The SAS-3 Delayed Command System

The fully redundant command system of SAS-3 is the first in a new generation. Its relay and data command services offer increased performance, and it has a new capability for delayed command operations.

by E. J. Hoffman

Introduction

While SAS-1 and -2, the first two Small Astronomy Satellites, had nearly identical control sections, SAS-3 required a new level of sophistication from its various support systems. The increased complexity of the power, attitude control, and telemetry systems and the requirement for delayed command capability meant that the SAS-1 and -2 command systems could not be expanded or modified. A completely new command system was designed to serve not only SAS-3 but future SAS-like missions.

The 36 bistate relay commands of SAS-1 and -2[1] were expanded to 56 in SAS-3, with capability for handling up to 64 commands in future versions. Of these, one command is used to power the experiment and 10 uncommitted commands are provided to the experimenter. The data command capability of SAS-1 and -2 consisted of shifting 24-bit words to either of two users. In SAS-3, this "short" data command service was expanded to five users and augmented with a "long load" data command service able to shift long strings (≤ 4080) of bits. The long load service is used to program the variable format telemetry system and the delayed command subsystems.

The most significant change for SAS-3 was the inclusion of a delayed command service, whereby a program of up to 30 relay or short data commands could be loaded for execution at designated times. This allows complex attitude control and other sequences to occur out of radio contact with the SAS ground station at Quito, Ecuador. The delayed command service also initiates special sequences of commands based upon on-board satellite events and executes the commands with a timing precision that is difficult to achieve from the ground.[2]

These added capabilities were to be accomplished with only minimal expansion in the allowed weight, volume, and power. In addition, the highly redundant configuration and fail-safe interfaces of the SAS-1 and -2 command system were to be retained. The operating ground rule was that no single component or subsystem failure should disable the command system. Finally, at each stage of the design and packaging, the possibility of expanding or modifying the system for use on other missions was to be kept in mind. The GEOS-3 command system was the first follow-on system to benefit from this flexible approach.

System Operation

Figure 1 is a block diagram of the command system. The real-time portion consists of two antennas, receivers, bit detectors, real-time logics, and relay matrices. The real-time system immediately executes relay and long or short data commands. Various methods of passive redun-

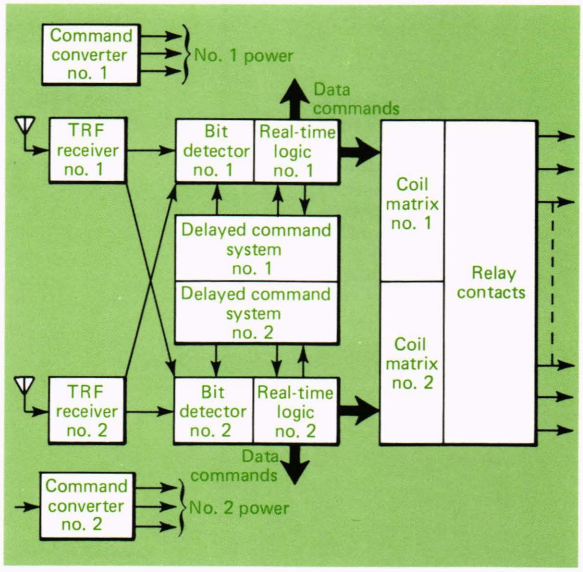

Fig. 1—Block diagram of the SAS-3 command system.

[1] E. J. Hoffman and A. L. Lew, "A Fully Redundant Command System for the SAS-A Satellite," *APL Technical Digest*, 10 Nos. 4 and 5, March-June 1971, 19–23.

The Small Astronomy Satellite Program is sponsored by the National Aeronautics and Space Administration and is managed by the Project Office at the Goddard Space Flight Center.

[2] "SAS-C Satellite Operator's Manual," APL/JHU Report S2-0-156, January 1975.

dancy are used to ensure complete command capability provided at least one antenna/receiver path and bit-detector/logic/matrix path are working. (The real-time command system, unlike other spacecraft systems, obviously cannot depend on relay commands to provide redundant protection.) The two Delayed Command Subsystems (DCS's) can handle up to 15 delayed commands each and operate only through their respective real-time logics. They can be powered individually by real-time relay commands, so that either one can run separately or both can run simultaneously. The use of relay commands to isolate a failed DCS is possible because the DCS's are not in series with real-time operations.

The basic 64-bit real-time command word, transmitted to the spacecraft at 64 bps, is shown at the top of Fig. 2. The SAS-3 command system is compatible with the NASA STDN PCM Instruction Command Format,[3] with each bit encoded as one of two tones in the vicinity of 11 kHz. The frequency shift keying (FSK) tone sequence is then 50% amplitude-modulated with a 64-Hz sine wave whose zero crossings facilitate rapid bit synchronization in the spacecraft. The composite baseband signal then amplitude-modulates a VHF command carrier for transmission to the spacecraft.

Reception of the command begins with the arrival of the radio frequency signal at the command antennas. One antenna is a simple dipole placed at the end of a solar panel. The other is a VHF turnstile placed on the opposite solar panel and shared with the VHF telemetry transmitter through a duplexer. The antennas are oriented so that both cannot be in a pattern null simultaneously. Since each antenna feeds one receiver, the output of each receiver can go through signal nulls as the satellite spins, but the summed outputs from both are relatively constant.

The SAS-3 command receivers are a new TRF (tuned radio frequency) design, replacing the earlier superheterodyne version.[4] In a TRF receiver, all necessary gain is achieved at the RF input frequency (without conversion), demodulated by a conventional diode detector, and buffered to the output through a video amplifier. Selectivity is provided totally by two high-Q crystal filters centered at the VHF input frequency. The SAS-3 receiver is one of the first designs to benefit from recent advances in crystal filter technology that make TRF designs of this bandwidth possible at VHF frequencies. The receiver's major technical advantage is its lack of frequency conversion, which eliminates the local oscillator (LO) and mixer, images, IF, and other spurious responses, and LO radiation. The active elements of the TRF receiver were packaged in four APL-built hybrid microcircuits of three types. This made the final package very light and easy to fabricate and test as compared to the superheterodyne unit.

The demodulated receiver output is a replica of the original amplitude-modulated FSK baseband sent from the ground. Two isolated outputs are cross-coupled from each receiver to the bit detectors, where they are summed. Since either receiver could be in an antenna null or even fail completely, the bit detectors must be able to operate under two extremes: both receivers outputting high signal, or one outputting low signal and the other generating full noise. Command threshold occurs at a received power of -110 dBm, which is typically exceeded by 15 to 35 dB in actual operation.

The bit detectors pass the summed receiver outputs through a filter double-tuned to the two FSK frequencies. The AM synchronizing sine wave is then stripped off in an AM detector and further filtered to reject noise. The detected sync sine wave is tested for proper timing and amplitude and, if within proper limits, a clock pulse is generated at the end of the bit time. This pulse applies power to the rest of the bit detector circuitry and to the real-time logic, most of which is unpowered in the quiescent state. Power is applied for only about 1½ bit times, so that a continuous stream of properly spaced clock pulses is needed for continuous power. Should noisy conditions cause a clock pulse to be dropped, the power is interrupted and the command aborted. With the detection of the first clock pulse (typically the second or third bit, due to ring-up in the filters), a conventional, tuned, noncoherent, envelope detector and comparator is powered to allow data bit decisions to be made. The bit detector also sets a flag indicating to the logic that the present command is of ground, rather than of DCS, origin.

With the arrival of switched power at the real-time logic, a power-up sequence sets all necessary

[3] "PCM/FSK Command Data System Standard," NASA/GSFC Report X-560-63-2, Part II, Section III, revised July 1, 1971.

[4] G. R. Seylar, "A VHF Receiver Designed for Fabrication as Hybrid Microcircuits," Government Microcircuit Applications Conference, Boulder, CO, June 1974, 224–225.

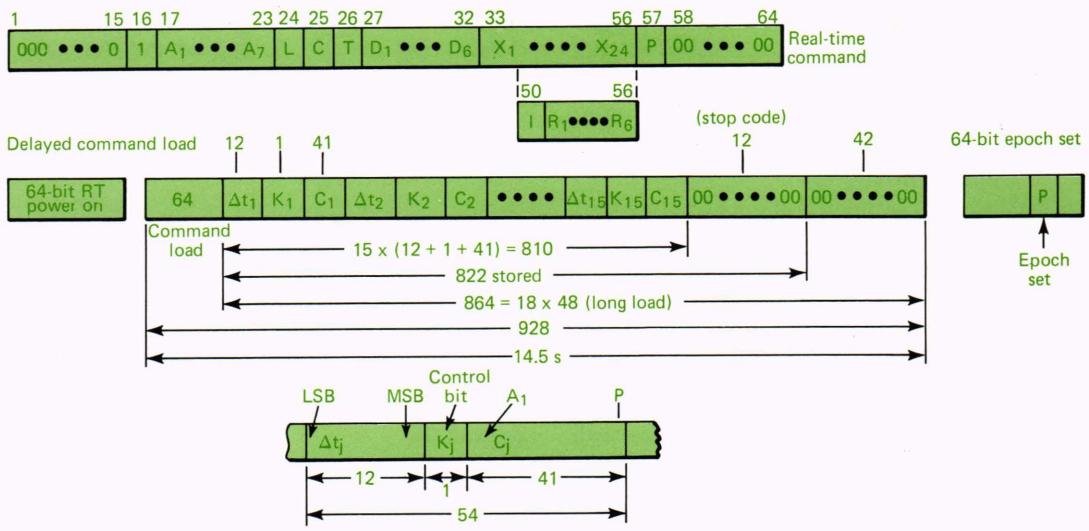

Fig. 2—Command word structures.

counters and registers to zero. As data bits begin arriving, both logics search for the leading sync consisting of 15 zeros and a one (see Fig. 2). To allow for bit detector ring-up, the logics look for 9 to 15 zeros followed by a one.

After leading sync is found, the logics test the next seven bits for satellite address and the following (L) bit for logic select. The address bits, A_1—A_7, are unique to SAS-3 and have good Hamming distance from other satellite codes on nearby frequencies. The L-bit is used by the ground operator to select the logic that will continue to process the command; the other logic is inhibited at that point.

Following the L-bit is a C-bit indicating the type of command, either relay or data; if the latter, a T-bit follows, indicating whether the data command is long or short. If a data command is indicated, the next six bits, D_1—D_6, indicate one of the five possible destinations of the data. In SAS-3, these users are the experimenter; the attitude control, power, and telemetry systems; and the DCS. For additional safety, the destination codes are selected to have a Hamming distance of at least three bits.

In the case of a short data command, the next 24 bits represent the data to be transmitted; they are sent immediately to the selected user along with clock and control signals. In the case of a long load data command, the data follow immediately after the real-time command word, beginning without a break at bit 65, as shown in the middle of Fig. 2.

For short data commands, bit 57 tests parity on all bits that follow leading sync; the test results are signaled to the data user to determine acceptance or rejection of the data. Parity failure inhibits a long load command from shifting any data to the user.

If the C-bit indicates a relay command instead of a data command, the T- and D-bits are ignored, and the last seven bits of the 24 data bits are further decoded. Immediately after a relay command is indicated, the capacitor bank of the relay matrix associated with the selected logic starts charging for eventual use in pulsing the relay coils.

Bit 50, the I-bit, gives the instruction regarding relay direction, i.e., ON or OFF. The next six bits select the row (R_1 R_2 R_3) and column (R_4 R_5 R_6) of the intended command in the 7 × 8 relay coil matrix. Receipt of proper parity permits execution of the command; parity failure aborts the command and lets the coil-switching energy bleed off from the charge bank.

Because a command can arrive from the ground and the DCS simultaneously, a "first-of" circuit is provided to accept the first to arrive and lock out the other. A "terminate" circuit is used to inhibit the logic when there is an indication of missed sync, wrong address, or some other suspicious condition. (In the command system business, it is usually better to do nothing than to do the wrong thing!)

Particularly critical functions, such as address and L-bit decode, are performed in a series-redundant manner even within each of the redundant

logics, to ensure that a failed logic cannot interfere with its surviving mate. Special interface designs are used for the data command outputs to allow the signals from the two logics to be combined in a redundant manner at the user's end. Special attention is also required at the DCS interface so that a DCS failure cannot lock out real-time commands. The logic is implemented by small- and medium-scale, low-power, TTL circuits in flat packs. Point-to-point welded wiring is used.

Eighty relays are arranged in a 7×8 relay matrix with as many as three at each intersection. Welded cordwood techniques are used in packaging the relay matrix, with most of the relay contacts internally wired for the specific satellite configuration. The relay matrix also provides special functions such as "reachbacks." In a reachback, a master relay controlled by the matrix drives a "slave" relay that switches the actual load. Provision is made to switch the slave independently when some on-board event occurs (e.g., sensing of bus undervoltage or completion of tape recorder playback). Another special case is to substitute pulse-forming circuitry for or add it to relay coils.

The relays are the dual-coil magnetic latching type, controlled by the direction of a current pulse in a single coil. One coil from each relay is placed in one of two completely independent matrices, the pulse direction determined by the I-bit. Thus each logic controls one coil of each relay, and redundant control is achieved. In most cases, the relay contacts are in parallel to provide additional failure protection. The relay coil pulse is obtained from the capacitor charge bank, which is started charging as soon as the C-bit decoding indicates a relay command is intended. As soon as the I-, R-, and parity bits are received and tested, the energy of the capacitor is steered through the chosen relay coil cluster by activating the proper pair of row and column SCR switches.

The states of the relays in the matrix are telemetered to the ground using "telltale" bits derived from judiciously selected voltages on the switched contacts. In this way, positive confirmation is received of proper command execution.

Delayed Command Subsystem

A block diagram of the DCS is shown in Fig. 3. The long load sequence required to program the DCS is shown in the center of Fig. 2. Each command to be executed is represented as a 54-bit

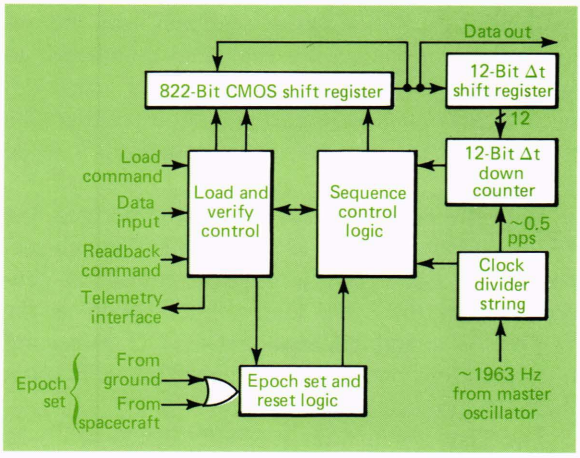

Fig. 3—Delayed command subsystem.

group containing 41 command bits (C_j), a control bit (K_j), and a 12-bit delay time (Δt_j). The 41 command bits consist of the real-time command word minus the leading and trailing sync. The control bit determines whether the command is to be executed or whether it is a "dummy." Dummy commands are used to increase delay times and to fill out the load so that there are always 15 commands loaded in the DCS. The Δt_j is a nonzero 12-bit number indicating the desired delay time between successive commands. Following the 15 command groups is a 12-bit "stop code" and some fill zeros used to satisfy NASA requirements for the length of the bit string.

To begin operation of the DCS, it is turned on by a real-time command through either logic. The long load command is then sent through the proper logic containing the 15 command groups, the stop code, and fill bits. The long load process takes 14.5 s, during which the first 822 bits are clocked into an 822-bit shift register storage in the DCS. (The trailing fill bits are ignored.) The static CMOS circuitry of the storage comprises 13 64-bit shift registers, 10 of which are operated as 63-bit registers by using a tap at the 63rd stage.

The DCS then requests a "verify" readout by the telemetry system. This involves circulating the DCS load three times for readback to the ground, a process that normally takes 3.3 s. The DCS counts the clock pulses and stops the readback after exactly three complete cycles, leaving the load properly aligned in the CMOS registers. The readback confirms that the data arrived correctly and that there are no gross problems in the CMOS storage circuitry. In addition to the automatic

readback after loading, a readback can be requested at any time, by real-time command, to verify the register contents. A "first-of" circuit is used to prevent interference should a readback be requested at the same instant a command is to be executed.

With the program now loaded and verified, the DCS passively awaits the arrival of an "epoch set" command. Epoch set is a pulse command generated from the same matrix intersection as the "DCS ON" command. Sending DCS ON the second time therefore constitutes an epoch set and starts the DCS execution phase.

The first action after receiving epoch set is to shift the first Δt into a 12-bit down counter and count it down using a stable clock of about a 2.1-s period derived from the spacecraft oscillator. When Δt underflow occurs, the K-bit is sensed to determine whether or not to shift the 41 command bits to the real-time logic. If $K = 1$, the DCS activates the power switch in the bit detector to apply power to the logic and shifts the 41 command bits to the logic at 61.3 bps. A control signal to the logic tells it that the incoming command is of DCS origin; it therefore skips the search for leading sync and begins immediately testing for satellite address, L-bit, etc. The inclusion of address and L-bits in the DCS load is redundant since each DCS interfaces to only one logic; however, this protects against data skew and other DCS failures very simply.

If $K = 0$, the command is still shifted out but is kept from activating the logic. In either case, the Δt-, K-, and C-bits are recycled back to the input of the 822-bit register to allow rerunning the program without reloading. That is why 15 commands must always be loaded, even if some are dummies.

After the C-bits are shifted out, the next Δt is shifted into the counter and the process repeated. This continues through all 15 commands until a Δt of 12 zeros (the stop code) is found that terminates the process. The DCS can now re-execute the program after another epoch set, or it can be reloaded or turned off.

The time between adjacent commands can be set in steps of about 2.1 s from a minimum of 2.98 s to a maximum of 2.37 h. The delay time to the 15th command can be as long as 35.61 h. Both DCS's can be operated simultaneously, each controlling its own real-time logic. Either DCS can turn itself or the other one off, and each can request a verify readback from itself. Each DCS can also epoch set the other DCS, but not itself. This feature permits subroutines, infinitely long command sequences, and other interesting possibilities limited only by the imagination and courage of the ground operators. SAS-3 experimenters have become highly skilled in using the two DCS's in sophisticated routines that greatly increase the efficiency of their data collection. In SAS-3, one of the DCS's receives epoch sets from the Solar Attitude Detector in addition to ground-originated epoch sets. This allows special attitude control sequences to begin at a known satellite attitude.

The DCS's and their interfaces are designed to fail safe so that DCS failure does not prevent real-time commanding. For missions that do not need delayed command capability, the DCS can simply be left out; its absence is transparent to the rest of the system.

Except for the CMOS storage registers, the DCS is implemented in low-power TTL. The 64-bit CMOS registers, considered state of the art at the time of the original design, have since been made obsolete by newer units capable of storing 1024 bits in a single chip. Also, since the DCS was designed, single-chip microprocessors have become available; we are anxious to re-examine the logic and DCS to see if they can benefit from a combined microprocessor redesign.

The entire SAS-3 command system, except for receivers and DC/DC converters, is packaged in five standard 6×6-in. SAS "books." The total volume is 360 in^3 including 18 in^3 for the dual receiver. Total weight, including the receivers but not the converters, is 11.6 lb, only 7% more than for the SAS-1 and -2 system. Quiescent power consumption is less than 350 mW for each half of the system with the DCS's off. Each DCS consumes about 420 mW when running.

Testing

A most unpleasant surprise was the amount of effort required to test the SAS-3 flight unit. This was partly the result of the large number of outputs: two types of data command to five different users from each of two logics, plus 437 active inputs and outputs from the relay matrix alone! Testing is complicated by the large number of redundant paths, fail-safe interfaces, addresses, check bits, and other safety features, all of which

require checkout under many different temperature and voltage conditions. The DCS is particularly tedious to check out, even when high-speed clocks are used to shorten the cycle times. In fact, the DCS Test Set was a much more difficult piece of equipment to design than the DCS itself.

The various blocks in the system were checked out in parallel as far as possible. Once combined into the complete system, a dual-track audio tape recorder was used to perform system-level tests. Command tone sequences were recorded on one track and voice instructions regarding the test on a second track. This procedure greatly speeded the system testing with little investment in hardware and none in software development.

Acknowledgment

Tim McAdams (formerly of APL and now with Allen-Bradley Systems) contributed many valuable suggestions to the system design and performed the detailed logic design. The TRF receiver was designed by George Seylar of APL.

The SAS-3 Programmable Telemetry System

by M. R. Peterson

The SAS-3 telemetry system was designed to be an extremely flexible means of collecting and returning to earth data generated by a satellite-borne X-ray experiment. This article describes the concepts used to implement the programmable format telemetry system.

Introduction

In the past, APL telemetry systems for near-earth spacecraft have had a fixed format (i.e., data sampling order) such that the sequence in which data were sampled could be changed only to a limited extent near and after launch. Obviously, it would be advantageous for an experimental satellite made up of several different scientific sensors to be able to change the data sampling format as the requirements of the experimenter vary. Such a programmable system would allow data from malfunctioning sensors to be omitted or would permit more frequent sampling of data of great interest. The programmable telemetry system developed by APL for SAS-3 allows almost limitless changes in format by a ground control station so that the greatest amount of useful information can be obtained from the orbiting satellite.

To ensure that the terminology used herein is understood, the following definitions are given. A "minor frame," shown in Fig. 1, is the basic repetitive sequence of data in a pulse code modulation