

Preliminary Results of the Autonomous Sciencecraft Experiment

Rob Sherwood, Steve Chien, Daniel Tran, Benjamin Cichy, Rebecca Castano, Ashley Davies, Gregg Rabideau
 Jet Propulsion Laboratory, California Institute of Technology
 4800 Oak Grove Dr.
 Pasadena, CA 91109
 Rob.Sherwood@jpl.nasa.gov

Abstract—The Autonomous Sciencecraft Experiment (ASE) will operate onboard the Earth Orbiter 1 mission in 2004. The ASE software uses onboard continuous planning, robust task and goal-based execution, and onboard machine learning and pattern recognition to radically increase science return by enabling intelligent downlink selection and autonomous retargeting. In this paper we discuss how these AI technologies are synergistically integrated in multi-layer control architecture to enable a virtual spacecraft science agent. We will also present the preliminary results from flight validation of this experiment. This software will demonstrate 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 will be simplified, thus reducing operations cost.¹

sand dune migration, growth and retreat of ice caps, cloud detection, and crust deformation. These onboard science algorithms are inputs to onboard decision-making algorithms that will modify the spacecraft observation plan to capture high value science events. This new observation plan will then be 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. This paper describes the specifics of the ASE and relates it to past and future flights to validate and mature this technology.

The ASE onboard flight software includes several autonomy software components:

- Onboard science algorithms that will analyze the image data to detect trigger conditions such as science events, “interesting” features, changes relative to previous observations, and cloud detection for onboard image masking
- 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 will replan activities, including downlink, based on science observations in the previous orbit cycles

The onboard science algorithms will analyze the images to extract static features and detect changes relative to previous observations. Prototype software has already been demonstrated on 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 will enable 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

TABLE OF CONTENTS

1. INTRODUCTION	1
2. THE EO-1 MISSION	2
3. AUTONOMY SOFTWARE ARCHITECTURE	3
4. ONBOARD SCIENCE ANALYSIS	3
5. ONBOARD MISSION PLANNING	5
6. ONBOARD ROBUST EXECUTION	5
7. FLIGHT STATUS	6
8. CONTRIBUTION TO FUTURE MISSIONS	6
9. IMPACT ON OPERATIONS	7
10. RELATED WORK & SUMMARY	7
REFERENCES	8
BIBLIOGRAPHY	9
ACKNOWLEDGEMENT	9

1. INTRODUCTION

In 2003, the ASE running on the EO-1 spacecraft will demonstrate several integrated autonomy technologies to enable autonomous science. Several science algorithms including: onboard event detection, feature detection, change detection, and unusualness detection will be used to analyze science data. These algorithms will be used to downlink science data only on change, and will detect features of scientific interest such as volcanic eruptions,

¹ 0-7803-8155-6/04/\$17.00 © 2004 IEEE

² IEEEAC paper #1551, Version 3, Updated December 16, 2003

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) will generate mission operations plans from goals provided by the onboard science analysis module. The model-based planning algorithms will enable 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 will accept 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 will be 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.

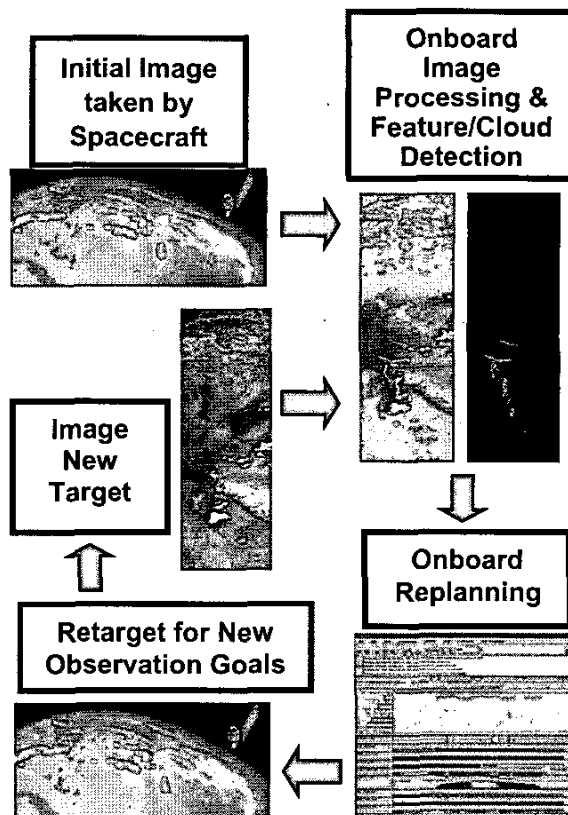


Figure 1. Autonomous Science Mission Concept

2. THE EO-1 MISSION

Earth Observing-1 (EO-1) is the first satellite in NASA's New Millennium Program Earth Observing series. The primary focus of EO-1 is to develop and test a set of advanced technology land imaging instruments.

EO-1 was launched on a Delta 7320 from Vandenberg Air Force Base on November 21, 2000. It was inserted into a 705 km circular, sun-synchronous orbit at a 98.7 degrees inclination. This orbit allows for 16-day repeat tracks, with 3 over flights per 16-day cycle with a less than 10-degree change in viewing angle.

For each scene, over 20-Gbits of data from the Advanced Land Imager (ALI), Hyperion, and Atmospheric Corrector (AC) are collected and stored on the onboard solid-state data recorder at high rates.

EO-1 is currently in extended mission, having more than achieved its original technology validation goals. As an example, over 5,000 data collection events have been successfully completed, against original success criteria of 1,000 data collection events.

The ASE described in this paper uses the Hyperion hyper spectral instrument. The Hyperion is a high-resolution imager capable of resolving 220 spectral bands (from 0.4 to 2.5 μm) with a 30-meter spatial resolution. The instrument

images a 7.5 km by 42 km land area per image and provides detailed spectral mapping across all 220 channels with high radiometric accuracy.

The EO-1 spacecraft has two Mongoose M5 processors. The first M5 is used for the EO-1 command and data handling functions. The other M5 is part of the WARP (Wideband Advanced Recorder Processor), a large mass storage device. Each M5 runs at 12 MHz (for ~8 MIPS) and has 256 MB RAM. Both M5's run the VxWorks operating system. The ASE software operates on the WARP M5. This provides an added level of safety for the spacecraft since the ASE software does not run on the main spacecraft processor.

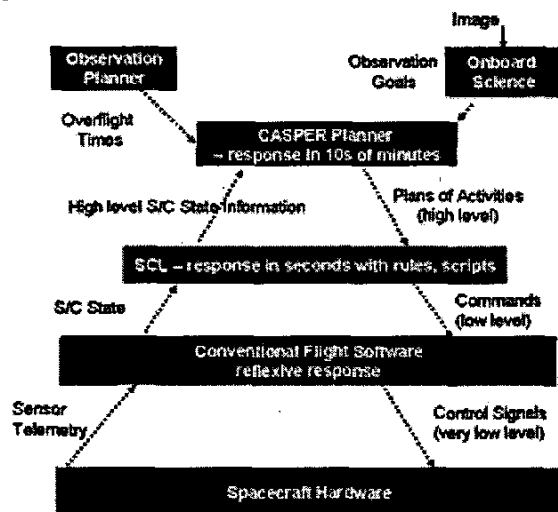


Figure 2. Autonomy Software Architecture

3. AUTONOMY SOFTWARE ARCHITECTURE

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.




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

4. 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 will be flown including:

- Thermal anomaly detection – uses infrared spectra peaks to detect lava flows and other volcanic activity. (See Figure 3.)
- Cloud detection [17] – uses intensities at six different spectra and thresholds to identify likely clouds in scenes. (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.
- Change detection – uses multiple spectra to identify regions changed from one image to another. This technique is applicable to many science phenomena including lava flows, flooding, freezing and thawing and is used in conjunction with cloud detection. (See Figure 5.)
- Generalized Feature detection – uses trainable recognizers to detect such features as sand dunes and wind streaks.

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

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.

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 will use 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 uses 6 spectral bands to characterize an image. These bands include two visible channels, a near-IR channel, and three shortwave infrared channels. These channels were chosen to provide enough information to analyze images while keeping processing costs to a minimum. The HCC algorithm can be 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 reimagine 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 will 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.




Figure 5. Change Detection Scenes indicating Ice Breakup in the Larsen Ice Shelf, Antarctica.

The University of Arizona developed flood scene classification algorithm 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.

Later flights will validate as many science analysis algorithms as resources allow. These flights will begin by validating change detection on multiple science phenomena, feature detection on Aeolian (wind) features such as sand dunes, sand shapes, and wind streaks, and the Discovery algorithm. Validating this portfolio of science algorithms will represent a valuable step forward to enabling future autonomous science missions [6].

5. 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 will modify 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 will initially perform management of 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. In the proposed follow-on experiment we will examine the possibility of migrating this function onboard.

6. 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 will also be 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). SCL will also verify that there is sufficient power so that the command does not trigger a low bus voltage condition and that there is sufficient energy in the battery. Using SCL to check these constraints (while included in the CASPER model) provides an additional level of safety to the autonomy flight software.

7. FLIGHT STATUS

The ASE software was integrated under the flight version of VxWorks in December 2002, and has been undergoing testing and integration with the WARP flight software. We are testing the individual software components to gain confidence before we perform an integrated flight test.

The cloud detection algorithms were tested onboard in March 2003. The SCL software was tested onboard in May 2003. This test involved starting up the SCL software, testing the software bridge between the SCL software bus and WARP software bus, testing the SCL message and telemetry logs, testing sending commands, and testing sending and executing commands that performed a dark calibration of the Hyperion instrument.

In July 2003, a ground version of CASPER generated several plans that were subsequently uplinked and executed onboard. These plans included image data takes, maneuvers, and telecommunication passes. The purpose of this test was to prove that CASPER could generate valid plans that could be executed by the satellite.

In August 2003, onboard decompression was tested. This capability is used to compress the software before uplink because the uplink rate is only 2 Kb/s. Without compression it would take more than a week to upload the entire ASE software. This test involved uplinking several compressed files, decompressing them onboard, and then downlinking them. The files were then checked for errors.

Onboard testing of the full ASE software has been delayed due to a failure in the EO-1 ground testbed. Additional testbeds have been built and will be used to ground test the entire ASE package (including CASPER) before flight. Meanwhile we have been limited to a ground testbed with only 32MB of memory, not enough to test the entire ASE software package. Using the 32MB testbed, we have been able to test the CASPER and SCL software together. In October 2003, we upload the CASPER/SCL software and ran a test. In this test, high level goals were uplinked to open and close the Hyperion instrument cover, to perform an instrument dark calibration, and to perform an xband downlink. CASPER autonomously developed a plan from these goals. SCL expanded the plan into spacecraft commands which were then executed.

The next step will be to test an autonomous instrument data take using CASPER/SCL. This test involves uplinking a high level goal that includes a target location and a few instrument mode parameters. Once this test has been completed, we will command 5 autonomous instrument data takes and autonomous downlinks of the data. That will complete 10 of 20 of our technology validation experiments. We will then complete testing of the full ASE software including science algorithms. After completion of testing, we will uplink the software and test at least 5 data editing and 5 response experiments. The data editing experiment

uses the cloud cover algorithms to discard scenes that are mostly cloudy. The response experiment performs onboard analysis of the science data followed by a CASPER autonomously planned image. The new image is triggered by either change in the data or discovery of an interesting feature in the data. We will continue to run additional data editing and response experiments until at least October 2004.

8. CONTRIBUTION TO FUTURE MISSIONS

The ASE enables demonstration of onboard science in an Earth-directed mission, but has direct relevance to a large number of Space Science missions throughout the solar system.

As described above, the ASE will monitor selected terrestrial environmental processes that directly impact human existence, but which, importantly, have extraterrestrial analogues. Onboard science data processing 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), Saturn Ring Observer (SRO) and Jupiter Icy Moons Orbiter (JIMO). 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 ASE onboard science processing 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 that have a high probability of lander survivability and where the crust is thin enough for deployment of a sub-crust submarine explorer

Mars is the target of a series of missions by NASA and other space agencies. These missions are summarized in Table 1. An imaging orbiter mission could 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 [11]. 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 Mars Network, which will yield a GPS capability. This would allow a low-cost second deployment to Mars of a variable-baseline interferometer SAR constellation.

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. ASE is a step on the road to achieving this level of autonomy.

Table 1. 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 Scout Mission
2008	Mars TelecomSat
2009	Mars Science Laboratory SAR-capable science orbiter
2011+	Sample return mission

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 have an existing or future Mars mission infuse the ASE software into their baseline flight software.

9. IMPACT ON OPERATIONS

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 computing limitations, the ASE architecture for EO-1 does not include an autonomous fault protection component. Although this wasn't included for EO-1, it's 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.

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 not even have to be downlinked. These algorithms can be particularly useful on bandwidth-limited missions by returning the most important science data.

10. RELATED WORK & SUMMARY

In 1999, the Remote Agent experiment (RAX) [13] executed for a few days onboard the NASA Deep Space One mission. RAX is an example of a classic three-tiered architecture [8], as is ASE. RAX demonstrated a batch onboard planning capability (as opposed to CASPER's continuous planning) and RAX did not demonstrate onboard science. PROBA [14] is a European Space Agency (ESA) mission that will be demonstrating onboard autonomy and launched in 2001. However, ASE has more of a focus on model-based autonomy than PROBA.

The Three Corner Sat (3CS) University Nanosat mission will be using the CASPER onboard planning software integrated with the SCL ground and flight execution software [3]. The 3CS mission was scheduled for launch on the Space Shuttle in late 2003, but has since been delayed indefinitely. The 3CS autonomy software includes onboard science data validation, replanning, robust execution, and multiple model-based anomaly detection. The 3CS mission is considerably less complex than EO-1 but still represents an important step in the integration and flight of onboard autonomy software.

More recent work from NASA Ames Research Center is focused on building the IDEA planning and execution architecture [12]. In IDEA, the planner and execution software are combined into a "reactive planner" and operate using the same domain model. A single planning and execution model can simplify validation, which is a difficult problem for autonomous systems. For EO-1, the CASPER planner and SCL executive use separate models. While this has the advantage of the flexibility of both procedural and declarative representations, a single model would be easier to validate. We have designed the CASPER modeling language to be used by domain experts, thus not requiring planning experts. Our use of SCL is similar to the "plan runner" in IDEA but SCL encodes more intelligence. The EO-1 science analysis software is defined as one of the "controlling systems" in IDEA. In the IDEA architecture, a communications wrapper is used to send messages between the agents, similar to the software bus in EO-1. In the description of IDEA there is no information about the deployment of IDEA to any domains, so a comparison of the performance or capabilities is not possible at this time.

ASE on EO-1 will demonstrate an integrated autonomous mission using onboard science analysis, replanning, and robust execution. The ASE will perform intelligent science data selection that will lead to a reduction in data downlink. In addition, the ASE will increase science return through autonomous retargeting. Demonstration of these capabilities onboard EO-1 will enable radically different missions with significant onboard decision-making leading to novel science opportunities. The paradigm shift toward highly autonomous spacecraft will enable future NASA missions to achieve significantly greater science returns with reduced risk and reduced operations cost.


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. Ortiz, 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] *Interface and Control Systems*, SCL Home Page, sclrules.com
- [11] M. C. Malin and K. S. Edgett, "Evidence for recent groundwater seepage and surface runoff on Mars," *Science*, 288, 2330-2335, 2000.
- [12] 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*, Toulouse, France, April 2002.
- [13] NASA Ames, Remote Agent Experiment Home Page, <http://ic.arc.nasa.gov/projects/remote-agent/>.
- [14] The PROBA Onboard Autonomy Platform, <http://www.estec.esa.nl/proba/>
- [15] 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.
- [16] 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.

[17] 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.

BIBLIOGRAPHY

Rob Sherwood is a Member of Technical Staff at the Jet Propulsion Laboratory, California Institute of Technology. He holds a B.S. in Aerospace Engineering from University of Colorado at Boulder, and a M.S. in Mechanical Engineering from the University of California at Los Angeles. He is currently pursuing an M.B.A. at Loyola-Marymount University. Robert has received 7 NASA Achievement Awards for his work in Spacecraft Mission Operations and Mission Autonomy Applications. He is currently working on several projects involving Planning and Scheduling technologies.



Dr. Rebecca Castaño, Supervisor, Machine Learning Systems group at JPL. She received her Ph.D. in Electrical Engineering from the University of Illinois with her dissertation in the area of computer vision. Dr. Castaño has been advancing the state of the art in onboard science analysis methods for the past five years and has been lead author on numerous publications in this field. She is currently the technology lead for science data processing for the ASE. Dr. Castaño is also the Team Lead of the Onboard Autonomous Science Investigation System (OASIS). Her research interests include machine learning, computer vision and pattern recognition.



Steve Chien is Technical Group Supervisor of the Artificial Intelligence Group of the Jet Propulsion Laboratory, California Institute of Technology where he leads efforts in research and development of automated planning and scheduling systems. He is also an adjunct assistant professor in the Department of Computer Science at the University of Southern California. He holds a B.S., M.S., and Ph.D. in Computer Science from the University of Illinois. His research interests are in the areas of: planning and scheduling, operations research, and machine learning.



Benjamin Cichy is a Member of the Technical Staff in the Flight Software Development and Technology Group at the Jet Propulsion Laboratory, California Institute of Technology. He holds B.S. and M.Eng. degrees in Computer Science from Cornell University where he focused on Artificial Intelligence. Ben develops, integrates, and tests autonomy software for NASA's New Millennium Program.



Ashley Davies is a Research Scientist in the Earth and Space Sciences Division of the Jet Propulsion Laboratory, California Institute of Technology. His main research focus is the remote sensing of volcanism, especially on the Jovian satellite Io. He is the Lead Scientist of the Autonomous Spacecraft Experiment. Ashley was a member of the Galileo Near-Infrared Mapping Spectrometer team from 1996-2003. He has a B.Sc in Geology and Astronomy, and a Ph.D from Lancaster University in the UK, where he specialized in extraterrestrial volcanism.



Gregg Rabideau is a Member of the Technical Staff in the Artificial Intelligence Group at the Jet Propulsion Laboratory, California Institute of Technology. His main focus is in research and development of planning and scheduling systems for automated spacecraft commanding. Projects include planning and scheduling for the first deep-space mission of NASA's New Millennium Program, and for design trades analysis for the Pluto Express project. Gregg holds both a B.S. and M.S. degree in Computer Science from the University of Illinois where he specialized in Artificial Intelligence.



ACKNOWLEDGEMENT

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