

Satellite-Enhanced Personal Communications Experiments

Deborah S. Pinck & Loretta H. Tong
Jet Propulsion Laboratory,
California Institute of Technology, M/S 161-241,
4800 Oak Grove Drive, Pasadena, CA 91109
Phone: +1-818-354-8041, Fax: +1-818-393-4643
c-mail: pinck@zorba.jpl.nasa.gov

Anthony J. McAuley & Michael Kramer
Bellcore
MRE 2N263
445 South Street, Morristown, NJ 07960-4698
Phone: +1-201-829-4698, Fax: +1 1-201-829-2504
e-mail: mcauley@bellcore.com

Abstract

Users of future generation wireless information services will have diverse needs for voice, data, and potentially even video communications in a wide variety of circumstances. For users in dense, inner-city areas, low power PCS technology should be ideal. Vehicular-based users traveling at high speeds will need high power cellular technology. For users in remote or inaccessible locations, or for applications that are broadcast over a wide geographic area, a satellite technology would be the best choice. Packet data networks provide an excellent solution for users requiring occasional small messages, while circuit switched networks provide more economical solutions for larger messages. To provide ubiquitous personal communications service, it is necessary to capitalize on the strength of each wireless technology and network to create one seamless internetwork including both current and future wired and wireless networks.

As an initial step in exploring the opportunities afforded by the merging of satellite and terrestrial networks, Bellcore and JPL conducted several experiments utilizing Bellcore's experimental Personal Communications System, NASA's Advanced Communications Technology Satellite (ACTS) and JPL's ACTS Mobile Terminal. These experiments provided valuable information on the applications, interfaces, and protocols needed for seamless integration of satellite and terrestrial networks. Looking at loss of bits, packets, and higher layer blocks over various satellite-terrestrial networks with mobile and stationary users under various conditions, our initial results indicate that the communication channel can vary dramatically, even within a single network. The effect of these conditions on error control protocols is highlighted, with a concentration on those that correct for losses of packets and higher layer blocks.

1 Introduction

The demand for personal information services has exploded in recent years and there seems to be no end in sight. This includes such fields as voice communications, data communications, image and video communications, multimedia, and interactive communications to name but a few. Although the technologies are available to support these applications, the interoperability among terrestrial, satellite, wired, wireless, packet, and circuit networks remains an important challenge.

Currently, many different, and sometimes conflicting, and non-interoperable telecommunications networks are deployed or planned. This is the case for both satellite networks as well as for terrestrial networks. Satellite system providers are planning and deploying the so-called "little LEOs (Low Earth Orbit satellites)" for data-only services, "big LEOs" for voice services, MEOs (Medium Earth Orbit satellites) and GEOs (Geosynchronous satellites) as well [1]. Terrestrial system providers are planning and deploying digital cellular systems, packet data networks, and micro-cellular PCNs using low-power hand-held communicators. As these systems will all co-exist, it is important to insure their compatibility with each other at the earliest opportunity.

In order to investigate the integration of these systems into a seamless network, Bellcore and JPL conducted a series of experiments [2] to demonstrate the joint use of satellites and terrestrial networks in the delivery of personal communications services. For these experiments, applications communicated over various combinations of networks including satellite, wireless packet data, the wired Internet, and the wired Public Switched Telecommunications Network (PSTN). This paper describes results from several field trials conducted during August 1994 and January 1995 in Los Angeles, CA and Morristown, NJ. Section 2 describes the experiments including the experiment goals, the hardware configurations, the satellite and terrestrial protocols, and the applications. Section 3 presents the results of the field trials in the following areas: radio network characterization, error characteristics, packet size, and ARQ parameters. Section 4 provides an analysis of the results and suggestions for future systems.

2 Experiment Description

This section describes the satellite-enhanced personal communications experiments. The goals

of the experiments are listed first, followed by a description of the experiment configuration. Next, the satellite and terrestrial protocols are explained, and lastly, each of the applications is described,

2.1 Experiment Goals

The experiment goals fell into three categories: a) demonstrate the delivery of personal communications applications via satellite, b) demonstrate interoperability of satellite and terrestrial networks, and c) evaluate protocol mechanisms and parameters that make efficient use of wireless links. This work is important because:

- Applications currently planned for terrestrial PCS need to be demonstrated on the satellite network to determine their behavior, and consequently their commercial viability, in a satellite-enhanced personal communications network.
- Various terrestrial wireless networks, including packet data and cellular networks, need to be integrated with the satellite network to evaluate overall end-to-end system performance.
- The protocol mechanisms (e.g., selective vs. cumulative acknowledgments) and parameters (e.g., packet size) developed for terrestrial PCS need to be tested and optimized for use over satellite and integrated terrestrial/satellite networks,

2.2 Experiment Configuration

These experiments utilized Bellcore's Experimental Personal Communications System (BEPCS), NASA's Advanced Communications Technology Satellite (ACTS), and JPL's ACTS Mobile Terminal (AMT). Figure 1 shows the experimental configuration involved in this series of experiments.

BEPCS consists of a Data Gateway PC located at the satellite ground station, the portable PCs (one belonging to the hypothetical Johndoe and the other belonging to the hypothetical Janedoe), the credit card scanner, the prototype application software (see 2.4), and an experimental transport protocol TPE (see 2.3). All the BEPCS hardware uses standard commercial components with custom applications and protocols. The Data Gateway has multiple serial ports allowing it to route all five connections simultaneously: ACTS Satellite link, RAM wireless packet data network, Internet, PSTN to any fax, and PSTN to any telephone. The portable PCs, however, have only a single serial port, so during internetworking experiments the RAM modem was directly connected

to the satellite equipment, bypassing Johndoe's PC.

ACTS is an experimental K/Ka-band satellite developed by Martin Marietta Astro Space under contract to NASA. The satellite, launched into geostationary orbit at 100 degrees west longitude in September 1993, operates in a virtually untapped frequency spectrum in the K- (20 GHz) and Ka-bands (30 GHz).

The AMT [3] is a proof-of-concept K/Ka-band mobile communications terminal intended to demonstrate the system techniques and high risk technologies needed to accelerate the commercial use of land-mobile systems at K/Ka-band. One such technique is the ability to integrate the mobile terminal for satellite communications with terrestrial personal communications equipment. To support this integration, the AMT Terminal Controller provides a digital interface between the portable PCS and the AMT. The data signal is then up/downconverted to the satellite frequencies for transmission over the K/Ka-band channel. The baseline AMT can support up to 128 kbps full-duplex communications; for this experiment, a 9.6 kbps link is used. For land-mobile experiments, the AMT is mounted in a customized Ford Econoline 350 van. The fixed station, located at JPL, utilizes a 2.4m antenna and a 10W high power amplifier (HPA).

Both the AMT and the fixed station are equipped with identical data acquisition systems (DAS). The DAS performs continuous measurements and recordings of a wide variety of propagation, communications link, and terminal parameters (e.g., received data, received power levels, and status of the pilot signal.) The signals are sampled at a rate of 4000 samples/second and are recorded on 5 Gbyte Exabyte tapes for off-line evaluation. The time, vehicle velocity, and position are derived from an on-board GPS system and updated at 10 Hz (time) and 100 Hz (velocity and position),

Figure 1 shows the communication channels. The primary channel used for evaluation and measurements was that between Johndoe's PC and the Data Gateway via the satellite channel. The most demanding configuration was the connection from Janedoe's PC located in New Jersey, via the RAM network, via the AMT van (effectively acting as a mobile base station), over the satellite, through the Data Gateway located in Los Angeles, and back over the Internet to a user back in New Jersey.

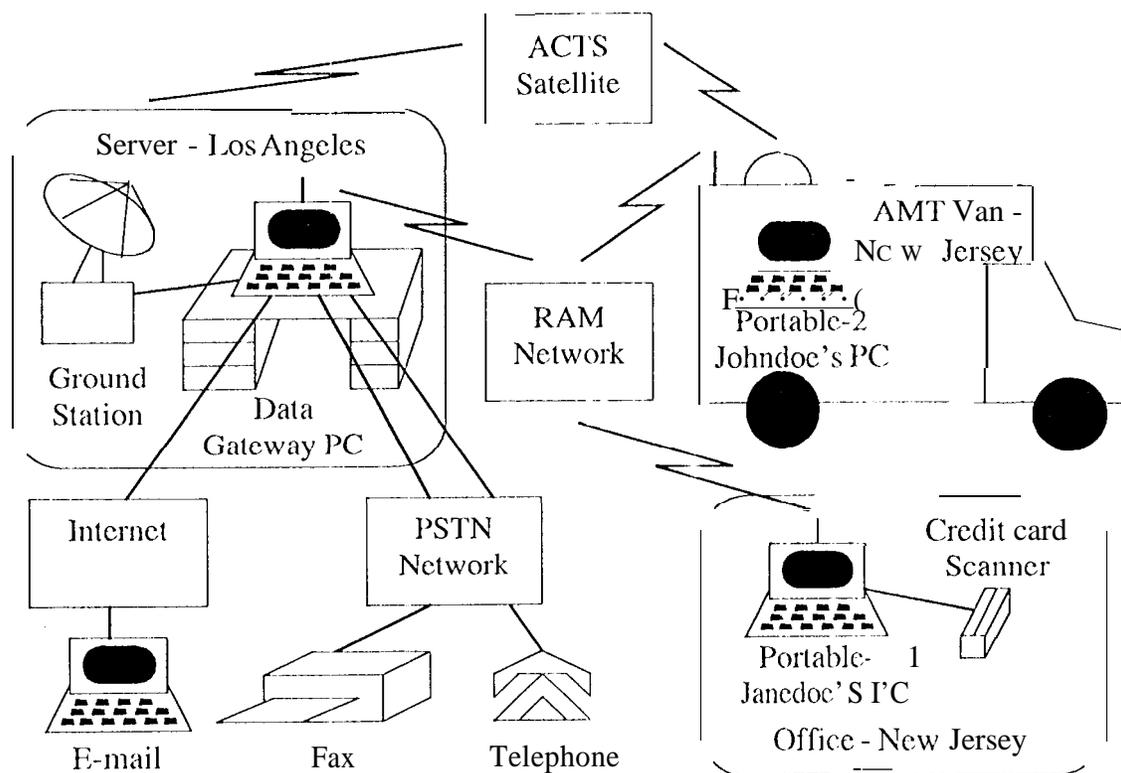


Figure 1 Satellite-enhanced PCS Experimental Setup

2.1 Protocol Description

An important question related to satellite-terrestrial interoperability is whether the data communications protocols can operate efficiently across multiple channels. Most protocols in use today (e.g., TCP/IP) have been optimized for wireline channels; their use over wireless networks presents significant new challenges. The radio and free space channels characteristic of these networks introduce noise, multipath interference, shadowing, and weather effects that can cause higher bit error rates and longer periods of increased bit errors than their wireline counterparts. Host motion creates time-varying communication paths that can cause packet delay, disordering, duplication and loss. In addition, the satellite-link channel can add significantly to the round-trip packet delay, especially when geostationary satellites are used. As currently envisaged, both terrestrial and satellite systems will operate in bandwidth-limited channels where efficient use of the spectrum takes on great importance.

These experiments used two independent protocols: the AMT communications protocol for the satellite link, and TPE located in the end systems.

2.3.1 AMT Communications Protocol Description

The AMT provided the satellite communications between terrestrial devices including the Data Gateway, the portable PCs, and the RAM modem. These devices interfaced with the AMT system by sending packets through a standard asynchronous RS232 serial port. However, the AMT system is based on synchronous data transmissions, with specially designated signals to control data flow. To minimize development time and costs, the discrepancies in data flow between the devices and the AMT system were handled by modifying the existing open-ended, unacknowledged, full-duplex AMT communications protocol (similar to that used for a voice link).

This modified AMT protocol maintained a link at all times (circuit connection); any data received was immediately transmitted over the satellite link with minimal buffering. In order to maintain the communications link during periods when there was no data from the device, the protocol transmitted uniquely identifiable "fill" bytes. On the receiving end, the protocol removed the fill bytes, passing only data bytes to the receiving device.

A satellite fill pattern was chosen that could not occur in the Serial Line Internet Protocol (SLIP) [4]. As SLIP uses an escape byte of 0xDB (i.e., 11011011), with escape-end being 0xDBDC and escape-escape being 0xDBDD, the uniquely identifiable fill pattern of 0xDBDF was chosen.

2.3.2 Wireless Protocol (TPE) Description

Although existing protocols (e.g. TCP) provide efficient error control over conventional wired networks, wireless networks present significant new challenges. The radio environment introduces noise and multipath interference that can cause bursts of errors, while host motion results in changing communication paths that can cause packet delay, misordering, duplication, and loss. Furthermore, scarce radio resources make bandwidth efficiency desirable. An experimental transport protocol was used, called TPE (details of TPE can be found in [5]), optimized for wireless networks. Although TPE is in the end systems, the location of the error control (e.g., whether solely on the wireless link or in the end system) was not addressed. The main objective was to understand the types of mechanisms and parameters that are desirable on wireless networks.

TPE uses the following flexible message framing mechanisms:

- Framing into **chunks** that fit inside a network packet.
- Transport Protocol Data Units (**TPDUs**) using an integer number of chunks.
- Signaling to reduce the chunk header size.

Figure 2 shows an example of how TPE segments an application PDU (**APDU**) into 5 data chunks. In addition to, the 5 data chunks, TPE adds 2 error detection (**ED**) chunks, because the application block was split into 2 Transport PDUs (**TPDUs**). TPE moves chunks directly to and from main memory - application data is never copied at the level of APDU or TPDU. Chunks contain only one type of information (e.g., a data chunk contains application data, while in an ED chunk contains only parities).

A datagram protocol, such as IP, transmits full address and Quality of Service (QOS) information in each packet. A virtual circuit protocol, like the TPE (or ATM), relies on advanced signaling to set up a virtual connection before sending data. Following this advanced signaling, a small virtual circuit identifier identifies the full address/QOS information. The chunk header sends the virtual circuit identifier along with other information that identifies exactly what information it contains; thus chunk headers identify the information, allowing out-of-order processing of chunks. This contrasts with “soft” virtual circuits, where a packet is dependent on previous packets. Such a scheme does not work well in a wireless environment (with frequent packet corruption, duplication, loss, misordering, and propagation delay), because if key segments are lost,

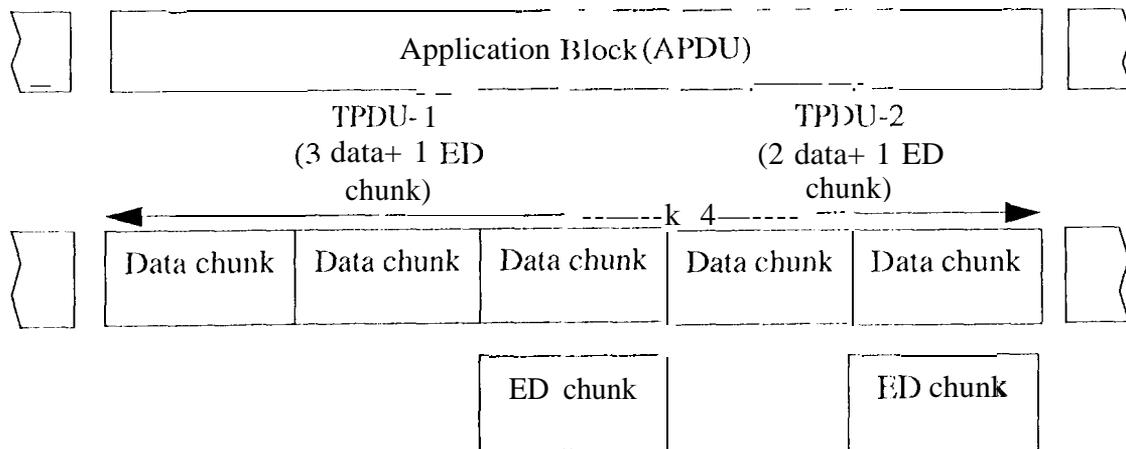


Figure 2 Segmentation of application data into chunks

subsequent packets must be dropped until end-to-end retransmissions re-synchronize state.

The small chunk header allows TPE to send a chunk in each packet, even with small packets. If a 'TCP/IP' TPDU is fragmented into IP packets, it adds 20 bytes/packet (for IP fragment header) and 20 bytes/TPIJU (TCP header). If the TCP/IP TPDU is fragmented into packets without a header, TCP/IP adds only 40 bytes/TPDU; however, if any packet is lost or misordered the entire TPDU is lost. In the field trials, TPE added 8 bytes/packet for chunk header, put exactly 1 chunk/packet, and used 4 bytes/TPDU for error detection parity.

TPE's error control mechanisms (operating on the TPDU's) are:

- Error detection using a powerful 32-bit code.
- powerful ARQ (Automatic Repeat reQuest).
- Combining data from multiple transmissions.

The error detection parities detect corrupted data. Using few parities (e.g., 2 byte parity) or a weak code (e. g., Internet checksum) leaves a significant chance of falsing, particularly in the hostile radio environment. TPE sends 4 parity bytes per TPDU, encoded using the Weighted Sum Code [5]. The Weighted Sum Code has similar error detection properties to the CRC, but can be calculated faster (within a factor of (wo of the Internet checksum).

TPE uses a computationally demanding form of ARQ: using selective-acknowledgments (receiver sends a list of correctly received TPDU's rather than a single cumulative acknowledgment) and selective retransmission (transmitter only resends those TPDU's that are unacknowledged after a time-out interval, rather than a Go-back-N scheme that resets the send pointer after retransmissions). The added complexity improves bandwidth efficiency and reduces latency variance by preventing unnecessary retransmissions [7].

TPE combines chunks from multiple partially received TPDU's to make up a complete TPDU. TPE combines data from different transmissions of the same TPDU. If the receiver gets a subset of a TPDU's chunks in the first transmission (and the remaining chunks were lost), then, when the TPDU is retransmitted, the receiver only processes the missing chunks (eliminating duplicate chunks). If chunk loss is more frequent than chunk corruption, combining data can significantly reduce the total number of retransmissions.

To allow fast set-up and robust connections TPE uses two key mechanisms:

- Fast opening and closing of unreliable connections.
- Reliable sessions.

Connections identify a path between two sessions. Using timer-based connection management, connections are set-up without a three way handshake, allowing data to be sent immediately after a set-up message. A reliable bidirectional session uses the unreliable single party, unidirectional connections, as required. The session layer may use multiple connections during a session, either for efficiency (to close the network connections during periods of inactivity) or for re-establishment of lost connections. The session layer could transparently re-establish failed connections. For our experiments, TPE kept connections open until the session was closed.

2.4 Applications

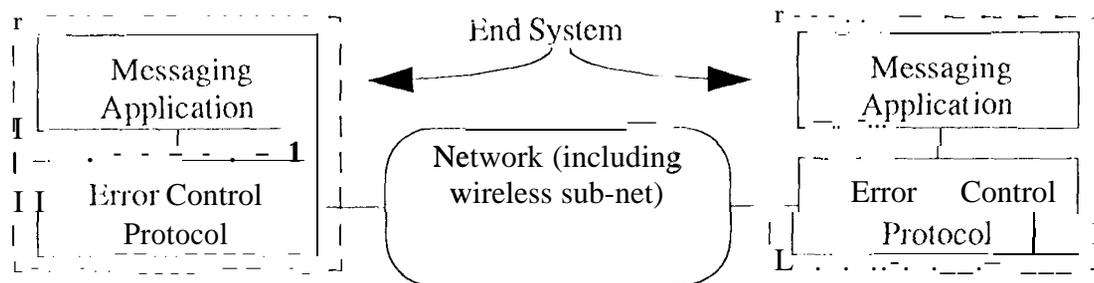


Figure 3 System Architecture

For these trials, four messaging applications designed to suit mobile users were developed. Figure 3 shows the relationship among the applications, the wireless networks, and the error control protocol. All the application and protocol software executes on PCS running UNIX¹. The server applications ran on the Data Gateway, and the client applications on the two portable PCs. The applications on the portable PC's had an X-windows²/Motif interface. The four messaging applications were: a) E-mail, b) Fax, c) Call Command and d) Credit Card Verification.

The E-mail application allows mobile users to exchange text messages with any machine on the

¹Unix is a registered trademark of Novell, Inc.

²X-Windows is a trademark of MIT.

Internet. The Fax application sends and receives standard Group III faxes between a mobile user and any standard fax machine connected to the PSTN. Both the E-mail and Fax application servers reduce bandwidth on the wireless link by sending only a small mail/fax header automatically to the mobile user. The user is informed of the incoming message and can decide whether to request delivery of the (potentially large) body of the message.

CallCommand [6] is a personal telephone management application that enables users to redirect telephone calls in real time. If someone calls Janedoe, the Call Command server (the Electronic Receptionist) tells the caller to wait while Janedoe is contacted. The Call Command server sends information about the incoming call to Janedoe's portable PC. Upon receiving the notification, Janedoe can route the call directly, either to any local phone, to voice mail, or to another person.

The Credit Card Verification application was designed for field service and sales people submitting charges for equipment or services anywhere within the radio coverage area. The user scans a credit card, causing a window to pop-up (on the portable PC) with the information on the card. When the user enters the amount of the transaction, it sends a message to the credit card server. A database at the server (the Data Gateway for this case) stores valid credit card numbers and available credit line. After checking the database, the server sends a message back either authorizing or rejecting the transaction.

2.5 Experiment Procedure

During the Los Angeles and Morristown field trials, experiments were conducted under varying channel conditions in order to characterize hybrid satellite-terrestrial personal communications. 130th stationary ant] mobile tests were conducted at various signal strength levels. For these tests the van was either parked (stationary communications) or driven (mobile communications) while the tests were being executed. Mobile conditions present more adverse and unique conditions for a communications network (particularly shadowing which disrupts the line-of-sight to the satellite).

Control over the signal level was accomplished by measuring the bit signal-to-noise ratio (bit SNR or E_b/N_0). The E_b/N_0 was measured by filtering the received signal with a known noise equivalent bandwidth (BW) filter. The resulting signal was supplied to a power meter for

measurement. With this measurement from the power meter, the pm-calibrated noise-only power reading, and the known bit rate, the DAS calculated and displayed the E_b/N_0 level. For the stationary tests, the target E_b/N_0 range was between 5 dB and 12 dB. This provides a wide range of bit error rates in which the performance of TPE can be examined.

For the mobile tests, the E_b/N_0 was set at a high level (-13dB) to allow characterization of shadowing and fading events. Tests were run in both clear weather (1 MA in August) and inclement weather (New Jersey in January). In addition, runs were taken on roads with clear line-of-sight (LOS) to the satellite as well as roads with foliage and/or buildings that disrupted the clear LOS to the satellite.

Each of the four messaging applications was tested under the various field conditions. These applications were used to determine the behavior of wireless messaging applications in a mobile satellite environment. In order to characterize the protocol, a simple file transfer application was used that sent a 100 kilobyte file from the portable PC to the Data Gateway and vice versa. No data compression was used. By sending the same file, we were able to gauge the relative communication throughput and delay at various locations. TPE was instrumented to collect a wide variety of packet and TPDU statistics at the transmitter and receiver as they affected packets and TPDU s .

Normally terrestrial radio errors are presented solely in terms of bit error rates and error control is looked at solely in terms of mechanisms that control detect and correct bit errors. The error rates on packets and TPDU, however, can be as important to the user as the fundamental bit error rate. Also, the error control mechanisms that work at the packet and TPDU level (using TPE in our experiments) can be as important as (and complimentary to) mechanisms that reduce the bit error rate. Although the results presented here are not exhaustive, they lead to some interesting initial findings about a) the characteristics of radio networks as seen from the level of packet and TPDU, and b) the performance of different ARQ mechanisms in reducing packet and TPDU retransmissions.

3 Experiment Results

This section summarizes the results from the field trials. As mentioned previously, the goal was

to characterize the error environment for satellite wireless networks and study its impact on the TPE protocol and “real-world” applications.

All the quantitative results presented below were obtained by sending the same 100 kilobyte file over both stationary and mobile links. The following sections describe the error distributions, first as they affect the bit error distribution within a packet, then as they affect the packet error distribution within the TPDU, and finally as they affect TPDU retransmission. Finally, this section describes the performance of the error control mechanisms.

3.1 Radio Network Characterization

The radio environment for a stationary satellite link is substantially different from that for a mobile satellite link. This can be clearly seen by comparing Figure 4 and Figure 5. Both figures show a representative time series of bit signal-to-noise (E_b/N_0). Figure 4 shows the radio environment for a stationary link where the bit signal-to-noise at the start of the test was set to 11.5 dB. Although there was some fluctuation, the maximum signal variation was less than 1 dB. Figure 5 shows the case where the initial bit signal-to-noise was set to 8 dB and the van was mobile. Here, the signal varied up to 3 dB for minor shadowing, and large drops of greater than 20 dB were present later in the run. Thus it is clear that there is a large difference in the E_b/N_0 characteristics between stationary and mobile runs,

3.2 Packet Error Characteristics

As expected, the difference in the E_b/N_0 characteristics translates into significantly different packet error distributions between the mobile and the stationary runs. Using results from experiments where each TPDU was composed of 5 small (24 byte or 32 byte) packets, Figure 6 and Figure 7 show how many of the 5 packets were destroyed within corrupted TPDU's. For the stationary tests, Figure 6 shows that the majority of TPDU's lost only one or two packets; Figure 7 shows that for the mobile tests, most corrupted TPDU's had all 5 packets destroyed. Interestingly, in stationary tests with larger (128 and 256 byte) packets (not shown), almost all TPDU's lost only a single packet.

3.3 TPDU Error Characteristics

The number of TPDU's lost or corrupted in transmission was significantly higher for mobile runs

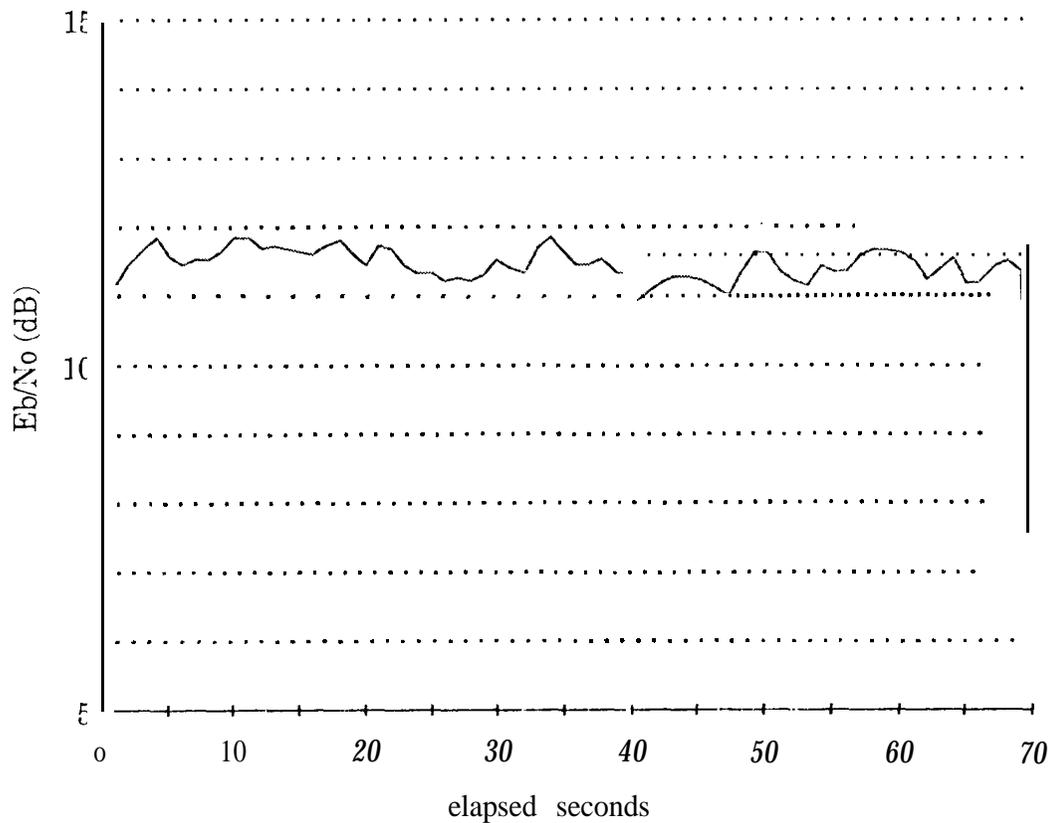


Figure 4 Bit Signal-to-Noise Ratio for Stationary Links

than for stationary runs due to the substantially lower E_b/N_0 's which exist for the mobile case. This result is easily predictable; however, of greater interest is the number of retransmissions which were required before a successful transmission was achieved. Figure 8 shows that for the stationary tests, the majority of the lost/corrupted TPDU's were sent successfully with only one retransmission. Few (under 20%) of the lost/corrupted TPDU's required more than a single retransmission, in contrast to this relatively small percentage, Figure 9 shows that for the mobile

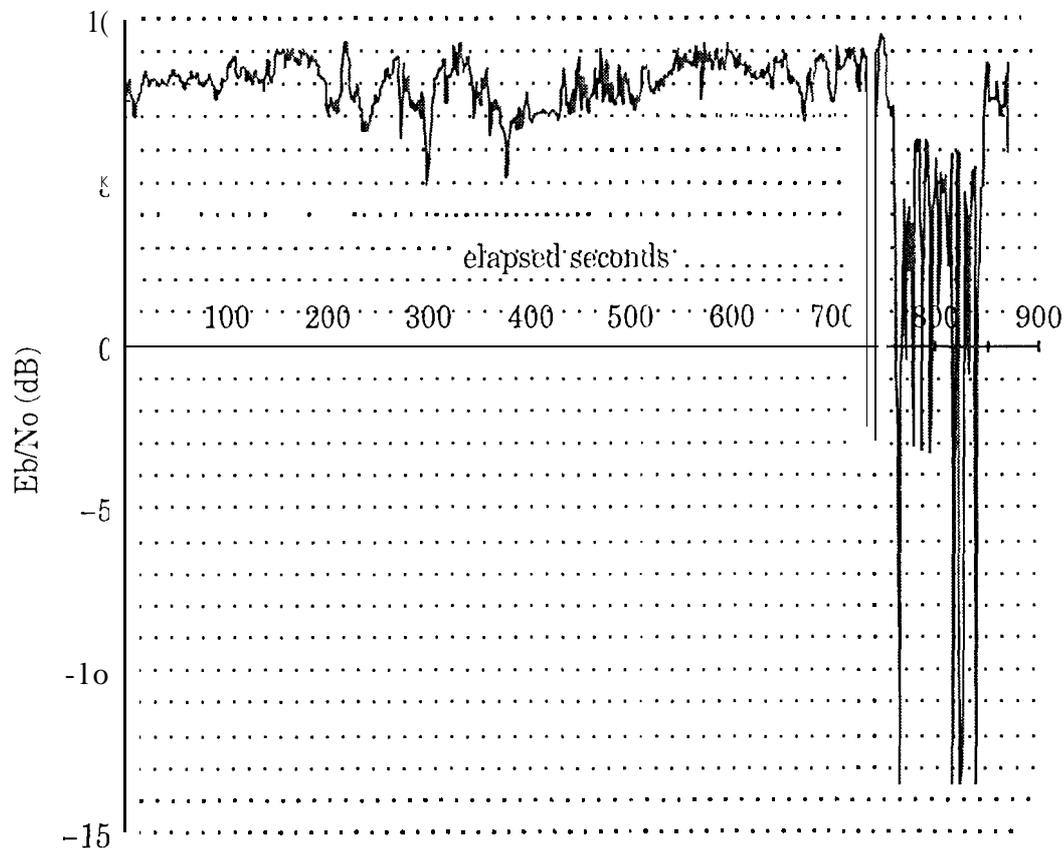


Figure 5 Bit signal-to -noise ratio for mobile links

tests, most (over 70%) of the lost/corrupted TPDU's needed several retransmission. Thus it appears that the mobile environment requires substantially more retransmission attempts to successfully transmit a TPDU than a stationary environment.

3.4 Packet Size

One error control parameter that was varied was the size of the packet. For this series of tests,

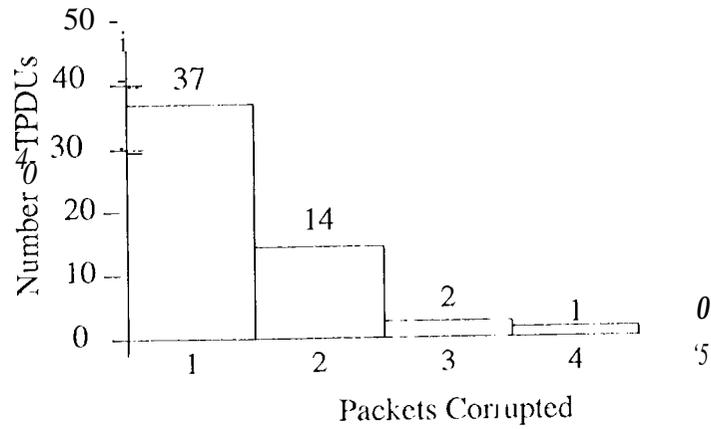


Figure 6 Packet errors distribution for stationary links (approximately 8000 5-packet TPDU's sent with 32-bytes/packet)

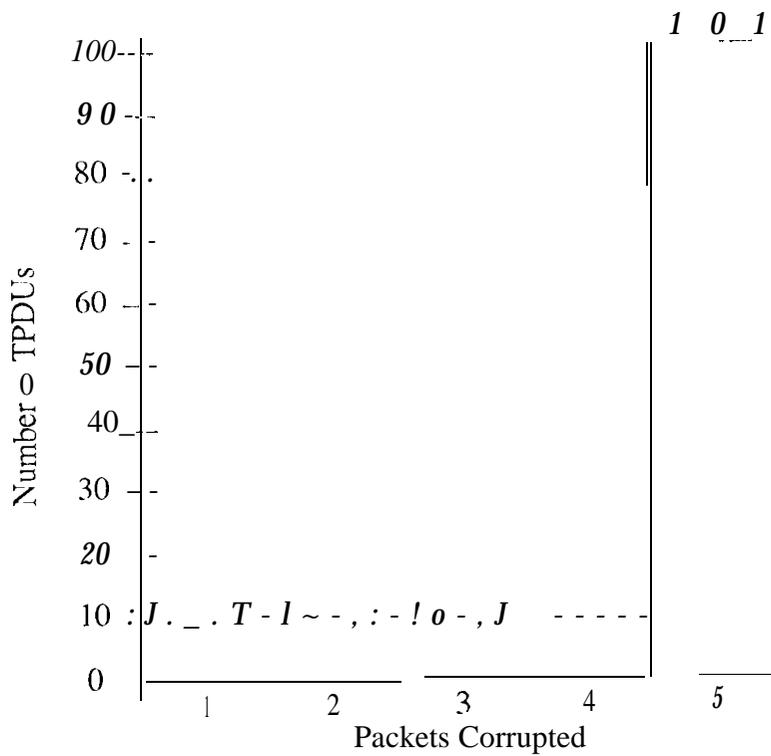


Figure 7 Packet errors distribution for mobile links (approximately 8000 5-packet TPDU's sent with 24-bytes/packet)

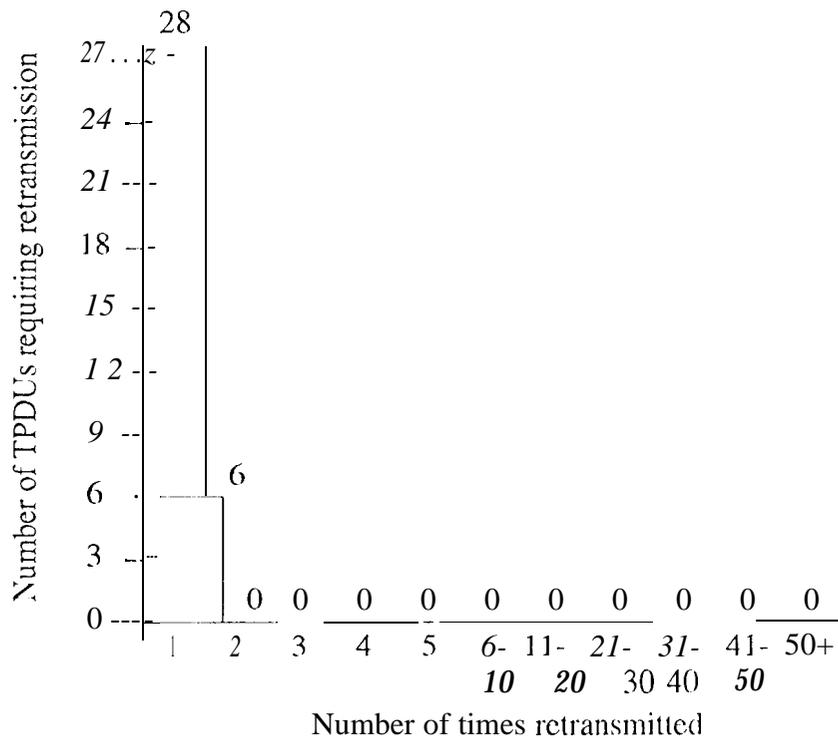


Figure 8 TPDU retransmissions for stationary tests (~15,000 192-byte TPDU's sent)

the packet size was varied from a small packet size of 24 bytes per packet to a large packet size of 256 bytes per packet. The same 100 kbyte file was sent for each packet size and the transmit time was measured.

Results indicated that under good conditions (that is, high E_b/N_0 values), the use of larger packets reduced the time needed to send the 100 kilobyte file as it lessened per-packet processing overhead. For example, it took approximately half the time to send the file using 256 bytes per packet as it did using 32 bytes per packet. Under poor conditions (with lower E_b/N_0 values), however, the larger packet size often took much longer than using smaller packets, and in extreme cases would result in no useful throughput.

3.5 ARQ Parameters

Figure 10 and Figure 11 show the effect of changing two ARQ parameters: the types of acknowledgment and the number of packets per TPDU

Figure. 10 shows the average time to transmit the 100 kbyte file using cumulative and selective

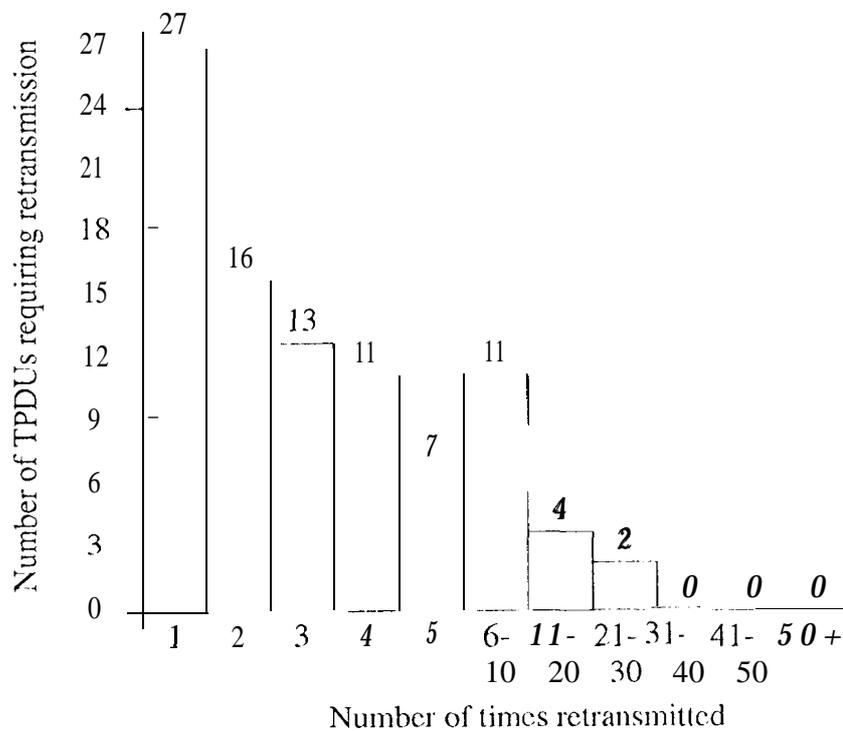


Figure 9 TPDU retransmissions for mobile tests (~15,000 192 byte TPDU's sent)

acknowledgments with both high and low signal-to-noise ratios. The results indicate that the type of acknowledgment had little effect on the file transfer time with high signal-to-noise ratios; however, the use of selective acknowledgments significantly reduced latency, and was therefore more bandwidth efficient with low signal-to-noise ratios.

Figure 11 compares the average time to transmit the same 100 kbyte file using different numbers of packets per TPDU. As expected, the processing and bandwidth overhead associated with a TPDU meant that very small TPDU's were less efficient. More surprising is that the very large TPDU's did not significantly degrade performance, even for runs with low signal-to-noise ratios.

3.6 Error Detection

The satellite-based PCS trials did not utilize link error control; therefore there were many packets with bit errors. However, all errors were detected by the powerful 32-bit TPDU error detection parities,

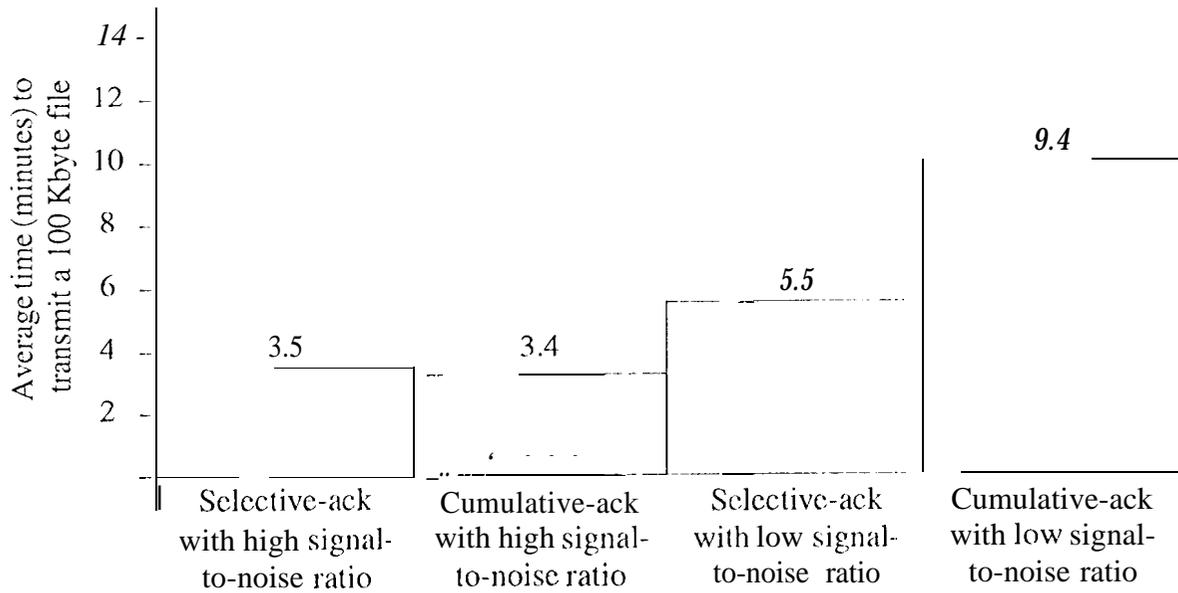


Figure 10 Transmission time using selective acknowledgment or cumulative acknowledgment

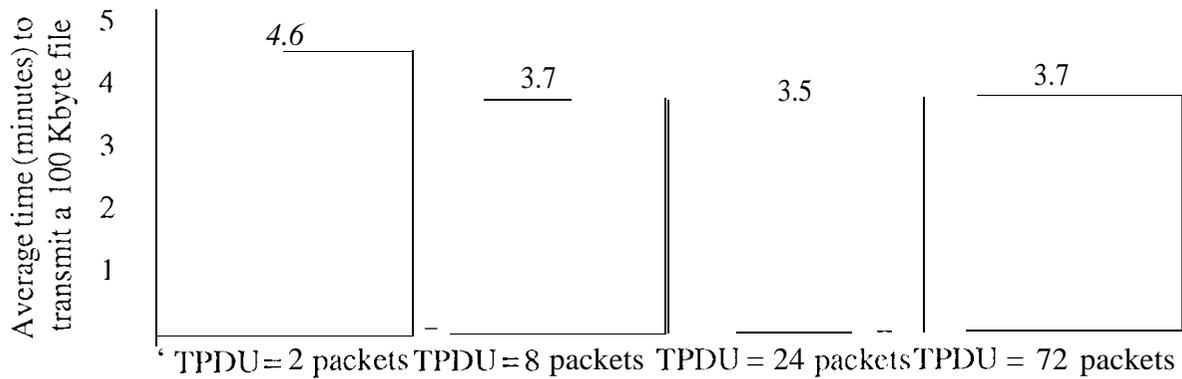


Figure 11 Transmission time using different TPDU sizes (with Selective Acknowledgment)

4 Analysis

4.1 AMT system

In order to simplify the integration between the JPL and Bellcore equipment and to reduce the

development time, a modified AMT protocol was designed. This protocol necessarily has some shortcomings (known a priori and observed during experimentation) which are described here.

There were no signal controls between the Bellcore equipment and the AMT for initiating and terminating a satellite link. Thus, the satellite link had to be manually established independent of and prior to the initiation of a data transmission. At the end of a data transmission, the satellite link had to be manually terminated. A dynamic link allocation which does not require operator intervention would be much more desirable.

The open-ended link meant that the communications channel was occupied at all times even when there was no "real data" being transmitted. Although this was not a problem during experimentation, it is an inefficient usage of an operational satellite channel.

The next weakness is related to corruption of the "fill" bytes. Since the satellite link was open-ended with no acknowledgments (the AMT system simply pipes data through with no additional error detection overhead), during the satellite link transmission "fill" bytes which suffered bit errors were not detected at the AMT receiving end. Any undetected "fill" byte was assumed to be data and passed to the Bellcore equipment. Thus, the data sequence received by the Bellcore equipment was likely to contain additional (erroneous) data bytes. This places more demands on the receiving processor of the Bellcore protocol.

During experimentation, the addition of corrupted "fill" bytes in to the data stream of the receiver was known to be occurring and seemed to have various undesirable impacts on the TPE protocol. Although these shortcomings were acceptable for this experiment, it is evident that a more efficient protocol would be beneficial.

As a result of our experiments, a new AMT communications protocol is being proposed that eliminates the above shortcomings. The key change is that when no data is present, the new protocol leaves the satellite link idle. By maintaining an appropriately sized data buffer, the protocol has time to resynchronize the connection. Although the additional buffering delay is unacceptable for real time voice or video applications, it is acceptable for most messaging applications, such as those used in our experiments, which can tolerate a small delay.

4.2 Channel Conditions

The communication channel between a satellite and a land-based mobile user presents many challenges, including signal fading. In general two different types of fading are encountered: multipath interference and shadowing. Multipath fading occurs when reflected (and/or delayed) copies of the transmitted pulse are coincident with the line-of-sight signal. Shadowed fading is the signal attenuation due to obstruction or partial obstruction of the transmitted signal by terrain, foliage, utility poles, trees, buildings, etc., surrounding the mobile terminal. Both fading modalities can cause severe variations in the received power, making reliable communications difficult. Shadowing effects are the most severe source of signal outages in a land mobile satellite system. Many operational terrestrial wireless systems operate at UHF (~800 MHz), while many satellite systems operate at L- and S-bands (1 - 3 GHz). The attenuation due to shadowing increases with frequency, and thus satellite systems operate on an inherently noisier channel than terrestrial systems.

The experiments described in this document were conducted using NASA's ACTS satellite which operates at Ka-band (20/30 GHz). This satellite was not chosen to suggest the use of Ka-band for satellite-enhanced personal communications. Rather, it was chosen because of the existing availability of the Ka-band ACTS satellite experimental platform. Although the Ka-band experimental results are not identical to expected results for a similarly configured L- or S-band system, they do provide a good indication of system level performance.

It is well known that shadowing in Ka-band systems is substantially worse than in L-band systems, which themselves are worse than terrestrial systems at UHF. The Ka-band experiments, therefore, serve as a "lower bound" on system performance. Various fading/shadowing countermeasures are currently being investigated, including error control coding techniques and antenna diversity. Given the fading effects of L-, S-, and Ka-bands, it is critical to design wireless error control schemes that work efficiently at these higher frequencies.

4.3 Error Control Mechanisms and Parameters

The results show the importance of powerful error control mechanisms such as selective acknowledgment and chunk combining [5]. Selective acknowledgments are very worthwhile for environments such as that in a mobile satellite link which are characterized by having low (and

varying) signal-to-noise ratios. Combining data from multiple transmissions caused the insensitivity to TPDU size. (Code combining, described in Section 2.3, means that the TPDU can be made up of packets from multiple TPDU transmissions.) The results also demonstrate that dynamic algorithms are required to change the error control parameters. For example, as the bit signal-to-noise ratio decreases so should the packet size.

Results obtained here also provide new insights into the possible use of packet Forward Error Correction (FEC) in a wireless environment. Although no experiments were performed with FEC, the results represented in Figure 6 and Figure 7 show that a packet FEC scheme, such as that described in [8], would have worked well for stationary tests, but would have been ineffective where errors were dominated by long periods of fading. The number of packets of overcode should vary depending on the packet size. For example, two overcode packets per TPDU would be necessary to ensure the most TPDU's with small packets got through; while, with larger packets, a single overcode packet would have been sufficient.

Thus, just as the different bit error characteristics translate into different packet error distributions, so the differences in packet error distribution lead to different TPDU error distributions between mobile and the stationary runs.

4.4 Applications

The four messaging applications consistently performed well in stationary tests. More surprising, however, was that the applications maintained reasonable performance in mobile tests, even when voice connections were unintelligible (messaging applications are tolerant of wide delay variation).

5 Conclusion

This paper describes results from satellite-enhanced personal communications field trials using the ACTS satellite with various wireless and wireline terrestrial networks. The interconnectivity among multiple wireless communications services shows how a mobile user can have personal services available even while outside the coverage area of a single communications provider.

The results of packet and higher layer block loss show that, even within a single wireless

network, conditions vary dramatically depending on the type of fading which dominates.

in order to deal with the varying radio conditions, an efficient general wireless error control protocol requires powerful error control mechanisms and sophisticated dynamic control algorithms. With the correct block and packet sizes, a selective ARQ protocol can provide efficient error control under virtually any radio condition. Also, with the correct packet overhead parameter, an FEC scheme could provide efficient low latency error correction under some radio conditions.

Acknowledgments

The research described in this paper was partially carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by Bellcore and the National Aeronautics and Space Administration.

References

- [1] K. G. Johannsen, M. W. Bowles, S. Milliken, A. R. Cherrette, & G.C. Busche, "Trends in Mobile Satellite Communication," Proceeding of the International Mobile Satellite Conference, Pasadena, CA, pp. 157-162, June 1993.
- [2] R. Wolff & D. Pinck, "Internetworking Satellite and Local Exchange Networks for Personal Communications Applications," Proceeding of the international Mobile Satellite Conference, Pasadena, CA, June 1993.
- [3] K. Dessouky & T. Jedrey, "The ACTS Mobile Terminal," SATCOM Quarterly, JPL, Pasadena, CA, No. 8, January 1993.
- [4] J. Romkey, "A Nonstandard for Transmission of 11' Datagrams over Serial Lines: SLIP," Internet RFC 1055, June 1983.
- [5] A. J. McAuley, "Error Control for Messaging Applications in a Wireless Environment," INFOCOM 95, Boston, MA, April 2-6, 1995.
- [6] M. Kramer, G. Ramirez, D. Turock, R. S. Wolff, "Use of Two-Way Wireless Messaging for Personal Telephone Management," 2nd Universal Personal Communications Conference, Ottawa, Canada, 1993.
- [7] S. Nanda, R. Ejzak, & B. Doshi, "A Retransmission Scheme for Circuit Mode Data on Wireless Links," IEEE Journal on Selected Areas in Communications, Vol. 12, pp. 1338-1352, October 1994.
- [8] A. J. McAuley, "Reliable Broadband Communication Using a Burst Erasure Correcting Code," Proceedings of ACM SIGCOMM, pp. 287-306, September 1990.