}$ .)
2. Establish a desktop-conferencing session.
3. Set the BER for the experiment via the bit-error simulator (see Tables 3.2 and 3.3).
4. Assess video and audio quality via subjective OS, noting the nature of any audio and/or video distortions. Record OS and observations.
5. If available in the implementation, establish an application-sharing session (e.g., whiteboard). If not available, skip steps 6 and 7. (Note: H.320 does not provide for application sharing.)
6. User 1 quickly moves the cursor from one corner of the screen to the diagonally opposite corner. User 2 moves the cursor to attempt to follow the motion. User 1 records the time from User 1's initial motion to when User 2's cursor arrives at the final position via a stopwatch (tracking response time). Record subjective observations on behavior and an OS for usability.
7. User 1 continuously moves the cursor in a random pattern around the screen. User 2 attempts to follow. User 1 notes the ability of User 2 to follow a similar route. Record subjective observations on behavior and an OS for usability.

Table 3.2 defines the set of BER configurations used in step 3 of the above procedure for steady-state experiments.

Table 3.3 defines the set of BER configurations used in step 3 of the above procedure for burst BER's. These experiments characterize the transient behavior (e.g., failure/recovery) of the application in response to changes in the BER. The burst duration and the burst gap were periodic (i.e., not random). This permits viewing the experiment as a collection of burst events that could be observed independently to note transient behavior to a single burst event. The background BER is the bit error density between bursts. The burst BER is defined as the bit error density during an error burst event. The burst gap is defined as the interval between the end of one error burst and the beginning of the next. The burst interarrival time is the sum of the burst duration and the burst gap.



Table 3.2. Steady State BER Configurations

Measurement Number	BER
1	0
2	$10^{-9}$
3	$10^{-6}$
5	$10^{-5}$
5	$10^{-4}$
6	$10^{-3}$
7	$10^{-2}$

Table 3.3. Burst BER Configurations

Measurement	Background BER	Burst BER	Burst Duration (s)	Burst Gap (s)
1	$10^{-6}$	$10^{-3}$	0.5	10
2	$10^{-6}$	$10^{-3}$	1.0	10
3	$10^{-6}$	$10^{-3}$	5.0	10
4	$10^{-6}$	$10^{-3}$	10.0	10
5	$10^{-5}$	$10^{-2}$	0.5	10
6	$10^{-5}$	$10^{-2}$	1.0	10
7	$10^{-5}$	$10^{-2}$	5.0	10
8	$10^{-5}$	$10^{-2}$	10.0	10

### 3.5 Expected Results

Changes in picture quality due to introduction of a satellite link are likely to be manifested in various ways. Compression algorithms tend to have characteristic visible error signatures. For example, bit errors in a discrete cosine transform (DCT)-compressed picture result in distorted blocks in the picture. Lost frames produce a “jerky” picture that may be observed as a “frozen” frame. Errors in a system using differential frame updates may produce a portion of the frame with persistent distortions (object retention) until that portion is updated again. A group of several events may be recorded as a single “severe” event, as that would be the user’s perception.

Satellite propagation delays may cause some difficulty in real-time whiteboard interactions. This difficulty arises from the time required for the parties to observe each other’s actions. Also, the voice and video information is transmitted with higher priority than the data. This may cause actions on the whiteboard to lose synchronization with the voice and video information.

The experiments involving error bursts characterize the transient behavior (e.g., failure/recovery or interruption) of the application in response to changes in the BER. Longer burst durations are likely to cause application and/or link failure due to expiration of timers, while short bursts will probably make the application unusable (i.e., audio dropout, video frame loss, and suspended data transmission) during the burst, but recover. In the case of the error bursts with recovery, the experiment identifies the recovery behavior of the system.

### 3.6 Results

The graphs in Figures 3.2 and 3.3 illustrate the usability opinion scores collected for proprietary Proshare® videoconferencing over both terrestrial and satellite (ACTS) links, respectively. These scores were based on experiments with a BER defined as the mean of Gaussian bit error arrival times, which simulates errors from many independent noise sources. The statistics for these errors are time-invariant for the duration of the experiment. These experiments describe the steady-state application performance for a given environment.




Figure 3.2. Terrestrial videoconference usability opinion scores.




Figure 3.3. ACTS videoconference usability opinion scores.

Figures 3.4 and 3.5 illustrate the usability opinion scores collected for H.320 Proshare<sup>®</sup> videoconferencing over both terrestrial and satellite (ACTS) links, respectively. These scores were based on experiments with Gaussian bit errors that simulate errors resulting from many independent noise sources. The statistics for these errors are time-invariant for the duration of the experiment. These experiments describe the steady-state application performance for a given environment.

No whiteboard data was collected for the H.320 experiments as the H.320 recommendation does not support application sharing, it only supports audio and video.

A subset of these experiments also was performed between a Proshare<sup>®</sup> H.320 system and a Picture Tel<sup>®</sup> system to test for differences in performance with different implementations and interoperability. No differences were observed in behavior of either H.320 implementation from a performance or an interoperability perspective.

Table 3.4 summarizes the results of the simulated error burst experiments conducted with the proprietary system over ACTS. In all cases, the burst intensity was severe enough to render the application unusable during the burst. The recovery time indicates the time needed after return to the background BER for normal performance at the background BER. No whiteboard measurements were reported since the measurement intervals could not be determined reliably. However, the recovery of the whiteboard appeared to be similar to that of the video (i.e., updates occurred within a few seconds of the end of the noise burst).




Figure 3.4. Terrestrial H.320 videoconference usability opinion scores.




Figure 3.5. ACTS H.320 videoconference usability opinion scores.

Table 3.4. Recovery Response to Noise Burst (Proprietary, ACTS)

Background BER	Burst BER	Audio Recovery Time (s)	Video Recovery Time (s)	Burst Duration for Total Failure (s)
$10^{-5}$	$10^{-2}$	5	5	5
$10^{-5}$	$10^{-3}$	5	5	5
$10^{-6}$	$10^{-2}$	1	5	>10
$10^{-6}$	$10^{-3}$	1	5	10 (only video failed)

Table 3.5 summarizes the results of the simulated error burst experiments conducted with the proprietary system over a terrestrial connection.

Table 3.5. Recovery Response to Noise Burst (Proprietary, Terrestrial)

Background BER	Burst BER	Audio Recovery Time (s)	Video Recovery Time (s)	Burst Duration for Total Failure (s)
$10^{-5}$	$10^{-2}$	<1	2	0.5
$10^{-5}$	$10^{-3}$	<1	2	5
$10^{-6}$	$10^{-2}$	<1	2	>10
$10^{-6}$	$10^{-3}$	<1	2	10 (only video failed)

Table 3.6 summarizes the results of the simulated error burst experiments conducted with the H.320 system over an ACTS connection.

Table 3.6. Recovery Response to Noise Burst (H.320, ACTS)

Background BER	Burst BER	Audio Recovery Time (s)	Video Recovery Time (s)	Burst Duration for Total Failure (s)
$10^{-5}$	$10^{-2}$	static during burst only	2	0.5
$10^{-5}$	$10^{-3}$	static during burst only	2	5
$10^{-6}$	$10^{-2}$	static during burst only	2	>10
$10^{-6}$	$10^{-3}$	static during burst only	2	10 (only video failed)

Table 3.7 summarizes the results of the simulated error burst experiments conducted with the H.320 system over a terrestrial connection.

Table 3.7. Recovery Response to Noise Burst (H.320, Terrestrial)

Background BER	Burst BER	Audio Recovery Time (s)	Video Recovery Time (s)	Burst Duration for Total Failure (s)
$10^{-5}$	$10^{-2}$	1	2	>10
$10^{-5}$	$10^{-3}$	1	2	>10
$10^{-6}$	$10^{-2}$	1	2	1
$10^{-6}$	$10^{-3}$	1	2	>10

Figures 3.6 and 3.7 illustrate the tracking response time data collected for proprietary Proshare® videoconferencing over terrestrial and satellite (ACTS) links, respectively. These scores were based on experiments with Gaussian-distributed errors. Response times were not measured beyond 15 s as such times were considered far outside reasonable real-time interactive limits for such a primitive action.

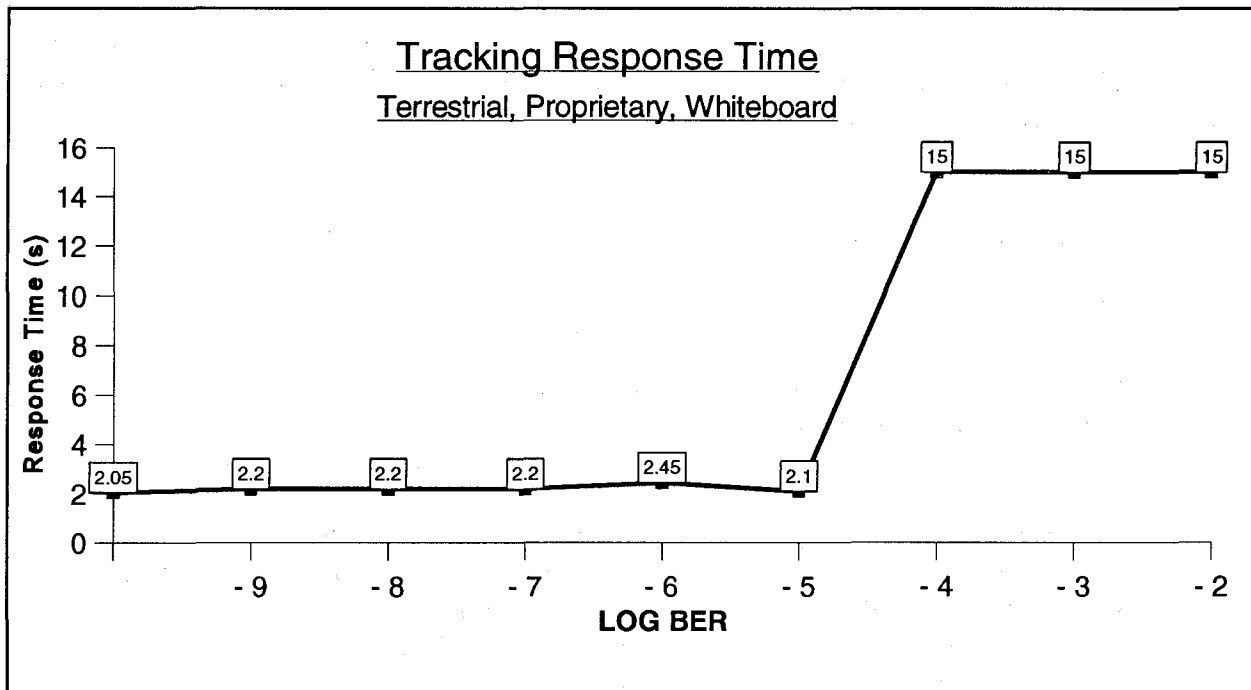


Figure 3.6. Terrestrial tracking response time.




Figure 3.7. ACTS tracking response time.

Audio delays for both terrestrial and ACTS experiments caused conversation conflicts. An experiment conducted between two units in the same room revealed that processing and transmission, using only a local ISDN switch, resulted in a one-way delay of about 1s.

### 3.6.1 Characteristics of Errors

Video:

The proprietary video processing appeared to detect frames with errors and did not display video frames containing errors. This resulted in a “frame freeze” for the last correct frame. This resulted in jerkiness, a still frame, or no image until the first good frame was received.

A typical H.320 video frame is shown in Figure 3.8. The H.320 video presented frames with errors. Two classes of errors were noted: error blocks and motion-related artifacts. Motion-related artifacts resulted in incorrectly translated portions of the frame (Figure 3.10) or “ghosting” (Figure 3.11). The block errors resulted in block distortion or tiling in the image or blurring typical of errors in a DCT-compressed image (Figure 3.9).



Additionally, when noise bursts occurred with a high ( $10^{-5}$ ) background BER, video recovery appeared to occur in multiples of 2 s. This was believed to be caused by background errors causing successive recovery cycles.



Figure 3.8. H.320 video frame with no errors.



Figure 3.9. H.320 video frame with block errors.



Figure 3.10. H.320 video frame with motion compensation distortion (translation).




Figure 3.11. H.320 video frame with motion compensation distortion (ghosting).

#### Audio:

Errors in the audio data stream of the proprietary system generally resulted in audio blanking. No noticeable noise of any other type was observed other than very minor loss of fidelity at lower error rates ( $<10^{-6}$ ).

In contrast, errors in the H.320 audio stream resulted in “static” which varied with the BER from an occasional “pop” or “crackle” beginning to be noticeable at error rates of  $10^{-6}$  to a roar of white noise that overwhelmed the audio signal at high error rates ( $>10^{-4}$ ).

The difference in the behavior appeared to be due to implementation choices, as both systems began to degrade at similar BER's.

#### Whiteboard:

Errors in the whiteboard data stream caused loss of interactive responses. Display updates became irregular, and in severe cases ceased, with a simulated BER in excess of  $10^{-5}$ .

#### General:

Service provided by the ISDN switch, the satellite signal processor, or FTS2000 failed when subjected to severe BER's. Connections were consistently lost with a BER of  $10^{-3}$ .

### 3.7 Analysis

The usability curves for proprietary and H.320 videoconferencing are very similar for both terrestrial and satellite configurations. For voice and video, both usability curves begin to degrade at a BER of  $10^{-6}$ . However, the video degrades more quickly, becoming unusable at a BER of  $10^{-4}$ . The voice remains usable at a BER of  $10^{-2}$  or greater.

The behavior of the whiteboard, similar to that of the video, begins to degrade at a BER of  $10^{-6}$  and becoming unusable at a BER of  $10^{-4}$ .

The difference in the audio usability between the proprietary and the H.320 systems is due to the different implementation approaches used. The proprietary system blanks the audio signal when noise is detected, causing audio gaps but preventing loud blasts. The H.320 system does not blank errors; the signal simply gets increasingly noisy until it is unintelligible.

The initial difference in the usability of the whiteboard component between satellite and terrestrial configurations was attributed by the users to increases in the interactive response time. This additional delay caused a minor increase in annoyance when attempting to perform highly interactive tasks as tracking the other user's cursor. However, the minimum interactive response time was 2 s. This is much greater than the increase in delay caused by satellite propagation (about 0.25 s).

Both proprietary and H.320 systems were extremely stable when subjected to severe simulated error bursts in either terrestrial or ACTS configurations. The videoconferencing systems rarely had fatal

errors even when subjected to the most severe simulated errors ( $10^{-6}$  background BER,  $10^{-1}$  burst BER for 10 s). In contrast, a burst sent into the terrestrial network with a BER of  $10^{-3}$  for greater than 5 s consistently resulted in a lost connection.

### 3.8 Interpretations

**1. BER, not delay, is the principal factor in videoconferencing usability.** The data indicate that all components of videoconferencing began to degrade with a BER in excess of  $10^{-6}$  and became essentially unusable with a BER of  $10^{-3}$ . The usability scores for satellite versus terrestrial communications services were essentially identical, indicating that satellite delay was not a significant factor. This was partly because the delay added by a satellite was small compared to the processing delays of the applications themselves. Therefore, a user that is accustomed to the delays incurred with the terrestrial services is not disturbed much by the increased delay incurred by the use of a satellite. A maximum BER of  $10^{-6}$  is recommended for reliable videoconferencing.

**2. Bandwidth management of various data streams may cause synchronization problems in highly interactive activities.** Voice and video take priority over data transfers in the Proshare<sup>®</sup> implementation and only BRI bandwidth is available for this application. Therefore, whiteboard activity tends to lag behind voice and video signals due to preemption of the whiteboard data. The voice and video signals must be synchronized to prevent very disturbing synchronization problems. The data are allowed to use any bandwidth that remains after voice and video requirements are satisfied. If major data transfers are required, the available data bandwidth is easily exceeded causing delayed updates in the data-based activities.

**3. Additional delay due to the satellite is an important factor in highly interactive activities.** Although the additional delay of a satellite link is only a fraction of the total transfer delay for this application, and therefore a minimal problem for most voice, video, and data uses, highly interactive uses (e.g., talking while using the cursor to point at objects with responses from the party at the other end) become increasingly difficult with increasing frequency of interaction.

**4. Terrestrial communications links may fail when sustained, severe BER's are present.** Data transmitted through the public terrestrial network with a BER in excess of  $10^{-3}$  that persist for several seconds resulted in lost connections. However, the applications themselves were able to remain stable with higher error rates when the errors did not propagate through the network, but were generated locally. This behavior may be attributed to an administrative decision in the terrestrial network to terminate connections that have persistent errors.



## **4. APPLICATION EXPERIMENT ON LAN BRIDGE**

These experiments evaluated the performance of TCP/IP-based applications where an ISDN LAN bridge was included in the communications link to transform two remote LAN's into a single logical LAN. TCP/IP applications involved both real-time interactive and noninteractive uses. Additionally, the communications service was based on virtual circuit or datagram service. In this experiment, each of these components was evaluated for usability and stability. A terrestrial baseline was established for evaluating changes in application performance when a satellite link was introduced.

### **4.1 Experiment Objectives**

The experiments were intended to identify the performance issues that should be considered when a satellite link is included in the communications system. The objectives of these experiments were not intended to evaluate the COTS implementations.

### **4.2 Experiment Methods**

A 66-MHz/486 personal computer and a Sun<sup>®</sup> workstation were used as the end systems for the LAN bridge experiments. These systems were selected because they represent commonly available workstations that were capable of sustained TCP/IP transfer rates significantly greater than could be supported by the LAN-bridge.

The experiments were conducted using the two equipment configurations illustrated in Figure 4.1. For terrestrial experiments, two workstations were connected to a local Ethernet segment in Boulder and Gaithersburg. The Ethernet segments were bridged via Combinet Interchange<sup>®</sup> LAN bridges communicating with each other via proprietary protocols over an ISDN BRI connection from a local ISDN PBX (Teleos Network Hub). The PBX's were interconnected via PRI ISDN service from FTS2000 (Network A). A bit-error simulator (Adtech SX-12) was connected between the FTS2000 PRI and the PBX at the Boulder end of the circuit to inject errors into the datastream in both directions independently.

For ACTS-based experiments, the same two workstations were connected to the same LAN bridges and PBX's in Boulder and Gaithersburg. The Gaithersburg PBX was connected via dedicated ISDN PRI service to a PBX in Clarksburg (COMSAT). The Boulder and Clarksburg PBX's were interconnected via PRI ISDN service from the ACTS T1-VSAT's in Boulder and Clarksburg. A bit-error simulator was connected between the T1-VSAT PRI and the PBX at the Boulder end of the circuit to inject errors into the datastream in both directions independently.

The bit-error simulator was used to simulate error scenarios that might be observed over a link in actual operations. Two scenarios were investigated. The first was based on a "steady-state" environment. For this scenario, the statistics of the random bit errors were uncorrelated and Gaussian

distributed for the duration of the experiment. This simulates the errors caused by noise from a variety of background sources. Thresholds for usability and application failure were determined. The second scenario was based on a transient error environment that causes severe error bursts for some period. These tests were intended to determine the stability/recoverability of the applications to transient error rates. A moderate error background was used with a periodic severe burst of random errors of a fixed duration. Thresholds for recovery/failure were determined by varying the burst length and intensity.


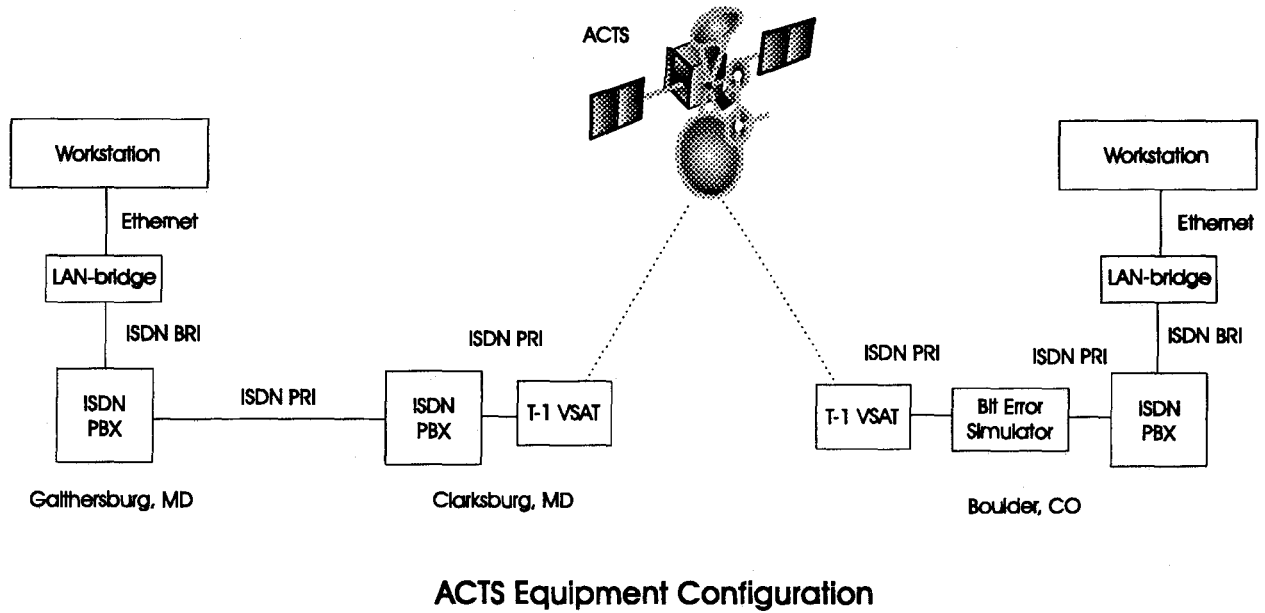


Figure 4.1. Experiment equipment configuration.

### 4.2.1 File Transfer Protocol

The file transfer protocol (FTP) experiments are designed to evaluate performance of bulk data transfers. These transfers are typically not time-critical, as in a real-time interactive application. However, correctness and completeness of the data are typically important.

A 278,507-byte image was used as the bulk data. To give some indication of the information that may be contained in such a transfer, the image is shown in Figure 4.2.




Figure 4.2. File used in FTP experiments (actual file is a color image).

### 4.2.2 Telnet

Telnet experiments are designed to evaluate performance of interactive keystroke-oriented applications. Such applications become “annoying” if the response time is excessively long. Correctness is important, but often not critical since rekeying can resolve the error.

### 4.2.3 Ping

The ping experiments are designed to evaluate the performance of datagram-oriented applications. A datagram is the underlying transport mechanism for most TCP/IP communication. These experiments determined the reliability of the datagram service by measuring the packet loss rate.

## 4.3 Metrics

**Opinion Score (OS):** An OS was used to evaluate the subjective usability for the FTP and Telnet applications. The rating scale was defined as found in Table 4.1 (fractional scores were not excluded from use).

Table 4.1. Definition of Opinion Score Values

Opinion Score	Description
5 - Excellent	Incapable of discriminating between LAN* and bridged LAN.
4 - Good	Subtle differences between LAN and bridged LAN performance.
3 - Fair	Obviously a bridged LAN, but not disturbing to the user in most cases.
2 - Poor	Somewhat disturbing differences in performance, but still usable.
1 - Unusable	Use of system is not practical or useful.

\* References to "LAN" performance assume a system where expected throughput typically exceeds 1 Mb/s and ping round-trip delays of about 2 ms.

**Packet Loss Percentage:** The fraction of packets that were transmitted, but not received.

**Throughput:** The number of bits per second (b/s) successfully transmitted to the destination.

#### 4.4 Experimental Procedure

The following sections describe the procedures used in the FTP, ping and TELNET experiments.

##### 4.4.1 FTP

1. Observe link signal quality via  $E_b/N_0$  at Earth stations to verify low link BER. (A terrestrial link is assumed to have a BER less than  $10^{-9}$ .)
2. Set BER for the experiment via the bit-error simulator (Tables 4.2 and 4.3).
3. Issue commands to initiate file transfer.
4. Record transmission statistics: elapsed time and transfer rate.
5. Check file transmitted for any errors in the data. Record number of bytes in error.
6. Assess FTP usability via subjective OS including any other observations. Record OS and observations.



Table 4.2. Steady State BER Configurations

Measurement Number	BER
1	0
2	$10^{-9}$
3	$10^{-6}$
5	$10^{-5}$
5	$10^{-4}$
6	$10^{-3}$
7	$10^{-2}$

Table 4.3. Burst BER Configurations

Measurement	Background BER	Burst BER	Burst Duration (s)	Burst Gap (s)
1	$10^{-6}$	$10^{-2}$	0.5	10
2	$10^{-6}$	$10^{-2}$	1.0	10
3	$10^{-6}$	$10^{-2}$	1.0	10
4	$10^{-6}$	$10^{-2}$	2.0	10
5	$10^{-6}$	$10^{-2}$	3.0	10
6	$10^{-6}$	$10^{-2}$	5.0	10
7	$10^{-6}$	$10^{-2}$	10.0	10
8	$10^{-5}$	$10^{-2}$	0.5	10
9	$10^{-5}$	$10^{-2}$	1.0	10
10	$10^{-5}$	$10^{-2}$	5.0	10
11	$10^{-5}$	$10^{-2}$	10.0	10

#### 4.4.2 Telnet

1. Observe link signal quality via  $E_b/N_0$  at Earth stations to verify low link BER. (A terrestrial link is assumed to have a BER less than  $10^{-9}$ .)
2. Set BER for the experiment via the bit-error simulator (Tables 4.2 and 4.3).
3. Issue command to establish Telnet session.
4. Type terminal commands, check data transmitted for any errors, and record number of bytes in error.
5. Assess Telnet usability via subjective OS including any other observations. Record OS and observations.

#### 4.4.3 Ping

1. Observe link signal quality via  $E_b/N_0$  at Earth stations to verify low link BER. (A terrestrial link is assumed to have a BER less than  $10^{-9}$ .)
2. Set BER for the experiment via the bit-error simulator (see Tables 4.2 and 4.3).
3. Issue commands to generate 100 pings, 64 bytes long.
4. Record transmission statistics: round-trip delay and percent packet loss.

Table 4.3 defines the set of BER configurations used in step 2 of the above procedures for burst BER configurations. These experiments characterize the transient behavior (e.g., failure/recovery) of the application in response to changes in the BER. The burst duration and the burst gap were periodic delays (i.e., not random). This permits viewing the experiment as a collection of burst events that could be observed independently to note transient behavior to a single burst event.

The background BER is the bit error rate between bursts. The burst BER is defined as the bit error rate during an error burst event. The burst gap is defined as the interval between the end of one error burst and the beginning of the next. The burst interarrival time is the sum of the burst duration and the burst gap.

### 4.5 Expected Results

The following summarizes the expected results for the experiments.

- Telnet and FTP usability will degrade and become unusable with increased bit error rate.
- Additional delay introduced by the satellite path may cause protocol timers to expire causing application failure.
- Undetected errors may be introduced into data transmitted.
- Satellite propagation delays may cause some difficulty in real-time Telnet response because character echos will be delayed.
- The experiments involving error bursts will characterize the transient behavior of the application to errors. Longer duration bursts are likely to cause application and/or link failure due to expiration of timers; short bursts will probably make the application unusable during the burst, but identify the recovery behavior of the system.

#### 4.6 Results

Figure 4.3 illustrates the throughput observations for FTP over both terrestrial and ACTS configurations. These scores were based on experiments with a BER defined as the mean of Gaussian bit error arrival times, which simulates errors from many independent noise sources. The statistics for these errors are time-invariant for the duration of the experiment. These experiments describe the steady-state application performance for a given environment.




Figure 4.3. LAN bridge FTP throughput results.

Table 4.4 summarizes the recovery behavior of FTP over a LAN bridge in the presence of error bursts over a terrestrial line.

Table 4.4. Burst BER Results Over the Terrestrial Line

Background BER	Burst BER	Burst Duration (s)	Burst Gap (s)	Behavior
$10^{-6}$	$10^{-2}$	0.5	10	Recovery
$10^{-6}$	$10^{-2}$	1.0	10	Recovery
$10^{-6}$	$10^{-2}$	2.0	10	Recovery
$10^{-6}$	$10^{-2}$	3.0	10	Failure
$10^{-6}$	$10^{-2}$	5.0	10	Failure
$10^{-6}$	$10^{-2}$	10.0	10	Failure
$10^{-5}$	$10^{-2}$	0.5	10	Recovery
$10^{-5}$	$10^{-2}$	1.0	10	Recovery
$10^{-5}$	$10^{-2}$	2.0	10	Recovery
$10^{-5}$	$10^{-2}$	3.0	10	Failure
$10^{-5}$	$10^{-2}$	5.0	10	Failure
$10^{-5}$	$10^{-2}$	10.0	10	Failure

Table 4.5 summarizes the recovery behavior of FTP over a LAN bridge in the presence of error bursts over ACTS.

Figure 4.4 illustrates the usability opinion scores for the FTP, Telnet, and Ping experiments in the ACTS configuration.

Table 4.5. Burst BER Results Over ACTS

Background BER	Burst BER	Burst Duration (s)	Burst Gap (s)	Behavior
$10^{-6}$	$10^{-2}$	0.5	10	Recovery
$10^{-6}$	$10^{-2}$	1.0	10	Recovery
$10^{-6}$	$10^{-2}$	2.0	10	Recovery
$10^{-6}$	$10^{-2}$	3.0	10	Recovery
$10^{-6}$	$10^{-2}$	5.0	10	Failure
$10^{-6}$	$10^{-2}$	10.0	10	Failure
$10^{-5}$	$10^{-2}$	0.5	10	Recovery
$10^{-5}$	$10^{-2}$	1.0	10	Recovery
$10^{-5}$	$10^{-2}$	5.0	10	Failure
$10^{-5}$	$10^{-2}$	10.0	10	Failure

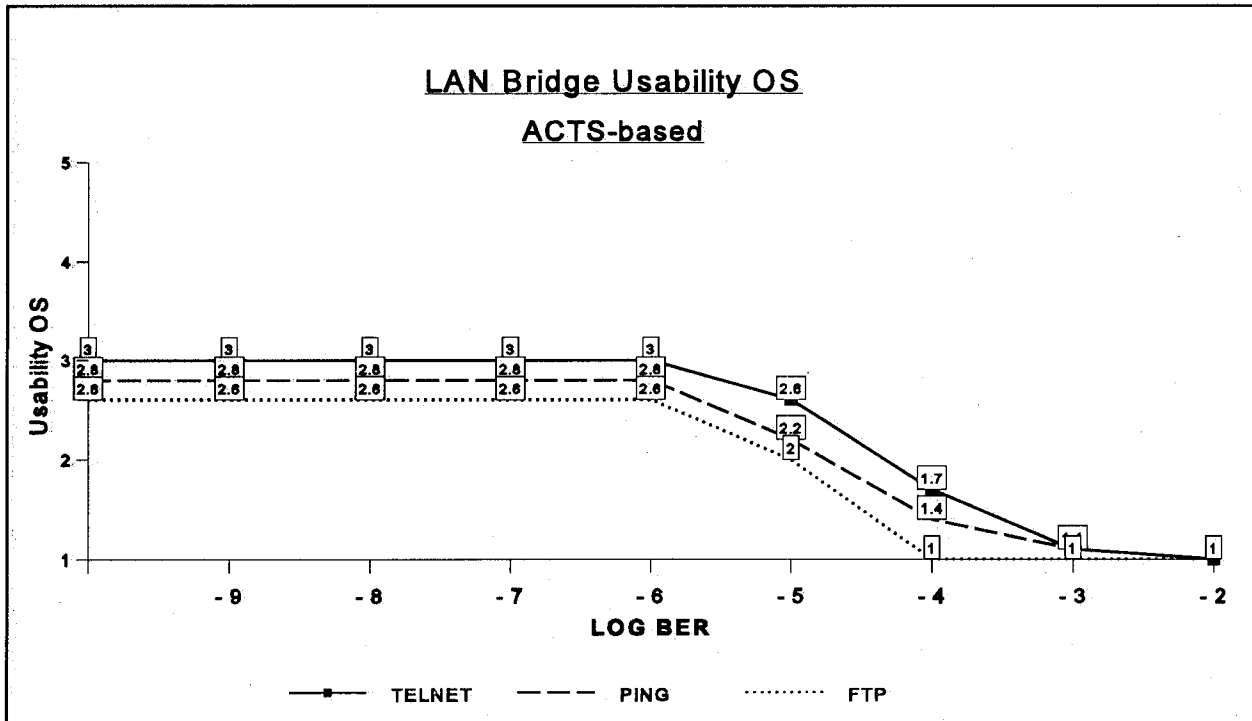


Figure 4.4. LAN bridge usability over ACTS.

Figure 4.5 illustrates the usability opinion scores for the FTP, Telnet, and Ping experiments in the terrestrial configuration.




Figure 4.5. LAN bridge usability opinion scores for the terrestrial line.

Figure 4.6 illustrates the percent packet loss for repeated pings as a function of BER for both terrestrial and ACTS configurations.

#### 4.7 Analysis

All applications tested began to degrade with error rates greater than  $10^{-6}$  and were not useful for error rates greater than  $10^{-3}$ .

##### 4.7.1 File Transfer Protocol

The curves for FTP throughput as a function of BER over both terrestrial and ACTS configurations exhibited parallel behavior. The throughput over ACTS was consistently less than that over a terrestrial line. This difference was due to the additional propagation delay causing throttling of the end-to-end protocols of the LAN bridge and FTP. This was confirmed through additional experiments

conducted with a delay simulator and no additional bit errors. No degradation of throughput was observed for one-way delays less than about 100 ms.




Figure 4.6. Ping packet loss results.

Other supplemental tests revealed that FTP becomes unstable with a BER of about  $10^{-4}$  or greater. The burst test revealed that FTP would recover if the BER improved to  $10^{-5}$  or less within 5 s. If the burst lasted longer, the transfer failed independent of the error density before, during or after the burst.

There were no errors in the transmitted data for any configuration in which the transmission completed normally. Abnormal termination resulted in a partial transfer of correct data.

The differences in the opinion scores for terrestrial and ACTS configurations were insignificant. Lower opinions resulted from bit errors causing extremely long transfer times, many times larger than that on an error-free connection. The main factor in usability was successful transmission, not differences in transmission speed.

#### 4.7.2 Telnet

Propagation delay in ACTS experiments caused some reduction in the opinion scores for Telnet, which resulted from increased delay for character echo. However, the character echo became

significantly larger for a BER in excess of  $10^{-4}$ . BER was the major factor that caused unusability. No transmission errors were observed in the echoed characters for any keystrokes, at any BER.

### 4.7.3 Ping

The ping experiments indicated insignificant packet loss differences between terrestrial and ACTS configurations. The packet loss statistics were only a function of BER.

## 4.8 Interpretations

**1. BER, not delay, is the principal factor in LAN bridge usability.** The data indicate that all applications used with a LAN bridge began to degrade with a BER in excess of  $10^{-6}$  and became essentially unusable with a BER of  $10^{-3}$ . The usability scores for satellite versus terrestrial communications services were essentially identical, indicating that satellite delay was not a significant factor for usability.

However, FTP throughput was impaired by the additional satellite delay at all error rates. This appears to be due to acknowledgment window sizes in the TCP/IP protocol and/or the LAN bridge protocol. This could be addressed through modifications in the protocols, as described in Section 5, to improve performance in the presence of delay. Experiments indicated that one-way delays less than 100 ms did not have any impact on throughput using the current protocols.

A maximum BER of  $10^{-6}$  is recommended for reliable LAN bridge service.

**2. Additional delay due to the satellite is an important factor in highly interactive activities.** While the additional delay of a satellite link is a minimal problem for file transfer and datagram delivery (ping), highly interactive uses (Telnet) become increasingly difficult with increasing frequency of interaction. If keystrokes are echoed from the remote system, the delay can become disturbing to a reasonably fast typist. This also could be a problem for other interactive applications such as remote graphic-based applications.

**3. While additional delay due to satellite reduces FTP throughput, a high BER makes FTP unusable.** While FTP throughput was reduced as a result of the additional satellite delay, FTP usability was impacted only minimally because of the noninteractive nature of bulk file transfers. The decreased performance was not significant to the user due to the already relatively long transfer time via terrestrial service. However, FTP was rendered unusable with a BER of  $10^{-4}$  or greater, independent of the presence of a satellite.

**4. Terrestrial communications links may fail when sustained, severe bit errors are present.** Data transmitted through the public terrestrial network with a BER in excess of  $10^{-3}$  that persists for several seconds resulted in lost connections. However, the applications themselves, except FTP, were



able to remain stable with higher error rates when the errors did not propagate through the network, but were generated locally. This behavior may be attributed to an administrative decision in the terrestrial network to terminate connections that have persistent errors.

FTP appeared to fail with severe, sustained error bursts due to expiration of FTP protocol timers during these bursts.

**5. Optimization of end-to-end protocols for channels with high propagation delay is needed for efficient use with satellites.** The FTP experiments indicated that propagation delays were causing degraded performance over a satellite channel. However, the degradation did not occur for delays less than 100 ms. This indicates that protocol parameters, such as window sizes, might be adjusted to improve performance over channels with longer delays. The protocol-oriented experiments, reported in Section 5, address this issue.



## 5. PROTOCOL EXPERIMENT ON TCP-LFN OVER FRAME RELAY

With the growth of the Internet, the TCP/IP protocol has become the most widely used protocol in use today. TCP/IP implementations are commonly found on almost every hardware platform/operating system with a wide variety of applications running over it. Originally designed in the 1960's, the protocol has evolved through the years to meet requirements of LAN's, WAN's, and other systems and networks.

When TCP is used over a satellite link, however, the large bandwidth-delay product can cause problems with throughput. TCP-LFN, an enhanced version of TCP for "long-fat" networks, attempts to rectify these problems [5]. The large bandwidth-delay product of a satellite channel requires larger window sizes to "keep the pipe full" in order to make full use of the channel capacity. Assuming the round-trip time over a satellite link is 0.6 s, the window size required to completely utilize a T1 link would be  $1,536,000 \text{ b/s} (0.6 \text{ s} / 8 \text{ b/B}) = 116 \text{ kB}$ . Most current TCP implementations have 64 kB for their maximum window size. TCP-LFN increases this limit to  $2^{31}-1 \text{ B}$ .

TCP retransmits packets when a retransmission timer expires. If this timer expires too soon, packets are unnecessarily retransmitted; if it expires too late, the pipe becomes empty during the intermediate period. Either way, this leads to wasted bandwidth. Current TCP implementations measure one round-trip time (RTT) per window in order to set the retransmission timer. As the window grows, the accuracy of this measurement degrades. TCP-LFN attempts to correct this by measuring an RTT per packet by time-stamping the packet. Another problem that arises for long delays is that the packet sequence numbers can wrap around (be reused). This means that a packet that arrives late has a higher probability of being mistaken for one that was transmitted later with the same sequence number, confusing TCP and delivering incorrect data to the application. TCP-LFN use of time-stamps eliminates the need for sequence numbers, providing protection against wrapped sequence (PAWS) numbers.

### 5.1 Introduction

The use of frame relay as a WAN protocol has grown significantly over the last few years. Frame relay can support several logical connections over the same physical connection. Although X.25 does this, it uses link-by-link recovery. This type of recovery is wasteful because it leads to lower throughput and requires more complex hardware and software at each network node. Also, with X.25, when the delay is large such as on a satellite link, packets are more likely to be lost and then retransmitted. Frame relay, on the other hand, performs end-to-end recovery leading to a simpler protocol, higher throughput and less complexity in each network node.

Older networks used static bandwidth management to allocate bandwidth to users. More recently, however, there has been a significant growth in the use of bandwidth-on-demand capability since that provides more efficient sharing of network resources, given the bursty nature of data. Also, users

observe less delay since the pipe expands as data flow increases. Users can be charged based on usage as opposed to paying a fixed charge decided at subscription time.

TCP-LFN running over frame relay with bandwidth-on-demand capability is a good solution for networks required for emergency operations that use satellite links. Satellite ground stations are easy to deploy during an emergency and the ability to adjust bandwidth to meet requirements at the emergency site is essential. Satellite networks themselves are less prone to failure since there are fewer physical links and it is possible to enhance link quality dynamically (by coding, etc.). Prioritization of messages is fairly simple to implement within a frame relay network, and is supported by most user devices (such as routers). Most of the user equipment is inexpensive and is commonly available from several vendors; this makes redundant sites more feasible to implement.

This experiment evaluated the performance of TCP-LFN over frame relay using the COMSAT - ACTS frame relay access switch (FRACS) over the ACTS network. The frame relay switches are capable of implementing bandwidth-on-demand over ACTS using the ISDN signaling interface on the ACTS Earth station.

## 5.2 Experiment Methods and Procedures

TCP packets generated by the application running on the Sun<sup>®</sup> workstation were encapsulated in IP packets and sent over the Ethernet to the WAN router, then they were encapsulated in frame relay packets and sent to the FRACS. The FRACS monitored data flow to each destination periodically and allocated bandwidth by obtaining usage information from the ACTS terminal. The packets were then sent to the ACTS terminal through the T1. The reverse process took place at the destination until the packets reached the application. Figure 5.1 shows the setup for the TCP-LFN over frame relay experiment.

Performance evaluation software tools like TTCP (public domain software) and TCPTGEN (COMSAT proprietary software) were run on the Sun<sup>®</sup> workstations. TTCP opens a TCP socket, sends a specified number of packets of specified size through it, and reports the average throughput measured at the end of the exchange. TCPTGEN is similar except that it measures throughput periodically and also reports the long-term average, therefore it is useful for observing TCP slow start and congestion control mechanisms in action.

The Solaris<sup>®</sup> (Sun<sup>®</sup> operating system) "netstat" command was used to obtain various TCP/IP statistics such as number of data packets sent, maximum segment size, number of retransmissions, retransmission timeout, and fast retransmissions. Another analysis program, SNOOP, was used to observe the following information about packet traces: the times at which each was sent, the outstanding transmit packets, and TCP/IP retransmission behavior.

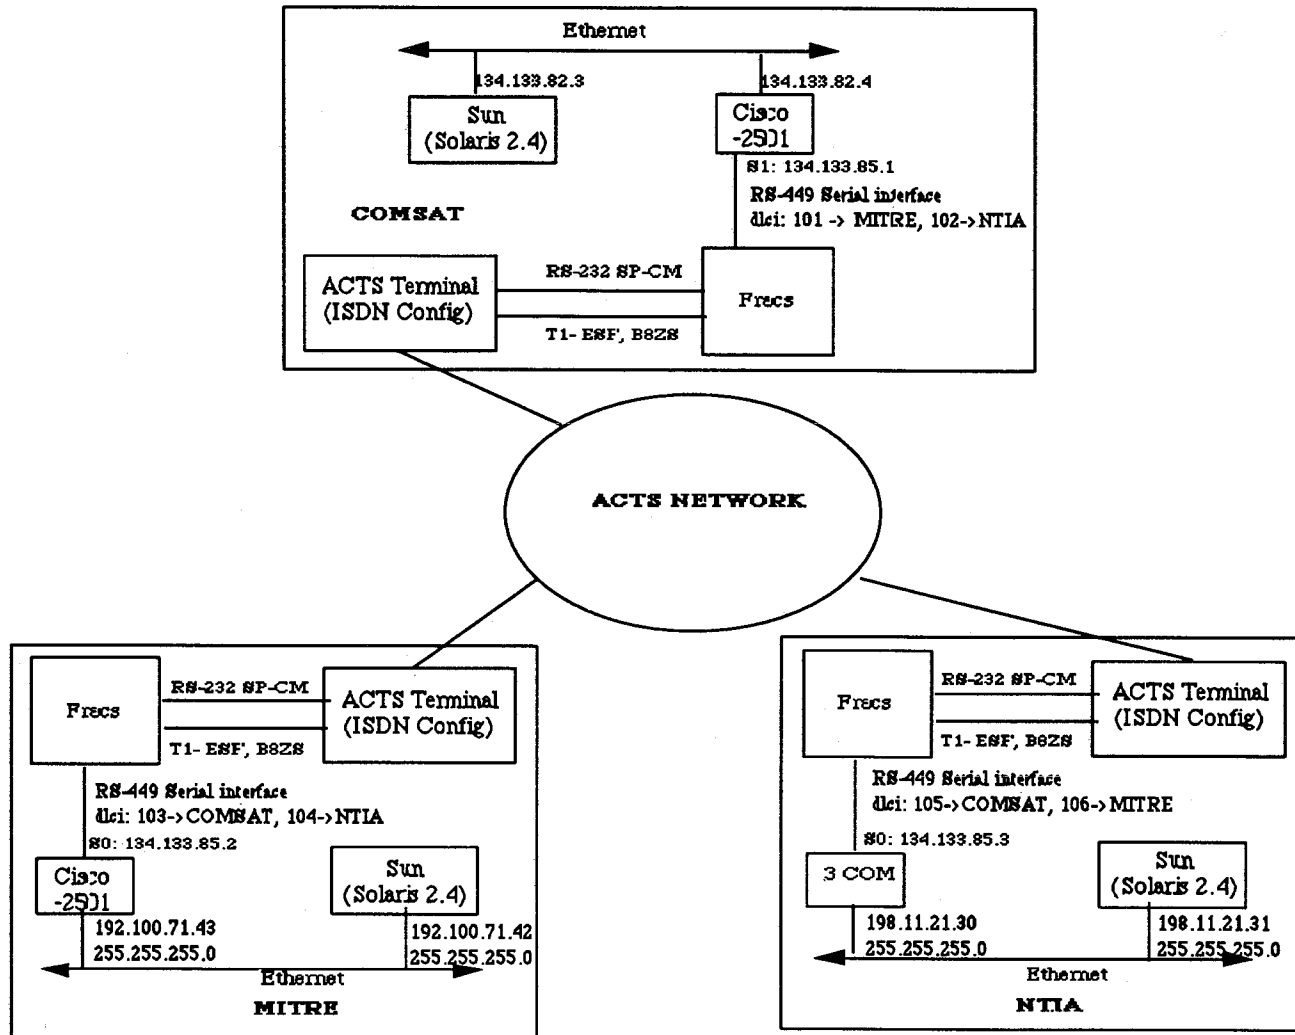


Figure 5.1. Experiment setup.

The FRACS itself also provided a number of valuable statistics, including traffic rate measurements, packet statistics at the data link and trunk levels, link quality measurements, bandwidth and call management statistics, and availability of buffers.

The WAN routers (CISCO and 3-COM) also provided some packet statistics at each of its interfaces, including transmitted and received packets, rate measurements, and buffer information.

### **5.3 Performance Analysis of TCP-LFN**

The performance capabilities of TCP and TCP-LFN were examined using both static and dynamic bandwidth management.

#### **5.3.1 Static Bandwidth Management**

Figure 5.2 shows the performance of TCP-LFN when static bandwidth management was used. Both the expected and the measured throughput are as seen by the application. The expected throughput was computed as follows for an application packet size of 1400 B: TCP adds 32 B of overhead, IP adds 20 B, the router adds 4 B, and the FRACS adds 48 B. Hence, the overhead is 104 B or 7%, and the expected throughput is 93% of the allocated bandwidth.

As can be seen in Figure 5.2, the measured throughput is fairly close to the expected throughput. The difference could be made even smaller by using TCP header compression and upgrading the FRACS to use a packet size of 1600 B.

#### **5.3.2 TCP-LFN vs. TCP**

Figure 5.3 shows the relative performance of TCP and TCP-LFN with static bandwidth management. The TCP default parameters curve shows the performance of TCP with the standard window size of 8 kB and a maximum (congestion) window size of 32 kB. The throughput is limited to 90 kb/s. Some TCP implementations (e.g., Solaris) allow the user to change some of the TCP driver parameters from their default values. When the transmit and receive window sizes were changed to 64 kB, and the maximum (congestion) window size was changed to 64 kB, the peak throughput increased to 577 kb/s. With TCP-LFN, however, the throughput increased in proportion to the allocated bandwidth.




Figure 5.2. TCP-LFN throughput for the static bandwidth management case.




Figure 5.3. Regular TCP vs. TCP-LFN throughput.

Note that several operating systems do not allow the user to change driver parameters from their default values and those that do often have limits lower than 64 kB (e.g., 50 kB in Solaris®) due to implementation/buffer constraints. Also, the application may have to be changed (e.g., “set socket” option in Solaris® when opening a TCP socket) or the operating system may have to be rebuilt to do this. Even if this were possible, the larger window sizes would be applied to every connection, not only to those being routed over the satellite. This would result in wasted memory at the hosts; it would also use a large number of buffers at the router, which would lead to longer delays for other connections through the router and less tolerance to short periods of congestion.

### 5.3.3 Dynamic Bandwidth Management – Preallocated Mode

Figure 5.4 shows the performance of TCP-LFN with dynamic bandwidth management. On the abscissa is the maximum allocated bandwidth. As can be seen in this figure, the measured and the expected TCP-LFN throughput (computed in the same way as in Section 5.3.1) are fairly close. This shows that TCP-LFN works well over a link where bandwidth is allocated on demand. During this measurement, there were no lost packets.

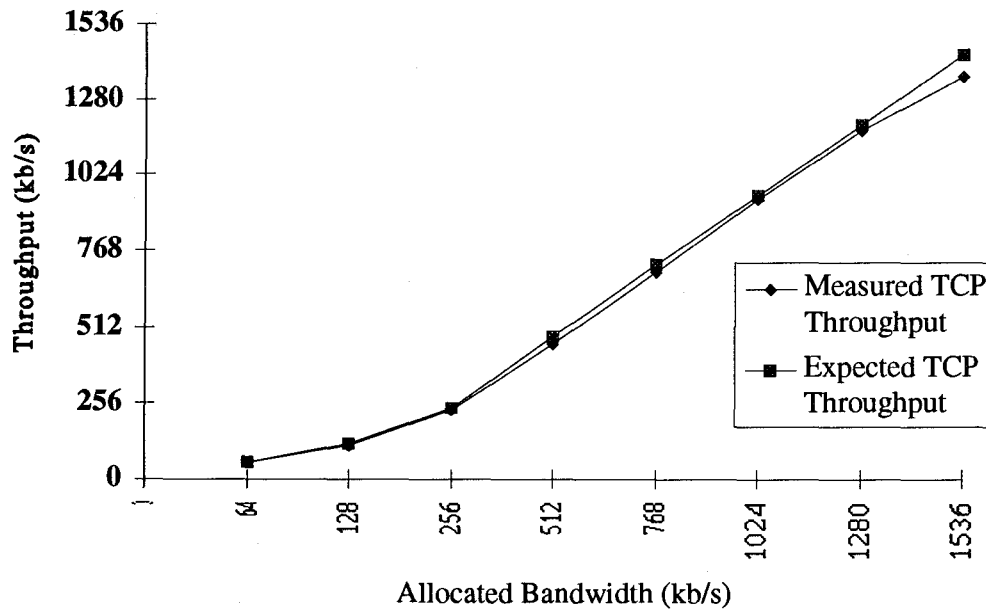


Figure 5.4. TCP-LFN preallocated mode throughput.



In the preallocated mode, a fixed number of channels were reserved from the source to the destination terminal by the FRACS, and these channels were made available to the TCP-LFN traffic as determined by the bandwidth management algorithm. There were several reasons for doing this as opposed to an actual bandwidth-on-demand operation. First, when the ACTS terminal informs the FRACS that a new call has been connected, this may not be true due to latency. Therefore, the first packet(s) sent through this channel may be lost. When this happens, TCP-LFN retransmits the lost packets and reduces the size of its congestion window. As a result the FRACS releases some of the allocated bandwidth. The cycle continues with the throughput rising to about 600 kb/s, and then dropping and rising again. Second, the time taken by the ACTS system to connect a new 64-kb/s call was fairly high (about 5 s). And finally, if there was more than one outstanding call request made to the system it would automatically stop. All of these effects taken together would result in a long delay before the TCP-LFN throughput would rise to its peak value. The preallocated mode alleviates these problems; the results obtained using this mode are the results expected if the above problems did not exist.

### 5.3.4 Effect of Slow Start

Slow start is a process in TCP-LFN that increases the window size each RTT if no retransmissions are required for the packets sent from the present window. Figure 5.5 shows the time taken by TCP to reach the peak throughput due to slow start (assuming no errors are observed during this period). The peak throughput information was obtained by reviewing the packet traces, finding the congestion window size, and cross-checking it with the periodic measurements done by TCPTGEN.




Figure 5.5. Effect of slow start on throughput.

In this case, several interesting observations were made. First, the round-trip time was 0.8 s, significantly more than the expected value of roughly 0.6 s. The FRACS uses 24 64-kb/s channels (since it may be splitting the T1 amongst several destinations). If a packet is segmented (the mean transmission unit (MTU) on the T1 is 256 B), all segments for the same packet are sent through the same channel. Given a packet size of 1504 B (1400 B of application data + 104 B of headers), the transmission delay is 0.18 s. The rest of the delay (a few milliseconds) is in the transmission delay for the acknowledgment, and the queuing delays in the router, Sun, and FRACS. In order to accommodate this delay, the window size was set to 154,000 B [1,536 kb/s (0.8s / 8 b/B)].

Assuming an initial congestion window and segment size of 1,500 B, a round-trip time of 0.8 s and assuming the congestion window doubles every round-trip time, the peak throughput of 1,536 kb/s (congestion window = 154,000 B) was expected to occur at about 5.6 s (7 round-trip times). As shown in the figure, it actually took 21.18 s for peak throughput to be reached. This is due to the way in which slow start is implemented. Every time the transmitter receives an acknowledgment, the congestion window is incremented by one segment size. This means that if the receiver sent one acknowledgment per data packet, the congestion window would double every round-trip time. However, the receiver sends one acknowledgment for every two (sometimes more) received packets when there is no data flowing in the reverse direction. As a result, the congestion window takes a lot longer to increase to its maximum value. If the transmitter accounted for the number of acknowledged bytes, the congestion window would reach its maximum much earlier. Given that the RTT of a satellite link is significantly larger, slow start severely limits the throughput for transfers that involve a small amount of information.

### **5.3.5 Effect of Link Errors**

Figure 5.6 shows the performance of TCP-LFN with an elevated BER. The errors were injected by putting a bit-error generator on the T1 link between the FRACS and the ACTS terminal.

When the BER was  $10^{-6}$  or worse, the throughput was degraded significantly. When a packet is lost, TCP's fast retransmit algorithm retransmits the packet immediately when a third duplicate acknowledgment is received. TCP then performs congestion avoidance, reducing its congestion window size and slow start threshold (and, hence, the throughput) to half the current value. It then slowly increases its throughput until it reaches the peak value or another error occurs. Using TCPTGEN, this fall and rise of throughput can be observed for each error.

## **5.4 Results of TCP-LFN over Frame Relay**

In this section, the performance of several applications running over TCP-LFN over frame relay, including web browsers, FTP, Telnet, and remote login (Rlogin) is presented.

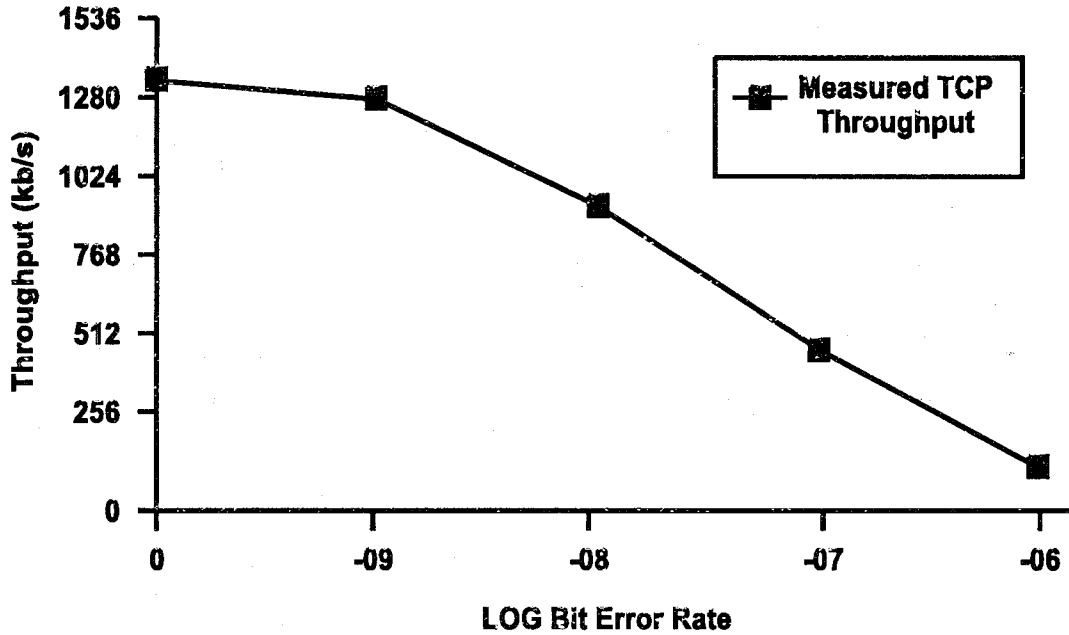


Figure 5.6. TCP-LFN throughput vs. BER.

#### 5.4.1 Web Browsers

Slow start severely limits the throughput seen by the user for small transfers. Since most browsers (like Netscape® and Mosaic®) typically open a connection, transfer data from the highlighted link, and then close the connection, the throughput seen by the user is much lower than the link speed. Better performance is seen while transferring large images that are at least several megabytes long.

#### 5.4.2 FTP Performance

Figure 5.7 shows the performance of FTP over TCP-LFN. As can be seen from the figure, the actual throughput at higher bandwidths was lower than that measured using TTCP and TCPTGEN. While the reasons for this were not completely understood, there are several possible reasons. Unlike TCPTGEN, the throughput could not be measured periodically, so the measured throughput was affected by slow start. The size of the file used for the measurements was 9 MB. If larger files had been used, higher throughput may have been observed. The accesses to the hard drive on the computer could also potentially have been a bottleneck at higher throughput.




Figure 5.7. FTP performance over TCP-LFN.

In comparison, when regular TCP was used with the window size set to 50 kB with an allocated bandwidth of 1.536 Mb/s, the measured throughput was 437 kb/s. When multiple FTP sessions were established, the following throughput values were measured.

Table 5.1. Throughput vs. Number of FTP Sessions

Number of FTP Sessions	Throughput (kb/s)
2	776
3	880
4	926

When the number of FTP sessions was increased to five, retransmissions occurred; at six FTP sessions, some packet loss occurred (due to lack of buffers). In both cases, the throughput was about 1 Mb/s.

### 5.4.3 Telnet and Rlogin

With both Telnet and Rlogin, the longer delay of the satellite link plays a significant role. Operation in the in-line mode can reduce the effects of delay; because this mode produces a local echo of the

characters typed at the terminal. This mode cannot be used for applications such as “vi” hence its usefulness is limited.

## 5.5 Conclusions

The following conclusions were drawn from the experiment:

The peak throughput that could be attained by TCP over a satellite link was limited to a few hundred kb/s. However, this occurred only when the TCP window sizes were changed to 64 kB. Using the default window size of 8 kB, the peak throughput was less than 100 kb/s. With TCP-LFN, the throughput scaled well with the allocated bandwidth up to 1.544 Mb/s (the limit for this experiment) and can be expected to do so into the gb/s range.

Slow start and the manner in which it is implemented introduces a significant delay in reaching the peak throughput. Most transfers from browsers such as Netscape<sup>®</sup> and Mosaic<sup>®</sup> will not reach peak throughput since they are usually of very short duration and involve setting up and tearing down a connection for each transaction.

TCP and TCP-LFN worked well with bandwidth-on-demand networks with no penalty in the delay or throughput as long as the bandwidth management algorithms were well tuned.

With the introduction of link errors (and no coding on the satellite link), TCP throughput was unaffected at error rates of  $10^{-8}$  or better but degrades significantly at error rates of  $10^{-6}$  or worse.

Interactive applications such as Telnet, which transfer very small amounts of data, suffer mainly from the longer delay. Some limited solutions (line mode) may help alleviate this problem.



## 6. ACKNOWLEDGMENTS

The authors of this report wish to acknowledge the efforts of a number of individuals who contributed to the progress and success of the work reported. Although not named as authors on the report, these individuals deserve recognition for their contributions.

First, we wish to thank Douglas Vandermade and Quoc Nguyen of MITRE and Paul Redman of COMSAT for their assistance in developing the ACTS Collaboration during its early stages. They worked with William Kissick and Marjorie Weibel of NTIA/ITS and Wayne McCoy and Mary Ruhl of NIST/CSL to develop the first program plans and assisted with early interactions with the sponsor.

Mike Nissley and Kamal Chahine of MITRE replaced Mr. Vandermade and Mr. Nguyen when the experiment phase began. The authors are greatly indebted to Mr. Nissley for ensuring the continued use of the ACTS terminal located at MITRE in Reston, Virginia, and for other administrative assistance. Mr. Nissley and Mr. Chahine both assisted the ACTS Collaboration with experiment planning. Mr. Chahine assisted NTIA/ITS with voice quality measurements and COMSAT with protocol performance measurements. His ability to operate the ACTS Earth station and other specialized, terminal equipment is recognized and appreciated.

Keith Junker of NTIA/ITS prepared the hardware and software for the Voice Quality Assessment System (VQAS). He provided training in its use and ensured that this capability continued to work as expected. He also assisted Dave Sutherland with data collection and reduction.

James Fox of NIST/CSL assisted the NIST authors with the desktop conferencing and LAN bridging experiment planning and execution. He assisted both NIST and NTIA authors with data collection and reduction.





## 7. REFERENCES

- [1] P. King and S. Voran, "Final report on the NCS voice communications interoperability standards project," Project Report, Sept. 1993.
- [2] J.L. Devore, *Probability and Statistics for Engineering and the Sciences*, Monterey, CA, Brooks/Cole, 1982.
- [3] International Telecommunications Union – Telecommunications Standardization Sector (ITU-T) Recommendation H.320, "Narrow-band visual telephone systems and terminal equipment," Mar. 1993.
- [4] American National Standards Institute (ANSI), American National Standard for Telecommunications – Digital Transport of Video Teleconferencing/Video Telephony Signals, "Video test scenes for subjective and objective performance assessment," Nov. 1995.
- [5] V. Jacobson and R. Braden, "TCP extensions for long-delay paths," RFC 1072, Internet Engineering Task Force, Oct. 1988.



## BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION NO. 98-354		2. Gov't Accession No.	3. Recipient's Accession No.
4. TITLE AND SUBTITLE National Security and Emergency Preparedness Communications Experiments Using the Advanced Communications Technology Satellite		5. Publication Date August 1998	
		6. Performing Organization Code NTIA/ITS.N4	
7. AUTHOR(S) William Kissick, et al.		9. Project/Task/Work Unit No.	
8. PERFORMING ORGANIZATION NAME AND ADDRESS National Telecommunications and Information Admin. Institute for Telecommunication Sciences 325 Broadway Boulder, CO 80303-3328		10. Contract/Grant No.	
		12. Type of Report and Period Covered	
11. Sponsoring Organization Name and Address National Communication System (NCS) Office of Programs 701 South Court House Road Arlington, VA 22204-2198		13.	
		14. SUPPLEMENTARY NOTES	
15. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.)			
<p>Many government telecommunications needs, especially those that support National Security and Emergency Preparedness (NS/EP) missions, are becoming increasingly dependent on commercially available equipment and services. This is consistent with the goals and concepts of the National Information Infrastructure. This report examines the use of an advanced satellite--in this case, NASA's Advanced Communications Technology Satellite (ACTS)--with ISDN and frame relay protocols to support NS/EP communications requirements. A network using three ACTS earth stations was established as a research facility. With this small network, several experiments were performed. Using new objective methods, voice quality was measured over the satellite and compared to other connections such as commercial, terrestrial lines. The performance of applications--desktop conferencing, file transfer, and LAN bridging--that are likely to be useful in NS/EP situations, was determined. The performance of TCP/IP running over frame relay was examined. The results indicate that advanced satellites can be very useful for emergency communications due to the rapidity that earth stations can be deployed, the ease of reconfiguring the satellite, and the practicality of using commonly available applications running over commonly used protocols. However, there are some limitations to the performance of some applications or parts of applications due to the propagation delay of a satellite channel. Telecommunications protocols such as TCP/IP must be significantly modified to perform well over a satellite channel and to take full advantage of bandwidth-on-demand capabilities of an advanced satellite.</p> <p>Key words: Advanced Communications Technology Satellite (ACTS); Integrated Services Digital Network (ISDN); Frame relay; National Security and Emergency Preparedness (NS/EP).</p>			
17. AVAILABILITY STATEMENT  <input checked="" type="checkbox"/> UNLIMITED.  <input type="checkbox"/> FOR OFFICIAL DISTRIBUTION.		18. Security Class. (This report) Unclassified	20. Number of pages 84
		19. Security Class. (This page) Unclassified	21. Price:

