Autonomous Sciencecraft Experiment
ST6 Technology Validation Report

Robert Sherwood, Manager
Autonomous Sciencecraft Experiment

Dr. Steve Chien, Technology Provider
Autonomous Sciencecraft Experiment

Approved:

Art Chmielewski, Manager
ST6 Project

Dr. Christopher M. Stevens, Manager
New Millennium Program

JPL Document Number D-34121

JPL Clearance Number CL 06-0165

Table of Contents

Executive Summary

The Autonomous Sciencecraft Experiment (ASE) has been operating onboard the Earth Orbiter 1 (EO-1) mission since January 2004. This software enables the spacecraft to autonomously detect and respond to science events occurring on the Earth. The package includes software systems that perform science data analysis, deliberative planning (CASPER planner), and run-time robust execution (SCL – Spacecraft Command Language). Through a detailed and extensive technology validation process, this software has demonstrated the potential for space missions to use onboard decision-making to detect, analyze, and respond to science events, and to downlink only the highest value science data. As a result, ground-based mission planning and analysis functions have been greatly simplified, thus reducing operations cost.

The technology was declared fully validated in May 2004, after all 20 onboard autonomy experiments were fully tested as described in Section 2. The overall system performed as expected and was considered a success. The validation consisted of the following onboard autonomy experiments performed 5 times each:

- Image planning and acquisition
- Downlink
- Data editing
- Image acquisition followed by image retargeting

Since the completion of the technology validation, over 3000 more autonomous data acquisitions have been completed (as of August 2005). The software is now running onboard EO-1 as the primary mission planning and control system operating 24/7. As such, the ASE is now at TRL level 9.

There were two important risks to our technology validation approach – one technical and one cultural. The technical risk was related to spacecraft safety. If the EO-1 satellite was lost due to the ASE software, that would have been a huge setback for onboard spacecraft autonomy.

This technical risk was mitigated by three factors.

1. We used rigorous software development processes, an extensive testing program, and phased deployment to ensure that the software would operate as expected.
2. The autonomy software implemented triple redundant safety checks via its 3-layered architecture.
3. The autonomy software was run on ther solid-state recorder CPU (WARP CPU) rather than the main spacecraft CPU.

The second risk was cultural. We needed to ensure that the technology validation of our software was convincing enough that scientists would use it on future missions. We had a multi-faceted approach to achieve this goal.

1. First and foremost, we involved (and funded) multiple science teams in the development of the experiment, software, and operations of the ASE software. If the scientists are involved from the start, they will help us develop a useful system and they will promote it to their peers.
2. We greatly exceeded the minimal set of validation experiments to show that this software is durable, maintainable, and can achieve increased science.
3. We started technology infusion early. This effort has so far paid off with infusion underway into the Mars Odyssey and Mars Exploration Rover missions.

The ASE flight tests were performed on the EO-1 satellite. At the conclusion of these tests, the ASE was advanced from TRL 6 to TRL 7. After further operations, in November 2004 ASE was adopted as the primary operations system for EO-1, elevating it to TRL 9. In the process of adopting ASE as its oeprations system, the EO-1 mission was able to significantly reduce its operations costs from \$3.6M in FY05 to \$1.6M in FY06. Of this reduction, approximately 50%, or \$1M/year was directly attributed to the use of the ASE software.

ASE represents a paradigm change in terms of enabling a spacecraft to respond autonomously to detected science events. This is in stark contrast to traditional labor-intensive ground-based operations. ASE enables benefits including:

1. Returning the most important science data
2. Fast reaction time to dynamic science events
3. Reduced operational costs
4. More responsive spacecraft to unknown environments

Autonomy software such as ASE enables new classes of missions at NASA in which the spacecraft conducts a highly interactive science investigation – dramatically increasing mission return. This technology is applicable to a wide range of missions including subsurface explorers, autonomous rovers, and coordinated systems of multiple spacecraft/sensors[1,2].

[1] The autonomous sciencecraft software is available for release for both government and external use. Please direct inquiries to the technology provider of the technology transfer office of the Jet Propulsion Laboratory.

[2] Significant performance data is available from the technology validation and ongoing flight of the ASE software. For further inquiries regarding this data please contact the technology provider or the New Millennium Program.

AUTONOMOUS SCIENCECRAFT EXPERIMENT ST6 TECHNOLOGY VALIDATION REPORT

I. INTRODUCTION

A. *What It Is/What is New About It*

From 2003-2005, the ASE running on the EO-1 spacecraft demonstrated several integrated autonomy technologies to enable autonomous science. Several science algorithms including: onboard event detection, feature detection, and change detection have been used to analyze science data. These algorithms are used to downlink science data only on change, and detect features of scientific interest such as volcanic eruptions, growth and retreat of ice caps, cloud detection, and flood tracking. These onboard science algorithms are inputs to onboard decision-making algorithms that modify the spacecraft observation plan to capture high value science events. This new observation plan is then executed by a robust goal and task oriented execution system, able to adjust the plan to succeed despite run-time anomalies and uncertainties. Together these technologies enable autonomous goal-directed exploration and data acquisition to maximize science return.

The ASE onboard flight software includes several autonomy software components:

- Onboard science algorithms that analyze the image data to detect trigger conditions such as science events, "interesting" features, changes relative to previous observations, and cloud detection for on-board image masking [19-21]
- Robust execution management software using the Spacecraft Command Language (SCL) [10] package to enable event-driven processing and low-level autonomy
- The Continuous Activity Scheduling Planning Execution and Replanning (CASPER) [5] software that replan activities, including downlink, based on science observations in the previous orbit cycles

The onboard science algorithms analyze the images to extract static features and detect changes relative to previous observations. These algorithms run using EO-1 Hyperion data 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, ice melt, and 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 on Europa.

The onboard planner (CASPER) generates mission operations plans from goals provided by the onboard science analysis module. 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. The onboard planner accepts as inputs the science and engineering goals and ensure high-level goal-oriented behavior.

The robust execution system (SCL) accepts the CASPER-derived plan as an input and expands the plan into low-level commands. SCL monitors the execution of the plan and 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 demonstration scenario involves monitoring of active volcano regions such as Mt. Etna in Italy. (See Figure 1.) Hyperion data have been used in ground-based analysis to study this phenomenon. The ASE concept have been applied 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 IR and near IR 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. Based on this detection the image is downlinked. Had no new lava flow been detected, the science software would

generate a goal for the planner to acquire the next highest priority target in the list of targets. (See Figure 1.) The addition of this goal to the current goal set triggers CASPER to modify the current operations plan to include numerous new activities in order to enable the new science observation.

5. The SCL software executes the CASPER generated plans in conjunction with several autonomy elements.
6. This cycle is then repeated on subsequent observations.

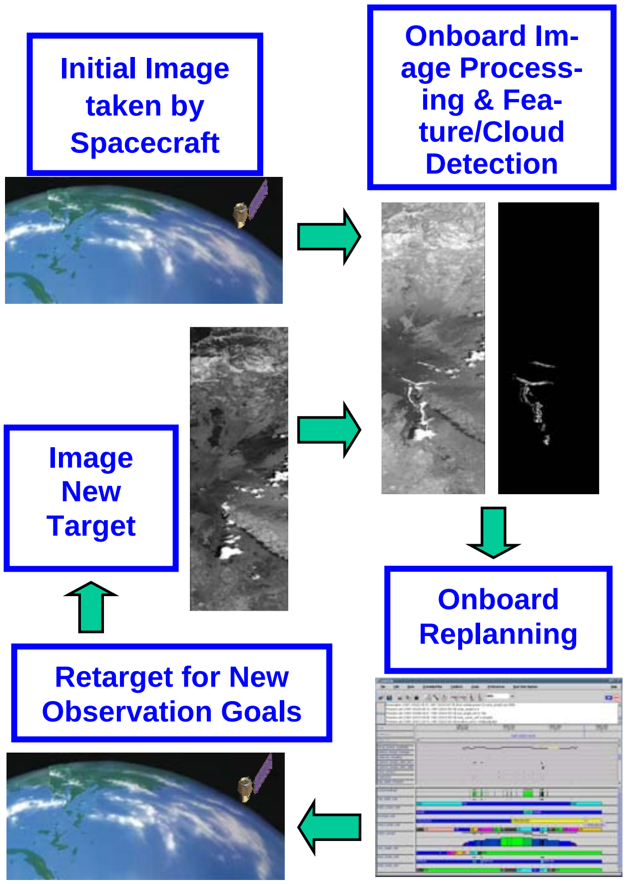

Figure 1. ***Autonomous Science Mission Concept***

B. Principles of Operation

The autonomy software on EO-1 is organized into a traditional three-layer architecture (See Figure 2). At the highest level of abstraction, the Continuous Activity Scheduling Planning Execution and Replanning (CASPER) software is responsible for mission planning functions. CASPER schedules science activities while respecting spacecraft operations and resource constraints. The duration of the planning process is on the order of tens of minutes. CASPER scheduled activities are inputs to the Spacecraft Command Language (SCL) system, which generates the detailed sequence commands corresponding to CASPER scheduled activities. SCL operates on the several second timescale. Below SCL the EO-1 flight software is responsible for lower level control of the spacecraft and also operates a full layer of independent fault protection. The interface from SCL to the EO-1 flight software is at the same level as ground generated command sequences. The science analysis software is scheduled by CASPER and executed by SCL in batch mode. The results from the science analysis software result in new observation requests presented to the CASPER system for integration in the mission plan.

The following sections give short descriptions of each of the software technology components.

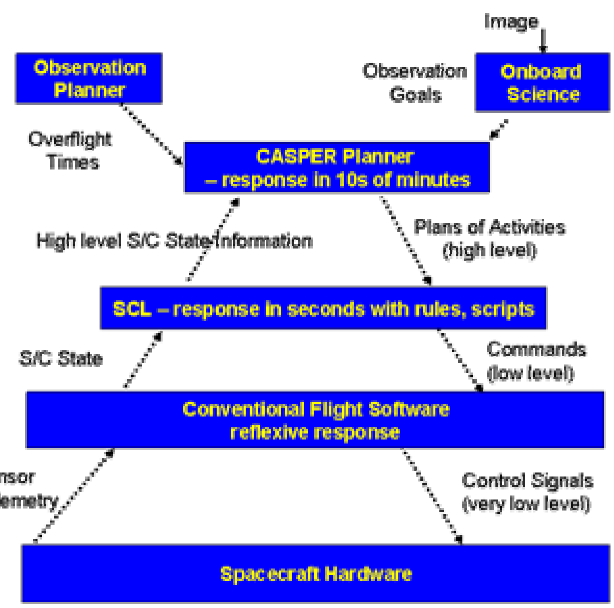

Figure 2. ***Autonomy Software Architecture***

1) Onboard Science Analysis

The first step in the autonomous science decision cycle is detection of interesting science events. In the complete experiment, a number of science analysis technologies have been flown including:

- Thermal anomaly detection – uses infrared spectra peaks to detect lava flows and other volcanic activity [19]. (See Figure 3.)
- Cloud detection – uses intensities at six different spectra and thresholds to identify likely clouds in scenes [18]. (See Figure 4.)
- Flood scene classification – uses ratios at several spectra to identify signatures of water inundation as well as vegetation changes caused by flooding [21]. (See Figure 5.)

- Cryosphere classification – uses multiple spectra to identify water, ice, and snow regions on the Earth [20]. This technique is used in conjunction with cloud detection. (See Figure 6.)

Figure 3 shows both the visible and the infrared bands of the same image of the Mt. Etna volcano in Italy. The infrared bands are used to detect hot areas that might represent fresh lava flows within the image. In this picture, these hot spots are circled with red dotted lines. The area of hot pixels can be compared with the count of hot pixels from a previous image of the same area to determine if change has occurred. If there has been change, a new image might be triggered to get a more detailed look at the eruption.

Figure 4 shows a Hyperion scene and the results of the cloud detection algorithm. This MIT Lincoln Lab developed algorithm is able to discriminate between cloud pixels and land pixels within an image. Specifically, the grey area in the detection results is clouds while the blue area is land. The results of this algorithm can be used to discard images that are too cloudy.

Figure 5 contains 4 images. The top two are detailed Hyperion images taken of the Larson Ice Shelf in Antarctica on 4/6/2002 and 4/13/2002. A large change in the ice shelf is seen in comparing the images. The bottom 2 images are results of the land-ice-water detection algorithm. The white area of the image is ice and the blue area is water. The ice and water pixels can be counted and compared with the second image to determine if change has occurred. If change is detected, the image can be downlinked and further images of the area can be planned.

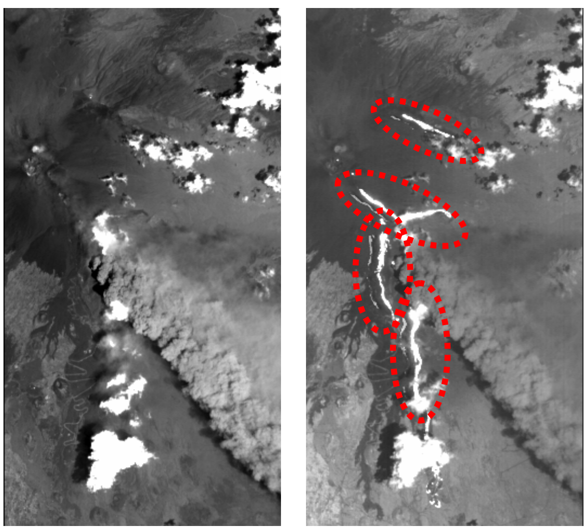

Figure 3. Thermal Anomalies associated with volcano activity at Mt. Etna, visual spectra at left and infrared at right.

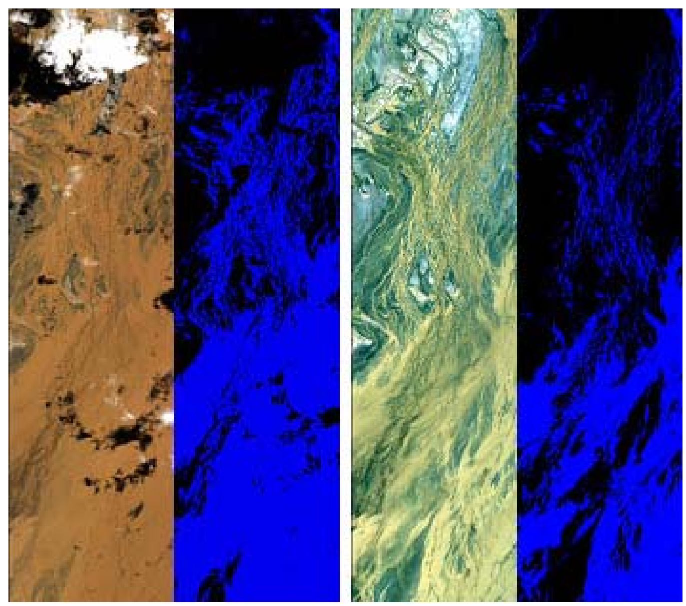

Figure 6. Flood Detection Scenes indicating receding flood region at the Diamantina River in Australia.

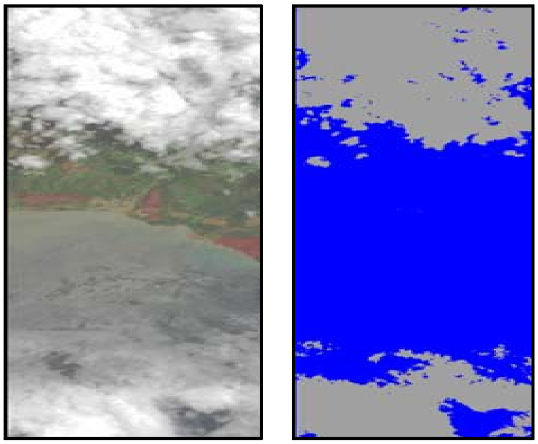

Figure 4. Cloud Detection of a Hyperion Scene – visual image at left, grey in the image at right indicates detected cloud.

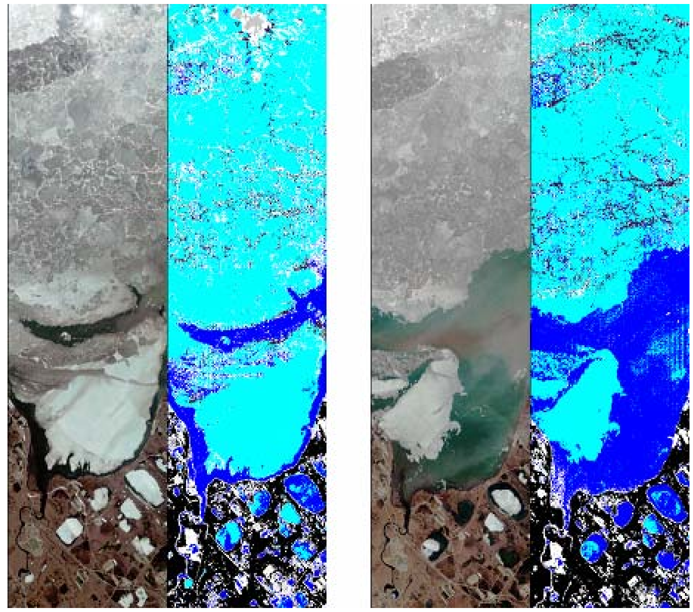

Figure 5. Change Detection Scenes indicating Ice Breakup in Prudhoe Bay, Alaska.

The onboard science algorithms are limited to using 12 bands of the hyperion instrument. Of these 12 bands, 6 are dedicated to the cloud detection algorithm. The other six are varied depending on which science algorithm is used. The images used by the algorithm are "Level 0.5," an intermediate processing level between the raw Level 0, and the fully ground processed Level 1. Each of the science algorithms except the generalized feature detection use simple threshold checks on the spectral bands to classify the pixels.

Initial experiments used the cloud detection triggers. The MIT Lincoln Lab developed cloud detection algorithm uses a combination of spectral bands to discriminate between clouds and surface features. The Hyperion Cloud Cover (HCC) algorithm have been run on all images acquired during ASE experiments. In the event of high cloud cover, the image could be discarded and a new goal could be sent to CASPER to reimage the area or image another high priority area. Images with low cloud cover can either be downlinked or analyzed further by other ASE science algorithms.

The JPL developed thermal anomaly algorithms use the infrared spectral bands to detect sites of active volcanism. There are two different algorithms, one for day time images and one for night time images. The algorithms compare the number of thermally active pixels within the image with the count from a previous image to determine if new volcanism is present. If no new volcanism is present, the image can be discarded onboard. Otherwise, the entire image or the interesting section of the image can be downlinked.

The University of Arizona developed flood scene classification algorithm (See Figure 6) uses multiple spectral bands to differentiate between land and water. The results of the algorithm include are compared with land and water counts from a previous image to determine if flooding has occurred. If significant flooding has been detected, the image can be downlinked. In addition, a new goal can be sent to the CASPER planning software to image adjacent regions on subsequent orbits to determine the extent of the flooding. We have noticed a few problems when ground testing this algorithm with existing hyperion data. The presence of clouds or heavy smoke within an image can cause the algorithm to fail.

The Arizona State University developed Snow-Water-Ice-Land (SWIL) algorithm is used to detect lake freeze/thaw cycles and seasonal sea ice. The SWIL algorithm uses six spectral bands for analysis.

1) Onboard Mission Planning

In order for the spacecraft to respond autonomously to the science event, it must be able to independently perform the mission planning function. This requires software that can model all spacecraft and mission constraints. The CASPER [5] software performs this function for ASE. CASPER represents the operations constraints in a general modeling language and reasons about these constraints to generate new operations plans that respect spacecraft and mission constraints and resources. CASPER uses a local search approach [15] to develop operations plans.

Because onboard computing resources are scarce, CASPER must be very efficient in generating plans. While a typical desktop or laptop PC may have 2000-3000 MIPS performance, 5-20 MIPS is more typical onboard a spacecraft. In the case of EO-1, the Mongoose V CPU has approximately 8 MIPS. Of the 3 software packages, CASPER is by far the most computationally intensive. For that reason, our optimization efforts were focused on CASPER. Since the software was already written and we didn't have funding to make major changes in the software, we had to focus on developing an EO-1 CASPER model that didn't require a lot of planning iterations. For that reason, the model has only a handful of resources to reason about. This ensures that CASPER is able to build a plan in tens of minutes on the relatively slow CPU.

CASPER is responsible for long-term mission planning in response to both science goals derived onboard as well as anomalies. In this role, CASPER must plan and schedule activities to achieve science and engineering goals while respecting resource and other spacecraft operations constraints. For example, when acquiring an initial image, a volcanic event is detected. This event may warrant a high priority request for a subsequent image of the target to study the evolving phenomena. In this case, CASPER modifies the operations plan to include the necessary activities to re-image. This may include determining the next over flight opportunity, ensuring that the spacecraft is pointed appropriately, that sufficient power, and data storage are available, that appropriate calibration images are acquired, and that the instrument is properly prepared for the data acquisition.

In the context of ASE, CASPER reasons about the majority of spacecraft operations constraints directly in its modeling language. However, there are a few notable exceptions. First, the over flight constraints are calculated using ground-based orbit analysis tools. The over flight opportunities and pointing required for all targets of interest are uploaded as a table and utilized by CASPER to plan. Second, the ground operations team manages the momentum of the reaction wheels for the EO-1 spacecraft. This is because of the complexity of the momentum management process caused by the EO-1 configuration of three reaction wheels rather than four.

2) Onboard Robust Execution

ASE uses the Spacecraft Command Language (SCL) [10] to provide robust execution. SCL is a software package that integrates procedural programming with a real-time, forward-chaining, rule-based system. A publish/subscribe software bus allows the distribution of notification and request messages to integrate SCL with other onboard software. This design enables both loose or tight coupling between SCL and other flight software as appropriate.

The SCL "smart" executive supports the command and control function. Users can define scripts in an English-like manner. Compiled on the ground, those scripts can be dynamically loaded onboard and executed at an absolute or relative time. Ground-based absolute time script scheduling is equivalent to the traditional procedural approach to spacecraft operations based on time. In the EO-1 experiment, SCL scripts are planned and scheduled by the CASPER onboard planner. The science analysis algorithms and SCL work in a cooperative manner to generate new goals for CASPER. These goals are sent as messages on the software bus.

Many aspects of autonomy are implemented in SCL. For example, SCL implements many constraint checks that are redundant with those in the EO-1 fault protection software. Before SCL sends each command to the EO-1 command processor, it undergoes a series of constraint checks to ensure that it is a valid command. Any pre-requisite states required by the command are checked (such as the communications system being in the correct mode to accept a command). Using SCL to check these constraints (while included in the CASPER model) provides an additional level of safety to the autonomy flight software.

C. Specific Performance Advance

The current state-of-the-art for flight missions is to:

- Develop a multiple day sequence of activities on the ground
- Uplink that sequence
- Collect science data using instruments
- Downlink all science data collected
- On the ground, analyze the science data, which will impact what science data is taken in the next muliple day sequence
- Repeat this process

Using ASE, the preliminary science analysis, updated goals, and reimaging are done onboard. The result of which is reducing the reaction time to science events from several days to several hours.

D. Rationale for Flight Validation

There are two important reasons why this software has to be flight validated.

1. It is very difficult to test software that interacts with unknown environments combined with tight real-time timing constraints on the ground.

2. No future mission/project manager would ever consider flying this autonomy software if it had not proved itself in a flight environment. In fact, the main reason we have met our success criteria by several orders of magnitude is to prove that this software can work operationally, not just as an experiment.

E. Benefits to NASA Missions

ASE as autonomous operations represents a revolutionary new approach to space exploration. ASE represents a complete paradigm change in terms of enabling a spacecraft to respond autonomously to detected science events. This is in stark contrast to traditional labor-intensive ground-based operations. This autonomous operations concept has been well proven in over one year of operations on EO-1 with over 3000 images acquired. This autonomy will enable a new class of missions at NASA including subsurface explorers, autonomous rovers, and coordinated systems of multiple spacecraft/sensors.

A summary of the significant benefits of ASE include:

- Using ASE, much of the extremely costly sequencing elements of spacecraft mission operations can be eliminated, dramatically reducing overall operations cost. (See Tangible Value Section for quantitative values)
- Using ASE, a goal-based spacecraft could perform opportunistic science and enable interactive science enabling a direct connection between the scientist and spacecraft.
- Autonomous operations is much more responsive to the environmental changes and uncertainty which can cause other systems to fail.
- Autonomous operations allow spacecraft missions, ground-based systems, and in-situ rovers to do more science for the same and in some cases lesser cost. On EO-1 ASE has documented a 100x increase in science return by enabling onboard tracking of dynamic science events.
- ASE enables some types of science that were previously impossible (e.g. unknown environments or no communications).
- ASE enables onboard data editing – allowing the most important science data to be returned. (For example, in the middle latitudes of the Earth, at any point in time, over 50% of the surface is obscured by clouds. Using ASE to filter these images and re-plan for other high priority images, the number of usable science images is increased by 100%. Using only this simple algorithm with onboard autonomy doubles the value of the mission.)

1) Cost Savings to NASA Missions

Automated planning systems all but eliminate the need for the mission operation's team to manually generate sequences, dramatically reducing costs. For example, using the automated planning system to command the Data-Chaser shuttle payload reduced commanding-related mission operations effort by 80% compared to manual sequence generation [Chien et al. 1999]. By combining automated planning with onboard science analysis and smart execution, an even more dramatic reduction is sequencing effort is obtained due to the reduction in sequences created in response to ground based science data analysis. These sequences are created by ASE onboard the spacecraft without ground interaction. Other cost savings are achieved by returning more valuable science for the original spacecraft investment.

As ASE was moved to operational status, it enabled automation of many mission planning and other operations efforts. From FY05 to FY06, the EO-1 missions cost was reduced from $3.6M/year to $1.6M/year. Approximately $1M/year of this reduction is directly attributed to the ASE software.

2) Other Measurable Value to NASA Missions

Increased *autonomy*. An autonomous spacecraft can more readily perform opportunistic science. When an unexpected opportunity occurs (such as a supernova or solar phenomena) the spacecraft can immediately perform appropriate measurements rather than wait until the ground operations team detects the event and uplinks commands to the spacecraft.

Increased *interactivity*. A goal-based autonomous spacecraft can facilitate interactive science. A self-commanding spacecraft can perform high-level science requests, such as "Perform an interferometry sweep with priority 5," A direct connection and faster feedback between scientist and spacecraft create a new model for scientific discovery in space.

Increased productivity. Autonomous systems technology has the potential to increase science return. It does this by producing operations plans that better optimize the use of scarce science resources or by quickly packing in science activities when the required turn-around time is small.

Simplified self-*monitoring*. Autonomous systems technology simplifies self-monitoring, onboard fault-management, and spacecraft-health tasks. Because the spacecraft can respond directly without waiting for ground communication, it can cover a greater range of faults.

II. Validation Objective and Approach

A. Relevant Environments and TRL Levels

There are 5 relevant environments and associated Technology Readiness Levels (TRLs) related to the validation of the ASE software:

- Workstation Environment – ASE software components are integrated and running on PC, Sun, or Linux workstations using test data. (TRL 4)
- Lab Testbed Environment – ASE components are integrated and running on a realtime CPU (Gespac) based testbed running VxWorks. (TRL 5)
- Flight Testbed Environment – ASE components are integrated and running on the EO-1 Mongoose V based flight testbed running VxWorks. (TRL 6)
- Flight Demonstration – ASE software is running onboard the EO-1 satellite. (TRL 7-8)
- Flight Operations – ASE software is used full time onboard the EO-1 satellite. (TRL 9)

B. Technology Validation Objectives and Success Criteria

The ASE flight validation objectives were captured in a signed Technology Validation Agreement between the ASE Team and the ST6 project.

1) Ground Test Objectives

The primary validation objective during ground testing and initial flight testing was to verify that the software subsystems (CASPER, SCL, science algorithms, bandstripping, software bus) were operating as expected. This was accomplished through a series of experiments to test the basic functionality of the software. The objectives for these experiment and success criteria are listed in Table 1. After each of these objectives is met in the lab testbed environment, the ASE TRL would change from 4 to 5. After each of these objectives is met in the flight testbed environment, the ASE TRL would change from 5 to 6

2) Flight Test Objectives

The flight test objectives and success criteria are listed in Table 2. There are 4 objectives, each comprising 25% of the technology validation success of ASE. After each of these objectives is met once, the ASE TRL changes from 6 to 7. After each of these objectives is met 5 times, the ASE TRL changes from 7 to 8. The following data are the expected observables used to measure the success criteria:

1. Image data and engineering telemetry returned as a result of ASE planned downlinks including ASE telemetry indicating proper functioning of the ASE software.
2. Image data returned as a result of ASE planned image data takes (we use the term data take to refer to the process of acquiring science data using one or more instruments).
3. Telemetry and associated image data for both positive and negative tests of onboard data editing. This includes the original image that was analyzed and the engineering telemetry containing the decision made as to whether the image is returned or discarded.
4. Telemetry and associated image data for both positive and negative tests of onboard image retargeting. This includes the original image, any follow up images, and the engineering telemetry indicating what features or change were detected in the science analysis of the original image.

The overall goals of the ASE software are to increase the value of science returned and decrease the operations cost. The questions to be answered by this experiment are:

1. Determine, through operational experiments, the amount of increased science achievable with ASE. This can only be achieved after multiple successful experiments (See Table 1.)
2. Determine, through operational experiments, to the extent possible, the exact level of expected savings in operations-staffing cost and on future NASA missions.

Table 1. Ground Technology Validation Objectives and Success Criteria

Ground Technology Validation Objectives (TRL 5-6)	Success Criteria
Test operation of CASPER planning software including creating plans using input goals, monitoring satellite state, and outputting high level commands to SCL.	CASPER is able to create plans using input goals (in less than 30 minutes), monitor satellite state, and output high level commands to SCL.
Test operation of SCL execution software including receiving CASPER commands, expanding them into spacecraft commands, sending satellite state information to CASPER, and monitoring command execution.	SCL execution software is able to receive CASPER commands, expand them into spacecraft commands, send satellite state information to CASPER, and monitor command execution. (All within timing constraints of spacecraft activities being performed.)
Test science algorithms and ground visualization environment using representative spacecraft data (EO1 archive data)	Science algorithms are able to accurately classify pixels and detect science phenomena using representative spacecraft data (EO1 archive data).
Perform unit-test and system test verification test runs in Gespac and EO-1 flight testbed environments for test of all ASE flight software capability.	Experiment scenarios run successfully (complete and meet performance goals) during unit-test and system test verification test runs in Gespac and EO-1 Flight testbed environments for all ASE flight software capability.

Table 2. Flight Technology Validation Objectives and Success Criteria

Flight Technology Validation Objectives (TRL 7-8)	Success Criteria
ASE shall autonomously plan, schedule and execute a payload data downlink onboard within one week of receiving a request. A request shall include a time window, a description of the data to be downlinked, and the view periods of the ground station.	Payload data is received successfully on the ground. Experiment repeated 5 times.
ASE shall autonomously plan, schedule and execute onboard a payload data collect (science observation) of a prescribed target area within two weeks of receiving a request. A request shall include the target location and imaging payload mode parameters.	Payload data is received successfully on the ground. Experiment repeated 5 times.
ASE shall control the content onboard of the downlink of payload data to retain only data of interest. Interest criteria include: a) change in the data as compared to a previous image, or b) detection of previously characterized features.	Payload data is either returned or discarded based on the results of a science algorithm. (5 times)
ASE shall perform onboard science analysis of the payload data to select and image (autonomously and onboard) a target in accordance with the following criteria: a) change in the data as compared to a previous image, or b) detection of previously characterized features. One of the five subsequent data collects shall be accomplished within 48 hours of the preceding data collect.	Instrument is retargeted based on the results of a science algorithm. (5 times, one within 48 hours)

I. TESTING

A. Ground Test

Gaining confidence in the safety of the EO-1 autonomy software requires extensive testing. We structured our testing methodology such that it would verify the protections provided by our layered architecture, and compliment the phases of the model development process. Specifically, the test plan was intended to validate the following system properties:

- CASPER generates plans consistent both with its internal model of the spacecraft and SCL's model and constraints (as checked by SCL).
- SCL does not issue any commands that violate the constraints of the spacecraft (as checked by our spacecraft simulator).
- Our model satisfies spacecraft operational and safety constraints enumerated by the model safety-review process.

We validated these three requirements by extensive testing of the autonomy software on generated test-cases, using checks at each layer to validate performance. The test cases described below address only the top-two levels of the onboard autonomy software (CASPER and SCL), with the flight software and spacecraft hardware replaced by a software simulator. Flight software testing and validation is addressed by a separate, more conventional, test plan

1) Test Case Parameters:

Each EO-1 test case covers seven days of operations containing multiple schedulable windows separated by a variable number of orbits. Each schedulable window represents an opportunity to schedule one or more science observations. The test cases must account for variations in the mission and science objectives (mission scenario parameters), initial state of the spacecraft (spacecraft state parameters), and changes to the spacecraft state during execution.

Since the autonomy software has no control over what happens outside of a schedulable window, we must be certain that our software performs reliably over a range of possible initial states. We cover these cases by using the simulator to vary the spacecraft state parameters as tracked within SCL and monitored by CASPER. The simulator also varies the spacecraft state parameters during execution to test the performance of our agent in the face of an uncertain environment.

Mission scenario parameters represent the high-level planning goals passed to CASPER. They are derived from a combination of the orbit of the spacecraft and the science objectives uplinked from the ground. They specify when targets have been available for imaging, as well as the parameters of a science observation (i.e. number of targets to image and science analysis algorithms we wish to execute).

The 22 spacecraft state parameters and 16 observation goal parameters used in the EO-1 test cases are shown in Table 3 and Table 4.

Table 3. Spacecraft State Parameters

Parameter	Expected Initial State
x-band ground station	unknown
x-band controller	enabled
ACS mode	nadir
target selected	unknown
WARP electronics mode	stndops
WARP mode	standby
WARP bytes allocated	0
WARP num files	0
fault protection	enabled
eclipse state	full sun
target view	unknown
hyperion instrument power	on
hyperion imaging mode	idle
hyperion cover state	closed
ali instrument power	on
ali active mechanism	telapercvr
ali mechanism power	disabled
ali fpe power	disabled
ale fpe data gate	disabled
ali cover state	closed
groundstation view	unknown
mission lock	unlocked

Table 4. Mission-Scenario Parameters

Parameter	Nominal	Off-nominal	Extreme
schedulable windows	0-3	3-5	5+
orbits between windows	2-7	1,8	0,8+
window start time	start of orbit	+/- 10 min	any

Parameter	**Nominal**	**Off-nominal**	**Extreme**
window duration	expected time of science analysis	+/- 10 min	any
target in-view start	anytime in orbit, 1 per orbit	1 per 3 orbits	any
target in-view duration	10 min +/- 1	+/- 3	any
groundstation in-view start	anytime in orbit, 1 per orbit	1 per 3 orbits	any
groundstation in-view duration	10 min +/- 1	+/- 3	any
eclipse start	60 min after orbit start	+/- 5	any
eclipse duration	30 min	+/- 5	any
imaging start	target view + 5 min	+/- 3	any
imaging duration	8 sec	+/- 4	any
science algorithm	any	Any	any
science goal start	fixed	not-specified	any
number of science goals	1 per orbit	1-3	>3
warp allocated	0	32K blocks	any

To exhaustively test every possible combination of state and observation parameters, even just assuming a nominal and failure case for each parameter and ignoring execution variations, would require a test set containing 2^{38} or 2.7×10^{11} test cases. Pruning the set of variations to just the sixteen observation parameters would still yield an impractically large set. The challenge thus becomes selecting a set of tests that most effectively cover the intractable space of possible parameter variations within a timeframe that allows for reasonable software delivery.

2) *Design of Test Cases*

Traditional flight software can be tested through exhaustive execution of a known set of sequences. Autonomy software however must be able to execute in, and react to, a much wider range of possible scenarios. As show above, testing all these possible scenarios would be intractable, however we can leverage the traditional nominal sequences and scenarios to baseline our tests – varying parameters off of a controlled scenario, and thus reducing the number of parameter variations our agent must consider. This is a similar approach to that used to validate the Remote Agent Planner for DS1. [11].

We started the design process by having spacecraft and operations personnel identify expected values for each parameter based upon the nominal mission scenario. Using these assignments we generated test cases by varying each of the parameters across three distinct classes of values – nominal (single value), off-nominal (range of acceptable values), and extreme (most likely failure conditions). For each parameter, based on this decomposition, we defined a set of five values at the boundaries of these classes – a minimum value, an "off-nominal-min" value at the boundary between the off-nominal and the extreme, a nominal value, an "off-nominal-max", and a maximum value (See Figure 7).

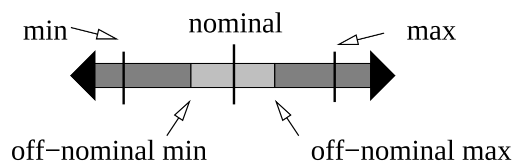

Figure 7. Parameter Decompositions

Using this decomposition of the test space, we generated three sets of test cases:

1. Coverage test cases that attempt to exercise a representative sample of all possible parameter-value assignments.
2. Stochastic test cases that verify nominal-operation scenarios.
3. Environmental test cases that evaluate how our agent performs in an uncertain environment.

a) Parameter-Coverage Test Set

Using the parameter decomposition we designed two sets of test cases, one that exercised the five values for each parameter while holding all other parameters within their nominal mission scenario, and another that exercised pair-wise combinations of parameter variations. Single-parameter variations allow for simple tests of off-nominal situations (variations that allow defects to be easily traced back to the source), while pair-wise combinations allow us to test the more complex interactions between parameters.

The single-parameter approach generates test sets that scale linearly with the number of parameters. Since we decomposed each of our parameters into five representative values, for N parameters, we have 5N test cases (or 4N+1 unique test cases as N of these have been the same nominal test set). For the EO-1 science agent this yields approximately 150 test cases. The pair-wise testing approach grows proportional to the number of pairs multiplied by the number of values for each pair or [k:2] * v * v for k parameters with v values. For EO-1, with 38 parameters each with 5 values, this gives us 17,575 test cases. Unfortunately even this number of test cases is impractical. Consequently we plan to use the method described in [6] and used by RAX [11], to reduce the number of pair-wise tests to a manageable level (under 100).

b) Stochastic Test Set

In all but one of the cases generated by our Coverage test set, some parameter has an off-nominal or extreme value. While these tests give us confidence in the robustness of our system, they do not provide much evidence as to the correctness of execution in nominal scenarios. In order to test more nominal scenarios, and also gain coverage in the off-nominal scenarios outside of the five representative values, we devised a scheme for generating stochastic test sets based on parameter value distributions.

Parameters were given normal distributions around their nominal value, with standard deviations half the width of the off-nominal range (such that 95% of expected values have been either nominal or off-nominal). Nominal test sets were then generated assigning values to parameters based on the defined distributions. Furthermore, by modifying the construction of the parameter distribution, we were able to create off-nominal and extreme test sets that would stochastically favor some parameters to choose values outside of their nominal range.

c) Environmental Test Set

We further extended the stochastic test sets described above to include execution variations based on the parameter distributions. The spacecraft simulator was modified to allow as input variations to expected parameter values. During the execution of activities the simulator simulates the change to each parameter of the current activity, and then varies the value returned based on the provided parameter distributions. Again nominal, off-nominal, and extreme test sets were generated that instructed the simulator to vary parameter values within the corresponding value class.

Finally we needed a way to test how the system responds to unexpected or exogenous events within the environment. These events could be fault conditions in the spacecraft or events outside of the CASPER model. Unlike the initial-state and execution-based testing described above, these events could happen at any time, and do not necessarily correspond to any commanded action or modeled spacecraft event. To accomplish this we added to our spacecraft simulator the ability to change the value of any parameter, at either an absolute time or time relative to the execution of an activity, to a fixed value or a value based on the distributions described above. We added small-variation events (within appropriate off-nominal and nominal classes) to our nominal and off-nominal stochastic test sets. Test cases are also currently in development that use this capability to exercise the fault scenarios outlined in the Spacecraft Safety document. [5]

3) Testing Procedure

The number of test cases we plan to run is limited by available testing resources and time. The EO-1 experiment had a compressed two month testing window with limited access to high-fidelity test beds. Two-thirds of testing time was spent evaluating output and running regression tests. The remaining one-third of our testing resources have been available for generated test cases. At two-hours per test run this gave us the capacity to run approximately 2400 tests.

Our automated test harness can detect “hard” test failures (i.e. crashes), and violations of system constraints (checked by the simulator). We also have “goal-detection” software that evaluates whether CASPER successfully executed the goals specified in the mission scenario.

As an additional complication, for EO-1 we had a number of testbeds with varying degrees of fidelity to the actual flight environment. The vast majority of tests were run on the Solaris and Linux testbeds, as they are the fastest and most readily available. However, these test the software under a different operating system, so are useful for testing of the model only. The operating system and timing differences are significant enough that many code behaviors occur only in the target operating system, compiler, and timing of interest. In order to validate aspects of the model dependent on precise timing we run tests on higher fidelity testbeds. These testbed tend to be much slower than the workstation testbeds. The testbeds available, fidelity, and TRL level at the conclusion of each set of tests are listed in Table 5.

Table 5. Testbeds Available to Validate ASE

Type	#	Fidelity	TRL
Solaris Sparc Ultra	5	Low – can test model but not timing	4
Linux 2.5 GHz	7	″	4
GESPAC PowerPC 100-450 MHz	9	Moderate – runs flight OS	5
JPL Flight Testbed RAD 3000	1	Moderate	5
EO-1 Flight Testbed Mongoose M5, 12 MHz	3	High – runs Flight Software	6

B. Flight Test

The flight tests were performed on the EO-1 satellite. At the conclusion of these tests, the ASE was advanced from TRL 6 to TRL 7. For each of the three experiment tests, the ASE code was uplinked several days before the test. Generally we would acquire several images using absolute time sequence (ATS) commands, with CASPER/SCL running in an idle state in the background in the hours leading up to the test. The CASPER input file is uplinked a few hours before the actual autonomous image data take. After the test is completed, the WARP processor is reset and restarts with the original flight software.

The next 3 sections describe the 4 flight experiment tests that were run to validate the ASE. The output of each of these tests included:

- Instrument image data
- Engineering log telemetry for CASPER and SCL actions
- The results of science algorithm including reduced scale (thumbnail) image of pixel classification and result of retargeting decision (for the Autonomous Retargeting Test)

1) Autonomous Data Take and Downlink Experiment Flight Tests

The objective of the autonomous data take experiment and resulting downlink test is to validate that the onboard software can develop and execute a plan to target the instrument, take an image, and return it to Earth. The input file contains activities to:

- Bias the reaction wheels
- Slew the satellite
- Acquire the image
- Turn back to Earth
- Re-bias the reaction wheels
- Downlink the engineering and image data (s-band and x-band)
- Clear the data off the WARP recorder

2) Data Editing Experiment Test

There were two types of scenarios for the data editing experiment test:

The first type is to validate that the onboard software can develop and execute a plan to target the instrument, take an image, analyze it onboard for cloud coverage, then either return it to Earth if the cloud coverage is low, or delete it if the cloud coverage is high. The input file contains activities to:

- Bias the reaction wheels
- Slew the satellite
- Acquire the image
- Turn back to Earth
- Re-bias the reaction wheels
- Run the cloud cover algorithm on the acquired image
- A dummy activity that would delete the image if that was required
- Downlink the engineering and image data (s-band and x-band). Note: The engineering telemetry contains the decision of whether or not the data would have been deleted.
- Clear the data off the WARP recorder

The second type is to validate that onboard software can develop a plan to target the instrument, take an image, analyze it onboard to detect areas of interest, and send down a summary product of that interesting area. The input file contains activities to:

- Bias the reaction wheels
- Slew the satellite
- Acquire the image
- Turn back to Earth
- Re-bias the reaction wheels
- Run the science classifier on the acquired image
- Write out a summary product for the area of interest
- Downlink the summary product on the next available ground contact.

3) Autonomous Retargeting Experiment Test

The objective of the autonomous retargeting experiment test is to validate that the onboard software can develop and execute a plan to target the instrument, take an image, analyze it onboard for science content, return it to Earth, then either reimage the same area if the science analysis warrants it, or image a different area. The input file contains activities to:

- Bias the reaction wheels
- Slew the satellite
- Acquire the image
- Turn back to Earth
- Re-bias the reaction wheels
- Run the selected science algorithm on the acquired image
- Downlink the engineering and a thumbnail version of the image (s-band). Note: The engineering telemetry contains the results of the science algorithm and the decision of which follow up image was selected.
- Return the full image
- Clear the data off the WARP recorder

- Bias the reaction wheels
- Slew the satellite
- Acquire the follow-up image
- Turn back to Earth
- Re-bias the reaction wheels
- Downlink the engineering and follow-up image data (s-band and x-band)
- Clear the data off the WARP recorder

C. Evaluating the experiment results

Prior to each flight experiment, a experiment scenario was devised by the ASE and EO-1 teams about 1 week prior to the test. Once the scenario had been defined, multiple test runs were conducted on the EO-1 flight testbed. The results of each test run were analyzed the ensure the following:

- CASPER was able to generate a conflict free plan.
- The set of requested goals scheduled by CASPER was consistent with the initial conditions and constraints modeled.
- The spacecraft commands issued by SCL were on time.
- The temporal sequence of spacecraft commands issued by SCL was considered safe.

We also collected data from each test run to compare the results of the ground testbed with the flight experiment. Specifically we collected data on:

- The number of repair iterations required by CASPER to generate a conflict free schedule.
- The amount of time required for each CASPER repair iteration.
- The time spacecraft commands were issued by SCL.
- The time each subsystem received the command as issued by SCL.
- The time SCL received notification of the change in the state of the spacecraft after issuing a spacecraft command.

After each flight experiment, the logs and telemetry from each ASE subsystem were checked against the collected on the ground testbed to ensure they were consistent. Each inconsistency was examined and the ASE software was updated as needed.

II. Technology Validation Summary

A. *Technology Validation Results Summary*

The technology was declared fully validated (TRL 8) in May 2004, after all 20 onboard autonomy experiments were tested as described in Section 2 and 3. TRL 7 was achieved in March 2004, after at least 1 of each of the 4 experiments were run successfully. The overall system performed as expected and was considered a success. The dates of each succesful test are listed in Table 6. Additional validation experiments have since been conducted, leading to a full flight/ground integration of the ASE software into EO-1 operations. The ASE software has been in use 24/7 on EO-1 since January 2005, raising the TRL level to 9.

Table 6. ASE Validation Summary

Experiments	When	# Additional Validations
1. Autonomous Downlink (25%)	01/08/2004 01/14/2004 01/22/2004 01/22/2004 01/29/2004	+1000 (as of 8/8/05)
2. Autonomous Image Acquire (25%)	01/08/2004 01/14/2004 01/22/2004 01/22/2004 01/29/2004	+2700 (as of 8/8/05)
3. Onboard Data Editing based on Science Analysis (25%)	03/25/2004 04/29/2004 05/07/2004 05/12/2004 05/14/2004	+107 (as of 8/8/05)
4. Onboard Retargeting based on Science Analysis (25%)	03/17/2004 03/25/2004 04/29/2004 05/07/2004 05/12/2004	+300 (as of 8/8/05)

B. *Summary and Discussion of Ground and Flight Tests*

1) *Autonomous Downlink Experiment Results*

For the autonomous downlink experiments, there were two types of downlinks performed on EO-1, an X-band and S-band downlink contact.

The X-band ground contact is used to downlink the science data from the solid state recorder. The X-band downlink was tested multiple times on the ground and in flight and each was successful. Comparisons of the timing of commands and receipt of spacecraft state updates were also consistent between ground and flight tests.

The S-band ground contact is used to downlink all the engineering data from the C&DH processor. During several ground tests of executing the S-band downlink goal, one major issue was encountered where the C&DH processor would perform a software reset when commanded by ASE; a condition that has never happened during the EO-1 mission. After checking the sequence of commands issued by ASE, it determined that they were consistent with the command sequence normally performed onboard EO-1. The reset was considered perplexing because the ASE software is hosted on the WARP processor and the only effects on the C&DH processor is through spacecraft commands. After more ground tests, it was determined to be a ground testbed issue, as software engineers were able to consistently reproduced the reset by ground commands, independent of the ASE software onboard the WARP processor. Once the problem was determined to be specific to the ground testbed, we continued with the flight experiment of the S-band contact, which was successful. Again, analysis of the sequence of the commands and timing of spacecraft state changes were consistent between ground and flight tests.

2) *Autonomous Data Take Results*

Data from the autonomous data take flight experiments were compared against the results acquired from ground tests. Each data point in the flight experiment compared well with the ground tests with one exception: the timing of receipt of several instrument commands issued from ASE.

Records from the flight experiment showed that the EO-1 science instruments would receive the commands to begin collecting data several seconds later after ASE would issue the command. For an eight second data take, this latency in the receipt is fairly significant. This differed from the results acquired on the ground because the flight testbeds are not equipped with instruments. Instead, on the flight testbeds, receipt of the commands and output of instruments telemetry were simulated by the VirtualSat software. Several more flight experiments confirmed that this was a repeatable and consistent timing error. To resolve the timing difference, the SCL model was updating to issue the command several seconds earlier than planned to account for the measured latency.

Determining, verifying, and updating this timing data was one of the many challenges in developing the CASPER and SCL model of the spacecraft. Within the SCL model, spacecraft state and resources are checked:

- Prior to executing the command to verify command prerequisites are met.

- After executing the command to verify the receipt of the command.
- After an elapsed time period when the effects of the command are expected.

Unfortunately, due to the varying rate of spacecraft data is available to SCL, it was not possible to perform the complete spacecraft state and resource check for each command. For time critical commands, several commands checks were completely omitted. Without these checks in SCL, the SCL loses its ability to issue "safe" commands and issues the commands similarly to the original flight software command sequence.

3) *Autonomous Data Editing Experiment Results*

For the two types of scenarios in the autonomous data editing experiment, there were no discrepancies between the flight and test runs. Each of the scenarios developed performed as expected and the data collected from each flight experiment were consistent with the ground tests.

In the scenario where the science classifier analyzed the collected image for cloud coverage, a dummy activity was inserted in the schedule to simulate the removal of the image from the solid state recorder. This was needed in order to downlink the image and verify the results of the cloud coverage algorithm.

For the scenario where the science classifier would output a summary product of the area of interest, the ASE team was constrained to write a maximum 20KB of data. The 20KB of data was the maximum number of bytes the EO-1 operations team could guarantee to be downlinked consistently during each ground contact. With ground contacts nominally scheduled every 3 hours, the summary product provided scientist with immediately relevant data for analysis.

4) *Autonomous Retargeting Experiment Results*

In this section, we discuss several of the differences between the ground and flight tests runs for an autonomous retargeting experiment, along with several of the issues that arose in developing the scenarios during mission planning.

One the major difference between the ground and flight tests runs for the autonomous retargeting experiment was determining the amount of time required to run the science algorithms. Several ground test runs of the science algorithms had been performed in previously collected data, but they were ran on the Linux testbeds, several orders of magnitude faster than the flight testbeds. On the flight testbed, a true solid state recorder is not available and the simulation of the science data was used, leading to a best guess for the execution time of the science algorithms during the flight experiments. After the flight experiments, the CASPER and SCL model for the duration of the science classifiers were updated to reflect the true duration for each of the classifiers.

We will now highlight a couple issues and challenges that occurred during flight testing of the autonomous data editing experiments.

a) Developing an autonomous retargeting scenario

One of the challenges encountered during the mission planning phase was attempting to introduce more complex autonomous retargeting scenarios. Currently, ASE scientists develop retargeting scenarios in the following manner:

- Image target X
- Playback data from the solid state recorder into RAM
- Run science event detector Y
- If science event is detected, re-image target X (the response scene) on the next opportunity

The other types of science campaigns we would have liked to specify are:

- If a science event is detected, re-image target X on the next 5 over-flights
- If a science event is detected, re-image target X on the next 5 over-flights, and on each over-flight, run science detector Y.
- If a science event is detected, continue to re-image target X until the science event is no longer present.

These types of scenarios, where one science scenario triggers multiple requests of scenes or other scenarios are difficult to specify for several reasons.

Each scenario to be executed requires the response scene to be uploaded prior to the science event detector being executed in order to place these new scenes into the schedule. However, inserting 5 new scenes into the schedule requires a projection of the over-flight time across multiple weeks. The further out the over-flight time is projected, the less precise the time. For example, to upload a target 2 weeks out may cause an error of a couple of seconds, and with a scenes nominally being 8 seconds long, this error is fairly significant.

Another issue with scheduling targets far into the future is handling any conflicts that may appear with other scenes or ground contacts. For example, if you have a high priority target in the schedule, this is should be accounted for when downlink contacts are scheduled. However, this requires knowledge of if a science event is detected onboard as an input into the scheduling meeting.

b) Mission planning an autonomous retargeting scenario

As ASE re-targeting scenarios are submitted for the weekly scheduling process, one of the issues that arose was how to consider the scenario as a whole, should one part of the scenario not be scheduled. For example, the three main activities for each scenario that may cause conflicts with other activities are the initial scene, the playback of the data from the solid-state recorder, and the response scene if a science event is detected.

This is currently handled in the following manner:

- If the initial target cannot be imaged, the scenario is rejected and no longer considered
- If the playback of data from the solid-state recorder will not fit within the schedule, the playback activity is removed from the schedule and initial scene is still considered.
- If the response scene will not fit within the schedule (i.e., a conflict with a downlink or oversubscribes the solid state recorder), the initial scene is still imaged, but the science event detector is no longer scheduled.

It would also be appropriate to remove the entire scenario from the schedule should the playback activity or response scene not fit within the schedule. However, this has caused several open gaps in the schedule that were not necessarily filled with other scenes. In order to populate the schedule as much as possible, we chose to keep the initial image in the schedule.

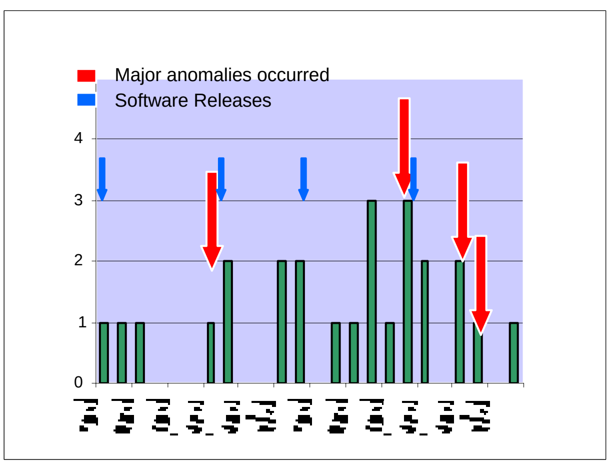

Figure 8. Timeline of Software Anomalies

C. Discussion of Onboard Anomalies

To date (July 2005), there has been approximately 25 total anomalies onboard EO-1, of which 4 resulted in the lost of scenes. The frequency of anomalies was generally higher after new releases of the ASE software were run on-board EO-1. A timeline of the anomalies is contained in Figure 8.

The anomalies that have occurred onboard can be classified into the following types: modeling, software, operator, and hardware.

- Modeling – This is the most common type of error, caused by an incorrect model of the spacecraft within CASPER and SCL. Many of these errors were not detected during testing and validation because the EO-1 mission did not have a high-fidelity testbed, requiring the development of simulators that made several incorrect assumptions of the spacecraft behavior.
- Software – These are your standard software implementation errors that occur with any large project.
- Operator – Commands are regularly issued from mission operators during ground contacts. These commands may modify the state of the spacecraft, so the ASE must be robust to these situations.
- Hardware – The Livingston 2 software component (which was uploaded to EO-1 in September of 2004) was designed to detect and diagnose this type of error. However, because hardware errors are rare on spacecraft systems, we chose not to focus on detecting these.

While 25 anomalies may seem to be a large number we would emphasize that they occurred over a long period of time (over 22 months) and this list includes many minor anomalies that did not impact operations or science data acquisition. Additionally, the experiment software was not intended to take over 24/7 operations, and thus it was not subjected to the more typical full suite of testing as more typical for mission operational software. However, at this point, through operations and subsequent testing, it has reached that level of maturity and reliability of mission operations software. For further details see Appendix C.

D. Operational Effectiveness Assessment

In traditional ground operations, a science operations team must select science targets, often weeks or months in advance. On EO-1, nominal scene selection was 7-14 days before over flight. The spacecraft operations team would then process these science targets, by building a low-level command sequence using a labor and knowledge intensive process. If last minute changes could be made to the observation plan, they would only be in extreme circumstances and would require significant additional effort by the operations team.

In contrast, ASE enables automation of the spacecraft operations flow both ground-based and onboard. On the ground, ASE enables rapid, automated generation of operations plans. In the event of anomalies, new operations plans can be generated without significant effort. This was tested recently when two of the ground stations were lost due to hardware issues (Svalbard and Poker Flats). These stations normally provide over half of the EO-1 ground tracks. Using the ASE automation, new schedules could be built within hours, enabling a minimal disruption of EO-1 science operations.

Onboard, ASE enables rapid decision-making to capture the highest priority science events without the delay incurred by traditional ground-in-the-loop operations. This decision-making can be used to acquire the highest priority scenes (based on onboard analysis) as well as respond rapidly to capture transient science events (such as volcanoes, floods, and sea ice breakup and formation).

The ASE software was designed to meet a number of unique requirements imposed by space autonomy.

- ASE had to be able to operate autonomously for significant periods of time. EO-1 has approximately 5 ground contacts per day, 10 minutes each. The remainder of the time there is no contact with the spacecraft. ASE was designed for deep space operations where commanding might occur every week or 2 weeks
- ASE had to operate despite limited observability. Because spacecraft have limited sensing, limited onboard storage, and limited downlink, ASE had to be able to accurately track spacecraft state and enable ground teams to do so as well.
- ASE had to model complex operations with multiple interacting science constraints and complex spacecraft operations constraints.
- ASE had to meet these requirements despite very limited onboard computing – on EO-1 we had only 4 MIPS and 128 MB RAM
- ASE had to be extremely reliable. With EO-1 mission cost over $100M, unreliable operations that risked the spacecraft are unacceptable.

This autonomous operations concept has been well proven in over one year of operations on EO-1 with over 3000 images (as of July 2005) acquired. The ASE software has been operating onboard the Earth Observing One Mission as their primary operations system since January 2005.

1) *Results of Increased Science Return*

ASE has achieved significant cost savings by returning more valuable science for the original spacecraft investment. For the specific case of using the ASE software on EO-1, science return per data downlink was increased by over 100x by rapid response and returning the most important science data. (See Table 7.) This compares with an original goal of a 10x increase in science return.

Table 7. Increased Science by Process

Process	Total Process Data Acquired	Data returned by ASE	Downlink Savings	Savings Factor (goal was x10)
Volcanism	33750 MB	294 MB	33456 MB	115
Cyrosphere (ice)	38100 MB	304 MB	37796 MB	125
Flooding	25500 MB	239 MB	25261 MB	106
Total	**97350 MB**	**837 MB**	**96513 MB**	**116**

2) *Determination of Cost Savings*

For specific cost savings (or value added) from increased science return from ASE, we submit the following analysis. To compute an economic value to the baseline EO-1 science return we use a conservative estimate based on what current customers pay for scenes.

$1000/image* x 8 images/day** x 25 days/month*** x 12 months/year = $2.4M/year

We take this as a conservative estimate of the value of the science return from conventional EO-1 Operations.

Assuming a conservative science increase of 10x (compared to the documented increase of over 100x), ASE has increased the science return of EO-1 as follows:

Science return with ASE – Science return without ASE =
10x $2.4M/yr – 1x $2.4M/yr =
$21.6M/yr

Reducing the ground operation team, which no longer has to prepare detailed spacecraft command sequence files, saves additional costs. In the case of EO-1, the labor costs for the ground operations team were reduced from $3.6M/year to $1.6M/year with approximately $1.0M/year directly as a result of using the ASE software. ****

Over a 5 year mission, using EO-1 as the example, the ASE software has the potential for saving:

Science value increase	$21.6M/yr x 5 yrs	= $108.0M
Operations cost reduction	$ 1.0M/yr x 5 yrs	= $5.0M
Total $ Savings	**$108.0M + $5.0M**	**= $113M**

* Conservative cost charged per image to EO-1 customers; actual costs ranges from $1000 - $2000.
** Typical number of paid images per day.
*** Conservative estimate of science operations days per month – roughly 3-5 days per month lost for engineering operations.
**** Estimate from EO-1 Mission Manager Dan Mandl at GSFC

III. Technology Infusion

ASE has produced a Technology Infusion Plan with specific action items on how the ASE technology is infused into future missions. The following list is a summary of tools that are being used to infuse ASE:

1. General Lectures. The advertising posters and the electronic announcements must be well thought out to assure the proper audience. One set of these lectures would be at NASA centers such as JPL, GSFC, ARC, and JSC. Another set would be at Universities and institutions with high space mission involvement such as ASU, University of Arizona, SWRI, etc.
2. Lunch seminars. These will be given at a rage of companies such as NGST, Lockheed or Ball as well as laboratories such as NRL, NRO, and APL.
3. Papers presented and published at powerful conferences. The papers are important because they leave a trail and allow interested parties to locate you. The presentations are more effective in getting the attendees interested and convinced.
4. Posters presented at conference poster sessions. Some scientists are turned off by glitzy presentations but like to meander around and ask detailed questions during the poster sessions at major conferences. The posters also attract a good share of post docs and fresh PhD's who work with PI's on mission proposals.
5. Individual presentations. These people need to be briefed one-on-one. This includes major scientists, Mars program scientists, managers of future advanced missions in early development stage and some NASA HQ decision makers.
6. Targeted audience short seminars. For example, MSL team, Mars science team, Exploration Office new technology briefing.
7. Magazine articles, Press Releases. The assault must be on all fronts to make an impression that this is a "now" technology. People need to see the ASE type of messages from all angles, on TV, in papers, conferences, etc. Magazines are a part of this assault.

A. Adaptability for Future Missions

ASE has applications for mapping and monitoring missions. Examples include, but are not limited to:

- Tracking lightning in planetary atmospheres, such as Jupiter and Saturn.
- Monitoring volcanism on Io from jovian or Io-dedicated mission: this would include identification of large, rare outbursts for preferential study and would afford auto-gain set or exposure duration change on detection of sensor saturation.
- Searching for volcanic or cryovolcanic thermal emission on Mars and icy satellites (especially Europa and Triton) from monitoring compressed datastreams from IR-sensitive instruments. Night-time data can be
- Tracking polar ice cap change on Mars.
- Searching for dust-devils and clouds on Mars in surface image data.
- Monitoring atmospheric weather conditions searching for polar volatile flow direction and detection of sand-storms.
- Detecting and tracking comet outbursts.

For further esriptions of the wide range of science applications of this technology please see the ASE Science Report nd the technology interdependencies section later in this report.

B. Impact on Future Missions

ASE can impact several aspects of spacecraft operations. The mission planning process is simplified because the operations team no longer has to build detailed sequences of commands. The spacecraft can be commanded using high-level goals, which are then detailed by the planner onboard. The processes of planning, build sequence, upload sequence, execute sequence, downlink data, analyze data, and build new sequence are entirely automated using ASE. For example, in the current EO-1 operations, a significant percentage of the images downlinked are of no value because they are mostly covered in clouds. Using ASE, these images can now be discarded onboard and the satellite can acquire another image of a different area. This saves time and labor for the mission planning team, science analysis team, ground station team, flight operations team, and data processing and archive team.

Due to initial funding limitations, the initial ASE for EO-1 did not include an autonomous fault protection component. However, fault protection is a natural fit for the ASE onboard autonomy software. In one example, CASPER generates a mission level plan that includes a sequence of behavior goals, such as producing thrust. The SCL executive is responsible for reducing these goals to a control sequence, for example, opening the relevant set of valves leading to a main engine. A device, such as a valve, is commanded indirectly; hence, SCL must ensure that the components along the control path to the device are healthy and operating before commanding that device. Components may be faulty, and redundant options for achieving a goal may exist; hence, SCL must ascertain the health state of components, determine repair options when viable, and select a course of action among the space of redundant options. Adding this level of fault protection autonomy to a future mission could in theory, eliminate the spacecraft analysis team. The team would no longer be required to monitor the spacecraft health because that would be done onboard using model-based mode estimation and mode reconfiguration. [16] The team would also not be required to respond to "safe-hold" periods because anomalies would be handled and reconfigured onboard. Using this software requires a greater up front investment in building the spacecraft models, but much of the underlying software has already been developed in research efforts.

In continuing flight, in partnership with the NASA Ames Research Center, we were able to integrate the Livingstone 2 Model-based mode identification system [10] with the ASE software, flying this package from Fall 2004 through August 2005 (as this report is filed).

Using the onboard science analysis software can also save time and labor for the science team. The feature detection algorithms can identify specific features of interest within the images. The spacecraft can then downlink the entire image when features are detected, only the detected features, or even a summary of the detected features. Scientists no longer have to analyze many different images to find a feature of interest. In fact, images that do not contain features of interest do no even have to be downlinked. These algorithms can be particularly useful on bandwidth-limited missions by returning the most important science data.

C. *Technology Interdependencies*

ST6 ASE was designed to have minimal impact on the operation of the revenue generating EO-1 science mission. There are, however, some important interdependencies to note for future missions that may be interested in deploying the technology. These are summarized in this section.

The planned science experiments meet many requirements of NASA's Earth Science Enterprise, primarily program aspects prescribed under the Natural Hazards Program [NASA ESE 2000]. We are primary concerned with Space Science applications, which are described in this section.

Onboard science data processing, as validated by ASE, has been identified by the NASA Space Science Technology Steering Group as an enabling technology for several Exploration of the Solar System (ESS) missions including Europa Orbiter (EO), Pluto Express (PE), Neptune Orbiter (NO), and Saturn Ring Observer (SRO). Specifically, the feature tracking and feature recognition technologies to be demonstrated through this report are considered highly enabling to these missions. In addition, eight Sun-Earth Connection (SEC) missions (GEC, ISP, MC, MMS, RAM, RBM, PASO, SN) and three Structure and Evolution of the Universe missions (ARISE, CON-X, OWL) have identified the need for this technology.

Specifically, the direct sciencecraft onboard science processing described in this report has numerous applications to Space Science Missions. For example, in Europa orbiter and lander missions, onboard science processing could be used to *autonomously:*

- Monitor surface change as function of changing tidal stress field
- Monitor areas of greatest tidal stresses
- Search for surface change, that is, evidence of recent activity
- Search for landing sites which have a high probability of lander survivability and where the crust is thin enough for deployment of a sub-crust submarine explorer

An Io volcano observer (Firebird/Argus, proposed under the Discovery program) would investigate volcanic phenomena, some of which have a direct bearing on understanding the evolution of the Earth, and derive the internal structure of Io, to better understand the transfer of tidal energy to Europa. This has direct bearing as to whether or not Europa is volcanically active: if so, this may increase the chances of a life-sustaining environment. Building on what has been learned from studying terrestrial volcanoes, a dedicated Io mission would map the changing shape of volcanism, measure tectonic stresses at global, regional, and local levels, and detect and quantify surface feature planform and topographic change. Additionally, a high degree of autonomous operation is necessary with an Io observer to allow real-time target switching if a high-priority, transient phenomenon occurs (e.g., an explosive fire-fountain event, with a lifetime on the order of minutes to hours).

As part of the NO mission, cryogenic volcanism and other phenomena could be studied on Triton using change recognition technology, e.g., study of active nitrogen plumes, resurfacing by flows, and changes in seasonal ice caps, as well as atmospheric changes on Neptune.

The highly successful Magellan radar-mapping mission imaged the surface of Venus, which is geologically young. A subsequent mission satellite constellation would initially search for evidence of change through comparison with Magellan data, with the option to reconfigure for high-resolution observation of areas where change has occurred, prior to the planned NASA Venus Surface Sample Return mission.

Table 8. Summary of Relevant Mars Missions

Launch Year	Mission
2001	Mars Odyssey
2003	Mars Exploration Rovers Mars Express Orbiter (ESA)
2005	Mars Reconnaissance Orbiter
2007	Phoenix Lander
2009	Mars Science Laboratory
2011+	Sample return mission

Mars is the target of a series of missions by NASA and other organizations. These missions are summarized in Table 8. Included are missions to monitor ice cap change, search for wind streaks, and changes in dune fields, as well as search for water-related change, such as mass-wasting and debris flow processes (Malin and Edgett, 2000). Of particular importance is the task of landing site selection. Selection algorithms can be pre-tested on terrestrial analogs. Also interesting is the gradual construction of the Mars Network, which will yield a GPS-like capability. This would allow a low-cost second deployment to Mars of a variable-baseline interferometer SAR constellation.

On a more speculative level, Cassini-Huygens may reveal if Titan has oceans or lakes. An ASE-like radar-mapping mission (as part of Titan Organic Explorer) can penetrate the cloud cover (as with Venus) and map the surface. Such information would be invaluable for insuring safe landing of in-situ packages on the surface. In addition, high-resolution interferometry could be used to monitor coastline/boundary changes.

A robot outpost on Mars has been proposed to pave the way for human exploration. The outpost may consist of a hundred rovers, functioning as a robot colony. Such an undertaking, with a wide range of rovers both on and above the surface, will by its nature need to operate autonomously. The massive amount of data generated will need autonomous processing to extract science content, which will in part be used to determine subsequent colony operations. ASC is a step on the road to achieving this level of autonomy.

The ASE Team has identified the NASA Mars Program as an ideal candidate for technology infusion of the ASE software. As a result, we have been working closely with the Mars Odyssey Project to identify and ground test science analysis algorithms that could be used for discovery of interesting science on Mars. The goal of this work is to infuse the science analysis component of the ASE software in the Odyssey extended mission [22]. Because of the limited computing power on Odyssey, the planning and execution components can not be included.

The ASE software have been particularly important the planned Mars robotic missions that will serve as a precursor to manned exploration of Mars as defined in the current NASA Exploration Vision. Because these missions have been extremely complex (i.e, sample return missions), they will require a high degree of autonomy similar to what ASE provides.

IV. Acknowledgments

The New Millennium Program at the Jet Propulsion Laboratory funded this work. 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.

We would like to acknowledge the important contributions of Nghia Tang and Michael Burl of JPL, Dan Mandl, Stuart Frye, Seth Shulman, and Stephen Ungar of GSFC, Jerry Hengemihle and Bruce Trout of Microtel LLC, Jeff D'Agostino of the Hammers Corp., Robert Bote of Honeywell Corp., Jim Van Gaasbeck and Darrell Boyer of ICS, Michael Griffin and Hsiao-hua Burke of MIT Lincoln Labs, Ronald Greeley, Thomas Doggett, and Kevin Williams of ASU, and Victor Baker, Felipe Ip and James Dohm of University of Arizona.

In addition, this experiment would not have been possible without the dedicated support of the ASE Team Members:

Victor Baker, Darrell Boyer, Rebecca Castano, Benjamin Cichy, Jeff D'Agostino, Ashley Davies, Thomas Doggett, James Dohm, Stuart Frye, Ronald Greeley, Jerry Hengemihle, Dan Mandl, Gregg Rabideau, Seth Shulman, Daniel Tran, Bruce Trout, Jim Van Gaasbeck

The following program management and project staff also provided critical support throughout development and operations:

Thomas Brakke, Art Chmielewski, Bryant Cramer, Simon Hook, Chris Stevens, Jack Stocky, Al Tao, Magdalene Chang, and Stephen Ungar.

V. References

[1] M.C. Burl, L. Asker, P. Smyth, U. Fayyad, P. Perona, J. Aubele, and L. Crumpler, “Learning to Recognize Volcanoes on Venus,” Machine Learning Journal, April 1998.

[2] S. Chien, B. Engelhardt, R. Knight, G. Rabideau, R. Sherwood, E. Hansen, A. Ortiviz, C. Wilklow, S. Wichman , "Onboard Autonomy on the Three Corner Sat Mission," Proc i-SAIRAS 2001, Montreal, Canada, June 2001.

[3] S. Chien, R. Sherwood, M. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, P. Zetocha, R. Wainright, P. Klupar, P. Cappelaere, D. Surka, B. Williams, R. Greeley, V. Baker, J. Doan, "The TechSat 21 Autonomous Sciencecraft Constellation", Proc i-SAIRAS 2001, Montreal, Canada, June 2001.

[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] 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)

[6] Davies, A. G., E.D. Mjolsness, A.G. Gray, T.F. Mann, R. Castano, T.A. Estlin and R.S. Saunders (1999) Hypothesis-driven active data analysis of geological phenomena using semi-autonomous rovers: exploring simulations of Martian hydrothermal deposits. EOS, Trans. Amer. Geophys. Union, 80, no. 17, S210.

[7] E. Gat et al., Three-Layer Architectures. in D. Kortenkamp et al. eds. AI and Mobile Robots. AAAI Press, 1998.

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

[9] 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.

[10] S. Hayden, A. Sweet, S. Christa, D. Tran, S. Shulman, “Advanced Diagnostic System on Earth Observing One,” Proceedings of AIAA Space 2004 Conference and Exhibit.

[11] Interface and Control Systems, SCL Home Page, sclrules.com

[12] M. C. Malin and K. S. Edgett, “Evidence for recent groundwater seepage and surface runoff on Mars,” Science, 288, 2330-2335, 2000.

[13] N. Muscettola, G. Dorais, C. Fry, R. Levinson, and C. Plaunt, “IDEA: Planning at the Core of Autonomous Reactive Agents,” Proceedings of the Workshops at the AIPS-2002 Conference, Tolouse, France, April 2002.

[14] NASA Ames, Remote Agent Experiment Home Page, http://ic.arc.nasa.gov/projects/remote-agent/.

[15] The PROBA Onboard Autonomy Platform, http://www.estec.esa.nl/proba/

[16] 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.

[17] B.C. Williams and P.P. Nayak, “A Reactive Planner for a Model-based Executive,” Proc. International Joint Conference on Artificial Intelligence, Nagoya, Japan, 1997.

[18] M. Griffin, Burke, H. Burke, D. Mandl, J. Miller, “Cloud Cover Detection Algorithm for EO-1 Hyperion Imagery,” 17th SPIE AeroSense 2003, Orlando FL, April 21-25, 2003.

[19] Davies, A. G., S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castaňo, B. Cichy, G. Rabideau, D. Tran and R. Sherwood (2005) Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE), Remote Sensing of the Environment, to appear.

[20] Doggett, T., R. Greeley, A. G. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm and F. Ip (2005) Autonomous On-board detection of Cryospheric Change, Remote Sensing of the Environment, to appear.

[21] Ip, F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, B. Castaňo, S. Chien, B. Cichy, R. Greeley, R. Sherwood (2005) Development and Testing of the

Autonomous Spacecraft Experiment (ASE) floodwater classifiers: Real-time Smart Reconnaissance of Transient Flooding, Remote Sensing of the Environment, to appear.

[22] Kiri L. Wagstaff, Rebecca Castano, Steve Chien, Anton B. Ivanov, and Timothy N. Titus. An Onboard Data Analysis Method to Track the Seasonal Polar Caps on Mars. Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space, September 2005.

VI. ASE Bibliography (as of July 2005)

2005

S. Chien, R. Sherwood, D. Tran, G. Rabideau, A. Davies, R. Castano, D. Mandl, S. Frye, B. Trout, J. D'Agostino, S. Shulman, D. Boyer, S. Hayden, A. Sweet, S. Christa, "Lessons Learned from Autonomous Sciencecraft Experiment," Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2005), Utrecht, The Netherlands, July 2005.

D. Tran, S. Chien, G. Rabideau, B. Cichy, "Preventing and Responding to Anomalies for the Autonomous Sciencecraft Experiment," Meeting on Safety and Security in Multi-Agent Systems, Utrecht, The Netherlands, July 2005.

D. Tran, G. Rabideau, B. Cichy, S. Chien, R. Sherwood, "Software Agents in Space: The EO-1 Autonomous Science Agent," Software Demonstration Track, Autonomous Agents and Multi-Agent Systems Conference, Utrecht, The Netherlands, July 2005.

D. Tran, G. Rabideau, B. Cichy, S. Chien, R. Sherwood, "Onboard Planning enables Autonomous Science: The EO-1 Autonomous Science Agent," Software Demonstration Track, International Conference on Automated Planning and Scheduling, Monterey, CA, June 2005.

B. Cichy, S. Chien, S. Schaffer, D. Tran, G. Rabideau, R. Sherwood, "Validating the Earth Observing One Autonomous Science Agent," Keynote Address by S. Chien, Meeting on Validation and Verification of Planning Systems, International Conference on Automated Planning and Scheduling, Monterey, CA, June 2005.

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau "Operational Lessons Learned from the Autonomous Sciencecraft Experiment," Proceedings of the 6th International Symposium Reducing the Costs of Spacecraft Ground Systems and Operations (RCSGSO), Darmstadt, Germany, June 2005.

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, & Communication, April 2005.

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau, "The ST6 Autonomous Sciencecraft Experiment", Proceedings of the 2005 IEEE Aerospace Conference, Big Sky, MT, March 2005.

Davies, A. G., S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castaňo, B. Cichy, G. Rabideau, D. Tran and R. Sherwood (2005) Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE), Remote Sensing of the Environment, to appear.

Doggett, T., R. Greeley, A. G. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm and F. Ip (2005) Autonomous On-board detection of Cryospheric Change, Remote Sensing of the Environment, to appear.

Ip, F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, B. Castaňo, S. Chien, B. Cichy, R. Greeley, R. Sherwood (2005) Development and Testing of the Autonomous Spacecraft Experiment (ASE) floodwater classifiers: Real-time Smart Reconnaissance of Transient Flooding, Remote Sensing of the Environment, to appear.

Davies, A. G., S. Chien, V. Baker, R. Castaño, B. Cichy, T. Doggett, J. M. Dohm, R. Greeley, F. Ip, G. Rabideau, R. Sherwood and D. Tran (2005) Maximizing Mission Science Return Through use of Spacecraft autonomy: active volcanism and the Autonomous Sciencecraft Experiment, Lunar and Planetary Sciences Conference XXXVI, abstract #1445. http://www.lpi.usra.edu/meetings/lpsc2005/pdf/1445.pdf

Chmielewski, A. B., R.M. Nelson, C.M. Stevens, S. A. Chien, A. G. Davies, R. L. Sherwood, W. Wyman (2005) NASA's New Millennium Program ST6 Project, Lunar and Planetary Sciences Conference XXXVI, abstract #1523. http://www.lpi.usra.edu/meetings/lpsc2005/pdf/1523.pdf

Ip F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, B. Castano, S. Chien, B. Cichy, R. Greeley, R. Sherwood (2005) Monitoring floods with NASA's Autonomous Sciencecraft Experiment, Lunar and Planetary Sciences Conference XXXVI, abstract #2263. http://www.lpi.usra.edu/meetings/lpsc2005/pdf/2263.pdf

2004

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau, "Safe Agents in Space: Lessons from the Autonomous Sciencecraft Experiment," Proceedings of the 17th Australian Joint Conference on Artificial Intelligence, Cairns, Australia, December 2004.

D. Mandl, R. Sherwood, S. Chien, et al., "Linking Satellites Via Earth “Hot Spots” and the Internet to Form Ad Hoc Constellations," Proceedings of the SPIE Remote Sensing of the Atmosphere, Ocean, Environment, and Space 2004, Honolulu, HI, November 2004.

R. Sherwood, S. Chien, A. Davies, D. Mandl, S. Frye, "Realtime Decision Making on EO-1 Using Onboard Science Analysis," Proceedings of the SPIE Remote Sensing of the Atmosphere, Ocean, Environment, and Space 2004, Honolulu, HI, November 2004.

S. Chien, R. Sherwood, et al., "Onboard Autonomy on the Earth Observing One Mission," Proceedings of the 1st AIAA Intelligent Systems Technical Conference, Chicago, IL, September 2004.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, R. Lee, D. Mandl, S. Frye, B. Trout, J. Hengemihle, J. D'Agostino, S. Shulman, S. Ungar, T. Brakke, D. Boyer, J. VanGaasbeck, R. Greeley, T. Doggett, V. Baker, J. Dohm, F. Ip., "The Earth Observing One Autonomous Science Agent," Proceedings of the Autonomous Agents and Multi-Agent Systems Conference, New York, NY, July 2004.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, R. Lee, D. Mandl, S. Frye, B. Trout, J. Hengemihle, J. D'Agostino, S. Shulman, S. Ungar, T. Brakke, D. Boyer, J. VanGaasbeck, R. Greeley, T. Doggett, V. Baker, J. Dohm, F. Ip., "The Earth Observing One Autonomous Science Agent," Proceedings of the Autonomous Agents and Multi-Agent Systems Conference, New York, NY, July 2004.

D. Tran, G. Rabideau, B. Cichy, S. Chien, R. Sherwood, “The EO-1 Autonomous Science Agent,” Software Demonstration Track, Autonomous Agents and Multi-Agent Systems Conference, New York, NY, July 2004.

D. Tran, G. Rabideau, B. Cichy, S. Chien, R. Sherwood, “An Autonomous Science Agent Onboard the Earth Observing One Mission,” Intelligent Systems Software Demonstration Track, National Conference on Artificial Intelligence (AAAI-2004), San Jose, CA, July 2004.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castaño, A. Davies, R. Lee, D. Mandl, S. Frye, B. Trout, J. Hengemihle, J. D'Agostino, S. Shulman, S. Ungar, T. Brakke, D. Boyer, J. Van Gaasbeck, R. Greeley, T. Doggett, V. Baker, J. Dohm, F. Ip, “The EO-1 Autonomous Science Agent Architecture,” International Workshop on Planning and Scheduling for Space, Darmstadt, Germany, June 2004.

B. Cichy, S. Chien, S. Schaffer, D. Tran, G. Rabideau, R. Sherwood, “Validating the Autonomous EO-1 Science Agent,” International Workshop on Planning and Scheduling for Space (IWPSS 2004). Darmstadt, Germany, June 2004.

D. Tran, S. Chien, G. Rabideau, B. Cichy, “Flight Software Issues in Onboard Automated Planning: Lessons Learned on EO-1,” International Workshop on Planning and Scheduling for Space (IWPSS 2004). Darmstadt, Germany, June 2004.

G. Rabideau, S. Chien, R. Sherwood, D. Tran, B. Cichy, D. Mandl, S. Frye, S. Shulman, R. Bote, J. Szwaxzkowski, D. Boyer, J.V. Gaasbeck, "Mission Operations with Autonomy: A preliminary report for Earth Observing-1", Proceedings of the 4th International Workshop on Planning and Scheduling for Space, Darmstadt, Germany, June 2004.

D. Tran, G. Rabideau, B. Cichy, S. Chien, R. Sherwood, “Onboard Planning for Autonomous Science: The EO-1 Autonomous Science Agent,” Software Demonstration Track, International Conference on Automated Planning and Scheduling, Whistler, BC, Canada, June 2004.

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau, "Operating the Autonomous Sciencecraft Experiment", Proceedings of the SpaceOps 2004, Montreal, Canada, May 2004.

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau, "Preliminary Results of the Autonomous Sciencecraft Experiment", Proceedings of the 2004 IEEE Aerospace Conference, Big Sky, MT, March 2004.

Davies, A. G., V. Baker, R. Castaño, S. Chien, B. Cichy, T. Doggett, J. M. Dohm, R. Greeley, R. Lee, R. Sherwood, Autonomous Sciencecraft Experiment (ASE) Operations On EO-1, Lunar and Planetary Sciences Conference XXXV (2004), abstract #1700

Doggett, T.C., R. Greeley, V. Baker, S. Chien, A.G. Davies, J.M. Dohm, T. Ferre, A. Hinnel, D. Rucker, and K. Williams, Airborne Radar Study of Soil Moisture at a Mars Analog Site: Tohachi Wash Little Colorado River, Lunar and Planetary Sciences Conference XXXV (2004), abstract #1326

Ip, F., J. M. Dohm, V. R. Baker, T. Doggett, A. G. Davies, R. Castano, S. Chien, B. Cichy, R. Greeley and R. Sherwood, ASE Floodwater Classifier Development For EO-1 Deployment, Lunar and Planetary Sciences Conference XXXV (2004), abstract #2142

Rucker, D. F., J.M. Dohm, T.P.A. Ferré, Felipe Ip, V.R. Baker, A.G. Davies, R. Castano, S. Chien, B. Cichy, T.C. Doggett, R. Greeley, and R. Sherwood, Avra Valley Storage and Recovery Project (CAVSARP) Site, Tucson, Arizona: Floodwater and soil moisture investigations with extraterrestrial applications, Lunar and Planetary Sciences Conference XXXV (2004), abstract #2114

Chien, S, A. G. Davies, R. Wright, A. Miklius, L. Flynn, B. Cichy, S. Frye, S. Shulman, D. Mandl, S. Grosvenor (2004) Streamlining Spacecraft Observation Response to Volcanic Activity Detection Using An Autonomous Sensor Web of Ground and Space-Based Assets, Eos. Trans., AGU 85(47), Fall Meeting. Suppl., Abstract V33C-1469

Ip, F., J. M. Dohm, V. R. Baker, R. Castano, B. Cichy, S. Chien, A. G. Davies, T. Doggett and R. Greeley (2004) Successful Detection of Floods in Real Time Onboard EO1 Through NASA's ST6 Autonomous Sciencecraft Experiment (ASE) Eos. Trans., AGU 85(47), Fall Meeting. Suppl., Abstract H23D-1162.

Davies, A. G., S. Chien, R. Wright, A. Miklius, L. Flynn, V. Baker, R. Castano, B. Cichy, J. Dohm, T. Doggett, R. Greeley, F. Ip, R. Sherwood, S. Frye, M. Pell, and the ASE and EO-1 Teams (2004) Rapid Response to Volcanic Activity using Autonomous Systems in an Integrated d Sensor Web, abstract, IAVCEI-GA 2004, Pucon, Chile.

Chmielewski, A. B., R.M. Nelson, C.M. Stevens, S. A. Chien, A. G. Davies, R. L. Sherwood, W. Wyman (2005) NASA's New Millennium Program ST6 Project, Bull. Amer. Astron. Soc. Vol 36 no 4, p S140. http://www.aas.org/publications/baas/v36n4/dps2004/511.htm

Davies, A. G., V. Baker, R. Castano, S. Chien, B. Cichy, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. J. Lee, G. Rabideau, and R. Sherwood (2004) The Autonomous Sciencecraft Experiment: Monitoring of Geological and Environmental Hazards With An Autonomous, Intelligent Spacecraft, International Geophysical Congress 32, Florence, Italy, abstract 247-8.

Lee, R. J., A. G. Davies, and the ASE Science and Flight Teams (2004) Autonomous Vegetation Cover Scene Classification of EO-1 Hyperion Hyperspectral Data, Lunar and Planetary Sciences Conference XXXV, abstract #1615
http://www.lpi.usra.edu/meetings/lpsc2004/pdf/1615.pdf

2003

R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, G. Rabideau, "Next Generation Autonomous Operations on a Current Generation Satellite", Proceedings of the 5th International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations (RCSGSO), Pasadena, CA, July 2003.

B. Cichy, S. Chien, "Validating the Autonomous EO-1 Science Agent," Keynote address, Meeting on Safe Agents, International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS03), Melbourne, Australia, July 2003.

R. Sherwood, S. Chien, D. Tran, R. Castano, B. Cichy, A. Davies, G. Rabideau, N. Tang, M. Burl, D. Mandl, S. Frye, J. Hengemihle, J. D'Augustino, R. Bote, B. Trout, S. Shulman, S. Ungar, J. Van Gaasbeck, D. Boyer, M. Griffin, H. Burke, R. Greeley, T. Doggett, K. Williams, V. Baker, J. Dohm, "DEMO: Autonomous Science Analysis, Planning, and Execution on the EO-1 Mission", Proceedings of the 13th International Conference on Automated Planning and Scheduling (ICAPS), Trento, Italy, June 2003.

S. Chien, R. Sherwood, D. Tran, R. Castano, B. Cichy, A. Davies, G. Rabideau, N. Tang, M. Burl, D. Mandl, S. Frye, J. Hengemihle, J. D'Augustino, R. Bote, B. Trout, S. Shulman, S. Ungar, J. Van Gaasbeck, D. Boyer, M. Griffin, H. Burke, R. Greeley, T. Doggett, K. Williams, V. Baker, J. Dohm, "Autonomous Science on the EO-1 Mission", Proceedings of the 7th International Symposium on Artificial Intelligence (iSAIRAS), Robotics and Automation in Space, Nara, Japan, May 2003.

R. Sherwood, S. Chien, R.Castano, G. Rabideau, "The Autonomous Sciencecraft Experiment," Proceedings of the IEEE Aerospace Conference, Big Sky, MT, March 2003.

Davies, A. G., S. Chien, R. Wright, P. Cervelli, L. Flynn, V. Baker, R. Castano, B. Cichy, J. Dohm, T. Doggett, R. Greeley, R. Sherwood, K. Williams, S. Frye, J. Jones, and the ASE and EO-1 Mission Teams, Streamlining Spacecraft Observation Response to Volcanic Activity Detection with a Ground and Space-based Sensor Web System, Abstract, AGU Fall Meeting 2003, San Francisco, CA

Chien, S., A. G. Davies, V. Baker, B. Castano, B. Cichy, T. Doggett, J. M. Dohm, R. Greeley, R. Sherwood and K. Williams, Autonomous Sciencecraft Experiment (ASE) test operations in 2003, Lunar and Planetary Sciences Conference XXXIV, abstract #1458, on CD-ROM

Dohm, J.M., S. Chien, G. Brakenridge, V. Baker, R. Castano, B. Caquard, A. Davies, T. Doggett, R. Greeley, S. Nghiem, R. Sherwood, K. Williams, D. Mandl, S. Ungar, S. Frye, J. Jones and S. Grosvenor, Ground and Space-based Sensor Web System: Streamlining Spacecraft Observation Response to Flood Detection, Eos Transactions American Geophysical Union, 84(46), Fall Meet. Suppl., Abstract H22D-0967, 2003

Davies, A. G., V. Baker, R. Castano, S. Chien, B. Cichy, T. Doggett, J. Dohm, R. Greeley, G. Rabideau, R. Sherwood, K. K. Williams, and the ASE Team, (2003) Autonomous Volcanic Activity Detection with ASE on EO-1 Hyperion: Applications for Planetary Missions, Abstract, DPS 35, Monterey, CA, BAAS vol 35, no 4, 41.14.

2002

S. Chien, R. Sherwood, G. Rabideau, R.Castano, A. Davies, M. Burl, R. Knight, T. Stough, J. Roden, P. Zetocha, R. Wainwright, P. Klupar, J. Van Gaasbeck , P. Cappelaere, D. Oswald, "The Techsat-21 Autonomous Science Agent," Proceedings of Agents 2002, Bologna, Italy, July 2002.

R. Sherwood, S. Chien, R. Castano, G. Rabideau, "Autonomous Operations through Onboard Artificial Intelligence", Proceedings of the SpaceOps 2002 Conference, Houston, TX, October 2002.

R. Sherwood, S. Chien, R. Castano, G. Rabideau, "Spacecraft Autonomy using Onboard Processing for a SAR Constellation Mission", Proceedings of the First International Workshop on Future Intelligent Earth Observing Satellites (FIEOS) at the (ISPRS) Commission 1 Symposium, Denver, CO, November 2002.

R. Sherwood, S. Chien, R. Castano, G. Rabideau, "Autonomous Planning and Scheduling on the TechSat 21 Mission", Proceedings of the Fifteenth Australian Joint Conference on Artificial Intelligence, Canberra, Australia, December 2002.

Doggett, T., A. G. Davies, R. Castano, V. Baker, J. Dohm, R. Greeley, K. Williams, S. Chien and R. Sherwood, “Autonomous Science Analysis with the NMP-Autonomous Sciencecraft Experiment. American Geophysical Union Fall Meeting 2002 abstract. San Francisco, CA. Dec. 2002.

Davies, A. G., R. Greeley, V. Baker, K. K. Williams, J. Dohm, R. Castano, E. Mjolsness, J. Roden, T. Stough, S. Chien, R. Sherwood, P. Zetocha, “Automating the operational and science analysis processes with the Autonomous Sciencecraft Constellation,” Lunar and Planetary Sciences Conference XXXIII, 2002.

2001

S. Chien, R. Sherwood, M. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, P. Zetocha, R. Wainwright, P. Klupar, P. Cappelaere, D. Surka, B. Williams, R. Greeley, V. Baker, J. Dohm, "The Techsat-21 Autonomous Sciencecraft Constellation," Proceedings of the 6th International Symposium on Artificial Intelligence, Robotics, & Automation in Space (ISAIRAS-2001), Montreal, Canada, June 2001.

S. Chien, R. Sherwood, M. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, R. Castano, T. Stough, J. Roden, P. Zetocha, R. Wainwright, P. Klupar, P. Cappelaere, J. Van Gaasbeck, D. Surka, M. Brito, B. Williams, M. Ingham, "A Demonstration of Robust Planning, Scheduling and Execution for the Techsat-21 Autonomous Sciencecraft Constellation," Proceedings of the 2001 European Conference on Planning (ECP2001), Toledo, Spain, September 2001.

I. Appendices

A. Experimental Results from Flight and Ground Tests

Date	Flight Experiment Type	Results
3/22/2003	First test of onboard cloud detection	Successful
5/23/2003	First test of onboard SCL execution software	Successful
7/21/2003	First test of automated ground developed plans running onboard EO1 (Yukon River)	Successful
7/27/2003	Second test of automated ground developed plans running onboard EO1 (Fire sensorweb)	Successful
7/31/2003	Third test of automated ground developed plans running onboard EO1 (Ganda Angola sensorweb)	Successful
8/13/2003	Test of decompression software onboard	Successfully decompressed file and downlinked
8/13/2003	Fourth test of automated ground developed plans running onboard EO1 (Montana fire sensorweb)	Successful
10/29/2003	Planner commanded the instrument (hyperion) covers to open/close, then to perform a dark calibration, then an X-band downlink (R-1 version)	Successful
11/19/2003	Planner commanded the instrument (hyperion) to acquire an image (R-1 version)	The WARP CPU performed a reset due to a critical task not running because of CPU over-utilization
1/08/2004	Planner commanded the instrument (hyperion) to acquire an image, then downlink that image (R-1 version)	Successful, 10% validation achieved
1/14/2004	Planner commanded the instrument (hyperion) to acquire an image, then downlink that image (R-1 version), 20% validation achieved	Successful, 20% validation achieved
1/22/2004	Planner commanded the instrument (hyperion) to acquire an image of Great Sand Dunes, then downlink that image (R-1 version)	Successful, 30% validation achieved
1/22/2004	Planner commanded the instrument (hyperion) to acquire an image of Lake Monona (Wisconsin), then downlink that image (R-1 version)	Successful, 40% validation achieved
1/29/2004	Planner commanded the instrument (hyperion) to acquire an image of Kokee State Park, Kauai, Hawaii, then downlink that image (R-1 version)	Successful, 50% validation achieved
2/20/2004	Planner commanded the instrument (hyperion) to acquire an image of Winnibigoshish Leech Lakes then downlink that image (R-1 version), (sensorweb, ASPEN on ground chose which image based on GOES cloud predict)	Successful, ground in the loop retargeting of spacecraft
2/25/2004	Planner commanded the instrument (hyperion) to acquire an image of Upper/Lower Red Lakes then downlink that image (R-1 version), (sensorweb, ASPEN on ground chose which image based on GOES cloud predict)	Successful, ground in the loop retargeting of spacecraft
3/17/2004 55%	Planner commanded the instrument (hyperion) to acquire an image of Great Sand Dunes, strip the relevant bands from the image, then run science analysis on the image (cloud detection), then retarget to Port au Prince based on the science analysis	Retarget successful (5%)
3/25/04 60%	Target: Lake Mendota (ice freeze) Retarget: Great Sand Dunes Algorithm(s): SWIL (retarget), thermal (file delete) (no hot spots)	Retarget unsuccessful (0%), (completely cloudy image) File delete successful (5%)
4/1/04 60%	Target: Prudhoe Bay (sea ice) Retarget: Barrow Lake Algorithm(s): SWIL (retarget), thermal (file delete) (no hot spots)	Failed, band stripping task timed out early
4/7/04 60%	Target: Stromboli Night (volcano) Retarget: Stromboli Night or Kokee Algorithm(s): thermal (retarget), thermal (file delete)	Failed - Warm reset of WARP processor occurred during 2nd thermal detection
4/29/04 70%	Target: Resolute Bay (sea ice) Retarget: Ward Hunt Ice Shelf Algorithm(s): SWIL (retarget), thermal (file delete), thermal (force retarget)	Retarget successful (5%), File delete successful (5%)

5/7/04 80%	Target: Erebus Night (volcano) Retarget: Either Erebus or Stromboli Night Algorithm(s): thermal (retarget), thermal (file delete)	Retarget successful (5%), File delete successful (5%)
5/12/04 90%	Target: Lake Monona (flooding) Retarget: Lake Monona Algorithm(s): thermal (file delete), SWIL (retarget)	Retarget successful (5%), File delete successful (5%)
5/14/04 100%	Target: Erebus Night (volcano) Retarget: Erebus daytime Algorithm(s): thermal (retarget), thermal (file delete)	Retarget successful (5%), File delete successful (5%)

Appendix B

The detailed listings of sequence of EO-1 spacecraft commands are removed due to ITAR/Export considerations. For further inquiries regarding this information (e.g. for non export disclosures) please contact the Technology Provider.

Originally included sequences included:

To command an autonomous data collect:
To command an S-band/X-band downlink

Appendix C: Details of Software Anomalies

We list below several of the ASE anomalies to date, classify each type, and described steps taken to resolve it.

1) *September 20, 2003 – Software anomaly*

In the first attempt to upload and test the ASE software, a benign command was issued to each subsystem on EO-1 in order to test the modifications to the routing tables. Originally, spacecraft commands originated from the C&DH processor. When using ASE, commands originate from the WARP processor. As a result, the software routing table needed to be updated. All commands were routed correctly except for the command to the ALI instrument, which was being rejected. After verifying the specific bytes generated for the command against a ground issued command, it was determined that several field of the commands are not set correctly.

Due to the limitations of the ground testbed, an actual ALI instrument had not been available for testing. Instead, we relied on a simulation of the instrument. The simulator had checked various fields on the commands, but a specific field had not been checked. Once it was determined this field was required by the ALI instrument software, the ASE software along with the simulation were updated.

2) *November 19, 2003 – Software anomaly*

In the middle of the data collect, a WARP reset occurred and the flight software restarted while it had been in record mode. This caused a TSM to trip on the C&DH, powering off the WARP causing the lost of code, reset log, and all other ASE logs. We were unsuccessful in replicating this problem on the testbed, but after inspecting the code, it was noticed that one of the flight software tasks, memory scrub was designated as critical. If this task does not receive the CPU every 8 seconds, it will cause a reset. The insertion of the ASE tasks placed one task at a higher priority than the memory scrub task. It was this task, along with the CPU intensive data collect, that caused the memory scrub task to be starved, resulting in the reset. After consulting with the EO-1 flight software engineers, it was decided that the memory scrub task did not need to be designated as critical, and the EO-1 flight software was updated.

3) *April 1, 2004 – Software anomaly*

During this early stage of the project, we were testing a single response scenario where the onboard science module would analyze an image, and issue requests to the onboard planner for more images of that target. The scenario went as follows:

- Image Prudhoe Bay, Alaska
- Playback data from the solid state recorder into RAM for image processing (band-stripping)
- Run image classifier to identify snow, water, ice, and land. This was scheduled to run immediately after the band-stripping process.
- Run image classifier to identify for cloud coverage. This was scheduled to run 60 minutes after the band-stripping process.
- If the classification of the scene was above a threshold defined by mission scientists, request additional images of Prudhoe Bay.

Several of the constraints modeled within CASPER and SCL were:

- The band-stripping process cannot begin unless the target Prudhoe Bay was successfully imaged
- The image classifiers cannot begin unless the band-stripping process was successful

During the first ground contact following this scenario, mission operators noticed several warnings from SCL telemetry and that the onboard science module did not perform any image processing. After collecting log files from SCL and CASPER, and replaying back telemetry collected during the test, it was determined that SCL had failed the band-stripping script because of a time out during the verification of the command completion. In actuality, this verification failure was not a result of band-stripping failing, but of a bug within the flight software time code. It is still interesting, however, to examine the response.

The failure of the band-stripping script resulted in a lack of change to a SCL database record. This record is continuously monitored by CASPER and a conflict with the scheduled image classifier algorithm was recognized. However, because the first image classifier algorithm was scheduled immediately after band stripping, CASPER had already committed to executing the classifier activity. When making this type of commitment, CASPER locks the activity (preventing any re-scheduling) and sends the execution request to SCL. The command was received by SCL, but failed the pre-requisite check, blocking the command from being sent to the science module.

The second image classifier was scheduled 60 minutes after the end of band stripping, and thus CASPER was able to modify the plan to resolve the conflict by removing it from the schedule.

This anomaly demonstrated how the layered architecture ensured the safety of the agent. CASPER was not responsive enough to prevent the first image classifier from being issued to SCL, but the SCL pre-requisite check failed and thus the command was not issued. However in the second case, CASPER was able to respond to this failure by removing the subsequent image processing activity from the schedule.

One possible modification of the architecture to prevent these false-positive anomalies from occurring would be to have redundant checks in the completion of the commands. In this example, a single SCL database item indicated that band stripping had failed when in fact it had succeeded. The model could have been updated to check multiple database records for the status of the band stripping, instead of relying on solely on a single data point to verify completion of the command.

4) *April 7, 2004 – Software anomaly*

In this anomaly, which caused the ASE software to be down for several days, a ground command to clear several command counters was issued that caused the WARP processor to reset. After careful analysis of the code, it was determined a buffer overrun of the flight software bus routing table had occurred, caused by the increase number of commands originating from the WARP processor.

When the original WARP flight software was developed, it was not intended to issue the large number of commands that was occurring during with ASE hosted onboard, leading them to specify an undersized buffer. The reason this problem was not caught during ground testing is that the WARP processor was not utilized in the same manner as in flight. During these flight tests, numerous commands to various subsystems are issued from the ground operators for normal operations and testing. This type of testing had not been performed on the ground and so the buffer overrun was not detected. Our lesson learned from this experience is during operations testing to interact with the testbed as closely as possible to what is expected to be done in flight.

5) *July 14, 2004 – Model anomaly*

This anomaly demonstrates how SCL was able to respond to a verification failure of command sequence. During this test, the anomaly occurred during normal operations for an X-Band ground contact. The scenario was:

- Using the X-Band transmitter, downlink all images from the solid state recorder
- Clear all images from the solid state recorder

Several of the constraints modeled were:

- The correct voltage/current level of transceiver must be met prior to operating X-Band activities.
- The downlink must complete successfully prior to clearing all the images from the solid-state recorder.

During the ground contact, mission operators noticed several warnings from SCL and also that EO-1 had not begun the X-Band downlink of images collected. The operators manually initiated the X-Band contact and completed dumping the data. After analyzing log files, it was determined that a prerequisite failure in the SCL model for the current/voltage of the transceiver prevented the contact from being initiated. As a result of the X-Band failure, SCL also rejected the command to clear all the images from the solid-state recorder.

This was actually an error within the SCL model. An early version of the model included a constraint that the transceiver cannot be powered on unless the current/voltage was at the correct level. However, the threshold values for the current/voltage in reality are not valid until the transceiver is powered on.

Unfortunately, this modeling error slipped through our testing and validation process because of the lack of a high fidelity testbed. The EO-1 testbed did not have a transceiver for testing and therefore, the current/voltage values were static (at the "on" levels) in the simulation. Without valid values on the current/voltage prior to powering on the X-Band transceiver, our resolution to this problem was to simply remove the current/voltage constraint from the SCL model.

6) August 11, 2004 – Modeling anomaly

After analyzing the command sequence for an ASE-commanded dark calibration, it was detected that the command to stop the WARP recorder (WRMEREC) was issued a couple of seconds prior to the instruments stop. The normal sequence is for the instruments to stop collecting stop, followed by the stop record command for the WARP.

A constraint in SCL was modeled to not issue the WARP stop record command if the instruments are still collecting data. However, because the WRMEREC command was scheduled to begin too close to the instrument stop command, the SCL database did not update the state of the instruments, and the WRMEREC command was allowed to be issued. This problem was resolved in the CASPER model by scheduling the WRMEREC later in the data collect sequence, ensuring that enough time has elapsed for the instruments to actually stop collecting data.

7) November 15, 2004 – Operator anomaly

This anomaly demonstrates CASPER robustness to malformed input requests. Goal files are uploaded daily to EO-1, containing the scene requests and contacts for the next day. However, they require being uploaded to RAM to a staging area, prior to being copied into the onboard ramdisk. Unfortunately, the procedure generated to upload to RAM and copy to the ramdisk specified different memory addresses, and the daily goal file loaded into CASPER contained random memory values. CASPER's response to the bad goal file was to reload the current plan back into the schedule, and removing any scenes that may have been inserted into the schedule from the bad goal file.

The current procedure to upload the goal file into the ramdisk required two procedures to be generated by mission planners, one to upload the goal file to the staging area, and the other to copy the memory to the ramdisk. However, this was prone to errors as previously describe. Our resolution to prevent this anomaly from occurring in the future is to combine the two procedures into one, and have the procedures auto-generated to ensure the memory addresses are consistent.

8) December 1, 2004 – Operator anomaly

In this anomaly, we lost a last minute replacement scene because of a supporting activity for a scene was not loaded into the schedule. During the goal loading process, any goals that start within 40 minutes of the current time are filtered out and not inserted into the schedule. This reasoning behind this is that there needs to be enough time for the CASPER planner to try and resolve any conflicts near-term activities may cause to the schedule.

When this supporting activity was filtered out the schedule, it caused the resulting scene to be in conflict and thus it was removed from the schedule. Our resolution to this problem was corrected operationally. Each goal file now contains the latest uplink time to the CASPER planner. If that latest uplink is not met on the ground, the entire goal file is rejected and not uploaded.

9) December 6, 2004 – Operator anomaly

For every scene imaged on EO-1, a base WARP file-id is used during recording on the solid-state recorder. This base file-id, must be unique until all scenes on the WARP are cleared. Mission planners do the assignment of a WARP base-id manually, and the daily reports received from GSFC contained two consecutive scenes with the same base WARP id. As a result, during execution of the second scene, the WARP flight software and SCL flagged several errors and warnings, indicating that it could not record the scene.

Representing this constraint in ASPEN requires the use of the Generalized Timeline module. However, this capability isn't part of the core ASPEN that has been used in many applications. To reduce the CASPER image size and reduce the risk, the Generalized Timeline module was not uploaded to EO-1. Instead, to ensure this problem does not occur in the future, we simply updated the script that converts the GSFC report to a CASPER goal file to check and verify that duplicate WARP base-ids are not assigned.

10) January 31, 2005 – Modeling anomaly

This anomaly describes CASPER's response to an unexpected change in the state of the spacecraft. During one of the scheduled ground contacts, the agent did not initiate the command sequence as requested from mission planners. An anomaly had occurred that removed the contact sequence from the mission plan. After analysis of collected telemetry, the cause of the anomaly was due to human intervention with the spacecraft several hours prior. Mission planners had initiated an unscheduled contact, which was performed externally from the onboard planner. The unscheduled contact required mission operators to perform a blind acquisition of EO-1 and manually power on the transceiver, changing the state of the onboard transceiver to "on". At the end of this contact, the operators manually powered down the transceiver.

The change to the transceiver state resulted in an update to the SCL database, which propagated to the CASPER schedule and created a conflict with the next ground contact activity. The conflict was with a constraint in the CASPER model that only allowed the transceiver state to transition from "on" to "off" or from "off" to "on". When the update to the transceiver state was received, it set the current state to the transceiver to "on". This created a conflict with the next scheduled contact that had planned to turn the transceiver on when the state was already "on". To resolve this conflict, CASPER removed the next contact from the plan. Once the mission operator powered down the transceiver at the end of the unscheduled contact, subsequent contacts were conflict free, but the deleted contact was not rescheduled due to the risk of inserting the goal too close to its scheduled time.

To prevent this anomaly for future operations, we simply removed the transition constraints from the CASPER model of the transceiver. While not ideal, it was determined that this presented no risk to the spacecraft, and allowed the ASE software to support real-time contact requests from mission planners without affecting the remainder of the schedule. In this anomaly, although the update to the state of the transceiver was short-lived as mission operators eventually powered it off, its affect on the planner propagated to the next scheduled contact, resulting in the contact being removed from the schedule. One possible solution to prevent this from occurring in the future is to delay resolving conflict until necessary. Some initial work has been started on CASPER to support time-sensitive repair heuristics, but is still experimental and was not deployed on EO-1.

11) February 7, 2005 – Software anomaly

In this anomaly, we had the WARP processor unexpectedly reset during nominal operations where no activities were scheduled for execution. After analyzing the activities leading up to the reset, we determined the cause to be an incorrect specification for the location of the thermal summary product. The location of the thermal summary product overwrote some of the data required by CASPER and after execution of the thermal summary algorithm, the software crashed.

Though the thermal summary algorithm was tested on the ground, due to testbed limitations, this error condition was not captured. The ground testbed does not contain a solid-state recorder and so valid science data was not available.

Our correction to the anomaly was to patch the location of the thermal summary product.

12) February 21, 2005 - Modeling anomaly

In this anomaly, the CASPER planner was continually deleting all scenes in the schedule. As a result, 11 scenes were lost. It was first recognized by mission operators when during ground contacts, no scenes were being downlinked as expected. At the time, we were unable to determine what was causing the scenes to be removed from the planner, and stopped the ASE software in order to troubleshoot the problem.

After replaying back some of the telemetry from the CASPER planner, it was determined that the problem was caused while CASPER was attempting to de-conflict two overlapping scenes. As a result, it placed the spacecraft into an off-nominal state that was not automatically correctable. This off-nominal state created conflicts with future scenes and they were deleted as result. Once we were able to determine the cause of the problem, a patch of was developed for the scheduling algorithm onboard to prevent this problem for occurring in the future.

Several other methods were considered in order to recover from this anomaly, but were not implemented due to the large change it would cause to the overall system. One method of recovery would be for another software module that would recognize the EO-1 was in an off-nominal state for a long period of time, and issue a request to CASPER to execute an activity to place it back into a nominal condition. This would eliminate any conflicts with future scenes and CASPER can begin scheduling scenes again.

13) May 2, 2005 – Software anomaly

This anomaly was first recognized when the ASE software did not initiate a scheduled ground contact. When ASE did not initiate the next scheduled contact as well, mission operators manually initiated the contact. When spacecraft telemetry was received, each subsystem was outputting telemetry values except for CASPER, which no longer was active in between the last contacts. Mission operators initiated a shutdown of the ASE software in order to determine the cause of why the CASPER task was no longer active.

After replaying back commands from the previous contacts, we determined that the goal file had been deleted from the ramdisk onboard, while CASPER was processing that file into its schedule. The procedure of removing files from the ramdisk was a new addition of the ground software automation that had been just been released.

This condition had not been tested on the ground prior to flight. Unfortunately, the software code to handle a loss of file was third party generated code and not developed for flight. The generated code was an integral part of the loading of the goal file and could not be patched to handle to error condition gracefully. Instead we updated the ground operations procedure to ensure this condition does not occur in the future.

Several lessons learned from this anomaly is all software modules onboard need to be robust enough to all types of commands that may originate from the mission operators and also any third-party generated software needs to be scrutinized to ensure it is ready for flight.

14) June 19, 2005 – Modeling anomaly

In this anomaly, as a result of the attempting to de-conflict two overlapping scenes, EO-1 was placed in an off-nominal condition when an activity to open the instruments covers were committed for execution, without its corresponding close cover activity. As a result of this off-nominal state, future scenes were in conflict and CASPER began removing the scenes from the schedule. However, associated with each scene is the command to place EO-1 back to earth pointing. This activity was not removed from the schedule and was committed to SCL for execution.

One of the pre-conditions required to execute this command was that the instruments cover must be closed prior to maneuvering the spacecraft. The SCL script that executes the earth pointing command recognized that the instruments covers were still open, and automatically issued the close cover command. Once the covers were closed, this satisfied the constraints for the earth pointing activity and SCL issued the command. With the instruments covers closed and the spacecraft earth pointing, this placed EO-1 in a nominal condition and future scenes could now be scheduled conflict free.

To resolve the issue of CASPER de-conflicting two overlapping scenes, a software patch to the scheduling algorithms was uploaded to prevent this condition for occurring in the future. This anomaly demonstrates how, if modeled correctly, SCL's ability to monitor the state of the spacecraft, and respond and recover from localized off-nominal situations.