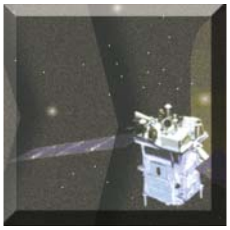

TIMED Autonomy System

Raymond J. Harvey

The TIMED onboard Autonomy System provides fault protection and safing of the spacecraft and performs a limited set of routine operations. The main requirements of fault protection and safing are to detect failed or improperly functioning spacecraft components and autonomously replace them in the operational configuration with properly functioning counterparts. The high degree of autonomy onboard TIMED and in the ground system greatly reduces the required post-launch mission operations staff by ≈66% compared with nonautonomous operations. To lessen the inherent risk of so much automation, a careful development process is followed that includes reviews and ground testing. Using embedded processes designed into processors and components and applying rule-based logic, the concepts for reducing the size of the post-launch operations staff are realized by the TIMED Autonomy System.

INTRODUCTION

The Thermosphere, Ionosphere, Mesosphere Energetics and Dynamics (TIMED) mission is the first in NASA's Sun–Earth Connections Program. (Complete information on the TIMED mission, including participants, status, science, etc., may be found at http://www.timed.jhuapl.edu/mission/.) TIMED is a low-cost mission to provide an understanding of the least explored and least understood region of the Earth's atmosphere—the mesosphere and lower thermosphere/ionosphere, which is located between 60 and 180 km above the surface.[1]

The TIMED onboard Autonomy System has two basic functions: (1) to ensure fault protection and safing of the spacecraft and (2) to perform a limited set of routine operations. The main requirements of fault protection and safing are to detect failed or improperly functioning spacecraft components and to autonomously replace them in the operational configuration with properly functioning counterparts. If a properly functioning spacecraft configuration cannot be achieved, the Autonomy System reconfigures the system into a low-power state. It then re-orients the spacecraft to maximize its power-generating capabilities so that a positive power balance can be maintained until ground support can intervene to diagnose and correct the problem. Examples of autonomous fault protection include monitoring component power consumption to verify that it is within prescribed limits; monitoring processor "heartbeats," which indicate proper instrument or component operation; and demoting the spacecraft mode from "operational" to "safe" when particular faults are detected. Since the goal of the safe mode attitude is to maximize available spacecraft power, it does not meet instrument-pointing requirements, and therefore instrument science data are not collected during a safe mode.

Autonomous routine operations include initiation of a transmitter downlink signal when in view of a ground station, battery charge control, and solar array positioning. By autonomously performing these functions on the spacecraft, the Mission Operations Team (MOT) on the ground is relieved from planning and performing these tasks, thereby enabling spacecraft operations with a smaller than usual team.

The Autonomy System is implemented in two ways: (1) through embedded software- and hardware-based processes designed within subsystems, processors, and components, and (2) by rule-based autonomy. The TIMED mission uses autonomy extensively, both onboard the spacecraft and in the ground system, to reduce the required post-launch mission operations staff by two-thirds[1] compared with staff needed for nonautonomous operations. To lessen the inherent risk of so much automation, a careful development process is used that includes a thorough review and Autonomy System ground testing.

AUTONOMY SYSTEM OVERVIEW

Cost-Savings Approach

In the past, many organizations adopted a technical approach that shifted satellite program costs from development Phase C/D (design, fabrication, and launch/checkout) into Phase E (mission operations and data analysis) to avoid budget cuts that typically affected Phase C/D. This shift of costs from prelaunch to post-launch was often done at the price of an increase in overall mission expenses. The TIMED program, cost-capped from the beginning, took a novel approach to reducing expenditures by focusing on reducing the high costs associated with post-launch operations staffing, the factor seen as having the potential for the greatest savings.[1] Thus the rationale behind the TIMED Autonomy System requirements was derived from the mission's system-level requirement to reduce life-cycle cost. Although this shift may have resulted in an increased expense during the development phase, the up-front time and effort reduced mission operations costs later. Cost savings were realized through the incorporation of efficient allocations of autonomous operations to reduce the amount of support needed by the post-launch operations team (today, a team of eight people operates TIMED). Both the spacecraft and ground system were examined for possible incorporations of autonomy to meet this cost-saving objective.

Design

As noted above, TIMED autonomously performs functions using embedded and/or rule-based autonomy. Much of the embedded autonomy for fault tolerance and system safing involves health monitoring and reporting to the rule-based autonomy for action.

Embdded Autonomy

Recall that the embedded autonomy consists of hardware- and software-based processes designed within subsystems, processors, or components that perform autonomous operations. These functions include closed-loop positioning of the solar arrays, momentum management, guidance and control (G&C) health monitoring, and others performed within the processors and components themselves. The embedded autonomy functions were developed and tested by the particular teams responsible for the subsystem in which they resided. The embedded functionality is used in both routine operations autonomy and autonomous safing logic.

Rule-Based Autonomy

The rule-based autonomy consists of the functionality designed within the command and data handling (C&DH) software, which allows for specification of rules used to monitor housekeeping telemetry and executes commands based on predefined Boolean functions. Autonomy rules compare specified telemetry values to predefined limits and trigger when the limits are exceeded for a certain period. The autonomy rule engine within the C&DH executes the rules themselves and provides arithmetic checks and storage variables to enhance the overall capabilities of the design. There are allocations for 512 autonomy rules per C&DH.

Fault Tolerance and System Safing

Since spacecraft status would be viewed only once per day, a more than traditional amount of fault tolerance and system safing had to be incorporated into the spacecraft design. In addition, because the possible time between spacecraft contacts could be up to 24 h, a significant amount of onboard safing autonomy was needed to react in real time to potential mission-threatening anomalies. The fault tolerance and system safing requirements were derived from a fault tree analysis that used a "top-down" approach, with the loss-of-mission fault as the top-level event. Under it were the events that could lead to the loss.

The high-level fault tolerance and system safing requirements became the events that could lead to a loss of mission. The TIMED Autonomy System combined with onboard spacecraft redundancy to meet these overall requirements.

Described below are some of the safing functions performed autonomously by the spacecraft.

G&C health monitoring and response: The G&C subsystem monitors the health of the G&C processors, sensors, and actuators. If any of these indicate a problem, the G&C subsystem flags the occurrence and the safing logic takes action to reset, power-cycle, or switch to the redundant unit.

Attitude mode management: The G&C subsystem monitors its inputs from other processors and its own components. If conditions warrant a change in attitude mode, this subsystem autonomously demotes to the safe mode attitude, in which the spacecraft can remain indefinitely. The G&C subsystem can also autonomously promote out of safe mode to an operational mode. This capability was tested on the ground but is not intended for in-flight use except in a dire situation that causes a demotion on a routine basis. Fortunately to date, there has been no need to look into the use of this functionality.

Sun avoidance: To protect the instruments from the Sun, "keep-out zone" logic is implemented in the G&C processor to avoid the Sun during maneuvers while in an operational mode.

Solar array rotation problem detection: If the G&C processor controlling the solar arrays detects a problem with commanding the arrays or with the arrays getting to the commanded position, it sets an indicator flag in its housekeeping telemetry. Safing logic then initiates a switch to the redundant processor.

Component soft short monitoring: If a component draws an excessive amount of power from the spacecraft bus (compared to a predefined limit), yet has not been removed by a blown fuse, then safing logic removes power from that component and, where possible, switches to the redundant unit. This is done to avoid a soft short from eventually dragging down the remainder of the bus below a critical threshold.

Response to safe mode entry: If the G&C demotes to safe mode, safing logic commands the instruments into a safe configuration in case the Sun is encountered during the process of slewing to the safe mode attitude. This is necessary since the keep-out-zone logic is not implemented in safe mode.

Processor heartbeat failure and response: Each onboard processor has a toggling parameter, or heartbeat, in housekeeping telemetry that indicates that the processor software is still functioning. Safing logic monitors each heartbeat and initiates a recovery response if it appears frozen.

Low-voltage shutdown and recovery: To avoid a critical hardware-based low-voltage shutdown, there are two levels of software-based checks at higher levels. These perform G&C processor switching and load shedding to alleviate the condition before triggering the more drastic hardware-based function.

Battery health monitoring and response: The battery is monitored for conditions indicating an overcharge or overtemperature status; if detected, charge rate control or load shedding is initiated.

Instrument safing: Under certain predefined conditions, the instruments may be power-cycled or powered down and placed in a safe configuration. Another function allows an instrument to set a bit in its spacecraft status message to request to be powered down by the spacecraft bus. Internal health monitoring functions within the instrument set the bit, and the spacecraft performs the ordered shutdown. The TIMED Doppler Interferometer instrument is the only one to have used this function.

INNOVATIVE CONCEPTS TO REDUCE POST-LAUNCH STAFFING

A key concept to enable the reduction in MOT staffing was the sizing of the end-to-end data system to allow the recorded science data onboard to be downlinked during a single daily contact. Other innovative concepts enabled a significant reduction in stored command load generation and the transfer of functions typically performed on the ground (e.g., orbit determination, spacecraft time management) to the spacecraft.

A number of requirements were derived from these concepts such as event-based commanding, automated station contact, decoupled instrument operations, and onboard fault detection and recovery. In addition to these onboard functions, ground-based requirements such as unattended or "lights-out" ground station contact operations were also implemented. The primary method used to implement these innovative concepts was the incorporation of a significant amount of autonomous functionality. No other missions at APL had incorporated and relied on so much automation. To mitigate the risk inherent to automated systems, a development process (detailed later in the article) was used that included a significant amount of review and test.

The ability to downlink the science data on a single ground station contact each day allowed the MOT to staff a single shift rather than the more traditional "24/7" control center operations. Limiting the team to a single shift reduced the staff by 66%.[1] To ensure that an instrument did not produce more data than could be downlinked in one contact, a feature was implemented within the C&DH system to limit the amount of data stored on the solid-state recorder. When the limit was reached, no more data produced by the instrument were recorded until the limit was reset by command.

Decoupled Instrument Operations

The TIMED mission design, based on orbit and attitude requirements, allows for decoupled instrument operations. This means that all instruments function independently from each other and from spacecraft resources such as power and attitude, i.e., they do not depend on "moving" the spacecraft. The spacecraft provides sufficient resources for unconstrained, independent instrument operations with 100% duty cycle.

Reference 2 describes a "bus paradigm" where the spacecraft represents the bus resource and the instruments correspond to the passengers. The bus provides the transportation along with the associated resources, and the individual passengers are responsible for scheduling and carrying out their business. The analogy is that the spacecraft provides the resources to allow the instruments to individually collect their science data.

Mission operations planning typically requires science objectives to be resolved by assigning priorities to the various instruments for a particular science acquisition sequence. For most spacecraft, the daily resolution of competing science and instrument requests is a complex personnel-intensive effort. With decoupled instrument operations, a science objective resolution process is not needed.

Decoupled instrument operations drove the need for the spacecraft power and thermal system's design to allow for the 100% duty cycle on the instruments. In addition, the instruments are gimbaled, enabling the G&C system to maintain nadir pointing as its primary mode of attitude control for the mission. This also eliminates the typically labor-intensive effort for the MOT to plan, schedule, and conduct attitude maneuvers for routine science data collection.

Onboard Orbit Determination and Event-Based Commanding

TIMED's low-Earth orbit permits the incorporation of a Global Positioning System (GPS) receiver, another enabling factor for reducing a significant amount of workload for the MOT. The GPS Navigation System (GNS) provides onboard orbit determination and onboard time management by using the GPS constellation. In addition, the GNS lies at the heart of the event-based commanding employed on the spacecraft. Knowing the spacecraft's ephemeris at any given time, the GNS sends indications of orbit milestones over the MIL-STD-1553 data bus so that subsystems and instruments respond to these "events" rather than more typical time-tagged commands. The instruments respond to these notifications by time stamping and changing configurations. These functions are usually done through onboard stored commands generated in the planning function of the MOT. As such, event-based commanding significantly reduces the amount of stored command loads to be generated by the MOT.

Autonomous Routine Operations

As noted earlier, another concept was to explore opportunities to reduce as much as possible the amount of routine operations performed by the MOT. The autonomous routine operations requirements were developed at the system level as functions that could be performed onboard the spacecraft itself to reduce the required post-launch support. An example of an autonomous routine operation is turning the spacecraft transmitter on and off over ground station contacts by utilizing the onboard GNS to determine when the spacecraft is "in view" of the ground station.

To better understand the operations of the spacecraft and help identify additional autonomy requirements, a state transition diagram was developed. To give a general idea of its format, the concept is depicted in Fig. 1. The diagram allowed the team to envision the spacecraft configuration requirements while in the various modes of operation. The ovals indicate a state of the overall spacecraft. Within the ovals, text describes nominal modes and configurations of the subsystems and instruments. Arrows indicate the direction of possible transitions from one state to another. Text on the transition lines indicates the functionality by which the transition occurs. Color-coding the transition lines indicates whether the transition is autonomous or initiated by ground command.[3]

The state transition diagram not only helped the MOT to gain a general understanding of spacecraft operation, it also helped to define design and testing requirements for the Autonomy System.[3] In addition to the automated transmitter turn-on for station contacts, described below are some other routine operations performed autonomously by the spacecraft.

Command history dump: The C&DH has a buffer that contains the last 150 commands carried out along with the time of their execution. Since this is a rolling buffer, the oldest commands roll out of the buffer once it is full. To capture all commands for future reference, onboard logic detects when the buffer is close to full and causes its contents to be dumped to the onboard solid-state recorder.

Processor reboot counters: Several of the onboard processors lack internal counters to monitor the number of times they reboot outside of ground contact. The Autonomy System implements logic onboard to detect whether the processor has reset, to take action to preserve information as to the cause, and to keep track of the number of resets since the last time it was re-initialized.

Solar array rotation: The G&C subsystem autonomously controls the positioning of the solar arrays. It contains a look-up table with the recommended array position for a given beta angle (angle of the Sun to the orbit normal). When the actual array position is outside the range of the recommended position, the G&C

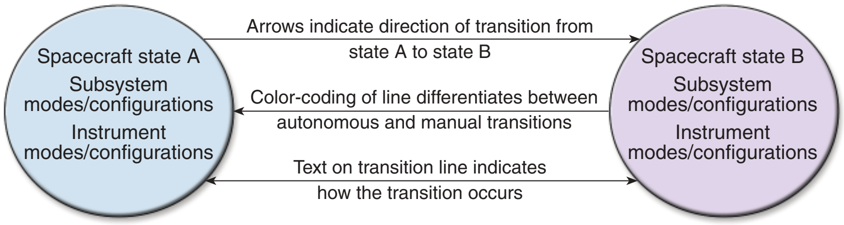

Figure 1. Diagram of the state transition concept.

automatically positions the array to the recommended position.

Momentum management: The G&C subsystem internally monitors overall system momentum. When it becomes greater than a threshold, the G&C autonomously commands the magnetic torque rods in order to remove or "dump" excess momentum in the wheels.

Unattended operations: The ground system is also configured for autonomous operations. To take advantage of the "off-shift" ground station contact time with the spacecraft, while still adhering to single-shift operations, the ground system can receive telemetry from the spacecraft during contacts when the Mission Operations Center is not staffed and can evaluate the spacecraft's state of health. If an anomaly is detected, a "text page" is sent out notifying a member of the MOT of the situation. Based on the information sent, the MOT member determines the urgency. To implement this functionality, the MOT creates scripts based on STOL (spacecraft test and operations language). Scripts are set up to wait until a certain time and then configure the ground system for receiving data. As the housekeeping and science telemetry are received, a rule-based expert technology process developed by NASA called the Generic Spacecraft Analyst Assistant (GenSAA) is used to analyze the housekeeping telemetry for predefined anomalous conditions and performs the required paging.[4]

ONBOARD AUTONOMY IMPLEMENTATION

Recall that the TIMED onboard Autonomy System incorporates two types of autonomy to perform the functionality described above: embedded autonomy and rule-based autonomy.

Autonomy Rules

The autonomy rule structure is set up to work in the form of an IF (conditions), THEN (action) construct. The conditions are based on housekeeping telemetry, and the action is based on commands. Each rule may have up to four individual checks. Up to four rules may be "chained" together, so in essence, a rule "chain" may consist of up to 16 individual checks. Each individual rule or chain of rules consists of only one "action," regardless of how many checks are incorporated. Each action is a single command (usually a command macro) consisting of a sequence of commands stored in a separate area of onboard command memory.[4] The following is a simple example of an autonomy rule with two checks:

> IF main bus voltage is less than 27 V OR
> battery pressure is less than 300 psi, THEN
> begin load-shedding sequence #1.

Arithmetic Checks

Arithmetic checks are used in conjunction with autonomy rules. They perform a mathematical calculation on two telemetry parameters and compare the result to a constant. A bit in housekeeping telemetry is used to indicate if the result of the comparison is true or false. The autonomy rules are set up to monitor these bits and take action based on the outcome of the arithmetic check expression.

Since the telemetry used by the arithmetic checks onboard consists of raw values, some constants must be uplinked as part of the arithmetic check structure to allow the C&DH to evaluate the expression based on engineering units. The arithmetic check parameters take the form of the following expression:

$$Axy + Bx + Cy + D >, <, =, /= E \ ,$$

where x and y are the telemetry points used in the computation; A, B, C, and D are constants used to take into account the engineering unit conversion of the telemetry points and the type of calculation; and E is the constant that is being compared to the result of the computation.[5]

A simple example of the use of an arithmetic check is a check to see if flight computer #1 is drawing an excessive amount of power. This takes the construct of

> IF FC #1 power > 14 W, THEN switch to
> the redundant unit.

In this case, an arithmetic check is used to calculate the flight computer power and determine if it is greater than 14 W. An autonomy rule is set up to monitor the result of that comparison. If true, the rule initiates the switch to the redundant flight computer. To compute power, the component current is multiplied by the spacecraft bus voltage. Therefore, in this arithmetic check, the flight computer current is multiplied by the main bus voltage and is checked to see if that result is greater than 14 W. In the expression above, x = flight computer current telemetry, y = main bus voltage telemetry, and $E = 14$. The calculations for A, B, C, and D take into account the engineering unit conversions for the telemetry points involved as well as the fact that the two parameters will be multiplied by each other.

Storage Variables

Storage variables allow a "snapshot" of a telemetry point or a constant to be stored in a variable and enable one to store an old value of a telemetry point for future reference.[5] An example of the use of a storage variable in conjunction with an arithmetic check and an autonomy rule is the case in which we want to detect if the C&DH command count has increased by 130. A storage variable is used to snap the original value of the total command count. An arithmetic check is then used to subtract the value contained in the storage variable from the current total command count and compare that result to see if it is >130. If so, the arithmetic check output is a true indication. An autonomy rule checks to see when that arithmetic check is true and performs the action to dump the contents of the C&DH command history buffer and "re-snap" the storage variable to re-initialize the process.

Design/Implementation Trade-offs

Trade-offs were done as to which autonomous functionality fit better as embedded or rule-based capabilities. In many cases, the spacecraft design dictated which capabilities were incorporated and in what manner. Because of certain configurations imposed by design, the rule-based autonomy is not functional in the lowest safe modes; therefore, safe mode attitude functionality had to be embedded in the processor that maintains it. Drivers for rule-based autonomy were for those things envisioned as potentially requiring change easily after launch. The changing of autonomy rules is accomplished through routine command loads to the spacecraft. Switching power to and from most components requires the rule-based system because it has access to the C&DH command path which includes power switching. On the other hand, the G&C main processor, the Attitude Interface Unit (AIU), controls power to certain G&C sensors. Because the AIU controls safe mode (in some states, without the aid of rule-based autonomy), functionality is embedded in the AIU to control power to those components. Since the AIU is a mission-critical component in all phases, its health must be monitored by the rule-based system whenever possible. When a problem is detected, a switch to the redundant processor is initiated.

DEVELOPMENT PROCESS

Because of the inherent risk of autonomous systems and the degree to which they would be employed on the TIMED spacecraft, a very thorough review and testing process was used. Figure 2 shows the various stages of development and the activities carried out.

Reviews

Several peer reviews were held during the design and development phase in a format that proved to be very successful. It was felt that a daylong presentation-type review was too formal and concise to guarantee a thorough examination. Instead, Red Team (an independent review team) and Blue Team (the TIMED Engineering Development Team) reviews were held. These reviews were organized as several half-day meetings spaced about a week apart over the course of a month or more. The spacing of the meetings allowed participants to exchange information via e-mail and respond to questions that could not be answered during the actual sessions. It also allowed for follow-up questions from the Red Team prior to the next meeting. Red/Blue Review #1 focused on the Autonomy System's architectural design and requirements; Red/Blue Review #2 was concerned with the Autonomy System design, where requirements were mapped into the various autonomous functionality designs. Also reviewed during this session was the test plan. Each review was conducted by people external to APL as well as APL engineers who were not involved with TIMED. These participants, having worked on similar systems, had very pertinent viewpoints.

Other reviews were held during the middle of developing the rules and responses and following some of the initial testing. These reviews, called the TIMED Autonomy Rule Reviews (TARR), were held in the Red/Blue format as well, but here the Blue Team consisted of only the autonomy engineer. The TIMED Engineering Development Team was used in this case to review the structure of the rules themselves and the command sequence responses. In addition to the members of the TIMED MOT, members of operations teams from other programs were included to take advantage of their expertise.

For the TIMED program, a "code walkthrough" of the rules and associated actions was to be held prior to final testing; however, staffing and schedule ran short. Instead, the walkthrough was held during a "lull" in the program following environmental testing when a launch slip was announced. This review included members of the TIMED MOT who were experts in the operation of the individual spacecraft subsystems as well as the spacecraft systems engineer. The walkthrough not only provided a thorough review of the autonomy rules but also gave the MOT invaluable insight into the design and implementation of the rules.

Testing

Tools

The most difficult part of the entire Autonomy System development process was the testing phase. Inherent in the testing of any system is the issue of knowing how and when it has been tested enough. This especially applies to rule-based systems, mostly because of possible interactions. On TIMED, approximately 200 h of dedicated spacecraft test time was used.[6]

Early autonomy rule testing was performed on a platform prior to onboard testing. The test bed for the C&DH was used in the early stages to work out initial bugs. Significantly, this test bed was also being used by the C&DH software developers at the same time, so test time was at a premium.

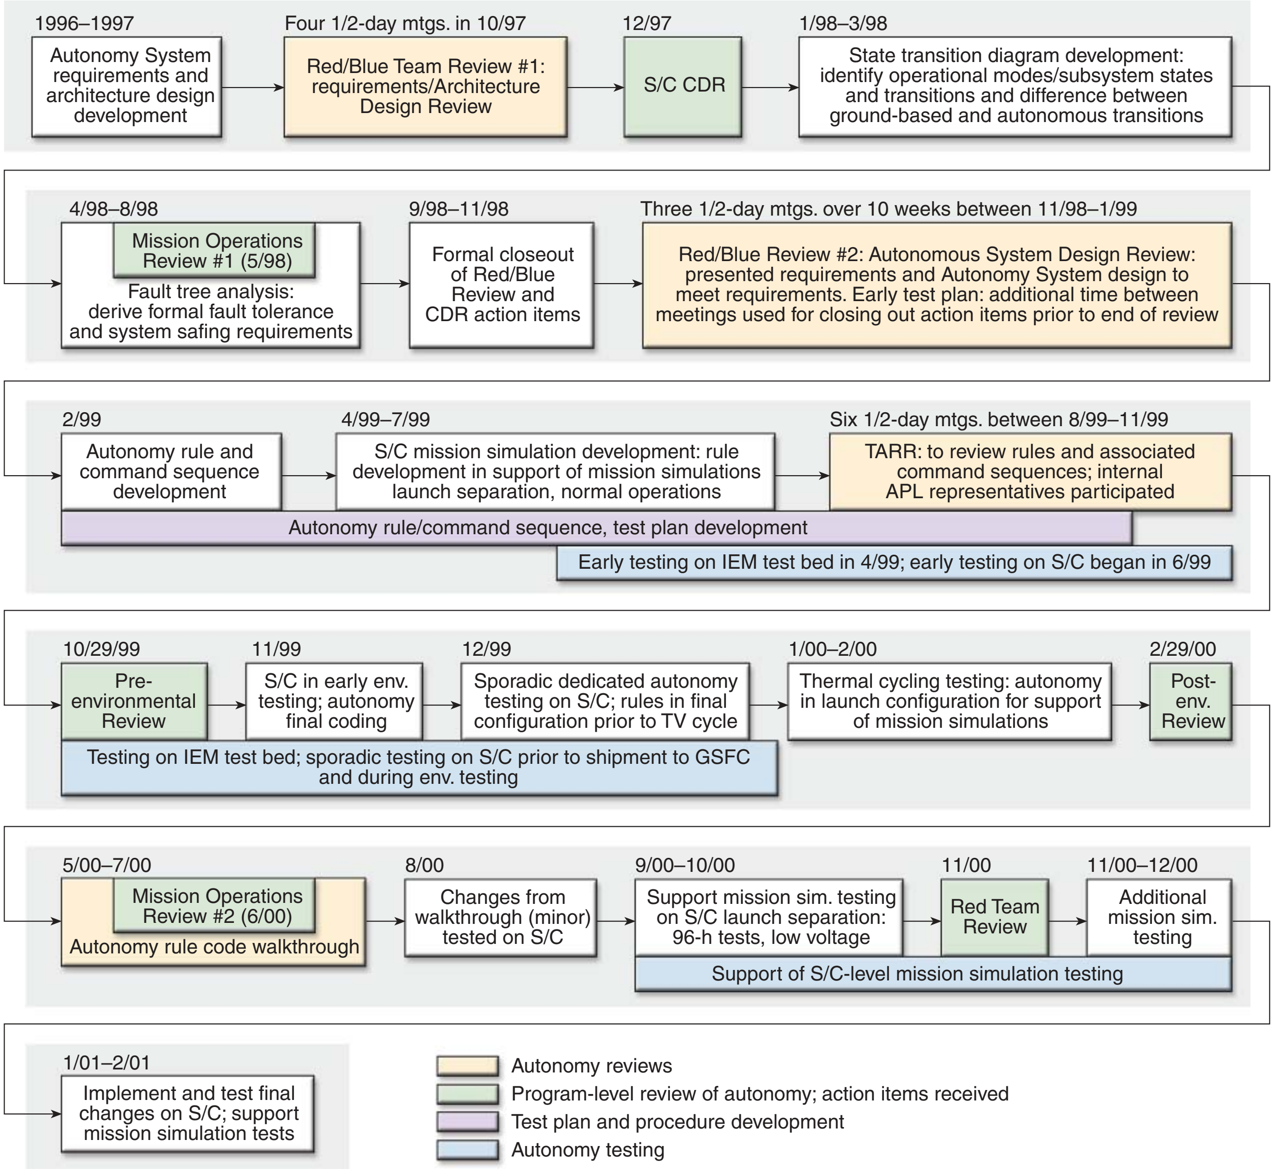

Figure 2. Flow of the TIMED Autonomy System development process (CDR = Critical Design Review, GSFC = Goddard Space Flight Center, IEM = Integrated Electronics Module, S/C = spacecraft, TARR = TIMED Autonomy Rule Review, TV = thermal vacuum).

One advantage of test-bed over spacecraft testing was that the C&DH test bed allowed for the modification of telemetry parameters to ensure that the rules triggered on the appropriate indications. This C&DH test bed was eventually included in the spacecraft hardware-in-the-loop simulator called TOPS (TIMED Operations Simulator), which became available later in the schedule. Although TOPS was an excellent tool for verifying much of the rule base off-line, spacecraft testing was still necessary to ensure proper verification.

Every command macro was tested on the spacecraft. Over 90% of the actual rules were tested on the spacecraft as well. At times, however, the actual rule could not be tested on the spacecraft because the test environment could not be set up to cause the rule to trigger. In these cases, the rule was tested on the TOPS and a "modified" version was used on the spacecraft to ensure that it would issue the desired response when it did trigger.

Types of Testing

As previously mentioned, the embedded autonomy testing was done within the host processor/component. This functionality was included in the Overall Autonomy System Requirements Verification Matrix, which was used to verify that the function was included in the design and that it was successfully tested. The matrix was also used to determine at which stage of development the function was tested. The rule-based autonomy was included in the matrix as well.

The rule-based Autonomy System was tested over many hours, both with and without the spacecraft. It was also tested during spacecraft-level mission simulations dedicated to simulating the spacecraft operation as it would be on-orbit. The mission simulations were designed to test the event-based commanding as well. Simulations included phases of the mission where the particular "events" occurred, such that verification of the appropriate responses could be performed. Mission simulation tests included multiple-day 24/7 "normal day-in-the-life" testing, launch and early orbit operations testing, and low-voltage shutdown simulations.[6]

Also included in mission simulations where the routine operations rules were exercised was an occasional fault rule when the spacecraft was inadvertently configured such that a failure was emulated. Early on, ensuring that the spacecraft was in an on-orbit–like configuration was a difficult aspect of testing. Particular components needed to be simulated and stimulated. If this was not done, the housekeeping telemetry would indicate a failure to the rule-based system.

There were also periods of dedicated autonomy testing when interaction testing was the focus. The most challenging issue in rule-based system testing was determining that no undesirable interactions had occurred between the rules and other embedded autonomous functions. Where these undesireable interactions could be envisioned to occur, tests were designed and implemented to uncover them. These tests included such things as triggering multiple rules concurrently and triggering safing rules with other rules that could possibly have counteracted the safing configuration. Also performed during this period was the "inadvertent" test which caused additional interactions as when the rules were set up as though the spacecraft was on-orbit but part of the spacecraft was not in an on-orbit–like configuration. These situations were considered "good tests" and were used to verify that the system was robust.

Another type of testing involved a level of regression assessment, which was used during spacecraft-level performance evaluation at various stages of the integration and test phase. Here, rule-based functionality was tested rather than the actual rules themselves.

The testing phase was the most difficult part of the Autonomy System development, but also proved to be the most gratifying. To see that the system worked as designed was very satisfying. The rigorously tested Autonomy System has to date performed as expected in orbit, and we are confident that it will continue to do so.

CONCLUSION

Using embedded processes designed into processors and components and applying rule-based logic, the concepts for reducing the size of the post-launch operations staff were brought to fruition by the TIMED Autonomy System.

REFERENCES

[1] Hersman, C. B., and Cameron, G. E., "Reducing Life-Cycle Costs by Effective Allocation of Autonomous Operations," in *Second Int. Symp. on Reducing the Cost of Spacecraft Ground Systems and Operations*, Oxford, England, pp. 75.1–75.7 (Jul 1997).

[2] Cameron, G. E., and Grunberger, P. J., "System Engineering of Cost-Efficient Operations," *Fourth Int. Symp. on Space Mission Operations and Ground Data Systems (SpaceOps 96)*, Munich, Germany (Sep 1996), available at http://www.esoc.esa.de/external/mso/SpaceOps/track8.htm, paper number SO96.8.03 (accessed 10 Mar 2003).

[3] Harvey, R. J., "Implementation of an Autonomous Spacecraft Designed to Reduce Life-Cycle Cost," in *Proc. Fourth Int. Symp. on Reducing the Cost of Spacecraft Ground Systems and Operations*, CD-ROM, Session 2, pp. 1–13, JHU/APL, Laurel, MD (Apr 2001).

[4] LaVallee, D. B., and Knopf, W. P., "TIMED Lights Out Operations," in *Proc. IEEE Aerospace Conf.*, paper #1202, Big Sky, MT (Mar 2003).

[5] Williams, S. F., *TIMED Command and Data Handling Software Command Specification Document*, TIMED Program Document 7363-9115, JHU/APL, Laurel, MD (Jun 1999).

[6] Harvey, R. J., "The TIMED Onboard Autonomy System," *Int. Onboard Autonomy Workshop*, Noordwijk, The Netherlands (Oct 2001), available at ftp://ftp.estec.esa.nl/pub/aut_work/Papers_2001/ (accessed 10 Mar 2003).

THE AUTHOR

RAYMOND J. HARVEY, a member of the APL Senior Professional Staff, has a B.S. degree in aerospace engineering from the University of Maryland. He has over 17 years of experience in space operations and systems engineering and has been involved in many missions in many roles, from flight controller to mission operations manager and from integration and test engineer to lead autonomy engineer. Mr. Harvey served as the TIMED mission autonomy engineer and as the CONTOUR lead autonomy engineer from development through launch and routine operations. His e-mail address is ray.harvey@jhuapl.edu.