



Control and Non-Payload Communications (CNPC) Prototype Radio—Generation 2 Flight Test Report

*Joseph A. Ishac, Dennis C. Iannicca, and Kurt A. Shalkhauser
Glenn Research Center, Cleveland, Ohio*

*Brian A. Kachmar
Vantage Partners, LLC, Brook Park, Ohio*

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.

- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, organizing and publishing research results.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at <http://www.sti.nasa.gov>
- E-mail your question to help@sti.nasa.gov
- Fax your question to the NASA STI Information Desk at 443-757-5803
- Phone the NASA STI Information Desk at 443-757-5802
- Write to:
STI Information Desk
NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320

NASA/TM—2014-218391



Control and Non-Payload Communications (CNPC) Prototype Radio—Generation 2 Flight Test Report

*Joseph A. Ishac, Dennis C. Iannicca, and Kurt A. Shalkhauser
Glenn Research Center, Cleveland, Ohio*

*Brian A. Kachmar
Vantage Partners, LLC, Brook Park, Ohio*

National Aeronautics and
Space Administration

Glenn Research Center
Cleveland, Ohio 44135

November 2014

Acknowledgments

The National Aeronautics and Space Administration thanks the faculty and staff of the Avionics Engineering Center at the Ohio University Russ College of Engineering for their generous logistical and institutional support, enabling us to successfully install and operate the Albany, Ohio ground station described in this report.

This report is a formal draft or working paper, intended to solicit comments and ideas from a technical peer group.

This report contains preliminary findings, subject to revision as analysis proceeds.

Trade names and trademarks are used in this report for identification only. Their usage does not constitute an official endorsement, either expressed or implied, by the National Aeronautics and Space Administration.

Level of Review: This material has been technically reviewed by technical management.

Available from

NASA Center for Aerospace Information
7115 Standard Drive
Hanover, MD 21076-1320

National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Available electronically at <http://www.sti.nasa.gov>

Control and Non-Payload Communications (CNPC) Prototype Radio—Generation 2 Flight Test Report

Joseph A. Ishac, Dennis C. Iannicca, and Kurt A. Shalkhauser
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

Brian A. Kachmar
Vantage Partners, LLC
Brook Park, Ohio 44142

Summary

NASA Glenn Research Center conducted a series of flight tests for the purpose of evaluating air-to-ground communications links for future unmanned aircraft systems (UAS). The primary objective of the test effort was to evaluate the transition of the aircraft communications from one ground station to the next, and to monitor data flow during the “hand-off” event. To facilitate the testing, ground stations were installed at locations in Cleveland, Ohio and Albany, Ohio that each provides line-of-sight radio communications with an overflying aircraft. This report describes results from the flight tests including flight parameters, received signal strength measurements, data latency times, and performance observations for the air-to-ground channel.

Introduction

NASA is currently working with the Federal Aviation Administration (FAA), RTCA Inc. (a Federal Advisory Committee), and other organizations to develop technologies and procedures to allow Government (public) and commercial (civil) unmanned aircraft (UA) to operate safely in the National Airspace System. An essential element of this new national capability is the radio communications channel linking each UA to a network of fixed ground stations (GSs). Data communication necessary for flight is referred to as control and non-payload communication (CNPC) and is exchanged over the UA and GS channel to ensure safe, reliable, and effective UA flight operation.

In Phase I, NASA Glenn Research Center (GRC) performed flight testing using Generation 1 prototype CNPC radios to characterize the performance of the air-to-ground (AG) channel operating in the 960-977 MHz frequency range (L-band). [1] Results indicated that a significant separation distance could be achieved between a single aircraft and a single ground station while maintaining a reliable data connection. Channel bandwidth, transmitter output power level, and the impact of terrain obstructions on line-of-sight (LOS) communications were studied under the Phase I flight testing campaign.

This report presents results on a Phase II system utilizing two spatially-separated ground stations communicating with one overflying aircraft. The principle objective of this test campaign was the evaluation and characterization of an in-flight “hand-off” event where the aircraft communication path transfers from one ground station to another with minimal interruption to the CNPC data flow. This transfer also required development of an automated system of ground terminals with the associated routing protocols to provide network mobility for a seamless hand-off. Other objectives of the Phase II test campaign were to demonstrate remote-controlled operation of multiple ground stations in a network environment. Data acquired from this test campaign will be used in subsequent UA flight test experiments of increasing complexity.

The Phase II flight tests utilize “Generation 2” software-defined radios in both the aircraft and ground systems to transmit and receive the CNPC signals. The radios were developed under cooperative agreement NNC11AA01A between the NASA Glenn Research Center and Rockwell Collins, Inc., of Cedar Rapids, Iowa. Descriptions of the radio modulation, signal waveforms, and operating modes are presented in [1]. The radios are prototypes only and were designed to address the initial “seed” requirements from RTCA, Inc., Special Committee 203 (SC-203)—a Federal advisory committee developing standards for UA performance requirements. For this series of flight tests, the Generation 2 prototype radio remains similar to the Generation 1 model, especially in terms of radio frequency (RF) operations. Changes made to the radio include modifications to the status and control interface, radio-to-radio link health reporting functions, and other “housekeeping” adjustments, resulting in only small changes to the over-the-air message overhead content.

This report presents results from the networking hand-off tests. The data presented are intended only as an early gauge of the general capabilities of the CNPC links. On-going iteration of system and radio software and firmware are underway to optimize CNPC data throughput rates. Those improvements, along with detailed discussion and test data from data header compression and data security flight tests, will be presented in subsequent reports.

Test System Description

The major elements of the Phase II CNPC radio test system are displayed in Figure 1. The system consists of two identical GSs, one installed at the GRC Lewis Field campus in Cleveland, Ohio and one installed at the Ohio University Gordon K. Bush Airport in Albany, Ohio (airport designation: KUNI). Each GS has a network connection to the Internet through hardwired connections for remote control, monitoring, programming, and test operations. Both GSs have automatic fail-over switching allowing them to use a 3G/4G cellular wireless connection if the Internet service at the GS site is interrupted. Using this arrangement, both GSs are networked to the GRC UAS laboratory where the test flights and data collection are centrally managed.

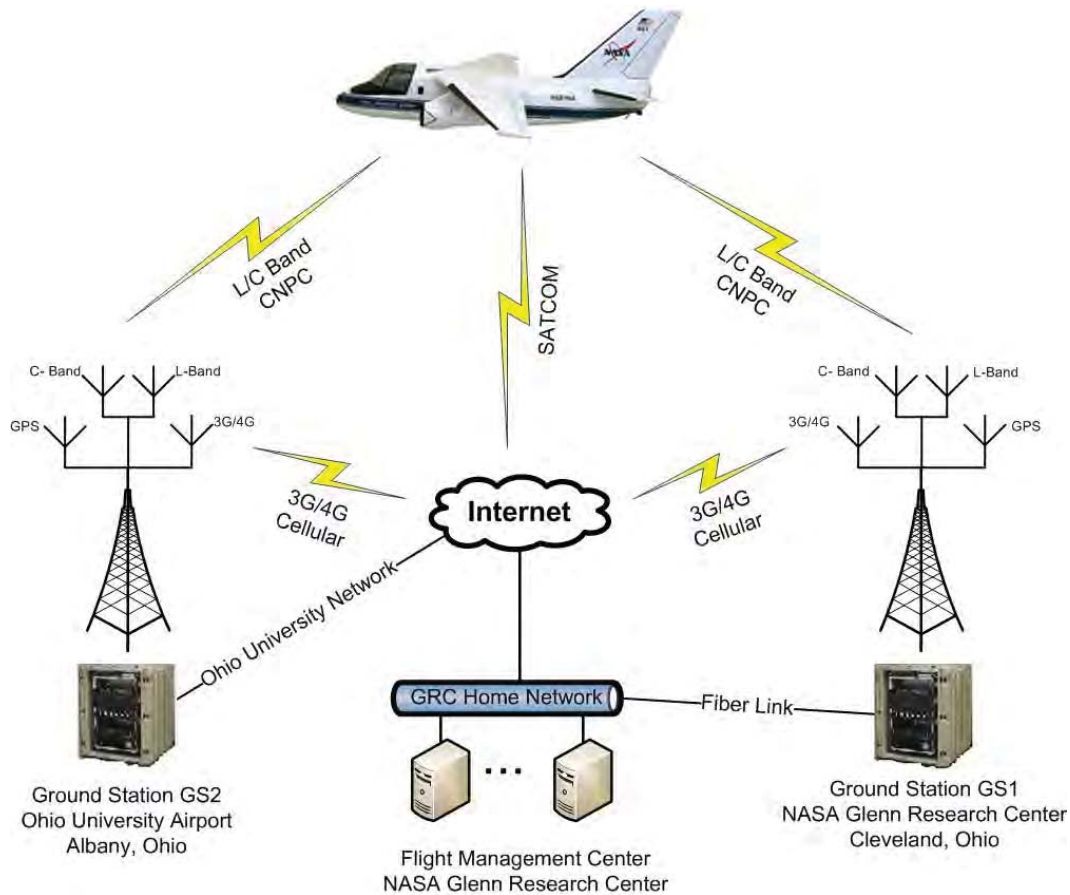


Figure 1. - Test system overview

Ground Stations

The ground stations used in the Phase II test campaign were designed to meet several primary requirements. First, the GSs had to be sufficiently compact, where each one could be transported by a small crew and installed into a potentially small or confined space. With ground stations possibly being installed at various locations around the country, surviving the shipping process and deployment into harsh environments is necessary. Thus, the ground stations are packaged in a ruggedized enclosure to protect the GS electronics during shipment and operation. The second requirement was for remote controllability. With the ground stations separated by 133 nautical miles (nmi), the ability to control stations remotely from one location would eliminate the need for deployed personnel. The third requirement was for expandability. Currently each prototype radio can only act as a single receiver. While the current tests utilize only one L-Band radio per ground station, future builds must support simultaneous operations of additional radios in order to support multiple aircraft.

Figure 2 illustrates the final arrangement of components in the ground station. To keep the GS physical size manageable but still provide for expandability, the smallest readily-available components were chosen. To provide uninterrupted electrical power to the GS electronic equipment, a 120V mains power supply was identified that would require only two “rack units” (2U) of height. A 28-volt direct current power supply is used to drive the radios and occupies only 1U in the equipment enclosure. A 1U computer was installed with sufficient processing power to satisfy the control and data-gathering needs of the GS. An Ethernet switch, firewall/gateway, and 3G/4G cellular modem were consolidated onto a single shelf occupying 1U of rack space. To accommodate all of the current Phase II GS electronics and to provide space for additional future equipment, a 14U tall equipment rack was required.

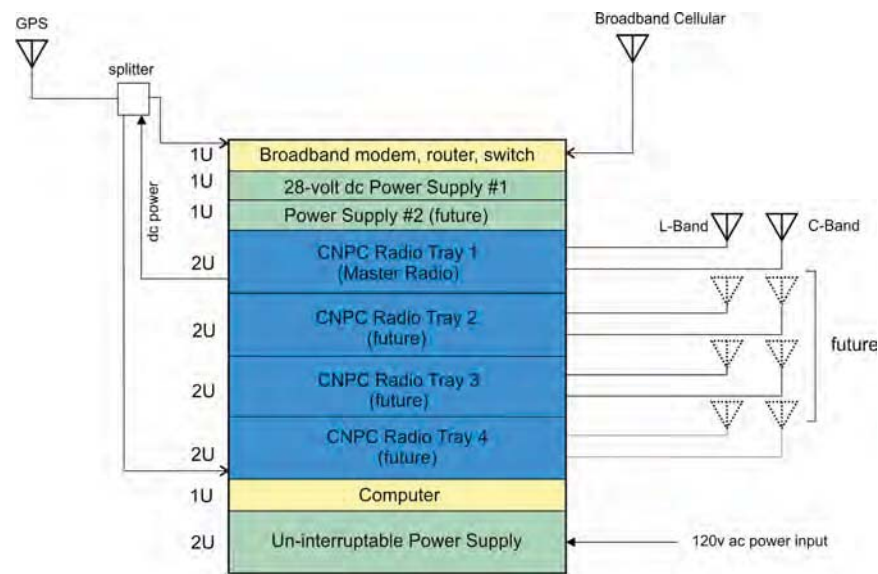


Figure 2. - Ground station component arrangement

The radios are installed onto the rack by using a radio tray. Each tray supplies power, cooling, and connectivity for up to two radios. Currently, each tray contains one L-Band and one C-Band radio. The 28V_{dc} power supply provides up to 25.7 amperes to easily drive up to four CNPC radios. A second, identical power supply will be added to the GS for later phases of flight testing when more than four radios will be in operation. An eight-port Ethernet switch isolates the radio network, allowing the master radio to communicate with any slave radios should they be present.

A ruggedized polyethylene transport case was selected to enclose the GS equipment. The case included a built-in rack-mount rail system on shock-absorbing mounts, waterproof covers, lifting and tie-down hardware, and wheels. The internal rack frame accommodates 14U of equipment height at the standard 19” equipment width.

Figure 3 shows the GS equipment in the transport case at a typical installation site. All antenna, power, and network cabling exiting the case to external locations is visible at the lower right. Waterproof covers are shown in stowed positions on the sides of the case, allowing front and rear ventilation for the GS equipment. Only one CNPC radio tray is installed in the case at this time.



Figure 3. - Transportable ground station installed at Ohio University Airport

Figure 4 shows electronic and electrical connections within each ground station. A majority of the equipment is electronically interconnected with RJ-45 Ethernet cables for control operations and data collection.

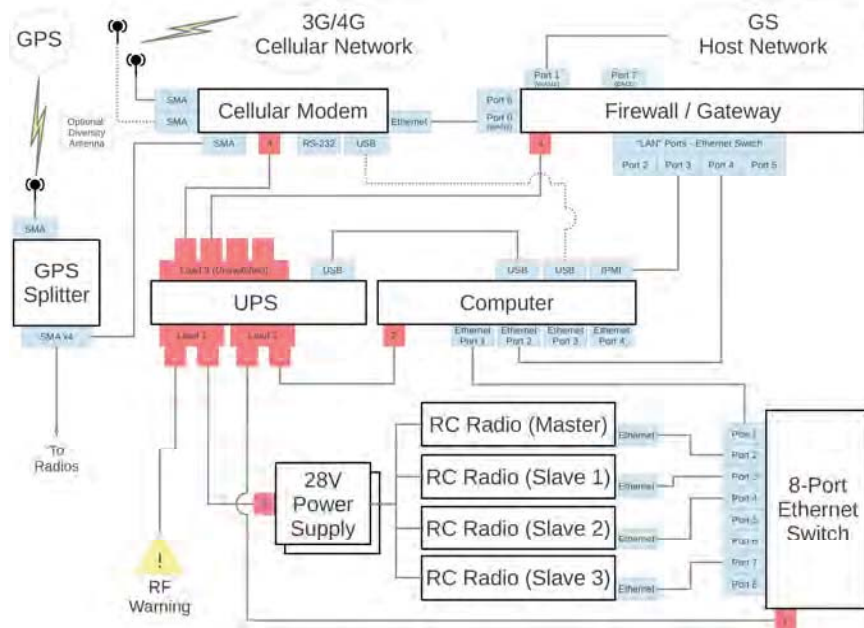


Figure 4. - Ground station internal connections

The ground station connects to the Internet through either a hard wired host-site network or a 3G/4G cellular network. The firewall/gateway is responsible for managing the networks provided by the 3G/4G connection and host network. In situations where both exist, the gateway will utilize the cellular network as a backup. Should the host network experience an outage or other disruption, the gateway will automatically switch to the cellular connection. The cellular modem is selected to provide data access using one of several cellular service providers, depending on service available at the specific GS location.

The uninterruptible power supply (UPS) delivers 120-volt power to three load circuits, two of which are independently controllable over the network. The Load 1 circuit feeds the power supply for the radios and a xenon safety strobe light. The light provides local visual indication that a radio transmitter is enabled and that personnel should maintain a safe distance from the antennas. The Load 2 circuit is connected to the Ethernet switch and the computer, which allows for power cycling (yielding a reboot of the computer) if necessary. The Load 3 circuit, which is not switchable, provides continuous power to the cellular modem and the firewall. In the event of a prime power (mains) loss at the GS installation site, the 1350 watt UPS has battery capacity to provide back-up power to the ground station for approximately thirty (30) minutes. During this time, operators would secure all test data and perform a controlled shut-down of the network and computer.

The Global Positioning Satellite (GPS) antenna feeds a power divider which provides a GPS signal to the CNPC radio(s) and cellular modem. The CNPC radio is used as a time server for the GS computer. This allows the GS to be synchronized with itself as well as with the other GSs and test equipment at the various test locations. This synchronization is necessary for measuring data flow latency throughout the GS network.

The computer has an Intelligent Platform Management Interface (IPMI) that permits remote monitoring and control of the computer without the use of the operating system. In the event of a computer failure, the GS systems can be powered down, rebooted, and have errors logged via IPMI.

Ground Station Site Selection

Demonstration of in-flight CNPC hand-off events required installation of a minimum of two ground stations separated by approximately 138 nmi (153 miles). This separation

corresponds to twice the recommended nominal GS coverage radius of 69 nmi [3]. This arrangement would allow each GS radio to be operated at its full transmit power level for communications throughout its individual coverage area, and would additionally create a flight region where LOS CNPC links could be established with both ground stations simultaneously.

Site selection for each GS involved multiple considerations to ensure that test requirements were met. Frequency authorizations for the area including most of Ohio, West Virginia, and western Pennsylvania have been obtained allowing some flexibility in site selection. The ability to operate the GSs by remote control also provided an increased number of placement options. Locations at or near airport property were desired (but not essential) in order to represent actual in-service operation. Flight planning also influenced GS site locations. Since test flights were anticipated to require non-standard and repeated flight maneuvers near the GS, isolated locations with low air traffic were preferred. This consideration holds true for the corridor between the two GSs, where straight and level flight was important for data collection. Terrain features were considered as well in order to avoid obstruction of the LOS signal path between the GS and aircraft. Once a site is selected, the GS equipment rack must be placed in a strategic location. The fundamental requirements included a 120V, 15-amp electrical power source, an Internet connection, a secure, indoor location near an exterior wall, and an antenna mounting location within 60 feet of the equipment rack.

GS1 was installed into an existing building in the western area of NASA's Glenn Research Center, adjacent to the Cleveland Hopkins International Airport (KCLE). Electrical power and Internet connections were readily available, as was access to an adjacent 50-foot tower for antenna installation. The close proximity of the GS installation site to the GRC UAS control center, approximately 0.25 miles, allowed for convenient GS servicing, if needed.

GS2 was installed into an existing aircraft hangar at the Ohio University Gordon K. Bush Airport in Albany, Ohio. An indoor mezzanine provided the necessary electrical connections and allowed easy access to the roof. Ohio University and the Avionics Engineering Center have a strong experience base in air-ground testing, so their facilities and staff were well-equipped to support the NASA GS installation. The Ohio University met all criteria for the GS equipment installation, including placement of the equipment in a well-secured area.

Figure 5 shows the relative position of the GSs within the state of Ohio. The GSs are separated by a distance of 133 nmi, which represents approximately two times the radius of a candidate CNPC coverage area [3]. The OU GS lies on a true vector heading of 187° relative to the GRC GS. The NASA S-3B aircraft traversed the flight path shown in Figure 5 repeatedly during the test campaign.



Figure 5. - Locations of ground stations within the State of Ohio. Separation distance 133 nmi

Four omnidirectional antennas are installed on masts at each location and connected to the GS. Each installation included one L-Band transmit/receive antenna, one C-Band transmit/receive antenna, one Global Positioning System (GPS) receive antenna, and one broadband cellular transmit/receive antenna. The antennas are all commercially available products with standard mounts and connectors. Each GS location allowed placement of the transportable electronics case to be within 60 feet of the antennas in order to minimize the length of the cable runs and amount of signal attenuation.

The antennas at the OU GS were mounted on an elevated platform above the main aircraft hangar, achieving an installed elevation of 47 feet above ground level (figure 6). The antennas at the GRC GS were mounted to an existing 50-foot, open-frame tower at an installed elevation of 42 feet above ground level (figure 7).




Figure 6. - Antenna installation at Ohio University Airport in Albany, Ohio.




Figure 7. - Antenna installation at NASA Glenn Research Center in Cleveland, Ohio.

Aircraft

The Lockheed model S-3B Viking aircraft owned by Glenn Research Center (registration number N601NA) was used for all Phase II flight tests. This airborne platform is the same one used in channel sounding measurements [2] and in the CNPC Phase I flight tests [1]. The aircraft offers significant electrical power, pressurized area for the radio and test equipment, significant flight duration (>4 hrs), room for two experimenters, and provides excellent flight stability and maneuverability.

The aircraft equipment consisted of a Rockwell Collins radio, spectrum analyzer, GPS time server, two computers, and networking equipment. These components were mounted in two 19-inch wide equipment racks in the rear of the aircraft. An aircraft-rated, low-profile, omnidirectional antenna was used for the CNPC radio. The antenna was connected to the radio by a custom-length, shielded coaxial cable with low insertion loss, allowing the radio system to be operated without the need for additional signal amplification.

The L-band antenna was mounted on the bottom surface of the aircraft. The CNPC antenna was physically separated from all other antennas by a distance of several wavelengths to reduce antenna-to-antenna coupling. When in level flight with the aircraft landing gear in stowed position, the flattened contour of the bottom of the aircraft provided a clear LOS path to the ground station with little or no airframe obstruction. Figure 8 shows the location of the L-Band CNPC antenna and underside geometry of the S-3B test aircraft.



Figure 8. - S-3B Viking test aircraft showing L-band antenna mounting location.
Image © 2008 – TarmacPhotos.com [9], used with permission

Network Description

Overview

The prototype CNPC network environment developed for the second generation radio was designed using Mobile IPv6 [4] with Network Mobility (NEMO) [5]. The architecture is comprised of three major components: the home network, the ground station, and the mobile node or aircraft (Figure 9).

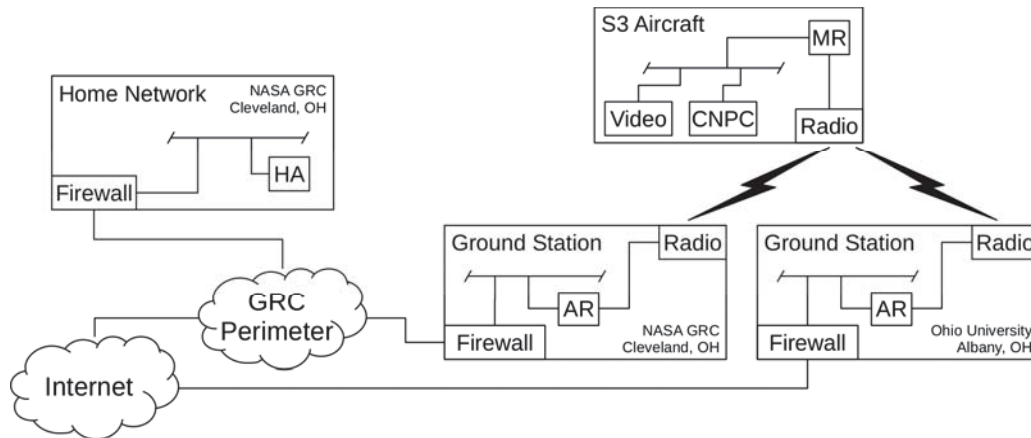


Figure 9. - Overview of network architecture used for phase II flight tests.

A home agent (HA) is placed in the agency's "home" network and is used as a central communications point for contacting the mobile node. The home agent serves to maintain a binding database of a mobile node's current foreign network addresses, or care-of address. The bindings map the home address, a known address in the home network, to the shifting care-of address in the foreign network in order to transparently reroute traffic from external users.

Each ground station (GS) serves as a broker providing a UAV with its foreign network IPv6 address. This is done by means of an IPv6 access router (AR) which communicates with the aircraft through the prototype radio. Router advertisements or other dynamic host configuration methods can then be used to provide the UAV with a network address. The ground station does not need to be owned by the same agency as the aircraft and will have address space associated with the terrestrial service provider used at that location.

The aircraft communicates to the ground station using the same prototype radio. A mobile IPv6 router attached to the radio allows the NEMO network to communicate

with the home network seamlessly. This allows the various components behind the mobile router to be addressed directly by using their “home” address. Users would not need to know the addressing schemes used at the various ground stations.

Internetworking

In a production environment, communication between endpoints would utilize IPv6 end-to-end, without the need for any encapsulation of traffic within legacy IPv4 packets for transmission across the Internet. However, the GRC perimeter network is currently unable to route IPv6 traffic natively.¹ In order to emulate a native IPv6 environment, “6in4” (IP Protocol 41) tunnels were created between the home agent and the ground station computer system to allow for utilization of the readily available IPv4 backhaul communications networks. While there are several alternative ways to configure the test architecture to transmit IPv6 packets across an IPv4 network, the 6in4 approach reduced configuration complexity and eliminated the need for devices between the home agent and ground station computer to be IPv6-compatible.

The home network and the ground stations are each protected by a dedicated firewall system that serves as both an IPv4 network firewall as well as a VPN gateway to secure the transmission of data across the Internet. Both ground stations were designed to support a primary wired WAN connection along with a backup WAN connection provided by a 4G LTE router. Each GS utilizes an IPv4 IPSec VPN connection to communicate with the home network and carry the 6in4 traffic between the home agent and the ground station computer.

Aboard the S-3B test aircraft are two research computers running Ubuntu Linux with multiple Ethernet interfaces connected into a switch. One system is configured as the mobile router and the other is a workstation statically configured with an address on the home network. The workstation routes all IPv6 traffic to the mobile router as a default gateway. By configuring the second research computer to a known home network address, users are able to communicate directly with the system regardless of what foreign network the aircraft is currently connected to. The aircraft computer will retain the same home network address as the mobile router’s care-of address moves from foreign network to foreign network.

Addressing Scheme

The Home Network and Ground Stations were designed to support a dual-stack IPv4/IPv6 environment in order to support direct communications with legacy devices

¹ NASA GRC is currently working on implementing IPv6 in response to industry and federal mandate [6]

such as UPS battery backup network interfaces or IPMI management interfaces on the ground station computers that may not support IPv6. In addition, an IPv4 addressing scheme was needed to serve as an endpoint for the 6in4 tunnels between the Home Agent and the Ground Station computer.

The ground station computers were configured to use the radvd daemon in order to support IPv6 neighbor discovery across the RF link by multicasting periodic Router Advertisement (RA) packets every 10 seconds.[7] The remote aircraft computer, upon receipt of one of these RA packets, would be able to self-configure its IPv6 Care-Of address and begin communicating with the ground station, and thus the home agent, at the network layer.

Network Mobility Concept

The Mobile IPv6 with NEMO design allows the aircraft to maintain seamless network-layer communications between the pilot's ground control station and the onboard avionics systems as the unmanned aircraft transitions between ground stations. This design may result in a few dropped packets, but established sessions should continue on without needing to reconnect. This method eliminates the costly procedure of tearing down and re-establishing active sessions.

In a fully implemented communications system, hand-offs would be performed automatically by software based on link conditions and other factors. For this phase of testing, all hand-offs were conducted manually onboard the aircraft by the flight researcher, as the exact conditions for automating the hand-offs are being explored. The process of performing hand-offs between ground stations was to first establish network communications with the new ground station, wait until an IPv6 address was acquired via the IPv6 neighbor discovery process, and finally terminate the network link to the original ground station. The mip6d process aboard the aircraft would sense that the active link was no longer available and use the newly established link to send a binding update to the home agent.

Flight Testing

The flight test campaign consisted of eight separate flights occurring over the period of April 11 through May 28, 2014. Each flight test consisted of multiple, preplanned aircraft maneuvers and flight path segments. The overall sequence of testing was to first examine the CNPC radio communications range and link performance using only the GRC-based ground station with an omnidirectional antenna. This data was used for

comparison to models and previous flight test data. Second, the radio communications range measurements were made with only the OU-based ground station. Once the RF link performance and slant range capabilities of the individual GSs were well understood, the GS's were then simultaneously operated and the combined radio environment was evaluated. From this information a "hand-off region" was established where CNPC signals could be simultaneously communicated between both ground stations and the aircraft, and where hand-off testing could be conducted.

Characterizing the digital data flow during the hand-off event was one of the primary objectives of the Phase II flight tests. The latency of the data packets through the CNPC radio link and the magnitude of any interruptions or loss of data during the hand-off event were of significant interest. As a result, the Phase II flight test plan included multiple periods of time where the aircraft would be flying in and through the hand-off region. Flight and ground data were sent to Rockwell Collins for a more detailed examination so that changes could be made before the Generation 3 radios are delivered. A summary of the Phase II flight tests is presented in Figure 10.

Test Number	Test Date	Ground Terminal Name (location)	Radio Configuration	Objectives	Comments
1	4/11/2014	GS1 (GRC)	1	Characterize L-band performance of GS1	Flight patterns selected to examine omnidirectional antenna performance.
2	4/18/2014	GS1 (GRC)	2, 3	Characterize L-band and C-Band performance of GS1	Flight patterns selected to examine omnidirectional antenna performance.
3	4/21/2014	GS2 (KUNI)	5, 6	Characterize L-band performance of GS2	Flight patterns similar to GS 1, selected to examine omnidirectional antenna performance at GS 2.
4	4/22/2014	GS1 (GRC) GS2 (KUNI)	6	Determine hand-off region, examine data latency and hand-off performance	Flight patterns directly between ground stations, through hand-off region.
5	4/23/2014	GS1 (GRC) GS2 (KUNI)	1, 4	Examine data latency and hand-off performance, investigate interfering signal, monitor other ops modes	Flight patterns directly between ground stations, through hand-off region.
6	4/30/2014	GS1 (GRC) GS2 (KUNI)	1, 11, 12	Test firmware updates to improve latency. UDP and TCP tests	Flight patterns directly between ground stations, through hand-off region.
7	5/28/2014	GS1 (GRC) GS2 (KUNI)	1, 11	Demonstrate multiple hand-offs and record full-system latency performance, perform pre-test of "sensed traffic" downlink, characterize C-band performance of GS2	Flight patterns directly between ground stations, through hand-off region.
8	5/29/2014	GS1 (GRC) GS2 (KUNI)	6	Integrated test of "sensed traffic" downlink into LVC, characterize C-band performance of GS1 and GS2	Flight patterns between ground stations, through hand-off region with moderate in-flight re-direction for demonstration of avoidance software.

Figure 10. - Flight test summary

A customized radio configuration was required for each flight test, depending on which ground stations were used or type of test. Each configuration explicitly arranged the radio's CNPC waveform parameters, including uplink and downlink frequencies, modes of operation (command and control (C2) data, weather data, video data), and frame timing. For example, test flight #1 utilized only GS 1 and the aircraft radios. The radios operated both the uplink and downlink modes at a center frequency of 968 MHz and interleaved narrowband (75 kHz) and wideband (875 kHz) streams at a 10 Hz rate. A complete listing of the radio configurations utilized in these tests is included in Appendix A. A more complete description of radio waveform and formatting is presented in [1].

Aircraft Flight Path

Flight plans were individually tailored to address the type of radio test being performed. Flights to characterize an individual ground station, such as that for test number one at Cleveland, Ohio, included a circular orbit at an approximately constant altitude and radial distance from the GS (Figure 11). This maneuver allowed a general assessment of antenna performance for a typical ground terminal installation, and identified the presence of any antenna pattern nulls or shadowing by ground obstructions. Once the circular orbit was completed, the pilot adjusted aircraft altitude and aligned for an over flight of the GS on an outbound vector heading of 195° (relative to magnetic north.) The aircraft then traveled directly towards GS2 along the vector at constant altitude and airspeed until the CNPC radio received signal strength dropped below sensitivity levels and the communications link was broken. The aircraft then reversed course and flew inbound along the same vector (magnetic compass heading 15°).

An orbital flight path and inbound/outbound flights were also performed for flight test three at GS2 in Albany, Ohio (Figure 12). Small, "racetrack" patterns were executed when the CNPC signals fell below the sensitivity level, and were used to verify the slant range distance where signal drop-out occurred. Narrowband waveform modes generally provided greater range than did the wideband modes, resulting in two distinct drop-out orbiting areas. The orbit nearer to the transmitting GS (labeled as "DLvideo" in Figure 12) confirmed the signal range for the downlinked video stream. The orbit farther from the transmitting GS (labeled as "DLC2") confirmed the signal range for the downlinked command and control data stream.

The air corridor between Cleveland, OH and Albany, OH was generally clear of air traffic, minimizing the need for en route flight path adjustments. The underlying terrain features and elevations along this path had been reviewed and modeled to determine necessary flight altitudes. Modeling and test flights confirmed that aircraft altitude of approximately 13 000 feet above mean sea level (MSL) would provide unobstructed LOS between the GS antennas and the aircraft over the entire flight path span. Weather and air traffic issues had only minor effects on the flights. This flight path shown in figure 13 was used repeatedly during the Phase II campaign.



Figure 11. - Aircraft flight track for characterization of ground station 1. Flight track created using GPS Visualizer [8], Map data from the National Atlas [10]



Figure 12. - Aircraft flight track for characterization of ground station 2. Flight track created using GPS Visualizer [8], Map data from the National Atlas [10]



Figure 13: Ground Station 1 - to - Ground Station 2 flight track. Flight track created using GPS Visualizer [8], Map data from the National Atlas [10]

Flight Test Results

April 11, 2014

The objective of the first flight test was to examine the general radiofrequency characteristics of the CNPC link between the S-3B and GRC ground terminal GS1. Flying the flight path shown in Figure 11, investigators could observe the general capabilities of the CNPC radio pair operating with L-band omnidirectional antennas. Figure 14 displays data for the entire flight period, from aircraft take-off through landing, with time annotated along the horizontal axis. Received signal strengths at the aircraft radio and at the GS1 radio are plotted along the vertical axis. Data traces for narrowband (75 kHz), wideband (875 kHz), uplink, and downlink combinations are shown to illustrate the similarity in channel performance. Taking off from Cleveland Hopkins International Airport (airport designation: KCLE), the aircraft ascended above the local ground obstructions (buildings) allowing the received signal strengths to reach their maximum level at the “spikes” indicated at 15:54 UTC. The signal strength curves quickly dropped to the -85 to -100 dBm range where it remained for approximately 30 minutes during the 25 nmi circular orbit.

The next four traces in figure 14 present data on the percentage of frame loss averaged over one-second at the ground and aircraft receivers. Individual traces for the UL1, UL20, DLC2, and DLvideo channels are shown. Where the CNPC communications path was transferring all data without error, the data are presented as 0-percent loss and no trace is visible on the grid. Where errors occurred in the radio link, the lost frame data created a visible trace ranging from 0 to 100 percent, with the latter representing a total loss of the radio link and clearly indicated by the heavy colored sections in the figure.

With the exception of a blockage period near time 16:06, the observed CNPC link streams encountered minimal frame losses during the 25 nmi circular orbit, the outbound flight towards GS2, and the inbound flight towards GS1. Similar to previous flight tests, the UL20 and DLvideo channels were lost near 70 nmi slant range and the UL1 and DLC2 channels were lost at a slightly greater distance.

The peaks, trends, and other perturbations displayed in the curves in Figure 14 can be explained by aircraft position, aircraft attitude, ground obstructions, antenna patterns, slant range, and other factors. Similar data traces are analyzed extensively in [1.] For this report the goal of Figure 14 is to display the overall similarity of the traces to one another, indicating both the uplink/downlink channel reciprocity and the narrowband/wideband (UL1/UL20) performance similarity.




Figure 14. - Flight Test 1 radio performance data and associated aircraft parameters at Glenn Research Center on April 11, 2014
GS, Ground Station; UTC, Coordinated Universal Time

One notable difference between the four traces in Figure 14 is the presence of multiple, closely-spaced “spikes” in the UL1 and UL20 uplink data traces. These are most visible in the UL1 trace during the period of 17:03 through 17:29. While the S-3B aircraft was flying at a range of approximately 70 nmi, the signal strength received at the aircraft radio was nearing its sensitivity limit near -119 dBm. While the underlying signal strength trace appears to trend further downward, the amplitudes of the spikes appear to be staying relatively constant. The April 11, 2014 flight is the first test where this effect has been recorded; previous tests have not displayed this effect. The spikes are not visible in the DLC2 or DLvideo “downlink” traces recorded by ground station radio.

Subsequent investigation revealed that the automatic dependent surveillance-broadcast (ADS-B) transponder and other equipment on the S-3B aircraft had been changed prior to the Phase II test series. The new ADS-B unit, equipped with an extended squitter transmitter at 1090 MHz, was suspected to be a probable cause of the spikes. Although the transponder is designed to operate outside of the 960-977 MHz CNPC radio band, it is apparent that the transponder transmissions are being received by the CNPC radio on the aircraft.

April 23, 2014

To investigate the possible source of the spikes, a flight test was planned in which the ADS-B transponder would be briefly turned off during flight. Radio operating conditions were identical to those of the April 11, 2014 test, including radio configurations, flight path, and downrange distance. In coordination with local air traffic controllers, the transponder was de-powered for approximately five minutes during the outbound pass and then reactivated for reversal and the inbound pass. Data for the test is presented in figure 15, showing the UL1 signal from ground-to-aircraft and the DLC2 signal from aircraft-to-ground.

The airborne ADS-B transponder was in operation from the time of take-off until the S-3B was approximately 80 nmi downrange from the GRC ground station. At this time the received signal strength at the aircraft was approximately -103 dBm which is above the sensitivity limit of the radio and at a level that the spikes would become visible in the data trace. From 14:48 to 14:53 the transponder output was deactivated while the aircraft continued on a straight and level flight path, extending to approximately 110 nmi downrange. The data trace in figure 15 clearly shows that the spikes have abated throughout the five-minute period, especially when compared to the same airspace and flight conditions on the inbound run (15:08 through 15:13) when the ADS-B transponder

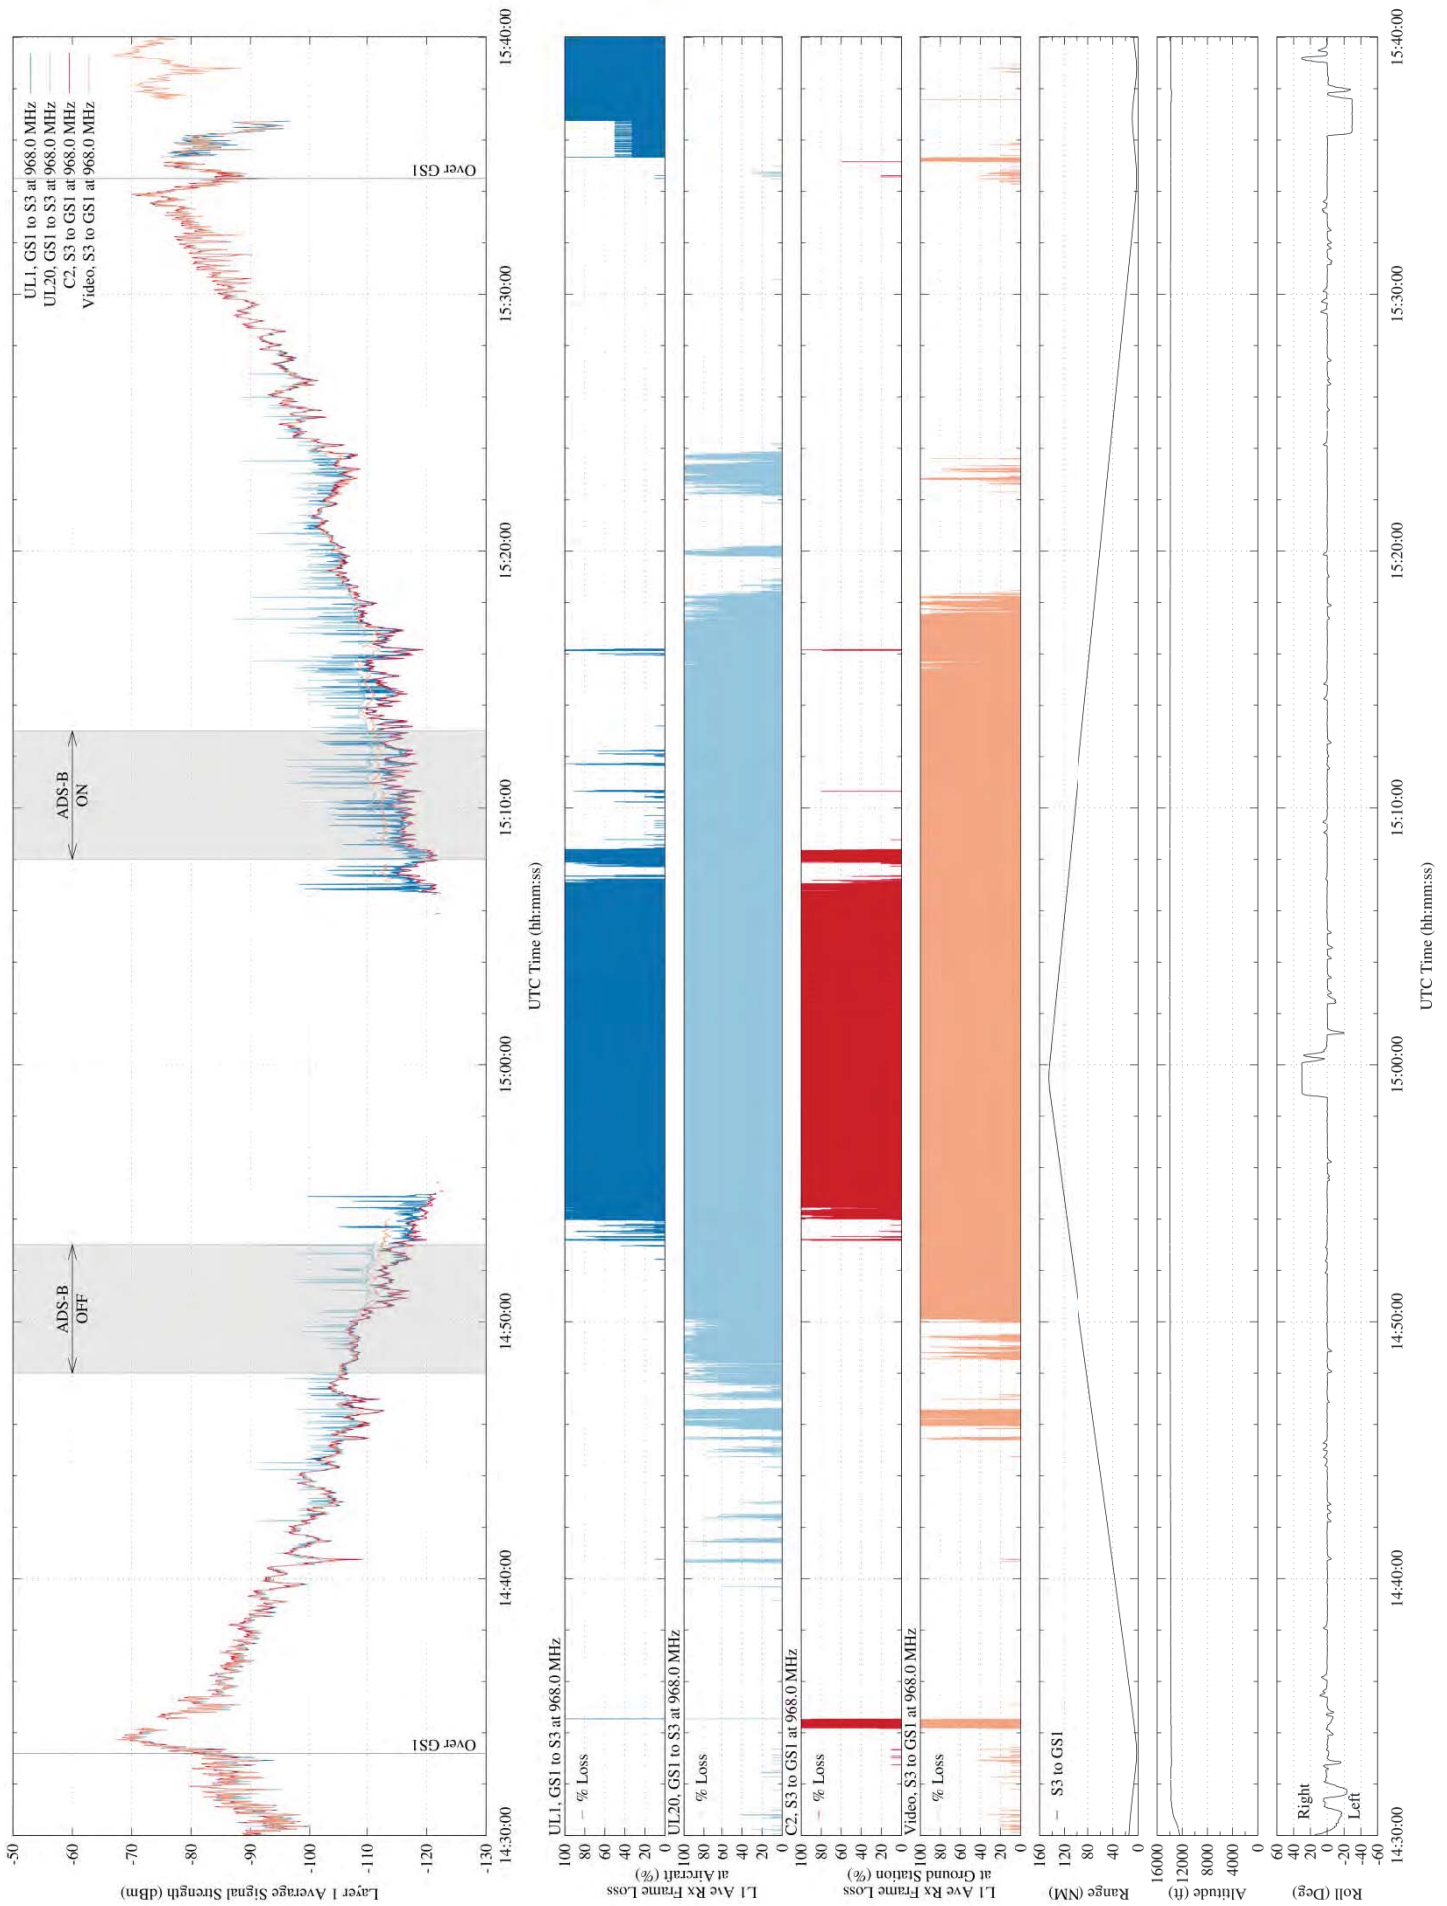


Figure 15. - Flight Test 5 radio performance data and associated aircraft parameters at Glenn Research Center on April 23, 2014
 GS, Ground Station; UTC, Coordinated Universal Time; ADS-B, Automatic Dependent Surveillance-Broadcast

was back in operation. The ADS-B deactivation appears to have greater impact on the UL1 trace, as it shows greater reduction in number of spikes. The wider bandwidth UL20 signal does not experience the same cleansing during the ADS-B off period, indicating that some other interferer is still present. Noise spikes were not visible in the downlink signals received by GS1. The overlap of the uplink and downlink data traces indicates the strong similarity of the characteristics of the channels, and also gives a forecast of how the uplink signal should be received once the interfering signal is removed.

The outbound portion of the flight test, beginning at 14:33 UTC, yielded high quality data transmission (0-percent losses) both ground-to-aircraft and aircraft-to-ground until a few spikes were recorded at the aircraft beginning at approximately 14:46. The ADS-B transponder was de-activated at 14:48 which caused the spikes in the UL1 trace to disappear. The ADS-B transponder was re-activated at 14:53, at which time the UL1 spikes returned. The UL1 frame loss trace shows signal deterioration prior to the onset of DLC2 frame loss. This indicates that the interfering signal is having negative impact on the uplink CNPC channel.

On the inbound portion of the flight, the impact of the interfering signal is visible in the UL1 frame loss trace from 15:08 through 15:13. The multitude of vertical lines indicates that the radio is experiencing frame losses. Over the same time period, the DLC2 frame losses are much lower, resulting in little visible trace on the grid.

The ADS-B interference issue described in this section and other co-site effects will be examined in future test activities. Factors including antenna placement, RF filtering, and radio waveform configuration/synchronization will be investigated in subsequent tests.

April 30, 2014

The primary objective of the early flight tests and one reason for characterization of the individual ground stations was to establish an environment to support networking and GS-to-GS hand-off testing. To perform the networking and hand-off tests, the aircraft would need to operate in a region where radio contact could be simultaneously made with both ground stations, preferably allowing the radio links to operate near the edge of their coverage range. The hand-off region for the Phase II tests was established as a sub-segment of the KCLE-KUNI flight path.

Figure 16 presents radio performance data for the complete Phase II flight test system, including both ground stations and the S-3B aircraft. Received signal strength traces

show the UL1 and DLC2 channels to and from both GSs as the aircraft flies along the KCLE - KUNI flight path. Reading the figure from left to right from time 14:30 represents the S-3B aircraft traveling northbound from overhead of GS2 at the KUNI airport. As expected, the received signal strength recorded at GS2 and on the aircraft decrease as the aircraft moves northward away from GS2 and the slant range between aircraft and ground station increases. In the same period, the signals between the aircraft and GS1 begin to increase in strength as that slant range distance decreases. Near time 14:40 on the plot, the signals from the two ground terminals are nearly equal at approximately -105 dBm. Further progression to the right on the plot shows the received signals between the aircraft and now-distant GS2 becoming lower in power until the CNPC link can no longer communicate without data errors. No data errors are measured from the aircraft-GS1 CNPC links as the aircraft progresses northward and overflies GS1.

The region of the flight path during which both ground stations were communicating with the aircraft without significant data errors is considered appropriate for the networking and hand-off testing. The first “ideal” test region occurs between times 14:35 and 14:44 in Figure 16, as the aircraft was traveling northbound from GS2 towards GS1. This region is readily observed as the period when the signal strength curves from the two GSs cross, and when the frame error percentages are zero or very low. Two other test periods were achieved when the aircraft repeated southbound and northbound runs on the KCLE-KUNI flight path; nominally 15:03 to 15:19 and 15:38 to 15:47. Test periods at 14,000 feet altitude lasted approximately 9 minutes during northbound flights and approximately 16 minutes during southbound flights. The difference in duration of the test periods was influenced primarily by differences in ground speed caused by winds aloft. Groundspeed during each of the flight segments was nominally 349 knots (northbound #1), 260 knots (southbound), and 344 knots (northbound #2). Network testing periods were often extended beyond the ideal periods in order to measure radio performance in the presence of frame errors.

The ADS-B extended squitter transponder was in operation during this flight test.

May 28, 2014

The May 28 flight provided the same test environment and flight conditions as in April 30, 2014 flight, offering constant speed and level flight along the GS1-GS2 route. The hand-off testing was conducted during the period when the CNPC communications path was established with both GSs, ranging from 16:28 to 16:42. The radio performance data for the hand-off test period is displayed in Figure 17.

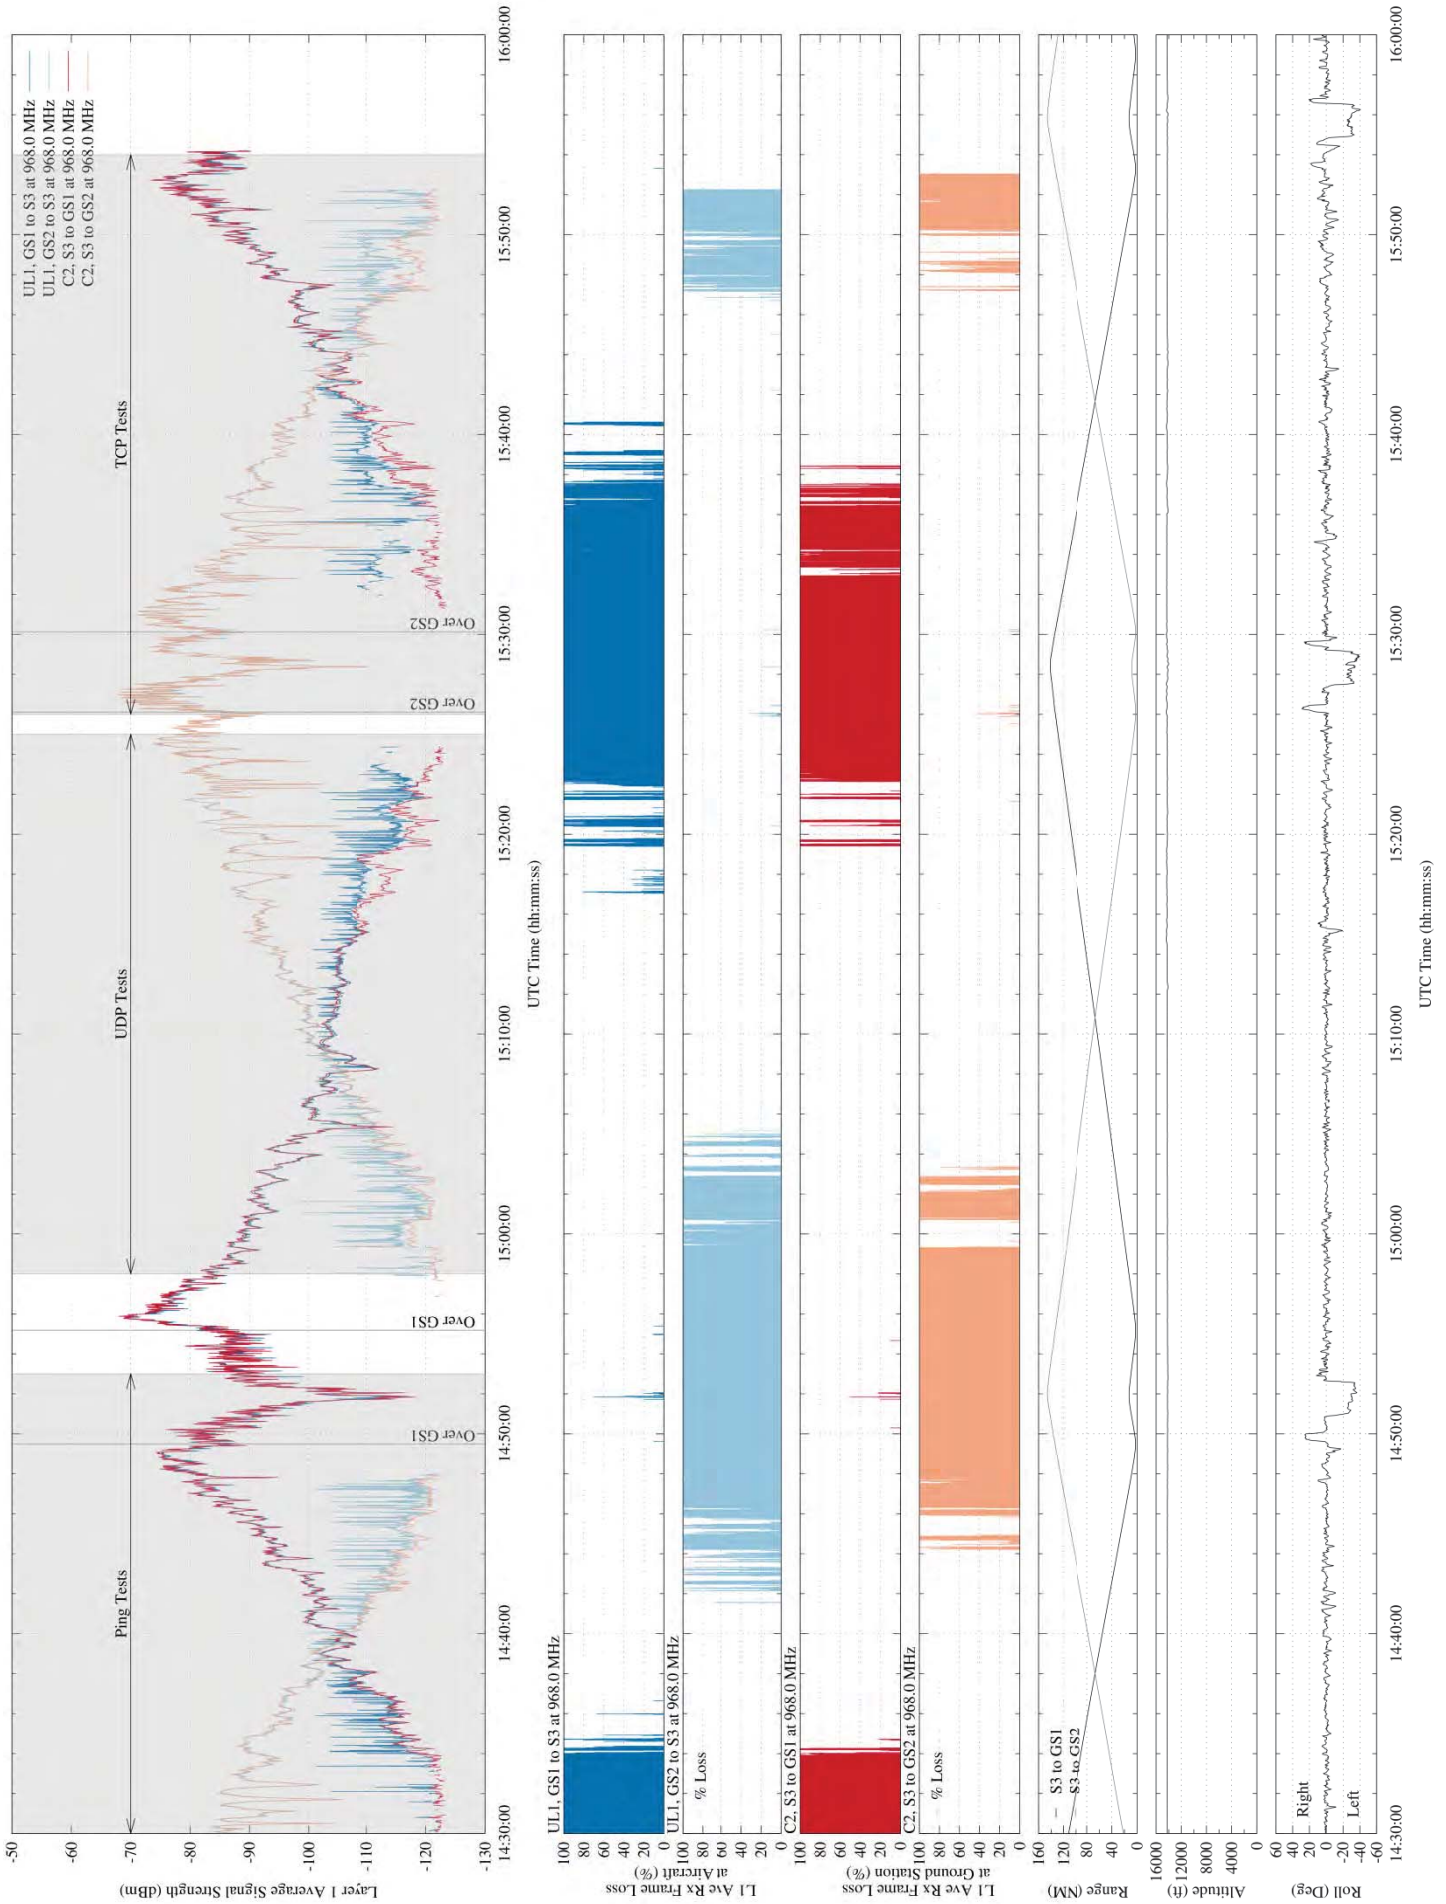


Figure 16. - Flight Test 6 radio performance data and associated aircraft parameters at Glenn Research Center and Ohio University Airport on April 30, 2014
 GS, Ground Station; UTC, Coordinated Universal Time

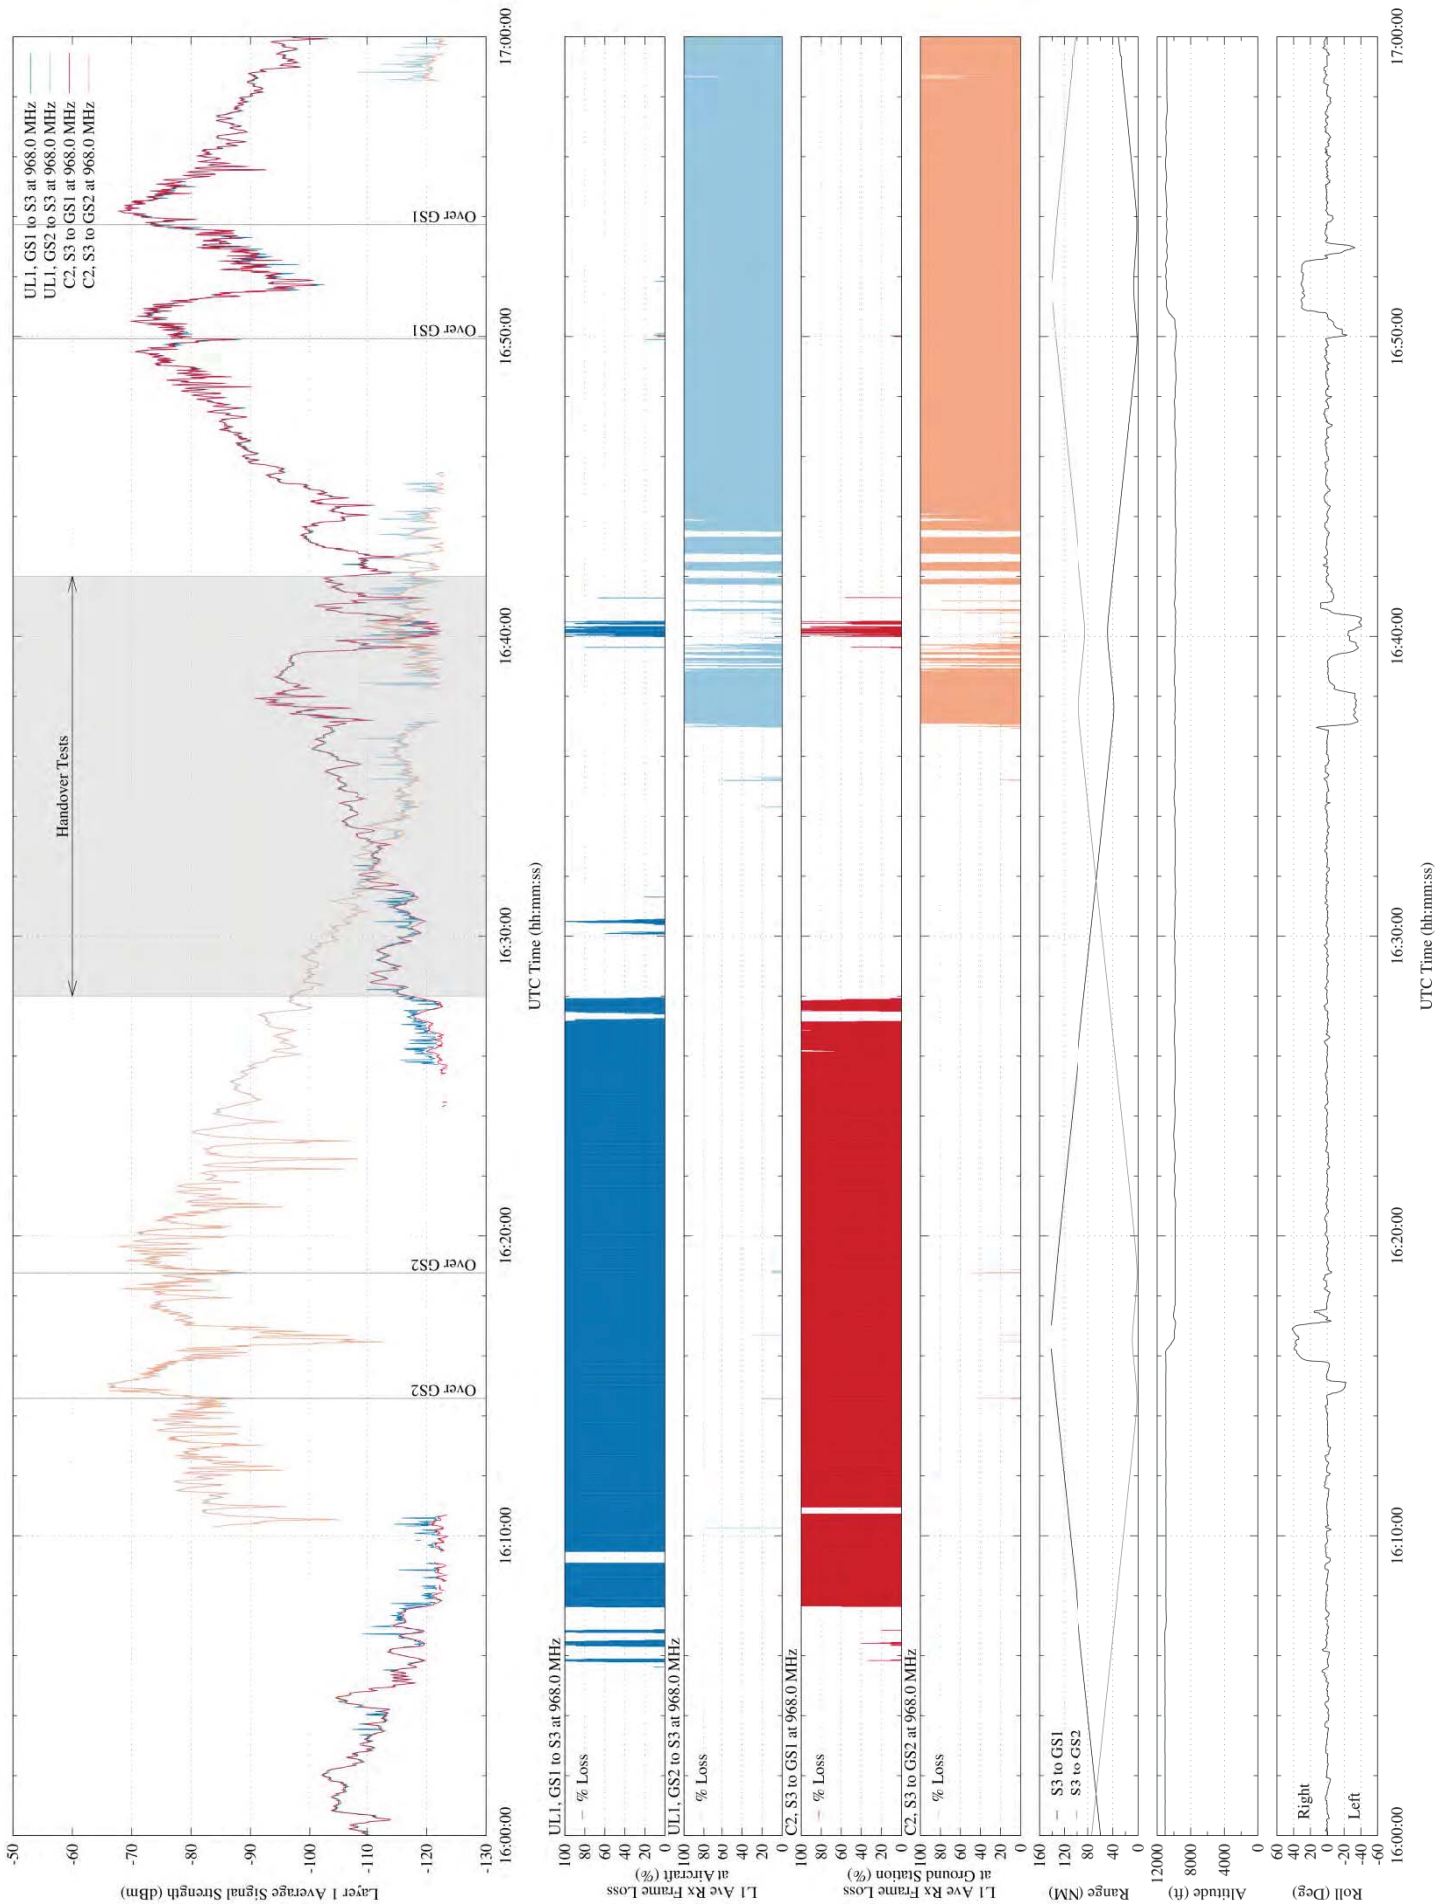


Figure 17. - Flight Test 7 radio performance data and associated aircraft parameters at Glenn Research Center and Ohio University Airport on May 28, 2014
GS, Ground Station; UTC, Coordinated Universal Time

Networking Tests and Results

Using the network configuration and architecture described earlier, tests were constructed to ascertain actual performance metrics of basic network statistics while in flight and with the CNPC radio system in operation. In addition, tests were performed to evaluate the ability to do mobile hand-offs while in range of both ground stations. Since these tests were done in flight and outside of the controlled laboratory setting, the results are subject to variations in the environment and operational conditions. Thus, it is important to reference the appropriate flight plan when discussing these results. Times discussed in the networking test results will directly correlate to figures found earlier in this document.

Delay and Jitter

One of the basic goals in these networking tests was to measure the time delay and delay variation (jitter) through the prototype radios while in a networked configuration. In order to obtain these values Internet Control Message Protocol (ICMP) echo requests, or “pings”, were sent directly from each GS to the aircraft. The aircraft-based test electronics would then reply to each request, sending the response to the appropriate GS. The total time between the initial request and its corresponding response, as measured at the GS, is the round trip time (RTT). Since the pings are generated directly at the GS, this test helped to isolate the delay through the radio as no other network queues or nodes were involved in the path. Figure 18 shows the components involved in this test. Measurements were made by capturing the ICMP traffic on the ground station interface to the radio.

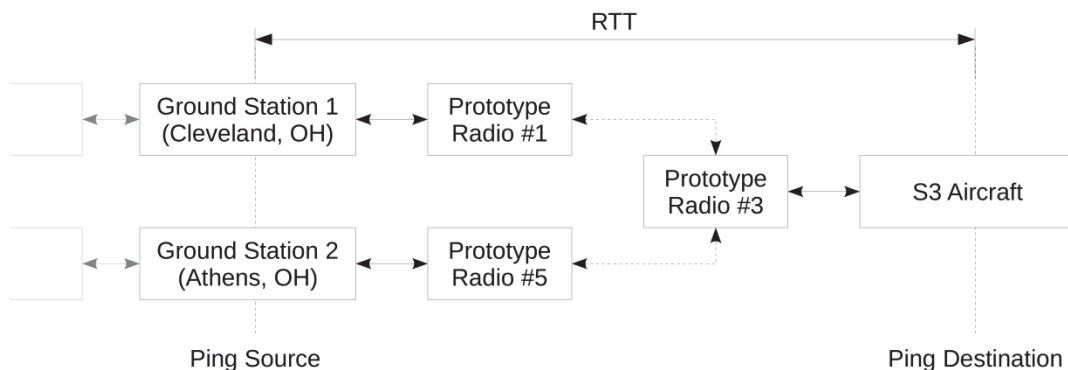


Figure 18. – Diagram showing round-trip path of ping message through CNPC flight test system.

The data for this test was collected during a segment of flight on April 30, 2014. During this time the aircraft was returning from GS2 in Albany, Ohio and heading towards GS1 in Cleveland. The UL1/C2 waveforms were used for these tests and were configured for a transmission rate of 10 Hertz. Testing consisted of sending ten ICMP echo requests every second from each ground station. Ten requests were sent every second in order to pair up with the 10 Hertz rate of the waveform. The payload size of each request was 20 bytes. With packet overheads of 8 bytes for ICMP and 40 bytes for IPv6, the total IP datagram size was 68 bytes, which is less than the allotted 103 bytes of user data allowed for each transmission event. However, the goal of these tests was not to saturate the link with data. Rather, it was to observe the latency and variations in delay if a user utilized a message frequency closely matching that of the underlying waveform.

Network mobility was not used for these tests. The echo requests were sent directly to the appropriate care-of address of the aircraft. This allowed for testing to measure the delays over just the radio portion of the link. Round trip times were calculated by processing a time-stamped packet trace at the ground station interface. Each request is tracked and when the corresponding response is seen in the trace the RTT is calculated by taking the difference in the two times. The time of each RTT sample is recorded with the time of the request. If a response was not received at the ground station, then the ping is marked as lost. Lost pings are not incorporated into the calculation of RTT metrics and simply reduce the number of samples in the time period the loss occurs.

The test had duration of roughly 23 minutes in which thousands of pings were transmitted. In order to better visualize the results, data was grouped into ten second ranges in order to obtain roughly one hundred samples per data point. This approach provided a good statistical basis for each data point while still allowing the metrics to be visualized over time and compared to the RF channel if necessary. For example, a marker at time 14:40:00 contains data from 14:39:55.00 to 14:40:04.99. The exact number of data points in each sample is shown in Figures 19 and 20, along with the number of losses encountered. Results falling just short of the expected 100 samples per data point will be discussed later in this section. The other major dips in the number of samples, such as the entry at 14:39:50, are due to briefly stopping and restarting the pings during the anticipated time a hand-off would normally occur.




Figure 19. -Number of Total Samples/Losses at the Glenn Research Center ground station.

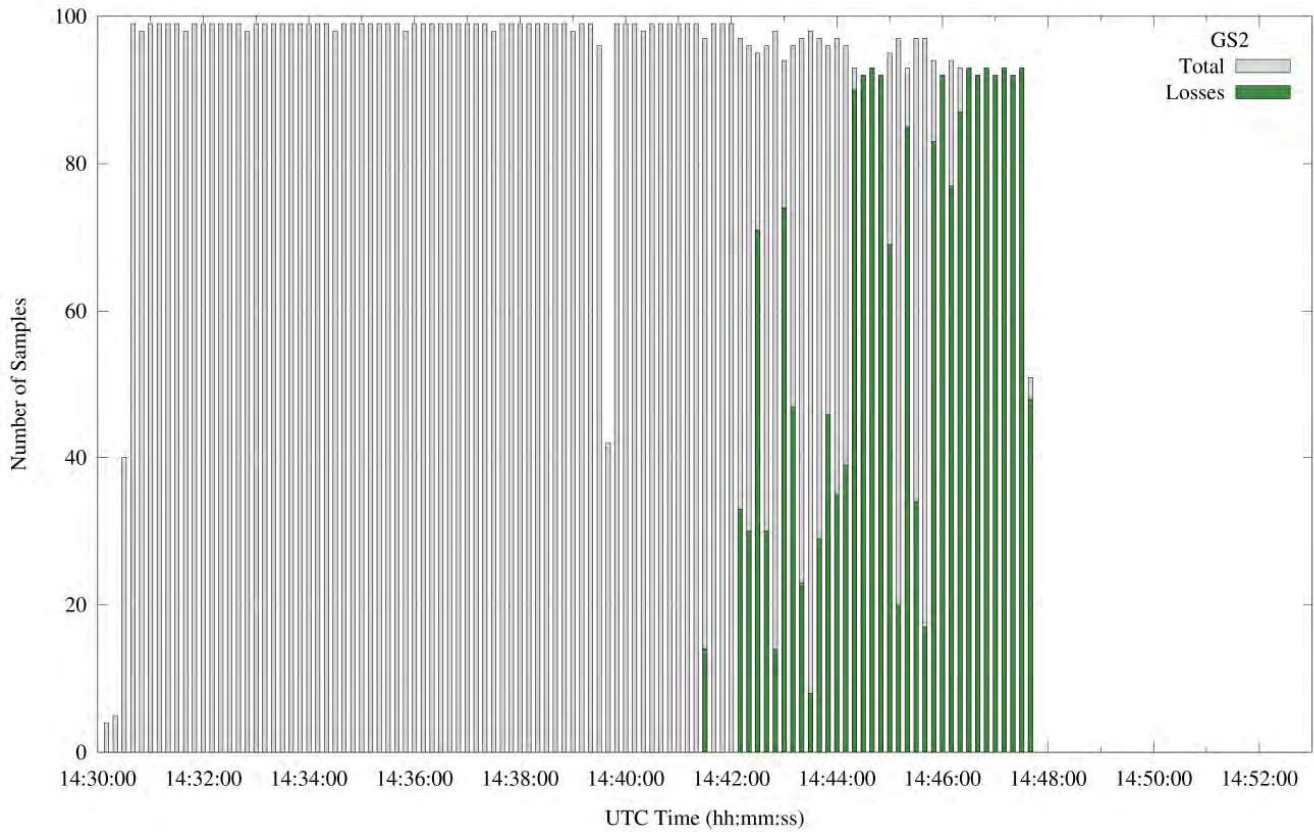


Figure 20. -Number of Total Samples/Losses at the Ohio University Airport ground station.

Figures 21 and 22 show the calculated round trip times at both ground stations by displaying a box-and-whisker plot of the samples for each data point. The darker “box” portion of each data set represents the first (Q_1) and third (Q_3) quartiles of the sample set. It is sometimes helpful to think of this as the “middle 50%” or “middle half.” The line within the box is the statistical median of the data. The “whiskers” of the plot extend to the farthest data sample that is not an outlier. For these tests, an outlier is defined as any point more than 1.5 times the interquartile range (Q_3-Q_1) from the end of the “box”. Each outlier is shown as a dot on the plot.

A few immediate conclusions can be drawn from figures 21 and 22. First, there are a large number of outliers which are often grouped together in clusters. It will be shown that the cause of these outliers is due to periodic IP traffic that is concurrently transmitted with the ping data. Second, no notable correlation is observed between RTT and time at a macroscopic level. Thus, as the radiofrequency link degrades or improves, there does not seem to be an impact on RTT. So, while figures 19 and 20 show an increase in lost pings as the link sufficiently degrades, the RTT of the pings that do succeed remain consistently with a range of 200 to 250 milliseconds.

With no apparent dependence on the RF link quality, the performance differences between ground stations were explored. To do so, all RTT samples were collected at each location and a cumulative distribution function (CDF) was generated. Figure 23 shows the CDF for each ground station. From the plot, it is shown that the distribution of ping travel times at each ground station is nearly identical. It is also confirmed that a majority of the data falls within 200 and 250 milliseconds, with 95% of all data having a RTT less than 282 ms. The outliers from earlier are visible producing the long tail at the top of the plot. Furthermore, there is a notable “plateau” between 190 and 200 ms, indicating that few samples fall into that range and warrant further analysis.

Due to the volume of data, analysis was based on aggregate ten second groups. In order to understand some of the observations, it was necessary to look closer at smaller portions of the data. Figure 24 shows individual RTT measurements as a function of time for a single minute of measurements. The figure shows a clear “saw-tooth” pattern in the behavior of the RTT delay. Furthermore, there are distinct spikes at various intervals throughout the data. The spikes correlate directly with ICMP6 router advertisement (RA) being sent during the test. The time of each RA is shown as a vertical dashed line on the figure. These advertisements are 128 bytes long and delay the pings following it as the radio tries to handle the data spike. Thus, the outliers in the data can be attributed to the periodic transmission of RA packets across the link.




Figure 21. - Round trip times to the Glenn Research Center ground station.




Figure 22. - Round trip times to the Ohio University Airport ground station.

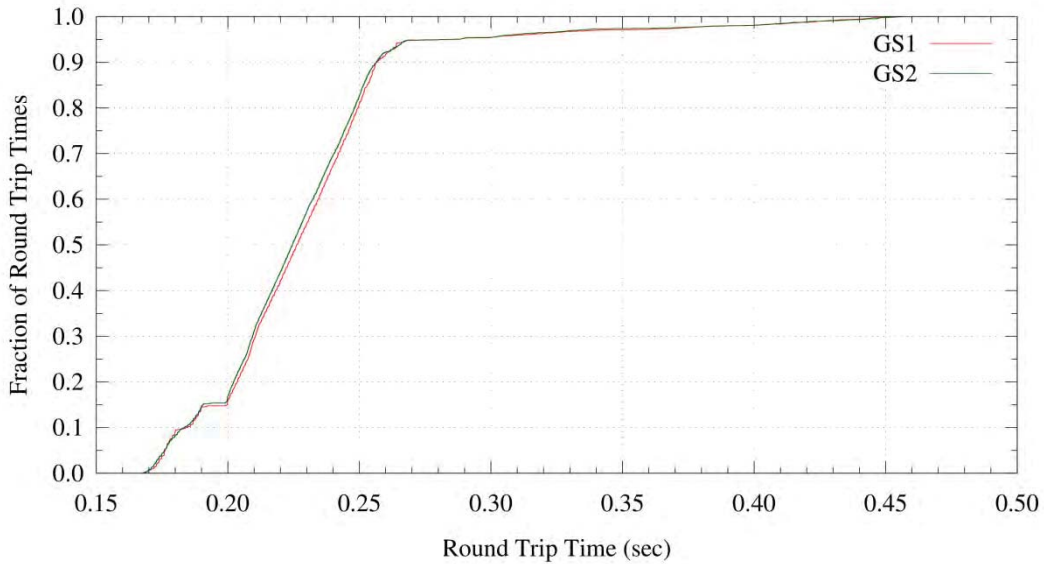


Figure 23. - Cumulative Distribution Function of Round Trip Time Samples.

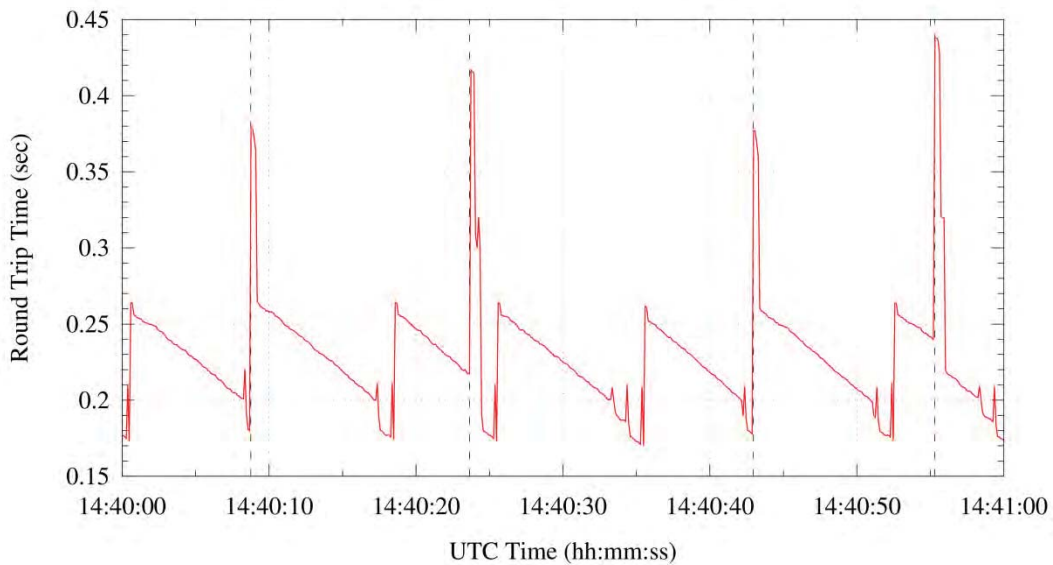


Figure 24. - RTT as a Function of Time at Ground Station 1.

Explanation of the trace’s saw-tooth behavior required analysis of the inter-arrival time of each ICMP echo request. Figure 25 shows a CDF of the delay between each ICMP echo request. The figure shows that although the test had initiated ten ping requests every second, a majority of pings actually took one to eight milliseconds longer than the expected 100 ms. This observation explains why there were less than 100 samples over ten seconds. While this extra delay may seem insignificant, it is suspected to be the cause of the cyclical RTT pattern.

Within the CNPC radio, data is queued and made ready for transmission before a new transmit window starts, otherwise it must wait until the start of the next window. The time spent in the queue, T_Q , impacts the overall round trip time metric. In the

configuration for this test, each transmit window was 100 ms long. Figure 26, shows the impact a few extra milliseconds can do the value of T_Q . Since the time between pings (period) is slightly longer than the transmission window, the value of T_Q will slowly decrease causing the RTT to also decrease. Eventually this leads to a ping request that arrives just prior to the transmission window. The resulting RTT is the “ideal frame” shown in the figure and has the lowest value. However, the very next request will “just miss” the next transmission event, leading to the worst possible value of T_Q and thus the longest delay.

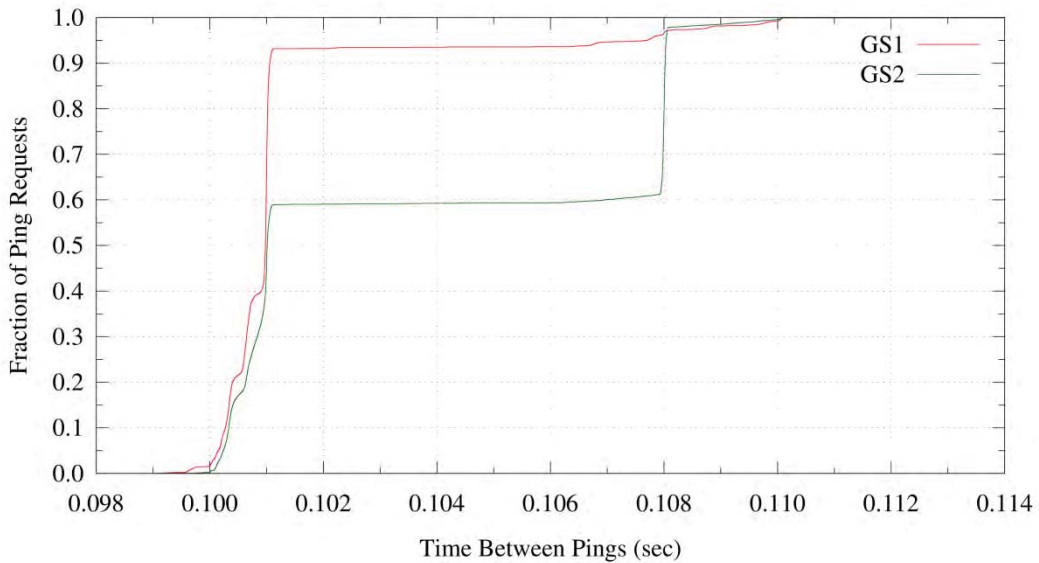


Figure 25. - Cumulative Distribution Function of Time between Ping Requests.

In Figure 24, as the RTT decreases, miniature spikes in the RTT can be observed. The erratic behavior is near 190 ms and is the reason behind the plateau observed in Figure 23. It is suspected that this effect is caused by an implementation detail within the radio. Specifically, a ping cannot truly arrive just prior to the transmission time. Rather, there is a processing delay that may vary and negatively impact some requests as T_Q is reduced towards zero.

Test data shows that the delay through the radio link is consistent over different RF characteristics. For the 10 Hertz update rate configuration, the median round trip delay through the CNPC link was 225 ms. Excluding outliers caused by other traffic, 99% of the RTT varied from 171 to 267 ms. The delay variation or jitter is due to messages being held in the radio's queue for different durations based on the messages' arrival time at the radio. This range of 96 ms is consistent with the 100 ms between transmission events in the waveform configuration that was applied at the time of the test. Thus, the amount of jitter is dependent on the configuration of the radio.

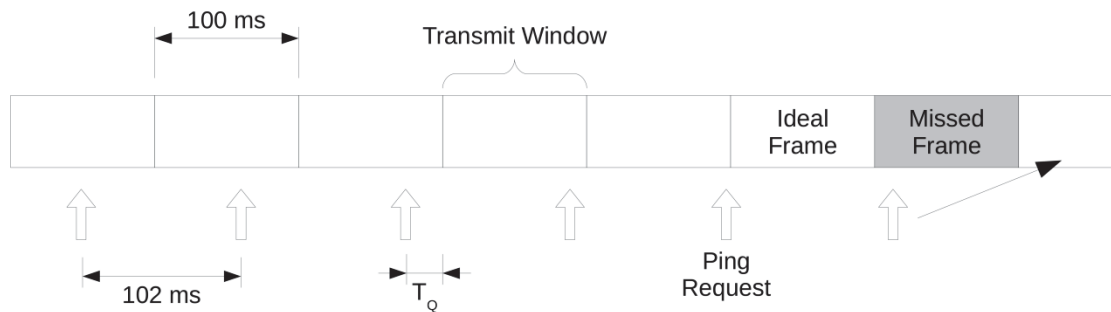


Figure 26. - Impact of Radio Queuing Delay for 10 Hz rate

TCP and UDP Testing

On April 30, 2014, after the completion of the time delay and jitter testing, a series of investigative tests were conducted. Network traffic was transmitted between the ground station and the aircraft using both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) transport formats. The tests were performed to provide an initial assessment of the throughput capacity of the radio pairs. The tools used in these tests were based on open-source utilities and were adequate for the initial assessment. However, additional software will need to be developed to complement those tools and extract additional metrics. The resulting testing suite will be utilized for upcoming NASA tests, which will measure network throughput capacity in the flight environment and study the effects of header compression and data security.

Hand-off Response Time and Loss

On May 28, 2014, a series of network hand-off tests were performed in order to characterize the performance of network mobility on the current radio waveform. In these tests, the aircraft flew to an area that was within range of both ground stations. The test began by first establishing a link to ground station 2. This active link would be considered the primary connection until an operator on the plane initiated a manual hand-off. Each manual hand-off consisted of first establishing a connection to the inactive ground station and then manually disabling the primary connection. The mobile router would detect that the primary link had failed and hand-off to the newly created connection. This procedure was repeated a total of 25 times as the aircraft attempted to remain between ground sites.

Figure 27 shows the three steps of the hand-off procedure. Even though the process first establishes the new link, the hand-off is “break-before-make” as the rebinding is done after the active link is destroyed. This will lead to some data loss and is considered the worst case for a hand-off when another link is ready and available. Even as hand-offs become more automated, such conditions may arise if an impending loss of link event happens very suddenly, such as severe shadowing. The figure also shows how these ping tests differ from those performed at the ground stations. During this test, ICMP echo requests are sent continuously from the HA to the aircraft. The requests are never paused and are dynamically re-routed through the correct path as the hand-offs occur. The frequency of ping requests matches that of the earlier test at ten requests every second. However, the size of each ping packet has increased due to the extra headers required for mobility. The end result is a ping packet whose total size is 93 bytes, slightly larger than the 68 bytes used in the previous test.

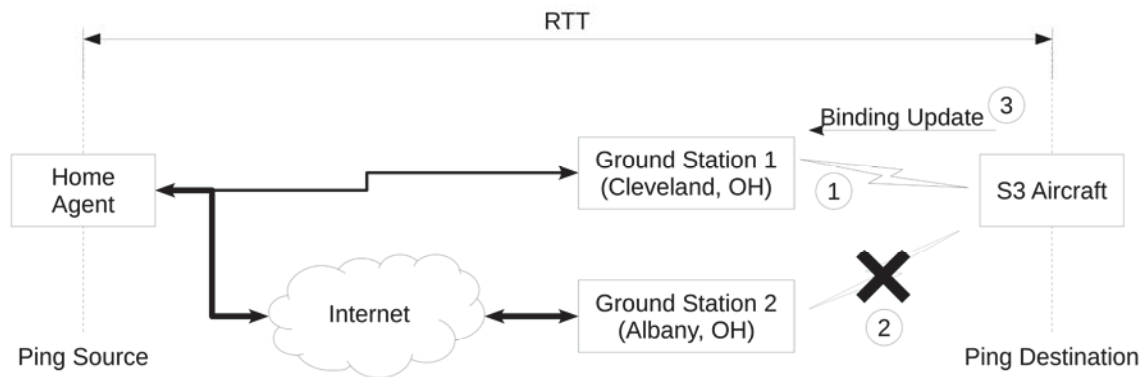
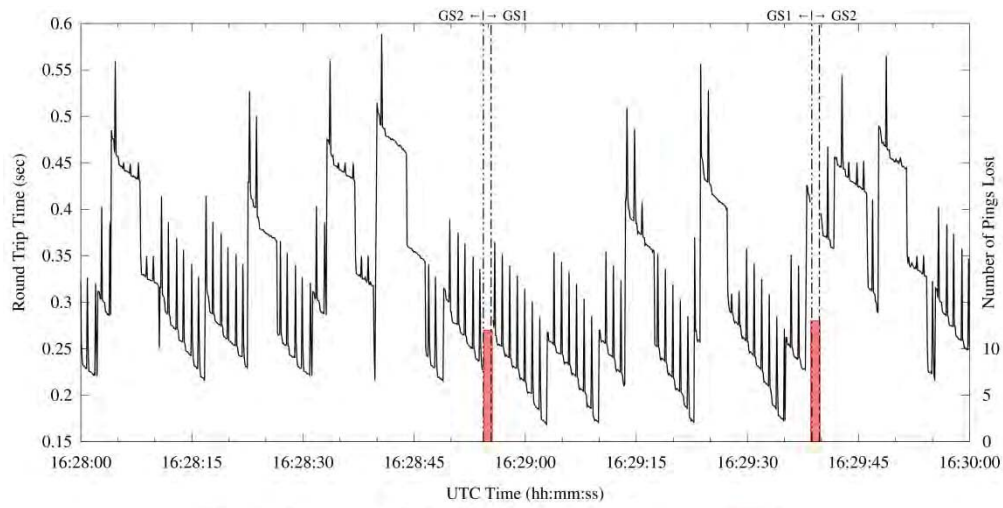



Figure 27. - Testing Points and Break-Before-Make Hand-off Process:
 (1) New Link is Established, (2) Old Link is Destroyed, (3) Binding Update is Issued.

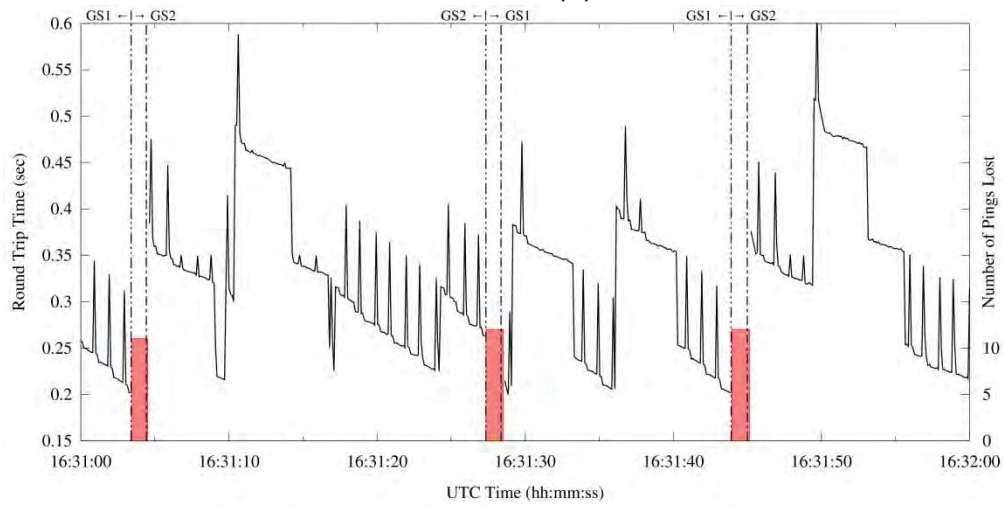
Figure 28 shows the first four minutes of this experiment and the resulting six hand-offs that were performed in that time. In these plots, the round trip time of the pings, as observed by the home agent, are plotted against time. The plots also show when the primary connection was terminated at the aircraft. These termination points are indicated by the double dotted dashed lines, and the labels above these lines indicate which ground station was acting as the primary link. Binding updated used by the mobile node to establish the new link to the home agent are shown with a single dotted dashed line. Pings that do not get a response at the home agent are considered lost. These losses are plotted as boxes on the figure, with the width of the box indicating the duration of the loss and the height of the box indicating how many consecutive pings were lost in that duration. The number of lost pings is plotted against a secondary y-axis in order to visualize both metrics on a single figure.



(a)



(b)



(c)

Figure 28. – Hand-off Timing and Signal Losses

Just as in the earlier ping tests, there are large plateaus in the round trip time caused by router advertisements being sent from the ground station. Also, the same overall sawtooth pattern, where the RTT slowly decreases, can be observed in all three plots. However, unlike the earlier tests, there are additional spikes along the sawtooth pattern, which are extremely consistent throughout the experiment and occur every second. The source of these spikes was traced back to a radio message that carries RF health information and is transmitted across the link every second. The radios transmit this health information in order to establish a database of towers and aircraft between radios. Since the queuing priority of these messages is higher than that of any user data, under loaded conditions it will induce extra delay into any user data. These spikes were not seen in the earlier experiment as the size of the ping packets was smaller. This allowed for both the health message and the ping packet to fit into a single transmission event. However, the size of ping packets increased to 93 bytes for these tests. Thus, when contention occurs with the health data, only a fraction of the ping packet can be transmitted. The remaining portion is delayed a transmission window adding 100 milliseconds to the round trip time.

Similar to the earlier ping tests, the observed round trip times fall into the same range. For example, the same 171 to 267ms sawtooth is observed for RTT values in the GS1 segments. This is due to their being very little latency between GS1 and the home agent as they are connected over a local area network. As expected, the RTT values at GS2 are 50 to 60 ms higher than those at GS1. This is directly attributed to the extra delay of traversing the terrestrial network between Cleveland and Albany, Ohio.

Figure 28 (a) and (c) show the first two minutes and fourth minute of the experiment, respectively. In this time, five “typical” hand-offs occur. In total, 23 of the 25 hand-offs follow the behavior shown in the two subfigures. The plots show that after the link is terminated, there is an amount of time before the mobile router reacts to the loss and initiates the first binding update to the home agent via the standby connection to the other ground station. In the 25 hand-off requests, the average time taken by the mobile router to initiate this first binding update was 1.054 seconds, with a standard deviation of 0.042 seconds. It is important to note that no special steps were taken to tune this response time for these tests.

Ping responses resume roughly one round trip time after the binding update is issued from the mobile router. This follows expectations as it is the amount of time needed for the binding update to reach the home agent, and for the home agent to send an acknowledgment confirming the change. Therefore, the total time to recover should be roughly $1.054 + \text{RTT}$ seconds. Table 1 shows the statistics for the total time to recover for all of the 23 “typical” hand-offs. It shows that the average time of 1.353 seconds

matches the expected result and is consistent with the round trip times seen earlier in the delay tests. The table also shows the average number of losses encountered. Since pings are sent ten times each second, it is reasonable that the number of packet losses is ten times the time needed to recover.

# Samples N = 23	Median	IQR	Mean	Std. Deviation
Time to Recover (seconds)	1.268	0.199	1.353	0.439
Number of Losses (packets)	13.0	2.0	13.7	4.1

Table 1: Hand-off Times

Figure 28 (b) shows one of two hand-offs that encountered losses during the start of the event. Looking back at the RF data (Figure 17), there is a loss event just after the 16:30:00 time marker that primarily impacts the uplink portion of the link. This implies that data sent to the aircraft was dropped, while data to the ground was not impacted. In this hand-off, the time needed to recover is much larger and longer and is due to the quality of the underlying RF channel. Figure 29 provides the actual packet trace as observed from the aircraft during this hand-off event. The trace spans 3.75 seconds, combines traffic from both ground stations, and highlights binding updates and acknowledgments in bold text. It begins with pings through GS2, identified by the use of the care of address, S3-COA-GS2. The link to GS2 is disabled 21.652 seconds into the trace, and the first binding update, with sequence number 42523, is sent out roughly one second later over the link to GS1. The home agent actually receives this update as echo requests begin to arrive over the GS1 care of address, S3-COA-GS1. However, the aircraft does not reply to these pings since it has not yet received an acknowledgment of the change from the home agent, which was lost due to the issues in the uplink RF path. During this time, there are 14 pings that arrive and do not get routed properly. Eventually, the mobile router retransmits another binding update which is successfully acknowledged and the echo replies resume as normal. Thus, the loss event was prolonged due to a critical loss of the binding acknowledgment. Figure 28 (b) shows that losses continued after the acknowledgment arrived. Those losses correlate with the larger of the two loss spikes in the RF plot.


```

21.551030 IP6 HA > S3-COA-GS2: IP6 HA > S3: ICMP6, echo request, seq 4078
21.551065 IP6 S3-COA-GS2 > HA: IP6 S3 > HA: ICMP6, echo reply, seq 4078
21.651806 IP6 HA > S3-COA-GS2: IP6 HA > S3: ICMP6, echo request, seq 4079
21.651844 IP6 S3-COA-GS2 > HA: IP6 S3 > HA: ICMP6, echo reply, seq 4079
22.730532 IP6 S3-COA-GS1 > HA: DSTOPT mobility: BU seq#=42523 AH
23.101219 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4093
23.201808 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4094
23.301144 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4095
23.401074 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4096
23.501922 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4097
23.801388 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4099
24.001369 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4101
24.101154 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4102
24.401343 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4105
24.501082 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4106
24.601722 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4107
24.700982 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4108
24.730623 IP6 S3-COA-GS1 > HA: DSTOPT mobility: BU seq#=42524 AH
24.801277 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4109
24.900965 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4110
25.001046 IP6 HA > S3-COA-GS1: srcrt mobility: BA status=0 seq#=42524
25.101001 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4111
25.101059 IP6 S3-COA-GS1 > HA: IP6 S3 > HA: ICMP6, echo reply, seq 4111
25.301948 IP6 HA > S3-COA-GS1: IP6 HA > S3: ICMP6, echo request, seq 4113
25.301991 IP6 S3-COA-GS1 > HA: IP6 S3 > HA: ICMP6, echo reply, seq 4113

```

Figure 29. - Network Trace During a Hand-off with Loss
(Column 1 time offsets are seconds from Wed May 28 16:30:00 UTC 2014)

Conclusions

The second phase of the Unmanned Aircraft System in the National Airspace System (UAS in the NAS) communications flight test campaign successfully demonstrated in-flight transfer (hand-off) of an aircraft radio communication link between two independent ground stations. Using remotely-controlled ground stations located at Glenn Research Center in Cleveland, Ohio and Ohio University Airport in Albany, Ohio, the test flights demonstrated sustained air-to-ground L-Band radio connectivity throughout a 133 nautical mile (nmi) flight path. The flight tests also characterized the use of hand-offs of the Control and Non-payload Communications (CNPC) data between ground stations. Test data was collected on roundtrip times through the communications network and on data flow interruptions during hand-offs.

The Phase II flight tests utilized “Generation 2” software-defined radios developed under cooperative agreement NNC11AA01A between the NASA Glenn Research Center and Rockwell Collins, Inc., of Cedar Rapids, Iowa. Although the radios are only prototypes, they have demonstrated very repeatable and reliable operation in the field. Like the performance recorded during previous test series [1], the radios are providing connectivity beyond the 69 nmi range initially planned. Omnidirectional antennas were used for the L-band communications at both ground terminals and on the aircraft.

The testing demonstrated the proposed network architecture will allow for smooth transitions between ground stations. The proposed architecture was designed to be scalable and allows for industry standard security controls to be incorporated in the future. Operators do not need to be aware of which network path is currently being utilized as the underlying architecture is designed to seamlessly direct messages to the correct endpoint as the mobile node traverses the network of ground stations.

The networking tests showed that the median round trip propagation time between the airborne and ground radios was 225 ms, when the CNPC radio was operating in a 10 Hz configuration. The observed round trip times had a significant amount of variation, with the majority of samples between 171 and 267 ms. This spread of 96 ms was consistent with the 10 Hz (100 ms) waveform configuration used for the tests. Thus, the maximum targeted 20 Hz (50 ms) configuration should yield the least amount of delay variation. Likewise, reducing the waveform configuration to values below 10 Hz would further increase the amount of jitter. These values are a result of small data packets that can fit into a single transmission event. Larger data payloads would increase the amount of time needed to transmit the full payload and decrease the amount of observed jitter. Finally, traversal over the terrestrial Internet added an additional 55 ms round-trip for the location in Albany, Ohio. This brought the total round trip delay from the home network location to the aircraft via the ground station to over 0.25 seconds in the average case.

The use of periodic router advertisements to assign IPv6 addresses, while sufficient in the lab environment, proved to add long-lived spikes to the observed round trip times of normal data by introducing extraneous traffic into an already severely constrained link. Future work will investigate transitioning IPv6 address assignment to using DHCPv6 or some other method where the additional bandwidth of periodic router advertisements can be suppressed.

Router advertisements were not the only extraneous traffic encountered in these tests. The radio would internally generate and transmit a small amount of “health” information every second. These messages would also significantly impact observed round trip times in numerous, short-lived spikes. The impact was quite significant and the utility and frequency of these messages will be revisited in future updates of the radio.

Network hand-offs are shown to be seamless and transparent to the end user. When forcefully breaking a connection, the airborne system waits a set period of time then issues a request to use the standby connection. This request is completed in one round trip time with a lossless link. The hand-off from one ground station to the next

nominally required 1.268 s. The set period of time for the mobile IPv6 software implementation to detect link failure is currently one second, which comprises the majority of the 1.268 s recovery time.

Several methods are being explored to further reduce this time. Attempting to make the one second timer more aggressive can reduce hand-off times in the current break-before-make scenario. However, additional analysis would be needed to demonstrate that more aggressive timers would not negatively impact performance. Another means of reducing this time is by utilizing a make-before-break hand-off. Such a hand-off would eliminate the one second timer completely, leaving only the one round trip time of roughly 268 ms and thus incur less data loss. Both of these techniques will be explored in future testing.

NASA will continue to investigate the performance of the CNPC radio links in an ongoing flight test program. Upcoming test activities will include but not be limited to: further investigation of the 5030- to 5091-MHz frequency range (C-Band) for either alternate or simultaneous CNPC operations, measuring user data throughput rates, demonstrating data header compression and data security, and evaluating multi-aircraft operations. Results from these tests will be presented in future reports.

References

1. Shalkhauser, K. et al: Control and Non-Payload Communications Generation 1 Prototype Radio Flight Test Report. NASA/TM—2014-218099, 2014.
2. Matolak, D. and Sun, R.:et al.: AG Channel Measurement and Modeling Results for Over-Sea Conditions. NASA/CR—2014-216674, 2014.
3. Wilson, W., “Terrestrial L-Band and C-Band Architectures for UAS Control and Non-Payload Communications”, SC203-CC019, RTCA SC-203, December 2010
4. Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6275.txt>
5. Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3963.txt>
6. Kundra, V., "Transition to IPv6", Federal Mandate, September 2010. [Online]. Available: http://www.whitehouse.gov/sites/default/files/omb/assets/egov_docs/transition-to-ipv6.pdf
7. Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc4861.txt>
8. Schneider, A., “GPS Visualizer”, Computer Software, <http://www.gpsvisualizer.com/>
9. Derrickson, W., “Lockheed S-3B N601NA NASA Glenn”, 2008, Cleveland National Air Show, <http://tarmacphotos.com/>
10. National Atlas of the United States, March 5, 2003, <http://nationalatlas.gov>

Appendix A

	Band	LSID /CID	Mode	Channel center frequency	Transmit Power level	Frame Rate (Hz)	Aircraft	GS1 (GRC)	GS2 (OU)
Configuration 1	L	1	UL1	968 MHz	63	10	X	X	
		2	DLC2	968 MHz	63	10	X	X	
		3	UL20	968 MHz	63	10	X	X	
		4	DLVideo	968 MHz	63	10	X	X	
	C								
Configuration 2	L								
	C	1	UL1	5060 MHz	63	10	X	X	
		2	DLC2	5060 MHz	63	10	X	X	
		3	UL20	5060 MHz	63	10	X	X	
		4	DLVideo	5060 MHz	63	10	X	X	
Configuration 3	L	1	UL1	968 MHz	63	10	X	X	
		2	DLC2	968 MHz	63	10	X	X	
		3	UL20	968 MHz	63	10	X	X	
		4	DLVideo	968 MHz	63	10	X	X	
	C	1	UL1	5060 MHz	63	10	X	X	
		2	DLC2	5060 MHz	63	10	X	X	
		3	UL20	5060 MHz	63	10	X	X	
		4	DLVideo	5060 MHz	63	10	X	X	
Configuration 4	L	1	UL1	968 MHz	63	4	X	X	
		2	DLWx	968 MHz	63	10	X	X	
		3	UL4	968 MHz	63	4	X	X	
		4	DLVideo	968 MHz	63	10	X	X	
		5	UL8	968 MHz	63	4	X	X	
		7	UL12	968 MHz	63	4	X	X	
		9	UL16	968 MHz	63	4	X	X	
	C	1	UL1	5060 MHz	63	4	X	X	
		2	DLWx	5060 MHz	63	10	X	X	
		3	UL4	5060 MHz	63	4	X	X	
		4	DLVideo	5060 MHz	63	10	X	X	
		5	UL8	5060 MHz	63	4	X	X	
		7	UL12	5060 MHz	63	4	X	X	
		9	UL16	5060 MHz	63	4	X	X	
Configuration 5	L	1	UL1	968 MHz	63	10	X		X
		2	DLC2	968 MHz	63	10	X		X
		3	UL20	968 MHz	63	10	X		X
		4	DLVideo	968 MHz	63	10	X		X
	C	1	UL1	5060 MHz	63	10	X		X
		2	DLC2	5060 MHz	63	10	X		X
		3	UL20	5060 MHz	63	10	X		X
		4	DLVideo	5060 MHz	63	10	X		X

Configuration 6	L	1	UL1	968 MHz	63	5	X	X	
		2	DLC2	968 MHz	63	5	X	X	
		3	UL20	968 MHz	63	5	X	X	
		4	DLVideo	968 MHz	63	5	X	X	
		5	UL1	968 MHz	63	5	X		X
		6	DLC2	968 MHz	63	5	X		X
		7	UL20	968 MHz	63	5	X		X
		8	DLVideo	968 MHz	63	5	X		X
Configuration 7	C	1	UL1	5060 MHz	63	5	X	X	
		2	DLC2	5060 MHz	63	5	X	X	
		3	UL20	5060 MHz	63	5	X	X	
		4	DLVideo	5060 MHz	63	5	X	X	
		5	UL1	5060 MHz	63	5	X		X
		6	DLC2	5060 MHz	63	5	X		X
		7	UL20	5060 MHz	63	5	X		X
		8	DLVideo	5060 MHz	63	5	X		X
Configuration 8	L	1	UL1	968 MHz	63	20	X	X	
		2	DLC2	968 MHz	63	20	X	X	
Configuration 9	C	1	UL1	5060 MHz	63	10	X	X	
		2	DLC2	5060 MHz	63	10	X	X	
		3	UL1	5060 MHz	63	10	X		X
		4	DLC2	5060 MHz	63	10	X		X
Configuration 10	C	1	UL20	5060 MHz	63	10	X	X	
		2	DLVideo	5060 MHz	63	10	X	X	
		3	UL20	5060 MHz	63	10	X		X
		4	DLVideo	5060 MHz	63	10	X		X
Configuration 11	L	1	UL1	968 MHz	63	10	X	X	
		2	DLC2	968 MHz	63	10	X	X	
		3	UL1	968 MHz	63	10	X		X
		4	DLC2	968 MHz	63	10	X		X
Configuration 12	L	1	UL20	968 MHz	63	10	X	X	
		2	DLVideo	968 MHz	63	10	X	X	
		3	UL20	968 MHz	63	10	X		X
		4	DLVideo	968 MHz	63	10	X		X
Configuration 13	L	1	UL20	968 MHz	63	10	X	X	
		2	DLC2	968 MHz	63	10	X	X	
		3	UL20	968 MHz	63	10	X		X
		4	DLC2	968 MHz	63	10	X		X
	C								

Hardware/Software Inventory

Home Agent / Network

- Intel Xeon E3-1230 CPU @ 3.30 GHz (Quad Core)
- 16 GB RAM
- 240 GB SSD (Seagate ST240HM000)
- 2 TB HDD (Seagate ST2000NM0033)
- 4x Intel 82574L Gigabit Ethernet Interfaces
- Ubuntu Linux 12.04.4 LTS Server
- Linux 3.11.10.4-mobility00 kernel (custom compiled to support Mobile IPv6)
- UMIP v1.0 (www.umip.org) (mip6d daemon)
- Radvd 1.8.3 (Router Advertisement Daemon)
- Cisco ASA 5510 Firewall

Ground Stations

- Intel Xeon E3-1230 CPU @ 3.30 GHz (Quad Core)
- 16 GB RAM
- 240 GB SSD (Seagate ST240HM000)
- 2 TB HDD (Seagate ST2000NM0033)
- 4x Intel 82574L Gigabit Ethernet Interfaces
- Ubuntu Linux 12.04.4 LTS Server
- Linux 3.8.0-38-generic kernel (Ubuntu distributed kernel)
- Radvd 1.8.3 (Router Advertisement Daemon)
- Iperf 3.0.3
- Custom Python and BASH shell scripts to communicate with radio
- Cisco ASA 5505 Firewall

Mobile Router (Onboard Aircraft)

- Intel Xeon E5-1660 @ 3.30 GHz (Hex Core)
- 16 GB RAM
- RAID-1 2x 256 GB SSD (Crucial M4-CT256M4SSD2)
- RAID-1 2x 2 TB HDD (Toshiba MK2002TSKB)
- 2x Intel I350 Gigabit Ethernet Interfaces
- Ubuntu Linux 12.04.4 LTS Desktop
- Linux 3.11.10.4-mobility00 kernel (custom compiled to support Mobile IPv6)
- UMIP v1.0 (www.umip.org) (mip6d daemon)
- Radvd 1.8.3 (Router Advertisement Daemon)
- Iperf 3.0.3
- Custom Python and BASH shell scripts to communicate with radio

