

Real Time GPS Data Transmission Using VSAT Technology

Michael E. Jackson^{1*}, Chuck Meertens¹, Oivind Ruud¹,
Spencer Reeder¹, Warren Gallaher¹, and Chris Rocken²

¹University NAVSTAR Consortium (UNAVCO),
University Corporation for Atmospheric Research (UCAR) Office of Programs

²GPS Science & Technology (GST) Program,
University Corporation for Atmospheric Research (UCAR) Office of Programs

Abstract

The University NAVSTAR Consortium (UNAVCO) Boulder Facility is assessing Very Small Aperture Terminal (VSAT) technology for near real-time transmission of GPS data from a remote receiver to a central processing facility. The study is motivated by the need for a robust, cost effective data communications solution to transfer GPS data from remote sites where no other communication alternatives exist. Future large-scale plate boundary deformation initiatives using spatially dense networks of GPS will require receivers to be located where the science dictates and not the power or communications infrastructure. For other applications, such as determining rapid GPS orbits and time transfer, there is a push towards reducing the latency in GPS data used to produce GPS data products and differential corrections (Talaya & Bosch, 1999; Jackson, Meertens & Rocken, 2000, Muellerschoen, Bar-Sever, Bertiger & Stowers, 2001), and to support upcoming Low Earth Orbiting (LEO) missions requiring low latency, 1 s GPS data.

In this paper we evaluate two Ku-band systems, the Nanometrics Libra VSAT and the StarBand 2-way satellite Internet VSAT. The Nanometrics system test results show that continuous, 1 s GPS data can be streamed from multiple remote stations within the VSAT footprint, quality checked, and delivered for processing with a <2.5 s latency (mean 1.2 s) and a 99.8% reliability. Benefits of the Nanometrics system include global coverage, control of bandwidth allocation and data hub and the low power draw of the system. Negatives include the cost of hub and remote infrastructure and the need to negotiate landing rights issues on a country-by-country basis. The facility currently operates a Nanometrics hub and 3 remote VSAT systems.

The StarBand system showed 98.9% reliability with a maximum latency of 10.2 s (mean latency 1.7 s) for 1 Hz GPS data and an average uplink speed of 31.7 kbps. Benefits of the StarBand system include the cost and small profile of the remote antenna. Negatives include coverage limited to

coterminous United States and high power draw of remote systems.

VSAT Introduction

Very Small Aperture Terminal (VSAT) networks use geostationary satellites orbiting in the equatorial plane of the earth at an altitude of 35786 km to transmit data from a network of remote sites to a data hub (Figure 1) (Maral, 1995).




Figure 1. A typical VSAT network depicting two-way communications from remote terminals through a VSAT satellite to a central hub.

VSAT satellite space segment providers offer three types of satellite beam: spot, hemisphere, and global. Spot beams are available in both Ku-band (12-16 GHz) and C-band (4-6 GHz) and are focused on population centers. Spot beams are generally high power, thus allowing smaller antenna dishes to be used at remote sites (for example 1.8 m for Ku and 3 m for C band). Hemisphere and Global beams have a much larger footprint and weaker signal strength. The majority of hemisphere and all global beams are C-band. C-band's lower frequencies are less sensitive to rain-induced signal degradation (attenuation of 1 to 2 dB) compared to the Ku-band signal (attenuation of 5 to 10 dB) but suffers from greater interference relative to Ku-band (Maral, 1995). Because the antenna gain depends on the frequency, the C-band signal has a smaller gain per unit area and thus requires a larger antenna diameter relative to Ku-band systems. For Ku-band transmissions, rain-induced signal degradation or rain fade, can cause significant loss

* Corresponding author address:
Michael E. Jackson, UNAVCO Boulder CO
e-mail: mikej@unavco.ucar.edu

of data. A number of options are available for mitigating this effect such as buffering data at each remote during outages, increasing the network transmit power on the satellite, or increasing the transmission power at the hub and remotes. Generally combinations of the above are used to reduce the impact of rain fade. A more predictable data outage can also occur when the geostationary satellite eclipses the sun. This is referred to as sun outage and generally lasts no more than 10 minutes once per year.

StarBand Two-way Satellite Internet

Introduction

The StarBand system is a two-way satellite Internet system consisting of a 0.5 m dish and a satellite modem that routes data from remote users through the GE-4 and Telstar 7 satellites to a StarBand hub and then to the Internet. The system is targeted for home users with the expectation of low uplink and high downlink bandwidth requirements and requires an onsite computer and AC power (modem draws 30 W power). StarBand offers a 2-way satellite based Internet service with download speeds close to those available via cable modem or DSL without the need for phone lines or an Internet Service Provider (ISP) (Figure 2).




Figure 2. The StarBand VSAT network currently provides two-way Internet service within the coterminous US with planned coverage for Alaska, Hawaii, and Mexico. Data are routed from a remote user to either GE-4 or the Telstar 7 satellite to the StarBand hub, which has a direct link to the Internet.

StarBand states consumers can expect download speeds up to 500 kbps (150 kbps minimum) and upload speed bursting up to 150 kbps (50 kbps average). The system consists of a 60 by 90 cm satellite dish mounted with a clear view of the southern sky. The dish is connected to a satellite modem with a dedicated IP address using standard

coaxial cables. For our tests we modified the satellite modem to accept input from an RJ45 connector attached to a PC running Linux. Internet commands are routed from the PC, through the satellite modem and dish to either the GE-4 or the Telstar 7 satellite. The satellite, in turn, communicates with StarBand's hub facility which has a direct connection to the Internet (Figure 2). Equipment costs for the StarBand system are approximately \$400 for a dish and satellite modem, installation is \$200 and monthly service is approximately \$70.

StarBand GPS data transmission tests

The StarBand system was tested to determine if it would be suitable for installation at remote GPS sites with no other communications infrastructure. To accomplish this the following tests were conducted:

1. Determine mean uplink speeds by repeatedly downloading a file from a remote computer using FTP.
2. Determine the magnitude of network outages by repeatedly sending a 'ping' command to the remote system and analyzing the response.
3. Determine real time data latency and link status using the JPL (Jet Propulsion Laboratory) developed RTNT software.
4. Analyze the average DC power draw of the system.

Mean uplink speed was tested by monitoring the transfer time (using FTP) to download a 1.5 Mb file, representing a ~24-hour GPS data set, from a remote Linux computer. Downlink speeds to the remote computer were not tested, as bandwidth requirements to a remote site are negligible for GPS permanent station installations. Figure 3 shows the transfer time for 102 transfers indicating a majority of the transfers occurred in ~500 s and a minority occurred in ~100 s.




Figure 3. FTP transfer time in seconds for a 1.5 Mb file from a remote station. Bimodal distributions of transfer times indicate burst transfer speeds of 159 kbps and sustained uplink speeds of 32 kbps.

The mean speed for the two distributions were divided into the number of bits transferred resulting in uplink speeds of 159 kbps with most transfers occurring near 32 kbps (minimum = 24.7 and maximum = 653.2 kbps). The fastest transfers were close to the uplink burst values advertised by StarBand while 32 kbps are lower than the advertised sustained uplink values.

The amount of time the StarBand system was operational was determined by sending a 'ping' command to the remote system every 60 s. If the remote computer responded successfully then the network is considered operational. Between February 21 and March 2, 2001 a total of 12306 ping commands were sent at a rate of once every 60 seconds to the remote (Figure 4).




Figure 4. Ping statistics indicate a 1.1% downtime for the StarBand network connection over a 10-day period. A value of 1 indicates the network was operational. A value of zero means the remote was out of contact.

The ping statistics indicate that out of 12306 attempts the StarBand network connection was up 12177 times or a 1.1% downtime for the StarBand network link. Note that downtime could mean the link was down anywhere between the host computer and the remote computer.

To test the reliability of streaming data through the StarBand system we configured a Linux computer running JPL's Real-Time Network Transfer (RTNT) software (Muellerschoen, Bar-Sever, Bertiger & Stowers, 2001) with an Ashtech Z12 GPS receiver streaming 1 s GPS observables. In this study, the GPS data is captured and compressed by the RTNT software, packaged into a UDP/IP protocol with a packet header and sent to a destination computer over the Internet via StarBand. On the

receiving end, the packets are stripped of UDP/IP protocol and placed in time sequence. Any missing packets are re-requested. The raw data are then sent to a real-time processing engine or combined into data files for further distribution. There were a maximum of 12 and a minimum of 6 satellites visible during the approximately 48-hour RTNT test period.

Data latency is calculated by comparing the time tag of each transmitted GPS observable against the time received at the local data hub. The data latency measured at the receiving computer indicates that most of data arrive in the 1-4 s range (Figure 5). Note that at least 0.5 s of latency are due to uplink and downlink satellite transmission (Maral, 1995). The horizontal banding at the approximately 2, 2.75, and 3.5 second latency are likely data re-requests. Periods of high latency values (for example between 4 and 10e⁴ seconds) likely indicate periods of high user activity on the StarBand network.




Figure 5. Plot of data latency in seconds versus time for a 48-hour period. Data latency vary over a range of 1-10 s and with periods of high data latency corresponding to periods of high user activity on the StarBand network.

A latency histogram (Figure 6) indicates a mean latency of 1.7 s (minimum 0.9 s, maximum 10.2 s) with a majority of the data arriving within 4 s of transmission.




Figure 6. Latency data from Figure 5 plotted as a latency histogram. Note that a majority of the data packets arrive within 4 seconds.

Nanometrics Libra System

The UNAVCO facility tested a low power, stand alone Libra VSAT from Nanometrics. The Libra system consists of a 3.8 m Ku band hub broadcasting to the GE-1 satellite and a low power (< 20 W), 1.8 m remote Cygnus satellite transceiver antenna connected to an Ashtech Z12 GPS receiver streaming GPS observables at 1 Hz. The remote stations are configured in a TDMA (Time Domain Multiple Access) network with each remote occupying ~100 KHz of satellite bandwidth at data rates between 64 and 112 kbps (kilobits per second). Data are transmitted to and from the hub in short bursts on satellite channels that are shared with up to 30 - 40 other remotes (Figure 7).




Figure 7. A simulation of the number of remote stations possible given either 64 or 112 kbps throughput for a range of streamed data types. For this calculation a hub overhead factor of 8000 bps, a frame length of 5 s, a transmission overhead gap of .072 s and a packet size of 2584 bits were assumed.

The TDMA configuration allows a discrete time in which each remote terminal is authorized to transmit data, and allows the dynamic assignment of more or less time to remotes based upon their individual data requirements ensuring maximum network efficiency and throughput is maintained.

Nanometrics Libra GPS data transmission tests

The Libra system was tested for suitability at GPS remote installations where no other communications infrastructure exists. To accomplish this the following tests were conducted:

1. To determine data latency we compare the time tag of each transmitted GPS observable against the time received at the local data hub. For this test we used a TDMA frame time of 1 second and evaluated data from a single remote.
2. To determine data reliability we look at the amount of data received compared to the amount expected over a ~74 day test period.
3. We verified the manufacturers reported values for power draw.

The test network consisted of a central data hub located in Boulder and 2 remote terminals, one ~20 km away at Marshall Field (MARS) and the other at the Nanometrics facility near Ottawa, Canada (KANA) (Figure 8). Figure 8 shows contours of Equivalent Isotropically Radiated Power (EIRP) for the GE1 satellite. EIRP is defined as the sum of output power from the satellite's amplifier, satellite antenna gain and losses.

Each station in the network (hub and remotes) was assigned a TDMA slot with a frame length of 3 seconds.



Figure 8. Coverage area for the GE-1 satellite showing contours of equal EIRP. Station MARS and KANA were used in this evaluation. Station GUAX is an operational station deployed since the evaluation.

To test the system we placed an Ashtech Z12 receiver streaming 1 Hz MBEN, PBEN, and SNAV raw GPS observables to the remote VSAT terminal.

The data were packed in UDP/IP protocol with a header (NMXP) which contains a unique sequence number. The data were transmitted to the central hub at an uplink frequency of 14.0-14.5 GHz. The satellite rebroadcast the data to the central hub antenna at a downlink frequency of 11.7-12.2 GHz. At the hub the sequence numbers are checked for continuity, and if data are missing, a retransmission request is sent from the hub to the remote. Data are stored in ring buffers at the remote sites and are retransmitted if requested. The remote systems used for this test had approximately 30 minutes of buffer storage at the remote site for 1 Hz GPS data. Transmission statistics are stored at the hub and summarized for analysis. Every 24-hours data were pulled from the hub ring buffers, translated into RINEX using teqc (Esty & Meertens, 1999), and reduced to baseline differences using the Bernese GPS processing software.

Figure 9 shows the amount of data transferred from the MARS and KANA sites during the test period. Note that on days 260-278 the system was shut down for maintenance leading to reduced data volumes for these days.




Figure 9. Data (in Mb) transmitted from the MARS and KANA remote sites. Significant periods of data loss are labeled and explained in the text.

Figure 9 also indicates we experienced loss of data due to power failures at MARS. By instrumenting the remote with a digital amp-meter, we were able to confirm that the system was drawing the manufacturers specified 18 ± 2 W of power. Further investigation revealed the failures were due to faulty solar panels that were removed and replaced. On day 271 we experienced data loss due to rain fade resulting from a very heavy wet snowfall. Attenuation of the radio signals causes an interruption in transmission from the remote to the hub. If the period of interruption is greater than the storage capacity of the ring-buffers (30 min) at the remote site, data will be lost. For sites in tropical climates a larger capacity ring buffer is

recommended. On day 322 and 323, 60-100 mph winds swept Boulder causing the hub antenna to misalign. In this case the remote stations were unable to authenticate with the hub so data were lost once the ring buffer size was exceeded. Note that because the hub was misaligned this caused data loss from both remotes. We suspect that the 3 meter diameter of the hub antenna which was exposed to the oncoming winds contributed to the misalignment.

In Figure 10 we look at the number of retransmission requests from the hub to each remote station. The hub sends out a retransmission request when it determines that a data packet has been missed based on a discontinuity in the packet sequence numbers. Discontinuities occur due to problems at the remote, a power loss for example, attenuation due to rain-fade, local LAN problems, or bandwidth limitations. Most of the retransmission requests seen in Figure 10 are for data housekeeping at the hub.




Figure 10. The percent of retransmissions calculated as a daily average. High retransmissions correlate with periods of data outage (Figure 9). The remaining retransmission requests are for data housekeeping at the hub.

Figure 11 shows the number of data packets lost which were not attributed to remote power issues. Figure 11 illustrates that of the majority of retransmission requests indicated in Figure 10, the actual number of lost data packets is small and can be traced to specific causes. Note that a packet of data can contain between 1 and 232 bits. Thus, for scale, the outage at the MARS on day 271 was 20721 packets or the equivalent of 24 Mb of data or about $\frac{1}{4}$ of the days data. The results from Figure 11 indicate that over a 72-day period the Libra network is 99.8% reliable.




Figure 11. Lost data packets not attributable to power outages. Significant periods of data loss are labeled and explained in the text.

The windstorm on days 322 and 323 that misaligned the hub resulted in data loss from all remote sites. For large-scale networks, a second hub should be considered, to eliminate this single point of failure.

To determine data latency for the Libra system we compare the time tag of each GPS observable at the time of transmission against the time the data packet was received at the local data hub. We used 24 hours of data from the MARS station with the TDMA frame time set to 1 second.




Figure 12. Latency values for 1 second GPS data from station MARS for a 24-hour period. Note that the majority of data packets arrive within 1.2 seconds.

Figure 12 indicates that, over a 24-hour period, there is a <2.5 s latency with a majority of the data arriving within 1.2 seconds.

Conclusions

This study has shown that the StarBand and Nanometrics Libra VSAT systems can be used to transfer GPS data in near real time in areas where

no communications infrastructure exists. With the Libra system the user is in control of the bandwidth and can configure the data burst duration at each remote individually, and thus can accommodate a variety of instrumentation and data rates at remote sites. Control of bandwidth comes at a financial cost in terms of purchasing and supporting a data hub. The Libra system is low power (<20 Watts) and can be used at sites where only DC power is available. Our experience with the Libra system indicates that 1 second GPS data can be delivered from remote sites with a <2.5 s latency and a 99.8% reliability.

The StarBand system is an alternative for permanent station applications with less stringent latency requirements where the field area is located within the coterminous US. The power draw on the StarBand system is high (30-40 Watts) and currently no DC only powered satellite modems are available. The system cost is lower (\$400 for dish and modem plus \$70/month for service) than the Libra system tested. However, the StarBand bandwidth is shared by an ever-increasing pool of users and may eventually lead to degraded service such as greater latency, slower uplink speeds, and network outages. Our experience with the StarBand system indicates that 1-second GPS data can be delivered from remote sites with 98.9% reliability and an average uplink speed of 31.7 kbps.

Acknowledgements

This work was performed as part of NSF grant 9910789 and NASA grant 2000-175. We thank Neil Spriggs, Brad Tavner, Emil Farkas, and Joseph Czompo of Nanometrics for their assistance and Ron Muellerschoen at JPL for assistance with the RTNT software.

References

Estep, L., & Meertens, C.M. (1999). TEQC: The multi-purpose toolkit for GPS/GLONASS data. *GPS Solutions*, 3, 42-49.

Jackson, M. C. Meertens & C. Rocken (2000). Real-time Data Streaming from GPS Networks, Paper presented at the International GPS Service Network Workshop 2000, Oslo, Norway.

Maral, G. (1995). *VSAT Networks*, Chichester England John Wiley & Sons.

Muellerschoen, R, Bar-Sever, Y.E., Bertiger, W.I., & Stowers, D.A. (2001). NASA's Global DGPS for High-Precision Users. *GPS world* (12)1.

Talaya, J. & Bosch, E (1999). CATNET: A Permanent GPS Network With Real-Time

Capabilities, ION GPS'99 Nashville,
Tennessee, EUA, 14-17 December 1999.
