Conference: Space Mission Challenges for Information Technology Symposium 2003

Paper: "High Capacity Ground Communications to Support Future Space Missions"

Author: Richard W. Markley, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA.

1 INTRODUCTION

NASA/JPL expect that mission requirements for many deep space missions will include an increase in downlink data rates by at least a factor of ten every ten years; i.e., approximately 10 Mbps in 2010 and 100 Mbps in 2020. Because of the high costs of supporting an international antenna network, such missions in their aggregate will have significant impact on the Deep Space Network (DSN) ground architecture. In addition, cooperation among distributed partners will continue, which increases data distribution costs to the various domestic and international operation centers and principal investigator teams.

In this report, high-rate ground communications are considered those that can support spacecraft data flow from the antenna to the mission operations center at rates higher than 10 Mbps.

Mission data flow from the antennas to the mission operations center (MOC) usually falls into at least the two following streams of data:

1. Real-time data: that which is required immediately for spacecraft monitor and control
2. Non-real-time data: that which can be delayed slightly, but is desired without errors as soon as possible for engineering and scientific analysis.

Real-time data is usually low volume, and should be given priority in the underlying ground communications network. Non-real-time data may be high volume as it may include images and other high-volume scientific data.

Due to the high rates, spacecraft processing systems must make a distinction between these data streams as it is passed to the ground communication system. This division between "real time" and "non-real time" is a major architecture driver. The delivery requirements for scientific data streams can be divided into two more categories: (1) spacecraft data that must be transmitted within a certain number of hours for ground processing, and 2) data that which is simply required before the next downlink to avoid overflowing intermediate buffers.

2 TRADITIONAL ARCHITECTURE

The traditional method for spacecraft data delivery has been over dedicated circuits from the ground station to some intermediate point such as a multimission operations center, and then to the mission operations center (MOC) and payload operations centers (POCs). Circuits provide an easily predictable data rate, and with multiplexing, can support concurrent synchronous services such as operational voice and video. The architecture can be designed to provide appropriate capacity and availability. However, the costs of dedicated circuits, especially for international services, are fairly high, and when the cost of supporting future missions is estimated, other types of wide area network services appear to be more cost effective.

Ground data system backup requirements are also important to consider, especially for real-time data. Again, the traditional architecture involves dual circuits, or perhaps one dedicated circuit and one dial up circuit. Usually the router will detect that the circuit has failed and will automatically connect over the alternate path. After recovery, the router will complete the transition back to the nominal configuration.

A technique used in the DSN, between the Network Operations Control Center and the stations, is to depend on at least two circuits. In the nominal case both circuits are available for Internet Protocol (IP) transmission. When one fails the other is sized to handle the expected backup requirements. This technique uses a common protocol within the routers called multilink point-to-point protocol (MLPPP). Recovery time in the event of an outage is on the order of seconds.

Another reliable architecture for transmission to remote locations such as MOCs involves a dedicated circuit and ISDN circuits. When the dedicated circuit fails, the ISDN circuit is established, and then disconnected upon circuit recovery. Multiple ISDN 64-bps lines can be automatically established depending the backup requirements.

The traditional configuration for voice and video are dedicated circuits, or channels within dedicated circuits, with proprietary protocols. This architecture requires manual intervention for service recovery, providing little or no automatic failover capability.

Lastly, for the IP data networks the traditional technique for monitoring the health of the ground network is usually an application that depends on the Simple Network Management Protocol (SNMP) to report service outages. Failure of non-IP services often depends on reports from customers of the outage, unless the circuit has a dual use, that is the channels are divided into both IP uses and non-IP uses, in which SNMP will report the loss of the IP services, and by deduction it can be assumed the non-IP services are lost.

3 HIGH-CAPACITY GROUND COMMUNICATIONS

Besides data, today the Internet Protocol (IP) supports all traditional synchronous services including voice and video. However significant management of quality of service is required because of the tendency for voice to be distorted when the IP network becomes congested. Thus, significant control is needed over quality of service. This has been achieved in the DSN and is planned for the near future.

Video broadcasts can also be supported over an IP network, and slight delays due to congestion are somewhat tolerable. This convergence is also planned in the near future.

Thus the goal in a high capacity ground network is to transmit data, voice, and video entirely using a reliable underlying IP network. That is not to say the Internet, but a secure intranet with reliable ground communications architecture. Also there must be significant management control of quality of service.

3.1 Dedicated Circuit Architecture

The obvious step to achieve a high-capacity network is to order multiple high-rate dedicated circuits from a telecommunications carrier. However this can be quite costly, especially in a global network such as the Deep Space Mission System.

Annual carrier costs to support the international Deep Space Mission System ground communications (Australia, US-California, and Spain) with multiple T1 circuits are declining, but at the time of this study the annual cost to all the stations concurrently ran approximately $1 million for every T1. Thus, for 10 Mbps, the annual costs could be as high as $7 million per year. This is a baseline to measure alternatives.

In our discussions with overseas carriers, we have discovered that for international circuits, the crossover point in recurring annual costs is about 3 T1s (4.5 Mbps) costs the same as 1 T3 (45 Mbps). Thus, if the demand is higher than 9 Mbps it is just as cost effective to order two T3s (aggregate at 90 Mbps). This crossover point is different for US circuits (and in fact in any country where the antenna complex is distant from a metropolitan area. In the case of the Deep Space Network, a dedicated T3 circuit to Goldstone, CA, out in the Mojave Desert has the potential to cost as much as a circuit to Australia. However, we have used an industry rule of thumb for the Goldstone case, six T1s cost the same as one T3.

Option 1: High-Rate Dedicated Circuits from Australia; Goldstone, California; and Spain to Pasadena

Configuration	Maximum Nominal Bandwidth	Worst Case Back-Up Bandwidth with Loss of 1 Circuit	Description	Estimated Annual Recurring Cost, $k
7 T1s	10.5 Mbps	9.0 Mbps	7 T1s. with 6 T1 Backup	$7,100k
1 T3, 2 T1s	48 Mbps	3.0 Mbps	1 T3 with 2 T1 Backup	$5,200k
2 T3s	90 Mbps	45 Mbps	2 T3s with 1 T3 backup	$6,200k

Thus for rates higher than 10 Mbps primarily to international locations, 45 Mbps circuits are very cost effective.

3.2 Public Internet Architecture

A candidate architecture has been evaluated that provides dedicated circuits for voice and real-time data traffic and best-effort IP services from the public Internet that are high rate but shared with other customers of the carrier. Public Internet access costs are significantly lower than end-to-end dedicated circuits. Basically the customer pays for dedicated service to the local Internet Service Provider and then shares in the costs of the Internet backbone through ISP charges.

Our studies revealed that Internet access charges were about 15%-25% of the dedicated circuit charges to the locations of the DSN. However, one of the drawbacks was that in some countries Internet users also pay a per-byte surcharge. This is modified in some countries so the surcharge is levied only on downloaded traffic. Since most of the DSN traffic is uploaded from the foreign country to the US, these charges are minimal. Another drawback was that to use VOIP, dedicated circuits may be necessary to control quality.

The main concerns with the public Internet are 1) unplanned congestion and 2) security. In some trial tests over several months between JPL and the Canberra antenna complex we found that congestion was never a problem and in addition the lowest rate link between the two locations was over 300 Mbps (!). Security was effectively dealt with by using encryption (virtual private networks-VPNs) on the path. Another security issue is the possibility of a denial of service attack by hackers. We are still addressing that issue. We came to the conclusion that this technique can be very cost effective, assuming the data was not time critical, such as large science files.

Option 2: Internet Communications from Australia; Goldstone, California; and Spain to Pasadena

Configuration	Maximum Nominal Bandwidth	Worst Case Back-Up Bandwidth with Loss of 1 Circuit	Description	Estimated Annual Recurring Cost, $k
7 T1s	10.5 Mbps	9.0 Mbps	7 T1s. with 6 T1 Backup	$7,100k
2 T1s and 1 T3 ISP	48 Mbps	3.0 Mbps	2 T1 for voice and backup, 1 T3 ISP	$2,900k

Outside of the drawbacks mentioned above, this is a cost effective approach to support missions with high-volume science content.

3.3 Carrier IP-Based Service Architecture

Some carriers provide an IP-based service in which they encrypt the data and guarantee quality of service parameters for the data. For example 30% of the data is guaranteed to be delivered in real-time (no congestion delays). This architecture concept is under investigation and test. Presumably the carrier can guarantee that certain data flows within the IP stream can be transmitted with no delay, while others may be best efforts. The unofficial costs seem to lie between the dedicated circuit architecture and the public Internet architecture.

4 OTHER ANTENNA REQUIREMENTS

NASA is committed to demonstrating optical space communications on an operating spacecraft by Year 2009. This space communications technique has the potential for communication rates up to 100 Mbps. However, one of the basic requirements for optical communications is that the area be dust-free. This implies high altitude and relative geographic isolation.

At the present time, we are investigating satellite communications from remote locations to nearby DSN antenna stations with high-capacity communications to JPL in Pasadena, CA. The new communications component will be communications satellites, and work remains to define capacity requirements and pricing.

5 CONCLUSION

This study investigated ground communications architectures that can support mission requirements of approximately 10 Mbps in 2010 and 100 Mbps in 2020. It was found that for the 2010 period standard 45-Mbps dedicated circuits look feasible and cost effective. Using the public Internet has some serious drawbacks including potential congestion and denial of service attacks. There are new emerging carrier services based on VPN and managed IP networks that may provide dramatically cheaper ground communications without the drawbacks of the public Internet.

The research described in this paper was carried out by the DSMS Strategic Planning Office, Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.