Communications Challenges in Field Research Activities

Ray Gilstrap
NASA Advanced Supercomputing Division
NASA Ames Research Center
Moffett Field, CA, USA
ray.gilstrap@nasa.gov

Abstract - *NASA conducts a number of science and exploration activities in remote field locations. Many of these activities require data, voice, or video connectivity from the field to external sites. These requirements have been met using hybrid communications networks that incorporate wired and wireless Ethernet networks at the field site, satellite links for connectivity outside of the site, and existing terrestrial networks for access to remote institutions and the Internet. However, the variety of activities and typical characteristics of field environments pose challenges for these networks. Meeting these challenges has yielded a number of useful lessons in the areas of system and network optimization, radio frequency engineering, and mission operations and management.*

Keywords: Networks, communications, field research, analog missions, surface exploration.

1 Introduction

NASA conducts a number of science and exploration activities in remote environments on Earth that approximate aspects of the lunar or Martian surface. For many of these expeditions, it is essential that researchers and equipment in the field be able to communicate with colleagues and resources at their home institutions. This communication can range from simple telephone conversations to relaying data gathered in the field to a remote supercomputer for analysis. Since these experiments generally take place in locations with no existing communications infrastructure, a communications network must be established as part of the initial setup of the science activities.

A typical field communications network is a hybrid of three communications technologies: on-site wired and wireless local area networks (LANs), satellite links to provide connectivity from the field site to a network gateway facility, and standard terrestrial networks based on the Internet Protocol (IP). Each of these technologies has a unique set of characteristics that affect the performance of the overall end-to-end communications network.

This paper explores the communications challenges encountered and the subsequent lessons learned over the course of a number of these field activities. Section 2 introduces the concept of field analog experiments and summarizes the types of science activities that our group, the NASA Research and Engineering Network (NREN), has supported. Section 3 explains the types of data that are normally generated during these field activities, and what requirements these different data types levy on communications networks. Section 4 illustrates a typical network architecture deployed in a field setting. Section 5 summarizes the lessons learned from experience operating these communications networks. Section 6 identifies areas for further research that could benefit field communications. Section 7 concludes the paper.

2 Field Analog Experiments

NASA's current space exploration plans call for establishing a manned lunar base, where astronauts can gain experience with long-term surface exploration activities in preparation for an eventual manned mission to Mars. Both lunar and Martian surface exploration, in turn, can benefit from experience gained in simulated missions that take place in terrestrial environments that mimic relevant features of the surface environment. NREN has supported communications at a number of sites in collaboration with other groups at NASA Ames, other NASA centers, and outside organizations. These include:

Ground Truthing trial, Vernal, Utah (2003)—Using soil-composition data gathered with on-site spectrometers to calibrate similar data gathered by remote sensing satellites [1].

Mobile Agents experiments, Hanksville, Utah (2003-2005)—Field trials of mobile agent software to facilitate interactions between human astronauts and robotic assistants at the Mars Desert Research Station (MDRS) [2].

Desert Research and Technology Studies (RATS), Meteor Crater, Arizona (2003-2004)—Testing of various spacesuit, robot, vehicle, and communications technologies [3].

Mars Astrobiology Research and Technology Experiment (MARTE), Rio Tinto, Spain and Santa Cruz, California (2004-2005)—Studying bacteria in the waters of Rio Tinto, which is highly acidic and has high mineral content. Water with similar characteristics is believed to exist elsewhere in the solar system, where similar bacteria could have lived [4].

Hurricane Katrina emergency response, NASA Michoud Assembly Facility, New Orleans, Louisiana (2005)—Providing emergency Internet and telephone connectivity to the facility after the hurricane disabled its existing communications infrastructure.

Ad Hoc networking trials, Moss Point, California and Meteor Crater, Arizona (2005 and 2008)—Field testing of protocols for mobile ad hoc networks (MANETs) [5].

K-9 Robot tests, Santa Cruz, California (2006)—Field testing of a prototype robotic planetary surface explorer.

Spaceward Bound, Atacama Desert, Chile (2006)—Teleoperation of robots in a Mars analog environment by students visiting NASA Ames [6].

Spaceward Bound Arctic field studies, Axel Heiberg Island and Devon Island, Nunavut, Canada (2006-2009)—Studies of environment, glacier movements, and native flora and fauna at McGill University's Arctic Research Station on Axel Heiberg Island and the Mars Arctic Research Station on Devon Island [7, 8].

The communications support provided for these activities ranges from providing simple telephony services over existing network infrastructure to full satellite-based connectivity from the field location to NASA networks and the general Internet. This wide range of activities illustrates the flexibility that a robust communications network can offer to on-site researchers.

3 Communications Requirements

Several different types of data that commonly generated in field exploration. To maximize interoperability with external systems and applications, all field traffic is based on the standard Internet Protocol (IP). Each traffic type is sensitive to particular network characteristics. The most critical of these characteristics are *throughput*, *loss rate*, *latency*, and *jitter*. Throughput is the rate at which data passes through the network, measured in bits (or bytes) per second. Loss rate is the fraction of data that is lost in transit, due to factors such as congestion in the network or radio frequency (RF) interference. Latency is the amount of time it takes for a given unit of data to travel through the network from sender to receiver. Jitter is a measure of the fluctuations in latency that happen over time, frequently calculated as the standard deviation of a set of latency measurements.

With these characteristics in mind, common field data types include:

Science data—A typical science instrument, such as the imaging spectrometer used in the ground truthing activity, generates large data files consisting either of binary data or of readable text. These data files can then be transferred to external computing resources for analysis. In this case, the reliability of the transfer is paramount; the loss rate must be zero to ensure the integrity of the received data. Throughput is a secondary concern, but is still important when dealing with a large number of large data files. Latency and jitter are generally less important; the data files are not useful until the last byte is received, so the arrival times of individual parts of the files are not critical.

Real-time telemetry—This includes continuously streaming data from sensors monitoring environmental conditions or the health of equipment or researchers. This data is not typically high-bandwidth, but it is sensitive to latency. Minimizing latency is especially important in control applications, where some action is taken in response to incoming sensor data. Acting on stale data may produce undesirable results. In some applications, it is better to simply discard data that is late, rather than try to use it.

Real-time video and voice—This includes webcams, telephony, and videoconferencing tools such as Skype. This type of traffic can tolerate a small amount of loss, since new data is being continuously generated. However, it is sensitive to jitter—network voice and video coding algorithms depend on having data arrive at regular intervals, and fluctuations in those intervals can lead to distortions in the rendered video or audio. Again, discarding late data may produce results that are less perceptible to a viewer or listener. Depending on frame rate and quality, video can sustain relatively high throughput, up to several megabits per second. Voice typically has a throughput of a few tens of kilobits per second.

General Internet usage—This includes routine Web browsing and e-mail activities. A user's web browsing experience is sensitive to latency and throughput. Remote logins (for example, via Secure Shell) are less sensitive to throughput because of their low data rates, but are highly sensitive to latency. E-mail does not place such specific requirements on the network because of its non-interactive nature—within reasonable limits, it is unimportant how long messages take to be sent or received.

These requirements must be supported by the network design and implementation for the applications to function well.

4 Field Network Architecture

A general network architecture is shown in Figure 1. A transportable ground station provides the satellite connectivity to external networks. Locally, the ground station is connected to wired and wireless LANs, which are designed and deployed according to the number and locations of systems that will be served. Often, activity focuses on a central location (such as the MDRS facility, shown in Figure 2), but can extend to sites that are several kilometers away. A system of long-distance wireless network links can provide coverage at these locations.

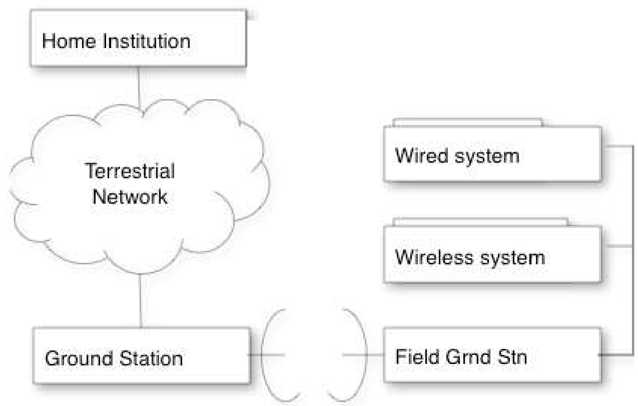

Figure 1. General field network architecture.

Figure 2. Mars Desert Research Station. Note the transportable ground station (yellow/white) to the right.

At the field site, Ethernet-based wired LANs are used to connect systems that are stationary, such as servers. Wireless LANs provide connectivity for mobile systems such as laptops, PDAs, and some science instruments. Originally, the wireless LANs employed the ubiquitous IEEE 802.11b (Wi-Fi) standard. More recent deployments have incorporated the now-common IEEE 802.11g standard and the more recent IEEE 802.11n standard. The newer standards provide improved throughput and better scalability to larger quantities of nodes. In addition to providing routine services, wireless technologies have also been a part of the research objective—testing has explored the effectiveness of new wireless equipment and protocols, and enabling access points to follow mobile systems around as they move in order to maintain optimal connectivity. From a network performance standpoint, wireless LANs generally have good throughput and loss rates over short distances, but throughput decreases and loss rate increases with distance.

The satellite link is the connection to the outside world. The transportable ground station at the field site consists of a dish, transceiver, satellite modem, and an IP router. Network traffic from the systems in the field is passed by the router to the modem, where it is converted to electrical signals that the transceiver then converts into RF and transmits through the dish. Incoming traffic follows the reverse path. Most activities that we have supported use satellite transponders with frequencies in the K_u band (12-18 GHz), but C band (4-8 GHz) and K_a-band (26-40 GHz) are also occasionally used. Performance-wise, satellite links have low loss rates and jitter, but latencies are high due to the amount of time it takes to transmit a signal from the field site to the satellite (virtually always in a geostationary orbit 22,000 miles above the Earth's equator) and back to the remote satellite dish. It commonly takes roughly 300 milliseconds for data to travel in one direction over the satellite link. (By comparison, a one-way latency across a terrestrial network from California to Maryland is nearly an order of magnitude smaller.) The amount of throughput possible on commercial satellites is scaled according to the price paid for the link; normally, field projects have the budget to support rates of one to five megabits per second. The satellite time is negotiated and purchased from the provider in advance of the expedition. Most contracts are for coverage 24 hours per day for the entire duration of the expedition, although shorter amounts of time may be purchased to lower costs.

The other end of the satellite link is a permanent ground station that has connectivity to one or more terrestrial networks. Some field activities have used NASA-operated ground stations with direct connections to institutional networks like NREN's nationwide research network, while others use commercial ground stations that route traffic over the commodity Internet. The advantages of the former are that the institutional network is relatively isolated from the unrelated traffic that is present on commercial networks, and the field site can be placed logically inside the institutional network, which simplifies configuration and security concerns. When commercial networks are used, competing traffic is a potential issue, and appropriate encryption standards must be used to ensure that confidential data is not exposed to unauthorized parties. Terrestrial networks offer high throughput, as well as low loss in the absence of congestion. Latency can vary depending on the exact path taken by the data, but for a given path jitter is normally low.

5 Lessons Learned

We have gained much experience in the years that we have worked in this area. These can be roughly categorized into the areas of system optimization, network optimization, RF engineering, and mission operations and management.

5.1 System Optimization

Tuning systems to accommodate the long latencies of the satellite link is essential to maximizing throughput for large file transfers. At issue is the Transport Control Protocol (TCP), which augments the base IP to provide reliable delivery of data. In IP networks, data to be sent is fragmented into chunks, or *packets,* that have a maximum size specific to the underlying network being used. (For example, 1500 bytes is a common maximum in Ethernet networks.) Roughly speaking, a sender using TCP does not transmit new packets of data until it receives an acknowledgement from the receiver that previously transmitted packets were successfully received. If the unacknowledged packets represent only a small amount of data, then on a high-latency link the sender spends most of its time awaiting acknowledgements and not transmitting anything. However, most computer operating systems (including as the Linux, Macintosh, and Windows systems generally used in these field experiments) will allow adjustments to the *window size,* or maximum allowed amount of unacknowledged data, so that the sender can transmit a larger amount of data before receiving an acknowledgement. This allows for better utilization of the satellite link and dramatically improves the throughput of applications that use TCP, including file transfers and Web browsing.

Another approach to mitigating the inefficient utilization of high-latency links is the use of commercially available, hardware-based *performance-enhancing proxies (PEPs).* In typical use, one PEP is placed on each end of the satellite link. Each PEP intercepts TCP traffic destined for the other side of the link and relays that traffic to its peer using a more efficient, and often proprietary, protocol. The advantage of using PEPs is that end systems do not need to be tuned to use them, which particularly saves labor when multiple hosts on each side of the link need to have optimized TCP connections. However, there is additional cost to acquire the PEP hardware, and field projects may not have the authority to insert a PEP into the permanent ground station.

We evaluated the SkyX PEP during one of the Desert RATS exercises in Arizona. Generally speaking, we did observe a marked increase, typically an order of magnitude, in throughput over the satellite connection. However, due to the constrained bandwidth of the satellite link, this performance turned out to be comparable to what we achieved with host tuning. It is anticipated that a PEP solution will scale better as link bandwidth is increased. While the SkyX product has been discontinued, there are other commercially available alternatives that we may evaluate in future deployments.

5.2 Network Optimization

If there are many types of traffic being sent over the network, there may be contention for the limited bandwidth of the satellite link. In this case, it may be necessary to prioritize mission-critical traffic, such as science data, so that it does not get crowded out by other types of traffic. This prioritization can be accomplished using *quality of service (QoS)* features common in routers, which enable the routers to treat particular types of traffic differently than others. For example, during periods of contention the router may be configured to delay or discard voice or video packets if packets of science data are waiting to be transmitted over the link. This particular QoS configuration was common in experiments that relied on exchange of science data with a remote facility, such as ground truthing.

Additionally, we have found it useful to monitor the network to ensure that there are no anomalies in traffic patterns. For example, during the ground truthing experiment, a network monitoring tool that we developed detected unusual network traffic coming from a particular user laptop. Further investigation revealed that the laptop was actually infected with a virus, which was attempting to contact the outside world. Those attempts were blocked by firewall rules put in place to protect the field network, a further reminder that security is an important consideration.

5.3 RF Engineering

Many of the most interesting findings have come at the RF layer. One important issue in deployments incorporating multiple wireless LANs is that there is some overlap between the eleven channels (frequency ranges) defined for use in the United States by the FCC. As a result, only three of these channels (1, 6, and 11) can be used simultaneously without interfering with each other.

In terrain that has many obstructions, such as the rocky Lith Canyon near MDRS, multipath is a problem. Multipath is a phenomenon where an RF signal reflecting off an obstruction interferes with itself. The effects of multipath can be partially mitigated through careful placement of the wireless LAN transmitters. Additionally, the recent IEEE 802.11n wireless networking standard utilizes multiple transmitting and receiving antennas to reduce the detrimental effects of multipath. Further testing of 802.11n in areas where multipath has previously degraded our observed 802.11b and 802.11g performance is needed.

Antennas are also an important issue on long-haul links. During the MDRS outings, a solar-powered 802.11b repeater, consisting of two back-to-back 802.11b access points with grid antennas attached to each, was deployed to connect the central facility to the remote canyon. These antennas produce a highly directional radiation pattern, which makes it essential that they are precisely aligned with the receiving antennas. This required careful surveying of the area using GPS and topological maps, as well as secure fastening of the antennas once they were correctly pointed to ensure that alignment would be maintained even in strong wind gusts. Newer, more expensive models have an audible tone to assist with alignment.

MANET protocols can also help in cases where RF connectivity to the central facility is marginal. These allow any system on the network to relay data on behalf of any other system. For example, if system A wishes to communicate with system C but is out of RF range, it can relay the data through a system B that is within range of both A and C. We have tested implementations of two different MANET protocols, Dynamic Source Routing (DSR) and Optimized Link State Routing (OLSR). Our experience has been that they do allow data transfers to proceed in areas where they otherwise couldn't, but the throughput, latency, and loss rate are all degraded. Nevertheless, the reduction in communications outages that these protocols offer will make them useful in terrain such as the lunar surface [9].

One unexpected but interesting finding occurred during the Axel Heiberg deployment. Because there were no available K_u band satellites capable of communicating with ground stations at such a high latitude (nearly 80 degrees north), K_a band was used instead. However, K_a band is very susceptible to attenuation from water vapor. As a result, when the afternoon sun melted and evaporated the snow in the vicinity of the dish, the resulting increase in water vapor degraded the satellite signal to the point where connectivity was lost. Unfortunately, a workaround for this could not be determined prior to the end of the deployment.

5.4 Mission Operations and Management

Although these areas may not have a direct technological component, the presence of communications networks does turn out to impact them. The ability of a researcher in the field to have voice, video, and data communications with a remote colleague is invaluable. In cases where the primary field activity is simulation of a mission, this communications capability allows Earth-based mission controllers to participate in the simulation. As a result, new mission procedures were specifically developed to take advantage of this capability, particularly during the Mobile Agents experiments.

6 Further Research

A number of promising areas of investigation could yield enhancements in the capabilities of field communications networks. The incorporation of disruption-tolerant networking capabilities would allow mobile systems to temporarily queue data destined for a remote system that is currently unreachable, then transmit it automatically when connectivity to that node is restored. Emerging mesh networking standards, such as IEEE 802.11s, allow wireless access points to automatically track connectivity among themselves and select the best path for relaying traffic to a given destination. The integration of sensor network technologies, such as new protocols for low-powered wireless sensors, into the field LAN will provide a good opportunity to streamline communication between the sensors and the applications that use their data. Determining the optimal coding algorithms for voice and video over the set of available wireless LAN standards would improve the perceived quality of these data streams. Finally, comprehensive testing of the IEEE 802.11n and IEEE 802.16 (WiMax) standards in field conditions will yield valuable information on their performance characteristics.

7 Conclusions

Field exploration activities are useful for a variety of purposes including mission simulation, scientific investigation, and technology evaluation. Communications networks are essential to those activities in that they provide voice, video, and data service to researchers in the field. Each activity and location poses its own unique set of challenges, but a properly designed and implemented communications network can account for these challenges and provide a critical link to the outside world.

Acknowledgment

We wish to acknowledge the invaluable assistance of the following collaborators in accomplishing the work described here: Rick Alena, John Ossenfort, Brian Glass, Bill Notley, Bill Hunt, Thom Stone and Mark Foster of NASA Ames Research Center; Mike Cauley and Dave Pleva of NASA Glenn Research Center; Marc Seibert of NASA Kennedy Space Center; and Kevin Bass of NASA Johnson Space Center.

References

[1] M. Johnson, K. Freeman, R. Gilstrap and R. Beck, "Networking technologies enable advances in Earth science," *Computer Networks,* Vol. 46, No. 3, Oct. 2004.

[2] W. Clancey, et al, "Mobile Agents: A distributed voice-commanded sensory and robotic system for surface EVA assistance," Proc. 9th Biennial ASCE International Conference on Engineering, Construction, and Operations in Challenging Environments, League City/Houston, TX, pp. 85-92, Mar. 2004, http://ti.arc.nasa.gov/m/pub/658h/0658%20(Clancey).pdf.

[3] National Aeronautics and Space Administration, "Desert RATS overview," http://www.nasa.gov/exploration/analogs/desert_rats.html, Last Updated: Oct. 1, 2009.

[4] National Aeronautics and Space Administration, "Mars astrobiology research and technology experiment," 2005, http://marte.arc.nasa.gov/.

[5] J. Toung, R. Gilstrap and K. Freeman, "A split implementation of the dynamic source routing protocol for lunar/planetary surface communication," Proc. IEEE Aerospace Conference, Big Sky, MT, 2006.

[6] National Aeronautics and Space Administration, "Spaceward Bound – Atacama 2006," http://quest.nasa .gov/projects/spacewardbound/atacama2006/index.html, Last Updated: Apr. 2007.

[7] National Aeronautics and Space Administration, "Spaceward Bound – Arctic 2008," http://quest.nasa.gov /projects/spacewardbound/arctic2008/, Last Updated: July 2008.

[8] Mars Institute et al., "Haughton-Mars Project," 2008, http://www.marsonearth.org/.

[9] T. Stone, et al, "A lunar communication satellite network architecture employing Internet protocol, laser communication technologies and small satellites," Proc. AIAA Conference on Satellite Communications, San Diego, CA, 2008.