Mission Operations of Earth Observing-1 with Onboard Autonomy

Gregg Rabideau, Daniel Tran, Steve Chien, Benjamin Cichy, Rob Sherwood
Jet Propulsion Laboratory, California Institute of Technology

Dan Mandl, Stuart Frye[1], Seth Shulman[2], Joseph Szwaczkowski[2]
Goddard Space Flight Center

Darrell Boyer, Jim Van Gaasbeck
Interface & Control Systems

Abstract

Space mission operations are extremely labor and knowledge-intensive and are driven by the ground and flight systems. Inclusion of an autonomy capability can have dramatic effects on mission operations. We describe the prior, labor and knowledge intensive mission operations flow for the Earth Observing-1 (EO-1) spacecraft as well as the new autonomous operations as part of the Autonomous Sciencecraft Experiment (ASE).

1. Introduction

EO-1 is the first satellite of the Earth observing series of NASA's New Millennium Program (NMP). Since its launch on November 21, 2000, under management of the Goddard Space Flight Center (GSFC), EO-1 has demonstrated many new technologies for space-based Earth observing satellites. These technologies include several advanced instruments for collecting multispectral and hyperspectral images, and a high-rate, high-volume solid state recorder (SSR) for storing the images.

As a part of the NMP Space Technology-6 (ST6) project, the Autonomous Sciencecraft Experiment (ASE) [10] was selected to demonstrate advanced concepts in flight software. ASE is a set of integrated technologies for autonomously collecting, processing and downlinking science data. These components include: onboard science analysis, onboard mission planning, and onboard robust execution. First, several different science algorithms are used to analyze science data onboard. These algorithms process the image to detect the presence of unique features in the image, such as clouds, flooding, ice formation, or volcanic activity. The detection of these features is used as a trigger to automatically submit requests for additional data. Next, these requests are serviced by the onboard planner, CASPER. CASPER processes the requests into a more detailed sequence of future spacecraft activities that is consistent with the extensive list of spacecraft and mission constraints. Finally, when an activity is imminent, CASPER submits a request for execution by the onboard executive, SCL. SCL initiates a set of scripts that perform the complete sequence of commands for the spacecraft and its payloads. Prior to executing each command, constraints are checked again to confirm the validity of the command as well as to ensure the safety of the spacecraft. After the command is sent, the executive checks for a successful initiation and completion of the command. When a full sequence for a data collection is complete, one or more of the science processing algorithms are triggered and the entire process repeats.

The ASE software provides many benefits to the science and engineering teams involved in operations. First, the onboard science algorithms can prioritize the data before it is sent to Earth. These priorities can then be used to increase science return by maximizing a spacecrafts limited resources. For example, if clouds are detected in an image, this image can be removed from the onboard recorder and from the downlink queue, freeing up these resources for higher quality data. Ultimately, this also results in fewer "junk" images that must be examined by the science team. The automated planning system, CASPER, also has several benefits. Scientists can use ASPEN, the ground version of CASPER, as a ground planner to select a consistent set of high priority observations from the full set of inconsistent target requests. When done manually, this task can be tedious and require extensive knowledge of spacecraft constraints. The same system can also be used for onboard planning. This not only reduces the inputs required by the flight system (only the goals need to be uplinked, rather than the full command sequences) but also enables faster re-planning responses to either engineering anomalies or changes in science priorities. Finally, the SCL onboard executive provides several benefits over traditional sequencing. By monitoring commands and spacecraft state, the executive can make quick changes to the more immediate parts of the

---

[1] MitreTek
[2] Honeywell

sequence. This can help ensure the success of the sequence, and possibly optimize its execution.

In the fall of 2002, the ASE team at JPL joined efforts with the EO-1 team at GSFC to demonstrate this advanced flight software system. After a successful demonstration, the ASE software became a part of the baseline operations of EO-1. In the following sections, we describe the past steps used in operating EO-1 and also the new methods that takes advantage of the autonomy software.

2. The EO-1 Spacecraft

The EO-1 spacecraft [1] is in an approximately 90 minute, 705 Km sun synchronous (98° inclination) orbit around the Earth. It has two main instruments for collecting science data. The Advanced Land Imager (ALI) is an advanced spectrometer with 30-meter resolution for multispectral pixels and 10 meter resolution for panchromatic pixels. The ALI generates data at a rate of about 102 megabits per second (102 Mbps). The Hyperion is a hyperspectral imager (HSI) that is capable of resolving 220 spectral bands from 0.4 to 2.5 µm with a 30 meter spatial resolution. The HSI generates data a rate of about 233 Mbps. A typical scene collects about 16 seconds of data including calibrations, so together the two instruments generate about 5 gigabits of data for each scene.

The EO-1 attitude control system (ACS) is used to point the spacecraft at the various targets. Reaction wheels are used to rotate the spacecraft and stabilize all three axes to a 0.03° pointing accuracy. Torque rods are used to desaturate, or bias the wheel speed. EO-1 also has a hydrazine propulsion system for correcting errors in orbit insertion, orbit maintenance, formation flying, and eventually de-orbiting.

EO-1 carries two Mongoose V R3000 12 MHz radiation-hardened flight processors with 256 MB of RAM. The primary CPU is used for Command and Data Handling (C&DH). The WARP uses the second CPU for data recording, processing and playback. To reduce risk and avoid contention for cycles on the primary C&DH processor, the ASE software runs on the WARP processor.

3. ASE

The ASE onboard flight software includes several autonomy software components (See Figure 1):

- Onboard science algorithms that analyzes the image data to detect trigger conditions such as science events, "interesting" features, changes relative to previous observations, and cloud detection for onboard image masking [2]
- Robust execution management software using the Spacecraft Command Language (SCL) [3] package to enable event-driven processing and low-level autonomy
- The Continuous Activity Scheduling Planning Execution and Replanning (CASPER) [4] software that replan activities, including downlink, based on science observations in the previous orbit cycles

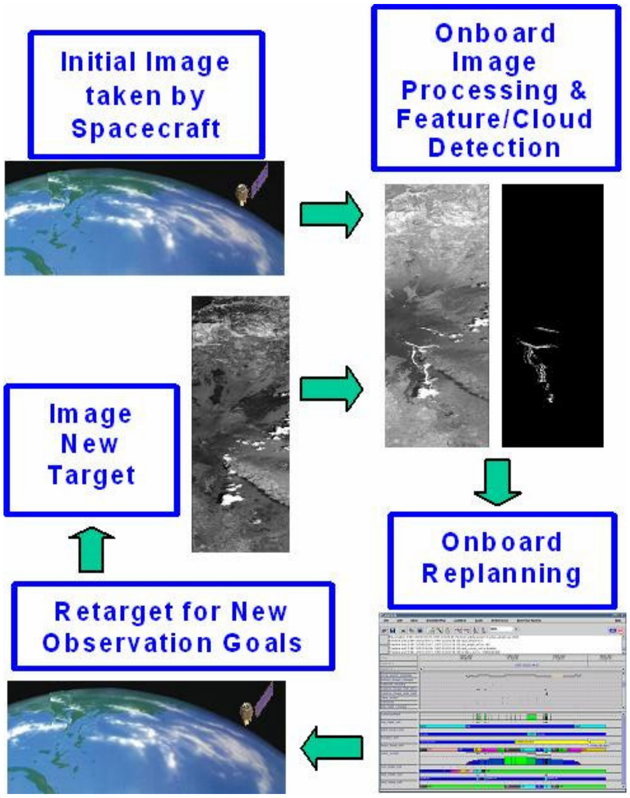

Figure 1: Onboard Science Scenario

The onboard science algorithms are used to automatically identify regions of interest including land, ice, snow, water, and thermally hot areas. Repeat imagery using these algorithms can detect regions of change (such as flooding and ice melt) as well as regions of activity (such as lava flows). Using these algorithms onboard enables retargeting and search, e.g., retargeting the instrument on a subsequent orbit cycle to identify and capture the full extent of a flood. On future interplanetary space missions, onboard science analysis will enable capture of short-lived science phenomena. These can be captured at the finest time-scales without overwhelming onboard memory or downlink capacities by varying the data collection rate on the fly. Examples include: eruption of volcanoes on Io, formation of jets on comets, and phase transitions in ring systems. Generation of derived science products (e.g., boundary descriptions, catalogs) and change-based triggering will also reduce data volumes to a manageable level for extended duration missions that study long-term phenomena such as atmospheric changes at Jupiter and flexing and cracking of the ice crust and resurfacing on Europa.

Responsible for long-term mission planning, the ASE planner (CASPER) accepts as inputs the science and engineering goals and ensures high-level goal-oriented behavior. These goals are provided by either ground operators or triggered by the onboard science algorithms. The model-based planning algorithms enables rapid response to a wide range of operations scenarios based on a deep model of spacecraft constraints, including faster recovery from spacecraft anomalies. CASPER uses repair-based techniques [5] that allow the planner to make rapid changes to the current plan to accommodate the continuously changing spacecraft state and science requests. During repair, CASPER collects a set of conflicts that represent violations of spacecraft constraints. Generic algorithms are used to select and analyze a conflict to produce a set of potential plan modifications that may resolve the conflict. Heuristics are used to select a potential modification, and the plan is updated and reevaluated for new conflicts. This process continues until no conflicts remain.

The robust execution system (SCL) is a software package that integrates procedural programming with a real-time, forward-chaining, rule-based system. SCL accepts the CASPER-derived activities as an input and executes the activities as scripts of low-level commands. Immediately before sending each command, the various constraints on the command are checked with the current spacecraft state. SCL uses a rule-based system to monitor spacecraft health and safety during the execution of the activities. When a safety risk is identified, SCL makes on-the-fly changes to the short-term plan to avoid potentially hazardous states or commands. SCL also has the flexibility and knowledge to perform event-driven commanding to enable local improvements in execution as well as local responses to anomalies.

A typical ASE scenario involves monitoring of active volcano regions such as Mt. Etna in Italy. (See Figure 1) The ASE concept is as follows:

1. Initially, ASE has a list of science targets to monitor that have been sent as high-level goals from the ground.
2. As part of normal operations, CASPER generates a plan to monitor the targets on this list by periodically imaging them with the Hyperion instrument. For volcanic studies, the infra-red and near infra-red bands are used.
3. During execution of this plan, the EO-1 spacecraft images Mt. Etna with the Hyperion instrument.
4. The onboard science algorithms analyze the image and detect a fresh lava flow, or active vent. If new activity is detected, a science goal is generated to continue monitoring the volcanic site. If no activity is observed, the image is not downlinked.
5. Assuming a new goal is generated, CASPER plans to acquire a further image of the ongoing volcanic activity.
6. The SCL software executes the CASPER generated plan to re-image the site.
7. This cycle is then repeated on subsequent observations.

4. Past Operations Flow

The EO-1 spacecraft is operated [6] out of the EO-1 Mission Operations Control Center (MOCC) at the Goddard Space Flight Center (GSFC). The Mission Operations Planning and Scheduling System (MOPSS) was previously used for long-term planning. Figure 2 shows the general operations flow.

A good approximation of the spacecraft's orbit can be predicted about a week in advance. Therefore, the schedule of activities was generated on a weekly basis. But because a 1-day orbit prediction is more accurate, the detailed commands were generated and uploaded on a daily basis.

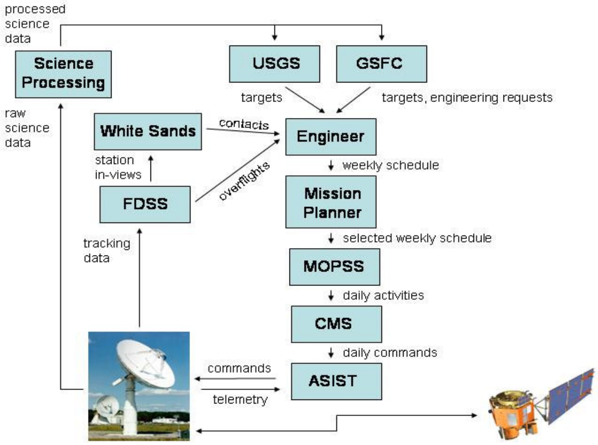

Figure 2: Past Operations Flow

4.1 Weekly Operations

The U. S. Geological Survey (USGS) previously managed the science requests for EO-1. These included standing and one-time requests from the EO-1 science team, the USGS, and from paying external customers. The first step in operations was to process the long term plan (LTP) of requests received by the USGS. This plan is a list of targets that will be visible for the upcoming week, including the orbits in which they will visible.

Currently, the groundstation contact support for EO-1 is managed out of the White Sands Complex (WSC). There are several stations in a ground network (GN) available to EO-1, including the primary sites in Alaska, Norway, and Virginia. The GN schedule, sent by the WSC, is a list of contact periods and is included into the weekly schedule.

Next, the Flight Dynamics Support System (FDSS) at GSFC uses groundstation tracking data to calculate the spacecraft ephemeris, predicting the spacecraft orbit through the upcoming weeks. This determines the approximate overflight times for science targets and ground station contacts.

In any given orbit (90 minutes long), many ground targets are visible, but only about one to two images can be acquired due to operations or spacecraft constraints. Therefore, conflicting scenes for a given week must be selected from the list of requests. Scene priorities are based on several factors including: who made the request, if it was paid for, and if it involves a fleeting science event. A USGS representative would manually prioritize and select the scene(s) with the highest priority in a given orbit. Also, the EO-1 science and engineering teams meet weekly with USGS to verify the selected requests and to make minor modifications to the plan for the following week.

After collecting several scenes, the WARP will reach capacity and commands must be scheduled to free up space for new requests. Before this can be done, an X-band contact must be scheduled to downlink the science data to Earth. These activities are selected at the same weekly science meeting when images are selected. About one X-band contact every other orbit is selected to keep the WARP from overfilling.

4.2 Daily Operations

After the weekly science meeting, the mission planner would use MOPSS to begin scheduling the 1-week set of activities. First, spacecraft maneuver commands are scheduled for each scene. Using the ephemeris, parameter values are calculated for the maneuver commands that point the instruments toward the target. After each scene is imaged, another maneuver command is added to the schedule to point the spacecraft at nadir. Next, because the maneuvers use reaction wheels, more commands are added to bias and de-saturate the wheels. When the wheels change directions, the spacecraft is less stable and may produce jitter during the observation. Therefore, prior to a group of scenes, the wheels are biased to a non-zero spin rate at the times when data will be collected for the scenes. After a group of scenes, the wheels are desaturated by biasing them to a zero spin rate. This provides the maximum flexibility for spinning the wheels in either direction for subsequent biasing.

While the original manual selection of scenes and contacts was done with the spacecraft requirements in mind, scheduling the details for these activities may still reveal conflicts. MOPSS was used to identify these conflicts and the mission planner resolved them manually. When all conflicts were resolved for the next day, the activities are sent from MOPSS to the Command Management System (CMS) where the command sequence is generated and prepared for uplink to the spacecraft. The commands for a given day are typically prepared the day before, then uplinked on the next available ground contact. This is performed at the latest reasonable time so that the most accurate ephemeris data can be used to generate the command parameters, and because the sequence is difficult to modify once it has been loaded onboard.

Re-planning for new science requests, while possible, is difficult in this scenario. After executing an original set of requests, the scientists must wait for the image products to be delivered. This includes waiting for the next science downlink, and often includes several days of waiting for the data (stored on tape) to be manually shipped from the ground stations. Once the data arrives, the scientists can run any number of manual or automated analyses on the images. The results of the analyses may suggest a change in priorities for the upcoming requests. For example, detecting a fleeting event such as a forest fire may increase the priority of a repeat scene of the same target. This request is then made at the next weekly meeting. However, if the meeting has already occurred, then the change may require manual rescheduling steps, and must be negotiated with the operations team. If the command sequence has already been uploaded, then the change is difficult and typically not worth the risk.

5. Current Operations Flow

The Autonomous Sciencecraft Experiment (ASE) team has developed advanced software for the EO-1 mission. Much of this software can be used both on ground workstations for mission operations and on the flight processor for autonomous operations. For example, we make use of the Automated Scheduling Planning Environment (ASPEN) [5], the ground version of the onboard CASPER planner. In this section, we discuss the modifications to the software on the weekly and daily operations of EO-1. Table 1 compares the current operations with modified steps that include ASE.

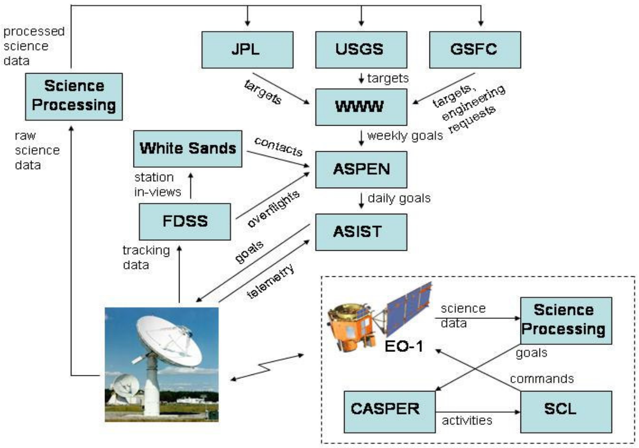

Figure 3: Current Operations Flow

5.1 Weekly Operations

For the weekly science planning, ASPEN lightens the workload of the scientists and engineers. Rather than *selecting* science targets, which requires knowledge of the spacecraft operations constraints, the scientists need only to *prioritize* the LTP for the upcoming week. ASPEN is then used to select the highest priority scenes while respecting spacecraft and operations constraints.

ASPEN is also schedules downlinks for the observations. The GN schedule identifies the set of X-band downlink opportunities and ASPEN uses its model of the WARP to predict when the memory will reach capacity. Using this model, it automatically adds X-band downlinks and file delete activities to free up space on the recorder. As with all activities, these are scheduled where allowed by the overflights and spacecraft constraints.

Finally, ASPEN interfaces with the Flight Dynamics Support System (FDSS) and generates the required maneuver and wheel bias commands. The FDSS software uses the spacecraft ephemeris to provide the required parameters for the commands. The ephemeris file is generated at GSFC using estimates of the spacecraft position and velocity vectors, and then pushed to JPL for processing. Using GPS data, these vectors are calculated onboard for the ACS, but the velocity vectors are not accurate enough to make long-term predictions. Therefore, weekly predictions are made on the ground using tracking data from the GN.

Web interfaces have also been developed as the main interface for processing scene requests and tracking the results of weekly operations. At the beginning of each weekly planning period, the observation LTPs from GSFC and USGS are submitted, which are automatically converted into input files for ASPEN. This is in contrast to observation LTPs being emailed to mission planners at GSFC for ingesting into MOPSS. JPL scientists are also able to request targets to be imaged without the need to generate LTP records for each.

Next, ASPEN selects the targets to be imaged from the total set of observation requests, while respecting the modeled operation and spacecraft constraints. After this step, any requested engineering activities such as instrument calibrations are inserted into the schedule such that they will minimize any lost of science data.

5.2 Daily Operations

With a planner operating onboard the spacecraft, we do not need to uplink the detailed command sequence, but only the high-level requests for scenes, downlinks, maneuvers, and wheel biasing. When CASPER receives these goals, it expands them into more detailed activities and schedules all activities at non-conflicting start times. Pending activities are continuously sent to SCL where the appropriate commands are executed and monitored.

This also means that we can uplink the entire week of goals rather than one day at a time. But, as the estimates for orbital parameters change, we need to send commands to the planner to change the relevant parts of the plan. This includes changes to scene start times and to parameters for maneuvers and wheel bias activities. When the planner receives these commands, it makes these and other changes necessary to maintain consistency in the plan. Using the ephemeris and the Flight Dynamics Support System (FDSS) on the ground, these commands must be uplinked, presumably on a daily basis. However, the onboard ephemeris is accurate enough for 1-day predictions and could be used to generate these daily plan updates. This would require additional work to port the FDSS (currently implemented in MATLAB) to the flight processor and operating system. Another alternative would be to skip the daily updates and use the less accurate (generated weekly) parameters for pointing, biasing, and image timing. This would result in slightly degraded science data, but possibly still within acceptable limits. Ultimately, this work would close the loop and allow us to fly autonomously for a full week. The final decision on the actual implementation will depend on available project resources.

Step	Past Ops	Current Ops
1	Process long term plan (LTP) requests	Process long term plan (LTP)
2	-	Process JPL requests
3	Process ground network (GN) schedule	Process ground network (GN) schedule
4	Process ephemeris and overflights	Process ephemeris and overflights
5	Manually prioritize science targets	Manually prioritize science targets
6	Manually select science targets	ASPEN selects science targets
7	Manually schedule downlinks	ASPEN schedules downlinks
8	Generate sequence and uplink	Uplink goals
9	Load time-tagged sequence into onboard queue	CASPER loads goals and generates plan
10	Execute sequence	SCL executes and monitors sequence
11	Manually reprioritize science targets	Science algorithms reprioritize science targets
12	Manually select replacement targets	Science algorithms select replacement targets
13	Manually reschedule	CASPER reschedules
14	Generate sequence and uplink	-

Re-planning scenarios become much easier with the addition of the ASE flight software. First, the science products are immediately available onboard after executing an observation request. The onboard science algorithms can start analyzing the data much earlier than if the analysis were done on the ground. The results of the analysis can then trigger new requests which are immediately sent to CASPER onboard. After receiving the new requests, CASPER will change the plan to accommodate the requests while maintaining consistency with spacecraft constraints. Onboard analysis and re-planning takes only minutes compared to ground-based operations which may take days.

To re-plan science activities onboard, we also need to re-plan the associated maneuver and wheel bias activities which were originally planned on the ground. The parameters for these activities, calculated by the FDSS, depend on the prior spacecraft orientation and wheel speed. However, by making a few simple assumptions, these activities can be scheduled onboard without requiring parameter recalculations. Specifically, if we assume that we always slew to nadir after a scene, then all maneuvers will begin at nadir and the parameters will remain constant regardless of the order of the scenes. If we also assume that the wheels are biased prior to each scene and de-saturated after each scene, again, the parameters remain constant. Therefore, CASPER can change the plan in flight using values pre-calculated by the FDSS. The disadvantage is that the plan will contain unnecessary activities and may not be optimal. However, we do not expect this impact to be significant.

5.3 Last minute requests

The implemented automation of the observation selection and mission planning portions of EO-1 operations also has benefits in operations flexibility. Because of the automation, it is possible to accept last minute requests (at latest before the last ground contact, 120 minutes before the overflight) and rapidly modify the schedule to accommodate an observation request. In this scenario, a user (e.g. from JPL, USGS, or GSFC) would input a last minute request at a web page. ASPEN then attempts to schedule the observation, pre-empting lower priority observations if necessary. If the request is accommodated, the new observation request is packaged for uplink by automation software at GSFC and the ground stations. Upon receipt onboard the spacecraft, the CASPER onboard planner replans, accommodating the new observation request. The uplink must be 120 minutes before the overflight to ensure that there is sufficient time for CASPER to replan as well as for the EO-1 spacecraft to complete preparatory activities such as slews prior to the observation.

This ground automation also permits autonomous observation requests to modify spacecraft operations. In this effort, the Earth Observing Sensorweb [9] links together a set of software agents which search for scientifically interesting events such as volcanic activity, flooding, wildfires, and cryosphere events. These events are detected by analyzing data provided from various data sources, which range from other NASA satellites such as TERRA and AQUA to volcanic in-situ sensors deployed by the Hawaiian Volcano Observatory. These software agents serve as proxies for scientists tracking recent science activities. After such a detection, an observation request is sent to the ASPEN planner. ASPEN integrates these requests into the weekly baseline schedule and evaluates scheduling that observation according to observation priorities and also mission and operations constraints. For observations that are feasible, that observation is uplinked to the EO-1 spacecraft identically to last minute manual requests. The Earth Observing Sensorweb has been operating autonomously since August 2003, and has enabled hundreds of observations of transient science events to be imaged by EO-1.

5.4 Impact on EO-1 mission

Ground and flight automation of EO-1 operations has had considerable impact on the EO-1 mission.

1. The automation has enabled much more flexibility in observation selection. In prior manual operations, nominally, science observations were planned 5 to 11 days in advance, with significant additional effort to change observations at a late date. In automated operations, changes in the observations are routinely made several hours before overflight.
2. The automation enables rapid response to spacecraft and ground station anomalies. In the spring of 2005 ground station anomalies removed about half of the scheduled ground contacts with little advance notice. Ground automation enabled resumption of relatively normal operations with new ground station support within a day, whereas prior similar anomalies caused several days of disruptions.
3. The automation has enabled dramatic cost reductions. As part of the EO-1 mission extension for October 2005 through September 2007 the EO-1 mission annual cost was reduced from \$3.6M/year to \$1.6M/year. Over \$1M/year of this reduction was attributed directly to the ASE automation.

6. Summary

The Autonomous Sciencecraft has enabled significant automation for the EO-1 mission. This automation includes both ground-based automation and onboard. The ground planner is used to reduce the workload involved in generating long-term plans as well as lower level daily activities. The onboard ASE software provides a high-level goal-oriented system capable of identifying and rapidly responding to several different types of science events. This advanced software suite is completely operational and represents a pathfinding approach to future space missions.

References

[1] Goddard Space Flight Center, EO-1 Mission page: http://eo1.gsfc.nasa.gov

[2] A.G. Davies, R. Greeley, K. Williams, V. Baker, J. Dohm, M. Burl, E. Mjolsness, R. Castano, T. Stough, J. Roden, S. Chien, R. Sherwood, "ASC Science Report," August 2001. (downloadable from ase.jpl.nasa.gov)

[3] Interface and Control Systems Home Page, www.interfacecontrol.com

[4] S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Using Iterative Repair to Improve Responsiveness of Planning and Scheduling," Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling, Breckenridge, CO, April 2000. (also casper.jpl.nasa.gov)

[5] G. Rabideau, R. Knight, S. Chien, A. Fukunaga, A. Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System," International Symposium on Artificial Intelligence Robotics and Automation in Space, Noordwijk, The Netherlands, June 1999.

[6] Dan Mandl, Joe Howard, "Lessons Learned on Operating and Preparing Operations for a Technology Mission from the Perspective of the Earth Observing-1 Mission," SpaceOps 2000, Toulouse, France, June 2000.

[7] J. Breed, S. Coyle, K. Blahut, C. Dent, R. Shendock, R. Rowe, "The MAP Autonomous Mission Control System," SpaceOps 2000, Toulouse, France, June 2000.

[8] M. Griffin, H. Burke, D. Mandl, & J. Miller, "Cloud Cover Detection Algorithm for the EO-1 Hyperion Imagery," Proceedings of the 17th SPIE AeroSense 2003, Orlando, FL, April 21-25, 2003.

[9] S. Chien, B. Cichy, A. Davies, D. Tran, G. Rabideau, R. Castano, R. Sherwood, D. Mandl, S. Frye, S. Shulman, J. Jones, S. Grosvenor, “An Autonomous Earth-Observing Sensorweb”, IEEE Intelligent Systems, May/June 2005.

[10] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, “Using Autonomy Flight Software to Improve Science Return on Earth Observing One”, Journal of Aerospace Computing, Information, and Communication . April 2005.

Acknowledgement

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.