MARILYN D. SEMPRUCCI

The First "Switchboard In the Sky": An Autonomous Satellite-Based Access/Resource Controller

Satellite communications systems can connect hundreds of dispersed users instantaneously. But when such a system is serving many small, widely dispersed users and is handling fluctuating traffic loads, the satellite channels must be allocated dynamically. A satellite-based access/resource controller, a "switchboard in the sky," has long been seen as the most efficient way to use the costly resources of a satellite. Two satellite communications packages, which will provide highly protected links to small mobile terminals, incorporate the first autonomous satellite-based switchboard.

Satellite communications offer a compelling advantage for military and civilian applications — instantaneous connectivity among hundreds of widely dispersed and mobile users. Military users, for example, use satellite communications to link ships, aircraft, submarines, even individual soldiers. Civilians communicate around the world and with remote villages via satellite.

But current satellite communications technologies have several serious deficiencies. Chief among these, for the military, is an extreme vulnerability to deliberate interference, "jamming." And both military and civilian users need a solution to the Demand Assigned Multiple Access (DAMA) problem.

When a satellite communications system is serving many small, widely dispersed users with widely fluctuating traffic loads, satellite channels must be allocated in a highly dynamic way (to utilize most efficiently the very expensive communications capacity of the satellite). This is the DAMA problem. The DAMA service available for today's satellite communications technologies is not totally satisfactory. A satellite-based access/resource controller, or "switchboard in the sky," has long been seen as the ideal solution to the DAMA problem.

In a new type of satellite communications package, called FEP (FLTSAT EHF Package), we have addressed the problems of jamming, DAMA, and other military-specific problems of satellite communications. The FEPs ride on two FLTSAT satellites, FLTSAT-7 and FLTSAT-8. These communications packages can provide highly protected satellite communications to many small mobile terminals at extremely high frequency (EHF).

The FEP features a true "switchboard in the sky" function. To our knowledge, it is the first of its kind.

CURRENT SATELLITE SYSTEMS

A typical communications satellite uses transponders, which simply receive uplink signals, amplify them, and retransmit them at a different frequency on the downlink. The use of different uplink and downlink frequencies prevents the satellite's transmitter from jamming its incoming signals. By using multiple access techniques, multiple terminals can share the communications resources of each of a satellite's transponders.

Two common types of multiple access are frequency-division multiple access (FDMA) and time-division multiple access (TDMA). In an FDMA system, the satellite's receiving frequency band is divided into narrower sub-bands (channels). Each user's terminal transmits data over a carrier frequency within its assigned channel. Data are transmitted by modulating the carrier frequency. In this way, the frequency "carries" the data. Since FDMA systems have multiple channels, they use multiple carrier frequencies.

A TDMA system uses only one carrier frequency. Access to the system is based on a time period or data frame, which is divided into timeslots. Terminals are assigned particular timeslots, during which they can transmit and/or receive data. The TDMA scheme requires users to be synchronized in time — so they know when their opportunities to transmit and receive data occur.

The FDMA frequency channels and TDMA timeslot channels make up a communications resource pool that must be allocated to user communications links. The allocation process, therefore, must include a mechanism for informing each terminal of its channel assignment.

The communications resources of satellites can be assigned statically or dynamically. Statically assigned satellites either have fixed channelization, which locks in spacecraft resources without regard to traffic conditions or are reconfigurable via commands from a central ground-control station (Fig. 1). Reconfigu-

ration from a ground-control station is relatively slow, but this method provides adequate service when traffic is fairly constant.

When communications traffic fluctuates widely, channel assignments must be made dynamically — a DAMA system is required. A DAMA system allocates and releases communications channels in response to varying user traffic demands. When the level of communications traffic is intermittent, a DAMA system can support many more users than a statically assigned system.

All DAMA systems have a control channel, which terminals use to request communications service and receive channel assignments. In a ground-based DAMA system using a transponding satellite, one terminal can act as a central controller or a distributed control protocol can direct terminals to available channels. The SPADE (Single Channel Per Access Demand Equipment) system currently in use on INTELSATs (International Telecommunications Satellites) uses a distributed control pro-
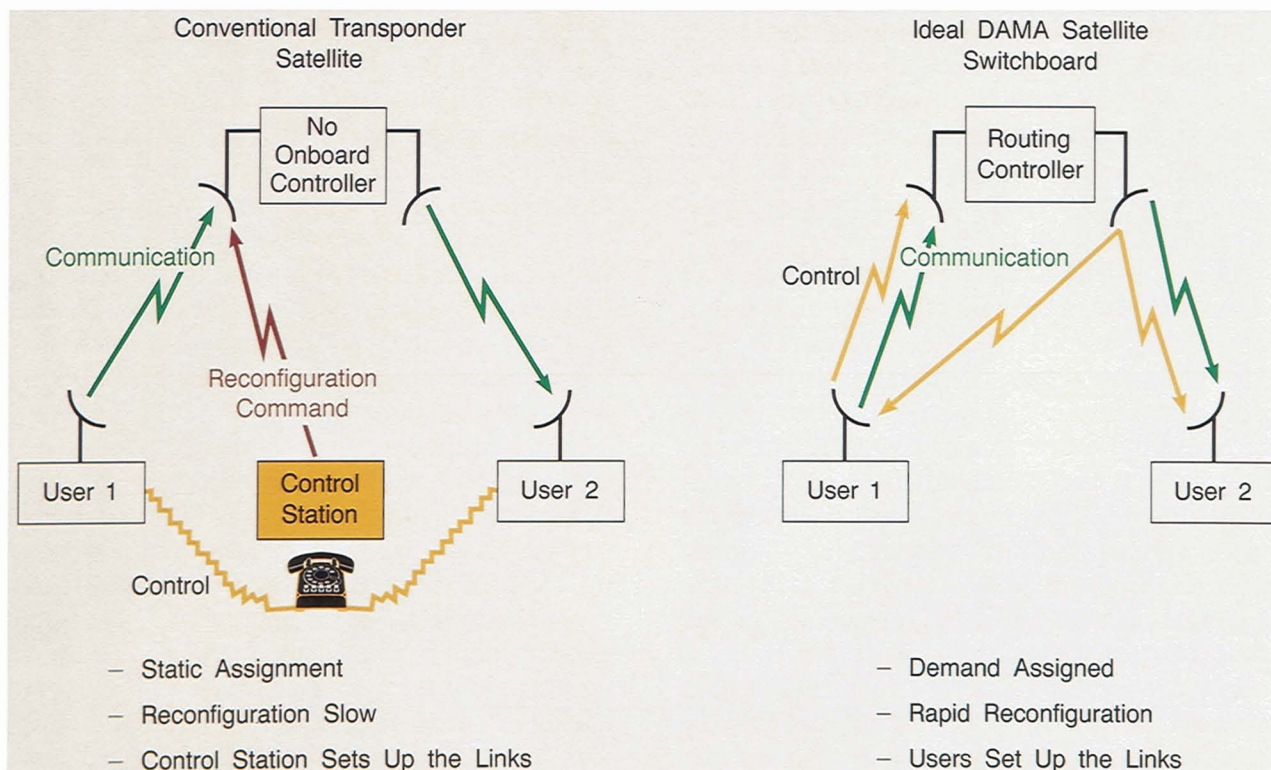

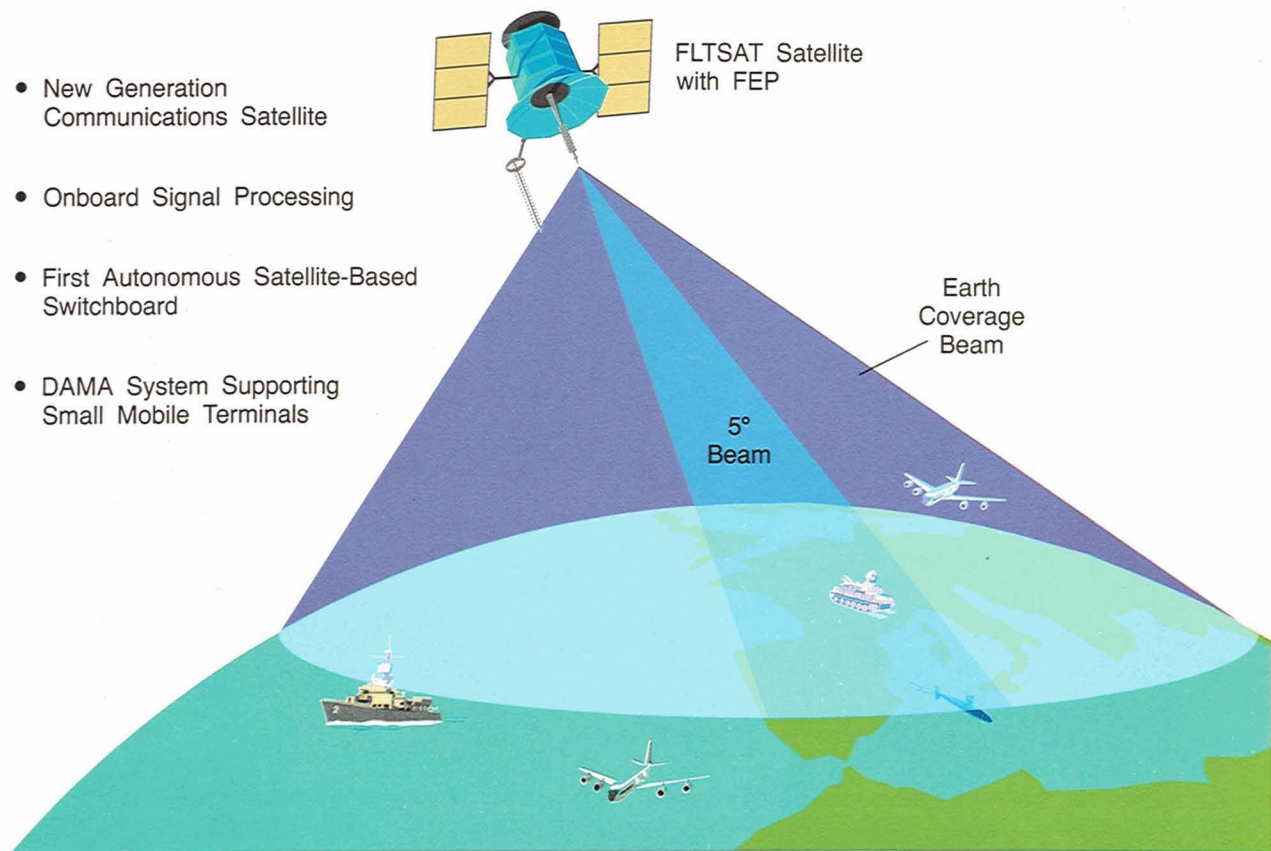Fig. 1 — Access control techniques for a satellite communications system.

- New Generation Communications Satellite

- Onboard Signal Processing

- First Autonomous Satellite-Based Switchboard

- DAMA System Supporting Small Mobile Terminals

FLTSAT Satellite with FEP

Earth Coverage Beam

5° Beam

Fig. 2 — FEP characteristics.

tocol. Central control is appropriate when allocating channels of different capacities. Such systems serve civilian needs quite well today, as long as all user terminals are in the antenna beam of the same satellite.

But if users are widely dispersed (as mobile users often are), multiple-beam systems make good economic sense. The most precious resource on a satellite is transmitter power. By using multiple narrow beams, transmitter power can be better concentrated on user terminals. As the traffic varies, however, a satellite's capacity must be allocated among multiple beams. In the transponding satellites in use today, entire transponders must switch among beams. But switching an entire transponder as traffic varies is inefficient, because it switches a large block (many channels) of a satellite's capacity.

In contrast with the transponding satellite,

the ultimate DAMA satellite is a "switchboard in the sky": a satellite that can quickly allocate individual circuits among its multiple beams in response to direct user requests (in the same way that a central telephone switchboard responds to user requests). In a satellite switchboard system, user terminals send control messages directly to the satellite (Fig. 1). The control messages may request communications service (much like dialing a telephone number). When a user terminal dials a link (by sending a control message), the switchboard allocates spacecraft resources; when the user terminals hang up (by sending another control message), the resources are released and made available to other terminals.

Until recently, such a satellite-based switchboard was infeasible. Now, it is not only feasible, we have implemented it in FEP and it has been functioning well on-orbit since December
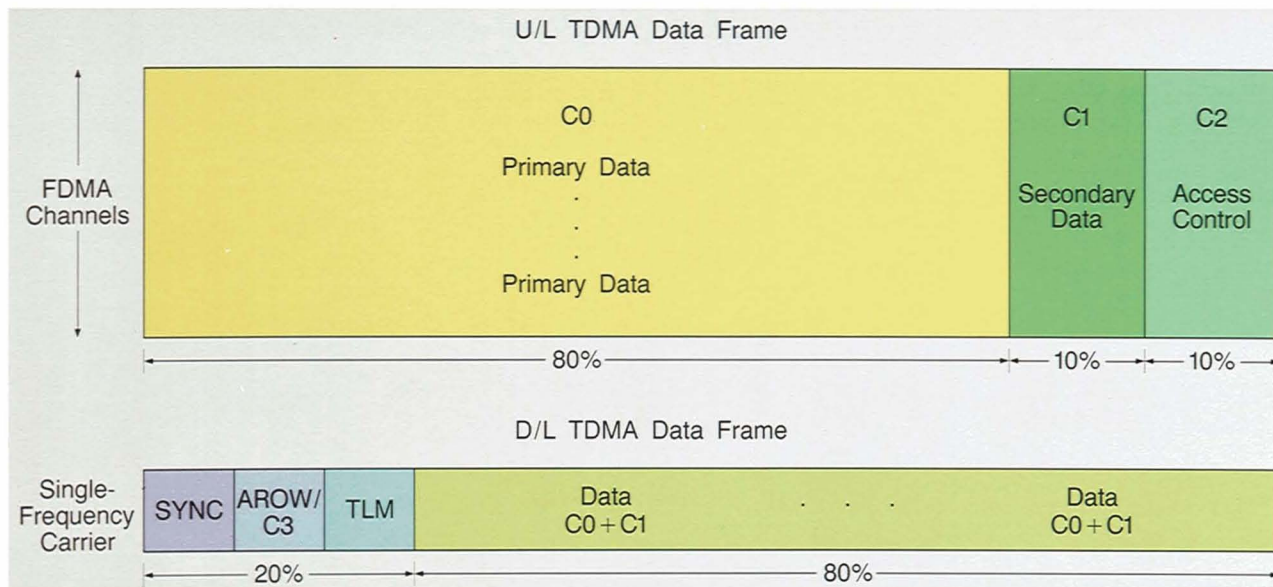
Fig. 3 — Channel resources (U/L = Uplink; D/L = Downlink).

1986. This switchboard has provided communications service to many small mobile terminals. And these terminals have both widely varying antenna sizes and widely varying communications needs.

FEP CHARACTERISTICS

A signal-processing satellite (unlike a transponding satellite) separates all user channels on the uplink, completely demodulates each channel, and repackages the channels on the downlink, usually in a different modulation format. The primary reason for this procedure is to separate jamming signals from valid user signals. However, this procedure gives the satellite the ability to control and assign communications resources on an individual channel basis. Therefore, an onboard signal-processing capability makes a satellite-based switchboard feasible.

The FEP is a signal-processing satellite. It has two beams; its communications links are configurable upon request. Unlike transponding satellites, which must switch large blocks of channels when switching capacity between beams, FEP can switch individual channels between beams. Data from an uplink channel can be routed on the downlink to the same beam, to the other beam, or to both beams.

Like all communications satellites, FEP's transmitter power is limited. But the amount of downlink resources allocated to a particular communication link can be tailored to fit the needs of individual users. The FEP supports widely different sized terminals and, since smaller terminals consume more downlink resources than larger terminals (because of the small size of their receiving antennas), FEP makes efficient use of its transmitter power.

As shown in Fig. 2, the FEP's beams are an earth-coverage beam and a 5° steerable spot beam. The spot beam illuminates an area approximately 2,000 miles wide. The spot beam can be steered via control messages sent directly from privileged user terminals to FEP's satellite-based switchboard.

Although FEP has only two beams, the principles demonstrated apply equally well to systems with many more beams. Each beam in FEP has a 44-GHz receiver. The uplink is frequency hopped for anti-jamming protection. It uses a combination of FDMA and TDMA techniques. The uplink frequency bands are divided into FDMA channels — 26 can be used for communications links.

8

Each of the 26 uplink FDMA communications channels is divided into three types of TDMA timeslots: C0, C1, and C2 (Fig. 3). The C0 timeslots are for user primary communications data. C1 timeslots are for user secondary communications (at low rates). The C2 timeslots are used for control messages that are sent from user terminals to the satellite-based switchboard. The C2 portion of each uplink channel is subdivided into additional timeslots, called "accesses," which let up to eight terminals share the C2 portion of one uplink channel.

The FEP's downlink is TDMA using a single hopped frequency carrier in the 20-GHz band. Data are transmitted over the earth coverage beam or the 5° spot beam on a timeslot-by-timeslot basis. Several timeslots on the downlink are fixed; they are devoted to such overhead functions as synchronization (SYNC) signals, telemetry (TLM) data, and access control (AROW/C3) messages (Fig. 3). The remaining timeslots (80% of the downlink) are available

for assignment via the switchboard for user communications links.

Each of FEP's 26 uplink communications channels can support up to 2,400 bps (vocoded voice) communications. Additional uplink channels are dedicated to command and control of the FEP itself. To provide communications service to the maximum number of users, FEP's uplink and downlink processors are configurable in a variety of modes, which are all under the control of FEP's autonomous switchboard.

OPERATION

Figure 4 shows the types of signals that pass between FEP and the user terminals. FEP is continuously transmitting a synchronization signal (1) in fixed timeslots on the downlink. At its turn-on, a user terminal must look for the synchronization signal and align its timing to match the satellite's TDMA cycle timing. This process is called "acquiring the downlink." At
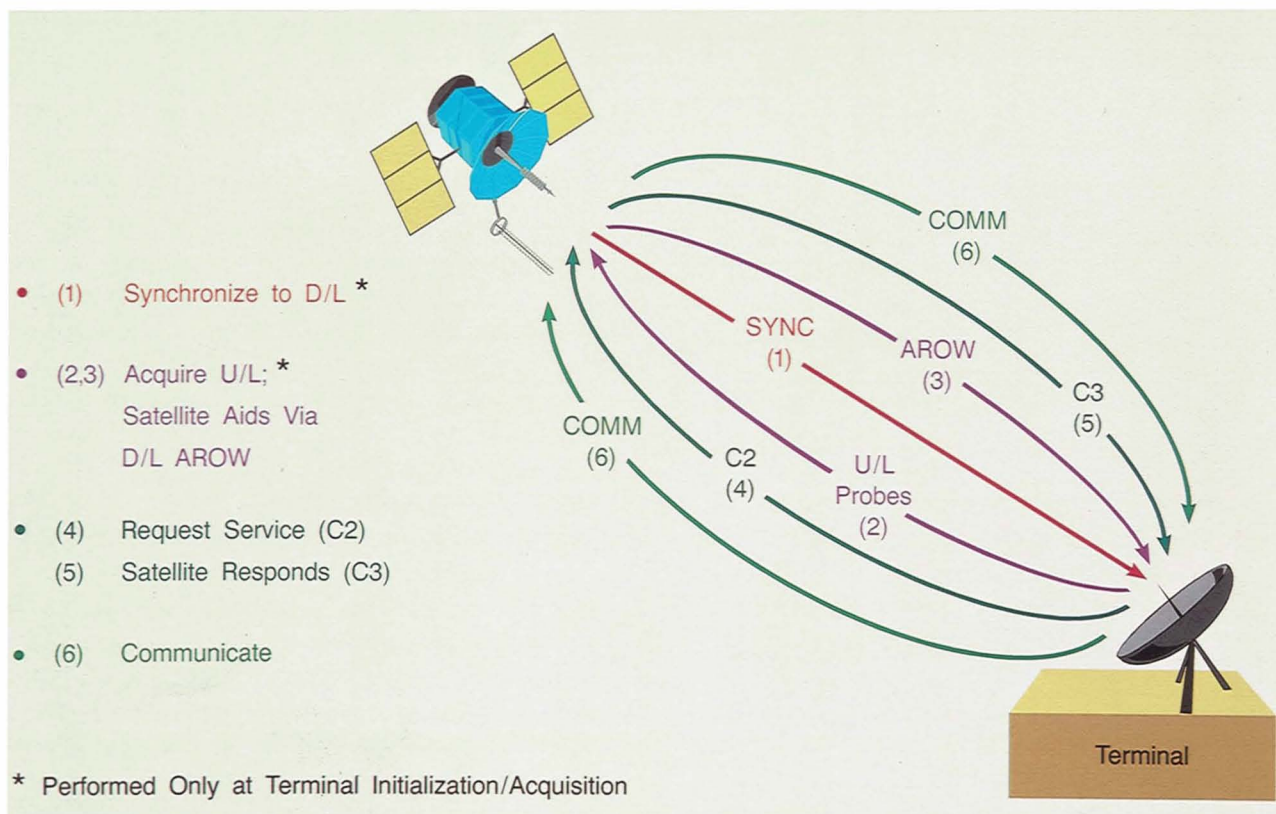

- (1)   Synchronize to D/L *

- (2,3) Acquire U/L; *
        Satellite Aids Via
        D/L AROW

- (4)   Request Service (C2)
- (5)   Satellite Responds (C3)

- (6)   Communicate

* Performed Only at Terminal Initialization/Acquisition
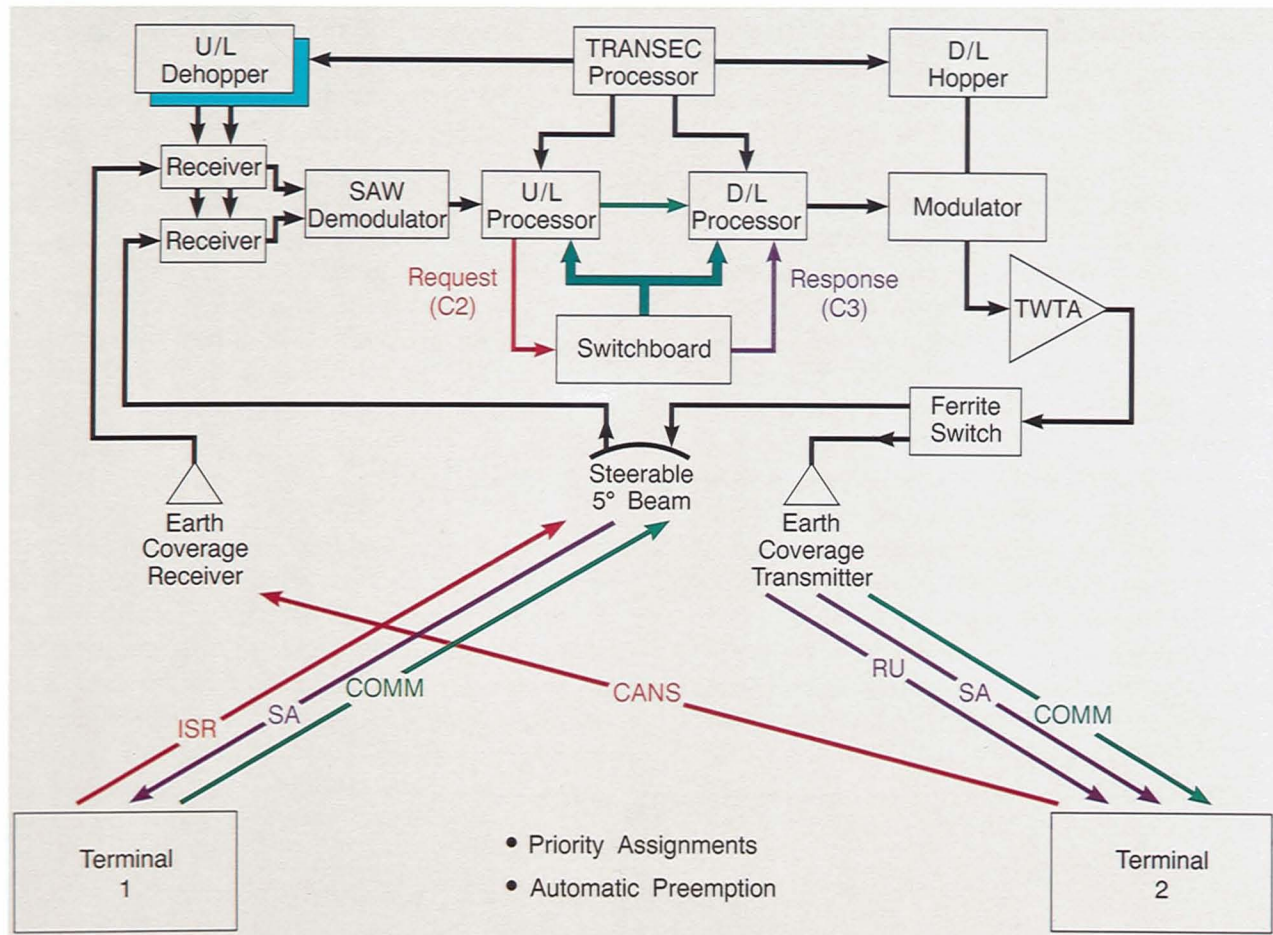
Fig. 4 — FEP signaling procedures.

Fig. 5 — The FEP allocates and releases channels as user traffic varies.

this point, the user terminal is able to be a passive listener on communications links. However, if the terminal wants to be an active participant in communications links, it must also align its uplink timing. To do this, it sends uplink probes (2) in a preassigned uplink channel and timeslot reserved for acquisition. The timing acquisition processor on FEP (which is contained in the uplink processor) responds with AROW control messages (3) in fixed downlink timeslots, telling the terminal if it is early or late with its probes. The terminal then adjusts its uplink probes and tries again. Once a terminal's uplink timing is correct, the timing acquisition processor on FEP notifies the switchboard that a terminal is acquiring. The switchboard communicates with the user terminals via C3 control messages (5) sent in fixed timeslots on the downlink. Following acquisition,

the switchboard sends a C3 control message to the terminal and assigns the terminal an uplink FDMA frequency channel and C2 TDMA timeslot. Now the terminal is enrolled in (logged on) the satellite.

An enrolled terminal communicates with the switchboard via C2 control messages (4) transmitted in its assigned uplink C2 channel and timeslot. For example, the terminal can send a C2 message up to the switchboard and request a communications link. The switchboard makes the necessary connections in the satellite and then sends the terminal a C3 message that specifies the uplink frequency channel and downlink timeslot assignments. The terminal can then begin communicating (6). These transactions are computer-to-computer; the protocol is transparent to the terminal operator and takes only a few seconds.

CONNECTIVITY AND CONTROL

The block diagram in Fig. 5 shows the overall design of the FEP. Data enter the FEP through either the earth-coverage receiver or the 5° spot beam receiver. The data are dehopped, demodulated and passed through the signal-processing package, where the data are rerouted, rehopped, remodulated and retransmitted over the 5° spot beam, the earth-coverage beam, or both beams. The uplink and downlink processors contain tables that determine which channels' data will be processed and how to process the data. The switchboard is a microprocessor that controls these tables via an IEEE-488 parallel data bus. If the IEEE-488 bus malfunctions, a slower backup serial bus can be activated.

The switchboard activates channels in the uplink processor and specifies the data rate and other mode information for processing data from the active channels. Until an uplink channel is activated by the switchboard, the uplink processor throws away data from that channel. Once a channel is activated, its C0 and C1 data pass to the downlink processor, where they are stored while they wait for transmission on the downlink. The switchboard then activates downlink timeslots for data transmission and assigns to the downlink processor a beam and a downlink burst rate for transmitting the data. By devoting more downlink timeslots to repeating data, a communication link can be made more robust on the downlink. Thus the switchboard connects uplink frequency channels to downlink timeslots. Until the switchboard makes this connection, the signal processor discards data from these uplink channels.

Figure 5 shows an example of how a terminal interacts with the switchboard to place a point-to-point telephone call to another terminal. The
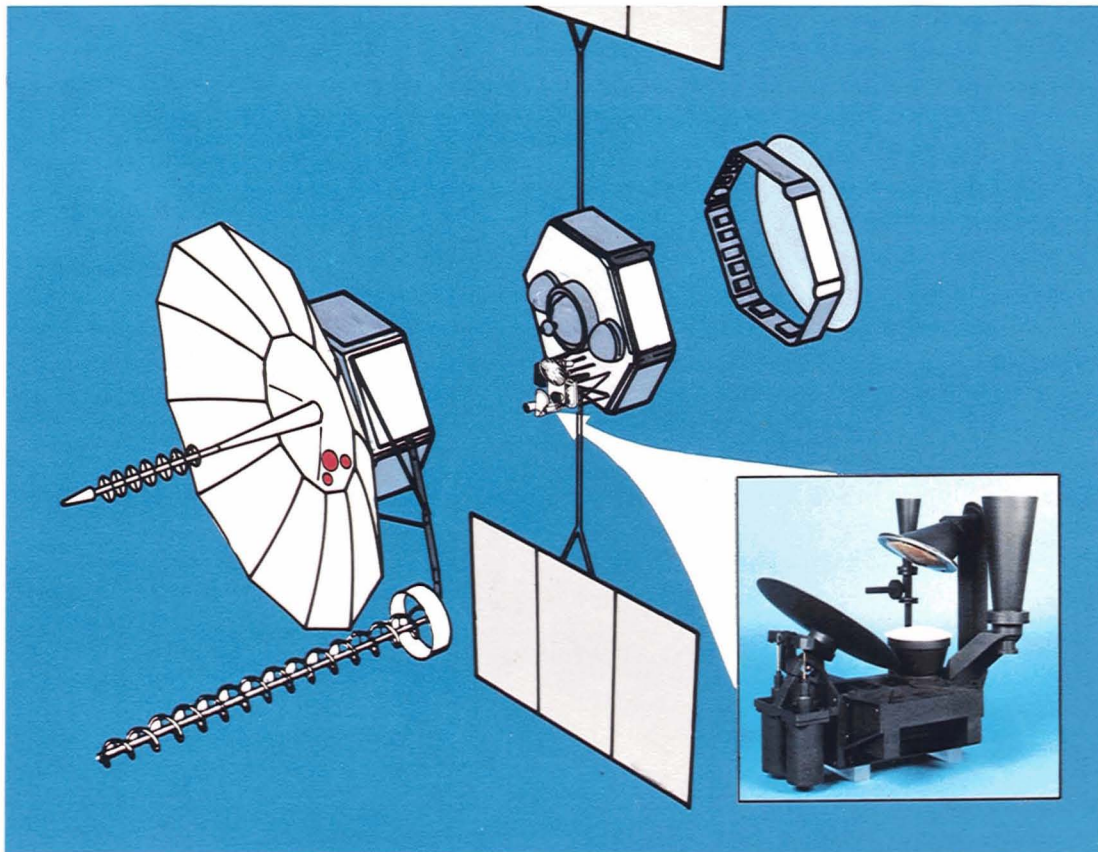

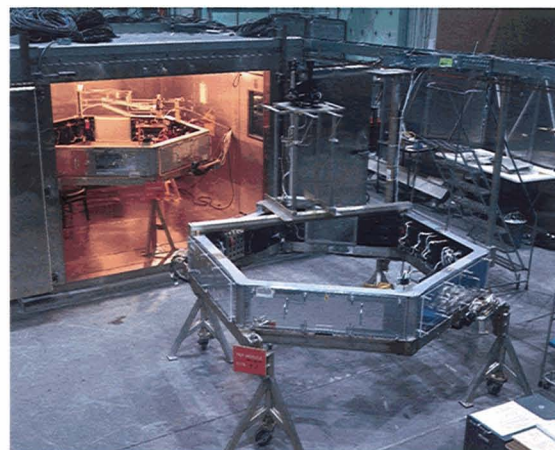*Fig. 6 — The FLTSAT spacecraft. FEP antenna system is shown in the inset.*

Fig. 7 — Implementation of FEP. The device includes seven boards, weighs 6 lb, and draws 6 W. It includes a real-time emulator, real-time Pascal, and 64K of memory.

other terminal can be in the same beam or in the other beam. Moreover, the caller doesn't need to know which beam the other terminal is in. The FEP's switchboard remembers the beam assignment from a terminal's log-on and forms the necessary connections in the satellite and sends the user terminals their channel assignments.

In this example, Terminal 1 and Terminal 2 are in different beams. If Terminal 1 wants to call Terminal 2, it sends a formatted C2 control message to the switchboard — an Initial Service Request (C2-ISR). This message specifies the data rate and tells the switchboard that Terminal 1 wants to talk to Terminal 2. The switchboard checks its data base to see if Terminal 2 is enrolled in the system. If it isn't enrolled, the switchboard sends a C3 message to Terminal 1, notifying it that Terminal 2 is not logged on. If Terminal 2 is enrolled, the switchboard sends a C3 message to Terminal 2 called a Ring-Up message (C3-RU), analogous to a telephone ring.

Terminal 2 responds to the "ring" with a C2 message called a Call Answer (C2-CANS) in which it tells the switchboard if it is accepting the call and how robust a downlink it needs. The switchboard then checks its data base to see if it has enough uplink frequency channels and downlink timeslots to form the link. If the channels and timeslots are available, the switch-board activates the channels in the uplink processor and the timeslots in the downlink processor. It then sends the channel assignments to both terminals via a C3 Service Assignment message (C3-SA).

Terminal 1 can now begin transmitting on the uplink. Data from Terminal 1 are processed by the uplink processor, passed to the downlink processor, and retransmitted to Terminal 2. Unless a reconfiguration is needed, the switchboard has no further involvement with the link. The data just pass between the uplink processor and the downlink processor. If a person at Terminal 2 wants to become the active talker, he selects his "push-to-talk" button and his terminal automatically sends a C2 message to the switchboard that requests a reconfiguration. The switchboard then reverses the link, sends both terminals a C3 message, and enables Terminal 2 to transmit on the uplink.

This example describes a half-duplex terminal-to-terminal call. Other protocols in FEP support full-duplex terminal-to-terminal calls and multiuser networks (which resemble conference calls).

The priority level of a call is set by the C2 request that initiates the call. If the switchboard does not have enough available uplink or downlink resources to set up a service, it will check its data base to see if any existing services have a lower priority. If a lower priority call

is found, the switchboard will preempt as many of these services as necessary (starting with the lowest priority), send C3 messages to those services that notify them of the preemption, and then configure the higher priority service. The terminal requesting the higher priority service isn't aware of the preemption, and the transaction takes only a few seconds. Thus communications resources don't have to be reserved in this system to ensure that important calls go through.

These protocols are computer-to-computer transactions. The user is no more aware of the details of the switching protocols than a person who makes a conventional long-distance call.

Besides accepting C2 control requests from user terminals, the switchboard can receive the following commands from the FEP Operation Center (FEPOC):

- Upload the parameters that the switchboard uses to make decisions.
- Activate/deactivate special control signals that disadvantaged terminals (especially submarines) use.
- Set up fixed communications links with specific uplink channels and downlink timeslots (like conventional satellites). This command preempts any services using those resources and takes them out of the assignment pool.
- Turn on/off the output of diagnostic information via telemetry.

FEP IMPLEMENTATION

Each FEP package weighs 245 lb and draws 305 W. The electronic boxes that make up the package are mounted on a hexagonal module. This module is attached to the rear of the FLTSAT spacecraft (Fig. 6). The FEP antennas
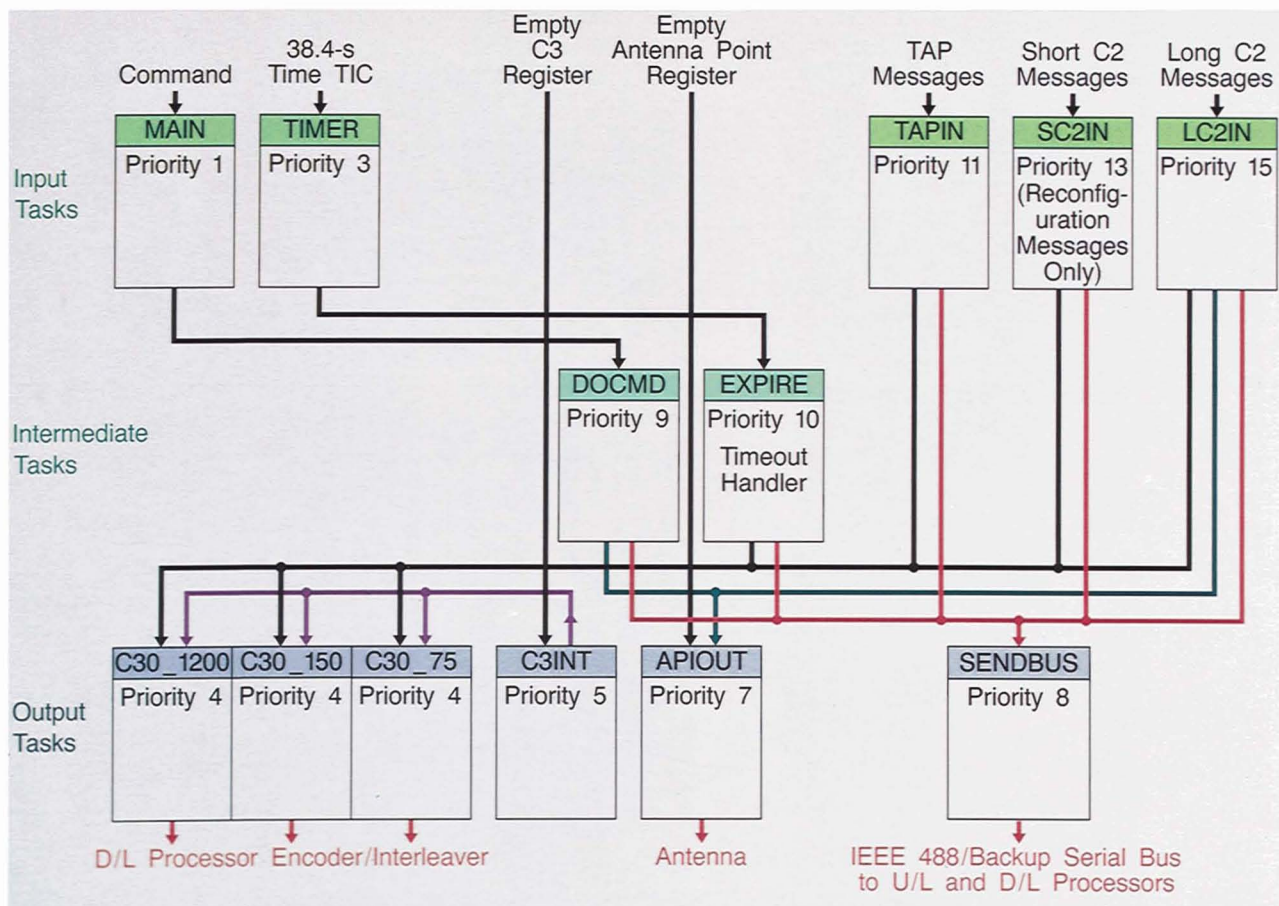

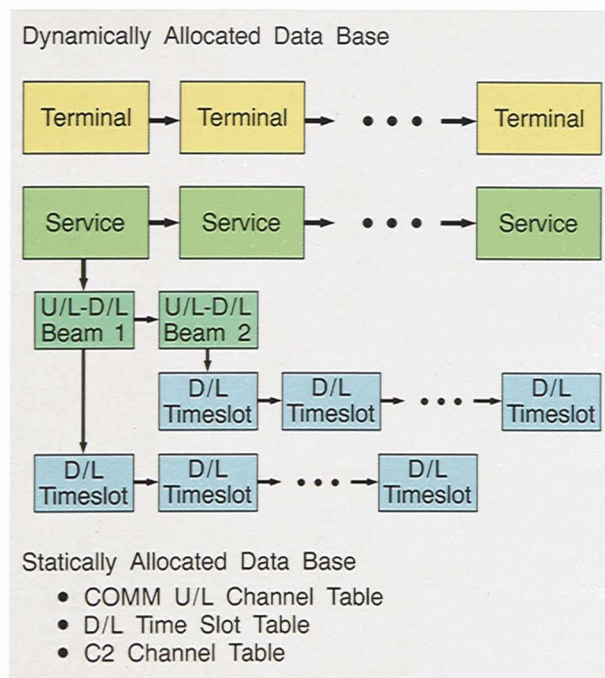*Fig. 8 — Switchboard software tasks.*

Dynamically Allocated Data Base

Statically Allocated Data Base
- COMM U/L Channel Table
- D/L Time Slot Table
- C2 Channel Table

Fig. 9 — Onboard data base.

look through openings in the ultrahigh frequency (UHF) transmit-antenna reflector.

Shown in the photographs in Fig. 7 are two FEP modules. The FEP-8 module (in temperature chamber) is getting ready for temperature stress tests; the FEP-7 module (in front) is undergoing testing with its antenna assembly. Figure 7 also shows a typical FEP digital electronics box. The FEP has a separate box for the uplink processor, the downlink processor, the TRANSEC processor, and the switchboard. The boxes are mounted on the inside of the FEP hexagon. The switchboard box consists of seven circuit boards: a processor board, two input/output boards, and four memory boards.

The FEP resides on a satellite, so power and weight limits are stringent. The heavier a satellite, the harder it is to launch. Moreover, power is limited. (You can't plug a satellite into a wall socket.) FEP receives its power from FLTSAT's solar arrays and uses battery backup during eclipse periods. Weight and power requirements were satisfied — the switchboard weighs only 6 lb and draws 6 W.

Because FEP is in synchronous orbit, it is exposed to radiation and therefore the switch-

board microprocessor had to be radiation hard. We ran radiation tests on several microprocessors and chose an I^2L version of the Texas Instruments (TI) 9900. This microprocessor is well established and well supported. It comes with a development system that includes a TI minicomputer (used to develop, compile, and unit-test the source code) and a real-time emulator. The emulator was attached to an off-the-shelf NMOS 9900 board. This configuration allowed real-time testing of the switchboard software while the switchboard hardware was being developed.

The switchboard software was written in Pascal, a programming language that imposes structured programming techniques. The source code was compiled under TI's MPP Pascal compiler; it runs under TI's RX operating system. The MPP Pascal compiler and its associated RX operating system include such extensions to the Pascal language as the ability to run a multi-tasking program in real time and to handle interrupts with Pascal processes. (For speed, some Pascal interrupt handlers call assembly language subroutines.)

The MPP system provided a reverse assembler, which showed the assembly language code that the Pascal compiler generated. All of the switchboard's object code was run through the reverse assembler. This output, coupled with the source code for the RX operating system, allowed us to know the contents of every location of the switchboard's program memory. We also wrote programs that analyzed the data base and provided the same location-by-location information for the switchboard's data base. Thus every location in the switchboard's memory is identifiable and can be monitored.

One major concern was that the TI 9900 microprocessor could only address a maximum of 64K. Therefore, both the program and its data base had to fit within the 64K limit. This was a challenge! Using the reverse assembler, we identified areas where the Pascal compiler was inefficient in converting Pascal code to assembly-language code. We then compensated for these inefficiencies in our Pascal coding. In the end, the program successfully fit into the available memory.

SOFTWARE OVERVIEW

Figure 8 gives an overview of the switchboard software. This software is a real-time multi-tasking system. Each box in the figure represents a separate task or process, which all operate in parallel. The boxes on the top are the input handlers; the ones on the bottom are the output handlers. When a terminal logs onto FEP, the timing acquisition processor (TAP) notifies the switchboard by passing a message. The TAPIN process handles the message. TAPIN looks in the data base for an available C2 channel. It activates that channel in the uplink processor by sending data over the IEEE-488 bus and sending the terminal a C3 message that specifies its channel assignment.

Any C2 messages sent to the switchboard are handled similarly — by the LC2IN process. If, for example, the C2 message is a request for a communication link, process LC2IN checks the data base to see if there are enough uplink channels and downlink timeslots to configure the link. If so, LC2IN activates the channels by sending data to the uplink and downlink processors over the IEEE-488 bus. Then the process sends a C3 message to the terminals that assigns them their channels. If the C2 message is a request to move the 5° spot beam and the message is from a privileged terminal, process LC2IN sends the data to the antenna-pointing interface, which moves the 5° dish. LC2IN then sends a C3 message to the terminal, which tells it that the spot beam has moved.

To avoid deadlock, output processes are generally given higher priority than input processes. The only exception to this rule is the MAIN process, which is the command input handler. If a failure occurs on-orbit, we must be able to issue a command to clear the failure and/or dump portions of the switchboard's memory over telemetry for analysis on the ground. No failures are expected, but the diagnostic capability is available and has been thoroughly tested.

DATA BASE

The switchboard's data base takes up 25% of the switchboard's memory. Two thirds of the

data base is devoted to the stack, which contains statically allocated data such as uplink channel tables and downlink timeslot tables.

The remaining data base, which is called the heap, consists of packets of dynamically allocated data. The heap contracts and expands in response to user terminal activity. It consists of terminal packets (which contain information for each logged-on terminal), service packets (which contain service parameters for each active communication link), timeout packets awaiting C2 responses, and I/O packets for the switchboard's input and output queues. Figure 9 shows examples of stack and heap elements.

We have divided the heap into a reserved section and a general unreserved section. During initialization, the switchboard software reserves sections of the general heap. Each of these sections, called subheaps, is devoted to one type of data base element. There is a subheap for each of the switchboard's input and output queues, a subheap for terminal packets, a subheap for service packets, and so forth.
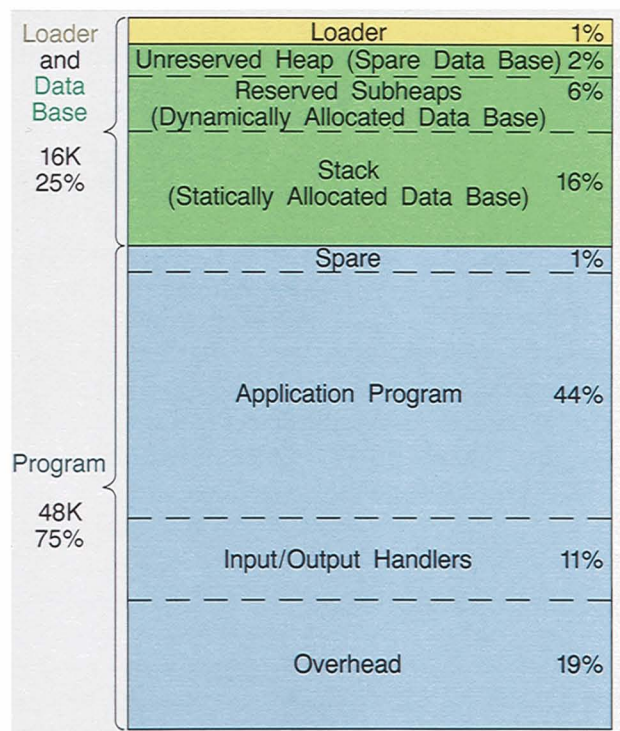

Fig. 10 — Switchboard memory organization.

The reserved heap serves two functions. First, it prevents fragmentation of the data base. The size of heap packets is variable; terminal packets are not the same size as service packets, and so on. Thus as the heap expands and contracts (in response to user activity), the data base memory can become fragmented so that the total memory is large enough for a packet but the contiguous memory is insufficient. By keeping identically sized packets in one subsection, we prevent the fragmentation.

The second advantage of the reserved heap — preventing the system from becoming unbalanced — is even more important than the first. If, for example, too many terminals log onto FEP, sufficient memory for holding service parameters may not be available. If this situation occurs, communication links cannot be set up. So we pre-reserve enough memory for 50 terminals, 26 service packets, and various sizes of input, output, and time-out queues. Then, if a 51st terminal logs onto the system, the terminal can overflow into the general unreserved heap. We keep the system flexible (and don't lock in memory resources any more than the communications resources), yet ensure a reasonable, balanced system. Figure 10 shows the allocation of the switchboard's memory.

SAFEGUARDS

In a satellite-based switchboard system, the switchboard processor must execute its tasks flawlessly. Therefore, we tested the processor exhaustively. But the system must operate in real time and we can't test every possible scenario. So we built an extensive telemetry reporting feature into the switchboard. Using the reporting feature, an operator in FEPOC can monitor the "health" of FEP and its switchboard.

The switchboard is constantly sending indications, via telemetry, of its current state: whether it's at "idle" or is processing an interrupt; its current interrupt level; the state of the IEEE-488-bus or backup-serial-bus handshaking signals; and whether the program is executing a portion of code that could cause a "hangup" if a hardware failure occurred. This last condition would occur if, for example, the bus output handler was active constantly. This activity might indicate a hardware failure of the IEEE-488 bus. The telemetry reporting enables us to detect such a failure on the ground and activate the backup serial bus.

The switchboard can echo any or all of its input and output data via telemetry. If the FEP is in a testing phase, a command can activate or deactivate this feature. When the echo is activated, FEPOC can monitor all the switchboard's interactions in real time. The interactions can also be stored for delogging at a later time. The echo data are useful not only for checking the switchboard functions, but also for diagnosing user terminal problems.

The switchboard can also be commanded to dump portions of its memory via telemetry. These dumps provide information about the state of FEP and can help us diagnose anomalous conditions of the switchboard.

The entire switchboard program is contained in read-only memory (ROM), but it is executed in random-access memory (RAM), which reduces power consumption and enables us to patch the program. The last 512 bytes of the switchboard's memory contain a loader program, resident in ROM. A command can activate the loader program, which then downloads the contents of the switchboard's ROMs to the RAMs. An alternate mode of the loader lets us repair errors in the switchboard software by uploading a patch or an entire new program to the RAMs. This alternate mode can also be used to upload new versions of program, possibly implementing different protocols and thus extending the useful life of FEP.

If the switchboard hardware fails, a command can set the switchboard into a dead state. In this state, special hardware reads switchboard commands and puts the command data directly onto the backup serial bus. Data sent to FEP in this mode can be used to reconfigure the uplink and downlink processors. FEP then operates as a "slowly reconfigurable" conventional communications satellite under direct control of a ground station.

As another backup option, the uplink and the downlink processors can be commanded into a fixed default configuration. Thus, if there is a

fatal failure of the satellite-based switchboard, FEP can still operate as a statically assignable communications satellite. The command channels used to activate these backup modes are protected from switchboard failures — the switchboard cannot reconfigure either of these (one in each beam) dedicated hardwired command channels.

SUMMARY

Meeting the communications needs of many small, geographically dispersed users with widely fluctuating traffic loads has always posed a problem for satellite communications systems. The FEP, with its satellite-based switchboard, provides one solution to this problem. It represents an ideal DAMA system, one that can dynamically allocate satellite communications resources among multiple beams in response to the varying traffic demands of many small, mobile terminals. These terminals have various sizes of antennas and, therefore, require different amounts of uplink and downlink capacity. The FEP can quickly switch communications channels between its beams and can tailor the amount of resources allocated to each user. It, therefore, can make highly efficient use of its communications resources.

The FEP can preempt communications services automatically and prevent important calls from getting busy tones. The switchboard features many safeguards that ensure proper operation.

On 4 December 1986, FLTSAT-7 and its FEP were successfully launched from Cape Canaveral on an Atlas/Centaur booster. After an initial check-out period, FEP began operating successfully as a satellite-based switchboard for the on-orbit testing of user terminals. The user terminals have included small submarine and airborne terminals, as well as larger shipboard and land-based terminals. A small manpack EHF terminal called SCAMP (Single-Channel Advanced MILSTAR Portable) developed for the Army by Lincoln Laboratory has also been successfully tested with the on-orbit FEP. The FEP-8 has been mated with the FLTSAT-8 satellite, which was originally scheduled for launch in May 1987. The launch is now set for 1989.

The advent of radiation-hard microprocessors and high-level structured programming languages has made it possible to use a communications satellite to handle increasingly complex control functions. The success of FEP has demonstrated that an autonomous satellite-based switchboard is a viable concept and, as of 4 December 1986, we have the first "switchboard in the sky."

ACKNOWLEDGMENTS

AUTHOR'S BIOGRAPHY

MARILYN SEMPRUCCI is an Assistant Group Leader in the Communication Technology Group. Her main interest at Lincoln Laboratory is the real-time control of communications satellites. In addition to working on the system described in this article, Marilyn has developed real-time software for the control of the Lincoln Experimental Satellites' K-Band ground terminal and cross-link antennas. Before coming to Lincoln Laboratory, she worked at Raytheon on the real-time control of the FAA's en route air traffic control system. Marilyn received a BA in mathematics and did graduate studies at Northeastern University. She has also taken courses in computer science and logic design. Marilyn enjoys skiing, tennis, and coaching her son's soccer team.