

The NASA Integrated Vehicle Health Management Technology Experiment for X-37

Mark Schwabacher^a, Jeff Samuels^a, and Lee Brownston^b

^aNASA Ames Research Center, MS 269-1, Moffett Field, CA 94035

^bQSS Group, Inc.

ABSTRACT

The NASA Integrated Vehicle Health Management (IVHM) Technology Experiment for X-37 was intended to run IVHM software on-board the X-37 spacecraft. The X-37 is intended to be an unpiloted vehicle that would orbit the Earth for up to 21 days before landing on a runway. The objectives of the experiment were to demonstrate the benefits of in-flight IVHM to the operation of a Reusable Launch Vehicle, to advance the Technology Readiness Level of this IVHM technology within a flight environment, and to demonstrate that the IVHM software could operate on the Vehicle Management Computer. The scope of the experiment was to perform real-time fault detection and isolation for X-37's electrical power system and electro-mechanical actuators. The experiment used Livingstone, a software system that performs diagnosis using a qualitative, model-based reasoning approach that searches system-wide interactions to detect and isolate failures. Two of the challenges we faced were to make this research software more efficient so that it would fit within the limited computational resources that were available to us on the X-37 spacecraft, and to modify it so that it satisfied the X-37's software safety requirements.

Although the experiment is currently unfunded, the development effort had value in that it resulted in major improvements in Livingstone's efficiency and safety. This paper reviews some of the details of the modeling and integration efforts, and some of the lessons that were learned.

Keywords: Diagnosis, model-based, IVHM, spacecraft, spaceplane, RLV, X-37, Livingstone

1. INTRODUCTION

Perhaps the most substantial single obstacle to the progress of space exploration is the cost of launching to, and returning from, space. The primary influence in the high costs of current systems is the operations, maintenance and infrastructure portion of the program's total life cycle costs. Incorporation of Integrated Vehicle Health Management (IVHM) technologies into 2nd and 3rd Generation Reusable Launch Vehicles (RLVs) will result in significant program life cycle savings by improving reliability and lowering operational costs:

Life cycle costs for the next generation RLVs include safety and mission assurance. They will be dominated by the cost of processing, operating and maintaining the vehicle and the time required for turning the vehicle around between missions, but safety is still a critical component. Advanced IVHM technologies are critical to the cost-effective management of the vehicle and for the determination of the maintenance actions required prior to the next mission.

This experiment was one of a number of flight and ground demonstrations planned by NASA to develop and mature critical IVHM technologies and to demonstrate the relevance and importance of IVHM technology to the future of the space transportation program. The eventual adoption of this technology will require a number of compelling experiments that demonstrate the ability of the technology to handle a broad class of failures and to perform robustly within an operational scenario. Furthermore, to gain acceptance the IVHM software must also demonstrate that it "plays well" in the avionics hardware and software environment of an operational vehicle without imposing a lot of additional requirements on the avionics hardware configuration.

1.1. Background

On xxxx date, NASA, the USAF, and The Boeing Company entered into a cooperative agreement to develop a new experimental space plane




Figure 1 - The X-37 spacecraft

called the X-37. The X-37 is intended to be the first of NASA's fleet of reusable launch vehicle experimental demonstrators to test future launch technologies in both the orbital and reentry phases of flight. The X-37 is planned to be an unpiloted vehicle that would be launched from the Space Shuttle's cargo bay, or from an expendable launch vehicle, and then orbit the Earth for up to 21 days before landing autonomously on a runway. The X-37 is shown in Figure 1.

There were originally a total of 32 technology demonstrations and 8 experiments on the X-37, including a NASA IVHM experiment. The IVHM experiment was intended to demonstrate IVHM technology developed at NASA Ames Research Center. The IVHM experiment was led by NASA Ames Research Center working in conjunction with Boeing. A Task Agreement between NASA Ames Research Center and Boeing was signed on March 6, 2000.

The X-37 IVHM experiment incorporated the Livingstone model-based health management system[1]. The Livingstone system was initially demonstrated as part of the Remote Agent Experiment on Deep Space One (DS1)[2]. Livingstone is a model-based inference engine that reasons about system-wide interactions to detect and isolate failures. It uses a high-level qualitative model of the components in which both nominal and off-nominal modes are modeled. The Livingstone system uses this model to track commands as they are sent to components and detects discrepancies between the available observations and the predictions of the model. When an anomaly is detected, Livingstone uses advanced techniques to efficiently search the space of both single and multiple-point possible failures to select the most likely failure that is consistent with the available observations. Once the failure is identified, Livingstone continues to monitor the system using its knowledge of the nominal and off-nominal behavior of the failed components.

Livingstone provides a number of potential advantages to 2nd and 3rd generation RLVs. Livingstone uses a high level, qualitative model identifying components and their interaction to perform system-level health management. By using a model of the actual system, Livingstone is able to reason about complex system interactions within a real-time monitoring and control loop, rather than requiring an engineer to reason through all possible interactions and then program in the appropriate response to a pre-defined set of failures. Also, as changes are made to the hardware design, updating and verifying the model is straightforward and less labor intensive than the task of identifying changes required in explicit procedural code. The model-based level of representation streamlines the software development process and maximizes code reusability across vehicles. Finally, the use of a model facilitates the generation of an explanation or justification of the diagnosis allowing the human operator to decide whether the diagnosis is reasonable before selecting or confirming the appropriate recovery action.

Note that Livingstone is only one of many technologies that are relevant to the broader task addressed by an IVHM system. This experiment was only designed to provide a limited demonstration of select IVHM capabilities. Specifically, this experiment focused on the real-time processing of component health information to provide fault detection, isolation, and recovery while in operation and during vehicle checkout. It was not designed to detect subtle degradations in component performance that requires a maintenance response prior to the next mission (prognosis). The experiment would only identify required maintenance actions when the characteristic of the failure could be defined a-priori and are observable within the available sensor data. The experiment was, however, a first critical step towards a more complete real-time and life-cycle-based IVHM system for enabling safe, cost-effective re-usable launch vehicle fleets.

On xxx, due to lack of funding, effort on this experiment was stopped before the first round of integration with the vehicle. Several development cycles were never implemented. The original experiment scope will be described below, but this document is primarily intended to describe the final state of the integrated flight software for the experiment when it was cancelled, i.e. when it was ready for initial integration efforts with Boeing.

1.2. Objectives

The goal of the X-37 IVHM flight experiment was to advance the technology readiness level of the Livingstone reasoning system within a flight environment and to begin its transition from experiment status to inclusion in future operational vehicles.

The main objectives of the X-37 IVHM experiment were to 1) provide a limited demonstration of the benefits of in-flight IVHM to the operation of an RLV, 2) advance the TRL of Livingstone 2 (L2) within a flight environment, and 3) demonstrate that the IVHM software could operate with the Vehicle Management Software (VMS) on the Vehicle Management Computer (VMC). To meet these objectives the flight software experiment would monitor sensor data from the above subsystems, perform real-time fault detection and isolation, and identify potential recovery actions of the X-37 vehicle throughout selected mission phases, while running on the X-37 VMC. This experiment focused on the real-time

processing of component health information to provide autonomous on-board fault detection, isolation and recovery functions during flight operations and during vehicle processing and ground checkout.

A key aspect of this experiment is that the IVHM software would be running in the same VMC as the flight-critical VMS. This experiment differed from other X-vehicle IVHM programs that have run the IVHM code on a separate computer, and represented an important step towards more efficient on-board computer utilization. By demonstrating the safe, integrated, operation of IVHM and VMS software combined in a single software module, running on a single computer, confidence would be gained for this type of operation on future reusable launch vehicles. Many of our challenges were tied to constraints resulting from our goal to implement IVHM on the existing flight computers.

Some of the specific functional requirements for the IVHM software experiment included:

1. The flight experiment shall be capable of identifying certain non-nominal states from sensor readings.
2. The flight experiment shall be capable of identifying sensor failures when possible from redundant sensors or consistency checks among multiple sensors.
3. When sensor failures are ruled out, the flight experiment shall be capable of declaring subsystem performance anomalies from non-nominal data. The goal shall be to isolate the failed component.
4. Diagnostic and reconfiguration/recovery recommendations shall be provided to the VMS for downlink to the ground.
5. The data downlinked shall be sufficient to determine the status and performance of X-37 subsystems/components and for determination of experiment status and verification of experiment performance.

1.3. Scope

The IVHM software was originally planned to assess the health of the X-37's control surface and nose wheel steering Electro Mechanical Actuators (EMA's) and associated electrical power system (EPS) status during all mission phases, with an emphasis on pre de-orbit checkouts, re-entry, and pre- and post- launch. The EMAs control the critical aerodynamic surfaces during re-entry and landing, and nose wheel steering during landing

There were two orbital flights (one of 2 days, another of 21 days) and several approach and landing tests (a.k.a. drop tests) planned. The X-37 IVHM experiment was to be part of all of these tests with an "IVHM day" in the 21-day orbital mission during which the potential of a real-time in-flight IVHM system was to be demonstrated. The initial scope of the model was intended for the drop tests with an understanding that there would be time to rescope/upscope the domain of the Livingstone model before the orbital flights.

The X-37 had a Vehicle Management Computer (VMC), which ran the Vehicle Management System (VMS) software as a VxWorks task. The VMS had several functions, including managing communications with the ground and managing the IVHM software. The IVHM software also ran on the VMC as a VxWorks task. The IVHM task received inputs (sensor values and commands) from the VMS task, and provided its outputs (diagnoses) to the VMS task to be downlinked to the ground.

To ensure vehicle safety, the experiment would be constrained to monitor only during flight operations (Shadow mode) due to technology readiness levels, and aggressive cost and schedule goals. While mission operators and processing personnel could look at the results of the experiment during a mission, the operating procedures for the X-37 would not depend upon the results of the IVHM experiment. Flight managers would not rely upon the information from the IVHM experiment during any mission operations, but the outputs could be used to assist in operational decisions if so desired.

There were five major task areas that NASA performed to get Livingstone ready to fly on X-37:

- 1 Port the Livingstone inference engine from LISP to C++ and meet Boeing safety requirements.
- 2 Develop the interface between Livingstone and the X-37 vehicle management software
- 3 Develop monitors to be used by the Livingstone model
- 4 Develop a Livingstone model of the X-37 vehicle subsystems
- 5 Integrate the model, the monitors, the interface code, and the inference engine, and perform testing of the resulting system.

1.4. Overview

The basic functionality of the software relies on vehicle sensor data and commands, which are provided to the IVHM software by the VMS. Changes in sensor state are detected by feature extraction monitors and used by higher-level system diagnosis or detailed diagnosis routines. The vehicle status, nominal or otherwise, is regularly provided to the VMS for downlink. The experiment status is likewise provided for downlink.

Figure 2 shows the architecture of the IVHM experiment. The Vehicle Management Software (VMS), written by Boeing, and the IVHM software would run on the same Vehicle Management Computer (VMC). Among its other functions, the VMS commands the power system, acquires subsystem status and sensor information, monitors commands issued by the Flight Management Computer (FMC), and stores this vehicle state information in a shared memory area. The IVHM software would retrieve the vehicle state data and commands from shared memory. Monitors translate the real-valued sensor data into a discrete representation to be used by Livingstone. These discrete values would then be fed into the vehicle model. The vehicle command stream would also be made available to the IVHM software in order to identify divergent behavior. The model would generate diagnoses. The IVHM software would never directly command any vehicle subsystems and under no circumstances would its recommendations be acted upon automatically by the vehicle. Rather, the IVHM outputs would be placed in the IVHM area of shared memory for insertion into the telemetry stream and relay to the ground station. The VMS would then telemeter this information to the ground, where ground software would process and display the information

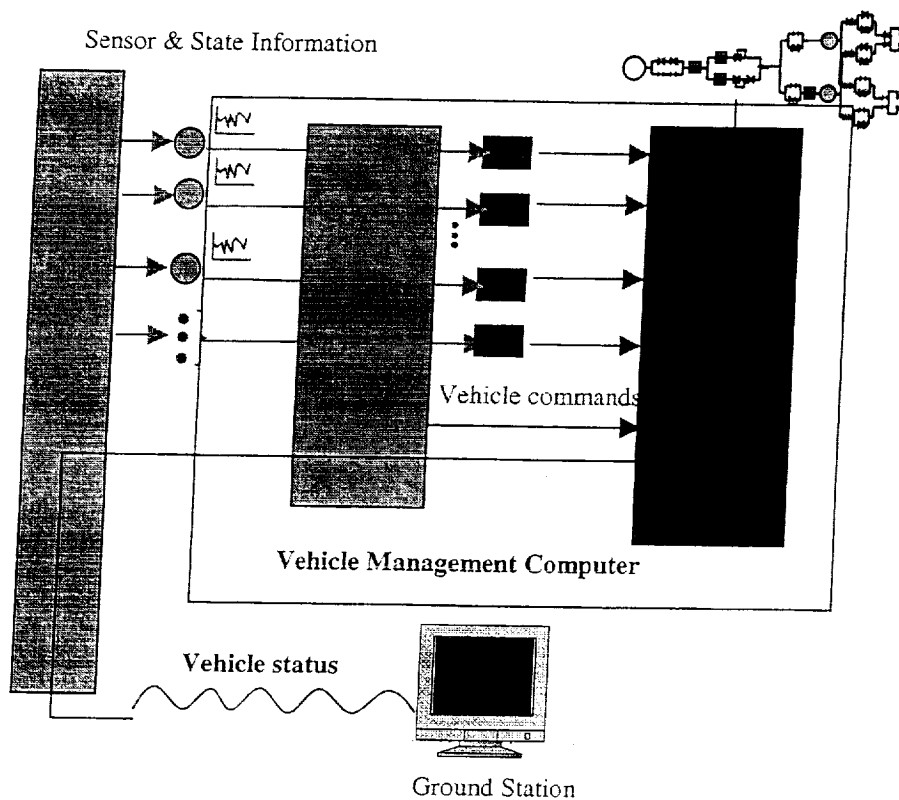


Figure 2 — X-37 Livingstone Overview

1.5. Challenges

Running Livingstone on the VMC presented some special challenges. First, the resources available to the experiment (CPU, memory, and telemetry bandwidth) were very limited. Second, the VxWorks 5.4 operating system does not provide any memory protection, so a bug in our software that caused it to write into the VMS task's memory space could have crashed the VMS and thereby resulted in the loss of the vehicle. Because of this risk, Boeing required our code to meet certain safety standards (such as no dynamic memory allocation and no pointer arithmetic), and to undergo very thorough testing.

2. SOFTWARE COMPONENTS

2.1. Livingstone

Livingstone is a software package that uses a search process to diagnose problems in a complex system that has been described using a declarative model[1]. Livingstone was developed at NASA's Ames Research Center. A previous

version of Livingstone, which was written in LISP, was successfully flown on board the Deep Space 1 spacecraft[2]. The current version of Livingstone (L2) is written in C++. In order to fly Livingstone on X-37, the following modifications to Livingstone were completed by the X-37 IVHM team:

1. Livingstone was ported to VxWorks
2. Livingstone was modified for flight code to satisfy the following X-37 safety standards:
 - No dynamic memory allocation
 - No recursion
3. Livingstone was interfaced with the VMS
4. We implemented routines to allow Livingstone to use a binary representation of the model, to avoid the need to fly the XML parser used in the current implementation

2.2. Monitors

Because Livingstone is only able to take discretized values as inputs, it is necessary to translate all sensor values into discretized values. This is the task of the monitors. There are command and sensor monitors. Command monitors are C functions that take as input the commands provided by the VMS (eg. 0 or 1), and translate these commands into the format needed by Livingstone (such as on or off). Sensor monitors are C functions that convert engineering-unit sensor values into discrete values. There are four types of sensor monitors, corresponding to different types of sensors in the X-37 subsystems covered:

- TranslateOneBit monitors interpret on-off sensor values. The monitors echo the on-off value.
- TranslateSSPCStatus monitors interpret sensors with three status bits. The monitors will return the bit pattern passed from these sensors.
- Bin/threshold monitors interpret floating-point sensors that need to be placed in a "bin" where a range of values corresponds to a bin. For example, a temperature value can be "low", "nominal", or "high" according to whether its value lies in the different ranges corresponding to these bins.
- Compare monitors interpret floating-point position sensors for the aerodynamic control surfaces. The desired values that the monitor should return are MOVING and NOT_MOVING. Thus, the compare function uses the difference between successive sensor values to determine whether its output should be MOVING or NOT_MOVING

Because the signal from the sensors on the spacecraft contains noise, some filtering is required to prevent Livingstone from interpreting signal noise as a sensor indicating a failure. Both persistence and confidence filtering were implemented in the monitors.

Persistence is related to the transition of a sensor signal from one value to another. It is the determination of when the new value should be reported by the monitors. In other words, it is the determination of when we are certain that the new value is correct and not merely caused by noise. The implementation is quite simple. We keep track of a counter for each possible value. We also keep track of the value that was previously reported by the monitors. Also, a persistence level is set as a parameter. The higher this parameter is, the more counts are required to switch the discrete value reported by the monitor. Each counter can have values ranging from 0 as a minimum to the persistence level as a maximum. Whenever a discrete value, calculated from the engineering value, is encountered that is not the same as the previous value, its counter will increase by 1. All the counters for all the other values will decrease by 1. If a value is not within the allowable range, all the counters for the values will decrease by 1. If one of these bin counters reaches the persistence level, then the monitor will report the new value.

Confidence is the determination of whether the monitors should report "unknown" when a sensor output is transitioning from one value to another, but cannot be determined with certainty as belonging to either value. In this implementation, confidence is determined independently from persistence. There are 3 states of confidence:

- Full confidence - equivalent to reporting the current bin consecutively n times, where n is the persistence level.
- Partial confidence - monitor reports previous value, but if next output from the monitors is not the current output, the monitors will report UNKNOWN.
- No confidence - the monitor reports UNKNOWN

Also, three parameters are set: 1) Max number of consecutive "not current discrete" values before switching to partial confidence, 2) Number of consecutive current values needed to change from a partial confidence state to full confidence,

and 3) Number of consecutive current values needed to change from a no confidence state to a partial confidence state. Thus there are 2 counters for each discrete value, one for consecutive "not current values" for switching to partial confidence, and another for restoration of partial and full confidences.

Monitors are "tunable" in that the transition ranges are set by parameters. Persistence and confidence filters are also "tunable" as described above.

When tuning the noise filtering in the monitors, there is a tradeoff between false positives and false negatives. If there is not enough noise filtering, there is the risk that L2 will diagnose a failure when there has not in fact been a failure in the vehicle. If there is too much noise filtering, there is the risk that a real failure will be missed because it is assumed to be noise.

2.3. Models

Livingstone makes use of a declarative model of the system being diagnosed. ARC completed an initial model of the X-37's electrical power system and electro-mechanical actuators. This model was compiled into a binary file, which was to be provided to Boeing for inclusion in the VMC's Flash RAM.

The initial scope of the model was intended for the drop tests with an understanding that there would be time to rescope/upscope the domain of the Livingstone model before the orbital flights. This paper describes the Livingstone model for the X-37 drop tests.

The Livingstone model for the X-37 consists of a single thread (no redundancy) from the high-voltage batteries to the electro-mechanical actuators (EMAs). Six of the nine actuators are modeled (the nosewheel steering and thrust vector control actuators are excluded). Redundancy and the low-voltage batteries were planned for subsequent versions of the model. Systems modeled include: high-voltage battery, high-voltage power relay, SSPCs (solid state power controllers), EMA controllers, and EMA motors and brakes. Figure 3 shows the top-level schematic of the model. Figures 4 through 7 highlight sub-modules within the top-level module.

The model consists of four main module types: High-voltage battery, Power Control and Distribution Unit (PCDU), actuator controller, and actuator. The first two types comprise the electrical power system while the last two comprise the flight actuation subsystem. In addition there are 18 distinct component types within those modules. The model consists of 9 instances of those modules, and 86 instances of components.




Figure 3 — X-37 Livingstone model.




Figure 4 — Battery module.




Figure 5 — Actuator module.




Figure 6 — LCU150V (PCDU) module.




Figure 7 — EMA Controller module.

3. INTEGRATION

3.1. Design Constraints

In addition to the modifications made to L2 for vehicle safety requirements, there were several other constraints on the experiment software. These included memory, CPU allocation, and telemetry bandwidth. These issues, along with L2 performance, have tremendous implications in terms of 1) L2's ability to perform real time diagnoses (if it were to generate recovery options for execution), 2) L2's ability to provide operators timely guidance on failures and appropriate recovery, and 3) the ability of experimenters to validate L2 diagnoses in real time and post-test.

3.1.1. Memory Constraints

The IVHM software was required to fit within a very limited amount of DRAM and Flash RAM on the X-37 vehicle. The size of the code and the amount of space it used for data had to be reduced substantially in order to fit within the requirements.

3.1.2. CPU constraints

The IVHM software was developed to run as a separate VxWorks task. The IVHM task was allowed a very small, fixed amount of CPU time on a fixed cycle. This constraint introduced the unknown element of how L2 performance (dependent on the model size and the actual sensor/command stream) would lag real time.

3.1.3. Telemetry Constraints

The X-37 was supposed to have telemetry on two bands: one low-speed, and one high-speed. The low-speed telemetry was to use well-known, redundant technology, which was expected to be reliable, but had a very low expected throughput. The high-speed band telemetry was an experiment, with no redundancy, but with a very high expected throughput. Though limited in throughput, the low-speed telemetry had expected reliability and high coverage during the orbital missions and was therefore reserved for high priority or mission critical data. Boeing constrained the IVHM task to very small number of bits per second on the low-speed telemetry. We were also to get an unspecified amount of bandwidth of the high-speed telemetry with intermittent coverage. This presented the IVHM team with a significant challenge. The interface between Livingstone and the VMS had to balance the desire to get information to the ground (operational and experiment status) with the limited RAM available for output queues.

3.2. Communication

All communications (receiving sensor input, receiving the command stream, and telemetering output) is done via the VMS. The IVHM software communicates with the VMS using three queues that are stored in shared memory. At a fixed frequency, the VMS uses the monitors to filter the relevant sensor values and process the relevant commands, and then places the filtered sensor values and processed commands into the input queue. Each time Livingstone completes a diagnosis, the diagnosis, represented as the incremental change in the state vector, is placed in the two output queues. Next, the IVHM software removes filtered sensor values and commands from the input queue until another diagnosis is needed. On each cycle through its main loop, the VMS removes as many items from the output queues as there is bandwidth available for downlink. NASA Ames developed a function to be called by the VMS at a fixed frequency that calls the monitors, adds the monitor outputs to the input queue, allows Livingstone to run for a fixed amount of time, retrieves Livingstone's outputs from the output queues, and returns these outputs to the VMS.

3.3. Integration Architecture

The IVHM software uses three queues: an input queue, a low-speed output queue, and a high-speed output queue. All three are first-in, first-out (FIFO). The VMS places monitor outputs (filtered sensor values and commands) into the input queue at a fixed rate. Thus a single input queue holds all of the monitor outputs that are input to Livingstone, in chronological order. The IVHM task removes these items from the input queue whenever it is ready for them (in between diagnoses). The IVHM software places Livingstone's outputs (diagnoses) into the output queues. The plan was for it to place a minimal description of its diagnoses into the low-speed output queue, and a more complete description of its diagnoses, and how it arrived at them, into the high-speed output queue. We only implemented the minimal description (low-speed output queue). The VMS removes items from the two output queues at a fixed rate and downlinks them on the appropriate bands. The integration architecture consists of the following components, shown in Figure 2:

1. **Input queue writer (IQW).** The VMS calls the IQW function, passing it engineering-unit sensor values and commands. IQW calls the monitors, which transform the engineering-unit sensor values into "monitor outputs," which are filtered sensor values that have been discretized. It also calls monitors to translate the commands into the format needed by Livingstone. IQW then places the monitor outputs into the input queue (only commands and discretized sensor values that have changed).

2. **Input queue reader (IQR).** This function removes monitor outputs from the input queue, passes them to Livingstone, and calls Livingstone as necessary to perform diagnoses (see diagnosis timing policy, below).
3. **Monitors.** The functions that take as input the engineering unit sensor values and commands provided by the VMS, and output discrete values that can be used by Livingstone. (M1, M2, ..., Mn in Figure xxx.)
4. **Livingstone.** The inference engine that performs a diagnosis.
5. **Output queue writer (OQW).** The function called by Livingstone to put a diagnosis and related information into the two output queues.
6. **Output queue reader (OQR).** The function that the VMS would use to remove diagnoses from the two output queues so that they can be downlinked.

If any of the queues becomes full, the experiment is terminated.

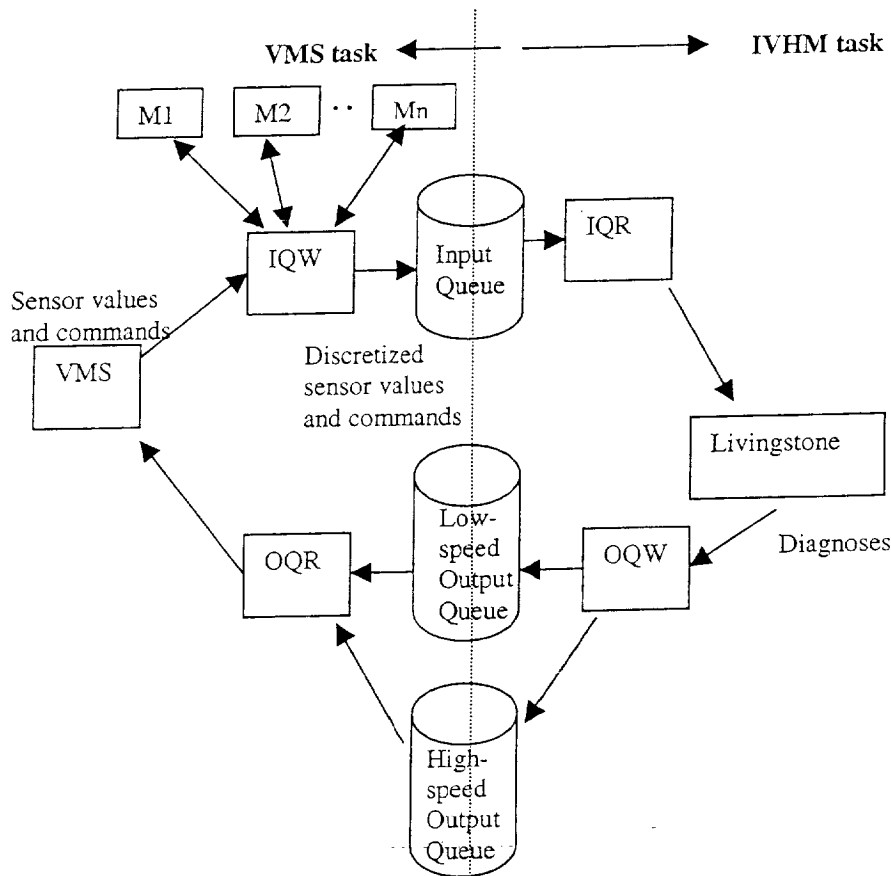


Figure 8: IVHM Experiment Software Architecture

3.4. Diagnosis timing policy

The policy engine is implemented within the IQW, which is within the VMS task. The policy engine decides when a diagnosis should be done, and then inserts a "find candidates" into the input queue. When the IQR removes this "find candidates" from the input queue, it instructs Livingstone to perform a diagnosis. The policy engine incorporates "tunable" parameters into the interface.

When a command is received, there is generally a small time delay before the effects of that command can be seen in the sensor values. If Livingstone performs a diagnosis immediately after receiving a command, then it will produce an incorrect failure diagnosis, since the effect of the command has not yet appeared in the sensor stream. One solution to this problem is to associate with each command a list of monitors that are affected, and the amount of time for each effect to appear. After a command, Livingstone ignores each specified sensor for the specified period of time. We

selected a simpler solution. In this solution, there is a global time delay for all commands. After a command is issued, the system waits for the specified amount of time, and then does a diagnosis, unless another command is received during the delay. So, in effect, after the first command is issued, the system waits until there is a period of the specified length without any commands, and then does a diagnosis. In addition, if there is an unexpected observation (one that does not follow a command), the system does a diagnosis (after a delay, assuming that there is no command during that delay).

In selecting the timeout after an unexpected observation, there is a tradeoff between alerting operators to the problem as soon as possible so that they can act on it, and waiting until enough data arrives to produce the correct diagnosis. We tried to implement a system that will obtain the correct diagnosis. It might be useful to do both: alert the operators of a failure as soon as the first evidence of the failure arrives, and then provide the diagnosis of the failure as soon as enough data has been received to diagnose it.

4. TESTING

Testing on the initial version of the software fell into safety, performance, accuracy, stability, and operational areas. We verified that the safety requirements were met by examining the source code, either manually or using tools that search the source code for prohibited items. We tested memory usage and statement coverage using CodeTEST.

Stability testing included nominal scenario and random failure testing on the PowerPC at Ames. These tests confirm code stability and memory use with and without failures. We successfully simulated a nominal 21-day orbital mission in 21 days of continuous execution on the ARC processor without exceeding memory allowances. Stability testing was also conducted by providing random inputs (Monte Carlo testing) to the integrated software package. Using the Monte Carlo approach, we simulated ten 21-day missions in faster-than-real time. The Monte Carlo tests did not allow us to verify accuracy, since we do not know what the correct diagnosis is for a given set of random inputs. They simply allowed us to verify that the code does not crash or exceed its memory allocation, given a very wide range of different inputs.

Accuracy testing was performed using a set of 28 hand-coded scenarios. No simulation data was available to test the accuracy of the Livingstone model and monitors in diagnosing faults in the electro-mechanical actuators and associated electrical power system. Instead, scenarios were constructed based on expected sequences of commands and observations for nominal and failure operating conditions. The system produced what we believe to be the correct diagnosis in every case. It should be emphasized that Boeing engineers did not confirm the applicability of the scenario files or verify the correctness of the diagnoses.

This version of the IVHM software is not ready for flight as of the date of this report, because some of the software deliverable acceptance tests either did not pass or were not tested. The next step in testing our software, had this effort continued, would have been a functional test conducted with Boeing to see if the IVHM software could have been compiled and executed on their system.

5. CONCLUSION

5.1. Accomplishments

Although the X-37 vehicle program was cancelled, a number of accomplishments were made towards the technical goals. First, Livingstone software was ported from LISP to C++ under VxWorks. During this process Boeing memory requirements were satisfied. Specifically, DRAM usage and Flash RAM usage were both greatly reduced. Most of Boeing's safety requirements were also met. For example, all dynamic memory allocation was removed from the code.

The software team completed the design and development of additional flight software components required for the integrated IVHM software package (an interface to Boeing's VMS and a testing environment). This included completion of version 1 (for atmospheric flight tests) of the X-37 subsystem models that involved single string of EMAs and partial EPS. The team had originally planned an expanded second version of the model (for orbital flight tests) before the project was cancelled. In addition, monitor routines were developed which act as an interface between vehicle data (in engineering units) and L2.

It was originally planned that 3 ARC internal and 3 Boeing versions of the integrated flight code would be developed (for integration with the vehicle, atmospheric flights, and orbital flights, respectively). Integration and testing of the first ARC internal flight code was completed in November 2001. This included integrating the subsystem model, monitors, test harness, I/O queues, and VMS interface code with the modified flight version of L2. Testing successfully demonstrated the stability of the software on a ground-based system similar to the flight hardware/software environment.

The next phase of development would have been working with Boeing to integrate and test this version with their hardware/software on the ground. However, X-37 delays prevented this integration.

No further work has been funded and the X-37 may be cancelled. An accepted approach to integration was successfully developed that would probably have led to safe integration with vehicle flight software. The work remaining if this code were to be flight-tested includes: 1) In-flight IVHM error handling and L2 initialization/restart capability, 2) Successful integration with flight hardware and software, and 3) Integration with ground station software.

5.2. Limitations and Future Work

There are a number of challenges and questions that arose during development of our software. We began to think of solutions for some of these items before this task was terminated. We did not have the opportunity to address these issues or even make serious recommendations. Many of them currently exist only as questions.

Much of the work in creating a model is acquiring the knowledge of the systems to be modeled. What are the interfaces to the system? What are the observable commands and sensor readings? What is the normal operation and how do faults of the system manifest themselves in the observable parameters? A severe limitation of the modeling effort for X-37 was lack of feedback from Boeing engineers. Assumptions had to be made about the system components and operation. Another difficulty was the dynamic nature of the design. As the systems change, so too must the model and monitors.

Systems must be fully characterized before applying Livingstone to diagnose them. Obviously, model-based reasoning tools can only diagnose what has been modeled. Unanticipated system behavior that has not been modeled will most likely lead to a diagnosis that is highly unlikely or misleading. It is important to have simulated data in order to debug the model and monitors. Unfortunately, no such data was provided from Boeing (only a few signals from X-40A and X-38 were provided). The model and monitors would undoubtedly have to be modified after running through simulated data scenarios and after further interaction with Boeing engineers.

Modeling was made more difficult by the high frequency response of X-37 subsystem components and interaction with embedded controllers and fault detection systems.

Choosing the domain of the model is important for a successful demonstration of Livingstone. Livingstone is a discrete, qualitative model-based reasoning engine. Continuous valued systems must be discretized into qualitative bins. Setting the number and threshold of the bins must be accomplished by examining system-wide interactions, not just local behavior. While this can be done for many cases, the problem becomes more difficult when the expected parameter values depend upon the mission phase or upon the duty cycle of the device being modeled. A goal of model-based programming is to reduce development time and costs by building a library of components that may be plugged into other models. In order to accomplish this, the model must not have inherent dependencies on the system being modeled (function in structure). Currently, defining a model without function in structure is difficult due to the interaction of the bin definitions, component failure modes, and system-wide behavior. Also, systems with high frequency feedback loop control present challenges to a discrete model. They must be abstracted to a "black-box" level to determine if they are functioning properly. For systems that have on-board health diagnostics, decisions have to be made about complementing and enhancing the current fault detection with Livingstone or demonstrating that Livingstone can replace the on-board system. The first approach was taken for X-37. In the model, the general rule of thumb was that when independent measurements could confirm or refute the status provided by a device, then the status (device) had a failure mode; otherwise, the status provided by the system was treated as infallible.

Successful integration requires experienced modelers and close cooperation with vehicle hardware and software developers, or extensive training of vehicle developers.

We implemented a simple policy engine, which uses a single global timeout value for commands, and another global timeout value for unexpected observations. We tuned these parameters to get the system to work correctly for all of our scenarios. It is not clear that the values we chose would work correctly with real flight data. Before the system could be flown, it would have to be tested with real flight data. The timeout values might have to be further tuned. It is also possible that we would discover that the simple approach we took to timing would not be sufficient, and that we would have to implement a more sophisticated approach, such as having a different timeout for each command and for each observation.

Our simple policy engine cannot handle interleaved commands – that is, the case where a second command is issued before the results of the first command can be observed. One way to extend it to handle certain interleaved commands

would be to have it ignore (for a time period) only the specific observations that should be affected by a particular command.

In the current system, there is no redundancy in our telemetry. Depending on the reliability of the telemetry (which depends in part on the adequacy of Boeing's error detection and correction), we might have needed to implement a system that had redundancy in the telemetry. For example, we could periodically downlink L2's full state, in addition to downlinking every incremental update to L2's state.

We recommend the development of a more formal approach to developing the monitors and the policy engine. Currently, these parts of the system are developed by writing C code from scratch for each new vehicle that is modeled.

Livingstone must be tuned differently for each application. Tunable components include the policy engine (interface timing), L2 operational parameters (history length, number of candidates, accuracy vs. speed/size, etc), and monitors (data smoothing, persistence, etc). Tuning requires experience, and perhaps art, on the part of integrators. We recommend the development of a more formal approach to "tuning" the various parameters of the system.

One difficulty we encountered is that the X-37 vehicle behaves very differently during different mission phases (such as on-orbit, reentry, and landing), but L2 uses the same model for every phase. It would be useful if L2 could support models that are aware of the mission phase and that model the differences among mission phases.

Our experience with X-37 has demonstrated the need for a system that can perform hybrid discrete/continuous diagnosis. Some parts of the X-37 systems that we modeled are inherently discrete, such as the switches in the EPS. Some parts are inherently continuous, such as the voltage in the batteries and the position of the EMAs. In some cases, such as battery voltage, it makes sense to use monitors to discretize the continuous values. But in other cases, such as EMA position, it is not practical to discretize the values. We found that the only way we could diagnose the functioning of the EMAs in an L2-based system would be to effectively perform the diagnosis within the monitors. We instead chose to rely on the X-37's sensors to tell us whether or not the EMAs were working. We believe that the portion of the X-37 that we selected to model in this experiment could be better modeled in a hybrid diagnosis system.

6. ACKNOWLEDGEMENTS

We would like to thank all of the members of the X-37 IVHM Experiment team, including Mina Cappuccio, Scott Christa, Scott Poll, Jeremy Ou, Jim Kurien, Benoit Hudson, Dan Clancy, and Erv Baumann. This experiment was funded by the X-37 Project Office at NASA Marshall Space Flight Center, and by the Design for Safety Program at NASA Ames Research Center.

7. REFERENCES

1. Brian C. Williams and P. Pandurang Nayak. A Model-based Approach to Reactive Self-Configuring Systems. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*. 1996.
2. Nicola Muscettola, P. Pandurang Nayak, Barney Pell and Brian C. Williams. *Remote Agent: to boldly go where no AI system has gone before*. Artificial Intelligence, 103 (1-2) (1998) pp. 5-47