Challenges and Lessons Learned in the Application of Autonomy to Space Operations

David J. Forrest

NASA Space Operations Management Office
Johnson Space Center / Mail Code TE
2101 NASA Road 1
Houston, TX 77058
david.j.forrest1@jsc.nasa.gov

Joseph I. Statman

Telecommunications and Mission Operations Directorate
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099
joseph.i.statman@jpl.nasa.gov

Abstract

NASA's Space Operations Management Office (SOMO) is working toward a goal of providing an integrated infrastructure of mission and data services for space missions undertaken by NASA enterprises. A significant portion of this effort is focused on reducing the cost of these services. We are interested in the potential of autonomy to reduce operations costs. Some attempts have already been made to apply autonomy and automation in these areas in the past with varying degrees of success. We present brief case histories and the lessons inferred from them. Combining this past experience with anticipated future needs, we attempt to clarify the challenges that must be met in order to realize the benefits of autonomy.

Introduction

NASA's Space Operations program, managed by the Space Operations Management Office (SOMO) is working toward a goal of providing an integrated infrastructure of mission and data services for space missions undertaken by NASA enterprises. A significant portion of this effort is focused on reducing the cost of these services. We are interested in the potential of autonomy to reduce operations costs. SOMO services support space missions, but are not part of the mission objectives; therefore the level of acceptable risk is very low. In fact, SOMO could be effectively prevented from applying autonomy if customers merely *perceive* it as adding risk to their mission(s). We are interested in this workshop from the standpoint of understanding what can be done to realize the potential cost savings due to autonomy while maintaining acceptable risk and serving the needs of our customers. We would like to present our lessons learned so far in adopting autonomy and automation, which we think will contribute to clarifying the challenges facing the use of such technology.

SOMO provides services to a diverse and ambitious set of mission customers. Many of these missions are groundbreaking missions for which communications, data, and other operations requirements sometimes cannot be clearly articulated early in the program. This motivates a need for systems that are robust in the face of unanticipated situations so that customer missions are not unreasonably constrained or impacted by "shortcomings" in SOMO services.

One of SOMO's primary goals is to realize a paradigm in which SOMO acts as a service provider to organizations that fly space missions for NASA, other government agencies, and even the commercial sector. These organizations purchase SOMO services "by the pound" as customers. We have to provide systems that are not experiments themselves, but rather stable bases from which to do bold experiments. To this end, SOMO also seeks to work closely with industry to see that robust autonomy technology gets infused into products and services for the space industry and beyond.

The potential for application of these technologies spans space-based communications networks (e.g. TDRSS) and ground-based assets including communication and tracking antenna systems, data networks, and control centers. There are several problems that are candidates for the application of autonomy, if it can be made reliable enough, including: antenna control, antenna scheduling, communication link scheduling and operation, navigation, attitude determination, fault detection, isolation, and reconfiguration (for spacecraft or ground assets), and mission-level planning and scheduling.

Some attempts have been made to apply autonomy and automation in these areas in the past with varying degrees of success. We will present relevant case histories and the lessons inferred from them. Combining this past experience with anticipated future needs, we can clarify the challenges that must be met in order to realize the benefits of autonomy.

Overview of Space Operations

The Space Operations program was initiated in 1995 as an agency-wide effort to provide an integrated, cost-effective approach to the delivery of routinely required communications and data-related services to NASA missions and other organizations that use NASA assets. Space Operations services are broadly divided into two classes: Mission Services and Data Services. Data Services include the fundamental telecommunications and data networking required to deliver commands to spacecraft, deliver spacecraft telemetry to control centers, and deliver payload data to customers. Mission Services include data processing and/or storage as well as mission operations and associated planning, analysis, design and development activities primarily for near-earth robotic missions.

Space Operations Assets

Space Operations operates NASA's antenna networks, control centers, and data networks. The antenna networks include the Deep Space Network (DSN), Space Network (SN) including the Tracking and Data Relay Satellite System (TDRSS), and a worldwide network of ground stations known collectively as the Ground Network (GN). An extensive system of data and telecommunication networks known collectively as the NASA Information Systems Network (NISN) connects the antenna networks, NASA centers, and research facilities to each other and to the global Internet. Finally, Space Operations provides control center facilities at Goddard Space Flight Center (GSFC), the Jet Propulsion Laboratory (JPL), Johnson Space Center (JSC), Kennedy Space Center, and Marshall Space Flight Center.

Operations, Development, Upgrades & Technology

Space Operations has three functional program components. The main function is Operations, which encompasses day-to-day operations, planning, and customer support functions. The Development program, which currently consists of the TDRS Replenishment Project, provides replacement for primary communications infrastructure systems, which are viewed as Agency capital investments. The Upgrades program includes modifications to existing assets including replacement of obsolete assets, enhancement of existing assets, or addition of new assets to provide new capacity or capability. Finally, the Technology program funds research and development into new technologies that have the potential to significantly reduce cost or meet anticipated future needs.

Potential applications of robust autonomy

There are two fundamental motivations for technology insertion in Space Operations: (1) enable new and needed capability and (2) reduce operations costs. This leaves a wide range of potential applications for sufficiently robust autonomy technology.

While the science community looks to onboard autonomy to enable missions that otherwise could not be flown – especially due to large distances, Space Operations can also benefit from onboard autonomy if it reduces the use of communications resources. If a spacecraft can operate with fewer ground station contacts or reduced data rates due to the use of onboard autonomy, Space Operations could see an associated reduction in operating costs. Such savings may be less than we would hope, however, because the "fixed costs" of supporting space missions must still include support for worst-case (i.e. emergency) communications and the various types of services required by different missions. These can swamp the costs (and hence marginalize the savings) of normal operations. Even in the case where spacecraft are fully autonomous, i.e. they can handle their own faults to a great extent, there will likely be a requirement to maintain a backup capability to do intensive communications and operations for troubleshooting and fixing a satellite when "everything goes wrong".

In ground-based systems there are many opportunities to reduce costs. Various efforts are already underway to streamline and automate historically manual processes. There is a great deal of effort to enable "lights out" operations and some success has been achieved in this. In the not-too-distant future, there will be a need for a planning and scheduling system that can optimize the delivery of a data services to a heterogeneous set of spacecraft via a heterogeneous mix of ground antennas and data networks. Moving further out, space operations services will have to be provided to more data-intensive missions and new kinds of missions involving things such as constellations and "Internet enabled" spacecraft. Since the budget trend is expected to continue to be downward, there is a strong incentive to find ways to do more for less.

Case Studies & Lessons Learned

Beacon Tones on DS-1

One of the most impressive steps forward in demonstrating autonomy and automation in space has been NASA/JPL's Deep Space-1 (DS-1) program. The DS-1 spacecraft carried several new technologies, including the Remote Agent Experiment (RAX), Autonomous Navigation (AutoNav) and the Beacon Monitor Experiment (BMOX).

RAX demonstrated onboard planning and scheduling and fault detection, isolation, and reconfiguration (FDIR). This was a milestone in the development of spacecraft that can operate at distances from the earth at which the round-trip light time is too great to support reconfiguration from the ground while still meeting mission objectives.

The BMOX successfully demonstrated the use of telemetry summarization and sub-carrier "side tones" to communicate simple status information without establishing a full telemetry link and dumping data. The result is that a system using this technology can reduce the use of communications assets and the labor involved in establishing communication and processing downlink telemetry just to find out that everything is normal.

In November of 1999 the spacecraft's star tracker failed. In order to continue the mission the team has developed a method to use an onboard camera to point at a guide star for steering while the ion engine is thrusting. The BMOX has been reprogrammed to provide status tones during periods of thrusting indicating whether or not the camera is still locked on to the guide star. The main (high gain) antenna cannot be pointed toward Earth during thrusting and regular telemetry cannot be detected if sent using the low gain antenna. The beacon tones, however, can be detected when transmitted using the low gain antenna. Thus the mission continues with the beacon tones contributing by reducing the risk of thrusting without lock on a guide star.

Involve the Operations Team

Historically, near-earth space missions have been supported by ground controllers in a "24x7" paradigm. Missions flown from NASA's Goddard Space Flight Center (GSFC) have reduced staffing from 24x7 to 1 or 2 (8 hr.) shifts per day, 5 days per week. These reductions have occurred through varying approaches to automation / autonomous operations. Most involve technology on the order of scripting languages and rule-based expert systems.

In (Cooter, et al. 2001), the case is made that successful automation efforts are correlated with involvement of the operations team in Integration & Test (I&T) and development (or re-definition in some cases) of the operations concept.

In the case of the MAP/IMAGE pair of missions, the operations personnel were involved in definition of the operations concept and in I&T. In fact, some of the operations personnel were certified as Test Conductors.

The GRO mission did not incorporate significant automation until after seven and a half years of operations. Automation became necessary due to funding constraints. A small team of developers built a suite of capabilities that automated nearly all operations aspects except for mission planning and transmission of stored command loads. The developers were co-located with the operations personnel. The operations personnel were directly involved in building a rule base for a CLIPS-based expert system.

Finally, in the case of Landsat 7, the operations personnel themselves automated monitoring functions using System Test and Operations Language (STOL) scripts. In combination with a mission that can operate overnight without receiving commands or needing attention to data recorders, they have been able to reduce operations to one shift per day.

By contrast, operations personnel were not involved in I&T or definition of the operations concept for the FUSE program. In that program, I&T proceeded with two different scripting languages in use: STOL was used for the spacecraft and Spacecraft Control Language (SCL) for the payload. Although the STOL scripts were translated to SCL before launch, operations were not significantly automated. The team is looking at ways to automate operations during the Extended Mission starting in 2002.

Simpler Missions are Easier to Automate

It is perhaps a truism that simpler things are easier to automate than more complex things. Manned missions and major science missions such as the Hubble Space Telescope (HST) or Terra - part of the Earth Observing System (EOS) - are staffed 24x7 by sizeable teams. These missions are much more sensitive to delays in response to an anomaly or losses in science data. More money and, in the case of manned missions, human life are potentially at risk. The complexity of the systems typically does not admit relatively simple technologies such as scripting and rule-based expert systems in anything but limited roles as aids to human operators. In order to adequately address the system, such approaches become intractable or prohibitively expensive for comprehensive automation.

These incentives lead to very risk-averse operations concepts based on past success and including reliance on the presence of highly trained operators at all times. This seems to present a terrible catch-22 situation to the proponent of advanced autonomy. Simple systems don't need advanced approaches, and so can't be used to prove the worth of them. Complex systems can't afford the risk of using an unproven technology.

SMEX Automation

NASA's SMEX (Small Explorer) missions are being supported from one control center, eight hours a day, five days a week. The functions that have been automated include configuration of the spacecraft and ground system for passes, dumping of the solid-state recorder, and monitoring of real-time health and safety data. Non-routine commanding and mission planning remain manual processes. Two of the SMEX missions experienced nearly three months of fully automatic operation (Maks, Breed and Rackley, 2001) (Cooter et al., 2001).

The SMEX control center also incorporates a notification system known as SERS (Spacecraft Emergency Response System) to page ground controllers in case of anomalies occurring during off-shift hours. The system contacts a controller via pager or e-mail with relevant messages extracted from the automatic monitoring functions.

"Keep an Eye on the Ball"

The biggest challenge in applying automation to the operations of the Deep Space Network (DSN) is the application of the old-fashioned principle stated in the title. The "ball" refers to clearly stated automation goals. In theory, the statement that "An automation effort must be driven by clearly stated automation goals" is just common sense. In practice, automation has often started with the goals not being very clear, or not agreed by all the stakeholders.

What could the automation goals be for the DSN? Here are a few candidates: reduce the operations cost by x%, or improve the quality of service by x%, or enable functions and capabilities that cannot be provided without automation. Since all these goals are attractive, at least to a significant number of the stakeholders, a tempting mistake is to state all three (and more) as goals, accept the derived requirements, and make the leap-of-faith that automation is the solution. In the DSN, the combination of operational realities in the area of cost distribution, unstable infrastructure and over-reliance on COTS have caused painful disappointments. We discuss the key factors in the next two sections.

On the other hand, when the goals were stated clearly, and automation applied in appropriate measures, automation of space operations proved central to enabling NASA to fulfill the agency missions. We discuss such an example, the saving of the Galileo Spacecraft mission, in "The Right Automation" below.

To illustrate the difficulty in "keeping the eye on the ball", let us assume that the goal is to reduce cost. Since most of the cost in DSN operations is workforce, it is tempting to zero on the real-time operations staff, those operators who staff computers terminals, entering commands and monitoring data, 24x7, a tempting target for cost savings. Imagine that we placed that automated system into operations. Could we significantly reduce or eliminate that workforce? First, the visible real-time operations staff represents only 20% of the total operations staff - so focusing all the automation effort of real-time operations may not bring about the desired cost reduction goal. Also, the unreliable subsystems discussed in the following section require operators to be present, so the prospects for reduction of workforce, without addressing the unreliable infrastructure are unrealistic. Finally, the realities of managing workforce must be recognized. If, for example, we require one operator and through automation are able to reduce his workload from 40 hours per week to one hour per week, we achieved no cost savings, unless we can find use for the other 39 hours of his time.

Thus it is crucial that automation in the DSN is applied as part of a thorough system engineering solution, responding to well-defined goals. In general, given that some real-time operations workforce must be present to address safety, unreliable equipment, and spacecraft critical events, automation has been a powerful ally in reducing repetitive data entry and routine monitoring. It was not yet successful in providing extensive fault recovery, partly for the reasons described below.

"Don't Build a Penthouse Without a Foundation"

When the term "automation" is mentioned in the context of Space Operations, one would envision a room full of whirring computers, lights flashing on-and-off, with no human in sight. The reality in the DSN is quite different. In the 1995 time frame, when the DSN approached a large automation project, the underlying subsystems were not ready for full automation. Examples ranged from (a) older subsystems that were not even computer-controlled, to (b) subsystems that were unreliable enough to require a significant number of manual resets, to (c) subsystems that provided minimal feedback to the central control system, insufficient for effective automation. Even the most effective automation engines or tools could not communicate with the first group and their effectiveness would be greatly diminished for the second and third groups.

The DSN has recognized the problem and has embarked on a broad effort to upgrade subsystems to a state suitable for automation. This effort is ongoing. With reliable subsystems, the DSN is starting to apply straightforward automation of routine processes, executing routine operations with minimal need for operator intervention.

"COTS Isn't"

In space operations, as well as across NASA and in other industries, there is a requirement to increase the reliance on Commercial-Off-The-Shelf (COTS) products. In two of our recent automation projects, we encountered the reality that COTS products must be applied with an eye to the potential pitfalls.

In the first case, the DSN Network Control Project (NCP) selected a COTS product as the underlying infrastructure connecting sites in the US, Spain, and Australia. While the product looked very promising in the beginning, the expectation that it would become the mature industry standard, with many supporting vendors and wide-area capabilities did not fully materialize. Even though NCP did not use it *per se* for automation, the efforts associated with attempting to make it work and its eventual removal have hampered the overall successful completion of NCP.

In another example, the DSN attempted to use an adaptation of a COTS product to automate the operations of the 26m antennas. While the automation function itself operated spectacularly well, the areas where adaptation to unique DSN and customer needs (e.g. unique old interfaces) required significant work on the COTS-based system, tainting and hampering the overall successful completion.

While in these two case studies, COTS was not used as the automation process, the same word of caution applies to COTS automation products. These must be selected carefully to meet the reliability, long-life, and resiliency of the overall system that is being automated.

"The Right Automation"

An example where automation was applied correctly to space operations is the support to the Galileo mission. That mission required a full redesign after the main spacecraft antenna failed to deploy in April 1991. JPL has redesigned the mission to rely on a low-gain antenna, adding significant complexity to the ground system. The goal for automation was thus crisply defined: enable operations with one real-time operator per DSN site. With such a crisp goal, the automation portion of the task could be scoped effectively:

- Most of the automation was via simple UNIX scripting
- There was a heavy emphasis of using reliable subsystems (the task had the rare luxury of being able to bypass many of the old, unreliable subsystems)
- While extensive fault detection was included, no automated fault recovery was deployed. Instead the operations relied on the real-time operations staff to take the corrective action. Because operational experience was being developed, operations staff worked with the developers to refine fault recovery procedures.
- COTS products were used, in carefully identified locations. While COTS hardware (e.g. SUN workstations) and utility software was used extensively, COTS application software was used sparingly, only where risk-benefit was well understood.

The Galileo mission, supported by this equipment, has been successfully unraveling the secrets of the Jovian system since 1995, including the inspiring discovery of the potential ocean inside Europa.

Implications for Autonomy Research and Development

What can the autonomy community draw from these experiences that will benefit both the researcher/developer and "end users" in the future? How can we help those who wish to infuse new approaches into space operations? For starters, there is what we can't do – we can't say that "If you do *x* and don't do *y* you will succeed." We do, however, hope to infer something from the successes and failures of past projects

Start Small

In space operations, the opportunity for a new approach to be an enabling technology is quite rare. Although NASA's New Millennium Program and the former Advanced Communications Technology Satellite (ACTS) mission provide wonderful opportunities of this kind, the vast majority of space missions are either about returning science data or delivering commercial services. In such cases, there are nearly always one or a few existing, tried and true, approaches for any element that is a potential application of autonomy. For a new approach to be selected, it must show a combination of acceptably low risk and cost reduction relative to the *status quo*.

In terms of acceptably low risk; there are levels of maturity that correlate to reduced risk. Approaches that are "flight qualified" i.e. have been successfully used in space, preferably by the same organization that is now looking for a solution, are perceived to carry the least risk. Approaches that have a successful track record in smaller or less critical aspects of a system carry more risk, but relatively more confidence than approaches that have not been so applied. There is relatively less confidence in approaches that have worked well outside of the space operations domain. This is both a cultural and a technical barrier, and may or may not be seen in a given situation depending on the persons and risks involved. Finally, there is technology with no track record. Fundamental ideas, newly developed technologies, etc. These are purely in the researcher/developer's domain.

NASA has recognized this progression in the Technology Readiness Level (TRL) scale. Technologies at TRL 1 are basic principles. At the highest level, TRL 9, they are "flight-proven". Programs within NASA have the explicit goal of advancing technologies along this scale, including the Cross-Enterprise Technology Development Program, the aforementioned New Millennium Program, and others. Within Space Operations, the Communications Technology Program (CTP) seeks promising technologies in the range of TRL 3 (proof-of-concept) through TRL 6 (prototype demonstration in a relevant environment) to develop. The CTP seeks technologies that either reduce costs or enable new capabilities needed by Space Operations.

We have observed that simpler missions have shown the greatest level of automation, partly due to the cost of failure or data loss in comparatively complex systems and the risk-averse approaches taken to avoid such consequences. A new, lower TRL, technology typically will only be applied where cost, risk, and benefit balance out. In complex systems this tends to occur in relatively limited, non-critical roles that can be crisply defined. While these applications are not particularly glamorous or lucrative, they result in a track record and movement along the TRL scale. A good track record in operational situations combined with cost benefits will improve the likelihood of progressively increasing the criticality and/or scope of the application.

The best way to advance a particular technology may not be directly in line with its ultimate application. It may be in providing a similar, smaller, and/or less critical function.

In addition, it is not uncommon to think of the spacecraft as the interesting system. The notion of a "thinking spacecraft" seems somehow crisp and appealing. Another view is to look at intelligent space systems – which include both a space segment (one or more spacecraft) and a ground segment. In most cases, an idea conceived of for onboard use can actually be used on the ground first – either on an evolutionary path toward onboard application or to provide a different function. In terms of risk, this has an advantage. Things on the ground can typically be turned off or disconnected if they are observed to be malfunctioning. This may not always be possible for onboard elements. Ground-based elements can also be repaired and replaced more easily, to allow for experience gained or updated technology.

Watch for COTS Pitfalls

There are (at least) two sides to COTS in this context – those who develop and sell it and those who buy and operate it.

For those who would use COTS products in the development of an autonomy technology, there is the risk that poorly chosen COTS can cause failure – whether the autonomy in question is good, bad or indifferent. This can be a problem in that often the only vendor in an autonomy-related domain may be small, new, or a "single-product" company with perhaps only one or a handful of other customers. From a business standpoint – especially for space missions that must fly for 5, 10 or 15 years, this is a risky bet. Risk-averse programs will not accept a dependency on such a COTS product with out strong need and risk mitigation alternatives. The alternatives can include contractual mechanisms such as "code-in-escrow" or, if time and funding permit, carrying multiple vendors through some decision point in the program. If the COTS dependency is in a non-critical element, there will be more tolerance for the risk as well.

Much more preferred is the case where widely adopted and available COTS products, such as workstations, databases, operating systems, etc. are used to provide significant portions of the system's functionality. These are typically well supported by stable companies. They may be superior to what could be custom-built, and don't have the maintenance burden that would come with a custom-built product. They are more likely to be associated with widely supported standards that will last for a significant period of time.

On the other hand, there is the autonomy researcher/developer who is bringing a COTS product or service to the marketplace. Often, this will be exactly the sort of company mentioned above - a new company with little more than Intellectual Property (IP) and perhaps some sort of startup funding. Regardless, there are probably intellectual property rights to be protected and the fundamental need to be on the road to profitability. Here the main insights to offer are these: be aware of the fact that many things have been over-hyped and that failures have left a bitter taste in the mouths of experienced potential customers. Good products are sought after. Seeking to sell to any potential customer may lead to bad situations, particularly if the customer doesn't have clearly defined autonomy/automation goals, etc. A track record must be established and grown before the product is likely to become a critical component in a successful complex system.

Think Globally, Act Locally

We have made an argument that the autonomy researcher/developer must start small, in isolated elements of a larger system. Now we claim that the same researcher/developer must think "at the system level".

Starting with the decision to pursue an opportunity, awareness of the larger context in which your work is being applied is invaluable. Mismatches during integration of elements of a system are common, if not inevitable, in systems built by human beings. The worst problems arise when specialists from different areas make assumptions about other elements of the system and those assumptions are not tested until late in the program. Theoretically, the role of system engineering is to keep this from happening and to a great extent it does – if a system engineer is present and communicating effectively with all of the parties involved. This can break down in any number of ways. We have seen cases of mismatched assumptions about things as basic as units of measurement. Other examples include the choices of two different scripting languages in the FUSE case. These disconnects nearly always have some cost or schedule impact, but in the case of a technology insertion such as a new autonomy approach, the solution to such a cost or schedule problem may be to forego the new technology. You can improve the chances of success for both your new approach and the system as a whole by understanding the larger context. , Asking questions about interfaces and interactions. Proactively adjust to match up with the larger system when it makes sense; raise a flag when it doesn't, so that accommodations can be made early – when they cost less.

Reduction by Enhancement

The notion of reducing cost by replacing people with automation can be somewhat misinterpreted. In the cases that we've cited above, the reduction in operating personnel came by increasing the capability of the individual controller. This can be somewhat subtle, but it is important. An attempt to literally replace a human being with intelligent hardware and/or software may involve giving away some of the human capability to respond to unforeseen circumstances. In the cases we've seen, success came when automation took care of routine and repetitive tasks, thereby allowing the human operator to easily monitor normal operations of one or more spacecraft and focus attention on planning, trends, and anomalous situations. The automation complements the human operator rather than eliminating them.

"But What About Validation?"

One of the questions that seems to inevitably come up in the context of autonomous systems is testing and validation. "How do we know that this thing won't go crazy on us and wreck the spacecraft?" This did not come up in the cases we've cited, with the probable exception of DS-1, because the efforts have not gone so far as to autonomously reconfigure a spacecraft or other system in response to a diagnosed fault. Nor have they incorporated reasoning capabilities that might be perceived to be capable of unexpected behavior.

Looking Into the Crystal Ball

Ground Network Evolution

NASA's Ground Network is evolving to include both NASA and commercially owned and operated assets. Space Operations will continue to support NASA and non-NASA space missions using this network for the foreseeable future. The cost of these operations might be reduced by a sufficiently powerful planning and scheduling approach that could handle the heterogeneous mix of spacecraft transponders and orbits, ground antennas and data networks, and control centers.

Such a system would be able to provide missions with ground contacts via the most cost-effective selection of ground station (antenna) and data delivery mechanism to the control center and/or Principal Investigator – for all requested ground contacts in a given time span. It would be capable of quickly re-planning in the case of spacecraft emergencies (especially for spacecraft using the beacon tones demonstrated on DS-1) or ground equipment failures with minimal impact to end users.

Budgets

If current trends continue, the budget for Space Operations will continue to decrease. Some of the automation work cited above has been commensurate with that expectation and provided real cost savings to meet the budget. In order to continue that trend, cost savings may have to be found in the more complex missions and systems. This will most likely require more sophisticated autonomy involving reasoning and knowledge capture and representation approaches not yet deployed.

There may be benefit in the application of performance modeling of entire systems (including human operators and organizations) with the objective of discerning what, if any, autonomy/automation approach to take and how to implement it for the greatest benefit. This and other means of getting to clearly defined goals should help reduce the number of failed attempts at cost savings.

New Kinds of Missions, Increased Data Return

Plans and concepts for future science missions involve new approaches to space missions including highly autonomous spacecraft and multiple spacecraft flying in formation as a 'virtual platform' for scientific observations. The implications of this for Space Operations are yet to be determined in detail, but there is a definite trend toward increased volume of science data returned.

Which Way Forward?

The benefits to Space Operations so far have accrued from relatively low-level automation technologies such as scripting and to some extent fault detection. The missions mentioned have, by and large, not automated functions such as planning and scheduling, fault isolation and recovery or trending. These particular functions are often thought of in the context of onboard software. Certainly they provide great potential for improved spacecraft robustness and, in the case of several deep-space missions, enabling technology.

We have shown the conditions under which ground systems will embrace and benefit from autonomy that helps to reduce cost or provides needed new capability. The next step is to apply that experience and knowledge to the remaining functions.

Current work on some topics not mentioned herein is reaching maturity. Among these are autonomous navigation, Ka-band telecommunications, the use of Internet Protocol and/or compatible approaches in space communications, and more sophisticated approaches to onboard planning and scheduling and FDIR.

Finally, we offer some suggestions as to the challenges to be met in order to realize the successful infusion of autonomy into Space Operations:

1) Gain infusion into critical elements of systems by building a track record of safe, reliable performance and possibly cost savings in less critical, less risk-driven elements.
2) Involve the end user of autonomy technology in the development of operations concepts as well as integration and testing in order to gain their experience and acceptance.
3) Apply autonomy to increase the capability or span of control of the individual such that a constant or decreasing staff of can control more or more complex systems.

Acknowledgements

The authors wish to acknowledge the contributions of Julia Breed and E. Jay Wyatt of NASA, and Charles Copeland of CSOC. In addition, several persons reviewed and commented on this text including Patrick Duffin, Jon Michael Smith, and Thomas A. Costello of the NASA Space Operations Management Office.

Part of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

Sherwood, R., Schlutsmeyer, A., Miles, S., and Wyatt, E. J., 2000. Lessons from Implementation of Beacon Spacecraft Operations on Deep Space One. In Proceedings of the IEEE Aerospace Conference, Big Sky, MT

Cooter, M., Davis, G., Mackey, J., and Rackley, M., 2001. Alternative Approaches to Mission Control Automation at NASA's Goddard Space Flight Center. To be published in *Proceedings of The 14th International FLAIRS Conference*, Key West, FL: AAAI Press

Maks, L., Breed, J., Rackley, M., 2001, Current Level of Mission Control Automation at NASA/Goddard Space Flight Center. To be published in *Proceedings of The 14th International FLAIRS Conference*, Key West, FL: AAAI Press