INFORMATION TECHNOLOGY AND CONTROL NEEDS FOR IN-SITU RESOURCE UTILIZATION

By

Anthony R. Gross[1], K.R.Sridhar[2], William E. Larson[3], Daniel J. Clancy[4], Charles Pecheur[5], and Geoffrey A. Briggs[1*]

[1]NASA-Ames Research Center, Moffett Field, CA, USA
[2]University of Arizona, Tuscon, AZ USA
[3]NASA-Kennedy Space Center, FL, USA
[4]NASA-Ames/Caellum Research, Moffett Field, CA, USA
[5]Research Institute for Advanced Computer Science, Moffett Field, CA, USA

Abstract

With the rapidly increasing performance of information technologies, a new capability is being developed that holds the clear promise of greatly increased exploration possibilities, along with dramatically reduced design, development, and operating costs. In addition, specific technologies such as neural nets will provide a degree of machine intelligence and associated autonomy which has previously been unavailable to the mission and spacecraft designer and the system operator. One of the most promising applications of these new information technologies is to the area of in-situ resource utilization.

Useful resources such as oxygen, carbon dioxide, methane, and water can be extracted and/or generated from planetary atmospheres, to be used for propulsion and life-support needs. This can provide significant savings in the launch mass and costs. This paper will present the concepts that are currently under investigation and evelopment for mining the Martian atmosphere, such as temperature-swing adsorption, zirconia electrolysis etc., to create propellants and life-support materials. This description will be followed by an analysis of the information technology and control needs for the reliable and autonomous operation of such processing plants in a fault tolerant manner. Finally, there will be a brief discussion of the software verification and validation process so crucial to the implementation of mission-critical software.

Introduction

When Europeans first began to explore the New World, they depended heavily upon the intelligence of the explorers and their ability to utilize the resources

---

Presented at the 50th International Astronautical Congress, Amsterdam, The Netherlands, October 4-8, 1999.

* Anthony R. Gross is Associate Director of the Information Sciences and Technology Directorate; K.R.Sridhar is a Professor at the University of Arizona; William E. Larson is Project Manager for the ISPP Project in the Technology Development Directorate; Daniel J. Clancy is a Research Scientist in the Computational Sciences Division; Charles Pecheur is a Research Scientist; Geoffrey A. Briggs is Scientific Director of the Center for Mars Exploration.

available within the newly explored territory. Thus, when Lewis and Clark first embarked on their voyage across North America, they left knowing little about what they would encounter on their voyage. To accomplish the exploration goals that lay ahead of them, they depended heavily upon their own intelligence and ingenuity to respond to the circumstances encountered while utilizing available resources.

These same capabilities have often proved critical in our on-going effort to explore the universe outside of our world. The successful safe return of the crew of Apollo 13 provides a compelling story of our ability to respond to unforeseen circumstances using the limited resources available at that time. Of course, the Apollo 13 incident demonstrated the ingenuity of both the crew and the large ground support team that worked around the clock to explore many different scenarios for the safe return of the crew to earth. Unfortunately, as we start to consider further manned exploration of our solar system, communication delays and budgetary constraints limit our ability to depend upon a large ground support team especially for the routine, day-to-day operation of the mission.

With exponentially increasing capabilities of computer hardware and software, including networks and communication systems, a new balance of work is being developed between humans and machines. This new balance holds the promise of greatly increased space exploration capability, along with dramatically reduced design, development, test, and operating costs. New information technologies, which take advantage of knowledge-based software and high performance computer systems, will enable the development of a new generation of design and development tools, schedulers, and vehicle and system health monitoring capabilities. Such tools will provide a degree of machine intelligence and associated autonomy which has previously been unavailable to the mission and spacecraft designer and to the system operator. These capabilities are critical as we look toward future exploration of our solar system due to both the requirements levied by these missions as well as the budgetary constraints that limit our ability to monitor and control these missions using a standing army of ground-based controllers.

In addition to the development of algorithms for monitoring and controlling the complex devices required to support further exploration of our universe, NASA is also pursuing the development of technology for the in-situ generation of propellant and life support gases from the planetary atmospheres. While there are a number of well-understood chemical processes by which this can be accomplished, a significant engineering challenge still exists to select the appropriate process and to develop a robust device that can operate for up to two years in a harsh environment such as the Mars. Furthermore, this device must operate autonomously over this period of time with limited ground interaction.

In the next section, we present the key concepts currently under consideration for an in-situ propellant production (ISPP) plant that would generate fuel from the Martian atmosphere for a return trip from Mars. This is followed by a presentation of advanced autonomous control techniques being developed here at NASA. These concepts are being developed and demonstrated within the context of a wide variety of mission concepts. In this paper, we will discuss how these ideas are being applied to the problem of autonomously controlling an ISPP plant.

Chemical Processes For In-Situ Resource Utilization On Mars

The processes currently being studied for the production of propellant, oxygen, and water on Mars are the following:

- Solid oxide electrolysis
- Sabatier reactor
- Reverse water gas shift reactor

These will be briefly described, in turn.

Solid oxide electrolysis

This process generates oxygen from the predominantly carbon dioxide atmosphere of Mars using solid oxide electrochemical cells. An oxygen ion conductor, such as yttria-stabilized zirconia (YSZ) electrolyte, is sandwiched between porous electrodes, e.g., platinum, to form an electrolysis cell. Carbon dioxide is split into carbon monoxide and oxygen, and the oxygen is pumped electrochemically from the cathode to the anode. This endothermic net cell reaction is as follows:

$$2CO_2 = 2CO + O_2 \qquad (1)$$
$$H = 566 \text{ kJ/mol}, \quad G = 514 \text{ kJ/mol}$$

The solid oxide electrolysis approach has the advantages of producing 100 percent pure oxygen, since the transport process in the electrolyte is solid state. Unlike the Sabatier process, oxygen is produced from the atmosphere without the need for any consumable or intermediary raw materials that are brought from Earth. This process produces oxygen and carbon monoxide in the proper stoichiometric ratio for combustion in a rocket motor. Initial design calculations by NASA engineers for Mars sample return missions indicate that the low specific impulse offered by a CO/ O_2 rocket does not warrant its use for ascent vehicle propulsion. However, recent works by Diane Linne at the NASA Glenn Research Center and Orbitec's solid CO / LOX rocket, seem to indicate that a CO/ O_2 rocket may indeed be suitable for a Mars ascent vehicle.

A brief summary of the solid oxide electrolysis process is provided here. Detailed descriptions of the principle of operation can be found elsewhere (Sridhar[1]). The oxygen generator works on the principle of solid oxide electrolysis. At elevated temperatures, solid oxide electrolytes such as yttria-stabilized zirconia become oxygen ion conductors. The basic configuration of the electrolysis cell is shown in Figure 1. At the cathode, CO_2 dissociates to form CO and O. The oxygen atom reacts with the incoming electrons from the external circuit to form an oxygen ion. The oxygen ion is conducted through the vacancies in the crystal structure of the electrolyte to the anode. At the anode, the oxygen ion donates the electrons to the external circuit to form an oxygen atom. Two oxygen atoms combine to form an oxygen molecule at the anode side of the cell.

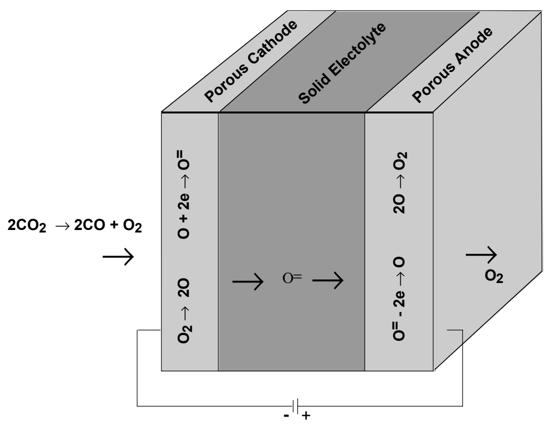

Figure 1. Principle of Operation of a Solid Oxide Electrolyzer.

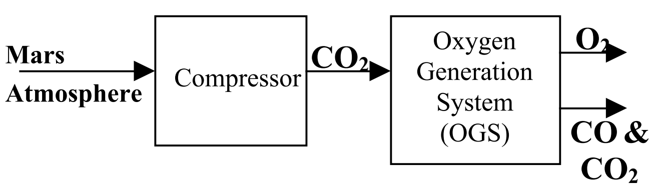

Figure 2. Simplified Oxygen Generation Plant (OGS) Flow Schematic

Figure 2 shows a simple schematic of an oxygen generation plant that utilizes the solid oxide electrolyzer technology. Since the Mars atmosphere is at 8 hPa ambient pressure, a front end compressor is used to get enough throughput in the electrolyzer. An oxygen generator based on this technology is manifested to fly on the 2001 Mars Surveyor Lander. Figure 3 is a photograph of the engineering model of the flight unit hardware.

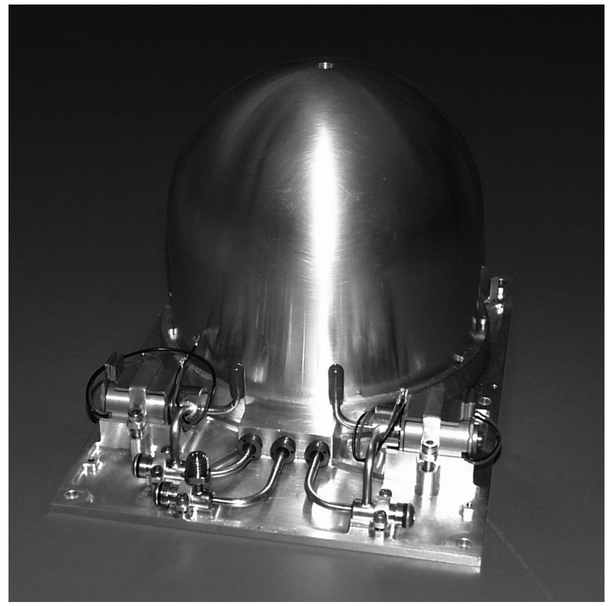

Figure 3. A Photograph of the OGS in the Development Unit Configuration.

Sabatier / Water Electrolysis

The other prime technology candidate for the production of in-situ propellant on Mars is the Sabatier/Electrolysis (SE) system. The subsystem components of the SE system have long been known to the chemical industry. Space-qualified components for the subsystems have been available through DoD and NASA funded programs, mainly for closed loop and/or regenerative life support systems. Such systems have been developed by Hamilton Standard in the sixties and later by Allied Signal, Boeing, and Dornier. There is also a rich heritage of such systems from the former Soviet space program. Using such systems for Mars propellant manufacture was first suggested by Ash et al in his seminal 1976 paper. Experimental work on integrated SE systems designed for Mars propellant manufacture began in 1993, with funding support from the New Initiatives Office at the NASA Johnson Space Center. In addition, a full scale working unit (for a MSR mission application) was built by Robert Zubrin, Steven Price, and Larry Clark at Martin Marietta Astronautics (now Lockheed Martin Astronautics) in Denver, Colorado.

The process works as follows: carbon dioxide acquired from the Martian atmosphere is reacted with hydrogen according to reaction (2)

$$4H_2 + CO_2 = CH_4 + 2H_2O \quad (2)$$
$$DH = -40 \text{ kcal/mole}$$

Reaction (2), known as the "Sabatier reaction," is highly exothermic, and has a large equilibrium constant ($\sim 10^9$) driving it to the right. It occurs spontaneously in the presence of either a nickel or ruthenium catalyst (nickel is cheaper, ruthenium is better) at temperatures above 250 C. (Typical reactors operate with peak temperatures around 400 C in the forward reaction zone, declining to 200 C at the exit). The methane and water produced by reaction (2) are easily separated in a condenser. The methane is then liquefied and stored, while the water is electrolyzed in accord with:

$$2H_2O = 2H_2 + O_2 \quad (3)$$
$$DH = +57 \text{ kcal/mole}$$

The oxygen so produced is liquefied and stored, while the hydrogen is recycled back into the Sabatier reactor to produce more methane and water.

The primary disadvantage of the SE system is the need to import hydrogen. This requirement is especially difficult on the MSR mission, where the relatively small tank sizes employed increases the tank surface area/volume ratio, increasing heat-leak and thus boil-off, making transport of the required hydrogen to Mars difficult. The SE process, operating alone, produces 2 kg of oxygen for every one kg of methane. But the optimal mixture ratio to burn O_2/CH_4 in a rocket engine is not 2/1 but about 3.5/1, where an engine specific impulse as high as 380 s can be achieved. If oxygen is also required for life support, then the ratio has to be greater than 3.5/ 1.0. If the SE process is acting alone, the only way to achieve this mixture ratio is to throw away some of the methane produced. This is undesirable due to the cost of carrying hydrogen from Earth and also the power required for propellant processing.

Reverse water gas shift reactor

The reverse water gas shift (RWGS) reaction has been known to chemistry since the mid-1800's. While it has been discussed in the literature as a potential technique for Mars propellant manufacture, there is very little experimental work done to date to demonstrate its viability for such applications. The RWGS reaction is given by equation (4).

$$CO_2 + H_2 = CO + H_2O \quad (4)$$
$$DH = +9 \text{ kcal/mole}$$

This reaction is mildly endothermic and will occur rapidly in the presence of an iron-chrome catalyst at temperatures of 400 C or greater. Unfortunately, at 400 C the equilibrium constant Kp driving it to the right is only about 0.1, and even at much higher temperatures Kp remains of order unity. There is thus a significant problem in driving the RWGS reaction to completion.

The practical difficulties of the RWGS scheme are challenging. It is important to note that the H_2O shift reaction is very effective in the forward direction, i.e.,

$$CO + H_2O = CO_2 + H_2. \quad (5)$$

It is exothermic with a 99% CO conversion in a single pass. Reversing the H_2O shift reaction requires special catalysts and temperature controls to prevent the endothermic reaction from reverting to the forward direction. A typical single-pass conversion of CO_2 and H_2 would be in the range of 10%. To obtain a conversion in the 90% range would require multiple cycles. While conversion rates could improve somewhat if the products could be separated and removed from the hot zone of the reactor, this separation and removal seem very difficult to achieve. Because of the low single-pass conversion, the typical reactor output would include a mixture of H_2, CO_2, CO, and H_2O. The H_2 and H_2O must be recovered to conserve the H_2 brought from Earth. The CO must be separated from the CO_2 to be recycled in order for the subsequent pass to reach the 10% conversion. The recovery of H_2O by the near-freezing condensation is quite effective for single-pass systems. However, the multiple recycles of the RWGS reactor will add a significant quantity of residual H_2O being rejected with cold products that, if not recovered by other means, will impact the H_2 conservation. The membrane recovery of H_2 from the product stream is based on a partial pressure differential diffusion. Thus a vacuum pump is required for separation and even then a 10% loss would be expected in any practical design.

In the next section concepts for autonomous control of such an in-situ propellant production plant, as well as an example of current practice will be discussed.

Autonomous Control Concepts

The ability to autonomously monitor and control complex devices such as an ISPP plant is critical to NASA's ability to accomplish many of its long term exploration goals. From the beginning of the Space Program in the late 1950's, control of spacecraft and systems have been managed by a large number of highly trained, ground control personnel. This has its roots in the limited capability/massive size of computers of that early period. Instead, sensor data were telemetered to the ground, where a room full of systems experts would monitor each individual system's health and send commands to the spacecraft, directly or via an astronaut. Over the past forty years there has been a radical shift in this paradigm, resulting from the advent of advanced computer technology. Automation has eased the burden of the ground controller and the astronaut, but often the tasks performed by the software are still quite rudimentary. This is because of both computational resource limitations and the difficulty encountered when trying to develop, test, and validate software that provides the required functionality. As we move outward in the solar system, beyond the Earth-Moon system, the physical and fiscal realities of space exploration will require new control technologies.

Conceptually, the task of controlling a device such as an ISPP is simply one of maintaining the system in a stable state while commanding transition of the device through a series of configurations designed to accomplish a sequence of goals in some optimal fashion. This task, however, is often quite difficult due to the normal variations that occur within both the process and the environment, limited observability into the current state of the device and the potential of abrupt failures and degraded component performance. Traditionally, these problems have been solved through the use of a tiered architecture comprised of three levels, as shown in figure 4:

1. ***analog and embedded feedback controllers*** to perform low-level regulatory functions,
2. ***higher-level system software*** to perform nominal command sequencing and threshold monitoring to detect and respond to off-nominal conditions, and
3. ***humans*** to generate the command sequences, monitor the state of the device to detect off-nominal conditions, diagnose failures when they occur and select recovery actions in response to these failures.

While the capabilities provided by the system level software have substantially increased over the past 40 years, the complexity of the missions undertaken by NASA has also increased. As a result, the requirements levied on the ground control team have increased, thus requiring larger ground support teams. This paradigm begins to break down for future planning, due both to time delays in communication with Mars as well as the cost of maintaining a large ground support team. As a result, the role traditionally performed by the ground support team is being shifted to the system level software,

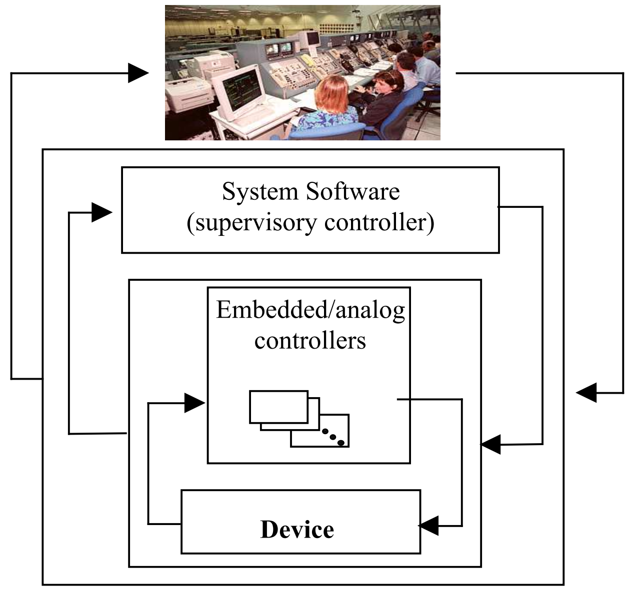

Figure 4: Tiered control architecture

thereby drastically increasing the functionality required of this component. Figure 5 shows how the functionality provided by these three different components has shifted over the years, and where it is expected that the responsibility will lie when supporting a human expedition to Mars.

Currently, the system level software is developed by engineers who use their commonsense understanding of the hardware and mission goals to produce code and control sequences that will allow a spacecraft, or other system, to achieve a particular goal while allowing for some (usually very small) amount of uncertainty in the environment. In developing this code, the engineer must reason through complex sub-system interactions to generate procedural code that can account for all the different combinations of failures and off-nominal conditions that might occur. As the functionality that is required of the system-level software increases, development, test, validation and maintenance of this software using this traditional approach becomes very difficult, if not impossible, due to the myriad of off-nominal conditions that the software is expected to handle. Furthermore, as the engineers gain a better understanding of how the device is behaving after deployment, it is often quite difficult to update the code to reflect the additional information that has been obtained.

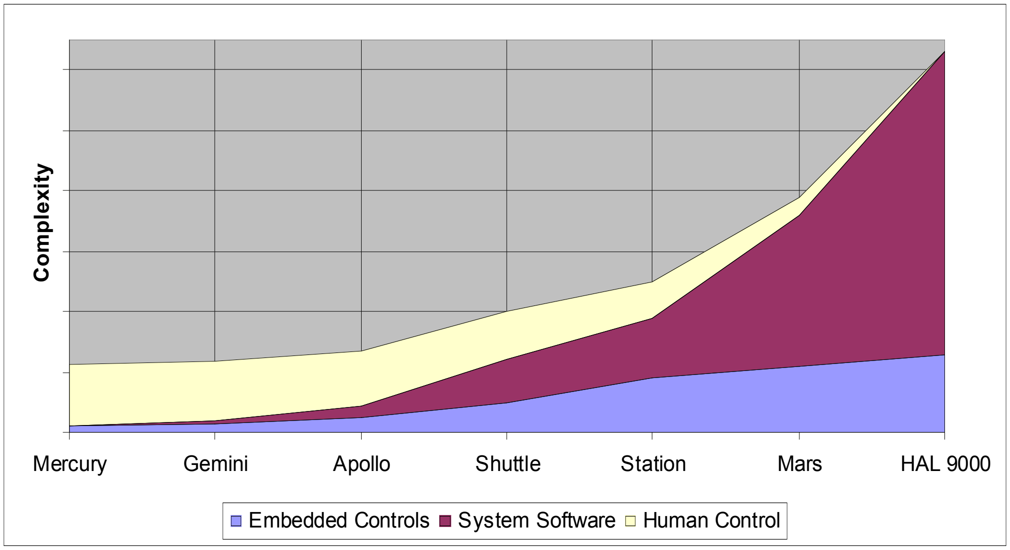

Figure 5: Progression across missions of the tasks performed by humans, system software and embedded controllers.

Artificial Intelligence and Autonomous Control

As attempts are made to automate the processes that are traditionally performed by humans when monitoring and controlling these devices, it is clear that it will often be necessary to replicate the sophisticated inferencing capabilities exhibited by humans when performing this task. Over the last 40 years, the field of artificial intelligence has been developing a variety of automated techniques that emulate a human's reasoning ability [15]. While the systems developed are far from performing at the visionary level exhibited by the HAL 9000 computer in *2001: A Space Odyssey,* recent accomplishments such as the victory of Deep Blue over Kasparov have demonstrated that sophisticated inferencing tasks can be automated.

One of the most notable recent information technology accomplishments within NASA is the development and demonstration of the Remote Agent (RA) autonomous control architecture. This is part of the Deep Space One mission within the New Millennium Program (NMP). The Remote Agent architecture, developed collaboratively between NASA Ames Research Center (ARC) and the Jet Propulsion Lab (JPL), combines high-level planning and scheduling, robust multi-threaded execution, and model-based fault detection isolation and recovery, into an integrated architecture that is able to robustly control a spacecraft over long periods of time [6, 12].

One of the primary components of the Remote Agent architecture is the Livingstone model-based health management system. Livingstone is an advanced inference engine that uses a high-level declarative model of a physical device to monitor the state of that device, detect off-nominal behavior, isolate failures to individual components, and reason about alternative recovery actions. The key benefit provided by Livingstone is the use of a first-principles model that describes the behavior of each component within the device and the interactions between the components [6,7,8]. By reasoning generatively about the behavior of the device using the model, Livingstone is able to detect failures whenever a discrepancy occurs between the observations and predictions. In addition, Livingstone is able to use the same model to generate the most likely hypothesis that is consistent with the observations and to select the optimal reconfiguration action for recovering from the failure. Thus, with the use of a model of the device, Livingstone is able to reason about novel combinations of failures and avoids the need to develop mission-specific code that must pre-enumerate all of the various failure combinations that might occur. Furthermore, the models used by Livingstone are easy to update and maintain and can often be reused across missions, thus further reducing the software development costs while increasing the functionality provided.

The Livingstone Model-based Health Management System

Recently, Ames Research Center and Kennedy Space Center have been investigating how the core ideas developed within the Livingstone system can be applied and extended to autonomously control an ISPP plant. In this section, a more detailed description of the Livingstone system will be provided by explaining how these ideas are being applied to the control of an ISPP plant.

Modeling Paradigm

As mentioned above, Livingstone uses a high-level, compositional model to identify the components within the device and the relationships between the components. This model is used for prediction, fault detection, isolation and reconfiguration. A Livingstone model is comprised of a set of components and connections between these components. Each component is modeled using a set of discrete valued variables. For example, a valve might be modeled using the variables *flow-in*, *flow-out, pressure-in* and *pressure-out* with values such as *zero, low, nominal, high.* For each component, a set of *modes* is defined identifying both the nominal and failure modes of the device. For each mode, a set of constraints is specified that restrict the values of the component variables whenever the component is in that mode. Thus, a valve might be modeled using modes such as *open, closed, stuck-open* and *stuck-closed* where the model of the valve in the *open* mode might be:

flow-in = flow-out
pressure-in = pressure-out

In addition, to the description of the behavior of the device for each mode, the model also includes transitions between modes with guard conditions describing when the transition occurs along with relative probabilities on the likelihood of the transitions. These probabilities are used to provide information about the relative likelihood of various failures. Figure 6 shows how a valve model might be represented as a finite-state automaton in which the labels on the links correspond to device commands.

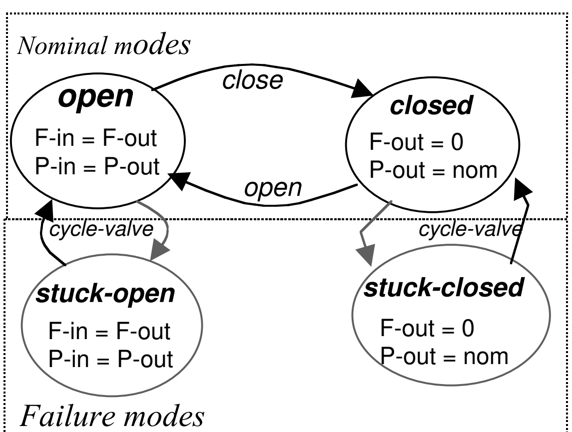

Figure 6: Valve model

One of the key benefits of this modeling paradigm is that the modeler is only responsible for describing the *local* behavior of each component and the relationships that exist between components. Livingstone then uses this specification to compose a larger, system model that can be used to reason about the *global* behavior of the entire system given the mode of each component. Furthermore, since the models are qualitative in nature it is often clear as to how to develop many of these models, even before the hardware design is complete.

Inference with Livingstone

Given a model of the form described in the preceding section, Livingstone performs two main tasks: 1) inferring the current state of the device given the limited available sensor information; and 2) identifying an optimal set of commands for system reconfiguration following a failure or external perturbation that transitions the system out of the desired state. At first glance, it might appear that the valve model described in the previous section is too simple to be of much use in performing these tasks. In practice, however, qualitative models of this nature have been found to be quite effective for detecting a wide range of likely failures. In fact, it is exactly these types of models that humans often use when reasoning about the current state of a device.

To estimate the current state of the system, Livingstone monitors the sequence of discrete commands that are issued to the ISPP plant. This allows the tracking of the expected state of the device and compares the predictions generated from its model against the observations received from the sensors. Once a discrepancy occurs, Livingstone performs a diagnosis by searching for the most likely set of component *mode assignments*[1] that are consistent with the observations. This is done using a search technique called *conflict-directed, best-first search,* developed within the model-based reasoning area of the artificial intelligence community. This search technique is able to efficiently search an exponentially large set of failure modes by focusing on the components whose state results in a conflict between the observations and the predictions. Within the Deep Space One experiment, this search technique was able to identify the most likely component failure within a couple hundred milliseconds for most device failures.

Once the state of the system is identified, the same search techniques can then be used when reasoning about reconfiguration commands to identify the lowest cost set of commands that can be issued to transition the system into a state that satisfies the current operational goals provided by a higher level executive.

[1] The *state* of the device is represented by identifying the current *mode* of each component in the model.

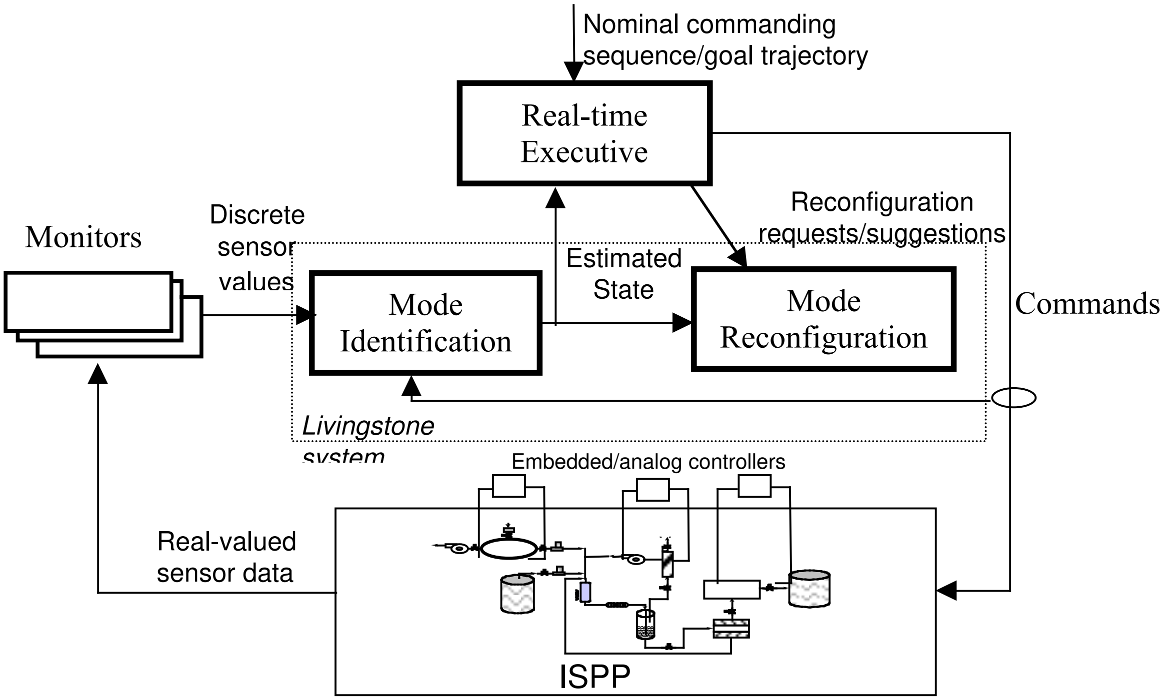

Figure 7: ISPP Control Architecture

Model-based Control of an ISPP Plant

The current system architecture being developed to control an ISPP plant combines the Livingstone health management system with a real-time executive for commanding the device. Figure 7 shows a block diagram of this architecture. At the lowest level, embedded and analog controllers are used to perform low-level regulatory functions. Nominal commanding of the ISPP is performed by a real-time executive. As these commands are sent, the *Mode Identification* component of the Livingstone system monitors the commands to identify the expected state of the plant. The real-valued sensor data is processed by a set of *monitors* that abstract the real-valued information from each sensor into a set of a-priori defined discrete values such as *high, medium, low* or *plus, zero, minus*. When a failure occurs, the real-time executive is notified. For failures that require a very fast response time, the real-time executive might respond reactively in a predefined manner. For other failures, however, the executive requests a sequence of reconfiguration commands from the *Mode Reconfiguration* component of Livingstone and then continues commanding the device.

Demonstration Testbeds

To support the development and evaluation of this technology, both a hardware testbed and a simulation-based testbed are currently being developed. For the hardware demonstration, a testbed that uses Reverse Water Gas Shift (RWGS) to generate C0 and O_2 from the Martian atmosphere will be developed.

As previously described, the RWGS reactor converts CO_2 and H_2 into CO and H_20 at a 10% efficiency rate. Thus, the outflow stream from the reactor contains liquid water, and gaseous CO, CO_2 and H_2. After exiting the reactor, a condenser is used to separate the water from the gases and then an electrolysis unit is used to separate the hydrogen from the oxygen. The oxygen is then stored while the hydrogen is fed back into the RWGS reactor. Similarly, the CO is extracted from the gas mixture and the remaining CO_2 and H_2 are routed back into the RWGS reactor. Control of an RWGS system is actually quite straightforward since the system has a limited number of components. A full-scale ISPP device, however, would require various other components along with redundant valves and flow controllers. As the number of components within the system increases, the probability of a failure increases and the discrete control problems become more complicated. The RWGS testbed, however, allows one to demonstrate how these techniques can be used to control a real physical device for an extended period of time.

At the same time, a simulation testbed that uses a hypothetical ISPP flight system is being developed. It is based upon a Sabatier/Electrolysis system for converting CO_2 and H_2 to CH_4 and O_2, coupled with a pair of zirconia cells for generating the extra O_2 that is required. The design of the system for this simulation tries to balance many of the design considerations (e.g. mass and power limitations) that must be satisfied by a true flight article, while also including redundant components and the additional margin that would be require to ensure a successful mission.

The simulation for the software testbed is based upon the hybrid concurrent constraint (hcc) programming language developed at Ames Research Center [13]. Hcc is a hybrid discrete/continuous modeling language that combines the benefits of a discrete event simulation with the precision provided by a dynamic simulation using ordinary differential equations. In addition, extensions to hcc are being developed that will allow the performance of a stochastic simulation that is able to inject faults, component degradation and variable process noise, thus allowing the testing of the control system under a broad range of conditions.

In addition to the demonstration of these techniques within the context of autonomous control of an ISPP plant, these techniques are also being applied to a variety of other missions. These missions include autonomous control of an advanced life support system and a space interferometer as well as integrated health management systems for two experimental reusable launch vehicles, the X-34 and X-37.

Verification and Validation of Livingstone Models for ISPP

Autonomous systems present difficult verification and validation (V&V) challenges. In contrast to conventional open-loop systems, they arbitrate many resources on-board using their own decision procedures. Thus the range of possible situations becomes very large and uncontrollable from the outside, making scenario-based testing much less efficient.

Model checking is an analytical V&V technique based on exhaustive exploration of all possible executions of a model of a dynamic system. It provides a much better coverage than traditional testing, and can be applied at an earlier design stage, thus reducing the costs of repairing errors. Model checking is limited by state space explosion, and the number of cases to be explored grows exponentially in the size of the system.

In collaboration with Carnegie Mellon University (CMU), NASA Ames is developing a translator that feeds Livingstone models into the SMV model checker from CMU [16]. SMV uses advanced symbolic computation techniques to represent and process huge state spaces in a compact and efficient way. The properties to be verified, expressed in a powerful temporal logic (CTL) or using pre-defined specification patterns, are added along the Livingstone model and similarly processed by the translator. The translator thus enables model checking of Livingstone models by their developers in their Livingstone environment, without requiring them to use or learn the input language of the SMV tool.

This model checking technology will be used, along with traditional scenario-based technology, by the ISPP Autonomous Controller team at the NASA Kennedy Space Center to support the development of the Livingstone model of ISPP. It is possible to find out, for example, whether the model allows a given configuration to happen. It takes less than a minute for SMV to answer this query, while symbolically analyzing a reachable space of the order of 10^{50} states!

Software verification and validation is thus a key element in the design and development of advanced software, as well as assuring the safety and robustness of space missions. The technology to facilitate the automated development of software systems for such applications must be a continuing priority.

Concluding Remarks

The coming decades will see many new opportunities to expand human presence in the solar system. As we penetrate deeper into space we must implement space exploration missions at lower cost, with greater safety, and achieve greater scientific return. Fortunately we are able to develop powerful new computer/information systems to meet these needs. To this end, NASA has embarked on the Intelligent Systems Program, which will research and develop new capabilities in the areas of automated reasoning, human-centered computing, intelligent use of data, and concepts for revolutionary computing, such as biomemetics and nanotechnology. Ultimately there will be a distinctly new balance of work between humans and intelligent machines, where tasks will reside with the entity having the best capability, irrespective of mechanism of origin. This will provide the total human-machine system with the capability to explore far beyond the limits of the present.

References

Some of the following papers may be found on the World Wide Web at http://ic-www.arc.nasa.gov/ic/projects/mba/

1. K. R. Sridhar, *J. Propulsion and Power*, **11**, 6 (1995).

2. R. Zubrin and R. Wagner. *The case for Mars: The plan to settle the Red Planet and why we must.* The Free Press, 1996.

3. S. J. Hoffman and D. I. Kaplan, editors. *Human Exploration of Mars: The Reference Mission of the NASA Mars Exploration Study Team*. NASA Special Publication 6107. July 1997.

4. B. C. Williams and P. Nayak, A Model-based Approach to Reactive Self-Configuring Systems, *Proceedings of AAAI-96*, 1996.

5. D. E. Bernard et Al. Design of the Remote Agent Experiment for Spacecraft Autonomy. *Proceedings of IEEE Aero-98*.

6. J. de Kleer and B. C. Williams, Diagnosing Multiple Faults, *Artificial Intelligence*, Vol 32, Number 1, 1987.

7. J. de Kleer and B. C. Williams, Diagnosis With Behavioral Modes, *Proceedings of IJCAI-89*, 1989.

8. J. de Kleer and B. C. Williams, *Artificial Intelligence*, Volume 51, Elsevier, 1991.

9. D.Schreckenghost, M. Edeen, R. P. Bonasso, and J. Erickson. Intelligent control of product gas transfer for air revitalization.. *Proceedings of the 28th International Conference on Environmental Systems* (ICES), July 1998.

10. R. P. Bonasso , R.J. Firby, E. Gat, D. Kortenkamp, D. Miller and M. Slack. Experiences with an architecture for intelligent, reactive agents. *In Journal of Experimental and Theoretical AI*, 1997.

11. B. C. Williams and P. P. Nayak. Immobile Robots: AI in the New Millennium. In *AI Magazine*, Fall 1996.

12. N. Muscettola. HSTS: Integrating planning and scheduling. In Mark Fox and Monte Zweben, editors, *Intelligent Scheduling*. Morgan Kaufmann, 1994.

13. V. Gupta, R. Jagadeesan, V. Saraswat. *Computing with Continuous Change*. Science of Computer Programming, 1997.

14. G. M. Brown, D. E. Bernard and R. D. Rasmussen. Attitude and articulation control for the Cassini Spacecraft: A fault tolerance overview. In *14th AIAA/IEEE Digital Avionics Systems Conference*, Cambridge, MA, November 1995.

15. S. Russel and P. Norvig *Artificial Intelligence: A Modern Approach*, MIT Press 1995.

16. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and J. Hwang, "Symbolic model checking: 10^20 states and beyond", *Information and Computation*, vol. 98, no. 2, June 1992, pp. 142--70.