

Deep Space Network Information System Architecture Study

C. A. Beswick

Information Systems Division

R. W. Markley (Editor), D. J. Atkinson, and L. P. Cooper

Advanced Information Systems Section

R. C. Tausworthe

Information Systems Division

R. C. Masline

Science Data Systems Section

J. S. Jenkins

Control Center Data Systems Development Section

R. A. Crowe, J. L. Thomas, and M. J. Stoloff

DSN Data Systems Section

The purpose of this article is to describe an architecture for the DSN information system in the years 2000–2010 and to provide guidelines for its evolution during the 1990s. The study scope is defined to be from the front-end areas at the antennas to the end users (spacecraft teams, principal investigators, archival storage systems, and non-NASA partners). The architectural vision provides guidance for major DSN implementation efforts during the next decade. A strong motivation for the study is an expected dramatic improvement in information-systems technologies, such as: computer processing, automation technology (including knowledge-based systems), networking and data transport, software and hardware engineering, and human-interface technology.

The proposed Ground Information System has the following major features: unified architecture from the front-end area to the end user; open-systems standards to achieve interoperability; DSN production of level 0 data; delivery of level 0 data from the Deep Space Communications Complex, if desired; dedicated telemetry processors for each receiver; security against unauthorized access and errors; and highly automated monitor and control.

I. Introduction

A. Background

The Deep Space Network (DSN) is the largest, most sensitive scientific telecommunications and radio navigation network in the world. Its principal responsibilities are to support automated interplanetary spacecraft missions and radio and radar astronomy observations in the exploration of the solar system and the universe. The DSN also supports high-Earth orbiter, lunar, and shuttle missions.

The DSN is managed, technically directed, and operated for the National Aeronautics and Space Administration (NASA) by the Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, California. This study responds to the following action item developed by the 1990 JPL Telecommunications and Data Acquisition (TDA) NASA Office of Space Operations (OSO) Planning Workshop: "Perform a study of the DSN information system architecture and recommend guidelines for its evolution."

It is timely that such guidelines be developed because of such improvements in information technology as:

- (1) Dramatic advances in computer hardware and software.
- (2) Reduced size of computers and other processing elements due to very large-scale integration (VLSI).
- (3) Dramatic increase in ground communication data rates using optical fiber.
- (4) Improved ability to internetwork local area and wide area networks.
- (5) Efforts by the U. S. Government and the international community to standardize network protocols.
- (6) Efforts by the U. S. Government to standardize software.
- (7) Advances in distributed computing.
- (8) Advances in providing assistance to decision makers by means of artificial intelligence (AI) and visualization techniques.
- (9) Advances in computer-human interfaces, including multimedia, interactivity, computer-supported cooperative work, and virtual reality.

New projects in NASA's interplanetary exploration program require increased communications and tracking capabilities that must be implemented while maintaining support for existing projects. There is also a requirement

to increase DSN availability beyond its already high level. Additional missions will have a tendency to increase development, implementation, maintenance, and labor force costs, and it is believed that these new technologies will enable significant reduction of these costs.

B. Purpose

The purpose of the study is to develop an architecture for the Ground Information System (GIS) in the years 2000-2010 and to provide guidelines for its evolution during the 1990s. The study also provides a forecast of information systems technologies that are pertinent to the GIS. Where forecasted technologies are inadequate for the projected architectural needs of the GIS, the study recommends appropriate pathfinder work. A flexible transition approach is developed, with the expectation that most improvements to the architecture will appear evolutionary, although they are concentrated on a clearly defined vision of the future architecture.

C. Scope

In order to allow for creativity in the selection of architectural candidates, the information systems boundaries for the study were deliberately chosen without consideration for current organizational boundaries. Specifically, the scope was defined to be the Ground Information System (Fig. 1), which extends from the ground side of the antenna front-end assemblies to the end user (spacecraft teams, principal investigators, archival storage systems, and non-NASA partners). This choice keeps the scope within bounds (i.e., only ground systems), but allows flexibility in defining where ground processing functions are performed. For example, the current interface between the DSN and the Space Flight Operations Center (SFOC) is not considered to be a constraint. After an appropriate technical solution is determined, organizational decisions may be made (outside the scope of this study) and appropriate interfaces defined.

The GIS functions include processing (telemetry, tracking, command, radio science, and very long baseline interferometry [VLBI]), monitor and control of the front ends and the GIS, and data delivery and management. The focus of the study is on multimission functions; project-unique activities will be dealt with only as external requirements. Issues that provide both focus and motivation for the study include:

- (1) The degree of automation and centralization of the monitor and control function.
- (2) The degree of distribution of computing and control functions.

- (3) The feasibility of implementing an open systems architecture.
 - (a) Cost and performance.
 - (b) Security.
- (4) The modularity and reconfigurability of the GIS.
- (5) Communication alternatives.
- (6) The choice of functions to be generic multimission functions.
- (7) The limits and capabilities of technology.

D. Evaluation Criteria

Architectural candidates have been developed that meet planned performance requirements and are evaluated with respect to the following criteria:

- (1) Performance, including the degree of margin for unplanned mission support.
- (2) Life-cycle cost, including development, implementation, maintenance, and workforce costs.
- (3) Operability, including manageability and human-interface simplicity.
- (4) Flexibility, evolvability, and growth potential.
- (5) Availability, including fault-tolerance, reliability, and maintainability.
- (6) Technical risk, including ease of transition.

Emphasis in this study is on general advantages and disadvantages of the candidate architectures relative to a baseline (the current architecture) or to the other competing candidates. Accurate measures of the above criteria (especially cost) require more design effort than is appropriate for a high-level architectural study. Sensitivity of the recommendations to technology advances is considered.

E. Organization of This Article

The remainder of this article is organized as follows:

Section II. Requirements: This section provides a functional partitioning of future requirements, and discusses the current functional architecture, system interfaces, future mission requirements and end-user interfaces.

Section III. Technology Forecast and Projected Impact: This section includes discussions of computers, software, data transport, human interfaces, and their potential impact on the DSN information system.

Section IV. Architectures: Processing, monitor and control, data transport, and software architecture discussions are included in this section.

Section V. Transition Approach: This section describes a plan for evolving in an orderly way from the current architecture to the envisioned architecture.

Section VI. Summary: This section highlights the most significant conclusions derived from evaluating the DSN requirements and technology, including recommendations for achieving the strategic objectives.

Appendix A. Acronyms: This section provides definitions of all acronyms used.

Appendix B. Glossary: This section provides more detailed descriptions of certain terms used.

II. Requirements

A. Functions

A functional partitioning of future requirements for the GIS is shown in Fig. 2. This figure illustrates functional interrelationships of the various subsystems and clarifies areas where functional commonality exists. An easy reference, it should help to ensure that any proposed design will address the functions needed to meet the long-term requirements. Several strategic goals for the GIS are described in Table 1. Some of these goals are included later (Section IV) as evaluation criteria when the architectures are discussed.

1. Processing. This section describes the existing general functions of the GIS. It is believed that future designs will embody most of the major functions that exist in the GIS today. There are three Deep Space Communications Complexes (DSCC's), each identical in terms of function, which are located in Goldstone, California; Madrid, Spain; and Canberra, Australia. Communication circuits that link JPL to the three global sites are managed by the Ground Communications Facility (GCF).

Each DSCC has a number of antenna stations distributed across its site with an associated set of electron-

ics. These antennas and their associated control rooms and electronics are called Deep Space Stations (DSS's). They perform a limited amount of analog signal processing before transmitting an analog signal to a collocated central area called the Signal Processing Center (SPC), shown in Fig. 3. If several antennas are arrayed to detect the usually weak spacecraft radio signal, the analog signals from several DSS's may be combined before further SPC processing. After the SPC has filtered and amplified the resultant analog signal, the telemetry, tracking, radio science, and VLBI data are digitally processed.

In the current system, the resulting digital data may be communicated to JPL over long-haul communications circuits, or stored on magnetic tape and shipped. Various types of processing take place at JPL in the SFOC, including deletion of duplicate data acquired by two DSCC's in the event of overlapping coverage during an encounter. Much of the current SFOC processing is mission-specific; however, it is expected that future missions will share identical standard data structures and coding techniques because of ongoing development of standards by the Consultative Committee for Space Data Systems (CCSDS).

a. Telemetry. Some spacecraft projects request that their data be immediately transmitted to JPL upon reception (i.e., real-time data). Rapid delivery depends on the availability of appropriate long-haul circuits. This type of data is delivered in the order that it is received at the antenna, delayed only by initial processing and the length of time it takes to traverse the network. In the event of errors in transmission (i.e., missing data blocks), retransmission may be requested later by the user—a process called post-pass replay. Other projects prefer that data be stored on magnetic tape and either mailed back to JPL or replayed at a slower rate over long-haul communication circuits (i.e., non-real-time data). Non-real-time delivery is more economical when the spacecraft downlink rates are very high or volume is very large.

The DSCC Telemetry Subsystem (DTM) currently acquires a digital telemetry baseband signal from the receiver and converts it into the format required for further processing at the SFOC to produce a data stream useful to the end user. This telemetry processing function is accomplished through a sequence of transformations that optimizes the amount and correctness of data acquired in the presence of noise. Deep space missions have significant variations in the specific transformations that are applied. There are also variations due to the different noise environments of Earth orbiters and deep space missions.

- (1) The Telemetry Source. Telemetry data originate in the spacecraft, where information from several sen-

sors is multiplexed into a single binary stream and is then modulated onto a subcarrier of the downlink signal. Spacecraft engineering information may be multiplexed into this signal or onto a separate downlink subcarrier. The multiplexed stream is formatted into logical groups, or frames, for synchronization purposes.

The telemetry data may be transmitted directly without further coding if signal-to-noise ratios are high on the spacecraft–Earth link; however, it is usually necessary in deep space missions to code the telemetry data with additional bits and reconstruct the correct data packets on Earth. The two principal coding schemes used in deep space missions are convolutional, such as Viterbi codes, and block, such as Reed-Solomon (RS) codes.

Both coding schemes reduce bit errors for a given signal-to-noise ratio over uncoded transmissions, but the bandwidth cost is the need to transmit additional bits from the spacecraft. For example, the RS code used by Mars Observer has a block code word structure of 10,000 bits, in which 8720 are information bits, and the remaining 1280 are parity (or check) bits. The rate of this code is $R = 8720/10,000 = 0.87$; thus, 87 percent of the bits received are information.

Convolutional codes, on the other hand, are codes in which n channel symbols are transmitted for each information bit. Since each information bit leads to n transmitted channel symbols, the rate is $R = 1/n$. Typical rates are 1/2 for high-Earth orbiters and lunar missions and 1/6 for deep space missions. Thus, in cases where $R = 1/6$ and the symbols are represented by 8 bits, 1 information bit may lead to transmitting and processing up to 48 bits.

Block and convolutional codes can operate on multisymbol sets. For example, the RS code above treats each 8 bits of the code word as one symbol because of the algebraic method used for correcting burst errors. Thus, the RS code above is also referred to as a (255,223) 8-bit symbol RS code.

The implementation of a convolutional coder is simpler, less costly, and more reliable than a block coder for spacecraft operation. However, the information may be RS-coded prior to convolutional coding to make the telemetry data resilient to both random noise and burst errors. The overall rate of information bits per transmitted binary symbol is the product of the two coding processes. A more quantitative discussion of the characteristics of the

telemetry stream (as well as deep space telecommunications in general) may be found in [1].

$$R_{rs} = 223/255$$

$$R_c = 1/6$$

The logic speed $S = 55$ Mbps

- (2) Subcarrier Demodulation and Symbol Synchronization. In order to extract the telemetry data at the station, it is necessary to reverse the processing steps performed on the spacecraft. First, the subcarrier is removed by standard heterodyne means; next, binary symbol synchronization is established. Each binary symbol is integrated for the duration of its symbol time. In the absence of noise, positive and negative integrated symbol values would be interpreted as ones and zeroes, respectively. However, in the presence of noise, such a decision may not always be correct. Rather than commit to a decision at this point, the convolutional decoders are designed to operate on soft symbols, i.e., the magnitudes of the outputs of the symbol integrator. Since the decoders work in digital logic, the soft symbols are quantized into 2^n levels per sample, which yields a number in the range $(-2^{n-1}, 2^{n-1} - 1)$ for each soft symbol.

In the future, subcarrier demodulation, symbol synchronization, and symbol-matched filtering will be accomplished by the planned-for Block V receiver. At that time, the input to telemetry processing will be a stream of soft symbols.

- (3) Telemetry Decoding. The next step in telemetry processing is to transform the soft symbol stream into an information bit stream. For uncoded data (typical of older spacecraft and some Earth orbiters), there is a one-to-one correspondence between symbols and bits; symbols are converted to information bits by simply taking the sign bit of each soft symbol. However, when both convolutional and RS coding have taken place, two more processing stages are necessary.

Current convolutional decoders use 3-bit symbols (eight levels), but performance improvements are possible through the use of 8-bit symbols (256 levels of quantization). The logic speed (S) required to process telemetry data must accommodate the product $QI(1/R_{rs})(1/R_c)$ where Q is the soft symbol quantization rate in bits per symbol, I is the information bit rate, R_{rs} is the Reed-Solomon rate, and R_c is the convolutional code rate. For example, in the case where

$$Q = 8 \text{ bits per symbol}$$

$$I = 10^6 \text{ bps}$$

- (4) Frame Synchronization. Each frame consists of a sequence of binary symbols; the beginning (or end) of a frame is identified by a unique bit pattern called the frame synchronization marker. Output from the frame synchronizer marks the data in such a way that subsequent telemetry transforms the data frame by frame.

For the DSN, frame synchronization is useful as a real-time data-quality monitor, and is a necessary step before doing RS decoding. Among the complications of frame synchronization are errors in the frame synchronization pattern, false synchronization (i.e., data in the frame that coincidentally match the frame synchronization pattern), reversed data (from spacecraft tape recorder dumps), and changes in frame length.

- (5) Output Formatting. Corrected frames from the RS decoder must be transmitted or recorded. Traditional transmission requires that the data be formatted into 4800-bit NASA Communications Network (NASCOM) data blocks. Data in the blocks include the NASCOM header and trailer, the telemetry header, telemetry partial status data, and the actual telemetry data. Some missions have placed non-NASCOM requirements for the content and organization of telemetry headers, partial status, and data. Thus, output processing must support a variety of output block formats.

Newer missions, such as Mars Observer, will require the output blocks to be formatted in accordance with CCSDS recommendations. Processing is required to recognize CCSDS-compliant telemetry packets and provide the capability to route different packets to different destinations based on the virtual channel identifier in the packet header.

b. Tracking. The DSCC Tracking Subsystem (DTK) has the following primary functions:

- (1) Control the transmitted uplink frequency to enable initial signal acquisition by the spacecraft receiver and maintain two-way communications in which the radio signal is transmitted from the ground to the spacecraft and relayed back to the ground receiver in a phase-coherent manner.

- (2) Measure the round-trip travel time of light to the spacecraft, thereby enabling precise calculation of the spacecraft range.
- (3) Measure the Doppler shift on the radio signal received from the spacecraft to give a precise indication of the spacecraft radial velocity.

Support information (generically called predictions or predicts), such as subsystem and instrument settings required for uplink signal control, is generated by the Network Support Subsystem (NSS) at JPL and sent to the DTK to support these functions; the DTK packages raw and computed Doppler shift data into data blocks for transmission back to the project navigation teams.

Frequency ramping commands go from the Metric Data Assembly (MDA) to a digitally controlled oscillator according to time-stepped predicts. The intention is to compensate for the Doppler shift at the spacecraft and keep the signal frequency received at the spacecraft within its closed-loop bandwidth. The MDA utilizes some predicts to construct acquisition profiles that are translated into frequency ramp commands aimed at establishing a two-way link with the spacecraft.

Ranging modulation from the current Sequential Ranging Assembly (SRA) is in the form of square waves or sine waves imposed on the subcarrier in a step-wise progression from highest to lowest frequency. The highest frequency now in use is 1 MHz and the lowest is 1 Hz. Capability is provided in the SRA to go to at least 2 MHz. The detection and exact timing of these signals coherently relayed back from the target spacecraft provides a precise estimate of the spacecraft range (via the round-trip travel time of light).

Doppler processing is done by the MDA to provide project navigation and network monitor personnel with a precise measurement of the Doppler shift at any given time. Counts timed to the thousandth of a cycle are provided for the S-band and X-band receivers as well as for the exciter reference frequency. Individual counts are taken with a granularity of 10 times per second. The MDA computes Doppler residuals from values provided by the NSS in a separate support-data file. Range, Doppler, and residual data are sent to the spacecraft navigation team.

In the future, the received signals will be digitized in the Block V receiver, and Doppler data will primarily contain the difference between the received frequency and a known reference frequency. The MDA will format the data and include the uplink frequency data and other ancillary data for transmission to the flight projects.

c. Command. The DSCC Command Subsystem (DCD) has as its primary function the transmission of commands to the spacecraft. As part of this function it receives spacecraft commands from JPL, developed by Project Operations Centers (POC's), processes those files, and transmits individual elements from selected files. The commands are received, verified, acknowledged, stored and subsequently transmitted to the spacecraft from the Command Processor Assembly (CPA) at each DSCC.

Actual spacecraft commanding is done by one of several methods of waveform modulation imposed upon the subcarrier by the Command Modulator Assembly (CMA), which has a maximum output of 2 kbps. The CPA, under direction from the POC, sends binary command elements to the CMA where they are modulated onto the subcarrier and transferred to the transmitter exciter. Commands are verified by a return line from the exciter to the CMA. Non-favorable comparisons result in a cessation of the command process and a notification sent to the POC.

d. Radio Science. Radio science is the analysis of scientific information extracted from the variation in radio signals coming from a spacecraft. Depending on the information being deduced, the variation of interest might be the Doppler shift, phase, polarization, or the frequency spectrum.

The received signal can have variations from three causes:

- (1) the motion of the spacecraft.
- (2) the influence of the media through which the signal passes.
- (3) deficiencies in the hardware used to capture the signal data.

Media that can affect the signal include the atmosphere, planetary rings, and solar corona. In order to detect the sought-for effects, instability in the ground system frequency standard must be minimized via ultrastable hardware and by the removal of identifiable noise during signal data processing.

Radio science signal processing performs the following actions:

- (1) Accepts commands from, and provides status to, the operator.
- (2) Obtains signal variation predictions from JPL.
- (3) Provides tuning information to the receiver.
- (4) Converts the analog signal to digital data.

- (5) Records signal data and ancillary data to storage.
- (6) Reconstitutes the analog signal for real-time monitoring.
- (7) Transmits partial signal data to JPL for validation.

e. VLBI. This technique uses extragalactic radio sources to determine UT1, Earth polar motion, the relative positions of stations on the Earth, clock synchronization among the stations, and clock stability. VLBI is also used to maintain and enhance the JPL catalog of extragalactic radio sources. Delta-VLBI alternates between extragalactic radio sources and spacecraft radio sources to accurately determine the position of the spacecraft.

VLBI processing measures the time difference between the arrival of a radio signal at two or more different receiving stations on the Earth (Fig. 4). The signal source may be either a spacecraft or an extragalactic source. The time difference is determined by correlating the two received signals.

Narrow channel bandwidth (NCB) capability is normally utilized. However, for cataloging extragalactic radio sources, wide channel bandwidth (WCB) capability is utilized. The radio signals are transmitted by the spacecraft and by extragalactic sources; received by two or more physically distant DSN antennas; tuned and translated for several simultaneous frequencies by the receivers; digitized and stored by the signal processors (except for WCB data); and transmitted to JPL for correlation. The VLBI signal processors perform the following functions:

- (1) Accept commands from, and provide status to, the operator.
- (2) Provide tuning and control commands to both receivers.
- (3) Convert the NCB analog signal to digital data.
- (4) Record the NCB signal and store ancillary data on a disk.
- (5) Store WCB ancillary data on tape.
- (6) Transmit stored NCB data to JPL for correlation.
- (7) Monitor the coherency of the signal and inserted tones.

2. Monitor and Control. The DSN Monitor and Control subsystem consists of JPL (Central Site) and DSCC components. A summary of the major functions at each site follows.

- (1) JPL. The Network Operations Control Center (NOCC) is the present central-site focus for monitoring and is the primary interface between the DSN

and the POC's. NOCC negotiates tracking schedules and generates sequences of events (SOE's), angle and performance predictions, and standards and limits. These support data are stored in databases and periodically distributed to the DSCC's. During spacecraft tracking, NOCC monitors the status, configuration, and performance of DSN equipment and subsystems. This function includes acquiring real-time data from each DSCC and monitoring all DSN missions. Monitoring is aided by real-time displays on workstations with alarm, event, and advisory messages, and access to historical logs.

- (2) DSCC Site. Prior to an acquisition, equipment is configured to provide a logical processing sequence. During an acquisition, operations personnel issue directives through a workstation to control the subsystems. The functions are divided into two hierarchical subsystems: the Complex Monitor and Control (CMC) and the Link Monitor and Control (LMC).

The CMC subsystem is responsible for unassigned equipment pool monitoring, link assignment management, support data management and distribution to subsystems, multilink equipment monitoring, configuration files maintenance, and equipment calibration.

Whereas the CMC covers all links at a Complex, the LMC subsystem is concerned with only one tracking antenna and its associated ground data system elements. The main LMC functions are subsystem control by operator directives, health and status monitoring of subsystems, receipt and presentation of subsystem configuration displays, and monitor data distribution to POC's (for missions launched prior to 1991).

B. System Interfaces

In the following sections, the input and output interfaces to the GIS are identified and the data transfer rates are estimated in order to provide a basis for performance trade-offs in Section IV.A.

1. Input Interfaces. At the Central Site, inputs to the GIS include mission schedules and sequences of events from the POC's to the NOCC and spacecraft commands from the POC mission-support teams.

At the DSCC's, the inputs to the GIS include the spacecraft signals from the Block V receiver and monitor and control data from many devices and subassemblies in the front-end area.

a. Central Site

Planning. Mission schedules and sequences of events from the POC's are communicated to the NOCC. The data rates are generally low, less than 10 kbps.

Commands. Spacecraft commands are generated by the flight project planning and sequence team and transmitted to the DSCC Command subsystem through a central-site interface. These data rates are also usually less than 10 kbps.

b. DSCC Site

Antennas. There are currently four primary antennas at each complex: one 70-m antenna, two 34-m antennas, and one 26-m antenna. A 10-m antenna is currently being added. A new antenna every three years is forecasted, for a total of about seven to ten antennas at each complex by the years 2000–2010.

Block V Receiver. The Block V receiver is a planned DSN upgrade that will eventually be used in all SPC's. Each SPC will require one Block V receiver with sufficient channel processors to accommodate the subcarrier signals from each spacecraft. These signals may be transmitted on various bands (i.e., S-, X-, and possibly Ka-band in the proposed time frame) and different polarizations. During a given acquisition, one channel processor must be logically connected to one telemetry processor. In this study, approximately 22 channel processors and telemetry processors are estimated for the year 2000 and 32 in 2010.

A functional block diagram that also includes the telemetry subsystem is shown in Fig. 5. Each channel processor of the Block V receiver has a planned output of 26.4 Msps (megasymbols per second, with 8-bit symbols) for current X- and S-band missions; however, advancing to Ka-band communications will enable a 165-Msps planned output rate.

As discussed earlier in Section II.A, redundant bits are usually added to spacecraft telemetry for forward error correction. Viterbi code may be used for high-Earth orbiter and lunar communications at a 1/2 rate (2 bits for every information bit) and deep space missions at a 1/6 rate (6 bits for every information bit). RS code may also be used to spread the impact of burst noise on the space link; a typical overhead is about 14 percent. Figure 6 illustrates the effect that this coding has on the receiver-to-telemetry communication rate. At the Block V planned output rate, which handles a spacecraft-information rate between 4 and 11 Mbps—depending on the coding—the transfer rate to the telemetry processor is 26.4 symbols per second, or about 211 Mbps.

Each Block V receiver also provides Doppler and ranging data at less than 1 kbps, monitor data at less than 50 kbps (including spectral data), and radio science data at less than 1.2 Mbps (100 kbps with either 8 or 12 bits per sample).

2. Output Interfaces.

a. Central Site. The outputs of the GIS are usually transmitted to flight project teams and other users, or stored temporarily on portable media for off-line delivery to users (see Section II.D). Some data, such as telemetry data, may be transferred to archival systems, such as the Planetary Data System, for later distribution on demand. Outputs usually emanate from JPL; however, provisions are made for direct delivery from the DSCC to non-NASA POC's.

Telemetry. Most present deep space missions have low data rates in the 10- to 100-kbps range. High-Earth orbiters have data rates ranging from 100 kbps to several hundred Mbps. Future missions, such as the Space Exploration Initiative (SEI), will range from 10 Mbps for Mars support to 100 Mbps for lunar missions.

Tracking. The tracking subsystem provides metric data in packet format to the flight projects navigation team. These formatted packets are relatively low-rate, that is, less than 10 kbps.

Radio Science. Radio science data are produced by the Block V receiver. It is forecasted that two or three concurrent receivers may each be producing up to 1.2 Mbps in the forecasted time frame.

VLBI. VLBI is an emerging technique for computing spacecraft range. Considerable celestial radio-source mapping is required to ensure its success. While the ranging function is performed only occasionally, the mapping is done on a daily basis. The VLBI subsystem typically acquires 20 to 30 minutes of high-speed data into a file (e.g., 130 Mbps in the NCB subsystem) and plays it back slowly for processing. In an effort to improve reliability of the acquisition, it is envisioned that the proposed VLBI subsystem will acquire data in a quick-look mode for 10–20 seconds, process it, and transmit real-time corrections back to the station.

b. DSCC Site At the DSCC's, the outputs of the GIS include monitor and control data to many devices and sub-assemblies in the front-end area.

3. Availability. The total loss of telemetry data in the DSN from 1986 through 1990 amounted to 2.9 percent,

or 4814 hours, of the scheduled support time, according to a recent study. (See Table 2; the items listed in that table account for 98 percent of the total lost hours.) Steady improvement is evident from the yearly data: 3.9, 3.5, 3.5, 2.4, and 2.2 percent for these five years (decreasing an average of 12 percent per year). These percentages are derived from planned hours of operation that have been lost and do not include such operational delays as calibration and scheduled maintenance.

More interesting are the causes of these lost hours as attributed to the subsystems and assemblies. The 4814 hours of loss were attributed to 22 subsystems located from the front-end areas to the NOCC, plus radio-frequency interference and undetermined origins.

If a major effort is made in the future to improve the availability of these data subsystems, control interfaces and control assemblies, and the way the front-end subsystems are controlled, it is conceivable that a 40- to 50-percent reduction of lost telemetry hours can be achieved. As compared with a 2.2-percent telemetry hour loss in 1990, a reasonable goal can be set at 1-percent loss, or 99-percent availability, which is consistent with the 12-percent annual decrease already noted. Another major improvement in overall availability can be made by reducing operational delays in precalibration, postcalibration, and scheduled maintenance.

C. Future Mission Sets

Table 3 summarizes the new DSN missions set and includes deep space and high-Earth orbiter (HEO) missions. (The Earth Orbiting System [EOS] is not shown.) Approval is expected for three to five of these missions. International cooperation missions may be supported. Also, the DSN may be requested to serve as a contingency network for near-Earth missions, including Space Station Freedom and shuttle missions. Support to the following missions are not mandatory requirements but were considered in evaluating the architectures:

- (1) SEI (Mars), 10 Mbps.
- (2) SEI (Lunar), 100 Mbps.
- (3) Tracking and Data Relay Satellite System (TDRSS) backup (e.g., EOS), up to 300 Mbps.
- (4) Advanced Tracking and Data Relay Satellite System (ATDRSS) backup, up to 650 Mbps.
- (5) International cross support, undefined.

Table 4 summarizes three views of the future data requirements. The architectures developed later in Section IV.A will use the moderate forecast for most of the

evaluation, and one can consider the aggressive forecast as a possible extension.

D. End-User Interfaces

End-user interfaces fall into the following four categories: mission-specific teams, principal investigators, data archives, and cross-support that the DSN provides to other agencies.

Figure 7 shows the current end-user interfaces between the GIS (including the DSN and the SFOC) and typical flight projects. Also shown is the type of information flowing to and from the end users.

With continued development of the SFOC capability under the Multimission Operations Systems Office (MOSO), the relationship between the end users and the GIS will probably change significantly. These changes may cause related interfaces to change. The driving force for the changes is to reduce mission operations and data analysis costs. The approach is to identify and replace mission functions that could be performed by multimission tools and teams. One of those changes, which can already be seen in the Mars Observer and CRAF/Cassini plans, is the use of the remote Science Planning Operations Computer (SPOC). The SPOC acts as a remote extension of the SFOC by providing command, telemetry, and navigation data access, and planning capabilities to the project scientists. SFOC-based multimission operations teams and remote extensions of the SFOC will continue to evolve.

The interfaces for each class of end user are shown in Table 5. The table is organized to show the direction of data flow (to and from the GIS), the type of interface, and a brief description of the content or purpose of the interface. Note that for some of the interfaces the interface medium is magnetic tape. The GIS in years 2000–2010 will likely use other physical storage media.

1. Mission-Specific Teams. JPL planetary spacecraft have a great deal of similarity in their mission operations design. The particular names given to the various teams may sometimes be different, but when the functions are identified, the teams that most spacecraft projects agree upon are: Mission Control Team, Planning and Sequencing Team, Navigation Team, Spacecraft Team, and Radio Science Team.

2. Principal Investigators. The term principal investigator is applied to the cognizant individual who leads a science experiment. The team may include the principal investigator, coinvestigators, and interdisciplinary scientists. The investigators of the future, especially for the

exploration missions, may do most of their work for the project at their home institution. From there, they will submit observation requests, receive virtual channel data from their instruments as well as ancillary data they need to analyze their data, and return their reduced data to the project for archiving. They may also receive data on tape (or other media) if an instrument's data rate is too high, or too expensive to transport.

3. Data Archives. The primary data archive of concern to this architecture study is the Planetary Data System (PDS). This archive is a distributed system that has a central computer which houses the master catalog and the archive operations functions. Data are archived and maintained at the distributed nodes that are defined by various planetary science disciplines (e.g., planetary geology, fields and particles, etc.). The interface with the GIS is still under development; for example, for the Magellan project, the interface has been negotiated as a CD-ROM interface in which the project produces the compact disks and delivers them to the PDS. The future interface will probably be in some other electronic form.

4. Cross-Support Services. The GIS provides several forms of cross-support services, for example:

- (1) Backup support for the TDRSS (i.e., for shuttle, Space Station Freedom [SSF], and EOS).
- (2) Reimbursable support to international space agencies (initial acquisitions, tracking support, etc., for National Space Development Agency [Japan], Institute of Space and Astronautical Science [Japan], European Space Agency, and others).
- (3) Arraying of antennas, such as was done for the Voyager encounters at Uranus and Neptune.

For backup and support to international agencies, the data exchange is through the use of CCSDS data packets.

III. Technology Forecast and Projected Impact

This section forecasts new and emerging technologies in the areas of computer hardware, software, data communications, and human interfaces that are of potential benefit to the DSN GIS in the 2000–2010 era.

A. A Look at the Future

The following perspectives from a research scientist, an operator, and a software developer are offered to stimulate

creative thinking and are not constraints on the architectural study.

1. A Research Scientist's Perspective.

"I had a personal computer on my desk 10 years ago, so that much hasn't changed. The personal computer that is on my desk now has dramatically higher performance, a larger screen, and larger storage capacity than before; however, the biggest change has come in the way I use it. For me, and for most of the people I deal with, the personal computer has replaced the telephone in all sorts of ways. I have the regular personal computer applications I use every day—the spreadsheets, a word processor, and a couple of small databases—but what really makes the system useful for me is that it is linked into the local network I share with my staff, into national and international computer networks, and into the flight project office, where I can access the really large databases we depend on for up-to-date scientific analysis and monitor the latest input from my planetary instrument.

"Today I often have no idea where the data I'm using is actually stored. I don't really need to know. Sometimes I'll pull summary information from the Planetary Data System, combine it with up-to-the-minute status reports on a current mission, pass the whole package to my science analysis team for some 'what if ...' projections—all without having to think about the mechanics of how things really link together. For all I know the data comes by way of Siberia. When I share the data with my colleagues we can discuss it as we analyze it on our screens and propose and conduct analysis during 'electronic meetings.' We hear everyone's voice and see team members and review the data—our attention is focused on the screen.

"Perhaps the most pervasive difference today is that we are much less dependent upon technicians than we were in the past. Most of our analysis is produced by our professional staff—the knowledge workers—with the assistance of sophisticated software tools. Compared to just a few years ago, there is much less manual analysis to be done, because of our reliance on 'electronic assistants.' We often annotate results with voice and video clips and share it with colleagues for their review and comments.

"I have a lot more time today to focus on doing the 'right things,' because we have learned to use the systems to avoid spending time doing the 'wrong things'—things that don't contribute to the accomplishment of our research. We make fewer errors because we are better able to monitor activities, to do reasonableness checks. We spend less time looking for related data, particularly pieces of paper, because more of our files are stored in electronic form,

and we can do extremely rapid electronic searches, comparisons, and associations. The problems associated with performing analysis in stages have been minimized because we have replaced much of the paperwork with electronic transactions. There just isn't that much 'float' in the process anymore.

"If you look around this office, you can see that science still uses paper, but in many senses it is a different kind of paper: It's what you might call secondary paper—reports, studies, and day-to-day reports—the 'real-time' information is almost entirely on the screen. And as I've said, that means a lot less looking and waiting. I just call up. It's there when I need it, it's in a format that I can use, and it's reliable. That's most important, of course—that I can trust the information to be correct and current."

2. An Operator's Perspective.

"The most significant change in my job is that all the equipment is remotely operated from here in Pasadena, and my job involves monitoring a highly automated process. If any subsystem needs to be adjusted, I can do it from my workstation. Every action has explicit, reliable feedback. I can also query the subsystems and get immediate responses. The subsystems themselves are smarter: I can now control them with much higher level instructions; they are self-monitoring and self-diagnosing and I can place one subsystem under the direct control of another. The biggest difference I see in operations is that overall operations seem more integrated, and the networks and tools available to me through my console have made the system more responsive to what I need to do to control it. I make fewer errors because I am better able to monitor activities, to do reasonableness checks.

"The DSN supports more missions now than 10 years ago, and I am able to maintain several spacecraft-to-ground links through multiple antennas at the Complex. JPL now arrays antennas routinely, and is progressing towards using one antenna to communicate with several spacecraft.

"I use automated monitoring tools to keep a constant check on the operational status. Graphics and animation on my console present information to me in such a way that I can effectively monitor several links at once. When things go wrong, fault detection, isolation, and diagnosis tools help me to narrow down the possible problems, suggest alternative solutions and quickly get a message to the maintenance personnel at the Complexes. Expert knowledge is available from my console so I can benefit from having the 'experts' available to me whenever I need them.

Maintenance people at each Complex have similar tools to help isolate and repair equipment.

"The procedures and manuals are now all available on the screen. Instead of having to sort through volumes of information to pull together the bits and pieces I need, I can now navigate through a library that has not only the documentation, but photographs, animated sequences, and voice notes all linked together. Getting the 'big picture' no longer involves a treasure hunt through documentation libraries and coworkers' notebooks. When I discover a problem in the documentation or procedures, I can instantaneously annotate the procedure, and even recommend changes.

"Embedded training programs enable me to practice the performance of difficult operations in an environment exactly like the operational one. Instead of sitting idly by when I'm in the middle of a long track, I can now place my console in a training mode. I'll run scenarios and practice procedures while background processes monitor the health and status of the space link. If any problem occurs, the system immediately notifies me and provides me with the necessary background information.

"Because of the training system, I am now proficient at performing operations tasks I would rarely work on a decade ago. I am better informed, understand the activities on a much deeper level, and am better equipped to deal with emergencies than before. I'm also more aware of what's important to our customers and can work with them to deliver the best possible data product.

"I now provide better problem reports to system developers with a quick and easy way to document problems: if a failure occurs, I dump all the appropriate information into a file and the file provides a history of events surrounding the failure and a clear picture of the health, status, and configuration of the station when it surfaced. This way, the developers and maintainers have first-hand information to work with—and problems can be evaluated to make sure they don't happen again."

3. A Software Developer's Perspective.

"International standards apply to the space and ground communications portion of the DSN, and this eliminates my need to create new communications software, and helps me focus more on what the user needs to see and control. International standards have made it easier for my division to assemble a substantial library of space-application modules, many developed by foreign programmers.

"Most of my software engineering efforts are focused on building maintainable code, and interfacing it to tested

and proven software. One major change is that there is greater emphasis on designing fault tolerance into the system. One way we do this is by integrating common monitor and control modules into the software architecture that enable remote monitoring of hardware and software performance.

“There is a mix of computer platforms in the Ground Information System; all use identical operating systems. Those that are roughly equal in capability can be dynamically assigned tasks. This makes the information system more dependable by making it tolerant to isolated computer failures. Network, computer, and software architectures are carefully designed to be interoperable and fault-tolerant. They must be—the DSN supports many more real-time users than it did ten years ago. These users are expecting portions of their data within seconds of capture and want to send real-time signals to their space instruments.

“The human interfaces that I design involve text, sound, and graphics—and sometimes include voice and video. The format of these data types has been standardized, and many human-interface application modules are available—either commercially or from software libraries. I work closely with human factors engineers and users to provide all the features they need to do their job at one location—at their workstation. The computer has also become the focus of communications between operators and engineers like myself, which simplifies their access to me.

“I obtain the majority of my application modules from software libraries. My development team tends to use object-oriented software modules; the porting of modules to other machines is easier and faster. One measure of my success as a programmer is how many people want to reuse my code for their applications; the usefulness of my code to other programmers is very important to me.”

B. Computer Hardware

The 1990s will see extremely high-speed electronic systems that will be increasingly compact, mobile, and lightweight. High-speed computer systems will include multiple high-speed microprocessors with increased parallelism; memories with ever-increasing densities; and specialized high-performance chips, such as video chips with both analog and digital memory interfaces and voice-processor chips with adaptive recognition capability.

The growing power of functions at the chip level enables greater degrees of parallelism. It will become more feasible to dedicate a processor of the proper kind to each element

of a system in subsumptive architectures with hierarchic arrangements of functions in increasingly detailed layers of computation, and systolic data-flow architectures with patterns of specialized processors arranged in data-flow orientation.

Initially, the DSN computers were JPL-built controllers, which gave way to minicomputers of several kinds, tending toward microcomputers today. Recently there has been some work to seek very-large-scale-integration solutions for some of the functions of the DSN.

1. Computers. The DSN will probably require microcomputer, minicomputer, and mainframe levels of capability in its future architecture. Mainframes, if used, will probably be located at the Central Site, with minicomputers and microcomputers distributed at the Central Site and DSCC's.

a. Microprocessors. Since the late 1970s, there has been a steady increase in the capability of microcomputer chips in execution speed, addressing capability, input and output (I/O) bandwidth, and instruction sets. This trend is expected to continue, with processing speed increasing by a factor of 10 or so every five years.

Compatibility between successive members of a manufacturer's line is reasonably good, at least at the level of users of commercial software packages and operating systems, although significant rework has often been necessary on the part of software suppliers to support a new version of a workstation microcomputer chip. The trend to graphical user interfaces (GUI's) has helped to disguise the extensive modification required to support new workstations.

Mainframes are emerging with higher speeds and larger storage capability, which solidifies the mainframe as the machine of choice for large numbers of users running database-oriented applications (e.g., in business). For scientific users running computation-oriented applications, the appearance of the new mainframes provides additional powerful analytic tools and the potential for new discoveries.

Today's microcomputers operate up to over 75 million instructions per second (MIPS), 22 million floating point operations per second (MFLOPS), with over 800 Mbytes of internal mass storage. External devices may extend the storage to over 10 Gbytes. So-called optical juke boxes extend the mass storage to over 300 Gbytes.

The points of demarcation among the performances of microcomputers, minicomputers, and mainframes, and be-

tween workstations and systems are not firm. NCube, Inc., for example, has introduced a computer capable of speeds up to 1 billion floating point operations per second (GFLOPS) using 512 processors, with up to 1 Gbyte of storage, all in one cubic foot of space and weighing about 60 pounds. The computer is housed in a portable chassis that fits into a personal computer carrying case.

Computer chip manufacturers predict that 1- to 2-GFLOPS capability will exist at the board level within ten years. The Intel Micro 2000, now in design, will contain 50-100 million transistors in four central processing units, two vector processing units, a human interface unit of 4 million transistors to provide high-resolution, full-motion graphics, and a 2-Mbyte cache to reduce off-chip reaches for data. In addition, this chip is planned to be upwardly compatible with the Intel 386 instruction set and would run all software that currently runs on that processor.

b. Parallel Processors. Parallel processors fall into two main categories: (1) single-instruction multiple-data (SIMD) machines, in which many processors work with the same program on vectors or arrays of data values, and (2) multiple-instruction multiple-data (MIMD) machines, in which each machine has its own program and the (more or less) independent machines communicate by means of I/O interconnections. The SIMD machines are particularly well-suited to calculation of multiple-degrees-of-freedom problems, such as orbit determination, while the MIMD machines are well-suited to simulation of complex events.

All these machines currently require extensive reworking of existing processing algorithms to take advantage of parallel processing; unfortunately, none of the existing implementations promises or even suggests compatibility with later versions of parallel machines.

2. Digital Signal Processors. Digital signal processing can be advantageously applied to several functions of the DSN. These include receiving, transmission, switching, encoding and decoding for error correction, security, storage, and retrieval.

Very-high-speed integrated circuits (VHSIC's) are expected to do complex signal processing dramatically faster with smaller chips. This development will enable advanced forms of modulation, synchronization, and coding. The realization of the Viterbi algorithm on a chip is an example of the application of this emerging technology.

Signal processors may operate as circuit boards managed by external computers. Current capabilities of such

board-level devices include 150-MFLOPS 32-bit vector processing, with 400-Mbps access to memory. The ability to do fast floating-point operations permits a very flexible signal-processing design. The speed of commercial digital signal processing boards is expected to keep pace with computing speeds, growing at least by a factor of ten over the next decade.

3. Storage. Storage technology is rapidly increasing the capability to rapidly and reliably store and retrieve large volumes of data. Three types of devices are all current areas of research: solid state, magnetic, and optical.

In 1990, the semiconductor industry announced the first 8-Mbyte dynamic random-access memory on a single 10- × 20-mm chip containing 40 million electronic components. At present, 32-Mbyte and even 128-Mbyte solid-state storage chips are in development. Access times, meanwhile, are falling steadily; whereas about 80 nsec is common now, about 60 nsec will be the norm within one year in standard complementary metal oxide semiconductor (CMOS) technology, and 35 nsec may be achieved using bipolar CMOS within another year. JPL is presently developing a 4-Gbyte, 320-Mbps vertical block line device, with an expected latency under 100 msec. By the years 2000-2010, it is not unreasonable to project solid-state devices reaching 1 Gbyte on a single chip, with 20-nsec access.

Magnetic device performance will also continue to rise dramatically. Magnetic tape is expected to reach 1-10 terabytes per reel, with a transfer rate ranging from 300 Mbps to 1 Gbps by the 2000-2010 time frame. Magnetic disks are expected to reach 0.15 Gbyte/cm² by then, with only about 1-μsec latency. JPL is currently conducting research in the application of tunneling microscopy to magnetic memory, which could reach densities of about 1 terabyte/cm² by the year 2020.

A less dramatic rise is being predicted for optical storage technology. Currently, digital optical tape systems are capable of 50 Gbytes per cartridge and a 24-Mbps sustained data-transfer rate, with seek times of less than 4 sec per Gbyte. The optical diffraction limit will permit only up to about 1 terabyte per cartridge, and only about a 30-Mbps readout is expected. Similarly, optical disks are projected to limit at about 45 Mbytes/cm² with 1-μsec latency. Because of these limitations and the projected 3:1 superior performance of magnetic technology, it has been questioned whether there will be a market for high-performance optical storage technology.

4. Hardware Integration. The Department of Defense (DoD) is developing multichip-module fabrication

and packaging technology that drastically decreases the volume, power, heat, and cost of high-performance computational systems, while increasing their speed, reliability, and effectiveness. The concept of three-dimensional (3-D) packaging is not new, but only recently have vendors made significant progress in implementing the process. The DoD has identified multichip-module technology as an enabling technology and plans to establish multiple manufacturing sources for each type. Figure 8 illustrates the concept.

A number of techniques are being explored. A Hughes Aircraft 3-D computer uses 4-in. diameter wafers. Other systems utilize bare chips (the dies themselves), some of which are trimmed to remove line drivers. Line drivers are rarely needed since the distances are reduced and no pinning is involved. The systems are considerably cooler, since as much as 85 percent of the power budget and heat budget for a chip is used by line drivers.

The multichip modules or wafers are interconnected by very-high-speed links. For example, a Hughes module has an internal rate of 160 Gbps, and a Texas Instruments network operates at 20 Gbps on the module and at 800 Mbps between modules. The various packaging techniques used by vendors often utilize off-the-shelf commercial chips, such as static random-access memory, dynamic random-access memory, central processors, communications processors, and digital signal processors. Repair of these systems will usually be done by replacing a module or wafer. Since there are typically only three to five types of modules, the parts problem could be minimal.

The systems currently being developed are for weapons control. They are programmable (sometimes in multiple languages) and reconfigurable. Other potential applications presently under consideration include avionics control, general-purpose supercomputer workstations, and sensor analysis systems.

A system being developed by the Environmental Research Institute of Michigan has performance goals of 30 MFLOPS and 15 MIPS. It will fit in 200 cubic inches and is expected to cost about \$85,000 in volume quantities. The same capability in 10 years is projected to cost less than \$10,000 (in current year dollars).

5. Reliability. The GIS has a functional availability goal of over 99 percent in the envisioned time period. This may be approached by a fault-tolerant design with appropriate fault detection, isolation, and recovery technology. One key need is for highly reliable components in each subsystem. This section summarizes a forecast of computer-oriented reliability issues.

a. Current Architectures. In 1975, the mean time between failures (MTBF) for disk drives was only about 3000 hours. Today, the MTBF for the newest physically small units is quoted in the 200,000-hour range. Much of the increased reliability has been attributed to the reduction in size. The 1991 disks are 3.5-in. units that hold as much as ten times more data and have transfer rates five times faster than the 1975 units.

Failure rates for tape units have not improved as dramatically, but helical scan devices have jumped in reliability from 28,000 hours MTBF to 40,000 hours MTBF within the last year, validated by field experience.

The next most frequent failures encountered in computer systems occur in the cabling and connection equipment and on the boards themselves. Flexing and stress, due to temperature and mechanical variations, create socket and tracing breaks. Years of age on connective and board components becomes a major contributing factor. The incidence of memory and processor failures, after initial so-called infant mortality, has been so rare that industry figures are not usually published.

b. Advanced Architectures. The environmental conditions that DoD's advanced integrated architectures (such as wafer and composite module devices) are intended to survive exceed those found in the conventional computer room: mechanical shock tests up to 168,000 g's (required for impact, and rail guns), thermal cycling from liquid nitrogen to boiling oil temperatures, vacuum tests, and radiation tests.

The failures that are experienced are inside the chips and dies themselves. Chip failures occur due to operating temperatures and radiation, usually not to mechanical stress in these packages. The design process can provide for cooling by convection or conduction flow where necessary.

Most of these dense packaging designs have built-in redundancy. Some have 50-percent spare components; others have up to 200-percent redundancy in connectivity (much of which can be switched dynamically); and still others have fully active switching to bypass components that are failing. These design features further enhance the reliability of the unit.

It appears that new, dense packaging architectures present a number of advantages over present traditional designs because the number of connections is significantly reduced. Shock failures due to physical or thermal events are virtually eliminated. All these attributes contribute

to significant increases in the reliability of the solid-state units and systems.

C. Software

The 1980s witnessed the advent of major efforts to standardize many aspects of computing, including programming languages, operating system interfaces, and network protocols. By the year 2000, most of these efforts will have resulted in wide adoption of international standards. The standards process will provide a common basis on which to build distributed computer systems.

1. Open Software and Standards. A major concept is that of open systems. Although not rigorously defined, in this context the term open refers to the willingness of a computer or software vendor to offer products compatible with those available from its competitors, usually because they mutually agree on standards. Open systems benefit users because they yield greater independence in the procurement of hardware and software.

Planning for an open system makes it possible to implement highly portable, hardware-independent, scalable computing systems through the use of existing software technology, an awareness of trends, and by adhering to a plan for compatibility.

The following paragraphs summarize various developments in the evolution of open systems that will have important implications for the DSN GIS. Large-scale computer applications often rely on capabilities beyond those specified by a programming language definition. For example, file system services, process creation, and network I/O are support functions that require additional specification. The universe defined by this broader specification is called an application environment.

Several related (and mutually consistent) efforts to set a standard application environment are under way. The following discussion is based on the Open Software Foundation (OSF)/Application Environment Specification (AES). AES is itself based on various other standards (both adopted and emerging) as noted. The discussion is intended to indicate current directions in the evolution of open systems.

a. Programming Languages. At the present time, the C programming language is the language of choice for systems software. C is a middle-level language originally developed to replace assembly code for the most demanding systems' programming tasks. C language deals with the same sort of objects that most computers do, namely,

characters, numbers, and addresses. It may be used in the GIS for applications that strongly manipulate operating system resources and the computer platform, such as operating systems, compilers, editors, and low-level I/O software.

C++, a superset of C, is a general-purpose language. C++ was designed for larger, more structured applications than C, and provides flexible and efficient facilities for defining new data types. C++ hides much of the application functionality behind object interfaces, hence it is referred to as an object-oriented language. It will be applicable to the GIS by providing a tool for the rapid development of highly portable, structured application programs. It is applicable to human-interface displays and database management systems.

Common LISP has proved to be a useful and stable platform for rapid prototyping and systems development in artificial intelligence and other areas. Artificial intelligence applications written in Common LISP are also successfully integrated with software written in other languages, including C, C++, Fortran, and Ada. The unique capabilities of Common LISP, which include memory and data structure management, flexible typing, symbolic representation, processing capabilities, and functional programming style, ensure that the language will remain the dominant language for most AI research and application software for the foreseeable future.

Extension languages support a new style of programming that has emerged in recent years called language-based programming. With this style, a programmer presented with requirements develops a special-purpose language in which a program that satisfies the requirements can be expressed succinctly and clearly. The ultimate program is written in this higher level, special-purpose language. This style allows easy experimentation, testing, and evolution of the application much more readily than do more traditional top-down, or structured, programming styles.

In the future, more and more programs will be designed and implemented as multilayer language systems, with the upper layers calling languages implemented in lower layers. High levels of abstraction will be provided by upper layers, allowing applications to be programmed in very few lines of code. Portability will be routinely achieved by defining one of the interlayer boundaries as a portability boundary, with machine-dependent code allowed only below the boundary. Any code above the boundary will port from one machine to another without change. Table 6 summarizes the languages that may be considered as open languages.

b. Operating Systems. The 1980s saw several important trends in scientific and engineering computing. The first and most important was the emergence of UNIX as the dominant operating system. The wide availability of UNIX also encouraged the development of highly distributed computing systems based on the client-server model.

The Institute of Electrical and Electronics Engineers (IEEE) has recently adopted Standard 1003.1-1990, the Portable Operating System Interface (POSIX.1). This document defines a C-language interface to an operating system (i.e., UNIX). POSIX.1 is also an International Organization for Standardization (ISO) standard (ISO 9945-1: 1990). The two major players in the open operating system field, UNIX International and the Open Software Foundation, have pledged compliance with POSIX.1 for their respective operating system products.

c. User Interfaces. The X Window System from Massachusetts Institute of Technology Project Athena has established itself as the dominant technology for windowing systems under UNIX. Standards for X Windows are developed under American National Standards Institute (ANSI) committee X3H3.

At a higher level, the most clearly drawn battle line in open systems is perhaps between the OSF/Motif and Sun OpenLook graphical user interfaces. At the present time Motif is proprietary software; OSF publishes the standard and the software itself must be licensed from OSF individually for each workstation. Efforts are under way to build a virtual interface supporting either technology.

d. Network Services. AES supports the Internet suite, plus other selected Internet services developed at the University of California, Berkeley, for UNIX networking. AES also includes support for selected Open Systems Interconnection (OSI) protocols.

From the standpoint of software architecture, the choice of network protocol is often incidental, since it is usually hidden by higher layers in the distributed computing environment (see Section III.C).

2. Distributed Computing. The concurrent explosive growth in the power of inexpensive computers and availability of network technology has led to widespread adoption of distributed systems in which large numbers of loosely coupled computers communicate over local area networks. The problems of designing software systems that can take advantage of diffuse computing power are many and derive mainly from the additional complexity of simply having more semi-independent components.

Any monolithic computer system can be distributed by the ad hoc insertion of system boundaries and corresponding communications support. A properly designed distributed application, by contrast, will be built upon a framework that provides a portable programming interface and subsumes much of the complexity that accompanies distributed computing.

OSF's Distributed Computing Environment (DCE) technology includes support for building distributed applications in the UNIX environment. The following discussion highlights the key capabilities of DCE, again to demonstrate the current state of open systems evolution.

a. Remote Procedure Call. Client-server computing, in which a process requests actions on the part of another process, is a useful abstraction for exploiting the power of distributed computing systems. With the necessary communications support, processes on different computers can be made to cooperate to bring greater processing power to a single task.

A remote procedure call (RPC) is a logical extension of the local procedure call model used by many programming languages. RPC technology allows the programmer to specify interactions among components of a distributed system through a convenient and consistent method. In addition to the obvious benefit of making distributed applications easier to write, RPC also insulates those applications from the details of specific processor architectures and network implementations. It therefore contributes to the portability of such applications. A well-written RPC application, for example, can operate among several vendors' systems, and use different network protocols (e.g., Internet or OSI) simultaneously without code changes.

b. Naming Service. The added flexibility that can be obtained from distributed computing leads immediately to a problem of information management. A client process, for example, needs to know how to contact a server process on another computer. If the binding of service to computer is deferred for as long as possible, reconfiguration can be accomplished, when necessary, with minimal side effects.

A naming service is a mechanism for storing and distributing such information throughout a network. Well-written applications should exploit the naming service for any information that may be subject to change. Source code is thereby insulated from such changes, and greater run-time configuration flexibility is afforded.

c. Threads. Figure 9 shows a process requesting data from an upstream server and responding to data requests

from a downstream client. The reading and writing activities are not directly coupled; a cache buffer is inserted so that a temporary decrease in output data rate does not cause a corresponding slowdown in input flow and processing. In addition to reading and writing data, the process also communicates with a monitor and control process. It has therefore, several semi-independent flows of control.

The general problem of multiplexing loosely synchronized activities has been addressed through a variety of mechanisms. Real-time systems typically provide a rich set of tools (scheduling primitives, interrupt service routines, etc.) but generally at the cost of increased complexity and reduced portability. Multitasking systems (e.g., UNIX) provide mechanisms based on interprocess communication, but these are often expensive in performance overhead.

Perhaps the most useful conceptual model for this problem is the thread. A thread is simply the flow of control within a process. Because all threads in a process exist in a single address space, they can communicate efficiently through shared memory. Synchronization among threads is accomplished via scheduling primitives provided by the threads implementation.

The threads model simplifies the design of multiplexed processes by decomposing them into a collection of threads whose temporal relationships are explicitly specified. Considering the example in Fig. 9, the decompositions might be as follows:

- (1) Read. Wait for server input. Read and process data. Wait for room in output cache. Queue output. Repeat.
- (2) Write. Wait for client request. Wait for data in output cache. Write data. Repeat.
- (3) Monitor. Wait for timer to expire. Report program status. Repeat.
- (4) Control. Wait for control input. Respond to directive. Repeat.

An IEEE standard for threads (POSIX 1003.4a) is currently in development.

d. Security Technology. Distributed computing systems need reliable methods for authentication, authorization, and the assurance of integrity and privacy of messages sent across the network. The DCE security technology is based on Project Athena's Kerberos (Version 5) service. Kerberos provides private authentication services and secure message service. The DCE authorization function is

provided independently via POSIX-compliant access control lists.

e. Distributed File System. Another useful abstraction for building distributed systems is the distributed file system (DFS). A DFS is an efficient method for application-transparent sharing of file-oriented data among computers. While not appropriate for stream transfers, a distributed file system can nonetheless simplify certain data exchange problems, particularly distribution of software across the network.

3. Database Technology. The DSN has traditionally developed custom storage and retrieval methodologies for most of its database management systems (DBMS's). Experience at JPL and throughout industry demonstrates the wide applicability of relational database systems, particularly when implemented in a client-server mode of operation.

Relational DBM's. Such systems have been in use for approximately 20 years and have achieved reliability, performance, recoverability, operability, and security levels acceptable to most applications. Furthermore, related computer-aided software engineering (CASE) tools, program generators, report writers, windowing interfaces, distributed operations, and parallel operation are becoming available.

The Structured Query Language (SQL) has been accepted as an ANSI-standard relational database interface language. Virtually all relational database systems, and many non-relational database systems, provide an SQL interface.

Object-Oriented DBM's. Objects are data that have semantic definitions and procedures attached to them. They are examples of abstract data types. Object-oriented DBM's are designed to support the storage, manipulation, retrieval, and safety of objects. However, there is currently no single, consistent, accepted model or standard for object-oriented languages and systems. Object-oriented DBMS's may take years to be standardized.

Text DBM's. Most of the world's data exist in textual form. Text DBM's have been around since at least 1980. They were initially developed for use in the intelligence community for searching the enormous volume of public information found in publications. Text DBMS's are now emerging as tools to be applied in office automation, publication, and documentation. Text DBMS's suffer from the same weaknesses as object-oriented DBMS's. They have no single codification, definition, theory, or de facto standards. Current efforts in industry are being expended to

broaden the capabilities of text DBMS's to include image and graphical data.

Based on history, it appears that at least 20 years of intensive use and development are required to provide a user with a mature DBMS. Object-oriented and text DBMS's have neither the single solid theoretical base that relational systems have nor do they have a champion. Without such cornerstones the systems may not progress, mature, and become reliable before 2010. Therefore, based on history and present reality, relational DBMS's are viewed as the most likely candidates for database management systems for the DSN in the 2000–2010 era.

4. Algorithms and Heuristics.

a. AI and Knowledge Systems. Expert and knowledge-based systems, more generally called knowledge systems represent specialized areas of AI that have developed rapidly over the last ten years. It has been estimated that there are currently over 10,000 deployed, operational knowledge systems in the United States in many different areas of industry, services, and government. For example, at JPL, the Spacecraft Health Automated Reasoning Prototype (SHARP) knowledge system has been successfully applied to Voyager and Magellan real-time link analysis operations and Galileo operations.

Knowledge systems are computer systems that provide automated support to human workers, or high degrees of autonomy to large-scale systems. This technology has opened many entirely new classes of aerospace applications to computer automation, such as fault diagnosis, fault isolation and repair, resource planning and scheduling, design, and intelligent interfaces to computer systems. Knowledge systems extend the capability of information systems by manipulating data at many levels of abstraction in ways that were previously impossible or achieved only by humans.

The combination of high-performance hardware, software algorithms, heuristics, and technology from many areas of computer science allows knowledge systems to address application problems that were previously not amenable to computer-based solutions, such as dynamic replanning and expert fault diagnosis. In the GIS, knowledge systems will play a major, enabling role in the automation of operations, data management and analysis, and in the achievement of higher degrees of autonomy in the next generation of DSCC's.

b. Neural Networks and Fuzzy Logic. A neural network is an AI technique that is used to mimic the way the human brain works. A significant benefit of neural

networks is their ability to exhibit acceptable computation performance in the presence of noisy or incomplete data. Although neural networks are most frequently implemented in software, the most impressive performance results from neural networks realized in hardware. Optics is a technology particularly well suited for implementing neural computers because of the relative ease with which a programmable, massive interconnection network can be optically synthesized. Experimental work along these lines is currently under way at the California Institute of Technology and elsewhere. At JPL, hardware neural network chips for resource allocation problems have been successfully developed using VLSI techniques.

Another interdisciplinary area of study related to AI and mathematics is fuzzy logic, originally proposed and mathematically formalized in the 1960s and only now undergoing widespread application. As an alternative to probabilistic reasoning, fuzzy logic provides methods for dealing with imprecise data and incomplete information in decision making. These techniques have been successfully used in a wide range of control applications, including space shuttle rendezvous and docking, elevator scheduling, and control of braking on subways. Like neural networks, fuzzy logic controllers can be implemented in special-purpose hardware, yielding impressive performance improvements over conventional controllers in many applications. Although these technologies are still experimental, it is expected that neural networks and fuzzy logic controllers in both hardware and software forms will find use in a variety of application problems in the GIS.

D. Data Transport

The data transport function transfers digital data among GIS computers over local and wide area networks. Government and industry are moving towards standard transport protocols, including message formats, for reliable and timely transmission. The goal is an environment where heterogeneous computers are capable of communicating among themselves by virtue of mutual implementation of ISO-standard protocols.

The protocols provide various services commonly organized into layers by the ISO OSI reference model (see Fig. 10). This section is roughly organized in terms of these layers. The protocols, while they are ISO standards, are often referred to as the OSI suite. For its part, the Federal government has chosen to require ISO-standard protocols. Within the large number of ISO protocols and options, the government has specified a subset of protocols for Federal networks and has published the list as the Government OSI Profile (GOSIP).

At the lower layers, protocols are required to move data reliably and accurately within local area networks (LAN's) and wide area networks (WAN's). The maximum transport speed of these networks depends on the physical characteristics of the medium, the transmitter and receiver technology, and the protocol processing speed. At intermediate layers, protocols must provide reliable process-to-process communication among different subnetworks. The goal of the intermediate layers is to interconnect multiple subnetworks to form a single logical network (an internetwork). The higher layers provide special services ranging from data compression to specific message and handshake procedures for file transfers and electronic mail.

The following sections summarize the current state of the art, forecast trends, and describe the impact of those trends:

1. Media. The physical properties of communications media limit the bit rate and transmission distance of digital links. Speed, limitations, and usage trends of several vital media in the GIS communications architecture are summarized in the following section. Common-carrier circuits are also considered in this section.

a. Wire. Twisted-pair wire is the most common medium-distance transmission medium (1-1000 m) and may be used to transmit data at rates up to 100 Mbps (depending on distance). It is often used for networks within buildings because twisted-pair wire is already a part of the installed telephone plant, relatively inexpensive, and easy to maintain. This medium will continue to be available because of the large installed base; however, it is primarily useful at older facilities.

It is expected that twisted-pair wire will be used in the future within locally instrumented areas to connect computers to peripherals and instrumentation by using such standard interchange standards as RS-232, RS-422, or IEEE 488. It will be an important medium for DSN monitor and control circuits.

Coaxial cable can support higher data rates than twisted pairs without emitting spurious radiation. Rates of hundreds of Mbps can be supported in both baseband and broadband implementations. Such current LAN standards as Ethernet 10 Base 2 (Thinnet) and 10 Base 5 (Standard) are relatively easy to implement. A popular monitor and control standard, the Manufacturing Automation Protocol (MAP), which includes a token-passing bus LAN standard (ISO 8802-4), is also implemented with coaxial cable.

The number of coaxial cable-based LAN's will likely decline because of the increased commercial availability

of high-rate, low-cost optical fiber components. New-generation LAN's based on optical fiber, such as the Fiber Distributed Data Interface (FDDI), will become the LAN's of choice.

b. Microwave. Terrestrial and satellite microwave systems provide high-capacity channels (hundreds of Mbps) and relatively rapid installation. Although subject to weather degradation and environmental disturbances, they are not encumbered with installation problems; they are ideal for backup or alternative routing. Recent development of commercial LAN's that use 18-GHz signals for communications may be applicable; however, the radiation of wireless networks may interfere with sensitive experiments at the DSS's.

One problem with satellite communications is the propagation delay associated with transmitting via geostationary satellites. At heights of approximately 35,000 km, the one-way delay is about 0.25 sec, depending on the elevation of the satellite above the horizon. This delay may significantly affect data throughput if the network uses error-correction methods usually associated with terrestrial communications.

Trends show increasing use of very small aperture terminals (VSAT's) for WAN applications to communicate with terminals in unreachable areas and for mobile communications. The primary usefulness of microwave communications in the DSN will be to provide redundant WAN routing paths for critical missions and temporary communications.

c. Optical Fiber. Optical fiber has the bandwidth to enable terabit-per-second transmission rates. Fiber not only has high bandwidth, but it also has the benefits of low noise and signal attenuation (see Fig. 11), with the result that high-rate data can be transmitted over long distances (e.g., 2 km with multimode fiber and over 40 km with single-mode fiber). Fiber has advantages in security, safety, immunity to electromagnetic interference, and reduced weight and size.

While multimode fiber is less expensive, single-mode fiber is the medium of choice for high-speed digital transport, and the number of installations will likely grow. Research is focused on signaling at higher rates. Current signaling records are 350 Gbps. Another research thrust is in wavelength multiplexing; the result is increased utilization of the optical bandwidth by the establishing of multiple communications channels at different wavelengths over a single fiber. The primary impact on the GIS architecture will be higher data rates over LAN's and common-carrier WAN circuits.

d. Common Carriers and ISDN. There are two major trends that will affect the services offered by the common carriers: the trend to convert long-haul links to optical trunks and the offering of digital interfaces directly to the users. The first trend has been spurred by the adoption of the International Telegraph and Telephone Consultative Committee (CCITT) Synchronous Digital Hierarchy (SDH, formerly SONET) standards for optical networks and the second by the increasing availability of Integrated Services Digital Network (ISDN) central office switches. ISDN is an international standard for providing digital circuits directly to the user and will eventually replace current analog circuits for voice, data, and image (i.e., facsimile and video) transmission.

Synchronous Digital Hierarchy. The SDH standard for optical phone company networks will enable the use of a hierarchy of service data rates significantly higher than the current CCITT-standard G.702 hierarchy (Table 7). At the present time, many corporate users, including JPL's Information Systems Division, are using 45-Mbps DS-3 trunks to replace multiple 1.5-Mbps DS-1 circuits. It is expected that OC-1 (51.84 Mbps) will replace many of these circuits when SDH is fully implemented. Note that the SDH hierarchy goes to 2.4 Gbps. Most of the data speeds in this hierarchy are for internal telephone company support networks. Availability of these high-speed trunks for other users will depend on market demand.

Terrestrial and undersea installation of fiber will make these higher rates available between the DSCC's and JPL. Also, replacing satellite links with fiber will eliminate the typical 0.25-sec propagation delay associated with each satellite hop and will significantly reduce bit errors. It is predicted that by the year 2000 all global links associated with the GIS will be fiber.

Table 8 summarizes the major differences between satellite and fiber circuits. Based on industry projections, optical fiber leased-circuit cost for the year 2000 time frame may be as little as 10 percent of today's rates.

ISDN. For lower data rates (< 1.5 Mbps, narrowband ISDN), user service is provided on three types of circuits (Table 9). The D channels are used for common-channel signaling and (optional) packet transmission; the D channel has its own OSI protocol stack. The basic user B channel may carry voice, high-speed fax data, slow-speed video, or computer data; the protocol is at the user's discretion. Three kinds of connections may be set up with B channels: circuit-switched (similar to today's dial-up service), packet-switched (user is connected to a packet switch), and semipermanent (user has the equivalent of today's dedicated circuit).

Current commercial services include basic and primary-rate service (Table 10). These digital services are direct improvements over slower 9.6- to 19.2-kbps links currently used over the GIS, and their implementation will enable a higher volume of data traffic between the DSCC's and JPL.

Future offerings built on optical fiber capability will make higher data rates available in a service called broadband ISDN (B-ISDN). B-ISDN will eventually permit switched services up to 2.4 Gbps. This service will also enable video conferencing and transmission of high-definition television on a switched, on-demand basis. It is also expected to enable a variable amount of bandwidth on demand, so that the user will pay only for the bandwidth required at a given time.

2. Lower Layer Technology. Network technology can be organized into three general classes, depending on the network dimensions. Each class has a need for lower layer protocols. At the lowest level are networks that interconnect computer boards in high-speed real-time applications. At the next level, computers within buildings may be interconnected through LAN's. At the highest level are WAN's that extend over longer distances. WAN's usually depend on common carriers for the links among computers.

a. Processor-Board Buses. Board-level computers may be linked with serial or parallel buses. Simplicity and relatively low speeds can be met with serial copper buses. High-speed links over optical fiber using such technology as FDDI may be common by the year 2000.

Tightly coupled computer systems with very high data rates among the computers are better served with parallel buses. Current buses are 32 bits wide; 64-bit buses are emerging, with future growth expected to 256 bits. Several standard buses are summarized in Table 11.

It is expected that the number of tightly coupled processors will continue to grow to take advantage of parallel-processing techniques and processor miniaturization. Further, 100+ MIPS workstation computers are emerging that will be used in distributed processing applications, such as on-site, front-end processing of spacecraft signals and back-end high-volume imaging data processing. Research efforts will continue to extend bus widths and integrate more functions on a chip.

b. Local Area Networks. LAN's provide loosely coupled communications between computers typically within a 2-km distance. Their numbers have climbed because of the

growth in personal computers and the need of users to share resources and information. Several standard LAN's are summarized in Table 12.

Current data rates range from 4 to 100 Mbps. The FDDI network, which runs at 100 Mbps, is designed for high-speed mainframe-to-mainframe and LAN-backbone applications. It is expected that the speeds of these types of networks will continue to climb and their useful distances will increase with improvements in fiber-optic technology.

The most significant standard is FDDI. It has a 2-km range between nodes. Multimode fiber is specified in the standard to minimize implementation costs; however, as discussed in Section III.D, the use of higher cost single-mode fiber may permit extensions up to 40 km. The FDDI network is ideal for interconnecting other packet-oriented networks, such as Ethernet and token rings. FDDI-II, a significant departure from FDDI, is emerging with unique capabilities. With the ability to integrate circuit-switched traffic with bursty LAN traffic, FDDI-II can concurrently support data, digital voice, and video traffic.

Research and development efforts are focused on higher speed optical LAN's. Some vendors are developing wireless LAN's to operate at radio frequencies and in the infrared spectrum; this effort is not leading to new breakthroughs in speeds, only to increased flexibility in placing workstations within an office.

c. Wide Area Networks. The High-Level Data Link Control (HDLC) protocol is a common ISO-standard link protocol for providing reliable service at the data-link layer. Errors are detected by using a 16-bit cyclic redundancy check (CRC) algorithm. Error recovery may be done in one of two ways:

- (1) Go-back-N: When an error is detected in a data frame, the receiver notifies the sender of the faulty frame; the sender retransmits it along with every frame that followed.
- (2) Selective retransmission: When an error is detected, the receiver notifies the sender; the sender retransmits only the frame in question.

Most terrestrial networks use the Go-back-N algorithm. However, noise on satellite circuits highly degrades the performance of high-speed satellite links when using this algorithm because of the 0.5-sec round-trip delays; selective retransmission is preferred in such cases.

GOSIP specifies CCITT X.25 as the WAN protocol of choice. It is a three-layer (physical, data-link, and

network) packet network standard. Its link-level protocol is link access procedure balanced (LAPB), a version of HDLC but without selective retransmission features. Therefore, it is limited to lower speed satellite links, or optical fiber links.

With X.25, data are packetized into basic 128-byte protocol data units (unless the overall network has been designed for larger packets). At the network layer, X.25 software establishes a path across multiple computers from the sender to the receiver through the use of a connection-request packet, and the data packets flow across this virtual circuit until the connection is terminated.

Frame relay is an emerging protocol that will enable higher throughput over common carrier circuits by reducing the number of acknowledgments, as compared with X.25. Research and development efforts are focused on minimizing delays due to protocol overhead and providing dynamic bandwidth allocation and optimal routing algorithms.

3. Internetwork Communications. When there is connectivity between LAN's and WAN's, additional users are able to share resources, access common databases, and exchange data. Standard internetwork protocols provide routing among the subnetworks and process-to-process communications, and lay the foundation for valuable application services, such as electronic mail, file management and file transfer. Research goals are to make details at this level of communication transparent to users, increase information throughput, and reduce delays with efficient software implementation architectures.

The most popular internetwork protocols today are two Internet protocols: Transmission Control Protocol and Internet Protocol. Together they are capable of providing interprocess communications among applications on widely separate computers and networks. TCP operates on the principle of virtual-circuit connectivity; that is, it establishes a path to the destination process with a connection request packet, transfers the data sequentially and reliably, and after completion of the transfer, terminates the connection.

GOSIP specifies a pair of ISO protocols that perform the same services as TCP/IP: Connectionless Network Protocol (CLNP) and Transport Protocol Class 4 (TP4). They are similar to TCP/IP but are not interoperable because of different addressing schemes and some other minor differences (see Table 13).

TCP/IP and TP4/CLNP protocols themselves are fairly efficient; however, many current implementations are in-

efficient because of how the code is executed and messages are buffered. Intermediate protocol layers constitute an area of continuing research to reduce delays in interprocess communications and increase performance.

4. Application Services. Protocols in the upper layers of the OSI model provide powerful application-oriented services for users of internetworks. They range from standard file transfer and electronic mail protocols to message formats for robot applications. Several are extremely important to the DSN and will be summarized here.

a. File Transfer. The Internet file management protocol is the File Transfer Protocol (FTP). It is a popular protocol for moving files from one host to another across a network. It is based on TCP/IP connections.

OSI provides a more powerful file-oriented protocol, File Transfer, Access, and Management (FTAM), that may be used for distributed database applications, document retrieval and updating (e.g., library information services), and special messaging systems that transfer long text messages.

FTAM is a two-party file transfer protocol; in other words, there is a controller of the file activity (initiator) who directs the action and a responder who responds to the initiator in a passive role. All file transfer and access operations occur between an initiator and responder over connection-oriented paths (virtual circuits, i.e., CLNP/TP4). Three-party file transfer is a future consideration that is expected to be available by the year 2000.

b. Electronic Mail. Electronic mail allows people to exchange text messages over networks. Internet uses the Simple Mail Transfer Protocol (SMTP). OSI specifies CCITT X.400 as its message-handling service (MHS).

MHS is composed of a series of message-transfer agents (intermediate computers) that are responsible for relaying a message from the sender to the receiver. The transfer agent serving the recipient need not be active when the message leaves the originator's transfer agent; the message can be stored at an intermediate location until the recipient becomes operational.

c. Network Management. At the present time, the Internet Simple Network Management Protocol (SNMP) is implemented on all new commercial bridges and routers. It is used to report the performance of network resources to a central monitoring location. The trend is toward implementing this protocol on all network elements, including workstations and personal computers.

The OSI equivalent is called the Common Management Information Protocol (CMIP). Although CMIP is specified in the Government Network Management Profile (GNMP), which is a network-management equivalent to GOSIP, OSI network management services are only in a state of definition and development today. The basic service provides a capability to report events remotely, read attributes of the network hardware, set a characteristic (an attribute that can be set to influence the operation of the network), set a characteristic if and only if the previous value was in a certain specified ranges, and demand actions of the resource.

SNMP and CMIP are used to report the status and performance of network elements (e.g., routers) to a network management system. Typically the network elements generate events that require immediate management attention. The network management system logs in the events, and depending on whether the option is requested, will keep track of their frequency, and notify a person when a threshold has been exceeded. The system will enable the following services: fault management, accounting management, configuration management, performance management, and security management. All these functions are expected to be fully developed by the year 2000.

d. Manufacturing Message Specification. Manufacturing Message Specification (MMS) is an OSI-unique application layer service that has no counterpart among the Internet protocols, and it illustrates the thoroughness of the OSI approach. MMS is a standard messaging service that may be used within application programs for monitor and control of instruments and robots. In a sense, it is the monitor and control equivalent of an electronic mail standard, only it provides a device interface (versus an electronic mail-human interface). MMS also uses the transport services of the OSI protocol stack.

5. Interfaces to Other Networks and GOSIP. The GIS, although internally protected, will pass data through communication servers to several external national networks, which in turn will connect to international, university, and research center networks. These include the NASCOM II, NASA Science Network, and Space Physics Analysis Network (SPAN). Because of international acceptance of the OSI protocols, it is expected that all these networks will interoperate with OSI protocols.

The CCSDS Advanced Orbiting Systems (AOS) architecture (CCSDS 701.0-B-1), which recommends standards for data transmission over special-purpose space data links and their associated interface networks also recommends OSI protocols. The CCSDS provides the following guid-

ance: "Wherever possible and technically appropriate, interfaces with the commercial network environment are provided so that the emerging infrastructure of commercial OSI may be exploited."

E. Human Interfaces

A variety of technologies that provide improved methods for allowing users to interface with computer systems will become available during the next decade. These interface technologies focus on improvement in the amount of information that the user can perceive from a given interface configuration. There are three general categories of interface technologies: graphic visualization, hypermedia, and AI.

1. Graphic Visualization. Graphic visualization may be of potential use in GIS human interfaces in the following areas:

- (1) Representing multidimensional data in images on computer graphics displays and in a form that allows people to perceive, amplify, and interpret the data.
- (2) Synthesizing data into animated models. Animation time may represent an actual time variable in the data, perhaps scaled, or it may be mapped from an arbitrary variable in the data, presenting novel opportunities for perception and discovery on the part of the user.
- (3) Interactively applying experimental transformations and algorithms to explore the effects on the data (e.g., filters, Fourier transforms, color maps, projections, and nonlinear geometric mapping).
- (4) Monitoring and debugging computer simulations of physical components and processes through the use of graphical representations of equipment (e.g., antennas, spacecraft, etc.).
- (5) Implementing virtual reality, where a computer and peripheral devices create a complete sensory environment for three-dimensional visualization applicable to such interactive activities as design, training, and remote control.

2. Hypermedia. Hypermedia software technology enables a user to retrieve data in various formats in one or more display windows. The formats may include text, graphics, animation, digital audio, and/or video. The data may be retrieved nonlinearly, that is, users may interactively traverse levels of description in a file from high-level overview to low-level details, interactively trace cross-references within a file and retrieve nonlocal references to

other files, or interactively connect with off-line databases and application programs.

Hypermedia may be useful in the following DSN areas:

- (1) Software development. The use of hypermedia may minimize time that programmers typically spend tracing cross-references in code by automating the tracing of such cross-references. Multimedia comments, such as text, diagrams, audio, and video, may be inserted.
- (2) Software engineering. Hypermedia enables such features as an encyclopedia of software components that supports software reuse, tracking of software versions and changes, hardware configuration, etc.
- (3) Requirements, design, training, and reference documentation. These documents are voluminous, expansive, and full of cross-references, the kind of complex documents for which hypermedia was originally devised.
- (4) Browsing and searching knowledge bases. Hypermedia is the natural technology for information retrieval from AI knowledge bases. The nonlinear structure of hyperdocuments can easily be made to mirror the inherently nonlinear structure of knowledge bases (as compared with the linear structure of relational databases).

3. Artificial Intelligence Support. AI techniques that help translate raw data into knowledge may be applied in the following ways:

- (1) Context sensitivity. Interpreting events and actions based on the current context. In a scientific domain, for example, context sensitivity can filter data by using dynamic thresholding.
- (2) Abstraction. Using AI techniques to abstract information and present summaries to the user. For scientists, AI can be applied to data sets to identify interesting features, summarize results and indicate areas that warrant additional review by the scientist. For engineering operators, AI can provide techniques for knowing what state the entire system is in and how ongoing activities are expected to affect that state.

IV. Architectures

The GIS architecture described and recommended in this section is intended to meet performance requirements in the 2000–2010 time frame and provide adequate margins; minimize development, implementation, and operational costs; take advantage of technology expected to be

available in the relevant time frame; and allow for a smooth transition from the current architecture.

The overall GIS architecture (Fig. 12) is designed as a single unified system from the front-end area to the end user. Processing activities and network resources (e.g., communication and database servers) are located at the most appropriate points to meet performance requirements, and standard networks are used to move data from one to another. Data buffers are provided at several points in the architecture, in each case insulating upstream parts of the architecture from the effects of failures downstream. In particular, buffering is provided at each DSCC to accommodate WAN outages.

Unbuffered real-time data exist only at the front end between the digital receivers and telemetry processors. Currently it is expected that the first level of buffering will be provided at the point where frames (or CCSDS packets) first exist, for example, at the output of the telemetry or radio science subsystem. Subsequent processing stages are based on client-server relationships. Data flow after the data are packetized may be temporarily buffered before transmission to control data flow. The goal at each stage is to perform a clearly defined transformation that minimizes overall operational costs, including data transport costs.

Multimission level 0 processing (LZP) is done at each DSCC, to the extent possible with locally available data. Final multimission level 0 merge (LZM) occurs at the Central Site (JPL), where it first becomes possible to consolidate data from all DSCC's into a single continuous data stream.

Project-specific processing is done at JPL (or at other locations, as negotiated with non-JPL DSN users). Data to be processed are made available to projects and other users in a uniform manner, using standardized protocols, procedures, and formats.

A monitor and control network complements the processing architecture, and its processors perform several major functions. In particular, a Network Control Facility (NCF) at the Central Site will remotely control the DSCC antennas through appropriate directives and maintain a management information base to enable status monitoring of all subsystems.

An appropriate version of the information base will be maintained at the DSCC's to automatically reconfigure antennas for each mission activity and automatically detect, isolate, and assist recovery from faults. A subset of

this information will be available at several locations at the DSCC to provide a human interface when necessary for manual intervention.

The following is a summary of the proposed GIS architecture:

- (1) Unified architecture from front-end area to end user.
 - (a) Client-server model for each processing stage.
 - (b) Unbuffered real-time acquisition only at the front end.
 - (c) Standardized interprocess communication.
- (2) Open-system standards to achieve interoperability.
- (3) Direct delivery of level 0 data to the user from the DSCC's, if desired.
- (4) Dedicated telemetry processors for each receiver.
- (5) Security against unauthorized access and accidental errors.
- (6) Highly automated monitor and control (M&C).
 - (a) Centralized control of entire network.
 - (b) Automated configuration of support activities.
 - (c) Vastly improved human interface.
 - (d) Automated decision support systems for network resource management, performance analysis, fault diagnosis, and contingency management.
 - (e) Embedded intelligent computer-aided training systems.

The Ground Information System is logically organized into four reasonably distinct high-level functions: processing, monitor and control, data transport, and software. For ease of analysis, each of these functional areas is considered separately in the following sections.

A. Processing

In the early stages of the study, several classes of processing architectures were considered (see Table 14). Each of these architectural candidates (in its pure form) was rejected with the better features of each subsumed into the recommended architecture. The DSCC aspect of this architecture and the Central Site aspect are detailed in this section. The architecture includes elements of these

candidates, organized to minimize implementation and operational costs, to take advantage of technology expected to be available in the relevant time frame, and to allow for a smooth transition from the current architecture.

1. DSCC Site. The DSCC processing architecture performs initial level 0 processing of telemetry, and initial processing of radio science, VLBI, ranging and tracking, and command data. The estimated data rates for each of these subsystems are summarized in Table 15. In the DSCC architecture (Fig. 13), downconverted signals are routed to the Block V receiver, VLBI subsystem, or special acquisition equipment through an analog distribution switch.

In some cases processing is initiated by the Block V receiver and data are passed directly to a dedicated subsystem (telemetry and ranging). In other cases, such as tracking and radio science, data are routed over a network to an available subsystem. The data paths depend on the real-time data rate, with functions with higher rates requiring dedicated processors. VLBI processing is a dedicated string from the point of signal downconversion and digitization to the stage of formatting the data for transmission to JPL for correlation with data from another antenna.

Each subsystem transmits its formatted packets to a network communications server over the DSCC backbone LAN. Reliable OSI protocols are utilized for communication to JPL. The M&C network, not explicitly illustrated in Fig. 14, complements the processing architecture and is the means used to communicate configuration data, status, and commands among the subsystems.

The command subsystem stores commands received from JPL. The subsystem modulates a subcarrier to form a baseband signal that is transmitted to an exciter under the control of the M&C subsystem.

The Ranging Subsystem, closely coupled to the Block V receiver, receives predict data from JPL and uplinks appropriate ranging signals through the exciter. Test data from JPL are sent to the test subsystem and it, in turn, modulates a signal to the exciter for telemetry testing.

a. Telemetry Processing. The Block V receiver and its successors will feed data to the telemetry process in the proposed time frame. For analysis, it is assumed that there will be one Block V receiver at each Complex, with the potential for at least 22 channel processors supporting seven antennas, and by the year 2010, at least 32 for 10

antennas. Two combiners are expected by 2000 and three in 2010, with dedicated channel processors.

The goal of the architecture is to pass telemetry data directly from a Block V channel processor to the telemetry processor with virtually no delays (real-time), process the data, and deliver it in sequence and with no detectable errors to JPL (near-real-time). Short delays may be required for buffering in the communication links to detect and correct packet errors downstream.

There are three fundamental architectures that are considered for the channel processor–telemetry processing operations: dedicated, switched, and distributed. All architectures may operate in a fully automated mode.

In the dedicated processor architecture, Fig. 14(a), a separate set of electronics is attached to each Block V channel processor. Two or more telemetry processors may be dedicated for increased availability. Processors may be reconfigured for different missions by downloading software. High availability is achieved through processor redundancy and component redundancy within each processor.

The switched architecture, Fig. 14(b), interconnects the receivers and processors via a matrix switch, similar to the way that it is done today. While mechanical switches today suffer from low reliability, this proposed architecture may be envisioned with highly reliable switches (e.g., asynchronous transfer mode switches) in the proposed time frame. The switched architecture provides flexibility in configuration management. High availability is achieved by telemetry processor redundancy.

The distributed architecture, Fig. 14(c), interconnects the receivers and processors via a common multiple access network (e.g., a LAN), and provides connectivity for passing data from any receiver to any processor on the network. This approach envisions extensible use of general-purpose computers and commercial network technology integrated into a DSCC-distributed processing system. The distributed network provides extra flexibility in configuration management and enhances options for load balancing among machines. As with the switched architecture, high availability is achieved through computer platform redundancy.

There may be occasions for exceptional support for which dedicated high-speed acquisition equipment may be used, and the data stored on portable media (e.g., tapes or disks) for off-line delivery to the users. This capability was illustrated earlier in Fig. 13.

Ancillary data, such as received power levels and system temperature, are required for real-time processing. Some of these data are generated by the Block V receiver, transmitted to a network file server and distributed to each telemetry processor. Also, software configuration data (e.g., the type of decoding, etc.) must be transferred to each telemetry processor by Monitor and Control during configuration of a link.

Output from the telemetry processor is transmitted over the DSCC backbone LAN to the communication server. The communication server sends the data to the Central Site for merging, deletion of duplicate data, and completion of multimission level-zero processing. The data are then available to users for further processing over mission-unique networks.

In addition, the DSN antennas may be arrayed with remote antennas (e.g., Goldstone, California, antennas with the Very Large Array [VLA], in Socorro, New Mexico). Soft-symbol combining must be performed within the DSCC combiners. Other occasions may require GIS to transmit these same symbols to other partners in an array; it is assumed that Block V receiver-to-array partner communications will be specially arranged, and there will be little impact on the information system.

Information rates above 10 Mbps may be expected for Earth orbiters and lunar missions, and these usually require a convolutional code rate of 1/2. Deep space missions with rates less than 10 Mbps usually require code rates of 1/6. Thus, for orbiters and lunar missions, one megabit per second of information translates into about 18 Mbps of Block V channel processor output (assuming Reed-Solomon [255,223] code), and with deep space missions, one megabit per second translates into approximately 55 Mbps (see Section II.A). Table 16 translates the proposed spacecraft telemetry profile into a telemetry processing rate profile. These rates have a major impact on the architecture.

Dedicated Processing Evaluation. The dedicated architecture (Fig. 15) features Block V channel processors and combiners with dedicated telemetry processors. It is assumed that communications between the channel processor and the telemetry processor are over a dedicated backplane bus, although a dedicated LAN may be used. In contrast to the distributed architecture, the data path is dedicated—with no sharing. In contrast to the switched architecture, the path is permanent.

Specially designed protocols are assumed for data transfer. The optimal performance is summarized in Table 17.

It can be seen that the potential for growth is excellent—even the 1980s bus technology meets the conservative mission-support forecast. The year 2000 bus technology meets the aggressive forecast.

The advantages of this architecture include:

- (1) The architecture meets all year 2010 requirements, up to 450-Mbps backup support.
- (2) Potential points of failure are the Block V channel processor, telemetry processor, and active interprocessor bus.
- (3) Dual telemetry processors within each string improve fault tolerance.

Major risks include:

- (1) Computers (I/O software) cannot keep up with bus speed.
- (2) Custom-designed interprocessor communication software may lead to higher life-cycle software costs.
- (3) Bus management (fault detection, recovery, etc.) will be nonstandard software, and may lead to high life-cycle costs.
- (4) The cost of many telemetry processors may be high.

Switched Processing Evaluation. A switched telemetry architecture is illustrated in Fig. 16. The architecture features multiple Block V channel processors and combiners that feed telemetry data to one of many switch-selectable telemetry processors.

The switch is computer-controlled and provides a temporary dedicated path on demand. The highest speed of the path is the serial I/O speed of the channel processor/telemetry processor link. Currently, the best technology is about 100–130 Mbps.

A major issue is the generally poor reliability of mechanical switch technology. It is assumed that more reliable technology, such as the asynchronous transfer mode switches, will evolve by the year 2000 and will provide high reliability.

The advantages of the switched architecture include:

- (1) Telemetry processors can be easily changed as mission-unique requirements are added.
- (2) A minimum complement of telemetry processors reduces costs.

- (3) Redundant telemetry processors increase availability.

The major risks include:

- (1) Forecasted telemetry rates may exceed I/O capacity.
- (2) The switches may be expensive.
- (3) The switch is a single point of failure.

Distributed Processing Evaluation. A distributed telemetry architecture is illustrated in Fig. 17. The architecture features multiple Block V channel processors and combiners, each feeding telemetry data to one of many software-selectable telemetry processors.

An ultra-high-speed network provides reliable data transfer among the receivers and telemetry processors. The network is shared—it provides a data path for all acquisitions that are under way at any moment. The optimal LAN performance is summarized in Table 18, with an indication of the instantaneous aggregate information rate. It can be seen that the best current technology (100 Mbps) provides an aggregate capability of 1.8 Mbps; a 1-Gbps LAN would provide aggregate capability of 18 Mbps.

The advantages of the distributed architecture include:

- (1) It may enable open interprocessor communications with commercial off-the-shelf (COTS) hardware and software interfaces.
- (2) Telemetry processors can be easily changed as mission-unique requirements are added.
- (3) FDDI, a dual counter-rotating ring, is available in 1990s technology that automatically recovers from faults.

The major risks include:

- (1) An aggregate telemetry stream may exceed LAN capacity—a 100-Mbps LAN can only support an aggregate of 1.8 Mbps of information.
- (2) A one-Gbps standard LAN may not be available.
- (3) Computers (I/O software) cannot communicate at the LAN speed. This is a 1990s area of research.
- (4) This configuration has low growth potential.

Evaluation. The alternative telemetry architectures were evaluated according to the criteria described in Section I.D, and the results are summarized in Table 19.

Based on performance, growth potential, and high availability, a dedicated telemetry architecture is recommended. Additional study may show that adding redundant telemetry processors to each string will significantly improve availability. Primary concerns will be the cost of custom interprocess communication software and the cost of many processors.

A hybrid of the switched and dedicated architectures was considered. The hybrid involves an addressable bus that permits each receiver or combiner to talk to any telemetry processor over its own dedicated bus. The hybrid is based on the ability to provide a backplane bus from any channel processor to any telemetry processor under program control. Further study will determine whether the additional cost of this technique justifies the improvement in availability. A cost study should be made before a final commitment is given to the dedicated architecture.

b. Additional DSCC Processes.

Tracking. Tracking is generally a low-speed process whereby Doppler and ranging packets are sent from a ranging processor associated with the Block V receiver to the tracking processor. Its primary function is to take these packets and format them into a suitable product for navigation processing by using ancillary data and predict data from the M&C database. After formatting, the packets are transmitted to users, such as the spacecraft navigation team, through the communication server.

Command. Command processing is another low-rate process that takes commands sent to the DSCC from JPL and sends command modulation signals to the exciter and transmitter for an appropriate antenna. The command subsystem also sends a verification back to the mission control team at JPL.

Radio Science. Radio science packets are generated by the Block V receiver and transmitted to a radio science subsystem for formation and transmission to the Central Site. The potential rate from each receiver is 1.2 Mbps, and several experiments may acquire data simultaneously. After formatting, the packets are transmitted to the radio science team through the DSCC communication server.

VLBI. VLBI ranging is a high-rate process (approximately 130 Mbps) that typically digitizes and stores X- and S-band signals from radio sources or special tones from a spacecraft. The data are temporarily stored on a high-speed, high-volume storage device and later played back at a slower rate (e.g., 1.5 Mbps) to JPL for correlation with data simultaneously acquired at another DSCC.

The results are passed to flight projects for navigation processing. Approximately 20–30 min of 130-Mbps data are acquired at one time.

If required, quick-look data may be acquired during the first 15–20 sec and sent to the radio science team for real-time processing. If anomalies are discovered in the data, adjustments to the instruments may be performed with real-time commands.

In summary, the overall processing architecture has the following features:

- (1) The architecture meets performance requirements in the forecasted time frame and provides an adequate margin. Each Block V channel processor is provided with at least two dedicated telemetry processors.
- (2) DSCC on-site processing is performed using distributed open systems.
- (3) Communications are based on open-systems standards.
- (4) Final-product processing traffic flows to the DSCC backbone LAN and to the DSCC communication server. Subsystem-unique requirements (e.g., VLBI) are accommodated with intermediate high-performance processing and storage devices.
- (5) Certain exceptional types of processing are acquired in real-time, but delivered off-line to JPL or direct to users (e.g., orbiting VLBI) via portable media.

2. Central Site (JPL). The Central Site performs two major GIS functions and provides an open systems interface to the multimission and mission-unique subnetworks (Fig. 18). A key characteristic is the high-speed DSN backbone LAN that permits client-server architectures.

The first GIS function is completion of level-zero processing, where duplicate data records from the DSCC's are deleted and the data streams are merged and prepared for delivery to higher level processing. This relatively high-rate function may involve distributed processing over a dedicated LAN (level-zero merge subnetwork). The output is delivered in near-real-time to multimission networks. Temporary buffering will be available for network flow control; i.e., to keep the level-zero merge processor from overrunning the multimission telemetry (SFOC) processor.

The second major Central Site function is monitor and control of the DSN. This is accomplished over a subnetwork performing network and multilink control. Network control includes the highest level functions in the control

hierarchy, such as resource management, scheduling, and data network management. Operators will control the entire DSN through multilink control workstations (see Section IV.B).

A key architecture requirement is a large operations database for schedules and resource management, accessible by each control subsystem. The database server will also make DSCC equipment status data available to the end users for later analysis.

The DSN backbone LAN interfaces through a local communications server (router) to the SFOC backbone and related local networks. This interface is designed to provide network security. Because of the extensive use of open standards, subsystems and users will use simple and relatively inexpensive communications software to exchange data across this interface. Through the use of standard protocols, data residing in DSCC or Central Site databases, such as resource status at the time of acquisition, may be retrieved by the users on demand, assuming that the network administrator has authorized such access.

B. Monitor and Control

1. Centralized Control. In the proposed M&C architecture (Fig. 19), the DSN is controlled from a central location, the NCF. The control structure is hierarchical: network control (the highest level), multilink control, link control, and subsystem control.

The two higher levels reside at the NCF and are integrated human-computer operations. The lower two levels, link and subsystem control, are located at the individual DSCC's and are automated. In addition, there are site-resource controllers at each of the DSCC's that autonomously monitor the health and status of station resources when they are not assigned to a link. Operations personnel are located primarily at the NCF, with only maintenance and backup operations personnel located at each of the DSCC's.

A generic model for all control layers is presented in Fig. 20. (The link controller is pictured as an example.) The controller receives assignments and the necessary support data from its superior controller, delegates control to its subordinate elements, receives monitor information (health, status, configuration, performance, and events) from each of the elements under its control, and consolidates and reports monitor information to its superior. The controllers themselves perform a variety of functions, including anomaly detection, isolation, diagnosis, and correction. If an anomaly cannot be handled at the current

level of control, the controller requests support from its superior. If more information is needed, the controller requests additional data from its subordinates.

Subsystem controllers interface with the subassemblies and devices that constitute the actual subsystem. They essentially perform the same basic functions that current subsystem controllers perform, but with a greater ability to execute control sequences. Rather than the explicit, low-level directives used to interface with the subsystems in the current DSN, control messages will be at higher levels (e.g., "Configure for Voyager 2 one-way"). The subsystems will have the intelligence to interpret these higher level directives and perform the necessary actions. Any problems encountered in executing a higher level function will be reported up to the next level of control (the link controller if the equipment is assigned to a link, otherwise the site-resource controller).

Subsystem controllers will have extensive built-in test equipment and self-monitoring functions that will assess the health and status of the subsystems internally and alert the higher level M&C systems to any problem.

The primary functional unit for the DSN is a link—the series of equipment necessary to establish and maintain communications with a spacecraft or to meet a scientific objective. In current DSN operations, setting up a link to support a given operation (referred to as precalibration) is a time-consuming manual process. In the proposed architecture, not only will this task be automated, but the monitoring of the link during the track and the performance of any postcalibration activity will also be automated.

The computer-based link controller will initiate, monitor, and control all the actions necessary to accomplish an activity. The link controller will activate the given subsystem controllers, pass down any required support data, monitor the progress, and respond to any problems that occur. Overall link health, status, and performance will be reported to the next higher level of control.

The functions currently performed by link monitor and control (LMC) operators at each of the DSCC's will be automated or will migrate to an NCF multilink controller (MLC). Multilink control is an integrated human-computer function. With the addition of automated background monitoring capability, the MLC operators will now be able to control and monitor the health, status, and performance of several links. There are three modes in which multilink controllers will operate: coordinated, consolidated, and combination.

Coordinated activities are those which require that two or more DSS's operate together to support a customer request. Examples of coordinated activities are antenna arraying and VLBI. During coordinated activities, the multilink operators must ensure that the DSS's and associated equipment are working together to provide the requested service. In some cases, such as VLBI, this requires coordinating antennas at different DSCC's. In a consolidated mode, the multilink controller is responsible for several links that are operating independently. The controller can work with several different types of antennas performing several different types of services. In the combined mode, the multilink controller supports both coordinated and consolidated control.

Site-resource controllers are automated processes that monitor and control equipment that is not operating in a link. In addition, the site-resource controllers are responsible for receiving and storing any support data that is to be maintained on site, scheduling site-specific resources, and resolving any conflict that may arise concerning the disposition of site resources.

The NCF, located at JPL, is the primary control point for the entire DSN. It provides the interface with the DSN customers. The NCF schedules DSN resources, processes customer requests and support data into a format usable by the DSN, and resolves any conflicts that cannot be resolved at a lower level (primarily those which require interaction with DSN customers, or which involve DSN resources from several DSCC's). The NCF network controller subsystem initiates all DSN operations and monitors the health, status, and performance of the DSN. The network controller delegates more specific monitor and control functions to the lower level controllers. The functions of the NCF are described in more detail in the operations concept (see Section IV.B).

2. Automation. The proposed architecture is highly automated. At each level in the control hierarchy, routine and contingency operations are automated, while functions requiring human decision making are supported by automated decision aids. User interfaces with embedded help and training facilities enable human operators to perform complex functions efficiently and effectively despite varying levels of experience and training. The proposed architecture incorporates automation techniques to support a variety of functions as detailed in the following sections.

DSN resources are currently controlled by text strings (directives) manually entered at the DSCC LMC consoles. The status is monitored by operators who interpret the various responses that are provided by the subsystems.

In positive-control systems, all control actions have corresponding explicit verifications that they have been received and processed. A major problem with the existing M&C architecture is that it does not provide positive control, that is, operators must interpret a variety of disassociated information in order to determine the true response of the link equipment.

In the proposed architecture, control functions for routine and specific contingency operations are automated. The higher level controller only needs to give the command, and the lower level controller will execute the appropriate control sequences to support the requested activity. The higher level controller coordinates actions among lower level controllers and has immediate feedback on the status, configuration, health, and performance of the control activities and associated equipment. In the proposed architecture, automated monitoring of health, status, performance, and configuration, with fast and efficient handling of exceptions, will support situational awareness at all levels of control.

Situational awareness means that the controllers (human or computer) have accurate knowledge of the health, status, configuration, and performance of the equipment and processes that are under their control. Situational awareness is especially important in the multilink operations modes because controllers need a clear picture of ongoing operations in order to assess the situation and deal with unexpected events. Situational awareness depends heavily on background monitoring capability, accurate and timely reporting from lower level controllers, and a user interface that ensures that the most important information is brought to the attention of the controller quickly and reliably.

The architecture incorporates embedded training and help facilities that enable human operators to learn new skills, and improve and maintain existing ones. Operators are able to use the help facilities to aid in performing tasks. Help facilities include instructions on how to perform certain activities, how to access and interpret support data, and how to retrieve desired information. The help facilities will also assist operators at a higher level of control to interact with lower level controllers.

The training facilities are accessible from the operations consoles. Training consists of two types: tutoring on declarative knowledge (such as the description of a subsystem, or discussion of the scientific principles behind VLBI) and procedural knowledge (such as how to perform a track). In the procedural training mode, the operator is able to run scenarios that simulate the procedure under

both optimal and anomalous conditions. The operator is able to activate a training session while performing actual operations and is alerted to any condition that requires human intervention.

The proposed architecture controls access to the DSN with tight security and safety features. In addition to the features provided by network security applications that prevent tampering with communications, an active security feature will be incorporated into monitor and control. Users (operators and experimenters who want to control DSN resources) will be modeled in a "client-server" relationship. Privileges to perform certain control functions will be restricted to only those users who are authorized. In addition, special security features will ensure that training activities are isolated from actual operations.

3. Operations Concept. The following paragraphs describe an operations scenario in terms of the control flow from the NCF to the DSCC's, and the monitor flow from the DSCC's to the NCF. This concept is summarized in Fig. 21.

a. Control Flow: Nominal Operations. There are two types of control flow: directions on what to do and authorization to do it. Standard operating procedures will be in place at the DSCC's and the NCF, and schedules will be transmitted in advance to the appropriate resource controllers. The authorization to perform operations will generally be given in real time. The following example describes the control flow for routine operations under nominal circumstances.

All DSN activities originate with a customer request for support. The NCF is responsible for coordinating these requests and iterating them with the various customers to develop a workable schedule. Scheduling is mainly the process of assigning an antenna (or group of antennas) to support a customer request.

Once the NCF has an acceptable schedule, the network controller assigns support activities to the multilink controllers. The multilink controllers, in turn, request the necessary link controllers and link equipment from the DSCC resource manager. The link controllers are configured to support the requested activities and the necessary support data are routed to them. The link controllers then perform an analogous function by configuring the subsystems assigned to their links. For example: The NCF may output a schedule of support activities that covers a 24-hr window. The network controller uses this schedule to create assignments for each of the MLC's. The assignments are distributed and authorization to execute the schedule is given.

MLC operations based on the NCF schedule are shown in detail in Fig. 22. In the figure, multilink controller A (MLC-A) has three activities to support: (1) a Voyager one-way track on DSS 15, (2) a Pioneer 10 ranging on DSS 43, and (3) a Galileo VLBI delta-differential one-way range (DDOR) on DSS's 14 and 63. MLC-A assigns four link controllers to configure the links for each of the activities. It also loads the appropriate procedures for performing its monitor functions and sets up the displays necessary to provide information to the human operator. For this particular set of assignments, MLC-A is working in a combined mode: It must coordinate the two DSS's assigned to the DDOR while consolidating control over the other two tracks.

The link controller (Fig. 23) assigned to the Voyager track loads the appropriate procedures from its procedure library and initializes the equipment assigned to its link. When it receives the authorization from MLC-A, it initiates the precalibration sequence in the subsystems. The link controller coordinates all the subsystem actions, based on the timeline associated with the particular track and the detailed timing information provided in the support data.

Each subsystem is initialized and, under the direction of the link controller, executes a series of steps necessary to configure for the Voyager track. The necessary procedures are in the subsystems' local memory. The link controller only needs to send a high-level instruction to the subsystem.

b. Monitor Flow: Nominal Operations. There are two fundamental types of monitor data: The first supports positive control by providing immediate feedback as to the results of control actions; the second provides health, status, and performance information regardless of any control actions. Therefore, monitoring has two components: knowing the status of ongoing operations and knowing the status of the resources. While control is a top-down function, monitoring is a bottom-up function. Information, in general, is sent from a lower level to its superior where it is interpreted, consolidated, summarized, and sent up the chain to the next level. The following example illustrates the flow of monitor data during normal operations continuing from the control scenario presented above.

A subsystem, e.g., the receiver, has completed a portion of its precalibration activities for the Voyager track. It sends a message to the link controller stating that it has completed the procedure and the resulting configuration, health, and status of the equipment is as desired.

The link controller receives this message, consolidates it with similar messages from other subsystems and determines that the precalibration sequence has been completed. The link controller then sends a message to the MLC-A reporting that precalibration has been completed and specifies the health, status, and configuration of the Voyager link as a whole.

MLC-A uses this information to create a summary health and status report for the network controller, which consolidates all the reports from the various MLC's and determines the overall health, status, configuration, and performance of the DSN as a whole.

c. Operations Under Anomalous Conditions. If a problem occurs during the course of an operational activity, it is handled at the lowest control level possible. For example, if a subsystem experiences a failed module, then the subsystem controller detects, isolates, diagnoses, and repairs or replaces the module (i.e., switches to a redundant unit). It then reports this event and subsequent change in configuration to the link controller.

If the subsystem cannot recover, it informs the link controller that the subsystem is not operating nominally. The link controller determines whether to proceed with the subnominal equipment or to replace the subsystem in the link, and informs MLC-A.

If the link controller decides to reconfigure the link, it requests resources from the site-resource controller. If they are available, then the change is made and reported to the MLC-A. If not, then the link controller requests a change from MLC-A.

MLC-A receives the status report from the link controller and notifies the network controller, as appropriate. If the link controller was unable to resolve the problem, MLC-A assesses the situation and determines whether to assign a new DSS to support the task, request that resources be pulled from some other task, or cancel the track. In the event that the decision is sensitive, MLC-A provides the necessary information to the network controller, which has ultimate decision-making authority over the DSN resources.

4. Backup Operations. It is important to provide for backup operations capability in the event of a natural disaster or other emergency. A complete failure-modes analysis was not performed; however, several types of failures were evaluated to identify how the proposed architecture could address such problems. The proposed architecture

makes use of redundancy, interoperability, and additional operations procedures for backup operations. To understand the nature of backup operations, it is important to first recognize some of the important features of the proposed architecture and some characteristics of operations:

- (1) The link and multilink controllers use the same basic architecture. The differences are in the databases used to perform their activities; therefore, the functionality of the controller can be changed by loading a new database.
- (2) The communications at the DSCC's will be very reliable. Redundant LAN's will ensure continuous communications at the DSCC's.
- (3) Storage is relatively inexpensive. Support data and procedural information can be duplicated and archived at each DSCC.
- (4) The amount of control information that must be sent in real time is minimal. In the event of an emergency, certain information can be sent over telephone lines.
- (5) The DSN (and the flight projects it supports) use established procedures which are known a priori. For most types of operations, the basics are defined well in advance, although specific parameters or sequences might change at the last minute.

The three primary types of failures—loss of a resource, loss of communications, and loss of the NCF—are addressed next. In the event that one of the controllers (network, multilink, link, or subsystem) fails, the DSN will rely on redundant units to replace the failed unit. The major issue then becomes one of graceful transition during ongoing operations. There are a variety of strategies that could be employed, including hot backups (redundant, active units performing the same function), warm backups (equipment on-line and running, but not necessarily performing the same function), and cold spares with context switching (spare equipment that can be brought up but with some level of time delay). An evaluation of the different redundancy management techniques is beyond the scope of this article. However, it is possible to envision an overall DSN strategy that uses a variety of these techniques based on the priority of the activity being supported (routine operations versus encounter).

A centralized architecture is extremely sensitive to WAN communication failures. If a communication problem cuts communications between the DSCC's and the NCF, the DSCC's will not be able to complete a command link with the spacecraft or feed back mission data in real time, and will lose contact with the MLC's. The MLC's

will be unable to coordinate activities, including those involving more than one DSCC, and the DSN will not be able to assess overall network performance and status.

In the event of a communications loss, the DSCC's will record incoming mission data for later playback, set up a link controller at each station to perform MLC functions for that station, call in backup operations personnel as required, and use telephone lines for voice and data communications with the NCF. Operations that have to be coordinated between more than one DSCC will be conducted using voice communications (if possible). In the event that no communications exist between the DSCC's and the NCF, and if voice links cannot be set up between DSCC's, then each of the DSCC's will operate individually until communications can be reestablished.

The greatest single threat to the coordinated operations of the DSN is loss of the NCF (a prime scenario involves a large earthquake that affects the Los Angeles basin). Loss of the NCF will result in the loss of the primary interface between the DSN and its various customers.

The proposed architecture supports near-real-time changes to operations due to customer requests. Therefore, in addition to precautions established to ensure operations through a major wide-area communications failure, special precautions must be taken to ensure that customers can still be supported. To preserve this capability, a limited-function backup NCF is proposed. The backup NCF will provide reactive scheduling capability and routing of support data. The backup NCF will also provide NASCOM connectivity in order to continue support to end users not affected by the same event that affected the NCF.

Each DSCC will have archives of spacecraft command sequences, adequate predicts, and other support data generation capability. As long as a voice link can be established between the DSCC's and the backup NCF, uplink support can continue if the necessary sequences are present in the DSCC archives.

C. Data Transport

At the DSCC, the GIS architecture depends on two sub-networks: the DSCC backbone LAN and the M&C LAN (Fig. 24). The backbone LAN is a key architectural feature; it provides connectivity between processes anywhere in the Complex.

The Central Site data transport architecture depends simply on a DSN backbone LAN. The data transport architecture features:

- (1) Total interconnection of all processing and monitor and control resources.
- (2) Standard protocols among networks.
- (3) Optical-fiber medium for the backbone LAN's and a high-speed global WAN network.
- (4) Dual, diverse circuits for the global WAN network.
- (5) A standard network management protocol to enable fault management, configuration management, performance management, and security of the network resources.

1. DSCC Backbone LAN. The DSCC backbone LAN provides connectivity among all the subsystems and network resources (see Fig. 13). At the DSCC, these include such devices as subsystems (e.g., telemetry, tracking, etc.), communication servers to the global WAN, an M&C LAN interface, test equipment, and storage devices.

A high-speed (i.e., 100-Mbps) LAN will enable the backbone to carry all forecasted traffic, with the possible exception of a continuous stream of VLBI quick-look data (the network should be capable of handling short 15-sec bursts). A subsystem buffer is proposed for continuous VLBI data to avoid congestion. A 1-Gbps LAN will remove this constraint for the 2010 time frame. Based on forecasted LAN technology, FDDI is the best candidate for a year 2000 backbone. It has the following desirable features:

- (1) 100-Mbps transmission.
- (2) Fiber medium, to enable later growth to higher data rates and ensure low electromagnetic emissions in the vicinity of sensitive station equipment.
- (3) Dual counter-rotating ring architecture for fault tolerance.

2. Monitor and Control LAN. The M&C LAN carries the following types of data (all generally at low rates): static parameters, such as directives to subsystems; dynamic parameters to processes, including received power, system temperature, weather, power meter, etc.; the health of processes, networks, and reports to the operator workstation; and fault-recovery directives to processes.

The forecasted LAN technology is not much different from what is available today, such as Ethernet and token-passing bus networks that use a coaxial cable medium. A 10-Mbps transmission rate will enable the LAN to carry all data in the forecasted time frame. The major improvements will be in the wide availability of application layer messaging services.

A likely candidate for the M&C LAN is MAP, a modern OSI-based network architecture designed for manufacturing automation. It incorporates MMS in the application layer (Section III.D). Other features include planned GOSIP 3.0 compliance, thus meeting open-systems interconnection requirements, and token-passing bus topology, with predictable performance as nodes are added, and a self-healing token-passing algorithm.

3. Global WAN. The Global WAN provides connectivity among the DSCC's, JPL, and potential international cross-support users. The flow across this network is lopsided—the flow of data inward from the Complexes to JPL is about three orders of magnitude greater than the outward flow.

Data from the DSCC's to JPL include the following:

- (1) Level-zero telemetry data from each DSCC to the level-zero merge function, with the volume and rate depending on number of current missions being supported; e.g., 1–3 Mbps (conservative), 1–10 Mbps (moderate), 1–100 Mbps (aggressive).
- (2) VLBI data, including both a short burst of quick-look data (130 Mbps) and slower playback data (1.5 Mbps).
- (3) Tracking data derived from Doppler and ranging data (<64 kbps).
- (4) Monitor and control data (<64 kbps).

Data from the Central Site to the DSCC's include the following:

- (1) Spacecraft commands (<64 kbps).
- (2) Monitor and control data (<64 kbps).
 - (a) High-level command language from multilink controllers to link controllers.
 - (b) Predicts and other directives for remote control.
 - (c) Database updates for hierarchical planning data.
 - (d) Fault-recovery commands.
 - (e) Test instructions.

Based on performance and bit error rates, optical fiber is an ideal medium for long-haul circuits; however, if disabled it may take a long time to repair. It is recommended that all global WAN circuits be dual, diverse circuits, with the diversity coming from either use of another medium (i.e., a satellite) or another physical fiber route from the DSCC to JPL. These two architectures are compared in

Table 20; it is concluded that a dual, diverse link provides significantly higher availability—an essential feature of the overall GIS architecture.

The following guidelines are suggested:

- (1) The communication server should support at least two diverse DSCC–Central Site circuits.
- (2) Circuit transmission rates should be negotiated on an as-needed basis, with opportunities to increase speeds on demand for certain missions.
- (3) Global WAN protocols should provide intermediate and upper layer OSI services.
- (4) Circuits must be full duplex to enable a path for acknowledgments.
- (5) Network management software should manage the status of international networks and recover automatically.

4. DSN Backbone LAN. The DSN backbone LAN at JPL provides connectivity among the level-zero merge network, NCF, user networks and network resources. It is assumed that users will interface with the DSN through a common router and use common OSI services, such as X.400 electronic mail and FTAM distributed database capabilities.

A 100-Mbps backbone LAN will be adequate to carry expected 1- to 10-Mbps telemetry traffic. A DSCC buffer is proposed for VLBI data to avoid congestion. A 1-Gbps LAN will remove this constraint and will enable the backbone to carry 1- to 100-Mbps telemetry traffic. Based on forecasted LAN technology, FDDI is the best candidate for a year 2000 backbone and has the desirable features described in Section IV.C.

5. Internetwork Architecture. A typical network path from JPL to a DSCC may include the SFOC LAN, DSN LAN, global WAN, DSCC backbone, and M&C LAN. In the DSCC-to-JPL path there may be just as many networks. As discussed in Section III.C, intermediate-level protocols are required to route the packets and ensure reliable delivery to the destination.

Because high bit errors increase the frequency of the acknowledgment process for reliable stream-oriented protocols, long-haul circuit error rates and the data-link layer recovery algorithm significantly affect internetwork throughput. Satellite propagation delays also affect throughput. For example, using TCP over a satellite link with a typical data rate of 1.5 Mbps, 10^{-7} BER (coded), results in a maximum throughput of only about 87 percent; with a

10^{-6} BER, it drops to 40 percent. An immediate solution is to use a data link protocol with the selective retransmission option (Section III.D). When this is done, TCP or TP4 will deliver significantly higher throughput rates.

The emergence of fiber circuits will significantly improve throughput because 10^{-9} error rates are possible and propagation delays from the DSCC's to the Central Site are only about 0.04 sec. The maximum throughput is much closer to 100 percent.

6. Network Management. Each subsystem and communication resource connected to the DSCC and Central Site backbone LAN's will have a network interface capable of being remotely monitored and controlled. The ISO-standard protocol for network management, CMIP, is a key protocol within the proposed GNMP. CMIP will enable the use of commercial off-the-shelf software to manage network faults, configuration, performance, and security.

- (1) Fault management alerts the network administrator at a DSCC or JPL when a fault is detected. It should provide fault isolation; it should accept and act upon error detection notifications, trace faults, and correct faults arising from abnormal operation.
- (2) Configuration management enables the network administrator to exercise control over the configuration of the DSCC LAN, Central Site LAN, or global WAN. It will enable an administrator to close down nodes should a fault occur or workloads change.
- (3) Performance management enables evaluation of the behavior of network and layer-entity resources (e.g., bridges and routers) and measures the effectiveness of communication activities. It can also adjust operating characteristics and generate network utilization reports by monitoring a resource's performance. An extension of this capability will enable monitoring subsystem (e.g., telemetry or tracking) performance.
- (4) Security management provides for protection of network resources by including authorization facilities, access controls, encryption when required, authentication, maintenance, and examination of security logs.

D. Software

It is possible to list some of the desired characteristics of a software architecture. These features are necessary to ensure that the design complies with the stated goals of openness, interoperability, and cost-effectiveness. A precise specification of software architecture requires a larger scale design effort than is undertaken in this study.

- (1) Open. The GIS should comply with national and international standards (when available) for programming languages, operating system interfaces, application environments, user interfaces, etc. This compliance has benefits in many areas, including portability and interoperability.
- (2) Distributed. More computing power per dollar is available in highly distributed configurations than centralized. Furthermore, there are robustness benefits to distribution.
- (3) Parallel. The design must minimize dead time in which a processor fails to service requests because it is blocking on some other incomplete action. The threads model provides a clean method for implementing parallelism.
- (4) Configurable. Insofar as possible, multiple configurations for missions, antennas, etc., should be designed to be selected at run time. The name "table-driven" is often attached to this feature, but it can be considerably more general.
- (5) Operable. Operators should be able to modify their interfaces, as appropriate. This also implies that the design should not be too tightly coupled with any particular user interface.
- (6) Regular. A standard interface syntax, particularly for monitor and control, should be obeyed by all subsystems. This encourages the development of general-purpose tools and software reuse.
- (7) Automated. Operators should not be burdened with routine commands.
- (8) Secure. System security must be incorporated in the design from the beginning to avoid inconvenience and operator frustration.

In this section, a simple end-to-end flow of data is used to illustrate the application of client-servers, a key element of the software architecture. The flow is illustrated in Fig. 25. It may be useful to review the discussion in Section III.C. Responding to a stored sequence of events, the LMC subsystem begins configuration for a tracking pass. It consults the name server (via an RPC) to find the names of all processes required to build a logical series of telemetry processes (i.e., a telemetry string). It then asks a load-managing process to allocate a set of available processors and maps the desired processes onto them. The control process gives the new string assignments a unique name and passes this definition to the name server.

The LMC then starts each remote process (again using an RPC). The remote processes self-configure, using

the name server to find the names of directly connected peers in the string. Each process also maintains its connection to the LMC. The RPC layers and connection-based network protocols provide reliable data transport between processes. In the event of network errors or other problems of short duration (approximately 0.1 sec), data integrity can be assured by protocol layers below application visibility. During normal operation, each process periodically saves its internal state to a monitor system (via an RPC or a distributed file system).

In the event of a longer failure (approximately 10 sec), a multithreaded process upstream of the failure continues to read and process input, and queues output; the writing thread is blocked. As the output queue enlarges, monitor messages are generated to alert both control processes and operators. When the problem clears, the writing thread unblocks and runs at maximum speed until the queue empties.

For a still longer failure (approximately 100 sec), the virtual circuit between the upstream server and its downstream client will expire. As before, input processing and output queuing continue. If the queue grows sufficiently large, it expands onto secondary (disk) storage, perhaps transparently via the host virtual memory system. The RPC layers contain error-recovery code that attempts to reestablish the connection; if successful, recovery takes place as described before.

In an even more extreme case, suppose that one computer in the string fails hard. When the complex controller determines that the problem is a hard system failure, it requests a backup assignment from the load manager and inserts its assignment into the string in place of the failed system. The replacement process negotiates a resynchronization with its peers (using the stored state information) and processing continues. At the end of the pass, the string assignments are dismantled and returned to the available pool.

V. Transition Approach

The transition strategy is to adopt the proposed architecture immediately and require all future implementations to conform. This strategy requires early development of requirements and design, and early demonstration of key elements.

In fiscal year (FY) 1992, a system design team should develop a top-level design that includes:

- (1) Technical and mission requirements.

- (2) Functional design.
- (3) Management approach, including estimated costs and a funding plan.
- (4) Specific standards to be employed (including protocols).
- (5) Identification of key elements for early implementation.

Guidance should be provided for the design of the Block V receiver and new antenna, microwave, and receiver controllers, in order to minimize future rework.

During FY 1992 and 1993, the SPC, GCF, and NOCC upgrades should be evaluated for compatibility with the new architecture. Changes consistent with current requirements should be implemented as soon as it is practical. These changes should be viewed as parallel activities to the main architectural implementation effort and not as essential elements of it.

In FY 1993 and 1994, lower level system requirements should be developed. An end-to-end design should be produced in sufficient detail to allow a specification of the architectural infrastructure necessary for implementation. This infrastructure includes the overall data transport and software architectures described in Section IV.C and IV.D. Development of this infrastructure should begin as early as possible because it requires a significant investment in developing and procuring common hardware and software. This effort will require extensive prototyping of standard protocols and other features of open systems in order to ensure adequate performance margins and compatibility with system requirements. In FY 1996 implementation of the processing architecture should begin.

The architecture should be partitioned into pieces small enough to allow incremental implementation. Two potential implementation strategies are:

- (1) End-to-end data flow strategy. Implement a segment of the GIS architecture, which includes a minimum capability for the entire end-to-end system.
- (2) Facility-based strategy. Implement the system one facility at a time (e.g., DSCC).

Both approaches should employ an incremental implementation strategy. That is, they should be partitioned into smaller segments for final implementation. For example, the end-to-end strategy might be to implement only one mission and one system (e.g., tracking) as a first end-to-end segment. Other systems and missions would follow. This incremental strategy will work only if the investment

in architectural infrastructure has been made as described above. In both approaches, the focus is clearly multimission; choosing one mission at a time for incremental implementation is only recommended to reduce the size of each incremental step.

The end-to-end approach is technically superior and is the recommendation of the study team because it allows the early testing and implementation of several end-to-end concepts that are crucial to the success of the proposed architecture. See Section V.D for further discussion of this approach.

The facility-based strategy has some advantages in terms of managing the scope of the implementation and the number of management interfaces for a given implementation segment. Cost issues may require consideration of this approach at later stages in the effort (e.g., after the top-level design has been completed and the architectural infrastructure has been developed).

Any implementation approach may necessitate that old systems operate for many years into the future. Using the end-to-end approach eliminates many of the requirements for gateways between old and new systems. Every effort will be made to minimize both the use of gateways and the number and duration of dual systems. Old missions and systems will be converted as soon as resources are available.

Specific transition approaches for processing, monitor and control, data transport, and software architectures are discussed in the following sections.

A. Processing

All the DSCC subsystems interface with a backbone LAN and communicate reliably to JPL through the use of standard protocols. Buffering requirements need to be identified for each of the subsystem data flows. The appropriate technology should be identified for each of the buffering requirements (0.1-, 10-, and 100-sec capability to meet the scenario described in Section IV.D), and a test-bed should be used to establish performance limits.

The processing architecture includes several high-speed systems that require special attention, including the telemetry and VLBI subsystems. Performance requirements need to be established for these subsystems concurrent with performance characterization of the networks (discussed in Section V.C).

A high-speed telemetry processor must be designed and prototyped to be integrated with the Block V receiver.

This is a critical requirement for the overall GIS architecture. This design effort should be undertaken as soon as possible so that feedback to the Block V effort will be timely (i.e., early 1992). The subsystem must have the following features:

- (1) Standard backplane interface with the Block V (e.g., Futurebus+) for modularity and flexibility.
- (2) Fault-tolerant architecture independent of the Block V (i.e., how many redundant systems are required? What are acceptable degraded modes? etc.).
- (3) Standard interprocess communications interface with the network database server.
- (4) Interface with the proposed M&C architecture for fault diagnosis and control.
- (5) Standard interprocess interface with the Central Site merge process.

For VLBI quick-look processing, requirements for the volume and rate of data must be established. Work done on high-speed formatting and buffering of telemetry data will apply to this application.

The real-time telemetry interface with the Block V may be developed in parallel with an end-to-end near-real-time architecture. The subsystem must be designed to buffer data during downstream outages as described in the software architecture (Section IV.D). Much of the downstream architecture depends on an infrastructure provided by the data transport architecture (Section IV.C).

B. Monitor and Control

Goals for the proposed Monitor and Control System have been described extensively in Section IV.B. The following section explores the transition to be made from the existing operator-intensive, primarily manual system to one that is automated and has a higher level of service.

The transition approach for monitor and control is complicated because both implementation strategies (see the introduction to Section V), require that the Monitor and Control Subsystem manage processes that reside in both old and new equipment. The changes recommended to achieve the proposed M&C architecture fall into four basic areas: message communications, user interfaces, automation, and centralization.

The following sections describe the transition paths for evolving each of these areas to the new architecture. They are summarized in Fig. 26.

1. Message Communications. There are two required changes in message communications (messaging): addition of new networks and adoption of new data formats to make use of higher level open-systems network application services. Messaging transition states are illustrated in Fig. 26(a).

The subsystem transition states are complicated because it is unlikely that all the subsystems will transition at the same time, therefore several different transition states are possible.

The transition options for monitor and control are (1) the new formats would be converted to provide the same types of information as the old formats; (2) a new software layer would be added to the M&C system that can make full use of both the new and old data formats; or (3) a separate M&C system would be developed for new data formats and the old ones would be phased out as subsystems convert.

Transition Option 1 removes any near-term advantages of moving to the new formats, but is probably simplest to implement; Option 3 may result in serious operations problems by requiring operators to work with two incompatible M&C systems to accomplish an activity.

Option 2 is recommended although it requires modifications to the structure of M&C itself to support each additional new subsystem. All new subsystems should be required to use the new network and the data formats associated with an open systems architecture.

2. User Interfaces. The transition-state diagram for user interfaces, Fig. 26 (b), is relatively straightforward. The transition approach is to integrate new workstation-based GUI's with existing hardware and displays, and phase out the old interfaces as early as possible. Since the user interfaces can be updated without affecting any other areas, work can begin now to make the transition to the goal state.

3. Automation. The transition-state diagram for automation, Fig. 26 (c), is deceptively simple. Implementing the desired level of automation requires not only increases in the types of functions automated (e.g., anomaly detection, schedule execution, replanning, or fault diagnosis) but also increases in the scope of the domain in which these functions operate (e.g., all antennas or just the 70-m antenna; telemetry and ranging, but not commanding or VLBI). As both the functionality and domain grow, the domain in which higher levels of automation are possible increases because of the additional available information.

While automation will eventually exist at all levels of control in the DSN, the area with the highest payoff in the near term is in link operations (multilink operations are currently special-purpose operations only). Link operations are operator-intensive and substantial gains can be made by making link operations more efficient.

4. Centralization. Centralization is characterized by the ability to exert multilink control from the Central Site. Multilink control will depend on two major developments depicted in the state transition diagram, Fig. 26 (d): multilink capability and networks to support data communications. The recommended transition from a distributed control environment to a centralized one is to develop the capability for multilink control at the DSCC's and then transfer the capability to the Central Site.

C. Data Transport

The present data transport configuration involves heavy reliance on JPL-developed protocols. One of the major goals for the year 2000 will be to convert to an open-systems architecture based on standard protocols. The present configuration also depends on satellite communications from the DSCC's to JPL; these should be converted to fiber circuits.

Many other companies and government agencies are presently converting from a vendor-specific, or proprietary, network to an open architecture. Some are basing their new architectures on the Internet suite (TCP/IP and their related protocols). Internet products are presently widely available and provide many internetworking capabilities, although at a reduced level from that which is envisioned for OSI. Installing an end-to-end capability with TCP/IP is a reasonable step—it bases the network on standards. However, the ultimate GIS goal is greater capability with the OSI protocol suite.

Since JPL's starting point is a suite of JPL-unique protocols, a logical protocol choice for the internetwork could be either one of the following: an Internet suite, followed later by a conversion to OSI, or an immediate direct conversion to OSI. The transition will depend on the results of a cost and performance evaluation.

1. Global WAN. The data circuits from the DSCC's to JPL are primarily over satellite paths. The current circuits are 244 kbps (simplex) for telemetry data and 56 kbps (duplex) for all other data. The 244-kbps link uses a JPL-developed data-link protocol that does not provide an acknowledgment (basically a datagram protocol). Higher level application programs at JPL identify later

(from several hours to several days) if any frames are missing, and request retransmission. The 56-kbps link uses a NASCOM data-link protocol. Both circuits encapsulate a JPL-developed network-layer protocol (JPL 890-201) that provides a header address for routing among facilities.

The GCF upgrade is in process, and it will affect the WAN circuits in the following ways:

- (1) An upgrade to 512-kbps aggregate rate is expected.
- (2) The simplex link-level protocol may be upgraded to include a selective retransmission feature to enable reliable transmission from a DSCC communications server to the Central Site communications server.

The year 2000 architecture focuses on the high-speed path from the DSCC's to JPL. This path is planned to be based on optical fiber, with its projected high speed, low error rate, and short delays. One of the most basic transformations to reach the proposed architecture will be the introduction of international optical fiber circuits.

Standard data-link and transport protocols that include error detection and correction are required for reliable (acknowledged) transmission. It is recommended that a data-link protocol with selective retransmission be implemented over current satellite links.

Another major recommendation is to introduce standard protocols at the intermediate layers, such as TCP/IP or OSI TP4 and CLNP. These may be installed simultaneously with the data link protocol. A key prerequisite to the effective use of these protocols is a relatively short time delay in the acknowledgment/retransmission process, which is provided by either optical-fiber circuits or the selective retransmission protocol with satellite communications (Section III.D). The high-speed WAN protocol implementation should be analyzed and thoroughly prototyped before installation. Once the underlying suite is installed, network-wide application-layer services can be implemented, such as file transfer and management, electronic mail, and network management. These services will provide a solid infrastructure for distributed applications.

2. DSCC Backbone LAN. The current DSCC backbone-equivalent LAN is an Ethernet that uses a JPL-developed protocol (890-131) for communication among the subsystems. The JPL protocol provides a user-defined quality of service that ranges from unreliable datagram service with no acknowledgments to reliable service with error detection and recovery.

An SPC upgrade is in process. It will affect the LAN in the following ways:

- (1) A high-speed LAN will be dedicated to telemetry data. This LAN is planned to be an Ethernet using JPL 890-131 protocol with unreliable datagram service. It will transport data from the telemetry processors to the communications server.
- (2) An SPC LAN will carry all other data using reliable service. This LAN is also an Ethernet using JPL 890-131.

The proposed DSCC backbone LAN architecture should be installed at the earliest opportunity. It will be the foundation on which the rest of the year 2000 DSCC architecture is based. FDDI is recommended because it is widely available now and meets the estimated performance requirements. A second recommendation is to introduce TCP/IP or OSI TP4/IP as the means by which to communicate among processes on the backbone. Conversion to standard intermediate-layer protocols will make other protocol services available, such as file transfer and network management.

3. Monitor and Control LAN. At the present time there is no dedicated M&C LAN; monitor and control data are communicated over the DSCC Ethernet with other scientific and engineering data. When the DSCC backbone LAN is installed, it is recommended that a parallel M&C LAN be simultaneously installed. A GOSIP-compliant architecture, such as MAP, is recommended. A MAP LAN similar to the DSS-13 M&C network should be extended to the SPC, and additional stations should systematically be brought on-line.

D. Software

Much of the software required to build the proposed data system is independent of the DSN's particular applications. For example, the low-level capabilities used to transport data (both spacecraft data and monitor and control messages) can be built generically and shared among subsystems. Common solutions are also possible for resource management, exploitation of parallelism, and other areas. These common system components should be prototyped and tested before detailed system design begins—subsystems can benefit from code-sharing only if the common software is available in advance. The following areas represent likely candidates for such an effort. Others are likely to emerge during the development process.

- (1) Resource management. The client-server architecture necessitates some method of locating services whose mapping into a particular processor is deferred until run time. This specific problem is one of a class addressed by directory services (or naming services).

- (2) Data transport. Processes in a distributed application can share common methods of connection management, data transport, and error recovery.
- (3) Threads management. In general, processes in the GIS will require one or more data sources, one or more data sinks, and one or more monitor and control connections. The threads model provides a useful framework for multiplexing tasks without loss of portability.
- (4) Monitor and control. The lower hierarchical levels of monitor and control are content-independent. All subsystems can share common methods for message passing and error handling, and generic message syntax rules.
- (5) Software development environment. As a parallel effort, a common software development environment should be assembled. The results of the prototyping efforts should be a set of standards for development, plus tools for code control, configuration management, documentation, failure reporting, etc.

The end result of this software transition will be two-fold: a software infrastructure that will allow DSN-specific applications to be built with minimal duplication of effort, and a development infrastructure that will assist application programmers in producing consistent, high-quality code in the common framework.

The early efforts should follow a spiral model of software development [2], and apply lessons learned in the design and implementation of the proposed GIS in an environment of rapid changes in the computer industry. For example, POSIX is relatively new (standardized in 1990), OSF DCE is only in a prerelease phase now, and software that implements OSI protocols is only beginning to appear on vendor price lists. A software development model that encourages prototyping will lend itself to more rapid convergence on a robust design.

When the system design team writes the lower level specifications, the quality of the resulting requirements will be enhanced by the experience gained during prototyping, and the subsequent implementations will be considerably streamlined by elimination of duplicated effort.

The system-wide consistency gained from a common software infrastructure will enable rapid and cost-efficient development of end-to-end implementations. This approach, therefore, lends itself more readily to the end-to-end approach discussed at the beginning of Section V. The facility-based approach, in contrast, would necessitate construction of gateways to interconnect old and new compo-

nents. These gateways lie outside the realm of open systems, and may divert resources from the effort to build an information system conforming with emerging standards.

E. Technology Development

This section highlights advanced development and engineering areas to be addressed during the transition period in order to implement the GIS as envisioned for the 2000–2010 era.

1. Advanced Development. The following is a summary of advanced development tasks that are required to achieve the envisioned architecture. In advanced development, JPL's role will be one of innovation, development, prototyping, demonstration, and incorporation into the DSN. Technology areas are summarized in Fig. 27.

Distributed Computing. Telemetry and VLBI subsystems require high-speed processing to decode and format the data. High-speed computing that can keep pace with input data rates is needed. Support data for antenna pointing and transmitter or receiver tuning may require intensive computation. Parallel computation techniques may be applied to these problems by sharing the workload over several processors with an appropriate software architecture. This approach involves coordinated use of multiple, possibly distributed, heterogeneous hardware and software computer systems on different aspects of a single problem. Additional investigation is needed into special cases where parallel processing should be used, and there must be prototyping to determine performance characteristics in those cases.

The availability, and projected availability, and capabilities of stable, portable parallel operating systems need to be continually tracked and evaluated as this area of computing matures. Research on distributed computing that incorporates bus I/O with dedicated telemetry processors is required. Investigation is needed into the use of multiple tightly coupled processors in 100+ MIPS computers for front-end signal processing in the GIS. Tools for custom-designed interprocessor communication software should be developed. Such tools are needed to help minimize any potential for increase in software life-cycle costs that could result from this special-purpose software.

VLSI and Optical Neural Processors, Including Hardware Packaging. Investigation of the application of high-density hardware integration and packaging techniques to DSN systems is required. Applications include board-level computing proposed for telemetry and VLBI real-time processing. The optimal field replaceable unit based on this

technology needs to be established through reliability testing and evaluation. Integrated packaging concepts should be pursued, standards should be established, and multiple sources should be developed.

Research is needed on the application of neural network hardware to GIS automation functions, such as pattern recognition, performance compensation, and device control. One area of possible application is adaptive antenna stabilization and performance compensation for parts that deteriorate.

Data Storage Technology. High-speed buffering using advanced storage technology, such as high-density semiconductors, magnetic memory, and optical storage techniques, needs to be developed. The performance characteristics in GIS-type applications of these technologies need to be evaluated.

Data Management. DBMS's will be required to manage the larger volume of data and provide users with rapid access to associated mission-related data from the monitor and control information base. High-reliability servers will be required for mission data, remote monitor and control functions, and reference data for operator assistance.

Applied research into object-oriented database management systems and text-based database management systems should be pursued. The applicability of these technologies to the DSN should be evaluated through prototyping. Through prototyping, the performance of relational database management systems should be compared with the other varieties in selected DSN applications.

Software Engineering. Improved ability to manage development and implementation of new software-intensive systems is required. Needed work in this area includes development of methodologies, standards, models, and CASE tools; and fabricating an infrastructure to bind these together. Software artifact reuse, software process management, automatic programming tools, and programming environments will benefit from AI support and are particularly compelling. Automation of software testing tools to improve tracking of requirements and specifications is required. Improved software configuration and version management tools are needed. These have the ability to reduce coding time significantly and to improve the reliability of software code.

Standard graphics specification languages that allow portable hypermedia and visualization applications to be written should be investigated. This is an enabling technology for other software engineering disciplines presented in this article.

Protocols. Internetwork communication transparent to users, increased information throughput, and efficient implementation architectures need to be developed. GOSIP-based application prototyping of DSN services is required.

High-Speed WAN. Protocols for reliable high-speed Global WAN circuits from the DSCC's to JPL need to be developed and evaluated in order to improve bandwidth utilization in both directions and minimize costs for the return link.

ISDN Prototypes. Prototyping of ISDN applications for the DSN is required for global voice and facsimile communications. Applications of ISDN to manage temporary data surges should be pursued.

High-Speed LAN. Research on 1-Gbps LAN technology is required, including system integration with parallel and multicomputers, to address system performance issues. This 1-Gbps LAN technology is needed to mitigate the backbone LAN risk identified in Section IV.B and provide technology options for future distributed systems in the DSN.

Network Management. Research and development of network management systems is needed to develop a fault-tolerant data transport architecture. This is an ideal area for prototyping by the DSN because of the number of networks, the immaturity of the commercial software, and the high availability requirements of the GIS. This topic is strongly related to AI and automation technology needs.

AI and Automation. Advanced automation techniques, including AI, neural networks, and fuzzy logic are needed to increase the level of automated monitor and control at the stations, including link and resource management, system diagnostics, fault isolation, and automated recovery. New AI technology is needed in such areas as fault tolerance, knowledge-based systems, constraint-based problem solving, and computer-aided training.

Human Interfaces. Providing uniform GIS user and operator interfaces throughout the GIS should improve M&C operability, user interfaces, and training. Standard graphical (and other) user interfaces are needed and include:

- (1) Multidimensional data in a form that allows people to perceive, amplify, and interpret the information content of the data.
- (2) Data represented in an animation format.
- (3) Transformations and algorithms introduced interactively to explore the effects on the data.

2. Engineering. The following is a summary of engineering capabilities required to achieve the envisioned architecture where JPL's role is to evaluate and adapt commercial technology for application in the GIS.

Engineering Laboratory. An Information System Engineering (ISE) Laboratory is required to evaluate commercial technologies for DSN utilization and prototype the key architectural features of the GIS. The scope should include computing environment, networking technology, automation architectures, expert systems, advanced workstations, and CASE tools. This laboratory will also provide a test-bed to evaluate platform-dependent characteristics of commercial processors; software; and data transport components, including such characteristics as WAN link throughput and LAN throughput, to help bridge the gap between theory and practice. This laboratory will be used in close conjunction with the Compatibility Test Area 21 to simulate the actual DSN operational environment.

Simulation. Simulation will be needed to estimate system loading on critical subsystems and networks. Some work on simulation is now in progress, which should provide the parametric answers needed to optimize the GIS configuration. The GIS will incorporate high-speed LAN's and WAN's. Efforts are needed to evaluate performance with different computer platforms and protocols and with their vendor-specific implementations.

VI. Summary

A unified architecture from the antenna front-end areas to the end user is proposed. Unbuffered real-time data will be processed only at the front end, and a client-server model will be employed at each subsequent processing stage. Open-system standards will be utilized to allow reliable (acknowledged) data transfer and achieve software interoperability.

These changes will substantially improve the interface between the SFOC and the DSN. Data will not be reacquired at SFOC and hence "post-pass replay" (Section II.A) will be eliminated for transmission problems of short duration. It is proposed that the DSN perform level-zero processing and merge multiple records prior to delivering the data to SFOC. Thus, the DSN would be responsible for eliminating artifacts associated with acquisition from the data stream (e.g., duplicate records).

The M&C architecture consists of four hierarchical layers: (1) network control, (2) multilink control, (3) link

control, and (4) subsystem control. The first two layers are assigned to the NCF at JPL and require limited operator support. The last two layers are typically performed at the Complex and are fully automated, requiring only backup and maintenance personnel.

The transition strategy will be to adopt the proposed architecture immediately. All future implementations should conform to the proposed architecture. This strategy requires early development of requirements and design, with a demonstration of key elements. The currently approved implementation program should be integrated with the proposed architecture.

In fiscal year (FY) 1992, a system design team should describe the system-level requirements, develop a functional design, estimate costs, and develop a funding plan. In FY 1993 and 1994, lower level system requirements should be developed. An end-to-end design should be produced in sufficient detail to allow a specification of the architectural infrastructure necessary for implementation. Development of this infrastructure should begin as early as possible. In FY 1995, the new Block V receiver, an-

tenna, and microwave controllers should be implemented with the proposed M&C architecture, and in FY 1996 implementation of the processing architecture should begin.

An end-to-end data flow approach is recommended for partitioning the processing architecture into pieces that are small enough to implement incrementally. This approach allows early testing and implementation of several end-to-end concepts that are crucial to the success of the proposed architecture.

The architectural vision presented here provides guidance for major DSN implementation efforts during the next decade. It also provides focus for flight project mission operations systems developments. The technology forecast and impact assessment provide the basis for an expanded program of advanced development and prototyping activities. Taken together, the technology plan and architectural vision can be a significant contributor to programmatic planning by the Office of Telecommunications and Data Acquisition and MOSO and to future development activities in the Technical Divisions.

Acknowledgments

Valuable contributions to the study were made by the following people: M. Bailey, B. Beckman, J. Breidenthal, R. Burt, M. Carlton, R. Coffin, J. Dundas, S. Eng, J. R. Hall, T. Handley, R. Heuser, R. Holzman, C. Hung, W. Hurd, J. Kahr, K. Kimball, J. Leflang, P. Li, P. Liggett, D. McCornock, M. Parisi, T. Peng, J. Perry, G. Resch, M. Rodrigues, A. Salazar, J. Smith, B. Wilkinson, P. Willis, and H. Younger.

Table 1. Summary of long-range goals.

Attribute	Requirements
Performance	The normal operating mode should be near-real-time processing of telemetry, radio science, and VLBI data; capability should be provided for very high-speed acquisitions, with later off-line processing.
Availability	The GIS should have high availability due to fault-tolerant design that minimizes the impact of hardware, design, operations, and environmental faults; fault isolation should be assisted by computer software to reduce the mean time to repair.
Human operability	The GIS should operate with a minimum number of real-time operational personnel; tools needed to successfully manage an acquisition should be available to an operator at one workstation.
Flexibility and evolvability	The GIS must have the ability to rapidly and easily configure (and reconfigure) for a mission support activity; adequate performance margins should be built in at all levels to ensure flexibility to support unexpected missions.

Table 2. Leading causes of lost telemetry data, 1986-1990.

Subsystems	Hours
DSCC telemetry ^a	1398
Antenna ^b	1038
Radio frequency interference	734
Receiver-exciter	521
Microwave	482
Facility	158
Transmitter	115
DSCC monitor and control	73
Undetermined	65
NOCC support	49
DSCC tracking	45
GCF digital communications	43

^a Major contributing assemblies are the telemetry processing assembly (1185 hr), the subcarrier demodulator assembly (110 hr), and the baseband assembly (42 hr).

^b Major contributing assemblies are the antenna control assembly, antenna servo controller, and the subreflector controller (558 hr); and the antenna pointing assembly (60 hr).

Table 3. Current advocacy DSN mission set.

Mission	Earliest launch date	Likely convolutional code rate	Maximum data rate, kbps	Remarks
Radio astronomy (HEO)	1994	1/2	144 Mbps	Tape and engineering telemetry
Cassini	1995	1/6	156	Science and engineering telemetry
VLBI Space Operations Program (HEO)	1995	1/2	123 Mbps	Tape and engineering telemetry
Comet Rendezvous Asteroid Flyby (CRAF)	1996	1/6	180	Science and engineering telemetry and video
Advance Composition Explorer (E, HEO)	1996	1/2	100	Science and engineering telemetry
Lunar Observer	1997	1/2	500	Science and engineering telemetry and video
Near-Earth Asteroid Rendezvous (D)	1998	1/6	10	Science and engineering telemetry and video
Mars Global Network/Communications Orbiters	1998/01/03	1/6	10	Science and engineering telemetry
Solar Probe	2001	1/6	50	Science and engineering telemetry
Mars Reconnaissance/Communications Orbiter	2001	1/6	100	Science and engineering telemetry and video
Space Infrared Telescope Facility (HEO)	2001	1/2	135	Science and engineering telemetry, IR imagery (4 hr/ day at 1.1 Mbps)
Lunar aeronomy	2002	1/2	100	Science and engineering telemetry
Submillimeter Intermediate Mission (SMIM)	2002	1/2	1.0 Mbps	Spectroscopy data, highly elliptical orbit, 30-sec burst
Pluto flyby	2003	1/6	20	Science and engineering telemetry and video
Mars aeronomy	2005	1/6	100	Science and engineering telemetry
Mars Small Rover	2007	1/6	30	Science and engineering telemetry and video
Mars Rover Satellite Relay (SR)	2009	1/6	200	Science and engineering telemetry and video
Main Belt Asteroid Rendezvous (MBAR)	2011	1/6	100	Science and engineering telemetry and video
Large Deployable Reflector (HEO)	2012	1/2	5.0 Mbps	Real-time imagery
Neptune Orbiter/Probe (O/P)	2015	1/6	100	Science and engineering telemetry and video
Mars Reconnaissance/Communications Orbiters	-	1/6	100	Science and engineering telemetry and video
Thousand AU or Stellar Probe	2020	1/6	0.1	Science and engineering telemetry
Total missions to 2020	22			

Note: HEO = high-Earth orbiter; D = Discovery series small spacecraft; E = Explorer series

Table 4. Views of future telemetry requirements.

View	Telemetry data range, Mbps	Description
Conservative	1-3	End-to-end real-time flow, with limited backup mission capability, no manned-mission support, such as SEI, etc.
Moderate	1-10	End-to-end real-time flow with capability for Mars SEI, added backup capability, etc.
Aggressive	1-100	End-to-end real-time flow that can support lunar and Mars Space Exploration Initiative, including additional backup capability, etc., over and above the moderate capability

Table 5. Ground Information System end users.

End-user interfaces	To/from Ground Information System	Data type	Content
Flight Project Mission Control Team	To/from From	Voice Data block	Real-time monitor and control DSN/SFOC configuration, performance, and status
Flight Project Planning and Sequencing Team	To	Data packet	Spacecraft commands
	To From	Project sequence of events Data file	Project sequence of events DSN system status, events, and fill tape
Flight Project Navigation Team	From	Data file	Radio metric data (range, Doppler, VLBI)
	To	Data file	Spacecraft trajectory
Flight Project Spacecraft Team	From	Data packets (virtual channel)	Spacecraft engineering data
Flight Project Principal Investigators	From	Data packet (virtual channel)	Instrument measurement data
	From	Data packet (virtual channel)	Instrument engineering data
	From	Tape-recorded data packet (virtual channel, high rate)	Instrument measurement data
Flight Project Radio Science Team	From	Data packet (digitized baseband)	Digitized spacecraft RF spectrum
	To	Data file	Trajectory information, planetary atmospheric model
Non-DSN antennas (cross support through NASA Communications Network) (e.g., TDRSS or Galileo German Space Operations Center)	From	Data packet	Spacecraft commands
	To	Data packet	Spacecraft telemetry (spacecraft engineering and instrument measurement data)
Non-DSN antennas (arraying) (e.g., Goldstone array with Very Large Array)	To	Base band signal, analog or digitized (real time)	Arraying of remote antennas with DSN
	To	Telemetry symbol, digital tape	
Planetary Data System	From	Data file	Archive of flight project data (including ancillary data)

Table 6. Summary of standard open programming languages.

Language	Standard	Comments
Ada	ANSI/MIL-STD (1815A-1983)	High-level language
C	ANSI (X3.150-1989)	Middle-level language for systems programming and real-time applications; also adopted by International Organization for Standardization
C++	ANSI standardization effort under way	Object-oriented language based on C
Fortran	ANSI 90	High-level language for numerical applications; new standard replaces Fortran 77 (ANSI X.39-1978)
Common LISP	X3J13	Major computer science research language; draft standard in public review
SCHEME	IEEE P1178	Used primarily for extension languages, computer graphics, and research

Table 7. Common-carrier circuit hierarchies.

Level	Current Asynchronous Digital Hierarchy, (CCITT G.702), Mbps		
	North America	Europe	Japan
1	1.544 (DS-1)	2.048	1.544
2	6.312 (DS-2)	8.448	6.312
3	44.736 (DS-3)	34.368	32.064
4	—	139.264	97.728

Synchronous Digital Hierarchy (CCITT G.707, G.708, and G.709), Mbps	
CCITT designation level	Data rate, Mbps
OC-1	51.84
OC-3	155.52
OC-9	466.56
OC-12	622.08
OC-18	933.12
OC-24	1244.16
OC-36	1866.24
OC-48	2488.32

Table 8. Satellite and optical fiber long-haul circuits.

Media	Typical rates	Typical bit error rate (uncoded)	One-way delay, sec	Comments
Satellite	9.6 kbps, 56 kbps, 64 kbps, 1.5 Mbps	10^{-4}	0.25	Uncoded BER typically 10^{-4} ; coded 10^{-7}
Fiber	Planned SDH: 51 Mbps, 155 Mbps, 622 Mbps	10^{-9}	0.04	Service depends on availability of undersea circuits and international tariffs

Table 9. ISDN channels.

Channel	Speed, kbps
D	16 or 64
B	64
H	384, 1536, and 1920

Table 10. ISDN services.

ISDN service	Number of D channels	Number of B channels
Basic	1 (at 16 kbps)	2 (at 64 kbps)
Primary	1 (at 64 kbps)	23 (at 64 kbps U.S., Canada, and Japan); 30 (at 64 kbps Europe)

Table 11. Processor-board bus standards.

Bus	Standard	Speed, Mbps	Comments
Fiber Distributed Data Interface (FDDI)	ANSI X3T9.5	100	Optical fiber
VME bus	IEEE 1014-1987	320	32-bit standard, 64-bit (VME 64) available
Futurebus+	IEEE 896	25,000	32-bit bus available now; up to 256-bit parallel
High-performance peripheral interface (HIPPI)	Under consideration as an ANSI standard by the ANSI X3T9.3 Committee	800	

Table 12. Applicable LAN's and their characteristics.

LAN	Standard	Network topology	Speed, Mbps	Comments
Carrier-sense multiple access with collision detection (CSMA/CD)	International Organization for Standardization (ISO) 8802-3	Bus	10	Widely available as Ethernet
Token-passing bus	ISO 8802-4	Bus	5, 10	Process control applications
Token ring	ISO 8802-5	Ring	4, 16	IBM PC applications
FDDI	ANSI X3T9.5	Ring	100	Fault tolerant; optical fiber; 2 km (multimode); 40 km (single mode)
FDDI-II	ANSI, Planned 1993	Ring	100	Circuit-switched traffic multiplexed with packet traffic

Table 13. Comparison of Internetwork protocols.

Open Systems Interconnection Protocol	ISO ^a standard	Internet protocol	Standard	Function
Connectionless Network Protocol	ISO 8473	Internet	MIL-STD-1777	Route data across different subnets
Transport Protocol Class 4	ISO 8073	Transmission Control	MIL-STD-1778	Provide reliable process-to-process data flow

^a International Organization for Standardization (ISO).

Table 14. Early processing architecture candidates.

Candidates	Description
Buffering at the complex	Telemetry data are buffered at the Deep Space Communications Complex (DSCC) and delivered to a central facility for level-zero processing (no reacquisition)
Level-zero processing (LZP)	LZP is done at the DSCC and delivered to a central facility, or in some cases directly to the principal investigator; this option has been the focus of a separate study (ongoing)
Central processing (bent pipe)	As much processing (as reasonable) is done at a central facility
Dedicated strings	Dedicated strings of equipment are assigned to each antenna with little need for reconfiguration

Table 15. Processing requirements.

Process	Rate per process	Maximum number of concurrent processes	
		Year 2000	Year 2010
Telemetry (conservative)	1-3 Mbps	7	10
Telemetry (moderate)	1-10 Mbps	7	10
Telemetry (aggressive)	1-100 Mbps	7	10
Tracking	10 kbps	3	3
Command	10 kbps	3	3
Radio science	1.2 Mbps	3	3
VLBI (quick-look)	130 Mbps	1	1
VLBI (playback)	1.5 Mbps	1	1

Table 16. Telemetry processing requirements per subsystem.

Mission support forecast	Upper limit in spacecraft information rate, Mbps	Typical convolutional code rate	Block V receiver-telemetry data path, Mbps
Conservative	3	1/6	165
Moderate	10	1/6	550
Aggressive	100	1/2 ^a	1830

^a assuming support for lunar SEI, etc.

Table 17. Dedicated telemetry subsystem performance.

	1980s	1990s	2000s
Backplane bus capability	320 Mbps (VMEbus)	1 Gbps (Futurebus+)	25 Gbps (Futurebus+ evolution)
Individual spacecraft information rate	5.8 Mbps	18 Mbps	454 Mbps

Table 18. Distributed telemetry subsystem performance.

	1980s	1990s	2000s
Local Area Network capability	10 Mbps (Ethernet)	100 Mbps (FDDI)	1 Gbps (future technology)
Aggregate spacecraft information rate	180 kbps	1.8 Mbps	18 Mbps

Table 19. Evaluation of the major telemetry architecture options.

Criteria	Dedicated telemetry architecture	Switched telemetry architecture	Distributed telemetry architecture
Performance, growth potential	Easily accommodates growth; meets 2000-2010 requirements up to 450-Mbps backup support	Future growth depends on evolution of serial input and output (I/O) technology	Future growth depends on evolution of Local area network (LAN) technology
Life-cycle cost	Major interprocess communication costs affect life-cycle costs; fault-tolerant software has impact on life-cycle costs	Reduces number of telemetry processors, but requires a high-availability switch	OSI systems generally have lower life-cycle costs.
Operability	May use a highly specific operating system and require specially trained programmers	Simple programming	Operating system interface can be "open," e.g., POSIX; large pool of programmers
Flexibility, evolvability	Hardware communication system is closed (i.e., non-standard); may be difficult to add new functionality	I/O protocols may be standard, easy to add new processors	System is "open," with standard protocols
Availability, reliability, maintainability	Each Block V receiver has access to its own telemetry processor; fault-tolerant architecture: redundant processors	Mechanical matrix switches have generally low reliability; will require new technology; architecture handles telemetry faults by switching to redundant processor	Each Block V receiver has access to any telemetry processor; fault-tolerant architecture: built into LAN architecture
Technical risk	Monitor and Control needs careful design (low design risk)	High-availability switch is moderate-high risk	Availability of 1-Gbps LAN is moderate risk

Table 20. Comparison of alternative WAN architectures.

Criteria	Single circuit	Dual, diverse circuits
Performance and growth potential	No impact, probably would be fiber	Slight improvement: second link, not necessarily the same data rate, will be available for special support
Life-cycle cost	Least expensive	More expensive (two circuits)
Operability	Easiest to manage	Requires semiautomatic fault management to select alternate circuit
Flexibility and evolvability	Straightforward to negotiate with carrier	Straightforward to negotiate with carrier
Availability, reliability, and maintainability	Generally high reliability (fiber); however, when it goes down, it goes down for hours; recourse is tape backup	Highly reliable and available because of route diversity; when primary goes down, alternate circuit backs it up
Technical risk	International circuits may take longer to reestablish than domestic circuits	Small technology risk in network management area: need to know health of both links all the time

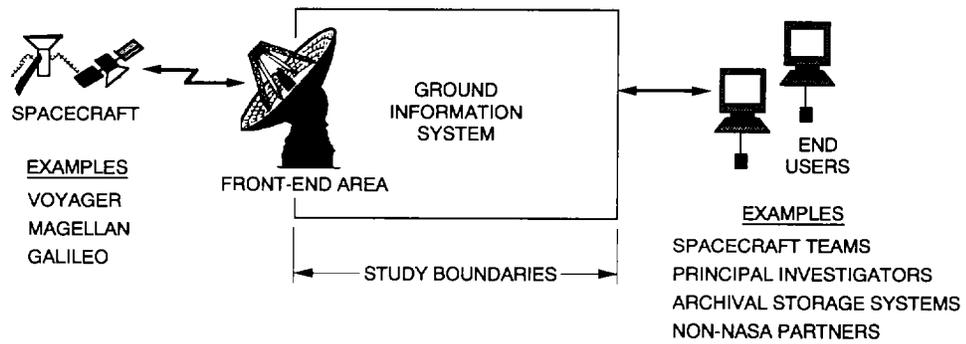


Fig. 1. Scope of architecture study.

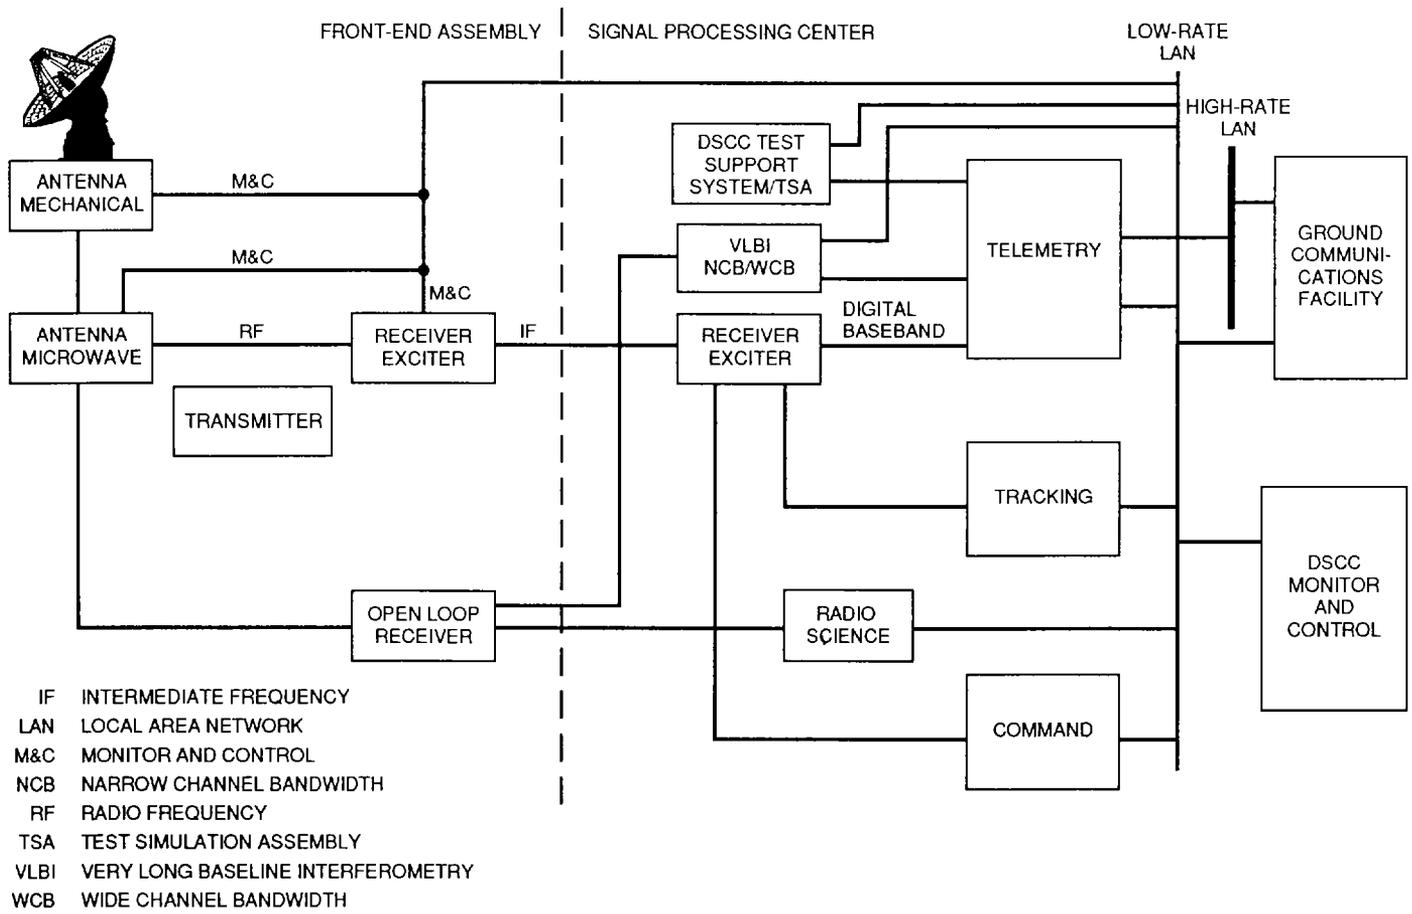


Fig. 3. Current Deep Space Station configuration.

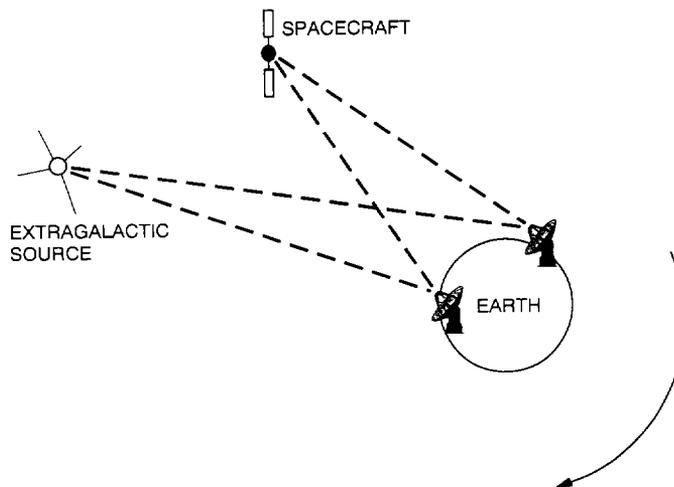


Fig. 4. Very long baseline interferometry.

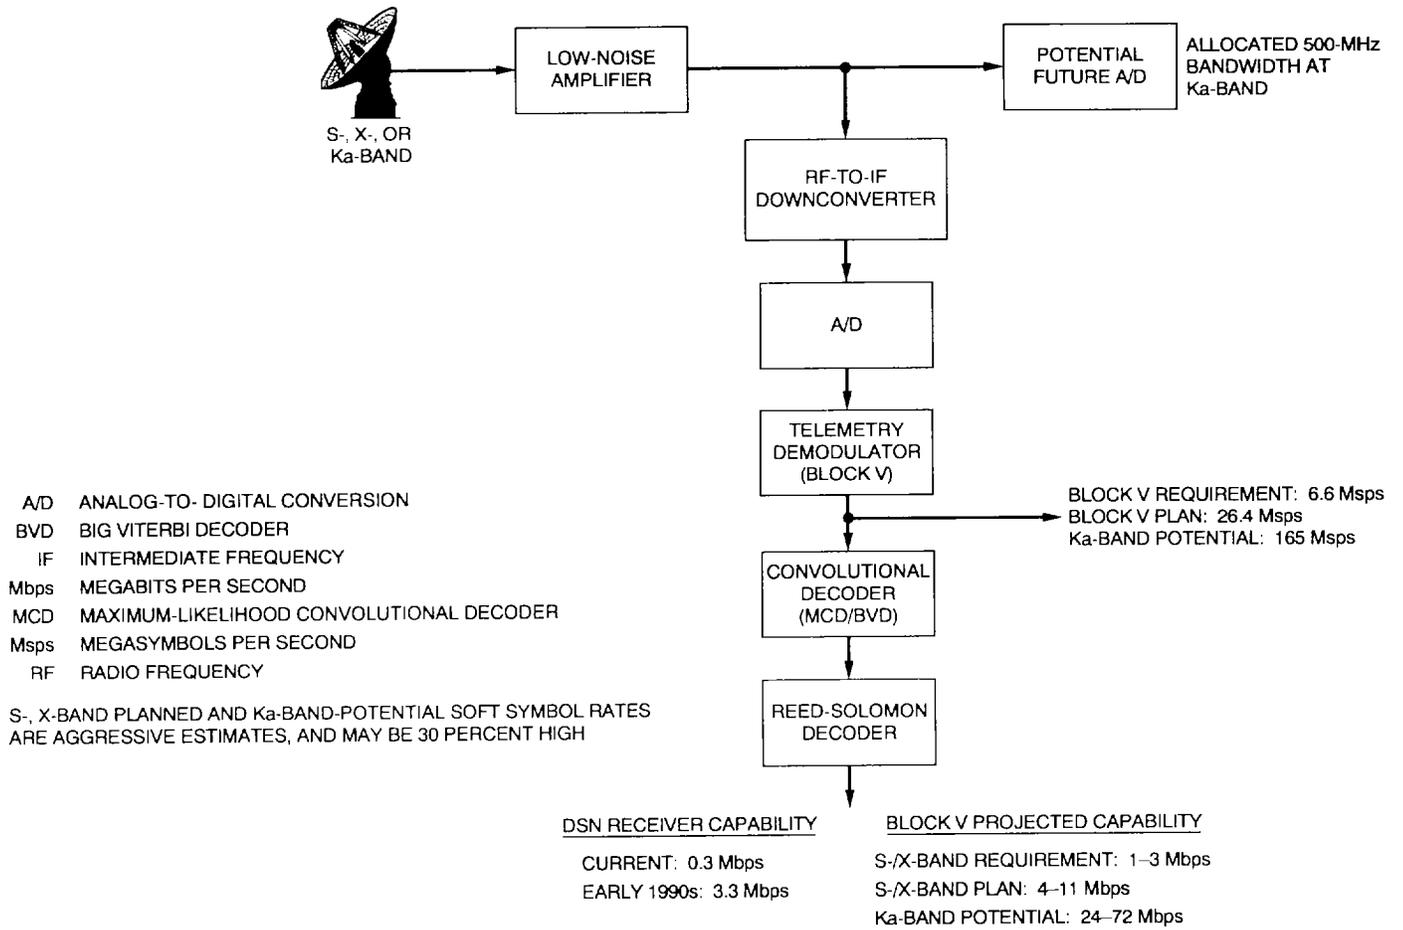


Fig. 5. Block V receiver and telemetry decoder.

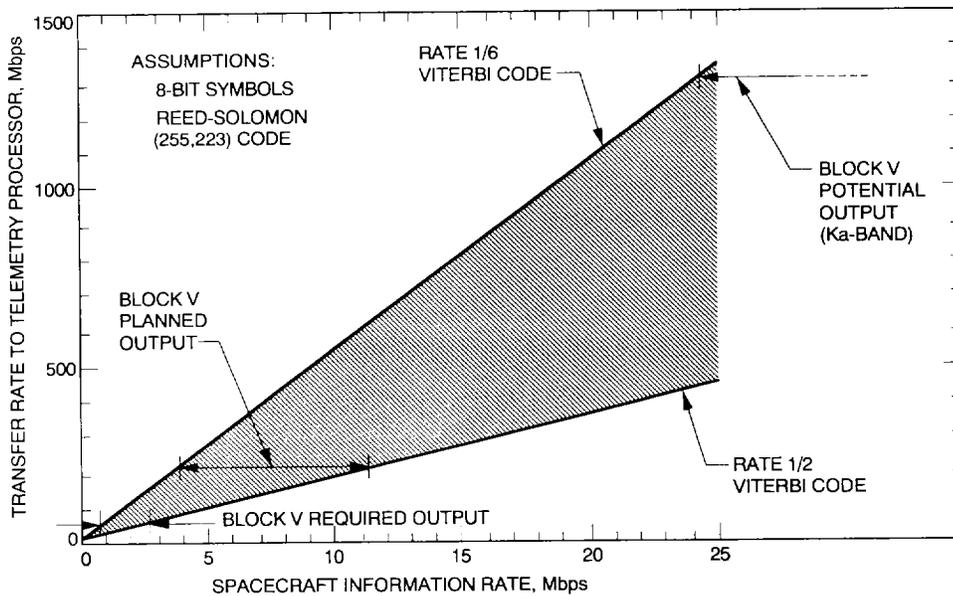


Fig. 6. Block V receiver transfer rate characteristics.

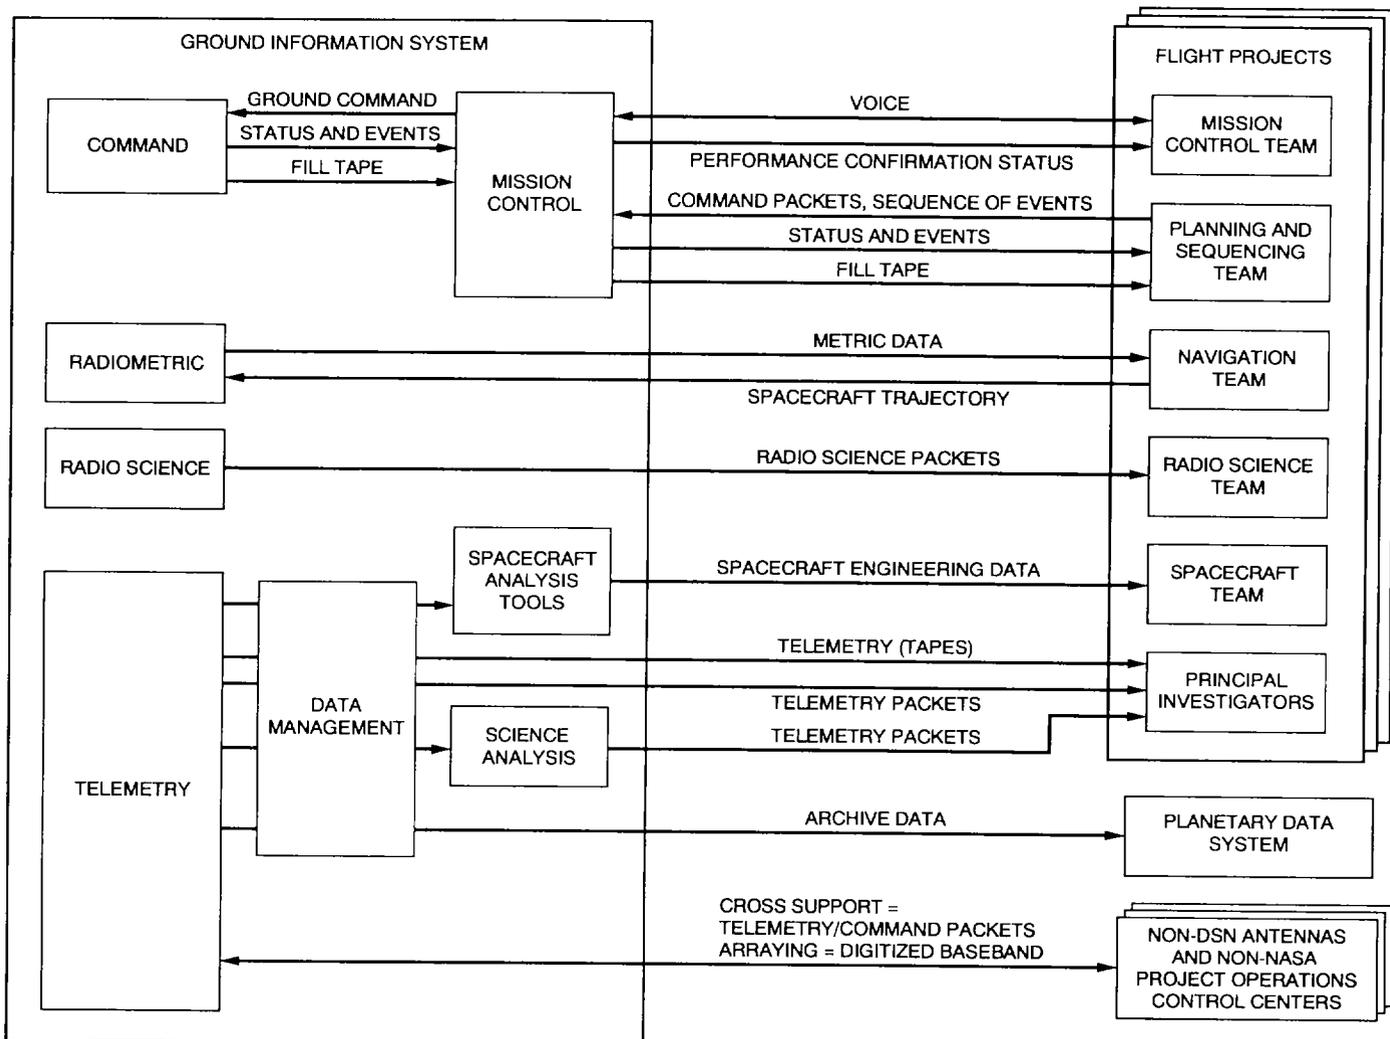


Fig. 7. External users of the Ground Information System.

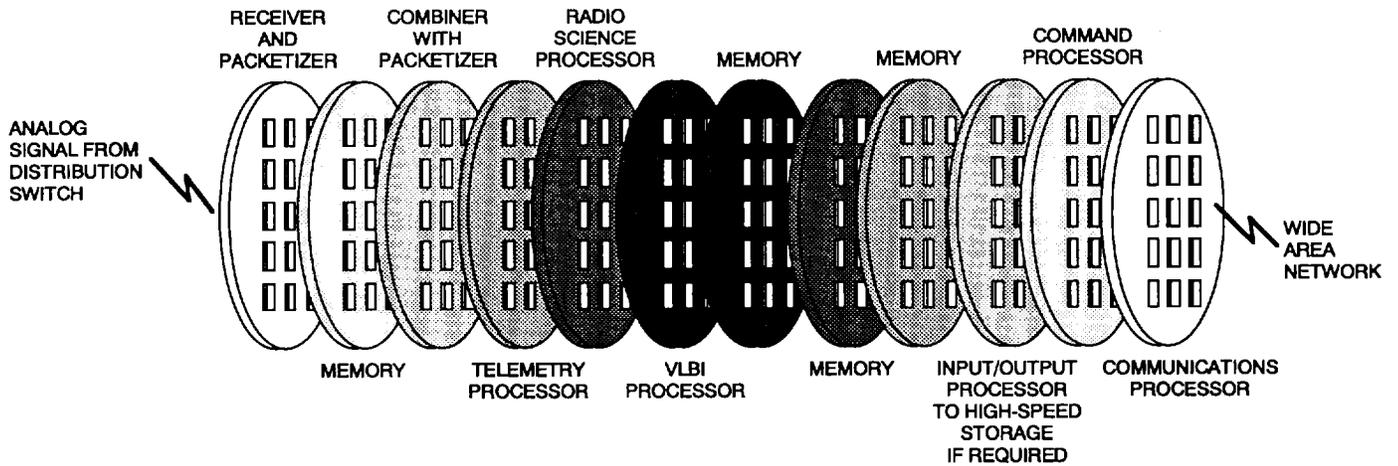


Fig. 8. Conceptual packaging of DSN processors.

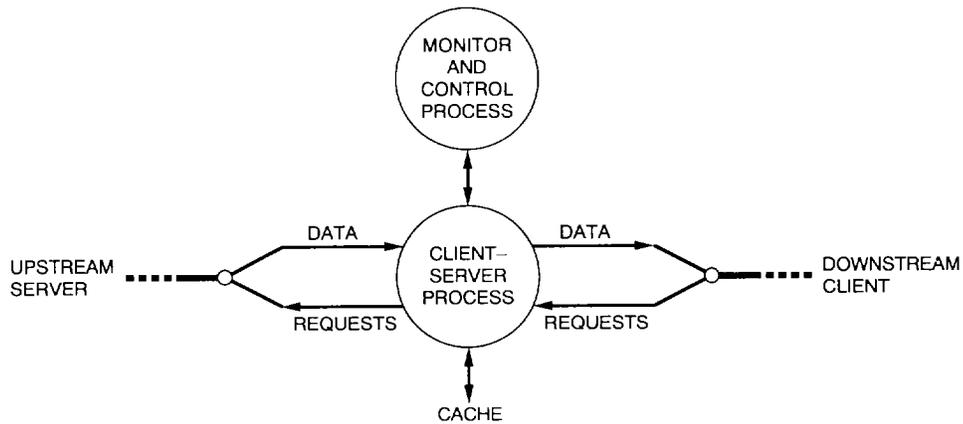


Fig. 9. Client-server model.

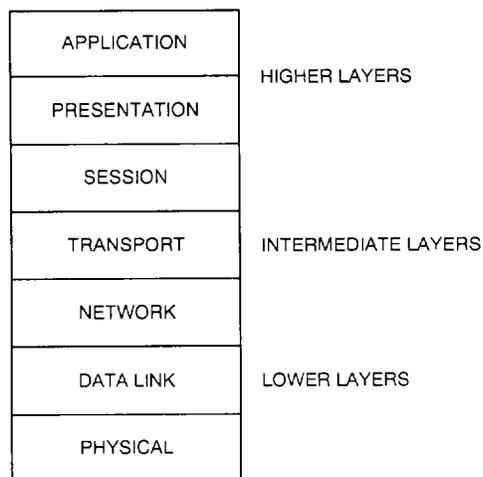


Fig. 10. Open-Systems Interconnection reference model.

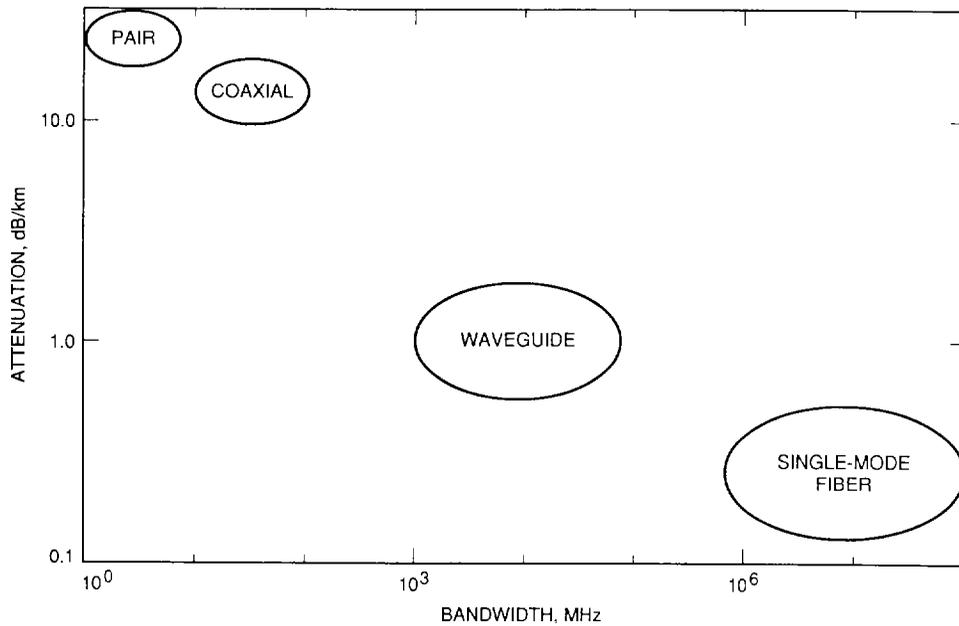


Fig. 11. Transmission properties of various media [3].

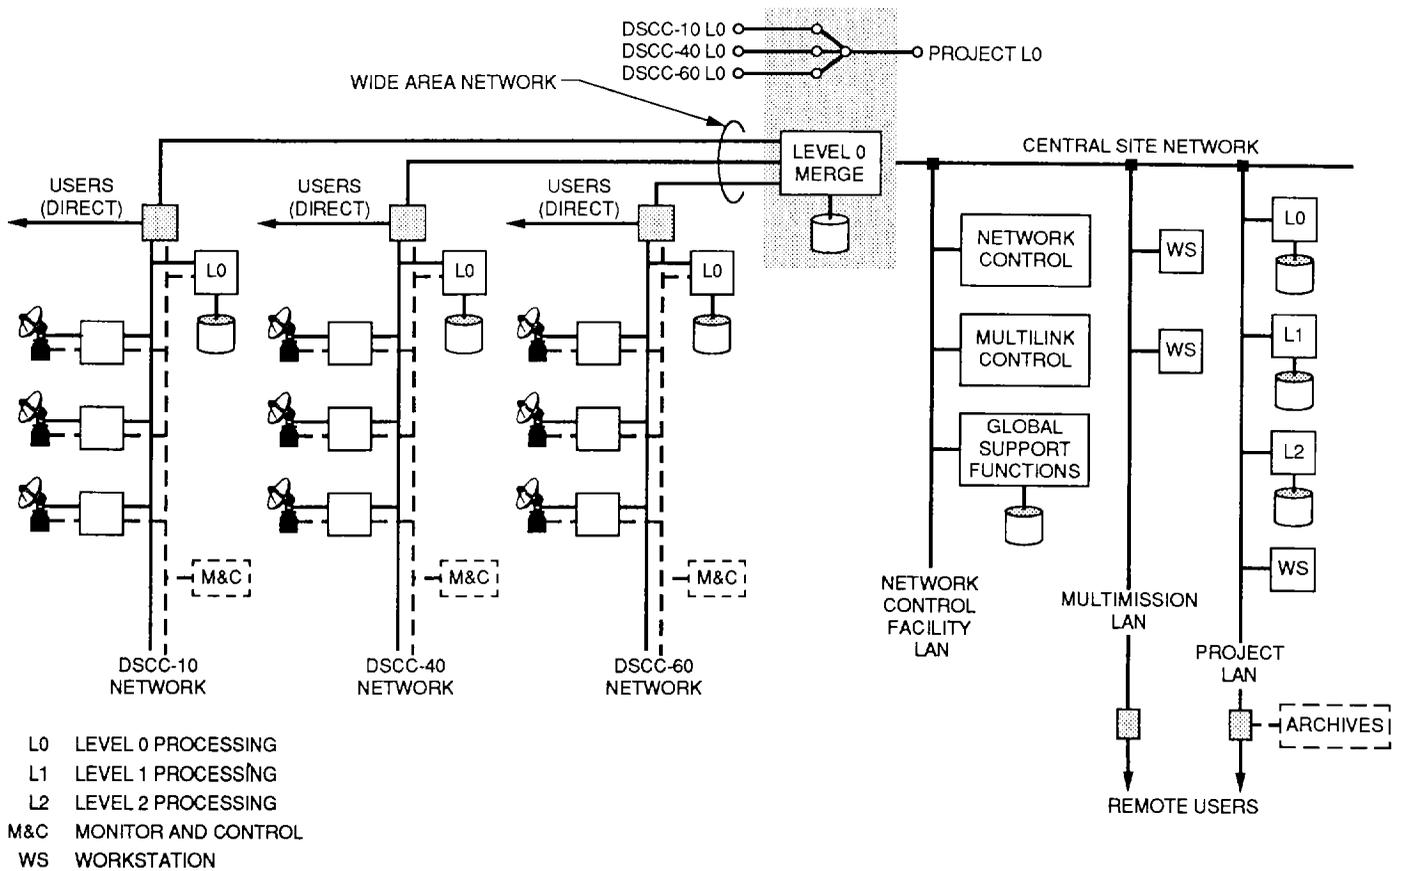


Fig. 12. Unified processing architecture.

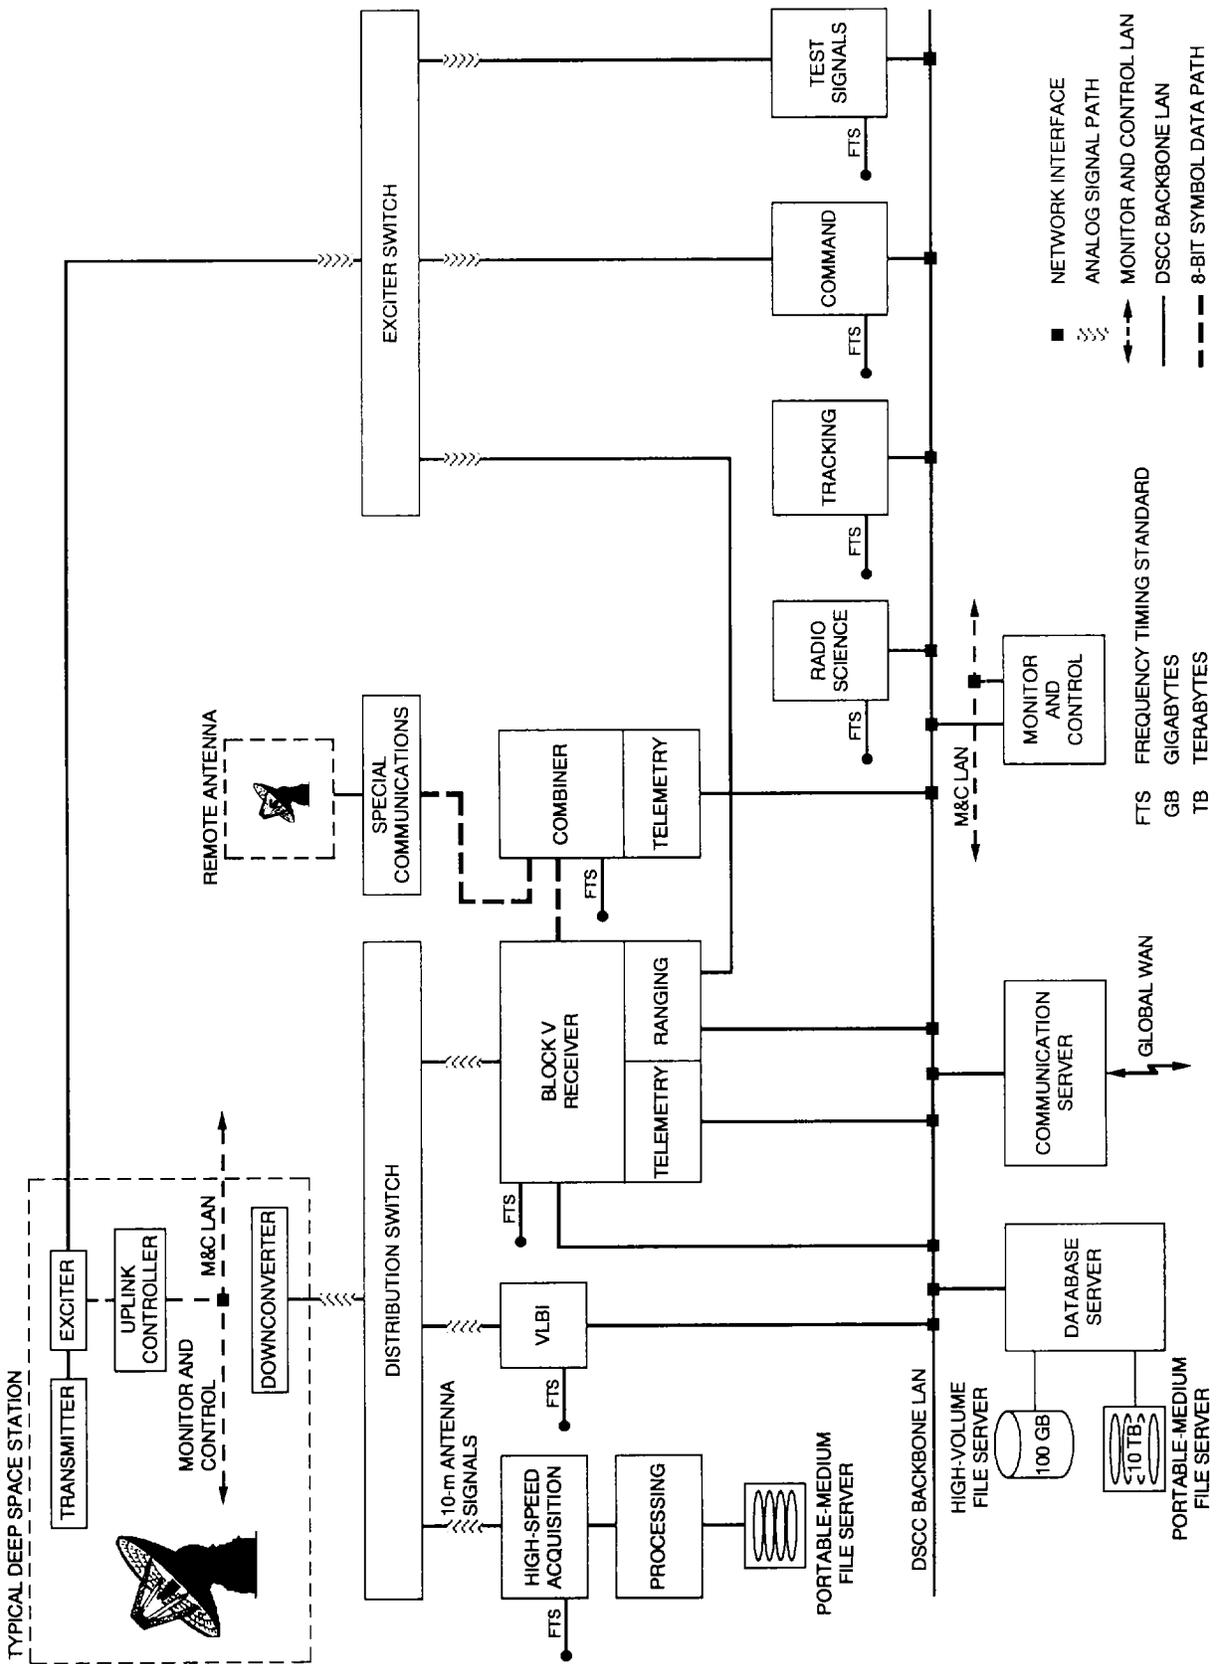


Fig. 13. DSCC processing architecture.

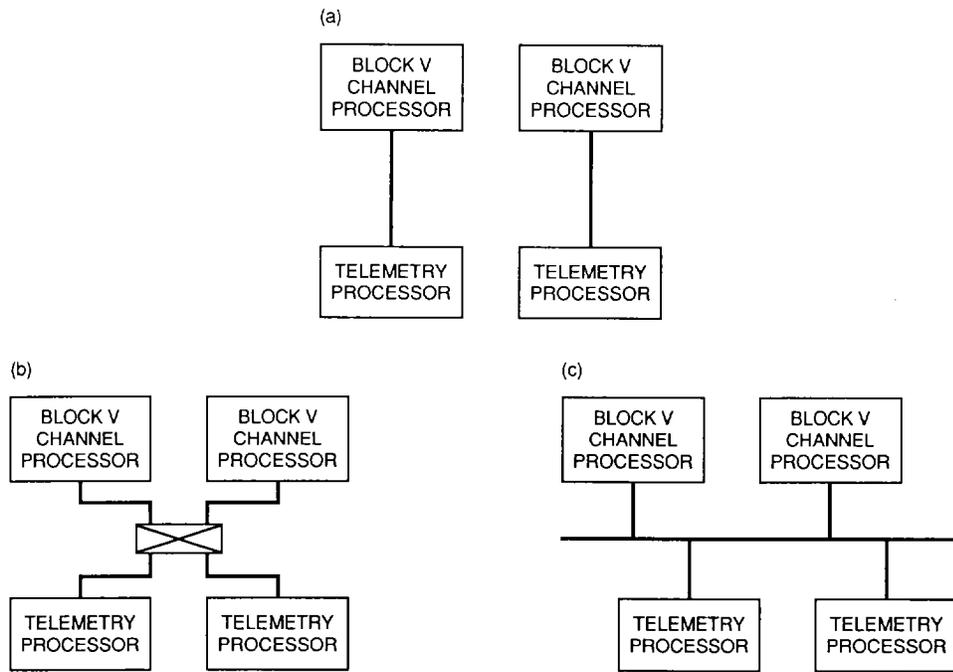


Fig. 14. Telemetry architecture candidates: (a) dedicated; (b) switched; and (c) distributed.

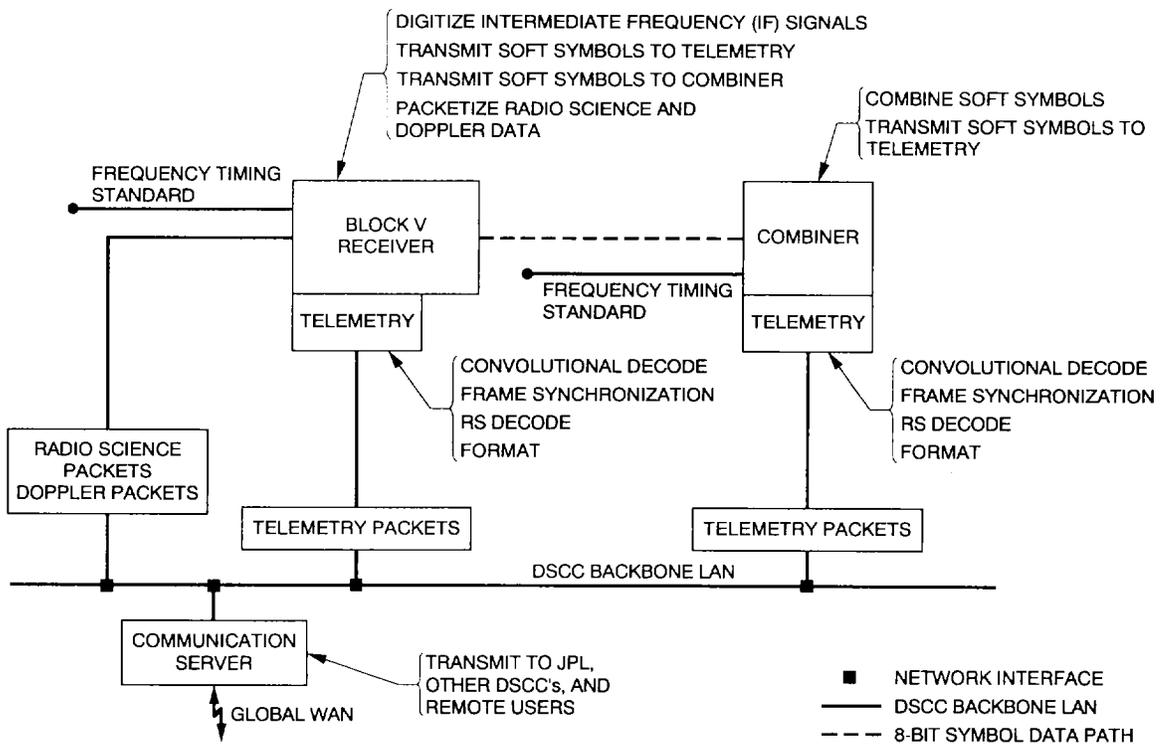


Fig. 15. Dedicated telemetry architecture.

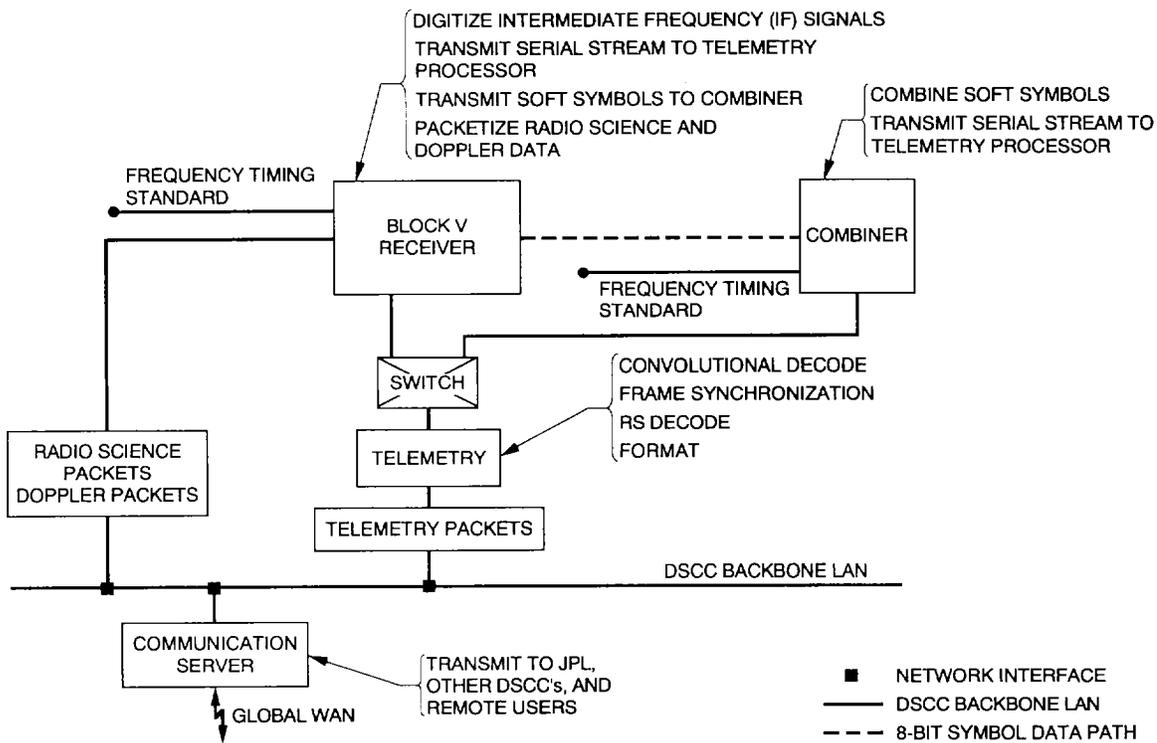


Fig. 16. Switched telemetry architecture.

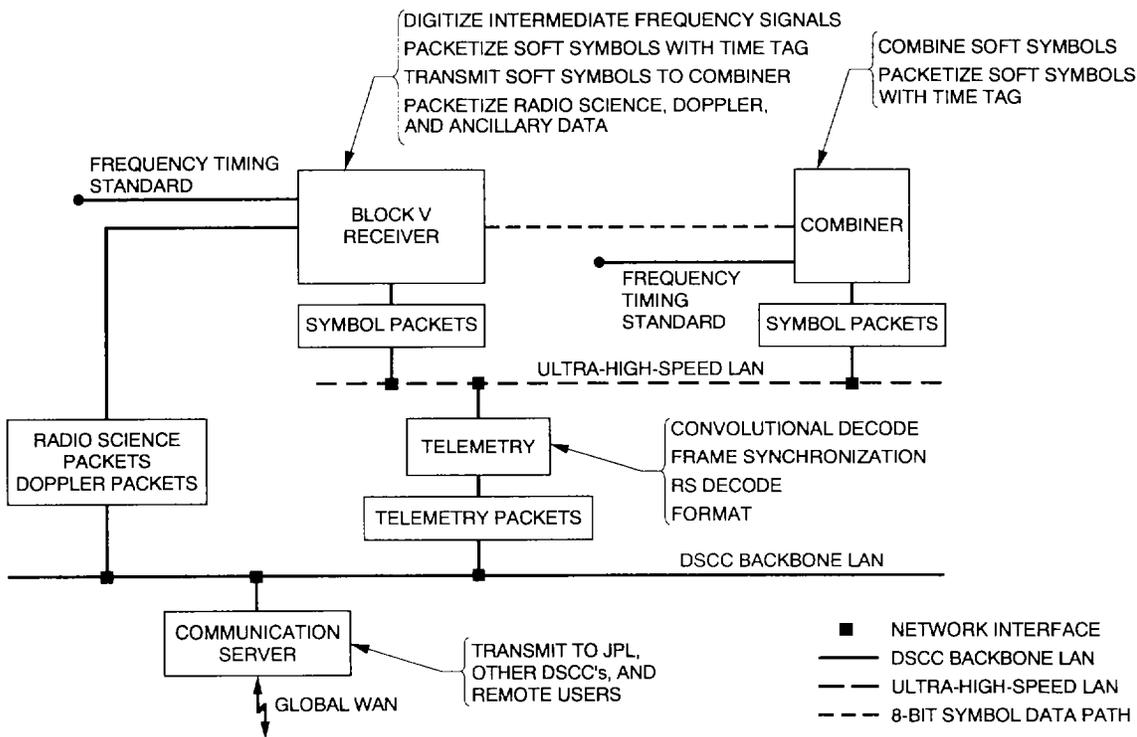


Fig. 17. Distributed telemetry architecture.

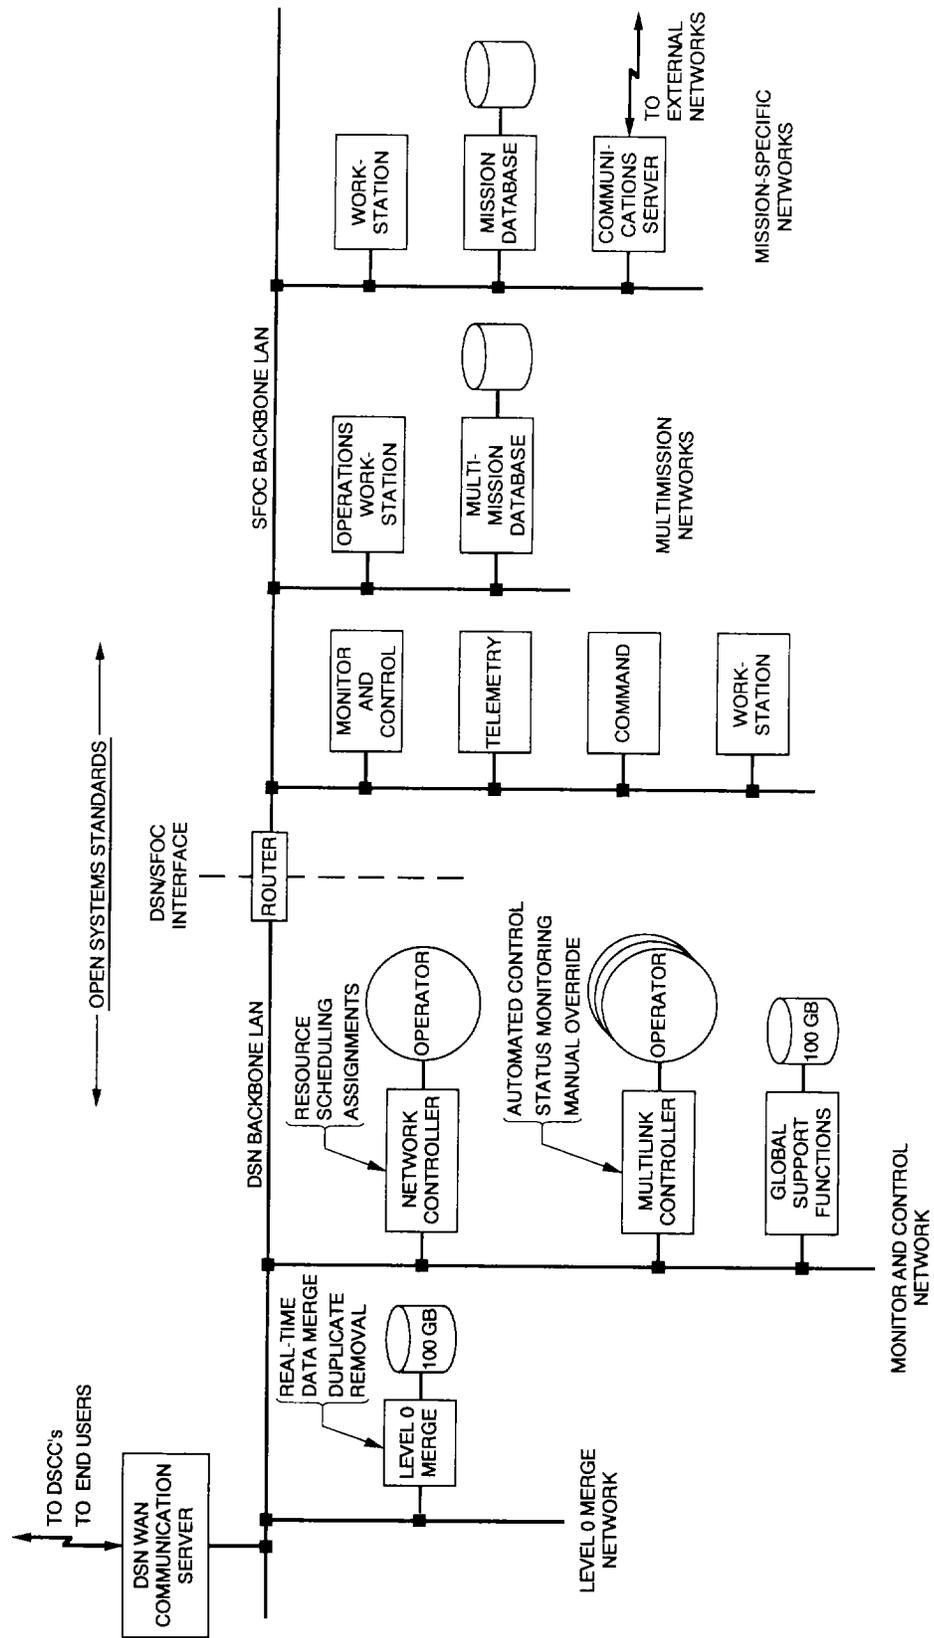


Fig. 18. Central Site architecture.

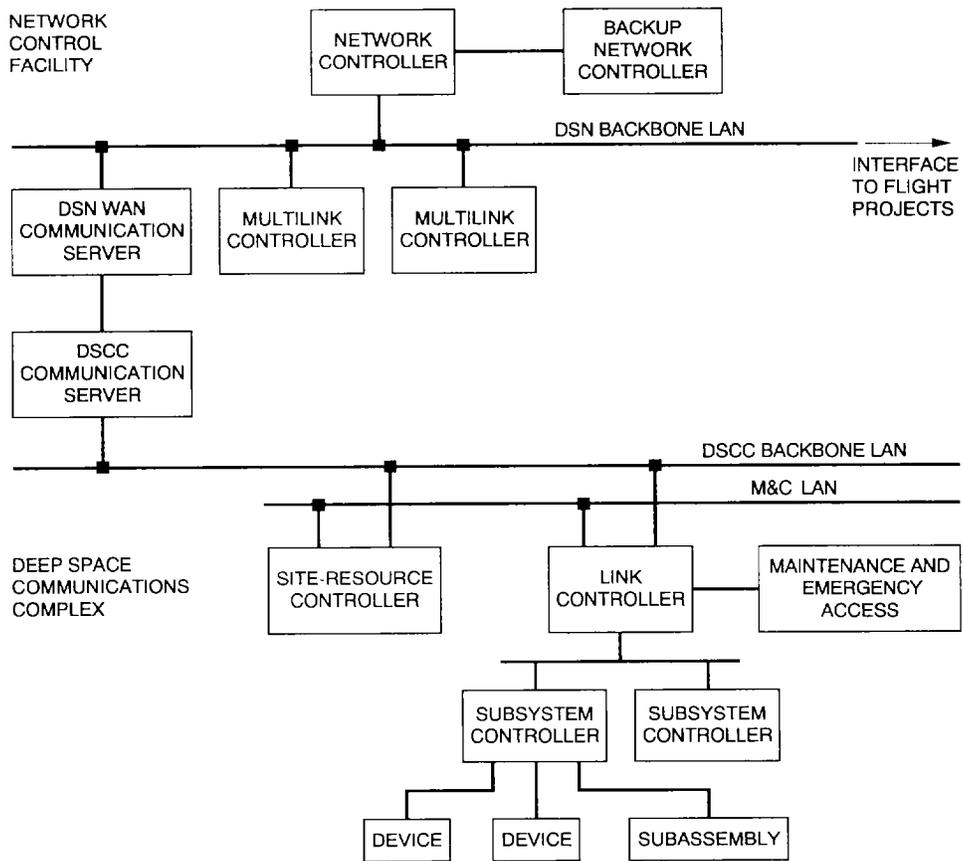


Fig. 19. Proposed Monitor and Control architecture.

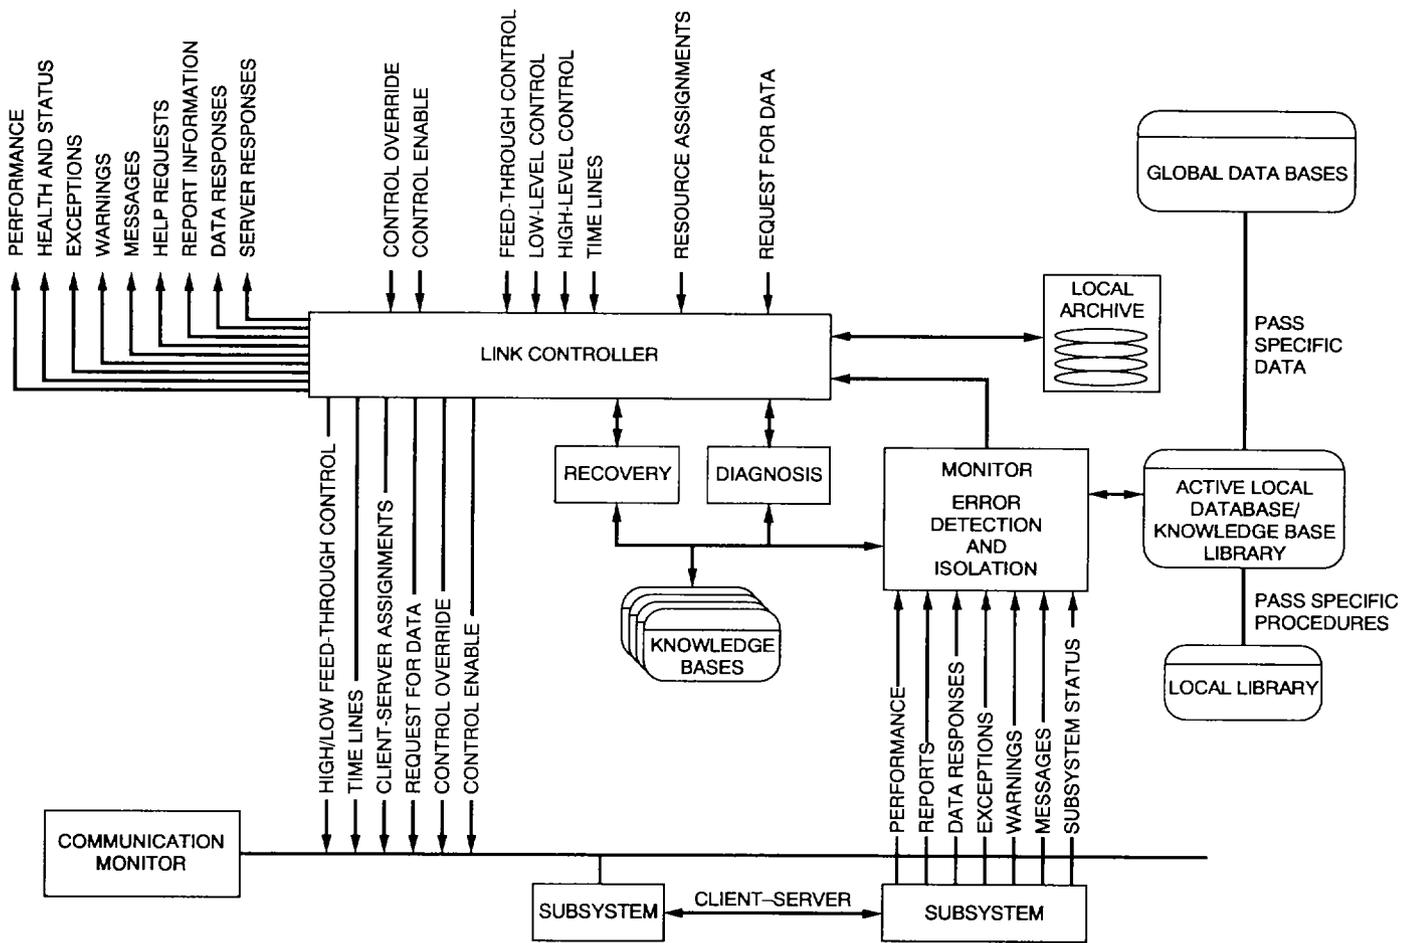


Fig. 20. A generic model for all control layers using the link controller as an example.

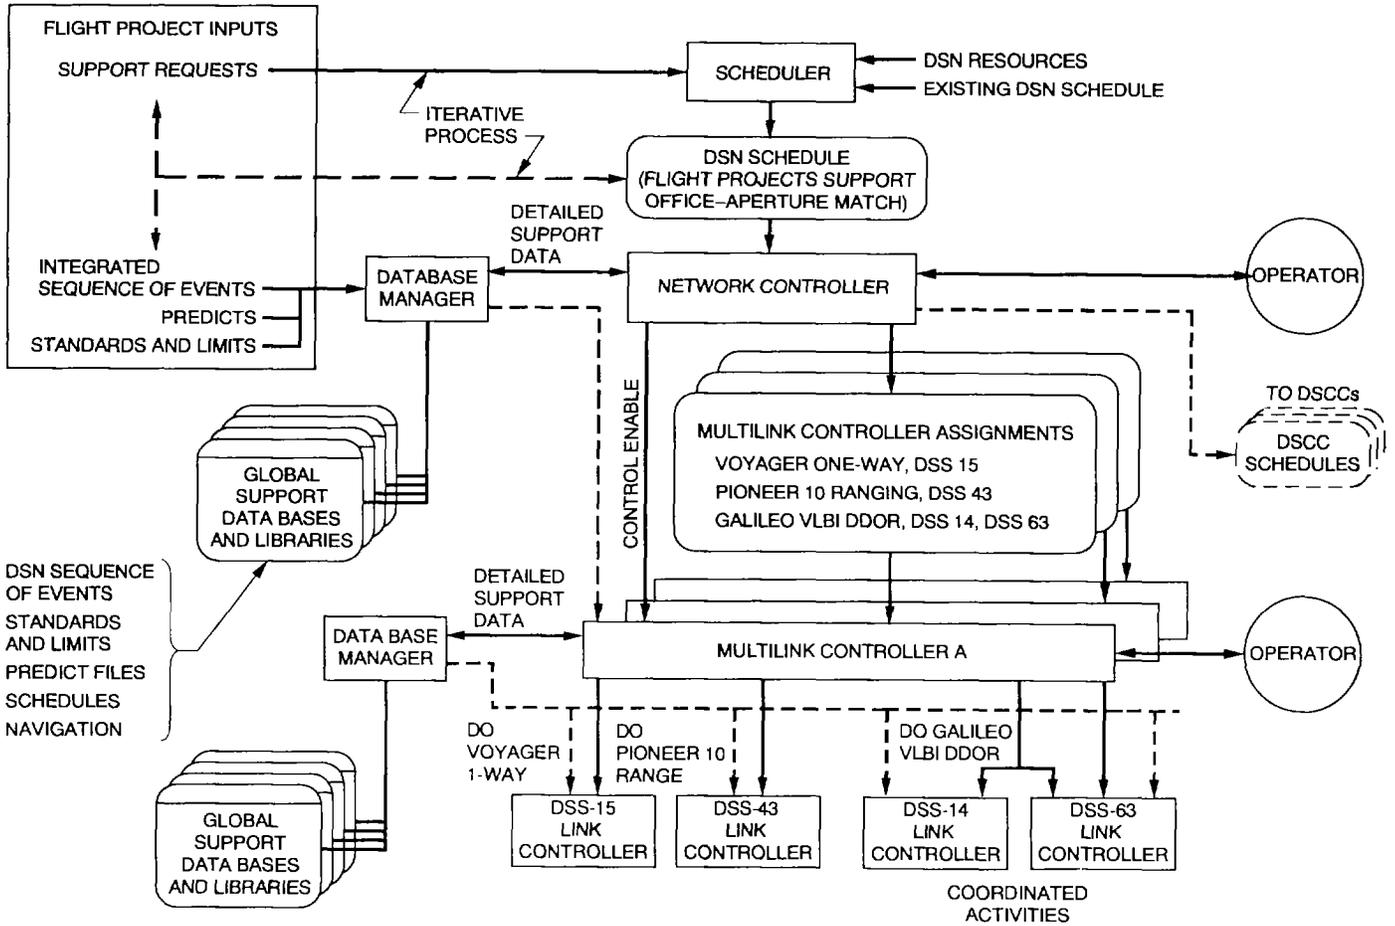


Fig. 21. Operations concept.

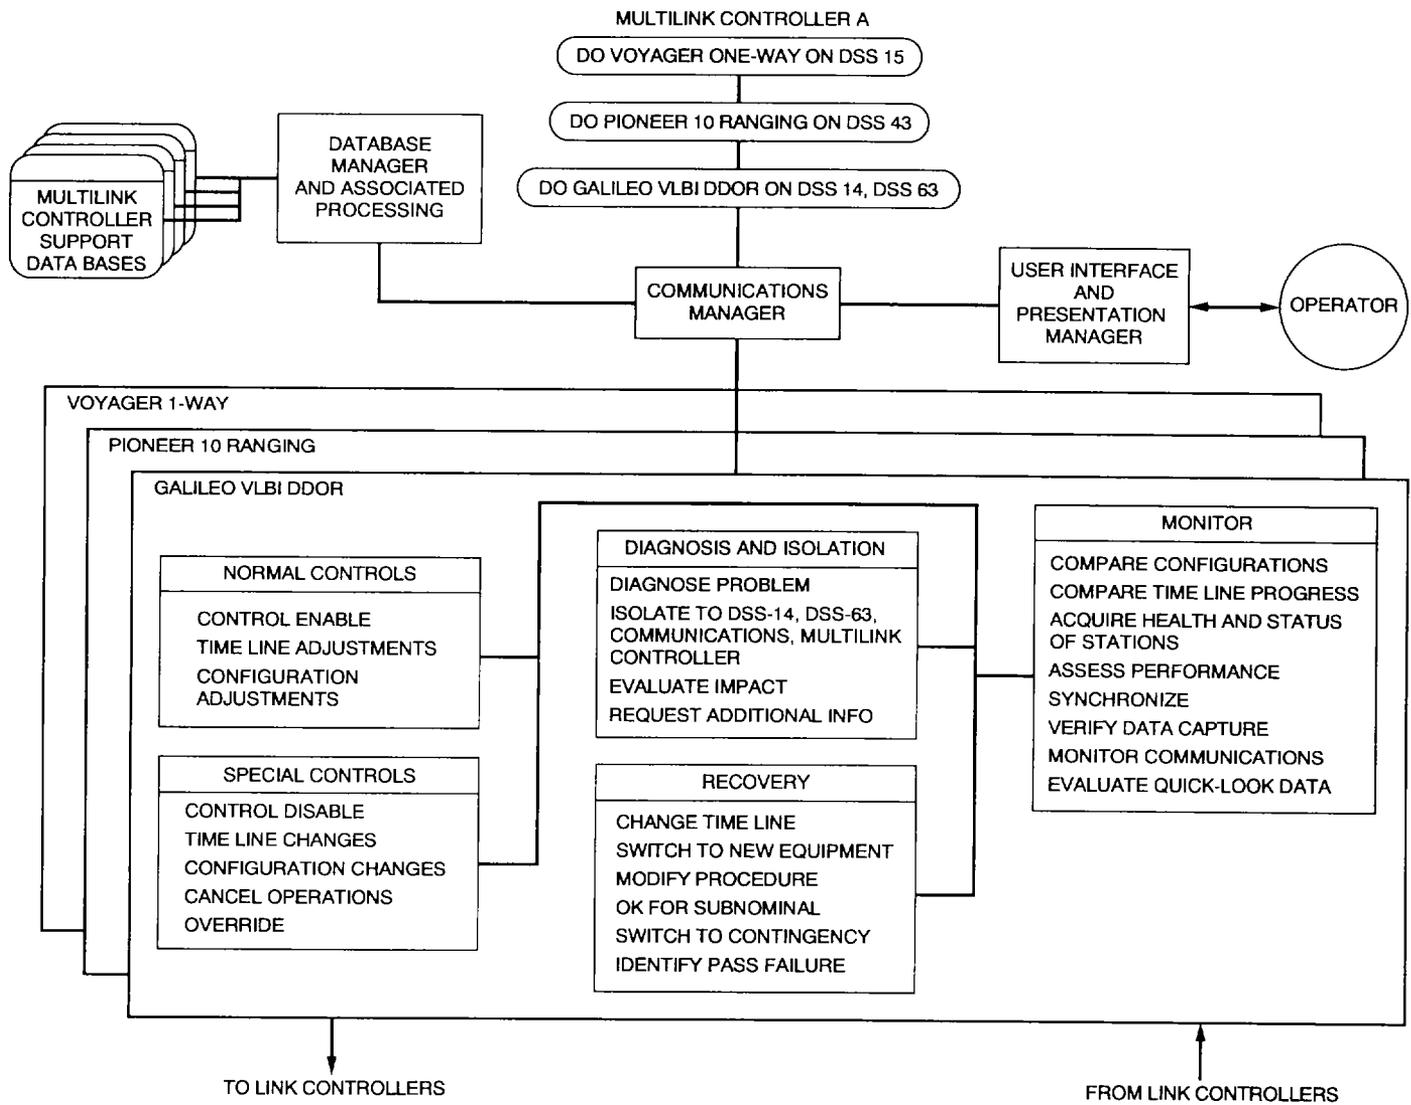


Fig. 22. Multilink controller operations.

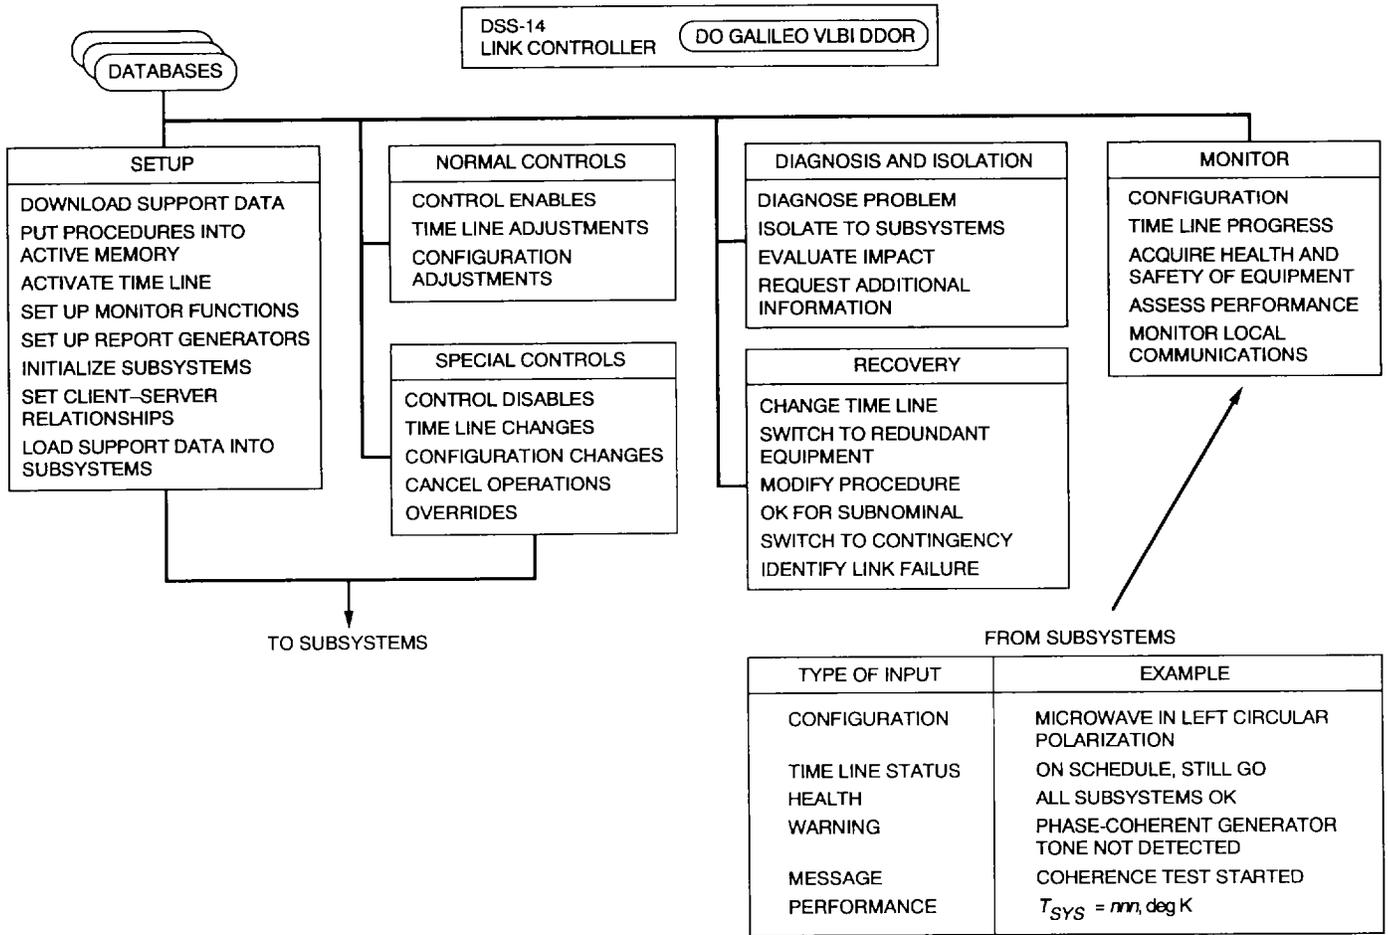


Fig. 23. Link controller operations.

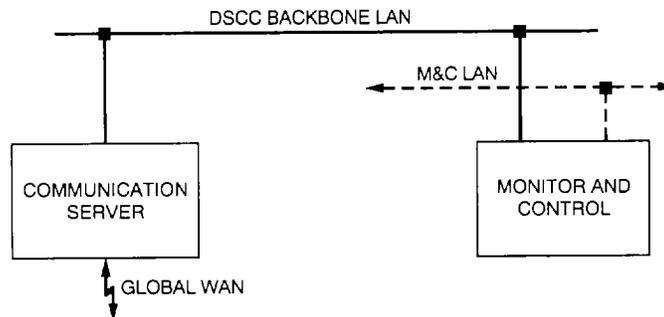


Fig. 24. Major DSCC subnetworks.

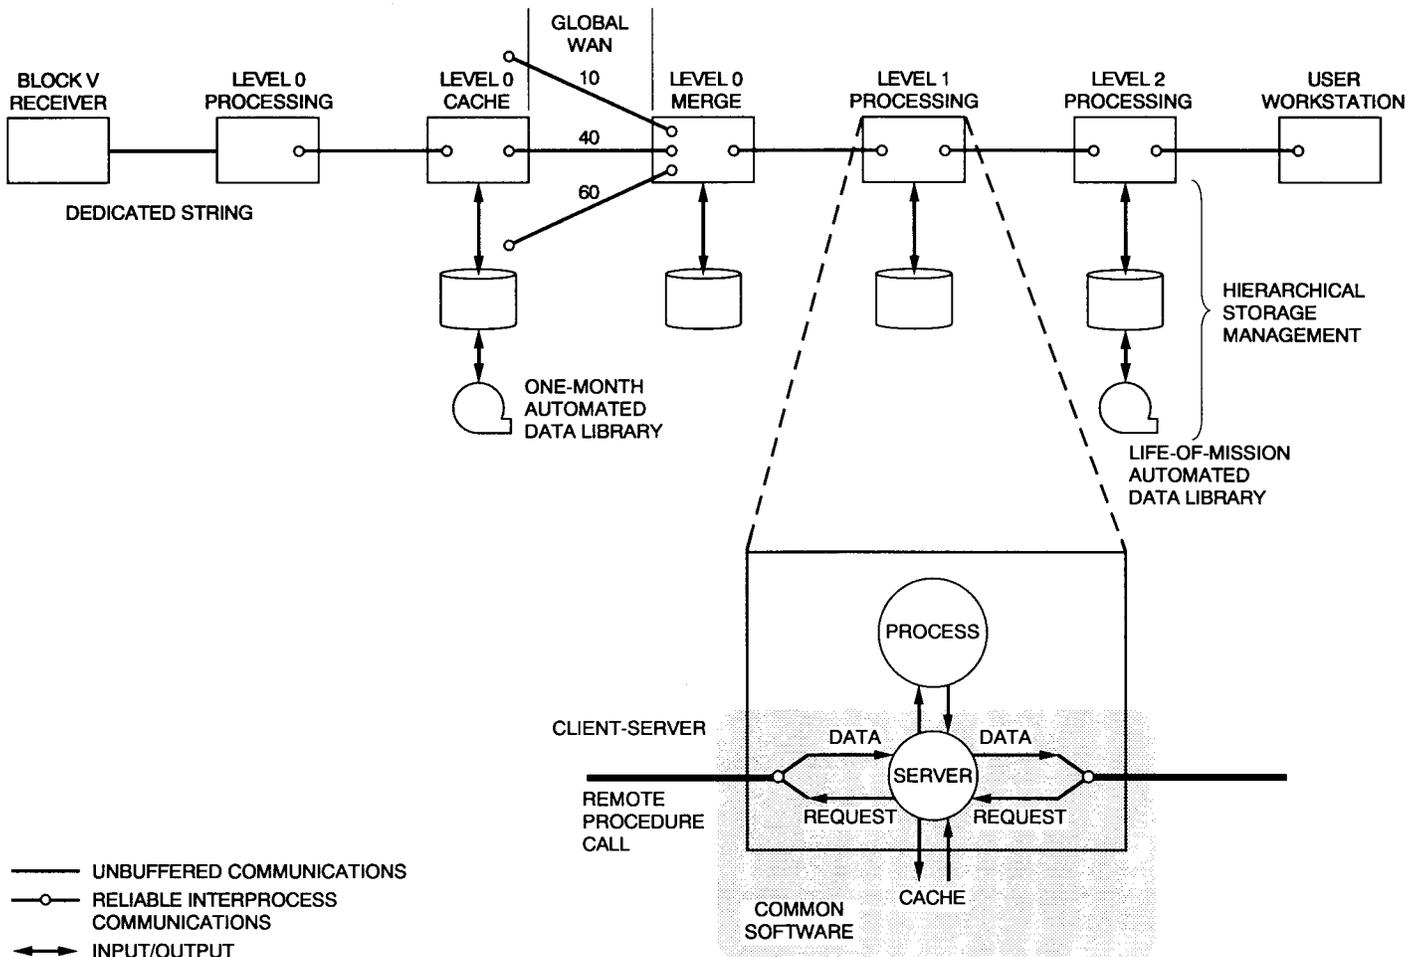


Fig. 25. Telemetry data flow with a client-server architecture.

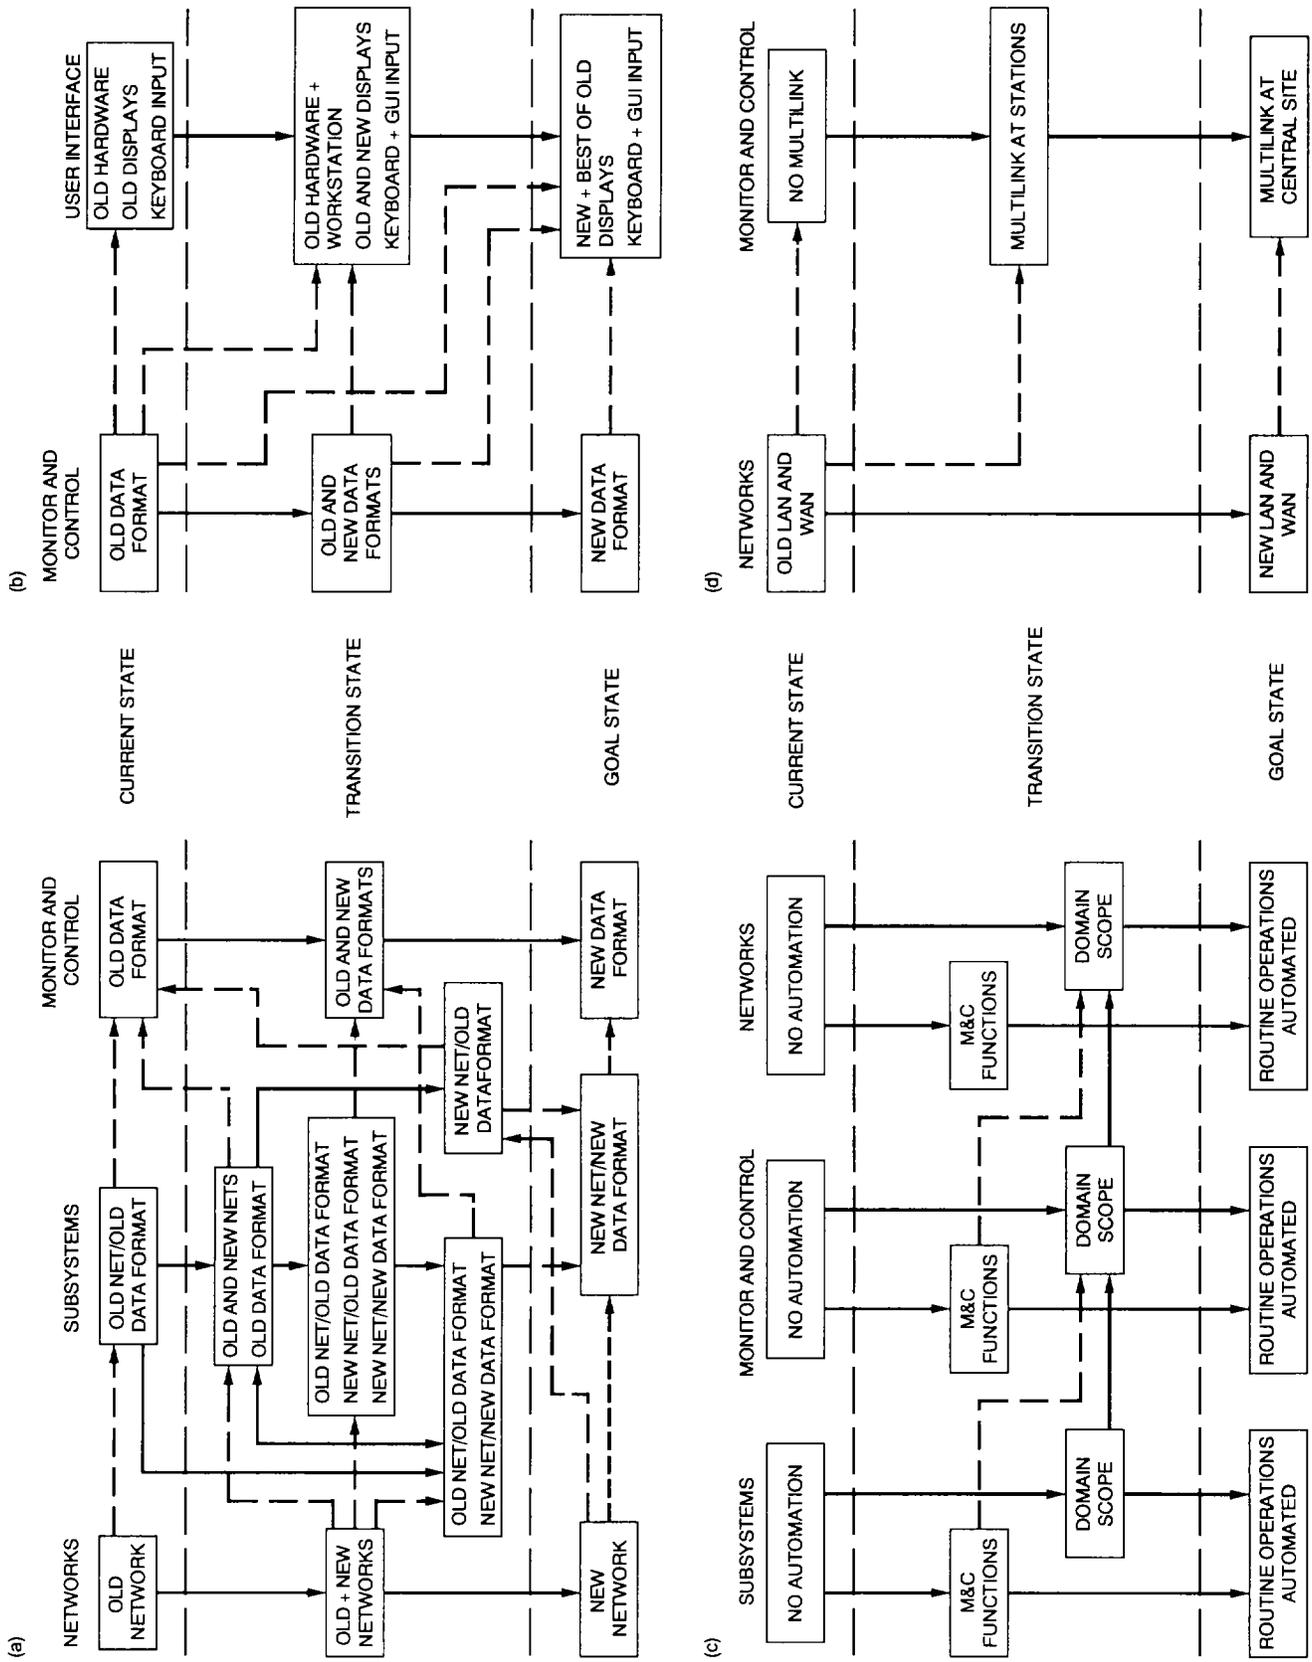


Fig. 26. Monitor and Control transition states: (a) message communications; (b) user interface; (c) automation; and (d) centralization.

<u>TECHNOLOGY DEVELOPMENT AREAS</u>	<u>ARCHITECTURE AREAS</u>				
	TELEMETRY PROCESSING	OTHER PROCESSING	MONITOR AND CONTROL	DATA TRANSPORT	SOFTWARE
HUMAN-COMPUTER INTERFACES			■		
ARTIFICIAL INTELLIGENCE AND AUTOMATION			■		■
LAN's, WAN's, PROTOCOLS, AND NETWORK MANAGEMENT			■	■	
SOFTWARE ENGINEERING	■	■	■	■	■
DATABASE MANAGEMENT SYSTEMS AND DATA MANAGEMENT	■	■	■		■
DATA STORAGE TECHNOLOGY	■	■	■	■	
VERY LARGE SCALE INTEGRATION AND OPTICAL NEURAL PROCESSORS, INCLUDING HARDWARE PACKAGING	■	■	■		
DISTRIBUTED COMPUTING	■	■	■		■

Fig. 27. Ground Information System architectures versus technology needs.

Appendix A

Acronyms

AES	Application Environment Specification	DSCC	Deep Space Communications Complex
AGC	Automatic gain control	DSS	Deep Space Station
AI	Artificial intelligence	DSN	Deep Space Network
ANSI	American National Standards Institute	DTK	DSCC Tracking Subsystem
AOS	Advanced Orbiting System	DTM	DSCC Telemetry Subsystem
ATDRSS	Advanced Tracking and Data Relay Satellite System	EOS	Earth Observing System
BER	bit error rate	FDDI	Fiber Distributed Data Interface
bps	Bit per second	FFT	Fast Fourier transform
CASE	Computer-aided software engineering	FIPS	Federal Information Processing Standard
CCITT	International Telegraph and Telephone Consultative Committee	FTAM	File Transfer, Access, and Management
CCSDS	Consultative Committee for Space Data Systems	FTP	File Transfer Protocol
CD-ROM	Compact disk/read-only memory	FY	Fiscal year
CLNP	Connectionless Network Protocol	GCF	Ground Communications Facility
CMA	Command Modulator Assembly	GFLOPS	Giga (10^9) floating point instructions per second
CMC	Complex Monitor and Control	GIPS	Giga (10^9) instructions per second
CMIP	Common Management Information Protocol	GIS	Ground Information System
CMIS	Common Management Information Service	GLL	Galileo
CMOS	Complementary metal oxide semiconductor	GOSIP	Government OSI Profile
COTS	Commercial off-the-shelf	GNMP	Government Network Management Profile
CPA	Command Processor Assembly	GSOC	German Space Operations Center
CPN	CCSDS Principal Network	GUI	Graphical user interface
CRAF	Comet Rendezvous Asteroid Flyby	HDLC	High-Level Data Link Control
CRC	Cyclic redundancy check	HEO	High-Earth Orbiter
CSMA/CD	Carrier-sense multiple access with collision detection	HIPPI	High-Performance Peripheral Interface
DBMS	Database management system	IEEE	Institute of Electrical and Electronics Engineers
DCD	DSCC Command Subsystem	I/O	Input/output
DCE	Distributed Computing Environment	IP	Internet Protocol
DCO	Digitally Controlled Oscillator	ISDN	Integrated Services Digital Network
DDOR	Delta differential one-way ranging	ISE	Information System Engineering
DFS	Distributed File System	ISO	International Organization for Standardization
DoD	Department of Defense	kbps	Kilobits per second (10^3 bps)
		KBS	Knowledge-based system
		LAN	Local area network
		LAPB	Link Access Procedure-Balanced

LMC	Link Monitor and Control	RF	Radio frequency
LZM	Level-Zero Merge	RFI	Radio frequency interference
LZP	Level-Zero Processing	RPC	Remote procedure call
M&C	Monitor and Control	RS	Reed-Solomon
MAP	Manufacturing Automation Protocol	SDH	Synchronous Digital Hierarchy
Mbps	Megabits per second (10 ⁶ bps)	SEI	Space Exploration Initiative
MBAR	Main Belt Asteroid Rendezvous	SFOC	Space Flight Operations Center
MDA	Metric Data Assembly	SHARP	Spacecraft Health Automated Reasoning Prototype
MFLOPS	Mega (10 ⁶) floating point instructions per second	SIMD	Single instruction-multiple data
MHS	Message Handling Service	SMIM	Submillimeter Intermediate Mission
MIB	Management Information Base	SMTF	Simple Mail Transfer Protocol
MIMD	Multiple instruction-multiple data	SNMP	Simple Network Management Protocol
MIPS	Million instructions per second	SOE	Sequence of events
MLC	Multilink Controller	SONET	Synchronous Optical Network
MMS	Manufacturing Message Specification	SPAN	Space Physics Analysis Network
MOSO	Multimission Operations Systems Office	SPOC	Science Planning Operations Computer
Msp	Megasymbols per second (10 ⁶ eight-bit symbols per second)	SPC	Signal Processing Center
MTBF	Mean time between failures	SQL	Structured Query Language
NASA	National Aeronautics and Space Administration	SR	Satellite relay
NASCOM	NASA Communications Network	SRA	Sequential Ranging Assembly
NCB	Narrow-Channel Bandwidth	SSF	Space Station Freedom
NCF	Network Control Facility	Tbyte	Terabyte (10 ¹² bytes)
NIST	National Institute for Standards and Technology	TCP	Transmission Control Protocol
NOCC	Network Operations Control Center	TDA	Telecommunications and Data Acquisition
NSS	Network Support Computer	TDRSS	Tracking and Data Relay Satellite System
OSF	Open Software Foundation	TP	Transport Protocol
OSI	Open Systems Interconnection	TP4	Transport Protocol Class 4
OSO	Office of Space Operations	TPDS	Tracking and Data Processing Subsystem
PAD	Packet assembler/disassembler	VHSIC	Very high speed integrated circuits
PCM	Pulse-coded modulation	VLA	Very Large Array (National Science Foundation, Socorro, New Mexico)
POC	Project Operations Center	VLBI	Very long baseline interferometry
POSIX	Portable Operating System Interface for Computer Environments	VLSI	Very large scale integration
RDA	Remote database access	VSAT	Very small aperture terminal
		WAN	Wide area network
		WCB	Wide-Channel Bandwidth

Appendix B

Glossary

Application environment: Refers to all the software necessary to run an application except for the application-specific code. It includes the operating system, I/O routines, and network services and may include such items as a DBMS and graphical user interface (see Section III.C).

Architecture: A specification that defines the design of a system, and describes the structure, functions, and interfaces of its cooperating parts (see Section I.C).

Distributed: In network technology, this term refers to geographically dispersed computers that share common applications and information via local-area networks and wide-area networks (see Section IV.B).

Fault-tolerant: Systems that are able to detect a fault, report it, mask it, and then continue service while the faulty component is repaired off-line.

Dedicated: Systems that perform only one primary function, e.g., telemetry processing on the output of only one receiver (see Section IV.B).

Artificial intelligence: The technology whereby computer systems can be programmed to emulate special properties and capabilities attributable to human intelligence (see Section III.E).

Backbone network: A network designed to interconnect lower speed channels (see Section IV.D).

Central Site: The control center of the DSN, located at Jet Propulsion Laboratory (JPL) facilities in Pasadena, California (see Section IV.B).

Client-server: In a communications network, the client is the requesting machine and the server is the supplying machine (see section III.C.3).

Operating conditions: The conditions under which a system performs its functions.

Anomalous: A situation where operational procedures do not exist.

Backup: A situation where operational procedures are used to accommodate a failure of the primary system.

Nominal: In the operations concept of the M&C architecture, nominal operations refers to a situation where existing operational procedures can handle any problem.

Codes: As used in this study, codes are error-correcting bit patterns that are used to reconstruct a spacecraft signal into a form so that departures from this construction in the received signal can be automatically detected. This permits the automatic correction at the receiving terminal of some or all of the errors. Such codes require more signal elements (or bits) than are necessary to convey the basic information [4] (see Section II.A).

Data: The normal definition of data in this article refers to representation of information in binary digits (bits). However, in estimating overall data rates, data may also include redundant bits added for error detection and correction.

Near-real-time: Data transmitted directly from a sender to a receiver with possible delays on the order of seconds due to temporary buffering (see Section II.A).

Non-real-time: Data transmitted from a sender to a receiver with significant storing and shipping delays (on the order of hours or days). (See Section II.A.)

Real-time: Data that is transmitted directly from a sender to a receiver with virtually no delays.

Datagram: Data packet that is transmitted with sufficient information to reach its destination, but with no expectation of an acknowledgment.

Data protocol: A standardized set of rules that specifies the format, timing, sequencing, and/or error checking for data transmission.

End users: Recipients of DSN-produced data products. Recipients include spacecraft teams, principal investigators, archival storage systems, and non-NASA partners (see Section II.D).

Front end: The antenna equipment and electronics collocated with the antenna.

Fuzzy logic: Logic based on probabilities that statements are true rather than whether they are true or false (see Section III.C).

Gateway: A device for translating application services from one set of protocol standards to another, e.g., electronic mail from Internet to OSI.

Ground Information System (GIS): The DSN information processing system. In scope, it extends from the ground side of the front end to the end users. GIS functions include processing (telemetry, tracking, command, radio science and VLBI), monitor and control (of front end and processing), and data delivery and management, with the focus of this study on multimission operations (rather than on project-unique activities). (See Section I.C, II.D.)

Hypermedia: Technology that enables a user to review data by following a nonlinear path, for example, by jumping to related data items. The formats may include text, graphics, animation, digital audio, and video (see Section III.E).

ISE Laboratory: Information Systems Engineering Laboratory; a proposed laboratory for the envisioned GIS architecture which will be required to evaluate, prototype, and adapt commercial technologies for DSN utilization (see Section V.E).

Internet: As normally used in this study, a short form of “internetwork,” a generic name for any network that interconnects two or more multivendor computers and networks (see Section III.D).

DoD Internet: A DoD packet-switched network that uses military standard protocols for such internetwork services as ensuring message reliability, routing, file transfer, and electronic mail (see Section III.C).

CCSDS Internet: Formal name for a spacecraft-to-user packet delivery service specified in CCSDS Advanced Orbiting Systems (AOS) architecture (CCSDS 701.0-B-1). (See Section III.D.)

Layers: Result of analyzing a system by dividing the system functions into groups of common services. Each layer specifies its own functions and assumes that lower level functions are provided. In “open systems” architecture, a collection of related functions that comprise one level of a hierarchy of communications functions.

Level-zero processing: Processing of spacecraft telemetry data to remove all artifacts of space data transport from the data and return it to the form it was in when it left the instrument aboard the spacecraft. Multimission level-zero processing (LZP) is done, to the extent possible,

at each DSCC; final multimission level-zero merge (LZM) occurs at the Central Site.

NCF: Network Control Facility; an envisioned organization located at the Central Site, that exercises the highest level of monitor and control; it is planned to be the primary control point for the entire DSN (see Section IV.C).

Neural network: Parallel, distributed computing system of simple computational elements connected in a fashion which mimics, at a low level, the connections between neurons or systems of neurons in a brain (see Section III.C).

OSF: Open Software Foundation; a consortium of computer-product vendors committed to developing products that provide common application environments (see Section III.C.1).

OSI: Open-systems interconnection; ISO’s reference model for enabling multivendor system to interoperate. The OSI protocols are a subset of the larger set of standard protocols established by the ISO.

GOSIP: Government OSI profile; a list of OSI protocols specified by the Federal government for use in Federal networks.

Open systems: A computer application environment based on products from a vendor compatible with those available from its competitors, usually because they mutually agree on standards (see Section III.C).

Quick-look data: The subset of the instrument science data stream for which access and processing is necessary earlier than that normally provided. These data may be used for instrument health or safety monitoring (see Section II.B).

Virtual circuit: A network model in which data is received from the transport layer and delivered over the perfect connection: no errors, no duplicates, and all packets are delivered in sequence (see Section III.D).

Virtual reality: An environment in which computer and peripheral devices create a complete sensory environment dynamically controlled by the user (see Section III.E).

References

- [1] J. H. Yuen, ed., *Deep Space Telecommunications Systems Engineering*, JPL Publication 82-76, Jet Propulsion Laboratory, Pasadena, California, July 1982.
- [2] B. Boehm, "A Spiral Model of Software Development and Enhancement," *Computer*, vol. 21, no. 5, pp. 61-72, May 1988.
- [3] P. S. Henry, "High-Capacity Lightwave LAN's," *IEEE Communications*, vol. 27, no. 10, p. 21, October 1989.
- [4] U. S. Department of Commerce, Federal Information Processing Standards, FIPS Publication 146-1, Washington, DC, August 1988.