# A FULLY REDUNDANT COMMAND SYSTEM for the SAS-A SATELLITE

E. J. Hoffman
A. L. Lew

*With satellites so critically dependent on commands from the ground, stringent demands are placed on command system reliability. This article outlines the special design techniques employed in the SAS-A command system, including the use of a heavily cross-connected redundant topology. Many of the design principles have broad application elsewhere.*

LIKE ANY UNMANNED SATELLITE, once SAS-A was placed in orbit all further control from the earth had to be accomplished through the spacecraft's command system. Typical command operations include turning other subsystems on and off, intricate attitude control maneuvers, setting parameters in the scientific sensors, and duty cycling the power to extend battery life and maintain reasonable temperatures. The satellite is so configured that, should any other subsystem fail, a ground command can usually isolate the failed unit to prevent its interfering with partial completion of the mission. The command system is one of the most critical systems on board, and its failure would almost certainly terminate satellite operations.

Because of this critical need for reliability, it was decided early in the design of SAS-A to employ a fully redundant command system. Redundancy in other subsystems is reasonably straightforward—simply build two black boxes instead of one and depend on the command system to switch between them. The command system itself, however, obviously cannot depend

upon commands to extricate itself from difficulty. Subtler methods of redundant design must be employed, based on techniques that have evolved over several years of design by the APL Space Telecommunications Group. Partially redundant command systems were flown on one or two previous APL satellites,[1] and a fully redundant system was successfully breadboarded several years ago for the Navy Navigation Satellites but was not flown. The experience with redundancy technique obtained from the Navy system contributed greatly to the SAS-A design. Nevertheless, SAS-A represents the first APL launch of a fully redundant command system.

## Command Requirements

Two types of command service are required for SAS-A: relay commands and data commands. Relay commands allow ground control of 36 groups of relays, with up to four double-pole, double-throw relays in each group. These magnetic latching relays can be directly commanded to either an "on" or "off" state, retaining this state indefinitely with no need for holding power until commanded to a new state. These relays provide the switching for power control, gain changes, subsystem turn-on, and any other commands requiring hard switching. The states of the relays are telemetered by means of "telltale" bits in the telemetry frame.

The data command service allows a 24 bit "word" to be transmitted to the satellite and then routed to the desired user. In SAS-A, this service is used by the experimenter and by the attitude control system. The experimenter uses most of the bits in this word to control additional functions in the experiment package, for example, to set gains in a sensor. Attitude control uses data commands to "program" an extended current-vs.-time function in the Z-coil, thus allowing Z-coil maneuvers to continue even after the satellite has dropped below the radio horizon.

Since SAS-A is supported by the NASA STADAN network of ground stations, there was a requirement to be compatible with the NASA PCM Instruction Command Format.[2] This dictated the use of a 64-bit-long command word sent at 64 bits/second using frequency-shift keying amplitude-modulated on to a VHF carrier. Two tones in the vicinity of 11 kHz are used: one tone for a binary "1," the other for a "0." The baseband signal, shown in Fig. 1, is amplitude-modulated 50% by a 64 Hz sine wave to facilitate recovery of bit timing in the satelite. The actual assignment of bit functions within the command word is not dictated by NASA, except for a requirement for leading and trailing sync patterns and the assigned 7 bit satellite address, which is unique to SAS-A.

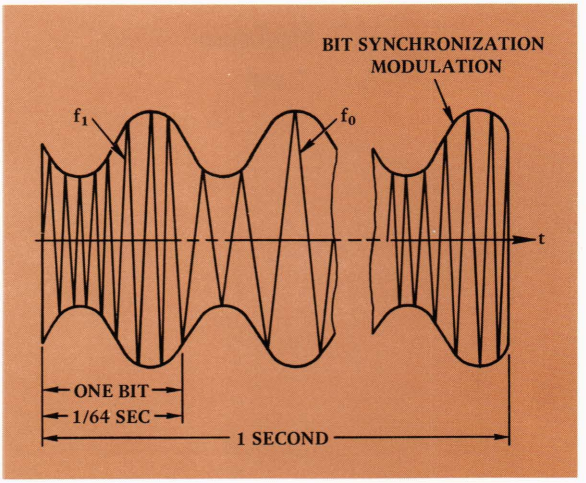


**Fig. 1—SAS-A command signal.**

## System Operation

Figure 2 is a block diagram of the redundant system, and Fig. 3 shows the SAS-A command word and its bit functions. Both decoding logics and roughly half of the bit detection circuitry are unpowered until a command arrives. Upon arrival of a command, the VHF superheterodyne receivers output a replica of Fig. 1 to both bit detectors. Each bit detector operates on the sum of the two receiver outputs to protect against receiver failures and to smooth out nulls in the antenna pattern. As soon as a valid clock signal is recognized, the bit decision circuitry in the bit detectors and the standby portions of the logic circuits are powered by transistor switches for a 2-second interval. Since a valid command is one second long, this provides sufficient tolerance on the interval timing and rejects overly long sequences.

The leading sync pattern of 15 "zeros" followed

[1] R. M. Rhue, *Design of Command Logic for a Near Earth Satellite*, APL/JHU TG 822, May 1966.

[2] Anonymous, *PCM Instruction Command System Standard*, Goddard Space Flight Center, Feb. 1965.

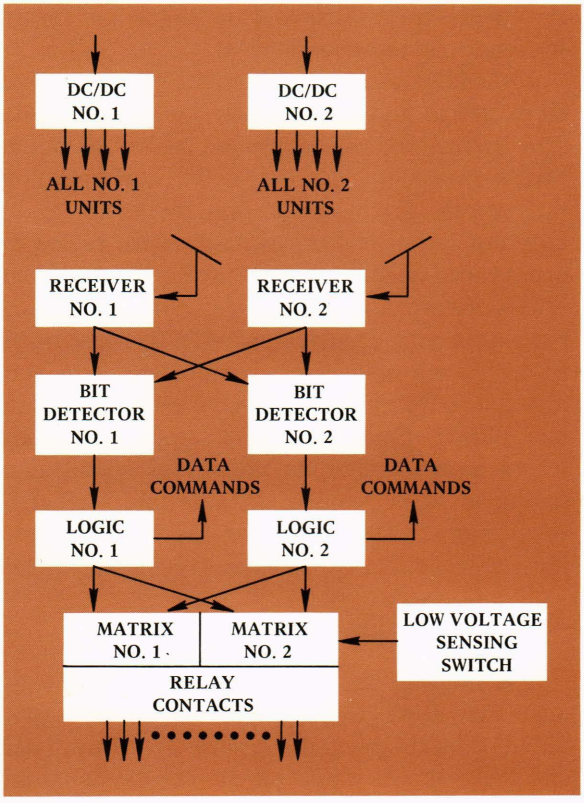


**Fig. 2—SAS-A command system.**

by a "one" is used to help the logic circuit find its place in the command word. The logic looks for nine or more zeros followed by a one, thus allowing the bit detector several bits to "ring up" and acquire bit synchronization. This is immediately followed by satellite address and the logic select (L) bit. Satellite address must be correct, or both logics are immediately inhibited from further decoding of the command word. The L-bit selects one of the two logics to continue decoding, the other logic being inhibited at this point. The L-bit is thus used to steer around a possible bit detector/logic failure. Following the L-bit are the command type (C) bit and matrix select (M) bits. The C-bit determines whether a relay or data command is intended. The M-bit designates, in the case of a relay command, which of the two independent coils in the relay will be used. Since the latching relays can be controlled by the direction of current flow in a single coil, the two coils of each relay are placed in independent 6 × 6 relay matrices. The M-bit thus provides protection against failure in the matrix current-steering circuitry or in the relay coil itself.

Following the matrix select bit, six destination bits indicate the destination of the 24-bit data command (if a data command is intended). On SAS-A only $D_1$ is used since there are only two destinations, but the word structure provides for up to 64 destinations. The next 24 bits ($X_1$ to $X_{24}$) constitute the data word in the case of a data command. These 24 bits are stripped off and shifted to the designated user. In the case of a relay command, only the last seven of these 24 bits are used. In this case one bit (the I-bit) provides the "on" or "off" instruction, $R_1R_2R_3$ selects a row in the relay matrix, and $R_4R_5R_6$ selects a column. Together $R_1$ to $R_6$ indicate the command to be switched.

Final action in the command system is not taken until receipt of bit 57, the parity bit. This bit is a parity check on bits 17 through 56 and guards against all odd numbers of bit errors in transmission. When proper parity is received, the relays are switched or an execute pulse is given to the data command user. Trailing sync is ignored by SAS-A. Approximately one second after command execution, the system is switched back to standby power with only the receivers, part of the bit detectors, and logic interface circuits powered.

Though properly part of the power system, the low voltage sensing switch (LVSS) was designed and packaged with the command system for convenience. Its function is to monitor the main power bus voltage for short circuits or excessive

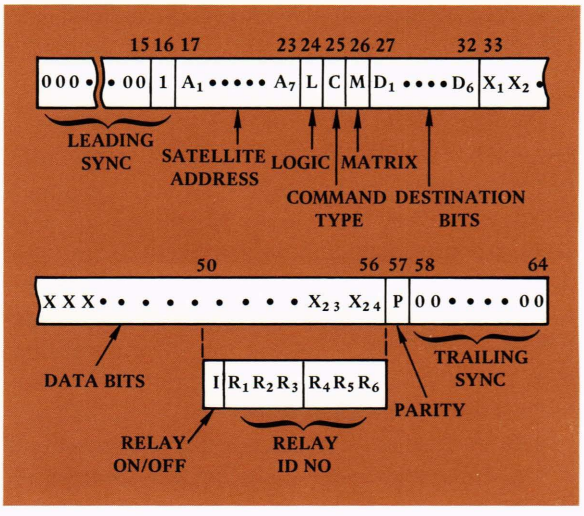


**Fig. 3—SAS-A command word structure.**

current indicated by a drop in voltage. Should this occur, the LVSS directly activates several commands in the matrices, removing all loads (except the command system!) and placing the battery on charge. Though not redundant, it can be removed by a redundant ground command.

From the above description and from Fig. 2 it is clear how redundancy is obtained. Each basic functional block has been included twice, and in addition they have been cross-coupled to each other where practical. This provides even more protection over a simple parallel redundant system by allowing a certain fraction of multiple failures. For example, in the present system, only one receiver, one bit detector/logic, and one matrix need be working to execute all commands. Power redundancy is achieved by using two independent DC/DC converters, one for each half of the system. The converter inputs are diode OR-ed from the battery, the main bus, and from an auxiliary solar array. Redundant interface lines for data commands are sent from each logic and combined at the user's end using special fail-safe combining circuits. Even the ground returns for each half system are run independently and combined only at the satellite ground bus.

Though redundant topology buys a lot in terms of overall reliability, there are many other considerations that are equally important. For example, every effort must be made to reduce parts counts in each block of the system since reliability goes down exponentially with component count. Quite a bit of effort was expended in devising the simplest possible circuits to do each function. The SAS-A bit detector, for example, uses less than one-third the number of parts used in a non-APL design for a similar system.

Another important consideration is the use of extremely conservative circuit design procedures, based on worst case assumptions about component performance and verified by liberal use of computer-aided circuit analysis and breadboard testing. Integrated circuits are used wherever possible to take advantage of their low failure rate per function and the reduction of interconnections they provide. Another useful technique is to remove power from as much of the command system as possible until the beginning of a command is sensed. In addition to reducing quiescent power consumption, this takes advantage of the reduced failure rate that occurs when parts are operated at zero stress. It also improves the chances of surviving long-term radiation in orbit.

Ultimately, of course, a system is only as good as its design and components. To help catch design errors and inadequacies, a formal design review procedure was set up for each block in the system. Each circuit was subjected to an in-group design review, an APL design review by experienced circuit designers, and a NASA design review. Every component used in the system was purchased and screened by APL's Space Reliability Group, and active components were subjected to a 300 hour burn-in to weed out "infant mortalities." Finally, the flight system was subjected to more than six months of qualification testing at every level of manufacture, including such things as temperature testing, vibration testing, and operation in a thermal-vacuum chamber.

## "Proving" Reliability

The original design goal for the SAS-A command system was to prevent any single component failure from disabling the system. To "prove" that this goal was met would require simulating failure in each of the 1858 component parts and testing the result. Even ignoring the fact that each part can have multiple failure modes (for example, a capacitor can short, open, or degrade), this approach is clearly impractical. The actual approach used consisted of heuristically searching out sensitive portions of the circuitry and examining these sections in detail. Interface circuits, as an example, warrant close inspection, especially those interfaces where cross-coupling between redundant units occurs. In this case failure in one unit must not be permitted to "drag down" both of the units it drives.

As an example of such a "malevolent" failure, consider the cross-coupling between receivers and bit detectors. The bit detectors operate on the sum of the receiver video outputs and are designed to handle the 3.3:1 range of outputs corresponding to both receivers degraded to high output or one receiver degraded low and the other having *no* output. In this case, a receiver failing to zero output is a "benevolent" failure. A malevolent failure would occur if the receiver failed in such a way as to not only provide no signal but also produce full noise output. This could result from an antenna failure or from a failure in the receiver RF stage. In this case, the bit detectors

TABLE 1
COMMAND SYSTEM CHARACTERISTICS

| | |
|---|---|
| Configuration | Fully redundant, multiply cross-connected |
| Command capacity | 36 bistate relay commands |
| Data commands | 24 bit words to 2 users, expandable to 64 users |
| Bit rate | 64 bits/second, STADAN compatible |
| Word length | 64 bits |
| Error detection | parity check bit |
| Modulation type | NRZ-FSK-AM-AM on VHF carrier |
| Power consumption | 800 mW quiescent (total, both halves, excluding converters), 1.6 watts peak during a command |
| Weight | 10.8 lb, excluding converters |
| Temperature range | −55° to +80°C |
| RF threshold | −107 dBm |

would have to operate with the sum of full noise and perhaps a low signal from the remaining receiver. The bit detectors have been designed to handle this worst case.

Other, more subtle, malevolent failure modes have been postulated for the logics and matrices. Wherever found, design countermeasures have been taken. After six years of experience with these subtleties, we cannot *guarantee* we've found every malevolent failure, but we can hope!

A more formal approach to reliability analysis is to compute the expected survival probability as a function of time. Of course, this does not guarantee that a system is truly redundant and in fact the *absolute* value of these numbers is debatable. However, they are very useful in comparing the relative reliability of two systems, for pointing out weak spots, and for computing the net gain achieved by redundancy. The Space Reliability Group has performed such an analysis,[3] yielding expected reliabilities of 0.9885 for 6 months operation and 0.9774 for 12 months. Compared with a nonredundant configuration of the same blocks, the redundant topology yields a 6.4:1 reduction in failure probability.

A summary of the SAS-A command system performance characteristics is shown in Table 1.

## Plans for the Future

As soon as a system has been launched, its designers can immediately think of ways to improve it. While the SAS-B command system is almost an exact copy of SAS-A's (except for command functions), the design will be completely new for SAS-C. Plans for SAS-C include expansion of the number of commands, use of Complementary — Metal Oxide Semiconductor (C-MOS) logic, a tuned radio frequency receiver design, smaller relays, and several other improvements. In addition, a delayed command service will allow the user to load up to 15 commands as a group for later execution at timed intervals over a span of 34 hours. Redundancy will be retained, we will continue to squeeze that last part out of a design, and the search for malevolent failure modes will enter its seventh year.

---

[3] D. B. Gilmore and G. San Lwin, *Reliability Prediction Analysis, SAS-A*, APL/JHU Memo SOR 2-70001, Jan. 1970.