AAS 20-603

COMMUNICATION-AWARE ORBIT DESIGN FOR SMALL SPACECRAFT SWARMS AROUND SMALL BODIES

Federico Rossi[*], Saptarshi Bandyopadhyay[*], Mark Mote[†], Jean-Pierre de la Croix[*], Amir Rahmani[*]

Exploration of small Solar System bodies has traditionally been performed by single monolithic spacecraft carrying a number of science instruments. However, science instruments typically cannot be operated simultaneously due to the instrument requirements including optimal viewing angle, surface illumination, altitude and ground resolution, power, and data constraints. This observation has motivated interest in multi-spacecraft architectures where a swarm of small spacecraft, each carrying a single science instrument, studies a small body after being deployed by a carrier spacecraft, which then collects data from the swarm and relays it to Earth. Such architectures hold promise to yield significant improvements in mission efficiency, increases in data quality, and reduced mission duration. A key difficulty in the design of such missions is the selection of orbits for the small spacecraft, which must satisfy not only instrument requirements, but also strict inter-spacecraft communication and on-board storage constraints. To address this, in this paper, we present a novel computationally-efficient optimization algorithm for *communication-aware design* of the orbits of a small-spacecraft swarm orbiting a small body. The proposed approach captures constraints including instrument requirements, inter-spacecraft communication bandwidths, and on-board memory usage, and it can accommodate highly irregular gravity field models and surface geometries. We propose an efficient algorithm for optimization of instrument observations and inter-spacecraft communications; we then leverage the differentiable nature of the proposed algorithm to accelerate a gradient-based global search algorithm. Numerical simulations of a six-spacecraft swarm studying 433 Eros show that the proposed approach successfully identifies high-quality orbits, and significantly outperforms communication-agnostic optimization techniques, resulting in a 10% increase in scientific returns and a 30% increase in the quality of the collected data.

INTRODUCTION

In recent years, a number of missions have explored small Solar System bodies such as asteroids and comets (e.g. Rosetta,[1] Hayabusa 1&2,[2,3] and Osiris-Rex[4]). The spacecraft performing these missions carry a number of science instruments to study the small body; however, only a small number of instruments can be operated simultaneously, due to constraints including viewing angle and sun angle requirements, range and ground resolution considerations, and power and storage constraints. Moreover, the orbital parameters best suited to collect data vary widely among instruments, resulting in complex operations, the need to execute costly orbital maneuvers, and inevitable sacrifices in the quality of the data collected by some instruments.

Swarms of small spacecraft, where each spacecraft carries a single science instrument, hold promise to perform science missions in this class more efficiently compared to monolithic architec-

[*] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
[†] Georgia Institute of Technology, 301 Ferst St., Atlanta, GA 30332, USA.

tures, resulting in better data quality, shorter mission duration, and additional robustness to failures of individual vehicles.

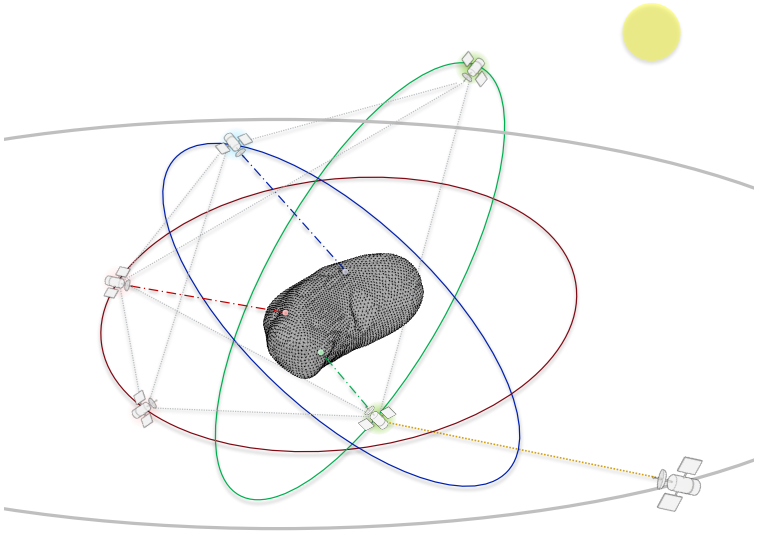

Figure 1: Notional representation of a mission configuration with 6 spacecraft orbiting around the asteroid 433 Eros. The small spacecraft conduct science observations of the asteroid, and communicate the data to the carrier spacecraft, which relays it to Earth.

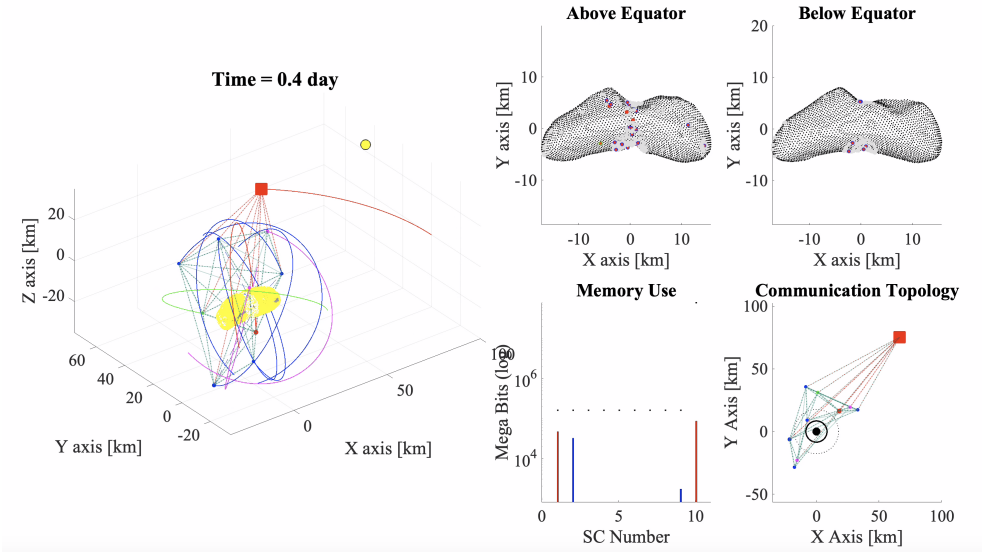

Figure 2: Output of the optimization framework. The proposed approach simultaneously optimizes observations collected by the spacecraft, inter-spacecraft communication flows, and spacecraft orbits, so as to maximize the amount of scientific data gathered and relayed to the carrier spacecraft (which, in turns, relays it to Earth).

One such mission concept is shown in Fig. 1. A carrier spacecraft, which is responsible for deploying the instrument-carrying spacecraft and relaying data from the entire swarm back to Earth,

Instrument	High Level Science Context	Measurement	Sun Angle [deg]	View Angle [deg]	Distance [km]	Surface Resolution	Data per pixel [bits]	Pixels per region	Data per region	Total Data
Imaging Spectrometer	Origin and formation of the solar system	Elemental composition of body	$-45^\circ, 0^\circ, 45^\circ$	$[-5^\circ, 5^\circ]$	$[0, 50]$	1m	360×16 ($0.4-4\mu$m range, 10nm resolution)	$1.103e5$	79.4 MB	0.81 TB
X-Ray Spectrometer	Origin and formation of the solar system	Elemental composition of body	$-45^\circ, 0^\circ, 45^\circ$	$[-5^\circ, 5^\circ]$	$[0, 50]$	10m ($1'$)	60×16 ($0.75 - 6.5$keV range,100eV resolution)	$1.103e3$	0.13 MB	1.34 GB
Camera	Chronology of the solar system, secondary processes	Photogeology (texture, color, size, shape)	$-45^\circ, 0^\circ, 45^\circ$	$[-10^\circ, 10^\circ]$	$[0, 50]$	10cm	3×16	$1.103e7$	66.2 MB	0.67 TB
Altimeter	Dynamics of small bodies	Topography	Any	$[-5^\circ, 5^\circ]$	$[0, 50]$	10cm	1×16	$1.103e7$	22.1 MB	0.22 TB

Table 1: Surface Science Instruments for a notional mission to 433 Eros. Each spacecraft carries a single science instrument. Here the Eros asteroid surface is partitioned into 10152 regions, where each point has $\approx 1.103 \times 10^5$ m^2 surface area. The total Eros asteroid surface area is 1106 km^2.

orbits near the small body. Multiple small spacecraft are deployed by the carrier in selected orbits around the small body; each spacecraft carries one science instrument, a radio allowing it to communicate with other spacecraft and with the carrier, and limited on-board memory. To collect data, each instrument-carrying spacecraft's orbit must satisfy the science instrument's viewing-angle, sun-angle, and altitude constraints (shown in Table 1 for a notional reference mission). As the spacecraft collect data, they relay it to the carrier, possibly through multiple communication hops. We assume that individual small spacecraft are unable to communicate directly with Earth, due to low-power radios and small-size antennas; therefore, all data collected by the spacecraft must be relayed to the carrier in order to be transmitted to Earth.

To realize this vision, it is necessary to devise orbit-design algorithms that can optimize the orbits of each instrument-carrying spacecraft, accounting not only for the individual instruments' requirements but also for communication constraints, i.e., the ability to return data to a carrier spacecraft in a timely manner, which is affected by inter-spacecraft bandwidths. Furthermore, it is desirable for these algorithms to be computationally efficient, so as to enable the redesign of orbits as more information is gathered about the small body's gravitational dynamics. The goal of this paper is to address this challenge by developing an efficient optimization framework for communication-aware multi-spacecraft orbit design tailored to small body dynamics.

State of the Art

The problem of orbit design for constellations of spacecraft has seen a significant amount of interest in the scientific community, and a number of techniques for constellation design *in Earth orbit* exist. We refer the reader to Ref. 5,6 for excellent overviews with historical perspective. However, constellation design techniques for Earth orbit tend to heavily exploit the regular form of the planet's gravity field; in contrast, the gravity field in close proximity of small bodies is highly irregular, making analytical and semi-analytical approaches infeasible. Small spacecraft constellation design using invariant manifolds has been proposed for lunar radio interferometer missions.[7] Numerical approaches are also used for orbit design for specific missions: for instance, particle swarm optimization is used to design the orbits of six smallsats for the Sun Radio Interferometer Space Experiment (SunRISE) mission in a GEO graveyard orbit.[8] Design of orbits around small bodies has seen significant interest in recent years, with works focusing on designing swarms of small spacecraft to estimate the gravity field around asteroid with very high accuracy,[9] and deploying the

spacecraft in passive collision-free orbits.[10] However, such approaches do not account for *communication* between the spacecraft, which introduces a critical coupling between the individual vehicles' orbits. In contrast, the approach proposed in this paper (i) allows for simultaneous optimization of *multiple* spacecraft's orbits, explicitly accounting for the coupling introduced by communications, and (ii) can accommodate arbitrary irregular gravity fields (so long as a state-transition matrix can be computed numerically).

Contribution

Our contribution is threefold. First, we propose an efficient communication-aware observation optimization framework that simultaneously optimizes the observations captured by the spacecraft and the inter-spacecraft data flows for a *given* set of orbits, maximizing the amount of scientific data relayed to the carrier spacecraft (and, in turn, to Earth). The optimization approach models Delay Tolerant Networking[11]-like inter-spacecraft communication with limited on-board memory, and it captures realistic instrument coverage, illumination, and geometry constraints. We also propose a linear relaxation of the problem, and show through numerical simulations that the relaxation yields high-quality solution close to the optimum and can be solved in seconds on commodity hardware.

Second, we turn our attention to the problem of orbit design. We leverage the differentiable nature of the linear relaxation of the observation optimizer, and the numerically-computed state transition matrix (efficiently computed through JPL's Small-Body Dynamics Toolkit (SBDT)[12]), to compute the *gradient* of the solution to the spacecraft' initial conditions with minimal computational overhead. Armed with this, we use a gradient-based global search algorithm to identify high-quality orbits for the swarm that approximately maximize the amount of scientific data returned to the carrier, while remaining collision-free over the optimization period of interest.

Third, we assess the performance of our optimization framework on a representative scenario of a multi-spacecraft, multi-instrument mission to 433 Eros; we show that the proposed approach can increase scientific returns by almost 10% and data quality by over 30% compared to a baseline communication-agnostic optimization approach.

Collectively, these results enable system designers to quickly assess the end-to-end performance of multi-spacecraft constellations, and represent a first step toward communication-aware design of multi-spacecraft missions for exploration of small Solar System bodies.

An implementation of our algorithm is available under a permissive open-source license at `https://github.com/nasa/icc`.

PROBLEM STATEMENT

Current small-body exploration missions are all build around one spacecraft, with multiple instruments gathering the scientific data over a period of time and relaying this data back to Earth through the deep space network when communication opportunities arise. In contrast, multiple small spacecraft concurrently collecting the scientific data not only have to coordinate their motion to maximize the amount of data collected, but also have to consider inter-spacecraft communication capacity and on-board memory required to off-load the collected data to the carrier spacecraft, the designated deep space network communicator with Earth. Hence, it is desirable to optimize for the largest joint coverage in the shortest time considering both instrument and communication constraints. We formally define this problem as:

Problem 1 (Communication-Aware Multi-Spacecraft Orbit Design)**.** *Efficiently design a set of orbits around a small body (with irregular gravity field) for a set of small spacecraft with hetero-*

geneous instruments, so as to maximize the amount of science delivered to the carrier spacecraft over a prescribed time horizon. The spacecraft have no propulsion other than required to enter their initial orbit; each instrument can only collect data from regions of the body where instrument constraints (i.e. sun angle, instrument range, and spacecraft view angle) are satisfied; spacecraft communicate according to a delay-tolerant networking framework and have limited on-board memory; and available communication bandwidth between the spacecraft depends on line-of-sight and distance between the spacecraft.

A MILP FORMULATION FOR SIMULTANEOUS OBSERVATION AND RELAY OPTIMIZATION

We are now in a position to describe the proposed framework to efficiently solve Problem 1. In this section, we describe a mixed-integer linear programming (MILP)-based problem formulation that optimizes the observations collected by the spacecraft and the inter-spacecraft communication activities to maximize scientific returns for a *given* set of orbits. We also propose a linear programming (LP) relaxation of the optimization problem, and show that the gradient of the LP's goal with respect to the spacecraft's initial conditions can be computed semi-analytically. In the next section, we will turn our attention to the problem of optimizing the spacecrafts' initial locations (and therefore their orbits), using the LP relaxation to efficiently assess the quality of a given set of orbits and a gradient-based optimization scheme to improve the selection.

We consider a set $\mathcal{A}$ of spacecraft orbiting a small body over a finite time period discretized by a set of time intervals $t \in \mathcal{T}$. The position of spacecraft $a \in \mathcal{A}$ along its orbit at time t is $\vec{x}_a = \vec{x}_a(\mathcal{S}_a(t_0), t)$. Each spacecraft's orbit is parametrized by its initial state $\mathcal{S}_a(t_0) = (\vec{x}_a(t_0), \vec{v}_a(t_0))$. Spacecraft are equipped with instruments in the set $\mathcal{I}$. Specifically, each spacecraft a is equipped with instrument i_a, and $\cup_{a\in\mathcal{A}} i_a = \mathcal{I}$. Multiple spacecraft may carry the same instrument; for simplicity, we consider the case where each spacecraft carries a single instrument.

It is of interest to observe a set of regions $\mathcal{R}$ of the small body. The observability function $O : (\mathcal{R}, \mathcal{I}, \vec{x}, \mathcal{T}) \mapsto [0, 1]$ captures whether it is possible to observe region $r \in \mathcal{R}$ with instrument $i \in \mathcal{I}$ from location $\vec{x}$ at time $t \in \mathcal{T}$, and the quality of the resulting observation. An observation with value 0 is infeasible or scientifically useless; an observation with value 1 is ideal. The observability function can encode a variety of geometric constraints including sun angle, spacecraft view angle, and ground resolution; the *Observability Function* Section below provides a formal description of the function and a discussion of relevant instrument constraints.

Observing region $r \in \mathcal{R}$ with instrument $i \in \mathcal{I}$ from location $\vec{x}$ at time $t \in \mathcal{T}$ yields a reward $U(r, i)O(r, i, \vec{x}, t)$ (which captures both the scientific interest of the region, with U, and the quality of the resulting data, with O) and produces an amount of data $D(r, i)$ which must be relayed to a carrier spacecraft. Spacecraft are able to communicate with each other through bandwidth-limited links. A bandwidth function $B : (\vec{x}_1, \vec{x}_2, \mathcal{T}) \mapsto \mathbb{R}$ denotes the available bandwidth between two spacecraft at locations $\vec{x}_1$ and $\vec{x}_2$ at time $t \in \mathcal{T}$. Spacecraft can also store their own data or other spacecraft's data on board and forward it at a later time. Each spacecraft has a limited amount of memory ${}_aM$. A spacecraft $c \in \mathcal{A}$ is denoted as the *carrier*. The goal of the problem is to maximize the amount of scientific data, weighted by the reward U and by the data quality O, that is collected by the spacecraft and delivered to the carrier.

Before we formally present the MILP model, we must define the observability and bandwidth functions.

Observability Function

We consider a differentiable observability function that captures the sun angle, view angle, and spacecraft-to-region range for each region, spacecraft, and instrument. Each region $r \in \mathcal{R}$ on the surface of the asteroid is characterized by a centroid $\vec{x}_r$ and a unit normal vector $\vec{n}_r$ perpendicular to the local surface of the region. We define the vector from the origin of the selected reference frame to the Sun position at time t as $\vec{x}_\odot(t)$.

We are now in a position to define the sun angle, view angle, and spacecraft-to-region range as follows, where $\|\cdot\|$ represents the ℓ^2-norm:

$$\begin{aligned} \text{Sun angle } \hat{s}_\odot(r,t) = &\quad \arccos\left(\vec{n}_r(t)\cdot\left(\vec{x}_\odot(t)-\vec{x}_r(t)\right)/\left(\|\vec{x}_\odot(t)-\vec{x}_r(t)\|\right)\right) \\ \text{View angle } \hat{s}_a(r,t) = &\quad \arccos\left(\vec{n}_r(t)\cdot\left(\vec{x}_a(t)-\vec{x}_r(t)\right)/\left(\|\vec{x}_a(t)-\vec{x}_r(t)\|\right)\right) \\ \text{Spacecraft-to-region range } d_a(r,t) = &\quad \|\vec{x}_a(t)-\vec{x}_r(t)\| \end{aligned}$$

The sun angle captures the angle between the local normal and the direction of the Sun, and therefore characterizes the illumination of the region; the view angle captures the angle at which an instrument on the spacecraft can observe the surface of the region; and the spacecraft-to-region range is simply the distance between the spacecraft and the centroid of the region.

For each instrument, the sun angle, view angle, and spacecraft-to-region range should lie within lower and upper bounds prescribed by science requirements determined on the basis of the science traceability matrix (STM). Table 1 reports a STM for a notional mission to 433 Eros, and the resulting constraints on sun angle, view angle, and spacecraft-to-region range.

In order to capture the sensitivity of observability to the spacecraft' orbits, we use a logistic function to enforce the upper and lower bounds. We denote the logistic function as $L(x) = \exp(x)/(1+\exp(x)))$. For a given upper bound $\overline{x}$, lower bound $\underline{x}$, and tolerance $\tilde{x}$, we define the *logistic window function* as

$$W(x,\overline{x},\underline{x},\tilde{x}) = L\left(\frac{x-\underline{x}}{\tilde{x}}\right)\left(1-L\left(\frac{x-\overline{x}}{\tilde{x}}\right)\right)$$

The function $W(x,\overline{x},\underline{x},\tilde{x})$ assumes values close to 1 between $\underline{x}$ and $\overline{x}$, and values close to 0 below $\underline{x}$ and above $\overline{x}$.

We denote the maximum and minimum sun angle, view angle, and spacecraft-to-region ranges for an instrument i as $\overline{s}^i_\odot, \underline{s}^i_\odot, \overline{s}^i_a, \underline{s}^i_a, \overline{d}^i_a$, and $\underline{d}^i_a$ respectively. We also denote the tolerances with respect to the sun angle, view angle, and spacecraft-to-region ranges as $\tilde{s}^i_\odot$, $\tilde{s}^i_a$, and $\tilde{d}^i_a$ respectively.

We are now in a position to define the observability function as

$$O(r,i_a,\vec{x}_a(t),t) = W\left(\hat{s}_\odot(r,t),\overline{s}^{i_a}_\odot,\underline{s}^{i_a}_\odot,\tilde{s}^{i_a}_\odot\right) W\left(\hat{s}_a(r,t),\overline{s}^{i_a}_a,\underline{s}^{i_a}_a\tilde{s}^{i_a}_a\right) W\left(d_a(r,t),\overline{d}^{i_a}_a,\underline{d}^{i_a}_a,\tilde{d}^{i_a}_a\right) \tag{1}$$

The observability function is close to one if the sun angle, view angle, and spacecraft-to-region range constraints are all satisfied; if any one of the constraints is not satisfied, the function decreases to zero.

We assume that, at each time step, the spacecraft can choose to observe any of the regions with non-zero observability, effectively abstracting away the spacecraft's attitude control and pointing problem. Future work will consider explicitly incorporating attitude and slewing constraints.

Communication bandwidth

We model the available communication bandwidth between every pair of spacecraft in absence of obstructions as a quadratic function with an upper bound. The model is justified by the assumptions that (i) the radio signal between every pair of spacecraft is only subject to free-path loss (hence, the received signal power decreases quadratically with inter-spacecraft distance); (ii) the spacecraft are equipped with omnidirectional antennas; (iii) the noise at the receiver is approximately constant; and (iv) an adaptive encoding scheme is available that changes the available data transmission rate to adapt to the SNR available at the receiver. Formally, we define the free-space bandwidth as

$$B_{\text{fs}}(a_1, a_2, t) = \min\left(\overline{B}, B_0 \cdot \left(\frac{d_0}{\|\vec{x}_{a_1}(t) - \vec{x}_{a_2}(t)\|}\right)^2\right), \tag{2}$$

where $\overline{B}$ is the maximum bandwidth between every pair of spacecraft (a function of the encoding scheme used), B_0 is a reference bandwidth, and d_0 is the distance at which the selected radios, antennas, and encoding scheme can achieve the reference bandwidth B_0.

We use a simple and conservative model to capture the effect of occlusions by the small body. For the purposes of communications, the asteroid is modeled by an outer sphere and an inner sphere, both centered at the asteroid's center of mass, characterized by an outer radius and an inner radius respectively. If the line-of-sight path between two spacecraft lies entirely outside the external sphere, the bandwidth is assumed to be the free-space bandwidth B_{fs}. If the line-of-sight path intersects the inner sphere, the bandwidth is assumed to be 0. If the free-space path intersects the outer sphere but not the inner sphere, the bandwidth is reduced proportionally to the closest distance between the free-space path and the center of the spheres.

Formally, the distance between the closest point on the line-of-sight path between spacecraft a_1 and a_2 and the small body's center of mass $\vec{x}_{cg}$ can be computed as

$$\underline{d}(a_1, a_2, t) = \min\Bigg(\|\vec{x}_{a_1}(t) - \vec{x}_{cg}\|, \|\vec{x}_{a_2}(t) - \vec{x}_{cg}\|, \sqrt{\|\vec{x}_{a_1}(t) - \vec{x}_{cg}\|^2 - \left[(\vec{x}_{a_2}(t) - \vec{x}_{a_1}(t)) \cdot (\vec{x}_{a_1}(t) - \vec{x}_{cg})\right]^2 / \|\vec{x}_{a_2}(t) - \vec{x}_{a_1}(t)\|^2} \Bigg)$$

where the first and second expression correspond to the cases where spacecraft a_1 or a_2 are the closest points to the center of mass on the line-of-sight segment, and the third expression is the distance between the line connecting the two spacecraft and the body's center of mass.

The level of communication obstruction between two spacecraft is then defined as

$$CO(a_1, a_2, t) = 1 - \min\left(1, \max\left(0, \frac{\underline{d}(a_1, a_2, t) - r_{min,co}}{r_{max,co} - r_{min,co}}\right)\right), \tag{3}$$

and the bandwidth between the two spacecraft is

$$B(a_1, a_2, t) = B_{\text{fs}}(a_1, a_2, t)(1 - CO(a_1, a_2, t)) \tag{4}$$

Remarks Some remarks on the communication model are in order. The communication model makes several strong assumptions, chiefly (i) quadratic loss model for the transmitted signal, (ii) availability of a perfect adaptive encoding scheme, (iii) availability of truly omnidirectional an-

tennas, (iv) neglecting interference between multiple spacecraft and (iv) simple (but conservative) modeling of obstructions. We argue that such a model is appropriate to capture the availability of communication bandwidth for *system-level design* of a multi-spacecraft swarm; in particular, in exchange for somewhat reduced fidelity, the proposed model provides analytical insight into how changes in each spacecraft's orbit affect available bandwidth, as discussed in the *Sensitivity* Section. More accurate communication models can then be used in *simulation* to validate the performance of the selected design and, if needed, refine the communication model. Two assumptions merit further discussion. First, the assumption that interference between the spacecraft is negligible is comparatively realistic if spacecraft use a time-division or frequency-division multiple access scheme to share the wireless medium; future work will address the inclusion of other adaptive multiple-access schemes. Second, the assumption that a perfect adaptive encoding scheme is available is a *continuous approximation* of existing adaptive encoding schemes that increase and decrease the transmitted symbol rate based on sensed signal-to-noise ratio, which are widely supported in terrestrial communication protocols such as 802.11.[13,14]

A mixed-integer linear programming formulation

We are now in a position to define the *communication-aware observation scheduling problem*.

Problem 2 (Communication-aware observation scheduling problem)**.** *Select the observations collected by the spacecraft and the inter-spacecraft data flows so as to maximize the reward for observations captured and delivered to the carrier, while satisfying bandwidth constraints on inter-spacecraft links and memory constraints on board each spacecraft.*

We pose the problem as a mixed-integer linear program. Specifically, we describe observations as binary variables that are true if and only if a given agent observes a given region at a given time, and we model data transfers between the spacecraft with a network flow formulation.[15] In the network flow formulation, the communication problem is modeled as a graph where every node represents a spacecraft at a given discrete time step; edges between pairs spacecraft represent the availability of a communication link (with an upper bound on the amount of data that can be transmitted between the spacecraft due to available bandwidth); edges between nodes representing the same spacecraft at different time steps represent the availability of on-board memory (with an upper bound capturing the overall amount of available memory); data sources correspond to the observations; and the carrier spacecraft acts as a sink for information flow. We refer the interested reader to Ref. 15 for a review of network flow formulations and their application to communication problems.

Variables We define the following variables.

- $o : (\mathcal{R}, \mathcal{A}, \mathcal{T}) \mapsto [0, 1]$ is 1 iff agent $a \in \mathcal{A}$ observes region $r \in \mathcal{R}$ at time $t \in \mathcal{T}$.
- $f : (\mathcal{A}, \mathcal{A}, \mathcal{T}) \mapsto \mathbb{R}^+$ denotes the amount of data transmitted from agent $a_1 \in \mathcal{A}$ to agent $a_2 \in \mathcal{A}$ during time interval $t \in \mathcal{T}$. If $a_1 = a_2$, the variable denotes the amount of data in agent $a_1 = a_2$'s memory during time interval t.
- $d : \mathcal{T} \mapsto \mathbb{R}^+$ denotes the amount of stored data received and stored by the carrier during timestep $t \in \mathcal{T}$.

Problem Formulation We formalize the *Communication-aware observation scheduling problem* problem as follows.

$$\max_{o,f,d} \sum_{a\in\mathcal{A}} \sum_{r\in\mathcal{R}} \sum_{t\in\mathcal{T}} o(r,a,t) \cdot U(r,i_a) \cdot O(r,i_a,\vec{x}_a(t),t) \tag{5a}$$

subject to

$$\sum_{t\in\mathcal{T}} \sum_{a\in\mathcal{A}} o(r,a,t) \leq 1 \qquad \forall r \in \mathcal{R} \text{ with dual } \mu_r \tag{5b}$$

$$\sum_{r\in\mathcal{R}} o(r,a,t) \leq 1 \qquad \forall a \in \mathcal{A}, t \in \mathcal{T} \text{ with dual } \mu_{a,t} \tag{5c}$$

$$\sum_{r\in\mathcal{R}} (o(r,a,t)D(r,a)) \; + \; \sum_{a_-\in\mathcal{A}} f(a_-,a,t \; - \; 1) \quad = \quad \sum_{a_+\in\mathcal{A}} f(a,a_+,t) \; + \; \mathbf{1}_{a=c} d(t)$$

$$\forall a \in a, t \in \mathcal{T} \text{ with dual } \lambda_{a,t} \tag{5d}$$

$$f(a_1,a_2,t) \leq B(a,b,t) \qquad \forall a_1 \in \mathcal{A}, a_2 \neq a_1 \in \mathcal{A}, t \in \mathcal{T} \text{ with dual } \mu_{a_1,a_2,t} \tag{5e}$$

$$f(a,a,t) \leq M_a \qquad \forall a \in \mathcal{A}, t \in \mathcal{T} \text{ with dual } \mu_{a,a,t} \tag{5f}$$

$$f(a_1,a_2,t_0) = 0 \qquad \forall a_1 \in \mathcal{A}, a_2 \in \mathcal{A} \tag{5g}$$

$$o(r,a,t) \in [0,1] \qquad \forall r \in \mathcal{R}, a \in \mathcal{A}, t \in \mathcal{T} \tag{5h}$$

Equation (5b) ensures that every region is observed at most once. Equation (5c) ensures that every agents observes at most one region per time step. Equation (5d) models data flows among the agents and ensures that data is only created when an observation is performed and data is only recorded at the carrier. Equation (5e) ensures that data flows on communication links do not exceed the available bandwidth, and Equation (5f) enforces the on-board memory capacity constraint. Finally, Equation (5g) sets the initial memory usage of the spacecraft and the initial data flows to zero, and Equation (5h) ensures that the observation variables assume binary values.

A Linear Relaxation

Problem (5) can be solved efficiently by commercial MILP solvers. However, the MILP formulation does not provide analytical insight into the *sensitivity* of the problem to the problem parameters (namely, the observability function and the bandwidths). In contrast, we are interested in assessing how changes to these parameters (which are, in turn, connected to changes in the selected initial conditions) can affect the solution to Problem (5).

To overcome this, we consider a *linear relaxation* of Problem (5), where Equation (5h) is replaced by its linear relaxation. The resulting problem is:

$$\max_{o,f,d} \text{(5a)} \tag{6a}$$

subject to

$$\text{(5b), (5c), (5d), (5e), (5f) and (5g)} \tag{6b}$$

$$0 \leq o(r,a,t) \leq 1 \quad \forall r \in \mathcal{R}, a \in \mathcal{A}, t \in \mathcal{T} \tag{6c}$$

The effect of the linear relaxation is twofold:

- first, at any given time step, a spacecraft may only observe a fraction of a region (notwithstanding the fact that the sum of all fractional observations of a given region should be at most one);
- second, at any given time step, a spacecraft may observe multiple regions, so long as the sum of the fractional observations is smaller than one.

The first effect is relatively benign: one can think of it as allowing the spacecraft to observe an entire region, but only storing and forwarding a fraction of it, if insufficient communication bandwidth is available. On the other hand, the second effect is highly problematic: allowing a spacecraft to observe multiple regions within the same time step would require it to slew to multiple points within a single time step, which is inconsistent with the problem formulation. To ensure that, at every time step, every spacecraft observes a single region, we *post-process* the solution of Problem (6) by setting to zero all but one of the variables o corresponding to a given agent and a given timestep, regaining feasibility but potentially losing optimality.

In the *Comparison of MILP, LP, and post-processed LP solutions* Section, we show through numerical simulations that (i) the solutions provided by the linear relaxation (6) are very close to the solutions to the MILP problem (5), and (ii) the post-processing step described above only results in a minimal loss of optimality. Accordingly, the LP relaxation appears to be a good proxy for the ILP problem. Crucially, as discussed next, the LP relaxation also provides analytical insight into the sensitivity of the problem to its parameters; we will exploit this insight to efficiently explore the space of feasible initial conditions and maximize the goal (5a).

SENSITIVITY

In this section, we show how the gradient of the objective function of the LP relaxation (6) with respect to the spacecraft initial conditions $\{\mathcal{S}_a\}_{a\in\mathcal{A}}$ can be computed semi-analytically.

The only parameters of Problem (6) that are affected by the spacecraft orbits (and therefore by their initial locations) are the observation rewards U and the bandwidths B.

Let us denote the objective function (5a) at the optimum as $K^\star$, and the corresponding optimal value of the observation variables as $o^\star$. We can express the gradient of $K^\star$ with respect to the spacecraft initial conditions $\{\mathcal{S}_a(t_0)\}_{a\in\mathcal{A}}$ as:

$$
\begin{aligned}
\frac{\partial K^\star}{\partial \mathcal{S}_a} = &\sum_{r\in\mathcal{R}}\sum_{t\in\mathcal{T}}\left(\frac{\partial K^\star}{\partial U(r,i_a,\vec{x}_a(t),t)}\cdot\frac{\partial U(r,i_a,\vec{x}_a(t),t)}{\partial \mathcal{S}_a(t_0)}\right)+\\
&\sum_{t\in\mathcal{T}}\sum_{a_-\in\mathcal{A}}\left(\frac{\partial K^\star}{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}\cdot\frac{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}{\partial \mathcal{S}_a(t_0)}\right)+\\
&\sum_{t\in\mathcal{T}}\sum_{a_+\in\mathcal{A}}\left(\frac{\partial K^\star}{\partial B(\vec{x}_a(t),\vec{x}_{a_+}(t),t)}\cdot\frac{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}{\partial \mathcal{S}_a(t_0)}\right)\\
=&\sum_{r\in\mathcal{R}}\sum_{t\in\mathcal{T}}\left(\frac{\partial K^\star}{\partial U(r,i_a,\vec{x}_a(t),t)}\cdot\frac{\partial U(r,i_a,\vec{x}_a(t),t)}{\partial \vec{x}_a(t)}\cdot\frac{\partial \vec{x}_a(t)}{\partial \mathcal{S}_a(t_0)}\right)+\\
&\sum_{t\in\mathcal{T}}\sum_{a_-\in\mathcal{A}}\left(\frac{\partial K^\star}{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}\cdot\frac{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}{\partial \vec{x}_a(t)}\cdot\frac{\vec{x}_a(t)}{\partial \mathcal{S}_a(t_0)}\right)+
\end{aligned}
$$

$$\sum_{t\in\mathcal{T}}\sum_{a_+\in\mathcal{A}}\left(\frac{\partial K^\star}{\partial B(\vec{x}_a(t),\vec{x}_{a_+}(t),t)}\cdot\frac{\partial B(\vec{x}_a(t),\vec{x}_{a_+}(t),t)}{\partial\vec{x}_a(t)}\cdot\frac{\vec{x}_a(t)}{\partial\mathcal{S}_a(t_0)}\right).$$

The derivative of the goal $K^\star$ with respect to the bandwidth $B(\vec{x}_a(t),\vec{x}_{a_+}(t),t)$ is $\mu_{a,a_+,t}$, by the definition of the dual. Likewise, the derivative of the goal $K^\star$ with respect to the bandwidth $B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)$ is $\mu_{a_-,a,t}$. The derivative of the goal with respect to the observability function $O(r,i_a,\vec{x}_a(t),t)$ is simply $o^\star(r,a,t)U(r,i_a)$, according to Equation (5a).

We recall that the sensitivity of a trajectory to its initial conditions is captured by the *state transition matrix*:[16]

$$\mathcal{S}_a(t)=\Phi_a(t,t_0)\mathcal{S}_a(t_0)$$

We denote the top three rows of the state transition matrix, representing the sensitivity of the position to the initial state, as $\Phi_a^x(t,t_0)$; then,

$$\frac{\partial\vec{x}_a(t)}{\partial\mathcal{S}_a(t_0)}=\Phi_a^x(t,t_0)$$

Accordingly,

$$\begin{aligned}\frac{\partial K^\star}{\partial\mathcal{S}_a}=&\sum_{r\in\mathcal{R}}\sum_{t\in\mathcal{T}}\left(o^\star(r,a,t)U(r,i_a)\cdot\frac{\partial U(r,i_a,\vec{x}_a(t),t)}{\partial\vec{x}_a(t)}\cdot\Phi_a^x(t,t_0)\right)+\\&\sum_{t\in\mathcal{T}}\sum_{a_-\in\mathcal{A}}\left(\mu_{a_-,a,t}\cdot\frac{\partial B(\vec{x}_{a_-}(t),\vec{x}_a(t),t)}{\partial\vec{x}_a(t)}\cdot\Phi_a^x(t,t_0)\right)+\\&\sum_{t\in\mathcal{T}}\sum_{a_+\in\mathcal{A}}\left(\mu_{a,a_+,t}\cdot\frac{\partial B(\vec{x}_a(t),\vec{x}_{a_+}(t),t)}{\partial\vec{x}_a(t)}\cdot\Phi_a^x(t,t_0)\right)\end{aligned}\tag{7}$$

The derivatives of the observability function (1) and of the bandwidths (4) can be computed analytically, as discussed next.

Gradient of the observability function

For ease of notation, let us call

$$\begin{aligned}W_{\odot}^i\left(\hat{s}_{\odot}(r,t)\right)=&W\left(\hat{s}_{\odot}(r,t),\overline{s}_{\odot}^{i_a},\underline{s}_{\odot}^{i_a},\tilde{s}_{\odot}^{i_a}\right);\\W_{\hat{a}}^i\left(\hat{s}_a(r,t)\right)=&W\left(\hat{s}_a(r,t),\overline{s}_a^{i_a},\underline{s}_a^{i_a}\tilde{s}_a^{i_a}\right);\\W_d^i\left(d_a(r,t)\right)=&W\left(d_a(r,t),\overline{d}_a^{i_a},\underline{d}_a^{i_a},\tilde{d}_a^{i_a}\right).\end{aligned}$$

The gradient of the observability function (1) with respect to the location of the corresponding agent is

$$\frac{\partial O(r,i_a,\vec{x}_a(t),t)}{\partial\vec{x}_a(t)}==W_{\odot}^i\left(\hat{s}_{\odot}(r,t)\right)\left(\frac{\partial W_{\hat{a}}^i\left(\hat{s}_a(r,t)\right)}{\partial\vec{x}_a(t)}W_d^i\left(d_a(r,t)\right)+W_{\hat{a}}^i\left(\hat{s}_a(r,t)\right)\frac{\partial W_d^i\left(d_a(r,t)\right)}{\partial\vec{x}_a(t)}\right).$$

Recall that the derivative of the logistic function is

$$\frac{\partial L((x-\underline{x})/\tilde{x})}{\partial x} = L((x-\underline{x})\tilde{x})\left(1-L((x-\underline{x})/\tilde{x})\right)/\tilde{x}.$$

The derivative of a generic logistic window function $W(f(\vec{x}_a(t)), \overline{f}, \underline{f}, \tilde{f})$ is

$$\frac{\partial W(f(\vec{x}_a(t)), \overline{f}, \underline{f}, \tilde{f})}{\partial \vec{x}_a(t)} = \left(\frac{\partial L((f-\underline{f})/\tilde{f})}{\partial f}(1-L((f-\overline{f})/\tilde{f}) - L((f-\underline{f})/\tilde{f})\frac{\partial L((x-\overline{x})/\tilde{x})}{\partial x}\right)\frac{\partial f(\vec{x}_a(t))}{\partial \vec{x}_a(t)}.$$

Finally, the gradient of the view angle with respect to the spacecraft position can be computed as

$$\frac{\partial \hat{s}_a(r,t)}{\partial \vec{x}_a(t)} = -\left(1-\left(\frac{(\vec{x}_a(t)\cdot\vec{x}_r(t))}{\|\vec{x}_a(t)\|\|\vec{x}_r(t)\|}\right)^2\right)^{-\frac{1}{2}}\left[\left(\frac{\vec{x}_a(t)}{\|\vec{x}_a(t)\|^3}\frac{1}{\|\vec{x}_r(t)\|}\left(\vec{x}_a(t)\cdot\vec{x}_r(t)\right)\right) + \frac{\vec{x}_r(t)}{\|\vec{x}_a(t)\|\|\vec{x}_r(t)\|}\right],$$

and the gradient of the spacecraft-to-region range with respect to the spacecraft position is simply

$$\frac{\partial d_a(r,t)}{\partial \vec{x}_a(t)} = \frac{\vec{x}_a(t)-\vec{x}_r(t)}{\|\vec{x}_a(t)-\vec{x}_r(t)\|}.$$

Gradient of the bandwidth function

The gradient of the bandwidth function (4) can be expressed as

$$\frac{\partial B(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)} = \frac{\partial B_{\text{fs}}(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)}(1-CO(a_1,a_2,t)) - B_{\text{fs}}(a_1,a_2,t)\frac{\partial CO(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)}.$$

The derivative of the free-space bandwidth is

$$\frac{\partial B_{\text{fs}}(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)} = \begin{cases} -2B_0\left(\frac{d_0}{\|\vec{x}_{a_2}(t)-\vec{x}_{a_1}(t)\|^2}\right)^2\cdot(\vec{x}_{a_2}(t)-\vec{x}_{a_1}(t)) & \text{if } B_{\text{fs}} < \overline{B} \\ \vec{0} & \text{if } B_{\text{fs}} = \overline{B} \end{cases},$$

$$\frac{\partial B_{\text{fs}}(a_1,a_2,t)}{\partial \vec{x}_{a_2}(t)} = -\frac{\partial B_{\text{fs}}(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)},$$

and the derivative of the communication obstruction is

$$\frac{\partial CO(a_1,a_2,t)}{\partial \vec{x}_{a_{\{1,2\}}}(t)} = \begin{cases} -\frac{1}{r_{max,co}-r_{min,co}}\frac{\partial \underline{d}(a_1,a_2,t)}{\partial \vec{x}_{a_{\{1,2\}}}(t)} & \text{if } r_{min,co} < \underline{d}(a_1,a_2,t) < r_{max,co} \\ \vec{0} & \text{if } \underline{d}(a_1,a_2,t) < r_{min,co} \text{ or } \underline{d}(a_1,a_2,t) > r_{max,co} \\ \text{undefined} & \text{if } \underline{d}(a_1,a_2,t) = r_{max,co} \text{ or } \underline{d}(a_1,a_2,t) = r_{min,co} \end{cases}$$

Finally, the derivative of the distance between the shortest-path segment from spacecraft a_1 to a_2 can be computed as:

$$\frac{\partial \underline{d}(a_1,a_2,t)}{\partial \vec{x}_{a_1}(t)} = \begin{cases} \left(\vec{x}_{a_1}(t) - \vec{x}_{cg}\right) / \|\vec{x}_{a_1}(t) - \vec{x}_{cg}\| & \text{if } \underline{d}(a_1,a_2,t) = \|\vec{x}_{a_1}(t) - \vec{x}_{cg}\| \\ \vec{0} & \text{if } \underline{d}(a_1,a_2,t) = \|\vec{x}_{a_2}(t) - \vec{x}_{cg}\| \\ 1/2\Big\{ 2\left(\vec{x}_{a_1} - \vec{x}_{cg}\right) - \Big[2\left(\vec{x}_{a_2} - \vec{x}_{a_1}\right) \cdot \left(\vec{x}_{a_1} - \vec{x}_{cg}\right) & \\ \left(\vec{x}_{a_2} - 2\vec{x}_{a_1} + \vec{x}_{cg}\right) \|\vec{x}_{a_2} - \vec{x}_{a_1}\|^2 & \\ +2\left(\left(\vec{x}_{a_2} - \vec{x}_{a_1}\right) \cdot \left(\vec{x}_{a_1} - \vec{x}_{cg}\right)\right)^2 \left(\vec{x}_{a_2} - \vec{x}_{a_1}\right)\Big] & \\ / \|\vec{x}_{a_2} - \vec{x}_{a_1}\|^4 \Big\} / \underline{d}(a_1,a_2,t) & \text{otherwise} \end{cases}$$

and

$$\frac{\partial \underline{d}(a_1,a_2,t)}{\partial \vec{x}_{a_2}(t)} = \begin{cases} \left(\vec{x}_{a_2}(t) - \vec{x}_{cg}\right) / \|\vec{x}_{a_2}(t) - \vec{x}_{cg}\| & \text{if } \underline{d}(a_1,a_2,t) = \|\vec{x}_{a_2}(t) - \vec{x}_{cg}\| \\ \vec{0} & \text{if } \underline{d}(a_1,a_2,t) = \|\vec{x}_{a_1}(t) - \vec{x}_{cg}\| \\ -1/2\Big\{ \Big[2\left(\vec{x}_{a_2} - \vec{x}_{a_1}\right) \cdot \left(\vec{x}_{a_1} - \vec{x}_{cg}\right) \left(\vec{x}_{a_1} - \vec{x}_{cg}\right) & \\ \|\vec{x}_{a_2} - \vec{x}_{a_1}\|^2 - 2\left(\left(\vec{x}_{a_2} - \vec{x}_{a_1}\right) \cdot \left(\vec{x}_{a_1} - \vec{x}_{cg}\right)\right)^2 & \\ \left(\vec{x}_{a_2} - \vec{x}_{a_1}\right)\Big] / \|\vec{x}_{a_2} - \vec{x}_{a_1}\|^4 \Big\} / \underline{d}(a_1,a_2,t) & \text{otherwise} \end{cases}$$

Remark The derivative of the bandwidth with respect to the spacecraft locations is undefined at certain locations. However, the set of points where the derivative is undefined has measure zero; accordingly, the likelihood that the undefined derivative will be encountered in execution is vanishingly low. Future work will consider using smoother transition functions in Equations (4) and (3) to side-step the issue altogether.

Putting it all together

We are now in a position to evaluate the gradient (7). The primal and dual variables $o^\star(r,a,t)$ and $\mu_{a_-,a,t}$ are computed numerically by the LP solver when solving Problem (5). The state transition matrix Φ_a^x, in turn, is computed numerically by the trajectory integrator. Finally, the derivative of the bandwidth function $B(\vec{x}_{a_1}, \vec{x}_{a_2}, t)$ and observation reward function $U(r,i,\vec{x}_a,t)$ with respect to the agent position $\vec{x}_a$ can be computed analytically, as discussed above. Hence, the gradient (7) can be computed numerically with negligible numerical overhead.

ORBIT OPTIMIZATION

Next, we turn our attention to the solution of Problem 1, i.e., selecting the set of initial conditions for the spacecraft that (approximately) maximizes the amount of science collected and delivered to the carrier. The problem is highly nonlinear; further, the irregular structure of the gravity field around small bodies makes it hard to gain analytical insight to reduce the size of the search space.

In the light of this, we turn our attention to global numerical optimization techniques; specifically, we advocate for the use of a multi-start gradient-based optimization scheme. In the proposed scheme, a number of collision-free initial conditions are sampled for the set of spacecraft. Initial conditions that result in a collision with the asteroid, or stray too far from its sphere of influence, are

rejected. For each set of collision-free initial conditions, a gradient-based trust region algorithm[17,18] attempts to locally optimize the initial conditions by following the numerical gradient (7), until it finds a local minimum. The process is summarized in Algorithm 1.

In the *Results* Section, we show that the trust region algorithm explores the neighborhood of the sampled points quite effectively, resulting in significant improvements compared to the sampled initial conditions; however, the problem presents a number of local minima, hence the multi-start approach is essential to effectively explore the space of possible orbits.

Algorithm 1 Multi-start gradient based orbit optimization algorithm

$\{\{\mathcal{S}_a^j(t_0)\}_{a\in\mathcal{A}}\}_{j\in[1,\dots,N]} \leftarrow$ Sample N sets of feasible initial conditions for all spacecraft $\mathcal{A}$
for $j \in [1,\dots,N]$ **do**
  $K^{j\star}, \{\mathcal{S}_a(t_0)^{j\star}\}_{a\in\mathcal{A}} \leftarrow$ TRUSTREGION(OPTIMIZEOBSERVATIONSANDCOMMS($\{\mathcal{S}_a^j(t_0)\}_{a\in\mathcal{A}}$))
end for
best_solution $\leftarrow \arg\max_{j\in[1,\dots,N]} K^{j\star}$
return $\{\mathcal{S}_a(t_0)^{\text{best\_solution},\star}\}_{a\in\mathcal{A}}$
function OPTIMIZEOBSERVATIONSANDCOMMS($\{\mathcal{S}_a(t_0)\}_{a\in\mathcal{A}}\}, \{U(r,i)\}_{r\in\mathcal{R},i\in\mathcal{I}}$)
  for $a \in \mathcal{A}$ **do**
    $\{(\vec{x}_a(t), \vec{v}_a(t)), \Phi_a^x(t,t_0)\}_{t\in\mathcal{T}} \leftarrow$ INTEGRATEORBIT($\mathcal{S}_a(t_0)$)
    if $\{\vec{x}_a(t)\}_{t\in\mathcal{T}}$ in collision *or* too far from small body **then**
      return
    end if
  end for
  $\{B(a_1,a_2,t), \partial B/\partial\vec{x}_{a_1}t, \partial B/\partial\vec{x}_{a_2}t\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}} \leftarrow$ COMPUTEBANDWIDTH($\{\vec{x}_a(t),\}_{t\in\mathcal{T}}\}_{a\in\mathcal{A}}$)
  $\{O(r,i_a,\vec{x}_a(t),t), \partial O/\partial\vec{x}_a(t)\}_{r\in\mathcal{R},i\in\mathcal{I},a\in\mathcal{A},t\in\mathcal{T}} \leftarrow$ COMPUTEOBSERVABILITY($\{\vec{x}_a(t),\}_{t\in\mathcal{T}}\}_{a\in\mathcal{A}}$)
  $K^\star$, $o^\star$, $\{\mu_{a_1,a_2,t}\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}} \leftarrow$ SOLVE PROBLEM 6 ($\{B(a_1,a_2,t)\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}}$, $\{O(r,i_a,\vec{x}_a(t),t)\}_{r\in\mathcal{R},i\in\mathcal{I},a\in\mathcal{A},t\in\mathcal{T}}$, $\{U(r,i)\})_{r\in\mathcal{R},i\in\mathcal{I}}$)
  $\{\partial K^\star/\partial\mathcal{S}_a\}_{a\in\mathcal{A}} \leftarrow$ COMPUTE GRADIENT 7($o^\star$, $\{\mu_{a_1,a_2,t}\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}}$, $\{\partial B(a_1,a_2,t)/\partial\vec{x}_{a_1}t\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}}$, $\{\partial B(a_1,a_2,t/\partial\vec{x}_{a_2}t\}_{a_1,a_2\in\mathcal{A},t\in\mathcal{T}}$, $\{\partial O(r,i_a,\vec{x}_a(t),t)/\partial\vec{x}_a(t)\}_{r\in\mathcal{R},i\in\mathcal{I},a\in\mathcal{A},t\in\mathcal{T}}$)
  return $K^\star, \{\partial K^\star/\partial\mathcal{S}_a\}_{a\in\mathcal{A}}$
end function

Remarks A few comments are in order. First, the proposed global search technique is *not* guaranteed to find the global optimum of the problem; this is unsurprising, given the nonconvex nature of the problem and the potential presence of a multitude of local minima. We remark that a number of more sophisticated gradient-based global search algorithms are available;[19] the application of such algorithms to effectively explore the space of feasible initial conditions is an interesting direction for future research. On the upside, the proposed technique is *anytime* and highly parallelizable. If it is desirable to obtain the best possible solution within a given allocated computation time, the trust region algorithm can be stopped at any point to yield a feasible solution. Furthermore, the sampling process in Algorithm 1 and the subsequent trust-region-based optimization can be executed in parallel for each sampled point. Therefore, the numerical performance of the the proposed approach scales linearly with the number of cores available for optimization.

Second, the proposed optimization scheme does not include constraints such as fuel consumption or ΔV budget; however, such constraints could be readily be incorporated as nonlinear constraints in the trust region algorithm.

Third, we remark that the local trust-region based optimization holds promise for *online re-optimization* of existing orbits: by only considering the gradient of the orbit with respect to the initial velocity, one can compute impulsive manoeuvers for one or multiple spacecraft to locally improve science returns. Applying the proposed approach to online re-optimization is an interesting direction for future research.

Finally, we remark that the proposed optimization scheme does *not* guarantee that the computed orbits will be stable beyond the time horizon of consideration. This is a major limitation of the proposed approach, since the optimizer may select orbits that guarantee good science returns over the time horizon of interest, but result in a collision with the asteroid or in a spacecraft leaving the asteroid's sphere of influence shortly thereafter. Analytical and numerical tools are available to assess the stability of orbits around small bodies;[16] the incorporation of such constraints in the global optimization algorithm is a critical direction for future research.

RESULTS

We assess the performance of the proposed approach on a notional mission to 433 Eros with multiple smallsats equipped with cameras, imaging and X-Ray spectrometers, and radio altimeters. The data rates and pointing requirements of the instruments are reported in Table 1.

First, we analyze the performance of the proposed approach to communication-aware observation scheduling; next, we turn our attention to the problem of communication-aware orbit optimization.

Communication-aware observation scheduling

The problem of communication-aware observation scheduling is a critical subroutine in the proposed optimization framework. To assess its performance, we analyze the suboptimality introduced by the LP relaxation (6) of Problem (5) and its feasible post-processed version, and assess the computation time of Problem (6) on commodity computing hardware.

Comparison of MILP, LP, and post-processed LP solutions Figure 3 shows the ratio between the optimal reward computed by the LP relaxation (6) and the optimal reward computed by the MILP (5) for 500 randomly-selected configurations of initial conditions.

The results in Figure 3 show that the LP relaxation is a high-quality proxy for the MILP problem: both the reward computed by the LP problem, and the post-processed feasible reward, are generally within 5% of the MILP solution, and only rare outliers fall beyond 10% of the MILP solution. These results suggest that it is appropriate to use Problem (6) as a proxy for Problem (5): the approach yields a solution with cost close to the optimum with reliably fast computational times (as discussed in the next paragraph) and, crucially, it allows to efficiently assess the sensitivity of the problem to the problem parameters.

Computation times Figure 4 shows the time required to solve Problem (6) (left) and the overall time required to compute observability and bandwidths, solve the problem, and compute the gradient (right), for 500 randomly-selected configurations of a seven-spacecraft system around 433 Eros. A 24-hour time horizon and 10-minute time step are considered. The problem is solved on a 10-core E5-2687W v3 Xeon workstation with 64 GBs of Ram. The Mosek LP solver is used.

The MATLAB implementation of the algorithm is not especially efficient; nevertheless, even a large seven-spacecraft problem can be posed and solved in an average of 10s, and the time required to actually solve the optimization problem with the Mosek solver is on average 2s. These results

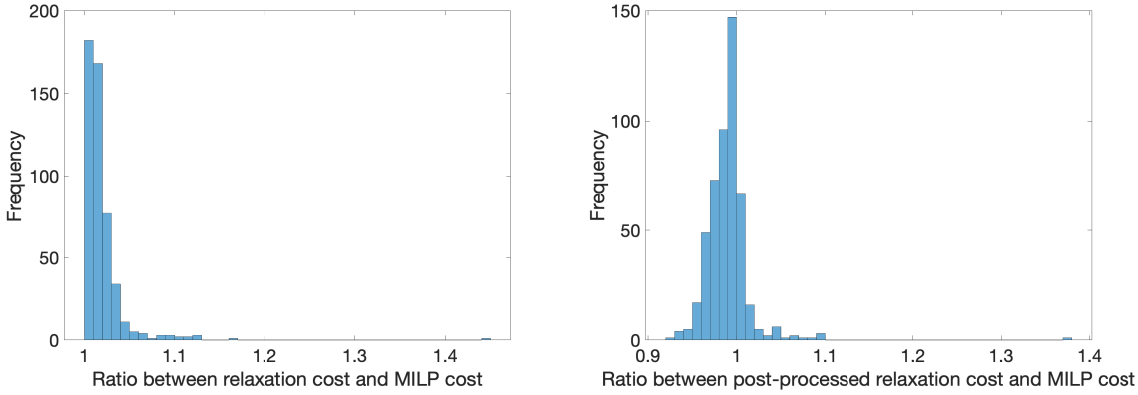

(a) Ratio between the optimal costs of Problem (6) (LP relaxation) and Problem (5) (MILP)

(b) Ratio between the optimal costs of the post-processed version of Problem (6) and Problem (5)

Figure 3: Rewards computed of LP relaxation, MILP relaxation, and post-processed feasible LP relaxation for 500 randomly-selected configurations of seven spacecraft around 433 Eros. A 24-hour time horizon and 10-minute time step are considered.

show that the proposed approach can be executed in seconds on commodity hardware, making it possible to effectively use it as the inner loop of a global search optimization scheme.

Communication-aware orbit optimization

Next, we turn our attention to the problem of communication-aware orbit optimization, and we assess the performance of Algorithm 1 compared to a communication-agnostic optimizer and its computation time.

Comparison to a communication-agnostic optimizer We compare the effectiveness of the proposed communication-aware orbit optimizer with a communication-agnostic optimizer that first approximately optimizes the orbits to maximize the amount of data collected, and then solves Problem (6) to determine the amount of data that can be effectively delivered to the carrier.

The problem of optimizing a set of orbits to maximize the amount of data collected is itself nontrivial. To approximately solve the problem, we employ a greedy heuristic where, for every instrument type, the orbits of the spacecraft carrying the instrument are optimized sequentially; each orbit approximately maximizes the amount of new data points observed, while ensuring that no data point observed by previous spacecraft is observed a second time. To achieve this, for each spacecraft a and for a given set of initial conditions for the spacecraft, we solve a simplified version of Problem (5) where the communication constraints are ignored, namely

$$\max_{o,f,d} \sum_{r\in\mathcal{R}} \sum_{t\in\mathcal{T}} o(r,a,t)U(r,i_a)\cdot O(r,i_a,\vec{x}_a(t),t) \tag{8a}$$

subject to

$$\sum_{t\in\mathcal{T}} \sum_{\hat{a}\in\mathcal{A}} o(r,\hat{a},t) \leq 1 \qquad \forall r\in\mathcal{R} \text{ with dual } \mu_r \tag{8b}$$

$$\sum_{r\in\mathcal{R}} o(r,a,t) \leq 1 \qquad \forall t\in\mathcal{T} \text{ with dual } \mu_{a,t} \tag{8c}$$

$$o(r,a,t)\in[0,1] \qquad \forall r\in\mathcal{R}, t\in\mathcal{T} \tag{8d}$$

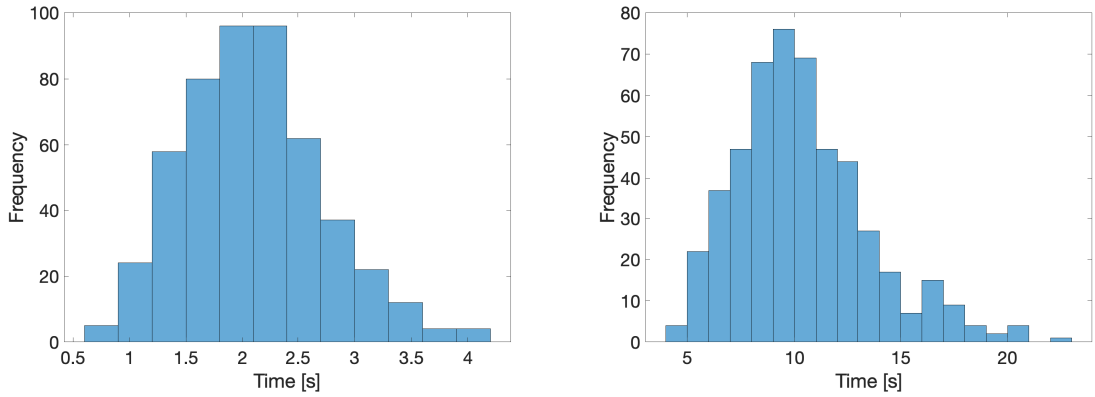

(a) Time required to compute an optimal solution to Problem (6)

(b) Time required to compute observability and bandwidths, solve Problem (6), and compute the gradient.

Figure 4: Computation time to solve Problem (6) for 500 randomly-selected configurations of seven spacecraft around 433 Eros. A 24-hour time horizon and 10-minute time step are considered.

When optimizing, *all* observations $o(r, a, t)$, including the observations of previously-optimized spacecraft, are considered in Equation (8b), ensuring that no point is observed twice. The coverage problem is submodular; therefore, if each spacecraft's orbit is selected so as to maximize the reward of Problem (8), the greedy algorithm achieves an $(1 - 1/e)$ approximation of the optimum[20] for the overall set of spacecraft.

To select (approximately) optimal orbit for an individual spacecraft, we simply sample a large number of candidate initial conditions, solve Problem (8) for each, and return the set of initial conditions that maximize the reward of (8a). The sampling approach is *not* guaranteed to be optimal, but it is a reasonable first-order approach for comparison purposes.

Comparison are shown in Table 2 and Figure 5. We consider a system with six small satellites carrying two imaging spectrometers, two cameras, an X-Ray spectrometer, and an altimeter (in addition to a carrier spacecraft in a circular 100km orbit) is considered. The reward $U(r, i)$ associated with imaging spectrometer observations, X-Ray spectrometer observations, camera observations, and altimeter readings is 3, 2, 1, and 0.5 respectively. The available bandwidth is $B_0 = 10$ kbps at a reference distance $d_0 = 100$ km.

The communication-aware approach results in a 10% higher reward compared to the greedy, communication-agnostic approach. Given the bandwidth bottleneck, the proposed approach prioritizes high-quality observations, resulting in an average 31% increase in the quality (as measured by the observability function) of the collected data; it also collects more high-reward data points from cameras and X-ray spectrometers, and almost twice as many lightweight altimeter readings. Conversely, the optimizer sacrifices the collection of large-size imaging spectrometer observations, which would be infeasible to relay to the carrier over the time horizon of interest, and elects to use the imaging spectrometer satellites primarily as relays. We remark that, if imaging spectrometer observations are deemed high-priority, their reward $U(r, i)$ could be increased to promote the collection of such observations; the key takeaway of these results is that the proposed algorithm is able to successfully optimize the user-specified reward, possibly by allocating use of the spacecraft in counterintuitive ways.

Overall, these results show that the proposed approach is able to successfully perform communication-

	$B_0 = 10$ kbps	
	Comm-aware	Comm-agnostic
Reward	679.15	618.32
Number of regions observed	288	295
by cameras	45	2
by imaging spectrometers	16	151
by X-Ray spectrometer	95	81
by altimeter	132	61
Avg. observability	0.14	0.107
by cameras	0.0083	0.0057
by imaging spectrometers	0.014	0.0107
by X-Ray spectrometer	0.34	0.34
by altimeter	0.053	0.062

Table 2: Comparison of communication-aware and communication-agnostic orbit optimization.

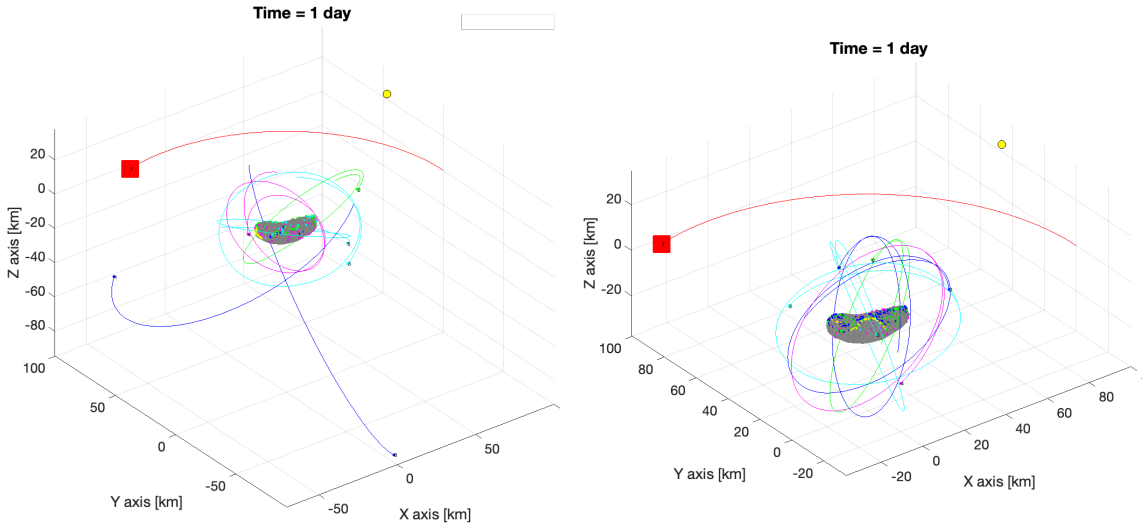

(a) Communication-aware orbits, $B_0 = 10$ kbps

(b) Communication-agnostic orbits, $B_0 = 10$ kbps

Figure 5: Comparison of optimized orbits: communication-aware vs. communication-agnostic optimizer, $B_0 = 10$ kbps.

aware optimization of observation-gathering orbits around bodies with irregular gravity fields, outperforming communication-agnostic approaches and enabling the efficient design of multi-spacecraft missions to small solar system bodies.

Effectiveness of local optimization Finally, we turn our attention to assessing the effectiveness of the gradient-based trust region optimization step. The proposed algorithm interleaves sampling and local optimization, and it is natural to wonder whether the local optimization has a significant effect on the overall solution quality. To explore this, we compute the ratio between the reward obtained with the sampled initial conditions and the locally optimal reward produced by the trust region algorithm for all sampled points. Figure 6 shows the ratio between the reward at the initial sampled point and the locally optimal reward returned by the trust region algorithm. We consider two scenarios, with reference bandwidths B_0 of 10kbps and 250kbps respectively at a reference distance $d_0 = 100$ km.

The gradient-based optimization yields significant improvements compared to the initial guess;

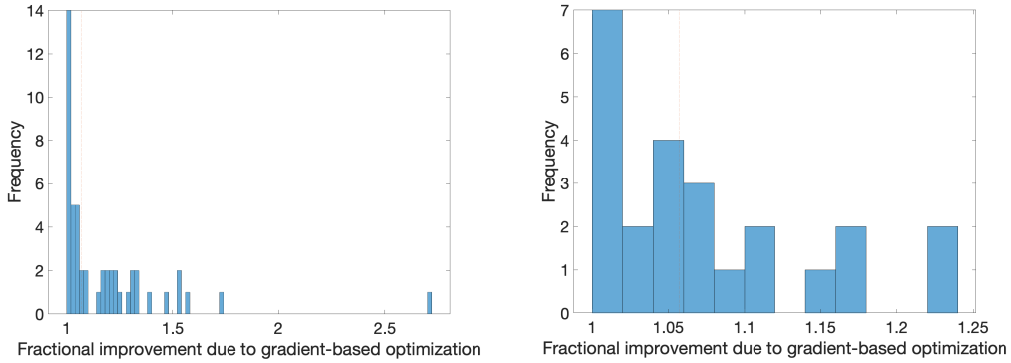

(a) Ratio between locally optimal reward and reward of initial sampled orbits. $B_0 = 10$kbps at 100km

(b) Ratio between locally optimal reward and reward of initial sampled orbits. $B_0 = 250$kbps at 100km

Figure 6: Ratio between locally optimal reward and reward of initial sampled orbits.

in particular, in the $B_0 = 10$kbps case, where bandwidth constraints are stricter, local optimization routinely yields double-digit percent improvements in the reward compared to the initial guess. This suggests that gradient-based optimization, deployed as part of a global search strategy, is an effective way of exploring the search space and identifying high-quality orbits for a multi-spacecraft mission around a small body.

CONCLUSIONS

Multi-spacecraft architectures hold promise to deliver increased science returns and shorter mission durations compared to monolithic mission architectures. Due to DSN capacity limitations and low available broadcast power, it is likely that such missions would require use of a dedicated "carrier" spacecraft to relay communications to and from Earth, and extensive inter-spacecraft communications to relay collected data from the instruments to the carrier. In this paper, we formulate the problem of communication-aware orbit optimization for such a spacecraft swarm, and we propose an efficient algorithm to solve the problem. Our simulation results demonstrate the feasibility of our proposed method, and sets the framework for future swarm orbit optimizations under instrument and communication constraints. An implementation of our algorithms is available under a permissive open-source license at `https://github.com/nasa/icc`.

A number of directions for future research are of interest. First, we plan to extend the proposed algorithmic formulation to capture long-term orbit stability constraints and ΔV budget constraints. Second, we plan to refine the proposed observability functions to capture other classes of instruments that introduce coupling between the spacecraft, e.g. radio science observations where a pair of spacecraft observes local transmissions to infer relative motion. Third, we plan to extend the problem formulation to capture data routing protocols other than Delay-Tolerant Networking, e.g. Ad-hoc On-Demand Distance Vector (AODV) routing;[21] And, finally, we plan to validate these results with a high-fidelity simulator that combines a popular orbital dynamic simulation software, GSFC's 42,[22] and a high-fidelity communication network simulator, ns-3.[23]

ACKNOWLEDGMENT

Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] K.-H. Glassmeier, H. Boehnhardt, D. Koschny, E. Kührt, and I. Richter, "The Rosetta Mission: Flying Towards the Origin of the Solar System," *Space Science Reviews*, Vol. 128, No. 1-4, 2007, pp. 1–21.

[2] A. Fujiwara, J. Kawaguchi, D. Yeomans, M. Abe, T. Mukai, T. Okada, J. Saito, H. Yano, M. Yoshikawa, D. Scheeres, *et al.*, "The rubble-pile asteroid Itokawa as observed by Hayabusa," *Science*, Vol. 312, No. 5778, 2006, pp. 1330–1334.

[3] Y. Tsuda, M. Yoshikawa, M. Abe, H. Minamino, and S. Nakazawa, "System design of the Hayabusa 2 – Asteroid sample return mission to 1999 JU3," *Acta Astronautica*, Vol. 91, 2013, pp. 356–362.

[4] M. C. Nolan, C. Magri, E. S. Howell, L. A. Benner, J. D. Giorgini, C. W. Hergenrother, R. S. Hudson, D. S. Lauretta, J.-L. Margot, S. J. Ostro, *et al.*, "Shape model and surface properties of the OSIRIS-REx target Asteroid (101955) Bennu from radar and lightcurve observations," *Icarus*, Vol. 226, No. 1, 2013, pp. 629–640.

[5] D. Mortari, M. P. Wilkins, and C. Bruccoleri, "The flower constellations," *Journal of Astronautical Sciences*, Vol. 52, No. 1, 2004, pp. 107–127.

[6] D. Mortari and M. P. Wilkins, "Flower constellation set theory. Part I: Compatibility and phasing," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 44, No. 3, 2008, pp. 953–962.

[7] S. Hernandez, J. R. Stuart, D. M. Garza, S. B. Broschart, S. J. Herzig, and S. A. Chien, "Satellite Constellation Orbit Design to Enable a Space-Based Radio Interferometer," *AAS/AIAA Astrodynamics Specialist Conference*, 2017.

[8] J. Stuart, A. Dorsey, F. Alibay, and N. Filipe, "Formation flying and position determination for a space-based interferometer in GEO graveyard orbit," *IEEE Aerospace Conference*, IEEE, 2017, pp. 1–19.

[9] N. Stacey and S. D'Amico, "Autonomous swarming for simultaneous navigation and asteroid characterization," *AAS/AIAA Astrodynamics Specialist Conference*, 2018.

[10] A. W. Koenig and S. D'Amico, "Robust and safe n-spacecraft swarming in perturbed near-circular orbits," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 8, 2018, pp. 1643–1662.

[11] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *IEEE Communications Magazine*, Vol. 41, No. 6, 2003, pp. 128–136.

[12] S. B. Broschart, M. Abrahamson, S. Bhaskaran, E. G. Fahnestock, R. Karimi, G. Lantoine, T. A. Pavlak, and L. Chappaz, "The small-body dynamics toolkit and associated close-proximity navigation analysis tools at JPL," *AAS Guidance and Control Conference*, 2015, pp. 1–12.

[13] R. Van Nee, G. Awater, M. Morikura, H. Takanashi, M. Webster, and K. W. Halford, "New high-rate wireless LAN standards," *IEEE Communications Magazine*, Vol. 37, No. 12, 1999, pp. 82–88.

[14] "IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC)and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput," *IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009)*, 2009, pp. 1–565.

[15] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.

[16] D. J. Scheeres, *Orbital motion in strongly perturbed environments: applications to asteroid, comet and planetary satellite orbiters*. Springer, 2016.

[17] R. H. Byrd, M. E. Hribar, and J. Nocedal, "An interior point algorithm for large-scale nonlinear programming," *SIAM Journal on Optimization*, Vol. 9, No. 4, 1999, pp. 877–900.

[18] R. H. Byrd, J. C. Gilbert, and J. Nocedal, "A trust region method based on interior point techniques for nonlinear programming," *Mathematical programming*, Vol. 89, No. 1, 2000, pp. 149–185.

[19] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2019, 10.1137/1.9781611975604.

[20] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—I," *Mathematical programming*, Vol. 14, No. 1, 1978, pp. 265–294.

[21] C. Perkins, E. Belding-Royer, and S. Das, "RFC3561: Ad hoc on-demand distance vector (AODV) routing," 2003.

[22] E. Stoneking, "42: An Open-Source Simulation Tool for Study and Design of Spacecraft Attitude Control Systems," `https://github.com/ericstoneking/42`, 2020.

[23] G. F. Riley and T. R. Henderson, "The ns-3 Network Simulator," *Modeling and Tools for Network Simulation* (K. Wehrle, M. Güneş, and J. Gross, eds.), pp. 15–34, Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.