Distributed Scheduling of Position Estimation Updates in Ad-Hoc Lunar Constellations

Jeremy Frank[1] and **Richard Levinson**[1] and **Eric Hillsberg**[2] and **Nicholas Cramer**[1] and **Roland Burton**[3]
[1] Intelligent Systems Division, NASA Ames Research Center
[2] Departments of Aerospace Engineering and Computer Science, U. Michigan
[3] IntelSat

Abstract

We describe the problem of scheduling position estimation updates for ad-hoc Lunar spacecraft constellations. Spacecraft communicate via directional antennas that must be simultaneously oriented towards each other in order to communicate. Each spacecraft only has two antennas, limiting communications and therefore localization performance. This problem poses challenges for existing approaches to multi-agent systems, since communications are neither pervasive nor free. In particular, spacecraft need to negotiate their 'partners' to perform position updates and localization, but must do so in the presence of these communications limitations. We show how a combination of a-priori shared knowledge and distributed coordination addresses this problem. We describe two approaches to this problem, one based on a greedy algorithm, and the other based on a matching algorithm. We present empirical results comparing the matching algorithm to a baseline in which communication is free and pervasive, and show the matching algorithm obtains acceptable localization performance.

Introduction

Existing approaches for multi-agent autonomous systems often assume the existence of pervasive, predictable, reliable, and free communications. These assumptions are the basis of the state of the art in many existing approaches to distributed planning and plan execution. If the environment is uncertain, agents can take advantage of the assumption of pervasive or predictable communications to exchange information about what they have learned and what they plan to do next. In this paper, we describe a problem that violates several of these assumptions, that of providing Lunar Position, Navigation and Timing (LPNT) services.

Upcoming decades will see a substantial increase in Lunar missions supporting and inspired by NASA's Artemis Program, including low-cost missions transported through NASA's Commercial Lunar Payload Services (CLPS) program. While Lunar missions need Position, Navigation and Timing (PNT) capabilities to ensure safe operations and meet their science objectives, the Moon does not currently have a dedicated system akin to Earth's GPS to provide localization services. The state of the art designs for Lunar positioning technology primarily utilize weak signal Global Positioning System (GPS) (Stadter et al. 2008) or high earth orbit (HEO) GPS (Ashman et al. 2018). Simulation of these techniques has shown positioning errors of 100 m ($3\,\sigma$)(Tian et al. 2014; Capuano et al. 2015; Wang 2014) which will not meet many mission localization requirements. Deep Space Network (DSN) based localization is available when missions are in view of the Earth, but low-cost surface missions may not be able to support the large power, mass, and weight requirements that these navigation solutions would entail. Additionally, DSN is heavily oversubscribed.

One alternative to provide Lunar PNT service is to create a dedicated Lunar PNT constellation, similar to GPS on Earth. However, it is not clear that there will be enough Lunar users to support the cost and resources this would require. Another alternative is to capitalize on upcoming orbital small-sat Lunar science and exploration orbiters to provide PNT services to these low-cost, surface asset Lunar missions. Examples of such small-sat science missions include Lunar Flashlight(Cohen et al. 2015), Lunar IceCube (Clark et al. 2016) and Luna H-Map (Kerner et al. 2016). Other NASA missions have begun to investigate the necessary technologies for a LPNT network, such as NASAs LunaNet (Israel et al. 2020). Creating a constellation where the PNT antenna payloads are always on (such as for LunaNet) may be too resource-intensive for the host spacecraft. Instead, these upcoming small-sat Lunar assets can be used to create an *ad-hoc, non-dedicated* PNT network capable of providing PNT services on-demand. This approach minimizes the resource usage of antennas while still providing high-quality PNT services to these low-cost, surface asset Lunar missions.

In order to minimize operating costs, this constellation should be as autonomous as possible, i.e. localization and PNT service provision should be done with as little interaction with Earth-based mission control as possible. Implementing even a partially de-centralized LPNT system requires solving a difficult multi-agent

systems problem. Communication between satellites is not necessarily pervasive; the orbits of the scientific missions permit some pairs of spacecraft to communicate directly either periodically, or not at all. Information exchange, therefore, must rely on establishing relays. Spacecraft whose primary mission is science will perform other tasks, or have constraints on how often they can perform LPNT-related duties. Orbital uncertainties lead to uncertainty in the ability to communicate at any specific time. Finally, limited time and resources require satellites to schedule communication activities to provide the best possible quality of PNT service; how to do so in a de-centralized manner, given all of the above assumptions, is a difficult challenge.

The paper is organized as follows. We first formally describe the problem of LPNT localization using ad-hoc constellations of spacecraft with limited resource availability. We describe a hypothetical design of Lunar scientific spacecraft forming these future constellations with low-energy directional antennas, and a hypothetical ad-hoc constellation. We describe the Distributed Extended Kalman Filter (DEKF) approach to state estimation, and constraints on future Lunar scientific spacecraft that lead to the need to schedule DEKF updates. Next, we describe a simple algorithm to reduce the resources needed to perform the position update step. While insufficient in its own right, this simple algorithm sets the stage for a more practical algorithm. We then describe an approach based on matching that satisfies the constraints our spacecraft's communication system imposes on communications. Both approaches require steps to distribute information throughout the constellation; we describe how this broadcast network topology can be computed on the ground and uplinked to the constellation periodically in order to drive the matching algorithm, which is performed completely onboard. We evaluate the matching approach and demonstrate that localization is acceptable when compared to the unconstrained position estimation approach. Finally, we conclude and describe future work.

Lunar Position Navigation and Timing: Challenges for Multi-Agent Systems

The Kalman Filter (KF) updates state information in the presence of uncertainty. The KF uses matrices representing estimates of satellite position and velocity, how sensor measurements reduce uncertainty in the state, how measurements correlate with each other, and linear models to simplify complex system dynamics to matrix manipulations. Kalman filters are ideal for systems which are continuously changing and have noisy sensors. They have the advantage that they are light on memory and they are very fast, making them well suited for real time problems and embedded systems.

The previously presented Lunar Autonomous PNT System (LAPS) (Hagenau et al. 2021) demonstrated the feasibility of orbital asset localization among ad-hoc Lunar small-sat constellations using a Decentralized Extended Kalman filter (DEKF), described in the following section. The DEKF uses pseudoranges[1] to, and relative velocities between, visible satellites as sensor values, or measurements. LAPS' initial feasibility demonstration assumed orbital and ground assets were continuously updating their position estimates using all available sources of information. This assumption will not hold in practice. In general, ad-hoc LPNT service cannot be provided at the expense of the nominal science missions. The DEKF update steps require measurements, which in turn require in-space links (ISLs) using the radios to perform two-way ranging operations with each other. Each such operation requires two spacecraft to communicate simultaneously. Continual updates require frequent communication between members of the constellation, which in turn, may not be feasible or allowed by the science missions. Providing service must use as few resource as possible, and may be precluded by high priority science or direct to earth transmission.

The spacecraft antenna design profoundly influences the acquisition of measurements for the DEKF update step. Omni-directional antennas are able to broadcast to all spacecraft in range; however, such antennas are power-intensive, and may not be available in the small, low-cost spacecraft comprising the ad-hoc networks. Lower-energy directional antennas are more likely on low-cost scientific space missions. However, directional antennas constrain both communications and two-way ranging operations, both due to antenna pointing and antenna slew and signal acquisition times. Furthermore, in a distributed multi-agent system, *both* satellites must schedule their actions to take place at the same time, otherwise the actions will not succeed.

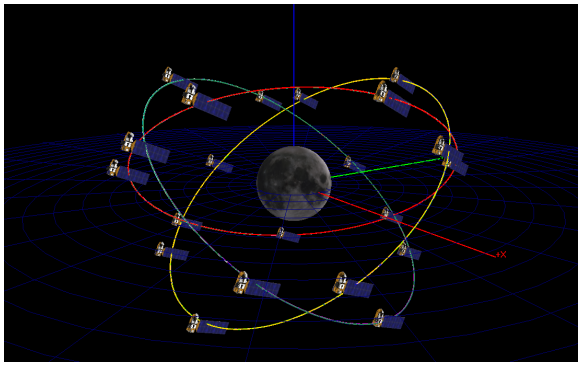

Figure 1: **High-Altitude, 21 satellite 'Frozen Orbit' Constellation.** The Frozen orbit constellation provides 24/7 global coverage but is an orbit that has never been flown at the Moon.

We evaluate our approach on a hypothetical *Frozen*[2] *orbit* constellation, which consists of 21 satellites at an altitude of 5500 km evenly spaced around 3 circular, 40°

[1] Approximation of the true range.

[2] Low or no propellant needed.

inclination orbital planes, as seen in Figure 1. Assuming all assets are always available, this constellation is designed to represent a constellation that provides global, continuous PNT coverage to all locations on the Lunar surface including the low-latitude polar regions (Pereira and Selva 2020). While frozen orbits have never been demonstrated at the Moon, this constellation provides a good test case for our approach.

The LPNT Distributed Extended Kalman Filter

We begin with describing the KF update from the perspective of a single spacecraft. Let x represent a single satellite's estimated state, and P the state covariance representing the uncertainty of the satellite's position and velocity after propagating its orbit. Φ and Q capture the state transition model, updating the state covariance. H is the measurement sensitivity matrix; it describes how the measurements (pseudoranges and velocities) inform the receiving satellites' estimate of its position and velocity. R, the measurement covariance, describes how likely combinations of measurements are. The Kalman Gain K is a function of P, H and R and represents the bias between between predicted state and the measurement. Each satellite's state x is a 6 dimensional column vector (3 position and 3 velocities) and P is the state covariance matrix,Φ is the state transition matrix and Q is the discrete time process noise covariance matrix as presented in (Hagenau et al. 2021) and all are 6×6 matrices . If a constellation has n satellites, then H is a $2(n-1)\times6$ matrix, R is a $2(n-1)\times2(n-1)$ matrix, and K is a $6\times2(n-1)$ matrix.

In the measurement update step of a KF, H and R are combined with measurements z, therefore modifying P to reflect reduced uncertainty. Each spacecraft in the constellation contributes a pseudorange and a velocity to the measurement; thus, z is an $2(n-1)$ dimensional column vector. Both H and R depend on the relative positions of all satellites, and thus must be updated after the orbits are propagated and globally shared, prior to incorporating the new measurement z. KFs come in many different varieties. The original DEKF update described in (Hagenau et al. 2021) is a *distributed* KF (each satellite performs its own updates) and *extended* (nonlinear instead of linear) KF. In this work, as in (Hagenau et al. 2021), we use the Increased Measurement Covariance DEKF of (Wen et al. 2019), whose update is of the form:

$$x_{k+1}^{-} = \int_{t_k}^{t_{k+1}} \dot{x}\, dt \tag{1}$$

$$P_{k+1}^{-} = \Phi_k {P_k}^{+} \Phi_k^T + Q \tag{2}$$

$$K = P^{-}H^T\left(HP^{-}H^T + R + \bar{H}\bar{P}^{-}\bar{H}^T\right)^{-1} \tag{3}$$

$$P^{+} = (I - KH)P^{-} \tag{4}$$

$$x^{+} = x^{-} + K(z - Hx^{-}) \tag{5}$$

Equation 3 for the Kalman gain is derived from Equation 52 in (Wen et al. 2019). (Hagenau et al. 2021) show the DEKF is able to localize effectively using all available measurements and updating every minute; since then, we have shown effective localization is possible when updating every 10 minutes.

Optimizing KF sensor updates in the presence of various constraints and costs has been studied previously. (Mourikis and Roumeliotis 2006) find the ideal update frequency for sensors in a multi-robot centralized EKF setting. (Chhetri, Morrell, and Papandreou-Suppappola 2003) solve the problem of selecting the single best sensor measurement to use in a single-agent EKF setting. (Chung et al. 2004) solve the problem of sensor selection, transmission and scheduling in a multiple sensor setting as a centralized KF. (Ny, Feron, and Dahleh 2011) consider optimal sensor selection and scheduling in an infinite-time horizon centralized KF. Our work is, to our knowledge, the first to be applied to the DEKF setting, combining inherently distributed systems, nonlinear EKFs, and measurement selection.

Greedy Scheduler

In the DEKF scheduler algorithms that follow, we will initially restrict the update step for each satellite i to use only a *single measurement* between satellite i and satellite j. We will use notation P_i instead of P^{-} to refer to i's pre-measurement covariance, and $P_{i\leftarrow j}$ instead of P^{+} to refer to i's post-measurement update with the measurement from j. Similarly, the Kalman gain K is denoted $K_{i\leftarrow j}$. Finally, the post-measurement state is $x_{i\leftarrow j}$. We want a DEKF scheduler that chooses the best two-way ranging operations to perform while respecting the directional antenna constraints of our LPNT constellations. As a stepping stone, we start by describing a *locally greedy* DEKF scheduler. A well localized single asset i will have a small state covariance $P_i = \mathrm{E}\left[(x_i - \bar{x_i})\,(x_i - \bar{x_i})^{\mathrm{T}}\right]$, where the state $x_i = \begin{bmatrix} r_i \\ v_i \end{bmatrix}$, $\bar{x_i} = \mathrm{E}[x_i]$ and E = expected value. A measure of the size of the state covariance (the magnitude of uncertainty) is the Frobenius norm of P_i, $\|P_i\|_{\mathrm{F}}$, which reduces a 6×6 positive semi-definite matrix to a scalar. A well localized constellation of assets will have small state covariances across all assets. An equivalent constellation-wide measure is the Euclidean norm of the sum of the asset Frobenius norms:

$$\sqrt{\sum_{i=1}^{n} \|P_i\|_{\mathrm{F}}^2} \tag{6}$$

The square of the Euclidean norm of the asset Frobenius norms is equivalent to the Frobenius norm of the full (decentralized) constellation state covariance matrix, i.e.

$$\sum_{i=1}^{n} \|P_i\|_{\mathrm{F}}^2 = \|P\|_{\mathrm{F}}^2 , P = \begin{bmatrix} P_1 & 0 & \cdots & 0 \\ 0 & P_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & P_n \end{bmatrix} \quad (7)$$

All measurement sensitivity matrices H have one row h_j for each available measurement j This row contains the unit vector of the measurement direction. The *locally greedy DEKF scheduler* assumes that at each time step the satellite can only utilize one measurement. The locally greedy DEKF scheduler chooses, for each satellite i, this one measurement in a way that attempts to minimize (in a Frobenius norm sense) the covariance of the satellite's state estimate. The measurement update for each satellite i is now:

$$K_{i \leftarrow j} = P_i h_j^T \left(h_j P_i h_j^T + r_{jj} + \bar{h}_j \bar{P}_i \bar{h}_j^T \right)^{-1} \quad (8)$$

$$P_{i \leftarrow j} = (I - K_{i \leftarrow j} h_j) P_i \quad (9)$$

$$x_{i \leftarrow z} = x_i + K_{i \leftarrow j} (z_i - h_j x_i) \quad (10)$$

Vectors h_j and $\bar{h}_j$ denote the single row of H and $\bar{H}$ (i.e. the single measurement) whose index is j, and r_{jj} is the (scalar) variance of that measurement. Let visibility graph $G = V, E$, with $e_{ij} \in E$ if satellite i and j are mutually visible and within radio range. h_j and $\bar{h}_j$ are determined using an exhaustive search through all n rows of measurement matrix H, solving

$$j = \mathrm{argmin}_{k=1 \ldots n} \|P_{i \leftarrow k}\|_{\mathrm{F}} \quad (11)$$

by first computing $P_{i \leftarrow k}$ for all $k \neq i$ such that $e_{ij} \in E$ using Equations 8 and 9. and then using $h_j, \bar{h}_j$ and $r_{j,j}$ in Equations 8 and 9 for the DEKF measurement update step. In selecting the measurement, the measurement itself (i.e. the pseudorange - velocity pair) is not used. Only the geometry (direction) of the measurement is required, and this only needs to be known approximately. Indeed, Equations 8 and 9 only use information based on the state prior to update.

As noted above, each 2-way ISL provides 2 measurements, pseudorange and velocity. To account for both measurements from a single ISL, we assume the two measurements from satellite j appear in consecutive rows of H and vector z, and interpret Equations 8, 9 and 11 as finding the 2×6 matrix $h_{j'}$ and 2×2 covariance sub-matrix $r_{jj'}$, and perform the same minimization; we then keep Equations 8, 9 and 11 unchanged (using only subscript j) for simplicity.

In order to both communicate and perform the DEKF update, two spacecraft must orient their antennas to point at each other, and simultaneously perform the two-way ranging operation. Such paired actions hearken back to required concurrency in temporal planning (Cushing et al. 2007). Since each satellite computes its own preferred partners, and since each spacecraft needs to orient its antenna at the same time in order to communicate, the spacecraft need to distribute their preferred partners globally across the constellation prior to the 2-way ISLs, and thus the DEKF update, meaning satellites must 'communicate in order to communicate'. We describe how to solve this seeming chicken and egg problem later in the paper.

The greedy measurement scheduler is optimal if each spacecraft can only use one measurement in the DEKF update step. To see why, see that each $\|P_{i \leftarrow j}\|_{\mathrm{F}}^2$ term in Equation 6 is minimized by the greedy approach.

The greedy scheduler solution has a significant drawback; some satellites may need to perform multiple two-way ranging operations, one to obtain its ideal measurement, and additional operations to provide the ideal measurement to another visible spacecraft. To see why, suppose satellite i has $\delta(i)$ neighbors in G. In the worst case, i is the preferred partner for all of its neighbors, and a communications schedule must be built after globally distributing the greedily selected preferred partners. Scheduling the 2-way ISLs and updates can (almost; see future work) be posed and solved as *Edge-Colorability* of the graph G_g, whose edges are the desired greedy matches, where $G_g \subset G$. The edge coloring constraint ensures no node has two edges of the same color incident on it. Each color of this graph corresponds to a 'time slot', meaning all edges of the same color correspond to 2-way ISLs performed at that time. Vizing's theorem shows that the number of colors needed to edge-color any graph G is either Δ or $\Delta + 1$ where Δ is the maximum degree of G. We know Δ, the maximum degree of G, bounds above Δ_g, the maximum degree of G_g. Thus, in the worst case, performing the greedy update may require multiple communication acts involving reorientation of the spacecraft antenna to perform all the needed 2-way ISLs. There may be insufficient time to perform all necessary antenna reorientations (called slews), and even if there were, it is very time and energy intensive to do so, and may not be allowed by the science missions.

There are other problems with the greedy scheduler. It is inefficient to perform multiple 2-way ISLs, only to discard all but the best. Furthermore, *both* spacecraft performing a 2-way ISL can update their covariances; the greedy criteria does not account for this in measuring the improvement of the state per Equation 6.

A Matching-Based Scheduler

The greedy scheduler algorithm is a useful way of introducing key concepts needed for a practical scheduler, but is inadequate on its own for the reasons described above. We now devise an algorithm that ensures each spacecraft performs only *a single 2-way ranging operation* per DEKF update cycle. The key features of the greedy scheduler we will carry forward to the next formulation are the use of the Frobenius norm as the measure of quality of a proposed update, and a global broadcast in preparation for performing two-way ISLs.

Our spacecraft have two antennas, so they could perform 2 ranging operations, which we will address later.

In order to limit all satellites to only a single two-way ranging operating per DEKF cycle, and also account for reductions in both satellites' covariances during the update step, we must find the best *set of pairs* of DEKF measurement updates to perform, according to the definition of the desirable constellation state in Equation 6. If i and j perform a two-way ISL and update, their contribution to the sum in Equation 6 is $f_{ij} \equiv \|P_{i\leftarrow j}\|_{\mathrm{F}}^2$. We desire this contribution to be as small as possible. The possible two-way ranging operations are constrained by visibility graph $G = V, E$, as with the greedy algorithm. If we annotate each edge $e_{ij} \in E$ with the sum $f_{ij} + f_{ji}$, we have defined the well-studied *Minimum Matching Problem*, which is solved by well-known algorithms. We use a $0-1$ decision variable y_{ij} to indicate i is matched with j due to the existience of edge e_{ij}. In order to account for both odd numbers of nodes and the possibility of visibility graphs with nodes of degree 1, self-matches are allowed; we let $f_{ii} = \|P_i\|_{\mathrm{F}}^2$. By definition, the post-measurement update can't increase the covariance, hence $f_{ii} \geq f_{ij}$ for $j \neq i$. Thus, in this matching problem, a pair of self-matches will contribute $f_{ii} + f_{jj}$, but a match of i with j will contribute some sum $f_{ij} + f_{ji} \leq f_{ii} + f_{jj}$, which by definition is more desirable. The matching problem is then defined as follows:

$$\min \sum_{i \neq j} (f_{ij} + f_{ji}) y_{ij} + f_{ii} y_{ii} \tag{12}$$

$$\text{s.t.} \sum_{j} y_{ij} = 1 \qquad \forall i \tag{13}$$

The constraints ensure each node is matched to exactly one other node, perhaps itself. If the optimal value of Equation 12 is $\hat{f}$, then the constraint ensures there is some j matched with each i and thus $\sqrt{\hat{f}} = \sqrt{\sum_{ij} \|P_{i\leftarrow j}\|_{\mathrm{F}}^2}$. Thus, the solution to the matching problem minimizes the desired objective (Equation 6) for the constellation.

As with the edge-coloring formulation, all satellites must pose and solve the *same* matching problem simultaneously, in order to determine the global, constellation-wide two-way ranging topology, and subsequently execute the same 2-way ranging operations. This means all satellites must know the visibility graph G for the time at which the matches need to hold. The visibility graphs can be computed on the ground hours or days in advance and uplinked to the spacecraft. However, all satellites must also exchange their f_{ij} values globally to build and solve the matching problem. The f_{ij}s cannot be computed and uplinked from the ground, because the covariances change after each DEKF update. Thus, in a similar manner to the greedy DEKF scheduler, satellites need to broadcast the f_{ij}'s globally in order for each satellite to pose and solve the identical matching problem. Since each satellite independently determines the same matching solution, they do not need to communicate the solution.

Each spacecraft has two antenna, so we can generalize this approach and construct a 'multi-match', in which each spacecraft is matched with two other spacecraft to perform 2-way ISLs. We generalize the notation, and let $P_{i\leftarrow j,k}$ refer to the post-measurement covariance of i with measurements j, k. We modify the formulation of this problem as follows: first, in addition to the single best ISL, we must also evaluate *pairs* (j, k) of ISLs:

$$K_{i\leftarrow j,k} = P_i h_{j,k}^T \left(h_{j,k} P_i h_{j,k}^T + r_{jk} + \bar{h}_{j,k} \bar{P}_i \bar{h}_{j,k}^T \right)^{-1} \tag{14}$$

$$P_{i\leftarrow j,k} = (I - K_{i\leftarrow j,k} h_{j,k}) P_i \tag{15}$$

$$f_{ijk} \equiv \|P_{i\leftarrow j,k}\|_{\mathrm{F}}^2 \tag{16}$$

We now pose and solve the '2-matching' as follows:

$$\min \sum_{ijk} f_{ijk} y_{ijk} + \sum_{ij} f_{ij} y_{ij} \tag{17}$$

$$\text{s.t.} \sum_{j,k} y_{ijk} + \sum_{j} y_{ij} = 1 \qquad \forall i \tag{18}$$

$$\sum_{k} y_{ijk} + y_{ij} = \sum_{k} y_{jik} + y_{ji} \qquad \forall i, j \tag{19}$$

Constraints 18 mimic the single-matching problem; Constraints 19 are symmetry constraints; if satellite i is matched to j, this constraint ensures j is matched to i, but allow both i and j to be matched with a second satellite k.

Ad-Hoc Broadcast Network Construction

As noted above, each satellite's DEKF scheduler needs all the f_{ij}'s to pose and solve the same matching problem. This means that in addition to the visibility graphs, an *ad-hoc broadcast network* must be centrally computed (on the ground) and provided to each satellite. A spacecraft with two antennas can communicate with two of its visible neighbors. If it is possible to construct a *Hamiltonian Path* G_H in the visibility graph G, once established, information can be communicated to all satellites via this network *without* expensive antenna reorientation operations.

We briefly describe how the complete DEKF update cycle is performed, using the single-match notation for simplicity. Visibility graphs G and broadcast networks G_H are assumed to be shared among all satellites i. The communication times are based on notional future radio performance; antenna reorientation time is based on analysis of CubeSats that demonstrates 3° per second (on average) body-mounted slew rate (Sin et al. 2021). First, P_i, x_i as computed from the last DEKF update must be globally broadcast over network G_H. We assume this takes 3 minutes; 1 minute to slew, 1 minute to join network, and 1 minute to transmit and

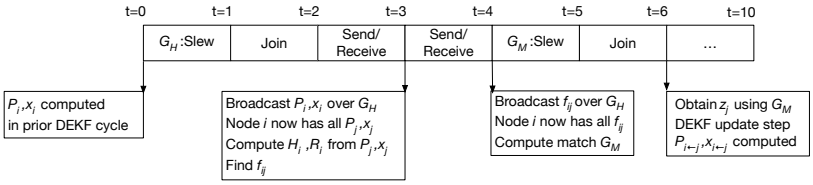

Figure 2: **DEKF Update Cycle** includes network setup, communication, and computations.

receive. Once this is complete, each satellite has P_j, x_j from all other satellites. Now satellite i can construct R_i, H_i from P_j, r_j, and then find all f_{ij}'s. Next, satellite i sends the f_{ij}'s over the same broadcast network. This takes 1 minute to transmit and receive. Now each satellite can pose and solve the matching problem using the f_{ij}'s (Equation 12). The solution forms the 2-way ISL network, which we call G_M, which takes 2 minutes to establish; 1 minute to slew and 1 minute to join network. This topology is only needed to generate z_j by commanding the radios, after which the DEKF update step (Equations 8,9) can be performed. This cycle is described in Figure 2.

The ad-hoc broadcast network G_H must persist long enough to enable global communication of the f_{ij}'s, which are needed to formulate the matching problem, and also to share state needed to update P, H and R to formulate the f_{ij}'s. These two communication acts are scheduled to occur one right after the other. Each P is a 6×6 matrix of doubles, and each x is a 6×1 vector of doubles, for a total of 2588 bits. If each f_{ij} is a double this is $64 \times 97 = 6208$ bits. Thus, the bandwidth requirements to update the matching exceed that required to propagate the DEKF. Our analysis of the notional 10 minute DEKF update cycle indicates broadcast networks must hold for 3 minutes to transmit both of these messages, not including the 1 minute antenna slew time, but that the expected bandwidth is sufficient to perform these transmissions.

For the Frozen orbit constellation of 21 satellites, we solved the general Hamiltonian Path problem to select an arbitrary G_H. We then checked to see if the resulting G_H held at the next time step; if it didn't, we would look for another G_H. Somewhat surprisingly, a *single* G_H holds for *all* G.

Comparing Approaches

We informally compare the locally greedy and matching formulations using different metrics: optimality, number of communication acts as proxies for expensive antenna slew time, and computation time.

The locally greedy approach is optimal with respect to Equation 6 if the DEKF update step constrains a satellite to *use only a single measurement*. The single matching algorithm is optimal if satellites are constrained to *perform at most one 2-way ISL*, and the 2-matching is optimal if satellites are constrained to *perform at most two 2-way ISLs*.

Both approaches broadcast information to set up a scheduling problem. Network traffic for the greedy algorithm is 1 double per satellite, while that for the matching algorithm is n doubles per satellite. However, the matching approach requires only one additional communication topology to perform the two-way ranging operations. The greedy DEKF update requires each spacecraft i to communicate with at worst all $\delta(i)$ of its neighbors, so the communication topology needs to be updated at most Δ times as a result.

The computation costs of both approaches are similar. To drive both algorithms, each satellite i needs to compute $\|P_{i\leftarrow j}\|_{\mathrm{F}}^2$ for all visible satellites j. Computing $\|P_{i\leftarrow g,h}\|_{\mathrm{F}}$ increases this time by a factor of n worst-case compared to greedy. For the greedy DEKF, polynomial-time approximations exist to edge color a graph G with $\Delta + 1$ colors, and the graphs are likely simple enough that such algorithms suffice. The visibility graph is not a bipartite graph, meaning that the matching problem is best solved by the polynomial time Edmond's Blossom algorithm, which is a modest computational cost (Edmonds 1965) [3]

Empirical Results

Table 1: **DEKF Performance Comparisons**

Fully-Sampled	One-Match	Two-Match
79.7m	10521.6m	127.8m

Based on our informal evaluation above, we choose to empirically evaluate the matching algorithm performance. Table 1 compares the DEKF of (Hagenau et al. 2021), which we refer to as 'fully-sampled', with the two matching-based scheduler approaches described above. The 21 satellite Frozen constellation was simulated for 7 days. A high-fidelity simulation provides ground truth, and error for each satellite is determined by comparing its position as determined by the DEKF with ground truth. The fully-sampled DEKF updates every 10 minutes but uses every measurement; the matching approaches are limited to either one or two ISLs, meeting the notional 2-antenna spacecraft design constraint.

Table 1 shows median error over all satellites, over all time. The performance of single matching is quite

[3]We do not report solver time results in our paper, only measurement performance. We use an MILP formulation for both matching algorithms.

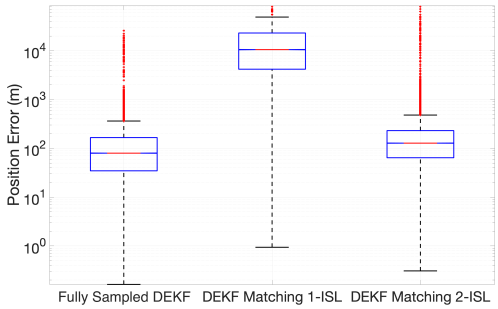

Figure 3: **Comparison between matching ISL performance.**

poor, but multi-matching leads to only 1.6× increase in median localization error compared to the unrestricted DEKF. Figure 3 shows more details; 1−matching simply is not able to achieve good localization performance, while two matching is comparable to the unrestricted DEKF. Figure 4 shows the localization error of all 21 satellites over the 7 day horizon. The fully-sampled DEKF, and by extension the 2−matching show good performance. The 1−matching is dominated by the systems geometry with only a few nodes dominating the localization.

Conclusions and Future Work

We summarize here some of the features of our approach. The complete solution to the LPNT problem still requires a centralized authority to provide the visibility graph G, and the pre-computed broadcast networks G_H, to each satellite. Thus, our solution does not address completely the problem of an unpredictable, non-pervasive communication foundation. However, our approach demonstrates the exchange of dynamically changing information (the f_{ij}'s) over these broadcast networks to perform the best possible two-way ranging operations at a fixed cadence. Once information is exchanged, rescheduling is completely distributed. The approach ensures all agents independently find the same schedule, and thus pairs of agents can reliably perform 2-way ISLs, demonstrating agents' ability to coordinate, even when communication is constrained. The comparable localization performance of 2−matching shows good localization performance can be achieved with significantly less communication than the approach that uses all measurements.

Our approach can be thought of as a classic 'replanning' approach in the presence of new information. One might ask if such a problem could be posed and solved as a DEC-POMDP (Bernstein et al. 2002). Because enumeration of the state space is possible on the ground, one could potentially plan for multiple DEKF cycles, rather than just a single cycle, as we propose. However, the state space is very large and complex to enumerate. Many such states are those in which no communication is possible; detecting and pruning these can reduce the state space but increases complexity. Finally, the policy could be very large, and potentially impractical to uplink, store and execute.

It is important to analyze other constellation designs. Variations in orbit architecture (number of planes, and spacing), orbit altitudes, and numbers of assets will influence visibility, existence of ad-hoc broadcast networks, and quality of the matching algorithms. Imposing constraints on the non-dedicated assets (e.g. spacecraft is performing science) also changes the problem.

The 3-uniform hyper-edge coloring problem (Obszarski and Jastrzebski 2017) accounts for the fact that spacecraft have two antennas when scheduling greedy algorithm measurements, suggesting an empirical evaluation of the greedy approach and comparison to the 2-matching approach.

Generalizing to other multi-agent localization problems will change the visibility graphs, existence of ad-hoc broadcast networks, and quality of the matching algorithms. Moving away from directional links to omnidirectional antennas may be feasible for, say, a terrestrial network of UAVs or robots, and change the flavor of the problem.

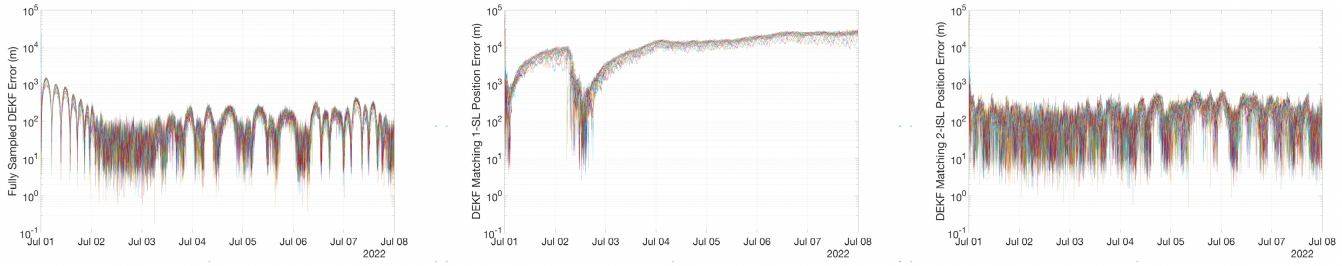

(a) Fully-Sampled DEKF (b) Matching DEKF with one ISL (c) Matching DEKF with two ISLs

Figure 4: **Comparison of the spacecraft localization errors** subject to the different forms of the DEKF. The fully sampled DEKF represents the best case scenario while the single ISL is the minimum functionality. The two ISL case and the fully sampled DEKF are similar magnitudes with the steady state peaks after July 4th being slightly larger than the fully sampled. The behavior of the single ISL in b) is dominated by the anchor nodes availability.

This work was funded by the NASA Game-Changing Development Program (GCD).

References

[Ashman et al. 2018] Ashman, B.; Bauer, F. H.; Parker, J.; and Donaldson, J. 2018. GPS Operations in High Earth Orbit: Recent Experiences and Future Opportunities. In *Proceedings of the* 15^{th} *International Conference on Space Operations*.

[Bernstein et al. 2002] Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The Complexity of Decentralized Control of Markov Decision Processes. *Math. Oper. Res.* 27(4):819 – 840.

[Capuano et al. 2015] Capuano, V.; Botteron, C.; Leclère, J.; Tian, J.; Wang, Y.; and Farine, P.-A. 2015. Feasibility Study of GNSS as Navigation System to Reach the Moon. *Acta Astronautica* 116:186 – 201.

[Chhetri, Morrell, and Papandreou-Suppappola 2003] Chhetri, A.; Morrell, D.; and Papandreou-Suppappola, A. 2003. Scheduling Multiple Sensors Using Particle Filters in Target Tracking. In *IEEE Workshop on Statistical Signal Processing, 2003*, 549–552.

[Chung et al. 2004] Chung, T. H.; Gupta, V.; Hassibi, B.; Burdick, J.; and Murray, R. M. 2004. Scheduling for Distributed Sensor Networks with Single Sensor Measurement per Time Step. In *Proc. IEEE Conference on Robotics and Automation*, 549–552.

[Clark et al. 2016] Clark, P. E.; Malphrus, B.; Brown, K.; Hurford, T.; Brambora, C.; MacDowall, R.; Folta, D.; Tsay, M.; Brandon, C.; and Team, L. I. C. 2016. Lunar Ice Cube: Searching for Lunar Volatiles with a Lunar Cubesat Orbiter. In 48^{th} *American Astronomical Society Division of Planetary Science Meeting*, 223–03.

[Cohen et al. 2015] Cohen, B. A.; Hayne, P. O.; Greenhagen, B. T.; and Paige, D. A. 2015. Lunar Flashlight: Exploration and Science at the Moon with a 6U CubeSat. *Proceedings of the American Geophysical Union Fall Meeting* 2015.

[Cushing et al. 2007] Cushing, W.; Kambhampati, S.; Mausam; and Weld, D. 2007. When is Temporal Planning *really* Temporal? In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1852 – 1859.

[Edmonds 1965] Edmonds, J. 1965. Paths, Trees, and Flowers. *Can. J. Mathematics* 17:449 – 467.

[Hagenau et al. 2021] Hagenau, B.; Peters, B.; Burton, R.; Hashemi, K.; and Cramer, N. 2021. Introducing the Lunar Autonomous PNT System (LAPS) Simulator. In *2021 IEEE Aerospace Conference (50100)*, 1–11.

[Israel et al. 2020] Israel, D.; Cooper, L. V.; Mauldin, K.; and Schauer, K. 2020. LunaNet: A Flexible and Extensible Lunar Exploration Communications and Navigation Infrastructure and the Inclusion of Smallsat Platforms. In *2020 Small Satellite Conference*. Utah State University.

[Kerner et al. 2016] Kerner, H.; Hardgrove, C.; Bell, J.; Amzler, R.; Babuscia, A.; Beasley, M.; Burnham, Z.; and Cheung, K.-M. 2016. The Lunar Polar Hydrogen Mapper (LunaH-Map) Cubesat Mission. In *2016 Small Satellite Conference*. Utah State University.

[Mourikis and Roumeliotis 2006] Mourikis, A., and Roumeliotis, S. 2006. Optimal Sensor Scheduling for Resource-Constrained Localization of Mobile Robot Formations. *IEEE Transactions on Robotics* 22(5):917–931.

[Ny, Feron, and Dahleh 2011] Ny, J. L.; Feron, E.; and Dahleh, M. A. 2011. Scheduling Continuous-Time Kalman Filters. *IEEE Transactions on Automatic Control,* 56(6):1381 – 1394.

[Obszarski and Jastrzebski 2017] Obszarski, P., and Jastrzebski, A. 2017. Edge-coloring of 3-uniform hypergraphs. *Discrete Applied Mathematics* 217:48–52. Combinatorial Optimization: Theory, Computation, and Applications.

[Pereira and Selva 2020] Pereira, F., and Selva, D. 2020. Exploring the Design Space of Lunar GNSS in Frozen Orbit Conditions. In *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 444–451.

[Sin et al. 2021] Sin, E.; Arcak, M.; Nag, S.; Ravindra, V.; Li, A.; and Levinson, R. 2021. Attitude Trajectory Optimization for Agile Satellites in Autonomous Remote Sensing Constellations. In *Proc. AIAA SciTech Forum*.

[Stadter et al. 2008] Stadter, P. A.; Duven, D. J.; Kantsiper, B. L.; Sharer, P. J.; Finnegan, E. J.; and Weaver, G. L. 2008. A Weak-Signal GPS Architecture for Lunar Navigation and Communication Systems. In *Proceedings of the IEEE Aerospace Conference*, 1 – 11.

[Tian et al. 2014] Tian, J.; Wang, Y.; Wang, W.; Shi, P.; Capuano, V.; Leclère, J.; Botteron, C.; and Pierre-André, F. 2014. Cross-Band Aided Acquisition on HEO Orbit. In *Proceedings of the International Astronautical Congress*.

[Wang 2014] Wang, Z. 2014. Review of Chip-Scale Atomic Clocks Based on Coherent Population Trapping. *Chinese Physics B* 23(3):030601.

[Wen et al. 2019] Wen, Y.; Zhu, J.; Gong, Y.; Wang, Q.; and He, X. 2019. Distributed Orbit Determination for Global Navigation Satellite System with Inter-Satellite Link. *Sensors* 19:1031–1046.