DATA SYSTEM OVERVIEW

Joseph R. Kahr
Jet Propulsion Laboratory/California Institute of Technology

ABSTRACT

The design of the SFOC data system is based on a “design for change” philosophy. It emphasizes standards throughout the implementation, allowing for reuse of software, for periodic changeout of hardware, and for an evolving network configuration. Commercial off-the-shelf hardware and software components are incorporated in a way that avoids dependencies on any single vendor. Multiple flight projects are supported by building upon the baseline system with a minimum of special purpose adaptations.

In addition to the multi-mission aspect of SFOC, it must also satisfy multiple users representing multiple disciplines. Data system operators monitor and control SFOC itself. Spacecraft team members keep a vigil to protect the health of the spacecraft. Mission planners and sequence designers control the spacecraft. Science investigators remotely calibrate and control their onboard instruments. SFOC provides near-realtime and non-realtime support to end-users for downlink (telemetry) and uplink (command) functions.

This paper provides an overview to the design of the overall SFOC system and describes the implementation of the current baseline SFOC. It summarizes the important design decisions that have been made, and explains the approach taken to meeting these challenging requirements.

INTRODUCTION

SFOC is a large scale distributed system for support of continuing and future space missions operated through JPL. The SFOC data system will ultimately be used in support of six or more planetary space missions, and will have over 200 analysis workstations. Portions of the system will extend coast to coast and even to other countries to support the global telecommunity of science investigators.

SFOC is a part of JPL’s end-to-end information system. Other parts include the Deep Space Network (DSN), the Ground Communication Facility (GCF), and the Planetary Data System (PDS). The DSN is the ground part of the space/ground telecommunications link -- both downlink (telemetry from spacecraft to ground) and uplink (telecommand from ground to spacecraft). The GCF links SFOC with the three geographically disposed radio telescopes of the DSN. SFOC serves as the operations center for the spacecraft teams, mission control teams, and science teams. PDS eventually receives data products from SFOC, usually after end of mission, and distributes them to the general science community.

To date the SFOC implementation has progressed through three stages: an early prototype, a generic baseline system, and a first adaptation to meet mission-specific requirements in support of the launch and early cruise phases of the Magellan mission.

Magellan is a mission to do high resolution mapping of the surface of Venus using synthetic aperture radar (SAR). To support orbital operations at Venus a high rate data subsystem will be added to SFOC for processing SAR data. For launch and early cruise the only spacecraft telemetry seen by SFOC is engineering data. An older mission control system will provide uplink related support for Magellan, so the current SFOC implementation only includes support for downlink related functions. Magellan uses time division multiplexed (TDM) telemetry; a packet telemetry capability will be added to SFOC for future missions.

DESIGN CONSTRAINTS

SFOC is intended to be a multi-mission system that can be adapted to meet new requirements and can evolve with changing technology. It is to serve as the primary spaceflight operations center at JPL throughout the 1990s and into the 21st century. There is a strong hope that new missions can be added to SFOC with relatively small cost increments, including the cost to develop and the cost to operate.

These constraints translate into a few basic guidelines. The system should be modular in hardware and software, with loose coupling between modules for easy replacement. It should enable users to perform multiple roles, either across missions or across disciplines. This implies multi-function workstations and consistent user interfaces. Also, keep the system flexible even if it may detract from simplicity or security.

Most important is the use of standards as a way to enforce commonality and to maximize reusability of hardware and software. The standards that have been adopted as part of the SFOC design and implementation are:

<u>Established or Emerging Computer Industry Standards</u>

UNIX	Operating System (4.2BSD with System V IPC)
TCP/IP	Network Protocol (ISO in the future)

C, Fortran	Programming Languages
X-Windows	Windowing Package
C-ISAM	File Access Method
SQL	Relational DBMS Language
NFS	Inter-Node File System (for non-critical use)
XDR	Data Representation
(TBD)	Graphical Display Package
(TBD)	User Interface Management System

Consultative Committee on Space Data Systems Recommendations

SFDU	Data Formats
Packet TLM	Downlink Telemetry (future)
Telecommand	Uplink Command (future)

JPL Standards

SPICE Kernels	Supplementary data records

UNIX and associated standards were seen as the way to meet three basic goals:

- Interchange of data between dissimilar machines,
- Interchange of programs between dissimilar machines, and
- Interchange of users between dissimilar machines.

Interchange of data is relatively easy when it works. Keeping it working requires system administrators skilled in networking and communication. Proper configuration of routing tables, permissions, versions, daemons, etc. is necessary. Non-ASCII data structures are exchanged between machines using either XDR or a SFOC developed global data representation (GDR). GDR is closer to spacecraft data representation and allows conversion of selected fields in a record. In addition, the case of data exchange between machines is related to the degree to which they adhere to industry standards. Eccentric machines are often trouble spots in the network.

Interchange of programs among dissimilar machines has been done in SFOC without massive rewrite, but it is not effortless. Again, the peculiar versions of UNIX usually causes problems. Even among similar machines proper configuration management and resource allocation is necessary to ensure portability of software. Users will have the ability to develop their own software as a way to customize analysis tools. This is anticipated to be a powerful and flexible capability, but it will aggravate portability concerns.

Interchange of users among the various SFOC machines has proven to be very effective. In our case “users” thus far have been programmers, integrators, administrators, and testers. They move freely between machines and feel comfortable with the system once they get enough experience with it. Because of the similar user interface throughout the system it behaves like a unified whole rather than a disjoint collection of machines. Training on UNIX now will save much retraining in the future.

HARDWARE ENVIRONMENT

The current SFOC data system is a distributed collection of workstations and dedicated computers. There are two local area networks, one at JPL in California, the other at a remote Mission Support Area (MSA) in Denver, Colorado. They are linked via two 56 kbps circuits into one logical wide area network.

The system design emphasizes allocation of functions to software; the allocation of software to hardware is kept rather fluid. Some machines, typically those used for “front-end” processing or large scale data basing, run only one software configuration. This is due to requirements for special purpose I/O or auxiliary processors. Other machines, however, are reconfigurable by system administrators and by users. In all cases the exact number of processes and their interconnectivity is defined by operators and users to meet their needs.

SYSTEM LAYERS

The SFOC system is partitioned into several layers. The layered organization results from the building block approach to the design. Each layer is built of pieces from lower layers. These pieces include subsystems, applications, jobs, processes, services and functions.

The system contains hardware, software, personnel, procedures and documentation. There is a core of common pieces around which the mission adaptations are built. Subsystems are implementation divisions of the system which are more easily managed from an administrative or accounting point of view.

An application is a portion of a subsystem that performs a well defined function that is under the control of an operator or user. A job is a collection of applications working on one logically connected stream of data. Jobs are treated as a single unit although they may span multiple machines.

Processes are the smallest piece of software that can be moved intact. They are implemented as UNIX processes. Porting a process to a different machine may require recompiling and relinking, but should not require coding changes.

Services are general purpose utilities, libraries, and daemons which can be linked or invoked by processes to perform specific functions. The same service is used to perform the same function in all processes, thus enforcing a standard across the system. There are services to access the network, access files and data bases, convert data types, receive user input, generate displays, etc.

Functions are the smallest piece of software that is under configuration control as a single entity. Functions build processes and include callable subroutines, shell scripts, macros, process code, etc.

Layering within the network protocol follows the ISO Open Systems Interchange model. The physical, data access, network, and transport layers are implemented using coaxial cable, Ethernet, IP, and TCP respectively. The session, presentation, and application layers are standard utilities (FTP and Telenet) plus custom software implemented as services.

SOFTWARE COMPONENTS

The software components in SFOC can be grouped into a few broad categories: basic system environment, generic services not directly related to uplink or downlink, applications common to both uplink and downlink, applications related to downlink, and applications related to uplink. Those subsystems which are not the current implementation of SFOC are marked with an asterisk.

System Environment

UNIX and its associated utilities provide the system-wide environment. The socket mechanism of 4.2BSD is the basis for inter-computer communication. Pipes and message queues are used within machines for interprocess communication.

Generic Services

DTS Data Transport Subsystem -- Provides data communications and transport services between processes and between machines. Completes the session, presentation, and application layers.

CDA Common Data Access -- Provides access methods for byte stream files, indexed files, and relational data bases. Also implements “spoolers” which are multiple access FIFO disk buffers.

CDB Central Data Base -- Provides general data management services to other subsystems. Loads, catalogs, queries, and archives data to a relational data base implemented using Sybase.

WSE Workstation Support Environment -- Provides the user interface and device independent display environment.

SMC SFOC Monitor and Control -- Provides job control, event logging, status monitoring, and time synchronization for operators of the data system.

TWS Test Workstation -- Provides general purpose data manipulation, inspection, and diagnosis tools for testing and troubleshooting.

EUA* External User Access -- Provides secure access by external users to a limited SFOC environment.

Common Applications

GIF GCF Interface -- Captures and routes incoming telemetry and outgoing command data. Reformats external data into internal structures.

DTV Digital Television -- Provides the digital-to-analog interface to the closed circuit television system.

NAV* Navigation -- Performs multi-mission navigation support functions, such as generation of ephemeris generation and optical navigation.

TAS* Telecommunications Analysis -- Analyzes and predicts spacecraft/ground telecommunications performance.

SIM* Simulation -- Generates simulated data for test and training purposes. Simulates the GCF, DSN and multiple spacecraft.

Downlink Applications

TIS Telemetry Input Subsystem -- Performs initial processing on telemetry frames and DSN monitor data. Telemetry processing includes frame synchronization, decoding for error correction, synchronous and asynchronous extraction, depacketization, decommutation, and channelization.

DMD Data Monitor and Display -- Performs channel processing including derivation and conversion, alarm checking, and display generation. Display types include plots, fixed matrix, variable matrix , and lists.

EAS* Engineering Analysis Subsystem -- Provides a set of tools such as graphics, trend analysis, statistics packages, etc., together with an environment for user- and project-supplied software for processing and analysis of spacecraft engineering data.

SAS* Science Analysis Subsystem -- Provides a set of tools such as graphics, trend analysis, correlation analysis, statistics packages, etc., together with an environment for user- and project-supplied software for processing and analysis of science data.

DPS* Data Products Subsystem -- Generates data products containing experiment data records (EDRs) and supplementary data records (SEDRs).

MIP* Multimission Image Processing -- Performs image processing and image product preparation and distribution. Certain high rate non-imaging data is also processed.

Uplink Applications

CMD* Command -- Controls transmission of commands to the DSN for radiation to the spacecraft.

SEQ* Sequence Generation -- Provides a set of software tools which support project command sequence design and generation. Performs constraint checking, command translation, and spacecraft memory management.

MAS* Mission Analysis Subsystem -- Provides a set of software tools for the support of mission planning and analysis activities.

TELEMETRY DATA FLOW

The engineering telemetry rate for Magellan is 1200 bps in real time, and 115.2 kbps replay from onboard tape recorder. Much of the tape recorded data is filler with the resulting replay rate equivalent to 18 kbps of engineering. Telemetry arrives from the DSN via the GCF (see Figure 1). It is first processed by GIF where it is unpacked from GCF blocks and packaged into Standard Formatted Data Units (SFDUs). GIF can also log the raw and the processed data. SFDUs are then routed on to TIS.

TIS frame synchronizes the incoming telemetry to identify minor frames within the bit stream. It then extracts the engineering data from the minor frames. Decommutation follows reassembly of engineering frames from the parent TDM telemetry. As part of decommutation each engineering measurement is located in the frame and its identifier is appended to the value. This process is called channelization and the resulting records are called channels. The decommutation and channelization process is table-driven; the table represents the commutation map used by the spacecraft to assemble the telemetry. TIS can log raw telemetry, minor frames, and channels. It broadcasts the channels to the network as they are built. Later the channels together with the original minor frames are sent to CDB for long term storage.

DMD applications running in several workstations are able to receive the broadcast channels in real time. Channels are collected by DMD and placed into a table of latest available data (LAD). Selected channels trigger derivation of other channels. Values are converted from data number to engineering units. Displays are driven from the LAD table. Alarm checking is also done to find those values which have exceeded early warning thresholds (yellow alarms) and those values which have exceeded their limits (red alarms).

Data sent to CDB is checked for consistency, loaded to the data base, and cataloged. Channels are compressed as they are loaded to remove repetitive values. CDB also removes overlaps by assessing which of multiple data has the better quality and loading that set. The assessment of data quality is based on signal-to-noise ratio plus human judgement. Queries on the data base are supported as well as browsing through the catalog. The amount of data stored on-line is a function of disk space and compression efficiency. Data will remain on-line for approximately 45 days before being moved off to archives

Special processing of stored telemetry data is done to meet Magellan unique requirements. Selected channels are formatted into data return files in ASCII format that are readable by analysis programs running on MS-DOS machines.

ACKNOWLEDGEMENT

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

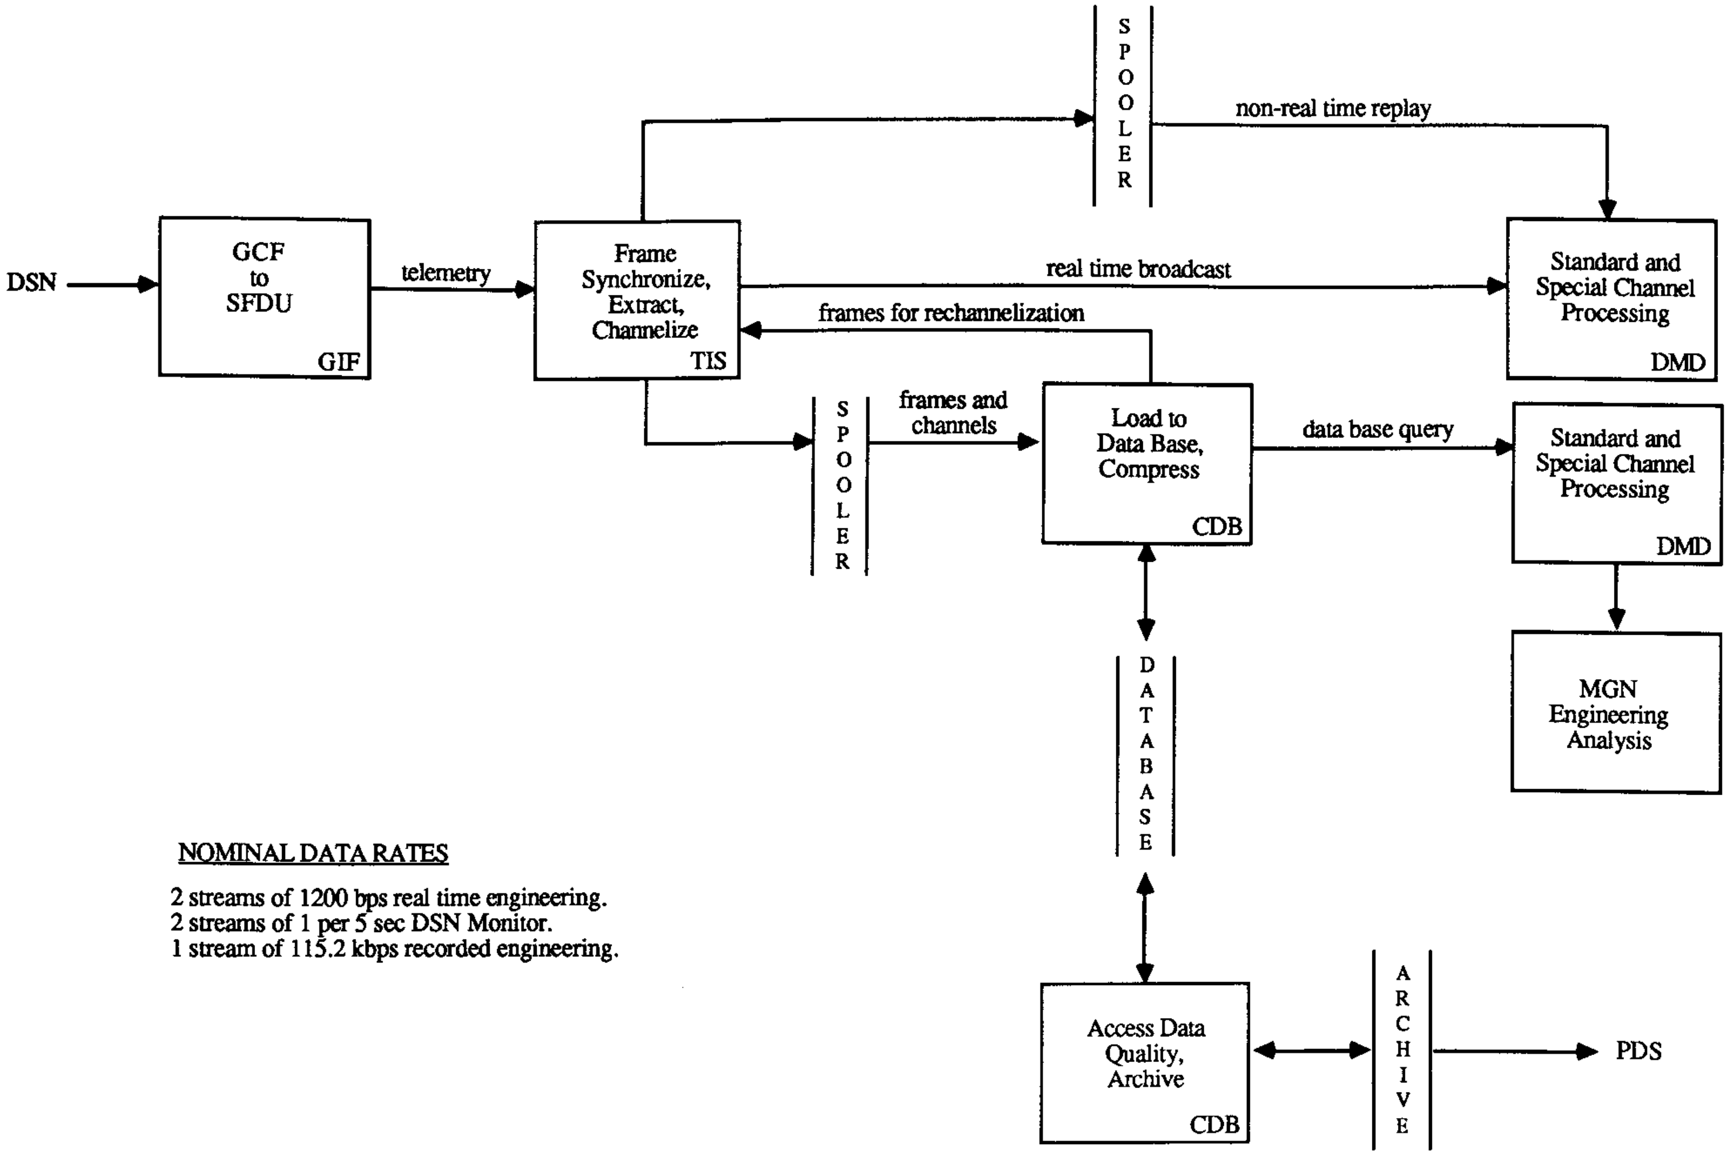

Figure 1. Engineering Telemetry Data Flow