

Technical Report
1087

EHF Satellite Communications on the Move: Experimental Results

J.B. Schodorf

15 August 2003

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Prepared for the Army under Air Force Contract F19628-00-C-0002.

Approved for public release; distribution is unlimited.

20030902 083

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by Army/CECOM under Air Force Contract F19628-00-C-0002.

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

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

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

FOR THE COMMANDER


Gary Tutungian
Administrative Contracting Officer
Plans and Programs Directorate
Contracted Support Management

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission has been granted by the Contracting Officer to destroy this document, when it is no longer required by the using agency, according to applicable security regulations.

Massachusetts Institute of Technology
Lincoln Laboratory

**EHF Satellite Communications on the Move:
Experimental Results**

J.B. Schodorf
Group 66

Technical Report 1087

15 August 2003

Approved for public release; distribution is unlimited.

ABSTRACT

This report summarizes the results of experiments conducted with an EHF SATCOM on the move (SOTM) terminal developed to work with LDR MILSTAR. The experiments were designed to measure characteristics of the EHF SOTM propagation channel, yield insight into the performance of an automatic repeat request error control protocol, and assess the pointing performance of the system's antenna positioner. With respect to the propagation experiments, the report contains statistical analyses of the data that yield probability density models for the received signal power, fade and nonfade duration distributions, and estimates of level crossing rates. Discrete models for the EHF SOTM channel are also explored and compared to the measured data. Results from the protocol experiments include packet error rate, throughput efficiency and packet latency as a function of protocol parameters (e.g., packet length) and channel conditions. As with the propagation experiments, the report contains analyses of the antenna pointing data recorded during the positioner experiments that result in statistical models for the LOS error under various terrain conditions.

ACKNOWLEDGMENTS

This work was supported by the Army's Communications and Electronics Command (CECOM), and by PM WIN-T. I am grateful for their support. I would especially like to thank Mr. Gerald T. Michael, of CECOM, for his commitment to this project and his patience with me in writing the report.

Of course, this report would not have been possible without the numerous contributions of others. The experiments themselves were long and tedious. I would like to thank Ted O'Connell and John Murphy for their professionalism in helping me set up the experiments and collect the data.

I was assisted on numerous occasions by Andrew Worthen and Ron Bauer in analysis of the data and interpretation of the results. Mike Boulet and Ed Corbett were also valuable contributors in the area of the positioner experiments. Andrew, Ron, and Mike also reviewed the manuscript and offered many helpful comments for its improvement. I would also like to thank Bill Moser for reviewing the manuscript.

Finally, I would like to thank Dave Snider and Mark Gouker for their support and leadership in this project. Mark was the original MILSTAR on the move project leader and his vision was largely responsible for facilitating the various experiments that are discussed here. Dave Snider ensured that whatever resources were necessary to conduct the experiments and process the results were made available. I am grateful for his support and his patience with me in preparing the report.

TABLE OF CONTENTS

ABSTRACT	iii
ACKNOWLEDGMENTS	v
LIST OF ILLUSTRATIONS	ix
LIST OF TABLES	xiii
1. INTRODUCTION	1
1.1 PROGRAM OVERVIEW	1
1.2 MILSTAR OVERVIEW	2
1.3 MOTM SYSTEM OVERVIEW	2
1.3.1 Antenna Positioner Subsystem	2
1.3.2 Testbed Terminal	3
1.4 EXECUTIVE SUMMARY	5
1.4.1 Communications Experiments	6
1.4.2 Positioner Characterization	9
2. EHF SOTM CHANNEL CHARACTERIZATION	11
2.1 EXPERIMENT OVERVIEW	11
2.1.1 Measurement Approach	13
2.1.2 Course Descriptions	16
2.2 STATISTICAL ANALYSIS	21
2.2.1 Probability Density Models for Received Signal Power	21
2.2.2 Discrete Models	27
2.2.3 Fade and Connection Duration Distributions	33
2.2.4 Average Fade and Connection Durations	39
2.2.5 Level Crossing Rates	41
2.3 BER RESULTS	42
2.4 SUMMARY	50
3. PROTOCOL PERFORMANCE	51
3.1 PROTOCOL OVERVIEW	52
3.1.1 QoS Management	53
3.1.2 User Applications	55
3.1.3 Packetization and Queueing	58
3.1.4 FEC codec	60
3.1.5 Framing and Deframing	60
3.1.6 Rate Control	62
3.2 EXPERIMENT OVERVIEW	62
3.3 PERFORMANCE EVALUATION	63

TABLE OF CONTENTS

(Continued)

3.3.1	Average Packet Error Rate and Throughput Efficiency	63
3.3.2	Packet Latency	67
3.4	SUMMARY	72
4.	POSITIONER CHARACTERIZATION	75
4.1	EXPERIMENT OVERVIEW	75
4.1.1	Measurement Approach	76
4.1.2	Course Descriptions	78
4.2	STATISTICAL ANALYSIS	81
4.2.1	Probability Density Models for LOS Error and Antenna Gain	81
4.2.2	Average Fade and Connection Durations	84
4.2.3	Level Crossing Rates	87
4.3	SUMMARY	87
5.	CONCLUSIONS	91
	APPENDIX A – GLOSSARY	93
	REFERENCES	95

LIST OF ILLUSTRATIONS

Figure No.		Page
1	Simplified block diagram of the MOTM terminal.	3
2	MOTM antenna positioner assembly.	4
3	Vehicle layout (top view).	5
4	Chassis layout.	6
5	Channel measurements and protocol performance experimental setup.	12
6	Typical velocity profile in the rural and urban environments.	13
7	Block diagram of ASCAMP hard-limiting front-end and P_r/N_0 estimator.	14
8	Probability density of the P_r/N_0 estimates when the terminal is in saturation.	16
9	Time series data segments from the three test environments.	17
10	Open and rural environments.	19
11	Urban environment.	20
12	Data description.	21
13	Probability densities for estimated received signal power in the three test environments.	24
14	Complementary cdfs for estimated received signal power in the three test environments.	26
15	Two-state Markov model.	27
16	Complementary cdfs and two-state models for fade and connection durations in the open environment.	35
17	Complementary cdfs and two-state models for fade and connection durations in the rural environment.	36
18	Complementary cdfs and two-state models for fade and connection durations in the urban environment.	37
19	Time share of connections and fades for the three test environments.	40
20	Average connection and fade durations as a function of threshold.	41
21	Depiction of shadow edge.	42
22	Received signal power level crossing rates for the three test environments.	43
23	Theoretical average BER as a function of LOS received power level.	44
24	BER scatter plots for the three test environments.	45

LIST OF ILLUSTRATIONS

(Continued)

Figure No.		Page
25	Uplink and downlink error results for the open channel.	47
26	Uplink and downlink error results for the rural channel.	48
27	Uplink and downlink error results for the urban channel.	49
28	Block diagram of the MEP implementation.	52
29	User interface to QoS management features.	54
30	User interface to packetized voice application.	56
31	Histogram of speech transmission times in a tactical scenario.	57
32	User interface to channel testing application.	59
33	MEP packet header structure.	60
34	FEC codec functionality.	61
35	Average packet error and success rates in the three test environments.	64
36	Throughput efficiency as a function of packet length.	66
37	Average packet latency.	70
38	Histogram of packet latencies.	71
39	Standard deviation of packet latency.	72
40	ACU position control loop.	76
41	Time series data segments for LOS error and antenna gain from the three test terrains.	77
42	Terrain types for positioner testing.	78
43	Comparison of vehicle dynamics over test terrains to Churchville B.	80
44	Probability densities for azimuth and elevation error over rough terrain.	82
45	Complementary cdfs for the three terrain types and two elevation angles.	83
46	Probability density for LOS error and antenna gain in improved terrain, 40° elevation angle.	84
47	Probability density for LOS error and antenna gain in unimproved terrain, 40° elevation angle.	85
48	Probability density for LOS error and antenna gain in unimproved terrain, 65° elevation angle.	85

LIST OF ILLUSTRATIONS

(Continued)

Figure No.		Page
49	Probability density for LOS error and antenna gain in rough terrain, 40° elevation angle.	86
50	Probability density for LOS error and antenna gain in rough terrain, 65° elevation angle.	86
51	Average fade and connection durations over the three terrain types.	88
52	Level crossing rates for the three terrain types.	89

LIST OF TABLES

Table No.		Page
1	Simplified, representative link budget for MOTM downlink.	7
2	Density model parameters derived from a least squares fit to the measured data.	25
3	Two-state Markov model parameters for the three test environments.	38
4	Gilbert model parameters for the three test environments.	38
5	Jointly optimized Gilbert model parameters for the three test environments.	38
6	Latency budget for MOTM and MEP implementation.	69
7	RMS LOS pointing error for each of the experiments.	75
8	Azimuth and elevation error density model parameters.	83

1. INTRODUCTION

Among the most critical aspects of the U. S. Army's drive to transform itself into an Objective Force (OF) is the need for robust, high data rate communications capabilities to mobile platforms that require connectivity at distances beyond line-of sight (BLOS). The combination of these communications needs, (i.e., robust, high data rates, BLOS, and mobile platforms) leads naturally to the identification of EHF SATCOM on the Move (SOTM)¹ as a technology critical to the success of Army transformation. As part of an ongoing effort supported by CECOM and PM WIN-T, Lincoln Laboratory has been engaged in the research and development of EHF SOTM terminals for use with the MILSTAR system. This work culminated in a low data rate (LDR) MILSTAR on the move (MOTM) terminal that was successfully demonstrated at Ft. Monmouth, NJ on 20 June, 2002. Since the demonstration, the MOTM terminal has been used extensively for experimental purposes to gain insight into fundamental issues such as EHF OTM propagation characteristics, error control protocol performance and antenna positioner behavior. The purpose of this report is to describe the results of these experiments.

The remainder of this section consists of background material and an executive summary of the experimental results. Section 2 presents a detailed discussion of results from the channel measurement experiments, including the areas of statistical analysis and modeling. In Section 3 the performance of the MOTM experimental protocol (MEP), an error control protocol based on a repeat request strategy, is examined. In Section 4 experimental results that quantify the pointing performance of the MOTM antenna positioner are discussed. A brief summary and a few conclusions are presented in Section 5. Finally, readers should note that the report contains a glossary, which is presented in Appendix A.

1.1 PROGRAM OVERVIEW

The LDR MOTM terminal was developed as part of CECOM's OTM Tactical Satellite Communications Technology STO, managed by Mr. Gerald T. Michael. The antenna pointing system was developed under the Antennas for Communication Across the Spectrum STO for integration into the Milstar OTM Terminal under the OTM Tactical Satellite Communications Technology STO. The Antennas STO was managed by Steven R. Goodall and the positioner contract was managed by Steven P. Lescrinier. Harris Corp. developed the antenna pointing system and MIT Lincoln Laboratory developed the MILSTAR terminal, took responsibility for the terminal/vehicle integration, and outfitted the system with data logging capability. Among the program goals were the development and successful demonstration of an LDR MOTM terminal, and the collection and analysis of performance data to support future terminal development and transition decisions.

¹The term *satellite communications on the move* is common in military circles while the academic community seems to prefer the term *land mobile satellite communications*. In this report, the two are used interchangeably.

1.2 MILSTAR OVERVIEW

While an in-depth treatment of the MILSTAR system is clearly beyond the scope of this report, a brief discussion of certain features is relevant. MILSTAR is a constellation of geosynchronous satellites that provides highly robust, survivable, and antijam (AJ) communications capability to a wide range of military terminal types. Global connectivity is achieved with the system via intersatellite crosslinks at 60 GHz, circumventing the need for retransmissions through intermediate ground stations. Strong AJ performance is achieved through a combination of techniques. First, MILSTAR is a processing satellite with on-board demodulation and switching, a configuration that is inherently more jam resistant than conventional transponding satellites. In addition, the system operating frequencies, 43.5-45.5 GHz on the uplink (UL) and 20.2-21.2 GHz on the downlink (DL), contain a great deal of bandwidth that is exploited by the system through a process of fast frequency hopping. These operating frequencies also facilitate terminals with narrow beams, which are difficult to jam, and also lend themselves to low probability of intercept (LPI), and low probability of detection (LPD) communications. Moreover, at 44/20 GHz, relatively high gains can be achieved with small aperture antennas (e.g., 1-2 feet), leading to small, portable and mobile terminal implementations.

The MILSTAR I satellites support LDR communications from 75 bps to 2400 bps, while the MILSTAR II satellites support medium data rate (MDR) modes with throughputs up to 1.544 Mbps. The advanced EHF system, scheduled to launch in the 2007-2008 time frame, is a follow-on to MILSTAR that will support rates up to 8 Mbps. In addition, the Transformational Communications (TC) system, currently in the planning stages, appears to be taking on many of the characteristics of the MILSTAR system, including on-board processing and EHF operating frequencies. The features discussed above, as well as the military's obvious commitment to MILSTAR and "MILSTAR-like" systems all suggest that research and development in the area of MOTM terminals is well justified.

1.3 MOTM SYSTEM OVERVIEW

The MOTM terminal discussed in this report was integrated into a 1097A2 high mobility multi-purpose wheeled vehicle (HMMWV). A high level block diagram of the terminal is given in Figure 1. Various subsystems are described below.

1.3.1 Antenna Positioner Subsystem

The antenna positioner subsystem is based on open loop pointing (i.e., no signal or beacon tracking is used) and includes an attitude and heading reference system (AHRS), an antenna control unit (ACU), a servo chassis, and an antenna positioner assembly. The AHRS consists of the GPS antenna, precision lightweight GPS receiver (PLGR), and the Tactical Advanced Land Inertial Navigator (TALIN) 4000, all of which were government furnished equipment. The remaining components were supplied by Harris Corp. under contract to CECOM. The ACU implements the control algorithms that command the antenna to the proper pointing angle based on vehicle attitude and heading information supplied by the AHRS, rate gyro outputs, and resolver information. The ACU also supports logging capability for positioner performance assessment. The servo chassis contains the driver amplifiers for the motors in the positioner assembly. The positioner assembly, depicted notionally in Figure 2, is an inertially stabilized, low friction-to-inertia design




Figure 1. Simplified block diagram of the MOTM terminal.

with an elevation over azimuth mount. Both stages rely on direct drive DC brushless motors. Fiber optic gyros are used on each axis for rate loop feedback, while resolvers are used for position feedback. A slip ring and coaxial rotary joint were also selected to provide continuous coverage in azimuth. Note that the transmitter power amplifier and receiver low noise amplifier are placed on the antenna side of the rotary joint to minimize losses, and hence minimize the required antenna aperture. The MOTM terminal uses a 12" circular reflector antenna. Finally, one well-known issue with the elevation over azimuth pedestal design is the singularity that occurs at zenith, leading to increased pointing errors in this region. Despite this limitation, the design represents the simplest mechanical configuration and the smallest swept volume, both important factors in the overall terminal design. The entire assembly fits in a 15" diameter, 15" tall radome, and extends 6" below the mounting surface.

1.3.2 Testbed Terminal

The terminal itself is basically a modified version of the ASCAMP terminal, developed originally as a portable (and stationary) LDR MILSTAR terminal. Changes to the terminal for OTM use include the following. The ASCAMP time tracking algorithm was modified so that updates in the tracking loop are suspended when the discriminator output falls below a certain threshold, indicating the terminal is in blockage. Suspension of the time-tracking loop during blockage events improves performance since discriminator




Figure 2. MOTM antenna positioner assembly.

outputs under this condition are inherently noisy and would typically drive the loop out of its tracking window, leading to loss of timing. Also, the open loop space tracking algorithm was modified to account for displacement of the vehicle from its acquisition point in the calculation of slant range and look angles to the satellite. In addition, modifications to the terminal antenna controller were made to accommodate communication between the terminal control processor and the ACU. Finally, several changes in the RF chain associated with "remoting" the antenna from the terminal were made.

A great deal of support for field experimentation was designed into the MOTM system. The primary means of this support comes in the form of data logging from subsystems that include the TALIN, ACU, and ASCAMP terminal. Fieldworks model 8000 computers were used for TALIN and ACU logging. From the TALIN, vehicle attitude information, including linear and angular velocities and accelerations was logged. From the ACU, control loop information like resolver and gyro outputs, as well as pointing information such as azimuth and elevation error, was logged. Zxi Stations supplied by Zmicrosystems were used for ASCAMP terminal control and logging. These machines also hosted the end-user applications and blockage mitigation software. Information logged from the ASCAMP terminal included time tracking metrics, received power estimates (i.e., channel state), ephemeris information, etc. In addition, all logging computers recorded GPS time so that data files could be time-synchronized offline. A keyboard/monitor multiplexer




Figure 3. Vehicle layout (top view).

was used so that all computers could be controlled with a single interface from the operator's position in the passenger seat of the HMMWV. Figures 3 and 4 illustrate the vehicle and chassis layouts, respectively.

1.4 EXECUTIVE SUMMARY

The experiments summarized in this report fall into two broad classes: communications experiments and positioner experiments. The communications experiments are further broken into two subcategories: channel measurements and MEP performance experiments. The primary issue of interest in the communications experiments was characterizing the signal blockage phenomenon and quantifying its impact on the communications performance of the terminal. These experiments were basically constructed as protocol performance experiments where the channel state (i.e., fade depth) information was recorded "in the background." Ease of presentation dictated that the channel measurements be discussed separately from the MEP performance data. The result is a statistical analysis of the MOTM channel in Section 2 followed by a discussion of MEP performance over these same channels in Section 3. The focus of the positioner experiments was the characterization of the MOTM terminal's LOS pointing error over various terrain. Previous experience with the MOTM positioner in a laboratory setting suggested that its performance was so good that degradations in communications performance due to LOS pointing errors would be negligible. This fact was the primary motivation for separating the communications and positioner experiments. By separating these experiments flexibility was achieved whereby sites for the communications experiments could be tailored to specific blockage scenarios, without regard for terrain conditions. Likewise, sites for the positioner experiments could be tailored to terrain, without regard for blockage conditions.

Crypto (2U)	Test Ports (2U)
End-User Comp (2U)	Antenna Control Unit (5U)
Data Logging Comp #2 (3U)	Servo Control (3U)
Data Logging Comp #1 (3U)	Spare (2U)
ASCAMP Chassis (6U)	Computer Monitor Switching (4U)

Figure 4. Chassis layout.

1.4.1 Communications Experiments

The communications experiments were performed in three blockage environments: an “open” environment, a “rural” environment, and an “urban” environment. The environments have a loose connection to those proposed in [1], which were chosen to be representative of typical terminal operating conditions. Of course, the parameters from the field environments do not match exactly those used in the simulations from [1]. However, this fact does not invalidate the earlier results since the real world is comprised of a virtually unlimited number of channel blockage conditions. In fact, it is important to realize that the results presented here are not the first or only results to be presented for the land mobile satellite channel. Numerous references to measurement campaigns and statistical modeling exercises are given in this report and should be explored by interested parties to obtain a more complete picture. The names given to the various test environments discussed here were chosen for their similarity to labels used in the open literature. However, it should be noted that the term “open channel” does not literally mean blockage-free.² Instead, the term open is meant to connote a situation, such as might be found on the open highway, where there are only occasional obstructions to the LOS path between the mobile terminal and the satellite. Likewise, the term “rural” is meant to imply tree-lined roads where communications is often disrupted by the foliage. The term “urban” is simply intended to suggest that frequent blockages occur and are due primarily to tall buildings. One final note is that the urban environment discussed in this report should not necessarily be considered a worst case scenario. Initial experiments were conducted in dense urban, or urban canyon environments, but

²Experiments were conducted in blockage-free environments and revealed virtually no loss in performance relative to a stationary terminal. The primary reason for this performance is the exceptional quality of the antenna positioner system.

Table 1
Simplified, representative link budget for MOTM downlink.

Satellite transmit power	$P_t = 20.0 \text{ dBW}$
Satellite antenna gain	$G_t = 40.7 \text{ dB}$
Satellite EIRP	60.7 dBW
Path loss	-211.9 dB
Terminal antenna gain	33.5 dBi
Received power	$P_r = -117.7 \text{ dBW}$
Effective noise temperature	$T = 300 \text{ K} \equiv 24.8 \text{ dBK}$
Terminal G/T	$G/T = 8.7 \text{ dB}$
Boltzmann's constant	$K_b = -228.6 \text{ dBW/K/Hz}$
Noise density	$N_0 = -203.8 \text{ dBW/Hz}$
Carrier to noise ratio	$P_r/N_0 = 86.1 \text{ dB Hz}$

blockage times in these environments were so long (i.e., on the order of minutes) that the MOTM terminal would often lose timing, forcing termination of the experiment. To avoid this problem, a sparser or "light urban" environment, where blockage times were on the order of tens of seconds, was selected. Experiments in more stressing urban environments are planned for the future, after improvements to the MOTM terminal's time tracking algorithms are made. For reference, a notional link budget for the MOTM system is given in Table 1.

Channel Characterization

Blockage in the open and rural channels was due exclusively to foliage. Mean blockage percentages for these environments were approximately 20% and 50% respectively. Blockage in the urban channel was dominated by tall buildings but also contained some foliage. The blockage fraction for the urban environment was roughly 35% - 45%. Despite their similarity in terms of blockage fraction, there are a number of differences between the rural and urban channels, statistically speaking. These differences are due primarily to the blockage mechanisms in the two environments. For example, fade depths in the open and rural channels were observed to follow a lognormal distribution with mean fade depths on the order of -12 dB to -15 dB. Fades in the urban environment, on the other hand, tended to be much deeper, typically -35 dB to -40 dB, which is essentially the noise floor of the ASCAMP terminal. A density mixture consisting of lognormal and a power law were found to yield a better fit to the urban fading data compared to a lognormal distribution alone.

With respect to second order statistics, fade and connection durations are presented in units of meters (i.e., the results are normalized by the velocity of the HMMWV, which was logged via the TALIN during the experiments), facilitating the use of this data with arbitrary velocity profiles in future analyses. As a

function of received power relative to LOS power, the average connection durations in the open and rural environments vary from 10 meters to 200 meters, with minimum values achieved at roughly -10 to -15 dB. Because fade depths in the urban environment tend to be “all or none” there is less variation in the average fade and connection durations as a function of received power level. Average connection durations in the urban environment are approximately 25 meters and average fade durations are approximately 10 meters. Average fade durations in the open and rural environments range from 1 meter to 15 meters.

Discrete models that represent the data were also investigated. The most common model used in the open literature is the two-state Markov model. This model yields a reasonably good approximation for the urban channel where there is little ambiguity between channel state (e.g., good or bad) and an observable event (e.g., a fading event below some predetermined threshold). On the other hand, the two-state Markov model is not very well suited to the open and rural environments where fading is caused exclusively by foliage, resulting in fade depths that are more random. For these environments, hidden Markov models (HMMs) are more appropriate because they essentially decouple the notion of channel state from observable event. For example, HMMs allow for a situation where the channel is in a “bad” state but does not necessarily produce a fading event below a threshold. Among the simplest HMMs proposed for communications channels is Gilbert’s model [25]. This model was examined within the context of the MOTM channel and found, with the appropriate parameter matching scheme, to yield good results for the connection and fade duration distributions in the open and rural environments. HMMs with more than two states are likely to produce even better results, but due to their increased complexity are not examined here in-depth.

Finally, bit error rate (BER) tests were performed over the channels. Due to limitations of the test, it is not possible to perform burst or error gap analyses on the data. Instead, simple scatter plots that illustrate BER in 20 second intervals vs. estimated received power were produced. These plots illustrate the severe penalty in error rate associated with the fading channel. In addition, by measuring the error rate at both the MOTM terminal and a stationary terminal at the other end of link, an estimate of the MOTM UL and DL BER was obtained. As expected, a high degree of correlation between these two links was observed.

Protocol Performance

The purpose of the protocol performance experiments, presented in Section 3, was to assess the impact of the fading channel on a packet-oriented communications scheme with a hybrid automatic repeat request (HARQ) error control protocol. To this end, an experimental protocol that implements the desired HARQ error control scheme, was developed for use over the MOTM channel. In addition to the error control algorithms, the protocol, referred to as the MEP, also consists of a small collection of end-user applications (e.g., packetized voice, file transfer, etc.) and configuration mechanisms to facilitate experimentation. The MEP runs as a stand-alone application over a point-to-point link (e.g., a MILSTAR channel). Reasons for implementing the MEP in this way are twofold. First, this approach was the easiest way, in the time allocated for protocol design and implementation, to address issues that would facilitate the use of the protocol as a learning tool. These issues included instrumenting the protocol for support of offline performance analysis and providing a high degree of configurability. The second reason for this approach was that given the relatively low data rate channels being used in the MOTM system, issues relating to the use of TCP/IP over

satellite links were not appropriate for investigation in this portion of the STO.³ Hence, providing support for a standard TCP/IP interface was not a requirement.

The MEP was used to collect data on packet error rate (PER), throughput efficiency, and packet latency in the test environments characterized in Section 2. To first order, PERs can be approximated by the blockage fraction for each of the test environments. However, this approximation is somewhat optimistic since the actual PER depends on the packet length and channel fade duration. For the open channel, PERs ranged from 25% to 50%. For the rural channel PERs ranged from 40% to 50%, and for the urban channel PERs ranged from 40% to 70%. These PERs are orders of magnitude higher than typical wireless channels, due exclusively to signal blockage. With respect to throughput efficiency, the effects of packet length were validated. The results demonstrate a trade-off whereby both “short” and “long” packets suffer from poor throughput efficiency. In the case of short packets, the poor efficiency is due to overhead associated with the fixed length packet header. On the other hand, while increasing the packet length effectively amortizes this header over a longer duration, longer packets are more likely to experience burst errors due to channel blockage. Hence, there exists an optimum packet length that balances these effects. In the case of the MOTM system, the optimum packet length was shown to be approximately 300 ms. Although the optimum packet length was consistent across each of the test environments, the actual throughput efficiency achieved with this packet length varied from channel to channel. In the open channel, a maximum throughput efficiency of approximately 70% was achieved. Maximum throughput efficiencies of approximately 60% and 50% were achieved in the rural and urban channels, respectively.

Packet latencies were more difficult to interpret within an analytical framework due to the difficulties associated with analyzing the queueing behavior of ARQ protocols and the effects of correlated errors. Hence, experimental results were compared to simplified analytical expressions. Depending on the channel and the time allowed between transmission attempts, average packet latencies ranged from 2 seconds to 20 seconds. Note, however, that one caveat to these results is that the MEP implementation is known to be responsible for approximately 450 ms of latency per transmission attempt, from an end-to-end latency budget totaling roughly 1 second per transmission attempt. An optimized implementation of the protocol could therefore be expected to reduce the latencies achieved in the experiments by almost 50%.

1.4.2 Positioner Characterization

The positioner experiments were performed over three terrain types: improved road, unimproved road, and rough terrain (i.e., off-road). Moreover, given the sensitivity of the two-axis design to elevation angle, experiments were performed at both 40° and 65° elevation angles. As opposed to the channel measurement and protocol performance data, there is little in the open literature with respect to the characterization of LOS pointing error in SOTM systems with directive antennas. Nonetheless, in Section 4 a statistical analysis similar to the one performed for the fading channel measurements is presented for the LOS pointing error data. The Laplacian density was found to be an excellent model for the azimuth and elevation mispointing errors. By extension, a probability density model for the LOS pointing error was derived and presented in

³TCP over satellite is currently under investigation in the wideband OTM portion of the CECOM STO.

the form of a definite integral. The models fit the data from all courses and elevation angles exceptionally well, albeit with different parameters. For completeness, estimated densities for the UL and DL antenna gains associated with the LOS pointing error were given assuming the 12" reflector antenna used by the MOTM terminal. Average fade and connection durations as a function of LOS error were also presented. In the case of connections, average durations increase linearly to their maxima, a value defined by the experiment duration. Likewise, as expected, average fade durations generally decrease as a function of LOS error to their minima, a value defined by the sampling rate of the experiment. A maximum level crossing rate of approximately 7 Hz was observed from the data. This maximum occurs anywhere from 0.02° to 0.1° , depending on the terrain and elevation angle of the experiment.

2. EHF SOTM CHANNEL CHARACTERIZATION

The focus in this section of the report is the characterization of random variations in received signal power that can be attributed to vehicle motion. Although stationary terminals also experience fluctuations in received signal power, these variations are typically due to weather and atmospheric effects that take place on a relatively long time scale. Vehicle motion, on the other hand, introduces fluctuations in received signal power that occur much more rapidly. Hence, an understanding of this phenomenon is critical to waveform design, error recovery schemes, performance analysis, and other EHF SOTM system engineering issues. Of particular interest here is signal fading caused by objects in the propagation path, such as buildings, trees, or other terrain features. Note that signal fading may also be caused by antenna mispointing due to an imperfect antenna positioner system. However, this phenomenon will not be addressed here. Instead, it will be examined separately in Section 4. From a modeling perspective it makes sense to treat these two types of fading separately, since their underlying physical mechanisms are completely different. However, from an experimental viewpoint, they may also be treated separately because, as will be seen in Section 4, the MOTM positioner delivers such excellent performance that in most cases the resultant signal fading is quite small. This is true on both improved and unimproved roads, where all of the experiments discussed in this section were conducted. In fact, the fade depths caused by antenna mispointing are so shallow that they are effectively hidden when MOTM terminal's hard-limiting front-end is in saturation, which is often the case. Hence, the effects of antenna mispointing are assumed to be negligible in the statistical analysis of the fading measurements discussed here.

The results presented here validate certain approaches used in modeling land mobile satellite (LMS) channels and call into question other common approaches. To first order, the probability density mixture models typically cited in the open literature appear to work quite well with the MOTM experimental data. Fade depth is shown to follow a lognormal distribution while a density model specific to the MOTM terminal implementation is used for the LOS (i.e., unfaded) condition. On the other hand, the two-state Markov model that has been proposed for LMS channels on numerous occasions does not always accurately depict the second order statistics of the MOTM channel data. Instead, a slightly more sophisticated hidden Markov model (HMM) is found to be more appropriate.

2.1 EXPERIMENT OVERVIEW

Approximately 80 experiments were conducted with the MOTM terminal over a three month period from the late summer to the early fall of 2002. Even when channel measurement was not the explicit purpose of an experiment, received signal power was always recorded so that this information was available for calibrating experiment results to the channel dynamics. Although technically not required for monitoring received signal power, a point-to-point call was usually initiated between the MOTM terminal and an ASCAMP terminal located on the roof of building B at MIT Lincoln Laboratory, as depicted in Figure 5. Experiments were always initiated "on the pause," and once proper operation was verified the vehicle would




Figure 5. Channel measurements and protocol performance experimental setup. Typically, a stationary ASCAMP terminal, located at Lincoln Laboratory, and the MOTM terminal were each simultaneously transmitting and receiving in a point-to-point call over MILSTAR.

begin moving.⁴ Experiments were coordinated between MOTM and ASCAMP terminal operators via cell phone. For consistency, all end-to-end experiments were conducted between the MOTM terminal and the same ASCAMP terminal, namely terminal ID number 3944. In addition to relevant experimental data, extensive logging of terminal diagnostics, including time tracking metrics, ephemeris information, etc. was conducted at both terminals. Also, comments were frequently entered in real time into the terminal log files to facilitate offline processing of the recorded data. In addition to terminal logging, vehicle attitude and heading information produced by the AHRS was logged via a separate computer interface. Vehicle velocity is another key parameter returned by the AHRS since it allows for the possibility of presenting certain channel statistics in units of distance (e.g., meters) so that their interpretation is not tied to a specific velocity. Figure 6 illustrates typical velocity profiles.

In addition to the protocol experiments, which are summarized in Section 3, end-to-end BER experiments were performed using the ASCAMP and MOTM terminals' built-in test. These experiments were conducted by first establishing a MILSTAR call between the two terminals and then initiating the test. Initiation of the built-in BER test is a two step process. First, a short (i.e., 63 bit) repeating pseudo-random bit pattern is transmitted by one terminal. The other terminal is configured to process the received pattern and record the number of errors that occur. The number of errors encountered are reported every 20 seconds. The

⁴It is worth noting that on numerous occasions the MOTM terminal achieved acquisition and re-acquisition OTM at speeds as high as 55 mph. However, these events were always in blockage free environments.




Figure 6. Typical velocity profile in the (a) rural environment and (b) urban environment.

MOTM BER tests were performed in both directions simultaneously (i.e., each terminal was both sending and receiving the pseudo-random bit patterns). With this test procedure it is possible to gain some insight into the MOTM terminal's uplink performance, as will be discussed in Section 2.3. Also, it is important to note that the BER tests were initiated while the MOTM terminal was stationary and the pseudo-random bit patterns were "locked" at each terminal before the MOTM terminal was allowed to move. Without this crucial step, many of the errors made during fading events would be ignored while the receiving terminal tried to resynchronize with the known pattern, producing overly optimistic results. Unfortunately, the framework of the BER test as described above does not lend itself to extensive analysis of second order statistics such as burst error lengths, etc. Instead, the relatively coarse sampling interval and minimal output parameters supported by the test yield insight only into the average BER of the channel. However, in many cases necessary error statistics may be inferred from the statistical analysis of the channel itself.

2.1.1 Measurement Approach

All experiments discussed here were conducted in the 2+54 DL mode and 1VV UL mode. In the future, additional experiments may also be conducted in other modes, but the 1VV/2+54 configuration represents a good trade between link robustness and satellite resource efficiency. Terminals used in the experiments were typically set to log the outputs from the processed high hop rate (HHR) course sync hops, which are processed further by the terminal to yield estimates of the received power to noise ratio, P_r/N_0 . The ASCAMP P_r/N_0 estimation procedure, developed originally in [2, 3] is summarized briefly below. Referring to the block diagram in Figure 7, the received signal during the k th chip, $x_k(t)$, is modeled as:

$$x_k(t) = A_k \cos(2\pi f_0 t + \phi) \quad (1)$$




Figure 7. Block diagram of ASCAMP hard-limiting front-end and P_r/N_0 estimator.

where A_k is the amplitude of the received signal during the k th chip, assumed to be of duration T_c , and ϕ is the signal phase, assumed to be constant over the duration of a hop, T_h . The signal is assumed to be corrupted by additive noise and filtered by a bandpass filter centered at frequency f_0 . The noise density at the output of the filter is assumed to be Gaussian:

$$p(n) = \frac{1}{\sigma_n \sqrt{2\pi}} e^{-n^2/2\sigma_n^2} \quad (2)$$

The filtered output is sampled on an interval $T_s \ll T_c$, and hard-limited to values ± 1 . Quadrature separation is performed digitally with simple reference signals yielding output sequences:

$$i_k = \dots, \text{sgn}(A_k \cos \phi + n_0), 0, -\text{sgn}(-A_k \cos \phi + n_2), 0, \dots \quad (3)$$

$$q_k = \dots, 0, \text{sgn}(A_k \sin \phi + n_1), 0, -\text{sgn}(-A_k \sin \phi + n_3), \dots \quad (4)$$

where $\text{sgn}(\cdot)$ is the sign operator. These outputs are then accumulated over a chip duration, yielding front-end outputs I_k and Q_k . During communications hops, these quantities are passed to a detector, summed, quantized and passed to the error correction decoder. For P_r/N_0 estimation, I_k and Q_k are accumulated over the duration of one hop. These accumulations, I_h and Q_h , contain $N = \frac{T_h}{2T_s}$ nonzero i and q samples, respectively. In an equivalent receiver without the hard-limiter, the mean values of these outputs would be directly related to $A_k \cos \phi / \sigma_n$ and $A_k \sin \phi / \sigma_n$ allowing for traditional detection methods to eliminate the phase dependence. Although the hard-limiter greatly reduces the complexity of the hardware, it changes the statistics of $E[I_h]$ and $E[Q_h]$, resulting in a nonlinear dependence on signal to noise ratio (SNR) and phase:

$$\frac{E[I_h]}{N} = \sqrt{2/\pi} \int_{-\frac{A_k}{\sigma_n} \cos \phi}^{\infty} e^{-n^2/2} dn - 1 \quad (5)$$

$$\frac{E[Q_h]}{N} = \sqrt{2/\pi} \int_{-\frac{A_k}{\sigma_n} \sin \phi}^{\infty} e^{-n^2/2} dn - 1 \quad (6)$$

The P_r/N_0 estimator compensates for this characteristic by applying an inverse transformation such that:

$$\lambda \left(\frac{I_h}{N} \right) \approx \frac{-A_k}{\sigma_n} \cos \phi \quad (7)$$

$$\lambda\left(\frac{Q_h}{N}\right) \approx \frac{-A_k}{\sigma_n} \sin \phi \quad (8)$$

where $\lambda(\cdot)$ is the transformation that attempts to invert (5) and (6). In the terminal a single lookup table is used to apply the divide by N , perform the inverse transformation, square the result and scale it by a constant, γ , defined to ensure that the output utilizes the available dynamic range. The sum of the I_h and Q_h table lookups are added together to yield a single number for the accumulation:

$$\Gamma = \gamma \left[\lambda\left(\frac{I_h}{N}\right) \right]^2 + \gamma \left[\lambda\left(\frac{Q_h}{N}\right) \right]^2 \quad (9)$$

where Γ is the accumulation value. An estimate for P_r/N_0 may then be formed as:

$$\frac{P_r}{N_0} \approx \frac{B_n}{2\gamma} \Gamma \quad (10)$$

where B_n is the noise bandwidth of the bandpass filter. The original P_r/N_0 estimator averaged 64 accumulations and calibrated the estimate based on laboratory measurements as follows:

$$\frac{P_r}{N_0} \text{ (dB)} = 30.11 + 10.369 \log \left(\frac{\sum_{i=1}^{64} \Gamma_i}{64} \right) \quad (11)$$

The first term above effectively sets the noise floor of the estimator. Given this noise floor, fade depths below approximately 40 dB to 45 dB can not be resolved by the estimator. Up to 75 dB-Hz, this estimator has a normal distribution with a standard deviation of 0.2 dB. Beyond this point, the estimator saturates, its distribution becomes bimodal and the standard deviation increases. In [4], the range of the estimator was extended to approximately 85 dB-Hz with additional processing and averaging. Unfortunately, even averaging over the 64 accumulations in the original estimator is undesirable from a time-varying channel measurement viewpoint, since the averaging reduces the time resolution of the estimates, making it more difficult to isolate individual fading events. Hence, only single accumulation values are used to estimate P_r/N_0 in this report. A probability density function of the P_r/N_0 estimates when the terminal is in saturation has been derived numerically using the actual ASCAMP lookup table and assuming a uniform distribution for the signal phase. This density is illustrated in Figure 8a and compared to measured data in Figure 8b. The bimodal nature of the density is readily apparent in the figure, and does influence the statistical analysis that follows. However, since signal blockage effectively takes the terminal out of saturation, the undesirable behavior of this “one-shot” estimator is restricted to the unblocked, or LOS regions and is believed to be a reasonably good estimator of fade depths down to approximately 40 dB below the average received LOS power level.

By logging the output of this estimation process, a time record of the received signal power is obtained. Figure 9 illustrates sample recordings for the various test environments. For convenience, results are normalized to an average received signal power of 0 dB in the unshadowed portions of the timeseries data. Fading in the open and rural channels is caused exclusively by foliage. The wide range of fade depths exhibited by these channels can be attributed to the variations in foliage depths and densities. The urban channel also contains some foliage, but the majority of blockage events can be attributed to buildings. For this reason, the urban channel contains a much larger fraction of deep fades (i.e., ≤ -30 dB) compared to the open and rural channels. This fact is also apparent in the results of the statistical analysis presented in Section 2.2.




Figure 8. Probability density of the P_r/N_0 estimates when the terminal is in saturation. (a) Theoretical density obtained via numerical analysis. (b) Density obtained from measured data.

2.1.2 Course Descriptions

Clearly, there are as many channel types in the world as there are combinations of roads, terrain, vehicle speeds, satellite look angles, etc. Moreover, the scope of the experiments summarized in this report is broader than just channel measurement. For these reasons, a relatively few representative channels were studied in depth, as opposed to conducting fewer experiments over a wider range of environments. Future efforts may concentrate on complementing the data presented here with additional measurements in other environments. The channels in this section of the report are referred to as open, rural, and the urban. As discussed in the introduction to this report, the term “open channel” is somewhat of a misnomer in the sense that this particular channel is not really blockage-free. However, the channel is open the vast majority of the time, as would be the case on the open highway, for example. The open and rural channels are defined by road courses at Ft. Devens, MA. These courses consist of tree lined roads that result in LOS obstructions to varying degrees. A variety of tree types exist on these courses, including, pine, birch, poplar, ash, and a few oak trees. Because the tests were all performed in the summer and early fall, the trees were all in-leaf. The urban channel is defined by a road course between Kendall Sq. and Central Sq. in east Cambridge, MA. Although LOS path obstructions on this course were caused mostly by tall buildings, there was some foliage as well. For all of the experiments the MILSTAR satellite was located in approximately geosynchronous orbit with a longitude of the ascending node equal to approximately -90° and inclination of approximately 3° , resulting in a satellite elevation of approximately 41° for the experiments discussed here. More detail on each of the courses is provided below.




Figure 9. Time series data segments from the three test environments. Each segment is two minutes in duration. (a) Open environment. (b) Rural environment. (c) Urban environment.

Open Course


Referring to the map of Ft. Devens, MA in Figure 10a, the open course begins at the intersection of Dixie and Harvard Rd and proceeds north on Dixie to the intersection of Dixie and Jackson Rd. At the intersection, a u-turn is made and the vehicle proceeds south on Dixie back to the intersection of Dixie and Harvard. The tree heights along this route are such that traversing the same route on opposite sides of the road yield slightly different blockage conditions. For the most part, the tree types consist of pine and a few oak. The trees are relatively sparse with foliage depths along the LOS path on the order of a few meters. The course is approximately 2 miles long and consists entirely of improved roads. Typical experiment durations in this environment were on the order of 5 to 6 minutes. Figure 10b contains a photograph of the HMMWV in the open environment.

Rural Course

Referring to the map of Ft. Devens, MA in Figure 10a, the rural course begins at the corner of Shirley and Bivouac Rd. The vehicle proceeds SW on Bivouac towards the Turner drop zone. Next, the vehicle turns left on the unmarked road that follows the tree line at the southern end of the drop zone. This road eventually crosses another unmarked road that cuts into the drop zone. The vehicle turns left on this road then right on the main road through the middle of the drop zone, which eventually crosses Shirley Rd. Next, the vehicle turns left on Shirley and proceeds north back to the starting point. Along the western and southern edges of the course the foliage consists of a combination of pine, birch and poplar, while along the eastern edge of the course the tree types include pine, ash, and oak. The foliage is much denser along this route with many layers of trees lining the roads. In several locations the trees hang over the road to form canopies. This course is approximately 3 miles in duration and consists of approximately 50% improved roads and 50% unimproved roads. Typical experiment durations in this environment were on the order of 7 to 8 minutes. Figure 10c contains a photograph of the HMMWV in the rural environment.

Urban Course

Referring to the map of Cambridge, MA in Figure 11a, the urban course starts at the corner of Third Ave. and Broadway, near Kendall Sq. The vehicle proceeds along Broadway and turns left on Inman St. Inman is followed past Cambridge City Hall to Mass Ave., where a left turn is made. Next, the vehicle moves east along Mass Ave. through Central Square to Vassar St. A right turn is made on Vassar and this road is followed to Memorial Dr. Next, the vehicle makes a left on Memorial and follows this road past the Longfellow bridge until it turns into Edwin Lands Blvd. From Edwin Lands, the vehicle turns left onto Binney St., then left onto Third Ave., back to the starting point. The segment from Central Sq. back to the starting point at Third and Broadway contains very long stretches of unobstructed views to the satellite. Because the purpose of the experiments in this environment was to characterize dense urban environments, this portion of the course was removed from the data sets offline. The remaining portion, from Kendall Sq. to Central Sq. is approximately 2 miles long and consists entirely of improved roads (numerous potholes notwithstanding). Because of the wide variations in vehicular and pedestrian traffic in this environment, experiment durations ranged from 6 to 12 minutes. Figures 11b and 11c contain photographs of the HMMWV in the urban environment.



(a)




(b)



(c)

Figure 10. Open and rural environments. (a) Partial map of Ft. Devens, MA with open course highlighted in green and rural course highlighted in blue. (b) MOTM system in the open environment. (c) MOTM system in the rural environment.



(a)



(b)



(c)

Figure 11. Urban environment. (a) Partial map of east Cambridge, MA with the urban course highlighted in red. The solid line represents that portion of the course extracted offline for analysis. (b) MOTM system near Kendall Sq. (c) MOTM system in Central Sq.




Figure 12. Data description. A time series data segment from the urban channel is used to help illustrate several parameters of interest.

2.2 STATISTICAL ANALYSIS

In this section statistical models for the LMS channel are briefly reviewed. Emphasis is given to approaches that are most suitable to EHF LMS systems. Experimental results for the MOTM system are presented within the context of these models and comparisons are made as appropriate. Statistical properties addressed here include probability density models, discrete models, and second order statistics such as time share of fades, average fade and connection duration, and level crossing rate. Although each of these terms are defined more rigorously in the sequel, Figure 12 is a useful reference that helps explain many of the parameters. The timeseries data in the figure were recorded with the MOTM system in the urban environment.

2.2.1 Probability Density Models for Received Signal Power

Numerous probability density models for received signal power in a LMS channel have been proposed in the open literature. These models tend to fall into one of two broad categories [5]: single models and mixture models. Single models are comprised of a single statistical distribution, and hence are appropriate for stationary channels where the statistics remain approximately constant over the observation interval of interest. On the other hand, mixture models are weighted sums of multiple statistical distributions. These models are developed for nonstationary channels where the statistics vary over a time period or area of interest.

In virtually all statistical models for received signal power in LMS systems, the dominant propagation phenomenon is taken to be signal fading. Typically, two classes of signal fading are cited. First, there is rapid fading due to the constructive and destructive combining at the receiver of randomly phased reflections of the transmitted signal. This type of fading is referred to as multipath fading and the resulting fluctuations in received signal power, S , are described by the central chi-square distribution:

$$p_{cc}(S|S_0) = \frac{1}{S_0} e^{-S/S_0} \quad (12)$$

where S_0 is the mean received power, due entirely to multipath. This process is also referred to as Rayleigh fading since the received signal envelope follows the Rayleigh distribution. The model in (12) holds well in practice for microwave cellular systems where typically no LOS path exists between the transmitter and receiver. In LMS systems where a LOS component is often present, the Ricean distribution characterizes the received signal envelope and a noncentral chi-square distribution is used to describe received power:

$$p_{nc}(S) = K e^{K(S+1)} I_0(2K\sqrt{S}) \quad (13)$$

where K is the direct-to-multipath signal power ratio, sometimes referred to as the Rice factor and I_0 is the modified zeroth order Bessel function of the first kind. Note that the distribution in (13) has been normalized so that the power of the LOS link is unity (i.e., the mean of the distribution is $E\{S\} = 1 + 1/K$). The second type of fading is referred to as shadowing and is caused by the complete or partial blockage of the transmitted signal by objects in the propagation path such as buildings, terrain, etc. Variations in signal power due to shadowing have been observed to follow a lognormal distribution [6–8]:

$$p_{ln}(S_0) = \frac{10}{S_0 \sqrt{2\pi\sigma \ln 10}} \exp\left(-\frac{(10 \log S_0 - \mu)^2}{2\sigma^2}\right) \quad (14)$$

where μ is the mean power decrease, in dB, and σ^2 is the variance of the received power due to shadowing. The effects of signal shadowing are dependent on a number of factors, including obstacle type, path length through the obstacle, elevation angle, carrier frequency, etc.

In [6], Loo proposed a single model whereby the received signal envelope follows a Ricean distribution (to account for multipath fading) with a lognormal mean (to account for signal shadowing). This model is quite popular in the open literature and has been applied to numerous frequency bands, including UHF [6,9], S-band [10], and Ka-band [11]. Note, however, that in many of these references the Loo single model is used as one density in a larger mixture model. In [5] an alternative to the Loo model is proposed whereby lognormal shadowing is modeled via a gamma distribution. The advantage of this approach is improved mathematical tractability with little or no sacrifice in fidelity. Numerical distributions for LMS systems at Ka-band are provided in [12].

In [13] one of the first mixture models for the LMS channel was proposed. The basis for the authors' approach was the observation that in LMS systems signal shadowing results in the removal of the LOS component from the received signal, yielding completely different channel statistics compared to the non-shadowed case where the LOS signal component is present. Hence, a time share approach is proposed where a Ricean distribution is used to describe the channel behavior during unshadowed times and a Rayleigh distribution conditioned on a lognormal random variable [14, 15] is used to characterize the received signal

during shadow intervals. The latter density is sometimes referred to as a Suzuki density and is popular in modeling terrestrial wireless systems where a LOS signal component is rarely present. The so-called Lutz model is then given by:


$$p_{lutz}(S) = (1 - B) p_{nc}(S) + B \int_0^\infty p_{cc}(S|S_0) p_{ln}(S_0) dS_0 \quad (15)$$

where $0 \leq B \leq 1$ is the time share of shadowing, or blockage fraction, and represents the fractional time that the LOS component is absent from the received signal. Although proposed originally for L-band systems the Lutz model has also been applied at K-band [16]. Further mathematical analysis of the Lutz model is contained in [17]. Other mixture models with two distributions have been proposed in [18, 19]. Mixture models with more than two distributions are proposed in [9, 20, 21].


Data collected with the MOTM system have been used to develop probability density models for received signal power. Figure 13 illustrates the density functions obtained from the data measured in each of the three test environments. In order to estimate the densities the time series data were first normalized to a LOS power level of 1. Next, a histogram with 10000 equally spaced bins between 0 and 10 was generated. The bin values were then normalized by both the total number of samples in the time series data and the bin size (i.e., 0.001). Data collected from approximately 20 separate experiments in each environment were used to produce the estimates in the figure. A mixture model is proposed here for characterizing the distribution of received signal power in the MOTM terminal. Like the Lutz model, this density uses a weighting factor, B , that is coupled to the fractional time that the LOS component is blocked:

$$p_{motm}(S) = (1 - B) p_{hl}(S) + B p_{ln}(S) \quad (16)$$


The first density, $p_{hl}(S)$, characterizes the behavior of the MOTM terminal in the presence of a LOS signal component. This density was obtained via numerical analysis of the MOTM terminal as discussed in Section 2.1.1 and illustrated in Figure 8a. As discussed in Section 2.1.1 the bimodal nature of this density is due to saturation effects of the terminal's hard-limiting front-end [2]. These effects can also be observed in the numerical pdfs plotted in Figure 13. However, around the 0 dB point, the densities in this figure do not match p_{hl} exactly as shown in Figure 8a. This difference is due to the presence of relatively shallow fades that bring the terminal just out of saturation. Note also, that the mixture models do a relatively good job of reproducing this effect. The second density, $p_{ln}(S)$, given by (14), is used to characterize received power levels in the absence of a LOS component. Use of the lognormal density in (16), as opposed to the Suzuki density that appears in the Lutz model, reflects a belief that virtually no multipath energy is present in the received signal during shadow events. Moreover, virtually no multipath energy is assumed to be present in the unshadowed regions either. Unfortunately, validating this assumption here is not possible since determining the contribution, if any, of multipath in the LOS regions is prevented by the saturation effects of the MOTM terminal's hard-limiting front-end. Nonetheless, multipath fading is believed to be negligible in the MOTM system for two reasons. The first reason is that the narrow beamwidth antenna (i.e., approximately 3.5° at 20 GHz) rejects spatial energy that might be present in directions other than the LOS. The second reason is that relative to the short wavelength of the MOTM received signal, most reflective surfaces will appear rough, resulting in diffuse scattering of the transmitted signal as opposed to the specular reflections typically associated with multipath propagation [7].



(a)



(b)



(c)

Figure 13. Probability densities for estimated received signal power in the three test environments. Each plot contains the density estimated from the measured data and a theoretical density based on the mixture model in (16). An alternate mixture model, based on three densities is also shown for the urban channel. (a) Open channel. (b) Rural channel. (c) Urban channel.

Table 2

Density model parameters derived from a least squares fit to the measured data.

Environment	$B_{(1)}$	μ (dB)	σ (dB)	B_2	α	β
open	0.227	-12.11	11.07	-	-	-
rural	0.484	-14.11	12.06	-	-	-
urban - 2 state	0.335	-37.09	11.75	-	-	-
urban - 3 state	0.136	-10.52	8.59	0.198	0.001	1.72

Least squares fits of the model in (16) are plotted in Figure 13 for each of the three test environments. Note from the figure that the model yields good results for the open and rural channels. However, the model is less accurate for the urban environment. This result could be due to the fact that blockages in the urban environment are caused not only by foliage, but also by buildings, yielding a high concentration of deep fades in addition to the less severe ones due to foliage. Hence, it is difficult for a single distribution to match accurately the range and associated probabilities of observed fading levels. For the urban channel measured by the MOTM system, a slightly more accurate representation can be achieved with a mixture model consisting of three terms:

$$p_{motm_3}(S) = (1 - B_1 - B_2) p_{hl}(S) + B_1 p_{ln}(S) + B_2(\alpha S^{-\beta}) \quad (17)$$

where α and β are constants, yielding a power law for the third term. The model proposed above has no real basis in the underlying shadowing process. Instead, it is based on the heuristic observation that the urban channel pdf appears to be more linear than its open and rural channel counterparts, suggesting that a power law approximation may be appropriate. Although mixture models with more terms may yield even better results, it is important to remember that as the number of densities increases these models quickly grow unwieldy and lose their utility as a simple and insightful characterization of the fading channel. The parameters associated with the least squares fit of the models to the experimental data are given in Table 2. Note the relatively large standard deviations associated with the lognormal density that describes the blockage fading, suggesting that for each of the three test environments a large variation in fade depths was experienced. This result is easily attributed to the wide variety of building heights, tree depths, etc., that were present in each of the test environments.

A few comments regarding the proposed mixture models in (16) and (17) are in order. First, because of the unique hard-limiting front-end in the MOTM terminal, the models in (16) and (17) do not have broad applicability. However, an alternate mixture model that replaces $p_{hl}(S)$, with a more suitable density could easily be used. Assuming again that very little multipath energy is contained in either the shadowed or unshadowed EHF received signal, the most appropriate density to replace $p_{hl}(S)$ is probably the gamma density [22] with a large shape factor, m . In situations where the gamma distribution describes the received signal power, the signal envelope follows the Nakagami distribution. Compared to the noncentral chi-square distribution, the gamma distribution is probably more appropriate for EHF SOTM channel modeling because of its ability to characterize situations in which there is little multipath fading. Theoretically, the




Figure 14. Complementary cdfs for estimated received signal power in the three test environments.

noncentral chi-square distribution can also model these scenarios, but the resulting Rice factors are large and lead to numerical problems with the Bessel function in (13). On the other hand, as the shape factor $m \rightarrow \infty$, the gamma distribution becomes an impulse (i.e., no fading). Recall that the gamma distribution has also been proposed as a suitable alternative to the lognormal distribution for analytical and numerical tractability reasons [5]. Hence, another alternative to (16) is a mixture model comprised of two different gamma distributions: one with a large shape factor m to model the channel behavior in the unshadowed state, and one with a shape factor appropriate for the expected fading conditions in the shadowed state.

Complementary cumulative distribution functions (cdfs) for the three test environments are given in Figure 14. Note from the figure that the probability of exceeding a received power level greater than roughly -12 dB (relative to the normalized LOS power level), is actually larger for the urban channel than for the rural channel. This fact is also reflected in the time share of shadowing parameter for the two channels, given in Table 2. Because the density that describes the received signal in the LOS region, p_{hl} , has a range of approximately ± 1.5 dB, the time share of shadowing for the rural channel is actually larger than for the urban channel. This result stands to reason in light of the previous discussion regarding the high concentration of deep fades in the urban channel. In fact, the heavier “tail” of the urban channel cdf is apparent in the figure. Thus, the probability of exceeding received power levels less than approximately -12 dB is greater for the rural channel compared to the urban channel.

Note that the cdfs in Figure 14 are a convenient way to assign fade margins in LMS systems. The concept of fade margin is introduced in [23] and defined simply as the threshold below which the channel is considered to be in a fade. In [23], the threshold is defined relative to the average received signal power, whereas here it is more appropriate to define the threshold relative to average received signal power when the LOS signal component is present, as in Figure 12. Presumably, the system is designed to achieve acceptable performance during nonfades. Hence, the cdfs in Figure 14 may be used to approximate system availability




Figure 15. Two-state Markov model.

based on a given fade margin. For example, given a fade margin of 15 dB, the figure shows that the open channel provides approximately 90% availability while the rural and urban channels yield roughly 76% and 73% availability, respectively. Conversely, if 95% availability is desired, the fade margin required for the open channel is 20 dB, while 30 dB is required for the rural channel and 40 dB is required for the urban channel. These results clearly indicate that the allocation of link margin to the problem of signal fading in the MOTM system is not feasible due to the excessive margins required. In certain cases, the application of forward error correction (FEC) coding can be applied to reduce the fade margin. However, as will be seen in Sections 2.2.3 and 2.2.4, a significant fraction of fading events last for several seconds or even longer, necessitating the use of prohibitively long block lengths and/or interleavers.

2.2.2 Discrete Models

Discrete models are a useful and convenient means for characterizing complex channels. Of particular interest are *generative* models [24], which when appropriately parameterized will produce outputs that match the statistics of the channel. First-order Markov models have been proposed by numerous researchers as a suitable discrete model for the LMS channel [9, 13, 16, 17, 21]. The so-called two-state version of this model is developed here for the MOTM channel. In addition, a simple extension to the conventional two-state model, proposed originally by Gilbert [25] for modeling burst errors in telephone channels, is explored and found to yield very good performance.

Two-State Markov Model

The two-state Markov model is especially popular due to its simplicity and good performance. Moreover, this model seems well suited to the MOTM channel because of the dominance that a single propagation effect, namely signal blockage, has in determining the quality of the received signal. The situation is depicted in Figure 15. According to the model, the channel is comprised of two states: a “good” (i.e., unblocked) state where the LOS signal component is present or acceptable communications performance is achievable, and a “bad” (i.e., blocked) state where the LOS signal component is absent and communications is either severely degraded or impossible. When the channel is in the good state, **g**, there is a probability, p_{gg} , associated with remaining in that state and a crossover probability, p_{gb} , associated with the transition to the bad state, **b**, such that $p_{gg} + p_{gb} = 1$. Likewise, there is a probability, p_{bb} , associated with remaining in

the bad state and a probability, p_{bg} , associated with switching from the bad state to the good state such that $p_{bb} + p_{bg} = 1$.

An equivalent matrix representation of the two-state Markov model is given by the following:

$$\mathbf{M} = \begin{bmatrix} p_{gg} & p_{gb} \\ p_{bg} & p_{bb} \end{bmatrix} \quad (18)$$

where the element in the first row and first column describes the probability associated with remaining in the good channel state, the element in the first row and second column describes the transition from the good to bad state, the element in the second row and first column describes the transition from the bad to good state and the element in the second row and second column describes the probability associated with remaining in the bad channel state.

Given the representations of the two-state Markov model defined above, several properties of interest may be defined as follows. Because the Markov model is ergodic, there is an associated “long run,” or steady state probability that describes the average time spent in each state. These quantities are found as the solution to the system [26]:

$$\begin{aligned} \pi \mathbf{M} &= \pi \\ \pi \mathbf{1} &= 1 \end{aligned} \quad (19)$$

where π is the 1×2 row vector of steady state probabilities and $\mathbf{1}$ is a 2×1 column vector of ones. The solution to (19) is given by:

$$\pi = [\pi_g \quad \pi_b] = \left[\frac{p_{bg}}{p_{gb} + p_{bg}} \quad \frac{p_{gb}}{p_{gb} + p_{bg}} \right] \quad (20)$$

Note that the dwell time for each state can be described by a geometric random variable [27] with a probability mass function:

$$\Pr \{D_g = n\} = p_{gb} p_{gg}^{n-1} \quad (21)$$

$$\Pr \{D_b = n\} = p_{bg} p_{bb}^{n-1} \quad (22)$$

where D_g , and D_b are the dwell times for the good and bad states, respectively. Given these mass functions, the mean dwell times are easily computed:

$$\overline{D_g} = \sum_{n=1}^{\infty} n p_{gb} p_{gg}^{n-1} = \frac{1}{p_{gb}} \quad (23)$$

$$\overline{D_b} = \sum_{n=1}^{\infty} n p_{bg} p_{bb}^{n-1} = \frac{1}{p_{bg}} \quad (24)$$

Finally, exploiting the connection with geometric random variables, complementary cdfs are given by:

$$\Pr \{D_g > n\} = (1 - p_{gb})^n = p_{gg}^n \quad (25)$$

$$\Pr \{D_b > n\} = (1 - p_{bg})^n = p_{bb}^n \quad (26)$$

Note that dwell times above are defined in terms of samples, n . To get the corresponding times in seconds, one must multiply by the sampling interval, T_s . In addition, converting to meters or bits is a simple process

of multiplying D_g and D_b by vT_s or R_dT_s , respectively where v is the vehicle velocity in meters/sec. and R_d is the data rate in bits per second (bps).

There are two ways to reconcile Markov models like the two-state model described above with density mixture models such as those discussed in Section 2.2.1. One interpretation is that there exists a one to one correspondence between the densities in a mixture model and the states in a Markov model such that the channel behavior in a Markov model state is completely described by a single density in the corresponding mixture model. In this case, the blockage fraction, B , in the Lutz model and the two-state Markov model parameter π_b , are equivalent. Note also that with this interpretation the densities in the mixture model are assumed to be disjoint. An alternate interpretation would allow for the densities in the mixture to overlap. In this case, the connection between the density mixture model and the Markov model is more complicated. Each state in the Markov model is actually characterized by a weighted mixture of densities where the weighting factor is a function of the fading threshold used to partition the Markov states. Moreover, the weighting factor for the mixtures in each state will typically vary from state to state. In this situation, the Markov model state can not be determined directly from the channel observations, since technically any value from the density mixture model is permitted in each state. In other words, the channel states are “hidden” from the observer. Markov models of this nature are therefore referred to as hidden Markov models (HMMs).

Gilbert’s Model

As described above, the distinguishing feature of HMMs is the decoupling of the model’s state (e.g., good or bad) from the observable event within that state (e.g., fade level above or below threshold). This distinction ultimately gives rise to a richer mathematical characterization of the channel. Among the simplest HMMs is the two-state model introduced by Gilbert [25]. In Gilbert’s application the good state was used to denote error free communications and the bad state represented communications with some nonzero probability of error. Note that with this model some “good” samples are still obtained in the bad state. Therefore, the state of the communications channel can not be determined directly from observations of the channel output. With respect to the LMS fading channel, this model can be applied with the good state representing the presence of a LOS signal component (i.e., unshadowed state) and the bad state representing the absence of this component. Although similar to the two-state Markov model discussed in the previous section, the Gilbert model more accurately reflects the situation at hand because it implicitly recognizes the variation in fade depth that is possible in the bad state. For example, suppose that a fading threshold, s_{th} , is of interest. In the bad state, as defined above, some fades will exceed this threshold and some will not. Hence, whereas with the conventional two-state Markov model the observation is synonymous with the channel state, Gilbert’s model decouples the observable event from the channel state.

Following the approach in [26], an diagonal event matrix $\mathbf{F}$ is used to quantify the probability of an event in a particular state:

$$\mathbf{F}_\gamma = \begin{bmatrix} 1 & 0 \\ 0 & h \end{bmatrix} \quad (27)$$

$$\mathbf{F}_\beta = \begin{bmatrix} 0 & 0 \\ 0 & 1 - h \end{bmatrix} \quad (28)$$

where γ represents a “good” event (i.e., received signal power above threshold), which occurs with probability one in the good state and probability h in the bad state. Likewise, β represents a “bad” event (i.e., received signal power below threshold), which never occurs in the good state, and occurs with probability $1 - h$ in the bad state.

The steady state probabilities for a good and bad event are then given by the following [26]:

$$\pi_\gamma = \pi \mathbf{M}_\gamma \mathbf{1} = \frac{p_{bg} + hp_{gb}}{p_{gb} + p_{bg}} = 1 - \pi_\beta \quad (29)$$

$$\pi_\beta = \pi \mathbf{M}_\beta \mathbf{1} = \frac{(1-h)p_{gb}}{p_{gb} + p_{bg}} = 1 - \pi_\gamma \quad (30)$$

where:

$$\mathbf{M}_\gamma = \mathbf{M} \mathbf{F}_\gamma \quad (31)$$

$$\mathbf{M}_\beta = \mathbf{M} \mathbf{F}_\beta \quad (32)$$

Dwell times for good and bad events are given by [26]:

$$\Pr \{D_\gamma = n\} = \pi \mathbf{M}_\beta \mathbf{M}_\gamma^n \mathbf{M}_\beta \mathbf{1} / \pi \mathbf{M}_\beta \mathbf{1} \quad (33)$$

$$\Pr \{D_\beta = n\} = \pi \mathbf{M}_\gamma \mathbf{M}_\beta^n \mathbf{M}_\gamma \mathbf{1} / \pi \mathbf{M}_\gamma \mathbf{1} \quad (34)$$

Equations (33) and (34) hold for arbitrary binary output HMMs. However, due to the simplicity of Gilbert’s model, these equations can be simplified further:

$$\Pr \{D_\gamma = n\} = A_1(1 - g_1)g_1^n + A_2(1 - g_2)g_2^n \quad (35)$$

$$\Pr \{D_\beta = n\} = (p_{bg} + hp_{bb})(1 - h)^n p_{bb}^n \quad (36)$$

where:

$$A_1 = \frac{(1-h)(p_{bg} - p_{gg} + p_{bb}g_1)}{(1-g_1)(g_1 - g_2)} \quad (37)$$

$$A_2 = 1 - A_1 \quad (38)$$

and $g_{1,2}$ are the eigenvalues of $\mathbf{M}_\gamma$, given by:

$$g_{1,2} = b \pm \sqrt{b^2 + h(p_{bg} - p_{gg})} \quad b = (p_{gg} + hp_{bb})/2 \quad (39)$$

The pmf in (35) is bi-geometric while (36) is geometric. The corresponding complementary cdfs are given by:

$$\Pr \{D_\gamma > n\} = A_1 g_1^n + A_2 g_2^n \quad (40)$$

$$\Pr \{D_\beta > n\} = (1-h)^n p_{bb}^n \quad (41)$$

As will be seen in the next section, the extra degree of freedom introduced by the bi-geometric distribution in Gilbert’s model yields superior performance relative to the conventional two-state Markov model.

In [28] Gilbert's model was extended further by allowing for the possibility of a "bad" event in the "good" state. In terms of the analysis above, this modification is accommodated by changing the element in the first row and first column of $\mathbf{F}_\gamma$ from one to some number, $0 < k < 1$, and the element in the first row and first column of $\mathbf{F}_\beta$ from zero to $(1 - k)$. Also, within the framework given above, extensions to more than two states is also straightforward. Of course, as with the density mixture models, more terms (states) may improve performance, but at the expense of the model's simplicity. In this spirit, both Elliott's model [28] and a three state HMM were examined only briefly and since their performance results relative to Gilbert's model were not significant enough to justify their complexity, they are not discussed further in this report.

Parameter Matching

One final issue of importance with discrete models is parameterizing the models based on recorded data. In the case of the two-state Markov model, there are at least two approaches. The brute force approach is to estimate the transition probabilities directly from the recorded data sequence. Another approach is to fit curves of the form given by (25) and (26) to complementary cdfs of the connection and fade dwell times computed from the experimental data. On a log scale (25) and (26) yield straight lines with slopes given by $\log p_{gg}$ and $\log p_{bb}$, respectively. In general, determining the optimal model parameters for a HMM based on observed data sequences is difficult. An iterative approach to the maximum likelihood solution is given by the Baum-Welch algorithm, described in [26]. Variations on this algorithm designed to satisfy alternate optimization criteria are also described in [26]. Luckily, Gilbert's model is simple enough that parameter matching can be accomplished with relative ease based on the raw data. This process is described in detail in [25]. Curve fitting as described above for the two-state model can also be used to parameterize Gilbert's model using the following relations [25]:

$$h = \frac{g_1 g_2}{g_1 - A(g_1 - g_2)} \quad (42)$$

$$p_{gb} = \frac{(1 - g_1)(1 - g_2)}{1 - h} \quad (43)$$

$$p_{bg} = A(g_1 - g_2) + (1 - g_1) \left(\frac{g_2 - h}{1 - h} \right) \quad (44)$$

where A , g_1 , and g_2 are derived from the process of fitting (40) to the complementary cdf of connection dwell times.

Although curve fitting is desirable for its simplicity, it must be done carefully if the generative nature of the model is to be preserved. One challenge associated with curve fitting is how to address the problem of a poor fit. In the case of the two-state model where a straight line is fitted to the data, this problem has been observed on numerous occasions. For example, in [13, 16] the reported dwell time distribution functions were comprised of at least two distinct components, a rapidly decaying component that represents short dwell times and a slowly decaying component that represents longer dwell times. The resultant curves resemble an "L" shape with an obtuse angle. In both papers the approach to fitting the straight line approximation was to ignore the rapidly decaying component and fit only to the slowly decaying component. Both papers cite the belief that the rapidly decaying component is due to multipath fading and the slowly decaying component is due to shadowing. Moreover, since shadowing was the phenomenon of interest, the authors felt justified in their approach. In any case, one problem with this selective curve fitting technique

is that the parameters produced by this approach do not yield the desired results when used in a generative two-state Markov model. In other words, since the approach taken in these papers effectively specifies a y-intercept for the complementary cdf that is not necessarily unity, a valid distribution function is not produced. Hence, use of these parameters in a generative two-state Markov model effectively normalizes the distributions which skews the probabilities and yields erroneous behavior in the model.⁵

An alternate approach to the selective curve fitting technique is to propose a different distribution function. For example, in [17] an exponential distribution is proposed for the dwell times. The lognormal distribution has also been observed to yield a good fit to measured fade duration data [29, 30] in LMS systems. In [16], the authors reconcile the exponential and lognormal descriptions of fade duration in LMS systems by suggesting that the lognormal distribution is a good model for fades of short to moderate length, while the exponential distribution is more appropriate for longer duration fades. Finally, it is interesting to note that in [29] a lognormal model and power-law representation are proposed for the fade and connection durations, respectively. The power-law representation takes the following form:

$$\Pr \{D_c > x\} = \alpha x^{-\beta} \quad (45)$$

where D_c is the connection duration, $\Pr \{D_c > x\}$ is the probability of the connection duration exceeding the abscissa value x , and α and β are constants. This model yields a straight line when plotted on a log-log scale. Of course, the problem with using distributions other than the geometric distribution is that they are not compatible with the two-state Markov model.

Parameterization in the Gilbert model can also occasionally be problematic. On one hand, the bi-geometric function given by the Gilbert model typically yields a good fit to the L-shaped distributions often encountered in LMS channels. However, as can be seen in (40) and (42) - (44), the Gilbert model is completely specified by a curve fit to the connection duration distribution. Unfortunately, in many cases, the model parameters so chosen do not yield a good fit to the fade duration distribution. This problem was noted in [31] where several different discrete Markov models were investigated for suitability in modeling the LMS channel at 20 GHz. The results of this investigation concluded that virtually none of the models are suitable in the generative sense because they fail to accurately reproduce multiple statistical properties of the data (e.g., fade and connection duration distributions) simultaneously. However, the curve fitting technique used in [31] (as originally recommended in [25]) where the model parameters are determined exclusively by a fit to the connection duration distribution is probably not the best approach if the end-goal is to produce parameters for a generative model. An alternate approach is to choose the model parameters based on a *joint* fit to the connection and fade duration distributions. This approach is explained in more detail in the next section where the discrete models discussed here are applied to the MOTM fade and connection duration distributions estimated from the field measurements.

⁵To avoid confusion, note that although the authors in [16] refer to their discrete model as the Gilbert model, they are in fact using the conventional two-state Markov model.

2.2.3 Fade and Connection Duration Distributions

Important second order statistics in the analysis of the LMS channel are the fade and connection durations. As seen in Figure 12, a fade duration is defined as the length of time that the received signal envelope is below a threshold, r_{th} . A connection is the length of time between successive fades (i.e., the length of time the received signal is above the threshold r_{th}). In the literature, it is common to see connections also referred to as nonfades or interfades. Note that fade and connection duration statistics are dependent not only on the fading environment, but also on the platform velocity. Fade duration information is an important parameter to consider in LMS system design since it provides information on outage periods and system availability. Other LMS system engineering issues influenced by fade duration statistics include frame length in packetized systems, choice of coding scheme, interleaver depth, buffer sizes, prediction algorithms and update schemes in adaptive transmission protocols, etc. Moreover, fade and connection duration information is a key parameter in LMS system performance analysis [17, 32–34].

Data collected with the MOTM system have been used to generate numerical complementary cdfs for fade and connection durations. Figures 16, 17, and 18 summarize the results for each of the three test environments. For convenience, fades and connections are actually defined relative to a power threshold, s_{th} , as opposed to the more traditional approach that uses an envelope threshold, r_{th} . Envelope fade durations can easily be obtained according to:

$$F_R(r_{th}) = F_S(r_{th}^2) \quad (46)$$


where $F_R(r_{th})$ is the envelope fade duration relative to an envelope threshold r_{th} , and $F_S(s_{th})$ is a power level fade duration relative to a power threshold s_{th} . For each of the three test environments two different power thresholds were used, $s_{th} = -3$ dB, and $s_{th} = -10$ dB. The first threshold, $s_{th} = -3$ dB was chosen to correspond to a situation in which a terminal is operating with only a small amount of link margin, as might be the case with future terminals that exploit rate adaptation schemes to continuously operate at the highest possible data rate. The second threshold, $s_{th} = -10$ dB, was intended to be representative of a more typical system where a moderate amount of link margin is budgeted to address propagation losses such as weather, etc. Duration results are presented in units of meters so that they can be interpreted independent of vehicle velocity. This process was accomplished by multiplying the duration values in samples by the factor $\bar{v}_i T_s$, where $\bar{v}_i$ is the average vehicle velocity during the i th fade as determined from velocity information recorded by the MOTM system's AHRS unit, and T_s is the sampling interval, which is constant. Converting the durations from meters to seconds or bits is accomplished by multiplying by the factor $1/\bar{v}_i$ or $R_d/\bar{v}_i$, respectively, where R_d is the data rate in bps. Data from 5-6 runs were used to generate each distribution, to which various discrete model fits were performed and superimposed on the plots. In each figure results for both the two-state Markov model and the Gilbert model are included. In all cases, the curve fits were performed by minimizing the Kullback-Leibler distance between the pmfs [35]:

$$D(u \parallel v) = \sum_x u(x) \log \frac{u(x)}{v(x)} \quad (47)$$


where $D(u \parallel v)$ is a measure of the distance between two probability mass functions $u(x)$ and $v(x)$. In the case of the Gilbert model, two fits are superimposed on each plot. One version of the Gilbert model is based on the parameters obtained from fitting a curve of the form given by (40) to the connection duration distribution. As mentioned in the previous section, one drawback to this approach is that there is no guarantee

that the model parameters obtained from fitting to the connection distribution curves will match the fade distribution curves. In fact, for the MOTM data a very poor match to the fade distribution curves is obtained. Hence, an alternate curve fitting process is proposed. With this new approach the Kullback-Leibler distance between (40) and the connection duration distribution is minimized jointly with the Kullback-Leibler distance between (41) and the fade duration distribution.


Several interesting observations can be made from the plots. First, note the difference in fade and connection distributions for a given environment as the threshold changes. In general, the fade distributions are steeper and the connection distributions are more gradual for $s_{th} = -10$ dB compared to $s_{th} = -3$ dB. However, these differences are less pronounced in the urban environment where virtually all fades are quite deep, and therefore less sensitive to thresholds in the range considered here. With respect to the discrete model approximations to the data it is clear that the Gilbert model offers by far the best approximation to the connection duration distributions. However, the resultant model parameters severely underestimate the fade durations. On the other hand, by jointly optimizing the Gilbert model parameters over both the connection and fade duration distributions the fidelity of the model with respect to the fade duration distributions is improved dramatically, albeit with a small loss in accuracy for the connection distributions. Finally, note that the straight line approximations derived from the two-state Markov model yield a relatively poor fit to the connection duration distributions for the open and rural environments. These distributions are characterized by numerous short duration connections and a relatively few long duration connections, which give rise to the “L-shaped” curves mentioned in the previous section. For example, in the rural environment there is a single connection that occurs roughly two-thirds into the course and lasts for approximately 1400 meters. On the other hand, fade distributions in all three environments are well approximated by the geometric distribution from the two-state Markov model. Finally, it is interesting to note that the two-state Markov model produces a reasonable fit to both the connection and fade duration distributions in the urban environment. This result is most likely due to the fact that fades in the urban environment tend to be “all or none,” producing a strong correlation between the channel state and a fading event, a property well suited to the two-state Markov model. The two-state Markov model parameters appear in Table 3, the Gilbert model parameters appear in Table 4, and the jointly optimized Gilbert model parameters appear in Table 5. Note from the tables that with the exception of the two-state Markov model in the urban environment, the parameter π_b (π_β) is consistently underestimated by the models. This result can be attributed to the fact that none of the models are able to match both fade and connection durations perfectly, perhaps suggesting that there is still room for improvement in the development of discrete models for the EHF SOTM channel.




(a)



(b)




(c)




(d)


Figure 16. Complementary cdfs and two-state models for fade and connection durations in the open environment. (a) Complementary cdf for connection duration with $s_{th} = -3$ dB. (b) Complementary cdf for fade duration with $s_{th} = -3$ dB. (c) Complementary cdf for connection duration with $s_{th} = -10$ dB. (d) Complementary cdf for fade duration with $s_{th} = -10$ dB. All figures contain discrete model fits to the data.




(a)



(b)




(c)




(d)


Figure 17. Complementary cdfs and two-state models for fade and connection durations in the rural environment. (a) Complementary cdf for connection duration with $s_{th} = -3$ dB. (b) Complementary cdf for fade duration with $s_{th} = -3$ dB. (c) Complementary cdf for connection duration with $s_{th} = -10$ dB. (d) Complementary cdf for fade duration with $s_{th} = -10$ dB. All figures contain discrete model fits to the data.




(a)



(b)



(c)



(d)

Figure 18. Complementary cdfs and two-state models for fade and connection durations in the urban environment. (a) Complementary cdf for connection duration with $s_{th} = -3$ dB. (b) Complementary cdf for fade duration with $s_{th} = -3$ dB. (c) Complementary cdf for connection duration with $s_{th} = -10$ dB. (d) Complementary cdf for fade duration with $s_{th} = -10$ dB. All figures contain discrete model fits to the data.

Table 3**Two-state Markov model parameters for the three test environments.**

Environment	Threshold (dB)	p_{gg}	p_{bb}	π_b
open	-3	0.9816	0.8767	0.13
	-10	0.9885	0.7794	0.05
rural	-3	0.9936	0.9542	0.12
	-10	0.9942	0.8901	0.05
urban	-3	0.9683	0.9473	0.38
	-10	0.9726	0.9465	0.34

Table 4**Gilbert model parameters for the three test environments.**

Environment	Threshold (dB)	p_{gg}	p_{bb}	h	π_β	A	g_1	g_2
open	-3	0.9888	0.9777	0.8903	0.04	0.1978	0.9907	0.8685
	-10	0.9949	0.9891	0.8833	0.04	0.0928	0.9953	0.8732
rural	-3	0.9986	0.9979	0.8809	0.05	0.0177	0.9986	0.8790
	-10	0.9991	0.9990	0.8799	0.06	0.0085	0.9991	0.8790
urban	-3	0.9795	0.9773	0.9036	0.05	0.2566	0.9837	0.8788
	-10	0.9748	0.9151	0.8148	0.04	0.3783	0.9821	0.7382

Table 5**Jointly optimized Gilbert model parameters for the three test environments.**

Environment	Threshold (dB)	p_{gg}	p_{bb}	h	π_β	A	g_1	g_2
open	-3	0.9913	0.8204	0.0010	0.05	0.1815	0.9913	0.0008
	-10	0.9943	0.8911	0.1276	0.04	0.1237	0.9944	0.1136
rural	-3	0.9977	0.9644	0.0128	0.06	0.0361	0.9977	0.0124
	-10	0.9974	0.9600	0.0808	0.06	0.0435	0.9974	0.0775
urban	-3	0.9938	0.9312	0.0004	0.08	0.0693	0.9938	0.0003
	-10	0.9959	0.9264	0.0003	0.05	0.0739	0.9959	0.0003

Another useful way in which to examine fade and connection duration information is in terms of the time share of fades (TSF) and time share of connections (TSC) [13, 19]. Given $p_c(\tau)$ and $p_f(\tau)$, the probability densities that a randomly chosen time instant lies in a fade or connection, respectively, of duration τ , the time share of fades and time share of connections are defined as:

$$TSF(D_f) = \int_{D_f}^{\infty} p_f(\tau) d\tau \quad (48)$$


$$TSC(D_c) = \int_{D_c}^{\infty} p_c(\tau) d\tau \quad (49)$$

It is illustrative to plot these quantities for a range of thresholds on the same figure, as is done for each of the three test environments in Figure 19. Whereas the fade and connection duration results presented in Figures 16 - 18 were based on distance, the results in 19 are in units of seconds. The time-based approach is used here to complement the previous distance-based duration statistics by providing insight into typical behavior of the vehicle, especially in urban environments. Also, as will be seen in Section 3, the TSC is a very useful metric for determining the performance of ARQ schemes in fading channels. Hence, expressing this information in units of seconds facilitates its use for this application.


The results in Figure 19 represent data averaged over 8 runs in each environment. Due to the wide range of values taken, the TSF and TSC curves are not plotted to the same scale for each channel. Nonetheless, careful examination of the plots is quite revealing. For example, note for the open environment that the probability of a randomly chosen time instant belonging to a connection of duration greater than 50 seconds is only about 10% when the fade threshold is -5 dB. On the other hand, increasing the fade margin by 5 dB (i.e., decreasing the threshold to -10 dB) increases this probability to the 30% level. In the rural environment there is a 30% probability of a randomly chosen time instant belonging to a fade of duration greater than 1 second when the threshold is -5 dB. However, a 20 dB increase in fade margin reduces this probability to almost 1%. There is very little separation among the TSC and TSF curves for the urban environment, reflecting the fact that most fades in this environment tend to be quite deep.

2.2.4 Average Fade and Connection Durations


Table 3 shows average fade and connection durations for two specific thresholds: -3 dB and -10 dB. In Figure 20, average fade and connection duration is presented as a function of threshold, s_{th} . As with the duration complementary cdfs in Figures 16 - 18, velocity information was used to yield average results in units of meters. The curves in the figure were computed numerically by measuring the duration of each fade and connection duration relative to a specified power level, s_{th} , and dividing the sum of these durations by the total number of fades and connections, respectively. Results were averaged over 5-6 runs. For the open and rural courses, average connection durations (ACDs) generally decrease and average fade durations (AFDs) generally increase with smaller threshold magnitudes. However, both channels exhibit a local minimum with respect to connection duration in the range of $s_{th} \approx -12$ to -15 dB, which is approximately the mean fade depth predicted by the least squares fit of (16) to the data, as summarized in Table 2. All three channels contain sharp drops in connection and fade durations at the -1.5 dB point. This characteristic is due to the variance of the received signal between -1.5 dB and $+1.5$ dB caused by the ASCAMP's hard-limiting front-end. The urban channel has relatively constant average fade and connection durations between approximately -25 dB and -2 dB. These data simply reinforce earlier observations that




(a)




(b)




(c)



(d)



(e)



(f)

Figure 19. Time share of connections and fades for the three test environments. (a) and (b) Time share of connections and fades, respectively, for the open channel. (c) and (d) Time share of connections and fades, respectively, for the rural channel. (e) and (f) Time share of connections and fades, respectively, for the urban channel.




Figure 20. Average connection and fade durations as a function of threshold. (a) Average connection duration. (b) Average fade duration.

fades in the urban environment tend to be “all or none.” Minima for both average connection and average fade durations in the urban channel exist around the -35 dB point, due probably to the proximity of the noise floor at this level. Another interesting characteristic of fade and connection durations in the urban channel is the presence of small “dips” that occur around -8 dB and -5 dB. These dips are caused by shadow “edges” that result in only a partial loss of signal power, due to refraction around and/or over buildings. Figure 21 illustrates the situation. This edge phenomenon gives rise to slightly shorter connection and fade durations around the mean signal, usually -8 dB to -5 dB.

As mentioned previously, the fade durations presented here seem to suggest that error control in the form of FEC coding and interleaving would lead to prohibitively long block lengths and/or interleavers. An alternate approach is to use a repeat request protocol to recover from channel fades. This approach was implemented as part of the MOTM system and performance results are discussed in Section 3.

2.2.5 Level Crossing Rates

The level crossing rate (LCR) is defined as the average number of times that the received signal envelope crosses a given threshold r_{th} , with a positive slope, per unit time (see Figure 12). The LCR is closely related to average fade duration. For this reason, the two quantities are often discussed jointly. Hence, the references cited above for fade duration statistics in LMS systems also include LCR information. Mathematically, LCR is defined as [8]:

$$LCR(r_{th}) = \frac{\Pr \{R \leq r_{th}\}}{AFD(r_{th})} = \frac{1}{AFD(r_{th}) + ACD(r_{th})} \quad (50)$$

where $AFD(r_{th})$ and $ACD(r_{th})$ are the average envelope fade and connection durations, respectively, relative to an envelope threshold r_{th} . Clearly, LCR is velocity dependent.




Figure 21. Depiction of shadow edge. Occasionally the vehicle was forced to move slowly through an area leading up to a blockage to avoid being forced to stop in a deep fade for prolonged periods of time.

LCR statistics as measured by the MOTM system in each of the three test environments are presented in Figure 22. As with the fade duration statistics, LCR is defined here relative to a power threshold s_{th} , instead of an envelope threshold. However, deriving the envelope statistics from these curves is straightforward, as shown in (46). The curves in Figure 22 were produced from the same data sets used to generate the ACD and AFD plots. Results are presented in units of inverse meters. Like the AFC and AFD plots, the fluctuations in received signal power between -1.5 dB and $+1.5$ dB caused by the ASCAMP's hard-limiting front-end result in an increased number of level crossings in this region. Ignoring this region, the peak LCR (due to fading) for the open and rural channels is seen to occur around -15 dB. In the urban environment the maximum occurs at -35 dB, due to the noise floor. Also, the small increases in LCR at -8 dB and -5 dB for the urban channel were expected given the discussion in the previous section regarding the presence of small decreases in the ACD and AFD data at these points.

Compared to published data [6, 8], the EHF SOTM channel LCRs are between one and two orders of magnitude smaller than those of typical UHF and L-band LMS and terrestrial wireless channels. Of course, the primary reason for this difference is that the UHF and L-band channels are dominated by multipath fading, an inherently different fading mechanism from the signal shadowing that dominates the EHF OTM channel.

2.3 BER RESULTS

BER experiments as described in Section 2.1.1 were conducted in each of the three test environments. In the results presented here, the MILSTAR FEC code was enabled, yielding a small probability of error ($< 1e - 5$) in the absence of fading. It is also possible that for sufficiently shallow fades the FEC coding




Figure 22. Received signal power level crossing rates for the three test environments.

was able to mitigate errors and maintain a low BER. On the other hand, the vast majority of signal blockage events overwhelm the FEC codec and burst errors result. Moreover, because the interleaving scheme used during the experiments was short relative to the duration of a most fades, the errors at the input to the convolutional decoder were typically bursty, yielding even worse BER performance at the decoder output. However, on the positive side, convolutional decoders are inherently self-synchronizing, so recovering from a blockage event occurs automatically and relatively quickly (within 10s of channel symbols). Nonetheless, it is important to remember when examining the results presented here that the FEC coding scheme used during the experiments was not optimized for use in blockage environments and may in fact contribute to a slight increase in the average BER. Experiments were also conducted without FEC coding, but these results are not presented here. Disabling the FEC codec results in a higher probability of error, even in the absence of signal blockage, making it difficult to distinguish errors due to fading from random errors.

Before discussing the experimental results, it is worth noting that given an expression for BER as a function of received power in the absence of fading, $\text{BER}(P_r/N_0)$, and given the pdf of the received signal power in the presence of fading, $p(S)$, the average BER for the fading channel is given by [36, 37]:

$$\text{BER}_{avg} = \int_0^\infty \text{BER}(S P_r/N_0) p(S) dS \quad (51)$$

where S is the instantaneous received signal power, normalized to the average received LOS power. This expression can be evaluated numerically for the MOTM system using the estimated pdfs for the test environments given in Figure 13 and downlink BER performance data measured in the absence of fading. Such reference data for the MOTM terminal's downlink BER performance were obtained with an ASCAMP block II terminal in a laboratory environment [38], where the power levels of the transmitting source could be easily controlled. The block II BER results are a useful reference since this terminal has essentially the same modem, hard-limiting front-end, and antenna configuration as the MOTM terminal. Figure 23 illustrates the





Figure 23. Theoretical average BER as a function of LOS received power level. These curves were calculated from (51), using numerical expressions for $BER(P_r/N_0)$ and the received signal power densities, $p(S)$, estimated for each of the three test environments.


results of a numerical calculation of (51) for each of the three test environments. The block II reference data are given in the same figure. Note the severe losses associated with the channel blockage.

As mentioned in Section 2.1, the BER tests do not lend themselves to the analysis of second order statistics such as burst error or error gap lengths. Hence, only average BER results will be discussed here. The experiments were structured so that both terminals were each sending and receiving BER patterns. With MILSTAR FEC coding enabled, it is reasonable to assume that the ASCAMP terminal's uplink and downlink are virtually error free. Hence, the BER results recorded at the ASCAMP terminal are a good indicator of the MOTM terminal's uplink performance. Likewise, the BER results recorded at the MOTM terminal are indicative of that terminal's downlink performance. Thus, for convenience the BER results measured by the ASCAMP terminal will be referred to as uplink results and the measurements collected at the MOTM terminal will be referred to as downlink results.


In Figure 24, "scatter plots" for each of the three environments are presented. In these plots, the error rate within each 20 second window of the BER test is given as a function of the average estimated received power within that same 20 second interval. Note that the "resolution" of the scatter plots is limited to approximately $2e-5$ due to the number of bits received in a 20 second time interval at a rate of 2400 bps. The ASCAMP block II terminal reference data discussed earlier are also included in the figure. Although comparisons between the scatter plots and the block II reference curve are insightful, they must



(a)



(b)



(c)

Figure 24. BER scatter plots for the three test environments. Each plot contains the 20 second average BER as a function of the 20 second average estimated P_r/N_0 . A reference curve measured with a block II ASCAMP in a controlled laboratory setting is also included. (a) Open channel. (b) Rural channel. (c) Urban channel.

be made carefully since there are important differences between the two. For example, as opposed to the reference curve, the scatter plot represents data collected over a nonstationary channel. Individual points are defined by 20 second average received power levels and 20 second average BERs. Another important difference between the reference curve and the scatter plots is that the reference curve is plotted against a known, calibrated received power level, whereas the scatter plots use the received power as estimated by the MOTM terminal. In comparing the scatter plots to the reference curve, note the significantly higher error rates on the scatter plot for the same received power. This behavior was expected given the degradations associated with blockage shown in Figure 23. Note that most data points on the scatter plot contain some fraction of blocked and unblocked states. The error rate in the blocked state is typically quite high, leading to relatively large average error rates, even for average received powers that don't seem very low. Another noteworthy characteristic of the scatter plots is the "brick wall" or "compression" effect where many of the data points fall on a nearly vertical line around 75 dB. This characteristic is due to the saturation effects of the MOTM terminal's hard-limiting front-end that effectively limits the terminal's ability to estimate received power levels beyond this range. Other techniques for extending the range of the P_r/N_0 estimator have been proposed [4]. However, these techniques involve a great deal of averaging which limits their ability to resolve short blockage events. Hence, they are not considered here.

Figures 25, 26, and 27 give additional BER results for the open, rural, and urban channels, respectively.

In these figures the number of errors counted within a 20 second window is plotted versus time for both the uplink (i.e., errors counted by the stationary ASCAMP terminal) and the downlink (i.e., errors counted by the MOTM terminal). Each error pattern is superimposed on a plot of received signal power as measured at the MOTM terminal during each of the experiments. The correlation coefficients between the uplink and downlink error patterns have also been estimated for each of the environments. As expected, the error patterns are highly correlated, suggesting that signal blockage has a similar effect with respect to the errors it introduces on both the transmitted and received signals in the MOTM terminal. Although the strong correlation between the uplink and downlink error patterns makes intuitive sense, it is harder to determine whether or not either of the links might be inherently more robust to the effects of signal blockage. On the one hand, due to the higher operating frequency, the MOTM terminal's uplink beamwidth is narrower than its downlink beamwidth. This fact might suggest that the uplink is more sensitive to blockage than the downlink. On the other hand, the uplink uses more robust signaling techniques and has a great deal more margin than the downlink, suggesting that it might be less sensitive to blockage than the downlink. Unfortunately, the results of the BER experiments are inconclusive. Figures 25c, 26c, and 27c show the cumulative errors measured by the two terminals during the BER test in each of the three test environments. In the open environment, many more errors are made on the uplink, suggesting that this link is more sensitive to signal blockage. However, in the rural environment, slightly more errors are made on the downlink compared to the uplink. Finally, in the urban environment virtually the same number of errors are made on each of the links. Hence, given the limitations of the BER tests conducted, perhaps the only result that has been established is the rather obvious fact that with respect to BER, both the MOTM terminal's uplink and downlink are highly correlated with one another and with the fading channel.




Figure 25. Uplink and downlink error results for the open channel. (a) Errors and estimated P_r/N_0 at the MOTM terminal vs. time. (b) Errors at the stationary ASCAMP terminal and estimated P_r/N_0 at the MOTM terminal vs. time. (c) Cumulative errors on the UL and DL over the duration of the experiment.




Figure 26. Uplink and downlink error results for the rural channel. (a) Errors and estimated P_r/N_0 at the MOTM terminal vs. time. (b) Errors at the stationary ASCAMP terminal and estimated P_r/N_0 at the MOTM terminal vs. time. (c) Cumulative errors on the UL and DL over the duration of the experiment.




Figure 27. Uplink and downlink error results for the urban channel. (a) Errors and estimated P_r/N_0 at the MOTM terminal vs. time. (b) Errors at the stationary ASCAMP terminal and estimated P_r/N_0 at the MOTM terminal vs. time. (c) Cumulative errors on the UL and DL over the duration of the experiment.


2.4 SUMMARY

Channel measurements from three test environments, referred to as open, rural, and urban, were conducted with the MOTM system. In all cases the dominant propagation phenomenon observed at the receiver was signal fading caused by objects in the propagation path. In the case of the open and rural channels, foliage was responsible for the fading, while in the urban channel the blockage was caused primarily by tall buildings. For fade depths on the order of 10 dB, the open channel was characterized by a blockage fraction of approximately 15% while the rural and urban channels yielded blockage fractions of approximately 30%. Despite the similarity in blockage fraction between the rural and urban channels, the fading behavior of the three channels is quite different. A statistical analysis revealed that fade depth in the open and rural channels was accurately modeled with a lognormal distribution while a combination of lognormal and a power law yielded good results for the urban channel. Fade and connection duration statistics were less sensitive to threshold selection in the urban channel since most fades in this environment tend to be quite deep. On the other hand, these distributions are more likely to change as a function of threshold in the open and rural environments. While the two-state Markov model commonly cited in the open literature produced reasonable results for the urban channel, this model proved to be inadequate for the open and rural channels. However, a simple extension to this model, proposed originally by Gilbert [25] improved the performance significantly. Finally, BER results essentially confirmed the high degree of correlation that one would expect between errors introduced by the fading process on the MOTM UL and DL.

3. PROTOCOL PERFORMANCE

The purpose of this section of the report is to characterize the performance of the MOTM experimental protocol (MEP), used in the system primarily for error control. Given the statistical properties of the EHF LMS channel discussed in Section 2, several possibilities for error control exist. In channels that contain only a modest amount of blockage, conventional FEC coding with interleaving may be used. However, as the blockage fraction increases the performance of this strategy degrades rapidly. Since interleaving requirements increase with average fade duration [39,40], even moderate length blockage events will have a substantial impact on end-to-end latency. While error control in the form of FEC coding and interleaving has been examined within the context of LMS systems [40–42], these investigations were focused on LMS channels dominated by Ricean fading, as opposed to signal blockage. Another strategy with FEC coding is to apply very long block codes (without interleaving) to groups of packets in such a way that corrupted or missing packets can be replaced by the code [43–45]. These schemes are usually proposed for high data rate broadcast or reliable multicast applications where large data transfers are required without the assistance of a return channel. While FEC coding is essentially the only choice for error control in systems that lack a return channel, ARQ strategies represent a practical and efficient choice for full duplex systems. There are three main classes of ARQ [46]: stop and wait, go-back- N , and selective repeat. Of these, selective repeat is the most efficient. In their simplest form, ARQ schemes make use of error detection coding to determine whether or not a packet should be accepted by the receiver. In the event that no errors are detected, an acknowledgment (ACK) is generated by the receiver and sent via the return channel to the transmitter, otherwise a negative ACK, or repeat request, is initiated and the packet is resent. Because of their support for the acknowledgment of data and their ability to deliver very small probabilities of undetected error, ARQ schemes are generally considered far more reliable than FEC coding. However, FEC coding and ARQ can be combined in the form of hybrid ARQ (HARQ) schemes [44, 46] to achieve very powerful and efficient error control. HARQ schemes fall into two broad classes. The first category, referred to as Type I schemes, apply FEC coding to each transmitted packet to reduce the probability of a packet error when the channel error rate is high. However, the overhead associated with the FEC coding is not necessary when the error rates are low, and hence represents an inefficiency. Type II HARQ schemes address this issue through the use of carefully constructed codes that allow for the transmission of incremental redundancy (i.e., parity) in response to each repeat request. The receiver can then combine the multiple retransmissions to obtain more powerful (i.e., lower rate) codes. The net result is that the code rate is effectively adapted based on the channel conditions.

HARQ schemes are a logical choice for error control in EHF LMS systems because of their ability to effectively mitigate random errors during connections with FEC coding and also ensure end-to-end reliability in the presence of long duration fades with ARQ. The MEP uses a Type I HARQ scheme for error control where the ARQ portion of the protocol is based on a selective repeat strategy. The protocol can be used either with or without the rate $R_c = 1/2$ MILSTAR FEC code. When in use, the MILSTAR code represents the FEC portion of the HARQ protocol and Reed-Solomon (RS) block codes are used for error detection in the ARQ portion of the protocol. Without MILSTAR coding, variable rate RS block codes are used for error correction and detection. More detail on the protocol is given in Section 3.1. The performance metrics investigated here are packet error rate (PER), throughput efficiency, and packet latency. Analytical



expressions for these quantities are presented and the experimental results show good agreement with the predicted performance.

3.1 PROTOCOL OVERVIEW

Before a detailed overview of the protocol is given, several important points must be made. As its name implies, the MEP was designed and implemented as an experimental protocol, rather than one to be deployed in a fielded system. Hence, in the time allocated for protocol design and implementation, the most important issues to be addressed were ones that would facilitate the use of the protocol as a learning tool. These issues included instrumenting the protocol for support of offline performance analysis and providing a high degree of configurability. On the other hand, issues such as optimizing the protocol for execution time, interfacing it to other standard protocols, etc., were considered less critical. Moreover, the multiple access problem, while very important, is beyond the scope of the current investigation. Hence, the MEP was designed for use over point-to-point links only.

A block diagram of the MEP is given in Figure 28. In this initial realization, the protocol runs in a stand-alone configuration. In other words, the MEP has been implemented as an application, rather than a

link layer protocol embedded in the kernel of a host computer's operating system. At the expense of interoperability with commercial applications written to a standard TCP/IP interface, this approach streamlined the development process, including testing and evaluation, significantly. Moreover, a TCP/IP interface to the protocol was not considered critical because many of the research issues associated with the performance of TCP/IP over satellite links are not suitable for investigation with low data rate channels, like those in the MOTM system. However, in the absence of a standard interface, several custom applications, or applets, had to be developed for use with the protocol to support for a very basic demonstration capability. Referring to the figure, data streams from these applications are passed to a packetization and queueing mechanism where they are packetized and placed on an outgoing queue. In addition to the outgoing queue, a retransmission queue is required to store a copy of packets that have already been sent. If an acknowledgment is received within the expected time-frame, the packet is purged from the retransmission queue. Otherwise, the priority value is increased and the packet is re-inserted into the outgoing queue. In situations where the MEP applies FEC coding (as opposed to, or in addition to the $R_c = 1/2$ MILSTAR code), Reed-Solomon (RS) block codes are used. The same RS codec is used for error detection in the protocol. Finally, the detection of asynchronously generated packets is accomplished through a process of framing. Each of these functions are explained in more detail in the following sections. An overview of the mechanism by which the various protocol parameters, including code rates, packet lengths, retransmission time, etc., are managed is also presented. The overhead associated with the protocol consists of only 9 bytes: 3 for the packet header, 4 for framing and a minimum of 2 bytes for error detection. Of course, the low data rate channels and point-to-point applications for which the protocol was designed were significant factors in achieving this low overhead.

3.1.1 QoS Management

Numerous protocol parameters can be configured within the MEP on an application-by-application basis. These parameters include:

- transmission priority
- FEC code rate
- packet length
- maximum number of transmission attempts
- wait-time between transmission attempts

and are collectively referred to as quality of service (QoS) parameters. A maximum of 8 preset combinations of QoS parameters for each application are stored in an initialization script, referred to as the attributes table, that is loaded by the protocol at startup. Once the protocol is running, users are restricted to one of the 8 choices contained in the attributes table. Of course, the attributes table may be configured offline to contain arbitrary values for the various parameters. Obviously, both ends of the link must be using the same QoS parameters for the MEP to function properly. Moreover, the protocol does not support adaptive QoS where parameters are changed dynamically on a packet by packet basis for a given application. Instead, QoS parameters are chosen by the user and must remain fixed while the application is active. Note, however, that QoS parameters can be changed without restarting the MEP. The QoS management interface to the MEP is depicted in Figure 29.

pull-down box for application selection

8 preset choices for QoS parameters

54

3.1.2 User Applications

A small collection of user-applications were written to demonstrate the utility and assess the performance of the protocol. These applications include a packetized voice application, file transfer, instant messaging, and channel testing application. Each of these applications are discussed briefly below.

Packetized Voice

The packetized voice application was built around the 1200/2400 bps MELP vocoder [47, 48]. Most often the 1200 bps configuration is used since this mode accommodates the overhead associated with packet headers when the MOTM terminal is running over a 2400 bps MILSTAR link. In the application, multiple speech frames produced by the vocoder are collected into a packet for processing via the MEP. Although it is possible to select QoS parameters for the packetized voice application that effectively disable the ARQ portion of the error control protocol, experience has shown that voice communications in heavily blocked environments is virtually impossible in this configuration. On the other hand, ARQ ensures that voice packets are eventually received in the presence of signal blockage. However, the delays associated with packet repetition severely degrade the subjective quality of the voice conversation. In order to address this problem several features have been built into the application. Referring to Figure 30, these features are described as follows. First, to address the increased likelihood of double-talk caused by the relatively long propagation delay and delays due to retransmissions, the packetized voice application defaults to a push to talk (PTT) discipline. Audio-visual indicators are used to enforce this discipline. For example, when a message is being received at the MEP, an incoming message indicator is displayed signifying that the receiving user should wait until the message has been received in its entirety before attempting to talk. In addition, a double-beep is played at the beginning of each received message and a single beep is played at the end of each received message. On the transmit side, a user first activates the PTT mechanism and waits for a “chirp” sound, which indicates that the link is available for transmission. When the PTT is released and the final packet has been acknowledged by the receiver, a single beep is played to indicate that the link is available for another transmission. Of course, because the vocoder and underlying satellite link are both full duplex, the PTT discipline is not required. In fact, the PTT mechanisms described above can be easily disabled from the QoS interface for the voice application. However, in blockage environments, experience has shown these features to be quite helpful in conducting an ongoing conversation. Next, the voice application contains an archive and replay mechanism so that received messages can be stored and played back with the pauses caused by retransmissions removed. Several escape mechanisms are also provided whereby incoming and outgoing voice messages can be aborted, allowing the users to communicate via alternate means, such as the instant messaging application. On the receive side, a simple “abort communications” button is provided that effectively resets the incoming voice buffers. This same capability is provided on the transmit side. In addition, the status of the outgoing transmission buffer is also displayed so that users are able to see in real time whether or not their outgoing voice packets are being acknowledged by the receiver. In situations where this buffer approaches its limit, users know that buffer overflow is imminent and should stop transmission.

Beyond standard acceptance tests that verify its operation, no formal experiments have been conducted or planned for the packetized voice application. Of course, packetized speech over satellite systems is not a novel concept, and numerous issues associated with this practice are discussed in [49]. However, the blockage channel does complicate matters significantly. So much so in fact that users may ultimately decide




Figure 30. User interface to packetized voice application.




Figure 31. Histogram of speech transmission times in a tactical scenario. These data were collected during training exercises at the NTC.

that voice communications over this class of channel is simply not practical. On the other hand, experience with the application indicates that while the pauses due to retransmissions do hinder real time voice communications, these limitations are not insurmountable. In fact, with a little practice users are able to communicate quite effectively, even in the presence of significant channel blockage. Also, it is important to note that military voice users communicate in a way that is altogether different from the commercial world. For example, tactical users are trained to minimize transmissions to avoid detection. Consequently, they use an extensive collection of special words and phrases designed to maximize information content with minimal transmission time. Hence, typical transmission times in tactical scenarios are quite short. Figure 31 reinforces this point. The figure is a histogram of utterance lengths collected from a mock tank battle conducted at the National Training Center (NTC) in 1999 [50]. The histogram is comprised of over 650 speech transmissions and yields an average duration of slightly less than 3 seconds. In addition, military voice users do not necessarily conduct real-time conversations. Instead, voice transmissions are often initiated to request further instructions or update status information, perhaps with no expectation of a response. These observations suggest that military voice users may be willing to accept the performance of an ARQ-based voice system in exchange for the required reliability mechanisms and the potential anti-jamming benefits associated with EHF SOTM systems.

File Transfer

The file transfer application is essentially the same as FTP, allowing users to select via a graphical interface various files from their host machine for transfer over the system. A graphical display for monitoring the upload/download status in terms of a completion percentage is also provided.

Instant Messaging

The instant messaging application allows users to exchange real-time or near real-time text messages with one another. This capability provides a convenient alternative to voice communications, especially in heavily blocked environments where conducting a real-time voice conversation is difficult. As with the voice application described above, the instant messaging application supports the archiving of messages so that this material may be referenced offline.

Channel Tester

The channel testing application is not so much a user application as it is a measurement tool provided for experimental purposes. The application can be used to generate on a regular basis packets that contain pseudo-random bit sequences and/or local time stamps. These packets are then evaluated at the other end of the link and information such as PER and packet latency are displayed to the end-users and logged for offline analysis. Figure 32 illustrates the interface to the channel test application. Note also in the figure that the graphical user interface (GUI) to the MEP contains a system status window and channel state indicator. The channel state indicator averages energy estimates produced by the MOTM terminal over approximately three seconds and uses this information to control a “bar meter” that gives the operator an idea as to the amount of near term channel blockage experienced by the terminal. As will be described in Section 3.2, the channel test application was used to conduct experiments to measure PER, throughput efficiency, average packet latency, and packet jitter in each of the three test environments. This portion of the MEP GUI is always visible regardless of which application is displayed in the main window.

3.1.3 Packetization and Queueing

Data from the applications are passed to a packetization and flow control mechanism where they are packetized and placed on an outgoing queue, as shown in Figure 28. Due to the relatively low data rates for which the MEP was designed to operate, the maximum packet length supported is only 256 bytes. Packet headers consist of an application ID, a QoS ID, and a sequence number, as depicted by Figure 33. The packets are assigned a priority value based on their source application. The priority value determines the packets' relative location on the outgoing queue [51,52]. In addition to the outgoing queue, a retransmission queue is required to store a copy of packets that have already been sent. If an acknowledgment is received within the expected time-frame, the packet is purged from the retransmission queue. Otherwise, the priority value is increased and the packet is re-inserted into the outgoing queue.

Counterparts to these functions on the receive side of the MEP include a packet demultiplexer and an ACK generation mechanism. The demultiplexer serves two main purposes. First, based on information contained in the packet header, it routes the data in the packet body to the appropriate application. Also, because the MEP is a self contained protocol it is responsible for delivering the data to the end-user applications in the proper order. Hence, the demultiplexer buffers received packets and only releases data to the applications in sequential order, as determined by the sequence numbers in the packet headers. For every packet received at the demultiplexer an ACK is generated and placed in the outgoing queue. Since the absence of an acknowledgment triggers a repetition, acknowledgment packets always receive the highest priority in the MEP. Otherwise, unnecessary retransmissions would occur, degrading the efficiency of the protocol at best,




Figure 32. User interface to channel testing application.




Figure 33. MEP packet header structure.

and leading to instabilities at worst. Acknowledgment packets have a very simple format consisting of only the packet header for the packet to be acknowledged with the first bit flipped from 0 to 1. Although it is less efficient, for simplicity the MEP was designed to acknowledge each received packet individually. A more resource efficient scheme that makes use of cumulative acknowledgments, or sliding window mechanisms would probably be more appropriate for a fielded system.

3.1.4 FEC codec

The FEC codec is responsible for encoding outgoing packets on the uplink and decoding packets received on the downlink. Note that this codec may be used either in addition to or instead of the $R_c = 1/2$ convolutional code that runs over the MILSTAR physical layer. The MEP codec employs Reed-Solomon (RS) block codes with configurable block lengths, n , code rates, $R_c = k/n$, where k is the number of information symbols to be encoded, and symbol sizes, m . These parameters are managed on an application-by-application basis via the attributes table described in Section 3.1.1. Figure 34 illustrates the FEC codec functionality. To facilitate decoding with the appropriate parameters, all packet headers are encoded with a (12,3) shortened RS code, with code symbols taken from $GF(2^8)$. At the decoder, the packet header is decoded first and the appropriate parameters as determined from the header information are used to decode the body. The penalty for independent coding of the packet header and body, as opposed to coding them together, or jointly as a single codeword, is a slight loss in error correction capability. However, the loss in performance that results from independently coding packet headers and bodies is negligible [53] and more than offset by the throughput gains that are achieved by using this approach to customize the FEC coding according to the needs of the application and channel conditions. In situations where the MILSTAR FEC code is used, the RS codec is still employed for error detection, an integral part of the repeat request mechanism in the MEP. In these cases a single byte is typically used for error detection in the packet header and another byte is typically used for error detection in the packet body.

3.1.5 Framing and Deframing

The detection of asynchronously generated packets is accomplished through a process of framing. The framing component places flags at the packet boundaries before transmission. These flags are detected at the receiver and used to identify packets, which can then be passed along for further processing. In the MEP, a framing process similar to the one employed by the well-known HDLC protocol is used [54]. This approach to framing consists of two basic steps: bracketing the packet with flags (i.e., special bit sequences known at both the transmitter and receiver), and disrupting natural occurrences of the framing flag that may exist within the packet through a process of bit-stuffing. Framing flags and stuff-bits represent overhead




Figure 34. FEC codec functionality. (a) Based on the application and QoS ID the appropriate FEC encoding parameters are retrieved from the attributes table and applied to the packet body. All headers are encoded separately with the same parameters. (b) The FEC decoder subfunction first decodes the packet header. Based on the application and QoS ID contained in the header, the appropriate decoding parameters are retrieved from the attributes table and used to recover the body.

that reduces the throughput efficiency of the protocol. Note that framing flags and stuff-bits are typically not protected by the MEP's FEC codec since framing occurs after FEC encoding. However, these bits are protected by the MILSTAR codec, which effectively operates within an end-to-end MEP connection over MILSTAR. At the deframer, failure to detect the framing flags and remove the stuff-bits results in errors that trigger a repeat request, adversely affecting throughput efficiency. Hence, on one hand short flags are desirable because they represent less overhead and are less likely to be corrupted by errors. On the other hand, short flags give rise to more stuff-bits and an increased probability that one or more will fail to be detected and removed by the deframer, leading to a corrupted packet. Within the context of the MEP, this tradeoff was examined in [53], where a two-byte framing flag was chosen because of the good performance obtained over a wide range of expected operating conditions. Moreover, working with an integer number of bytes offers implementation advantages.

3.1.6 Rate Control

The rate control mechanism, not pictured in Figure 28, ensures that packets are not produced at a rate higher than the physical channel can support. This component works by simply pulling packets off of an outgoing buffer at a specified rate. In the event that the outgoing buffer is full, individual user-applications are responsible for adjusting their output accordingly.

3.2 EXPERIMENT OVERVIEW

MEP protocol experiments were conducted in each of the three test environments described in Section 2.1.2. The experimental setup was exactly as depicted in Figure 5. The experiments were designed to measure PER (and by extension throughput efficiency), as well as packet latency statistics. Moreover, all experiments were conducted over both the forward and reverse paths simultaneously using the channel test application described in Section 3.1.2. The PER experiments consisted of generating at a constant rate packets that contained pseudo-random bit sequences and logging via the MEP at both the base (i.e., the stationary ASCAMP) and mobile (i.e., the MOTM terminal) statistics that included:

- the total number of packets transmitted
- the total number of packets received
- the total number of acknowledgments generated
- the total number of acknowledgments received
- the total number of retransmitted packets

This information was then processed offline to determine average PER and throughput efficiency, as will be discussed in Section 3.3.1. These experiments were repeated for several different packet lengths. Latency experiments were executed in a similar manner. However, the packets used in the latency experiments contained a timestamp based on the current value of the local clock at the host computer where the packet was generated. This timestamp was compared to the clock value of the host computer at the receiver and the difference was logged as the end-to-end latency for that packet. In addition to packet latency, the parameters listed above were also logged during the latency experiments. Prior to each latency experiment, the host computer clocks were synchronized manually by system operators who communicated via cell-phone. This

process typically resulted in clock discrepancies on the order of a few hundred milliseconds, which was sufficiently close for monitoring the experiment as it was being conducted and ensuring that it was progressing as expected. Offline, the offset between the clocks was determined exactly by comparing each of the host computer clock values to MILSTAR time as recorded at each of the terminals. This offset was then applied to the recorded latency values for improved accuracy. Latency experiments were conducted for two different packet lengths and for two different values of waiting time between retransmission attempts. Also while experiments without MILSTAR coding were conducted with the MEP, only results from the experiments where MILSTAR coding was enabled are presented here. The original intent of the experiments without MILSTAR coding was to demonstrate the improved throughput efficiencies that could be achieved in this case. Unfortunately, the existence of an error floor in the ASCAMP [38] and MOTM terminals limited the protocol's performance in the absence of stream-based FEC coding, such as provided by the MILSTAR waveform. Hence, these results will not be discussed here. Finally, note that for the experiments described here overhead associated with maintaining synchronization between the source and KIV-7 clocks effectively reduced the channel data rate by a factor of 8/11 (i.e., one start and two stop bits had to be applied to every data byte processed by the KIV-7). This overhead constrained the data rate over the MILSTAR channel to a maximum of approximately 1745 bps, or about 218 bytes per second.

3.3 PERFORMANCE EVALUATION

In this section several important figures of merit are discussed within the context of experimental results obtained via the MEP. These include average PER, throughput efficiency, and latency statistics. Recall that experiments were conducted within the framework depicted in Figure 5, where one terminal was mobile and one was stationary. Hence, the experimental results and supporting analysis that are presented in the sequel should be interpreted in this context. However, extending this work to evaluate the performance of a system where multiple terminals are on the move and subjected to signal fading is relatively straightforward.

3.3.1 Average Packet Error Rate and Throughput Efficiency

Because the MEP relies on FEC coding for the correction of random errors due to channel noise as opposed to fading, blockages will typically introduce more errors than can be corrected by the FEC decoder during a packet. Hence, the PER during a fade is assumed to be approximately 1. Moreover, assuming the PER during a connection to be approximately 0, a reasonable approximation to the average PER is simply the blockage fraction, B . However, this approximation will generally be optimistic since it fails to account for the interactions between fade durations and packet length, the effects of blockage on the return channel, etc. In the MEP, average PER and packet success rate (PSR) are measured directly as:

$$P_f = \frac{N_r}{N_g + N_r} = 1 - P_s \quad (52)$$

$$P_s = \frac{N_g}{N_g + N_r} = 1 - P_f \quad (53)$$

where P_f is the probability of a packet error, or failed transmission attempt, P_s is the probability of a successful transmission attempt, N_g is the number of new or original packets generated by the test application over the course of an experiment, and N_r is the total number of packets retransmitted by the protocol during




Figure 35. Average packet error and success rates in the three test environments. (a) Average packet error rate as a function of unblocked fraction. (b) Probability of a successful packet transmission as a function of blockage fraction.

the experiment. Figure 35a is a scatter plot of average PER for the various experiments as a function of the fractional time spent in a connection, $(1 - B)$. For each experiment, the parameter $(1 - B)$ was computed directly from the data based on the estimated received signal power relative to a threshold, s_{th} . The threshold used in Figure 35 is $s_{th} = -5$ dB, a value that was observed empirically to yield $P_f \approx 1$ during a fade in most of the data sets. Note, however, that one experiment produced a PER less than B . This may be explained by noting that the appropriate threshold will vary somewhat from experiment to experiment depending on external operating conditions, such as weather, location of the satellite beam, etc. The average PSR is plotted in Figure 35b as a function of blockage fraction B .

The PSR is an important parameter in the determination of throughput efficiency, which for a Type I HARQ protocol like the MEP is given by:

$$\eta = \frac{L_i}{L_t} P_s \quad (54)$$

where L_i is the number of information bits contained in a packet, L_t is the total number of bits in the packet, and P_s is the probability of a successful packet transmission. The fraction L_i/L_t will be referred to as the *overhead efficiency*. In the MEP, $L_t = L_i + L_h + L_p + L_f$, where L_h is the length of the packet header, in bits, L_p is the number of parity bits contained in the packet, and L_f is the number of bits in the packet associated with framing. In general, the PSR is influenced not only by channel conditions like random BER and fading, but also by protocol parameters like FEC coding, framing, etc. Typically, a tradeoff will exist where a certain amount of overhead introduced by the protocol (i.e., a decrease in the overhead efficiency) can be used to compensate for poor channel conditions and yield a higher probability of successful packet transmission. FEC coding is an excellent example of how protocol overhead (i.e., parity in this case), if used


correctly, can significantly improve throughput efficiency. Clearly, coding parity decreases the overhead efficiency. The question is under what circumstances is this reduction more than offset by a gain in the probability of a successful transmission. Optimizations of this nature, conducted within the framework of (54), have been examined in [53, 55].

With the aforementioned assumption that the $R_c = 1/2$ MILSTAR code has reduced the channel BER to negligible levels, an expression for PSR that accounts for packet length and the actual distribution of connection durations is given by:


$$P_s = \int_{L_t}^{\infty} \frac{\tau - L_t}{\tau} p_c(\tau) d\tau \quad (55)$$

where L_t is the packet length (expressed now in units of seconds), and $p_c(\tau)$ is the probability density that a randomly chosen time instant lies in a connection of duration τ . A comparison to (49) reveals (55) to be the time share of connections with duration greater than or equal to L_t , conditioned on the probability that the entire packet is contained in the connection (i.e., no overlap with a blockage event is allowed). Note that (55) reduces to B as $L_t \rightarrow 0^+$. Combining (55) with (54) yields an expression for throughput efficiency as a function of packet length. Note that the analysis thus far has implicitly assumed an error free return channel for the transmission of acknowledgments. However, in practice this is typically not the case. One simple way to account for fading on the return channel is to require that a connection be long enough to support the receipt of a packet *and* the generation and transmission of an ACK. This is accomplished by replacing L_t in (55) with $L_t + \delta_a$, where δ_a is the time required to generate and transmit an ACK. This approach neglects the possibility that there could be a very short fade between the receipt of a packet and the transmission of the ACK. However, given the highly correlated nature of the fading process in the channels considered here, this probability should be quite small.


Experiments designed to measure throughput efficiency as a function of packet length were conducted with the MEP and summarized in Figure 36 along with the theoretical expression described above. In the figures, experimental results for throughput efficiency were obtained by using (53) and (54). For convenience, the overhead associated with the MILSTAR code is ignored. Strictly speaking, this overhead would scale the throughput efficiencies in the figure by 0.5. Also, from a throughput efficiency perspective, the overhead associated with the start and stop bits required for synchronization with the KIV-7 encryption devices is also ignored. However, this overhead is reflected in the conversion of packet lengths to seconds via L_t/C_{eff} , where L_t is the packet length in bytes, and $C_{eff} = 218$ bytes per second is the effective channel rate, constrained by the start and stop bits associated with maintaining synchronization with the KIV-7, as described in Section 3.2. Packet lengths are expressed in units of seconds to facilitate comparisons with the theoretical expression that makes use of the TSC information for the channel, which is typically expressed in units of seconds. All other overhead associated with the protocol, including the $L_h = 3$ byte packet header, $L_f = 4$ bytes for framing, and $L_p = 2$ bytes for error detection, is included in the overhead efficiency calculation. Four different packet lengths were used in the experiments: $L_i = 32$ bytes, $L_i = 64$ bytes, $L_i = 128$ bytes, and $L_i = 232$ bytes, yielding overhead efficiencies of approximately 78%, 88%, 93%, and 96%, respectively. The theoretical curves in the figure were obtained by computing the quantity $p_c(\tau)$ in (55) for each experiment with a threshold of $s_{th} = -5$ dB and averaging these results to obtain a single density for each environment. This density was then used in (55), with L_t augmented by $\delta_a = 150$ ms to



(a)



(b)



(c)

Figure 36. Throughput efficiency as a function of packet length. (a) Open environment. (b) Rural environment. (c) Urban environment.

account for blockage on the return channel, to obtain the predicted PSR which was then be used to obtain throughput efficiency according to (54).

The figures reveal a tradeoff where increasing packet length initially yields improved throughput efficiency, due to the gain in overhead efficiency. However, as packet length increases beyond a certain point, the likelihood that these longer packets are corrupted by signal blockage increases, outweighing the improvements in overhead efficiency and resulting in reduced throughput efficiency. In general, the measured data points follow the same trend predicted by the analysis, but some of the data points do not line up exactly with the theoretical curves. One possible reason for the discrepancies is statistical variation. Only one experiment was conducted for each packet length resulting in two data points, one for the forward path and one for the reverse path. Moreover, only a single theoretical curve is generated for each environment, based on the average statistics from the four experiments in that environment. Hence, variation in the statistics due to things like vehicle speed will give rise to differences between the measured and predicted values. In any case, the optimum packet length for each of the environments appears to be in the range of 300 ms to 400 ms. Packet lengths shorter than this suffer from poor overhead efficiency while packets longer than this are more likely to be corrupted by channel fading. Finally, note that the optimum packet length clearly depends on the channel rate. In general, the optimum packet length, in seconds, will decrease with increasing channel rate. For a fixed packet duration (in seconds), higher channel rates imply more bits per packet and hence improved overhead efficiency. For lower channel rates, high overhead efficiencies can only be obtained with longer packets. However, longer packets are more likely to suffer the effects of fading. In other words, higher channel rates allow for packets to be simultaneously “short” (in seconds), which is desirable in the presence of channel blockage, and “long” (in bits), which is desirable with respect to fixed overhead.

3.3.2 Packet Latency

In selective repeat schemes, such as implemented in the MEP, there are three main sources of latency: input queueing, transmission delay, and output buffering. Determining analytically the queueing delay associated with the transmit buffer is difficult and depends on the channel load, which also includes packet retransmissions. This problem is addressed in [56–58] and will not be discussed further in this report. Transmission delay includes the propagation time between the transmitter and receiver, δ_p , delay associated with sending a nonzero length packet over a fixed rate channel, L_t/C , where L_t is the packet length in bits and C is the channel rate in bps, and delays associated with retransmission attempts. Transmission delay depends on the average number of transmission attempts, given by [46]:

$$A_t = P_s + 2P_s(1 - P_s) + 3P_s(1 - P_s)^2 + 4P_s(1 - P_s)^3 + \dots \quad (56)$$

$$= \frac{P_f}{1 - P_f} + 1 \quad (57)$$

where $P_s = 1 - P_f$ is the probability of a successful transmission and P_f is the probability of a failed transmission, and both probabilities are assumed to be independent from transmission to transmission. In this situation, average transmission delay is given by [59]:

$$\delta_t = \delta_p + \frac{L_t}{C} + \left(2\delta_p + \delta_a + \frac{L_t}{C}\right) \frac{P_f}{1 - P_f} \quad (58)$$

where δ_t is the average transmission delay, and δ_a is the time required to generate the ACK, including the delay associated with sending it over the fixed rate channel, L_a/C , where L_a is the length of the ACK in bits. The first two terms in (58) essentially represent the one way transmission time, while the third term represents the product of the round-trip transmission time (in parenthesis) and the average number of retransmissions (i.e., $A_t - 1$). If the channel can be modeled by a two-state Markov process then the fact that the average number of transmission attempts is the inverse of P_s can be used to rewrite (57) as:

$$A_t = \frac{P_{s|f} + P_{f|s}}{P_{s|f}} \quad (59)$$

$$= \frac{P_{f|s}}{1 - P_{f|f}} + 1 \quad (60)$$

where $P_{s|f} = 1 - P_{f|f}$ is the probability of a successful transmission, given that the transmission one round trip earlier was a failure, and $P_{f|f}$ and $P_{f|s}$ are similarly defined. Since the Markov model that describes the propagation channel is usually defined with a sampling interval much shorter than the round trip time, the relationships between $P_{s|f}$ and P_{bg} , $P_{f|s}$ and P_{gb} , etc. must be determined with the K -step transition matrix, $\mathbf{M}^K$ [60, 61], where $K = T_r/T_M$, T_r is the round trip time, and T_M is the sampling interval of the Markov model described by the transition matrix $\mathbf{M}$. Using (60), the average transmission delay of a selective repeat ARQ scheme over a Markov channel is given by:

$$\delta_t = \delta_p + \frac{L_t}{C} + \left(2\delta_p + \delta_a + \frac{L_t}{C} \right) \frac{P_{f|s}}{1 - P_{f|f}} \quad (61)$$

The expressions in (58) and (61) are only applicable in selective repeat ARQ implementations where in-order of delivery of packets is not assured. Such implementations may be appropriate for the delivery of IP packets where, in general, out of sequence delivery is not a concern. However, from an application's viewpoint in sequence delivery typically does matter. Hence, if it is not addressed by a link layer protocol it must be handled elsewhere, such as TCP or within the application itself.

In [62], average packet latency from the time of transmission until it is delivered to the end-user (i.e., resequencing delay is included but transmit queueing delay is ignored) is examined for a selective repeat ARQ protocol under the assumption of independent transmission attempts. Details of the derivation will not be repeated here. Instead, the final expression for the average latency experienced by a packet from the time of transmission until its delivery to the end-user is simply taken to be [62]:

$$\delta_t = \delta_p + \frac{L_t}{C} \sum_{j=1}^K (-1)^j \binom{K}{j} \frac{P_f^{2j}}{(1 - P_f^j)} + \frac{(K+1) \frac{L_t}{C}}{(1 - P_f)} \sum_{j=1}^K \frac{(-1)^{j+1}}{(j+1)} \binom{K}{j} \frac{P_f^j (1 - P_f^{j+1})}{(1 - P_f^j)} \quad (62)$$

where δ_p is the one-way propagation time, in seconds, and K is the round trip propagation time, in units of packets. Because correlated packet errors were not considered, (62) serves only as an approximation to the expected latency in the MEP.

Experimental results for end-to-end (i.e., application-to-application) average packet latency were obtained with the MEP channel test application. Experiments were designed to minimize the impact of queueing delay. This was accomplished by generating packets at the relatively modest rate of 1 per second. Of

Table 6
Latency budget for MEP and MOTM implementation.

<u>Source</u>	<u>End-to-End Latency (ms)</u>
MEP execution time (x2)	250
MEP buffering for rate control (x2)	200
terminal buffering and interleaving (x2)	150
propagation time	250
satellite processing and buffering	150
Total:	1000

course, queueing delay could not be eliminated entirely, but by attempting to minimize it, comparisons to (62) could be accommodated. Also, to facilitate comparisons with (62) a latency budget of the various “implementation-related” delays introduced by the system is given in Table 6. These delays are simply lumped into a single value and used as δ_p in (62). The estimate for MEP execution time includes potential operating system delays and the fact that the MEP consists of 3 major process threads. MEP buffering for rate control is a configurable parameter where the primary tradeoff is buffer length and processing requirements. Buffers sizes of 100 ms were used for all of the latency experiments discussed here. Figure 37 summarizes the results for two different packet lengths, $L_t = 39$ bytes, or ≈ 190 ms, and $L_t = 137$ bytes, or ≈ 630 ms, and two waiting times between transmission attempts, $T_w = 3.5$ sec and $T_w = 1.5$ sec. In Figure 37a, PER was estimated according to (53). The relatively long waiting time for the experiments summarized in this figure was chosen for two reasons. First, it is slightly larger than the round trip propagation time for the longest packet length, given by $2(\delta_p + L_t/C) \approx 3.25$ sec. Second, longer waiting times will yield reduced correlation between transmission attempts, facilitating comparisons to (62). In general, the results in Figure 37a agree reasonably well with the predicted latencies. Note that in generating the theoretical curves contained in both figures, a packet length of 410 ms (i.e., the average of the two packet lengths used in the experiments) was used. For the experiments summarized in Figure 37b, the waiting time of $T_w = 1.5$ sec. was less than the round trip time, leading to at least one retransmission for every packet, regardless of whether or not it was received correctly at the other end of the link. In situations where the packet was received correctly on the initial transmission, the duplicate packet was simply ignored by the MEP. In estimating PER for these experiments, the “extra” retransmissions (i.e., N_g , or one extra retransmission for every packet generated) were subtracted out, yielding:

$$\widetilde{P}_f = \frac{N_r - N_g}{N_r} \quad (63)$$

where $\widetilde{P}_f$ is a somewhat pessimistic estimate of PER given the waiting time specified. In comparing Figures 37a and 37b, it can be seen that the decreased waiting time between transmission attempts did improve end-to-end latency, but not by as much as predicted by (62). There are several reasons for the discrepancy





Figure 37. Average packet latency. (a) Time between transmission attempts = 3.5 sec. (b) Time between transmission attempts = 1.5 sec.

between the experimental and theoretical results. First, the extra retransmissions caused by the relatively small value for T_w increased the queueing delay, which is not predicted by (62). Second, the shorter waiting time yields increased correlation between the transmission attempts, another condition not accounted for in (62). Finally, note that both figures contain a curve label "optimized." This curve was generated with (62), assuming a packet length of 410 ms, and $\delta_p = 0.5$ sec., a factor of two improvement over Table 6 that could probably be obtained by optimizing the MEP and certain terminal buffers.


In addition to average packet latency, the data recorded during the experiments were used to generate histograms for packet latency. Typical results for each of the three test environments are illustrated in Figure 38. Bin sizes of 200 ms were used to generate the histograms. Moreover, each histogram was normalized by its sample size so that it actually represents an estimate of the probability mass function (pmf). Note from the figure the relatively high probability of a latency around 1 to 1.5 seconds or so, representing packets that were received on the first transmission attempt. The primary difference among the pmfs in Figure 38 is the duration of the "tail." In the open environment virtually every packet is received within 20 seconds. On the other hand, the rural and urban environments exhibit nonzero probabilities of packet arrival out to 60 seconds. Visually, the pmfs look as though they might be approximated by an exponential distribution, with standard curve fitting procedures yielding the distribution parameters. However, in the absence of a more thorough analytical framework, this exercise seems premature.

Finally, the standard deviation of the packet latencies from each of the experiments was estimated according to:


$$\sigma = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\lambda_i - \bar{\lambda})^2} \quad (64)$$



(a)



(b)



(c)

Figure 38. Histogram of packet latencies. (a) Open environment. (b) Rural environment. (c) Urban environment.




Figure 39. Standard deviation of packet latency. (a) Time between transmission attempts = 3.5 sec. (b) Time between transmission attempts = 1.5 sec.

where σ is the estimated standard deviation, λ_i is the latency of the i th packet, $\bar{\lambda}$ is the estimated mean latency, and n is the sample size, (i.e., the number of packets collected during an experiment). Results are depicted in Figure 39 for each of the two wait times. It is interesting to note that the standard deviations are reasonably close to the mean packet latencies illustrated in Figure 37, perhaps lending credence to the idea that the distribution of packet latencies might be approximated with an exponential distribution.

3.4 SUMMARY

An overview of the MEP, including the ARQ-based error control mechanism as well as brief discussion of the various supported applications, was given. The MEP was used in numerous experiments in the various test environments described in Section 2 to measure PER, throughput efficiency, and packet latency. With respect to PER and throughput efficiency, analytical expressions were obtained and general agreement between these expressions and the measured results was observed. Packet length is one of the key parameters that effects PER and throughput efficiency for a given fading environment. On the one hand, shorter packet lengths are less likely to be corrupted by fading, yielding smaller PERs and increased throughput efficiency. On the other hand, with fixed length packet headers, shorter packet lengths suffer from poor overhead efficiency. Higher data rate channels lead to a more acceptable trade space in the sense that they allow for packets to be simultaneously “short” (in seconds), which is desirable in the presence of channel blockage, and “long” (in bits), which is desirable with respect to fixed overhead. The packet latency results are harder to interpret for several reasons. First, analytical expressions for packet latency in ARQ systems with correlated errors are quite complicated due to the analysis of the input queueing delays. Secondly, the MEP implementation was not optimized for latency, leading to slightly inflated measurement values. The first problem was addressed by structuring the MEP experiments to reduce channel loading and therefore

minimize the effects of input queueing. Moreover, results were compared to a simple analytical expression derived for the case of independent packet errors, a situation typically not satisfied during the MEP experiments. The second problem was addressed by presenting analytical results that assume an optimized MEP implementation. Experimental and analytical results show that average packet latency can range from 1 to 5 seconds for lightly shadowed channels to as high as 5 - 20 seconds for heavily shadowed channels.

4. POSITIONER CHARACTERIZATION

In this section of the report an overview of the results from experiments designed to characterize the MOTM positioner is presented. The primary performance measures considered here are LOS error, in degrees, and the associated loss in antenna gain, measured in dB. A statistical analysis of these quantities, similar to the one conducted in Section 2.2 for the MOTM fading channel, is given in the sequel. With respect to statistical *modeling*, emphasis is given to the mispointing data since these results are more general, allowing interested parties to apply arbitrary antenna patterns to predict the gain losses. However, it is important to note that even the LOS error results have limited generality since the various hardware components, control algorithms and other positioner design choices that affect performance will vary widely from system to system. In other words, the results and models presented here are probably not extensible to arbitrary positioner designs.

The format of this section is similar to that of Section 2, consisting of an overview of the experiments followed by an analysis of the results. However, compared to the fading channel, there is far less to draw from the open literature for modeling the effects of antenna mispointing. The Laplacian density was found to be a good model for the azimuth and elevation error densities, but suitable models for the second order statistics are not explored here. Instead, these results are simply presented as is and may therefore serve as a reference in the space of possible positioner performance data. In general, the MOTM positioner performance is quite good. Although the rough terrain course described below is probably not as severe as the worst case military scenario, the LOS error rarely exceeds 0.5° . The RMS pointing error over the rough terrain course was 0.11° at a 40° elevation angle and 0.14° at a 65° elevation angle. Table 7 summarizes RMS pointing performance for each of the experiments.

4.1 EXPERIMENT OVERVIEW

In general, the MOTM positioner experiments were simpler to execute than the channel characterization and protocol performance experiments since use of the MOTM terminal itself was not required. Instead, only the positioner system, including the antenna control unit (ACU), servos, and AHRS was enabled and configured to point to various fixed reference angles while the vehicle traversed different terrain types. Extensive logging of ACU and AHRS data was conducted to facilitate offline processing.

Table 7
RMS LOS pointing error for each of the experiments.

Environment/elevation angle	RMS pointing error (deg.)
improved/ 40°	0.04°
unimproved/ 40°	0.09°
unimproved/ 65°	0.12°
rough/ 40°	0.11°
rough/ 65°	0.14°




Figure 40. ACU position control loop.

4.1.1 Measurement Approach

The ACU azimuth and elevation control loop position errors are used as the measurement of mispointing angle for subsequent analysis. The ACU calculates these errors by subtracting the gimbal's inertial orientation from the inertial pointing command, as depicted in Figure 40. The gimbal's inertial orientation is determined based on knowledge of the relative gimbal angles, as measured by resolvers on each axis, and a coordinate transformation of the vehicle's inertial attitude, as measured by the TALIN. These quantities are calculated internally by the ACU and logged at a rate of 50 Hz, yielding timeseries data like those in Figure 41. Note that although the control loop position errors are the best possible measurement of pointing error available from the MOTM system, it is not perfect. Sources of possible error with this measurement approach include misalignment between the TALIN and the antenna's RF LOS angle as well as limited accuracy in the TALIN and resolvers. These errors are effectively ignored in the analysis that follows.

From the azimuth and elevation mispointing angles the LOS error can be approximated for small azimuth and elevation errors as:

$$\theta \approx \sqrt{\theta_{az}^2 + \theta_{el}^2} \quad (65)$$

where θ is the LOS error, θ_{az} is the azimuth error, and θ_{el} is the elevation error. Finally, loss in antenna gain can be approximated by:

$$\alpha \approx e^{-2.76(\theta/\theta_{3dB})^2} \quad (66)$$

where θ_{3dB} is the 3 dB beamwidth of a circular reflector antenna:

$$\theta_{3dB} \approx \frac{70\lambda}{D} \quad (67)$$

where λ is the signal wavelength and D is the antenna diameter. For the 12" reflector in the MOTM system $\theta_{3dB} \approx 1.5^\circ$ on the UL and $\theta_{3db} \approx 3.5^\circ$ on the DL. The approximation in (66) is quite good in the range $-\theta_{3dB}/2 \leq \theta \leq \theta_{3dB}/2$, a condition that is always satisfied by the MOTM positioner in the experimental results discussed here.




Figure 41. Time series data segments for LOS error and antenna gain from the three test terrains. In these examples the elevation angle was 40° . Each segment is one minute in duration. (a) Improved road. (b) Unimproved road. (c) Rough road.




Figure 42. Terrain types for positioner testing. (a) Improved road course in Harvard/Ayer, MA. (b) Unimproved road course (highlighted in blue) and rough road course (highlighted in red) at Ft. Devens, MA.

4.1.2 Course Descriptions

As with the blockage channel experiments, a representative collection of environments consisting of three courses was chosen for the positioner tests. The courses include an improved road course, an unimproved road course, and a rough terrain (i.e., off-road) course. The courses are located at Ft. Devens, MA, and the surrounding area. Because pointing performance with the MOTM positioner depends on the elevation angle, experiments in the unimproved and rough road courses were conducted at both 40° and 65° look angles.

Improved Road Course

Referring to the map of Harvard/Ayer, MA, in Figure 42a, the improved road course begins at the intersection of Route 2 and Ayer Rd. (exit 38 B), due north. The vehicle follows Ayer Rd. for approximately 2 miles traveling at speeds between 25 and 55 mph before reaching a rotary approximately 350 feet in diameter (just before the rotary, Ayer Rd. changes names to Harvard Rd.). The vehicle enters the rotary

from the south, completes two trips around the rotary at approximately 25 mph and exits the rotary from its entry point due south on Harvard Rd. The vehicle follows Harvard/Ayer Rd. back to the starting point. Road conditions for this course consist entirely of paved roads. The experiment duration for this course was approximately 6 minutes.

Unimproved Road Course


Referring to the map of Ft. Devens, MA in Figure 42b, the unimproved road course begins at the intersection of Bivouac Rd. and an unmarked trail near the northwest corner of the Turner drop zone. The vehicle proceeds south on Bivouac until making a left on a fine gravel road on the southern border of the drop zone. Next, the vehicle turns left onto an unmarked trail near the middle of the drop zone and proceeds north until making another left on the unmarked trail that returns to the starting location. The unimproved road course consists of approximately 75% well-traveled trails and 25% fine gravel road. The course length is approximately 2.5 miles and vehicle speeds ranged from approximately 15 to 35 mph. Experiment durations were approximately 4 minutes.

Rough Road Course


Referring to the map of Ft. Devens, MA in Figure 42b, the rough road course begins on an unmarked trail off Old Turnpike Rd. heading southwest. The beginning of the trail is characterized by a relatively steep hill. The vehicle remains on the trail for approximately 0.5 miles before making a sharp right onto Combat Rd. The vehicle then follows Combat Rd. for approximately 0.5 miles at which point it stops at the intersection of Combat Rd. and Old Turnpike Rd. The course consists entirely of rough terrain and poorly maintained trails. Vehicle speeds ranged from 10 to 25 mph and experiment durations were approximately 3 to 4 minutes.

Comparison to Churchville B Course


Simulated HMMWV motion data for Course B in the Churchville Test Area at the Aberdeen Proving Grounds were obtained from the U. S. Army Tank-automotive Armaments Command (TACOM) for experimental purposes. The vehicle motion data were calculated for a point 60" above the ground and 120" aft of the forward frame rail of the HMMWV (i.e., the approximate location of the MOTM gimbal) assuming that the vehicle was traveling at a constant 20 mph along a "straightened" version of the course. The course itself consists of approximately 4 miles of rugged hilly terrain and is considered to be representative of a typical rough, off-road environment for a military vehicle. A comparison of the angular velocities and linear accelerations logged during the positioner experiments over the courses described above and those given by the Churchville B data is provided in Figure 43. In general, the Churchville B proves to be more severe than the rough course used in the field experiments. With respect to azimuth velocity and forward acceleration, the Churchville B data appear somewhat mild by comparison to other courses. However, these data are probably artifacts of the simulation parameters. Specifically, the fact that a straightened version of Churchville B was simulated probably gives rise to the relatively small azimuth velocities, and the assumed constant platform speed is probably responsible for the moderate forward accelerations. In any case, with the possible exception of vertical acceleration, the severity of the unimproved and rough road courses is not too far removed from that of Churchville B.




(a)




(b)




(c)



(d)



(e)



(f)

Figure 43. Comparison of vehicle dynamics over test terrains to Churchillville B. (a) Roll velocity. (b) Forward acceleration. (c) Pitch velocity. (d) Lateral acceleration. (e) Azimuth velocity. (f) Vertical acceleration.

4.2 STATISTICAL ANALYSIS

In this section, a statistical analysis of the pointing error and loss in antenna gain due to mispointing for the MOTM positioner is given. Judging from the open literature, similar analyses with comparable systems are rare and focused solely on probability density models. Statistical properties addressed in this section include probability density models, fade duration statistics and level crossing rates.

4.2.1 Probability Density Models for LOS Error and Antenna Gain

Very few probability density models for LOS error in LMS systems with directional antennas have been proposed in the open literature. One example where such work has been done involves the Advanced Communications Technology Satellite (ACTS) system, where mobile propagation experiments were conducted at 20 GHz using the ACTS mobile terminal (AMT). The AMT uses a mechanically steered elliptically-shaped reflector antenna with dimensions of approximately 8 inches by 3 inches. Because of the wide elevation beamwidth, no satellite elevation tracking is necessary. For azimuth steering, a hybrid open-loop/closed-loop system is used. An inertial yaw rate sensor is used for open-loop control and a mechanical dithering scheme, where the antenna is swept sinusoidally 1° in each direction about the azimuth pointing angle at a rate of 2 Hz, provides the closed-loop control. More details on the AMT antenna and tracking system can be found in [63, 64].

In [65] mispointing error for the AMT over smooth road conditions is characterized experimentally. Measurements of the vehicle pitch, roll, and heading were taken at 0.1 mile intervals along a specific route traveled by the terminal. The error associated with these measurements was used to upper bound the azimuth and elevation angle mispointing errors at 3.9° and 3.3° , respectively. Finally, through logarithmic interpolation of the antenna gain pattern data, the loss in received signal power due to antenna mispointing was calculated to be on the order of 1.5 dB. These results were then used to rationalize the 1 to 2 dB variations in received signal power observed during previous AMT runs.

In [66, 67] data from a wider range of experiments were used to develop a statistical characterization of the loss in antenna gain due to mispointing with the AMT. Specifically, a bimodal density function was proposed since this model was observed to fit the experimental data. In [68], another statistical analysis of antenna mispointing and the associated loss in antenna gain was conducted. In this work, Gaussian distributed mispointing errors in the azimuth and elevation directions were assumed. Closed form expressions for the corresponding pdfs of the LOS mispointing error under the assumptions of equal and unequal variances are given. In addition, the antenna mispointing pdf is used to generate the pdf for received signal loss, from which the average BER is easily computed.

Data collected with the MOTM positioner have been used to develop probability density models of the pointing error and associated loss in gain. In all cases azimuth and elevation error densities were estimated via histogram using 100 equally spaced bins between -1° and 1° . Interestingly, the azimuth and elevation errors are well approximated by Laplacian densities:

$$p_{az}(\theta_{az}) = \frac{1}{2} \lambda_{az} e^{-\lambda_{az} |\theta_{az}|} \quad (68)$$

$$p_{el}(\theta_{el}) = \frac{1}{2} \lambda_{el} e^{-\lambda_{el} |\theta_{el}|} \quad (69)$$




Figure 44. Probability densities for azimuth and elevation error over rough terrain. The elevation angle was 65° . Each plot also contains the least squares fit of a Laplacian density to the estimated pdfs. (a) Azimuth. (b) Elevation.

where λ_{az} and λ_{el} are the azimuth and elevation density parameters, respectively, and are not necessarily equal. Figures 44a and 44b illustrate the azimuth and elevation mispointing densities, respectively for the rough terrain course with an elevation angle of 65° . While only these results are depicted graphically, the model yielded excellent results for all measured elevation angles and terrains. Another interesting aspect of these results is that the random variables that describe the azimuth and elevation pointing error are almost completely uncorrelated. Table 8 summarizes the density parameters and correlation coefficients between azimuth and elevation error for each of the experiments. Note from the Table that λ decreases as the terrain conditions worsen. This behavior is expected since smaller values for λ are indicative of a greater spread in the density (the variance of a Laplacian random variable is inversely proportional to λ^2).

Exactly why the Laplacian density is such a good model for the azimuth and elevation mispointing error is not clear at this time. Relative to the Gaussian density, the Laplacian density is characterized by a higher peak around the mean and heavier tails. Typical applications of the Laplacian density are in the realm of signal processing where it serves as a good model for prediction error and transform coefficients in speech and image coding systems. With respect to the MOTM positioner it is interesting to note that the Laplacian model seems to hold regardless of operating terrain.

LOS error densities were estimated by forming the histogram of (65) with 100 equally spaced bins between 0° and 1° . The pdfs were also used to produce complementary cdfs. The complementary cdfs are depicted in Figure 45 and the pdfs are shown in Figures 46, 47, 48, 49, and 50. Assuming that azimuth and elevation errors are independent, the pdf of the LOS pointing error, $\theta = \sqrt{\theta_{az}^2 + \theta_{el}^2}$, can be found by

Table 8

Azimuth and elevation error density model parameters. The parameters were derived from a least squares fit to the measured data.

Environment/elevation angle	density parameters		correlation coeff.
	λ_{az}	λ_{el}	ρ
improved/40°	52.74	32.05	0.103
unimproved/40°	30.70	23.00	0.095
unimproved/65°	18.73	20.10	0.155
rough/40°	22.66	17.90	0.105
rough/65°	14.80	16.47	0.168




Figure 45. Complementary cdfs for the three terrain types and two elevation angles.




Figure 46. Probability density for LOS error and antenna gain in improved terrain, 40° elevation angle. (a) LOS error. (b) Antenna gain relative to perfect pointing.

applying the standard techniques [69]:

$$p_{LOS}(\theta) = \frac{\lambda_{az}\lambda_{el}}{4} \int_0^\theta e^{-\lambda_{az}\sqrt{\theta^2-w^2}} e^{-\lambda_{el}w} \frac{\theta}{\sqrt{\theta^2-w^2}} dw \quad (70)$$

$$= \frac{\theta\lambda_{az}\lambda_{el}}{4} \int_0^{\frac{\pi}{2}} e^{-\theta\sqrt{\lambda_{az}^2+\lambda_{el}^2}\cos(\phi-\phi_0)} d\phi \quad (71)$$

where $\phi_0 = \arccos\left(\frac{\lambda_{az}}{\sqrt{\lambda_{az}^2+\lambda_{el}^2}}\right)$. Although it may be possible to simplify (71) further, it is probably just as easy to carry out the integration numerically. This was the approach used in generating the density models for θ that appear in Figures 46 - 50. Note the close agreement between the histograms and the Laplacian-based model in (71), lending further support to the observation that azimuth and elevation mispointing errors are uncorrelated. Lastly, densities for the loss in antenna gain were estimated by generating histograms of the result in (66) with 100 equally spaced bins between 0 and -1 dB. These results are depicted in Figures 46 - 50.

From the figures, the heavier tails that coincide with rougher terrain and higher elevation angles are readily apparent. Also, the difference in UL and DL gain due to mispointing manifests itself as a heavier distribution tail due to the narrower beamwidth on the UL.

4.2.2 Average Fade and Connection Durations

Average fade and connection durations as a function of LOS error and antenna gain relative to perfect pointing were estimated from the data and depicted graphically in Figure 51. For the LOS error data, a fade is defined as an event where the LOS error exceeds a threshold and a connection is defined as an event where




Figure 47. Probability density for LOS error and antenna gain in unimproved terrain, 40° elevation angle. (a) LOS error. (b) Antenna gain relative to perfect pointing.






Figure 48. Probability density for LOS error and antenna gain in unimproved terrain, 65° elevation angle. (a) LOS error. (b) Antenna gain relative to perfect pointing.




(a)




(b)

Figure 49. Probability density for LOS error and antenna gain in rough terrain, 40° elevation angle. (a) LOS error. (b) Antenna gain relative to perfect pointing.



(a)



(b)

Figure 50. Probability density for LOS error and antenna gain in rough terrain, 65° elevation angle. (a) LOS error. (b) Antenna gain relative to perfect pointing.

the LOS error remains below a threshold. Because antenna gain is considered relative to perfect pointing, the thresholds, in dB, are negative. Hence, fades and connections are defined as in Section 2.2.3. Moreover, only UL gain is shown since this represents the worst case compared to DL gain. From the figure it can be seen that connection durations increase (decrease) to their maxima (minima) as a function of LOS error (gain relative to perfect pointing). Likewise, fade durations decrease (increase) to their minima (maxima) as a function of LOS error (gain relative to perfect pointing). Maximum durations are given by the experiment duration and minimum durations are given by the sampling interval.

4.2.3 Level Crossing Rates

Level crossing rates as a function of LOS error and antenna gain relative to perfect pointing were estimated from the data. Results for each of the experiments are depicted graphically in Figure 52. The peak LCR is approximately 7 Hz and occurs between 0.02° and 0.1° , depending on the terrain and elevation angle. As LOS error increases and antenna gain decreases, the LCR decreases, reflecting the fact that these thresholds are crossed less frequently.

4.3 SUMMARY

An overview of experiments designed to measure the performance of the MOTM positioner in various terrain and elevation angles was given. Antenna pointing error and the corresponding loss in antenna gain were quantified in terms of probability density and cumulative distribution functions. Azimuth and elevation pointing error were observed to follow a Laplacian distribution, a result that appears independent of terrain type or elevation angle. Moreover, azimuth and elevation pointing errors were shown to be nearly uncorrelated. This fact was used to derive a Laplacian-based model that describes the LOS pointing error density. The model (and data) generally behave as one would expect, producing heavier tails as terrain conditions degrade and elevation angle increases. The average duration of mispointing events (i.e., “fades”) was also examined and observed to decay rapidly as LOS error increases. Similarly, the average duration of events where pointing performance is acceptable (i.e., “connections”) increases as LOS error increases. Finally, the peak LCR was observed to be approximately 7 Hz, and occurs in the range of 0.02° to 0.1° , depending on terrain and elevation angle.






Figure 51. Average fade and connection durations over the three terrain types. (a) Average connection duration as a function of LOS error. (b) Average fade duration as a function of LOS error. (c) Average connection duration as a function of UL gain relative to perfect pointing. (d) Average fade duration as a function of UL gain relative to perfect pointing.



(a)



(b)

Figure 52. Level crossing rates for the three terrain types. (a) Level crossing rate as a function of LOS error. (b) Level crossing rate as a function of UL gain relative to perfect pointing.

5. CONCLUSIONS

This report contains an analysis of results from experiments designed to characterize the EHF SOTM propagation channel, assess the communications performance of a selective repeat HARQ protocol implementation over this channel, and quantify the antenna pointing performance delivered by the MOTM positioner over various terrain. The primary application of these results is intended to be in the area of channel modeling. Of course, suitable channel models are an important component to the processes of waveform design, protocol development, performance analysis, etc. The contributions of this report include the following:

- derivation of probability density models for received signal power in different blockage environments
- investigation of simple discrete blockage models for the EHF SOTM channel and performance comparison to the field data
- derivation of second order statistics, including average fade and connection durations, fade and connection duration distributions, and level crossing rates for the EHF SOTM channel
- BER performance results for the EHF SOTM channel
- packet error rate performance results for a selective repeat HARQ protocol when applied to the EHF SOTM channel
- throughput efficiency performance results for a selective repeat HARQ protocol when applied to the EHF SOTM channel
- average packet latency and packet latency distributions associated with a selective repeat HARQ protocol when applied to the EHF SOTM channel
- derivation of probability density models for LOS error and antenna gain associated with the MOTM positioner over various terrain
- statistical analysis of the fading behavior induced by antenna mispointing with the MOTM system

The details of these results can be found within the body of the report and are not discussed further in this section. However, there are a number of related issues that are worth discussing further.

With respect to the EHF SOTM propagation channel the point was made earlier in this report that the data collected with the MOTM system represents only a small fraction of possible channel conditions. Numerous references to other measurement campaigns are also cited here. One compelling issue regarding this data is how to use it in the form of a requirements specification to drive future system designs. For example, specifying the BER at the output of a codec over an additive white Gaussian noise channel is relatively straightforward. On the other hand, how does one specify the communications performance of a system over the EHF SOTM propagation channel? Obvious performance metrics to consider include PER, throughput efficiency, and packet latency. However, these performance metrics are tightly coupled to the channel itself, implying that what must also be specified are expected channel conditions that include statistics such as average and maximum fade durations, and level crossing rate. Moreover, one set of parameters

will probably not be sufficient. Instead a small collection of representative channels should be specified. Within this collection will likely be a "typical" scenario and a "worst-case" scenario. The work presented in this report represents an important first step toward these goals, but only a first step. Certainly, more work needs to be done in the area of developing appropriate blockage scenarios, specifying the relevant channel statistics, and determining the appropriate performance metrics and reasonable values for these metrics over the various channels. As the Army attempts to develop transformational communications capabilities with support for land mobile platforms, such efforts will be key to their success.

Equally important to the success of transformational communications efforts is an understanding of how the underlying EHF SOTM channel affects higher layer protocols and applications. The MEP experiments summarized in this report provide insight into the essence of the numerous challenges associated with internetworking over EHF SOTM channels. These challenges include accommodating PERs and latencies that are orders of magnitude greater than experienced with typical wired, and even wireless channels. To the extent that certain protocols, like multiple access, resource allocation, etc. will likely be custom-designed for military systems anyway, these issues can be addressed as part of these design efforts. Of course, the preceding statement is predicated on the discussion in the previous paragraph: that "standard" blockage models and scenarios be developed to assist in waveform and protocol design efforts. In addition, there are certain widely deployed and standard protocols, most notably TCP/IP, that can not be customized to the EHF SOTM channel. It is important to assess the impact of the EHF SOTM physical and link layers on these protocols in order to anticipate end-to-end system performance. Moreover, to the extent that legacy applications will be used over networks that contain EHF SOTM links, it is important to understand how these applications will respond to intermittent connectivity and longer latencies. Presumably, new applications will be made robust to these characteristics. In any case, it should be clear that a concerted "educational" effort among proponents of EHF SOTM must be brought to bear on the networking, protocol, and application development communities to help them understand the implications of the EHF SOTM channel.

With respect to the positioner performance data several obvious issues arise. First, given the superior performance of the MOTM positioner, can pointing requirements be relaxed somewhat with the end-goal of reducing cost? This, in turn, begs the question: how well does a typical positioner have to perform? Where does one draw the line between performance that must be delivered by the positioner versus mispointing mitigation solutions such as FEC coding and interleaving, or HARQ? These questions are extremely important, but also quite difficult to answer definitively due to the large number of variables involved. However, it seems reasonable to assume that some distribution of LOS errors could be accommodated by the system. In light of this, channel statistics and models similar to those presented in this report must be developed for future positioner designs. In addition, joint models that account for blockage and antenna mispointing are desirable. To this end, the discrete models (specifically the HMMs) discussed in this report could potentially be extended to include more states, distinguish between more than two fade depths within each state, or both.

APPENDIX A

GLOSSARY

ACK: ACKnowledgment	HDL: High-level Data Link Control
ACTS: Advanced Communications Technology Satellite	HHR: High Hop Rate
ACU: Antenna Control Unit	HMM: Hidden Markov Models
AFD: Average Fade Duration	HMMWV: High Mobility Multi-purpose Wheeled Vehicle
AHRS: Attitude and Heading Reference System	Hz: Hertz
AMT: ACTS Mobile Terminal	IP: Internet Protocol
ARQ: Automatic Repeat reQuest	Kbps: Kilobits per second
ASCAMP: Advanced Single Channel Anti-jam Man Portable (MILSTAR terminal)	LDR: Low Data Rate
BER: Bit Error Rate	LMS: Land Mobile Satellite
BLOS: Beyond Line of Sight	LOS: Line Of Sight
bps: bits per second	LPD: Low Probability of Detection
cdf: cumulative distribution function	LPI: Low Probability of Intercept
CECOM: U. S. Army Communications-Electronics COMmand Research, Engineering and Development Center, Fort Monmouth, NJ	Mbps: Megabits per second
COTM: Communications On The Move	MELP: Mixed Excitation Linear Prediction (vocoder)
DC: Direct Current	MEP: MOTM Experimental Protocol
DL: DownLink	MILSTAR: MILitary Strategic and Tactical Relay
EHF: Extremely High Frequency (30-300 GHz)	MIT: Massachusetts Institute of Technology
EIRP: Effective Isotropic Radiated Power	mph: miles per hour
FEC: Forward Error Correction	MOTM: MILSTAR On The Move
GF: Galois Field	ms: milliseconds
GHz: GigaHertz	NTC: National Training Center
GPS: Global Positioning System	OF: Objective Force
GUI: Graphical User Interface	OTM: On The Move
HARQ: Hybrid ARQ	pdf: probability density function
	PLGR: Precision Lightweight GPS Receiver
	pmf: probability mass function

PM WIN-T: Program Manager, Warfighter
Information Network - Tactical

PER: Packet Error Rate

PSR: Packet Success Rate

PTT: Push To Talk

QoS: Quality of Service

RF: Radio Frequency

RMS: Root Mean Square

RS: Reed-Solomon (class of FEC codes)

SATCOM: SATellite COMmunications

SNR: Signal to Noise Ratio

SOTM: SATCOM On The Move

STO: Science and Technology Objective

TACOM: U. S. Army Tank-automotive Ar-
maments Command

TALIN: Tactical Advanced Land Intertial Nav-
igator

TC: Transformational Communications

TCP: Transmission Control Protocol

TSC: Time Share of Connections

TSF: Time Share of Fades

UHF: Ultra High Frequency (300-3,000 MHz)

UL: UpLink

REFERENCES

1. J. B. Schodorf, "EHF satellite communications on the move: Baseband considerations." MIT Lincoln Laboratory Technical Report TR-1055, February 2000.
2. R. J. Figucia, "Advanced SCAMP's demodulator design for performing envelope recovery from hardlimited samples," in *IEEE Military Comm. Conf.*, 1992.
3. R. J. Figucia, "Downlink acquisition and tracking procedures for the ASCAMP satellite communications terminal," Tech. Rep. 984, MIT Lincoln Laboratory, 1993.
4. T. C. Marek, "Extending the range of the ASCAMP Pr/No estimator," in *IEEE Military Comm. Conf.*, 1995.
5. A. Abdi, W. C. Lau, M. Alouini, and M. Kaveh, "A new simple model for land mobile satellite channels: First- and second-order statistics," *IEEE Trans. Wireless Comm.*, pp. 519–528, May 2003.
6. C. Loo, "A statistical model for a land mobile satellite link," *IEEE Trans. Vehicular Tech.*, pp. 122–127, August 1985.
7. D. Parsons, *The Mobile Radio Propagation Channel*. New York, NY: Halsted Press, 1992.
8. W. C. Jakes, *Microwave Mobile Communications*. Piscataway, NJ: IEEE Press, 1993.
9. B. Vucetic and J. Du, "Channel modeling and simulation in satellite mobile communication systems," *IEEE J. Select. Areas Comm.*, pp. 1209–1218, October 1992.
10. F. Perez-Fontan *et al.*, "S-Band LMS propagation channel behaviour for different environments, degrees of shadowing and elevation angles," *IEEE Trans. Broadcasting*, pp. 40–76, March 1998.
11. C. Loo and J. S. Butterworth, "Land mobile satellite channel measurements and modeling," *Proc. IEEE*, pp. 1442–1463, July 1998.
12. E. Kubista *et al.*, "Ka-band propagation measurements and statistics for land mobile satellite applications," *IEEE Trans. Vehicular Tech.*, pp. 973–983, May 2000.
13. E. Lutz *et al.*, "The land mobile satellite communications channel – recording, statistics and channel model," *IEEE Trans. Vehicular Tech.*, pp. 375–386, May 1991.
14. H. Suzuki, "A statistical model for urban radio propagation," *IEEE Trans. Commun.*, pp. 673–680, July 1977.
15. F. Hansen and F. I. Meno, "Mobile fading - Rayleigh and lognormal superimposed," *IEEE Trans. Vehicular Tech.*, pp. 332–335, November 1977.
16. M. Rice *et al.*, "K-band land-mobile satellite characterization using ACTS," *Intl. J. Sat. Comm.*, pp. 283–296, January 1996.
17. J. P. Castro, "Statistical observations of data transmission over land mobile satellite channels," *IEEE J. Select. Areas Comm.*, pp. 1227–1235, October 1992.

18. G. E. Corazzo and F. Vatalaro, "A statistical model for land mobile satellite channels and its application to nongeostationary orbit systems," *IEEE Trans. Vehicular Tech.*, pp. 738–742, August 1994.
19. H. Ernst and S. Scalise, "Land mobile DVB-S based broadcast system in Ku-band: Analysis of a measurement campaign," in *Proc. IEEE Intl. Conf. Comm.*, 2003.
20. H. Lin and W. J. Vogel, "Photogrammetric satellite service prediction in a roadside tree shadowing environment," in *Proc. IEEE Vehicular Tech. Conf.*, 1998.
21. F. Perez-Fontan *et al.*, "Statistical modeling of the LMS channel," *IEEE Trans. Vehicular Tech.*, pp. 1549–1567, November 2001.
22. G. L. Stuber, *Principles of Mobile Communication*. Boston, MA: Kluwer Academic Publishers, 2001.
23. C. K. Siew and D. J. Goodman, "Packet data transmission over mobile radio channels," *IEEE Trans. Vehicular Tech.*, pp. 95–101, May 1989.
24. L. N. Kanal and A. R. K. Sastry, "Models for channels with memory and their applications to error control," *Proc. IEEE*, pp. 724–744, July 1978.
25. E. N. Gilbert, "Capacity of a burst noise channel," *The Bell System Technical Journal*, pp. 1253–1265, September 1960.
26. W. Turin, *Digital Transmission Systems: Performance Analysis and Modeling*. New York, NY: McGraw-Hill, 1999.
27. S. Ross, *A First Course in Probability*. Macmillan Publishing Co., 1988.
28. E. O. Elliott, "Estimates of error rates for codes on burst noise channels," *The Bell System Technical Journal*, pp. 1977–1997, September 1963.
29. Y. Hase, W. J. Vogel, and J. Goldhirsh, "Fade durations derived from land mobile satellite measurements in Australia," *IEEE Trans. Commun.*, pp. 664–668, May 1991.
30. S. Yamamoto, E. Okamoto, and T. Ikegami, "Land mobile satellite propagation measurements in the S-band with non-geostationary with ETS-VI satellite," in *Proc. IEEE Vehicular Tech. Conf.*, pp. 31–35, 1997.
31. J. Slack, "Finite state Markov models for error bursts on the ACTS land mobile satellite channel." MSEE Thesis, Brigham Young University, Provo, UT. (Also available as NASA Technical Report ACTS-96-046), 1996.
32. J. S. daSilva and S. Mahmoud, "Capacity degradation of packet radio fading channels," in *Proc. 6th Data Comm. Symp.*, pp. 96–101, 1979.
33. L. Lugand and D. J. Costello, Jr., "A comparison of three hybrid ARQ schemes using convolutional codes on a non-stationary channel," in *IEEE Global Telecomm. Conf.*, 1982.
34. L. F. Chang, "Throughput estimation of ARQ protocols for a Rayleigh fading channel using fade- and interfade-duration statistics," *IEEE Trans. Vehicular Tech.*, pp. 223–229, February 1991.
35. T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: Wiley, 1991.

36. D. Cygan, "Analytical evaluation of average bit error rate for the land mobile satellite channel," *Intl. J. of Sat. Comm.*, pp. 99–102, April 1989.
37. J. G. Proakis, *Digital Communications*. New York, NY: McGraw-Hill, 1989.
38. C. J. Karpinsky, "SCAMP Block II baseline 0 characterization: Methodology and results." Proj. Rep. SC-115, MIT Lincoln Laboratory, 2000.
39. G. D. Forney, Jr., "Burst-correcting codes for the classic bursty channel," *IEEE Trans. Commun.*, pp. 772–781, October 1971.
40. M. Richharia, *Mobile Satellite Communications: Principles and Trends*. Addison-Wesley, 2001.
41. E. Lutz, "Code and interleaver design for data transmission over fading channels," in *IEEE Global Telecomm. Conf.*, 1984.
42. J. Hagenauer and E. Lutz, "Forward error correction coding for fading compensation in mobile satellite channels," *IEEE J. Select. Areas Comm.*, pp. 215–225, February 1987.
43. B. Schein and S. B. Bernstein, "A forward error control scheme for GBS and BADD," in *IEEE Military Comm. Conf.*, 1997.
44. J. J. Metzner, *Reliable Data Communications*. San Diego, CA: Academic Press, 1998.
45. J. W. Byers, M. Luby, and M. Mitzenmacher, "A digital fountain approach to reliable asynchronous reliable multicast," *IEEE J. Select. Areas Comm.*, pp. 1528–1540, October 2002.
46. S. Lin and D. J. Costello, *Error Control Coding: Fundamentals and Applications*. Englewood Cliffs, NJ: Prentice Hall, 1983.
47. A. V. McCree and T. P. Barnwell III, "A mixed excitation LPC vocoder model for low bit rate speech coding," *IEEE Trans. Speech and Audio Proc.*, pp. 242–250, July 1995.
48. E. J. Daniel and K. A. Teague, "Federal standard 2.4 Kbps MELP over IP," in *IEEE Midwest Symp. on Cir. and Sys.*, pp. 568–571, 2000.
49. C. J. Weinstein and J. W. Forgie, "Experience with speech communication in packet networks," *IEEE J. Select. Areas Comm.*, pp. 963–980, December 1983.
50. J. B. Schodorf, "Error control for Ka-band land mobile satellite communications systems," in *Proc. IEEE Vehicular Tech. Conf.*, 2000.
51. D. E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching*. Reading, MA: Addison Wesley, 1998.
52. P. Ferguson and G. Huston, *Quality of Service: Delivering QoS on the Internet and Corporate Networks*. New York, NY: John Wiley and Sons, Inc., 1998.
53. J. B. Schodorf and M. A. Gouker, "Performance evaluation of a hybrid ARQ protocol implementation for EHF SATCOM on the move systems," in *IEEE Military Comm. Conf.*, 2001.
54. D. Bertsekas and R. Gallager, *Data Networks, Second Edition*. Upper Saddle River, NJ: Prentice Hall, 1992.

55. E. Lutz, "Simulation of FEC/ARQ data transmission using stored land mobile satellite channels," in *Proc. IEEE Vehicular Tech. Conf.*, 1986.
56. D. Towsley, "A statistical analysis of ARQ protocols operating in a nonindependent error environment," *IEEE Trans. Commun.*, pp. 971–981, July 1981.
57. R. Fantacci, "Queuing analysis of the selective repeat automatic repeat request protocol for wireless packet networks," *IEEE Trans. Vehicular Tech.*, pp. 258–264, May 1996.
58. J. G. Kim and M. M. Krunz, "Delay analysis of selective repeat ARQ for a Markovian source over a wireless channel," *IEEE Trans. Vehicular Tech.*, pp. 1968–1981, September 2000.
59. Q. Yang and V. K. Bhargava, "Reducing time delay of ARQ protocols by simple error correcting codes," *Elec. Letters*, pp. 1447–1448, August 1990.
60. D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes*. New York, NY: Chapman and Hall, 1984.
61. J. B. Schodorf and P. Wu, "Diversity reception for blockage mitigation in EHF land mobile satellite communications systems," in *Proc. IEEE Vehicular Tech. Conf.*, 2001.
62. D. M. Chitre, "Analysis of throughput efficiency and delay in ARQ systems," *COMSAT Technical Review*, pp. 345–367, 1981.
63. A. C. Densmore and V. Jamnejad, "A satellite tracking K- and Ka-band mobile vehicle antenna system," *IEEE Trans. Vehicular Tech.*, pp. 502–513, November 1993.
64. A. Densmore *et al.*, "K- and Ka- band land mobile satellite-tracking reflector antenna system for the NASA ACTS mobile terminal," in *Intl. Mob. Sat. Conf.*, 1993.
65. M. Rice, B. J. Mott, and K. D. Wise, "A pointing error analysis of the ACTS mobile terminal," in *Intl. Mob. Satellite Conf.*, 1997.
66. M. Rice and B. Humpherys, "Statistical models for the ACTS K-band land mobile satellite channel," in *Proc. IEEE Vehicular Tech. Conf.*, 1997.
67. M. Rice and B. Humpherys, "A new model for the ACTS land mobile satellite channel," in *Intl. Mobile Sat. Conf.*, 1997.
68. J. B. Schodorf, "A probabilistic mispointing analysis for land mobile satellite communications systems with directive antennas," in *Proc. IEEE Vehicular Tech. Conf.*, 2001.
69. A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. New York, NY: McGraw-Hill, 1991.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 15 August 2003		3. REPORT TYPE AND DATES COVERED Technical Report
4. TITLE AND SUBTITLE EHF Satellite Communications on the Move: Experimental Results			5. FUNDING NUMBERS C—F19628-00-C-0002	
6. AUTHOR(S) J.B. Schodorf				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Lincoln Laboratory, MIT 244 Wood Street Lexington, MA 02420-9108			8. PERFORMING ORGANIZATION REPORT NUMBER TR-1087	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Space Technology Branch CECOM AMSEL-RD-ST-SS-SP Ft. Monmouth, NJ 07703			10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2003-051	
11. SUPPLEMENTARY NOTES None				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report summarizes the results of experiments conducted with an EHF SATCOM on the move (SOTM) terminal developed to work with LDR MILSTAR. The experiments were designed to measure characteristics of the EHF SOTM propagation channel, yield insight into the performance of an automatic repeat request error control protocol, and assess the pointing performance of the system's antenna positioner. With respect to the propagation experiments, the report contains statistical analyses of the data that yield probability density models for the received signal power, fade and nonfade duration distributions, and estimates of level crossing rates. Discrete models for the EHF SOTM channel are also explored and compared to the measured data. Results from the protocol experiments include packet error rate, throughput efficiency and packet latency as a function of protocol parameters (e.g., packet length) and channel conditions. As with the propagation experiments, the report contains analyses of the antenna pointing data recorded during the positioner experiments that result in statistical models for the LOS error under various terrain conditions.				
14. SUBJECT TERMS			15. NUMBER OF PAGES 108	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Same as Report	