

SCIENCE OPERATIONS AND DATA HANDLING FOR THE FAST SATELLITE

J. P. McFADDEN, R. E. ERGUN, C. W. CARLSON, W. HERRICK, J. LORAN,
J. VERNETTI, W. TEITLER, K. BROMUND and T. QUINN
Space Sciences Laboratory, University of California, Berkeley, CA 94720, U.S.A.

Abstract. The University of California, Berkeley, is the primary science analysis center for the Fast Auroral Snapshot Explorer (FAST). It is responsible for processing and routing the science data to the various investigators and serves as an on-line data archive and science command center for the project. Raw data, received from NASA within a few hours of its transmission to the ground, is processed into CD-ROMs for archive and distribution to the Co-Investigators, and summary data is generated for viewing at the FAST Web site or for retrieval as key parameter data from NASA or the FAST Web site. Daily science command loads and real-time commands are generated to optimize science data collection. This paper presents a description of the FAST ground operations performed at Berkeley. Also included is a discussion of the data analysis software and tools that allow the large data volume (>1 terabyte) to be accessed quickly and efficiently by the scientists. These tasks are more extensive than those performed in the past by science institutions and can be considered as a model for future small programs to reduce cost and maximize science return.

1. Introduction

The Fast Auroral Snapshot Explorer (FAST) was designed to investigate the plasma physics of auroral phenomena at extremely high time and spatial resolutions. It is the second of the Small Explorer (SMEX) Missions which were established to provide rapid, low cost opportunities to the space science community using a single Principal Investigator approach. The primary science objectives for this mission necessitated high data volume, real-time command capability, and control of science data collection on suborbital time scales. Since these requirements could increase costs to the project beyond the SMEX budget, the project adopted a nonstandard approach to its ground operations. This approach includes the use of semi-automated ground equipment to reduce personnel costs, and the offloading of specific tasks to the science institutions to shorten command turnaround time and reduce data collection costs. Descriptions of the NASA semi-automated ground stations, data packet processing equipment, and the Project Operation and Command Center (POCC) can be found in NASA/Goddard publications (NASA Document No. 822-DEM-1724, 1996; Shi *et al.*, 1993; SMEX FAST POCC System Users Guide, 1994). This paper presents a description of the FAST ground



operations performed at Berkeley, including a discussion of the commanding, data archiving, real-time software, and data analysis tools.

The Science Operation Center (SOC) at the University of California, Berkeley, is the primary science analysis center for the Fast Auroral Snapshot Explorer (FAST). The SOC is responsible for processing and routing the science data to the various investigators and serves as an on-line data archive and science command center for the project. The SOC receives level 0 (raw, time ordered) data from the NASA Goddard Space Flight Center within a few hours of transmission to the ground, processes the data into CD-ROMs for archive and distribution to the Co-Investigators, and generates summary data for viewing at the FAST Web site or for retrieval as key parameter data from the National Space Science Data Center or the FAST Web site. The FAST Key Parameters (KPs) are ingested into the NASA International Solar-Terrestrial Physics (ISTP) data base where they are included on both the ISTP KP CD-ROMs and on the ISTP CDAWeb site (<http://CDAWeb.gsfc.nasa.gov>). Daily science command loads and real-time commands are also generated at the SOC to optimize science data collection.

The paper is organized into four parts. It begins with a discussion of the science commanding, including real-time commanding and daily command loads performed by the SOC. The second section describes the data archive processing, including the generation of key parameter data. The third and fourth sections describe the real-time and data analysis tools developed for the scientists. It is hoped that this paper provides not only an overview of the FAST specific ground operations and software, but also a general outline of SOC tasks and science software design which can be used as a model for future small programs to reduce cost and maximize science return.

2. Science Commanding

The FAST spacecraft requires that daily commands be loaded into the stored command memory buffer for normal operations. These command loads include spacecraft commands generated by NASA such as attitude control system commands to preserve the nominal reverse cartwheel attitude, transmitter on/off commands, and telemetry channel enable/disable commands. Science commanding consists primarily of science instrument mode commands, trigger commands, and memory management commands. Due to power and data volume constraints imposed on the FAST spacecraft, data collection and instrument power are cycled. This cycling is performed by the science mode commands which selectively turn on and off the instruments during the orbit. These stored science commands are limited to 110 per day (10 per orbit).

2.1. SCIENCE MODE COMMANDS

Science mode commands are single commands that trigger execution of a command set stored on the spacecraft that completely configures the instruments. Science modes are specified by two one-byte codes that specify the field and particle experiments separately. Eight particle and eight field modes are stored onboard at any one time allowing up to 64 combinations. Although real-time mode commanding can be used, selection of instrument modes is generally performed on the ground where they are included in the daily stored command load. As new modes are developed, they are uploaded, replacing the older modes. Mode status is tracked on the ground to assure that a discontinued mode is not called out in any command set.

Instrument modes are used to configure the instruments to emphasize a particular science observation or to select the data collection rates to match data volume constraints. Particle modes are used to configure the 2 ion and 14 electron electrostatic analyzers (Carlson *et al.*, 2001 this issue) and the TEAMS mass spectrometer (Klumpar *et al.*, 2001, this issue). The mode sets the instrument sample rate (1.6–6.5 ms), high voltage sweep rate (64–16 sweeps spin^{-1}) and energy range (typically 2 eV to 12 keV for the mass spectrometer, 3 eV to 25 keV for the ion spectrometer, 4 eV to 30 keV for the electron spectrometer, and 6 fixed energies between 64 eV and 16 keV for the high time resolution electron spectrograph). Field modes determine electric (Ergun *et al.*, 2001, this issue) and magnetic sensor (Elphic *et al.*, 2001, this issue) measurement combinations. There are 10 electric field and 6 magnetic field sensor measurements which can be combined or multiplexed in various ways for the low rate survey data, the swept frequency analyzers, 8 high rate burst A/D channels, and 4 high frequency waveform capture A/D channels. The field mode also determines filtering, selection of electric field probe configuration (voltage or current), and sample rates (8 Hz to 2 MHz sampling). Details of the more than 100 FAST particle and field modes can be found on the FAST web site (<http://sprg.ssl.berkeley.edu/fast>). The mode also determines four data rates that are part of the FAST science optimization scheme. These four rates are referred to as Slow Survey, Fast Survey, Burst, and High Speed Burst Memory (HSBM). Slow survey is the slowest data collection rate and these data are used to provide a large-scale context for the higher resolution measurements. Slow survey data collection rates are comparable to previous satellite measurements with particle and wave spectra gathered on spacecraft spin time scales, and with low frequency waveforms having 30 ms resolution. Fast survey data collection is typically 8 to 32 times the slow survey rate and is triggered onboard in active regions, replacing slow survey data collection. Fast survey rates (230–1100 kbit) can be comparable to the highest burst memory collection rates on previous satellites. During orbits having consecutive memory dumps, fast survey data collection can be maintained through most of the auroral zone crossings.

Burst data collection is reserved for the most interesting events and is triggered by onboard algorithms that sample a variety of field and plasma parameters. Waveforms are sampled at 32 kHz, high frequency wave spectra are generated at 16 Hz, high resolution (48 energy \times 32 angle) electron and ion distributions are measured with 78 ms resolution, and coarse electron distributions (6 energy \times 16 angle) are measured with 1.6 ms resolution. These data are typically 1–2 orders of magnitude higher in time resolution than the fast survey data, thus they quickly fill the available onboard memory. Onboard trigger and event capture evaluation algorithms are used to identify the best events for selective transmission to the ground. Finally the HSBM data capture is used to measure high frequency waveform (2 MHz sampling) and also includes a separate trigger and event evaluation algorithm.

2.2. DATA COLLECTION TRIGGERS

In addition to the instrument mode commands, trigger commands are also sent as part of the daily science stored command loads. Eight trigger algorithms are stored onboard providing a selection of criteria for optimizing both fast survey and burst capture. Table I lists the inputs to the trigger algorithms and includes five current non-trivial algorithms. The simplest algorithms use the raw inputs such as the total low frequency filtered power (Ch4 or Ch5) which triggers on lower hybrid and whistler waves, the DC (<10 Hz) electric field magnitude (Ch0 or Ch1) that triggers on electrostatic shocks, or electron energy flux (Ch22 to Ch29) which triggers in enhanced auroral precipitation. More sophisticated algorithms that involve two or more inputs can be tailored to provide more reliable triggers for specific science (bottom of Table I). For example an algorithm that uses the density (Ch3) and high frequency wave power (Ch6) is an excellent trigger for selecting events with Auroral Kilometric Radiation and the associated density cavities and ion beams. The stored command loads allow selection of various triggers during an orbit or day to emphasize different science.

Figure 1 shows an example of FAST data gathered at different rates and the usefulness of the triggering algorithms. The top four panels show the electric and magnetic fields, electrons, and ions across a portion of the auroral oval. Although difficult to discern from the figure, a transition from slow survey (48 kbit s⁻¹) to fast survey (432 kbit s⁻¹) data rate occurred at 20:48:24 as FAST entered the auroral electron precipitation region, and a transition back to slow survey rate occurred at 20:55:00 as the spacecraft left the nightside aurora and entered the polar cap. These data rate transitions were triggered by the onboard algorithm that measured the electron energy flux. The second transition occurred about 30 s after entering the polar cap due to onboard filtering which prevents rapid toggling between data collection rates. The middle three panels show burst data captured by the electrostatic shock trigger algorithm. The large electric fields (panels 1 and 5) that initially triggered the burst are associated with intense upgoing electron beams (panel 6) and heated ions (panel 7) in the downward current region (positive

Table I
List the inputs to the trigger algorithms and 5 typical multi-input algorithms

Electric field inputs			
Channel	Quantity	Frequency	Science trigger
0	V1-V4 DC	< 10 Hz	Electrostatic Shocks
1	V5-V8 DC	< 10 Hz	Electrostatic Shocks
2	ELF V1-V4	10 Hz - 1 kHz	Ion Cyclotron waves
3	Ne_6	< 10 Hz	Density cavities
4	LFF V1-V4	1 kHz - 16 kHz	Whistlers, Lower Hybrid
5	LFF V5-V8	1 kHz - 16 kHz	Whistlers, Lower Hybrid
6	BBF V1-V2	>200 kHz	AKR, Langmuir
7	BBf V3-V4	>200 kHz	AKR, Langmuir
Particle inputs			
Channel	Quantity		Science trigger
8-9	i+ ESA1/2 max counts		Ion beams, conics
10-11	e- ESA 3/4 max counts		Arcs, Upgoing e-
12	i+ ESA 1/2 peak energy		Ion beams, conics
13	e- ESA 3/4 peak energy		Arcs, Upgoing e-
14-17	i+ ESA1 Energy Flux, 4/sweep		Ion beams, conics
18-21	i+ ESA2 Energy Flux, 4/sweep		Ion beams, conics
22-25	e- ESA3 Energy Flux, 4/sweep		Arcs, Upgoing e-
26-29	e- ESA4 Energy Flux, 4/sweep		Arcs, Upgoing e-
Typical algorithms			
Algorithm			Science trigger
AKR = Ch6 + (255 - Ch3)/2			AKR in density cavity
ESS = abs(Ch1 - 128) × 2			Electrostatic shocks
VLF = (Ch5 + Ch4 + 255 - Ch3)/4			VLF in low density
MAX = Max(AKR, ESS × 10, VLF)			General field trigger
EF = (Ch24 + Ch28)/2			e- Energy Flux, > 500 eV

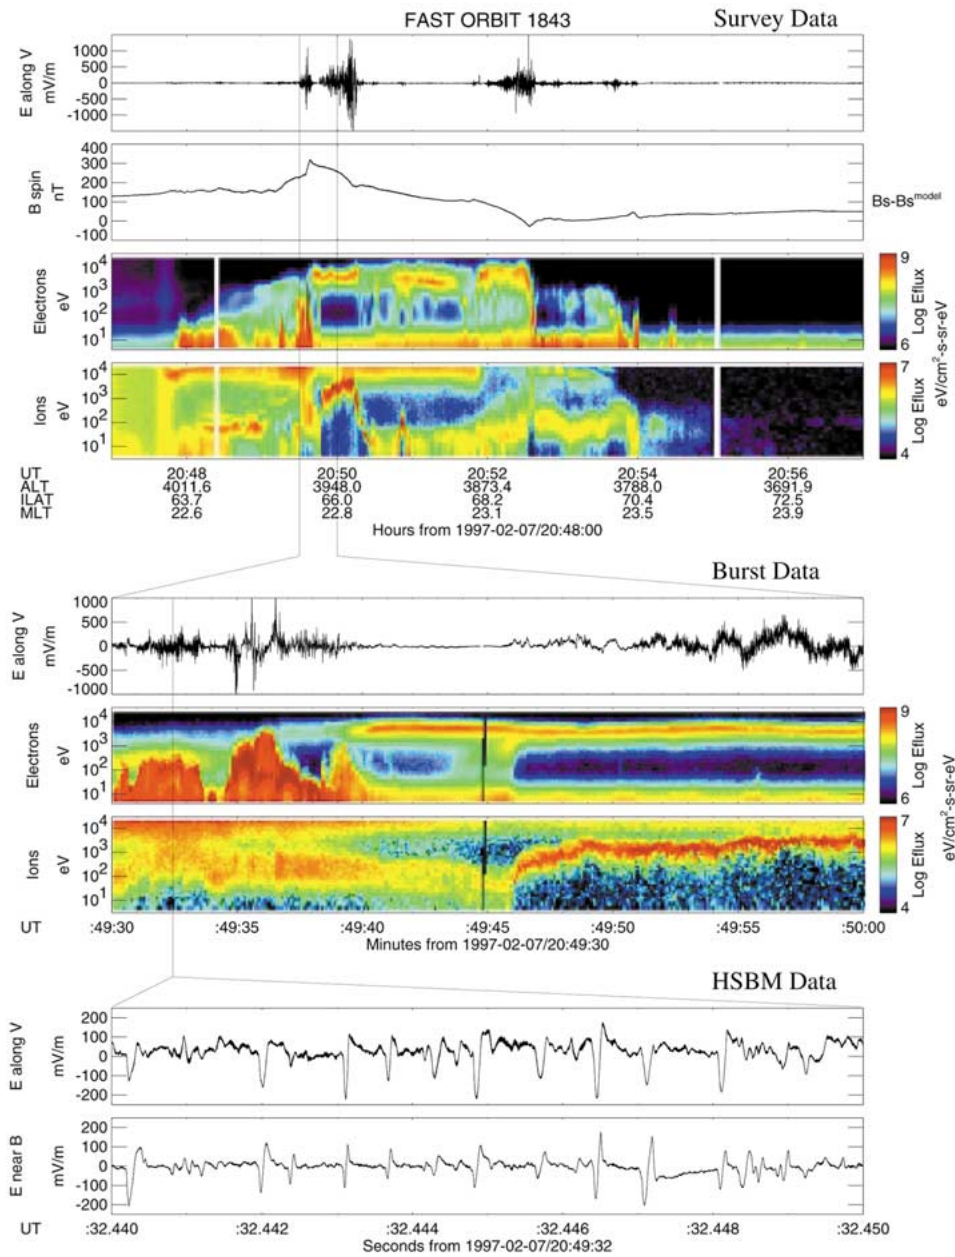


Figure 1. FAST data captured during a nightside auroral crossing. The top four panels show field and particle data collected at survey data rates. The middle three panels show burst data whose collection was triggered by an onboard algorithm sensitive to large electric fields. The bottom panels show high speed burst memory waveform capture of fast electron solitary waves that are common in downward current regions.

slope in the spin axis or eastward magnetic field component in panel 2). These downward current regions have been shown to contain fast electron solitary waves (or electron holes) that travel up the field line at the electron drift speed (Ergun *et al.*, 1998). The bottom two panels show the HSBM capture of the solitary waves showing a unipolar perpendicular (panel 8) and bipolar parallel (panel 9) electric field signal. The burst capture continued to trigger during the large electric fields associated with the upgoing ion beam spanning 20:49:46 to 20:50:18. Burst data (not shown) was also captured during the large electric fields at 20:52:30 amid intense electromagnetic ion cyclotron waves and at 20:53:42 during a second upgoing electron beam. The multiple data rates and onboard triggers have allowed the FAST spacecraft to capture scientifically interesting data on nearly every orbit.

2.3. DATA MANAGEMENT

Since the majority of data is generated by the instruments, the FAST solid state data recorder (125 MBytes) was incorporated into the science electronics package. To accommodate the various data products, the FAST data recorder is segmented into five partitions. These partitions contain packet headers, housekeeping, survey data, burst data, and HSBM data. To simplify data collection and readout, a standard 1024 byte packet is used for science data, with 40 bytes of packet header stored separately. Housekeeping packetization is handled by the spacecraft electronics developed by NASA, but housekeeping data is stored in the mass memory. The housekeeping data partition is treated as a FIFO (First-In-First-Out) but takes multiple (>6) orbits before it rolls over due to low housekeeping data rates, assuring that data is not lost even during days with only a few contacts. The burst and HSBM memory segments use selection criteria to overwrite events with poorer evaluation thus preventing memory overflow. The survey memory segment, however, is a FIFO and the high fast survey data rate may cause it to overflow.

To prevent survey memory overflow, ground software is used to manage this memory partition with a 'Fast Survey Allocation Command' which limits fast survey data collection. Onboard data transmission priority assures that the survey memory partition is emptied at the end of each nominal contact. Ground software is used to calculate the fraction of fast survey data collection that can be sustained without overflowing the memory before the next contact. The program uses as inputs the contact schedule provided by NASA, the mode data rates, and the data collection times. The Fast Survey Allocation Commands are added to the stored command load and executed following each contact.

2.4. POWER MANAGEMENT

In order to minimize spacecraft weight and achieve a higher apogee, the FAST power system was matched to science requirements and does not support continuous operation of all onboard systems throughout the mission. Figure 2 shows the prelaunch estimate of solar array power during the first six months of the mission.




Figure 2. The calculated solar array power during the first 6 months of the mission. The average spacecraft power is about 70 watts which exceeds the orbit averaged solar array power during most of the mission. Instead of using a larger solar array which would have added mass and reduced the FAST apogee, the scientists chose to power cycle the experiments. The instruments are turned on for about 30%–40% of the orbit, during auroral zone crossings, reducing the orbit averaged power to about 50 watts.

Variations in the curve are due to changes in the spacecraft spin vector with respect to the Sun and due to periods of eclipse as the orbit precesses rapidly about the Earth. Available power will also degrade with time due to radiation damage.

The spacecraft power budget for all the systems is shown in Table II. With all systems powered and assuming a nominal 30-min transmitter contact and 10 min of magnetic torquing for attitude control, the spacecraft requires about 70 W. This level is well above the available orbit averaged power for much of the mission, necessitating power cycling of the experiments.

Power cycling is achieved by using stored command loads to tell the experiments when to turn on/off. The primary purpose of these commands is to limit data collection to those periods when auroral science can be optimized, namely the high altitude auroral zone crossings. However, a secondary purpose essential to the mission is to limit power consumption by the experiments. This is achieved by using a low power ‘back-orbit’ mode for about $\frac{2}{3}$ of the orbit. The back-orbit mode requires only 7.7 W from the experiments, as opposed to 37.1 W for normal science modes. During back-orbit mode only low rate magnetic field data is collected and most experiments are turned off. Normal science modes are typically selected only for the auroral crossings in the Earth hemisphere containing apogee.

2.5. SUN–NADIR TABLE LOADS

In addition to the stored command loads, ground software is also required to generate a Sun–Nadir Table of angles. Particle data on FAST is sorted by spin phase

Table II
The spacecraft subsystem power budget

Spacecraft subsystem	Nominal power	Back orbit* power	Duty cycle	Average power
Mag S/C	0.5		1.0	0.5
Sun sensor	0.6		1.0	0.6
HCI	0.7		1.0	0.7
Spin coil	5.1		0.**	0
Precession coil	4.3		0.063***	0.27
MUE	13.5		1.0	13.5
Shunt drivers	0.5		1.0	0.5
Receiver	4.5		1.0	4.5
Transmitter	28.3		0.23****	6.5
Battery heaters	1.0		1.0	1.0
Other heaters	4.0		1.0	4.0
IDPU	18.3	6.6	0.3	10.11
E-field booms	2.0	0	0.3	0.60
Mag flux gate	0.08	0.08	0.3	0.08
Mag search coil	0.50	0	0.3	0.15
ESA stacks	11.75	0	0.3	3.53
TEAMS	4.5	1.0	0.3	2.05
	—			—
	100.13 W			48.59

*Backorbit instrument power is defined as non-science data collection modes.

**Spin coils are used primarily during early orbit operations to spin up spacecraft during boom deployments.

***Precession coil is typically run for 45 min, twice a day during Backorbit.

****Assuming one 30 min contact.

relative to the Sun. In order to accurately maintain this spin phase sorting during eclipse, onboard algorithms were developed to calculate the sun spin phase pulse based upon horizon sensor pulses, the midpoint of which give the Earth nadir direction. In order to calculate the sun pulse, a table of Sun-to-Nadir angles and delta angles, where delta angles are the change in sun-to-nadir angle per orbit, are uploaded daily. Due to the highly elliptical orbit and rapid precession rates, the table is non-trivial and changes significantly over a period of a few days. This Sun-to-Nadir Table allows the calculation of the sun-to-nadir angle, projected onto the spacecraft spin plane, as a function of time. Use of the onboard algorithm provides a smooth transition for the particle sorting when entering or leaving eclipse.

2.6. COMMAND GENERATION SOFTWARE

The mode, trigger, and fast survey allocation commands and the Sun-to-Nadir tables are generated for the daily command load by an automated program. This program takes as inputs a contact schedule (weekly file), a predicted orbit vector (daily file), a predicted spacecraft spin vector (weekly file), and a power budget (weekly file) all provided by NASA. The program also uses a file of rules generated at the SOC to customize the commands. The rules file allows specification of the spacecraft mode and trigger as a function of invariant latitude, magnetic local time and altitude. These parameters allow some flexibility in the automated mode selection. For example, different modes for midnight and cusp auroral crossings can be used to optimize different science in these regions.

At the start of a command generation session, the orbit vector is propagated out for three days and used to build three days of Sun-to-Nadir tables and stored command loads. Three days of commands and tables are generated and sent to the POCC to provide continuity in case network lines go down. During the second year of the mission, a five-day work week was implemented to reduce mission costs. To accommodate this change, the daily (24-hour) spacecraft command loads were replaced by 48-hour loads.

The automated command generation program works as follows. Data collection periods, and their associated Instrument Mode Commands, are determined from the calculated orbit and the rules file. Typical data collection is limited to periods above 60 deg invariant latitude in the hemisphere containing apogee. Once data collection periods are determined, the power is calculated for each orbit and compared to the available power budget. If the calculated power exceeds the available power, the data collection period is reduced appropriately. Burst Trigger and Fast Survey Trigger Commands are also generated and inserted in the load. Next, the Fast Survey Allocation Commands are calculated, based upon contact schedule and instrument mode data rates, and inserted into the stored command load. Sun-Nadir table loads are also automatically calculated and included. At this point, an opportunity is given for a SOC operator to check the loads before they are sent to NASA. The SOC operator also has the option of editing the command load, which may require a recalculation of the power allocation and Fast Survey Allocation Commands. The automated command generation program provides a convenient commanding method to generate daily command loads, and includes the flexibility needed to allow commands to be tailored to the orbit in order to accommodate special events such as multiple spacecraft conjunctions or coordination with ground-based operations.

2.7. REAL TIME COMMANDING

Although stored command loads are the primary means of controlling the FAST experiment, real-time commanding is available through the FAST Project Operation and Command Center. Individual commands or batch command procedures

can be initiated with a phone call. Commands are checked for validity and format with the FAST command database before transmission. As new commands are developed, they can be directly entered into the command database from Berkeley. Turn around time for the development of a new command procedure, incorporation into the database, and execution can be reduced to a few minutes, allowing changes in instrument operation to be initiated during a single orbit. During the FAST campaign, turn around time for commanding was less than 1 min and allowed rapid changes in instrument configuration based upon ground observations or real-time science analysis.

3. FAST Data Archive

The FAST SOC is the primary data collection and data archive processing facility for the project. The FAST science team took on this task, normally provided by NASA, to reduce costs to the project and to create an on-line data center that could handle the large volume of data (>1 terabyte). The large data volume is inherent to the mission, which specifies the collection of high rate measurements in the auroral regions each orbit. The 125 MByte solid state memory on the FAST satellite, together with a contact nearly each orbit (133 min orbital period), has produced a data volume approaching a terabyte during the first 2 years. In order to handle the data at a low cost, an automated archive system was developed that requires minimal personnel.

3.1. DATA COLLECTION

The data path for FAST begins on the spacecraft where the onboard solid state memory is used to collect high rate data and buffer it before transmission to the ground. Some of these data are preselected by onboard trigger and 'goodness' factor algorithms as higher quality data with priority for transmission to the ground. A table of data priorities and data rates is shown in Table III. Onboard software selects the data with the highest priority for first transmission, but limits the data rate for each quantity to its maximum rate. In general all data are sent to the ground and the memory emptied for nominal contacts. However, on those contacts where data collection is simultaneous with transmission, not all the Burst and HSBM data (VC 2) will be transmitted since these data are filled faster than can be read out. Here the selection criteria, based upon the goodness of each burst, is used to select priority for burst data readout. For more information on data selection, see Harvey *et al.* (2001, this issue).

Data are received by a number of ground stations. The primary ground stations are a Transportable Orbital Tracking Station (TOTS, see NASA Document No. 822-DEM-1724, 1996) located at Poker Flat, Alaska, and the NASA Wallops Flight Facility located at Wallops Island, Virginia. NASA Deep Space Network stations

Table III
FAST virtual channel encoding including data priorities and data rates

Virtual channel	Data type	Data priority	Bandwidth at 900 kHz	Bandwidth at 1500 kHz	Bandwidth at 2250 kHz
VC0	Real time housekeeping	1	Unlimited	Unlimited	Unlimited
VC1	Stored housekeeping	2	49.2	500	563
VC2	Burst data	6	Unlimited	Unlimited	Unlimited
VC2	HSBM data	7	Unlimited	Unlimited	Unlimited
VC3	Rate limited survey data	4	49.2	49.2	49.2
VC4	Nominal survey data	5	900	1000	1000
VC5	Engineering data	3	Unlimited	Unlimited	Unlimited
VC7	Fill data	8			

at Canberra, Australia and Santiago, Chile are also used. An ESA Station near Kiruna, Sweden was used during the FAST 1997 ground coordinated observing campaign, and a new TOTS at McMurdo, Antarctica, became available in mid-1997. The primary ground stations send the data, rate buffered in real-time, over dedicated lines to the Packet Processing System (PPS, see Shi *et al.*, 1993) located at the Goddard Space Flight Center. Other stations record data on tape and mail the tapes back to the PPS.

The PPS is a semi-automated system developed to handle data for the FAST project. The PPS repackages and stores the telemetry as it is received. Once the contact is completed, it sorts the data by data type into separate files, ordering it in time, eliminating duplications, and removing untrustworthy packets. The resulting data set is transmitted to secure computers in the SOC over a dedicated network. Once the data has arrived at the SOC, it is checked for size and the result posted on the FAST Web page. This page is viewed by PPS personnel for comparison with their records to assure that the SOC receives all the data.

3.2. DATA PROCESSING AND ARCHIVE

Figure 3 shows a block diagram of the automated processing that is triggered once the data is received at the SOC. During this processing, the data is held on the secure network until enough data has been received to fill a CD-ROM (650 MBytes). A disk image is then built and transferred to an open data analysis net-

FAST Archive Data Flow

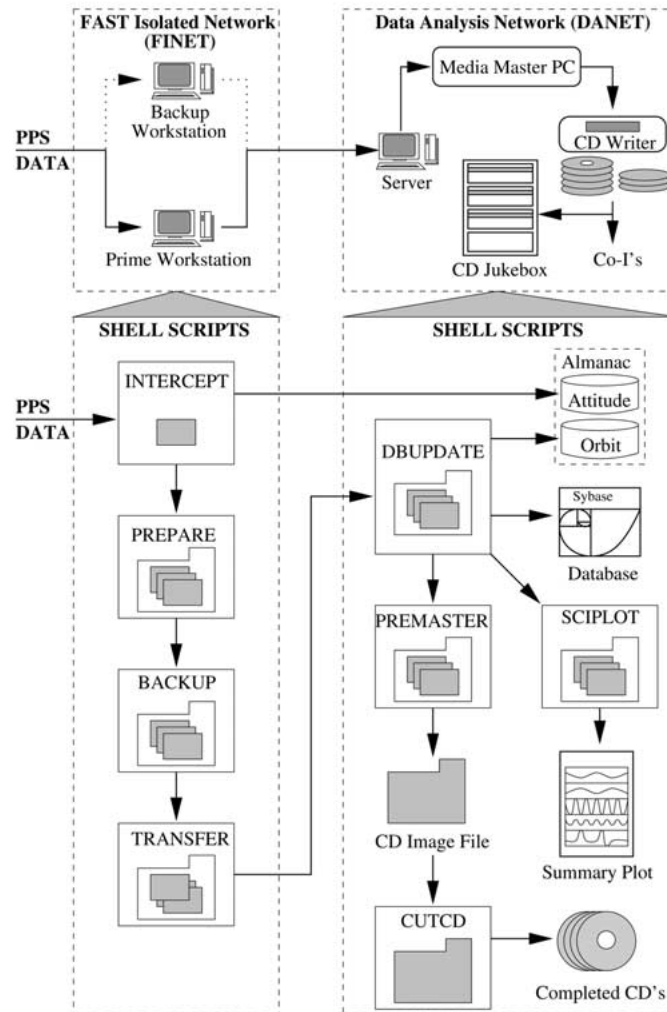


Figure 3. A block diagram of the data archive process. Scripts are run on both the Fast Isolated Network (FINET) and on the Data Analysis Network (DANET) to complete the archive. The INTERCEPT script senses a new data set has been received and processes the data generating a 'fastcount file' that is stored with the raw data. The fastcount file is used by data analysis tools to determine memory (RAM) allocation requirements for analyzing the data and by the database update (DBUPDATE) script. In addition, fastcount information is posted on the FAST Web page for internal use and to allow PPS personnel to check data volume received. INTERCEPT launches the PREPARE script which determines if the volume of data currently on FINET exceeds that of a CD-ROM (650 MBytes). INTERCEPT performs further processing as well as launching the BACKUP and TRANSFER scripts when a CD-ROM data set is ready. The PREPARE scripts generates a UNIX directory structure that will mirror the CD-ROM, BACKUP copies across the 'firewall' software that isolates the FINET and DANET networks. Automated scripts are run on DANET to complete the data archive once each CD-ROM file set is received. DBUPDATE updates the FAST database and the orbit-attitude almanac files. The PREMASTER script forms an 'image' (single file) of the cd. The SCIPLOT script generates summary plots of the data. The CUTCD script transfers the CD 'image' to a Young Minds Media Master which directs a CD writer to make the CD-ROMs.

work where the CD-ROMs are actually created. Typical contacts have data volumes that range from 50 to 300 MBytes, so multiple data sets are required to fill a CD-ROM. Routines are run on both of these networks in support of the archive process (see Figure 3). In the final step, a Young Minds Media Master (Young Minds Inc., Redlands, CA, <http://www.ymi.com>) directs a CD writer (Kodak CD Writer 600) to make the CD set (ISO 9660 with Rock Ridge Interchange and System Use Sharing Protocols that provides UNIX-style file and directory names with complete ISO 9660 compatibility). The CD writer includes robotics that hold 75 blank CDs and automatically loads and unloads CDs during the processing, and an image printer which labels the CDs after they are written. Eight CDs are created for distribution to the CoIs, to the NASA archive, and for loading into the on-line data archive at Berkeley.

The on-line data archive consists of an array of jukeboxes (Kubic model CDR-240) attached to a computer on the open data analysis network. Each jukebox holds 240 CDs and contains 4 CD readers which can be independently addressed by the jukebox software. The entire system of jukeboxes is controlled by a software package provided by Luminex (Fire Series Software, Luminex Software Inc., Riverside, CA., <http://home.luminex.com>), which runs the internal robotics, controls SCSI communications and handles disk access contention, making the system appear as a UNIX file system. CDs are entered or removed from the jukeboxes by programs developed at Berkeley that interface to the jukebox software. The jukeboxes are rather compact (36 in \times 17 in \times 7 in) and the entire on-line archive fits in two standard 19 inch racks.

Additional processes are run on the open network to complete the data archive once each CD-ROM file set is received. The primary task is to update the FAST database. The large volume of data files, over 1 million files organized by contact and data type, require a sophisticated database to keep track of files and to allow selection of data based on criteria other than time. A commercial product (Sybase, Sybase Inc., Emeryville, CA, <http://www.powersoft.com>) was chosen and a baseline set of tables were created to organize the database. Figure 4 shows the table structure and elements of each table. In addition to the tables containing file information, the database contains an events_data table with ephemeris, attitude, and model magnetic field. Events_data entries exist for all operational_events (science mode changes) and ephemeris_events (orbital highlights) entries. An operational_events entry is also generated every 20 s to assure minimal resolution in the database for searches based upon ephemeris. Science tables will be added to the database as they are developed to allow searches based upon auroral characteristics, such as energy flux, maximum wave power, and convection flow.

3.3. SUMMARY DATA

The data archive processing also includes the production of summary plots and key parameter data sets. The summary plot routine determines which orbits are

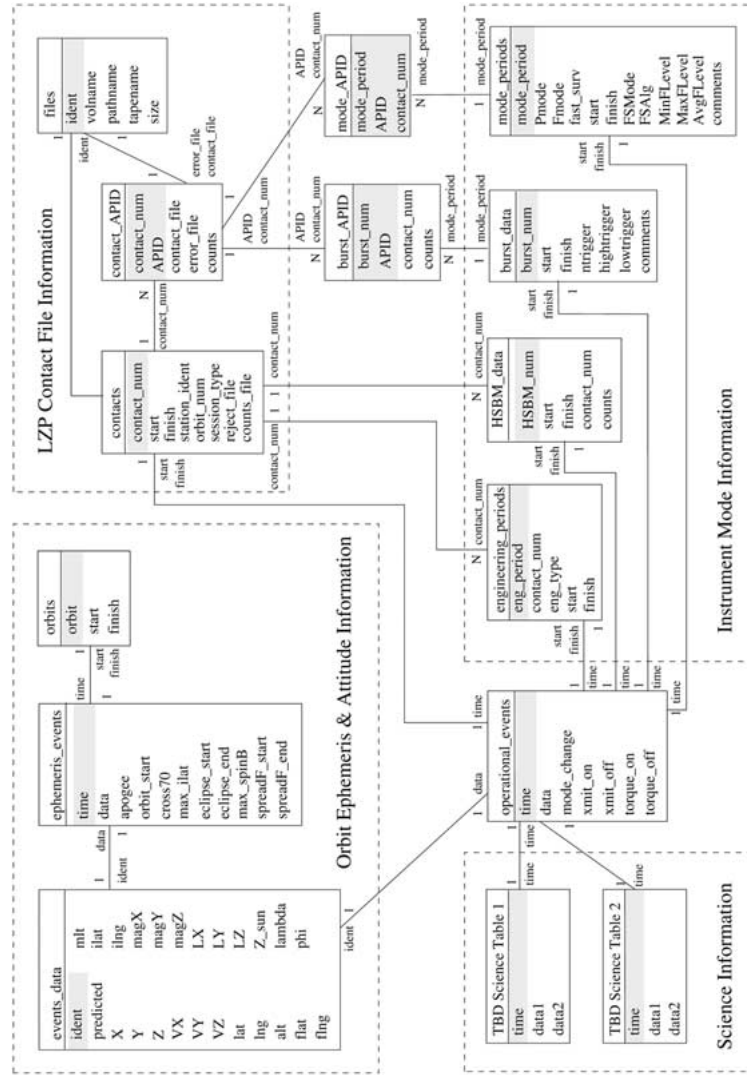


Figure 4. FAST database table structure and elements of each table. Sybase is used as the database platform and the database resides at Berkeley. Database tables are automatically updated as part of the data archive process. Scientists access the database primarily through queries generated by the data manager program. This program can be run locally or from a remote site to obtain a list of data files based upon data type, time, ephemeris, or spacecraft operating mode. The data manager program also interfaces to the data analysis program, automatically generating queries based upon the selected plot types, and can be used to download files from the on-line data archive at Berkeley. The database can also be accessed directly using the WISQL, a convenient user interface for entering database queries.

currently present in the CD data set, then launches a batch process that generates summary plots and CDF (Common Data Format, <http://nssdc.gsfc.nasa.gov/cdf>) files for each orbit. The batch process uses both the FAST data analysis program (SDT, Science Data Tool, developed by the FAST science team and discussed below) and IDL (Interactive Data Language, Research Systems Inc., Boulder, CO, <http://www.rsinc.com/idl>) to generate 1 spin average data sets for five summary plot data types: Ions, Electrons, Mass Spectrometer, AC Fields, and DC Fields. GIF plots are made and placed immediately on the FAST Web site for viewing. Plots cover 20-min periods prior and post the highest invariant latitude for each polar crossing during science data collection, so typically two or four GIF files are generated for each data type each orbit. In addition, 5 CDF files are generated containing summary data for each plot type, which can be downloaded from the FAST Web site. A description of the summary plot and CDF file contents is given in Table IV.

Key parameter data sets are produced for inclusion in the ISTP Key Parameter (KP) data base from the CDF files described above. Key parameter files are made by combining all (typically 11) of the single orbit CDF files of each data type for a day into one large file. Most collaborative studies, such as spacecraft or ground conjunctions, require only a single orbit to be downloaded over the net. Thus scientists may prefer to use the FAST Web page as a faster means of obtaining a single orbit's data. For those interested in obtaining complete key parameter data sets, we recommend using the NASA centers. ISTP KP values in CDF format can be accessed through the ISTP CDHF (<http://www-istp.gsfc.nasa.gov/istp/science/kpi>) and the NSSDC Spycat (<http://nssdc.gsfc.nasa.gov/space/ndads/spycat.html>) and CDAWeb (<http://CDAWeb.gsfc.nasa.gov>) systems for on-line retrieval and plotting.

3.4. EPHEMERIS AND ATTITUDE DATA

A final task of the data archiving procedures is the production of ephemeris and attitude data sets. To handle ephemeris data, an orbit propagator and field line tracing program was developed that uses a NASA supplied orbit propagator, a field line tracing program using the IGRF 1995 magnetic field model with secular variation (<http://www.ngdc.noaa.gov/IAGA/wg8/wg8.html>), and takes an orbit vector supplied by NASA (one per day) as an input. The program generates spacecraft GEI (Geocentric Earth Inertial) position and velocity, latitude and longitude, invariant latitude, Magnetic Local Time (MLT), and latitude and longitude of the magnetic footprint (at 100 km altitude). Rather than store ephemeris data with time resolution that allows linear interpolation, the FAST team decided to rely on a small ephemeris file which contains one vector (position, velocity in GEI coordinates) for each orbit. The orbit propagator program is then used to quickly calculate ephemeris data from the nearest time in the file.

Similarly, an attitude propagator program was developed to calculate spacecraft attitude from the spacecraft attitude housekeeping data and magnetometer data. An

Table IV
FAST summary plot and CDF file contents

EES – Electron Electrostatic analyzer Survey data
Electron Energy Spectra, 0–30 deg pitch angle
Electron Energy Spectra, 60–120 deg pitch angle
Electron Energy Spectra, 150–180 deg pitch angle
Electron Pitch Angle Distribution, 0.1–1.0 keV
Electron Pitch Angle Distribution, 1.0–30. keV
Electron Energy Flux, mapped along B to 100 km altitude
Electron Number Flux, mapped along B to 100 km altitude
IES – Ion Electrostatic analyzer Survey data
Ion Energy Spectra, 0–30 deg pitch angle
Ion Energy Spectra, 40–140 deg pitch angle
Ion Energy Spectra, 150–180 deg pitch angle
Ion Pitch Angle Distribution, 0.05–1.0 keV
Ion Pitch Angle Distribution, 1.0–25 keV
Ion Energy Flux, mapped along B to 100 km altitude
Ion Number Flux, mapped along B to 100 km altitude
TMS – Time-of-flight Mass spectrometer Survey data
Hydrogen Energy Spectra 0–360 deg pitch angle
Helium Energy Spectra, 0–360 deg pitch angle
Oxygen Energy Spectra, 0–360 deg pitch angle
Hydrogen Pitch Angle Distribution, 0.01–1.0 keV
Hydrogen Pitch Angle Distribution, 1.0–15 keV
Helium Pitch Angle Distribution, 0.01–1.0 keV
Helium Pitch Angle Distribution, 1.0–15 keV
Oxygen Pitch Angle Distribution, 0.01–1.0 keV
Oxygen Pitch Angle Distribution, 1.0–15 keV
ACF – AC Fields Survey data
HF Electric Field Spectra (0–2 MHz)
VLF Electric Field Spectra (0–16 kHz)
ELF Electric Field Spectra (0–2 kHz)
HF Magnetic Field Spectra (0–2 MHz)
VLF Magnetic Field Spectra (0–16 kHz)
ELF Magnetic Field Spectra (0–2 kHz)
HF Band Integrated Power
VLF Band Integrated Power
ELF Band Integrated Power
Spacecraft mode (slow survey, fast survey, burst)
DCF – DC Fields Survey Data
Electric field perpendicular to spin plane projection of B, roughly equatorward
Electric field along spin plane projection of B
Langmuir probe current
Delta B perpendicular to spin plane projection of B, roughly equatorward
Delta B \pm spin axis, roughly West
Delta B along spin plane projection of B
Spacecraft potential
Angle between spin plane and B

attitude almanac file is generated using multiple orbits. Later calls for attitude data use an attitude propagator to calculate the attitude from the nearest data point in the almanac file. The FAST science team took on the task of determining the spacecraft attitude, normally performed by NASA, in order to reduce project costs and since accurate attitude determination is required for precise magnetic field solutions.

4. Real Time Software

One of the central objectives of the FAST program was the inclusion of a 'campaign' in which real-time data analysis would be used to coordinate spacecraft operations with ground based observations. The need for a real-time data analysis tool precluded the typical development path where instrument ground station equipment (IGSE) used for prelaunch testing and commanding is separate from the data analysis programs. Instead a series of programs that run on the data analysis platform (Sun SPARC stations running the Solaris 2 operating system) were developed to interface with the spacecraft GSE and the data analysis tools, producing a single combined IGSE analysis package.

Since the data analysis platform was UNIX based, it was not configurable as an event driven platform such as DOS based PCs. Instead, networking techniques were used to connect to the spacecraft GSE and draw data to the data analysis platform. A feature of the Solaris 2 operating system (threads: coordinated, cooperating subprograms) was used to duplicate the fast response inherent in event driven systems. In addition, the same protocol for communication between the spacecraft GSE and the IGSE were adopted for the other ground systems (TOTS and PPS), providing a single interface for all real-time science operations. This software package allows the scientist to monitor spacecraft operations and display science quantities simultaneously.

4.1. IGSE TO GSE INTERFACE PROGRAM

A block diagram of the real-time programs is shown in Figure 5. The interface to the spacecraft GSE (or TOTS or PPS) is provided by a relay program. This program has the following functions:

- Real-time data input from the spacecraft GSE, TOTS, or PPS.
- Playback data from a disk file created in a previous session.
- Data capture to a disk file.
- Data relay to monitor workstations on a secondary network.
- Data relay to a remote workstation on a wide area network such as the internet.
- Data set monitoring including data volume, data rate, and data packet statistics.

Each of these functions are implemented as separate threads (subprograms) in the relay software. The actual data resides in a systemwide shared memory area,

FAST Real Time Software

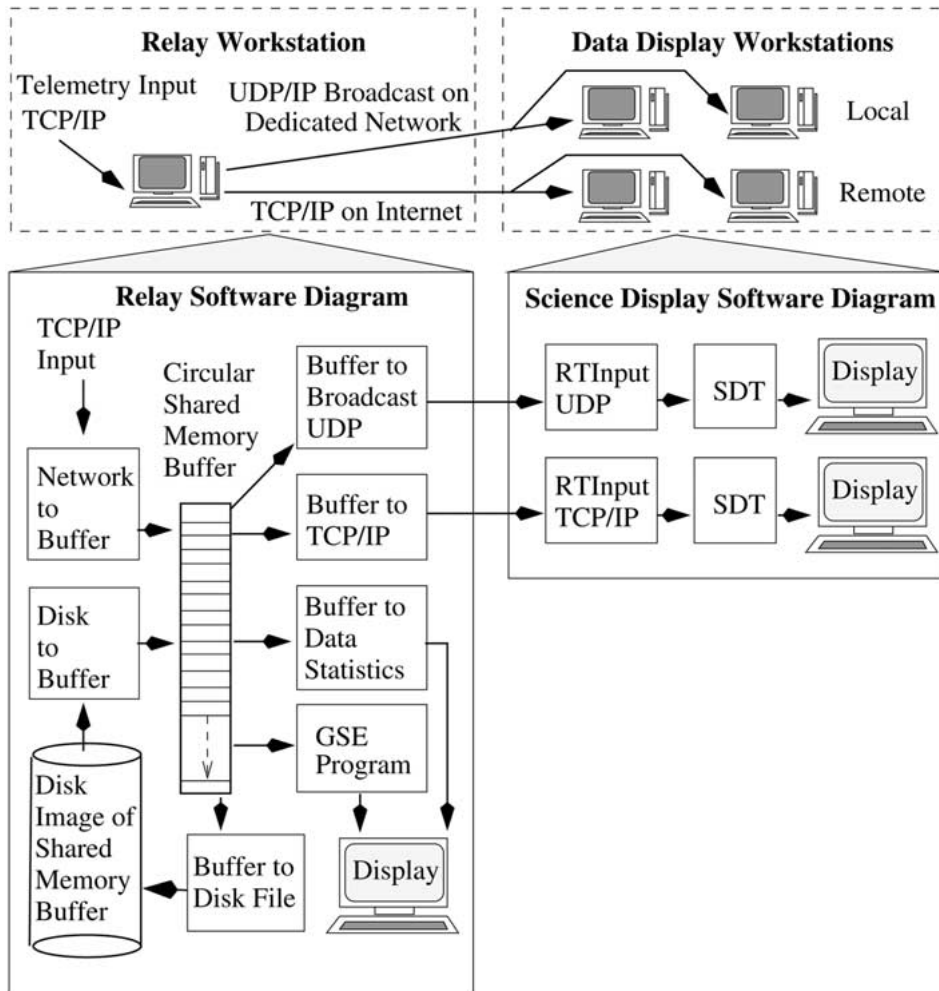


Figure 5. A block diagram of the real-time data capture and display programs. The relay workstation gathers telemetry from a NASA ground station into a circular memory buffer where the data can be copied to disk, broadcast to workstations on a local network, or transmitted over the internet to a remote workstation. Data statistics are displayed for the various data types, and a GSE program is usually run to display instrument housekeeping. Workstations attached to a local network or at remote sites can use the RTInput program to access selected subsets of the data for display using the FAST data analysis program, SDT.

which enables other programs to access the relay data flow. The data memory subsystem monitors how many threads are active, and when the last thread exits the shared memory buffer is removed.

The implementation of the shared memory subsystem is optimized for performance. All data flow functions are concurrent and independent, allowing the various computer subsystems, such as disk, network I/O and video to update in parallel. Access to the shared memory is not computationally intensive: threads reading the memory only acquire read locks on a single data block (approximately 1/1000 of the total shared memory) at one time, and the number of concurrent read locks on a given block are unlimited in number. Only threads writing to the memory acquire exclusive locks, and again only on one block at one time. A simple conditional locking scheme handles thread positioning and data availability within the data buffer.

A simple one character command interface is employed for easy control of all of the relay program's functions, and the display works entirely within a UNIX terminal window, so that the program can be efficiently executed from a remote system. The screen update rate can also be controlled via a start-up option so that slow network links to a remote display can work properly. The ability to remotely monitor the instruments prevented key personnel from being tied up for long periods at remote sites during instrument integration.

4.2. HOUSEKEEPING DISPLAY AND REMOTE MONITORING

One program designed to read the memory buffers of the relay program is the FAST GSE program which displays housekeeping data, thereby allowing real-time monitoring of instrument health. This multi-threaded program reads data packets containing housekeeping data and generates a multipanel display which can be configured to display different spacecraft health and safety data. The user interface allows the scientists to change individual panels to check on specific instruments, or choose from pre-configured combinations of displays geared for specific tests. The GSE program is most useful for monitoring instrument health and safety during instrument integration, launch and early orbit, and can also be used to diagnose instrument problems by replaying telemetry from a disk file.

The broadcast feature of the relay program allows other nearby computers to listen in on the captured data stream. In this way each individual computer need only decode and display a limited portion of the data stream. Computers connect to the broadcast network using a data capture program (RTInput, Figure 5) that opens a UDP (User Datagram Protocol) connection with the relay program. The relay program broadcasts different science or housekeeping data types on individual network channels, so that the other computers can receive a limited data set by listening on the corresponding set of channels.

The data capture program reads packets from the UDP connection and writes them into its own shared memory buffer with the same format as is written by the

relay program. Hence other programs, like the GSE program, can read these data as it would on the relay machine. The data capture program allows the user to select which data to receive from the data stream, and is usually used in conjunction with the FAST science display program, SDT. When a scientist selects data to be displayed in SDT, SDT automatically sends the data capture program a list of the data types which are used to configure the UDP connection. If a computer cannot keep up with the data it is trying to display, then UDP packets can be dropped without affecting the ability of any other computer to receive data.

The data capture program (RTInput, Figure 5) can also be used to store data to a disk file, to read a disk file, or to connect to a remote site over the internet. The disk read feature can be used to replay telemetry in near real-time, and is useful for diagnosing instrument performance or health and safety in conjunction with GSE and/or SDT programs. The scientist can select playback speed, pause, or step through the data forwards or backwards one telemetry packet at a time. The remote access capability allows computers at remote sites, not connected to the relay broadcast network, to connect and listen in on the data stream captured by the relay machine. This enables scientists to monitor their instruments without being present at the test site. In this mode, the program uses a TCP/IP connection to send the relay machine a list of data types it wants to receive. An independent thread (subprogram) running in the relay program then directs the requested subset of the data to the TCP connection. If the remote machine or the network cannot handle the data rate, the relay program will drop packets without affecting other functions of the relay machine.

4.3. REAL TIME SCIENCE DATA DISPLAY

The primary output interface of the data capture program is to the FAST data analysis program, SDT. SDT is a science display software package, developed at Berkeley initially for Geotail and upgraded for several other projects including CRRES, Polar and FAST. SDT presents a menu driven interface to the user that allows rapid selection of data quantities to be plotted, and quick changes to the plot attributes (plot limits, line or points, integration limits, labels) and display attributes (plot selection, zoom, pan, plots/page, print, color table, etc.). SDT can exhibit time-series (multi-line) and spectrogram data, can run as an off-line 'batch' program, and can be used to feed data into IDL programs. It can be used to view data either in post-orbit data analysis mode, or in real-time mode for prelaunch testing, instrument turnon, and the FAST campaign. In real-time mode, SDT uses a circular RAM buffer to store the most recent data for capture and display on its screen in wrap around mode (oscilloscope). The circular buffer prevents memory overflow during long viewing periods, but prevents the user from viewing all the data with arbitrary plot limit changes. Since it is written using XWindows, it can also be used remotely (over the internet) as well as on stand alone machines.

SDT Software Diagram

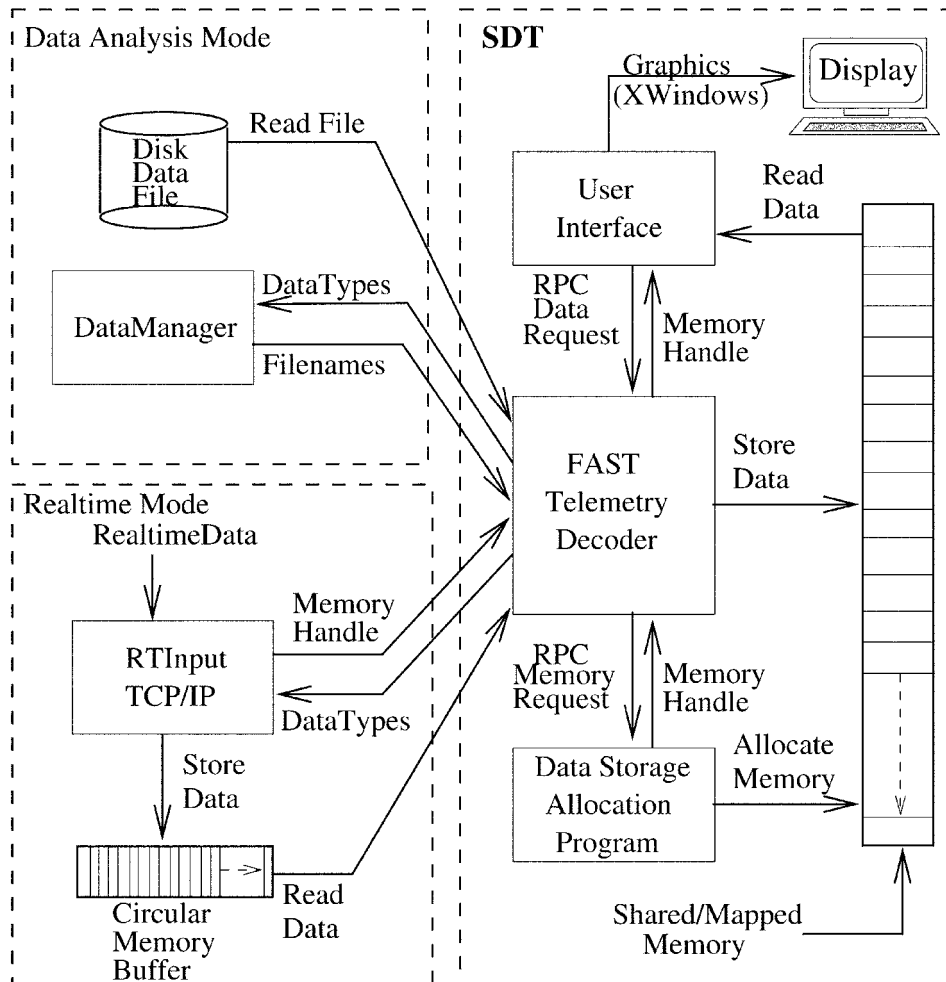


Figure 6. A block diagram of the real-time and data analysis program, SDT. The SDT program allows the same display software to be used for integration, launch and early orbit, and during the data analysis phase. SDT locates real-time data buffers through an interface with the RTInput program, or locates data files through an interface with the data manager program. SDT's telemetry decoder decommutates the data and stores it in memory in a format readily displayed by the user interface. The user interface allows rapid changes in data quantities to be plotted, and quick changes to the plot attributes (plot limits, line or points, integration limits, labels) and display attributes (plot selection, zoom, pan, plots/page, print, color table, etc.).

Data Analysis Software and Hardware

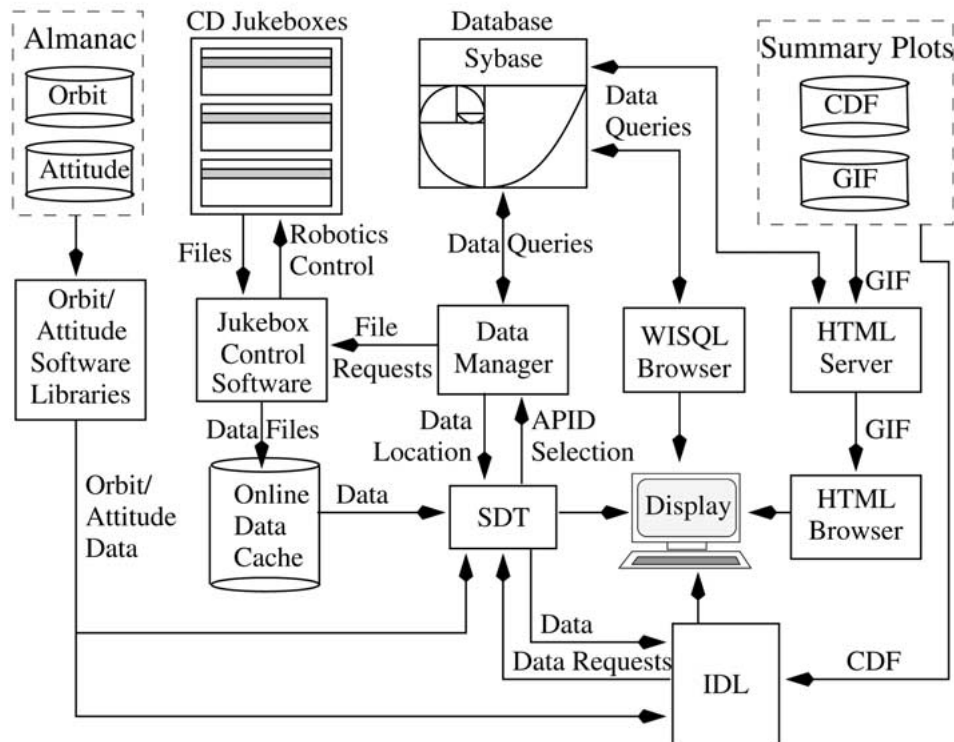


Figure 7. A block diagram of the data analysis software. The central programs are the data manager program and data analysis program, SDT. A typical data session would begin with the selection of science data plots within SDT followed by a call to the data manager to find a particular orbit or time. The data manager queries SDT to determine data types, accesses the database to determine file names, searches for the files on local hard disks, downloads those files not found from the data archive, and provides SDT with a list of file paths. SDT reads the files, decommutates the data placing it in memory, and plots the data on the fly. Once in memory, data can be further accessed through IDL. Browser programs are also available to provide database and key parameter searches.

A block diagram of SDT is shown in Figure 6. SDT consists of a user interface program, a data storage allocation program, and project specific programs, called 'decommutators' of which the FAST telemetry decoder is one, which decode the corresponding telemetry streams into data buffers which the user interface program can display. Use is made of multi-threading and various UNIX inter-process communication facilities such as Remote Procedure Calls, TCP Sockets, Shared-Memory, and Mapped-Memory files. For real-time FAST usage, data is read from the shared memory buffer which is written to by the data capture program (RTInput).

5. Data Analysis Software

FAST data are analyzed through a set of programs, both commercial and in-house, which allow the scientists quick and easy access to the data. Figure 7 illustrates the primary programs and hardware used by the FAST science team. Several browsing tools (WISQL Browser allows database queries, HTML browser allows GIF display of summary plots via FAST Web page, IDL displays CDF file summary data) are available to help select events for detailed investigation. Jukebox access software allows the user fast access to the entire data set. Once events are selected, the data can be loaded into RAM by the FAST data analysis tool, SDT, for detailed viewing as time series, or the data can be further accessed by IDL for detailed analysis. These programs and tools are outlined below.

5.1. INTERFACE TO ON-LINE DATA STORAGE

The FAST data is archived on CD-ROMs and stored on-line in jukeboxes (see Data Archive section). A commercial software package (Luminex Fire Series Software) interfaces to the jukeboxes making the system appear as a UNIX file system. Data analysis software requires that the files be downloaded to a local hard disk before being accessed by SDT or IDL. The database (Sybase) keeps information about which CD-ROM volume a data file resides on, and also keeps information about what data is stored in each file, such as timespan, instrument modes, and ephemeris data. Scientists can write SQL (Standard Query Language) queries to find interesting data, however, the primary tool for finding data is the data manager program described below. If specific data queries are required, the WISQL program provides a useful platform for initiating queries or browsing the database. (WISQL was downloaded from the web as part of 'sybtcl', an interface between Sybase and TCL developed by Tom Poindexter, <http://www.nyx.net/~tpoindex/tcl.html#Sybtcl>).

The data manager program provides the primary interface between the database, the jukeboxes, the data stored on-line on hard disk drives, and SDT. A scientist using SDT can select the data quantities he wants to analyze, and then select a timespan or orbit number in the data manager program. This program makes a database query for the files and CD-ROM volumes needed. The query can be further restricted by instrument mode, data type (application identifier or APID in Figure 7), and invariant latitude. The program obtains a list of data types from SDT before it makes the query, so that only the files needed by the scientist for the SDT session will be retrieved. When the data manager program receives the query results, it locates the files on hard disks, or allows the scientist to download the files from the jukebox. Upon completion of file downloading, the program sends SDT a list of file locations, and SDT decodes the data. Each scientist has a personal disk drive to which data can be downloaded and the data manager program keeps track of data files stored on all disks on the local network. The intermediate buffering of data files to the user's local hard disks improves the response time of the analysis

programs for any data accessed more than once by the entire group (as is typical for interesting data), but also prevents any single user from overwhelming the system by requiring that the individual remain within his or her disk space limits.

Co-Investigators can use the data manager program at remote sites to download data they need from the jukeboxes at Berkeley via the internet. They can also use it to access their own jukebox, which could be loaded with a subset of the FAST data set. In both cases, they access the Berkeley database using Sybase client software bundled with the data manager program, so that it is not necessary to purchase and maintain Sybase for sites other than at Berkeley.

5.2. SCIENCE DATA DISPLAY

Once files are located by the data manager program, SDT loads the data for display. Operation of SDT in data analysis mode is much like its operation in real-time mode, except that the data set is loaded into a linear memory buffer rather than the real-time circular buffer. A 'fastcount file', generated as part of the archive process and stored with the data, is used to determine memory requirements. This use of memory makes the program's response to changes in plot parameters quite fast, but requires the user's machine to have a large amount of RAM (> 100 MBytes) to view most data sets. When the data are available on local hard disk, SDT makes an excellent browsing tool. However, a few minute wait is initially required (depending upon the data volume) if downloading from the jukebox.

SDT can also be used as the first stage of data processing. Its 'batch' mode is used to create summary plots of housekeeping and basic instrument operation. For more advanced analysis, SDT can make data available to an IDL Batch session or produce disk files, which can then be read into other programming environments. SDT also has the ability to perform advanced analysis internally through the development of 'Science Modules'. These are subprograms that attach to SDT for the purpose of creating new data quantities (such as despun electric and magnetic field data) from one or more of the raw data quantities. Science modules must be written in C (or C++) and SDT has the appropriate interfaces and libraries to allow the addition of these programs.

In addition to viewing data using SDT, scientists require a data analysis tool that allows them to easily access and manipulate the raw data. To permit this, an interface between data stored by SDT and IDL was developed. The SDT to IDL interface consists of a library of C and C++ routines, and a group of IDL routines which extract the various generic SDT data types. These are used as the building blocks for data specific access routines which exist for each of the FAST data types. Once the data is extracted to IDL, the scientist is free to manipulate the data in a simple environment. The IDL programs are used in a command line environment rather than a menu-driven, point and click, program. Menu-driven programs are, in general, limiting since they require the developer to foresee any manipulation the user will consider. Since this is an impossible task, FAST scientists operating in

the IDL environment are expected to learn the general data analysis tool routines. In practice this method can be much faster than menu driven programs since the user can quickly develop a 'crib-sheet' of custom commands and use cut and paste to execute multiple commands with a few clicks of a mouse.

5.3. DATA VIEWING ON THE WEB

The data analysis effort is also supported by the FAST web site (<http://sprg.ssl.berkeley.edu/fast>). This site was created to provide the science team quick, timely access to numerous information products generated during the FAST mission and to provide the worldwide space physics community access to science data and general information on the mission. In general this information ranges from specific operational items (stored command loads, ground station contact schedules, orbit information, spacecraft and instrument status, instrument modes, and FAST documentation) to FAST science summary data which can be viewed and downloaded by the user. The science summary data link presents the user with a form which allows him or her to build a query to search for specific summary data by date/time or orbit. The user can further define the query for data type and auroral region. The user is presented with links to GIF, Postscript, and CDF files matching the query information. The FAST FTP site also contains the up-to-date data analysis and computer system software releases available to the FAST science team and associated support documentation.

6. Conclusion

The Science Operations Center for the FAST mission was developed to process and quickly route FAST data to the scientists, and to serve as an on-line data archive center for the project. The SOC also serves as an interface between the FAST investigators and NASA ground operations, providing daily command loads to spacecraft that optimize the science data collection and supporting near real-time commanding during spacecraft testing. A low cost operations center was developed by relying on automated programs and the latest available hardware. Software was developed that provides a smooth transition from pre-launch ground testing to post-launch data analysis. Extensive data analysis tools are available that allow the scientists to quickly access the large volume of data expected over the mission life (>1 terabyte). It is hoped that the above description not only provides an introduction to the FAST mission for those scientists wishing to collaborate in the future, but also serves as a general outline on how to design a low cost, science operation center with fast, convenient data analysis tools.

Appendix

A/D	Analog to Digital
ALT	Altitude
APID	Application Identifier
CD-ROM	Compact Disc - Read Only Memory
CDAWeb	Coordinated Data Analysis Workshops Web site run by NSSDC
CDF	Common Data Format
CDHF	Central Data Handling Facility
CRRES	Combined Release and Radiation Effects Satellite
FAST	Fast Auroral Snapshot Explorer
FIFO	First In First Out
FTP	File Transfer Protocol
GEI	Geocentric Earth Intertial coordinate system
GIF	Graphics Interchange Format
GSE	Ground Station Equipment
HSBM	High Speed Burst Memory
HTML	Hyper Text Markup Language
IDL	Interactive Data Language software from Research Systems Inc.
IGRF	International Geomagnetic Reference Field
IGSE	Instrument Ground Station Equipment
ILAT	Invariant Latitude
ISTP	International Solar-Terrestrial Physics
KP	Key Parameter
MLT	Magnetic Local Time
NASA	National Aeronautics and Space Administration
NSSDC	National Space Science Data Center (NASA/GSFC Code 633 Facility)
POCC	Payload Operations and Control Center
PPS	Packet Processing System
RAM	Random Access Memory
RTInput	Real Time Input program developed by FAST science team
SCSI	Small Computer System Interface
SDT	Science Data Tool
SMEX	Small Explorer Missions
SOC	Science Operations Center
SQL	Structured Query Language
TCL	Tool Command Language
TCP/IP	Transmission Control Protocol / Internet Protocol
TEAMS	Time-of-flight Energy, Angle, Mass Spectrograph
TOTS	Transportable Orbital Tracking Station
UDP	User Datagram Protocol
UT	Universal Time
VC	Virtual Channel
WISQL	Windowed Interactive Structured Query Language

References

- Carlson, C. W., McFadden, J. P., Turin, P., Curtis, D. W., Magoncelli, A.: 2001, 'The Electron and Ion Plasma Experiment for FAST', *Space Sci. Rev.* (this issue).
- Elphic, R. C., Means, J. D., Snare, R. C., Strangeway, R. C., Kepko, L., and Ergun, R. E.: 2001, 'Magnetic Field Instruments for the Fast Auroral SnapshoT Explorer', *Space Sci. Rev.* (this issue).
- Ergun, R. E., Carlson, C. W., McFadden, J. P., Mozer, F. S., Delory, G. T., Peria, W., Chaston, C. C., Temerin, M., Roth, I., Muschietti, L., Elphic, R., Strangeway, R., Pfaff, R., Cattell, C. A., Klumpar, D., Shelley, E., Peterson, W., Möbius, E., and Kistler, L.: 1998, 'FAST Satellite Observations of Large-Amplitude Solitary Structures', *Geophys. Res. Lett.* **25**, 2041–2044.
- Ergun, R. E., Carlson, C. W., Mozer, F. S., Delory, G. T., Temerin, M., McFadden, J. P., Pankow, D., Abiad, R., Harvey, P., Wilkes, R., Primbsch, H., Elphic, R., Strangeway, R., Pfaff, R., and Cattell, C. A.: 2001, 'The FAST Satellite Electric Field and Magnetic Field Instrument', *Space Sci. Rev.* (this issue).
- Harvey, P. R., Curtis, D. W., Heeterds, H. D., Pankow, D., Rauch-Leiba, J. M., Wittenbrock, S. K., and McFadden, J. P.: 2001, 'The FAST Spacecraft Instrument Data Processing Unit', *Space Sci. Rev.* (this issue).
- Klumpar, D. M., Möbius, E., Kistler, L. M., Popecki, M., Hertzberg, E., Crocker, K., Granoff, M., Tang, L., Carlson, C. W., McFadden, J., Klecker, B., Eberl, F., Kuneth, E., Kastle, H., Ertl, M., Peterson, W. K., Shelley, E. G., and Hovestadt, D.: 2001, 'The Time-of-Flight Energy, Angle, Mass Spectrograph (TEAMS) Experiment for FAST', *Space Sci. Rev.* (this issue).
- NASA Document No. 822-DEM-1724, System, Tracking, and Acquisition Handbook for the Transportable Orbital Tracking System (TOTS), 1996.
- Shi, J., Mao, T., Chesney, J., and Speciale, N.: 1993, 'Fast Auroral SnapshoT Explorer (FAST) packet processing system', *Proc. Inter. Telemetry Conf.*, 449–459.