

Stanford University



Final Report NASA NIAC Phase II Study, Grant #NNX14AT49G

SPACECRAFT/ROVER HYBRIDS FOR THE EXPLORATION OF SMALL SOLAR SYSTEM BODIES

Prepared for
Jason Derleth
Program Executive, NASA Innovative Advanced Concepts Program

Submitted December 15, 2017

Submitted by
Prof. Marco Pavone, PI
Department of Aeronautics and Astronautics
Stanford University
496 Lomita Mall, Room 261, Stanford, CA 94305-4035
Phone: (650)723 4432
Email: pavone@stanford.edu

Dr. Julie C. Castillo-Rogez, Co-I
Jet Propulsion Laboratory
California Institute of Technology
Email: julie.c.castillo@jpl.nasa.gov

Andreas Frick, Co-I
Jet Propulsion Laboratory
California Institute of Technology
Email: andreas.frick@jpl.nasa.gov

Prof. Jeffrey A. Hoffman, Co-I
Massachusetts Institute of Technology
Department of Aeronautics and Astronautics
Email: jhoffma1@mit.edu

Dr. Issa A. D. Nesnas, Co-I
Jet Propulsion Laboratory
California Institute of Technology
Email: issa.a.nesnas@jpl.nasa.gov

Executive Summary

This study investigated a mission architecture that allows the systematic and affordable in-situ exploration of small solar system bodies, such as asteroids, comets, and Martian moons (Figure 1). The architecture relies on the novel concept of spacecraft/rover hybrids, which are surface mobility platforms capable of achieving large surface coverage (by attitude-controlled hops, akin to spacecraft flight), fine mobility (by tumbling), and coarse instrument pointing (by changing orientation relative to the ground) in the low-gravity environments (micro-g to milli-g) of small bodies. The actuation of the hybrids relies on spinning three *internal* flywheels. Using a combination of torques, the three flywheel motors can produce a reaction torque in any orientation without additional moving parts. This mobility concept allows all subsystems to be packaged in one sealed enclosure and enables the platforms to be minimalistic. The hybrids would be deployed from a mother spacecraft, which would act as a communication relay to Earth and would aid the in-situ assets with tasks such as localization and navigation (Figure 1). The hybrids are expected to be more capable and affordable than wheeled or legged rovers, due to their multiple modes of mobility (both hopping and tumbling), and have simpler environmental sealing and thermal management (since all components are sealed in one enclosure, assuming non-deployable science instruments).

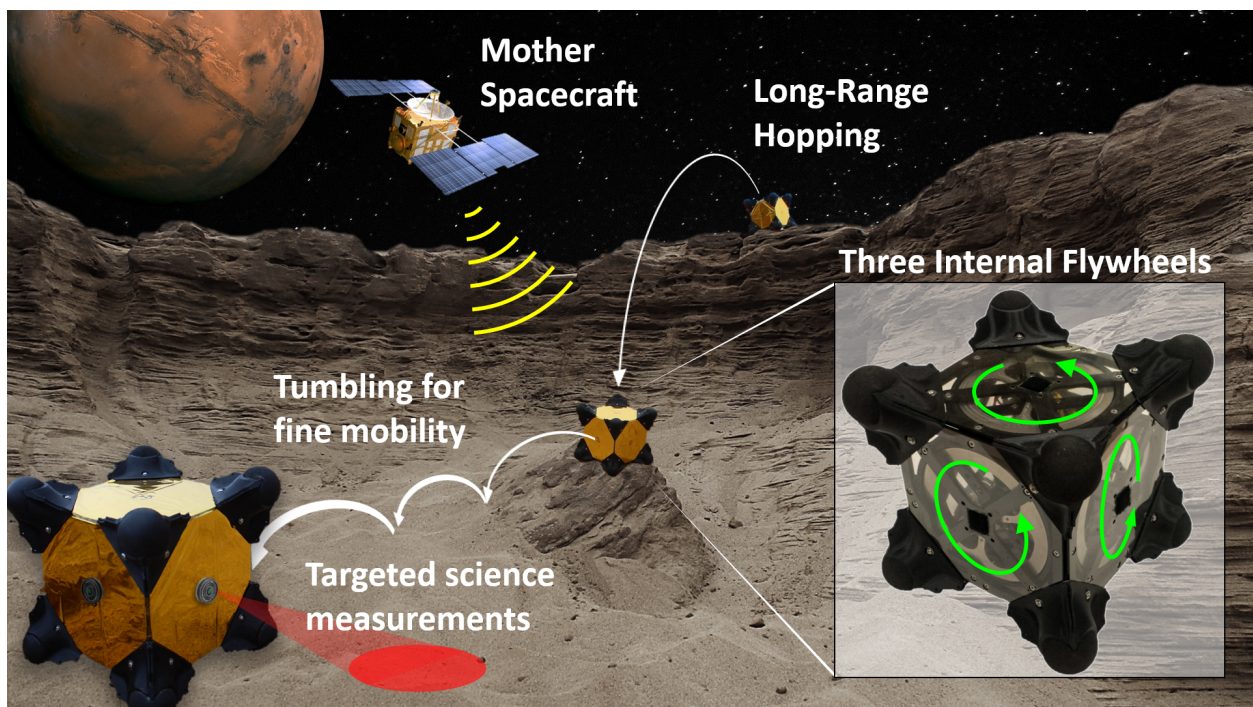


Figure 1: The mission architecture: one mother spacecraft would deploy on the surface of a small body one (or more) spacecraft/rover hybrids (from dm- to m-scale). Once deployed, the hybrids would perform attitude-controlled hops for long-range traverses (on the order of 10-100 m per hop) and would tumble to reach *specific* (as opposed to random) locations. Each hybrid is sealed in one enclosure and internally actuated through three mutually orthogonal flywheels. Synergistic mission operations would ensure precise planning and control of the hybrids, while keeping their end-to-end design minimalistic.

A NIAC Phase I project demonstrated that the bounding assumptions behind the pro-

posed mission architecture are reasonable, with a sound scientific and engineering basis. The NIAC Phase II had three main objectives:

1. To advance from Technology Readiness Level (TRL) 2 to TRL 3.5 the mobility subsystem of the hybrids¹, in particular, their mechanical design;
2. To advance from TRL 2 to TRL 3.5 the autonomy subsystem of the hybrids, in particular, motion planning and localization strategies.
3. To study at a conceptual level (TRL 2) system engineering aspects, in the context of a mission to Mars' moon Phobos.

The main results of this study can be summarized as follows:

- **Mobility subsystem:** The internally actuated mobility system of spacecraft/rover hybrids has been designed, analyzed, and tested in a relevant environment. Starting from a “first-principles” modeling approach, we proposed simplified dynamics models for studying various motion primitives, including hopping, tumbling, and twisting. These simplified models facilitated an analytical characterization of control strategies for *targeted* mobility. These models also allowed us to extract insights that inform the *design* of the mobility subsystem, for example, that a cubic chassis generally yields 45° hops or that the hop efficiency is closely related to the relative inertias of the flywheels and chassis, see Figure 2. From these insights, several prototypes have been developed, each one improving upon deficiencies of the previous. Finally, we spent a considerable effort to construct *first-of-a-kind* experimental environments for observing the dynamics that would be expected in microgravity; namely, a custom microgravity offloading test bed (Figure 3), and a parabolic flight campaign (Figure 4). Such experimental environments are of independent interest and might benefit the future development of microgravity mobility platforms, beyond those studied in this project. Experiments generally showed a strong agreement between observed trajectories and those predicted by our models.

There are several aspects of dynamics and control that would benefit from further study. For design, a more thorough trade study of the rover geometry is required. In particular, we only considered symmetric configurations, whereby the rover can maneuver from any orientation. However, it may be more advantageous to only have one larger flywheel for energetic hopping and smaller auxiliary flywheels for “self-righting” or reorientation prior to hopping. Second, more detailed models (supported by tests on terrain simulants) are required to understand the complex interactions of hopping on loose granular regolith.

- **Autonomy subsystem:** The semi-autonomous exploration of small bodies with hopping rovers requires both motion planning and localization. We have developed a motion planning framework that treats the hopping dynamics as a stochastic system and leverages offline simulations to generate closed-loop control policies that are both safe and efficient. In other words, although hopping dynamics may be subject to uncontrolled bouncing, we have shown that through a series of subsequent “corrective”

¹In the this report, the term “spacecraft/rover hybrid” and its abbreviated form, namely “hybrid,” will be used interchangeably.

maneuvers, a spacecraft/rover hybrid is able to achieve target precision on the order of one body length (roughly 20 cm), even when traversing large distances over highly irregular terrain (Figure 5) – *this is one of the key breakthroughs of this project*. We have also developed a collaborative localization and mapping approach between a primary spacecraft and a deployed hybrid, that uses visual SLAM techniques adapted from the state-of-the-art ORB-SLAM2 algorithm. We have tested the algorithm and demonstrated its capabilities to work in an asteroid-like environment, and in particular, the ability to localize the hybrid with cm-scale precision while remaining robust to changes in lighting conditions and large image scale changes during hops. Collectively, these results show that spacecraft/rover hybrids would indeed be capable of *targeted mobility, a key advancement over the state of the art that might enable a new class of scientific investigations*.

Future work should address two key aspects: (1) making the autonomy subsystem increasingly robust, for example, against modeling and localization errors and in the context of highly dynamic maneuvers (e.g., highly energetic hops), and (2) testing the overall autonomy subsystem (with planning and localization operating in closed-loop) in a high-fidelity simulation environment that captures cm-scale contact interactions as well as realistic environmental disturbances (e.g., dust).

- **Reference mission architecture: spacecraft and operations concept:** Considering a reference mission to Phobos, this study has determined that most subsystems of the hybrids (including the science payload) could be implemented with existing technologies leveraging recent developments driven by deep space CubeSat missions. In particular, key lifetime drivers, thermal and power, can be accommodated in a small form factor at levels commensurate with science-grade and data-intensive investigations. In this regard, a spacecraft/rover hybrid, and in particular the 25 kg prototype studied in this project (Figure 6), represents a counterpart to science-grade CubeSat platforms and could be considered as a secondary payload for future large planetary missions [1]. However, the minimalistic configuration of the hybrids will limit certain types of investigations, especially those that involve material sampling and processing, deployable elements (e.g., antennas for radar), and observations that bear stringent attitude requirements. This limitation is compensated by the prospect that a large mission (e.g., a Discovery class mission and beyond) could carry and deploy many low-mass hybrids for extended spatial coverage.

Key open questions left for future work include defining interfaces with the mothership, defining scenarios that leverage networking and coordination among multiple hybrids and with the mothership, and advanced mission design work on deployment and landing strategies.

In summary, this NIAC Phase II study has significantly increased the TRL of the mobility and autonomy subsystems of spacecraft/rover hybrids, and characterized system engineering aspects in the context of a reference mission to Phobos. Future studies should focus on improving the robustness of the autonomy module and further refine system engineering aspects, in view of opportunities for technology infusion.

This study led to seven publications:

- S. Chiodini, R. G. Reid, B. Hockman, I. A. D. Nesnas, S. Debei, and M. Pavone. Robust visual localization for hopping rovers on small bodies. In *Proc. IEEE Conf. on Robotics and Automation*, 2018.
- B. Hockman, R. G. Reid, I. A. D. Nesnas, and M. Pavone. Gravimetric localization on the surface of small bodies. In *IEEE Aerospace Conference*, 2018.
- B. Hockman and M. Pavone. Stochastic motion planning for hopping rovers on small solar system bodies. In *Int. Symp. on Robotics Research*, 2017.
- B. Hockman, A. Frick, I. A. D. Nesnas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. *Journal of Field Robotics*, 34(1):5–24, 2016.
- B. Hockman, R. G. Reid, I. A. D. Nesnas, and M. Pavone. Experimental methods for mobility and surface operations of microgravity robots. In *Int. Symp. on Experimental Robotics*, 2016.
- B. Hockman, A. Frick, I. A. D. Nesnas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. In *Field and Service Robotics*, 2015 (**best student paper award**).
- R. G. Reid, L. Roveda, I. A. D. Nesnas, and M. Pavone. Contact dynamics of internally-actuated platforms for the exploration of small Solar System bodies. In *Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, 2014.

Technical papers, related presentations, movies of the experiments, etc. can be found at the project’s website: <https://asl.stanford.edu/projects/surface-mobility-on-small-bodies/>. The investigators would like to acknowledge the contributions of Ben Hockman (Stanford, chief engineer for the hybrids), Alessandra Babuscia (JPL), Robert Reid (JPL), and several undergraduate and graduate students at Stanford and JPL whose internships focused on this project.




Figure 2: A spacecraft/rover hybrid prototype shown without avionics, covers, or solar panels. The cubic chassis encloses three orthogonal flywheels and is surrounded by eight compliant spikes on its corners.




Figure 3: Left: 6 DoF microgravity test bed CAD rendering. The powered gantry tracks the translational motion of the hybrid in x , y , and z within a volume of $3\text{ m} \times 1\text{ m} \times 1\text{ m}$ respectively, while allowing for free fall in z at sub-milli-g levels. The gimbal frame allows the hybrid to rotate in all three axes. Right: Image of the test bed.




Figure 4: Experimental setup for parabolic flight campaign. Left: The hybrid (A) is held in place on the test surface (C) by an actuated arm (B). An array of five cameras (D) capture its motion as it hops within the container. Right: Photo of our experiments on NASA's C9 aircraft.




Figure 5: Three rollouts of a learning-based, closed-loop control policy for targeted mobility. The surface color map shows the value function Q^* under such a policy, roughly representing the difficulty to reach the target region from any given point on the small body (the higher Q^* is, the easier is to reach the target region).




Figure 6: Example configuration for a 25 kg spacecraft/rover hybrid.

Contents

1	Introduction	9
2	Spacecraft/Rover Hybrid: Mobility Subsystem	12
2.1	Dynamics and Control	13
2.1.1	Hopping Dynamics and Control	14
2.1.2	Tumbling Dynamics and Control	20
2.1.3	Twisting Dynamics and Control	21
2.2	Mobility System Design	22
2.3	Mobility Experiments	23
2.3.1	Experiments in Custom 6 DoF Microgravity Test Bed	24
2.3.2	Parabolic Flight Experiments	28
2.4	Summary	31
3	Spacecraft/Rover Hybrid: Autonomy Subsystem	33
3.1	Motion Planning	34
3.1.1	Preliminaries	35
3.1.2	Single-Hop Planning	37
3.1.3	Sequential Hop Planning	41
3.2	Localization	46
3.2.1	Related work	47
3.2.2	Proposed Localization Method	48
3.2.3	Experimental Evaluation	50
3.3	Summary	55
4	Reference Mission Architecture: Spacecraft and Operations Concept	57
4.1	Scalability	57
4.2	Reference Design Overview	58
4.3	Subsystem Descriptions	59
4.3.1	Mechanical Configuration and Mass Equipment List	59
4.3.2	Electrical Power System	59
4.3.3	Telecom Subsystem	61
4.3.4	Thermal Design	62
4.3.5	Guidance, Navigation, and Control	62
4.3.6	Command & Data Handling (C&DH)	62
4.4	Concept of Operations	63
4.5	Example Timeline	65
4.6	Summary	66
5	Conclusions	67
6	References	68

1 Introduction

This project targets the robotic in-situ exploration of small bodies (such as asteroids and Mars’ moons) whose scientific significance pertains to all three science themes highlighted in the recent Decadal Survey Report [9] as well as to the goals of the Human Exploration program. While some measurements of small bodies can be obtained with remote platforms, such as space telescopes, several other measurements require close proximity to the surface at *multiple* locations for an *extended* period of time (e.g., characterization of regolith shape and porosity via microscopic observations, ground-truth thermal measurements, and observation of surface electrostatic field) [10]. This is also the case for precursor science enabling human exploration, which requires the characterization of regolith mechanical properties, dust dynamics, electrostatic charging, etc. [11]. Hence, in-situ exploration of small bodies at multiple designated locations is a critical need in the scientific community and requires surface mobility. On the other hand, low gravitational fields, which are characteristic of small bodies, hamper the adoption of “traditional” mobility platforms, which are designed for bodies with significant gravity like Mars or Earth’s Moon [10, 12, 13]. This calls for the development of new technologies for both *targeted* surface mobility and surface operations.

In general, mission architectures for the targeted in-situ exploration of small bodies can be *monolithic*, i.e., involving a single spacecraft, or *mother-daughter*, where a mothership deploys one or several daughter mobility platforms on the body. In monolithic architectures, a spacecraft lands multiple times by firing its thrusters, as in the Comet Hopper mission architecture [14]. Such approach only allows for limited *discrete* sampling, versus spatially dense sampling which is critical to understanding, e.g., the nature of the interface between two spectral units. Also, this approach might lead to surface contamination (which would severely limit science examination of volatile content), and might involve high risks during each surface sortie. This motivates the investigation of mother-daughter architectures. Mobility mechanisms for the daughter platforms can be divided into four classes: mobility via thrusters, wheeled mobility, legged mobility, and hopping mobility. Mobility via thrusters tends to be quite complex and has the disadvantages of limited sampling and surface contamination, as discussed above. Because of low friction, wheeled vehicles are limited to extremely low speeds (1.5mm/s per previous JPL studies [15]). Legged systems are quite complex in low-gravity environments. For example, anchoring the legs would highly depend on soil properties that are largely unknown. Alternatively, hopping systems use the low-gravity environment to their advantage. Space agencies such as NASA [15, 16], RKA [17], ESA [18], and JAXA [19] have all recognized this advantage and have designed a number of hopping prototypes. However, such prototypes do not appear to allow precise traverses to designated targets in low gravity environments.

This report describes a mission architecture that allows the systematic and affordable in-situ exploration of small solar system bodies (Figure 7) – a significant advancement with respect to the state of the art. The architecture relies on the novel concept of spacecraft/rover hybrids, which are surface mobility platforms capable of achieving large surface coverage (by attitude-controlled hops, akin to spacecraft flight), fine mobility (by tumbling), and coarse instrument pointing (by changing orientation relative to the ground) in the low-gravity environments (micro-g to milli-g) of small bodies. The actuation of the hybrids relies

on spinning three *internal* flywheels. Using a combination of torques, the three flywheel motors can produce a reaction torque in any orientation without additional moving parts. This mobility concept allows all subsystems to be packaged in one sealed enclosure and enables the platforms to be minimalistic. The hybrids would be deployed from a mother spacecraft, which would act as a communication relay to Earth and would aid the in-situ assets with tasks such as localization and navigation (Figure 7). The hybrids are expected to be more capable and affordable than wheeled or legged rovers, due to their multiple modes of mobility (both hopping and tumbling), and have simpler environmental sealing and thermal management (since all components are sealed in one enclosure, assuming non-deployable science instruments).

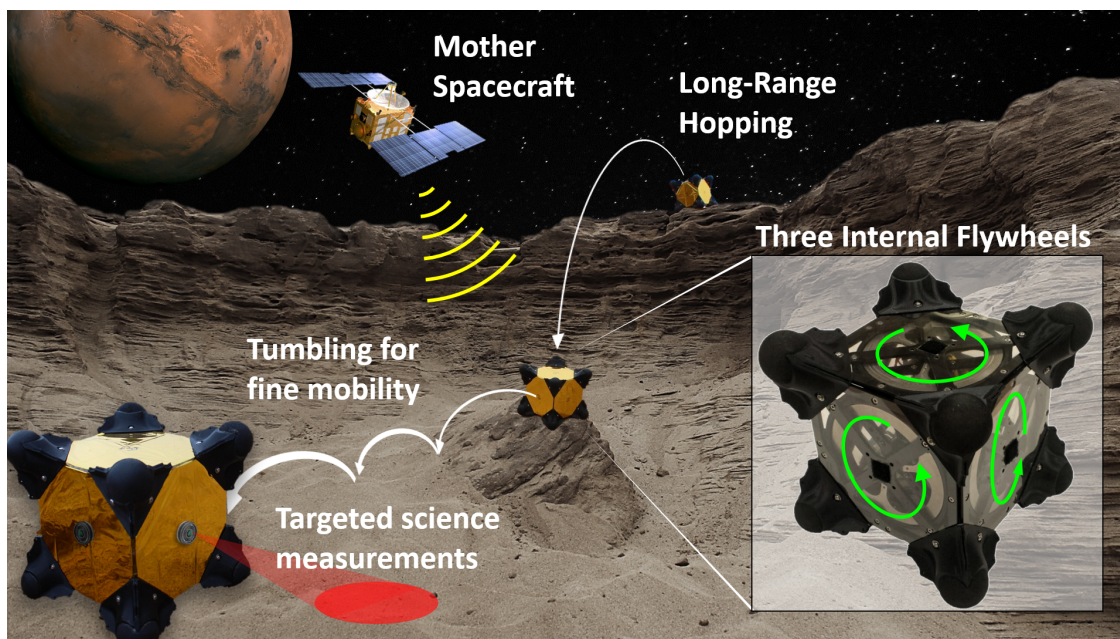


Figure 7: The mission architecture: one mother spacecraft would deploy on the surface of a small body one (or more) spacecraft/rover hybrids (from dm- to m-scale). Once deployed, the hybrids would perform attitude-controlled hops for long-range traverses (on the order of 10-100 m per hop) and would tumble to reach specific locations. Each hybrid is sealed in one enclosure and internally actuated through three mutually orthogonal flywheels. Synergistic mission operations would ensure precise planning and control of the hybrids, while keeping their end-to-end design minimalistic.

A NIAC Phase I project demonstrated that the bounding assumptions behind the proposed mission architecture are reasonable, with a sound scientific and engineering basis. The NIAC Phase II had three main objectives:

1. To advance from Technology Readiness Level 2 to TRL 3.5 the mobility subsystem of the hybrids, in particular, their mechanical design.
2. To advance from TRL 2 to TRL 3.5 the autonomy subsystem of the hybrids, in particular, motion planning and localization strategies.
3. To study at a conceptual level (TRL 2) system engineering aspects, in the context of a mission to Mars' moon Phobos.

This report is structured as follows. In Section 2 (corresponding to objective #1, and based on papers [5, 6, 7, 8]) we discuss the mechanical design of the hybrids, their dynamics and control to achieve a desired hopping velocity vector, and experimental results in various microgravity test beds, including a custom 6-degree-of-freedom offloading gantry and a parabolic flight campaign. In Section 3 (corresponding to objective #2, and based on papers [2, 3, 4]) we present an algorithmic framework for motion planning and on-board localization for the hybrids, along with experimental results validating the proposed techniques. In Section 4 (corresponding to objective #3) we present trade studies and system engineering analyses in the context of a reference mission to Phobos, which enable estimates of critical subsystems such as electrical, telecom, and thermal. Finally, in Section 5, we draw our conclusions and discuss directions for future work.

2 Spacecraft/Rover Hybrid: Mobility Subsystem

A spacecraft/rover hybrid is a small (≈ 0.25 m geometrical diameter, ≈ 15 kg even though the design is scalable) multi-faceted geometric solid that encloses three mutually orthogonal flywheels and is surrounded by external spikes or specialized contact surfaces (see Figure 8, where we consider a cube geometry). Specifically, there is no external propulsion. The combination of the flywheels with the enclosure- and spike-geometry enables controlled tumbles, hops, and high-altitude ballistic flight.

The basic principle behind internally-actuated mobility is the conservation of angular momentum, which ensures that angular momentum can be swapped between the platform and the flywheels. Specifically, a flywheel consists of a spinning mass with a substantial amount of inertia. Due to the presence of the flywheels, the total angular momentum of the platform is given by (vectors and matrices are represented in boldface):

$$\mathbf{H} = \mathbf{I}_{\text{platform}} \boldsymbol{\omega}_{\text{platform}} + \sum_{i=1}^3 \mathbf{I}_{\text{flywheel},i} \boldsymbol{\omega}_{\text{flywheel},i},$$

where \mathbf{I} denotes the inertia matrix and $\boldsymbol{\omega}$ denotes the angular velocity vector. Since, in absence of *external* torques, the total angular momentum stays constant, by controlling the *internal* torque between the flywheels and the platform one can control both magnitude and direction of the angular rotation of the platform. In turn, this angular rotation can give rise to (controllable) surface reaction forces at contact points, which lead to either tumbling (i.e., pivoting around a tip) or hopping (when the reaction forces are large enough). The JAXA's MINERVA *hopper* included a related actuation mechanism (specifically, a *single* flywheel mounted on a turntable), which, however, did not allow for precise traverses to designated targets. Unfortunately, MINERVA did not succeed during its deployment [19].

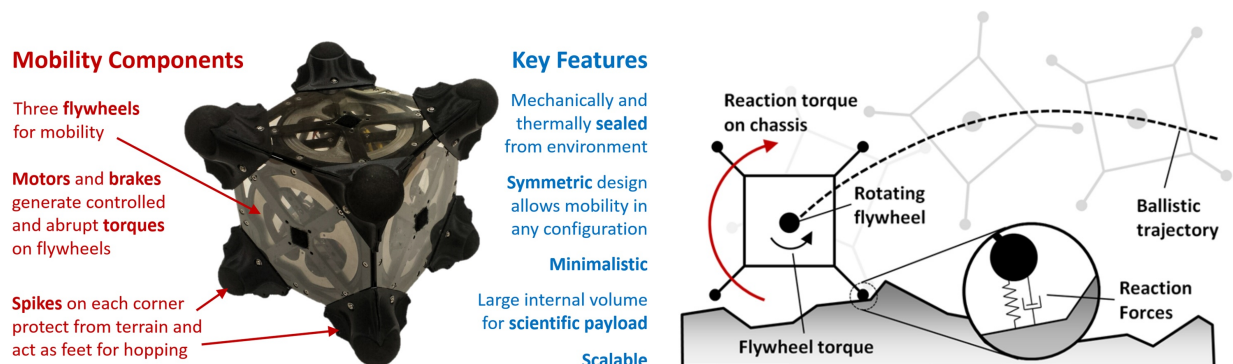


Figure 8: Left: A spacecraft/rover hybrid prototype shown without avionics, covers, or solar panels. The cubic chassis encloses three orthogonal flywheels and is surrounded by eight compliant spikes on its corners. **Right:** By accelerating internal flywheels, surface reaction forces cause the rover to tumble or hop.

Achieving targeted mobility with such a hopping system involves a number of important considerations. This section focuses on the platform itself, and in particular, the dynamics and control for achieving a desired hopping velocity vector (Section 2.1), considerations for the mechanical design of the mobility subsystem (Section 2.2), and experimental results in various microgravity test beds, including a custom 6-degree-of-freedom offloading gantry and a parabolic flight campaign (Section 2.3).

2.1 Dynamics and Control

When in free flight, the hybrid simply behaves as an attitude-controlled spacecraft, for which the dynamics are well-known. A number of interesting behaviors arise, however, when the hopper actuates its flywheels *while in contact with the surface* of a small body. In general, the three orthogonal flywheels allow for any arbitrary torque to be applied to the hopper—large abrupt torques via mechanical brakes and smaller, more controlled torques via motors. It is convenient to *classify* the rather large space of possible maneuvers into a set of *motion primitives*.

		Direction with respect to surface		
		Parallel	Angled	Normal
Magnitude	Low	Tumble	Twisting tumble	Twist
	High	Hop	Twisting hop	Tornado

Table 1: Motion primitives classified by magnitude and direction of angular velocity.

Table 1 proposes a classification of motion primitives based on the magnitude and direction of the hopper’s angular velocity. Focusing on motion primitives allows for closed-form insights to be derived from simplified analytical models. In particular, three of the six motion primitives are especially useful for mobility:

- **Hopping** occurs when sufficient force is generated against the surface to launch the hopper forward in a ballistic trajectory, anywhere from one body-length to hundreds of meters. Thus, it is the primary mode of mobility for large distance coverage.
- **Tumbling** is simply a less energetic form of hopping, whereby the hopper rotates about a pair of spikes without losing ground contact. Tumbling is less likely to induce random bouncing, and thus offers a much more precise, albeit slower, form of mobility.
- **Twisting** is not a direct form of mobility, but rather a way of changing orientation by turning in place. Such a maneuver can help to point instruments or enable subsequent hops/tumbles from a different orientation.

Hopping and tumbling about axes parallel to the surface is more controllable than twisting hops/tumbles about arbitrary inclined axes (as we will discuss in Section 2.1.1) and are more amenable to an analytical treatment. A “Tornado” maneuver (inspired by the way in which it flings nearby surface regolith) is an energetic twist that typically also hops up due to surface imperfections—not a particularly controllable motion, but potentially useful for escaping from being stuck or buried. In this section we study the dynamics and control of these three primary motion primitives, establishing a suite of *controllable* maneuvers whose concatenation can provide targeted mobility.

2.1.1 Hopping Dynamics and Control

As the primary mode of mobility, understanding the fundamental dynamics that govern *hopping* trajectories is a critical step towards characterizing the hopper’s mobility capabilities. In this section we consider two analytical models to derive useful control laws, (1) first examining a simplified 2D model for the case of single-axis hops, and (2) a more general 3D model of the hopper pivoting about only one of its spikes.

Single Axis Hopping

In the case of hopping and tumbling about a single axis (i.e., with just one flywheel), the hopper pivots about a *pair* of spikes, allowing for a *planar* (2D) representation of the dynamics. The hopper is modeled as a disk with equispaced rigid spikes attached to it, similar to the model commonly used in the field of passive dynamic walking. At the center of mass, a motor drives a single flywheel, producing a torque on the hopper (see Figure 9). This model can be thought of as a 2D vertical slice of the hopper, orthogonal to the applied torque.

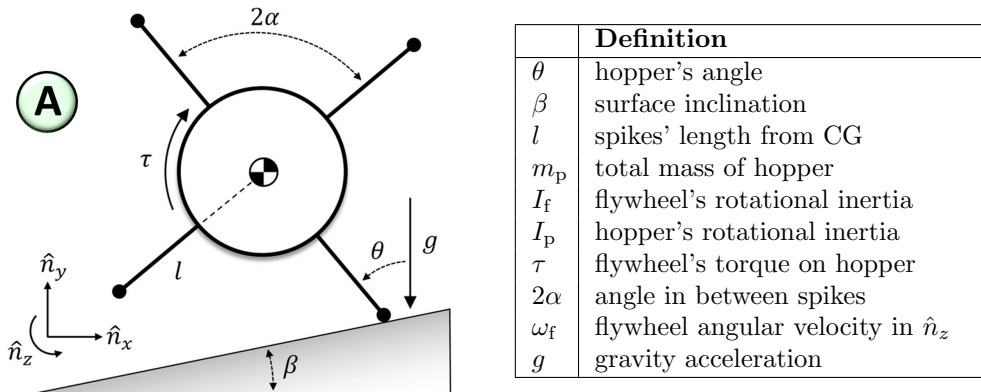


Figure 9: 2D model: A hopper is modeled as a rigid body that pivots on an inclined surface. This planar representation is also the basis for the tumbling analysis in Section 2.3.2.

The key assumption in our study is that the stance spike acts as a pin joint and does not slip. Under this assumption, the 2D model of the hopper is uniquely described by two states, θ and $\dot{\theta}$. See Figure 9 for a detailed description of all parameters. One can show that the no-slip assumption is a reasonable approximation for coarse spike geometries where $(\theta - \beta) > \tan^{-1}(1/\mu_d)$, where μ_d is the coefficient of dynamic friction. For the rubber spike tips on our current prototype, $1 < \mu_d < 1.5$, which, as validated via simulations in Section 2.1.1 and experiments in Section 2.3, is high enough to justify this no-slip assumption. This assumption, however, would not hold in cases where the hopper operates on non-rigid surfaces (i.e., loose regolith), whereby the slip properties are governed by frictional interactions with granular media. This aspect is left for future research.

A hopping maneuver consists of a stride phase, when the system is supported by a single stance spike (in the 2D world), and a flight phase when the stance spike leaves the ground. We study the flywheel’s torque needed to cause the hopper to hop to the right at a desired speed v_h and angle θ_h (the subscript “h” denotes quantities evaluated at the hopping instant).

Assume that the hopper starts at rest on the inclined surface and applies a sufficient torque $\tau(t)$ to induce a clockwise rotation about its stance spike. For the stride phase (i.e., before ground contact is lost), the equations of motion are those of an inverted pendulum and can be easily written as

$$\ddot{\theta}(t) = \frac{m_p g l \sin \theta(t) - \tau(t)}{I_p + m_p l^2}. \quad (1)$$

The first thing to note is that there exists a minimum torque that will cause the hopper to initiate a clockwise rotation (i.e., $\ddot{\theta} < 0$) from rest:

$$\tau_{\min} = m_p g l \sin(\alpha + \beta). \quad (2)$$

For typical gravity levels of interest ($10 \mu g - 1000 \mu g$), small motors of only a few Watts would be sufficient to exceed this torque. Secondly, by studying the free body diagram of the system, one can readily show that in order to obtain a negative normal force (i.e., loss of ground contact) it is required that

$$|\dot{\theta}(t_h)| > \sqrt{\frac{m_p g \cos \beta + \frac{\tau(t_h)}{l} \sin(\theta(t_h) - \beta)}{m_p l \cos(\theta(t_h) - \beta)}}. \quad (3)$$

For a level terrain (i.e., $\beta \approx 0$) and with no input torque, $|\dot{\theta}(t_h)|_{\min} = \sqrt{g/[l \cos \theta(t_h)]}$, which corresponds to a *minimum hop distance* on the order of $2l$.

In this section we study a control strategy that leverages (2) by slowly spinning up the flywheels with motor torque $\tau < \tau_{\min}$, such that the hopper remains grounded. When the desired flywheel speed is achieved, a high-torque mechanical brake (discussed in Section 2.2) is applied and a hop is initiated. This approach is attractive as it is simple, does not cause momentum build up in the flywheels, and rapidly transfers energy for more aggressive hops. Therefore, a control strategy of particular interest assumes “instantaneous” momentum transfer from the flywheel to the hopper².

By equating the initial angular momentum of the flywheel ($I_f \omega_f$) to the resulting angular momentum of the hopper about the spike tip ($\dot{\theta}(I_p + m_p l^2)$), and assuming that a hop is initiated immediately after momentum transfer (i.e., $v_h = l\dot{\theta}(0^+)$), the resulting hop velocity, angle, and lateral distance are given by, respectively,

$$v_h = l\omega_f \left(\frac{I_f}{I_p + m_p l^2} \right), \quad \theta_h = \alpha + \beta, \quad d_h = \frac{v_h^2}{g} \sin(2\theta_h). \quad (4)$$

A few interesting observations can be made from these results. First, in this regime, the hop angle is governed exclusively by the spike geometry and surface inclination. To maximize the lateral distance of the parabolic trajectory (which scales as $\sin(2\theta_h)$), a 45° hop is desired. This is one of the reasons why our current design has a cubic spike geometry (i.e., $\alpha = 45^\circ$), see Section 2.2. Second, we define the energy transfer efficiency as

$$\eta := \frac{E(t^+)}{E(t^-)} = \frac{I_f}{I_p + m_p l^2}, \quad (5)$$


²The experimental results in Section 2.2 validate this assumption.

where $E(t^-)$ is the energy of the system just before actuation (flywheel kinetic energy), and $E(t^+)$ is the energy just after actuation (hopper's kinetic energy). Interestingly, the efficiency is given by the ratio of flywheel inertia to hopper's inertia *about the spike tip*, which depends quadratically on the length of the spikes. Hence, there is an important trade-off between the capability of negotiating obstacles (that would require long spikes) and the actuation efficiency (that prefers short spikes). For our current prototype (augmented with dead mass as stand-in for scientific payload), $\eta \approx 0.01$. This result is critically enabled by angular momentum arguments.


Now, Eqs. (4) and (5) can be combined to develop an expression for the flywheel speed (ω_f) required to cover a lateral distance d_h :

$$\omega_f(d_h) = \sqrt{\frac{d_h g}{\eta^2 l^2 \sin(2(\alpha + \beta))}}. \quad (6)$$

For a square geometry, this expression is minimized for flat terrain, but tends towards infinity as $\beta \rightarrow 45^\circ$. This motivates the potential utility of controllable friction brakes, which can extend the duration of the stride phase and thus control the hop angle.



(a) Hopping angle (θ_{hop}) as a function of input torque ($\bar{\tau}$). The horizontal line marks the 45° ‘‘ideal’’ hop angle.



(b) Lateral hop distance (d_h) as a function of input torque ($\bar{\tau}$).

Figure 10: Resulting hop angles and distances as functions of input torque for three initial flywheel speeds: $\omega_f = 2000, 5000,$ and 10000 rpm (the x -axis is in logarithmic scale). Each curve corresponds to a particular surface inclination β . The vertical line on each graph marks the minimum torque at which the flywheel can be fully stopped before a hop is initiated (see Eq. (3)). Results are based on Phobos' gravity level (0.0058 m/s^2) and parameters of our prototype (see Section 2.2)

In the more general case when momentum transfer is *not* assumed to be instantaneous, one can regard the initial flywheel speed (ω_f) and constant braking torque ($\bar{\tau}$) as the *two*

control variables. In bringing the flywheel to a full stop, the control variables are related by $\bar{\tau}\Delta t = I_f\omega_f$, where Δt is the time duration of braking. In the limit as $\Delta t \rightarrow 0$, the impulsive torque corresponds to the case of instantaneous momentum transfer. To study the case when Δt is finite, the nonlinear differential equations of motion given by (1) must be solved numerically. However, for aggressive hops, one can assume that $\bar{\tau} \gg m_p g l \sin(\theta)$, so (1) can be well approximated by the linear second order ordinary differential equation, $\ddot{\theta}(t) \approx -\bar{\tau}/(I_p + m_p l^2)$. For high enough torques, the hop criterion in (3) is not met until immediately after actuation (i.e., at $t_h = \Delta t = \omega_f I_f / \bar{\tau}$), so the hop velocity, angle, and resulting distance can be determined by integration:

$$v_h = l\eta\omega_f, \quad \theta_h = \alpha - \frac{\eta I_f \omega_f^2}{2\bar{\tau}}, \quad d_h = \frac{v_h^2 \sin(2\alpha - \eta I_f \omega_f^2 / \bar{\tau})}{g}. \quad (7)$$

Solving for the control inputs, $\bar{\tau}$ and ω_f , now requires solving this set of nonlinear algebraic equations. To better visualize these results and validate the pivoting assumptions, numerical simulations were generated based on a full 6 DoF model, including normal spring/damper and tangential Coulomb friction contact forces.

The plots in Figure 10 illustrate the hopping angle and distance that result from a constant braking torque, $\bar{\tau}$. Each plot represents a different flywheel speed (2000, 5000, and 10000 rpm). The kink in each curve marks the threshold of an “early hop”—the torque level ($\bar{\tau}_s$) below which surface contact is lost before the flywheel is fully stopped. In other words, for a given flywheel speed, $\bar{\tau}_s$ is the minimum braking torque that should be applied to convert all of the flywheel’s available kinetic energy to forward motion. This threshold (marked by a vertical line) is in very close agreement with predictions based on (3).

Figure 10(b) shows that for $\beta \leq 0$ (i.e., hopping downhill), travel distance increases as the torque is increased. However, the situation is different when considering inclined poses ($\beta \geq 0$), whereby high torque inputs result in high angle arching hops—an undesirable effect for distance coverage but potentially useful for getting out of pits. The peaks in these distance curves are in agreement with (6).

Hopping Directional Control

The analytical model presented in Section 2.1.1 is well suited for *single-axis* hops because the out-of-plane degree of freedom is locked by having two spikes in contact with the surface. For a cubic design, this only allows for hops in 90° directional increments, analogous to a “Manhattan-style” mobility approach. There are two ways in which a finer resolution of directional control can be achieved: (1) using a twist maneuver as discussed in Section 2.1.3 to point a primary axis towards the desired heading and execute a single-axis hop, or (2) hopping directly over a single spike about an “oblique” axis. The later approach is faster, as it avoids an intermediate twist maneuver and may also be useful in situations where an accurate twist maneuver is difficult due to terrain roughness. In this section we will consider a 3D model of the hopper to study the dynamics and control of such oblique hops (see Figure 11).

In general, when pivoting about a single spike, the hopper behaves like a ball-jointed inverted pendulum with three degrees of freedom and complicated nonlinear dynamics. However, just as in Section 2.1.1, the assumption of instantaneous momentum transfer (i.e., high

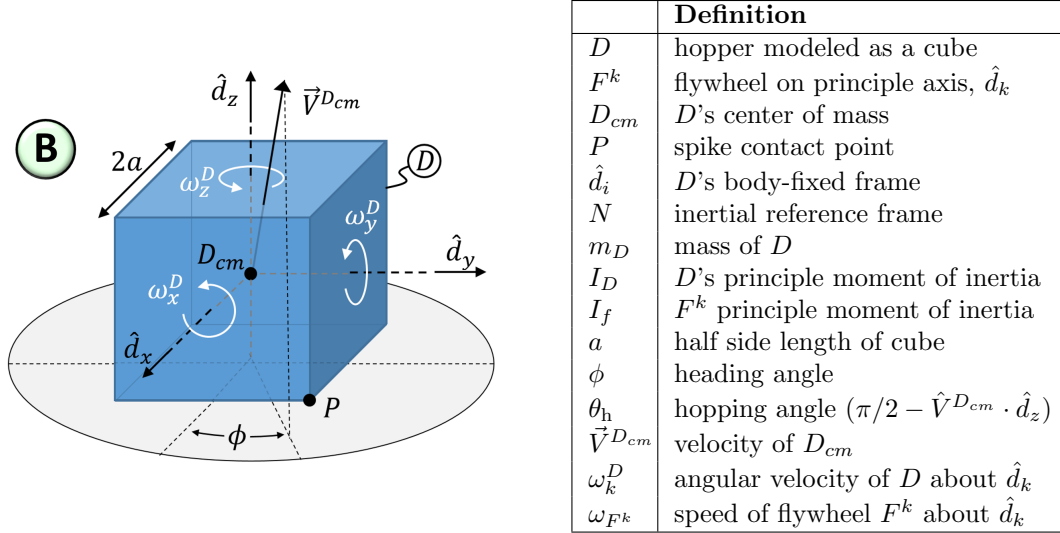


Figure 11: 3D model: A hopper is modeled as a uniform cube that pivots about a corner.

friction brakes) can again be leveraged to collapse the dynamics of interest to an instant in time. In this way, conservation of angular momentum can again be invoked (this time about the single spike contact) to relate the hopper's rotation rate (${}^N\vec{\omega}^D$) immediately after braking, to the initial flywheel speeds (ω_{F^k}):

$$\left(\sum_{k=x,y,z} {}^N \vec{H}^{F^k/P} \right) = \left(\vec{I}^{D/D_{cm}} + m_D \left[(\vec{r}^{P/D_{cm}} \cdot \vec{r}^{P/D_{cm}}) \mathbf{I} - (\vec{r}^{P/D_{cm}})(\vec{r}^{P/D_{cm}})^T \right] \right) \cdot {}^N \vec{\omega}^D,$$

$$\begin{bmatrix} I_f & 0 & 0 \\ 0 & I_f & 0 \\ 0 & 0 & I_f \end{bmatrix} \begin{bmatrix} \omega_{F^x} \\ \omega_{F^y} \\ \omega_{F^z} \end{bmatrix} = \begin{bmatrix} I_D + 2m_D a^2 & -m_D a^2 & m_D a^2 \\ -m_D a^2 & I_D + 2m_D a^2 & m_D a^2 \\ m_D a^2 & m_D a^2 & I_D + 2m_D a^2 \end{bmatrix} \begin{bmatrix} \omega_x^D \\ \omega_y^D \\ \omega_z^D \end{bmatrix} \quad (8)$$

for $\vec{r}^{P/D_{cm}} = [a \ a \ -a]^T$.

This equation uses Steiner's theorem to express the inertia tensor of the hopper about the spike tip (P), which produces non-negligible products of inertia. Here, the principal moments of inertia are assumed to be equal due to the symmetric geometry of the hopper. Refer to Figure 11 for definitions of all symbols.

When pivoting about point P , the hopper's velocity (${}^N\vec{V}^{D_{cm}}$) is constrained such that D_{cm} moves along the arc defined by $\vec{r}^{D_{cm}/P}$; that is, ${}^N\vec{V}^{D_{cm}} = {}^N\vec{\omega}^D \times \vec{r}^{D_{cm}/P}$. Thus, given a desired heading angle (ϕ), the resulting hop angle (θ_h) is geometrically constrained to be

$$\theta_h = \tan^{-1}(\cos \phi + \sin \phi). \quad (9)$$

For $\phi = 0^\circ$ (akin to a single-axis hop), this predicts a hop angle of $\theta_h = 45^\circ$, which indeed agrees with Eq. (4) for a cubic geometry. For hopping directly over a single spike (i.e., $\phi = 45^\circ$), Eq. (9) is maximized, predicting a hop angle of $\theta_h = 54.7^\circ$. Although not as efficient for long distance hops, a steeper hop angle may be useful for ascending steep terrain or escaping from pits. Incidentally, the force vector imparted by the ground is exactly parallel to ${}^N\vec{V}^{D_{cm}}$, so hop angles ranging between 45° and 55° generate enough normal force to validate the no-slip assumption, as discussed for the 2D model in Section 2.1.1. Now,

given a desired heading angle (ϕ), a prescribed hop angle (θ_h), and a desired hop speed (v_h), the velocity vector is given by,

$${}^N\vec{V}^{D_{cm}} = \frac{v_h}{\sqrt{2(1 + \sin \phi \cos \phi)}} \begin{bmatrix} \cos \phi \\ \sin \phi \\ \sin \phi + \cos \phi \end{bmatrix}. \quad (10)$$

So far, (8) relates ${}^N\vec{\omega}^D$ to the control inputs (ω_{Fk}), and (10) relates ${}^N\vec{V}^{D_{cm}}$ to the desired trajectory, governed by ϕ and v_h . However, to formulate a control law relating flywheel inputs to a desired trajectory, ${}^N\vec{\omega}^D$ must be calculated as a function of ${}^N\vec{V}^{D_{cm}}$, which is an ill-posed problem, in general. In other words, ${}^N\vec{V}^{D_{cm}}$ does not uniquely define ${}^N\vec{\omega}^D$. Intuitively, it makes sense that the angular velocity component about $\vec{r}^{D_{cm}/P}$ does not influence ${}^N\vec{V}^{D_{cm}}$. There are many ways to constrain ${}^N\vec{\omega}^D$ such that it is uniquely defined by ${}^N\vec{V}^{D_{cm}}$. For the case of “pure” hopping (refer to Table 1), the angular velocity should be parallel to the surface³, which forces $\omega_z^D = 0$. Thus, for $\vec{r}^{P/D_{cm}} = [a \ a \ -a]^T$ and $\omega_z^D = 0$, it follows that

$${}^N\vec{\omega}^D = \frac{v_h}{a\sqrt{2(1 + \sin \phi \cos \phi)}} \begin{bmatrix} -\sin \phi \\ \cos \phi \\ 0 \end{bmatrix}. \quad (11)$$

Finally, Eqs. (8) and (11) can be combined to calculate the control inputs, ω_{Fk} . Interestingly, even though $\omega_z^D = 0$, $\omega_{Fz} \neq 0$ in general, which can be attributed to the effects of products of inertia in $\vec{I}^{D/P}$. Note that Eqs. (8) – (11) are derived for a particular spike contact point, P , although the same arguments can be used to derive similar expressions for any other spike. Furthermore, by constraining the hopper’s rotation to be parallel with the surface, the heading angle is bounded to $0 \leq \phi \leq \pi/2$ for pivoting about P . Equation (11) can be modified to perform twisting hops, which expand this range to as much as $-\pi/4 \leq \phi \leq 3\pi/4$. However, this also produces shallower hop angles, which may violate the no-slip contact assumption for low friction surfaces.

Travel Speed

The analysis of Section 2.1.1 can be extended to answer the question, “How fast can the hopper travel?” More specifically, the speed we care about is the *gross* speed at which the hopper can traverse a relatively flat and unobstructed terrain. The duration of a single hopping maneuver can be thought of as the sum of the time to spin up the flywheels (T_{spin}), and the time of flight (T_{flight}), where

$$T_{\text{spin}} = K_S \left(\frac{\sqrt{2}\omega_f I_f}{m_p g l} \right), \quad T_{\text{flight}} = K_B \left(\frac{\sqrt{2}\eta l \omega_f}{g} \right), \quad d_{\text{hop}} = K_D \left(\frac{(\eta l \omega_f)^2}{g} \right). \quad (12)$$

These equations result directly from (1) and (4), and assuming $\theta_{\text{hop}} = \alpha = 45^\circ$. Here, K_S represents a safety factor to prevent counter rotation during flywheel spin-up, K_B accounts for

³Other types of constraints can be conceived for twisting hops, but care must be taken to ensure that they do not force neighboring spikes into the surface, which would violate the single-point pivoting assumption.

residual bouncing as a proportional correction factor on the flight time of the first parabola, and K_D is also a proportional correction factor on hop distance to account for bouncing as well as for deviations in heading. Based on observations from simulations, conservative estimates are $K_B = 2$, and $K_S, K_D = 1.2$. Combining (12) and (6) yields the average expected speed:

$$\bar{V} = \frac{d_{\text{hop}}}{T_{\text{flight}} + T_{\text{spin}}} = \frac{\sqrt{2d_h g}}{2} \left(\frac{K_D \eta m_p l^2}{K_B \eta m_p l^2 + K_S I_f} \right) \approx \frac{\sqrt{2d_h g}}{2} \left(\frac{K_D}{K_B + K_S} \right). \quad (13)$$

The above approximation assumes $I_p + m_p l^2 \approx m_p l^2$, which is reasonable for our prototype ($m_p l^2 = 0.065$ and $I_p = 0.013$). Interestingly, \bar{V} depends on the square root of hop distance and gravity, indicating that farther hops result in faster net motion, and motion on bodies with weaker gravity is slower. On Phobos ($g = 0.0058 \text{ m/s}^2$), with the parameters of our current prototype, the parameters K_S , K_B , and K_D defined above, and for an average 10 m hop, we can expect a net speed of about 7 cm/s. However, for longer excursions, hops of 100m are reasonable (i.e., $\omega_f = 6000 \text{ rpm}$), and could increase net speed to over 20 cm/s.

2.1.2 Tumbling Dynamics and Control

Tumbling refers to a rotational maneuver, whereby the hopper pivots about one or two spikes and lands on an adjacent face without losing surface contact. Although tumbles about a single spike are possible, the dynamics of the hopper about a single point of contact have 3 DoF and are relatively uncontrollable compared to pivoting over a pair of spikes. Therefore, the analysis here will only consider single-axis tumbles, again using the 2D model presented in Figure 9. In this context, the goal of a tumbling maneuver is to pivot to the right and land on the next consecutive spike such that its orientation is incremented by -2α .

To characterize actuation bounds for tumbling, the actuation is regarded as an instantaneous transfer of momentum, similar to the hopping analysis in Section 2.1.1. Accordingly, the initial kinetic energy of the hopper at $t = 0^+$ can be equated to the gravitational potential energy at the tumbling apex ($\theta = 0$). This yields an expression for the *minimum* flywheel velocity required to vault the hopper over its leading spike:

$$\omega_{f, \min} = \sqrt{\frac{2m_p g l (1 - \cos(\alpha + \beta))}{\eta I_f}}. \quad (14)$$

To characterize the *maximum* flywheel velocity for tumbling, consider the hop criterion given by (3) and a zero torque input for $t \geq 0^+$. It follows that $\theta(t)$ and $|\dot{\theta}(t)|$ both decrease with time. Thus, if surface contact is lost, it will occur just after momentum transfer when $\theta(0^+) = \alpha + \beta$, and $|\dot{\theta}(0^+)| = \eta \omega_f$. This yields the maximum flywheel velocity to perform a tumble *without* hopping:

$$\omega_{f, \max} = \sqrt{\frac{g \cos \beta}{\eta^2 l \cos \alpha}}. \quad (15)$$

Interestingly, there exists an inclination angle, β_{\max} , for which $\omega_{f, \min} = \omega_{f, \max}$ and tumbling is impossible. For a square geometry ($\alpha = 45^\circ$), $\beta_{\max} \approx 30^\circ$. Also, as expected, $\omega_{f, \min} = 0$

when $\beta = -\alpha$, which corresponds to the declination angle at which the hopper freely tumbles without actuation. Practically, a robust control input for tumbling is simply the average of these upper and lower bounds.

2.1.3 Twisting Dynamics and Control

In addition to hopping and tumbling for mobility, another useful maneuver for changing orientation on the surface is a turn-in-place, or “twist” maneuver. Instead of using the flywheels to spin about a horizontal axis (i.e., for hopping and tumbling), the idea here is to spin about an axis *normal* to the surface such that the hopper rotates a desired angle, ψ_d (see Figure 12). The ability to control orientation could be extremely useful for pointing cameras and science instruments, or even to simply get in a more favorable configuration for a subsequent hop.




Figure 12: A twisting maneuver: the hopper turns in place to change orientation.

In the nominal resting configuration, the hopper is oriented such that three or four of its spikes are contacting the surface with one flywheel axis dominantly pointing “up.” Thus, a twisting maneuver can be executed with a single flywheel. Although the motors should be capable of applying the small torque required to initiate a twist, a more controllable and robust approach is to spin the flywheel up slowly (such that the hopper does *not* rotate) and hit the brakes—similar to the hopping control strategy discussed in Section 2.1.1. Assuming the hopper is resting on a flat, inclined surface (angle β) with a Coulomb friction contact model (i.e., $F_\mu = \mu_k F_N$) and a spike-to-spike base length of s , the resisting friction torque is given by $\tau_\mu = -K_s s \mu_k m g \cos \beta$, where the “mean contact radius”, $K_s = \sqrt{2}/2$ for a cube. This frictional torque corresponds to the maximum rate at which the flywheel can be accelerated prior to braking, and the rate at which the hopper decelerates after braking. Conservation of angular momentum can again be invoked to derive an expression for the flywheel speed required to rotate an angle ψ_d :

$$\omega_f(\mu_k, \psi_d) = \frac{I_p}{I_f} \sqrt{\frac{-2\tau_\mu \psi_d}{I_p}} = \frac{1}{I_f} \sqrt{(2I_p K_s s m g \cos \beta)(\mu_k \psi_d)}. \quad (16)$$

This equation for the control input is written explicitly as a function of μ_k to emphasize that it is strongly dependent on the surface friction—a parameter that is, in general, unknown a priori. However, an estimate for the friction parameter ($\hat{\mu}_k$) can be refined by observing the response to a known input. In this way, a *sequence* of twist maneuvers can be executed until

the desired precision in ψ_d is achieved, even in the presence of uncertainty and variability in the environment. Furthermore, the assumption of Coulomb friction on a smooth surface is not true in general, and there could even be cases in which the local terrain is so uneven that twisting causes the hopper to lose ground contact altogether. Equation (16) can be modified to account for other proposed contact models (e.g., adding a viscous damping term), but Coulomb friction is a simple *single-parameter* model that can capture a wide range of surface conditions.

2.2 Mobility System Design

First and foremost, the flywheels are one of the most important components of this rover concept. The shape, sizing, and placement of the flywheels represents an important trade space to consider, which is influenced by the target body of interest, the scale of mobility desired, and mass/space allocation requirements for other subsystems. For example, higher surface gravity requires more momentum to be generated for a given hop distance, as indicated by Eq. (6). One of the most important metrics to consider is mobility *efficiency* (see Eq. (5)), which is highly dependent on the flywheel inertia. More precisely, the efficiency favors a high flywheel inertia *and* low total inertia—an inherent tradeoff because the mass and inertia of the flywheel also contribute to the total inertia. We have extensively studied the flywheel shape-optimization problem in the context of efficient mobility with mass and size constraints; the key finding is that efficiency favors dense, thin-rimmed, large diameter flywheels located as far from the CM as possible. Accordingly, the flywheels on our current prototype (which was designed to optimize efficiency for mobility in 1 g) are 5 in. diameter, steel rimmed disks with an inertia of 9500 g cm^2 and mass of 300 g, or about 1/8th the total mass (see Figure 13, left).




Figure 13: Left: CAD model of prototype, annotating key components. Right: Schematic of *band brake* mechanism.

Equally important to the flywheels are the actuators used to manipulate them—namely, motors and brakes. Because the motors are not directly responsible for mobility under the control regime proposed in Section 2.1, they can be arbitrarily small, solely constrained by the time required to build momentum and possibly by the torque requirements for in-flight attitude control. A direct-drive is the simplest approach, but geared or belt-driven

couplings can also be implemented for more flexibility over motor placement. The prototype in Figure 13 has three 50 W brushless motors with hall sensor feedback (Hobby People 2208/15), and electronic speed controllers (Emax BLHeli-12A).

The braking system requires special attention, as it is directly responsible for executing the motion primitives discussed in Section 2.1. We have experimented with various types of braking mechanisms including friction disks and impulsive braking via impact hammers, but by far the most effective architecture utilizes a *band brake*, as illustrated in Figure 13, right. Consisting of a Kevlar-reinforced neoprene belt wrapped around the circumference of the flywheel, the band brake applies a friction-generated torque directly on the flywheel by pulling the belt taut. In the retracted state, the belt is robustly held against an outer, concentric retaining ring by small springs, which maintains a 1 – 2 mm radial clearance for the flywheel to freely rotate. The applied torque is governed by the well-known capstan equation, $T_1 = T_2 e^{\mu\alpha}$ (for $\omega_f > 0$), whereby the exponential gain ($e^{\mu\alpha}$) allows very high torques ($\tau = (T_1 - T_2)/R$) to be achieved for relatively low force inputs. Thus, this presents two possible forms of actuation: (1) *high* torque by pulling on P_2 and fixing P_1 to the frame or (2) *low* torque by pulling on P_1 and fixing P_2 to the frame (the reverse is true for CCW rotation). For our prototype, with a micro servo motor (DSM44 by Power HD) generating tension up to 15 N, $\alpha = 3\pi/2$, and $\mu \approx 1$, option (1) can generate torques up to 100 Nm and option (2), up to 1 Nm. Thus, this architecture has the capability of achieving a wide variety of hopping maneuvers (refer to Figure 10). Furthermore, it has proven to be robust, showing no appreciable signs of wear on the belt or structural degradation after hundreds of cycles.

The overall structure and frame consists of a cube with a 15 cm edge constructed out of lightweight laser-cut and 3-D printed parts ($m_p = 2.3$ kg, $I_p = 0.013$ kg m²), and one spike on each corner. Previous prototype iterations have included more spikes, but it has been determined through experimentation and insights from dynamic analysis (see Section 2.1) that a cubic geometry with 8 spikes offers the best balance of protection and mobility performance. Each spike is fitted with a rubber foam tip to mitigate impact stresses and increase surface friction. However, even for large hops over 100 m, the impact speeds are typically low (< 1 m/s) and structurally tolerable even without extensive padding. In other words, due to the typical *slow* dynamics in microgravity, repeated impacts do not generally pose structural concerns. Also, although not considered in the analysis of Section 2.1, a larger spike surface area has better performance in loose granular media.

2.3 Mobility Experiments

Testing the dynamics and control of a rover in microgravity is a classical challenge. Methods for effectively emulating microgravity conditions can be broadly characterized into three categories: (1) designing an effective 1g analogue rover (e.g., how “scarecrow” is a slimmed down version of the Curiosity rover), (2) gravity offloading systems that attempt to simulate reduced gravity by applying an upward force on the system (e.g. buoyancy tanks, air-bearing tables, etc.), (3) and free fall chambers such as drop towers and parabolic flights. Although a mobility prototype was developed that is capable of performing small hops in 1g, we relied on two test beds for observing the “true dynamics” in microgravity: (1) a custom offloading

test bed that allows for full 6 degree-of-freedom (DoF) motion of the rover (Section 2.3.1), and (2) a parabolic flight campaign (Section 2.3.2).

2.3.1 Experiments in Custom 6 DoF Microgravity Test Bed

To the best of the authors’ knowledge, no preexisting test beds are capable of meeting the stringent requirements for emulating the hopper’s dynamics in microgravity: (1) effective gravity levels on the order of 10^{-3} g’s, (2) allowing for full 6 DoF motion with (3) minimal exogenous dynamic interference, (4) over an extended period of time (say, more than 20 seconds), and (5) within an extended workspace (say, more than 1 m^2). ARGOS, a gravity offload system developed at NASA’s Johnson Space Center, may come the closest [20]. Used primarily for human testing in zero-g environments, ARGOS consists of an actively-controlled overhead 3-axis gantry crane that tracks the motion of the suspended subject, enabling the “free-floating” behaviors observed in space. A similar gantry system, aimed at reproducing zero-gravity conditions, was developed at NASA’s Ames Research Center for testing “Personal Satellite Assistants (PSAs)” [21]. At Stanford, we have extended this idea to create a novel 6 DoF test bed for operating hoppers in *microgravity* conditions (see Figure 14).




Figure 14: **Left:** 6 DoF microgravity test bed CAD rendering. The powered gantry tracks the translational motion of the hopper in x , y , and z within a volume of $3 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$ respectively, while allowing for free fall in z at sub-milli-g levels. The gimbal frame allows the hopper to rotate in all three axes. **Right:** Image of the test bed.

Similar to ARGOS, our test bed is built on a powered gantry crane that permits the tracking of translational motion. The 3-axis rotational motion of the hoppers is achieved by mounting it within a lightweight rigid gimbal frame (see Figure 14). The gimbal-mounted hopper is suspended by a cable from an overhead attachment point on the gantry crane so that it can swing freely. By accurately measuring the relative deflection of the pendulum, the x and y axes are actuated using feedback control techniques to keep the pendulum in a vertical state. In this manner, external lateral forces that act on the hopper cause the whole system to accelerate as Newton’s second law predicts. To enable microgravity free fall, the control of the vertical axis utilizes force feedback with a series-elastic coupling to apply a

precise offloading force on the hopper, even in the presence of external forces. More details on the dynamics and control of the test bed can be found in [5].

To further characterize the dynamics and *controllability* of the hopper and to validate the models presented in Section 2.1, various experiments were conducted on the microgravity test bed. Specifically, trajectory data was collected while performing (1) hops about a single axis, (2) oblique hops over a single spike, and (3) twisting maneuvers. The position feedback from the gantry was used in conjunction with the force and displacement signals to determine the translational trajectory of the hopper. At this time, no measurements of orientation are taken, which is sufficient for comparison with the translational trajectory predictions from Section 2.1. However, encoders on the gimbal axes or visual tracking techniques could be implemented in future experiments to achieve accurate pose estimates. Here, we also restrict our study to maneuvers on a flat and level surface, leaving mobility on rough and uneven surfaces for future work. The surface material was a rigid tile, and the friction between it and the rubberized spike tips was high enough to prevent slip during hopping.

Single-Axis Hopping

As a first set of experiments, simple hopping and tumbling maneuvers were executed about a single axis as a basis for comparison with 2D analytical model from Section 2.1.1. In each experiment, the control approach discussed in Section 2.1.1 was executed, whereby the flywheel is slowly accelerated until a target angular velocity is reached, at which point the band brake is applied and the hopping sequence ensues unactuated. For a desired hop distance of 1 m in an emulated gravity level of 0.001 g's, the target flywheel velocity was calculated using (6) to be 550 rpm. A total of 30 single-axis hops were executed, 10 of which are plotted in Figure 15.

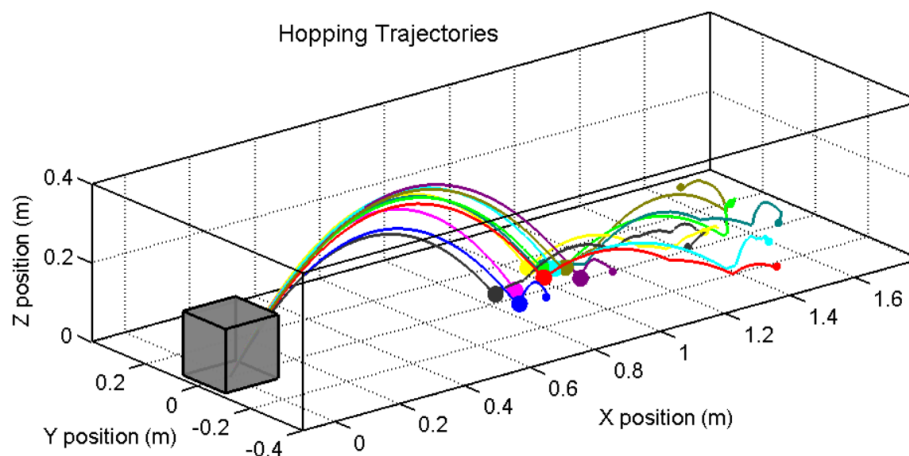


Figure 15: Trajectories of the hopper within the microgravity test bed. The gravity level of these experiments was set to 0.001 g's, and the flywheel was commanded to 550 rpm. Position data for each experiment was shifted to start at the origin. A z position of zero corresponds to a flat stance where four spikes are in contact with the ground. Thus, bounces above zero indicate collision at a tilted orientation. Initial bounces are highlighted with a large point, and the final resting location, with a smaller point.

An interesting observation from Figure 15 is the variability in bouncing. Although the “landing envelope” for the first bounce is relatively small, subsequent trajectories are highly

sensitive to the spin and orientation on impact, which results in a much wider “resting envelope”. However, the hopper’s rotational dynamics are corrupted by the gimbal in flight, so these bouncing trajectories do not exactly capture the “true” expected distribution. On the other hand, hopping angle and distance measurements exhibit a more consistent trend and are in close agreement with the predictions of (7) and (4). The mean hop angle for the 30 experiments was 50° with a standard deviation of 3° . This is marginally higher than the 45° prediction, likely due to the elastic rebound of the spike tip. Indeed, numerical simulations with a spring/damper contact model also exhibit higher hop angles. Remarkably, without any adjustments in the model parameters, the mean hop distance of 0.94 m was also very close to the target of 1 m, with a standard deviation of 0.07 m. Here, the variability likely comes from imperfect control of flywheel’s speed and variability in the hop angle. Finally, as expected for single axis hops, only minor deviations in heading were observed, with a mean of just 1.5° and std. dev. of 2° . See <http://web.stanford.edu/~pavone/movies/hop.mov> for sample videos of hopping experiments.

Oblique Hopping

A similar set of experiments was performed to demonstrate directional control by hopping over a single spike, as discussed in Section 2.1.1. In this case, the flywheel speeds were calculated based on Eq. (8) and augmented to account for the gimbal’s effective mass and inertia. Twenty hops were executed over a range of desired heading angles ($0 \leq \phi_d \leq 45^\circ$) and at a nominal speed of 5 cm/s (which roughly corresponds to a 1 m hop). The observed heading (ϕ) and hop (θ_h) angles are compared with predictions based on the analytical model from Section 2.1.1 and with simulated trajectories based on a spring/damper contact model (see Figure 16). The “predicted range” is derived from numerical simulations, and it highlights the trajectory variability that can be induced by changes in the surface contact model, ranging from a highly damped to purely elastic. Numerical results from a “best-fit” surface contact model are also indicated by the green dashed line.

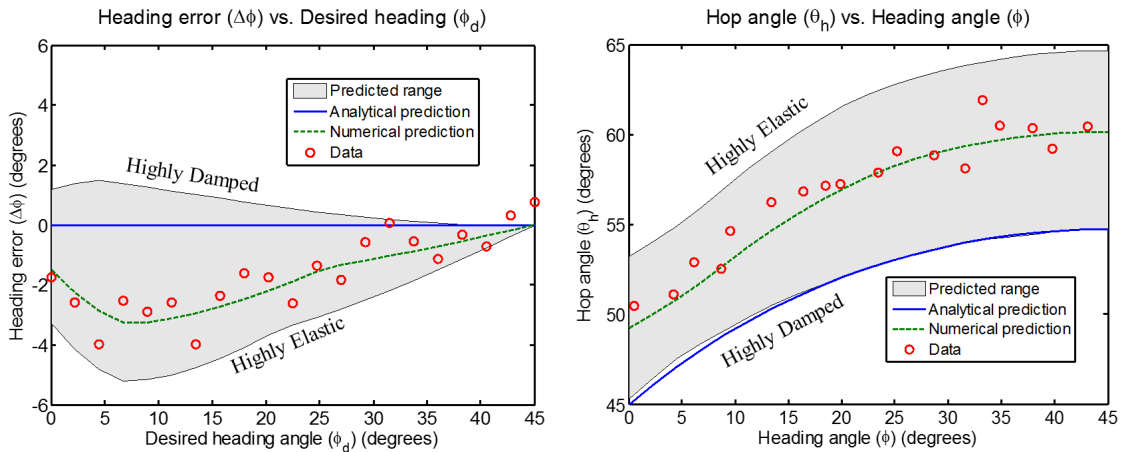


Figure 16: Oblique hops were executed over a range of desired heading angles and compared with predictions based on analytical and numerical models. **Left:** The error in heading angle ($\Delta\phi$) as a function of intended heading angle (ϕ_d). **Right:** Hop angle (θ_h) as a function of intended heading angle.

Figure 16 left illustrates that, overall, the directional control law is highly effective at achieving the desired heading angle, with a mean deviation of only 2° . Furthermore, the data and simulations both indicate a bias towards *undershooting* the intended heading, which can be corrected for with a simple shift in the control law. Nonetheless, even heading deviations up to 10° (say) still constitute a stark improvement in directional control over single-axis hops.

The hop angle data shown in Figure 16 right follows the trend predicted by Eq. (9) and also reiterates the observation from the single-axis hop experiments: surface elasticity increases the hop angle. Indeed, a sweep of numerical simulations shows that θ_h can vary by up to 10° depending on the contact elasticity, whereby the lower bound corresponds to a highly damped surface and strongly agrees with the analytical model. Although the surface properties of small bodies may be highly variable, for a rigid surface, the contact elasticity is predominantly governed by the compliance of the spike tip itself, which may be roughly characterized a priori and incorporated into the predictive model.

Twisting Experiments

Unlike hopping and tumbling, a twist is the only motion primitive that is not critically enabled by microgravity, as the CM (nominally) remains fixed. Thus, Eq. (16) should work equally well to predict the response in a 1g environment and the microgravity test bed is not needed. So the hopper was simply placed on a flat and level tile surface, and the friction coefficient was first estimated from a number of sliding force measurements to be $\mu_k = 1.1$. Equation (16) was then used to calculate the required flywheel speed to achieve a range of target twist angles ($0 < \psi_d < 180$) based on the empirically determined μ_k and measured hopper parameters. The results in Figure 17 show a very strong agreement with the predictive model, whereby the slight shift in the data could be attributed to any combination of imperfect estimates in the model parameters.

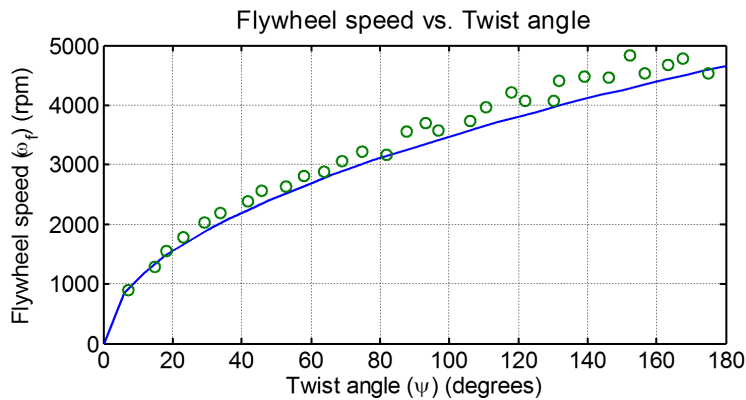


Figure 17: A series of 30 twist maneuvers were executed over target range of $0 < \psi_d < 180$.

2.3.2 Parabolic Flight Experiments

While the microgravity test bed discussed in Section 2.3.1 allows for small scale maneuvers in *simulated* microgravity, it has a few key limitations. In particular, (1) the added mass and inertia of the gimbal slightly corrupts the hopping dynamics, (2) it cannot track more aggressive maneuvers, and (3) it is fundamentally incapable of offloading the regolith with which the rover interacts. These key limitations motivated a parabolic flight campaign in which we were less constrained in terms of the types of mobility we could test.

However, the variable *quality* of microgravity provided by the aircraft posed a significant challenge for testing a rover that requires strictly positive gravity. A great deal of effort was invested in designing a test system that could accommodate such variability by appropriately timing maneuvers during portions of the parabola that were deemed acceptable by an onboard accelerometer.

Our experimental setup is as follows (see Figure 18): the hopper prototype sits on the test surface and is restrained by a retractable arm that applies a gentle downwards force. An accelerometer, rigidly mounted to the floor of the aircraft, measures the transient accelerations and is used to automatically retract the arm and initiate each experiment when the resulting gravity conditions are deemed acceptable. An array of small cameras fixed inside the payload container track the hopper’s motion (position and attitude) with millimeter precision and high frame rates (240 Hz) via body-mounted fiducial markers. Some cameras were also focused on the surface to observe contact interactions.

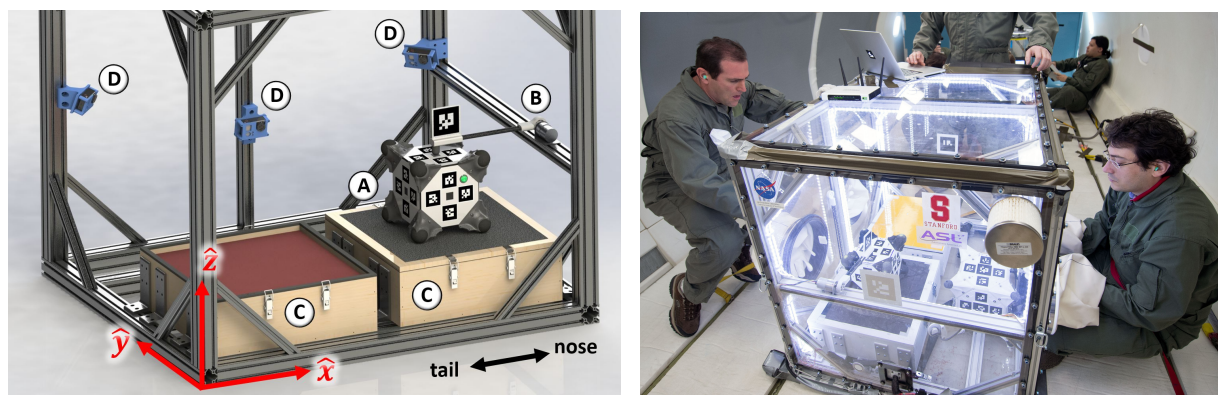


Figure 18: Experimental setup. **Left:** The hopper (A) is held in place on the test surface (C) by an actuated arm (B). An array of five cameras (D) capture its motion as it hops within the container. **Right:** Photo of our experiments on NASA’s C9 aircraft.

To emulate a range of surface properties that rovers may encounter on small bodies, the experimental payload container in Figure 18 was fitted with two boxes (labeled “C”) that provided a total of four different test surfaces: a low-friction *Kapton tape* covering the (rigid) lid of one box, a high-friction *grip tape* on the other lid, a cohesive *comet regolith simulant*, and a low-cohesion *garnet sand* (see Figure 19). The two rigid surfaces, “A” and “B,” aim to mimic rocky or icy surfaces with varying degrees of “traction.” The simulant “C,” is a crushable material consisting of an aerated cement that mimics cohesive (yet friable) regolith. Its cohesive properties allowed it to be exposed during negative g’s, unlike the garnet sand, “D.” More details on the test setup and procedure can be found in [6].



Figure 19: Test surfaces: (A) low-friction and (B) high-friction rigid surfaces, (C) crushable comet regolith simulant, and (D) granular, low-cohesion garnet sand.

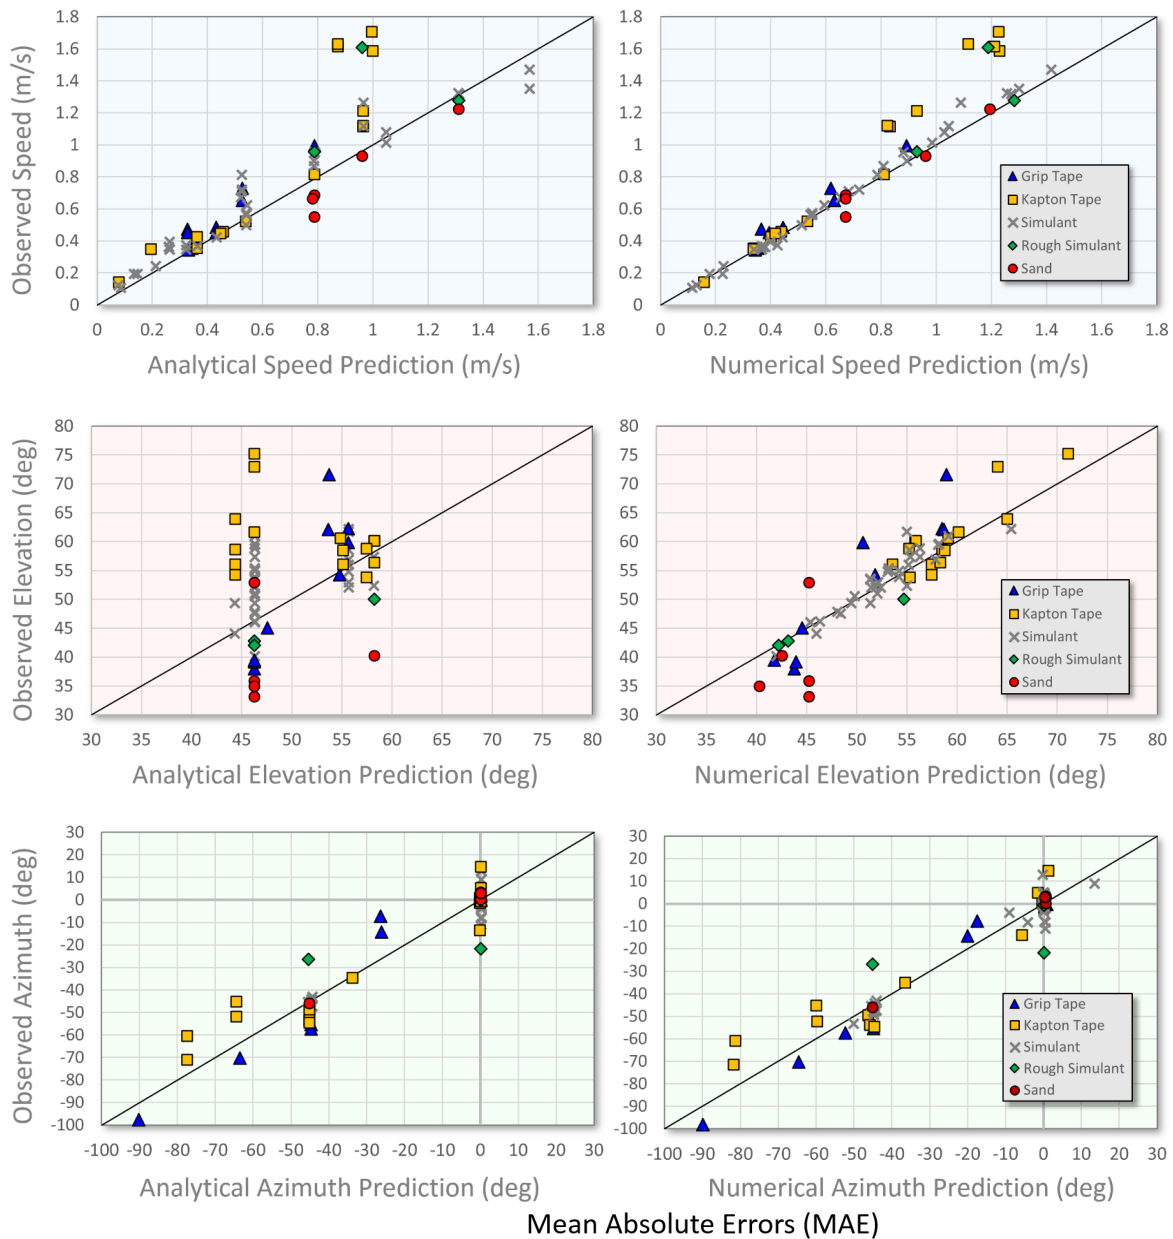
This experimental technique was used to evaluate the controllability of two hopper prototypes performing maneuvers on various surfaces. Over the course of four flights, 74 of 190 parabolas resulted in successfully triggered mobility experiments. Of those, 64 were performed on three different surfaces in zero-g parabolas (0 ± 0.02 g), while 10 were performed on the garnet sand in positively-biased parabolas (0.03 ± 0.02 g). The remaining parabolas experienced unfavorable gravity conditions or technical difficulties (timing error of the arm release and hop trigger, operator error, software bugs, and wireless interference from aircraft communication). Most parabolas were utilized for hopping experiments, but a few were also used to test more precise maneuvers such as tumbling and twisting. A video compilation of several maneuvers can be found at <http://web.stanford.edu/~pavone/iser16>.

Hopping experiments

Since the dynamics of a hopping rover in ballistic flight are deterministic (for an airless body with known spin and gravity model), we can characterize the resulting trajectory with three parameters describing its initial launch velocity vector: speed (v_h), elevation angle (θ_h), and azimuth angle (ϕ_h). These parameters are extracted from the visual tracking data by fitting a parabola to the time-series position measurements of the hopper’s mass center for the first 20 cm of its trajectory after takeoff. The observed hop vectors can then be compared with predictions obtained by inputting the observed flywheel speeds into our analytical models from Section 2.1 as well as predictions derived from forward simulation with a numerical contact model.

The results in Figure 20 show predicted values on the horizontal axis and measured trajectory data on the vertical axis. Each data point represents a trajectory resulting from a particular set of flywheel speeds. Overall, there was strong agreement between the experimental and model-generated data with mean absolute errors of about 10% for speed, and 5° for elevation and azimuth angles. It is important to note that hopper experienced slight drift before actuation on many of these maneuvers, such that its initial state was not exactly grounded and stationary, as assumed by the analytical model. Therefore, it is not surprising that the numerical model, *which is simulated from the actual measured initial states* and accounts for the varying surface properties, exhibits stronger agreement with the data than the less-informed analytical model.

Examining systematic bias in the data can help to identify unmodeled effects and make improvements. For example, the clustering of analytical elevation predictions at 45° reflects



		Speed		Elevation		Azimuth		
#		Analytical	Numerical	Analytical	Numerical	Analytical	Numerical	
Grip Tape	▲	9	17.7 %	9.6 %	6.2°	4.8°	7.8°	5.0°
Kapton Tape	■	15	24.3 %	16.5 %	9.8°	2.3°	7.2°	6.6°
Simulant	×	33	16.7 %	5.1 %	5.1°	1.5°	3.7°	3.3°
Rough Simulant	◆	3	22.2 %	11.2 %	4.6°	1.6°	13.2°	12.2°
Sand	●	5	17.1 %	5.7 %	9.7°	6.9°	1.9°	1.8°
Total		65	18.4 %	7.8 %	6.6°	2.4°	5.3°	4.7°

Figure 20: Experimental trajectory data (speed, elevation, and azimuth) for 65 hopping maneuvers on 5 surfaces (see Figure 19) compared with predictions based on an analytical model (left plots) and a numerical model (right plots). “Rough simulant” corresponds to a few experiments in which the comet regolith simulant was highly fractured and uneven. Predictions and observations that are in agreement lie along the black lines with slope 1. The table of mean absolute errors summarizes these results.

the no-slip and instantaneous momentum transfer assumptions for single-axis hops, which are not realizable for lower friction surfaces and brakes with limited torque. If, however, information about the surface friction is known *a priori*, the control law can be adjusted to reflect the higher expected elevation ($\theta_h \approx \cot^{-1} \mu$). Also, contact interactions with loose granular regolith, which is essentially “fluidized” by microgravity, does not adhere well to either the pin-jointed spike contact assumption or the numerical contact model (such as over-estimated hop elevation on sand in Figure 20); it will be the subject of future work. Finally, it is suspected that the high-speed outliers can be attributed to a temporary hardware issue with one of the prototype’s braking mechanisms.

Tumbling Experiments

For a tumbling maneuver, an upper and lower bound on the control input are given by Eqs. (14) and (15), which correspond to the speed at which the hopper would rotate too fast and lose surface contact and the speed at which it would just barely tip over, respectively. These bounds are functions of the hopper’s inertial and geometric properties ($m_p, I_f, \eta, \alpha, l$), its initial pose (β), and gravity (g). However, due to the need to maintain continuous ground contact over a longer time period, tumbling maneuvers could not exploit brief gravity transients and were therefore restricted to positively-biased parabolas. Table 2 summarizes data for the two tumbles performed.

Trial	Surface	Inclination*	ω_{\min} (rpm)	ω_{\max} (rpm)	ω (rpm)	Success?
1	sand	-10.7°	1451	2937	1968	Yes
2	sand	-22.7°	255	837	274	Yes

Table 2: Data from two successful tumbling experiments on sand at about 0.035 g’s. Note that the measured flywheel speed (ω) is indeed between the predicted minimum and maximum bounds (see Eqs. (14), (15)). *Negative inclination indicates a “downhill” tumble.

While the data is sparse, a few insightful observations were made. For one, on loose granular media, the leading spikes tend to sink into the surface, which shifts the pivoting axis inward and effectively shortens the modeled spike length (l). Also, faster tumbles have a higher chance of producing undesirable rebounds upon impact. However, both of these incidental effects can be mitigated by operating in the lower speed range (e.g. 10% higher than ω_{\min}).

2.4 Summary

The internally actuated mobility system of spacecraft/rover hybrids has been designed, analyzed, and tested in a relevant environment. Starting from a “first-principles” modeling approach, we proposed simplified dynamics models for studying various motion primitives, including hopping, tumbling, and twisting. These simplified models facilitated an analytical characterization of control strategies for targeted mobility. These models also allowed us to extract insights that inform the *design* of the mobility system, for example, that a cubic chasis generally yields 45° hops or that the hop efficiency is closely related to the relative inertias

of the flywheels and chassis. From these insights, several prototypes have been developed, each one improving upon deficiencies of the previous. Finally, we spent a considerable effort to construct experimental environments for observing the dynamics that would be expected in microgravity; namely, a custom microgravity offloading test bed, and a parabolic flight campaign. Experiments generally showed a strong agreement between observed trajectories and those predicted by our models.

There are several aspects of dynamics and control that would benefit from further study. For design, a more thorough trade study of the rover geometry is required. In particular, we only considered symmetric configurations, whereby the rover can maneuver from any orientation. However, it may be more advantageous to only have one larger flywheel for energetic hopping and smaller auxiliary flywheels for “self-righting” or reorientation prior to hopping. Second, more detailed models (supported by tests on terrain simulants) are required to understand the complex interactions of hopping on loose granular regolith.

3 Spacecraft/Rover Hybrid: Autonomy Subsystem

In addition to controllability, hopping rovers operating on distant bodies require a high degree of *autonomy*, as communication suffers from long light-speed delays and relies on a mothership relay, which may be infrequent. On Mars, wheeled rovers are equipped with visual perception, terrain classification, and path planning algorithms for autonomous mobility [22]. However, in contrast to wheeled rovers, which operate through continuous interaction with the environment, hoppers can only apply forces from rest on the surface and have no control of their trajectory mid-flight. Thus, autonomy for hopping systems requires a more discrete and sequential structure, which can be decomposed into four phases: (1) localization, (2) trajectory planning, (3) hop execution, and (4) ballistic dynamics (see Figure 21).




Figure 21: The high-level autonomy architecture for hopping rovers consists of four phases: (1) localization, (2) trajectory planning, (3) hop execution, and (4) uncontrolled, ballistic dynamics.

Localization on the surface of small bodies is an open area of research, which has been addressed from two perspectives: (1) assuming an “eye-in-the-sky” mothership [23], and (2) mothership-independent, on-board perception, primarily through visual odometry [24]. In Section 3.2, we discuss our vision-based approach to localization for hopping rovers.

Hop execution addresses the problem of pushing off from the surface to achieve some desired velocity vector and, as discussed in Section 2, is highly dependent on the rover architecture (e.g. actuation mechanisms, surface interaction, etc.). Generally, for a given hopper, the target velocity can only be executed approximately—due to environmental uncertainty and control errors—and must obey certain constraints, such as speed and direction limitations.

Dynamics about small bodies has a rich body of literature for orbiting spacecraft [25], and to a lesser extent for surface interaction [26, 27]. The key challenges lie in accurate modeling of the gravity field about irregularly shaped bodies and the physical interactions with the surface. However, even the highest fidelity gravity and contact models rely on information that may not be available a priori via remote observations, such as internal mass distribution and surface structure.

Motion planning for hoppers seeks to answer the question, “*What is the next best hop to perform, given a set of mission objectives, an estimate of the rover’s location, an understanding of its capabilities, and a model of the world?*” In other words, planning represents the rover’s decision-making module and is tightly intertwined with all other elements in the mobility pipeline. It is important to note that although we assume a specific configuration for the mobility system of the hybrid in Sect. 2, the motion planning problem largely deals with the ballistic dynamics of the hybrid *in motion* (i.e. after hopping), and thus, is a more

general problem that applies to any hopping system. Planning is addressed in Section 3.1.

3.1 Motion Planning

The semi-autonomous operation of hopping rovers requires a decision making module that can determine a safe and effective hopping trajectory. Not only does this require an understanding of how a hopper should move given a perfect model, but importantly, it also requires an understanding of the uncertainty associated with potential actions. However, previous studies of planning for hopping rovers have assumed highly simplified and deterministic models of the dynamics, localization, and the hopper itself. Bellerose, et. al derive analytical control laws for a spherical hopper on a smooth, spherical asteroid with a Coulomb friction contact model, and exact localization [28]. This work is extended to the case of smooth, *ellipsoidal* asteroids in [29], which also derives approximate speed constraints to prevent escape trajectories. However, these results are founded on oversimplifications of the body shape, gravity, and contact models, as evidenced by remote observations of highly uneven and rocky surfaces (see Figure 22), and the chaotic bouncing pattern of the Philae lander on comet 67P [30].




Figure 22: High-resolution images and shape models of two small bodies. **Left:** Asteroid 25143 Itokawa (535 m, $6 - 9 \mu g$), **Right:** Comet 67P Churyumov-Gerasimenko (4.3 km, $140 - 300 \mu g$).

Instead, we propose a planning architecture that directly accounts for various sources of uncertainty, and thus, produces control policies that are more robust to modeling, control, and localization errors. The key idea is to shift from a “first-principles” approach, to a data-driven approach, whereby high-fidelity dynamics models can be used to simulate *instances* of hopping trajectories from various uncertainty distributions.

First, in Section 3.1.2 we solve the problem of planning a single hop under uncertainty and assuming minimal bouncing. We develop a robust and efficient algorithm for solving the “Lambert’s orbital boundary value problem”—the problem of finding an initial velocity to intercept a target—in highly irregular gravity fields (Section 3.1.2). We then forward propagate control and gravity uncertainty (modeled as Gaussian mixtures) through the dynamics to compute trajectory funnels and landing distributions for a given nominal trajectory (Section 3.1.2). Then, we compute myopic policy gradients based on these landing distributions to derive time/energy-optimal single-hop control policies (Section 3.1.2).

Second, in Section 3.1.3 we extend the planning problem to the case where multiple hops are required and structure the problem as a Markov Decision Process (MDP). We apply a variant of Least Squares Policy Iteration (LSPI) to derive approximately optimal control policies that are safe, efficient, and amenable to real-time implementation on computationally-constrained rover hardware. The performance of these policies is evaluated on a high fidelity

dynamics simulator and compared to that of a greedy heuristic policy. Collectively, the methods presented here constitute the first ever study of uncertainty-aware motion planning for hopping rovers—a crucial component of autonomy for future exploration missions to small solar system bodies.

3.1.1 Preliminaries

In this section, we present the dynamics models used to simulate the trajectories of a hopping rover. First, we detail the forces acting on the rover during ballistic flight, and then, we discuss the model used for the contact dynamics of bouncing on the surface.

Ballistic Dynamics

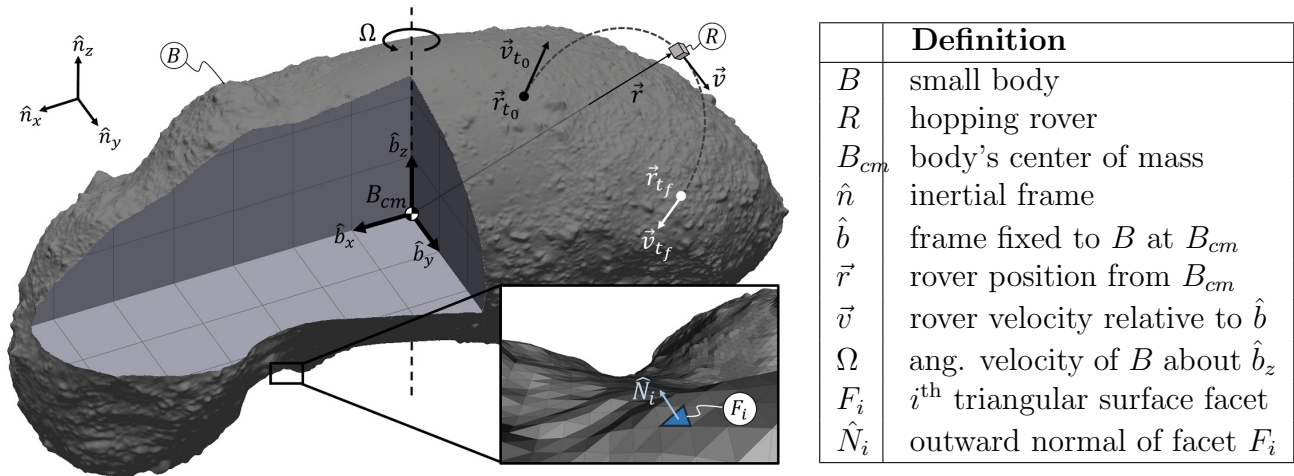


Figure 23: Dynamic model of rover R hopping on body B , which is rotating at $\vec{\omega}_B = \Omega \hat{b}_z$ and is represented by a closed surface mesh consisting of k triangular facets, F_i .

A rover, R , at position \vec{r} with mass m and velocity \vec{v} relative to the asteroid body, B , hops from rest at \vec{r}_{t_0} with velocity \vec{v}_{t_0} and impacts at \vec{r}_{t_f} with velocity \vec{v}_{t_f} . The body is represented as a closed polygonal mesh with k triangular facets, where facet F_i has outward normal \vec{N}_i . The asteroid rotates at a constant angular velocity $\vec{\omega}_B = \Omega \hat{b}_z$.

The external forces (\vec{F}_e) acting on a hopping rover include gravitational forces (of the primary and possibly tertiary bodies), solar radiation pressure (SRP), electrostatic forces, and contact forces. It will also be convenient to represent the rover's dynamics in the rotating body frame, thus introducing effective centrifugal and Coriolis forces. The total effective force is expressed as,

$$\vec{F} = \vec{F}_e - m\vec{\omega}_B \times (\vec{\omega}_B \times \vec{r}) - 2m\vec{\omega}_B \times \vec{v}. \quad (17)$$

In general, \vec{F}_e (expressed in \hat{b}) is a function of not only position and velocity, but also time. In this paper, we focus on the *stationary* case to derive time-invariant control policies, which excludes forces that are periodic when expressed in the rotating body frame such as

SRP and third body perturbations. This assumption is not overly restrictive since SRP is typically three to six orders of magnitude weaker than gravity in close proximity to small bodies, and most small bodies of interest for exploration are either gravitationally isolated (e.g., Itokawa, Bennu, and Psyche) or tidally locked secondaries (e.g., Phobos and Deimos). However, if periodic force models are important, the methods in this paper can be generalized by augmenting the rover’s “state” with a temporal state (e.g., body phase). Nutation of the body’s spin axis can also be accounted for in this way, although most asteroids are believed to be in stable spin about their major axis [31].

Gravity on small bodies may be orders of magnitude weaker than on Earth, but it still represents the dominant force on rovers, so accurate modeling is essential. With only shape information, the most accurate gravity model is a polyhedral model [32], which leverages the divergence theorem to exactly model the gravitational potential (U), attraction ($\vec{g} = \nabla U$), gradient ($\nabla \nabla U$), and Laplacian ($\nabla^2 U$) of a constant density polyhedron as a summation over all facets and edges of the surface mesh. This representation, while highly computational, is especially critical in close proximity to irregular bodies (see, e.g., Figure 22), where conventional models such as harmonic expansions and mascons yield large errors [32]. We alleviate the computational burden of evaluating the surface integral at each time step by precomputing the gravity field at regular grid points within the vicinity of the body *offline*, and then interpolating within this precomputed field online. Validation tests comparing this approximate interpolation to exact evaluation suggest errors on the order of 0.01% for a 5m field discretization on Itokawa using a 5000-facet shape model.

However, most small bodies are likely *not* constant density, and internal density variations are, in general, not known a priori. The polyhedral model may be superimposed with harmonics or mascons if in situ measurements are taken, but a model of *uncertainty* is perhaps more important for robust policy generation. We consider an arbitrary parametric uncertainty model,

$$\vec{g} = \bar{\vec{g}} + \delta\vec{g}, \quad \delta\vec{g} \sim P_\theta, \quad \theta \in \mathbb{R}^k, \quad (18)$$

where $\bar{\vec{g}} \in \mathbb{R}^3$ is the nominal modeled gravity vector, and $\delta\vec{g}$ is a random perturbation from distribution P_θ . As a simple example, Gaussian uncertainty on the total mass of the body could be encoded with a single parameter, $\delta\vec{g} \sim \mathcal{N}(0, \sigma_M^2) \bar{\vec{g}} / |\bar{\vec{g}}|$.

Contact Model

Finally, we require a model for the dynamics of the rover upon impact with the surface. In some cases, it may be sufficient to assume that the rover can achieve a dead-stick landing (e.g., if it is equipped with a damping mechanism), but in general, uncontrolled impact will cause the rover to rebound somewhat randomly. The dynamic response of an impact event is a complex function of the physical properties of both the surface and rover, and the speed and orientation of the rover upon contact. Tardivel and Van wal, et al. developed high-fidelity small body lander simulations to model the rebound and settling behavior of landers using rigid-body contact models [26, 27], and Murdock, et al. have studied low-velocity impact dynamics in granular media [33]. However, for the motion planning problem, we only need

a model of the rebound *distribution*,

$$\vec{v}_{t_f^+} \sim P_\zeta(\vec{v}_{t_f^-}, \zeta(\vec{r}_{t_f})), \quad (19)$$

where the rebound velocity, $\vec{v}_{t_f^+}$, is a random variable dependent on the pre-impact velocity, $\vec{v}_{t_f^-}$, and a parametrized description of the surface properties, $\zeta(\vec{r}_{t_f})$ (e.g., surface friction and elasticity). This allows us to abstract away the detailed contact physics in a general way, whereby the rover can be modeled as a dimensionless particle. In practice, P_ζ should be fit to the rebound dynamics observed on a higher fidelity model, including surface shape uncertainty. For the data used in Section 3.1.3, we use a kernel density estimator to fit P_ζ to the rebound dynamics of a cube impacting a flat surface with some friction and elasticity. The simulation is stopped when $\|\vec{v}_{t_f^-}\| < v_{\min}$. An example of 20 Monte Carlo simulations is shown in Figure 24.

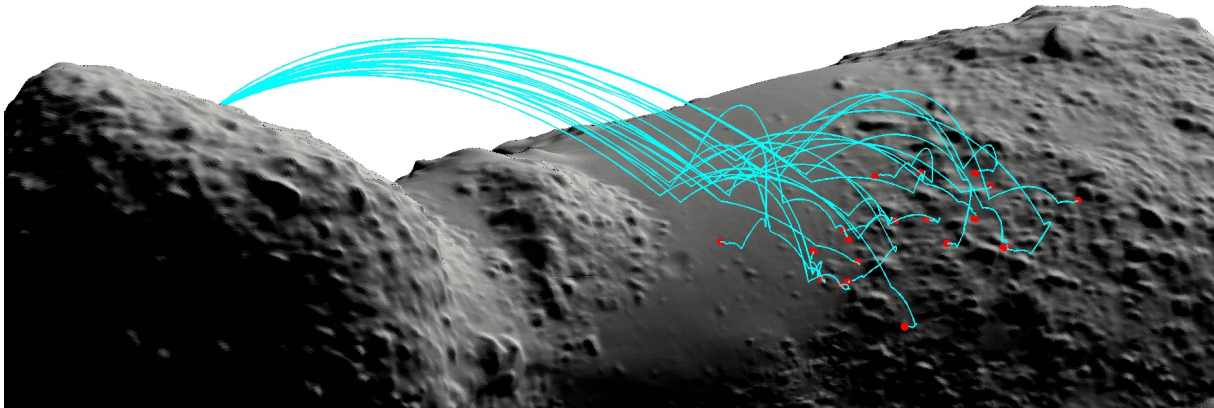


Figure 24: Monte Carlo simulation of a single hopping trajectory subject to control, gravity, and rebound uncertainty. The rover is modeled as a particle and is subject to forces based on Eq. (17).

3.1.2 Single-Hop Planning

Before addressing the inherently sequential planning problem, we first consider the simpler problem of planning a single hop. However, even this problem is far from trivial, as the rover must contend with highly nonlinear dynamics (discussed in Section 3.1.1) and many sources of uncertainty. In this section, we develop a framework for studying stochastic hopping trajectories and extracting approximately optimal (myopic) control policies.

Impact Targeting

In some cases, it may be desirable for a hopper to target a specific touch down location in a single hop (e.g., if the impact rebounds are expected to be minimal). The problem of computing the launch velocity (\vec{v}_{t_0}) to intercept a target location (\vec{r}_{t_f}) at time $\tau = t_f - t_0$ is the well-known ‘‘Lambert orbital boundary-value problem,’’ and has efficient numerical

solutions for spherical [34] and perturbed [35] gravity fields. However, for a polyhedral gravity model, a shooting method is required, which relies on good initial guesses for convergence.

Accordingly, we propose an algorithm that procedurally solves for the set of initial velocities (\vec{v}_{t_0}) that correspond to a *range* of flight times $\tau = [\tau_1, \dots, \tau_n]^T$. Algorithm 1 leverages three key insights to robustly and efficiently compute the solution set. First, the dynamics model in the shooting solver (line 3) *ignores collisions* with the surface, which leverages the fact that the polyhedral gravity model is also valid inside the body [32]. Avoiding collision checks makes the dynamics continuous and differentiable and drastically speeds up integration. Surface penetration is checked for feasibility only after the solution has converged (line 6). The second key insight is that gravity has a second order effect on position, and thus has more influence the longer it is integrated. Thus, by setting τ_1 sufficiently small, the solution is close to a straight line between \vec{r}_{t_0} and \vec{r}_{t_f} , and, although it may often be infeasible (or impractical), this serves as a robust initialization for subsequent solutions for larger τ . Finally, we leverage the differentiability of the dynamics to make a good initial guess of $\vec{v}_{t_0}(\tau_{i+1})$ given the solution for $\vec{v}_{t_0}(\tau_i)$ and the Jacobian, $J(\tau_i)$ (lines 4-5).

Algorithm 1 Procedural Lambert Solver for Irregular Gravity Fields

Input: $\vec{r}_{t_0}, \vec{r}_{t_f} \in \mathbb{R}^3$, $\tau = [\tau_1, \dots, \tau_n]^T \in \mathbb{R}^n$ s.t. $\tau_{i+1} > \tau_i$, gravity field $\vec{g}(\vec{r})$

- 1: initialize guess for $\vec{v}_{t_0}(\tau_1)$ (e.g., as Lambert solution in spherical gravity field)
- 2: **for** $i = 1, \dots, n$ **do**
- 3: Use shooting solver to solve for $\vec{v}_{t_0}(\tau_i)$ and Jacobian, $J(\tau_i)$:
 $\vec{v}_{t_0}(\tau_i), \vec{v}_{t_f}(\tau_i), J(\tau_i) \leftarrow \text{Solve}(\vec{r}_{t_0}, \vec{r}_{t_f}, \vec{g}(\vec{r}), \vec{v}_{t_0}(\tau_i))$
- 4: Compute first-order estimate of partial \vec{v}_{t_0} w.r.t. τ : $\frac{\partial \vec{v}_{t_0}}{\partial \tau} \leftarrow J(\tau_i)^{-1} \vec{v}_{t_f}(\tau_i)$
- 5: Initial guess for τ_{i+1} : $\vec{v}_{t_0}(\tau_{i+1}) \leftarrow \vec{v}_{t_0}(\tau_i) + \frac{\partial \vec{v}_{t_0}}{\partial \tau}(\tau_{i+1} - \tau_i)$
- 6: isValid(i) \leftarrow Check for trajectory surface penetration
- 7: **end for**

Output: $\vec{v}_{t_0}, \vec{v}_{t_f} \in \mathbb{R}^{3 \times n}$, $J \in \mathbb{R}^{3 \times 3 \times n}$, isValid $\in \{0, 1\}^n$

Lambert’s problem is known to have multiple solutions for a given τ : two for each integer number of orbits. Although Algorithm 1 will always return the most direct family of solutions (i.e., shortest path with zero orbits), other families of solutions may be found through a bisection search on \vec{v}_{t_0} . Figure 25 illustrates an example of three such families of solutions for a given $(\vec{r}_{t_0}, \vec{r}_{t_f})$ pair, which vary in duration from 25 minutes to 4 hours (blue trajectories represent the nominal, most direct solution family). Interestingly, the fact that *three* families of solutions exist within a single orbit—albeit, not all feasible—contradicts the Lambert solution for spherical gravity, thus illustrating the importance of high-fidelity gravity models.

Uncertainty Propagation

Computing the nominal Lambert solutions for planning a single hop helps to inform what trajectories may be beneficial for targeting a specific impact location, but it assumes a perfect gravity model, perfect control accuracy, and perfect state information. It is also

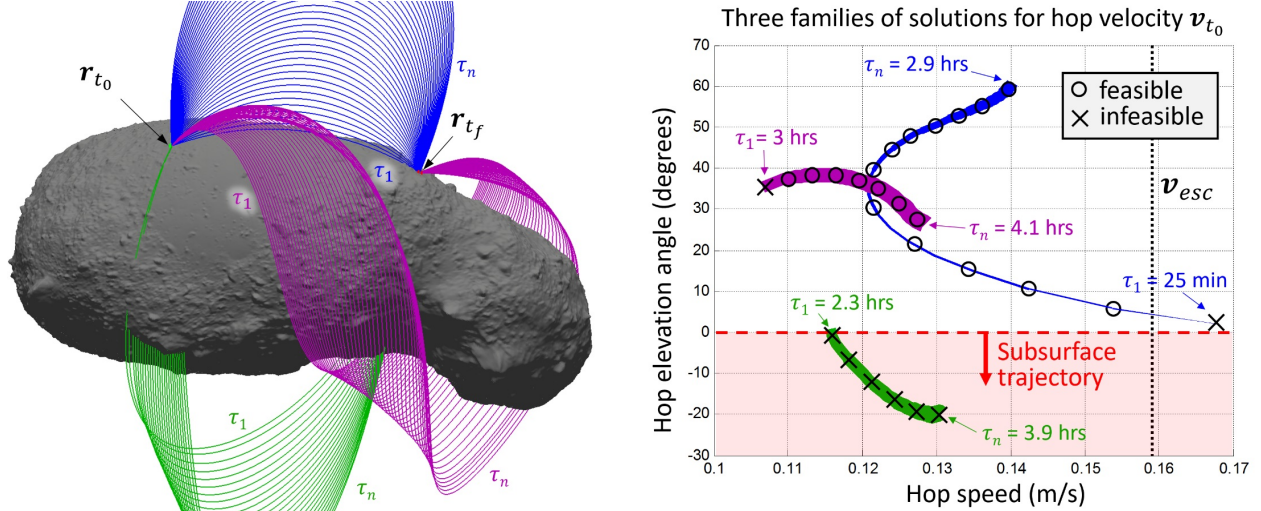


Figure 25: **Left:** Three families of hopping solutions computed by Algorithm 1, ranging from short and direct to long and winding. **Right:** Plot of hopping speed, $\|\mathbf{v}_{t_0}\|$, vs. elevation angle (with respect to the local surface plane). Line thickness is proportional to τ for that solution.

important to understand how sensitive these solutions are to various sources of uncertainty. In other words, we would like to predict the impact *distribution* by treating the gravity field, control accuracy, and state estimate as random variables rather than known quantities.

In general, there are two approaches to density estimation for nonlinear functions of random variables: (1) analytical propagation of simplified uncertainty models through linearized dynamics, and (2) sampling-based techniques. While sampling-based techniques (e.g., kernel methods) are amenable to arbitrarily complex dynamics and uncertainty models, they typically assume some measure of “local smoothness” and do not scale well to higher dimensions. On the other hand, analytical methods can be much more sample efficient for high-dimensional uncertainty models (e.g., by approximating gradients), but are often restricted to simple uncertainty models and locally linear dynamics. The ballistic dynamics of hopping are indeed smooth and linearizable, but also depend on the collision with a highly irregular surface. Accordingly, we decompose the density estimation problem into two phases: (1) Gaussian error propagation through the linearized ballistic dynamics, and (2) projection onto the irregular surface mesh.

For error propagation through the ballistic dynamics, we assume Gaussian uncertainty on the gravity, \vec{g} , according to Eq. (18) and Gaussian uncertainty on the control (\vec{v}_0) and initial state (\vec{r}_0) according to $\vec{v}_0 \sim \mathcal{N}(\mu_{\vec{v}_0}, \Sigma_{\vec{v}_0})$, and $\vec{r}_0 \sim \mathcal{N}(\mu_{\vec{r}_0}, \Sigma_{\vec{r}_0})$. More generally, the joint uncertainty of \vec{g} , \vec{v}_0 , and \vec{r}_0 may be expressed as,

$$\begin{bmatrix} \delta \vec{g} \\ \vec{v}_0 \\ \vec{r}_0 \end{bmatrix} \sim \mathcal{N}(\mu_{\delta \vec{g}, \vec{v}_0, \vec{r}_0}, \Sigma_{\delta \vec{g}, \vec{v}_0, \vec{r}_0}), \quad \mu_{\delta \vec{g}, \vec{v}_0, \vec{r}_0} = \begin{bmatrix} \vec{0} \\ \mu_{\vec{v}_0} \\ \mu_{\vec{r}_0} \end{bmatrix}, \quad \Sigma_{\delta \vec{g}, \vec{v}_0, \vec{r}_0} = \begin{bmatrix} \Sigma_{\delta \vec{g}} & \mathbf{0} \\ \mathbf{0} & \Sigma_{\vec{v}_0, \vec{r}_0} \end{bmatrix}, \quad (20)$$

where $\Sigma_{\delta \vec{g}} \in \mathbb{S}_+^k$ is the covariance of the gravity model, and $\Sigma_{\vec{v}_0, \vec{r}_0} \in \mathbb{S}_+^6$ is the joint state-control covariance. A two-sided finite difference approximation of the Jacobian, $J_{\vec{v}_0, \vec{r}_0} = [\partial \vec{r}_f / \partial \theta \quad \partial \vec{r}_f / \partial \vec{v}_0 \quad \partial \vec{r}_f / \partial \vec{r}_0] \in \mathbb{R}^{3 \times (k+6)}$, can be approximated with $2(k+6)$ simulations (note that $\partial \vec{r}_f / \partial \vec{v}_0$ is obtained for free from Algorithm 1). With this linear approximation, the impact covariance about \vec{r}_f can be computed as $\Sigma_{\vec{r}_f} = J_{\vec{v}_0, \vec{r}_0} \Sigma_{\delta \vec{g}, \vec{v}_0, \vec{r}_0} J_{\vec{v}_0, \vec{r}_0}^T$.

To get an impact distribution *over the surface*, we then project this covariance along \vec{v}_f , whereby the probability of impact on any facet can be computed as,

$$P(\vec{r}_{t_f} \in F_i) \approx a_i \left(\frac{e^{-\frac{1}{2} \vec{r}_{F_i}^T \Sigma_{uw}^{-1} \vec{r}_{F_i}}}{-2\pi \sqrt{|\Sigma_{uw}|}} \right) \max \left\{ \hat{N}_i \cdot \hat{c}_v, 0 \right\}, \quad (21)$$

$$\text{where } \Sigma_{uvw} = {}^c R^b \Sigma_{\vec{r}_f} {}^c R^{bT} = \begin{bmatrix} \Sigma_{uw} & \sigma_u \sigma_v \\ \sigma_v \sigma_u & \sigma_w \sigma_v \\ \sigma_v \sigma_u & \sigma_v \sigma_w & \sigma_v^2 \end{bmatrix}. \quad (22)$$

Here, R_c^b is the rotation matrix from \hat{b} to \hat{c} , $\vec{r}_{F_i} \in \mathbb{R}^2$ is the vector from \vec{r}_{t_f} to the center of facet F_i (in $\langle \hat{c}_u, \hat{c}_w \rangle$), and a_i is the area of facet F_i (see Figure 26, left).

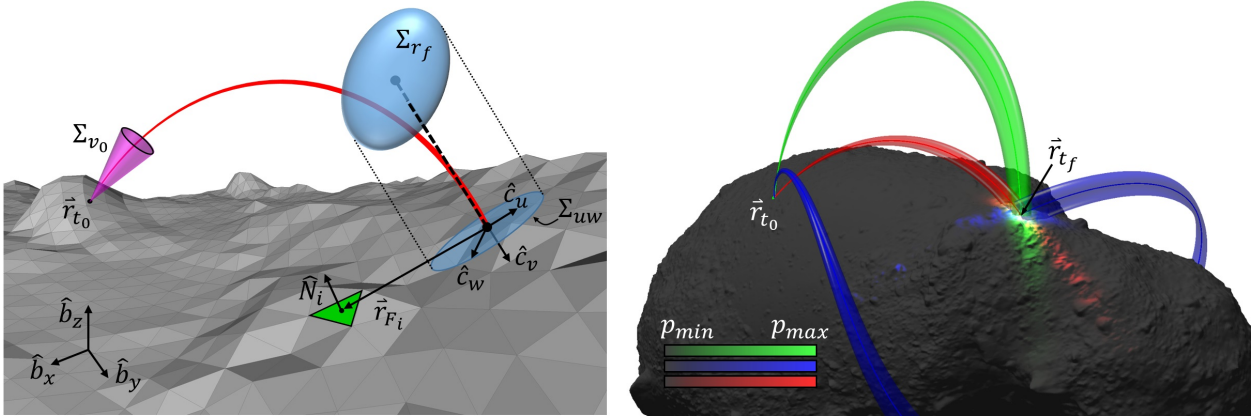


Figure 26: **Left:** Landing distributions are computed using covariance propagation and projection. **Right:** Error propagation for three trajectories on Itokawa, corresponding to position uncertainty, $\Sigma_{r_0} = I\sigma_{r_0}^2$, $\sigma_{r_0} = 1\text{ m}$, velocity error $\Sigma_{v_0} = I\sigma_{v_0}^2$, $\sigma_{v_0} = 0.03\|\vec{v}_0\|$, and gravity error $\sigma_g = 0.03\|\vec{g}\|$.

Figure 26 (right) illustrates this two-phase error propagation method on three trajectories from the example in Figure 25, where trajectory funnels bound a 90% confidence envelope, and surface color denotes projected uncertainty distribution, Σ_{r_f} . Interestingly, these funnels are not always divergent along the trajectory as one might expect (e.g., the blue trajectory). While the total magnitude of the error ($|\Sigma_{uvw}|$) does generally grow, its projection along the trajectory (i.e., $|\Sigma_{uw}|$) can shrink, indicating that errors can grow in *time* while shrinking in the cross-track directions. Also, while this analysis assumed Gaussian uncertainty, it can be trivially extended to Gaussian mixture models by simply taking a weighted sum of the components.

One drawback of the projection in Eq. 21 is that it does not capture “shadowing” effects, which is the case whenever $\hat{N}_i \cdot \hat{c}_v > 0$. For surfaces that are highly irregular in the vicinity of \vec{r}_{t_f} or for shallow, glancing impacts, Eq. 21 may be augmented with a “ray-tracing” collision checker to accurately project shadows—an established, albeit more computational technique commonly used for graphics rendering.

Optimal Hop Selection

Equipped with an algorithm to compute exact solutions for Lambert’s two-point boundary value problem (Section 3.1.2) and a method for propagating uncertainty to compute

approximate landing distributions (Section 3.1.2), we can now formulate an optimization problem for selecting the best (nominal) hop velocity, $\mu_{\vec{v}_0}^*$. We consider the following optimization problem:

$$\begin{aligned}
 & \underset{\mu_{\vec{v}_0}}{\text{minimize}} && J(\mu_{\vec{v}_0}) = \mathbb{E} [\lambda_u \vec{v}_0^T S_u \vec{v}_0 + \lambda_T \tau(\vec{r}_0, \vec{v}_0) - V(\vec{r}_f | \vec{r}_0, \vec{v}_0)] \\
 & \text{subject to} && \mu_{\vec{v}_0} \in \mathcal{A}(\vec{r}_0) \\
 & \text{where} && \vec{v}_0 \sim \mathcal{N}(\mu_{\vec{v}_0}, \Sigma_{\vec{v}_0}), \quad \vec{r}_0 \sim \mathcal{N}(\mu_{\vec{r}_0}, \Sigma_{\vec{r}_0})
 \end{aligned} \tag{23}$$

where the additive cost function, J , represents an expectation of the control effort (weighted by λ_u), flight time, τ (weighted by λ_T), and the negative “value,” V , of impacting at location \vec{r}_f . The surface value map, V , may encode various mission objectives including the distance to the goal, possible hazards, or even the expected future rewards in the context of sequential hopping (to be addressed in Section 3.1.3). The constraint on $\mu_{\vec{v}_0}$ belonging to the action space of the rover, \mathcal{A} , at position estimate, $(\mu_{\vec{r}_0}, \Sigma_{\vec{v}_0})$, can be quite naturally constructed as the intersection of two convex sets representing the *speed*, $\|\mu_{\vec{v}_0}\| \leq v_{\max}$, and *direction*, $\cos^{-1}(\mu_{\vec{v}_0} \cdot \hat{N}(\vec{r}_0) / \|\mu_{\vec{v}_0}\|) \leq \pi/2 - \theta_{\min}$, which is a cone about the local surface normal vector, $\hat{N}(\vec{r}_0)$, bounded by the minimum elevation angle, θ_{\min} .

Assuming that V primarily encodes a “distance-to-goal” metric, choosing one solution from each homotopy class of trajectories (e.g., the blue and pink solution families in Figure 25) can serve as good initializations for a gradient descent solver. An estimate for the global minimum can then be obtained by comparing the local minimum obtained from each homotopy class. For example, in the case where $\lambda_T \gg \lambda_u$, $\mu_{\vec{v}_0}^*$ belongs to the most direct class (blue trajectories in Figure 25), whereas for $\lambda_u \gg \lambda_T$, $\mu_{\vec{v}_0}^*$ belongs to the longer, albeit less energy intensive solution class (pink trajectories in Figure 25). The details of the gradient descent solver are omitted for brevity, but at a high level, $\mathbb{E}(V)$ can be derived from Eq. (21) and finite difference estimates of the cost gradient, $\nabla_{\mu_{\vec{v}_0}} J$, can be obtained in a similar fashion to the Jacobian in Section 3.1.2.

3.1.3 Sequential Hop Planning

Recall the sequential autonomy architecture illustrated in Figure 21. For mobility tasks that require multiple hops (e.g., traversing long distances or correcting for unfavorable bouncing), we must extend the myopic strategies developed in Section 3.1.2 for planning over a longer horizon. This requires some notion of how immediate actions facilitate future actions and how this sequence of actions achieves certain mission objectives. A natural framework for modeling this inherently discrete and stochastic planning problem is a Markov Decision Process (MDP). In contrast to “classical” open-loop motion planning algorithms (e.g., combinatorial and sampling-based) that search for feasible (or even “optimal”) reference trajectories, MDPs provide a more explicit representation of uncertainty and a powerful reward structure for encoding more complex mission objectives (i.e., not just “steering towards the goal”). The next section outlines how the planning problem can be structured as an MDP. We then discuss a sample-efficient reinforcement learning method for learning approximate state-action value (Q-) functions and implicitly, approximately optimal policies. Finally,

we compare the performance of learned control policies to heuristics proposed in [13], and evaluate performance robustness to modeling errors.

Hopping as an MDP

In accordance with the “classic” infinite horizon MDP formulation, we cast the sequential hop planning problem as the five-tuple, (S, A, T, R, γ) —the state space, action space, transition model, reward model, and discount factor. However, unlike most planning problems in robotics that force an MDP structure by temporally discretizing an inherently continuous-in-time process, hopping has a natural sequential decomposition, where transitions are marked by the eventual settling of each hop.

State space: At rest on the surface, the rover’s state is simply its position and orientation. However, assuming that a lower level controller can reorient the rover as needed (as is the case for a spacecraft/rover hybrid), we can collapse the state to just the surface position—an irregular manifold in \mathbb{R}^3 . For nearly-spherical bodies, spherical coordinates (i.e., latitude/longitude) may be sufficient to uniquely parametrize the surface, but for highly irregular bodies such as those in Figure 22, more elaborate map projections may be required. More generally, other state formulations might also include the rover’s internal state (e.g., battery charge), its state history (e.g., in the context of a coverage problem), or its *belief* state (in the case of partial observability).

Action space: In its most fundamental form, the action space of a hopping rover can be described by raw actuators (e.g., three motors and brakes for a spacecraft/rover hybrid). However, for motion planning it is more convenient to consider the rover’s action as its velocity immediately after hopping—a higher level abstraction of the action space that leverages a lower level hopping controller (e.g., the controller discussed in Sect. 2) and is amenable to the ballistic particle simulator presented in Section 3.1.1.

Moreover, it is critical that this velocity vector be expressed in a *global* reference frame—rather than a local surface frame—to maintain continuity in the transition dynamics; informally, $T(\cdot|s, a) \approx T(\cdot|s + \delta s, a)$. In other words, action descriptions that depend on the local surface slope (e.g., “spin flywheel number 2” or “hop left”) can yield sharp changes in T for small changes in state on irregular terrain, whereas global descriptions (e.g., “hop north”) are unaffected by local changes in topography.

However, local surface properties impose critical constraints on the feasible action space of the rover, $\mathcal{A}(s)$ (e.g., that the velocity vector must lie within some “friction cone” about the local surface normal). Thus, expressing actions in a global frame helps to “smooth” the dynamics (and consequently, the Q-function) but comes at the cost of requiring sharp discontinuities in $\mathcal{A}(s)$, suggesting that it may be advantageous to store a policy *implicitly* through a Q-function approximator (i.e. $\pi(s) = \arg \max_{a \in \mathcal{A}(s)} \hat{Q}(s, a)$) rather than explicit function approximators on π .

Reward model: The reward model is a mission designer’s tool for encoding various mission objectives, such as “visit sites A, B, and C under time and energy constraints while avoiding hazards D and E.” Thus, a reward function can take many forms and, in general, may be updated based on new information gathered or new objectives. We consider a general

formulation that penalizes the time and energy required for each hop, incentivizes a set of n_g goal regions, and penalizes a set of n_h hazardous regions. In summary, the state, action and reward models we consider here are:

$$\begin{aligned}
 s &\in S^2, \quad a \in \mathcal{A}(s) \subset \mathbb{R}^3, \\
 R(s, a) &= -\mathbb{E}[\tau(s, a)]/\tau_{\max} - \lambda_u \mathbb{E}[u(s, a)], \quad R(s_{g_i}) = r_i, \quad R(s_{h_j}) = r_j,
 \end{aligned} \tag{24}$$

where $\mathbb{E}[\tau(s, a)]$ and $\mathbb{E}[u(s, a)]$ are the expected time and energy required to execute action a at state s . States $s_{g_i} \in S_{g_i}$ and $s_{h_j} \in S_{h_j}$ are states within the goal and hazard regions, with associated rewards $r_i > 0$ and $r_j < 0$, respectively. τ_{\max} is a maximum travel time, and λ_u weights the control effort.

Reinforcement Learning Method

The transition model, T , is unknown. In the case of minimal bouncing, approximations of the dynamics such as those discussed in Section 3.1.2 may work, but in general, a series of elastic bounces (e.g., Figure 24) induces chaotic, highly non-Gaussian, multimodal transition dynamics. Without discretization of the state or action spaces, T is extremely difficult to approximate. Accordingly, model-free methods are better suited for this domain, whereby simulations (discussed in Section 3.1.1) can be used to generate large sets of transition data offline.

One popular technique for such batch, off-line, off-policy, model-free RL is Least Squares Policy Iteration (LSPI), which, as originally described in [36], uses a linear function approximator for the Q-function, and an exact, fixed-point projection for policy evaluation (LSTD-Q):

$$\text{solve: } (\hat{A}_n - \gamma \hat{B}_n)\theta = \hat{b}_n, \quad \text{where} \quad \hat{b}_n = \frac{1}{n} \sum_{i=1}^n \rho_i [\phi(s_i, a_i) r_i], \tag{25}$$

$$\hat{A}_n - \gamma \hat{B}_n = \frac{1}{n} \sum_{i=1}^n \rho_i \phi(s_i, a_i) [\phi^T(s_i, a_i) - \gamma \phi^T(s'_i, \pi(s'_i))],$$

$$\pi(s) = \operatorname{argmax}_{a \in \mathcal{A}(s)} [\phi^T(s, a)\theta]. \tag{26}$$

Here, $\phi(s, a)$ is the state-action feature vector, and ρ is a weight vector that sums to one. Like most approximate RL methods, LSPI is not guaranteed to yield optimal policies, but it is a stable algorithm. That is, it will either converge or it will oscillate in an area of the policy space where policies have suboptimality bounded by some approximation error, ε , which is highly dependent on the richness of the feature space and the coverage of the sampling distribution [36]. In practice, this bound is often quite conservative and LSPI typically converges in very few iterations.

One of the most important features of LSPI for deriving hopping policies is its amenability to off-policy exploration strategies, which provides the ability to reuse large data sets, and thus, relearn a policy for a new reward structure on the fly. The weight vector, ρ , provides a convenient way to preferentially bias previously collected samples via *importance*

weighting (i.e., $\rho_i = p(s_i, a_i)/q(s_i, a_i)$, where p is the desired distribution and q is the sample distribution); q may be approximated directly from samples via kernel density estimation, and p may be chosen arbitrarily (e.g., a uniform distribution).

Practical Considerations

Data Collection: An off-line simulator provides large flexibility for data collection. For a given initial state, s_0 , and policy, π , state-action samples can be biased towards more likely regions (e.g., through direct Monte Carlo sampling, or importance sampling/variance reduction techniques). However, in the more general case when s_0 is uncertain, or the reward structure may change (and thus, π), we would like a good fit of \hat{Q} over a much *broader* range of the state-action space. Thus, a mostly “pure exploration” strategy is preferred, with a combination of full episode rollouts and periodic restarts, perhaps with some bias towards “hard-to-reach” regions.

Feature Engineering: Linear function approximation relies on a rich set of features over the state-action space to produce good estimates for the Q-function (i.e., small $\|\phi^T(s, a)\theta^* - Q^*\|_2$). At the same time, the feature set must be amenable for computing the argmax in Eq. (26) for policy extraction. Accordingly, we decouple state and action features such that $\phi(s, a) = \phi_s(s)\phi_a(a)$, where ϕ_s can be arbitrarily complex in the state while ϕ_a remains “simple” enough for optimization. We construct ϕ_s from a set of “hand-crafted” features that leverage domain knowledge (e.g., local geopotential and surface slope) and a set of distributed basis functions—namely, a Fourier basis (similar to [37]), which is more naturally suited to spherical state domains than say, polynomials or RBFs. The action features are simple monomials of the form $\phi_a(\vec{v}) = v_x^i v_y^j v_z^k$, where $i + j + k \leq m$, such that a polynomial root solver can compute all local minima exactly.

Evaluation of Learned Policies

As a preliminary case study, we consider a notional mission scenario on Asteroid Itokawa in which the hopper must reach a target location in minimum time. The reward model (Eq. (24)) is defined as $R(s_g) = +1$ and $R(s, a, s') = -\tau/\tau_{\max}$, where $\tau_{\max} = 10$ hrs, and $\gamma = 1^4$. Approximately five million trajectories were simulated over a broad range of the state-action space, sampling from a Gaussian distribution on the gravity field ($\sigma_g \sim 5\%$) and a rebound distribution as discussed in Section 3.1.1, with a mean restitution of 0.65. Uniformly random mini-batches of size 100,000 were used in each iteration of LSPI. The action space, $\mathcal{A}(s)$, is speed constrained by $v_{\max} = 10$ cm/s and direction constrained by $40^\circ \leq \beta \leq 50^\circ$, where the elevation angle, β , must lie within an annular cone about the local surface normal (in accordance with the hopping constraints for a spacecraft/rover hybrid, derived in [5]). With this problem definition, LSPI was able to converge to a small error, $\|\theta_i - \theta_{i-1}\| < 10^{-6}$, in only a few tens of iterations and within tens of minutes on a laptop (though, significant speedups may be achieved with a more efficient implementation).

⁴Non-discounting is stable since the reward is always negative and each episode must terminate.

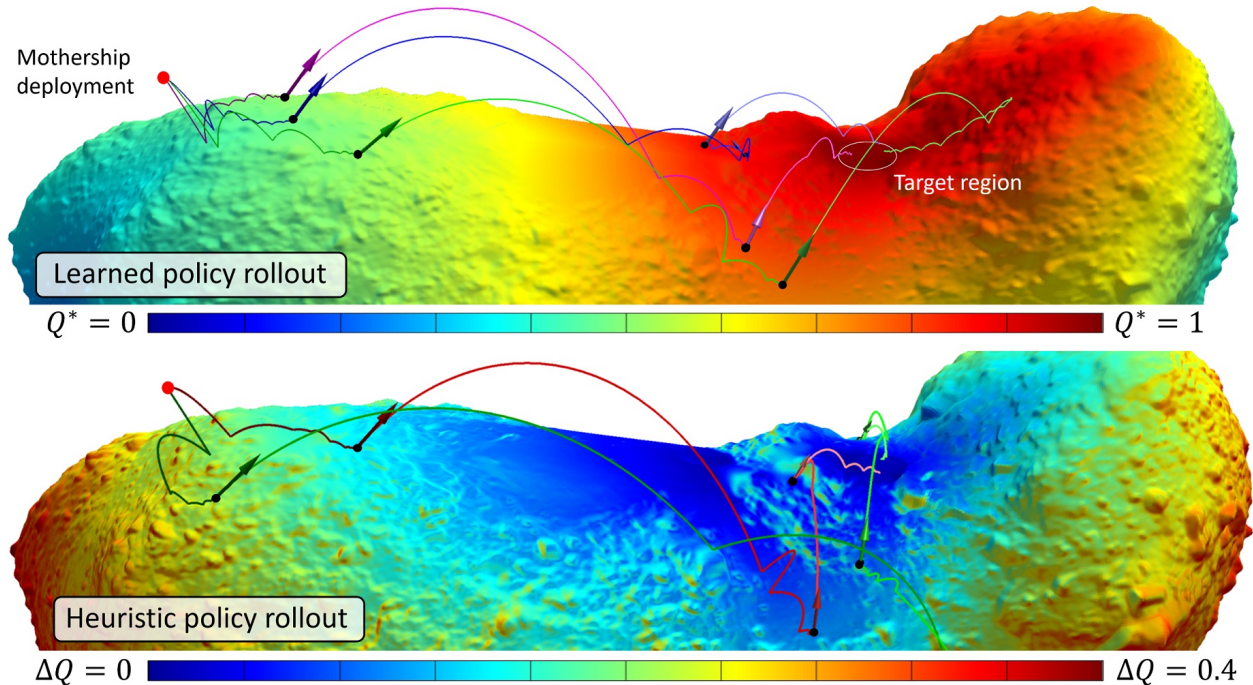


Figure 27: Top: Three rollouts of the learned policy. The surface color map shows the optimal value function under π^* . **Bottom:** Two rollouts of the “hop-to-the-goal” heuristic policy, where the color map shows the *difference* between the learned and heuristic value functions ($\Delta Q = Q^{\pi^*} - Q^{\pi_h}$).

Figure 27 shows a few example trajectories comparing the performance of the learned policy, π^* , with a “hop-towards-the-goal” heuristic policy, π_h , that attempts to take the most direct path to the goal. From 1000 policy rollouts, the mean time to reach the goal from deployment was 5.1 hours for the learned policy, and 7.6 hours for the heuristic policy. The color map in the top figure shows the optimal value, Q^* , at each point on the surface, and due to the reward structure defined above, it also represents the expected time to reach the goal (as a function of τ_{\max}). Not surprisingly, Q^* decays away from the goal region. The color map in the bottom figure shows the value *margin* of the optimal policy over the heuristic policy, suggesting that the heuristic policy has difficulties on sloped surfaces and states farther from the goal. The second hop of the blue trajectory in the top figure illustrates how the learned policy enables the hopto perform local adjustments to better position itself for future hops—in this case, by performing a small backwards hop off of a sloped region. The last hop of the green trajectory highlights another interesting learned behavior: the rover hops uphill from the goal region, “understanding” that it is likely to tumble downhill, thereby *leveraging* the dynamics of the environment without ever having explicitly learned a model.

An important consideration when learning in simulation and executing in the real world is robustness to modeling errors. This “transfer learning” problem for asteroid environments may have three major types of modeling errors: (1) the asteroid’s shape, (2) its surface properties, and (3) its density/gravity. While this learning method does require an accurate shape model at a global scale, it is insensitive to smaller scale variations due to the global-frame representation of the action space. That is, unanticipated deviations in local topography only affect the constraints for policy extraction, not the optimal value function itself. As a preliminary study of robustness to contact modeling errors, we rolled out

the learned policy in an environment with a different contact model—specifically, one with higher surface elasticity and one with lower elasticity. For the *more* elastic case, 1000 policy rollouts exhibited significantly longer traverse times, with a mean of 8.2 hours. These trajectories often bounce off course or overshoot the goal, requiring major corrections. On the other hand, policy rollouts in an environment with *reduced* elasticity actually exhibit better performance, with a mean traverse time of only 4.7 hours. This result suggests that one should err on the side of overestimating the surface elasticity for simulations. Finally, although density/gravity models are likely to be fairly accurate from preliminary surveying by the mothership, simulations can use a conservative *underapproximation* of gravity for safe policy transfer (i.e., so that the hopperdoes not overshoot its target or reach escape velocity).

3.2 Localization

On-board localization, which is required for autonomous surface activities, is particularly challenging for hopping rovers on small bodies. Vision-based localization, which has seen extensive success in terrestrial applications, must contend with (1) high-contrast shadows that move due to changing illumination conditions associated with fast-rotating, airless bodies, (2) large image scale variations (centimeters to kilometers) associated with ballistic hops, (3) determining absolute scale from sensors with limited observability (impulsive acceleration measurements and narrow baseline stereo provide minimal information during large hops), (4) detecting loop closures, or *relocalizing*, between wide and narrow FOV images (the rover uses wide-FOV cameras to maximize observability, while the primary spacecraft uses narrow-FOV cameras to maximize distances and minimize risks to spacecraft), (5) lack of variety in surface features (e.g., rocks are often locally self-similar), (6) rapidly rotating camera views from tumbling rovers, which also results in frequent off-nadir pointing, and (7) frequent visual occlusions.

We investigated a collaborative visual localization method for hopping rovers that overcomes these challenges by adapting ORB-SLAM2, a state-of-the-art visual SLAM algorithm [38]. This approach consists of two phases. First, a global map of 3D landmarks is generated by an orbiting primary spacecraft with a narrow FOV camera, and this prior map is downloaded to the rover. Second, after deployment to the surface, the rover uses a wide-FOV camera to perform on-board visual SLAM to (1) provide localization, (2) map its surroundings, and (3) relocalize to the global prior map.

In Section 3.2.1, we review related approaches for localization on small bodies along with the state of the art in monocular visual SLAM. Section 3.2.2 describes how we adapted ORB-SLAM2 [38] to handle the visually challenging problem of rovers hopping on small bodies. In Section 3.2.3, we evaluate the robustness and accuracy of our method with a series of experiments. These experiments use realistic images of a mock asteroid with Sun-analog illumination, as well as both narrow and wide FOV cameras executing orbital and hopping motions.

3.2.1 Related work

In recent years, several techniques have been proposed that combine visual and inertial sensors for the relative localization of hopping/tumbling rovers on the surface of small bodies. While the MINERVA rover never validated its localization approach, Yoshimitsu et al. [39] describe how the rover estimates its attitude by fusing observations of the Sun and integrating gyroscope measurements. Relative velocities during hops were to be estimated using optical flow from surface images.

In [40], Fiorini et al. propose a localization algorithm for hopping robots that estimates *a posteriori* trajectories and the landing area of a robot by fusing camera images with accelerometer, gyroscope, and contact sensor measurements that were recorded during jumps.

Optical flow and visual odometry were explored by So et al. in [41] for the relative localization of a hopping rover that explores the surface of an asteroid. Their proposed algorithm works with tumbling camera motion and without continuous tracking of surface features. Estimates for the rover's ballistic motion are computed from visual odometry measurements captured at the start of each hop. A guided search of landmarks is used to estimate the rover's position along extrapolated trajectories.

Absolute metric scale is relatively unobservable using on-board sensors on hopping rovers. For terrestrial applications, accelerometers are frequently combined with visual odometry techniques to determine absolute scale [42]. However, during ballistic motions, hopping rovers are in constant free-fall and experience no acceleration. In [41], So et al. also describe how a stereo pair of cameras help recover absolute scale at the start of each hop. However, the small baselines afforded by hopping rovers (on the order of 10 cm) provide almost no scale information when hops exceed 10 meters in height. Thus, for the vision-based techniques described in the literature, hopping rovers are highly prone to position and scale drift while hopping and tumbling across the surface of small bodies.

Loop closures (the ability to recognize the same place when revisited) allow visual SLAM algorithms to correct for drift (accumulated errors) while exploring an environment. A review of the state of the art in visual SLAM can be found in [43]. Early approaches estimated pose and landmarks online using extended Kalman filters, while newer approaches use incremental batch optimization techniques such as bundle adjustment [44] and pose graph optimization [45]. ORB-SLAM2 [38], the visual SLAM algorithm modified for this work, integrates both bundle adjustment and pose graph optimization approaches to create an implementation that is both robust and efficient. Furthermore, ORB-SLAM2 uses Sim(3) to represent *constraints* in its pose graph optimization, which allows it to explicitly parameterize and correct for scale drift.

For robust tracking and estimation, ORB-SLAM2 uses Oriented Binary (ORB) features [38] to detect and track landmarks in camera images. These ORB features require minimal storage and are efficient to calculate, making them well-suited for space applications where computation is often limited. ORB features describe landmarks with 3D position estimates that are compressed into *words* (here integers) and placed in a *bag-of-words* database that is used to quickly detect loop closure hypotheses [46]. We refer the reader to [38] for more details.

3.2.2 Proposed Localization Method




Figure 28: (a) In a prior mapping phase, the primary (mother) spacecraft images the body as it rotates in view. Over many weeks or months, the spacecraft maneuvers slowly to image the entire surface. Throughout this section, longitude around the body’s rotational axis is denoted θ_a , while the relative surface illumination angle is denoted α_s . After the rover is deployed to the surface (red) its local illumination angle continues to vary from sunrise ($\alpha_s \simeq -90^\circ$) to noon ($\alpha_s = 0^\circ$) and to sunset ($\alpha_s \simeq 90^\circ$). (b) Graphical model for the prior mapping phase. K_j are keyframes, P_k are keyframe priors (e.g., ground-based pose estimates), X_l are 3D landmarks and $S_{j,j+1}$ are measurements of the relative Sim(3) transformation between keyframes.

This section describes how to adapt the ORB-SLAM2 algorithm to enable collaborative localization for a typical mission scenario to a small body (see Figure 28). During a prior mapping phase, the primary (mother) spacecraft images the body as it rotates in view. Over many weeks or months, the spacecraft maneuvers to image the entire surface from all longitudes, θ_a , and from multiple relative surface illumination angles, α_s . After deployment to the surface, the rover hops and tumbles accumulating both position and scale errors (the rover’s attitude is constrained by gyroscope sensors and occasional observations from a star tracker).

To reduce these errors, the rover occasionally performs large hops to capture views of the surface that it matches to the prior map (see Figure 28). Relocalizing to the prior map creates loop closure constraints in the pose graph. After maximum likelihood optimization, this corrects the rover’s on-board pose estimates and reduces its pose uncertainty to approximately the sum of the uncertainties in the prior map and the loop closure constraint (in practice, these uncertainties are difficult to recover accurately [47]).

Both the primary spacecraft and the hopping rover’s maps are composed of sets of keyframes K_j and landmarks X_l . Each keyframe stores: the SE(3) pose relative to the world frame ${}^w\mathbf{T}_j$, the camera intrinsic parameters, and the ORB descriptors and their image positions $\mathbf{x}_{j,l}$, which are possibly associated with landmarks. Each landmark includes its world position ${}^w\mathbf{X}_l$, the viewing direction (averaged over the camera viewing directions that observed it), the ORB descriptor and the maximum d_{max} and minimum d_{min} distances at which the point is likely to be observed. Landmarks can be augmented with an estimate of the local illumination angle α_s from when they were observed.

Prior mapping from the primary spacecraft

In a typical mission to a small body, a navigation team on the ground uses additional

measurements such as radio doppler to estimate the primary spacecraft’s orbital trajectory. These measurements allow absolute scale to be estimated, and when combined with a star tracker can provide SE(3) priors on the spacecraft’s pose during the creation of the global map. These priors are added every n camera frames, as shown in the graphical model in Figure 28.

On the primary spacecraft, a *prior adjustment* thread has been added to ORB-SLAM2 that promotes camera images with SE(3) priors to keyframes and performs additional pose graph optimizations. These pose graph optimizations are performed over the *essential graph*. The essential graph connects the keyframe’s poses with estimated transformations, or graph constraints, that are either derived by 2D-to-3D matching (e.g., in the tracking front-end or during relocalization) or 3D-to-3D matching (e.g., loop closures). These constraints are effectively derived from large sets of ORB features that have been co-observed between keyframes, and they are represented by a Sim(3) Lie group. $\mathfrak{sim}(3)$ is the corresponding Lie algebra, represented by a 7-vector (ω, ν, σ) , where $\omega = (\omega_1, \omega_2, \omega_3)$ is the axis-angle representation of the rotation, ν is the rotated version of the translation \mathbf{t} and $\sigma = \log s$, where s is the scale [47]. The relation between $\mathfrak{sim}(3)$ to Sim(3) is given by:

$$\begin{pmatrix} \omega \\ \nu \\ \sigma \end{pmatrix} = \log \text{Sim}(3) \begin{bmatrix} s\mathbf{R} & \mathbf{t} \\ 0 & 1 \end{bmatrix}. \quad (27)$$

The pose graph optimization distributes residual errors between the various constraints in the essential graph. A residual is defined as:

$$\mathbf{e}_{j,j+1} = \log \text{Sim}(3)({}^w\mathbf{S}_{j,j+1} \quad {}^w\mathbf{S}_{j,w} \quad {}^w\mathbf{S}_{j+1,w}^{-1}), \quad (28)$$

where ${}^w\mathbf{S}_{j,j+1}$ is the relative Sim(3) transformation between the connected keyframes computed before the pose graph optimization expressed in the world reference frame w , and ${}^w\mathbf{S}_{j,w}$ is the Sim(3) transformation between the frame j and the world reference frame.

During the pose graph optimization, a “virtual” constraint is added between the tracked pose K_j and the prior pose P_k , and the relative transformation between them is set to identity, ${}^W\mathbf{S}_{j,k} = \mathbf{I}$. The constraint error between a tracked keyframe and its prior pose is defined as:

$$\mathbf{e}_{j,k} = \log \text{Sim}(3)(\mathbf{I} \quad {}^w\mathbf{S}_{j,w} \quad {}^w\mathbf{P}_{k,w}^{-1}), \quad (29)$$

where ${}^w\mathbf{P}_{k,w}^{-1}$ is the transformation between the prior pose and the world reference frame. The cost function in the ORB-SLAM2 essential graph optimization is augmented with the pose priors:

$$\chi^2({}^w\mathbf{S}_{2,w}, \dots, {}^w\mathbf{S}_{m,w}) = \sum_j \mathbf{e}_{j,j+1}^\top \mathbf{e}_{j,j+1} + \sum_j \mathbf{e}_{j,k}^\top \Omega_{\Delta\mathbf{s}_{i,k}} \mathbf{e}_{j,k}, \quad (30)$$

where $\Omega_{\Delta\mathbf{s}_{i,j}}$ is the diagonal information matrix for each prior:

$$\Omega_{\Delta\mathbf{s}_{i,k}} = \text{diag}(\sigma_r^2, \sigma_r^2, \sigma_r^2, \sigma_t^2, \sigma_t^2, \sigma_t^2, \sigma_s^2).$$

Here, σ_r , σ_t , σ_s are the priors’ rotation, translation, and scale standard deviation estimates, respectively. These values are manually tuned to ensure pose graph convergence (they are

sensitive to the scene scale and camera parameters—FOV, resolution, etc.). After the essential graph optimization, landmark positions are adjusted, and a global bundle adjustment is performed to refine the keyframe poses and landmark estimates.

Rover relocalization on the prior map

The hopping rover’s relocalization approach is very similar to ORB-SLAM2’s; the approach is described here for completeness. First, a *bag-of-words* technique, [46], provides a set of candidate keyframe matches via a fast database search. These candidate keyframes are filtered using two parameters k_{words} and k_{score} , which provide a lower bound on the number of shared words (ORB features), and a similarity score, respectively.

For each keyframe candidate that passes the filters, the ORB image features are matched to the current frame. A geometric consistency test is then performed between the 2D image features and the 3D landmarks using RANSAC and a PnP algorithm. If enough inliers are founded, the current frame is considered as “relocalized” and the relative pose is further optimized by considering all inliers. A search is performed via the essential graph to establish additional feature/landmark correspondences, and the keyframes’ and landmarks are optimized using bundle adjustment.

To adapt ORB-SLAM2 to enable relocalization between cameras with different focal lengths (e.g., wide and narrow FOVs) the following changes were implemented:

- The number of pyramid levels in the ORB feature extractor was increased from 8 to 11.
- The threshold on the minimum number of inliers in the RANSAC test was decreased from 50 to 10.
- The scale filtering was disabled to ensure new frames are tracked after global relocalization.

3.2.3 Experimental Evaluation

To demonstrate the proposed adaptation of ORB-SLAM2, a series of experiments were performed around a 1.4 meter diameter mock-asteroid in the JPL Robodome. To simulate realistic illumination changes, a collimated light was placed 5 m away from the asteroid and the asteroid was mounted so it could rotate around a single axis. Three configurations were tested: (1) both prior mapping (primary spacecraft) and visual SLAM (deployed rover) with static illumination, (2) prior mapping with the asteroid rotating (changing illumination) and visual SLAM with fixed illumination and (3) both prior mapping and visual SLAM with the asteroid rotating (most realistic). The prior mapping phase was performed with the primary spacecraft orbiting close to the asteroid’s equatorial plane. While the mock asteroid had a diameter of 1.4 m, all measurements have been normalized here to an asteroid with an effective diameter of 1000 m.

The rover’s wide FOV camera had a FOV of $100^\circ \times 82^\circ$ and a resolution of 640×480 pixels. The primary spacecraft’s narrow FOV camera had a FOV of $60^\circ \times 35^\circ$ and 1920×1080 pixels resolution. Both cameras were calibrated with Zhang’s calibration method [48]. The motion of both cameras and the asteroid were tracked using a Vicon motion capture system, which in this experimental configuration gave a 2 mm tracking accuracy.

In order to analyze algorithm performances and generate the prior map the timestamps of the camera and the Vicon system need to be synchronized. Moreover the two trajectories are captured in two different reference frames: the Vicon reference frame w and the camera reference frame c . The ORB-SLAM2 trajectory is expressed with reference to the first keyframe pose. In order to align the two timestamps, an initial guess of the time-shift was roughly estimated by identifying some interest point in the time-displacement plot. Then the Vicon data is interpolated on camera time stamp. At this step of the alignment phase there are two sets of 3D points that have the same dimension, which represent the trajectory in the two different frames of reference. We estimate the transformation rotation matrix ${}^w_c\mathbf{R}$, translation ${}^w\mathbf{t}_{c,w}$, and scale $s_{c,w}$ between the camera and the Vicon reference frame, and translation between the camera center and the camera rig ${}^r\mathbf{t}_{c,r}$ by solving the following non-linear optimization problem:

$$\underset{({}^w_c\mathbf{R}, {}^w\mathbf{t}_{c,w}, s_{c,w}, {}^r\mathbf{t}_{c,r})}{\operatorname{argmin}} \sum_i \| {}^w\mathbf{r}_i - s_{c,w} {}^w_c\mathbf{R} {}^c\mathbf{r}_i - {}^w\mathbf{t}_{c,w} + {}^w\mathbf{R} {}^r\mathbf{t}_{c,r} \|^2 \quad (31)$$

where ${}^w_r\mathbf{R}$ is the pose of the camera rig in the Vicon frame of reference.

Performance in an asteroid-like environment

In order to show that ORB-SLAM2 works in an asteroid-like environment the following video sequence has been processed with the algorithm. The narrow FOV camera is mounted on a cart that is manually moved around the asteroid mock up in an ‘orbit like’ way. The light direction, kept fixed, is parallel to the asteroid orbital plane. The RMS error is about 0.44 cm, the accuracy of the ground truth system (Vicon) is around 0.2 cm. The RMS angular error is equal to 0.2 degrees along the three reference frame axis.

Vicon pose is used in order to create the prior maps that will be used to localize the rover’s camera. Two maps with fixed illumination conditions have been realized: the first one with a sequence taken at a distance of 160 cm from the asteroid mock up (*far mapping*), and the second one is the result of three consecutive sequences closer and closer to the mock up, respectively 160 cm, 120 cm and 80 cm (*three distance mappings*). Figure 30 shows the position of the keyframes and the landmarks of one of the two maps with fixed illumination conditions. During the mapping phase, we have changed the prior insertion rate n and the information matrix values. We extracted around $N = 4000$ ORB features for each image.

Figure 31 shows the relocalization rates from prior maps with different prior insertion rates n and information matrix values, σ_r , σ_t and σ_s .

Relocalizing a wide FOV camera to narrow FOV priors

A video sequence with the wide FOV rover camera has been captured by moving it from




Figure 29: Example relocalization. Left: an image from a hopping rover’s low-resolution wide-FOV camera (640×480 pixels, $f = 274$ pixels). On the right, this image has been successfully matched to a high-resolution narrow-FOV camera image (1920×1080 pixels, $f = 1527$ pixels) from the primary spacecraft’s prior map.

the asteroid surface up to the mapping orbit distance. By using the approach described in Section 3.2.2 it is possible to localize the wide FOV frames over a map built with narrow field of view images. Figure 30 shows the wide FOV frame poses which have been localized over the map. Figure 29 shows feature matches between a wide FOV rover camera frame which has been localized on a narrow FOV map keyframe. In this example fixed illumination conditions have been considered.

Robustness to large scale changes

The localization accuracy of the wide FOV camera on the saved map has been evaluated as a function of the distance from the asteroid. The distance from the asteroid is the Euclidean distance from the frame position and its projection in the asteroid point cloud along the camera optical axis, as shown in Figure 30.

Figure 31 shows the percentage of localized frames as a function of the normalized distance from the asteroid. Localization performance has been evaluated by changing the number of map candidate keyframes returned by the bag-of-words library. The candidate keyframes are filtered using the two parameters k_{words} and k_{score} . As one can see from Figure 31 if we filter the candidate map keyframes, at least 17% of wide FOV camera frames are globally localized at a normalized distance of 250 m, this value grows up to 22% if we take all the keyframes in the map. At an height above a normalized distance of 250 m, the percentage of localized frames is 40% for the *far distance* map. The percentage of localized frames grows up to 60% if the map is built by extracting twice as many (4000) ORB features for each image.

Figure 32 shows how accurate the algorithm’s localization performance is as a function of the distance from the asteroid mock-up. The localization accuracy mean value goes from 16 m to 26 m.

Robustness to off-nadir camera pointing

The localization accuracy of the wide FOV camera on the saved map has been evaluated as a function of the off-nadir pointing direction. As off-nadir pointing angle we have considered

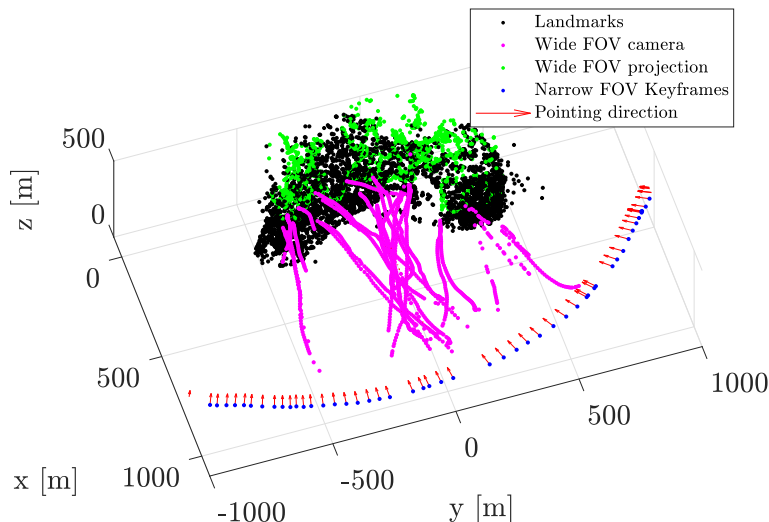


Figure 30: Successful rover relocalizations in the *far sequence* prior map. Keyframes from the primary spacecraft’s prior map are shown in blue, with red vectors indicating the camera’s optical axis. Dozens of simulated descents to the surface were performed in this image sequence; magenta markers indicate where single images from the rover’s wide FOV camera were successfully relocalized in the primary spacecraft’s prior narrow-FOV images. In this 6100 image sequence, about 40 percent of the rover’s wide FOV images were relocalized. The projection of the camera’s optical axis onto the body (green) allows the camera’s off-nadir pointing angle to be estimated. Distances are normalized to a 1000 m asteroid.




Figure 31: Percentage of successful rover relocalizations vs. the normalized distance from the asteroid surface for the simulated descents shown in Figure 30. These histograms indicate the percentage of the rover’s wide-FOV camera images that were successfully relocalized in the primary spacecraft’s narrow FOV image priors in the *far map* and *three distance map* priors. For both map priors, the best relocalization rates occurs when the number of ORB features per frame is increased to $N = 4000$. The parameters used in the bag-of-words keyframe candidate selection and RANSAC geometry consistency tests are also varied. Distances are normalized to a 1000 m asteroid.

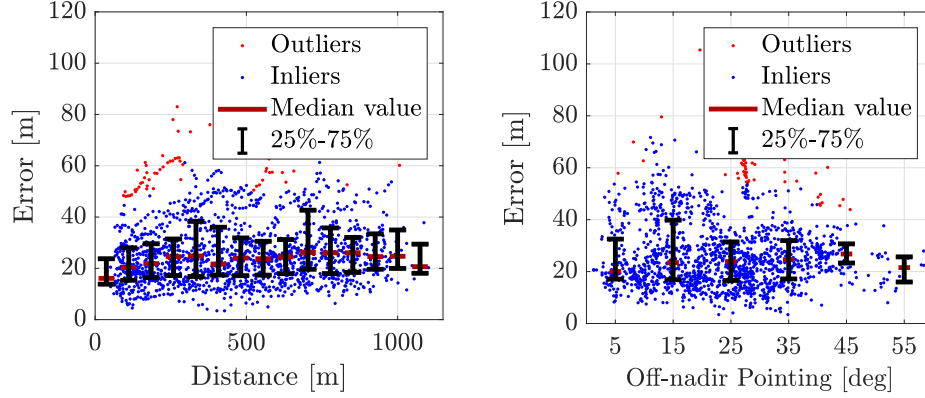


Figure 32: Relocalization error for the rover’s wide-FOV camera in the simulated descents shown in Figure 30 and the *three distance map* priors with $n = 20$ and $N = 4000$. Distances are normalized to a 1000 m asteroid. (a) Relocalization error vs. the normalized distance to the asteroid surface. (b) Relocalization error vs. the off-nadir pointing directions, which occurs when the rover tumbles.

the angle difference between the optical axis of the localized frame and the reference map keyframe optical axis. As one can see from Figure 32 (right) ORB-SLAM2 handles viewing direction changes up to 55 degrees.

Robustness to illumination changes




Figure 33: Histogram of successful wide-FOV camera relocalizations in the simulated descents (from Figure 30) vs. the rover’s longitudinal position θ_a above the asteroid’s surface for three prior maps. The three prior maps were captured with different surface illumination angles as measured between the primary spacecraft and the sun while the asteroid rotated ($\alpha_s = 23^\circ$, 53° and 82°). The wide-FOV simulated descents were recorded with a fixed $\alpha_s = 82^\circ$.

A series of experiments under different illumination conditions have been performed. The experimental setup allows us to rotate the asteroid and change the light incidence direction. During the tumbling rover operations illumination conditions could be very different from those during the mapping phase. The narrow FOV camera is mounted on a tripod, and the distance between the tripod and the asteroid is kept constant. The asteroid is rotated by keeping the light source angle of incidence constant. This allows the formation of a

circular orbit around the asteroid, where the local surface illumination is a fixed angle α_s . Meanwhile, the pose of the asteroid is recorded with the Vicon system. In addition, the pose of the illumination source is tracked during the tests.

First, a series of maps was generated by changing the illumination source position. α_s is the angle belonging to the orbital plane between the illumination source direction and the mapping camera position. Figure 33 (d-e-f) shows three of the generated maps. For each keyframe of the map there is a corresponding asteroid rotation angle θ_a , which is referenced to an inertial reference frame. Since the lighting conditions change as the asteroid rotates, two keyframes with the same asteroid rotation angle (θ_a) but different illumination direction (α_s) have different lighting conditions.

By localizing a wide FOV sequence that is captured with fixed illumination while rotating the asteroid, it is possible to estimate the limit value of the illumination angle difference until which ORB-SLAM2 becomes robust to lighting changes. The light source direction of the wide FOV sequence is the same as in the $\alpha_s = 82^\circ$ prior map. For these conditions when $\theta_a = 82 - \alpha_s$ the wide FOV frame has the same lighting conditions as the mapping frame. Figure 33 shows the histogram plot of the number of localized wide FOV camera frames as a function of the map keyframe position θ_a . As one can see, the map keyframe angles θ_a are spread in a range of 30 deg for all the generated prior maps. This suggests a localization invariance to lighting condition of $\Delta\theta_a = 15$ degrees.

While ORB features are relatively invariant to the illumination angle, for a body with a 12 hour rotation rate, $\pm 30^\circ$ translates into the requirement that the primary spacecraft should image the surface every hour. In order to extend the localization range, different prior maps generated with different illumination angles can be fused together. A map has been generated by overlapping a series of maps with $\Delta\alpha_s$ steps equal to 5 degrees. In this experiment the wide FOV sequence has been captured by rotating the asteroid, thus illumination conditions vary for each frame. At the bottom of Figure 34 we can see the number of localized frames for each illumination conditions and at the top we can see the corresponding average relocalization errors. Figure 34 suggests that some illumination conditions are unfavorable for mapping purposes. For example, if α_s is close to 0 degrees the amount of shadows on the asteroid surface decrease, thus decreasing the number of possible feature matches. On the other hand if α_s is close to 90 degrees a considerable part of the asteroid is characterized by shadows, which also limits correct feature matches. Figure 34 (b) shows the possibility to localize wide FOV frames all around the asteroid by exploiting its rotation and different illumination conditions.

3.3 Summary

The semi-autonomous exploration of small bodies with hopping rovers requires both planning and localization. We have developed a motion planning framework that treats the hopping dynamics as a stochastic system and leverages offline simulations to generate closed-loop control policies that are both safe and efficient. In other words, although hopping dynamics may be subject to uncontrolled bouncing, we have shown that through a series of subsequent “corrective” maneuvers, a hopper is able to achieve target precision on the order of one body length (roughly 20 cm), even when traversing large distances over highly irregular

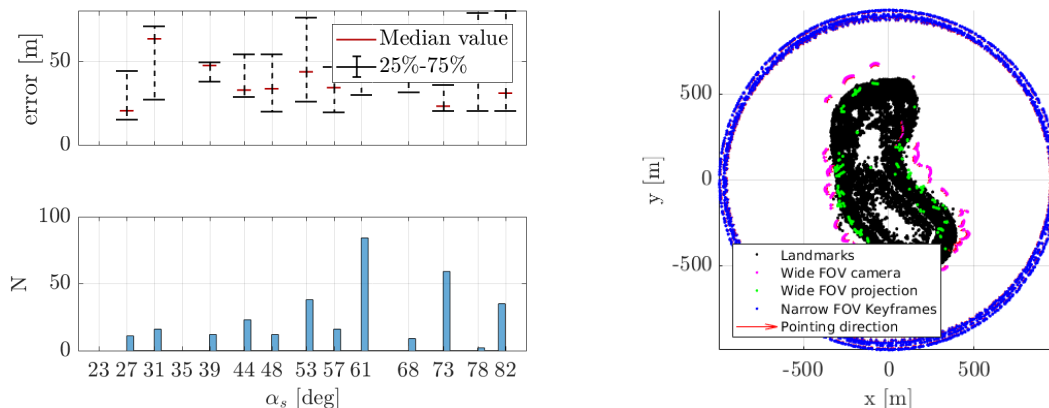


Figure 34: Wide FOV camera relocalization on a large map generated by fusing individual maps illuminated every $\alpha_s = 5^\circ$. Top: relocalization errors as function of α_s . Middle: the number of localized frames for each α_s map. Bottom: The final map and successfully relocalized wide FOV frames.




Figure 35: Pose errors for the long hopping sequence, where the rover uses its wide FOV camera to relocalize multiple times to the prior map while hopping across the mock asteroid's surface.

terrain. We have also developed a collaborative localization and mapping approach between a primary spacecraft and a deployed rover, that uses visual SLAM techniques adapted from the state-of-the-art ORB-SLAM2 algorithm. We have tested the algorithm and demonstrated its capabilities to work in an asteroid-like environment, and in particular, the ability to localize the rover with cm-scale precision while remaining robust to changes in lighting conditions and large image scale changes during hops. Future work should address two key aspects: (1) making the autonomy subsystem increasingly robust, for example, against modeling and localization errors and in the context of highly dynamic maneuvers (e.g., highly energetic hops), and (2) testing the overall autonomy subsystem (with planning and localization operating in closed-loop) in a high-fidelity simulation environment that captures cm-scale contact interactions as well as realistic environmental disturbances (e.g., dust).

4 Reference Mission Architecture: Spacecraft and Operations Concept

In order to demonstrate the applicability of the mobility approach to a feasible mission concept, a reference architecture was developed for a spacecraft/rover hybrid that would be deployed from a mothership. Several trade studies and systems engineering analyses were conducted as part of the trade space exploration. The resulting reference design serves to establish initial feasibility and as a starting point for more detailed designs with more specific requirements determined by a specific primary mission. As such, the architecture remains scalable and adaptable to many different mission scenarios or technology infusions beyond the reference design.

For the purposes of the reference architecture, Phobos was chosen as the primary target. Although the mobility concept is applicable to a wide range of solar system bodies and gravity regimes, Phobos was chosen as a body of a high scientific and human exploration interest, and of a gravity regime (around $581 \mu g$) suitable for traverses using the flywheel-based mobility approach (see Figure 36).




Figure 36: Notional illustration of the trajectory that a hybrid could execute in order to sample both the chemical and the physical diversity on Phobos (close to the Stickney crater).

4.1 Scalability

A hybrid can be readily scaled to spacecraft of various sizes, ranging from “CubeSat-class” technology demonstrations or miniaturized distributed sensor platforms (1-10 kg) to high-capability spacecraft with sizable payloads (100+ kg). For the purposes of this study, 25 kg was chosen as the target mass, as this was considered a “sweet spot” between payload capability, mission duration, and ease of accommodation on larger primary missions. Table

3 shows the reference design in the context of potential miniaturized or augmented versions of the spacecraft/rover hybrid.

	“Mini” Hybrid	“Baseline” Hybrid	“Enhanced” Hybrid
System Mass	~1-2 kg	~25 kg	~50-100 kg
Form Factor	~10x10x10 cm (1U)	~25x25x25 cm (~8U internal)	Flexible, possibly rectangular with dissimilarly sized flywheels
Avionics	CubeSat/COTS-Class highly integrated avionics	Miniaturized, radiation tolerant avionics (e.g. JPL “Sphinx” LEON3-FT-based C&DH)	
Electrical Power Generation/Storage	Primary Battery or Solar/Battery	Primary Battery	Deployable Solar Panels (stowable for mobility)
Telecom	UHF or other proximity network to relay spacecraft	UHF to relay spacecraft	S/X/Ka-Band Direct-to-Earth” with deployable, gimbaled antenna or phased array
Mission Duration	hours or days	days to weeks	weeks to months
Payload/Applications	Technology Demonstration or Miniaturized Sensor Network, small cameras	APXS, science-grade imaging/mapping	Advanced spectrometers, seismic instruments, microscopes, science-grade imaging/mapping

Table 3: Possible reference designs for a spacecraft/rover hybrid.

4.2 Reference Design Overview

The resulting architecture is an approximately 25x25x25 cm cubic spacecraft with eight spikes (one at each corner) and three flywheel/brake assemblies, analogous to the configuration of the prototypes tested as part of this study. This configuration and size category allows for the use of CubeSat-class components (albeit selected for relatively higher standards of interplanetary missions from JPL’s interplanetary CubeSat projects, e.g., INSPIRE, MarCO, and Lunar Flashlight/NEA Scout) while allowing for flexibility in overall configuration and payload. Flywheels are sized for the Phobos gravity environment and placed in a mutually orthogonal configuration. The spacecraft is powered by primary batteries and assumes the availability of a relay asset in relatively close proximity at reasonable intervals (e.g., non-Keplerian ellipses around Phobos within 300 km).

The hybrid can achieve a broad range of scientific objectives by accommodating a variety of instruments, which benefit from recent advances in sensor and electronics miniaturization and packaging. In response to the Planetary Science Decadal Survey [9], investigations encompass elemental and mineralogical composition of the surface and physical characterization of the regolith (e.g., see [10]). Investigations geared toward retiring strategic knowledge gaps for human exploration also include the quantification of surface electrostatic charging, dust dynamics, and the search for subsurface water. A representative payload includes a multispectral camera, the IntelliCam context imager developed for the Near Earth Asteroid Scout mission [49], a small alpha particle and X-ray spectrometer (APXS), such as flown on Rosettas Philae comet lander [50], and the MicrOmega microscope developed for the Phobos-Grunt mission [51]. Mass and power specifications for these instruments are summarized in

Table 4. In addition to the science payload, a complement of accelerometers, sun sensors, star-field cameras, and wide-angle context cameras aid in localization and navigation, as well as providing context for science investigations.


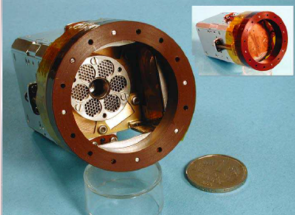

Instrument	IntelliCam	APXS	Microscope
Science Objective	Context imaging, surface navigation	Elemental composition	Regolith physical properties
Mass	500 gm	640 gm	500 gm
Power	2.5 W (peak)	1.5 W (peak)	2 W
			

Table 4: Example of representative payload for a spacecraft/rover hybrid. Integration times vary from a few hundreds of milliseconds (intelliCam and microscope) to several hours (APXS).

4.3 Subsystem Descriptions

4.3.1 Mechanical Configuration and Mass Equipment List

The spacecraft primary structure consists of a “roll cage” designed to withstand impacts from initial deployment and hopping/tumbling maneuvers. Three of the cube’s sides are housing the flywheels, while the remaining volume is roughly divided into four 10x10x20cm sections, equivalent to a “2U” CubeSat volume, for a total of $\sim 8U$. These sections are allocated as follows (see Figure 37):

- Avionics (Command & Data Handling, Interface Boards, IMU, Transceiver): 2U.
- Two Battery Packs: 2 x 2U.
- Payload Allocation: 2U.

A benefit of this configuration is that batteries and flywheels provide radiation shielding for sensitive avionics on most sides. A mass breakdown is shown in Table 5, showing healthy margins relative to the 25 kg allocation.

4.3.2 Electrical Power System

Primary batteries were chosen for the reference mission, in part due to the energy storage density, but also to make the design physically robust (avoiding fragile solar panels on a tumbling rover, or dust accumulation) as well as facilitating a concept of operations (CONOPS) that is largely independent of local lighting conditions or landing orientations. An added




Figure 37: Configuration for a 25 kg hybrid spacecraft/rover.

benefit is relatively easy, passive thermal management, allowing the external surfaces to be selected based on the desired thermal properties on Phobos. Solar panels may present a more viable option for Near-Earth Asteroids, however, due to the increased solar energy available at solar distances near 1AU, compared to up to 1.66 AU for Phobos, thus potentially allowing smaller spacecraft and longer duration missions.

The reference mission makes use of 54 Lithium-Thionyl Chloride (Li-SOCl₂) D-cell primary batteries, such as the SAFT LSH-20, in a 3s18p configuration. While the nominal capacity of the LSH-20 is 13 Ah, this was derated to 10 Ah to account for potentially off-nominal temperatures and current draws. A further 5% loss was applied for storage and depassivation losses, resulting in an overall capacity of approximately 1700 Wh. This capacity could potentially be increased by adopting Lithium Carbon Monofluoride (Li-CFx) batteries, such as the EaglePicher LCF-129 or similar cells, which have a nominal capacity of 16 Ah at a lower mass, although this chemistry may yet need to be qualified for spacecraft use.

A power budget was developed using average subsystem power estimates over the duration of high-level modes, which is summarized in Table 6. Separately from the aforementioned derating applied to the batteries, a 30% contingency was added to the power estimates, as well as a 15% electrical power system efficiency and line loss penalty.

Mech/Thermal (3120 g)	Primary Structure	2000	1	2000	30%	2600
	Thermal Hardware	150	1	150	30%	195
	Fasteners/Secondary Structure	250	1	250	30%	325
Power (6558 g)	Primary Batteries (Pack 1)	100	27	2700	3%	2781
	Primary Batteries (Pack 2)	100	27	2700	3%	2781
	Battery Packaging	250	2	500	30%	650
	Protection Circuitry	80	2	160	30%	208
	EPS Controller Card	115	1	115	20%	138
Avionics (276 g)	CDH Board	115	1	115	20%	138
	Interface Board	115	1	115	20%	138
Sensors (781.65 g)	Sun Sensors	8	6	48	5%	50
	Navigation Cameras	85	6	510	5%	536
	GNC/ACS Processing Board	115	1	115	20%	138
	IMU	55	1	55	5%	58
Mobility (4212 g)	Flywheels	1000	3	3000	30%	3900
	Motors	50	3	150	30%	195
	Motor Controllers	30	3	90	30%	117
Telecom (539 g)	UHF Transponder	250	1	250	30%	325
	Coax/RF Hardware	100	1	100	30%	130
	UHF Antenna	80	1	80	5%	84
Payload Allocation (4000 g)	Payload Allocation	3000	1	3000	33%	4000
Cabling/Harness (260 g)	Cabling/Harness	200	1	200	30%	260
Total Dry Mass and Average Contingency (kg)				16.40	17%	19.75
Mass Allocation and System Margin (kg)					21%	25.00

Table 5: Mass equipment list for a 25 kg hybrid spacecraft/rover. “CBE” stands for “Current Best Estimate,” and “MEV” stands for “Maximum Expected Value.”

Average Power Consumption (W)	Descent & Landing	Imaging/Pat hfinding	Mobility	Science Integration	Sleep/Safe Mode
Avionics	2.0	4.0	4.0	4.0	2.0
Telecom	4.0	4.0	0.5	1.5	4.0
Mobility (Actuators + Controller)	0.0	0.0	6.0	0.0	0.0
GNC (IMU, Sun Sensors, Cameras)	3.0	3.0	3.0	1.5	1.5
Payload	0.0	2.0	0.0	4.0	0.0
CBE Power Consumption by Subsystem	9.0	13.0	13.5	11.0	7.5
EPS Efficiency and Line Loss Penalty (15%)	1.4	2.0	2.0	1.7	1.1
MEV Power Consumption (30% Contingency)	13.5	19.4	20.2	16.4	11.2
Survival Heater Power (when needed)	2.8	0.0	0.0	0.0	5.0
Max Power by Mode incl. Heaters (when needed)	16.2	19.4	20.2	16.4	16.2

Table 6: Power breakdown for a spacecraft/rover hybrid.

4.3.3 Telecom Subsystem

The hybrid would use a UHF radio operating near 900 MHz to communicate to the primary spacecraft acting as a relay, similar to the Child-COM/Parent-COM (CCOM/PCOM) systems used to communicate between the Hayabusa-2 spacecraft and MASCOT [52]. Scaling the RF output power to 1W (est. 4W electrical power consumption) results in sufficient data rates to cover likely distances between a spacecraft/rover hybrid operating on Phobos and a primary spacecraft operating in the vicinity. The 900 MHz frequency also enables physically smaller, omnidirectional patch or loop antennas that can easily fit on a face of the hybrid without additional protrusions that would inhibit tumbling. In order to cover larger distances, or even Direct-to-Earth links from Near-Earth Asteroids, an S-Band or X-Band option would be feasible at the probable expense of greater power consumption for the same overall data volume. For the UHF option, while achievable data rates will depend on the

primary (relay) spacecraft distance and antenna configuration, a 32 kbps link is assumed for the reference mission.

4.3.4 Thermal Design

The reference mission makes use of a passively cooled thermal design that is actively heated when needed. As an initial estimate, the spacecraft is assumed to be isothermal, and maintained at the same temperature throughout, subject to the average power dissipations as established in the electrical power budget. The worst case hot (WCH) condition was estimated for daytime on Phobos with Mars at aphelion, and a Phobos surface temperature at $\sim 270\text{K}$, and a maximum projected area to the sun. The previous power budget established a worst case dissipated power (on average) of $\sim 20\text{W}$, which occurs during mobility or imaging/pathfinding modes. An average emissivity and absorptivity of the spacecraft was then determined that maintains the energy balance at a maximum allowable flight temperature (AFT) of $+40\text{C}$, corresponding to $\sim 75\%$ of the surface covered in standard multi-layer insulation (MLI), with the rest assumed to be covered in white paint.

Using this as a going-in assumptions for the worst case cold (WCC) condition, the required survival heater power could be determined. The worst case cold scenario assumed night on Phobos, with a surface temperature of 170K . In order to maintain the minimum allowable flight temperature of -20C , a minimum internal power dissipation of approximately 16W is necessary. Since the Sleep/Safe modes and Descent & Landing modes fall below this value, electrical heater power may be required during these modes during night time conditions.

4.3.5 Guidance, Navigation, and Control

To determine attitude and facilitate localization on a planetary surface, multiple sensors are employed by the hybrid. The reference design assumes the following:

- A small (e.g., Sensoror STIM300) IMU to resolve the gravity vector and determine angular rates and accelerations.
- CubeSat-class analog Sun Sensors with nearly spherical coverage.
- Six navigation cameras that can image constellations or surface features to facilitate localization.

Localization strategies are described in more detail in section 3.2.

4.3.6 Command & Data Handling (C&DH)

The reference mission would take advantage of JPL's interplanetary CubeSat developments, which balance the need for reliable, radiation-tolerant processing capabilities with small size

and low power demands. Based on a LEON-3FT processor, the JPL “Sphinx” board is under development for the Lunar Flashlight and NEA Scout interplanetary CubeSats, and offers sufficient capability and interfaces for the reference mission. Specifically, the Sphinx technology is based on the rad-hard Aeroflex GR712 dual-core SoC (System-on-a-Chip) device [53]. The LEON 3FT processors are capable of supporting a clock frequency of up to 100 MHz with 2×134 DMIPS (Dhrystone Million Instructions Per Second). A Microsemi ProASIC FPGA connects to the GR712 via the memory I/O bus. It provides additional computation and control logic as needed for the hybrid.

4.4 Concept of Operations

The following represents an example scenario for using the reference hybrid as defined above for a surface mission on Phobos. The primary activities of the spacecraft can be summarized as Descent & Landing, Science Integration, Imaging/Pathfinding, Mobility, and Sleep/Safe Mode. Figure 38 illustrates the sequential relationship between these modes.

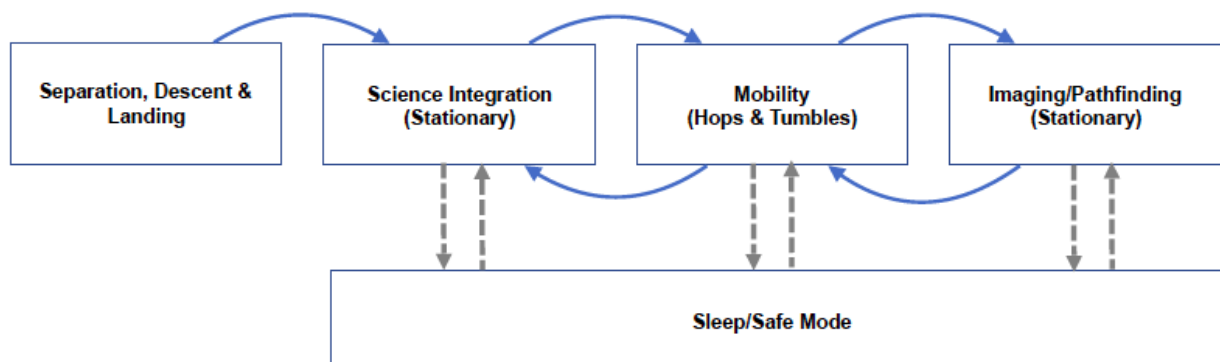


Figure 38: Five main operational modes for the operation of a spacecraft/rover hybrid.

After an initial science integration, pre-planned mobility sequences (i.e., a combination of hops and tumbles) alternate with imaging/pathfinding stops to adjust the path towards the next primary science target as needed (as illustrated in Figure 39). Once the primary science location is reached, the next science integration occurs. It should be noted that the “science” conducted by the spacecraft/rover hybrid is not limited to primary science targets or science integration – while this mode is reserved for long spectrometer integrations (and to transmit bulk data), other images and measurements can be taken at any other phase, including pathfinding waypoints and during mobility (which may offer a more distant vantage point for wider-scale imaging).

Each mode is described in more detail by the following:

- **Descent and Landing:** The spacecraft/rover hybrid is deployed from a primary spacecraft near the surface. During this phase, the spacecraft could take images and relay telemetry, but remains otherwise passive until it comes to rest on the surface. If necessary, small CubeSat-class propulsion systems could be used to provide a de-orbit ΔV or guidance towards a specified landing site. As one example, the VACCO “cold




Figure 39: Typical operations of a hybrid spacecraft/rover. Mobility sequences can be composed of long-range hops and short, precise tumbles, with periodic imaging and science integration.

gas” system developed for the MarCO CubeSats (see Figure 40) has a total impulse of 755 N-sec in a < 3.5 kg wet mass package, although small solid rocket motors or monopropellant systems could also be used.

- **Science Integration:** In regions of high scientific interest, the spacecraft/rover hybrid will remain stationary for a complete scientific investigation. This will generally be the first activity after descent & landing, in order to guarantee a measurement before conducting mobility, although self-righting and minor adjustments may occur immediately after landing. The science integration phase duration is driven by a 7.5 hour APXS integration duration. This also roughly corresponds to a day/night cycle on Phobos, allowing potential instruments to analyze the surface at a variety of lighting and thermal conditions. While stationary, the lander intermittently transmits data for a total of 250 Mbits per site visited. This results in a radio transmit duty cycle of 29% at 32 kbps (4W), with the remaining time in receive mode (0.5W), hence the average telecom power consumption of $\sim 1.5W$ during this mode.
- **Mobility:** As discussed in Section 2, mobility is defined by a series of hops (for long-range traverses) and tumbles (for more precise short-range traverses). In either case, the flywheels are spun up to a predetermined velocity slowly, and abruptly halted via brakes, thus transferring the desired momentum to the body along the resultant vector of the combined flywheel actuations. Depending on flywheel and motor configuration as well as surface properties, mobility on Phobos is estimated to be characterized by a “hopping” capability of 80m to 240m, and average traverse speeds of 14-24 cm/s (including flywheel spin-up, flight, and bouncing). During hops, the same flywheels could be used as reaction wheels to provide a stable, nadir-pointed attitude for imaging at a distance.
- **Imaging/Pathfinding:** Mobility sequences are punctuated by brief stops at pre-

determined or “ad-hoc” waypoints. This mode serves to transmit any data collected during the mobility sequences and gather local information for science or localization purposes. The spacecraft can then integrate this data to plan the next mobility sequence, either autonomously or by awaiting command from the operations team on the ground.



Figure 40: A VACCO cold gas system.

4.5 Example Timeline

Table 7 provides an example allocation of the hybrid resources and time spent in each mode to a specific scenario. Crucially, this example should be considered as a mere guideline, as the mission design is flexible – actual allocations and timing can be adjusted (e.g., trading science for mobility or vice versa) depending on science planning, risk posture, environmental factors, or observed performance. Surface operations are mostly independent from the day/night cycle, and LED lighting could be provided for imaging in the vicinity of the hybrid. Localization using surface features for reference would naturally be limited at night, but may still be facilitated using celestial imaging, inertial measurements, or ranging/imaging provided by the primary spacecraft. The reference CONOPS assumes average power figures from the power budget described above, subject to electrical heaters being used 50% of the time for modes that require them.

The resulting reference mission enables six science integrations, each at a different site, assuming the first site investigated is the initial landing site. On average, a traverse from one science integration site to another is assumed to take five mobility sequences (of one hour each), with four imaging/pathfinding waypoints in-between. A generous allocation is made for Sleep/Safe mode, which would allow for one additional hour spent at each

Duration, Energy Management, and Data Allocations by Mode	Descent & Landing	Imaging/ Pathfinding	Mobility	Science Integration	Sleep/Safe Mode (Ops. Margin)
Planned Number of Instances	1	20	25	6	20
Duration per Instance (hrs)	2	0.25	1	7.5	1
Total Time Spent in Mode (hrs)	2	5	25	45	20
Total Time Spent in Mode (% of Total Mission Duration)	2.1%	5.2%	25.8%	46.4%	20.6%
Avg. Power Consumption in Mode (W) ¹	14.8476528	19.435	20.1825	16.445	13.7264028
Energy Consumed by Instance (Wh)	29.7	4.9	20.2	123.3	13.7
Total Energy Consumed by Mode (Wh) ²	29.7	97.2	504.6	740.0	274.5
Total Energy Consumed by Mode (% of Total Capacity)²	1.7%	5.7%	29.4%	43.1%	16.0%
Transmittable Data Volume per Instance (Mbits @ 32kbps)	230.4	28.8	n/a	0	health/status
Total Data Transmitted by Mode (kbits)	230.4	576	n/a	0	health/status

¹ Assumes 50% duty cycles for survival heater use during night time, if mode requires them

² Remaining state of charge at end of mission as defined is approximately ~4.2% (relative to the de-rated capacity of non-rechargeable primary batteries)

Table 7: Example timeline for a spacecraft/rover hybrid operating on Phobos.

imaging/pathfinding waypoint. This would allow for at least one ground-in-the-loop cycle at the worst 1-way light time from Earth to Mars (between 4 and 24 minutes), or else serves as margin (or opportunity for additional science or mobility sequences) if mobility sequences are planned autonomously.

4.6 Summary

This study has determined that most subsystems could be implemented with existing technologies leveraging recent developments driven by deep space CubeSat missions. In particular key lifetime drivers, thermal and power, can be accommodated in a small form factor at levels commensurate with science-grade and data-intensive investigations. Also, the concept presented here builds on miniaturized instrument developments achieved over the past decade. Most types of instruments are becoming available in a small form factor and low power requirements, which will enable a new class of scientific investigations at relatively low resource cost. In this regard, a spacecraft/rover hybrid, and in particular the 25 kg example presented in this section, represents a realistic counterpart to science-grade CubeSat missions and could be considered as a secondary payload for future large planetary missions [1]. However, the minimalist configuration of the hybrids will limit certain types of investigations, especially those that involve material sampling and processing, deployable elements (e.g., antennas for radar), and observations that bear stringent attitude requirements. This limitation is compensated by the prospect that a large mission (e.g., Discovery class and beyond) could carry and deploy many low-mass hybrids for extended spatial coverage. Key open questions left for future work include defining interfaces with the mothership, defining scenarios that leverage networking and coordination among multiple hybrids and with their mothership, and advanced mission design work on deployment and landing strategies.

5 Conclusions

This study investigated a mission architecture that allows the systematic and affordable in-situ exploration of small solar system bodies, such as asteroids, comets, and Martian moons. The architecture relies on the novel concept of spacecraft/rover hybrids, which are surface mobility platforms capable of achieving large surface coverage (by attitude-controlled hops, akin to spacecraft flight), fine mobility (by tumbling), and coarse instrument pointing (by changing orientation relative to the ground) in the low-gravity environments (micro-g to milli-g) of small bodies. During this NIAC Phase II investigation, we raised to TRL 3.5 the mobility and autonomy subsystems (two key feasibility aspects left open by the NIAC Phase I effort) and studied at a conceptual level system engineering aspects.

Future investigations should focus on the feasibility and maturation aspects identified during Phase II, chiefly:

- **Mobility subsystem:** There are several aspects of dynamics and control that would benefit from further study. For design, a more thorough trade study of the rover geometry is required. In particular, we only considered symmetric configurations, whereby the rover can maneuver from any orientation. Second, more detailed models are required to understand the complex interactions of hopping on loose granular regolith.
- **Autonomy subsystem:** Future work should address two key aspects: (1) making the autonomy subsystem increasingly robust, for example, against modeling and localization errors and in the context of highly dynamic maneuvers (e.g., highly energetic hops), and (2) testing the overall autonomy subsystem (with planning and localization operating in closed-loop) in a high-fidelity simulation environment that captures cm-scale contact interactions as well as realistic environmental disturbances (e.g., dust).
- **Mission architecture and technology infusion:** Key open questions left for future work include defining interfaces with the mothership, defining scenarios that leverage networking and coordination among multiple hybrids and with the mothership, and advanced mission design work on deployment and landing strategies.

Ultimately, the goal is to seek opportunities for technology infusion, for example, by seeking flight opportunities in the context of small bodies missions that can accommodate a secondary payload and whose science return would be increased by an in-situ mobility platform.

6 References

- [1] B. Sherwood, K. Reh, R. Jones, Castillo-Rogez J. C., S. Spangelo, A. Frick, and A. Klesh. Planetary cubesats come of age. In *Proceedings of the 66th International Astronautical Congress*, number IAC-15,A3,5,8,x30103, 2015.
- [2] S. Chiodini, R. G. Reid, B. Hockman, I. A. D. Nesnas, S. Debei, and M. Pavone. Robust visual localization for hopping rovers on small bodies. In *Proc. IEEE Conf. on Robotics and Automation*, 2018.
- [3] B. Hockman, R. G. Reid, I. A. D. Nesnas, and M. Pavone. Gravimetric localization on the surface of small bodies. In *IEEE Aerospace Conference*, 2018.
- [4] B. Hockman and M. Pavone. Stochastic motion planning for hopping rovers on small solar system bodies. In *Int. Symp. on Robotics Research*, 2017.
- [5] B. Hockman, A. Frick, I. A. D. Nesnas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. *Journal of Field Robotics*, 34(1):5–24, 2016.
- [6] B. Hockman, R. G. Reid, I. A. D. Nesnas, and M. Pavone. Experimental methods for mobility and surface operations of microgravity robots. In *Int. Symp. on Experimental Robotics*, 2016.
- [7] B. Hockman, A. Frick, I. A. D. Nesnas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. In *Field and Service Robotics*, 2015.
- [8] R. G. Reid, L. Roveda, I. A. D. Nesnas, and M. Pavone. Contact dynamics of internally-actuated platforms for the exploration of small Solar System bodies. In *Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, 2014.
- [9] Decadal Survey Vision and Voyages for Planetary Science in the Decade 2013–2022. Technical report, National Research Council, 2011. Available at <http://solarsystem.nasa.gov/2013decadal/>.
- [10] J. C. Castillo Rogez, Marco Pavone, Issa A. D. Nesnas, and Jeffrey A. Hoffman. Expected Science Return of Spatially-Extended In-situ Exploration at Small Solar System Bodies. In *IEEE Aerospace Conference*, pages 1–15, Big Sky, MT, March 2012.
- [11] M. Wargo. Strategic knowledge gaps. 2012. Available at http://science.nasa.gov/media/medialibrary/2012/05/04/HEOMD_Strategic_Knowledge_Gaps_--_Mike_Wargo.pdf.
- [12] M. Pavone, J. C. Castillo-Rogez, J. A. Hoffman, I. A.D. Nesnas, and N. J. Strange. Spacecraft/rover hybrids for the exploration of small Solar System bodies. In *IEEE Aerospace Conference*, March 2013.

- [13] R. Allen, M. Pavone, C. McQuin, Issa Nesnas, Julie C. Castillo-Rogez, Tam-Nguyen Nguyen, and Jeffrey A. Hoffman. Internally-actuated rovers for all-access surface mobility: Theory and experimentation. In *Proc. IEEE Conf. on Robotics and Automation*, 2013.
- [14] NASA Comet Hopper mission. Technical report, NASA, 2011. Available at www.lpi.usra.edu/meetings/acm2008/pdf/8131.pdf.
- [15] R. M. Jones. The MUSES-CN rover and asteroid exploration mission. In *22nd International Symposium on Space Technology and Science*, pages 2403–2410, 2000.
- [16] P. Fiorini and J. Burdick. The development of hopping capabilities for small robots. *Autonomous Robots*, 14:239–254, 2003.
- [17] R. Z. Sagdeev and A. V Zakharov. Brief history of the Phobos mission. *Nature*, 341:581–585, 1989.
- [18] C. Dietze, S. Herrmann, F. Kuß, C. Lange, M. Scharringhausen, L. Witte, T. van Zoest, and H Yano. Landing and mobility concept for the small asteroid lander MASCOT on asteroid 1999 JU3. In *61st International Astronautical Congress*, 2010.
- [19] JAXA Hayabusa mission. Technical report, JAXA, 2011. Available at <http://hayabusa.jaxa.jp/e/index.html>.
- [20] P. Valle, L. Dungan, T. Cunningham, A. Lieberman, and D. Poncia. Active Response Gravity Offload System. 2011.
- [21] K. Nicewarner and G. Dorais. Designing and Validating an Adjustably-Autonomous Free-Flying Intraspacecraft Robot. In *Space 2006*, page 7395. 2006.
- [22] M. Bajracharya, M. W. Maimone, and D. Helmick. Autonomy for mars rovers: Past, present, and future. *IEEE Computer*, 41(12):44–50, 2008.
- [23] S. Higo, I. Nakatani, and T. Yoshimitsu. Localization over small body surface by radio ranging. In *Proc. of Space Sciences and Technology Conference*, 2005.
- [24] E. W. Y. So, T. Yoshimitsu, and T. Kubota. Relative localization of a hopping rover on an asteroid surface using optical flow. In *SICE Annual Conference*, 2008.
- [25] D. J. Scheeres. Orbit mechanics about asteroids and comets. *AIAA Journal of Guidance, Control, and Dynamics*, 35(3):987–997, 2012.
- [26] S. Tardivel, D. J. Scheeres, P. Michel, S. Van wal, and P. Sánchez. Contact motion on surface of asteroid. *AIAA Journal of Spacecraft and Rockets*, 51(6):1857–1871, 2014.
- [27] S. Van wal, S. Tardivel, and D. J. Scheeres. High-fidelity small body lander simulations. In *Int. Conf. on Astrodynamics Tools and Techniques*, 2016.
- [28] J. Bellerose and D. J. Scheeres. Dynamics and control for surface exploration of small bodies. In *AIAA/AAS Astrodynamics Specialist Conference*, 2008.

- [29] A. Klesh, J. Bellerose, and T. Kubota. Guidance and control of hoppers for small body exploration. In *Int. Astronautical Congress*, 2010.
- [30] E. Hand. Philae probe makes bumpy touchdown on a comet. *Science*, 346(6212):900–901, 2014.
- [31] I. Sharma, J. A. Burns, and C.-Y. Hui. Nutational damping times in solids of revolution. *Monthly Notices of the Royal Astronomical Society*, 359(1):79–92, 2005.
- [32] R. A. Werner and D. J. Scheeres. Exterior gravitation of a polyhedron derived and compared with harmonic and mascon gravitation representations of asteroid 4769 castalia. *Celestial Mechanics and Dynamical Astronomy*, 65(3):313–344, 1996.
- [33] N. Murdoch, I. A. Martinez, C. Sunday, E. Zenou, O. Cherrier, A. Cadu, and Y. Gourinat. An experimental study of low-velocity impacts into granular material in reduced gravity. *Monthly Notices of the Royal Astronomical Society*, 468(2):1259–1272, 2017.
- [34] R. H. Gooding. A procedure for the solution of lambert’s orbital boundary-value problem. *Celestial Mechanics and Dynamical Astronomy*, 48(2):145–165, 1990.
- [35] R. M. Woollands, A. B. Younes, and J. L. Junkins. New solutions for the perturbed lambert problem using regularization and picard iteration. *AIAA Journal of Guidance, Control, and Dynamics*, 38(9):1548–1562, 2015.
- [36] M. G. Lagoudakis and R. Parr. Least-squares policy iteration. *Journal of Machine Learning Research*, 4(Dec):1107–1149, 2003.
- [37] G. Konidaris, S. Osentoski, and P. Thomas. Value function approximation in reinforcement learning using the Fourier basis. In *Proc. AAAI Conf. on Artificial Intelligence*, 2011.
- [38] R. Mur-Artal, J. M. M. Montiel, and J. D. Tards. Orb-slam: A versatile and accurate monocular slam system. *IEEE Transactions on Robotics*, 31(5):1147–1163, Oct 2015.
- [39] T. Yoshimitsu, T. Kubota, I. Nakatani, T. Adachi, and H. Saito. Micro-hopping robot for asteroid exploration. *Acta Astronautica*, 52(2–6):441–446, 2003.
- [40] P. Fiorini, C. Cosma, and M. Confente. Localization and sensing for hopping robots. *Autonomous Robots*, 18(2):185–200, 2005.
- [41] E. W. Y. So, T. Yoshimitsu, and T. Kubota. Hopping odometry: Motion estimation with selective vision. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2009.
- [42] A. I. Mourikis and S. I Roumeliotis. A multi-state constraint Kalman filter for vision-aided inertial navigation. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 3565–3572, 2007.

- [43] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, Dec 2016.
- [44] G. Klein and D. Murray. Parallel tracking and mapping for small ar workspaces. In *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 225–234, Nov 2007.
- [45] K. Konolige, J. Bowman, J. D. Chen, P. Mihelich, M. Calonder, V. Lepetit, and P. Fua. View-based maps. *The International Journal of Robotics Research*, 2010.
- [46] D. Galvez-Lopez and J. D. Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, Oct 2012.
- [47] H. Strasdat, J. M. M. Montiel, and A. J. Davison. Scale drift-aware large scale monocular SLAM. *Robotics: Science and Systems VI*, 2010.
- [48] Z. Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, Nov 2000.
- [49] L. McNutt, L. Johnson, D. Clardy, J. Castillo Rogez, A. Frick, and L. Jones. Near-Earth Asteroid Scout. In *AIAA Space 2014 Conference*, San Diego, CA., August 2014.
- [50] G Klingelhöfer, J Brückner, C Duston, R Gellert, and R Rieder. The Rosetta alpha particle X-ray spectrometer (APXS). *Space science reviews*, 128(1-4):383–396, 2007.
- [51] C Pilorget, J-P Bibring, and M Berthe. MicrOmega: an IR hyperspectral microscope for the Phobos Grunt Lander. In *Lunar and Planetary Science Conference*, volume 42, page 1930, 2011.
- [52] C. Loisel, C. Dudal, O. Bompis, H. Guillon, F. Rousseau, G. Liabeuf, and J-L. Issler. InterSatellite Links for Rosetta/Philae, Hayabusa2/Mascot, and next-gen miniaturized ISL in S and Ka band. In *European Test and Telemetry Conference*, June 2017.
- [53] T. Imken, J. Castillo-Rogez, Y. He, J. Baker, and A. Marinan. Cubesat flight system development for enabling deep space science. In *ac*, pages 1–14, March 2017.