



NRL/FR/5554--08-10,168

High Altitude Relay and Router (HARR)

MICHAEL A. RUPAR

RAYMOND MEREISH

IVAN CORRETJER

BRIAN VOREES

JONATHAN DOFFOH

*Transmission Technology Branch
Information Technology Division*

November 20, 2008

Approved for public release; distribution is unlimited.

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 this 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 Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 20-11-2008		2. REPORT TYPE NRL Formal Report		3. DATES COVERED (From - To) Dec. 1, 2005 – Nov. 1, 2006	
4. TITLE AND SUBTITLE High Altitude Relay and Router (HARR)				5a. CONTRACT NUMBER DOC #N0001406WX20615	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Michael A. Rugar, Raymond Mereish, Ivan Corretjer, Brian Vorees, and Jonathan Doffoh				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER 55-8882-B-6-5	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320				8. PERFORMING ORGANIZATION REPORT NUMBER NRL/FR/5554--08-10,168	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Code 31 875 Randolph Street Arlington, VA 22203-1995				10. SPONSOR / MONITOR'S ACRONYM(S) ONR	
				11. SPONSOR / MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report describes the development and test of processing and transponding payload configurations for the High Altitude Relay and Router (HARR) project. This work was performed by the U.S. Naval Research Laboratory (NRL) to support a viable long-range end-to-end mobile ad hoc wireless network in a tactical environment. HARR has the potential to provide closed network communications and other tactical capabilities between nodes separated by up to 200 miles at a reduced cost as compared to other existing technologies. HARR achieves these results by flying its payloads in untethered balloons at near-space altitudes around 20,000 m (approximately 65,000 ft), providing a relay capability over a substantial area of operation. The report describes the design and integration of the airborne and ground node systems that make up this network, and analyzes test data collected using unicast and multicast transport protocols in an IP-based environment. The field test data discussed in this report was collected at Lubbock, Texas, in June 2006. Additional followup testing was conducted through the summer and fall of 2006 at NRL in Washington, DC.					
15. SUBJECT TERMS High altitude Network relay Transponder Tactical relay RF communications 802.11b					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Michael A. Rugar
Unclassified	Unclassified	Unclassified	UL	49	19b. TELEPHONE NUMBER (include area code) (202) 767-3155

CONTENTS

INTRODUCTION AND OBJECTIVES	1
PAYLOAD DESIGN AND COMPONENTS	1
802.11b Payload (Relay and Router).....	2
UHF Payload (Relay)	3
Airborne System Integration	5
GROUND NODE DESIGN AND COMPONENTS	5
802.11b Ground Nodes.....	5
UHF Ground Nodes.....	7
MGEN — NETWORK LAYER PERFORMANCE.....	7
802.11b PAYLOAD TESTING.....	8
802.11b Flight 1.....	8
802.11b Flight 2.....	13
802.11b Flight 3.....	17
Applications Testing, 802.11b Payload.....	23
Summary of 802.11b Testing Results.....	24
UHF PAYLOAD TESTING.....	26
UHF Flight 1.....	27
UHF Flight 2.....	31
POWER AND TEMPERATURE PERFORMANCE	33
SUMMARY AND CONCLUSIONS	36
FUTURE WORK.....	37
ACKNOWLEDGMENTS	38
APPENDIX — IP Network Performance Testing of the Harris PRC-117F Radios.....	39

HIGH ALTITUDE RELAY AND ROUTER (HARR)

INTRODUCTION AND OBJECTIVES

In 2005, the Satellite and Wireless Networking Section (Code 5554) of the Naval Research Laboratory (NRL) was tasked with determining the feasibility of developing and using communications relays held aloft by balloons at near-space altitudes (between 20,000 and 21,300 m, or approximately 65,000 and 70,000 ft, MSL). The overall goal of this program, titled High Altitude Relay and Router (HARR), is to extend line-of-sight (LOS) connectivity between ground terminals at distances of 100 nautical miles or greater. The HARR program objectives are as follows:

- Develop a balloon-borne communications payload to accommodate a number of different capabilities.
- Develop payloads capable of doing on-board routing of network traffic between multiple ground nodes within and out of line of sight.
- Develop payloads capable of extending tactical communications in the UHF band for both data and voice communications for point-to-point links.
- Assess commercial off-the-shelf (COTS) and government off-the-shelf (GOTS) resources for developing the airborne payload and the ground nodes.

NRL has developed payloads and ground node configurations that achieve these objectives, utilizing COTS and GOTS equipment whenever possible. This report describes the airborne and ground node systems, and analyzes test data collected using unicast and multicast transport protocols in an Internet Protocol (IP)-based environment.

This effort was a three-year Office of Naval Research program to identify a short-term capability for extending LOS links over long ranges. Field testing was carried out in Roswell, New Mexico, in November 2005, and Lubbock, Texas, in June 2006, and follow-up laboratory testing was conducted at NRL in Washington, DC, in the summer and fall of 2006. The objectives for our experiments were to determine the maximum slant range (distance from ground node to airborne payload) for each payload to be operational, and to maximize the throughput for data communications at those ranges.

PAYLOAD DESIGN AND COMPONENTS

Two communications payloads were developed, each with the objective of extending the LOS range of links between communications assets already used in the Fleet and by Marine Corps ground forces. The two payloads are a processing payload with a router function that utilizes the 802.11b Wireless Local Area Network (WLAN) standard, and a surrogate Fleet Satellite (FLTSAT) UHF transponder.

802.11b Payload (Relay and Router)

Processing waveforms such as the Wideband Networking Waveform (WNW) and Subnet Relay are under development for Fleet and Marine Corps communications, but have not yet transitioned. Since these waveforms and protocols were not available for the HARR testing, a COTS shared channel network protocol was used to demonstrate the relevant concepts.

The 802.11b WLAN standard was selected for the processing payload because of its lightweight hardware and reasonable power consumption. 802.11b solutions can be pursued with readily available hardware and implemented over an already defined protocol.

Despite the convenience of the already defined protocol, some adaptation of 802.11b was required to meet the objectives of HARR. For example, the 802.11b WLAN is typically implemented in a “managed mode” where users communicate with each other through at least one Access Point (AP), usually a more powerful device in a fixed location that manages users. While managed mode is highly scalable and offers a simple configuration, it is inflexible, being more difficult to implement when both the AP and the users are mobile rather than stationary. In addition, managed mode offers no alternate path in the event that the AP malfunctions. Therefore, NRL implemented “ad hoc mode,” where each user manages itself and participates equally on the wireless network. Ad hoc mode is cheaper to implement because it requires less hardware. There is also no need for an AP, which makes it highly flexible since each user can communicate directly with its neighbor without first going through a fixed AP. This allows the use of multiple airborne relays, with as few or as many as necessary being used to complete a link. Additionally, each node can route traffic for its neighbors, thereby ensuring the most direct communication path.

One complication with using the 802.11b standard for HARR was that there is a Federal Communications Commission (FCC) limitation on the radiated power levels for the 802.11b standard. This restricts the ranges that are possible with that standard in its normal operating frequency range (~2.4 GHz). To avoid interference issues, special hardware was developed to both amplify and downconvert the 802.11b waveform to operate around 1.8 GHz, a part of the spectrum that allows government use for long-range terrestrial communications in the United States.

To enable a processing payload, an on-board computer with command software, Global Positioning System (GPS) interfaces, and control interfaces were integrated using the PC104 standard. The PC104 form factor provided for small, lightweight packaging for each of the payloads. The PC104 stack was enclosed in an aluminum case to reduce electromagnetic interference and provide strength and structure to withstand shock and vibration during landing and recovery. This payload design was integrated into a platform for flight operations (see Fig. 1) and met the requirements for a simple, cost-effective design for the airborne element of HARR. The platform was provided by a team from the Air Force Research Laboratory (AFRL), who was responsible for the launch and recovery of the balloons and payloads.

Figure 1 shows schematically the 802.11b payload, which fundamentally consists of a PC104 single board computer running NRL’s Optimized Link State Routing (OLSR) on a Linux operating system, and an amplifier/converter that radiates a 13 W signal in the vicinity of 1.8 GHz. Also shown in Fig. 1 is the alternate telemetry link that utilized an Iridium modem set to broadcast the position data for the balloon down to the command center and the various ground nodes.

The 802.11b relay and router payload was equipped with a single omnidirectional antenna hung from the bottom of the payload. The antenna provided coverage with a maximum gain of approximately 1.5 dBi. A telemetry payload from AFRL was also included.

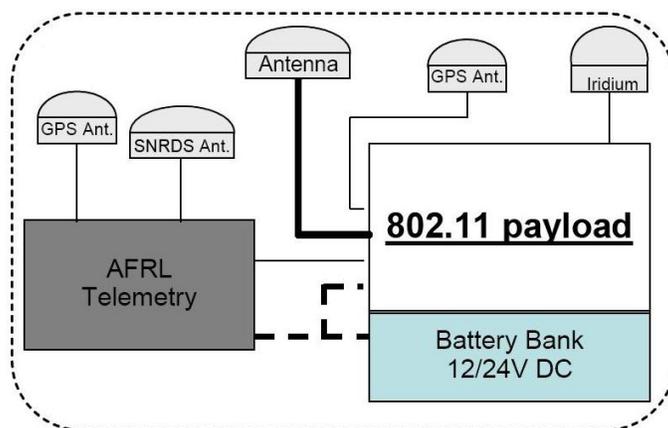


Fig. 1 — Processing payload architecture for the HARR 802.11b payload. The AFRL telemetry system provided telemetry and GPS data over UHF, while the processing payload carried an alternate path to the ground stations via the Iridium modem.

UHF Payload (Relay)

For extending tactical links, NRL designed a payload to act as a surrogate FLTSAT 25 kHz transponder. This “bent-pipe” relay operating at UHF can support many tactical radios, such as the PRC-117F, that operate in SATCOM or LOS mode over similar channels. Moreover, it is rare that forward forces are issued a FLTSAT channel for their use. Therefore by providing this capability, the ground forces gain a greater utility out of their current equipment, as well as the ability to set up extended links independently.

Both single-channel and dual-channel transponder payloads were developed from discrete components, utilizing a single, class-A, 2-W power amplifier. The single-channel version is shown in Fig. 2. The transponder utilized two independent frequency sources for its two local oscillators, which allowed for some flexibility in choosing FLTSAT bands with the uplink and downlink separated by either 41 MHz or 33.6 MHz. The oscillators (direct digital synthesizers, or DDS) were pre-programmed to support actual FLTSAT channels that were not being used by satellites in view from the western continental United States.

Two separate antennas were used for the uplink and downlink functions of the transponder. The antenna selected was a crossed loop design (referred to as an “eggbeater”) with a ground plane that provided 2.15 dBi gain, hemispherical coverage, and good voltage standing wave ratio (VSWR) performance over the frequencies of operation (250 to 320 MHz). The antennas were suspended on a 3.6-m-long yoke that was part of the payload’s rigging to the parachute (for recovery) and balloon. This yoke provided the necessary separation between the transmit and receive elements (see Fig. 3).

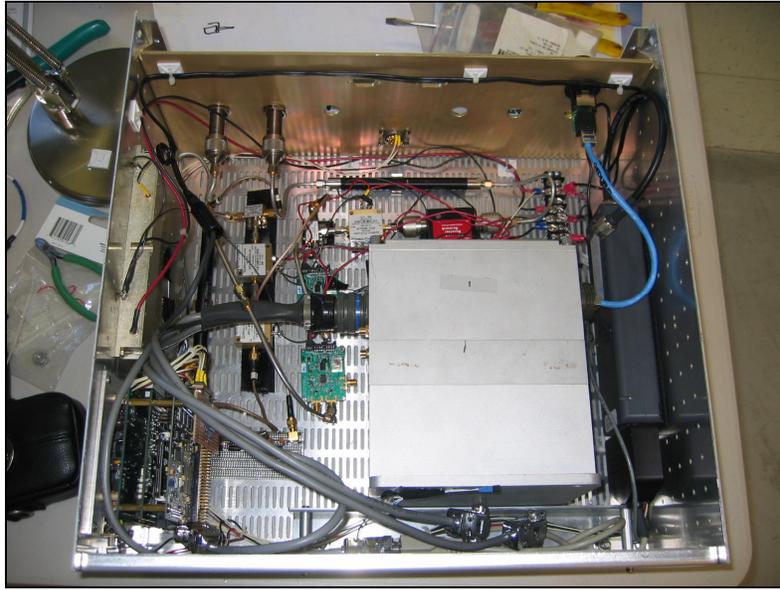


Fig. 2 — HARR UHF payload, showing layout of transponder components, along with NRL's PC104 controller stack (with aluminum enclosure), and an Iridium modem (far right). The 2-W power amplifier is mounted on the sidewall at the far left.



Fig. 3 — UHF payload integrated into the launch package on the truck-mounted launch assembly. The 3.6-m horizontal beam provides a stable separation between the two “eggbeater” UHF antennas, which are suspended at that distance to prevent coupling between the cables, overpowering of the receive antenna or AFRL's signals with the transmit antenna, and interference with the housed electronics in the payload.

Airborne System Integration

The Air Force Research Laboratory, based at Kirtland AFB in Albuquerque, New Mexico, provided the launch and recovery support and telemetry for the experimentation. AFRL developed the truck-mounted launch assembly shown in Fig. 3, which allowed for launching in winds up to 10 knots. AFRL also provided the telemetry package and release mechanism for the payload, including the Federal Aviation Administration (FAA) transponder necessary for the balloon and assembly to safely pass up and down through controlled airspace (below 60,000 ft).

To mitigate electromagnetic interference complications, a specific layout of the multitude of antennas was established, utilizing separation and blockage whenever possible. The greatest concern was interference between the UHF transponder payload and the Synthesized Netlink Radio Data System (SNRDS, a UHF receiver used for telemetry and GPS data) utilized by AFRL for their telemetry link. Filters were installed on the telemetry link, and the SNRDS antenna was suspended from the package body diametrically opposite the transmitting antenna for the payload. This was sufficient to allow simultaneous operation.

GROUND NODE DESIGN AND COMPONENTS

The payload nodes provide a router/relay or relay-only capability to dispersed ground nodes. To exercise this functionality, ground nodes were developed in conjunction with the various payloads. The 802.11b and UHF ground nodes were integrated on sport utility vehicles, with the antenna controllers mounted on the roof of each vehicle (Fig. 4).

802.11b Ground Nodes

The drawing in Fig. 5 depicts ground node configurations for the 802.11b communications testing. The payload GPS data was relayed back to the command center via telemetry from the balloon (provided by AFRL) or the Iridium link. For the 802.11b communications, Yagi antennas with a gain of approximately 13 dB were used. These antennas have beamwidths of 35 degrees and therefore require a relatively accurate means of pointing. To automate and simplify this process, single-axis antenna controllers were developed at NRL to provide a steerable azimuth motor to accurately point the antennas to the balloon.

Each antenna system was controlled by a process that calculated the heading, elevation angle, and range to the balloon. This process, run on a Linux PC, accepted manual inputs for local (ground node) heading and GPS position, as well as the GPS telemetry stream sent via the AFRL telemetry package. There was an intermediate process implemented at the ground nodes for accepting, parsing, and distributing GPS data received over the telemetry link as well as from other sources.

The 802.11b ground node required two antennas, one for transmit and one for receive, as the power amplifier/converter had separate input/output ports and could not be accommodated on a single antenna.



Fig. 4 — Antenna controllers on top of the sport utility vehicles. Each vehicle has two directional antennas, one for transmit and one for receive.

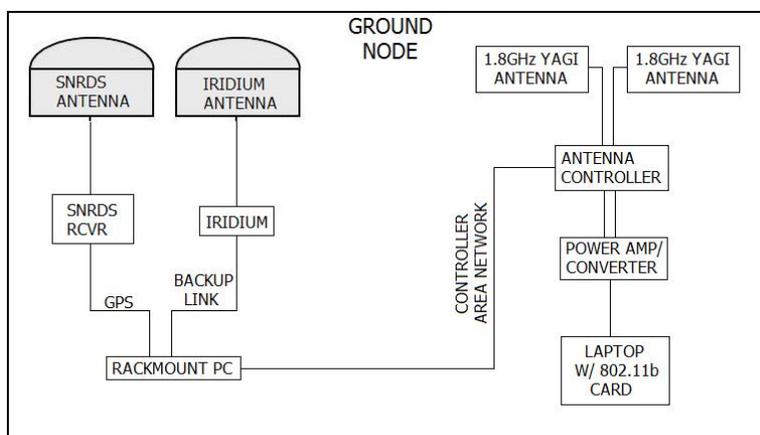


Fig. 5 — 802.11b ground node configuration. The power amplifier/converter had separate ports for transmit and receive, and so required two antennas. The SNRDS system provided the AFRL-generated telemetry from the payload.

UHF Ground Nodes

The antennas used on the UHF ground radios were either “sector” antennas with a gain of approximately 6 dB and a beamwidth of approximately 90 degrees, or omnidirectional antennas provided with the PRC-117F radios.

The PRC-117 is the primary UHF radio for the U.S. Marine Corps, and the WSC-3 is the legacy UHF transceiver for the U.S. Navy. Therefore, our ground nodes primarily carried PRC-117 transceivers for transmitting and receiving voice and IP data. Additionally, WSC-3 UHF transceivers were set up at the base station (launch site) and operated over the UHF transponder payloads. The WSC-3 transceivers used an external modem and transmitted simplex data (unidirectional) with a bit error rate test unit, or BERT. One WSC-3 was designated as a permanent transmit unit and the second as a permanent receive unit. Voice was also transmitted over the radios as a connection check before each test to verify proper operation.

The sector antennas were always pointed manually, either by using a visual observation of the balloon (which was surprisingly effective in the clear west Texas sky out to about 30 miles) or using a compass and calculating the look angle to the balloon (the NRL telemetry software package did so automatically). An omnidirectional antenna was always used first at one of the nodes in the PRC-117F testing. It was replaced by a sector antenna when communications were lost, thereby continuing the tests to greater ranges. A diagram of the end-to-end configuration (using the PRC-117Fs) is shown in Fig. 6.

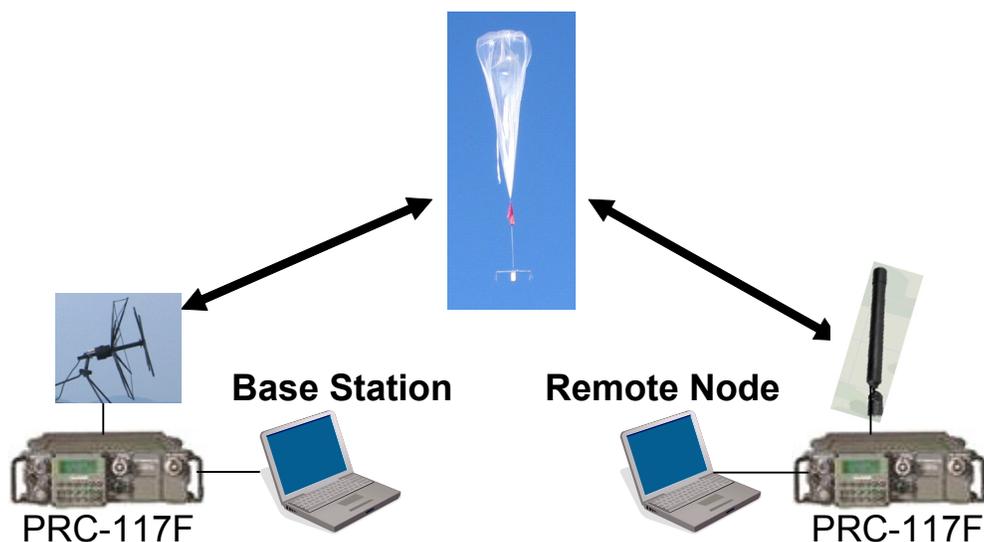


Fig. 6 — Test configuration for HARR UHF transponder relay. Frequencies were selected that required the relay to complete the RF link. The remote node operated with the omnidirectional whip antenna until communications were no longer possible. It was then replaced with a sector antenna identical to the base station antenna, and testing was continued.

MGEN — NETWORK LAYER PERFORMANCE

NRL’s Multi-Generator (MGEN) software was used to generate IP-based traffic flows for the HARR wireless network performance tests. MGEN is open source software provided by the NRL PROTOcol Engineering Advanced Networking (PROTEAN) Research Group. MGEN provides the ability to perform IP network performance tests and measurements using User Datagram Protocol (UDP)/IP traffic.

The traffic generated by MGEN was logged and used for performance analysis. MGEN input files consisted of UDP unicast and multicast traffic generation between sites. Dell laptop computers with 802.11b wireless cards were used. The laptops collected signal-to-noise ratios (SNR) for the link that were provided by the 802.11b cards. NRL's Optimized Link State Routing and Simplified Multicast Forwarding (SMF) daemons were utilized to pass multicast traffic through the relay. The use of SMF in conjunction with OLSR was required for our HARR network because in most ad hoc networks, no matter what routing protocol is used, multicast is not forwarded beyond any nodes in the local area.

To characterize link performance, the following MGEN input files were created for use between the ground nodes:

- Data rate incremental step-up: Traffic flow was slowly increased with time by increasing the packets-per-second rate, but keeping the packet size constant to 512 bytes. This was done for multicast and unicast traffic. Tests were performed unidirectionally and bidirectionally, for 5 to 10 minutes.
- Link saturation: Traffic flow at a rate of 70% of the maximum throughput was sent over periods of 30 minutes and 60 minutes. The objective was to characterize the link over a longer period of time than 5 minutes to see how the link performance changed as the balloon/relay drifted farther away from the ground nodes. The test was performed using unicast and multicast traffic unidirectionally and bidirectionally.

802.11b PAYLOAD TESTING

Field testing was conducted near Lubbock, Texas, from June 23 through June 29 to collect data to characterize the link between the ground nodes and the payloads. Also, applications such as chat, video, FTP and VoIP were attempted on the 802.11b link to prove the functionality and versatility of the systems under test.

802.11b Flight 1

The path of the first test flight (802.11b Flight 1) is shown in Fig. 7. The circles, of 50 and 100 statute miles radius, are centered on the launch site at Terry County Airport (Brownfield). One ground node (node A) was located at the airport. The second ground node (node B) was located far enough away to ensure that all traffic would flow through the payload rather than traveling directly between the ground nodes.

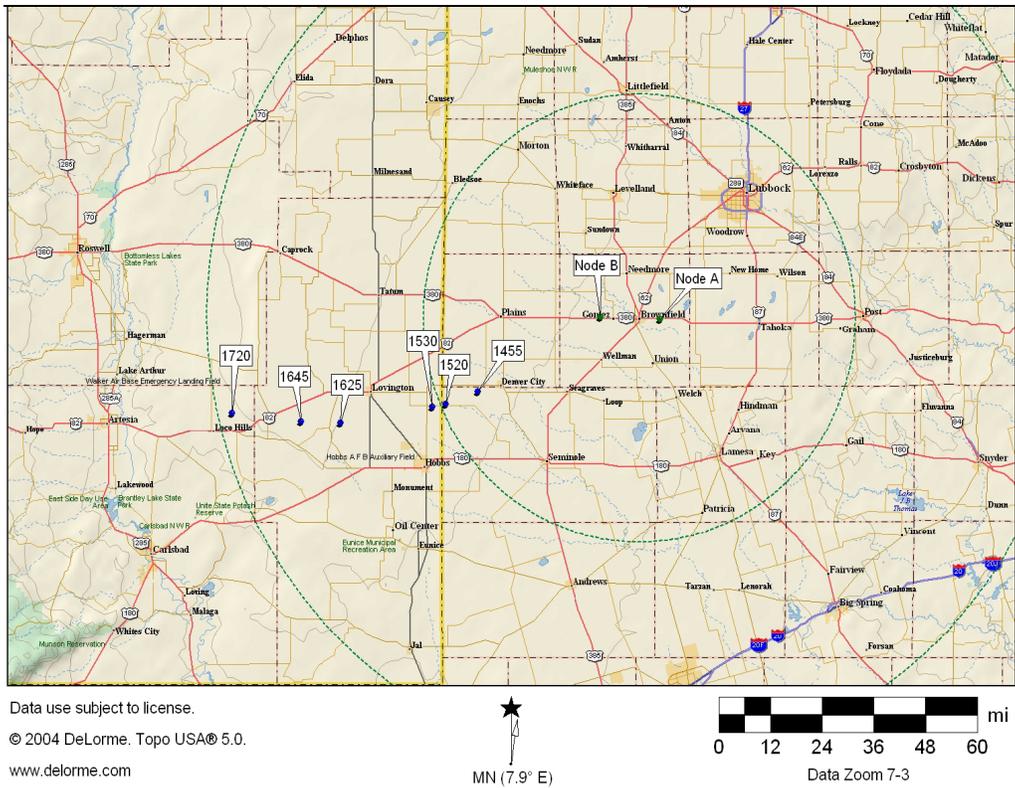


Fig. 7 — Flight path for the first launch of the 802.11b payload on June 23 (802.11b Flight 1). Ground nodes A and B are noted, and the balloon’s movement is indicated by the hour the coordinates were taken. The center point of the circles is Terry County Airport (base of operation). The smaller circle has a 50-mile radius, the larger circle has a 100-mile radius.

The first set of MGEN tests consisted solely of multicast traffic due to problems with the unicast input files, which were fixed later and tested on 802.11b Flight 2. Figures 8 through 12 show results collected at node A, node B, and the relay. The arrows indicate the direction of traffic flow and the SNR (dB) from the ground node to the relay and vice versa. The slant distance between each node and the relay is presented along with the throughput (Kbps) recorded. For unidirectional traffic flow, the maximum sustained throughput is presented, whereas for bidirectional traffic flow, the average sustained throughput is presented.

Figures 8 and 9 show a basic multicast test between the two nodes (A and B). In the first test (Fig. 8), node A started sending to node B and the maximum achieved throughput was 350 Kbps. The balloon (relay) was receiving node A at an SNR of 19 dB, and node B was receiving the balloon at 13 dB. The SNR values were generated by the 802.11b wireless cards, then collected and recorded by scripts running in the background. The distances between the ground nodes and the balloon were calculated by mapping software using the coordinates collected at the beginning of the test. In the second test (Fig. 9), node B was sending to node A and the maximum achieved throughput was 250 Kbps. The difference in throughput between the two tests is due to the difference in SNR values for the uplink (node A to the balloon in Fig. 8, node B to the balloon in Fig. 9). The uplink SNR value at the relay from the transmitting node plays a major role in determining the achieved throughput.

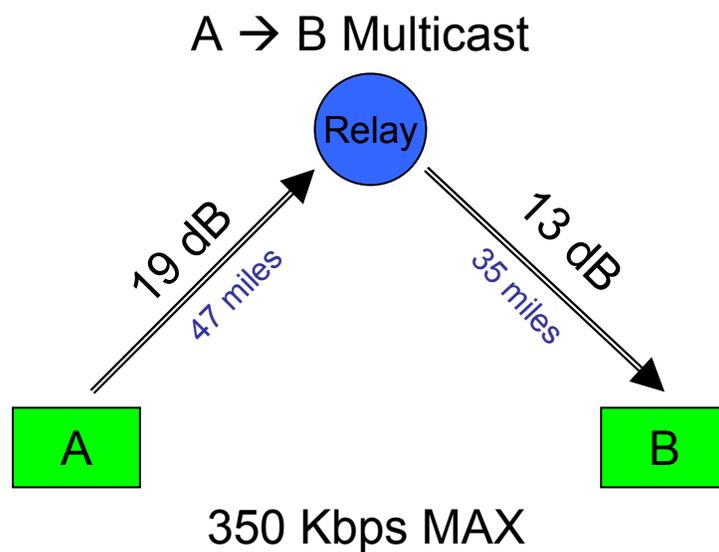


Fig. 8 — Test results, 802.11b Flight 1. Unidirectional data test of multicast traffic. Node A sent to node B through the relay in the air. Node B was receiving data from node A at a maximum rate of 350 Kbps. The SNR from node A to the relay was 19 dB, and from the relay to node B was 13 dB. The slant distances from node A to the relay and node B to the relay were 47 and 35 miles, respectively.

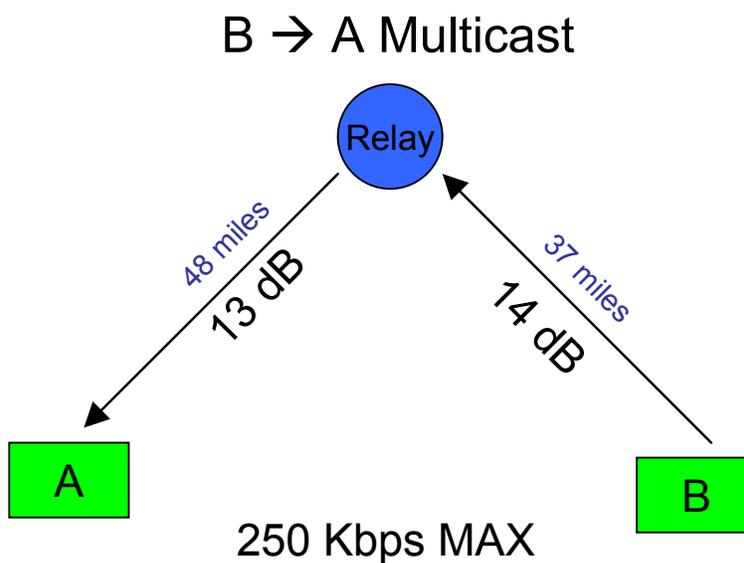


Fig. 9 — Test results, 802.11b Flight 1. Unidirectional test of multicast traffic. Node B sent to node A. Notice the data rate dropped as the SNR for the links dropped, compared to Fig. 8.

Figure 10 shows a bidirectional test of multicast traffic: the MGEN input files on each node started transmitting data at the same time to test the link in a bidirectional scenario. In this test, the SNR on the uplink for node A (17 dB) was higher than the SNR for the uplink for node B (14 dB), which is believed to be the reason why node A was able to transmit more data than node B. Average throughput is presented here rather than maximum throughput; due to frequent drops and peaks in the results, calculating the average throughput achieved is a better representation of the data than is the maximum (in bidirectional tests).

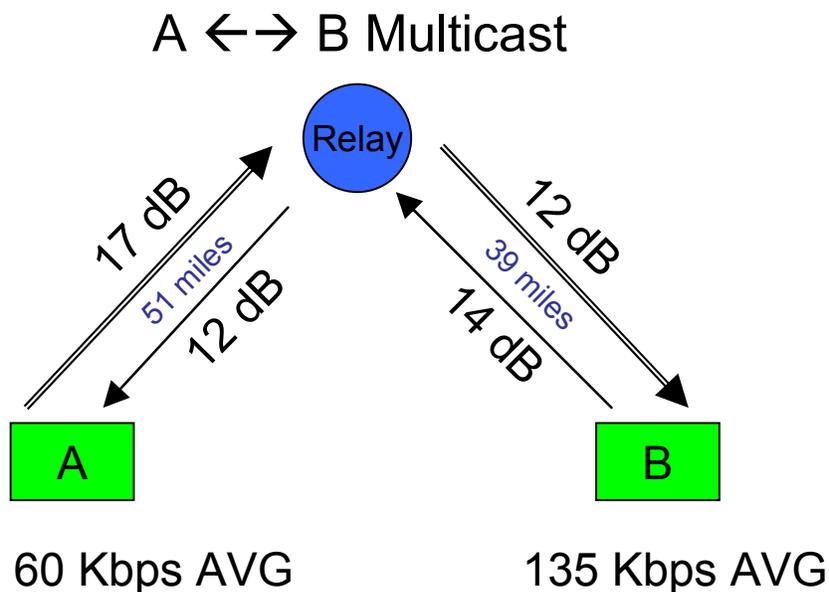


Fig. 10 — Test results, 802.11b Flight 1. Bidirectional test of multicast traffic; both nodes were sending and receiving. Node A was receiving data from node B at an average of 60 Kbps, while node B was receiving at an average of 135 Kbps.

Figures 11 and 12 show link saturation tests, in which the MGEN input files were run for about 60 minutes. The purpose of the saturation tests was to characterize the link over a longer period of time to see how the link was affected as the payload drifted away from the nodes. The MGEN input files transmitted data at a constant rate that was 70% of the maximum throughput calculated from the previous bidirectional test (Fig. 10). Once again, these tests show that SNR plays a role in the amount of traffic sent across the link.

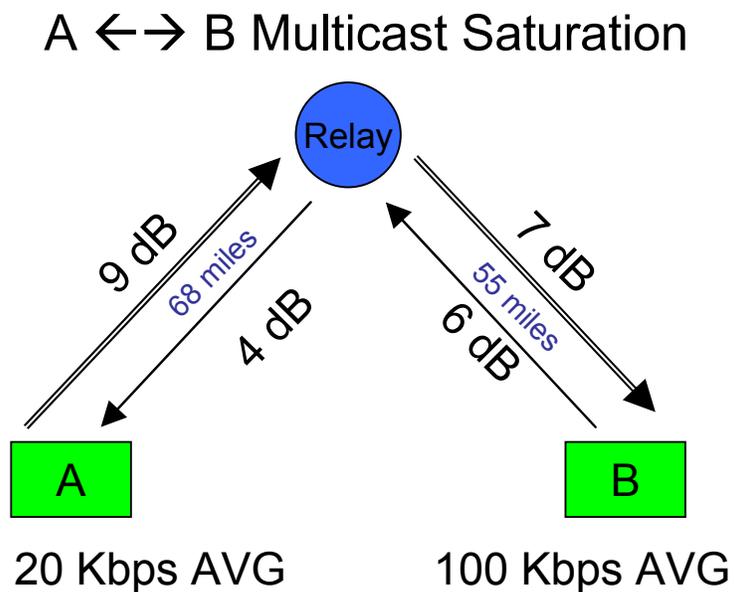


Fig. 11 — Test results, 802.11b Flight 1. Bidirectional test of multicast traffic. Both nodes were sending to each other. Saturation test input files were set to transmit 70% of the maximum throughput calculated from previous tests.

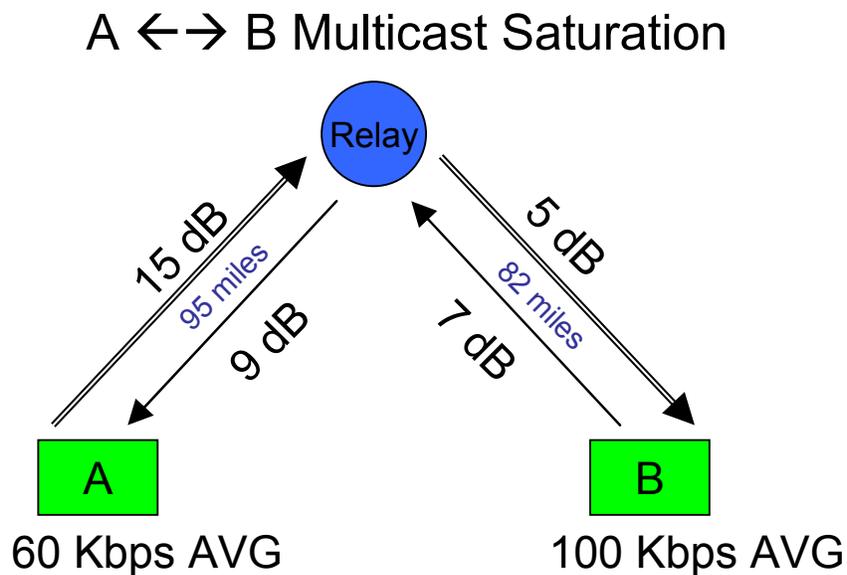


Fig. 12 — Test results, 802.11b Flight 1. Bidirectional test of multicast over longer distances than shown in Fig. 11. The low SNR in signals from/to node B are likely due to pointing errors for the ground station antennas.

802.11b Flight 2

The path of the second flight (802.11b Flight 2), performed on June 27, is shown in Fig. 13. The circles, of 50 and 100 statute miles radius, are again centered on the launch site at Terry County Airport (Brownfield). Low wind speed on this day resulted in a rather slow drift of the balloon, effectively causing the payload to circle around and hover close to the ground nodes. This resulted in higher look angles and shorter path lengths to the balloon than on Flight 1, in turn resulting in higher received SNR values at the nodes. Figures 14 through 19 show the test results from 802.11b Flight 2.

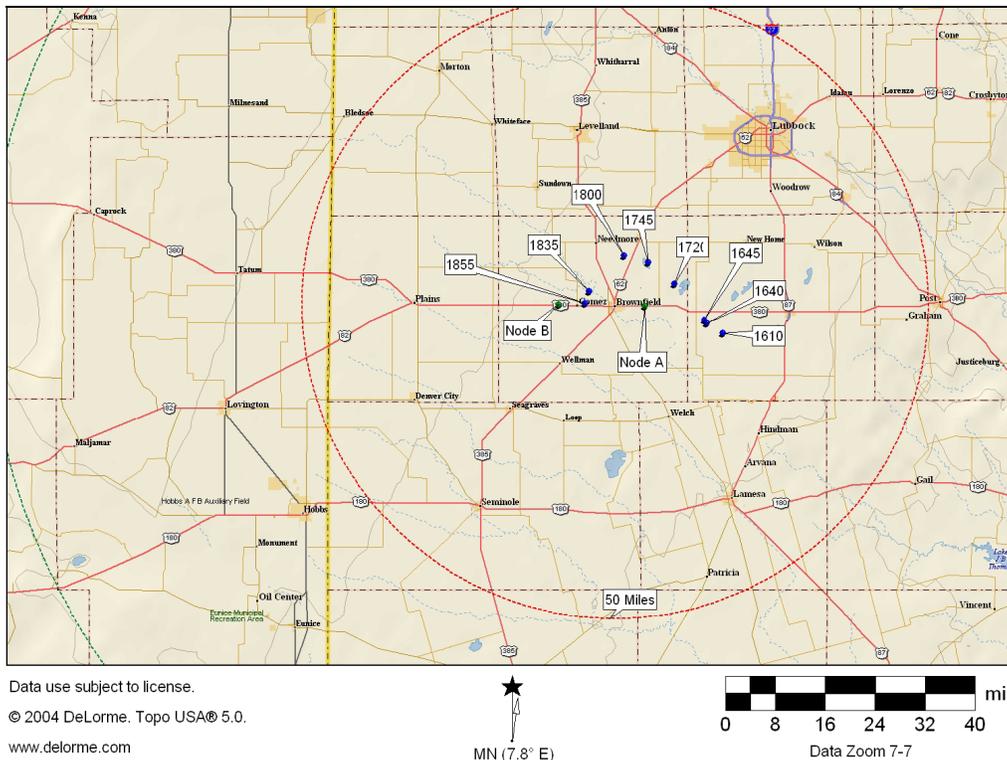


Fig. 13 — Flight path for the second launch of the 802.11b payload (802.11b Flight 2), on June 27. Ground nodes A and B are noted, and the balloon’s movement is indicated by the hour the coordinates were taken. The center point of the circles is Terry County Airport (base of operation). The smaller circle has a 50-mile radius, the larger circle has a 100-mile radius. Light winds aloft resulted in a rather slow drift of the balloon around the ground nodes.

Figures 14 and 15 show tests following the same test plan as for Flight 1, but using unicast data instead of multicast. As expected and proved from previous tests conducted in the laboratory and field, the performance of multicast is better than unicast. This is due to the lack of Acknowledgements (ACKs) in multicast traffic.

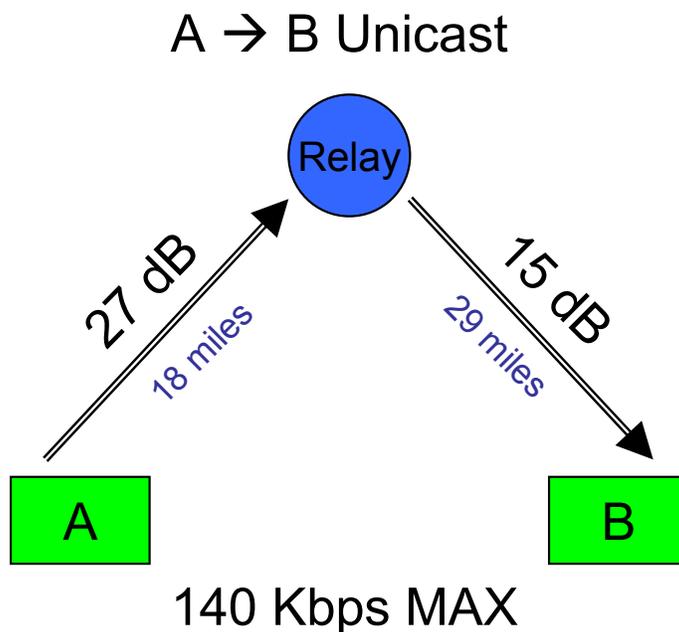


Fig. 14 — Test results, 802.11b Flight 2. Unidirectional test of unicast traffic. This illustrates how multicast performs much better than unicast (compare to 350 Kbps throughput in Fig. 8). This improvement is attributed to there being no use of Acknowledgements (ACKs) in multicast.

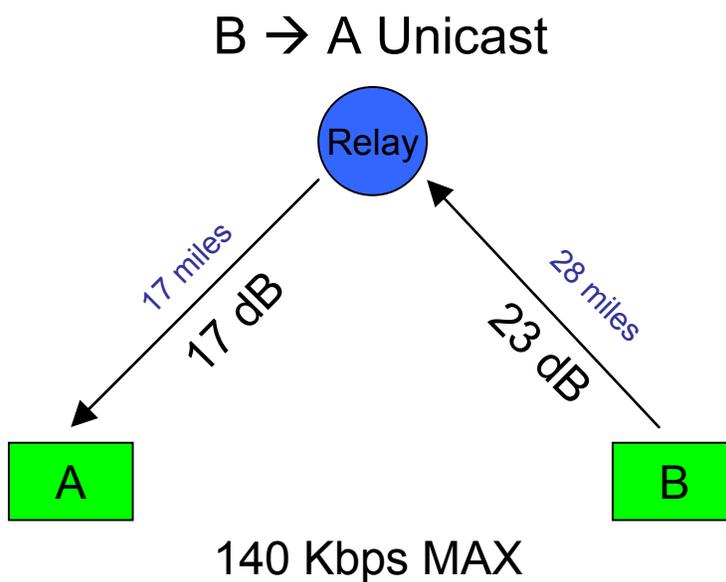


Fig. 15 — Test results, 802.11b Flight 2. Unidirectional test of unicast traffic. The performance here is poor compared to multicast (compare to 250 Kbps in Fig. 9).

Figures 16 and 17 are both bidirectional scenarios. The test presented in Fig. 16 was run for 5 minutes, flooding the link completely to determine the maximum throughput possible. The test presented in Fig. 17 (saturation test) was run for 60 minutes at a constant rate that was 70% of the maximum throughput calculated in the previous test (Fig. 16). This was done to characterize the link over a longer period of time.

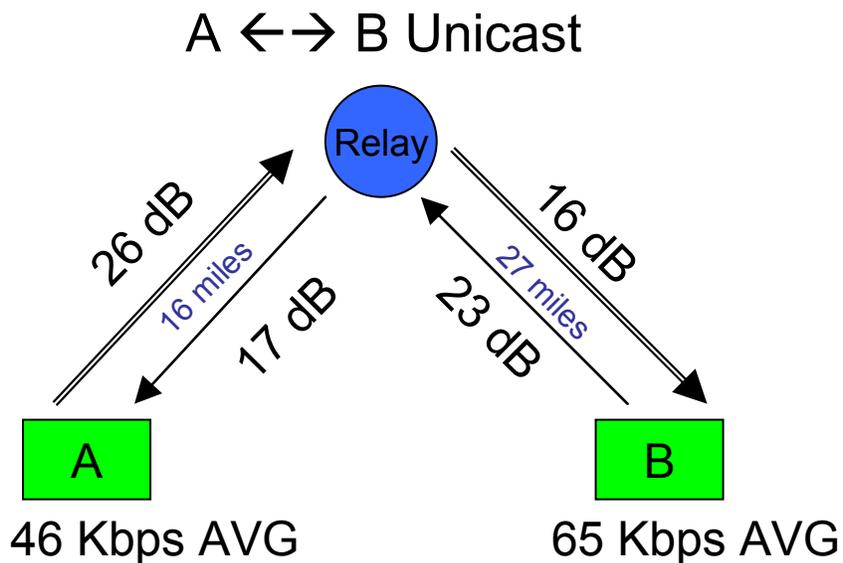


Fig. 16 — Test results, 802.11b Flight 2. Bidirectional test of unicast packets. As also seen in the multicast tests, the link with higher SNR values has a higher throughput.

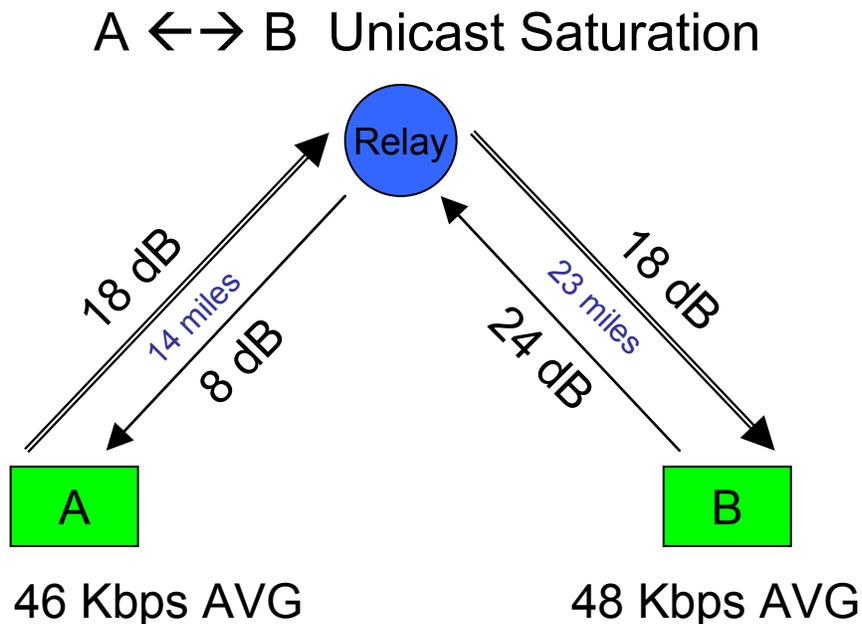


Fig. 17 — Test results, 802.11b Flight 2. Bidirectional test of unicast packets in a saturation test.

After a full suite of tests in unicast was conducted on Flight 2, the tests were repeated in multicast (as on Flight 1) under the new ranges and SNR values. Figures 18 and 19 show an improvement in performance for multicast over the results gathered from Flight 1, due to better link quality, hence less packet loss.

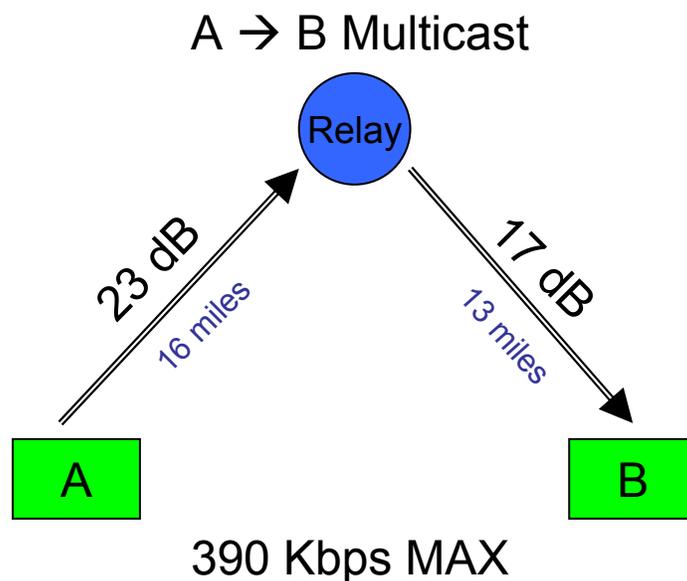


Fig. 18 — Test results, 802.11b Flight 2. Unidirectional test of multicast traffic. (Compare to Fig. 8.)

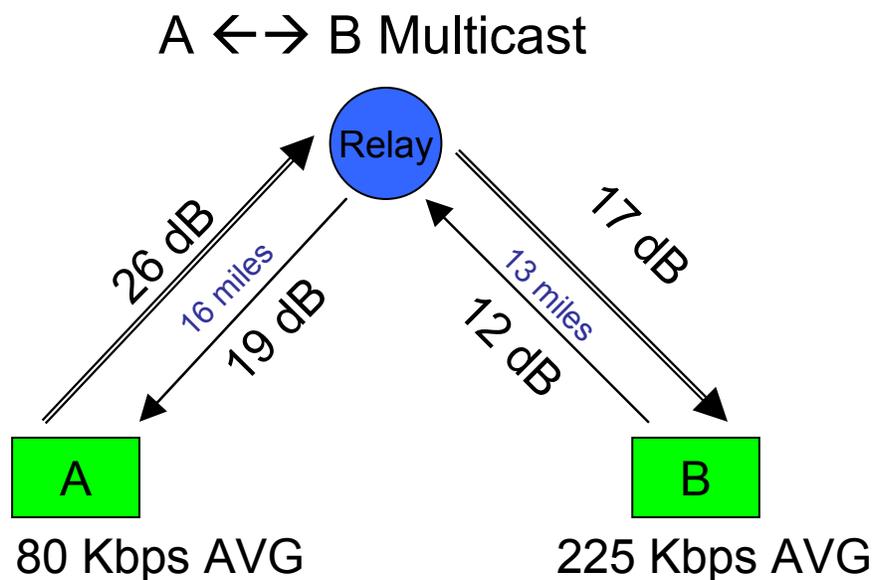


Fig. 19 — Test results, 802.11b Flight 2. Bidirectional test of multicast traffic. Higher data rates were achieved compared to the tests conducted on 802.11b Flight 1 due to the higher SNR values. (Compare to Fig. 10.)

802.11b Flight 3

The path of the final flight of the 802.11b payload (802.11b Flight 3), performed on June 29, is shown in Fig. 20. One factor affecting results during this flight was that it was difficult to accurately point the antenna in the direction of the balloon and payload, especially at longer ranges when one could not see the balloon aloft (those with better vision could see the balloon under clear skies at ranges up to 30 miles). The telemetry link from the payload was rather poor, and the high surface winds that day overwhelmed the motors used to steer the antenna positioner units, to the point where the antennas had to be secured (locked in place) frequently during the test. This resulted in a severe “lesson learned”: to identify directional antennas with much less wind loading!

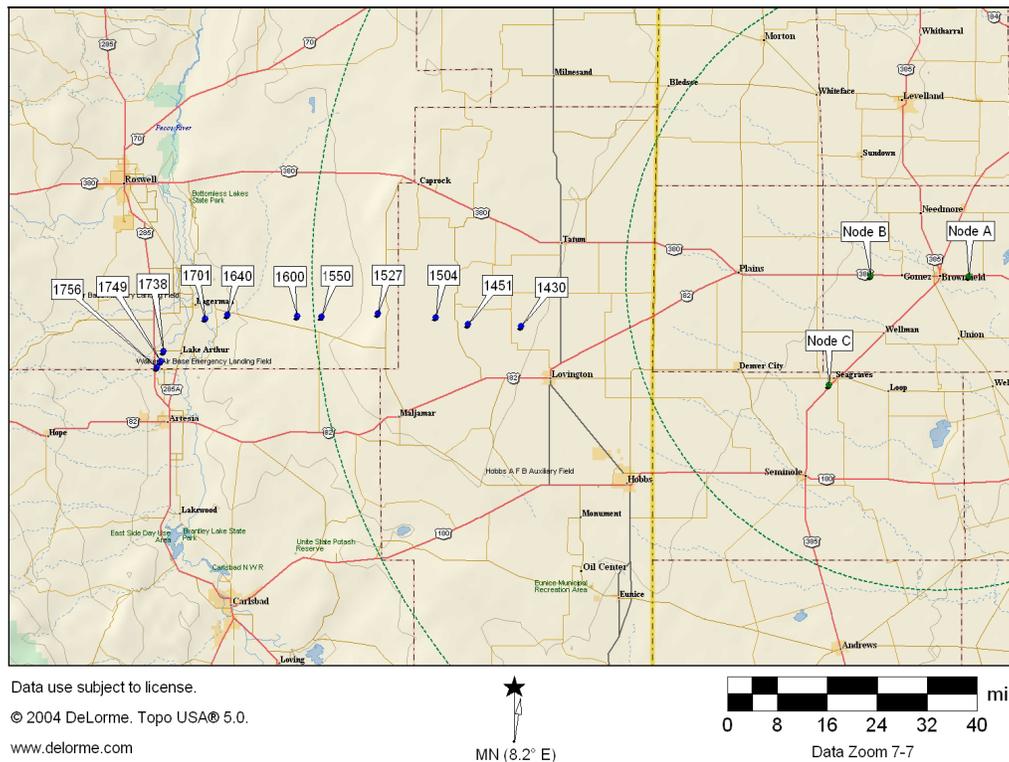


Fig. 20 — Flight path for the third launch of the 802.11b payload (802.11b Flight 3), on June 29. Ground nodes A, B, and C are noted, and the balloon’s movement is indicated by the hour the coordinates were taken. The center point of the circles is Terry County Airport (base of operation). The smaller circle has a 50-mile radius, the larger circle has a 100-mile radius.

On 802.11b Flight 3, the same unicast and multicast tests were executed as on Flights 1 and 2, but a third node (node C) was added to exercise more of the utility of the router function of the payload. The results of the tests are shown in Figs. 21 through 30. Figures 21 through 26 show tests of unicast and multicast traffic between nodes A and C. Figures 27 through 30 show tests using all three nodes.

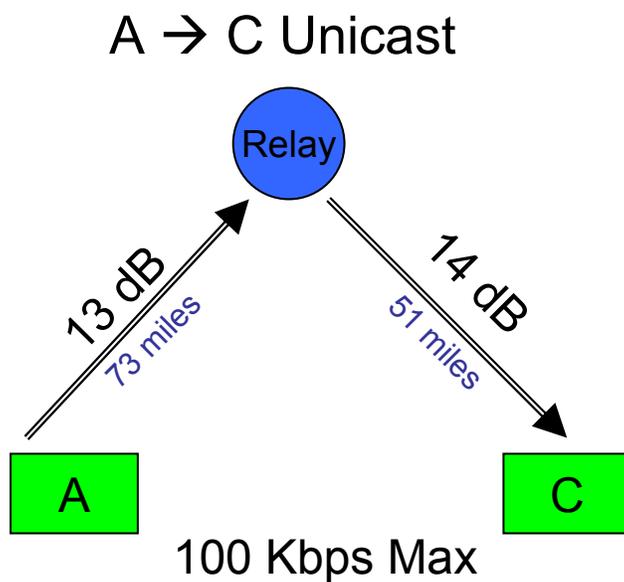


Fig. 21 — Test results, 802.11b Flight 3. Unidirectional test of unicast data from node A to node C. This result again supports the observation that lower SNR values result in lower throughput.

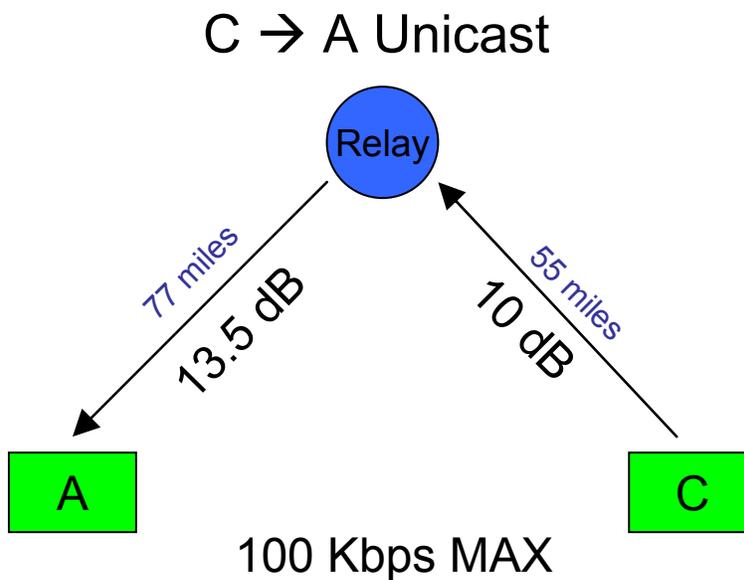


Fig. 22 — Test results, 802.11b Flight 3. Unidirectional test of unicast data from node C to node A. The performance is comparable to the traffic achieved between nodes A and C shown in Fig. 21.

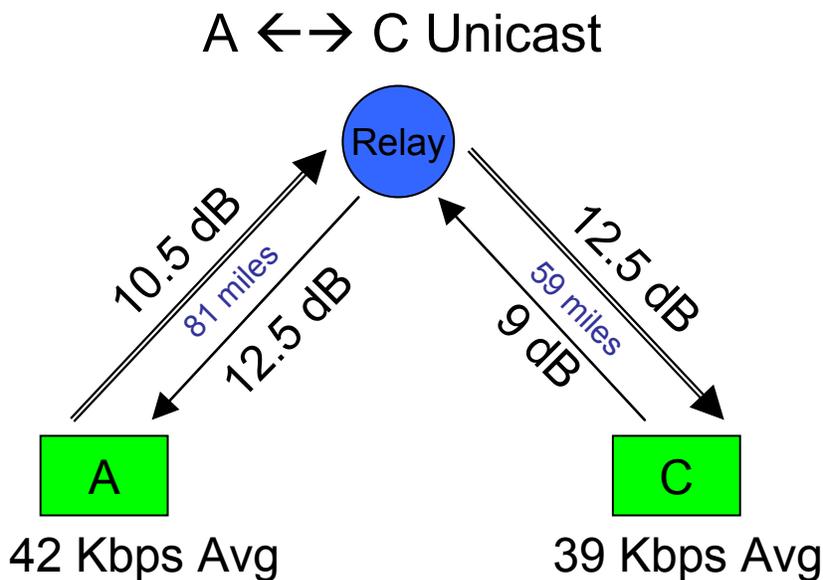


Fig. 23 — Test results, 802.11b Flight 3. Bidirectional test of unicast data.

These unicast results are similar to those from 802.11b Flight 2, showing a limited throughput. It should be noted, however, that the reduction in throughput does not appear to increase significantly at greater ranges.

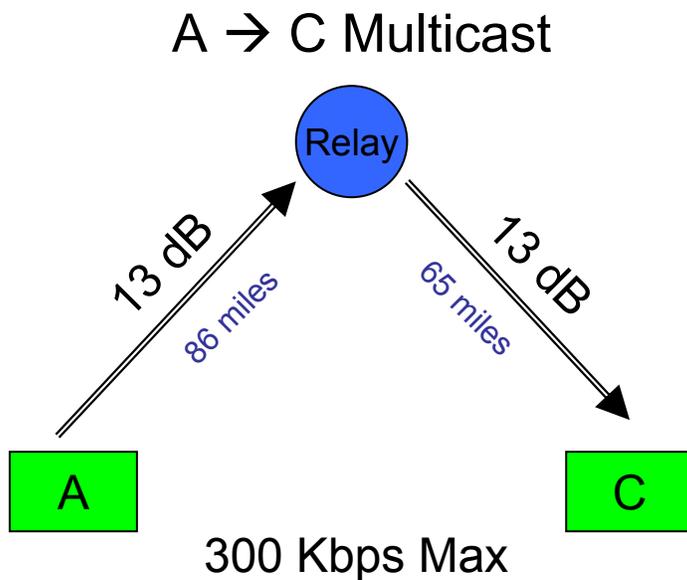


Fig. 24 — Test results, 802.11b Flight 3. Unidirectional test of multicast data.

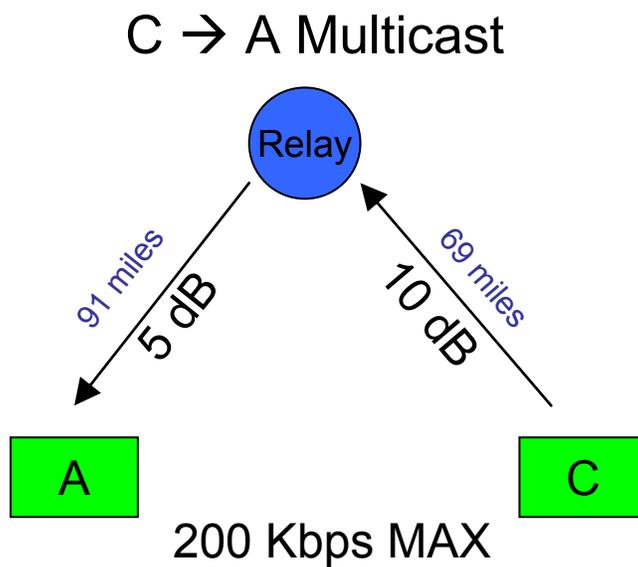


Fig. 25 — Test results, 802.11b Flight 3. Unidirectional test of multicast data. This test, like the others, demonstrates that a reduced SNR value results in lower throughput.

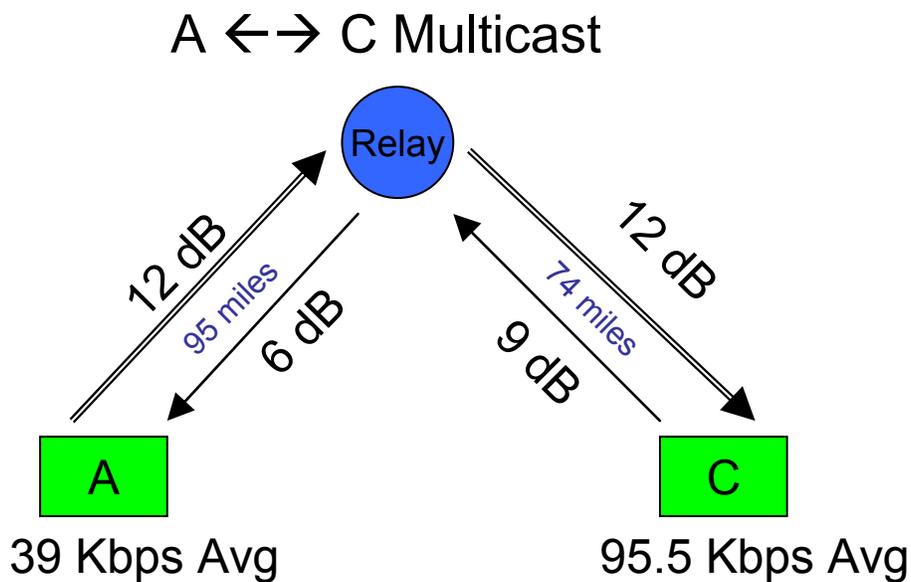


Fig. 26 — Test results, 802.11b Flight 3. Bidirectional test of multicast data. This can be compared to Fig. 19 for short-range results.

Figure 27 shows a test in which the third node was added to further stress the network and the relay. In this test, all three nodes were sending and receiving at the same time. Each node was broadcasting its multicast data while receiving data from the other nodes. It was again observed that a higher SNR value on the uplink of the sending node dominates the network. Node C had 11 dB SNR to the balloon compared to 8.5 dB and 6.5 dB for nodes A and B, respectively. Therefore, node C was able to send out more data than the other two nodes, and nodes A and B were able to receive more data from C than from each other. All results shown are average throughputs.

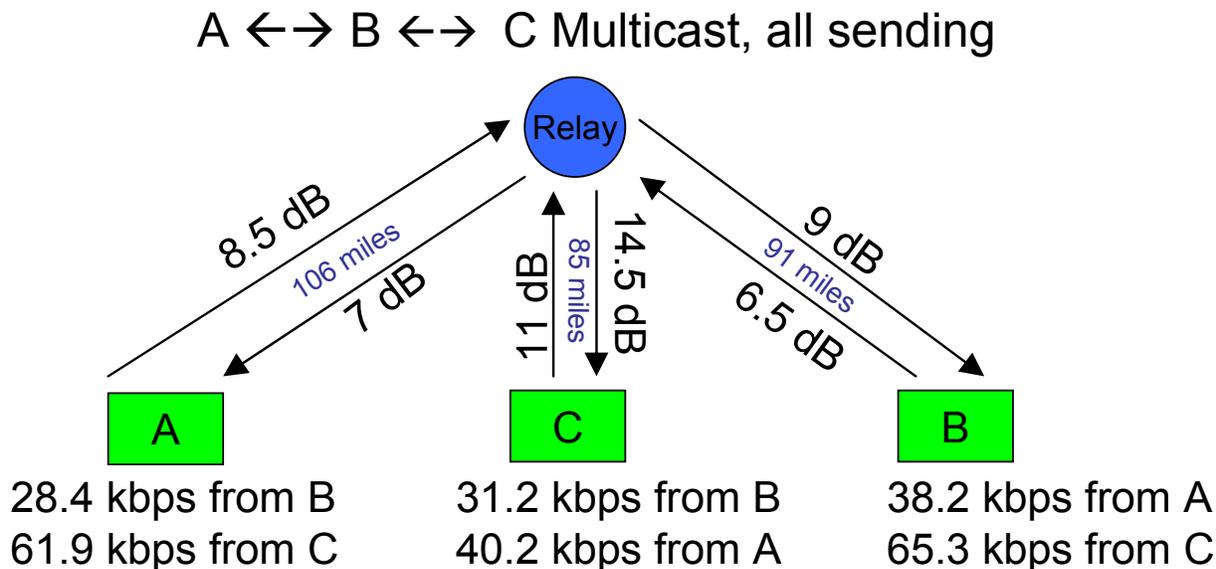


Fig. 27 — Test results, 802.11b Flight 3. Bidirectional test of multicast data for each node. A third node was added to the ad hoc network. All three nodes were sending and receiving simultaneously.

Figure 28 shows a test of unicast flow from only two senders, node A and node B. Node A had a higher SNR value than node B for its uplink; therefore, it was able to push more data through the network compared to node B. This tendency was also observed in the test with all three nodes sending unicast traffic, illustrated in Fig. 29: the throughput from node C to node A and from node A to node B was higher than that from node B to node C, with the latter having the poorest SNR at the relay. The results in Figs. 28 and 29 are averages.

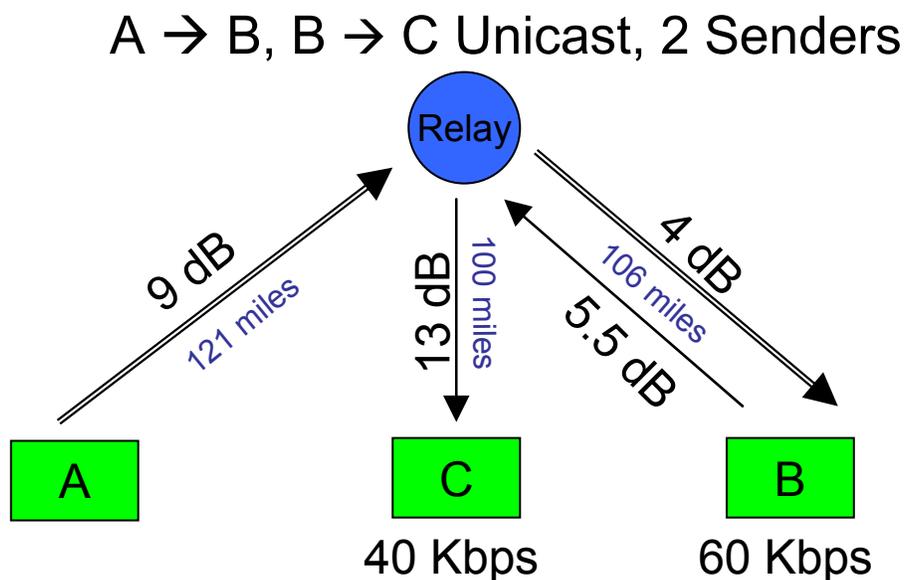


Fig. 28 — Test results, 802.11b Flight 3. Unicast test from two senders in a three-node ad hoc network.

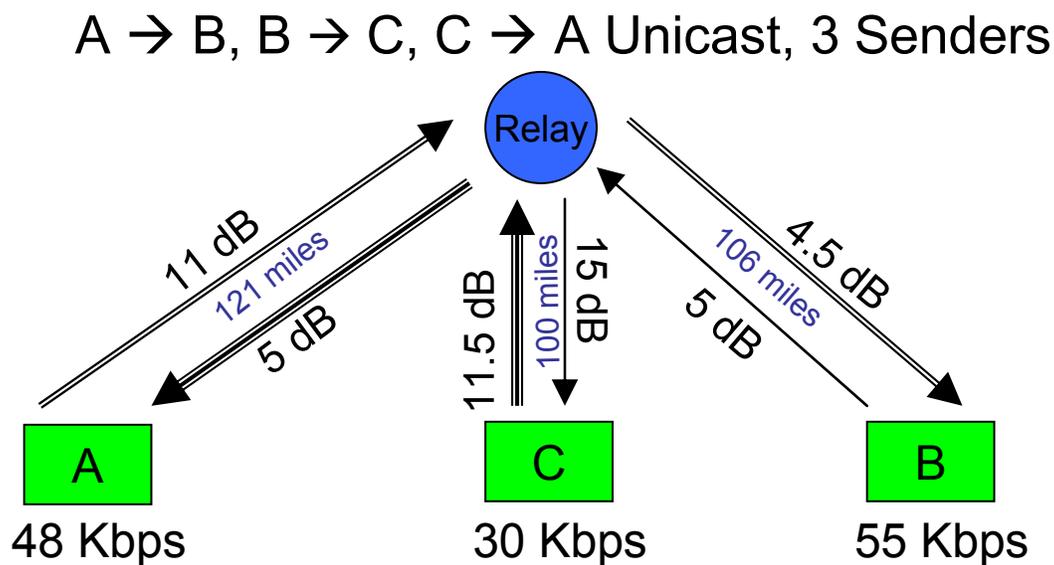


Fig. 29 — Test results, 802.11b Flight 3. Unicast traffic flow from three senders in a three-node ad hoc network over large distances. As the arrows show, A sent to B, B to C, and C to A. Again, the larger SNR from A to the relay compared to that from B to the relay resulted in a greater throughput of packets to node B than to node C.

Comparing these results with the similar unicast tests conducted earlier, shown in Figs. 16, 17, and 23, shows that the performance of the link has very little to do with range, and that adding the third node also appears to have a minimal impact on networked throughput.

Figure 30 shows a multicast test with only two senders in a three-node ad hoc network. As expected and shown in previous figures, node A could not push as much data as node B due to the smaller SNR value on the node A uplink compared to node B. The results show some increase in performance in operating with multicast traffic (compared to unicast).

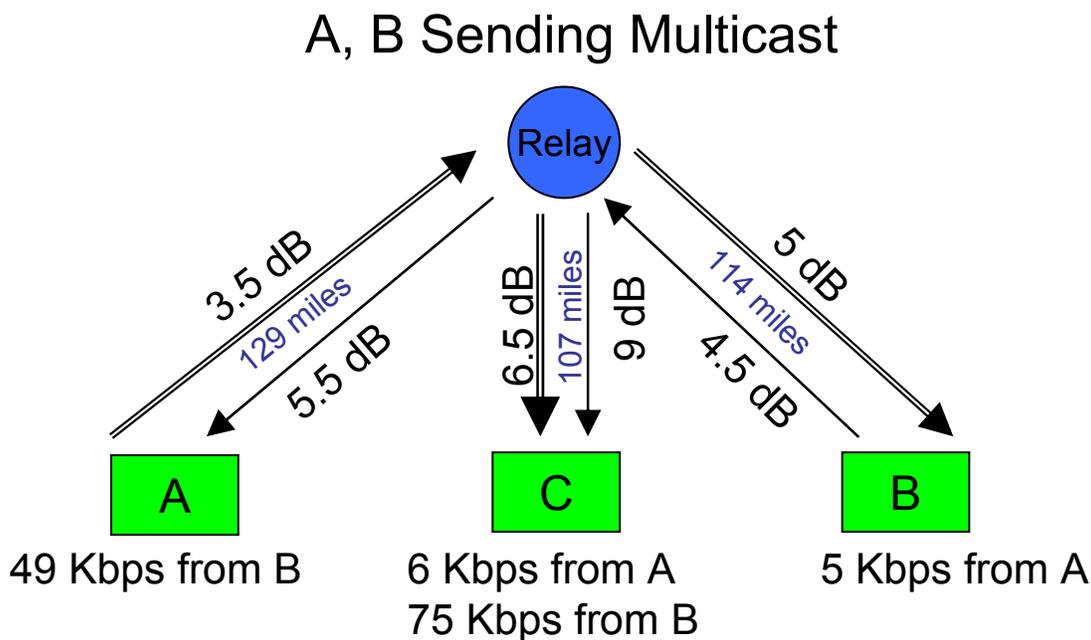


Fig. 30 — Test results, 802.11b Flight 3. Multicast traffic flow from two senders (A and B) in a three-node ad hoc network over large distances.

Applications Testing, 802.11b Payload

On Flight 3, a number of freeware applications were tested over the networks created over the HARR 802.11b router payload.

Chat Testing

A Java-based multicast chat program called Mucacha was used for testing the chat feature of the system. This program does not require a host server and can connect several nodes on an ad hoc network. However, it was realized in the field that the Time-to-Live (TTL) parameter for the packets generated by the program is set to 1. Therefore, since data was passing through a relay, those packets were dropped as soon as received by the relay. NRL attempted to increase the TTL by modifying the code; however, due to limited resources and time during field testing, the modification was not a success. In future testing, different chat software will need to be used.

Video Testing

VLC media player was used to test streaming video across the network. One node acted as the server, while the other node viewed the streaming video. Video was successfully transferred over the network. However, the codec bit rate needed to be dropped to the lowest possible option (16 Kbps). Even then, the video quality was still slightly choppy.

VoIP Testing

Figure 31 illustrates the setup used for VoIP testing. An analog phone was connected to the test laptop through a Cisco ATA 188 analog telephone adaptor. The test was successful. Phone calls were made and a conversation was held between two nodes, but the voice quality was slightly choppy. For future testing, methods to make the audio quality clearer will be investigated.



Fig. 31 — Setup diagram for VoIP testing.

Applications Testing Summary

The chat, video, and VOIP applications were successful enough to prove the link's utility in supporting such traffic. However, the applications were not vetted sufficiently prior to testing, so the results were limited.

Summary of 802.11b Testing Results

The HARR 802.11b payload test results are collated in Tables 1 and 2. Based on these results, a number of observations can be made.

- The SNR is not the sole arbiter of throughput, although it definitely has an effect, particularly when it drops below a particular threshold (approximately 7 dB).
- As demonstrated in the 2005 testing, multicast is much more effective than unicast at moving traffic through the network, even at relatively short ranges.
- Network performance for unicast does not degrade appreciably after ranges exceed 20 miles, even to greater than 100 miles, nor when the number of nodes is increased from two to three.
- The effectiveness of data transfer appears to be most dependent on the first leg of the relay, from the transmitter to the payload. This is analogous to satellite communication links through a bent-pipe satellite—the carrier-to-noise ratio set on the uplink sets the effectiveness of the link, and cannot be improved upon no matter how favorable the downlink conditions.

Table 1 — 802.11b Payload Test Results, Unicast MGEN Traffic

Test Date	Nodes		Distance		SNR		Throughput (Kbps)	
	Sending Node	Receiving Node	Sender to Balloon (miles)	Receiver to Balloon (miles)	Uplink SNR (db)	Downlink SNR (db)		
6/27/2006	A	B	18	29	27	15	140	MAX
6/27/2006	B	A	28	17	23	17	140	MAX
6/29/2006	A	C	73	51	13	14	100	MAX
6/29/2006	C	A	55	77	10	13.5	100	MAX
6/27/2006	A	B	16	27	26	16	65	AVG
6/27/2006	B	A	27	16	23	17	46	AVG
6/27/2006	A	B	14	23	18	18	48	AVG/Sat
6/27/2006	B	A	23	14	24	8	46	AVG/Sat
6/29/2006	A	C	81	59	10.5	12.5	39	AVG
6/29/2006	C	A	59	81	9	12.5	42	AVG
6/29/2006	A	B	121	106	9	4	60	AVG
6/29/2006	B	C	106	100	5.5	13	40	AVG
6/29/2006	A	B	121	106	11	4.5	55	AVG
6/29/2006	B	C	106	100	5	15	30	AVG
6/29/2006	C	A	100	121	11.5	5	48	AVG

Table 2 — 802.11b Payload Test Results, Multicast MGEN Traffic

Test Date	Nodes		Distance		SNR		Throughput (Kbps)	
	Sending Node	Receiving Node	Sender to Balloon (miles)	Receiver to Balloon (miles)	Uplink SNR (db)	Downlink SNR (db)		
6/23/2006	A	B	47	35	19	13	350	MAX
6/27/2006	A	B	16	13	23	17	390	MAX
6/29/2006	A	C	86	65	13	13	300	MAX
6/23/2006	B	A	37	48	14	13	250	MAX
6/27/2006	B	A	13	16	15	17	270	MAX
6/29/2006	C	A	69	91	10	5	200	MAX
6/23/2006	A	B	51	39	17	12	135	AVG
6/23/2006	B	A	39	51	14	12	60	AVG
6/27/2006	A	B	16	13	26	17	225	AVG
6/27/2006	B	A	13	16	12	19	80	AVG
6/29/2006	A	C	95	74	12	12	95.5	AVG
6/29/2006	C	A	74	95	9	6	39	AVG
6/23/2006	A	B	68	55	9	7	100	AVG/Sat
6/23/2006	B	A	55	68	6	4	20	AVG/Sat
6/23/2006	A	B	95	82	15	5	100	AVG/Sat
6/23/2006	B	A	82	95	7	9	60	AVG/Sat
6/29/2006	A	B, C	129	114, 107	3.5	5, 6.5	5, 6	AVG
6/29/2006	B	A, C	114	129, 107	4.5	5.5, 9	49, 75	AVG
6/29/2006	A	B, C	106	91, 85	8.5	9, 14.5	38.2, 40.2	AVG
6/29/2006	B	A, C	91	106, 85	6.5	7, 14.5	28.4, 31.2	AVG
6/29/2006	C	A, B	85	106, 91	11	7, 9	61.9, 65.3	AVG

UHF PAYLOAD TESTING

The primary means of exercising the UHF relay was via a Falcon II AN/PRC-117F Multi-band Tactical Radio by Harris Communications on each end of the link. Test procedures were enacted that tested the voice and IP data capability of the relay as it drifted away from the launch site. The radios operated in Harris's High Performance Waveform (HPW) IP mode, as shown in Fig. 6, to support MGEN testing of network performance. Also, WSC-3 radios were configured for a unidirectional high data rate link using SDM-300 satellite modems and Fireberd BERT sets. The ground nodes originated at the launch site at Terry County Airport; due to the offset frequencies for transmit and receive, the terminals could operate in close proximity to each other and still require the payload to complete the link. Frequency plans utilized unused FLTSAT channels so as to not interfere with pre-existing FLTSAT infrastructure.

Voice communications were also tested to slant ranges of 95 miles with great success, and therefore this section focuses primarily on the HPW IP mode operation of the PRC-117F radios.

The setup for each PRC-117F was identical, except that each had a unique IP address. The radios were operated in plain text (PT) mode, and set up with the configuration below. The x and y in the addresses are for the particular nodes: one base station node (172.016.000.001) and two remote nodes (172.016.000.002 and 172.016.000.003) were used in the experimentation.

COMSEC	None	
DATA/VOC		
	MODULATION	hpw
	BAUD RATE	42.6 or 64 Kbps
	ENABLE TCPIP	yes
	WIRELESS IP ADDR	172.016.000.00x
	SUBNET MASK	255.255.255.000
	GATEWAY ADDR	000.000.000.000
	ARQ	ENABLED
GENERAL		
	SELECT IP ROUTES	
	DESTINATION ADDR	172.016.000.00y
	SUBNET MASK	255.255.255.255
	NEXT NODE ADDR	010.000.000.001
	INTERFACE	wireless
PORTS		
REMOTE		
	INTERFACE	data port
	PROTOCOL	RS-232
	ASYNC RATE	9600
	DATA BITS	8
	PARITY	none
	STOP BITS	1
	FLOW CONTROL	none
	ASYNC ECHO	on
DATA		
	ASYNC	
	ASYNC RATE	9600
	DATA BITS	8
	PARITY	none
	STOP BITS	1
	FLOW CONTROL	none
	ASYNC ECHO	on
	ASYNC ENHANCED	enable
PPP		
	ENABLE PPP PORT	yes
	IP ADDRESS	010.000.000.001
	PEER IP ADDRESS	172.016.000.00x
	DATA RATE	115K

UHF Flight 1

The first UHF payload test took place on June 24 (UHF Flight 1). The PRC-117s were configured to operate at 42.6 Kbps in HPW IP mode. The track of the balloon-borne payload is shown in Fig. 32.



Fig. 32 — Flight path for the first UHF relay payload test, June 24 (UHF Flight 1). Two ground nodes were located near the launch site at Terry County Airport, and a third was flown out to Hobbs Airport, New Mexico.

One PRC-117F was flown out to Hobbs Airport, New Mexico, approximately 70 miles away, and voice links were established between the Hobbs node and the two ground nodes at the Terry County Airport launch site. Data links between the two locations were not set up due to the termination of the flight—the balloon (which had been aloft for a significant period of time) had to be brought down due to telemetry and recovery concerns.

The results of the data tests between the two nodes at Terry County Airport are presented in Figs. 33 through 37 as pairs of plots, the left plot showing the actual data throughput based on the packets received, and the right plot showing the “loss fraction,” i.e., the number of packets received successfully compared to the number of packets sent. For all of the tests run, the MGEN input files set the number of packets per second (pps) at 33, and the packet size (bytes per packet) was incrementally stepped up from 36 to 234 bytes to show the performance of the channel as the effective bit rate was approached.

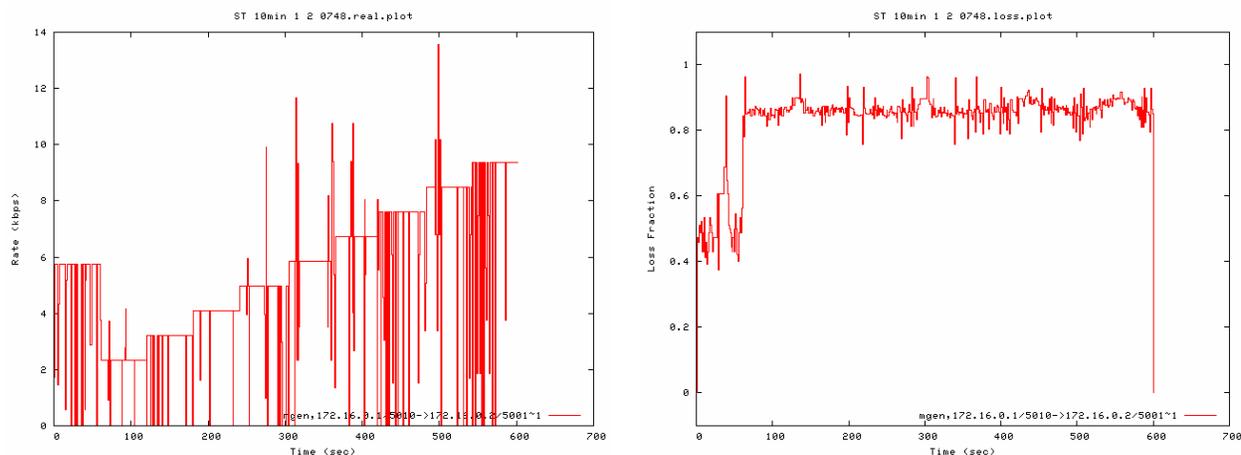


Fig. 33 — Throughput (left) and loss fractions (right) for unicast one-way packetized traffic from node 2 (remote node) to node 1 (base) over an airborne transponded UHF relay at a range of 10 miles. Packets were incrementally stepped up in size from 36 to 234 bytes; packet rate was constant at 33 pps.

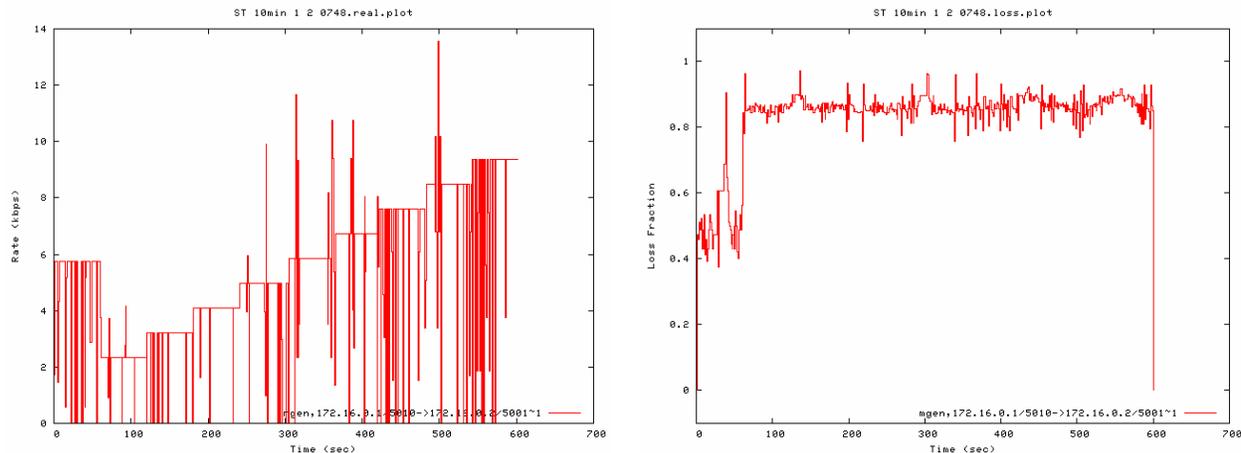


Fig. 34 — Throughput (left) and loss fractions (right) for unicast one-way traffic from the base to the remote node over an airborne transponded UHF relay at a range of 11 miles. Packets were incrementally stepped up in size from 36 to 234 bytes, with a constant rate of 33 pps.

At the short ranges indicated in Figs. 33 and 34, it can be seen that the throughput increased steadily with the increase in bytes/packet. However, the total throughput was quite poor, peaking at 9600 bps, far below the 42.6 Kbps setting for the radios.

In addition, the radio seemed to perform at a higher efficiency initially, but the throughput dropped once the packet size increased to more than 36 bytes. From there, the throughput climbed with increasing packet size, but the efficiency (loss fraction) remained constant. This occurred in every UHF data set collected in Lubbock. One theory to explain this is that the packet size is so small that the protocols for the PRC-117Fs can successfully process the traffic, even at a packet rate that is too rapid for the radios to handle efficiently.

Figure 35 shows unidirectional traffic sent over the UHF relay at a mid-range distance from balloon to node of 56 miles. There were frequent dropouts in throughput that are believed to have been caused by switching the radios from transmit to receive and back. (NRL has not been granted access to the link layer/physical layer protocols of the radios, and so cannot do more than conjecture as to the cause of some phenomena.) The loss fraction shows a similar difficulty in transferring IP packets as seen in the short-range examples in Figs. 33 and 34. It should be noted that the received signal level during these tests appeared to be more than sufficient for tests even out to 90+ miles from the nodes to the balloon: voice checks were five by five, and the indicated signal strengths on the front panels of the radios appeared to be substantial.

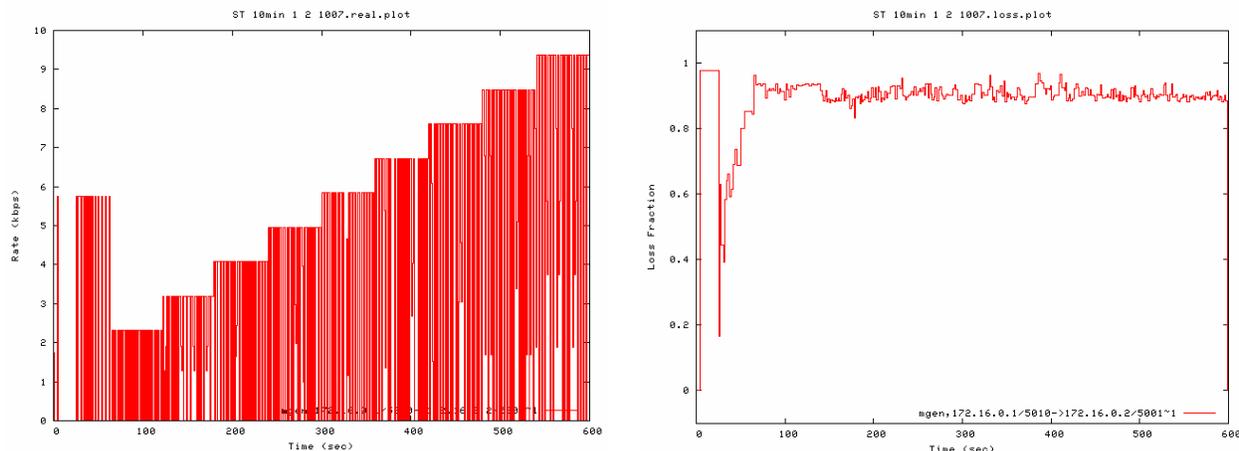


Fig. 35 — Throughput (left) and loss fractions (right) for unicast one-way traffic from the base to the remote node over an airborne transponded UHF relay at a range of 56 miles. Packets were incrementally stepped up in size from 36 to 234 bytes, with a constant rate of 33 pps.

Figure 36 shows bidirectional traffic sent simultaneously over the link at a range of 66 miles. The throughput and loss fraction data show a consistently high difficulty in successfully passing traffic. The longer, more defined dropouts are assumed to have been caused by both radios sending traffic and Acknowledgements at the same time.

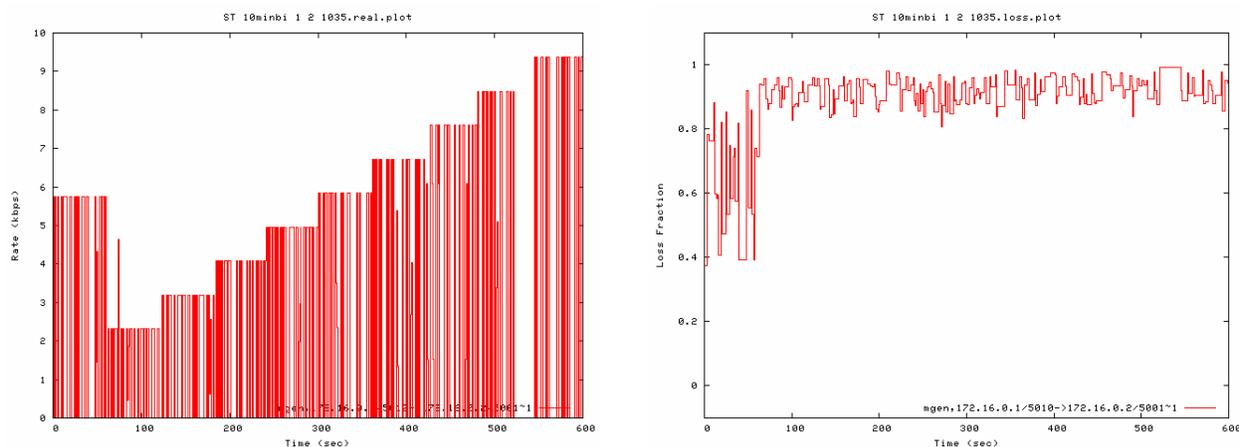


Fig. 36 — Throughput (left) and loss fractions (right) for unicast bidirectional traffic received at the remote node (using an omnidirectional whip antenna) over an airborne transponded UHF relay at a range of 66 miles. Packets were incrementally stepped up in size from 36 to 234 bytes, with the packet rate held constant at 33 pps.

In one interesting test, one of the nodes operated in a mobile capacity using the Harris omnidirectional antenna. Even more notable was the configuration: the radio and antenna were located in the passenger seat of the vehicle during the test, as no roof mounts were available, and so the antenna was in a rather compromised location on the vehicle. Despite this, the test results (Fig. 37) show similar performance to having the antenna clear of the chassis of the vehicle. There were occasional complete

dropouts in data that cannot be correlated with any particular events, but are believed to have been caused by blockage due to passing trucks or roadside buildings. The loss fraction increases to near 100%, presumably due to the bidirectional traffic, the extending range between the relay and the nodes, and the IP packet limitations of the PRC-117.

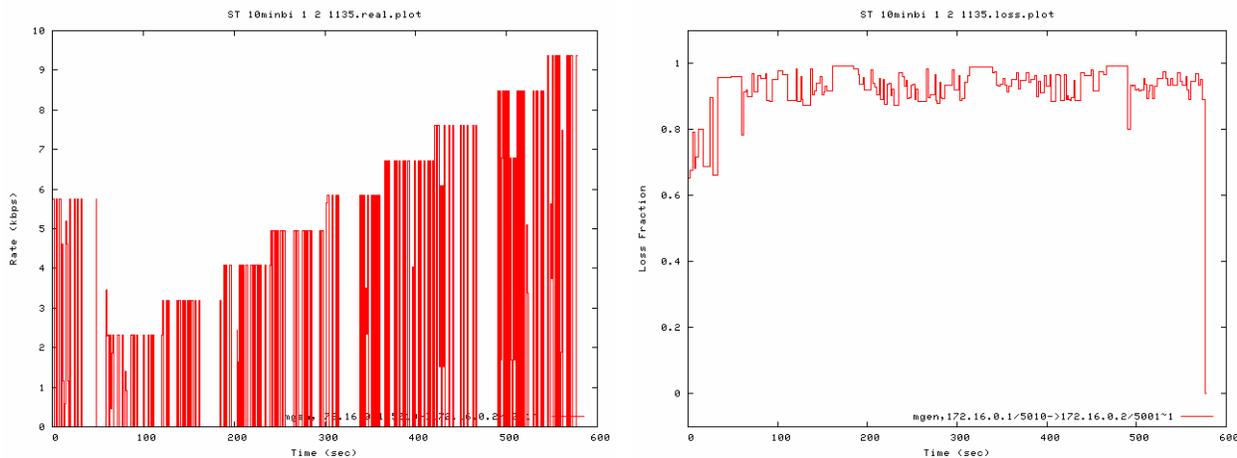


Fig. 37 — Throughput (left) and loss fractions (right) for unicast bidirectional traffic received at a mobile remote node (using an omnidirectional whip antenna) over an airborne transponded UHF relay at a range of 86 miles. Packets were incrementally stepped up in size from 36 to 234 bytes, with the packet rate held constant at 33 pps.

UHF Flight 2

For the second UHF flight, on June 28 (UHF Flight 2), a dual-channel analog transponder was tested. One channel had a narrowband 20 kHz bandpass filter and the other channel contained a wideband 100 kHz bandpass filter. For all tests during this flight, the PRC-117F tests (voice and data) were done over the 100 kHz channel. The narrowband channel carried a simplex set of data generated from a satellite modem at 19.2 Kbps QPSK (3/4 rate encoding). The PRC-117F terminals were now configured to operate at 64 Kbps, the maximum throughput rating on the terminal. Over the wideband channel, the MGEN input files run previously were replicated.

The simplex data went from one WSC-3 to another, both located at the launch site. Both utilized the directional sector antennas, each with a gain of approximately 6 dB. The receiving WSC-3 was outfitted with a single 20 dB gain low noise amplifier (LNA) to improve that transceiver's poor noise figure, but that was still found to have insufficient gain: at extended ranges, the SNR of the waveform relayed back to the receiving WSC-3 was more than sufficient, but the absolute signal level was too low for the WSC-3 to adequately process.

Both channels in the analog transponder shared the LNA and power amplifier, reducing the overall gain of the relay by 7 dB. Signals were continuously run through the narrowband channel to ascertain whether the collocation of two channels in the same payload would degrade the relay's performance. As with the tests done on UHF Flight 1, MGEN input files were created that would steadily increase the generated throughput by incrementally stepping up the packet size from 36 to 234 bytes and holding the packet rate constant at 33 pps. All data was collected at the "remote" node (often no more distant than a mile from the base node) utilizing an omnidirectional antenna on its PRC-117F terminal. The transmitting

PRC-117F continued to use the sector antenna, manually pointed toward the balloon. Initially, the PRCs were operating at 2 W output power; the signal level appeared to be quite sufficient for data.

Figure 38 shows the real packet throughput and packet loss performance over the wideband 100 kHz channel at a range of approximately 5 miles from the launch site. It should be noted that the loss fraction is quite high.

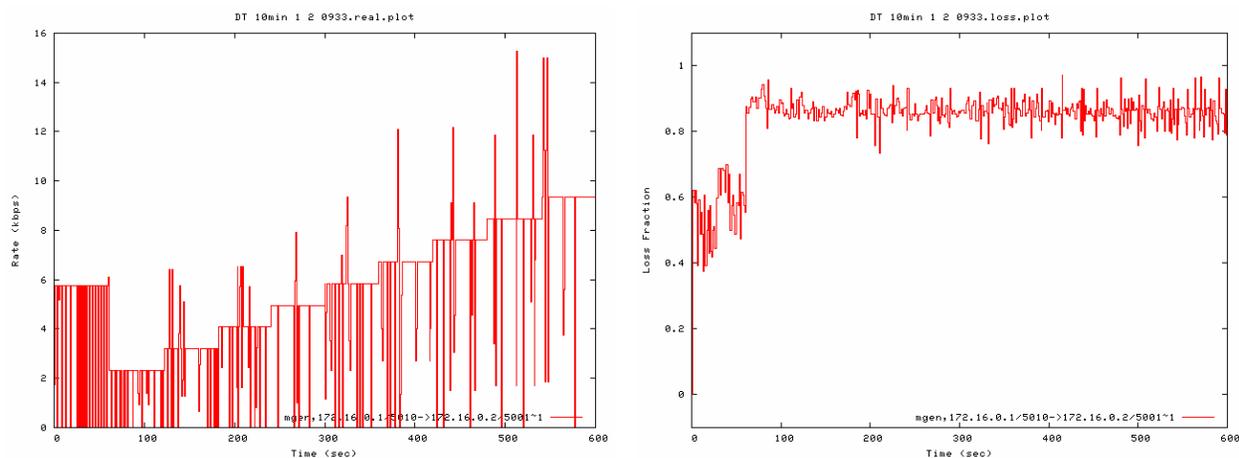


Fig. 38 — Throughput (left) and loss fractions (right) for unicast bidirectional traffic received at the remote node over a transponded UHF relay at a range of 5 miles.

Figure 39 shows the performance over the same wideband channel at a range of approximately 10 miles from the launch site. The complete dropouts in packet throughput have become more pronounced.

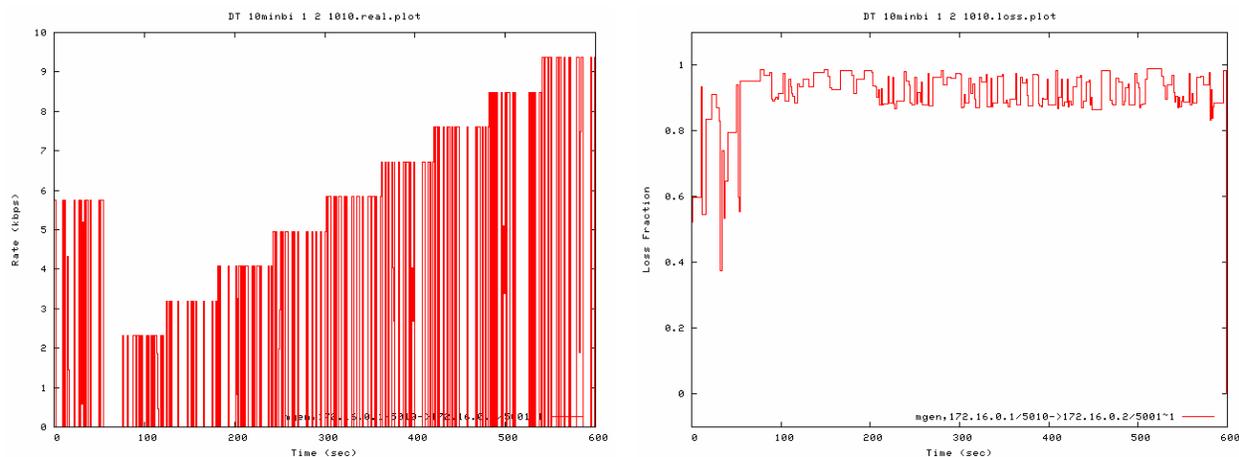


Fig. 39 — Throughput (left) and loss fractions (right) for unicast bidirectional traffic received at the remote node over a transponded UHF relay at a range of 10 miles.

On the narrowband channel of the UHF relay, the WSC-3 terminals passed data over a simplex 19.2 Kbps link with Reed-Soloman encoding. At ranges less than 10 miles, no errors were reported. As the range extended to approximately 14 miles, the Fireberd performance degraded to a bit error rate (BER) of 1×10^{-3} for transmitting 78229 blocks (approximately 120 minutes). It was observed that the energy per bit/noise (E_b/N_o) displayed on the receive modem (an SDM-300) was 5 to 6 dB when the local PRC-117F was communicating (a marginal level of performance, even with Reed-Soloman), but 16 dB E_b/N_o when PRC-117s were silent. This issue was discovered to be a collocation problem at the launch site rather than an issue with the transponders.

At a range of 35.5 miles, the performance degraded severely over even the previously low levels achieved. The PRC-117 power settings were increased to 5 W and the omnidirectional antenna on the remote node was switched to a sector antenna. The WSC-3 performance was relatively steady at 1×10^{-3} BER.

The WSC-3s were then shut down due to interference with the AFRL telemetry link which was also operated at UHF. At 1513 hours, at a range of 47.5 miles, the MGEN results shown in Fig. 40 were obtained, again over the wideband channel. This 5-minute MGEN input file fixed both the packet rate (still at 33 pps) and the packet size at constant values. It is interesting to note the steady increase in the loss fraction with time, as if the backlog was growing faster than whatever buffers are in use in this version of the HPW firmware.

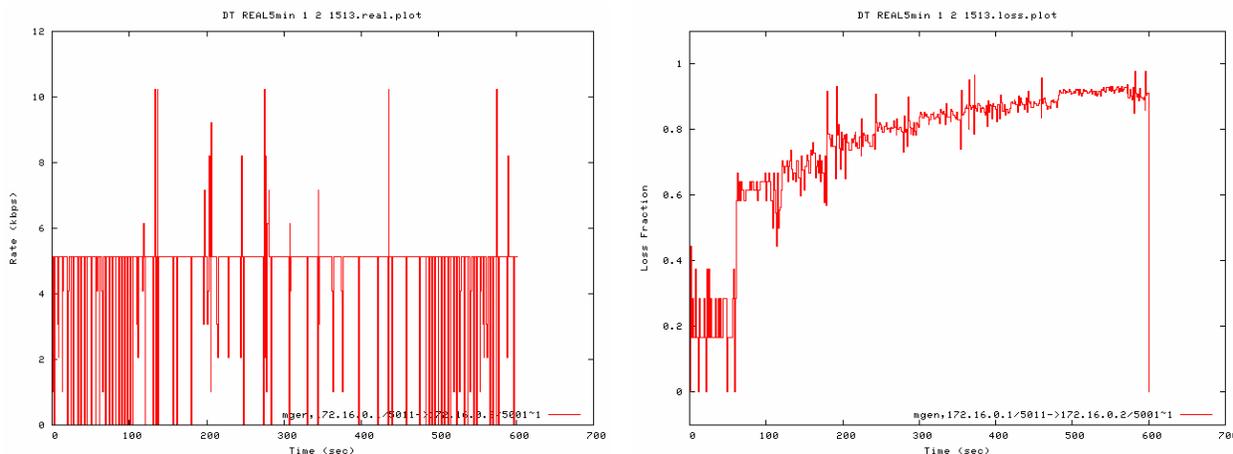


Fig. 40 — Throughput (left) and loss fractions (right) for unicast bidirectional traffic received at the remote node over a transponded UHF relay at a range of 47.5 miles. Input files were run with a fixed effective bit rate of 64 Kbps. With the PRC-117 waveform the sole user of the HARR relay power amplifier, the results improved.

POWER AND TEMPERATURE PERFORMANCE

Operating the payloads at altitudes exceeding 20 km exposes the electronics to ambient temperatures of -40 °F or lower. The low air density at that altitude results in very low convection in extracting heat from a source. Low temperatures also affect the performance of batteries, whose output capacity is diminished as the temperature drops to those levels.

Concern for those aspects of high altitude operations was paramount in the packaging and assembly of the overall payload. The payload electronics were insulated from the outside environment using solid

insulating sheets cut to size. Heat generated from the equipment, particularly from the amplifier/converter in the 802.11b payload and the power amplifier in the UHF payload, was contained within to keep all of the electronics in temperatures warm enough for normal operation. The battery rack (seven 2590 rechargeable batteries with foam insulation to hold them in place) was located above the payload rack to take advantage of any heat radiating off the equipment.

Both sets of payloads were outfitted with temperature and voltage sensors so that the insulation and hardware layouts could be assessed. One set of sensors was lost in a crash of one of the UHF payloads, but two other data sets were collected. The first data set, shown in Fig. 41, was taken during 802.11b Flight 1. It shows temperature readings from four sensors. The ambient reading indicates the temperature of the air inside the payload rack (standard 19" rack), and the case reading refers to the temperature of the metal chassis of that same rack.

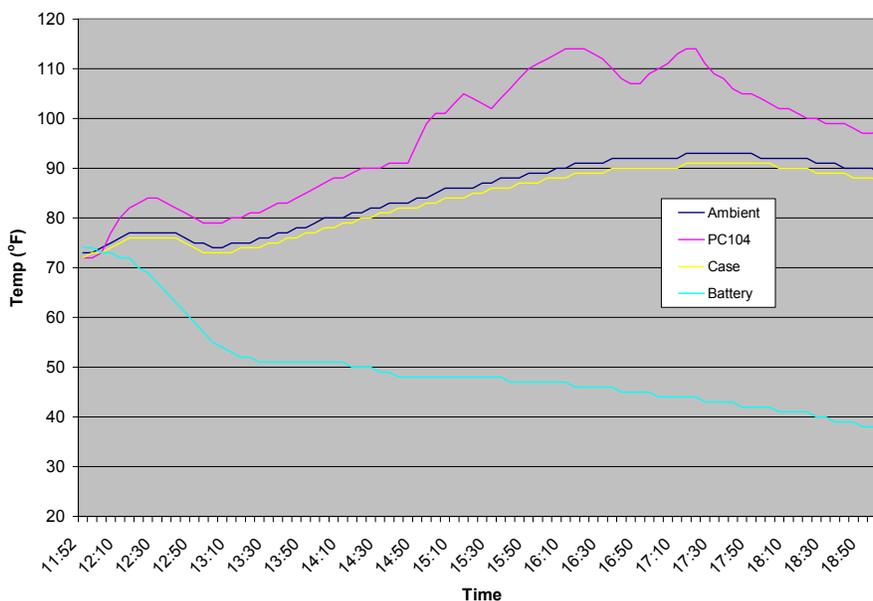


Fig. 41 — Temperature readings for the payload and battery racks, 802.11b Flight 1, June 23. The ambient reading refers to the temperature of the air inside the payload rack. Outside air temperature at this altitude is estimated at -40 °F. Time is local (MDT).

The payload functioned well during the entire flight and the payload temperatures remained within their operating temperatures over a seven-hour period, so it appears that the layout and insulation were effective.

The battery voltage for the same flight is shown in Fig. 42, reflecting the performance of the batteries over time. The figure shows a small drop in voltage, which is believed to be due to the draw from the electronics, since the temperature of the batteries remained above freezing throughout the flight. Even with the large current draw of the PC104 chassis (~ 4 A), the voltage did not drop significantly over a seven-hour period.

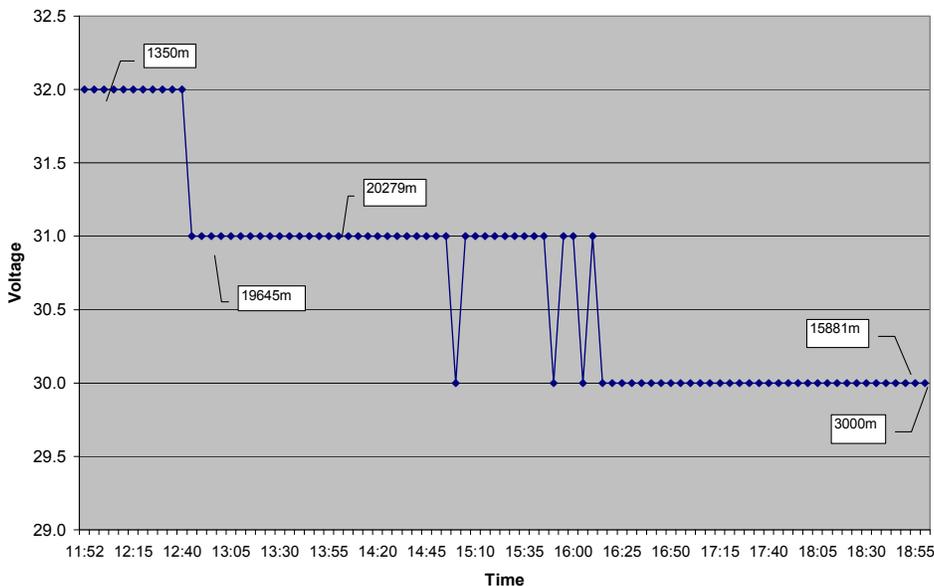


Fig. 42 — Voltage readings for the battery rack output, 802.11b Flight 1, June 23. The sensor resolution is in 1 V increments. Altitudes of the payload during the flight are indicated on the plot. The payload cruised at approximately 20,000 m throughout the flight. Data was also collected during ascent and descent. Time is local (MDT).

A second set of temperature data was collected on 802.11b Flight 3; see Fig. 43. This time, one of the sensors was moved to the amplifier/converter, which is the unit with the second largest current draw in the payload (after the PC104 assembly). There is an unusual dip in temperature around 1820 hours when the balloon was still at an operational altitude (the balloon did not begin its descent until after 2200). This is followed by a sudden and large spike in battery temperature. Although it did not affect the performance of the payload, it is some cause for concern. No comparable voltage output data was successfully collected for this flight, so only the fact that the payload continued to operate throughout the flight gives an indication that the temperature spike was not severe enough to cause a problem.

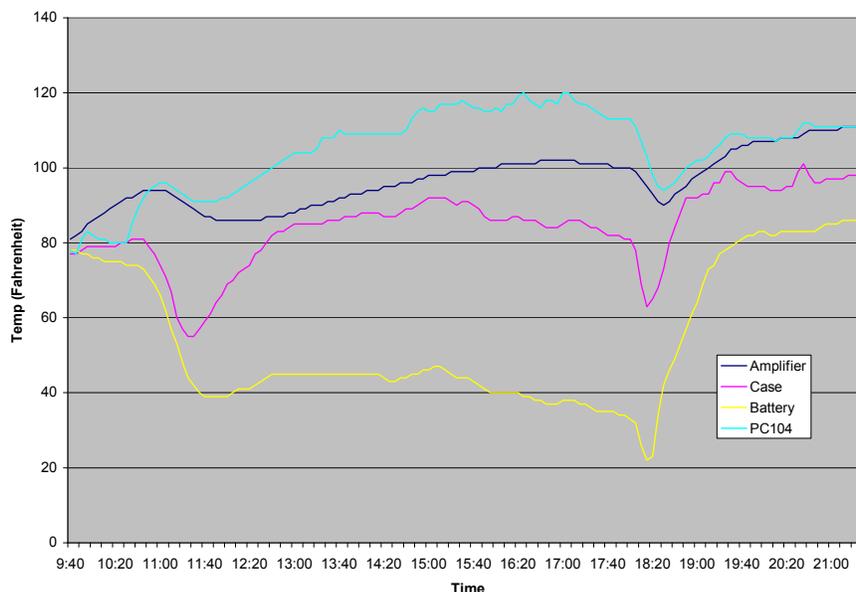


Fig. 43 — Temperature readings for the payload and battery racks, 802.11b Flight 3, June 29. Outside air temperature at this altitude is estimated at -40°F . Time is local (MDT).

SUMMARY AND CONCLUSIONS

NRL's Satellite and Wireless Networking Section has successfully demonstrated two functional line-of-sight communication relay payloads in its High Altitude Relay and Router (HARR) program for extending terrestrial line-of-sight communications over hundreds of miles. These two relays, one an 802.11b router, the other a transponded UHF (surrogate FLTSAT) bent-pipe relay, were borne aloft on a free-floating balloon to altitudes between 65,000 and 70,000 ft for operational periods up to 8 hours. The payloads supported line-of-sight communications from the launch point outside of Lubbock, Texas, via balloon to one or two other nodes in the operational area.

This technology provides the capability for Fleet/land forces to establish and maintain local area network connectivity over extremely long ranges. These payloads can be either recoverable or disposable depending on the scenario, and although demonstrated in a free-floating, zero-pressure (stable altitude) balloon, it can also be deployed in station-kept airships (manned or unmanned) or in unmanned aerial vehicles (UAVs). This technology provides a link extension capability that can be locally controlled and maintained, as opposed to satellite channels, which are usually unattainable and require significant advance planning and process.

In short, the High Altitude Relay and Router can provide in-theater forces a locally controlled means of significantly extending their ranges for line-of-sight communications between themselves and coalition partners with packages that are easily launched and, if necessary, disposable.

Testing in 2005 on this program emphasized ground distance between the nodes, while this 2006 exercise focused on performance with respect to slant range from the individual nodes to the balloon-borne payload. The payload antennas for both the 802.11b and UHF packages were omnidirectional, so that the ground nodes could be expected to exact the same performance if located anywhere on a circle with the same slant range to the balloon. For example, two ground nodes, each with a slant range of 90

miles to the payload, could give the same performance up to 180 miles distant if the balloon were directly between them.

The 802.11b router was successfully demonstrated at slant ranges greater than 100 miles and between two and three nodes on the ground. Multicast was the primary type of network traffic simulated, although some unicast traffic scenarios were also tested (unicast does not perform as well at ranges over 20 miles). The latency created by the extended path lengths for the links resulted in degradation of network performance, but after a major decrease in performance at very short ranges (10 to 20 miles), the degradation was much more gradual as path length increased, and network throughput rates in excess of 60 Kbps could be achieved at slant ranges greater than 90 miles. These extended ranges do come with the caveats that a) directional antennas used on the ground required some tracking of the payload in order to point the antenna, and b) the nulls for the semi-omnidirectional antenna for the payload were at the horizon, so that at extended ranges the gain of the overall router was significantly compromised (over 10 dB). The latter was somewhat overcome by the large power amplifiers (50 W or greater) used at the ground nodes.

Over the UHF network, the results were severely compromised by what turned out to be limitations in the military radios used in the testing, the PRC-117F. (These limitations were confirmed in the laboratory after the conclusion of the Lubbock field testing, and are documented in the Appendix.) Throughput rates of only 9600 bps were achieved, both at very short slant ranges (10 to 11 miles) and at ranges exceeding 85 miles. Voice communication at all ranges through the PRC-117F was never a problem through the relay, and was demonstrated at slant ranges exceeding 90 miles. It should be noted that in the November 2005 exercise conducted out of Roswell, New Mexico, a voice link between two stations approximately 245 statute miles distant was demonstrated.

A dual-channel analog transponder was demonstrated, with traffic run simultaneously over both channels. The performance of the dual-channel unit compared to the single-channel unit was diminished due to a reduction in the overall gain of the UHF relay payload. However, there were no observed effects of crosstalk between the channels; that is, when the link margins were sufficient for the communications link, the relay performed well on both channels. This validated the future design that will be implemented in a digital version of the UHF payload, utilizing digital signal processing and FPGA-based designs that will allow for multiple channels at a reduced weight and size.

FUTURE WORK

The experiment documented here demonstrated the feasibility of the HARR payloads in supporting LOS communications over great distances. The next step for both payloads is to reduce their size, weight, and power (or SWaP). A reduction in size makes for easier handling and launch. A reduction in weight not only simplifies handling, but also results in reducing the size of the balloon required and the amount of helium required to launch it. This is graphically illustrated in Fig. 44, where a drop from the current payload weight of over 120 lbs (with AFRL telemetry and rigging) to less than 20 lbs results in a drop from over 50,000 cubic ft of helium to less than 8,000 cubic ft. The latter balloon and payload can be launched by two people from land or sea.

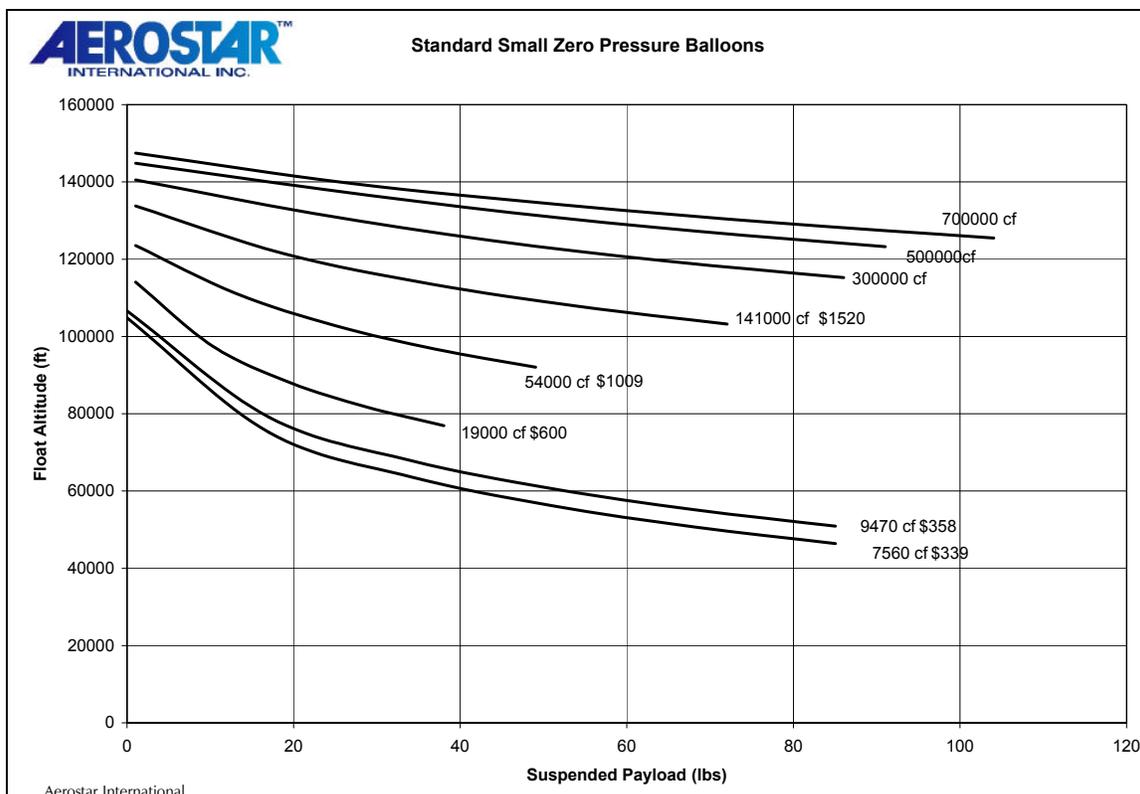


Fig. 44 — Zero-pressure balloons with corresponding payload capacity and float altitudes. Given a target of <20 lbs for future payloads, it can be seen that the required cubic feet of helium to lift such a payload drops dramatically from the current requirement of over 50,000 cu. ft. for the 120-lb payload configuration described in this report.

ACKNOWLEDGMENTS

The authors would like to acknowledge a number of individuals who significantly contributed to this effort: our sponsors, Mr. Douglas Crowder and Mr. Adrian Eley of the Office of Naval Research, whose support and vision were critical to our success; Mr. Reed Porada, formerly of NRL, whose engineering efforts contributed significantly to the startup of this effort; Mike Smith of Aerostar International, Inc., who provided the means to get this technology aloft; U.S. Navy reservist Lt. Dan Campos and Allen St. Jean, who supported our testing in Texas; and Bill Flynn and the Air Force Research Laboratory team out of Kirtland AFB, Albuquerque, New Mexico, who provided the telemetry and launch and recovery support for our exercise, as well as technical guidance in the development of our payloads. NRL would also like to recognize the city of Brownsville, Texas, and the good people at Terry County Airport, who were our most gracious hosts during the field tests described in this report.

Appendix

IP NETWORK PERFORMANCE TESTING OF THE HARRIS PRC-117F RADIOS

INTRODUCTION

The Naval Research Laboratory (NRL) Satellite and Wireless Networking Section has been supporting the Office of Naval Research (ONR) in their High Altitude Relay and Router (HARR) program. One of the payloads developed by NRL is a surrogate UHF transponder that is capable of relaying UHF line-of-sight (LOS) communications over 200 miles. This capability has real-world utility, and therefore it was desired to determine not only the performance of the relay itself but also the performance of the radios likely to be utilizing the relay.

The U.S. Marine Corps has made a substantial commitment to the Harris PRC-117F (V)(C) Manpack radio (hereafter referred to as simply the 117), a tactical radio operating in voice and a multitude of data modes for both VHF and UHF. This radio can operate in LOS mode or over the DoD FLTSAT satellites. In addition, the 117 offers a wireless Internet Protocol (IP) capability that allows operators to use IP-based personal computer (PC) software applications. This is done with the radios in the High Performance Waveform (HPW) IP mode, and can be done in SATCOM or LOS. The 117 was therefore a most suitable candidate for testing the HARR UHF relay.

The following describes a condition of the Harris PRC-117F transceiver when it is set to operate in its HPW mode for transferring IP data on a point-to-point link. This condition, in which throughput is severely compromised due to a large rate of packets per second (pps), was first realized during extended LOS testing of the radios in the field, as documented in the main body of this report.

Once the June 2006 field exercise was completed, the results were plotted out and analyzed, and it was realized that there was a ceiling to the realizable throughput that seemed to be independent of range (and therefore of latency). It was therefore important to determine whether the performance results of the 117s in HPW IP mode were a function of the relay or of the radios themselves (and our operation of them). In addition, it was necessary to ascertain whether such high loss fractions are prevalent in all communication links operating through the radio set, or if it was due to the test configuration used.

The following testbed was set up in the laboratory at NRL. Originally, tests utilized whip antennas at both ends, but these were replaced almost immediately with the attenuator configuration shown in Fig. A1, to remove any possible multipath interference.

The 117 configuration remained identical to that utilized in the June field exercises, with the exception that the power output was reduced to 1 W. Testing was originally conducted based on results obtained from a Marine Corps contractor (CenGen, Inc.) that reported that performance of the 117F radios favored large packet sizes with a low packet rate. Unless otherwise specified, the radios were fixed to operate at 64 Kbps in HPW IP LOS mode.

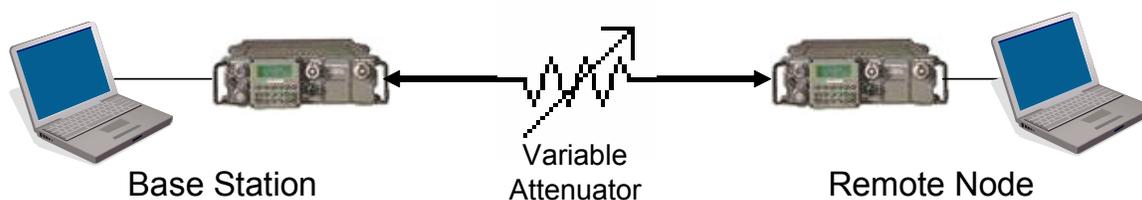


Fig. A1 — Verification of field results with PRC-117F radios, operating in HPW IP mode. A direct cable was installed between RF antenna ports (J1) on the radios to eliminate multipath. Power output = 1 W, with each radio receiving -85 dBm. Radios configured for plain text (PT) mode, variable attenuation set at 115 dB.

MGEN

As in the field tests, the PC-based application used for testing the performance of the 117 radios was the Multi-Generator toolset, or MGEN. MGEN is open source software provided by the NRL PROTOCOL Engineering Advanced Networking (PROTEAN) Research Group. MGEN provides the ability to perform IP network performance tests and measurements using User Datagram Protocol (UDP)/IP traffic.

LABORATORY RESULTS

MGEN input files were executed to locate and determine the optimum combination of packet size and packet rate (pps) that would allow the radios to pass the most traffic possible. Results are shown below in pairs of plots with packet throughput on the left and loss fraction on the right. Shown in Fig. A2 are the bidirectional results for the lab testing between the two radios. The radios were programmed for 64 Kbps, and traffic was generated at a constant rate for three successive 5-minute periods, maintaining a data rate of 45, 50, and then 55 Kbps, respectively. To achieve that rate, the packet size was ramped up from 625 to 1125 bytes for 45 Kbps, 694 to 1250 bytes for 50 Kbps, and 763 to 1375 bytes for 55 Kbps. The packet rate was correspondingly dropped from 9 to 5 pps for all data rates.

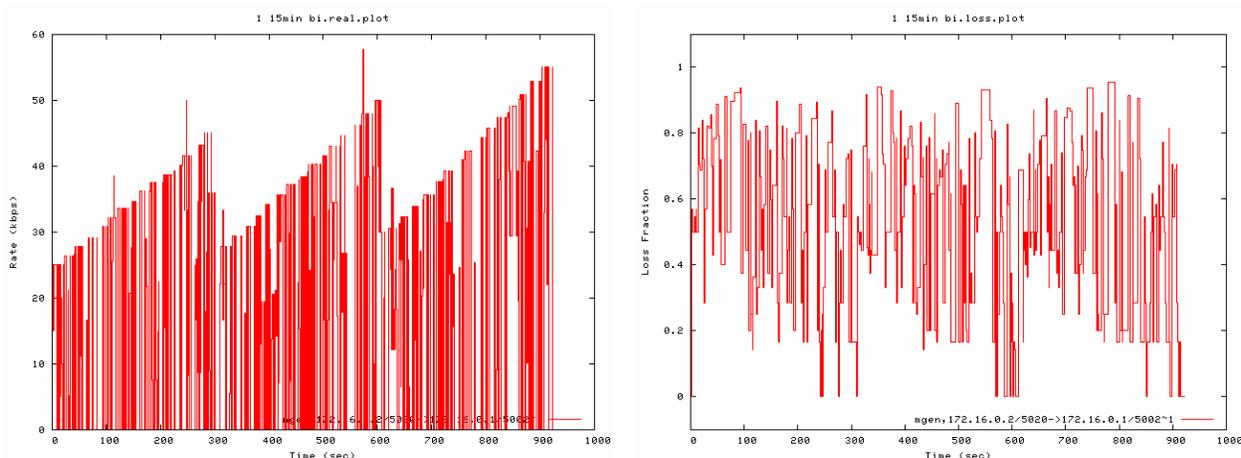


Fig. A2 — Bidirectional unicast traffic (throughput and loss fraction) between two PRC-117s connected as shown in Fig. A1. Traffic was sent at data rates ramping up to 45, 50, and 55 Kbps in 5-minute intervals. During each interval, the packet rate transitioned from high to low and the packet size from small to large, maintaining a constant traffic rate.

Unidirectional tests using the same MGEN input file at one end only were also executed as shown in Fig. A3. Although it is difficult to glean from the bidirectional loss fraction results (Fig. A2), it is clear to see in the unidirectional results that as the throughput is more successful (corresponding in every case to larger packets and fewer packets per second), the loss fraction improves substantially.

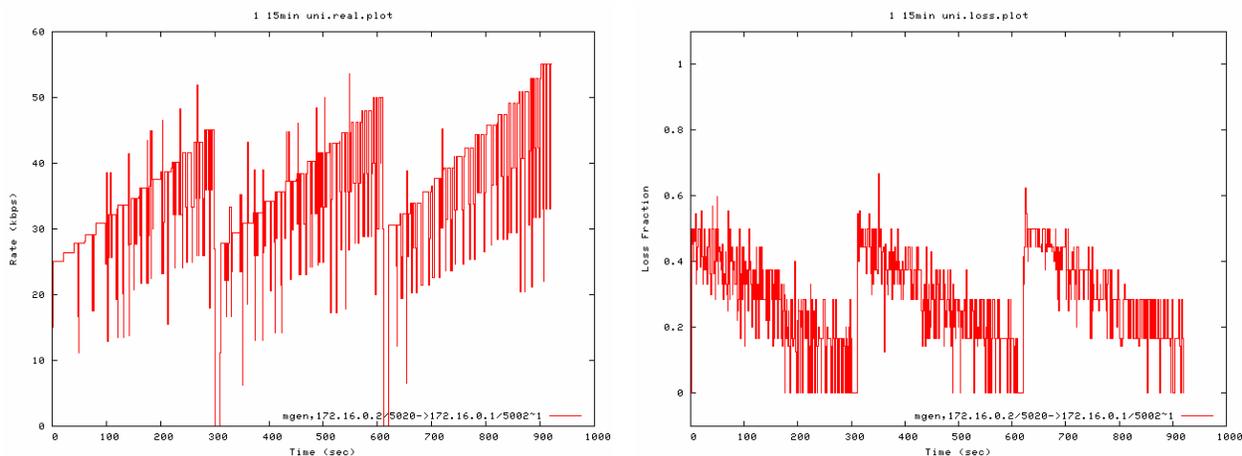


Fig. A3 — Unidirectional unicast traffic (throughput and loss fraction) between two PRC-117s connected as shown in Fig. A1. Traffic was sent at constant data rates of 45, 50, and 55 Kbps in 5-minute intervals. During each interval, the packet rate transitioned from high to low and the packet size from small to large, thus maintaining the constant traffic rate. The loss fraction is approximately 20% when near the optimum combination.

Next, it was attempted to push the 64 Kbps as hard as possible for both unidirectional and bidirectional traffic. MGEN input files were executed with the same general plan as above—increasing the packet size and decreasing the packet rate for the maximum throughput and minimum loss fraction. Figure A4 shows the unidirectional results, and indicates that the peak performance of the radios occurs just before the radios are “overloaded” due to over-large packet sizes, and the throughput drops substantially. Figure A5 shows the bidirectional results, which show a similar effect in reaching peak traffic just before a collapse in performance. It is more difficult to determine much from the loss fraction plot of Fig. A5 (as compared to the unidirectional scenario in Fig. A4), but the throughput shows both an achievement of peak performance and more pronounced outages in the traffic as the radios switch from transmit to receive and back.

The results were correlated with the MGEN input files used to generate the data, and it was determined that the optimum packet rate is approximately 5.4 pps, with an optimum packet size of 1489 bytes per packet at that rate. The corresponding total throughput was about 59.5 Kbps.

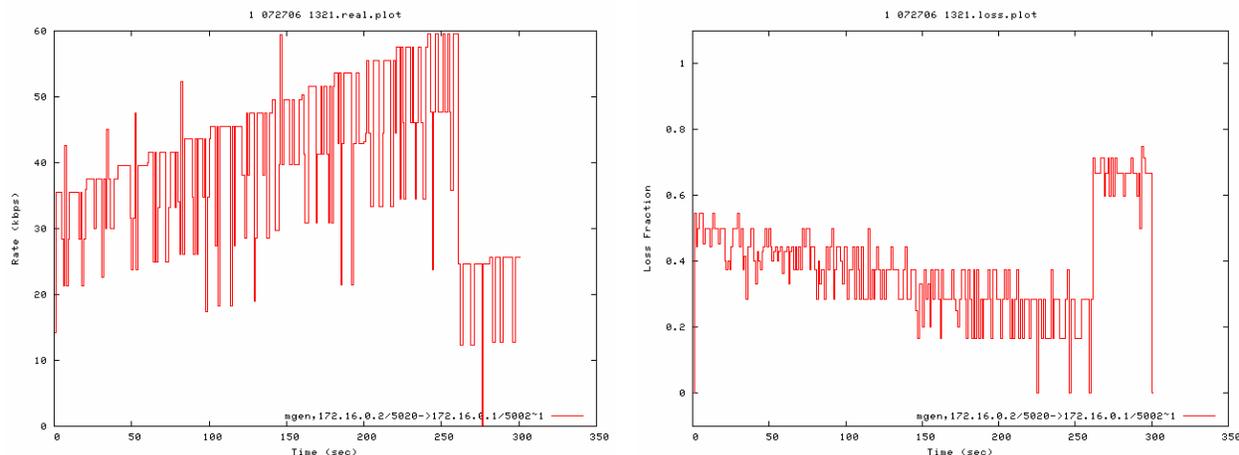


Fig. A4 — Unidirectional unicast traffic (throughput and loss fraction) between two PRC-117s connected as shown in Fig. A1. Traffic was sent at a constant data rate of 64 Kbps, and the pps (decreasing) and packet size (increasing) were adjusted to maintain that traffic level throughout the 5-minute test.

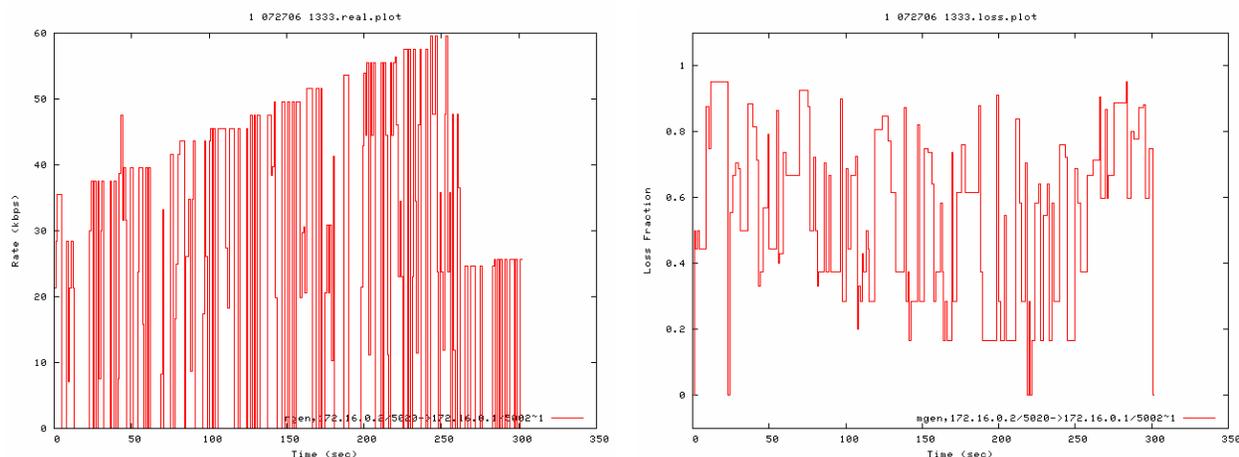


Fig. A5 — Bidirectional unicast traffic (throughput and loss fraction) between two PRC-117s connected as shown in Fig. A1. Traffic was sent at a constant data rate of 64 Kbps, and the pps (decreasing) and packet size (increasing) were adjusted to maintain that traffic level throughout the 5-minute test.

After it was demonstrated that performance was dependent on both packet size and, more importantly, packet rate, the tests conducted in the field in Lubbock were replicated in the laboratory, under the same operating configuration shown in Fig. A1. The identical MGEN input files used in the field were run again, with the results shown in the figures below. Once again, the first period shows a throughput performance better than expected for the first packet size throughout the testing.

Figures A6 and A7 are MGEN unicast packets sent in one direction (from node 1 to node 2 and vice versa), and their performance can be directly compared to that shown in Figs. 33 and 34 in the main report. This laboratory version shows fewer complete dropouts of performance as well as occasional spikes in higher throughput, but on the whole shows that regardless of the channel bandwidth (here being infinite), received signal strength, or latency, the radios perform in a very similar fashion.

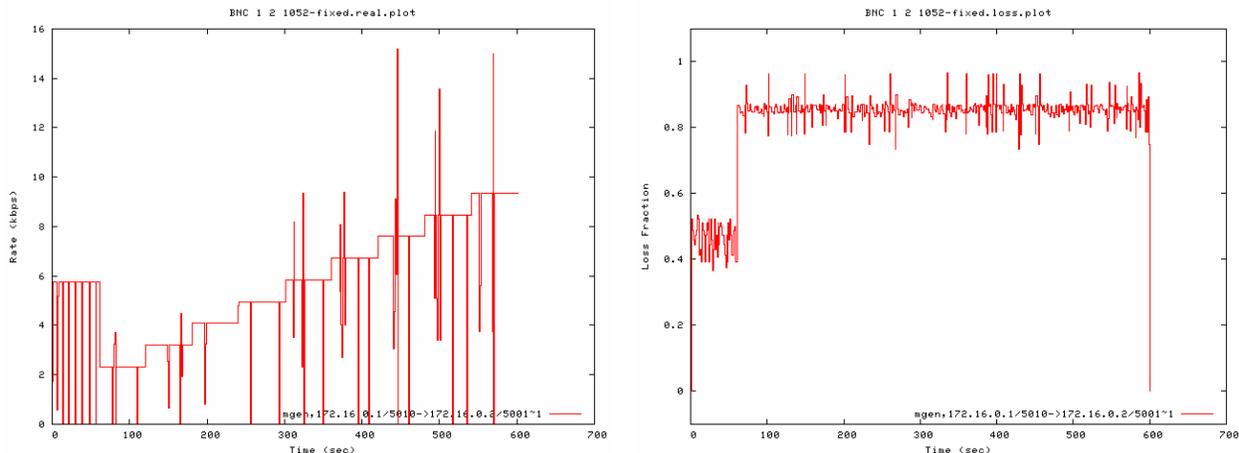


Fig. A6 — Unidirectional unicast traffic between two PRC-117s connected in the lab, with node 1 sending to node 2. Packet rate was held constant at 33 pps, packet size was increased to fill the 64 Kbps channel. These were generated using MGEN input files that were used for field testing with the HARR relay in June 2006.

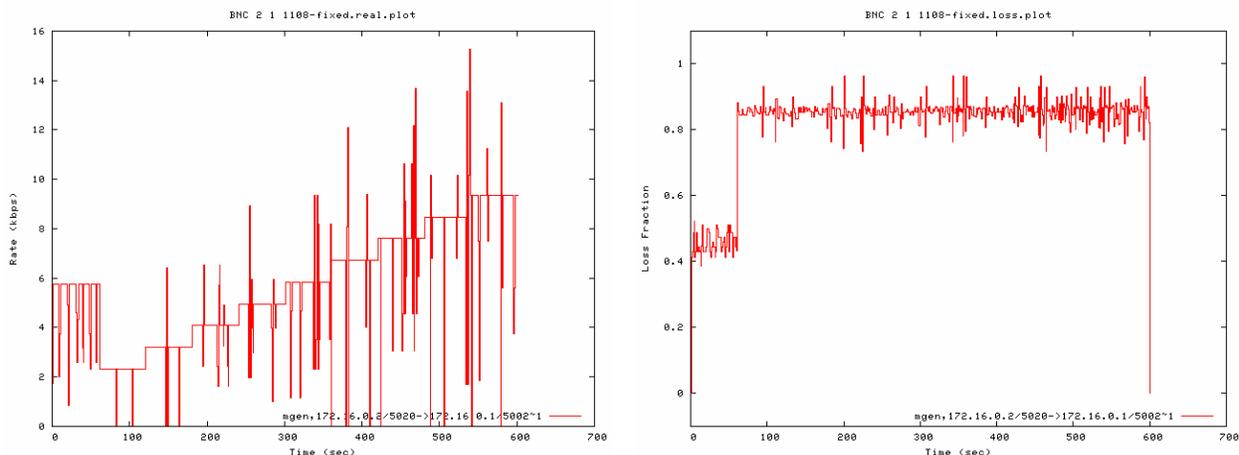


Fig. A7 — Unidirectional unicast traffic between two PRC-117s connected in the lab, with node 2 sending to node 1. Packet rate was held constant at 33 pps, packet size was increased to fill the 64 Kbps channel. These were generated using MGEN input files that were used for field testing with the HARR relay in June 2006.

Figures A8 and A9 are the results for bidirectional unicast traffic, shown at both ends of the link, and can be directly compared to the results presented in Figs. 36 and 37. Again, there is minimal difference in IP UDP performance as indicated in these results, and it should be noted that the large dropouts in traffic movement occur in the lab as well as in the field, and therefore it is reinforced that latency was not the issue for the earlier results.

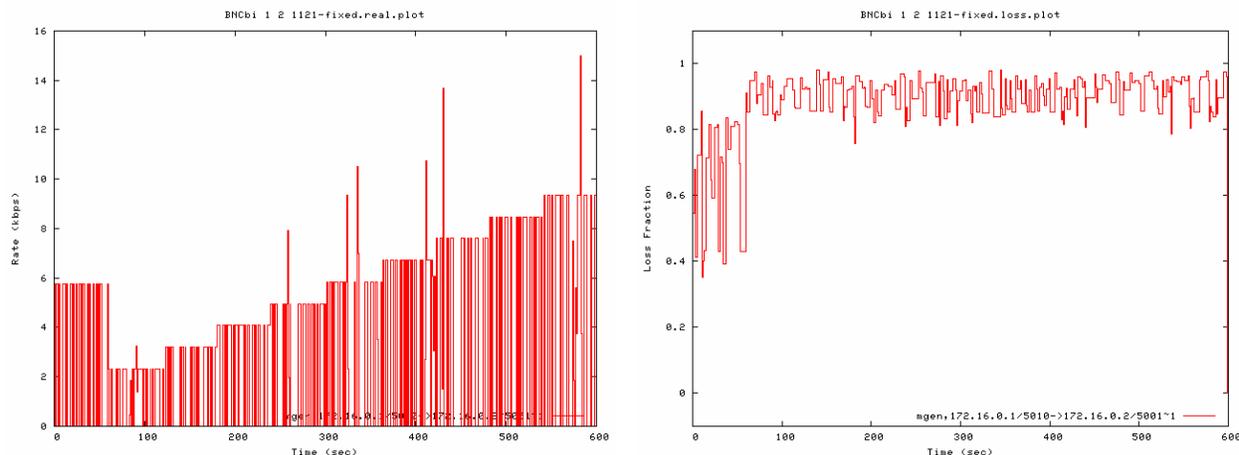


Fig. A8 — Bidirectional unicast traffic between two PRC-117s connected in the lab, data collected at node 2. Packet rate held constant at 33 pps, packet size increased to fill the 64 Kbps channel.

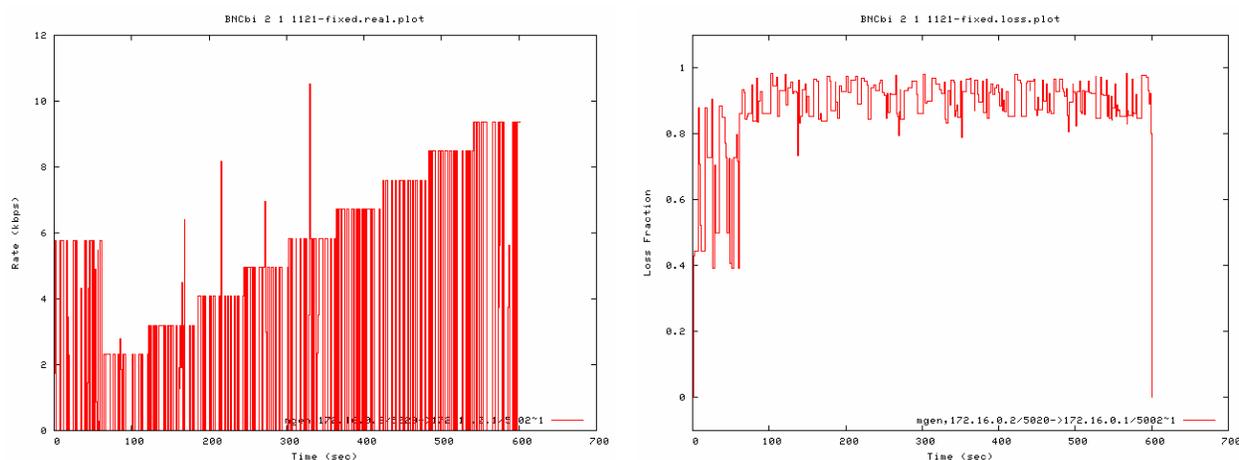


Fig. A9 — Bidirectional unicast traffic between two PRC-117s connected in the lab, data collected at node 1. Packet rate held constant at 33 pps, packet size increased to fill the 64 Kbps channel.

FURTHER TESTING — SHORT PACKET PERFORMANCE

After further analysis, the decision was made to look more closely at the scenario in which the packet rate was ramped up, but the packet size was held steady at a small value, equivalent to constant bit rate (CBR) traffic such as voice. Packets were held constant at one value, and the packet rate was ramped up from 2 to 25 packets per second. Each pps value was held for one minute as traffic was sent and received between the two PRC-117F radios, in the same laboratory configuration shown in Fig. A1.

The data is again presented with the packet throughput on the left in each figure and the loss fraction on the right. A packet size of 30 bytes/packet was used to generate Fig. A10; Fig. A11 shows the results with a packet size of 35 bytes; Fig. A12 shows the results for a packet size of 50 bytes; and Fig. A13 shows the packet size increased to 75 bytes.

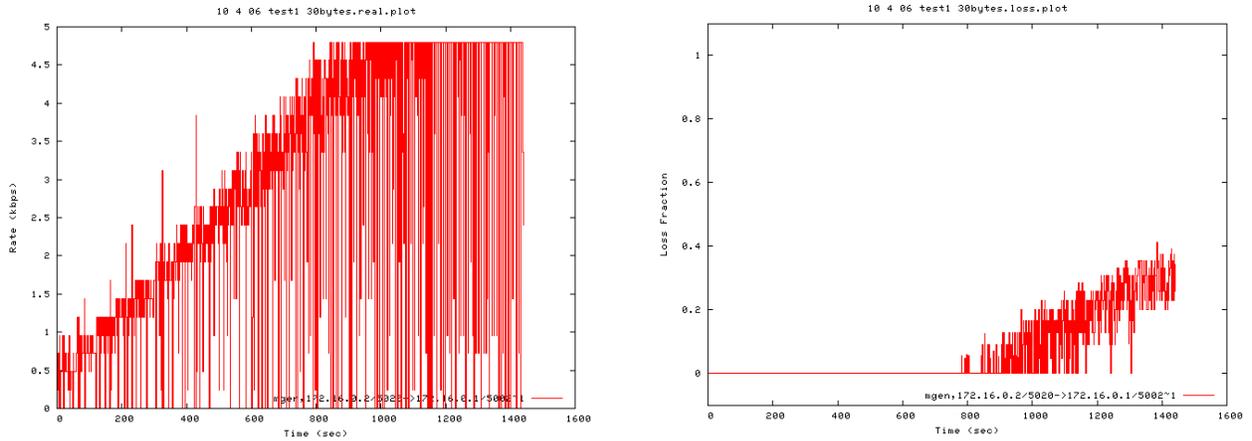


Fig. A10 — MGEN plots of throughput and loss fraction for unidirectional unicast traffic between two PRC-117s connected in the lab. A constant packet size of 30 bytes per packet was held throughout the test. The packet rate was increased from 2 to 25 pps, with each interval held for one minute.

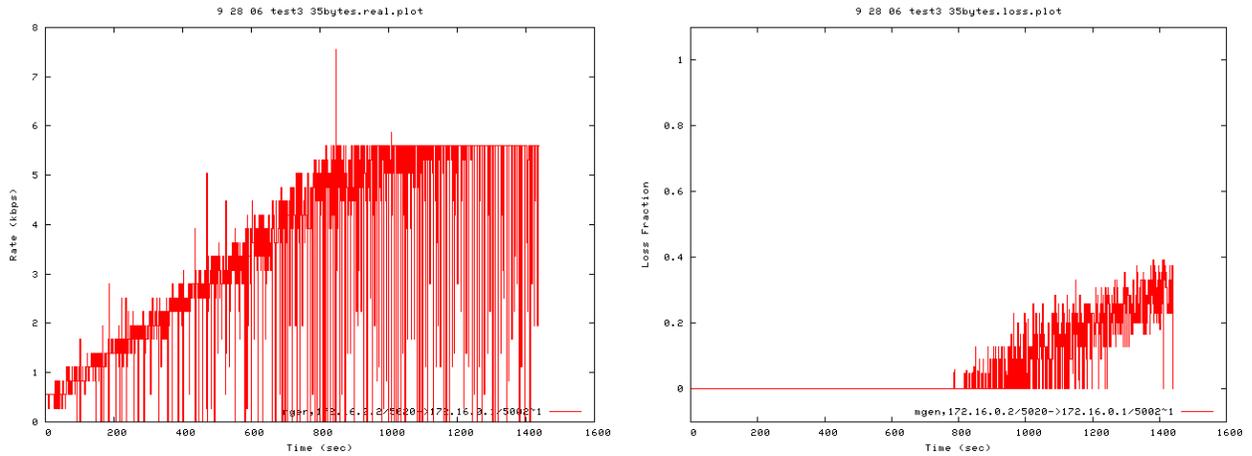


Fig. A11 — MGEN plots of throughput and loss fraction for unidirectional unicast traffic between two PRC-117s connected in the lab. A constant packet size of 35 bytes per packet was held throughout the test. The packet rate was increased from 2 to 25 pps.

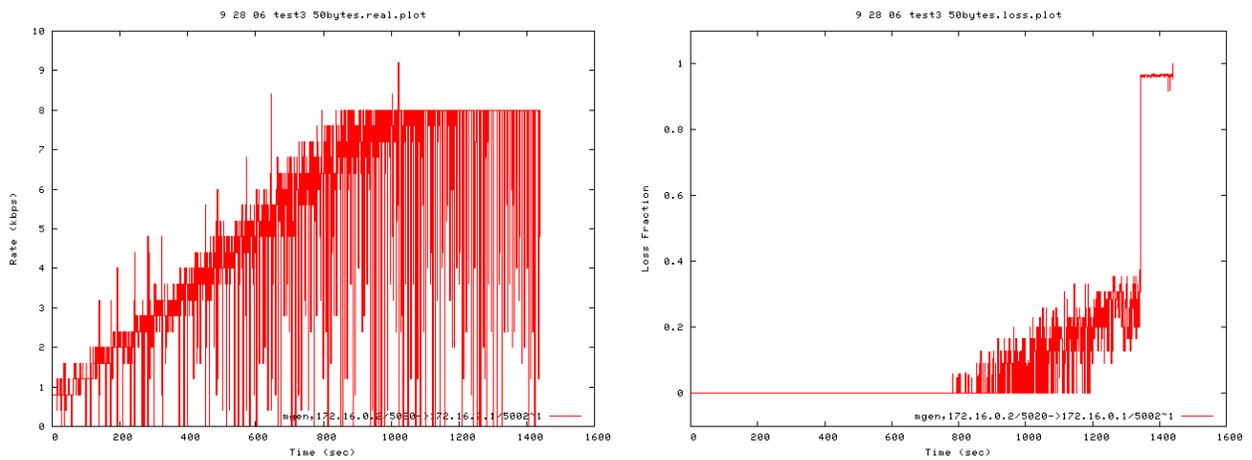


Fig. A12 — MGEN plots of throughput and loss fraction for unidirectional unicast traffic between two PRC-117s connected in the lab. A constant packet size of 50 bytes per packet was held throughout the test. The packet rate was increased from 2 to 25 pps.

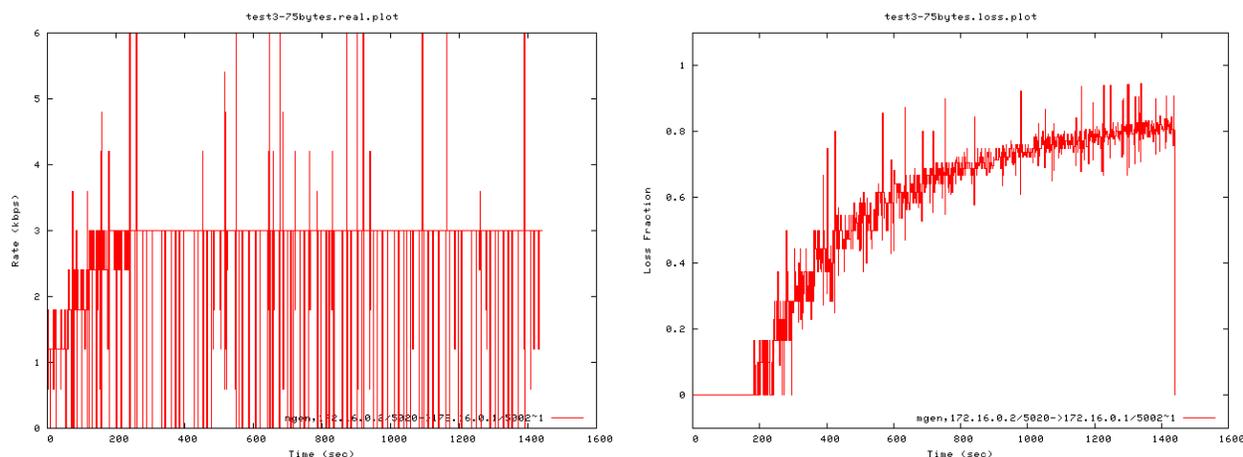


Fig. A13 — MGEN plots of throughput and loss fraction for unidirectional unicast traffic between two PRC-117s connected in the lab. A constant packet size of 75 bytes per packet was held throughout the test. The packet rate was increased from 2 to 25 pps.

In each instance, the loss fraction jumps from zero at or near the instant that the “wall” is hit in throughput performance. The packet rate at which that occurs changes little between the 30-, 35-, and 50-byte packet scenarios, with the loss fraction jumping from zero at the onset of the 13th minute, where the packet rate goes to 15 per second. The loss fraction jumps rather severely for the 75 byte per packet condition from about 5 minutes on, and it is believed that this is where the packet size begins to have a greater impact. Each of these tests was conducted more than once to verify/validate the results shown, although the subsequent tests at 75 bytes per packet indicated an even worse performance than shown in Fig. A13 (with a nearly identical throughput performance).

CONCLUSIONS

It has been identified and confirmed that the Harris PRC-117F radios do indeed pass IP traffic, but the effective throughput rate of that traffic depends considerably on the packet size and packet rate output of the application generating them. Without optimization, the effective throughput of the links will be severely compromised. Moreover, this limitation is believed to be exclusively the result of the radio's operation in the physical and link layer implementations in the 117 units.

It should also be noted that the radios perform at significant ranges in LOS mode, as made possible by the UHF relay developed by the HARR program, and that these radios have been demonstrated to pass data and voice over total ranges exceeding 180 miles (240 miles for voice). Moreover, it is believed that the high loss fractions in IP packets experienced in the field was a condition of the packet rate at which the testing was conducted, and that operations at range could be optimized for performance just as was done by NRL in the laboratory testing documented here.

Harris Corporation has been made aware of these test results, and has said that the issues set forth here will be addressed in a future software modification to the PRC-117F firmware.