N89-19839

AUTONOMOUS SATELLITE COMMAND AND CONTROL: A COMPARISON WITH OTHER MILITARY SYSTEMS

Robert J. Kruchten and Wayne Todd, Capt, USAF
Rome Air Development Center/COAD
Griffiss Air Force Base, New York 13441-5700

ABSTRACT

Existing satellite concepts of operation depend on readily available experts and are extremely manpower intensive. Areas of expertise required include mission planning, mission data interpretation, telemetry monitoring, and anomaly resolution. The concepts of operation have evolved to their current state in part because space systems have tended to be treated more as research and development assets rather than as operational assets. These methods of satellite command and control will be inadequate in the future because of the availability, survivability, and capability of human experts. Because space systems have extremely high reliability and limited access, they offer challenges not found in other military systems. Thus, automation techniques used elsewhere are not necessarily applicable to space systems. RADC has developed a program to make satellites much more autonomous using a variety of advanced software techniques. The purpose of this paper is to present the problem the program is addressing, some possible solutions, the goals of the RADC program, the rationale as to why the goals are reasonable, and the current program status. Also presented are some of the concepts used in the program and how they differ from more traditional approaches.

1. Introduction

Operation and control of satellites can be divided into two major areas: health/status and mission. Health/status is a broad definition covering all activities not directly concerned with executing the primary mission of the satellite. This includes power control, thermal control, attitude control, telemetry collection/formatting, station keeping, overall monitoring to insure nominal operation, and anomaly resolution. Anomaly resolution consists of detecting and diagnosing a real or apparent anomaly(ies), developing recommended courses of action to resolve the anomaly(ies), executing one or more courses of action, and observing the results of that action(s). These health/status activities are somewhat generic between different satellites although the details will differ.

The mission activities may involve situation assessment, scheduling, tracking, processing preplanned and real-time user requests, and mission data interpretation. These mission activities are not as generic as health/status activities because they are more dependent on the type of satellite. For example, mission activities of a communications satellite may be significantly different than a surveillance satellite or a navigation satellite.

Both health/status and mission are potential targets for greater autonomy. Each relies on ground support from experts with highly specialized knowledge. The basic goal of satellite autonomy is to transfer this knowledge to the satellite and employ it so the satellite achieves much greater independence from the experts and the ground support infrastructure they require. The result will be a command and control system that is both affordable and survivable for future operations.

2. Problem

Current methods of commanding and controlling both the health/status and mission activities have limitations that must be addressed to insure effective satellite operations in the future. The problem areas considered herein are cost, time delays, survivability, the increasing dependence on experts, and some unique aspects of this problem compared to other systems.

2.1. Cost

The command and control cost of satellite assets is measured in terms of the resources consumed, including manpower and facilities. Current operations are costly and the cost will likely become prohibitive in the future without significant changes in operational concepts.

As the number of space systems grows, it will become increasingly difficult, if not impossible, to find enough qualified personnel to operate and control these systems. This is especially true in the case of the technical experts who evaluate situations that are unusual, complex, and difficult to diagnose. These experts must also recommend actions; such actions may be novel solutions for unanticipated problems. A significant portion of the dollars expended in operating and maintaining space systems is for personnel. Again, the key problem lies with those experts used to deal with difficult situations. Thus, even if enough people could be found to perform the expert functions, the money required may be prohibitive. This could have a major impact on an effort such as the Strategic Defense Initiative (SDI) which may involve many more satellites than are currently supported. One way to reduce to number of experts required for satellite support is to centralize command and control facilities. However, this can also create more time delays and reduce survivability as described below. Also, centralized command and control doesn't ameliorate the need for numerous geographically-dispersed transmit/receive ground stations which experts need to maintain frequent satellite contact during difficult situations.

2.2. Time Delays

One of the concerns with centralized command and control is that it can cause delays by requiring individual users to communicate directly or indirectly with the central control. This becomes necessary because the ground control must resolve conflicts in user requests and insure the system constraints are always satisfied. Even ignoring the delays entailed by requiring central coordination of user requests, the current methods of satellite control are not particularly fast. This is especially true when unusual situations arise. That is, situations that are not covered by routine operating procedures require extensive expert analysis before resolution. Unique satellite failures are one example of an unusual situation. Many satellite failures take weeks or, in some cases, even months to totally resolve. During this time, the satellite may be safed and often cannot fully meet its mission goals. As new requirements come into play (such as SDI), response time will become more critical. Not only will it be necessary to take an immediate response to reach a safe state, but it will also be necessary to quickly reach a final decision on a situation. Fortunately, efforts to improve response time can also increase efficiency. By speeding the decision process, resources on-board the satellite can be used more effectively since more windows of opportunity are available. The improvement could be faster recovery from an anomalous condition or it may even be redirecting mission functions faster than ground controllers are able.

2.3. Survivability

Most of today's space systems utilize a centralized command and control concept for health/status and mission functions. Although this can reduce cost, it can also create survivability problems. That is, a centralized command and control facility in a fixed location is easier to target (destroy or electronically jam) than several mobile targets. Multi-node, mobile, distributed control systems for health/status have been suggested to improve survivability because such systems would be difficult to target or jam. Unfortunately, it is not possible to simply relocate all personnel controlling today's satellites into mobile systems. Many people would simply be unwilling to work in a mobile (especially remote) facility. Also, the savings in personnel and facility costs, gained through centralized control, would be lost.

2.4. Increasing Expert Dependence

Current space systems are controlled by technical personnel assisted by experts (usually from the manufacturer of the satellite) who evaluate complex situations and recommend actions. These are highly technical people with long experience in satellite design and operation. Interestingly, some efforts to introduce more automation into satellites may actually compound the need for these experts. This is because, although the automation reduces the efforts required by technicians, it can greatly increase the complexity of the satellite system in both the health/status and mission areas. This, in turn, can create a higher dependence on experts for difficult situations or problems.

Another problem is beginning to occur in space systems with long life spans, a characteristic which is generally desirable for an expensive, inaccessible asset. The problem is that a satellite can "live" long enough for the original experts to retire or move on to other programs. Incoming people are not as familiar with the satellite's design or history (i.e., heuristics of operation). This creates a greater need to capture this knowledge before the experts leave the program. Some NASA programs are recognizing this problem and attempt to create historical records of the experts knowledge.

2.5. Different Type of Problem

Satellite anomaly resolution has significant differences from the anomaly resolution of most other systems such as aircraft. By design, satellites are highly reliable and have no physical access. Whereas most anomaly resolution for an aircraft involves isolating a common failure, most anomalies in a spacecraft are unique and unanticipated. Expert systems that utilize a knowledge base built up from experience and heuristics are inadequate for most satellite problems because the typical problem is new.

2.5.1. High Reliability

Many studies have been performed to attempt to anticipate problems before they occur, thereby decreasing the dependence on the skilled experts. Unfortunately, in spite of these efforts to anticipate problems and provide procedures and equipment to resolve them, as many as 90% of all major anomalies experienced by current satellite systems have been unanticipated and required experts for resolution (Figure 1). This is probably due to the great emphasis on system reliability. That is, if a failure mode is identified, the system is designed to make the failure extremely unlikely. Thus, the failures that occur tend to have not been identified before hand.

2.5.2. Limited Access

Even though today's satellites are very complex, the limited number of actions that can be taken for any one situation creates a bounded problem. Although permutations cause this number to be large, it is much smaller than the number of possible actions that could be taken with a system such as an aircraft because the aircraft is physically accessible. Thus, space systems allow the use of techniques that would be inappropriate in other domains.

3. Potential Solutions

All of the problems discussed can be alleviated if satellites çan be made more autonomous and thereby reduce the dependence on available human experts. The need for more autonomous satellites has been acknowledged by the Satellite Control Architecture Study sponsored jointly by Air Force Space Command and Air Force Systems Command Space Division. Actually, satellite autonomy is a function of system design, hardware reliability, redundancy, the environment, and a capability to analyze and act on changes in situations.

Fortunately, new software techniques associated with Artificial Intelligence (AI) research offer the unique hope of addressing the problem. The goal is to have a system that will act as an expert would in the same situation. To be truly successful, such a system must not be limited to predefined responses to predefined situations. It must, instead, be capable of reasoning about a new situation as a real expert would.

3.1. Redundant Components

Redundant components are an essential part of any satellite autonomy concept. Without redundancy, few, if any options exist for dealing with satellite anomalies. The traditional approach for employing redundancy is to provide automatic anomaly detection and automatic switching from the anomalous component to an equivalent backup. For some anomalies, automatic switching without understanding the nature of the problem may be necessary because of the time-critical response required. One example might be switching to a backup voltage regulator because of a loss of power. Unfortunately, automatic switching is based upon assumptions which may not be valid at the time the anomaly occurs. As a result, switching to a backup component can further complicate the situation. Thus, it's important to minimize switching to backup components when the anomaly is not fully understood. Because of limited ground support, the tendency for today's satellites is to employ automatic switching whenever possible, as long as further complications are not expected. In contrast to this method of using redundant components, an on-board self-reasoning software system could approach the problem at various levels to understand the cause of the problem much as a human expert would.

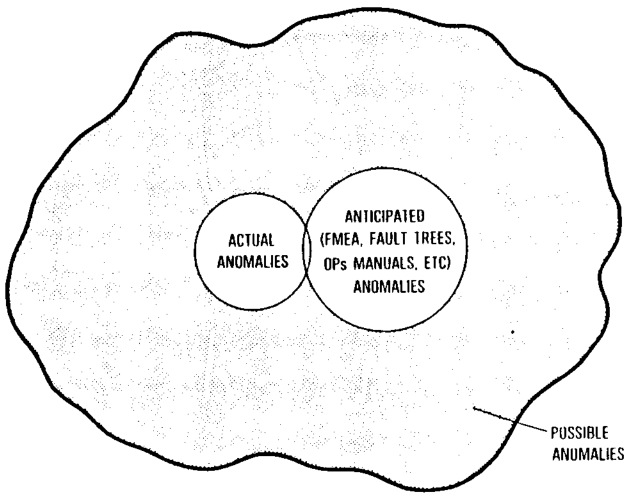

Figure 1.

3.2. Reliable Components

Today's satellites are generally composed of highly reliable components, and further advances in reliability will certainly enhance autonomy. If fact, if a component is considered reliable enough, the redundant component may be left out of the design. Nevertheless, some components are too critical to avoid redundancy no matter how reliable they are.

Another aspect of reliability is the impact on the expert supporting the satellite. A primary source of knowledge for the expert is the technical documentation for the satellite. However, it's also essential that the expert have operational experience in dealing with satellite anomalies in order to remain competent. The irony is that, while both a reliable satellite and a competent expert are desirable, the expert's competency depends, in part, on satellite failures!

3.3. Expert Systems

Many expert systems are in use today whose problem-solving performance matches or exceeds that of a human expert within a limited domain. By incorporating heuristics, or "rules of thumb," derived from a human expert(s), expert systems are able to mimic the behaviour of the human expert. These heuristics are the aggregate of experiences from which the expert has learned (sometimes the hard way). Often the expert is unable to precisely delineate a heuristic, or, even worse, will falsely delineate a plausible heuristic. Even without these problems, heuristics axiomatically reference predefined problems. Thus, an expert system by itself cannot provide autonomy for a satellite suffering unexpected problems. At best, an expert system can provide a reasonable initial attempt to isolate a problem.

The human expert, on the other hand, can fall back upon more general knowledge when simple heuristics don't solve the problem.

3.4. Neural Networks

In the area of health/status, Neural Networks (Neural Nets) have been used to diagnose satellite anomalies. The diagnosis can be almost instantaneous if the Neural Net is implemented with parallel processing. However, Neural Nets require many training examples of predefined problems and, like Expert Systems, cannot diagnose a problem that was never defined. Also, the cost of generating numerous real or precisely-simulated anomalies for the Neural Net to train on may be prohibitive. Even if the Neural Net correctly diagnoses the anomaly, it doesn't offer the potential for prescribing a solution.

In the mission area, Neural Nets may prove valuable in certain types of mission data interpretation (e.g., satellite photos). However, they don't appear to have the potential for autonomously allocating mission sensor resources to meet dynamic requirements.

3.5. Model-based Reasoning

It was stated above that AI promises significant promise to solving the described problems. Even though, the term "Artificial Intelligence" means different things to different people, several AI concepts appear to be well suited to the problem of satellite autonomy. The first of these is the idea of model based reasoning. Model based reasoning uses causal models of the system and its environment to reason about situations. The models are built using the object-oriented programming techniques. This allows concise models to be built for different physical aspects of the system (e.g. electrical models, structural models, thermal models ,etc.). Each of these models would include a deep basic understanding of the specific satellite design as well as basic physical principles. Although the final models would be specific for an individual satellite, their architecture and much of their basic knowledge is generic. Finally, they will be built with tools making it easy to adapt them from one application to another. Model based reasoning uses special techniques to determine the cause of conflicts between model predictions and actual observed events. In addition, the models are used to construct solutions to problems and to try the solutions (through simulation) before actual commands are given to the satellite. Overall the model based concept provides (for the first time) the capability to deal with unanticipated events. The system would be able to reason about variances between the observed world and the world predicted by its models. The models themselves could then be modified to more closely match the observed world. They would also serve to evaluate a goal-oriented search toward resolving any problems. The resolution would not be limited to only predefined actions but could also be new, novel actions. It would also be much more flexible. Thus, when a power system was degraded, the models would automatically adjust themselves to the new situation and would then serve to revise the constraints in power utilization procedures/schedulers. Viewed another way, the model based system tries to capture how the system should work as opposed to the traditional method of attempting to capture all possibilities of all problems. It's important to note that model-based reasoning is not limited to naturally occurring events. Hostile events that degrade or impact the system performance can be handled the same way. The system need not know all possible hostile actions, but merely that the world is not as it should be.

3.6. Natural Language

Another AI technique that is applicable to autonomous satellites is natural language. In its purest form, Natural Language programming would allow a user to converse with a machine in the same manner in which people communicate. Actually the current state-of-art for natural language is not yet that advanced. Although current language parsers are extremely useful in some applications, they are still only valid for relatively small domains. Natural language research is, however useful for applications other than natural language processing. This research has developed useful knowledge representation schemes, search techniques, and problem solutions. Some of its techniques for treating sequential events are of particular interest to satellite operations. Scripts is a programming technique that was developed for natural language understanding. It was found to be impossible to understand language without understanding the context of the situation. Words, sentences, and whole thoughts require an understanding of the situation to eliminate ambiguities. Scripts can be written for common situations to establish a loose relationship between events. These scripts can be referenced to understand the dialog. A similar concept can be used to understand the telemetry data of a satellite. For example, a script could be written to cover the events that usually occur when going into an eclipse or for some possible hostile events. These scripts then serve as possible references to resolve ambiguities in the telemetry data. A key feature of the scripts is that they serve as a reference framework and not as a mandatory sequence of events. This gives them great flexibility.

4. RADC Satellite Autonomy (SA) Program

4.1. Background

RADC has undertaken a major effort in the area of Artificial Intelligence, initially concentrating on tactical and intelligence applications. During 1985-86 RADC worked with the Air Force Satellite Control Facility and Space Division to fund some studies and research into the application of AI into satellite systems. In addition, Aerospace Corporation has done some research in the same area. Finally, many aerospace contractors have also been conducting independent research and development in this same field. All of these studies concluded that artificial intelligence techniques combined with advances in computer speed, memory density, and architecture promise significant progress toward solving satellite autonomy problems.

Studies by Space Division and Space Command also defined roadmaps for future satellite control systems and thus help show how to incorporate the technology advances. These studies advocate a phased approach toward achieving autonomy beginning first with systems that interact with humans.

4.2. Description

Rome Air Development Center (RADC) has worked with the Strategic Defense Initiative Organization (SDIO), Air Force Satellite Test Center (AFSTC), Air Force Space Technology Center (AFSTC), Space Division (SD), Air Force Astronautics Laboratory (AFAL), and Air Force Space Command to create a program that uses AI techniques to achieve satellite autonomy. The program is a multi-phased effort that first shows the feasibility of the concept and then builds a prototype. This will be done using limited ground systems and then full ground prototypes for an existing satellite. The long term goal of this program is an on-board autonomous design. This goal includes both health/status and mission functions. The output of the program will be a system design that can be given to a System Program Office for incorporation into a new satellite design. This program is not designed to extend the AI technology, but it will make use of the most current technology methods and ideas. This is not simply a program to develop better built-in-test or data compression/data analysis, but is a program to develop a generic system capable of reasoning about itself or its environments.

4.3. Goals

4.3.1. Handle Unanticipated Situations

The ability to handle the unanticipated and the ability to generate novel solutions are key issues. This can best be explained by example. A typical function currently performed by people might be the scheduling of power utilization. Certainly, an algorithm could be made to perform this function, but as time progresses the satellite capabilities might change (systems fail, etc.) or the operational requirements may change. Thus, the real key to autonomy in this case is the ability to modify the scheduling procedure as events change. That is, a system should be able to develop and implement new scheduling constraints as the situation dictates. This is particularly true for situations where multiple users are tasking the same system. Therefore, a good scheduling system should be a flexible, reasoning type of system.

4.3.2. Generic Solution

The RADC SA program will be developed and demonstrated initially for three different types of satellites. However, the emphasis is to prove the viability of particular advanced software techniques that can be applied to any satellite. These techniques can then be incorporated into the initial design of a satellite so that, once deployed, elaborate ground support is not required.

4.3.3. Cost Reduction

Autonomous satellites that can, among other things, recover to a maximum extent from anomalies, respond to contention for resources from users, and adjust nominal operations based on degrading components will provide a tremendous cost savings in terms of manpower and facilities. Note that the total numbers of people and facilities to support satellites will likely increase in the foreseeable future, even with highly autonomous satellites. This is because the total number of operational satellites will dramatically increase. However, the cost reduction goal of the RADC SA program is to reduce the manpower/satellite and facilities/satellite ratios (Figure 2). This will make future operations both possible and affordable.

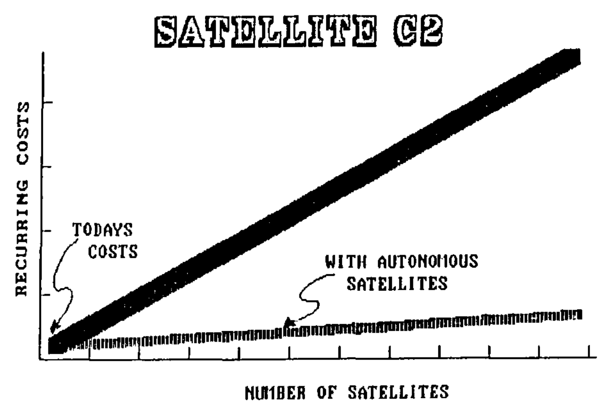

Figure 2.

4.4. Approach

Even though health/status and mission operational concepts can be considered separately, some integration and dependency exists. For example, mission sensors have telemetry which must be analyzed in the same manner as support function telemetry. Thus, both health/status and mission must be considered if greater autonomy is to be achieved.

One can reasonably expect components comprising the health/status and mission areas to improve in both reliability and redundancy. However, the RADC SA program will pursue greater autonomy through a synergistic application of expert systems, model-based reasoning, and natural language concepts as previously discussed.

4.5. Unresolved Issues

As described above, advanced software

techniques appear to offer promise toward solving the difficult problem of autonomy. However, there are some issues that must be explored before these techniques can be dictated for use on future satellites. Some of these issues are described below.

4.5.1. Scaling

The first of these issues is scaling. Although small self reasoning model-based systems have been built it is not yet proven that the results on these small systems can be scaled to larger systems such as an entire satellite. Fortunately, as stated above, the satellite problem is bounded. In addition, hierarchical structures and parallel processors may be used if scaling becomes a major problem.

4.5.2. Satellite Configuration

Another issue concerns the configuration of the satellite. A model-based system must always track the current configuration of the satellite including the status of expendables. Unfortunately, the exact status of all systems on the satellite may not be known (especially if anomalies have occurred). Thus, the model-based system may have to infer the status of systems. Since multiple inferences are often possible, the model-based system must be able to track multiple configurations simultaneously. As further information becomes available, the models must be revised and incorrect representations must be eliminated.

4.5.3. Data Archiving

Data archiving is yet another issue. Proper interpretation of events requires the system to archive data for later use. This is especially necessary for detecting gradual changes in system performance. Archiving can become extremely expensive in memory and obviously not all data can be archived. The problem lies in deciding how much to archive, how to compress it, and how to purge it.

4.5.4. Environment Modeling

Modeling the external environment will be necessary for a robust model-based system design. However, the external environment is not nearly a straightforward model as the satellite itself. Ideally, the external environment would be represented in enough detail to resolve environmentally induced situations, but not in so much detail as to unnecessarily complicate the system. Early studies have shown the advantage of supplementing a model-based system with an expert system (to heuristically handle anticipated situations) and this will probably be necessary for models of the environment.

4.6. Status

Currently, the Satellite Autonomy program is in its first phase. The three prime contractors are Boeing, Ford Aerospace, and TRW. These contractors are developing a system to prove the feasibility against real satellites. In this first phase, a software system capable of doing all health/status for three major satellite subsystems will be developed. In addition, a system will be developed to perform one major mission function in this phase. These systems will be tested against existing robust satellite simulators and evaluated by current operational users. The exact designs for satellite autonomy are currently being developed, but several things have been shown:

4.6.1. Goal Attainment

The perfect software that can totally mimic the human expert is not yet in sight. However, it does appear possible to achieve most aspects of each of the goals described above. In addition, the software developed will exceed the human expert in many situations. This occurs because the systems are more thorough and can reason faster than people. Preliminary systems are addressing the key problem areas. They have found novel solutions to problems and have been able to reason about unforeseen situations.

4.6.2. AI versus Conventional Processing

An autonomous system will ultimately use both conventional and AI techniques. The problem is large and complex, but it still appears to be generally solvable.

5. Summary

The problems with operation and control of satellites are real, here today, and getting worse. Advanced techniques offer unique promise in this area. Although these techniques have limitations and concerns, they will go a long way toward solving these problems. The RADC Satellite Autonomy Program is a challenging program to use advanced techniques to develop a self reasoning system that can reduce the experts/system ratio currently necessary to operate and control satellites.