# Agile science operations: A new approach for primitive bodies exploration

David R. Thompson[1], Julie C. Castillo-Rogez[2], Steve A. Chien[2], and Richard Doyle[2], Tara Estlin[2], David McLaren[2]
*Jet Propulsion Laboratory, California Institute of Technology. 4800 Oak Grove Dr., Pasadena, CA 91109*

**Primitive body exploration missions such as potential *Comet Surface Sample Return* or *Trojan Tour and Rendezvous* would challenge traditional operations practices. Earth-based observations would provide only basic understanding before arrival and many science goals would be defined during the initial rendezvous. It could be necessary to revise trajectories and observation plans to quickly characterize the target for safe, effective observations. Detection of outgassing activity and monitoring of comet surface activity would be even more time constrained, with events occurring faster than round-trip light time. "Agile science operations" would address these challenges with contingency plans that recognize the intrinsic uncertainty in the operating environment and science objectives. Planning for multiple alternatives could significantly improve the time required to repair and validate spacecraft command sequences. When appropriate, time-critical decisions could be automated and shifted to the spacecraft for immediate access to instrument data. Mirrored planning systems on both sides of the light-time gap would permit transfer of authority back and forth as needed. We survey relevant science objectives, identifying time bottlenecks and the techniques that could be used to speed missions' reaction to new science data. Finally, we discuss the results of a trade study simulating agile observations during flyby and comet rendezvous scenarios. These experiments quantify instrument coverage of key surface features as a function of planning turnaround time. Careful application of agile operations techniques could play a significant role in realizing the Decadal Survey recommendations for primitive body exploration.**

## I. Introduction

**FUTURE primitive bodies missions would demand an innovative approach to mission operations.** Mission concepts recommended for NASA's *New* Frontiers program include a Comet Surface Sample Return [CSSR] or a Trojan Tour and Rendezvous [TTR]. These missions would have little information about the targets before arrival, so key science goals would be defined in concert with initial data collection. Such missions would require a period of characterization in the target environment to understand how best to operate a mission effectively from both a science and an engineering perspective [2,3]. Mission success would require operators to quickly climb this learning curve and revise observation plans. Further complications include multifaceted physical constraints from diverse instrument suites, varied orbits, and proximity operations. Even fast reactions by ground operations may not always be sufficient, since events like flybys and transient cometary activity would occur faster than the Round Trip Light Time (RTLT). This work investigates new operations strategies that could address these challenges.

**We use the term *agile science operations* to describe this new approach**. Agile science exploits *contingency* planning that recognizes uncertainty is unavoidable, and aims for rapid and reliable responses to the evolving objectives. Revisions can take the form of decision point opportunities for followup data collection, selective downlink, or transfer to alternative orbits. Alternatively, a fast *replan* revises or reorders entire portions of the spacecraft activity sequence. When these decisions can be automated, the spacecraft itself can react faster than RTLT using automatic data analysis. Such analyses can recognize key events or targets of opportunity for followup, or select from a menu of priority subtasks. In the extreme, mirrored planning systems on both sides of the light-time gap can transfer authority as needed. A mirrored planning system plays to the best strengths of both onboard action (fast reaction time) and ground operations (human expertise to prioritize science objectives).

**Science goals and physical constraints could drive operations either toward predetermined or flexible operations.** For example, the unique constraints of the Cassini mission lead to a conformant planning approach that

[1] Research Technologist, Machine Learning and Instrument Autonomy 306-463, *david.r.thompson@jpl.nasa.gov*
[2] *{firstname.lastname}@jpl.nasa.gov*

completely predetermines the spacecraft activities in advance.[4-6] Gravitational and pointing constraints demand trajectories crafted years in advance to achieve very specific observation objectives. In contrast, the Mars Exploration Rovers (MER) study many localized targets that cannot be anticipated in advance such as rocks, outcrops, or other targets of opportunity.[7] Redundant actuation and rover mobility provide more flexibility, leading MER to regular tactical decisions and an ongoing cycle of discovery with new targets driving each next day's activities. Over the course of its extended mission the MER team has refined this process through hundreds of iterations and now achieves a sub-24 hour cycle from the initial tactical data analysis to a vetted command sequence. Recently they have incorporated onboard data analysis for science-driven instrument targeting that bypasses RTLT altogether.[8] We find several primary drivers of operations strategy:

- *Foreknowledge of instrument targets,* which determines what *could* be planned in advance.
- *The flexibility to revisit targets*, since redundancy reduces the risk of opportunistic changes.
- *Downlink intervals and bandwidth*, which affect autonomy requirements and tactical situational awareness.

**Future primitive bodies missions would require the flexibility of MER operations**. Like MER, small bodies missions would benefit from comprehensive high-resolution measurements with multiple instruments at localized targets of opportunity.[9] Such targets would include surficial units visible during flyby, transient features on active bodies, and candidate landing sites during rendezvous. Especially in the case of flyby encounters, the science that could be achieved at an unknown target might not be fully known in advance. This situation is broadly analogous to the rover that cannot predict the outcrop waiting over the next hill. However, flyby opportunities in, e.g., the Trojan clouds would likely be limited in comparison to the diversity of properties the Trojan asteroids are expected to exhibit that would reflect diverse origins and evolutionary paths. Hence each flyby must be optimized to access key science information that may be unique to each target. Probably the most interesting yet challenging feature to capture at primitive bodies is outgassing activity, the occurrence of which would likely be fortuitous. Other important science would come from the detection and characterization of localized patches of volatiles and/or organic-rich material. As a matter of comparison, Saturn's satellite Phoebe displays ice-rich material over very restricted areas – mainly along crater rims (see Castillo-Rogez et al.[10] for a review). Its dark surface also presents spectral variations at the regional scale whose origin is undetermined.

**However, the unique challenges of small bodies missions would prevent the MER model from translating directly.** Small body missions do not generally use redundant sensor or antenna actuation. Also, as noted above, the properties of flyby targets may significantly differ between each other so that the experienced learn from one flyby may not necessarily inform the strategy to implement for the next flyby. Mobility around the target would also differ from the MER scenario. Microgravity trajectories could access arbitrary surface locations, implying some MER-like flexibility in the ordering of key activities, but free trajectories are subject to asymmetric multibody gravity fields and the strong geometric interdependencies between instrument pointing, power, communications, and thermal issues in a rigid spacecraft. Moreover, adjustments would deplete limited fuel. Situational awareness is more difficult since fixed high-gain antennas would require spacecraft slewing. This would leave a paucity of timely context information needed for tactical replans. Further complicating situational awareness is the intrinsic challenge of visualization and navigation in a three-dimensional environment. and the potential need to manage various forms of risk inherent to multiple asteroid systems (i.e.,asteroid satellites or dust ring not detected from ground-based observations – A similar issue recently arose for the *New Horizons* mission after unexpected satellite P4 was discovered.[11]

**The remainder of this paper examines technological advances that could help close this gap.** The following section describes technology and operations solutions for improving reaction time to new science data. We then simulate several specific small-bodies exploration scenarios to show how these improvements would affect the science yield. We conclude that careful application of agile operations could play a significant role in realizing the Decadal Survey's recommendation for small bodies exploration.

## II. Approach

For our purposes the mission planning process incorporates data acquisition, downlink, science interpretation and target prioritization, planning, sequencing and validation, uplink and execution. We evaluate this interplay with the Critical Path Method (CPM), a technique from operations research with a long history in aerospace, defense, and software development.[12] We will explore how planning and execution techniques can speed various components in the critical chain. Fundamentally, all such improvements expedite responses to some new piece of information. Specifically, we seek operations strategies that improve *reaction time*, which we define as the *minimum time required for a new measurement to affect appropriate changes to some subsequent action by the spacecraft*. One can posit assumptions about the spacecraft and mission conditions to analyze how modifying specific aspects of the critical path affect reaction time, and as a consequence, potential science return.  Such models are highly simplified but they provide a formal model for describing changes to the planning process.

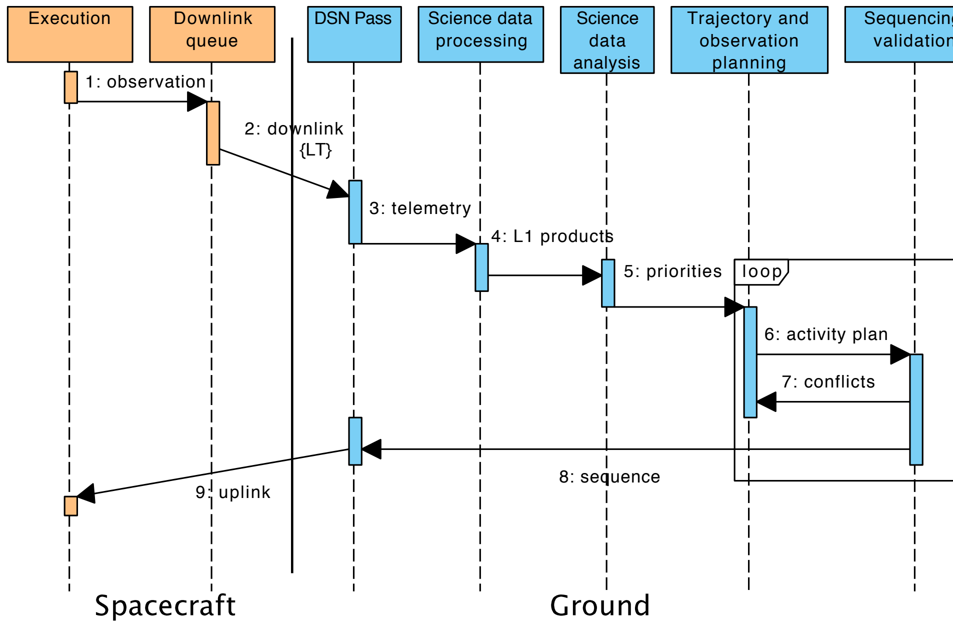


**Figure 1. Activity sequence diagram for a generic planning and execution cycle.**

The reaction time is the total of all segments in the chain from the first acquisition of the new information until the desired response. Figure 1 illustrates a generic case.  Each activity on the critical path occupies a separate vertical column, with Earth-referenced time advancing downward.  This "activity sequence" diagram uses a Universal Modeling Language standard.[13] The pipeline includes the following chain of dependencies:

- *Execution, runout, and downlink queuing* The spacecraft acquires the new information and finishes the rest of its command sequence, possibly performing other navigation and data collection prior to downlink.
- *Downlink.* A Deep Space Network (DSN) communication incurs a light-time delay.
- *Science data processing.* The telemetry is unpacked, processed into human-understandable form and delivered to analysts.
- *Science data analysis.* New understanding causes revision to observation priorities.
- *Trajectory and Observation Planning.* Decision makers make appropriate changes to the plan, producing new science goals and the corresponding data collection activities.
- *Sequencing and validation*.  This reconciles the new data collection priorities and spacecraft activities with the physical constraints on the spacecraft.. The result is a valid command sequence for onboard execution.
- *Uplink.*  This block corresponds to the time before the next scheduled uplink opportunity, the subsequent DSN pass and corresponding light delay before the new commands arrive to the spacecraft computer.

The rest of this section presents several techniques that might speed this process.

**Multi-goal planning** Multi-goal (or Multi-objective) planning improves reaction time by *planning in advance*, with redundant plans to handle contingency cases.[14-15] One defines sets of valid plans that cover all possible contingencies.   Science priorities are described by objective functions that map specific science or measurement objectives to scalar values.  For example, the goal "collect close-range images of active areas near the south pole" could measure the image coverage in this region.  Alternatively, objective functions could describe goals related to depletable resources like electrical power or hydrazine. A pool of candidate plans can provably cover any linear combination of coefficient weightings ascribed to these science priorities.   Planners can err on the side of caution, including goals they do not expect to exercise without ever paying a run time cost. The selection of chosen plans after updating the science goals is trivial, the actual work having been done in advance.  This can reduce the planning/validation loop to a simple feedforward selection and adjustment step (Fig. 2), yielding reaction time improvements of hours to days. At runtime, whenever new information appears, mission planners select the appropriate sequence template from the current pool of Pareto-optimal possibilities[16]. A demonstration of this strategy was recently performed using MUSE, the Multi-User Scheduling Engine.[15] MUSE specifically advances this capability for the Deep Space Network and Cassini applications. Similar techniques could select from different

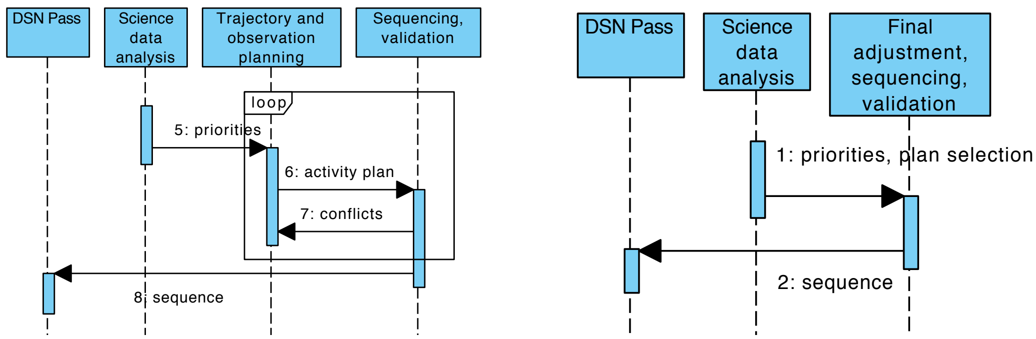


**Figure 2. Multi-goal planning.** *Left: Traditional trajectory and observation validation loop. Right: Multi-goal planning can shorten the timeline by collapsing the iterative process to a feedforward step,*

trajectory/observation combinations during a small bodies mission.

**Automated science data analysis** Faster planning cannot overcome the time required to interpret new instrument data. Key compositional information derived from data products like spectra is challenging to use in a tactical setting due the intrinsic data volume (which limits spatial coverage) and the subtlety of interpretation (which requires extensive calibration and possibly mixture modeling). Often discoveries are made months or years after the first availability of the data. Hyperspectral imagery challenges the limits of a computer screen, and it is impossible to visualize the entire data cube. New research aims to make these products more suitable for tactical use, providing "quicklook" summary products for the spectral domain. This can potentially reduce the time required to yield actionable changes in observation priorities. For example, the HiiHAT Hyperspectral Analysis tools use numerical techniques to identify physical features in the scene corresponding to the purest examples of each material class.[17-19] These endmember features act as a concise visual summary of the scene, and facilitate mapping.

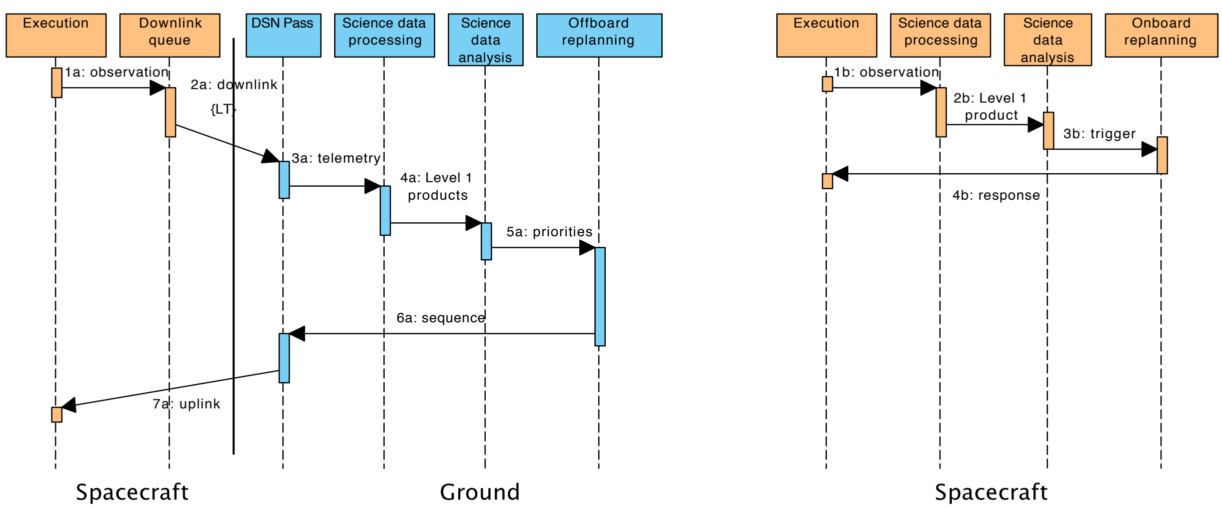


**Figure 3. Left: Traditional planning cycle.** *It incurs the time delay of the downlink queue, RTLT, the DSN telemetry pass and unpacking.* **Right: Onboard replanning provides an alternative decision path**.

**Onboard data analysis and response** Onboard decision can significantly speed reaction time; it offers an alternative path to ground-based analysis that forgoes the downlink interval and light time delay altogether. For the simplest autonomy scenarios, such as the *Deep Impact* encounter,[20,21] time and bandwidth resources are pre-allocated with only targeting geometry adjusted on-line. A monitoring mode could reallocate these resources dynamically (for example, the event no target is found). The most advanced onboard replanning includes a fully resource-aware planner that can arbitrate between science priorities while also preserving resources. These systems, such as CASPER have an explicit model of activity conflicts, dependencies, and resource usage. and can re-shuffle plan activities if planned observations become redundant. [22,23] An extreme case is a *parallel planning system* that mirrors the ground system as closely as possible. We can represent the onboard decision process by an alternative processing pipeline. Figure 3 shows the new reaction loop that takes place onboard the spacecraft.

| | Objective | Data analysis | Trajectory Generation | Activity Planning | Followup execution | Total reaction time | Reference |
|---|---|---|---|---|---|---|---|
| **ASE (EO-1)** | Prioritize downlink (thermal detection) | ~2hr | - | 30m | 90m | ~4hr | [Chien 2005] |
| **HiiHAT (EO-1)** | Prioritize downlink (spectral mapping) | 5hr | - | - | - | - | [Bornstein 2011] |
| **Autonav (Deep Impact)** | Trajectory updates during encounter | 1.5-8h | 10-200m | - | 1m | 2-10h | [Ridel 2001] |
| **AEGIS (MER)** | Target detection, followup | 10-20m | - | 2m | <1m | <25m | [Estlin 2011] |

**Table 1: Timing of activities for autonomous analysis, replanning, and followup (historical missions).**

## III. Case Study: Smart Flyby

Here we simulate intelligent data collection and transmission to characterize science yield possible from a *smart flyby* capability. In our scenario the spacecraft would explore a spectrally ambiguous Main Belt Asteroid and pass it on a fixed trajectory. This is broadly analogous to missions such as Deep Impact or the upcoming *New Horizons* flyby of Pluto and Charon in which the main encounter is shorter than RTLT. Currently flybys typically use one of two measurement strategies. Either the instrument field of view subtends a small fraction of the target, with coverage based on geometry or best guess prior knowledge of surface feature locations and distribution; or the instrument covers the entire target but sacrifices spatial resolution. It is generally desirable to get non-uniform coverage that hits key features of interest with high-resolution measurements. This can provide both breadth and depth of sampling.

A *smart flyby* would begin with onboard spectrometer or telescope targeting demonstrated in the Deep Impact encounter.[24] It would add onboard image interpretation to target morphological features (rather than a pre-planned target point). It would use a telescopic imager similar to the Deep Impact High Resolution Imager with an integrated Visible and Near Infrared (VNIR) or Short Wave Infrared (SWIR) spectrometer similar to the Deep Impact encounter, but more specific target selection and a closer approach would permit a higher spatial resolution and improved signal-to-noise for sampling spectrally diverse surfaces.

The simulation assesses the surface features that could be measured for various onboard reaction times, which constrains the additional data that could be acquired during a smart flyby. Except where stated, all geometry and trajectory information is computed using data and software from the NASA Navigation and Ancillary Information Facility.[25] We model the asteroid as a triaxial ellipsoid. We borrow trajectory geometry from Rosetta's flyby of asteroid 21 Lutetia.[28] The physical parameters of this encounter, such as the approach trajectory and Lutetia's rotation period of ~8 hours, are broadly representative of many small body targets of interest. Figure 4 shows the encounter timeline including the distance between spacecraft and the target. We simulate the span from July 10, 2010, 13:00 to 17:00 GMT. During this time the *Rosetta* spacecraft approached the target to within 3200 km and performed a "turn and track" maneuver pointing the instrument suite at the target center. This viewed a large fraction of the surface, providing high-resolution imagery of the target (Fig. 5). Naturally many features were only visible for very short periods due to illumination and geometric constraints. As the spacecraft passes, features quickly disappear over the local horizon or become too distant to resolve.

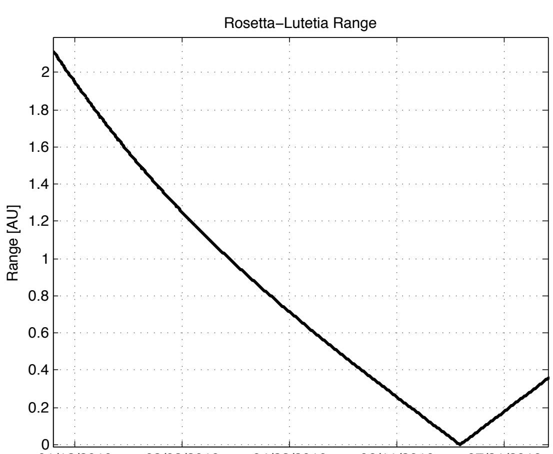


**Figure 4: Timeline for Rosetta encounter of 21 Lutetia.** *The simulation spans about 4 hours during closest approach.*

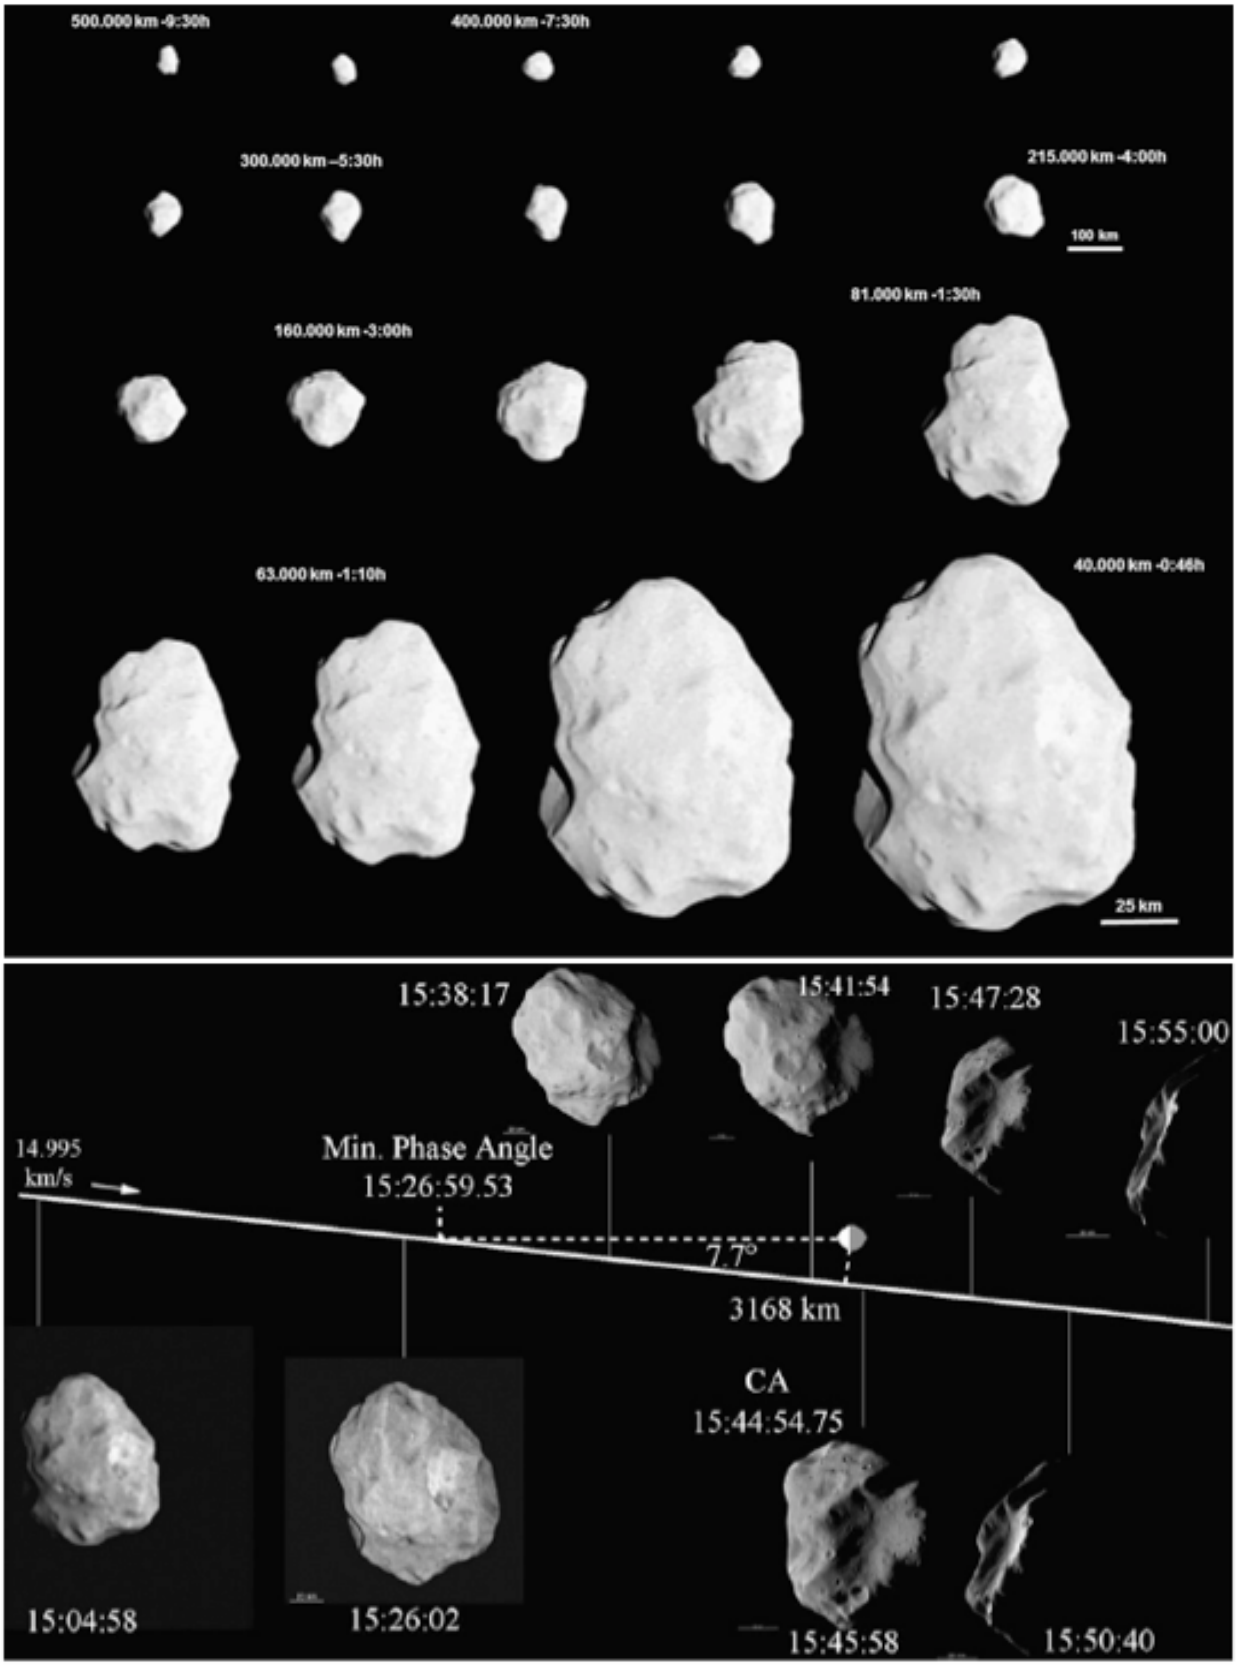


**Figure 5.** *Geometry and Illumination of Lutetia as viewed by Rosetta prior to (top)*[26] *and during (bottom)*[27] *the encounter on July 10, 2010.*

We simulate targets as points at random locations. These features represent potential targets such as ridges, craters of various sizes that could provide insight into the deeper, fresh material, boulders, and deposition or mass wasting regions. They range in size from 100m to 5km. This simulation examines the onboard reaction times needed to achieve coverage of these surface features. We enforce physical geometric, illumination, and range constraints on visibility. Specifically, the constraints are (1) that the surface feature is not obscured by cast shadows, which for our purposes means having a solar zenith angle is at least 30 degrees above the local horizon; (2) that the illumination provides sufficient shadow contrast to see the features, requiring a solar zenith at least 30 degrees from vertical; (3) the surface features must be viewed without significant affine distortion, which demands the angle between the spacecraft and the surface normal to be no greater than 60 degrees; and (4) the surface feature can only be *discovered* once it can be imaged by the context camera at a resolution of at least 60m/pixel; this was the minimum size of resolved features in the actual context imagery during the actual Rosetta flyby. After this initial imaging its location is known and it could still be captured from longer distances using narrow Field of View telescopes.

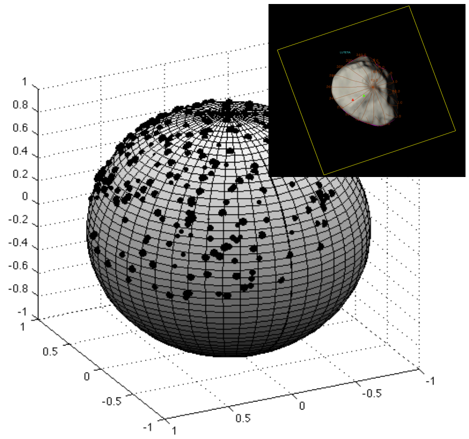


**Figure 6. Example of simulated target features.** *Targets appear randomly on the surface of an ellipsoidal target model.*

Applying all of these constraints leave a subset of the target visible during the flyby (Fig. 6). Coverage regions vary somewhat based on the size of the feature detected. We track each simulated feature to discern when it comes into view by the context imager, and then compute the time window over which the feature is visible until it passes over the spacecraft horizon or violates one of the other constraints on range, geometry or illumination. Based on this visibility window, we can calculate how quickly an onboard system would have to react in order to collect additional image or spectrometer data.

Figure 7 shows the simulation results, charting the fraction of surface features that are targetable for various reaction times, again defined as the entire time interval between the initial context image and followup data collection. For comparison, the average run of AEGIS target selection onboard the Mars Exploration Rovers requires approximately 10-20 minutes, while EO-1 retargeting typically occurs in tens of minutes. We see two processes represented in the performance curves. An initially graceful degradation represents the slowly changing geometry of the target's spacecraft-relative rotation. As reaction time slows, a greater fraction of the targets disappear beyond the horizon of permissible observation angles. We also see a sharp drop-off in performance due to the resolution of the context instrument. If the initial range upon detection is too short, it would be impossible to schedule a followup before closest approach and loss over the horizon.


**Figure 7. *Smart Flyby* Performance.** *The fraction of total targets captured under reasonable illumination and geometry constraints, as a function of plan*

The resolution constraint becomes dominant for the smallest features, suggesting that early detection is important in these cases. It would be important to tune the context resolution images appropriately to the speed of onboard data analysis to ensure that it falls on the proper side of this threshold. Good capture rates for most feature sizes suggest that timely followup would be possible for contemporary feature detection strategies. For an AEGIS-relevant timeframe of 10-20 minutes, a large fraction of the possible surface features are achievable at sizes below 500m. This offers significant benefits for very high-resolution instruments *vis a vis* random data collection. It suggests that the timescales of existing onboard detection and response systems already operate at sufficient speeds to be of relevance for a smart flyby mission.

## IV. Case Study: Landing Site Characterization

Agile science operations could enhance a time-limited characterization of candidate sites as a precursor to proximity operations. We consider a scenario based on the initial 3-week period during the rendezvous between Rosetta and the comet 67P/Churyumov-Gerasimenko. This target, trajectory and payload resemble many potential small bodies missions including Comet Sample Return or Cryogenic Comet Sample Return mission concepts. The *Rosetta* mission includes landing a small platform, *Philae*, with on-board ten instruments that will characterize the physical and chemical properties of the comet surface as well as perform sounding of the deep interior with magnetometry and radar sounding. A survey of the surface is required to identify a site that is safe for the lander but also offers the opportunity for the most compelling science (e.g., evidence for volatile or organic enrichment). Here we focus on the former aspect.

One proposed operations plan calls for *Rosetta* trajectories to be specified at three-week intervals with more regular weekly updates to the observation activity plan.[2] It is believed that there will be approximately three to five trajectory updates prior to delivery of the *Philae* spacecraft. In this short time, the science teams must identify the most interesting landing site and adequately characterize its safety (considering factors like roughness, fissures, or active outburst regions). The spacecraft will almost certainly begin this characterization with a generic trajectory to view many locations on the surface; subsequent trajectories would focus on top candidates of greatest interest.

The trajectories for characterization would be non-Keplerian, but absent more specific knowledge a simple elliptical orbit segment provides a reasonable approximation of the coverage that might be expected during this mission stage. Figure 8 shows an example; this reference trajectory lasts three weeks between updates and maps the surface with an elliptical path encircling the nucleus along the terminator from a distance of approximately 30km. When the first data arrives, *Rosetta's* path will probably remain fixed along the original observation trajectory for at least the first three weeks of its encounter. Here we can meaningfully simulate agile observation selection while bypassing the more complex and mission-specific challenge of trajectory selection. Specifically, we will evaluate the effect of fast updates to the observation schedule during this first trajectory, and their potential to improve the information at hand for selecting the subsequent "close approach" target sites. Mission planning requires mapping and characterization of all potential landing sites, detecting activity level and drawing high-resolution images to facilitate topographic measurements at 10cm/pixel (to ensure that there are no serious hazards within the *Philae* lander's landing ellipse), and spectral measurements at tens of meters per pixel (to provide compositional context to later in situ data). The range of angles under which a site can be viewed constrains stereo estimates of surface topography, so multiple measurements from widely spaced locations are important.

We examine the effect of fast observation replanning, simulating on high-resolution targeted observations of specific candidate landing sites. We envision operators identify a candidate site in context images and then wish to follow up with a targeted measurement during the same trajectory. Our computation of reaction time assumes a reference Rosetta planning scheme detailed in Koschny et al.[2] A one-week downlink interval determines the length of the runout/downlink queue. We integrate over all possible time gaps between the initial imaging of the target and the next downlink, assuming discovery is equiprobable at any time during the one week interval. The reaction time also includes a runout period, a 2-hour RTLT delay and ground planning time before the new plan can be uplinked. We simulate replanning times ranging from instantaneous replans to three weeks.

As with 21 Lutetia, we model the comet surface as a triaxial ellipsoid. We simulate a trajectory during the period from Nov. 15 to Dec. 6 2014. Targets are points distributed randomly on the surface. We compute target visibility windows with geometry and illumination constraints as before. However, the illumination constraint differs; in the reference trajectories the spacecraft orbited the comet in a plane perpendicular to the vector to the sun, and we did not consider viewing and illumination angle constraints. Instead, we use give events a finite lifetime to represent the fact that they may be occluded by surface activity or topography of a bumpy surface. We consider static features that are always visible as well as those

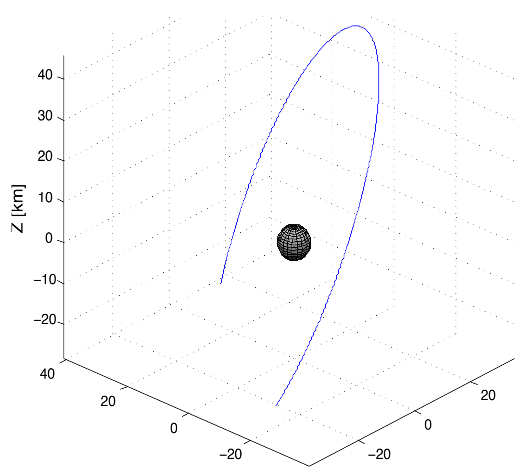


**Figure 8: Example mapping trajectory segment.** *Actual microgravity orbits would be non-Keplerian, but the coverage here approximates an initial characterization or*

which disappear after 7 or 14 days.

Followup observations achieve several science goals, but the most relevant metric is measurement coverage of surface features. We evaluate the total fraction of features for which followup observations are possible, as well as the time span of followup observations that could be achieved within the remaining time period (Fig. 9.). The former measure is relevant to followup spectroscopy as well as mapping. The right panel shows the average number of days during which the "captured" features would be visible for followup. This is important because more time implies greater flexibility in planning followup observations and could relieve pressure on the oversubscribed final days of the trajectory. In other words, catching an observation early enough and turning around a plan fast enough to provide a large time window for followup possibilities makes it more probable that such followup actions actually fit within the constraints already imposed by the trajectory. Figure 10 (left) shows the range of spacecraft angles above the feature's local horizon that one could achieve in followup images after the initial identification. This affects stereo reconstruction accuracy as well as the likelihood of finding good unobstructed visibility of the target. The right panel shows the range of solar zenith and spacecraft angles, representing the range of distinctive lighting conditions that one could achieve. Larger ranges increase the likelihood of seeing diverse stages of activity.

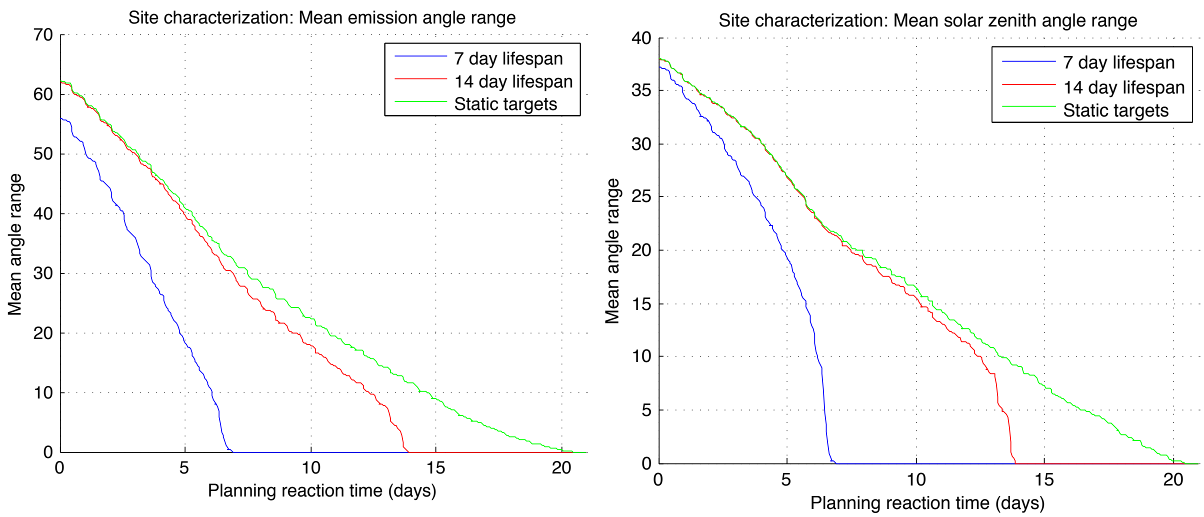


**Figure 9.** ***Landing Site Selection*** **performance.** *Left: Viewing angle range of captured targets. Right: Average solar zenith angle range for different lifespan assumptions.*

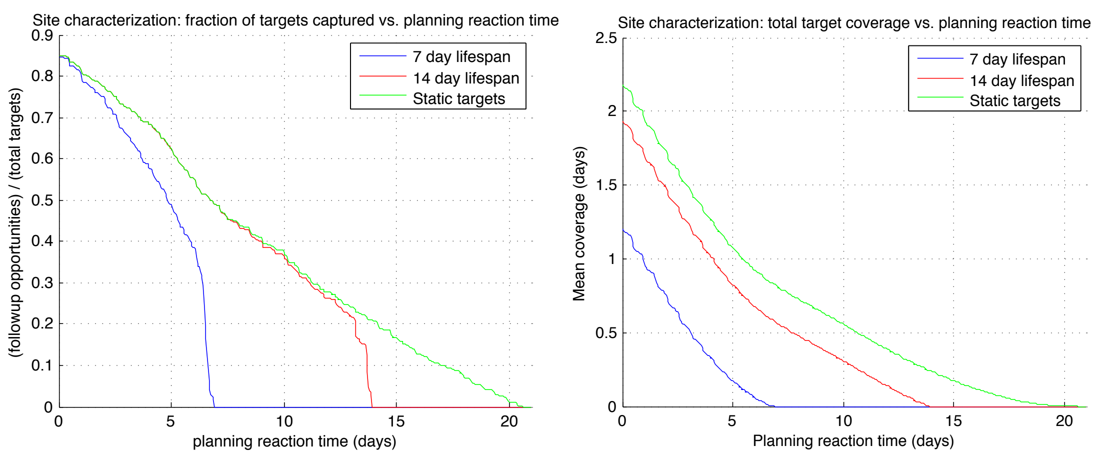


**Figure 10.** ***Landing Site Selection*** **performance.** *Left: Fraction of total targets captured after they become visible. Performance is plotted against reaction time (after the initial detection). Right: Time during which captured targets remain visible for different lifespan assumptions.*

These simulations suggest that faster planning reaction time can significantly improve the return of high-resolution observations. Simulations reveal these measures to be *highly sensitive* to plan turnaround time, and especially to differences shorter than the one-week standard downlink interval. Results suggest several general performance regimes. A sub-week planning capability (achievable through a MER-style daily operations or onboard replanning) could provide good coverage of most targets. A weekly turnaround could produce revised sequences in a week and meaningfully follow up on some fraction of targets. This might be achievable with ground-side planning such as multi-goal "plan in advance" strategies. A two-week turnaround must generally wait for a subsequent pass before collecting followup measurements of candidates. Missions like *Rosetta* will only have a few trajectories to devote to characterization activities; early downselection can preserve the time close to delivery.

## V. Acknowledgments

We thank the *Dawn* and *Rosetta* teams for all their help, guidance, and support. Key members of our advisory committee include Drs. Claudia Alexander, Artur Chmielewski, Bob Pappalardo, Carol Polanskey, Carol Raymond, and Kim Reh. We thank Rob Witoff, Barbara Streiffert and Anna Marie Aguinaldo for assistance with trajectories and the Science Opportunity Analyzer visualization software. This research was performed at the Jet Propulsion Laboratory, California Institute of Technology.  We acknowledge U.S. Government Support. This study was conducted with sponsorship by the JPL Solar System Exploration Office.

## VI. References

[1]Committee on the Planetary Science Decadal Survey; National Research Council, "Vision and Voyages for Planetary Science in the Decade 2013-2022." National Acadamies Press, 2011. ISBN-13: 978-0-309-20951-9. http://www.nap.edu/.

[2]Koschny, D. and V. Dhiri, K. Wirth, J. Zender, R. Solaz, R. Hoofs, R. Laureus, T. M. Ho, B. Davidsson, G. Schwehm. "Scientific Planning and Commanding of the Rosetta Payload." *Space Science Reviews* 2007. 128: 167–188.

[3]Russell, C. T. and C. A. Raymond, T. C. Fraschetti, M. D. Rayman, C. A. Polanskey, K. A. Schimmels, S. P. Joy. "Dawn Mission and Operations." International Astronomical Union Symposium on Asteroids, Comets, and Meteors, 2005. 229: 97-119.

[4]Russell, C. T., editor. *The Cassini-Huygens Mission: Overview, Objectives and Huygens Instrumentarium*. Springer:Netherlands, 2003.

[5]Paczkowski, B. G. and T.L. Ray. "Cassini Science Planning Process," *Proceedings of SpaceOps 2004.*

[6]Paczkowski, B.G. and B. Larsen and T. Ray, "Managing complexity to maximize science return: Science planning lessons learned from Cassini," *IEEE Aerospace Conference*, March 2009. pp.1-14. *doi: 10.1109/AERO.2009.4839700.*

[7]Crisp, J. A. and M. Adler, J. R Matijevic,. S. W. Squyres, R. E. Arvidson, and D. M. Kass, "Mars exploration rover mission," Journal of Geophysical Research, 108:E12, p. 8061, 2003.

[8]Estlin, T., Bornstein, B., Gaines, D., Anderson, R.C., Thompson, D. R., Burl, M., Castano, R., and Judd, M., "AEGIS Automated Targeting for the MER Opportunity Rover," *ACM Transactions on Intelligent Systems and Technology,* 3:2, 2012.

[9]Castillo-Rogez, J. C., Pavone, M., Nesnas, I. A. D., Hoffman, J. A., "Observational Strategies for the Exploration of Small Solar System Bodies," *Proceedings of the IEEE Aerospace Conference, 2012.*

[10]Castillo-Rogez, J. C., Johnson, T.V., Thomas, P.C., Choukroun, M., Matson, D.L., Lunine J.I. "Geophysical evolution of Saturn's satellite Phoebe, a large planetesimal in the outer Solar System," *Icarus* 219:1 2012, p.86-109.

[11] Showalter, M. R., Hamilton, D. P., Stern, S. A., Weaver, H. A., Steffl, A. J., Young, L. A., New Satellite of (134340) Pluto: S/2011 (134340) 1, IAU Circ., 9221, 2011, 1, Edited by Green, D. W. E.

[12]Kelley, J. "Critical Path Planning and Scheduling: Mathematical Basis." *Operations Research*, 9:3, May–June, 1961.

[13]Odell, J. J. and H. V. D. Parunak, and B. Bauer, "Representing Agent Interaction Protocols in UML." *Agent-Oriented Software Engineering. Lecture Notes In Computer Science* 2001, Vol 2001/1957. p. 201-218.

[14]Johnston, M. D. "Multi-Objective Scheduling for NASA's Deep Space Network Array." in IWPSS-06. Baltimore, MD, Space Telescope Science Institute, 2006.

[15]Johnston, M. D. and Giuliano, M. E. "MUSE: The Multi-User Scheduling Environment for Multi-Objective Scheduling of Space Science Missions," IJCAI Workshop on Space Applications of AI, 2009. Pasadena, CA. July 17-19.

[16]Giuliano, M. and Johnston M. D. "Multi- Objective Evolutionary Algorithms for Scheduling the James Webb Space Telescope", In *International Conference on Automated Planning and Scheduling*, Sydney Australia, 2009. p.107-115.

[17]Gilmore, M. and D. R. Thompson, L. J. Anderson, N. Karamzadeh, L. Mandrake, R. Castano. "Superpixel segmentation for analysis of hyperspectral datasets, with application to CRISM data, M3 data and Ariadnes Chaos, Mars." *Journal of Geophysical Research*, Vol. 116, E07001, 2011.

[18]Thompson, D. R. and L. Mandrake, M. S. Gilmore and R. Castano, "Superpixel Endmember Detection." *Transactions on Geoscience and Remote Sensing*, 48(11): 4023-4033, Nov. 2010.

[19]Bunte, M. and D. R. Thompson, R. Castano, S. Chien, R. Greeley. Enabling Europa Science Through Onboard Feature Detection in Hyperspectral Images. *Lunar and Planetary Sciences Conference.* March 2011.

[20] A'Hearn M. F. et al., *Science* 310, 258 (2005); (10.1126/science.1118923).

[21] Riedel, J. E. and S. Bhaskaran, S. Desai, D. Han, B. Kennedy, G. W. Null, S. P. Synnott, T. C. Wang, R. A. Werner, E. B. Zamani. T. McElrath, D. Han, M. Ryne. "Autonomous Optical Navigation (AutoNav) DS1 Technology Validation Report. " JPL Publication 00-01, 2001. Jet Propulsion Laboratory, California Institute of Technology.

[22]Chien, S. and R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, and others, "Using autonomy flight software to improve science return on Earth Observing One," Journal of Aerospace Computing, Information, and Communication, vol. 2, 2005, pp. 196–216.

[23]Chien, S., Knight, R., Stechert, A., Sherwood, R., Rabideau, G. "Integrated planning and execution for autonomous spacecraft," IEEE Aerospace and Electronic Systems Magazine, 24:1. 2009. pp. 23—30.

[24]Hampton, D. and J. W. Baer, M. A. Huisjen, C. C. Varner, A. Delamere, D. D. Wellnitz, M. F. A'Hearn, K. P. Klaasen. "An Overview of the Instrument Suite for the Deep Impact Mission." *Space Science Reviews* (2005) 117: 43–93.

[25]Acton, C. H. et al., Ancillary data services of nasa's navigation and ancillary information facility, *Planetary and Space Science* 44:1, 2996, p. 65-70.

[26]Schulz, R., Sierks, H., Küppers, M., Accomazzo, A., Rosetta fly-by at asteroid (21) Lutetia: An overview, Planetary and Space Science, in press.

[27]Thomas, N. and 29 co-authors, The geomorphology of (21) Lutetia: Results from the OSIRIS imaging system onboard ESA's Rosetta spacecraft, Planetary and Space Science, in press.

[28]Barucci, MA, M. Fulchignoni and A, Rossi. Rosetta asteroid targets: 2867 Steins and 21 Lutetia, *Space Science Reviews* 126:1, 2007. p 67-78.