516-32
185876
p-7

N94-14385

Asynchronous Transfer Mode Link Performance Over Ground Networks

E. T. Chow and R. W. Markley
Advanced Information Systems Section

This article reports the results of an experiment to determine the feasibility of using asynchronous transfer mode (ATM) technology to support advanced spacecraft missions that require high-rate ground communications and, in particular, full-motion video. Potential nodes in such a ground network include Deep Space Network antenna stations, the Jet Propulsion Laboratory, and a set of national and international end users.

The experiment simulated a lunar microrover, lunar lander, the DSN ground communications system, and distributed science users. The users were equipped with video-capable workstations. A key feature was an optical fiber link between two high-performance workstations equipped with ATM interfaces. Video was also transmitted through JPL's institutional network to a user 8 km from the experiment. Variations in video depending on the networks and computers were observed, and this article reports the results.

I. Introduction

The traditional role of the DSN has been to support deep-space robotic missions with relatively low data requirements (40 bps to 300 kbps). NASA is considering future missions that require high-rate (>1 Mbps) space communications and ground support from the DSN. These missions include a large number of low-cost Earth orbiters (0.1 to 10 Mbps), lunar missions with rovers and human exploration (10 to 100 Mbps), and Mars missions (1 to 10 Mbps). In order to support higher rate missions, the telemetry, command, and ground communications architecture must be significantly upgraded. This article reports on one potential upgrade to the ground communication system.

The most recent upgrades to the DSN include 2.4-Mbps telemetry processing and 865-kbps real-time ground communications, somewhat short of the requirements for future missions. The command uplink is limited to 2 kbps. The focus of this experiment was to evaluate an advanced ground communications architecture to support a lunar rover mission (Fig. 1).

The present ground communications architecture is designed to (1) transmit data directly to the user in real time, (2) buffer the data and meter it back slowly to the user (near-real time), and/or (3) store the data on transportable media for shipping (nonreal time). As an example, small amounts of data may be transmitted to users

in real time to support mission planning. Most telemetry data, such as data from the recently launched Mars Observer, are delivered in near-real time. Very-high-volume data, such as data generated by the Magellan Venus radar mapping mission, are handled in nonreal time in order to reduce communications costs.

In order to support a lunar rover mission, the DSN must support devices with radically different communications requirements than those for deep-space probes. A rover can be expected to simultaneously downlink video and telemetry, and receive commands from the rover operator on an uplink channel. It is highly likely that the operator, presumably a principal investigator, will want to be able to control the rover in real time, tolerating only a 1- to 2-sec propagation delay due to the distance of the Moon. The data rates for the video and telemetry could range from 1 to 10 Mbps, depending on the degree of video quality required. The command stream would be low (2 to 10 kbps), but with an additional requirement to be virtually real time with no ground communications delays.

One of the keys to success from the user's perspective in such a mission is an architecture that provides telepresence. Telepresence may require multiple displays to enable the rover operator to sense—at least visually—the immediate environment and take advantage of the rover's maneuverability in a scientifically interesting location.

A colony of microrovers would require several video channels for each microrover. Therefore, the scientists' workstations should be capable of managing multiple video sessions on one screen. In addition, real-time video conferencing among scientists would enable cooperative experiments.

The communications architecture should also support other services in addition to high-speed downlinks and real-time connectivity. For example, video, telemetry, and command data should be received with no detected errors. To reduce costs, the video channel should be compressed and decompressed with no loss of data (nonlossy). Greater circuit efficiency should be possible by dynamically allocating bandwidth to data and video in order to take advantage of lapses in transmission.

New network technologies enable high-speed data communications among local users and widely distributed users. For example, commercial fiber local area networks (LAN's) are capable of speeds of up to 100 Mbps. Traditional commercial services for wide-area networks are available at speeds up to 45 Mbps. Switched-based architectures based on the International Telegraph and Telephone Consultative Committee (CCITT)-standard asynchronous transfer mode (ATM) protocol will extend these rates into the range of gigabits per second. This switching technology, coupled with virtually unlimited bandwidth in optical fiber media, justifies consideration of a new approach to a ground network infrastructure to support space missions.

II. Technology

In 1990, the CCITT chose ATM as the standard switching technique for transport within the Broadband Integrated Services Digital Network (BISDN). BISDN is an emerging digital network planned by common carriers for integrating high-speed data, voice, and video services at transmission speeds starting at 155 Mbps. The following is a summary of ATM; a more complete, highly readable description may also be found in [1].

At the core of the ATM switching technique is the concept of cell switching. Cell switching refers to multiplexing of multiple logical data flows on a single physical interface, where the information from each source is in the form of fixed-size packets called cells. CCITT members, including both European and North American standards-setting groups, agreed on a standard ATM cell size of 53 eight-bit bytes. Five bytes are used for header information and 48 bytes carry the user information.

The ATM standard comprises three layers in the BISDN protocol stack. The BISDN stack is slightly different from the traditional Open System Interconnection (OSI) stack, but uses the same layering concepts. In the BISDN model, the three lower layers are occupied by ATM (Fig. 2). They correspond roughly to the bottom two layers of the OSI model (physical and data link layers). The three lower BISDN layers are (1) the physical layer, (2) the ATM layer, and (3) the ATM adaptation layer (AAL). The physical layer defines physical interfaces and framing protocols for the network. Physical connections in ATM are based on synchronous digital hierarchy (SDH) principles, but are not restricted to SDH framing protocols, and are thus very flexible. Because the functions are layered, the physical interface may be changed without impacting the higher layers. For example, the physical layer may be changed from Synchronous Optical Network (SONET) to T1 or T3 with no impact on the overlying ATM layer.

The ATM layer defines the information within each cell, and how cells flow over the logical connections in an

ATM network. Figure 3 shows the structure of those cells. Cells belonging to different virtual paths are distinguished by their virtual path identifier (VPI) and virtual channel identifier (VCI). Virtual paths can contain multiple virtual circuits. Each virtual circuit can be either permanently established or set up dynamically. Other fields in the 5-byte header of the ATM cell include a generic flow control (GFC) field, a payload type (PT), a cell loss priority (CLP) bit, and header error control (HEC) field. As mentioned earlier, the 48-byte information field carries user information.

While the AAL layer is a part of the overall communication process, it is not a network process, but is a process performed by the user's equipment. By keeping AAL functions separate, the ATM network becomes primarily concerned about delivering 48-byte payloads based on the information contained in the header. The adaptation layer plays a key role in making it possible for multiple types of services to use the ATM network. There are five types of services:

(1) Constant bit rate (CBR) services for isochronous traffic, such as voice and video.

(2) Variable bit rate (VBR) timing-sensitive services for features such as compressed video.

(3) Connection-oriented VBR Data Transfer for intermittent, large, long-period data file transfers over preestablished ATM connections. AAL 3 provides error detection.

(4) Connectionless VBR data transfer for short, bursty messages over non-preestablished ATM connections where call setup time could be longer than the transfer time.

(5) Simple and efficient adaptation layer (SEAL), for services similar to AAL 3 but with the assumption that the user's higher-layer processes provide error control.

The fundamental parts of the ATM standard have been defined by CCITT. The user network interface (UNI) (which incorporates the AAL layer) is being developed by the ATM Forum, a consortium of interested users and developers. The ATM Forum will also address switch-to-switch interfaces, and signaling specifications across the user interface, and establish network management specifications.

III. Experiment Configuration

The experiment simulated a lunar rover exploring the surface of the Moon (Fig. 1), with a downlink to science operations on the Earth. The major links in a full simulation would have included a lunar radio link between the rovers and a lander, a radio link from the Moon to the Earth's deep-space stations, overseas communications to JPL, and a link to science operations. The focus of this experiment was very limited; it was to determine the feasibility of using ATM networks on the Earth for transmitting data from the antenna stations to science operations. Therefore, design and prototyping of the space segment were postponed for future work.

Figure 4 illustrates the laboratory configuration. The rover and simulated lander were in a laboratory equipped with a rock-strewn floor. The principal investigator (PI) environment was in an adjacent room. Another investigator was in an office 8 km away. The experiment was run using the resources of the DSN Information Systems Engineering (ISE) Laboratory.

A. Microrover

The microrover was equipped with a miniature television camera, a radio transmitter, and a 2.4-kbps radio modem for transmitting telemetry and receiving commands.

B. Computers and Networks

The computer in the lunar environment (A, Fig. 4) was a Sun Sparcstation 10/41 rated at 96 million instructions per second (MIPS). The PI environment computer (B) was a Sparcstation 2 (24 MIPS). Each workstation had an Ethernet interface, ATM interface, and a video digitizer. The computers had the following capabilities:

(1) Video processing and compression were accomplished with Parallax video digitizers. Compression was based on the Joint Photographic Experts Group (JPEG) standard. Each board was capable of processing two video signals simultaneously; this enabled the science users to simultaneously oversee the rover and its immediate terrain.

(2) The Internet suite Transmission Control Protocol/ Internet Protocol (TCP/IP) was the common protocol suite among both systems.

(3) The ATM interfaces were Fore Systems SBA-100 S-Bus ATM computer interfaces. The ATM physical layer was a 140-Mbps optical fiber interface. The user-level application interface was a Berkeley socket interface. Underlying software in the drivers utilized AAL 4 protocols, providing cell-level error detection (but no correction). If a bit error was detected by the driver, the cell was dropped and TCP performed error recovery.

(4) Computer A had an RS-232 interface to a radio modem for sending commands to and receiving instrument telemetry from the rover.

(5) The rover video was also transmitted from workstation A to workstation C, another Sparc 2, over an institutional network. Simple teleconferencing software provided by Parallax made this capability possible, and the software was integrated with other X Windows software to enable remote control of the rover. The path to C was a series of networks of varying types, most notably Ethernet and serial channels, which yielded low-rate, intermittent stop-action video.

(6) A voice/video teleconferencing demonstration was performed among users at workstations B and C using a commercial videoconferencing package called Communiqué. This was also low-rate video.

IV. Results

A. ATM Network Performance

The raw TCP/IP throughput across the ATM network was approximately 14.9 Mbps (11 percent of the bandwidth). This illustrates the limitations of the TCP/IP software implementation in the ATM board and Sun operating system. When video was generated by the Parallax board, the throughput dropped to 3.5 to 4.5 Mbps, depending on the complexity of the images.

Comparable data across a dedicated Ethernet connection was 8.5 Mbps (85 percent of the Ethernet bandwidth) for raw TCP/IP and 1.4 to 1.5 Mbps for video data.

Visually, the SPARC 10 transmitted 8 to 10 frames/sec (fair quality video) over the ATM network.

Two suggestions to improve the video and use the complete bandwidth of ATM and the associated optical fiber networks were (1) an alternate implementation to TCP/IP and (2) special ATM onboard processing hardware to reduce the central processing unit overhead in the workstations.

B. Compressed Video

JPEG compression can produce a good quality picture at a 20:1 to 30:1 compression ratio. At above 30:1 compression, blocking effects start to show up. Improved compression software may also make an improvement in the quality of the video.

C. ATM Versus Ethernet

This experiment illustrated the advantage of ATM over Ethernet for multimedia communication. The ATM full-motion video was relatively smooth and continuous. When video was sent from workstation A to C, some frames were delayed because the Ethernet was being shared with other users. This video was of poor quality—only 1 to 2 frames/sec. This type of performance has been noted elsewhere [2]. The isochronous nature of an ATM channel is very important for multimedia communications.

D. Teleoperation

It was demonstrated that potential teleoperation of NASA rovers by the general public is feasible. Numerous individuals, ranging from children to higher management personnel, were each able to learn how to drive the rover in less than 3 min. With the help of multimedia networking, it is very easy for anyone to do teleoperation.

E. Video Teleconferencing

Teleconferencing was attempted between a PI in the DSN ISE Laboratory and the remote investigator using commercial software. The video frame rate was low, typically 1 to 2 frames/sec, because of intervening institutional networks. For many teleconferences, such a low rate may be good enough. One technical problem was voice feedback. Two techniques were tried to improve the voice quality: (1) use of an echo canceller (this alone did not do the job) and (2) use of a unidirectional microphone. By combining (1) and (2), acceptable, but not really good, voice quality was achieved.

Acknowledgments

The authors wish to express their appreciation to Charles Goodhart and Julia George, who designed and built the software for this experiment. They also express sincere thanks to Rajiv Desai, who designed and built the microrover used in this demonstration; without his help and cooperation the demonstration would not have been possible.

References

[1] *Asynchronous Transfer Mode: Bandwidth for the Future*, Norwood, Massachusetts: Telco Systems, Inc., 1992.

[2] B. Cole, "The Technology Framework," *IEEE Spectrum*, vol. 30, no. 3, pp. 32–39, March 1993.

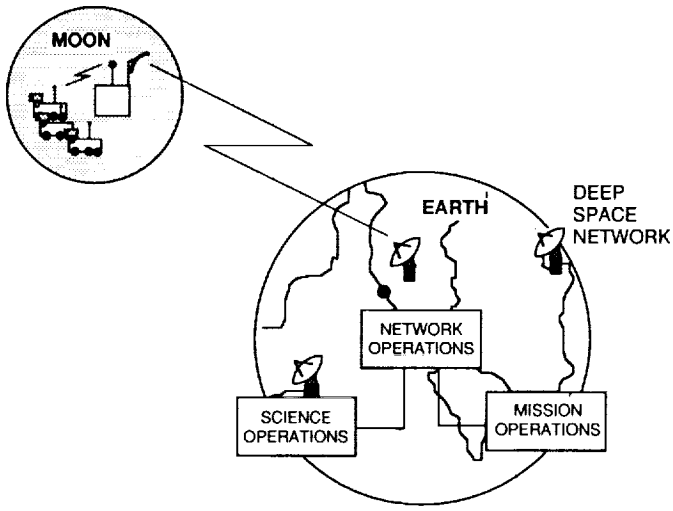

Fig. 1. Lunar microrover exploration.

APPLICATION LAYER	
HIGHER LAYERS	
CONVERGENCE SUBLAYER	ATM ADAPTATION LAYER
SEGMENTATION AND REASSEMBLY SUBLAYER	
ATM LAYER	
TRANSMISSION CONVERGENCE SUBLAYER	PHYSICAL LAYER
PHYSICAL MEDIUM DEPENDENT SUBLAYER	

Fig. 2. BISDN protocol stack.

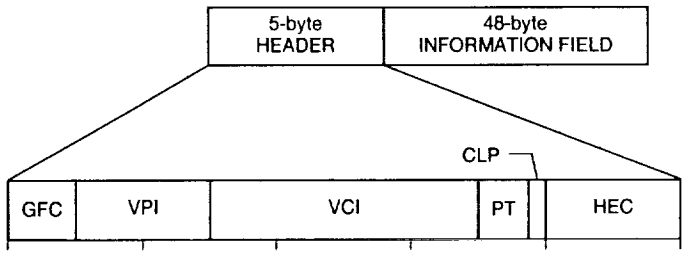

Fig. 3. ATM cell structure.

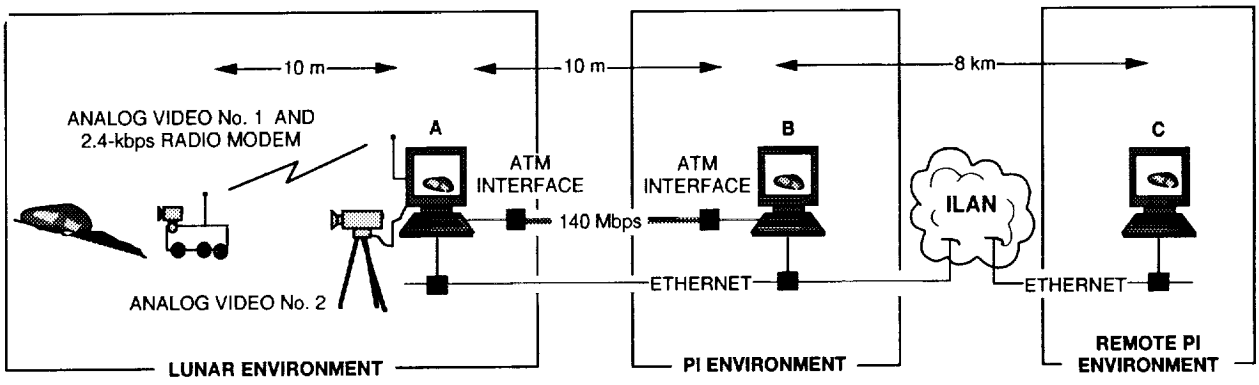

Fig. 4. Experiment configuration.