Acta Futura 9 (2014) 83-91

DOI: 10.2420/AF09.2014.83

Acta Futura

Procedural Onboard Science Autonomy for Primitive Bodies Exploration

Steve Chien*, Gregg Rabideau, Dero Gharibian, David Thompson, Kiri Wagstaff, Brian Bue and Julie Castillo-Rogez

Jet Propulsion Laboratory, California Institute of Technology, 91109-8099, Pasadena, CA (USA)

Abstract. Future missions to primitive bodies will have limited time to explore these unknown bodies. Because of long round trip light times, if a mission commands the spacecraft at a detailed, time-sequenced level, there will be no opportunity to dynamically change the mission to respond to science opportunities. In order to address this issue, we are developing flight software to enable onboard science target detection and onboard response technologies to enable closed loop autonomous response for primitive bodies missions. These response methods must be able to predict the future opportunities to view the newly detected target using predicted spacecraft trajectory, target position and rotation, and future illumination conditions. These types of geometric reasoning for observation planning have traditionally been performed on the ground by highly skilled operations personnel. We describe the software under development and its application to future primitive bodies missions.

1 Introduction

Primitive bodies offer a unique view into the early solar system. Many of these objects (e.g. asteroids, comets) are unevolved since the early solar system and therefore present a unique view of the early solar system. Current space missions are exploring primitive bodies - most notably Dawn which is exploring Vesta and Ceres and Rosetta which is visiting the comet Churyumov-Gerasimenko. Additionally, future missions are under study to further explore these unknown primitive bodies.

Many primitive body missions have limited durations at a target. Because primitive bodies have a lesser gravitational field and there are a large number of asteroid primitve bodies (in the asteroid belt) it is possible to visit multiple bodies in a single mission. Therefore missions with multiple flybys and orbits are quite feasible. During a flyby the approach velocity is likely to be quite high so that the entire encounter might be of short duration (hours) - not enough time for the ground to be in the loop to modify the plan based on observed science.

We have adapted and extended a science event detection and procedural response capability to enable onboard autonomous mission response for primitive body exploration. This procedural response capability uses a resource-aware reasoning system [9] developed as an extension to the Virtual Machine Language (VML) flight executive [5] flying on numerous spacecraft. Future plans include adapting the CASPER model-based onboard planning system [1, 2] as well as compari-

*Corresponding author. E-mail: steve.chien@jpl.nasa.gov

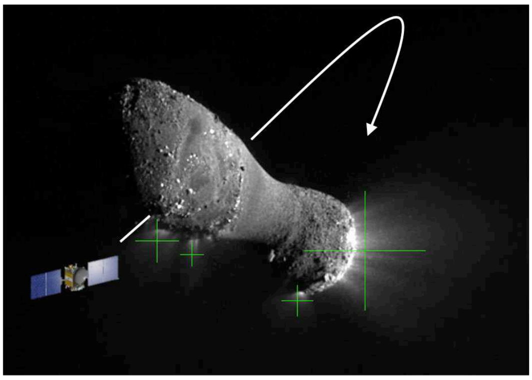

Figure 1: An Agile Science Spacecraft investigating high albedo areas on an asteroid.

son and evaluation of the two approaches. Our current VML adaptation runs in a software simulation of an embedded platform.

This prototype interfaces with onboard science detection software to enable rapid detection of science phenomena, and with navigation and geometric reasoning libraries to enable accurate planning and re-planning of followup observations. In the implemented scenario, the executive has a default observation plan of mapping the asteroid as it part of a pre-planned flyby. However early imagery is processed onboard to detect targets of interest, such as an outgassing event, an area of compositional interest, or detection of a satellite. Predeveloped science priorities indicate that such an even is higher priority than the pre-defined observation campaign, spawning new science goals. The onboard navigation/geometry software calculates potential new slews to support possible followup imaging activities to enhance science. The resource-aware VML accepts these new observation requests, and incorporates them as possible within the science priorities and operations constraints. The new plan is then executed and the spacecraft is able to acquire the preferred science imagery. This scenario is depicted in the graphic shown in Figure 1 in which a spacecraft investigates high albedo areas on an asteroid.

The software prototype is currently running with ad hoc interfaces (e.g. file) and a software simulation. As part of our continuing effort this software will be deployed to a hardware embedded platform to further mature it and ready it for future mission use. Future efforts also include expanding the range of operations scenarios.

Autonomous science provides the spacecraft with the capability to:

- detect science phenomena onboard, and
- respond by altering the original mission plan to take key observations to increase science.

Many operational scenarios exist where this onboard capability could enhance science. For example, onboard software could enable detection of an outgassing event at an asteroid or comet using an imaging instrument. Additional onboard software might then respond by commanding the spacecraft to acquire additional imagery, thereby exploiting a brief science opportunity otherwise missed.

In the remainder of this paper we first describe some of the unique challenges for primitive body missions. We then describe the ways in which agile science tehnologies can address these challenges through adaptive, onboard autonomy. Next we describe the overall concept of operations of agile science. We then provide examples of onboard science event detectors under study. We then describe the onboard procedural response component of the onboard autonomy system. Finally we discuss status, related work, future work, and conclusions.

2 Primitive Bodies Exploration - Challenges and Opportunities

Exploring primitive bodies present a number of challenges. First, the science features and events being detected include varied and subtle signatures:

- Plumes and outgassing events can be quite faint and may present in orientations that challenge detection (e.g. a plume erupting towards the spacecraft).
- The relative position of the Sun (illumination) with respect to the target and observer may not be ideal (e.g lighting behind the target).
- The morphology of the target may also present illumination challenges. If the target body has a very irregular shape, the exact illumination and observer viewing geometry may not be easily predictable.
- The target body may have unkown geology, making estimation of reflectance and other parameters more challenging.

 DOI: 10.2420/AF09.2014.83

		Missions									
		Asteroid / inert					Comet / active				
		Dawn-style mapping	Hayabusa II	OSIRIS-Rex	Trojan Tour	Chiron Orbiter	Rosetta	Comet Hopper	CSSR	CNSR/CCSR	Coma Sampler
Mission and science unknowns	Morphological units	x	x	x	x	x	x	x	x	x	x
	Surface composition, mineralogy	x	x	x	x	x	x	x	x	x	x
	Localized targets (boulders, crater walls, etc)	x	x	x	x	x	x	x	x	x	x
	Satellites	x									
	Plume activity, distribution over space and time						x	x	x	x	x
	Gravity field	x									
	Location of site for sampling/landing		x	x			x	x	x	x	
	Surface conditions at sample site		x	x			x	x	x	x	
	Rotation rate and pole location	x	x	x			x	x	x	x	x
	Spacecraft performance / faults	x	x	x	x	x	x	x	x	x	x
Applicable ground ops technologies	Single-cycle trajectory/observation selection	x	x	x	x		x	x	x	x	x
	Fast instrument data processing	x	x	x	x		x	x	x	x	x
	Fast instrument data interpretation	x	x	x	x		x	x	x	x	x
Applicable onboard technologies	Trajectory replan (fault or hazard recovery)		x	x			x	x	x	x	x
	Observation replan (opportunistic targeting)	x	x	x	x		x	x	x	x	x
	Morphological pattern recognition	x	x	x	x		x	x	x	x	x
	Spectral pattern recognition	x	x	x	x		x	x	x	x	x
	Plume/change detection						x	x	x	x	x
	Satellite detection	x									
	TRN / optical navigation for prox. ops		x	x			x	x	x	x	
	Onboard planning / execution for prox. ops		x	x			x	x	x	x	

Figure 2: Table highlighting (top) unknowns for primitive bodies missions and (bottom) applications of agile science technologies to these primitive bodies missions. Table from [12] reproduced with author permission.

In addition, primitive bodies exploration often present challenging timescales. Target bodies in the asteroid belt imply round trip light times to the Earth of approximately 1 hour. Given flyby durations of approximately 1 hour ground analysis and response to downlinked science data by a ground team is not possible.

Exploring primitive bodies is challenged by many unknowns.

- The gravity field of the target body is typcially now known except with very poor estimation. This poor gravity model will add uncertainty to any projected trajectory.
- Gas fields (e.g. for a comet) and outgassing events for comets and asteroids are unpredictable and can change the science environment as well as trajectory (due to changes in drag) with little or no warning.
- Unkown satellites. It may not be known if there are satellites prior to arrival. Satellites are both science targets and spacecraft safety hazards.

Finally, there are a limited number of close-up datasets for primitive bodes. This is particularly important in the context of the diversity of primitive body objects.

Agile science tchniques are applicable across a wide range of primitive body missions/concept either flying or under study. Figure 2 (top) shows the many unknowns that primitive bodies missions encounter. Figure 2 (bottom) shows the many relevant agile science technologies for each of the target primitive bodies missions.

3 Agile Science Scenario: Flyby

Our driving scenario for onboard autonomy aka "Agile Science" is a primitive body flyby scenario. Consider the 2010 Rosetta Orbiter flyby of the Lutetia asteroid. The timeline of the flyby is shown in Figure 3. With a relatie velocity of approximately 15 km per second, the flyby lasts less than an hour, far too short to involve the ground in the loop to command the spacecraft with

DOI: 10.2420/AF09.2014.83

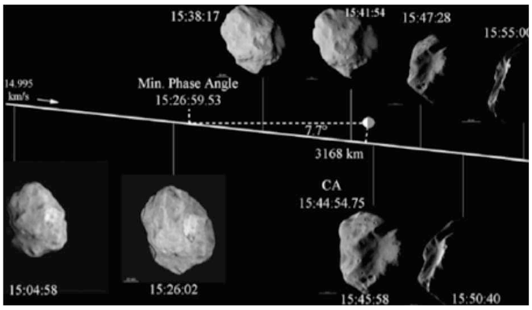

Figure 3: Timeline of the Rosetta spacecraft flyby of the asteroid Lutetia in 2010.

round trip light times being around an hour.

Traditionally, flybys wold be painstakingly planned by the ground using best estimated locations of expected targets of highest science interest. These ground planned observation sequences would be time-based sequences and would be executed open loop. Specifically, early acquisitions of science data would not be able to inform later observations.

In the Agile Science paradigm, the spacecraft and flight software enable onboard analysis of acquired science data to inform the subsequent actions of the spacecraft. Specifically, the Agile Science paradigm of execution would be as follows.

- Acquire science data
- Analyze science data
- Generate new data acquisition/target requests with priorities as pre-specified by the science team
- Assimilate new target requests into operational plan as appropriate based on prioritization.

This paradigm is highlighted by the operations scenario shown in Figure 4 which provides further detail on the autonomous target detection, prioritization and response.

4 Automated Target Identification

An important aspect of the Agile Science methodology is the ability to analyze data autonomously onboard the spacecraft to detect high priority science targets. While the general concept of Agile Science applies to a wide range of instruments initially we have focused on imaging instruments because they are central to most space missions. For primitive bodies exploration there are a wide range of science phenomena that can be discovered upon arrival at the target that warrant followup observations. A number of these we discuss below.

- Satellite search
- Outgassing/plume detection
- Volatiles search, materials search

Figure 5 shows two promising onboard processing analysis products. At left is shown a High Albedo detection in imagery of the Hartley asteroid as acquired by the Deep Impact spacecraft. In this algorithm bright areas on the target body are extracted as these are areas of high science interest because of possible presence of volatile substances. At right is shown a morphology-based detection of a plume in imagery of Enceladus acquired by the Cassini spacecraft. In this algorithm the body of the target (The moon Enceladus) is fitted to an ellipse and the algorithm is searching for an area of brightness outside of the estimated target body (ellipse). Such an area if found may be a plume which is of high scientific interest.

The onboard detection algorithms produce images with a (possibly empty) set of detections. Next, based on pre-specified science priorities defined by the science team, targets are determined from these detections. For example, a target may be generated only if a plume appears in N consecutive frames of imagery. Or a bright albedo algorithm may only produce a target if the area of bright albedo exceeds a given threshold of brightness and exceeds an area (size) threshold. The output of the target detection algorithm is a set of prioritized targets in the acquired imagery.

5 Geometric Computation

The target identification produces a set of targets (with associated priorities) in the imagery (e.g. specified in the image space, e.g. line, sample). This image space coordinate must be transformed into a target space coordinate (e.g. lat/lon, altitude on target body). Next calculations based on the spacecraft trajectory must be combined to determine legal viewing times (in effect accounting for solar position, rotation of target body, etc.). This will produce a set of possible re-imaging opportunities. Each of these can be considered a tuple of (opportunity-type-ID, priority, start time, end-time)

 DOI: 10.2420/AF09.2014.83

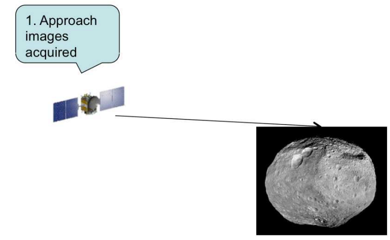

(a) Initial images acquired on approach.

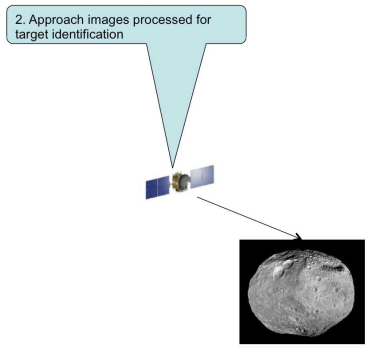

(b) Approach images processed onboard to determine followup targets.

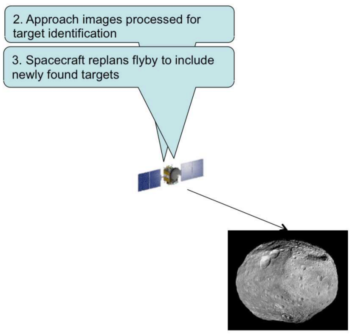

(c) Remaining flyby replanned in light of new detected targets.

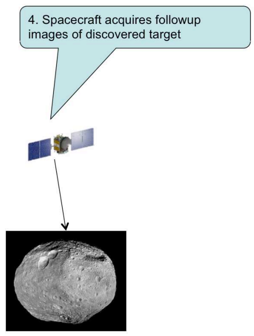

(d) Followup images acquired if science priorities warrant.

Figure 4: Agile Science Flyby Scenario

which implies a required spacecraft pointing (via the opportunity type and associated observation and target location. These observation opportunities are then passed onto the onboard response system as new requested science goals with appropriate prioritization. An important point is that this geometric reasoning involving the relative positioning and trajectories of the target body, spacecraft, sun, and other bodies is typically done in a time and knowledge intensive ground-based observation planning process. One of the unique aspects of this work is to migrate this functionality onboard the spacecraft.

For many of these geometric calculations the SPICE library [7] is the common standard used for spacecraft operations. In our implementation we have used a combination of libraries from SPICE as well as some custom code. One element of future work will be to ensure that these calculations can fit within limited flight software computing resources.

Once the timing of the re -observation opportunities has been computed they can be passed to the procedural response system which can then attempt to (schedule followup observations as warranted by the science priorities.

6 Procedural Onboard Response

For the current Agile Science software prototype we utilize a procedural response system. Specifically we have implemented our response system on top of the Virtual Machine Language (VML) flight executive [5]. VML is flying aboard numerous missions including Mars Odyssey, MRO, Genesis, Pheonix, Grail, Spitzer, Juno, Dawn, and others. VML enables layered, modular autonomous response organized by Virtual Machines (VM's). We utilize a goal and resource manager [9] layered on top of base VML that enables efficient reasoning about prioritized goals and resource conflicts. Central

DOI: 10.2420/AF09.2014.83

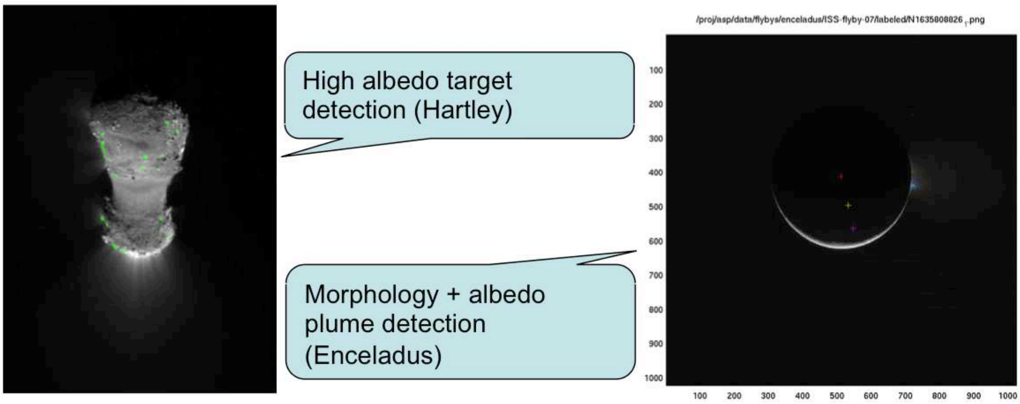

Figure 5: Example target detection and identification as shown by [11]. Left: High Albedo Target detection on Deep Impact/Hartley data Right: Morphology-based plume detection using Cassini/Enceladus data.

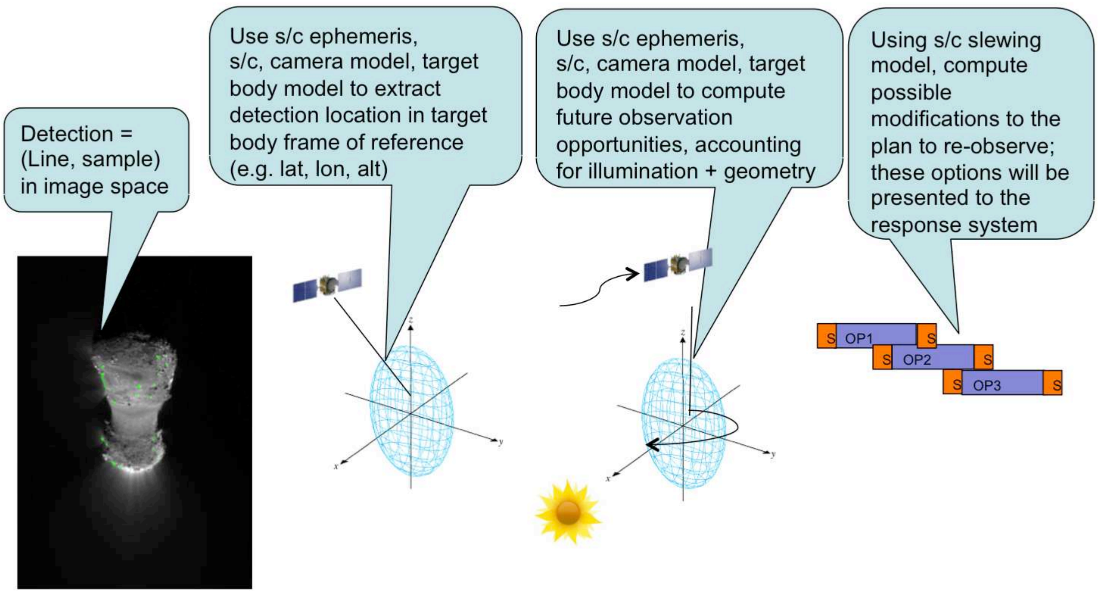

Figure 6: The process of determining when a target can be re-observed requires mapping the image space target into the target frame of reference coordinate system and then acounting for spacecraft position and pointing, target position and rotation an dother relevant features (such as the position of the sun).

 DOI: 10.2420/AF09.2014.83

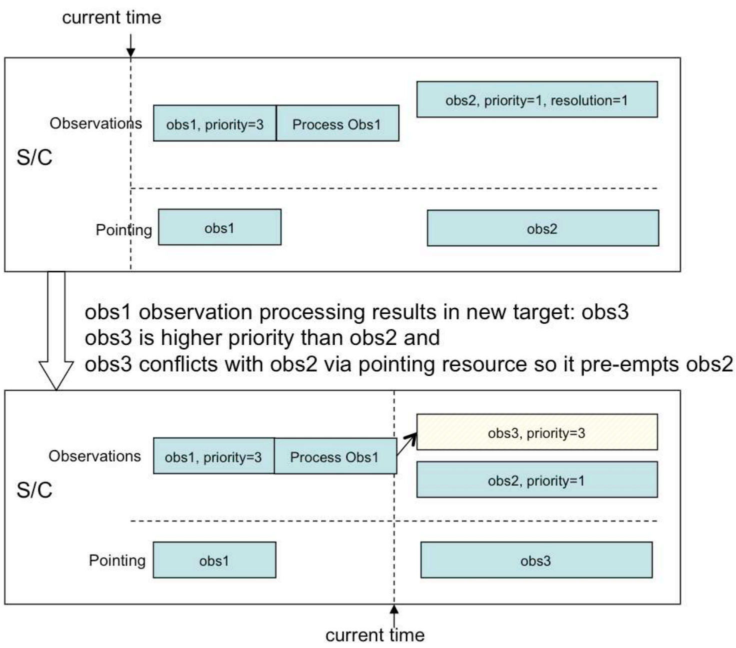

Figure 7: Dynamic goal selection in response to a newly generated high priority observation goal.

to the goal and resource manager is the concept of a flexible goal set. In our goal manager, goals represent a fixed set of activities to accomplish some purpose. For example, the set of activities required to take an image of the comet may represent a single goal even though many subparts are required to point the spacecraft, prepare the camera, acquire the image, write the image data to mass memory, repoint the spacecraft, and shutdown the camera. The goal manager also understands state and resource needs of the goals and uses this information to detect and track interactions (conflicts) between goals. At the core of the goal manager is an efficient computational algorithm to select the highest priority goals that do not conflict for inclusion into the current baseline plan at any point in time.

Figure 7 illustrates the operation of flexible goals sets. At first, the plan includes Observation 1, processing the data from observation 1, then taking observation 2. However, processing the data from observation 1, results in the creation of a new goal observation 3. Observation 2 and Observation 3 conflict, as they occur at the same time and require different pointings. The goal manager therefore inserts Observation 3 into the baseline plan and removes Observation 2 from the baseline plan. The baseline plan is then executed. A key assumption of the goal selection algorithm is that goals do not have temporal flexibility (i.e. their start and end times are fixed). This assumption is key to the computational tractability of the goal selection algorithm. The goal selection algorithm is an incremental algorithm. It must be run to re-select goals whenever goals are added or removed from the goal set. A goal is changed by removing it and adding the updated version. Goals are maintained in sorted sets. Shared resource and state interactions are maintained. When adding or removing a goal, only goals of equal or lower-priority to the added/removed goals need to be re-selected. Re-selection is worst-case $O(N^2 \lg N)$ and average case $\Theta(N \lg N)$. The worst case is when each goal has a constraint on every available resource (i.e. maximum interaction) and all goals are selected (i.e. no conflicts). In the average (typical) case each goal has a constant number of interactions via state/resource. The goal selection algorithm is described in greater detail in [9].

7 Discussion and Conclusions

7.1 Related Work

Considerable prior work has investigated spacecraft autonomy using procedural methods.

- The Autonomous Sciencecraft (ASE) [1] utilized a planner (CASPER) in concert with the Spacecraft Command Language (SCL) executive and has been used for primary operations of Earth Observing One 2004 to the present (2013).
- V AMOS [13] is an onboard executive in development by DLR that validates branching plans on the ground and then selects execution branches onboard for operational flexibility.
- GOAC [3] is a goal-oriented architecture developed by ESA for future onboard use.
- The Remote Agent [6] utilized the batch planner RAX-PS and the procedural executive ESL [4] to control the Deep Space One mission for 48 hours in 1999.
- T-Rex [10] – is a planning and execution architecture that has been deployed for control of autonomous underwater vehicles.
- SCL [8] in addition to ASE has been used on the Tacsat-2 mission to offer onboard procedural rule-based automation.

7.2 Conclusions

Onboard autonomy can enable dynamic science for primitive body missions. Many science events can be detected via instrument processing techniques that are amenable to onboard computation.

We have demonstrated a capability to perform:

- Target Detection
- Target extraction and geometric computation required for re-observation opportunity analysis
- Modification of the existing observation plan to incorporate the new observation if warranted by science priorities
- Execution of the new plan

This capability is currently implemened in a software testbed and is being matured for future NASA missions.

Acknowledgements

This work was perfomed by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

References

[1] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davis, D. Mandl, B. Trout, S. Shulman, et al. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication*, 2(4):196–216, 2005.

[2] S. A. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau. Using iterative repair to improve the responsiveness of planning and scheduling. In *AIPS*, pages 300–307, Beckenridge, CO, April 2000. AAAI.

[3] S. Fratini, S. Martin, N. Policella, and A. Donati. Planning-based controllers for increased levels of autonomous operations. In *Proc 12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2013)*, Noordwijk, The Netherlands, May 2013.

[4] E. Gat. Esl: A language for supporting robust plan execution in embedded autonomous agents. In *Proceedings of the IEEE Aerospace Conference*, volume 1, pages 319–324, Snowmass, CO, 1997. IEEE.

[5] C. Grasso and P. Lock. Vml sequencing: Growing capabilities over multiple missions. In *Proceedings of SpaceOps*, volume 8, Heidelberg, Germany, 2008.

[6] N. Muscettola, P. P. Nayak, B. Pell, and B. C. Williams. Remote agent: To boldly go where no ai system has gone before. *Artificial Intelligence*, 103(1):5–47, 1998.

[7] NAIF. About the spice library, August 2013.

[8] B. L. Prumo, B. Buckley, M. Long, and C. Finley. Unmanned spacecraft operations. In *Proc AIAA Infotech@Aerospace,*, Seattle, WA, 2009.

 DOI: 10.2420/AF09.2014.83

[9] G. Rabideau, S. Chien, and D. Mclaren. Tractable goal selection for embedded systems with oversubscribed resources. *Journal of Aerospace Computing, Information, and Communication*, 8(5):151–169, 2011.

[10] K. Rajan, F. Py, and J. Barreiro. Towards deliberative control in marine robotics. In *Marine Robot Autonomy*, pages 91–175. Springer, 2013.

[11] D. R. Thompson, M. Bunte, R. Castano, S. Chien, and R. Greeley. Image processing onboard spacecraft for autonomous plume detection. *Planetary and Space Science*, 62(1):153–159, 2012.

[12] D. R. Thompson, J. C. Castillo-Rogez, S. A. Chien, R. Doyle, T. Estlin, and D. McLaren. Agile science operations: A new approach for primitive bodies exploration. In *Proceedings of SpaceOps*, Stockholm, Sweden, June 2012.

[13] M. T. Wörle and C. Lenzen. V amos verification of autonomous mission planning onboard a spacecraft. In *Proc Intl Workshop on Planning and Scheduling for Space*, Moffett Field, CA, March 2013.