

Heuristic Search and Receding-Horizon Planning in Complex Spacecraft Orbit Domains

David A. Surovik and Daniel J. Scheeres

University of Colorado at Boulder
ECNT 320, 431 UCB
Boulder, CO, 80309, USA

Abstract

Spacecraft missions to small celestial bodies face sensitive, strongly non-Keplerian dynamics that motivate the employment of automated sampling-based trajectory planning. However, the scarcity of onboard computing resources necessitates careful formulation of heuristics for efficiently searching the reachable sets, which exhibit complex and finely-detailed structure. We examine a global search heuristic that combines aspects of simulated annealing and hill-climbing to locate sparse regions of the planning domain that simultaneously satisfy numerous geometric and timing constraints associated with remote sensing objectives for points of interest on the central body surface. Subsequently, we demonstrate the use of a receding-horizon implementation of this maneuver-planning strategy to produce mission profiles that fulfill sets of such goals.

Introduction

Traditional approaches to space mission design are rooted in Kepler's discovery that orbiting bodies travel along conic sections with the central gravitating body positioned at one focus. The integrability of Keplerian motion allows for any orbital state to be immediately associated with its conic section, an easily-described one-dimensional path through the six-dimensional orbit state space. These unchanging orbits can then be pieced together using impulsive thrust maneuvers to take a spacecraft from one desired state to another.

In many scenarios, the difference between the actual dynamical environment and the ideal Keplerian approximation is small: in Earth-orbiting missions, these discrepancies may be accurately modeled via linear perturbations and averaging theory; in preliminary design of interplanetary missions a "patched conic" approach, where each leg of the journey is modeled as a Keplerian problem with a different central body, can be applied effectively.

However, missions geared to the close study of small celestial bodies such as asteroids and comets — important targets for planetary science, planetary defense, and eventual resource exploitation — face constant exposure to strong perturbations from Keplerian motion, driven by highly non-spherical gravity fields and proportionally large third-body

gravitation and solar radiation pressure; in these conditions, conic sections are no longer an apt basis for even short-term trajectory design (Russell 2012; Scheeres 2012).

An alternate class of mission design strategies for non-Keplerian orbit environments exists: though the state space is no longer filled with conic sections, it nevertheless contains sparse sets of periodic orbits, which are bordered by stable and unstable manifolds that can be leveraged to travel between disparate orbit regimes using only minuscule expenditures of control energy (Parker and Anderson 2014). The complexity of these system dynamics has motivated automated mission design approaches that apply graph search techniques to exhaustive pre-computed databases of known periodic orbits and low-energy transfers (Tsirogiannis 2012; Trumbauer and Villac 2014).

Yet, the practicalities of small-body missions introduce further difficulties: a lack of accurate a priori system knowledge, when combined with large navigation uncertainties and dynamical sensitivity, can rapidly result in highly off-nominal conditions or even mission failure. Further, the matter of relating system dynamics to ultimate mission goals and constraints is equally nontrivial, and poses an increasingly difficult optimization problem when not conducted in a smaller and more restrictive precomputed dynamical framework (Noton 1995). The sum of these complications motivates yet another distinct approach to mission design: sampling-based planning in the form of an efficient heuristic search of the spacecraft's complexly structured reachable set (Komendera, Scheeres, and Bradley 2012).

Applied onboard in a receding-horizon fashion, this trajectory design strategy could allow flexible and opportunistic responses to off-nominal conditions and meet spaceflight autonomy requirements seen as pivotal to enabling ever more ambitious exploration missions (Wood et al. 2012; Pavone et al. 2014), in complement to the equally crucial development of autonomous planners for managing science instrument operations and other spacecraft subsystems on such endeavors (Knight et al. 2001).

Dynamical Model

The challenging non-Keplerian environments presented by the small Martian moon Phobos, subject of many mission design studies (Wallace et al. 2012), and asteroid Itokawa, visited by the Hayabusa mission (Scheeres et al. 2006),

are used to demonstrate the automated planning scheme. A spacecraft is modeled as a massless point, whose kinematic state \mathbf{x} is described within the rotating, body-fixed frame to facilitate natural expression of gravitation and science goals.

$$\mathbf{x} = (\mathbf{r}, \dot{\mathbf{r}}) = [x \ y \ z \ \dot{x} \ \dot{y} \ \dot{z}]^T \quad (1)$$

The equations of motion evolve the state through time via rotating-frame accelerations governed by the constant angular velocity vector $\boldsymbol{\omega}$, by the gravitational acceleration \mathbf{g} of the body's irregular shape, approximated as a triaxial ellipsoid, and in the case of Phobos also by the tidal accelerations \mathbf{g}_M resulting from the different Mars-relative altitudes of Phobos and the spacecraft (Scheeres 2012).

$$\ddot{\mathbf{r}} = -\boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{r}) - 2\boldsymbol{\omega} \times \dot{\mathbf{r}} + \mathbf{g}(\mathbf{r}) + \mathbf{g}_M(\mathbf{r}) \quad (2)$$

These equations produce trajectories through state space; next, various measures of objective fulfillment are taken.

Remote sensing of each point of interest (POI) on the body surface may occur when four topocentric constraints are satisfied. To observe the i -th POI, located at \mathbf{r}_i , the spacecraft must fall within prescribed bounds of azimuth θ , elevation φ , and range ρ relative to that point (shown as colored regions in figure 3) while the Sun elevation φ_s falls within prescribed bounds — a timing requirement. Since these constraints might only rarely be met simultaneously, they are each tracked continuously to aid the heuristic search for such occurrences.

For generic observation parameter a with ideal value a^* , the quality $q_a \in [0, 1]$ is a function of its actual, dropoff, and cutoff deviations: $\delta(a) = |a - a^*|$, δ_d , and δ_c respectively.

$$q_a(\delta) = \begin{cases} 0 & \text{if } \delta > \delta_c \\ \frac{\delta - \delta_d}{\delta_c - \delta_d} & \text{if } \delta_d < \delta < \delta_c \\ 1 & \text{if } \delta < \delta_d \end{cases} \quad (3)$$

This produces acceptable bounds $[a^* - \delta_d, a^* + \delta_d]$ outside of which the quality metric linearly decreases to zero. The set $\{q_a\}$ then defines the overall quality metric $Q_i \in [0, 1]$ using either the product L_Π or n-norm L_n of all q_a .

$$Q_i = L(q_\theta, q_\varphi, q_\rho, q_{\varphi_s}) \quad (4)$$

$$L \in \left\{ L_\Pi = \prod_j^{j_f} q_{a_j}, L_n = \left(\frac{1}{j_f} \sum_j^{j_f} q_{a_j}^n \right)^{1/n} \right\} \quad (5)$$

These functions can be used to alter the character of the gradients leading up to the sparse goal-fulfilling regions.

Duration g_i of successful observation of the i -th target is accumulated at a constant rate $\dot{g}_i = 1$ while all constraints are satisfied and the target value $g_{i,cap}$ has not yet been met:

$$\dot{g}_i(\mathbf{r}, t) = \begin{cases} 1 & \text{if } (Q_i == 1) \text{ and } (g_i < g_{i,cap}) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

All mission outcome data are stored within the vector \mathbf{y} , with each POI's parameters grouped together as \mathbf{g}_i

$$\mathbf{y}(t) = [\mathbf{g}_1 \ \mathbf{g}_2 \ \cdots \ \mathbf{g}_{n_g}] \quad (7)$$

$$\mathbf{g}_i(t) = [\{q(a_j)\} \ Q_i \ \hat{Q}_i(\tilde{\mathbf{x}}) \ \dot{g}_i \ g_i(\tilde{\mathbf{x}})] \quad (8)$$

where the historical maximum quality \hat{Q}_i , which is used to identify near misses of goal regions, and the accumulated observation time g_i both depend on the trajectory history $\tilde{\mathbf{x}}$.

Planning Domain

Despite the starkly atypical orbit environment, the classic control input paradigm of intermittent impulsive-thrust maneuvers remains most appropriate: a continuous-thrust approach would further exacerbate the curse of dimensionality in the search of the mission design space, while also complicating the spacecraft's ability to simultaneously operate science subsystems; a frequent-impulse scheme, i.e. inertial or body-fixed hovering, would incur excessive fuel costs over time. Furthermore, both of these alternatives would introduce undue noise into the process of orbit determination and system parameter estimation.

Thus, the planning domain for a single leg of the mission is given as $(\Delta t, \Delta \mathcal{V})$, where Δt is the time elapsed since previous maneuver's occurrence at t_{k-1} and $\Delta \mathcal{V}$ describes a velocity-space sphere of radius Δv_{max} centered at $\dot{\mathbf{r}}(t)$.

$$(\Delta t, \Delta \mathcal{V}) = \{ [\Delta t_{min}, \Delta t_{fail} - \Delta t_{cushion}] \} \times \cdots \\ \{ \Delta \mathbf{v} \in \mathbb{R}^3 \mid \|\Delta \mathbf{v}\| < \Delta v_{max} \} \quad (9)$$

The margins Δt_{min} and $\Delta t_{cushion}$ enforce practicality of plan implementation by restricting the temporal proximity of selected maneuvers to critical events: i.e., a sufficient duration must elapse to accommodate the planner's runtime between maneuvers. This same restriction, with additional built-in margin-of-error, is applied to the end of the temporal planning domain to allow sufficient time for action in avoidance of failure scenarios, e.g. impact of the central body. If failure does not occur within the prediction horizon Δt_{max} , which for this investigation is twice the Keplerian approximation of the orbit period at an altitude equivalent to the central body's mean radius, then $\Delta t_{fail} = \Delta t_{max}$.

The state-space dynamics F (Eq. 2) map from this planning domain to full sets of reachable trajectories $\tilde{\mathcal{X}}$ and subsequently to associated mission result sets \mathcal{Y} via the mission dynamics G (Eqns. 3–6; observation metrics given visually in figure 3). Lastly, an automated planner H must evaluate these results to assign scores $\mathcal{S}(t)$ that determine which action within the planning domain will be taken.

$$(t_k, \mathbf{x}(t_k) + \Delta \mathcal{V}) \xrightarrow{F} \tilde{\mathcal{X}}(t) \xrightarrow{G} \mathcal{Y}(t) \xrightarrow{H} \mathcal{S}(t) \quad (10)$$

The reachability map M used to inform trajectory planning can thus be described in condensed form as below, with the time horizon variably truncated to maximize score.

$$M : (\Delta t_s, \Delta \mathcal{V}; t_{k-1}, \mathbf{x}_{k-1}) \longrightarrow \mathcal{S} \quad (11)$$

As an illustrative example, we select an initial position \mathbf{r} at the first Lagrange point of the Mars-Phobos system, where their gravitational pulls precisely null out the acceleration of their co-rotating frame. Figure 1 plots a level surface within the map $(0, \Delta \mathcal{V}) \longrightarrow \hat{Q}_{max}$, which bounds a subdomain of maneuvers whose resultant trajectories come close to fulfilling remote sensing requirements of one or more targets. True goal regions, where $\hat{Q}_{max} = 1$, compose only a small fraction of the displayed region.

Heuristic Search

The nonintegrability of the dynamical model necessitates a sampling-based approach for charting the reachability map

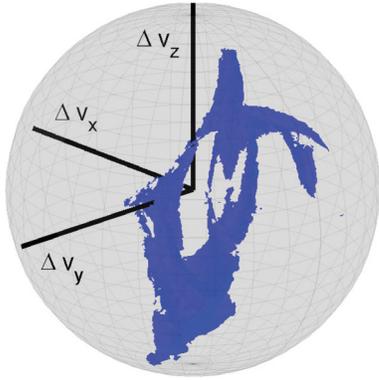


Figure 1: Occurrences of large \hat{Q}_{max} in planning domain.

M . A heuristic-based mesh refinement strategy is used to drastically improve the efficiency of this process relative to a simple grid search (Komendera, Scheeres, and Bradley 2012; Surovik and Scheeres 2015).

Given no a priori knowledge of map content, the first set of sampled maneuvers $\Delta \mathbf{v}_i \in \Delta \mathcal{V}$ is distributed evenly and numerically propagated over time horizon Δt_{max} through associated trajectories $\tilde{\mathbf{x}}_i(t) \in \tilde{\mathcal{X}}(t)$ to mission results $\mathbf{y}_i(t) \in \mathcal{Y}(t)$ and score trajectories $s(t) \in \mathcal{S}(t)$ which are lastly reduced to their maximal values and associated lifespans $(\hat{s}, \Delta t_{\hat{s}})$. Next, Delaunay Triangulation is used to build this small sample of points in M into a mesh, i.e. a set of tetrahedral “simplex” volume elements each bounded by four vertices, to allow inference of its continuous structure.

The search heuristic then operates on the mesh, assigning each j -th simplex a weight W_j based upon the extent to which the numerically propagated results of its vertices indicate that it contains relevant missing detail. Each simplex’s largest interior sphere is computed and its center and radius are used to define a 3D normal distribution from which an additional maneuver is sampled each time the simplex’s ID is drawn from a W -weighted random sample. Weights are assigned such that the random sampling of simplex IDs functions similarly to a sampling of the control domain via a volumetrically-specified probability density function.

$$W_j = V_j \left(\text{mean}_j \{ \Delta t \} \right) \left(\max_j \{ s + \hat{Q}_{max} \} \right)^{2^{8\tau}} \quad (12)$$

The first factor of the weight function, simplex volume V , accounts for uneven sample distribution while the second, the mean trajectory lifespan, scales sample probability density in accordance with the varying temporal depth of M . A final factor leverages the observation quality metric \hat{Q} to simulate a smooth gradient, whose shape depends on the choice of quality scalarizing function L , leading up to the potentially sparse plateaus of large s toward which the search is biased. This factor is raised to an exponent that varies with the search progress parameter τ , which increases from 0 to 1 as the sampling process iterates, progressively amplifying the bias toward high-scoring areas in a manner akin to simulated annealing.

Result: Monte Carlo mean results of heuristic searches of a planning domain at Itokawa are plotted below in figure 2 for three varieties of L , in addition to a heuristic that does not exploit quality gradients \hat{Q} and a standard grid search. Searches were seeded with 400 vertices before undergoing 50 progressively smaller refinement iterations, resulting in a final resolution of 2000 vertices within 30 seconds of runtime on a 2.2 GHz Intel Core 2 CPU.

All gradient-augmented heuristics performed well, exhibiting minor trade-offs between speed and score-attainment upper bound, associated with gradient sharpness. The gradient-free heuristic fared less favorably but nonetheless outperformed a standard grid search. Also plotted is a metric of sparseness for each augmented score, illustrating how favorable scores compose a vanishingly small fraction of the domain’s total volume. This indicates that final gains in performance are very hard-won, and shows how gradient-augmented heuristics smooth out the search for these gains.

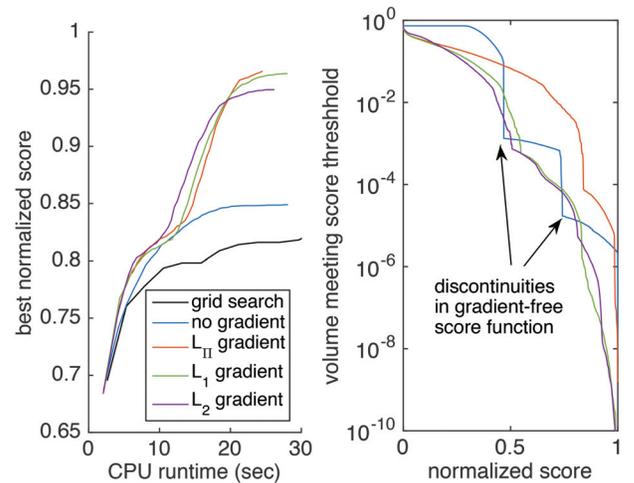


Figure 2: Left: Monte Carlo mean performance of search heuristics. Right: Sparseness of high-scoring regions.

Receding-Horizon Planner

The preceding demonstration shows how an automated planner can select an advantageous maneuver that completes a single science objective during a single planning cycle. To complete an entire mission that consists of many such objectives, a receding-horizon implementation of this scheme is applied (Surovik and Scheeres 2014), as has frequently been deemed appropriate for aerospace vehicle motion planning scenarios in highly dynamic environments (Lee, Longo, and Kerrigan 2012; Morgan, Chung, and Hadaegh 2014; Goerzer, Kong, and Mettler 2010).

Mission progress may be straightforwardly described with the mission objective function $s_m(t)$, a normalized measure of fulfillment of the set of n_g science goals:

$$s_m(t) = \frac{\sum_i^{n_g} g_i(t)}{\sum_i^{n_g} g_{i,cap}} \quad (13)$$

During each k -th planning cycle, a maximal increase $s_m(t) - s_m(t_{k-1})$ is desired. However, two additional subtleties must be accounted for in order to ensure effectiveness of the receding-horizon scheme: the timing of maneuvers, i.e. the selection of trajectory lifespan Δt in the planning domain, and the augmentation of the objective function to account for long-term consequences. The implementations of these two points are closely interconnected.

First, the mission score increase is combined with a *horizon promise function* $p_h(\mathbf{x}(t))$, similar in purpose to a cost-to-go function, which defines a state-space field associating each prospective final state of a sampled trajectory with the likelihood that it will produce a favorable velocity-space planning domain $\Delta\mathcal{V}$ for the next planning cycle. While the mission score time series plateaus when the spacecraft trajectory departs an observation region, the planning horizon's objective function $s(t) = s_m(t) - s_m(t_{k-1}) + p_h(\mathbf{x}(t))$ gives unique maxima \hat{s} that can be used unambiguously identify designate the optimal trajectory lifespan $\Delta t_{\hat{s}}$, subject to the bounds of practicality given in Eq. 9.

This process decouples the temporal and velocity-space planning operations, greatly lightening the computational burden by pruning one dimension of the search space. The matter of avoiding over-pruning of favorable solutions is synonymous with the selection of an effective p_h . The formulation presently employed,

$$p_h(\mathbf{x}) = q \left(\delta(r, 0), \max_i \{r_i + \rho_i + \delta_d(\rho_i)\}, r_{esc} \right) \quad (14)$$

enforces a preference for arc termination at altitudes in the range associated with the observation regions, helping to ensure that they do not appear vanishingly small in M . Each planning cycle thus uses a heuristic search to identify an available action and trajectory lifespan $(\Delta t, \Delta \mathbf{v})$ that maximize $s \in \mathcal{S}(t; \Delta\mathcal{V})$, providing a balance of short-term progress with long-term prospects. Repeated application of this scheme produces mission profiles that fulfill all goals.

Result: Figures 3 and 4 show body-frame and inertial-frame views of one such mission profile obtained via these methods for each observation mission. Additionally, a Monte Carlo analysis revealed the planning scheme to always fulfill 95% of objectives at Phobos without ever experiencing three successive planning horizons devoid of progress; when instead using $p_h = 0$, only 75% of trials met this success rate. Total propulsion costs of about 50 m/s for the sample mission were in line with the cost-per-day found by previous studies for far more restricted operations that did not approach the surface as closely (Wallace et al. 2012).

Conclusions

Despite the complex structure of reachable sets for spacecraft in strongly non-Keplerian orbit environments, sampling-based planning appears an effective tool for pursuing elusive opportunities to conduct science operations. Augmentation of a the objective function with a gradient-smoothing component allowed an adaptive mesh refinement strategy to effectively transition from coarse global search to fine local optimization via a simple mechanism akin to simulated annealing, producing a compelling advantage over a

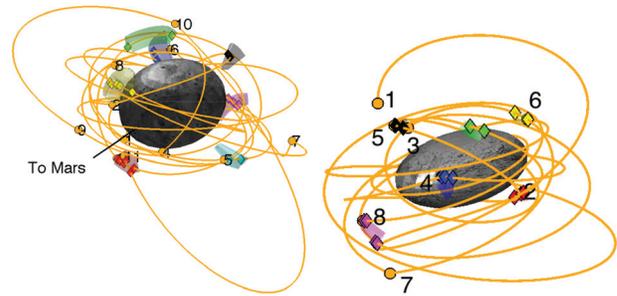


Figure 3: Typical mission solutions, shown in rotating body-fixed frame, for Phobos (left) and Itokawa (right). Numbered orange circles indicate maneuvers. Colored diamonds bound trajectory segments that permit observations.

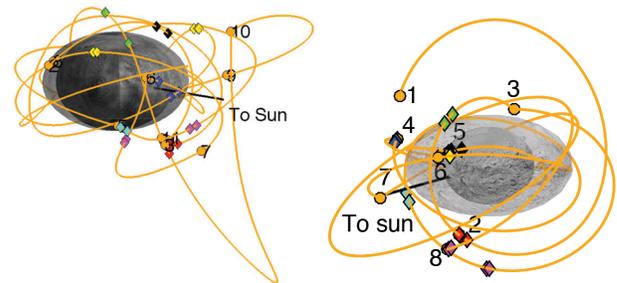


Figure 4: Inertial-frame view of above solution; sun-side location of diamonds (observation segments) illustrates fulfillment of timing-based surface lighting requirements.

simple grid search. The low computational demands of the demonstrative case suggest ample room for growth of complexity in the dynamical and mission goal models while still maintaining a footprint appropriate for onboard implementation.

Future Work

Alternate formulations of the horizon promise function $p_h(\mathbf{x})$ will be applied to potentially enhance planner performance by mapping and exploiting dynamical sensitivity, e.g. via Fast Lyapunov Indicators, to improve reachability prospects. The description of performance will also be extended beyond mere mission success by incorporating fuel cost and mission duration into the score functions.

Moreover, as the use of receding-horizon control was motivated not only by the complexity of the mission planning domain but also by the considerable levels of uncertainty to be expected in estimates of the spacecraft state and of dynamical model parameters at previously unvisited small-body systems, robustness to these sources of error will be a vital final aspect for demonstrating basic viability of this planning paradigm for autonomous onboard mission design applications. Incorporation of this aspect will be imperative before exhaustive tuning and characterization of algorithm performance can be effectively conducted.

Acknowledgments. This research is supported by NASA Space Technology Research Fellowship grant #NNX12AM40H.

References

- Goerzer, C.; Kong, Z.; and Mettler, B. 2010. A survey of motion planning algorithms from the perspective of autonomous uav guidance. *Journal of Intelligent and Robotic Systems* 57:65–100.
- Knight, R.; Rabideau, G.; Chien, S.; Engelhardt, B.; and Sherwood, R. 2001. Casper: space exploration through continuous planning. *Intelligent Systems, IEEE* 16(5):70–75.
- Komendera, E.; Scheeres, D. J.; and Bradley, E. 2012. Intelligent computation of reachability sets for space missions. *Proceedings of the 24th Conference on Innovative Applications of Artificial Intelligence*.
- Lee, D.; Longo, S.; and Kerrigan, E. C. 2012. Predictive control for soaring of unpowered autonomous uavs. *4th IFAC Nonlinear Model Predictive Control Conference* 4(1):194–199.
- Morgan, D.; Chung, S.-J.; and Hadaegh, F. Y. 2014. Model predictive control of swarms of spacecraft using sequential convex programming. *Journal of Guidance, Control, and Dynamics* 37(6):1725–1740.
- Noton, M. 1995. Orbital strategies around a comet by means of a genetic algorithm. *Journal of Guidance, Control, and Dynamics* 18(5):1217–1220.
- Parker, J. S., and Anderson, R. L. 2014. *Low-Energy Lunar Trajectory Design*. Wiley.
- Pavone, M.; Acikmese, B.; Nesnas, I. A.; and Starek, J. 2014. Spacecraft autonomy challenges for next generation space missions. In *Lecture Notes in Control and Information Systems*. Springer.
- Russell, R. P. 2012. Survey of spacecraft trajectory design in strongly perturbed environments. *Journal of Guidance, Control, and Dynamics* 35(3):705–720.
- Scheeres, D.; Gaskell, R.; Abe, S.; Barnouin-Jha, O.; Hashimoto, T.; Kawaguchi, J.; Kubota, T.; Saito, J.; Yoshikawa, M.; Hirata, N.; et al. 2006. The actual dynamical environment about itokawa. *AIAA paper* 6661.
- Scheeres, D. J. 2012. *Orbital Motion in Strongly Perturbed Environments*. Praxis Springer.
- Surovik, D., and Scheeres, D. J. 2014. Autonomous maneuver planning at small bodies via mission objective reachability analysis. *AIAA/AAS Astrodynamics Specialist Conference*.
- Surovik, D. A., and Scheeres, D. J. 2015. Adaptive reachability analysis to achieve mission objectives in strongly non-keplerian systems. *Journal of Guidance, Control, and Dynamics* 38(3):468–477.
- Trumbauer, E., and Villac, B. 2014. Heuristic search-based framework for onboard trajectory redesign. *Journal of Guidance, Control, and Dynamics* 37(1):164–175.
- Tsiogiannis, G. A. 2012. A graph based methodology for mission design. *Celestial Mechanics and Dynamical Astronomy* 114(4):353–363.
- Wallace, M.; Parker, J.; Strange, N.; and Grebow, D. 2012. Orbital operations for phobos and deimos exploration. In *AIAA/AAS Astrodynamics Specialist Conference, Guidance, Navigation, and Control and Co-located Conferences*. American Institute of Aeronautics and Astronautics.
- Wood, L. J.; Bhaskaran, S.; Border, J. S.; Byrnes, D. V.; Can-gahuala, L. A.; Ely, T. A.; Folkner, W. M.; Naudet, C. J.; Owen, W. M.; Riedel, J. E.; Sims, J. A.; and Wilson, R. S. 2012. Guidance, navigation, and control technology assessment for future planetary science missions, part 1: Onboard and ground navigation and mission design. Technical report, Strategic Missions and Advanced Concepts Office, NASA JPL.