

110779

2/6-21

31751

P-6

ECONOMICAL GROUND DATA DELIVERY

Richard W. Markley, Russell H. Byrne, and Daniel E. Bromberg

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California, 91109

ABSTRACT

Data delivery in the Deep Space Network (DSN) involves transmission of a small amount of constant, high-priority traffic and a large amount of bursty, low priority data. The bursty traffic may be initially buffered and then metered back slowly as bandwidth becomes available. Today both types of data are transmitted over dedicated leased circuits.

The authors investigated the potential of saving money by designing a hybrid communications architecture that uses leased circuits for high-priority network communications and dial-up circuits for low-priority traffic. Such an architecture may significantly reduce costs and provide an emergency backup. The architecture presented here may also be applied to any ground station-to-customer network within the range of a common carrier. The authors compare estimated costs for various scenarios and suggest security safeguards that should be considered.

INTRODUCTION

The DSN is a geographically distributed antenna network with antenna complexes in Canberra, Australia; Goldstone, California; and Madrid, Spain. The DSN is managed, technically directed, and operated for the National Aeronautics and Space Administration (NASA) by the Jet Propulsion Laboratory (JPL) of the California Institute of Technology in Pasadena, California. Data communications between the complexes and JPL include telemetry, command, tracking, radio science, and monitor and control information. Downlink telemetry data are usually acquired at the remote complexes and transmitted to JPL for further processing, and ultimately delivered to customers located anywhere in the world.

GROUND NETWORK TECHNOLOGY

Spacecraft data are usually delivered over carefully engineered data networks because of their high scientific value and irreplaceability. The DSN is in the midst of upgrading its ground networks to use the Transmission Control Protocol/Internet Protocol (TCP/IP) suite of networking standards, and intermediate buffers. This new architecture provides useful services such as automatic error detection, recovery, flow control, and fault-tolerance. This transition to TCP/IP makes it possible to use commercial, off-the-shelf network devices such as routers and bridges to interconnect local and wide area networks. In addition, the architecture enables NASA to potentially use emerging cost-saving technologies. One such technology that we have investigated provides dial-up bandwidth-on-demand. The enabling devices are dial-up routers and inverse multiplexers, which are an advancement of dial-up router technology.

Dial-up routers are very similar to traditional routers, only they include a network interface to a switched circuit. Whenever the user attempts to send data to a predefined

site, the router signals its interface to dial a dial-up router at the remote site and the connection is established. At the completion of the call, the connection is terminated. The user only pays for the time that the call takes place, plus a relatively small monthly fee (similar to telephone service).

Inverse multiplexers have the additional capability of aggregating multiple independent switched circuits to create a single higher-rate channel. An inverse multiplexer segments the data in the outgoing data stream and sends the streams out over the individual channels. At the receiving end, the inverse multiplexer accepts the data from these channels, reorders the segments, and compensates for variances in channel transit times. Inverse multiplexers can also add or remove channels from the aggregated connection without terminating the connection. This allows the total amount of bandwidth between the two sites to vary according to real-time bandwidth requirements—for economies of operation. This feature is sometimes referred to as dynamic bandwidth allocation. One of the penalties of using this approach is delay associated with establishing phone circuits (5-10 s for digital circuits and up to 30 s for analog circuits).

Interoperability is another important issue. Early inverse multiplexers implemented proprietary protocols to combine digital channels to form a transparent aggregate stream of data. Units had to be bought in pairs from the same vendor in order to achieve connectivity.

In September, 1991, the Bandwidth on Demand Interoperability Group (BONDING) was formed. Version 1.0 of the BONDING standard was published in September of 1992, and the first conformance event was held in April, 1993. The specification defines a frame structure and procedures for establishing an aggregate channel by combining multiple switched channels. It is now possible to implement networks using inverse multiplexers from several different vendors (there are 31 equipment manufacturers represented in the BONDING group).

The Integrated Services Digital Network (ISDN) is still unavailable in many areas, and just beginning to be supported by several of the BONDING manufacturers. An alternative technology, which is more widely available and supported, is the 56-kbps switched type (or "Switched-56") provided in most cities in the U.S. by commercial phone carriers. Since such circuits are entirely digital, they have low bit error rates and provide an economical, reasonably sized increment of bandwidth. Bandwidth-on-demand devices also work with analog modems. These modems can run wherever analog (Plain Old Telephone Service—POTS) phone service is available (i.e. almost anywhere in the world). There are several disadvantages: 1) circuit quality can vary widely, from virtually error-free to unacceptably noisy in which much bandwidth is wasted on error-correction, 2) the analog lines are only guaranteed for transmitting and receiving 4800 bps by the local service provider, and 3) calls take much longer to establish because of low-speed protocol negotiation and carrier detection. While compression standards such as V.42 *bis* create a virtual maximum throughput of 56 kbps, this maximum is rarely, if ever attained, and in practice throughput varies widely depending on the compressibility of the data.

While installation and monthly line costs are substantially cheaper on an analog phone line versus a Switched-56 digital line, the serious disadvantages discussed above make the analog option impractical except for (1) maintaining a single analog backup line should the digital system fail, and (2) in the event of a power outage, the analog system can be used during the period of time that access equipment is powered by

Uninterruptable Power Supplies (UPS). (The digital lines are powered by the customer, while the analog lines are powered by the service provider.)

3 GROUND ARCHITECTURES

3.1 Remote Antenna Complexes to Pasadena, CA

DSN ground communications from the antenna complexes to Pasadena are currently over dedicated satellite circuits that are exceptionally clean (error-free 99.5% of the day), secure, and dependable. The overseas links are very expensive because of the distance and generally higher cost of telecommunications in foreign countries.

An example of the nature of customer traffic can be deduced from the aggregated spacecraft downlinks at Goldstone, CA illustrated in Figure 1. These are the data that must be delivered to customers such as spacecraft teams, principal investigators, and non-NASA partners. Some of the traffic is "real-time," and must be delivered as quickly as possible. The real-time traffic from the stations to JPL usually totals less than 200 kbps. This traffic includes spacecraft engineering data, quick-look data, and other critical data. These data tolerate very little additional latency (over and above the expected 270 ms satellite propagation delays). They are not candidates for dial-up bandwidth, nor error correction techniques made possible by TCP/IP. This traffic requires dedicated circuits.




Figure 1 Aggregated Spacecraft Downlinks at Goldstone on January 8, 1994

The remainder of the traffic may be delayed to provide additional communications services such as automatic error correction or to balance the load on the ground circuits to Pasadena. The telemetry delivery system is capable of prioritizing the data and handling it appropriately.

This data may be buffered at the station or at Pasadena before being passed on to the customer. As shown in Figure 1, it is bursty (the result of brief spacecraft visibility for Earth orbiters). Switched circuits are ideal for delivering these bursty streams because: (1) the streams occur for brief periods of time, (2) there is no critical latency requirement, and (3) the streams are delivered with TCP/IP protocols, which provide appropriate flow control mechanisms and are compatible with inverse multiplexers.

Leased circuits make sense when the circuit is utilized most of the time. When considering the option of using leased versus dial-up lines, the monthly and per-hour cost of dial-up lines and the monthly cost of the leased lines can be expressed in an equation which is linear in dial-up hours. This can then be solved for the "break-even" number of hours per month between the two approaches. If the circuit will be used more than this value, it makes more sense to lease.

The costs involved depend on many things: the distance between the endpoints of the communications channel, whether or not this distance spans local service provider areas, the long-haul carrier (if any) used, the discount program used for leased lines (the longer the lease, the better the monthly rate), and whether or not data transmission can be scheduled to take advantage of lower evening and night time toll charges.

The resulting network architecture (Figure 2) has a limited amount of dedicated bandwidth for real-time traffic and optional "elastic" bandwidth for the lower-priority traffic. The traffic flows initially to the router where its priority is determined and the low priority traffic is shunted to the inverse multiplexer. The inverse multiplexer establishes circuits as required. In addition, in the event of losing the dedicated channel, the router may reroute high priority data to the inverse multiplexers to provide emergency communications channels.




Figure 2 Complex-to-Central Site Network

3.2 Pasadena, CA to Customers

Once the non-real-time data is processed at JPL, it is transmitted to the individual customers. A typical example is illustrated in Table 1, which lists the preliminary plans for supporting the upcoming Cassini mission. Table 1 identifies the locations of the customers for the Cassini down-link data and the expected data rates. None of the traffic is in the real-time category.

Table 1 Candidate Cassini Customer Locations

Customer	Location	Required Bandwidth (kbps)
European Space Operation Center	Darmstadt, Germany	56
Goddard Space Flight Center	Greenbelt, MD	56
Southwest Research Institute	Phoenix, AZ	112
University of Arizona	Tucson, AZ	56
University of Heidelberg	Heidelberg, Germany	56
University of London	London, England	56
Johns Hopkins University	Baltimore, MD	56
University of Iowa	Ames, Iowa	56
University of Colorado	Boulder, Colorado	56

There are several options for data delivery. The first is to use traditional dedicated circuits, which may or not be cost-effective depending on the volume. The second is to transmit the data from JPL to the customer over the Internet since these particular customers are on the Internet. Security safeguards are necessary, such as secure local area networks at the customer sites for hosts that perform spacecraft data processing.

The third option is to use dial-up routers and inverse multiplexers, and establish dial-up circuits as required. Security safeguards available with inverse multiplexers include: (1) encrypted password protection, (2) dial-back features, and (3) data encryption.

3.3 Remote Antenna Sites

In addition to the DSN architecture, dial-up routers and inverse multiplexers may support remotely located antenna sites, assuming that there are common carrier services in the area. In this case there may be both monitor and control and telemetry data. If the station is used as a transmitter, there may also be command uplink data. The volume of data will determine the data rate required for the individual channels.

Assuming that the volume of data is relatively low, a low-cost architecture could involve one leased circuit and one dial-up circuit (Figure 3). We estimated communications costs for such a system between Goldstone and a customer site in Pasadena with 56-kbps circuits. Such a configuration could support volume up to 605 Mbytes per day over the dedicated circuit and cost-effectively support up to an average of 20.8 Mbytes per day (625 Mbytes/mo.) over the switched circuit. Above 605 Mbytes/mo. a second leased circuit would be more advantageous.

The details of the crossover volume of data calculation are as follows: A leased 56k line from Barstow to Pasadena costs \$538 per month. Switched-56 service is \$77 per month plus \$18.60 per hour in toll charges. So the equation gives a value of 24.8 hours




Figure 3 Remote Site-to-Customer Network

of connectivity per month as the crossover point. At a data rate of 56 kbps, this corresponds to an average volume of 20.8 Mbytes of data per day.

4 SUMMARY

This paper proposes a hybrid communications architecture that uses inverse multiplexers and dial-up circuits in addition to traditional leased circuits for spacecraft ground communications. Such an architecture may significantly reduce costs. In some cases it may significantly reduce the delivery time by providing additional bandwidth on demand. With appropriate security safeguards, non-critical data may be sent directly from the antenna complex to the end user. Therefore, the network architecture presented here may be applied not only to the DSN, but to any ground station within the range of a common carrier.

The research described in this paper was carried out by JPL, California Institute of Technology, under a contract with NASA.