THE BELL SYSTEM TECHNICAL JOURNAL
SAFEGUARD SUPPLEMENT
Printed in U.S.A.

SAFEGUARD Data-Processing System:

Maintenance and Diagnostic Subsystem

By J. R. HAHN, JR. and F. E. SLOJKOWSKI

(Manuscript received January 3, 1975)

The SAFEGUARD Maintenance and Diagnostic Subsystem (M&DSS) is a unique, independent, hardware group within the data-processing system through which the nonreal-time functions of fault detection and isolation are performed. In this paper, the M&DSS hardware and fault detection software are described and system performance is reviewed.

I. INTRODUCTION: AN OVERVIEW OF SAFEGUARD MAINTENANCE OPERATIONS

The specific tactical mission for which the SAFEGUARD system has been designed is of extremely short duration compared to the life of the system. Once such a mission has begun, fault isolation and repair are of no concern; at this point, mission success in the face of hardware failures is totally dependent on real-time fault detection and, when necessary, the automatic execution of system recovery. Thus, the fault detection and isolation features of the Maintenance and Diagnostic Subsystem (M&DSS) are oriented primarily toward the goal of maximizing system availability, the probability that, at any random point in time, a complete set of fault-free Data-Processing System (DPS) resources exists.

The M&DSS contributes to maximizing system availability in two ways. First, M&D tests are periodically run on critical DPS equipment to supplement real-time fault detection methods in minimizing the mean-time-to-awareness of hardware faults. These tests are automatically scheduled by real-time software in the green partition and the test requests are sent to the M&DSS over a special interface through the status unit. In this way, every processor in the DPS is switched into the amber partition and tested once every hour; the complete amber partition is tested once each hour; and the green I/O controller with its slaved peripheral controllers is switched amber and tested once every four hours. The M&DSS passes test results back to green system software again via the status unit interface.

Second, and more important, the M&DSS minimizes the mean time to repair of faulty racks by rapidly identifying a minimum set of replaceable or easily repairable modules in which the fault is located. These fault isolation functions may be initiated in response to fault symptoms detected either in real time or during the nonreal-time scheduled tests described above. In either case, fault isolation takes place with the failed rack isolated from the rest of the DPS.

The M&DSS accomplishes this goal through the unique integration of two significant maintenance concepts. First is the use of a special two-way maintenance data path into each DPS digital unit, which bypasses normal data paths. Second is the use of a small general-purpose computer dedicated to system testing, which applies tests over the maintenance paths and interprets test results.

The communication interface between the green partition status unit and the M&DSS provides a rapid and flexible means for bringing maintenance resources to bear on any DPS fault indication. Nonetheless, until a specific faulty rack has been identified, the particular response to be made to any given fault indication often involves judgments based on the total status of DPS resources. Thus, normal SAFEGUARD maintenance operations involve a significant degree of manual interaction. In general, two primary maintenance management functions are performed manually:

(*i*) Monitoring and response to overall system status as reported by green system real-time software and hardwired displays.

(*ii*) Direct control of maintenance testing: The M&DSS will not honor any scheduled test request unless manual "permission" is granted, any test in progress may be manually aborted, and alternate tests may be requested via green system software and the status unit interface.

II. THE SAFEGUARD MAINTENANCE TASK

In its largest configuration, the SAFEGUARD DPS consists of as many as 50 digital racks, each containing up to 100 logic chassis. Each chassis can have between 500 and 600 logic gates. A total installation can have over 2000 chassis with over 500 unique chassis designs. Approximately two million distinguishable faults can occur distributed over these 2000 logic chassis in the typical installation.

The primary goal of the SAFEGUARD M&DSS is to provide rapid fault isolation for the largest, most common class of faults likely to occur.

Other, more subtle faults will involve longer isolation times, but by optimizing isolation for the most common faults, the required overall mean time to repair will still be met. Several assumptions are made concerning this major class of faults which must be handled by the M&DSS:

- (*i*) Only hardware faults are considered.
- (*ii*) Only permanent faults are considered. Transient and intermittent faults, when they occur in the green partition, are handled by real-time error response mechanisms.
- (*iii*) All faults have equal probability of occurring.
- (*iv*) Only one fault will occur at a time: Measured device failure rates support this assumption.

These assumptions, along with further assumptions regarding real-time fault detection capabilities and the distribution of the various classes of faults expected, provided input to a series of parametric studies designed to arrive at specific M&DSS design objectives. The studies led ultimately to the goal of a four-hour mean time to repair for 90 percent of all DPS faults. The mean time to repair includes the time to:

- (*i*) Isolate the fault to a reasonable number of suspect chassis.
- (*ii*) Remove these chassis and test them on an automatic test set that identifies the specific faulty chassis and the failed circuit pack.
- (*iii*) Repair the chassis.
- (*iv*) Replace all chassis and verify the repair.

An analysis of the possible trade-offs of time between these activities led finally to the requirement that the M&DSS be capable of isolating 90 percent of the class of faults defined by the assumptions above, to three or less logic chassis within 15 minutes of their detection.

III. M&DSS HARDWARE

The conventional approach to digital fault diagnosis involves applying a set of input data to the particular circuit under test and, by comparing the output of the circuit to an expected value, deducing the location of the possible circuit faults that could have caused any observed differences. Obviously, the larger and more complex the circuit between input and output, the greater the number of circuit faults that could cause any specific output error, and the greater the ambiguity in the final fault resolution. The primary design feature of the M&DSS (Fig. 1) is aimed at overcoming this problem.

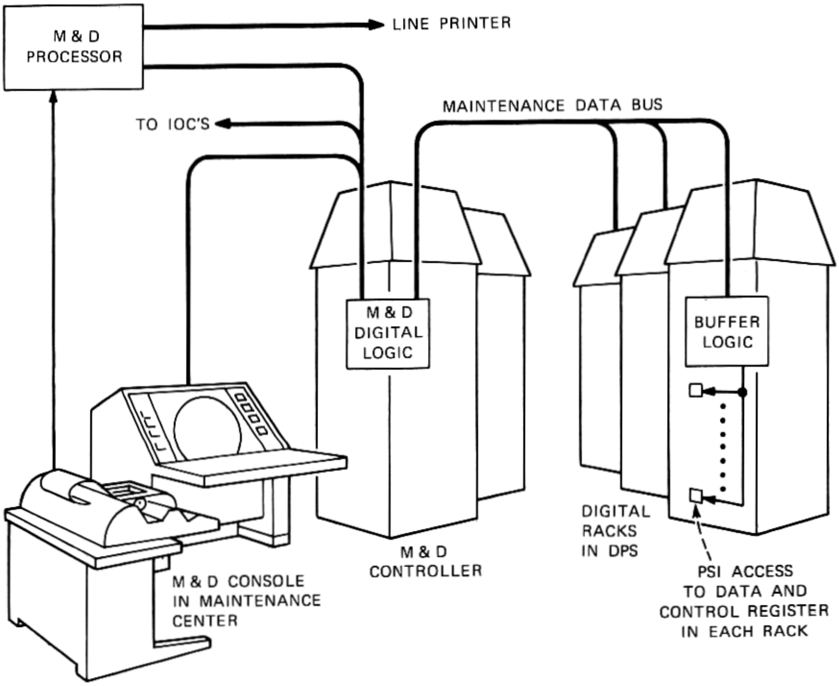

Fig. 1—Maintenance and Diagnostic Subsystem.

Every digital rack within the SAFEGUARD DPS is equipped with a unique internal logic interface to the M&DSS. This interface consists of special programmable Pulsed-Set-and-Indicate circuits (PSIs) connected to most data and control registers within the rack. These circuits provide the means to read from or write into these registers independent of normal data paths. The PSIs are connected via an internal data bus to a maintenance buffer chassis within the rack through which the PSI'd registers may be selectively accessed.

The proper placement of PSIs was an integral part of the logic design process for each SAFEGUARD digital rack. Through PSI access, large blocks of sequential logic are effectively dissected into smaller combinational blocks, each having a number of inputs and outputs accessible via the M&DSS. This not only makes it quite simple to implement system recovery, as will be explained later in this paper, but also results in two important advantages related to fault isolation. First, it makes possible considerably greater fault resolution than can be had in standard logic design. Second, it makes practical the simulation approach to fault dictionary construction.[1]

Testing a digital rack, therefore, involves the repetitive execution of a simple four-step "program":

(*i*) "Set" data onto one PSI-accessible register.
(*ii*) "Set" bits in one or more control registers to enable circuit operation.
(*iii*) "Indicate" (read) the contents of another PSI-accessible register.
(*iv*) "Compare" the result to an expected value.

The execution of such programs is one of the primary functions of a digital rack called the M&D controller. The M&D controller receives maintenance programs from one of several program sources, translates and executes the program in a unit called the sequencer, and communicates with the rack being tested through fan-out logic called a data tree. The data tree is connected to the buffer chassis of each digital rack in the DPS through a separate maintenance channel.

Once the communication channel to a particular rack has been established, the sequencer uses this channel to set data into and read data from selected registers within the rack. Data returned through the "read" instructions can be compared within the sequencer to an expected value and the results of the comparison will be returned to the program source. Again, these three operations, write, read, and compare, are the essence of the sequencer function. The sequencer can also specify up to two additional channels to allow interface maintenance tests between racks.

DPS recovery is implemented through the M&DSS via sequencer "write" instructions stored in a protected core memory (part of the M&DSS itself) and designed to accomplish two functions:

(*i*) Set the appropriate partition bits in the status unit to configure a minimum DPS.
(*ii*) Initialize operational registers in selected DPS racks to boot-load a simple DPS control program and pass control to it; this program then completes the recovery operation.

When recovery is initiated, the M&D sequencer automatically switches to the recovery memory as its program source.

Since the M&DSS is used for both fault diagnosis and system recovery, it must be extremely reliable. The M&D controller, the heart of the M&DSS, can overcome most single faults within itself. It has built-in redundancy, built-in fault detection logic, and PSI access that permits the application of M&D tests to one of the redundant sequencers via another. The chassis involved in system recovery are duplicated, as are the stores containing the system recovery programs.

IV. NONREAL-TIME MAINTENANCE SOFTWARE

The M&D test program itself is the most basic unit of nonreal-time maintenance software. Conceptually, the design of an M&D test is quite straightforward, in keeping with the limited command repertoire of the M&D sequencer described above. Design begins at the level of "micro" tests, each oriented toward a single logic circuit path. Each consists of a number of set-up instructions that set a test vector into a register via PSI access, further instructions which toggle the necessary control bits to cause the test vector to propagate through the logic path to an "output" register, and finally an instruction to compare the output data to an expected value.

From 200 to 2000 such "micro" tests might be designed to cover all the circuits within a logic block. The size of a logic block depends on functional boundaries of logic within a rack. Five to ten such logic block tests typically make up the total test for a single SAFEGUARD digital rack; over 300 block tests are involved in the maintenance facility for the largest SAFEGUARD DPS configuration.

Three independent means exist for applying M&D tests to the digital equipment. The first and most direct means employs a mobile console that is used only during installation of a site. This console, containing a simplified version of the main M&D controller, has its own control panel and associated tape machine. The mobile console connects to the normal M&D buffer chassis in each rack to verify the operation of the rack before the installation of system cabling.

After system cabling is installed, the M&D controller has direct access to each rack, and the second means of applying tests is made available. This consists of the M&D console (shown in Fig. 1) through which tests are transferred to the M&D sequencer from magnetic tape, and test results are displayed on a cathode-ray tube (CRT).

Both the mobile console and the CRT console, however, are extremely slow, depending on magnetic tape as a test program source. Moreover, both return test results to the user in the form of an identification of the compare instructions that failed and the resulting error patterns. Fault isolation then requires a fairly knowledgeable maintenance man to interpret test results. Thus, while the CRT M&D console is a part of the tactical maintenance center, it exists primarily as an emergency backup to the third and most important test facility, the M&D Processor (MDP).

The MDP is a modified CDC Model 1700 general-purpose digital computer. It provides the means for fully automatic high-speed selection and transfer of tests to the M&D sequencer and the automatic interpretation of test results.

The total collection of M&D logic block tests is stored on MDP disc along with all MDP operating software, including a test control program that accepts commands ranging from a request to test a single logic block to a request for a test of an entire digital subsystem.

These test commands may be sent to the MDP automatically from green partition software or manually from its own TTY. In this latter mode, which is normally used for fault isolation, the test program saves the error symptoms (M&D noncompares) encountered and then requests that the fault dictionary tape for the logic block test which detected the fault be mounted on one of the MDP tape transports. Another MDP program then searches the dictionary to find fault lists for the noncompares detected. After the lists are processed, the result is printed out as a list of suspect chassis.

The MDP provides the additional bonus of extending the diagnostic capabilities of the M&DSS beyond PSI-accessible boundaries. The use of fault dictionaries is limited to SAFEGUARD digital logic, but faults in other equipment may be diagnosed by applying functional tests through PSI-accessible registers in a digital unit that interfaces with the unit being tested. An MDP program controlling the test analyzes test results as they occur and branches to other tests along a program path that terminates with the identification of one or more likely faulty circuit cards, or the output of an error code pointing to a written manual procedure to be followed for a final fault resolution. This approach has been successfully applied to the main SAFEGUARD memories and CRT consoles and their supporting equipment.

V. M&DSS APPLICATIONS AND PERFORMANCE

Any evaluation of overall SAFEGUARD M&DSS performance must, of necessity, consider the entire maintenance concept, not only the M&DSS itself, but also the role of the partitionable DPS, its status unit interface with the M&DSS, and the function of system recovery. All play a significant part in achieving the required system availability/reliability product.

At this time, however, the full-scale system tests that will eventually yield specific maintenance system performance data are just beginning. Nonetheless, data do exist in two categories. Extensive testing has been done on the detection and dictionary-isolation capabilities of the basic M&D tests.[1] The M&DSS has also been used extensively in the maintenance of the DPS equipment at the tactical sites during the installation and test period. Maintenance experience in this environment, while not directly translatable to the tactical situation, has produced considerable insight into M&DSS performance.

More than anything else, experience to date has demonstrated the fundamental power and flexibility inherent in the primary M&DSS feature, the extensive maintenance data interface with the entire DPS, in concert with the general-purpose computing capability of the maintenance data processor. Just as encouraging, however, has been the performance of a set of extended M&DSS capabilities developed during the early phases of installation and operation, before the widespread availability of M&D tests and dictionaries. A brief description of these capabilities is instructive as background for the quantitative performance data to be discussed later.

Central to all the extended capabilities of the M&DSS is a set of MDP programs known as Digital Unit Exercisers (DUX). One such program exists for each unique DPS rack type. Each DUX program provides the capability to control the functional operations of a rack on a macroscopic level and to "dump" the contents of individual registers or groups of related registers within the rack. DUX perform these functions by accepting commands in a functional language, translating these commands within the MDP into appropriate M&D sequencer "write" commands, and transferring these to the sequencer for execution. Subsequent "read" commands are used to dump the desired registers, and the results are output on MDP peripheral devices.

In actual hardware maintenance operations, DUX have been used primarily to provide manual interaction, via the M&DSS, with a set of real-time programs originally developed to verify the complete functional capabilities of the DPS.* Data currently being gathered at SAFEGUARD sites show that this mode of fault detection and isolation continues to play an important role.

Table I shows the results of data that have been gathered on the actual use of all MDP resources for a three-month period at the tactical sites. As mentioned earlier, the basic M&DSS and MDP software capabilities were designed to optimize fault detection and isolation on the most common class of faults anticipated, namely, single "hard" device failures. This class is shown in the table under the heading Hard Faults. The Other category includes timing and intermittent failures, design errors, and a variety of miscellaneous failures, largely mechanical in nature. It is important to note that these data were gathered midway during the site test and integration period, a time when design errors are indeed expected to be uncovered, and when frequent handling of the equipment, because of change activity, directly contributes to a greater number of intermittent and mechanical problems.

In view of these facts, the data shown in Table I are extremely encouraging. They show that, for the period covered, the M&DSS success-

* Though not the subject of this paper, it is worth noting that the various DUX capabilities also provide an extremely powerful means for system *software* debugging by allowing dumps and snaps of otherwise inaccessible DPS registers without perturbing the very condition being probed. This capability has found extensive use throughout SAFEGUARD software development.

Table I — MDP performance (July–September 1973)

Total Faults*		Fault Type		Grand Total (75)
		Hard Faults (51)	Other (24)	
M&D tests only	Detect.	96% (49)	83% (20)	92% (69)
	Isol.	92% (47)	54% (13)	80% (60)
DUX/ITPs required	Detect.	8% (2)	0% (0)	3% (2)
	Isol.	17% (4)	0% (0)	11% (4)
All MDP resources	Detect.	96% (49)	92% (22)	95% (71)
	Isol.	100% (51)	71% (17)	91% (68)

* In those cases where isolations exceed detections for a given capability, the fault was usually first detected by a user program. The CDC 1700 was then used to gather enough additional data to achieve isolation.

fully achieved its design goals with respect to the Hard Fault class. Moreover, through use of the MDP extended capabilities, the M&DSS achieved at least its detection goals with respect to all faults.* Finally, the M&D tests alone come very close to achieving design objectives for all faults. Experience, then, supported by the data shown above, leads to a number of specific conclusions regarding M&DSS performance.

Maintenance considerations must be an integral part of logic design. SAFEGUARD development schedules did not allow two or three iterations of the PSI placement-simulation-evaluation cycle. As a result, during test design, cases were discovered where additional PSIs, or a more efficient distribution of existing PSIs, would have produced significant improvements in fault detection, isolation, or both. In particular, more PSI access to control circuits and within logic feedback loops would have made it possible to define smaller and more independent logic blocks. In the most serious cases, hardware change orders were processed to add or rearrange PSIs. Nonetheless, nonoptimum PSI placement remains as the single most significant limitation on detection and isolation.

Increasing the speed of the entire M&DSS would significantly extend its fault-detection capabilities. In its present design, the M&DSS executes a complete read-write-compare cycle in approximately 35 μs, more than two orders of magnitude slower than many internal logic events in the DPS. In the design of the M&DSS, speed was sacrificed for reliability; for example, communication between the M&D controller and each DPS rack is in serial form to minimize the number of con-

* Isolation times using DUX are significantly longer than for M&D tests. Thus, we cannot conclusively say whether or not the goal of 15-minute isolation for 90 percent of all faults has yet been achieved.

nectors, relatively low-reliability components, in the entire path. As a consequence of this design decision, however, the M&DSS is limited in its ability to detect failures that only affect logic timing. A compare instruction can verify whether or not the expected value eventually appeared in a PSI's register, but not whether it arrived there on time. If, however, the M&DSS operated at system speed, it would be more effective in diagnosing this class of faults.

The extended capabilities of the M&DSS described earlier in this section are effective, however, in compensating for both the shortcomings owing to M&DSS speed and those owing to insufficient PSIs. By using M&D access to load and set into execution the more complex real-time functional test programs, the effects of timing faults and faults in complex control circuits can be detected. DUX capabilities can then be used to sample various PSI'd registers along the more elaborate functional path exercised by the test program, and the results can be interpreted to obtain fault isolation to a functional level. In fact, there are very few DPS fault conditions that cannot be handled by one or another of the maintenance tools available through the M&DSS. It is this aspect of experience that leads to a final conclusion on M&DSS performance.

The total M&DSS concept offers great power and versatility as a digital maintenance facility. "Total concept" means the integral combination of PSI access and general-purpose computational control of the PSIs. On-line dictionary search makes possible the rapid isolation of the largest class of common device failures, while the extended capabilities available through the MDP allow the remaining faults to be dealt with in such a manner that the only limitation is the ingenuity of the maintenance man.

In retrospect, the full range of M&DSS capabilities has yet to be fully explored. For example, again because of project schedule constraints, the logic block partitions originally defined have not been changed; but different partitions, chosen perhaps with timing faults specifically in mind, might allow timing faults to be handled via straight M&D test/dictionary methods. Conversely, the real-time DPS capability verification tests that have proven to be so useful in conjunction with the DUX might themselves be restructured with fault isolation more in mind (they were not originally designed for this purpose); it would then be possible to use the MDP to analyze the fault symptoms obtained through PSI access to yield on-line chassis level isolation information.

REFERENCE

1. C. J. Rifenberg, "SAFEGUARD Data-Processing System: The Dictionary Approach to Digital Maintenance," B.S.T.J., this issue, pp. S73–S85.