Optimally Convergent Autonomous and Decentralized Tasking with Empirical Validation

Samuel J. Fedeler[*], **Marcus J. Holzinger**[†]
University of Colorado at Boulder
William Whitacre
Draper Laboratory

ABSTRACT

This paper presents a novel solution to the space domain awareness (SDA) sensor tasking problem using decentralized, autonomous planning. Decentralized methodologies for sensor tasking are necessary because of an increasing number of active sensors and resident space object populations. The proposed approach describes the problem as a Markov game, a schema considered in detail in multi-agent reinforcement learning literature. Decentralized Monte Carlo Tree Search (MCTS) is applied to optimize rewards with communication between agents assumed over regular random directed graphs. The use of random graphs enables robust communication between agents and guarantees upper bounds on communication gaps. The proposed approach is validated empirically and in simulation. First, a simulation of a cislunar tracking problem tasking both space-based optical sensors is illustrated, demonstrating capabilities to autonomously track space objects through maneuvers. The Vision, Autonomy, and Decision Research (VADeR) observatory is then automated and utilized for a geostationary tasking campaign. This study is the first known methodology to fully decentralize the SDA sensor tasking problem with empirical validation. As such, the research presented in this paper is an important step towards meeting the demands of an evolving near-Earth environment.

1. INTRODUCTION

Over recent years, surging investment in space operations has greatly increased the population of active spacecraft in the near-Earth environment. While spacecraft in low-Earth orbit (LEO) and medium Earth orbit (MEO) may largely be tracked by technologies such as phased-array radar, additional challenges arise when objects are placed in more distant orbits. At and above geosynchronous orbit distances, it generally becomes advantageous to utilize optical sensors for space object detection and tracking, but because such sensors offer comparatively small fields of regard, only a small number of objects may be captured in a single observation. Optical sensors must therefore be tasked to track subsets of a considered population. This problem of "catalog maintenance" is well-studied in literature and is formulated in a large variety of manners.

It is most common to consider the catalog maintenance problem in an information sense, utilizing knowledge of the underlying dynamical system and the applied estimators. Classically, these features are applied by sequentially prioritizing covariance features [1, 2] or the immediate nonlinearity of the dynamics [3]. Such methods are often myopic, and more recent literature has recognized a critical feature of the catalog maintenance problem is when to observe a candidate object that is high value in such metrics. Often, for reasons as such as an object being occluded, nearing a threshold probability of detection, or achieving a future close approach, there is a decision trade space in time. These factors motivate more recent methodologies incorporating decision making over time horizons, in particular the use of the Markov decision process (MDP) formalism.

An MDP applies stochastic transitions between states, which fully encapsulate current information, over which an actor may make sequential decisions, receiving a reward at each epoch. The goal of an MDP is to maximize long-term reward. Most prominently, the reinforcement learning paradigm utilizes MDPs as the underlying problem setting, and such literature has been incorporated into the catalog maintenance problem [4, 5]. Alternatively, a decision-time planning algorithm called Monte Carlo Tree Search (MCTS) can be applied to solve MDPs in an online, anytime manner, with great success in partially observable [6] and continuous [7] environments. As such, it is an excellent candidate for the decision space presented in the catalog maintenance problem, and was previously applied to cislunar space object tracking [8].

[*]Johns Hopkins University Applied Physics Laboratory
[†]Associate Professor, H.J. Smead Faculty Fellow, Smead Department of Aerospace Engineering Sciences

The goal of Monte Carlo Tree Search is to determine an ideal immediate action to take given initial knowledge of the environment; this is a clear differentiation from reinforcement learning methods, where the problem goal is to estimate the value function of a problem or state-action value across any state of the environement, determining a policy for maximizing long-term value. The specificity of MCTS is quite useful in large problem settings that often require a large amount of offline learning for reinforcement learning-based methods, and MCTS is easily adaptable to scenarios in which the environemnt changes from prior assumptions. Examples for such settings in space object tracking include the coordination of multiple observers, tracking maneuvering space objects, or the sudden loss or denial of an observer.

While MCTS was initially considered for centralized decision making tracking objects that remain in periodic orbits, recent work has been submitted extending the MCTS methodology for maneuvering space object tracking [9] and decentralizing the MCTS methodology for many-observer settings [10]. It is also critical to empirically validate the MCTS methodology and the key outcome of this paper is application of MCTS to the Vision, Autonomy, and Decision Research (VADeR) observatory at the University of Colorado at Boulder.

The remainder of this paper is outlined as follows. Section II presents an overview of the sensor tasking methodology as is utilized in [10]. Section III continues with an outline of autonomous and distributed observatory operations for the VADeR observatory. Results using the tasking methodology are then presented in Section IV, and two cases are considered. An overview of submitted decentralized tasking results in cislunar space [10] is first presented, followed by empirical validation of the tasking methodology alongside Two-Line Element information over a geostationary tasking campaign over the VADeR observatory. These results demonstrate that MCTS is a robust methodology for autonomous and decentralized space domain awareness operations.

2. DECENTRALIZED MONTE CARLO TREE SEARCH FOR SPACE DOMAIN AWARENESS

Generically, a multi-agent decision problem is any scenario in which many agents operate over a Markov game [11], a problem that may be defined by the set of elements $\{K, S, \{A^i\}_{i\in\{1,\dots,K\}}, P, \{R^i\}_{i\in\{1,\dots,K\}}, \gamma_i\}$. These definitions may be outlined as

- K: the number of agents considered in the problem.
- S: the probabilistic state of the environment shared across agents.
- $^i\}$: the action space for agent i
- $P : S \times A$: a transition function between states over time conditioned on global sets of actions.
- R^i: a scalar reward function for the ith agent.
- $\gamma_i \in [0,1]$: the discount factor over time for agent i.

When narrowing the problem scope to space object tracking, we consider the following problem design.

- K: the number of observing agents considered in the problem.
- S: An ensemble of Gaussian space object state estimates, with objects labeled $1,\dots,N$.
- $\{A^i\}$: $[1,\dots,N]$, describing the object an agent should observe at a given epoch
- $P : S \times A$: Kalman updates are performed for each observation an agent makes, then nonlinear propagation is performed using the underlying dynamical system.
- R^i: Generally, reduction in state uncertainties across the catalog.
- $\gamma_i \in [0,1]$: the discount factor over time for agent i.

In this problem setting, each agent may individually generate a search tree that informs the long term value of feasible observational actions for that agent. Detailed discussion of tree search methods is presented in [8], and the developed methodology combines convergence guarantees of MCTS with domain knowledge that may be utilized in sampling heuristics. With the assumption that actions sampled using domain knowledge achieve a value within a neighborhood ε of the optimal value, improved claims on convergence may be made [12]. Extended action sampling heuristics are developed in [13], and a combined heuristic utilizes features such as the hypervolume of the projection of state uncertainty into measurement space, threshold detection of prior manuevers, and local analysis of the dynamical flow about the first order state estimate. As each agent generates its own search tree, this methodology scales in constant time onboard each agent as more observers are considered.

Decentralizing the decision-making process introduces several new challenges to this problem setting. As each agent plans locally, it must consider the impact of actions that other agents take. Due to combinatorial explosion, it is infeasible for agents to search over the full action space of all other agents, and this information must be obtained from the plans that other agents generate. A key component of this problem, then, is the formulation of an efficient and robust architecture for communicating decision information between agents. In [10], we identify that agents may efficiently communicate the best local action trajectories found via MCTS to other agents, and if an agent has knowledge of the action trajectory another agent expects to take, local rewards may then be conditioned on that action trajectory. When no information is available, local rewards are evaluated with stochastic action selection for other agents.

Ideally, an observing agent could pass information to every other agent in the problem, but such a communication paradigm is prohibitively costly as the number of cooperating agents grows large. An efficient communication architecture such as a hub and spoke system may be utilized, but such an architecture is susceptible to denial of operations along any line of communication. In [10], agents are assumed to intermittently augment a set of downstream agents to which all known information on local action trajectories and action trajectories of other agents are passed. A directed random graph formulation of a communication architecture is developed that offers guarantees on the existence and maximal length of lines of communication to any other agent. Realistic communication schemes for many-agent decision problems could hope to combine the guarantees of randomized algorithms with known efficient architectures that closely place agents with overlapping decision spaces and minimize necessary lines of communication.

It is worthwhile to visualize this discussion and consider how these decentralized methods may be applied to a Markov game. Algorithm 1 presents the methods by which a search tree is generated locally and results are passed to other agents in the problem. The Dec-MCTS [14] selection methodology is applied with modifications, and agents alternate between search tree generation and communication steps. An agent communicates its current best action trajectory, and in addition, it also propagates any other actions sets it has received from other agents through the graph. With timestamped trajectories, each agent then has downstream access to recent likely actions for all other agents with the assumption of a fully connected graph.

During simulation, an agent applies the actions specified for other agents to inform expected rewards; if no action trajectory is available for another agent, an action is randomly sampled. Progressive widening [15] is applied for large action spaces to simulate new actions within the search tree. A random rollout is applied if a new action is desired, and domain knowledge may be applied to generate new actions. Simulation proceeds until a time or computation budget is exhausted, at which point the D-UCB [16] maximizing action is applied.

The outlined tasking methods may be applied in simulation or asynchronously with observatory control, imaging, and processing pipelines toward the overall goal of maintenance of a large catalog of space objects. To better integrate tasking techniques with the VADeR observatory, a Python interface to the underlying C++ models implementing decentralized MCTS was developed using SWIG [17]. In simulation, looser assumptions on necessary observations may be made, allowing for a shorter decision cadence. During observatory operation, the tasking procedure is performed over a receding horizon while data collection for the prior tasking decision is performed. Tasking decisions are assumed over a longer horizon of two minutes while shorter exposures are taken within the observatory. This admits redundancy for cases in which poor imagery is collected and allows for observation of space objects over longer orbital arcs.

Algorithm 1 Dec-MCTS over random graphs.

Input: belief b, objective g, t_{max}, τ, d
Output: optimal action for local agent i

$\mathscr{T} \leftarrow$ initialize D-UCT tree
$t \leftarrow 0$
$a_{(i)} \leftarrow []$ ▷ Received action trajectories for other agents
$a_i \leftarrow []$ ▷ Optimal action trajectory
$comm \leftarrow []$ ▷ outward edges for communication
while $t < t_{max}$ **do**
  for τ iterations **do**
    $\mathscr{T} \leftarrow \text{simulate}(\mathscr{T}, b, d, a_{(i)}, g, t)$
    $t \leftarrow t+1$
  $a_i \leftarrow I_t(\mathscr{T})$ ▷ D-UCB applied recursively
  if increased outdegree needed **then**
    $comm \leftarrow [comm, rand([1,...,K] \setminus (i \cup comm))]$
  transmit$(comm, a_i \cup a_{(i)})$
  return a_i

Algorithm 2 Dec-MCTS Simulation Loop.

function SIMULATE$(\mathscr{T}, b, d, a_{(i)}, g, t)$
  if $d = 0$ **then return** $\mathscr{T}, 0$
  $a \leftarrow$WIDEN$(\mathscr{T})$
  $a' \leftarrow$DRAW$(a_{(i)}, |\mathscr{A}|)$ ▷ Take received or random action for other agents
  $b' \leftarrow T(b, a \cup a')$
  $r \leftarrow g(b, a \cup a', b') - g(b, a', T(b, a'))$
  if $a \notin C(\mathscr{T})$ **then**
    $C(\mathscr{T}) \leftarrow \{a, b', r\}$
    $\mathscr{T}a, r_\gamma \leftarrow$ROLLOUT$(\mathscr{T}a, b', d-1, a_{(i)}, g, t)$
    $R(t) \leftarrow r + \gamma r_\gamma$
  else
    $\mathscr{T}a, r_\gamma \leftarrow$SIMULATE$(\mathscr{T}a, b', d-1, a_{(i)}, g, t)$
    $R(t) \leftarrow r + \gamma r_\gamma$
  $N(\mathscr{T}) \leftarrow N(\mathscr{T}) + 1$
  visits$(\mathscr{T}) \leftarrow$visits$(\mathscr{T}) \cup t$
  $Q(a) \leftarrow \bar{X}_{a,t_a}(\gamma)$ ▷ Discounted empirical average reward **return** $\mathscr{T}, Q(a)$

function WIDEN$(\mathscr{T})$
  if $C(\mathscr{T}) \leq N(\mathscr{T})^\alpha$ **then**
    $a \leftarrow$NEW$(b, C(\mathscr{T}))$ ▷ Simulate new action
  else
    $a \leftarrow I_t$ ▷ Take maximized D-UCT arm.
  return a

3. VADER OBSERVATORY OPERATIONS AND AUTOMATION

It is worthwhile to provide an overview of the VADeR observatory as a whole with specific focus on the necessary automation steps taken to empirically validate the decision-theoretic tasking methods. The VADeR observatory, installed at the University of Colorado at Boulder in July 2021, fields a suite of optical, infrared, and event sensors. The primary telescope in the observatory, "SITH", is a 0.6 meter f/6.5 corrected Dall-Kirkham optical tube with an approximate diagonal field of view of 44 arcminutes that operates in the visible spectrum using an array of Sloan filters. Alongside the primary telescope, the observatory also operates a coaligned array, "PANOPTICON", of four 0.2 meter f/3 Riccardi-Honders astrographs of approximate diagonal field of view of 2.82 degrees. Each astrograph on the coaligned telescope utilizes a different supporting sensor, and the observatory may use this system for observation with Sloan filters in the visual spectrum, optical spectroscopy, observation in the infrared spectrum, and rate observation using an event sensor. This instrument may also be reconfigured to stack imagery in the visual spectrum or yield a mosaiced field of view of 5 degrees. All instruments in the observatory are operated using Planewave L600 mounts, allowing for slews of up to 25 degrees per second and accurate rate tracking of space objects.

A variety of developments were required to ensure safe autonomous operation over a variety of systems. A visualization of the observatory architecture is provided in Figure 1. In order to safely utilize observatory sensors, constraints are placed on mount control and dome operation. In order to open the VADeR dome, weather APIs are utilized to ensure cloud cover, precipitation probabilities, wind speeds, and humidity remain below desired thresholds. In the future, automated visual inspection of any snow cover using security cameras will also be applied. To autonomously operate observatory mounts, the VADeR dome must also be opened.

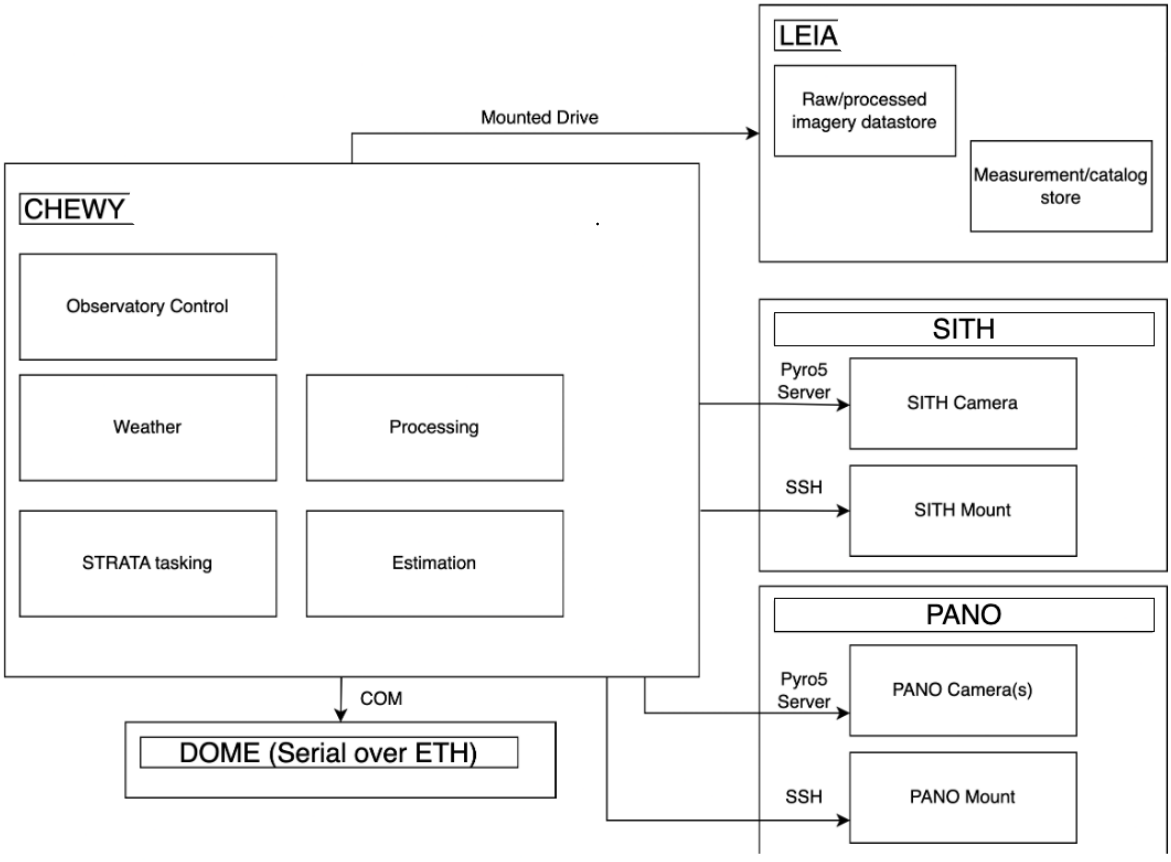

Fig. 1: Operations diagram for the VADeR observatory

Autonomous operation of the VADeR observatory also requires coordination between a variety of systems. The

VADeR CHEWY server acts as the central node in autonomous operation, managing any control input, image processing, tasking, and estimation threads. The server stores generated data over a mounted drive and interfaces remotely with telescope control systems. A client-server architecture is utilized for imaging with each telescope camera, and both the SITH and PANOPTICON mounts are tasked via SSH. Finally, the VADeR dome is operated over a networked serial interface.

3.1 Data Processing and Imaging

This section outlines the processing and source extraction pipelines utilized within the observatory. Each processing step is abstracted such that custom methodologies may be applied, and methods are exemplified specifically for the sCMOS sensors utilized within the observatory for visual spectrum applications. A visualization of the processing pipeline as a whole is presented in Figure 2.

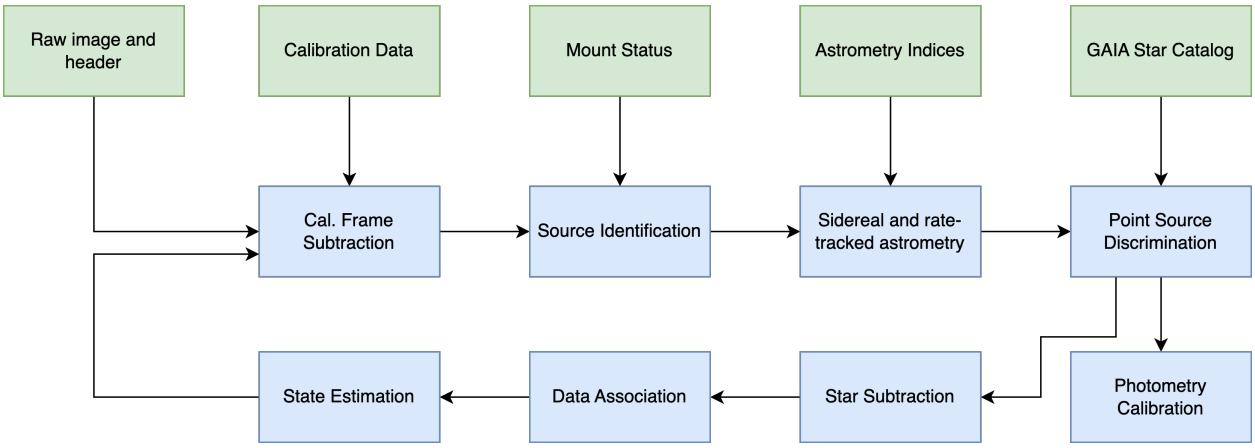

Fig. 2: Image Processing pipeline for the VADeR observatory

Within the observatory, bias and dark frames consistent with the exposure length are captured each campaign prior to observation. Because the sCMOS sensors used are relatively nonlinear in exposure time, dark frames at the exposure time are necessary for accurate subtraction of dark current. In addition to these frames, flat field images are necessary for correction of vignetting within images, a feature that is especially prominent for the 0.2 meter optical systems.

A variety of methods may utilized for subtraction of background noise from imagery, and a mix of open source and custom methodologies may be applied, including iterative sigma clipping and mode estimation [18], sigma clipping and box interpolation, and iterative polynomial fitting [19]. The resultant background-subtracted image may then be expected to be zero-mean and adequately prepared for source identification. Optionally, additional procedures such as registration of hot pixels and correction of edge pixels may be applied.

Largely, the source identification process makes use of the well-documented Source Extractor library [20]. Within Source Extractor, a pyramidal convolution is applied to threshold and merge sources above a specified photometric signal to noise ratio. For point sources within the image, this methodology may be used to accurately identify the point source centroid, and for streaking sources in an image, the methodology may be used to identify streak centroids corresponding to the location in the image of the streaking object at the midpoint of exposure. Streaks act as a somewhat challenging problem for deblending potentially separate objects, but Source Extractor is often used for such processes with use cases for objects such as extended galaxies. Generally, saddle points are utilized to make deblending decisions, and since one may expect that a bright streak is resolved in an image as a multivariate combination of error functions [21], Source Extractor most commonly identifies streaks as a single source.

The resultant pixel locations of source centroids are then applied to generate an astrometric solution for the telescope field at the midpoint of exposure. This result is generated using the Astrometry.net library [22] and the 2MASS star catalog [23]. With information sourced from mount statuses on approximate pointing and a priori knowledge of approximate sensor field of view, this solution is generated on the order of milliseconds. As a result, pixel information at the midpoint of exposure can be translated into angular information on the celestial sphere. Augmenting this process with accurate timing and high-frequency information on mount slew rates, one may then describe pointing throughout the exposure.

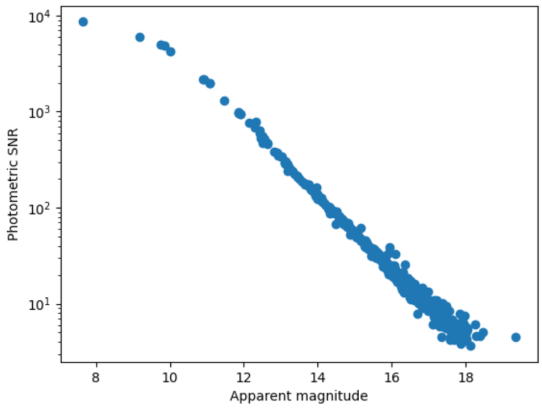

(a) Photometric SNR of detections against star apparent magnitudes.

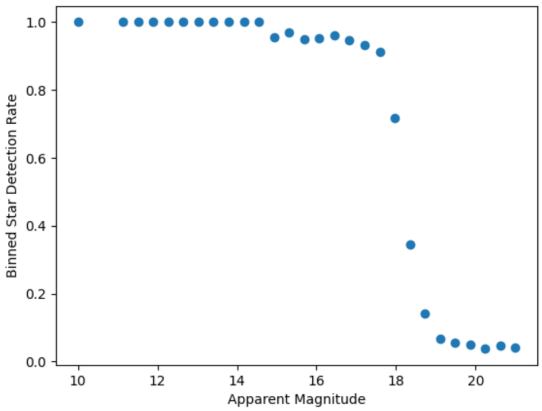

(b) Binned star detection rate by star apparent magnitudes.

Fig. 3: Apparent magnitudes plotted against photometric SNR and detection rates for the SITH telescope.

Because the 2MASS catalog generally only includes stars with apparent magnitude of 15 or brighter, the GAIA DR3 catalog is applied for star subtraction processes and photometry estimation. Currently, the data release consists of known apparent magnitudes for approximately 1.8 billion sources with a limiting magnitude of 21 [24]. Sources in the catalog are matched using k-d trees, and a regression between known apparent magnitude and image photometric SNR may be formed. The resultant regression may then be used to estimate the apparent magnitude of detected space objects. Additionally, matched objects may then be subtracted from the image. Subtracting stars from imagery is relatively straightforward when stars are realized as point sources, but is challenging when stars streak through the sensor field of view. In such cases, many artifact detections can be left in imagery, necessitating multi-target tracking methods for data association. The star registration process is visualized in Figure 3 and used to demonstrate an approximate limiting magnitude of 18 for a 30 second exposure using the SITH telescope.

A source may then be associated with angular and photometric information, and it remains to associate observations with prior state estimates or new tracks. For this problem, a variety of methods may be incorporated, including joint probabilistic data association, [25], multiple hypothesis tracking [26], or finite set statistics [27]. Currently, for rate-tracked imagery, a Gaussian Mixture Probability Hypothesis Density Filter is used to confirm tracks from sequential imagery [27]. Detections may then be gated for a specific track, and the resultant measurement set can then be associated with a catalog object. The effects of this process are visualized in Figure 4, where 11 images are taken centered on the ECHOSTAR 11 satellite with two other geostationary objects in the field of view. The dynamical evolution of tracks and measurement updates in the full state space may then be considered in a variety of manners, most commonly using the Unscented Kalman filter [28]. Orekit is utilized within this context for state and uncertainty propagation [29].

4. CISLUNAR CATALOG MAINTENANCE

We now consider application of the decentralized MCTS methodology to a tasking problem in cislunar space akin to prior cislunar catalog maintenance simulations [8]. Here, we narrow the focus of the problem to a set of 100 SOs placed in Halo orbits about the Earth-Moon L1 and L2 Lagrange points. Each object intermittently performs stationkeeping maneuvers that are utilized in the truth trajectories for these states, and four agents are tasked to maintain state estimates on all objects in the catalog. Observers may make detections every 120 seconds, and a combination of two lunar surface observers and a space-based observer placed in a L2 Northern Halo orbit are utilized. In addition, the sensing architecture is augmented with a space-based observer placed in a L1 Northern Halo that is also 1:1 resonant with the Earth-Moon synodic period. The L2 Halo observer is instantiated at perilune, while the L1 Halo observer is instantiated at apolune, such that each observer occupies a largely distinct subset of state space throughout simulation, since observer periods are equivalent. The lunar observers are placed at the lunar north and south poles. The simulation

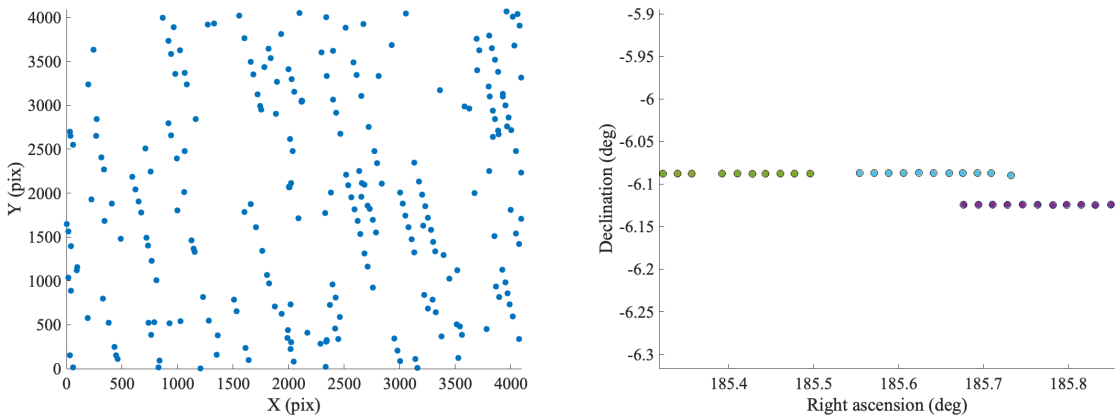

(a) Initial satellite and clutter detections across 11 images of ECHOSTAR 11.

(b) Final geostationary tracks plotted in right ascension and declination.

Fig. 4: Successful data association using Gaussian Mixture PHD filters.

Specification	**Lunar**	**Space-based**
Aperture (m)	0.2	0.5
f-number	3	7
Pixel pitch (μm)	5	5
QE	0.8	0.9
Read noise (pix/s)	2.0	2.0
Optical Transmission	0.756	0.756
Atmospheric Transmission	1.0	1.0

Table 1: Space and ground-based sensor specifications for decentralized cislunar catalog maintenance.

is approximately instantiated at new moon, with an epoch Julian date of 2459153.5. The sensor models are outlined in Table 1.

We first consider the scenario with all sensors operating in a decentralized manner, with sensor communication every 500 tree search iterations. A MCTS rollout heuristic is applied for each sensor that incorporates state uncertainty alongside detected maneuver information and dynamical knowledge of maneuver potential [9]. Specifically, the trace of projected state uncertainty into the observer field of regard is utilized as a measure. If the trace of the projection is significantly larger than that of measurement uncertainty, the observation offers significant information gain. In addition, information from the Cauchy-Green Stress Tensor is incorporated to inform maneuver utility, and the resultant sample weight is rescaled if the studied object previously maneuvered [9]. Each component weight may be expressed as

$$\omega_{y,i} = \frac{\mathrm{tr}(H_i P_i H_i^T)}{\sum_j \mathrm{tr}(H_j P_j H_j^T)} \omega_{\rho,i} = \frac{\rho_i(t+T_i,t)}{\sum_j \rho_j(t+T_j,t)} \tag{1}$$

with the full sampling weight

$$\omega_i = \xi_i \left(\nu \omega_{y,i} + \eta \omega_{\rho,i}\right) \xi_i = \{\, 1 \; t - t_{M,i} > \tau \omega_M t - t_{M,i} \leq \tau. \tag{2}$$

Reduction of space object covariance traces is applied as a reward. In the presented results, a depth $d = 20$ is used for a search time t_s of 4 seconds for each observer. A discount $\gamma = 0.99$ is used. Positional and velocity covariance traces across the catalog are first presented in Figure 5. We find that the ensemble of sensors utilized are sufficient to maintain state estimates across the catalog of SOs, with approximate median $3-\sigma$ positional uncertainties of 2 kilometers and velocity uncertainties of 10 millimeters per second across 27 days of simulation. Note that one object is lost after a large maneuver at approximately 17 days into simulation.

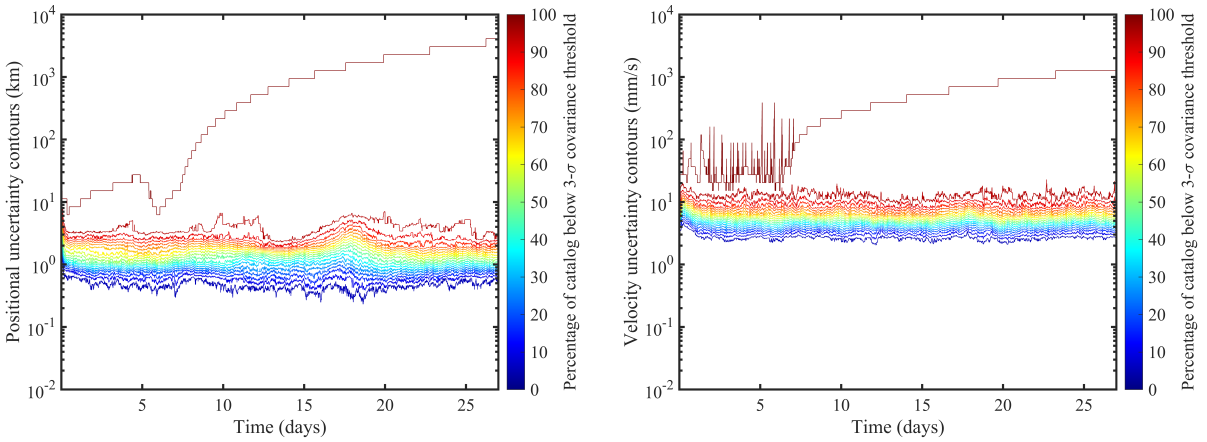

(a) Percentage of $3-\sigma$ positional uncertainties below contours over time in cislunar space.

(b) Percentage of $3-\sigma$ velocity uncertainties below contours over time in cislunar space.

Fig. 5: Catalog uncertainty contours over a month of observation with a suite of four cislunar observers.

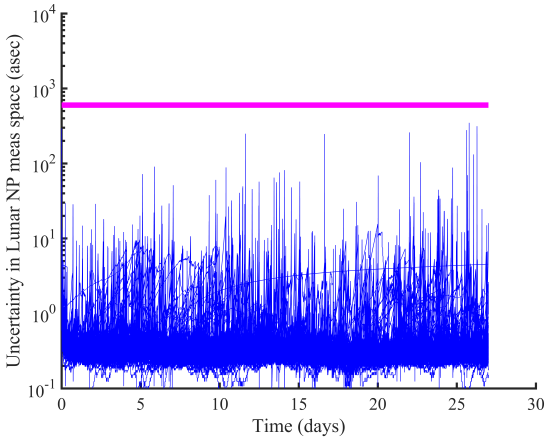

Fig. 6: State uncertainties projected into the field of regard of an observer at the lunar north pole.

It is also worthwhile to discuss the structure of the velocity uncertainty contours presented in the Figure. At any point in the simulation, the largest uncertainty contours evolve on the order of a meter per second; these spikes correspond with maneuver epochs for each epoch, and the largest velocity uncertainties in the catalog at a given time are generally associated with objects that have recently maneuvered. These structures are also clearly visible in Figure 6, in which state uncertainties are projected into the field of regard of the lunar north pole observer. Generally, these uncertainties evolve on the order of 10 arcseconds $3-\sigma$, but clear spikes are visible on the order of several arcminutes as maneuvers occur. These structures are a direct result of the adaptive process noise methods that enable correction of state estimates as measurement residuals increase in norm.

Very interesting structures may be noted when studying the objects tracked by each sensor in greater detail. Figure 7a first visualizes the number of unique objects that each sensor tracks. The full catalog is observed by approximately 2 hours of real time, and the L1 Northern Halo observer is found to be most impactful in the initial periods of the simulation. This is a direct result of the solar phase angle leading to little lunar illumination from the perspective of the L1 observer. Because the simulation is instantiated at new moon, the moon is quite illuminated from the perspective of the L2 observer, and thus, that observer struggles to observe many objects that evolve about L1. On the other hand, as the simulation nears full moon at approximately 14 days into simulation, the L1 Northern Halo observer is challenged to observe objects evolving about L2. This structure is clearly represented in Figure 7b, in which the

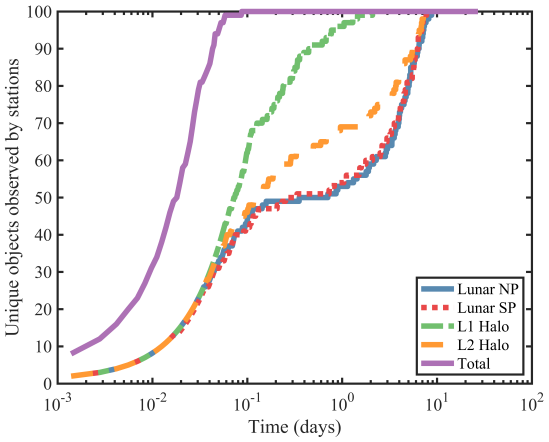

(a) Unique objects tasked by each sensor in cislunar space.

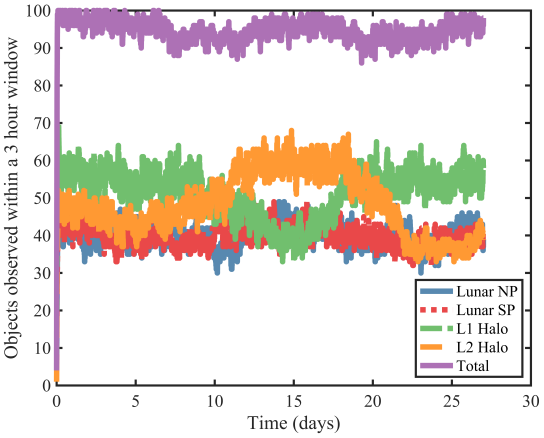

(b) Unique objects tasked by each sensor over sliding three hour windows.

Fig. 7: Tasking data features tracking maneuvering objects in cislunar space.

percentages of the total population the space-based observers detect over a sliding window switch near full moon. It is also worthwhile to briefly discuss the lunar surface observers placed at the lunar north and south poles. Both observers consistently track a subset of the population, and each detects all initially visible objects relatively quickly, on the order of several hours. Other objects then slowly evolve into each sensor's field of regard, and all objects are observed by each sensor by the maximal orbital period in the catalog, on the order of 9 days.

Comparing the presented results to prior research [9] studying cislunar space objects through maneuvers, the combination of lunar surface and space-based observers admits an increase in positional resolution, while velocity uncertainties are relatively comparable between cases. This suggests that increasing the number of observers utilized is less of a factor in the reduction of velocity uncertainties; an alternative observer feature that may be more successful in this regard would be the utilization of a sensor with much greater pixel resolution. If such a sensor were incorporated, subsequent observations would offer greatly increased velocity information, leading to a greater reduction in velocity uncertainties. Another critical factor when considering velocity uncertainties is the presence of many maneuvers. Velocity uncertainties are greatly increased around maneuver epochs, likely leading to long-term challenges in reduction of uncertainties in velocity space.

Finally, it is useful to briefly demonstrate that estimators remain consistent in this scenario. Figure 8 visualizes state uncertainties alongside estimation error for a random object in the catalog, demonstrating successful state estimation. Note that some estimation error may be expected, since the object must be observed to initiate the correction process. This briefly occurs between 14 and 16 days, but is corrected after a maneuver at approximately 17 days into the simulation. It it also interesting to note several periods during which it appears the object was weakly observable or unobservable. Most prominently, no observation appears to occur between 5 and 9 and between 11 and 12 days into the simulation. Maneuvers during the sample case largely appear to be accounted for.

4.1 Robustness to Communication Failures

We now additionally consider a modification to the prior scenario in which communication failures occur over a subset of the simulation. In this case, no communication is possible with the L2 Northern Halo observer between 3 and 9 days into the simulation; during this period, the L2 observer must operate in isolation, while the lunar and L1 observers coordinate without the isolated observer. It is first useful to characterize the catalog uncertainties as a whole across the simulation. Between 3 and 9 days, the connected observers and the isolated observer have a differing notion of state estimates across the catalog, and when communication with the isolated observer is resolved, observations from all observers are processed sequentially over the period of isolation to form a common catalog between all observers. A visualization is given of the catalog used by connected observers between days 3 and 9 in Figure 9 and the isolated observer in Figure 10.

First considering Figure 9, we observe catalog uncertainties that are quite similar to those presented in the prior

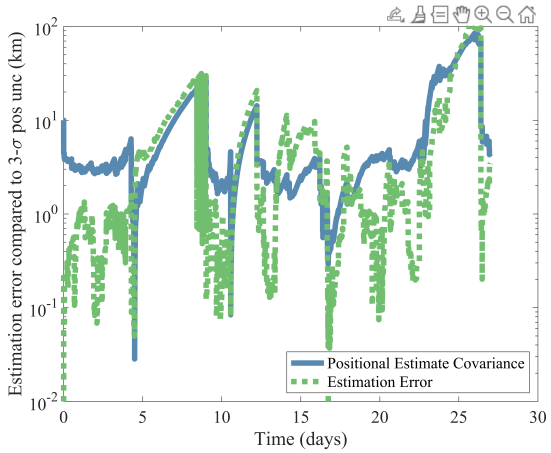

(a) Positional state estimates for a L2 Southern Halo tracked using decentralized MCTS.

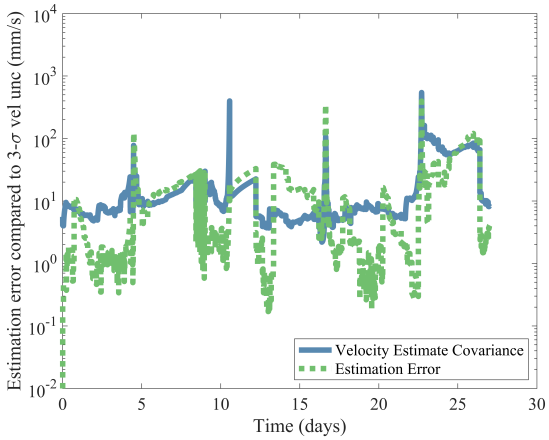

(b) Velocity state estimates for a L2 Southern Halo tracked using decentralized MCTS.

Fig. 8: Successful state estimation with decentralized sensor tasking.

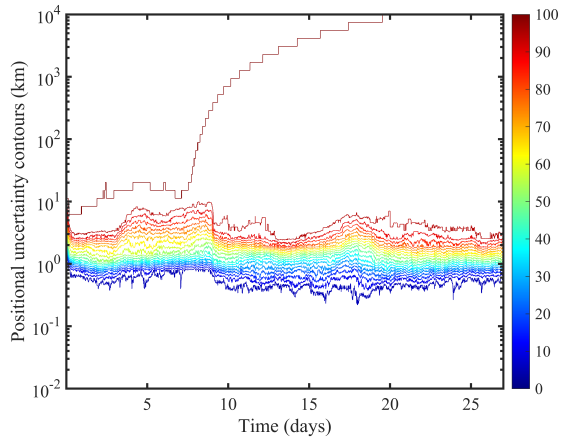

(a) Percentage of $3-\sigma$ positional uncertainties below contours over time in cislunar space.

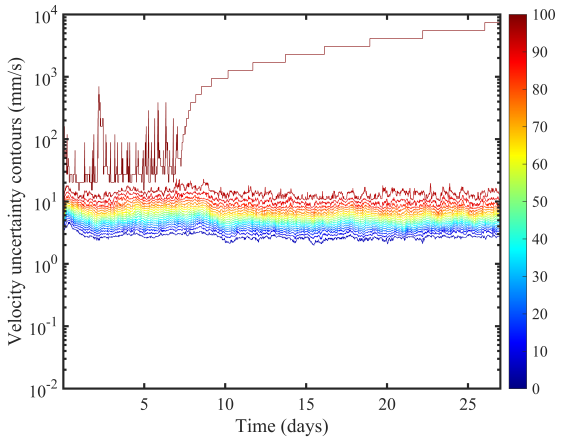

(b) Percentage of $3-\sigma$ velocity uncertainties below contours over time in cislunar space.

Fig. 9: Catalog uncertainty contours in the view of the three observers that remain in communication.

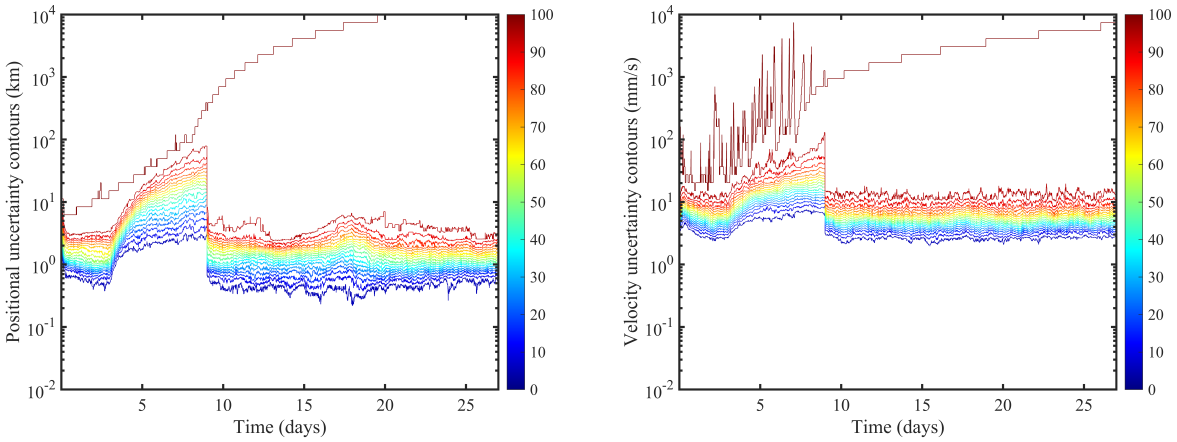

(a) Percentage of $3-\sigma$ positional uncertainties below contours over time in cislunar space.

(b) Percentage of $3-\sigma$ velocity uncertainties below contours over time in cislunar space.

Fig. 10: Catalog uncertainty contours for the isolated observer over a lunar synodic period.

section. In this case, a single object is lost earlier in the simulation, while the remainder of the catalog maintained to approximate median uncertainties of 1 kilometers $3-\sigma$ in position and 5 millimeters per second $3-\sigma$ in velocity. During the period in which the L2 Northern Halo observer is isolated from the remainder of the sensing architecture, slight changes from the original case may be noted. Between 3 and 9 days, a slight increase in positional and velocity uncertainties across the catalog occurs, with median uncertainties reaching approximately 3 kilometers in position by 9 days. When information on the tasking history of the isolated observer suddenly becomes available at this time, the catalog immediately shifts back to covariance levels that are quite comparable to the fully connected case. Interestingly, there is little shift in velocity uncertainties, a feature that suggests it would be quite likely that lunar sensors and a L1 Northern Halo observer were sufficient to maintain the catalog over the period of isolation. This result further demonstrates the robustness of decentralized MCTS to challenges in communication, even in scenarios where space objects are maneuvering.

We next may consider structures in the catalog uncertainties from the perspective of the isolated observer. Figure 10 visualizes these uncertainties, and it is especially critical to study the Figure between 3 and 9 days into simulation. One may note that uncertainties increase somewhat quickly over this gap, with median positional uncertainties for the isolated observer reaching approximately 20 kilometers $3-\sigma$ by 9 days. While the observer continues to detect objects during the period of isolation, the observations it makes are not sufficient to maintain the catalog, though a subset of the catalog appears to stabilize at approximately 6 days into simulation, a feature that is more apparent in the velocity contours. The behavior of the isolated observer during this period may also be compared to the prior case in which the L2 Halo agent was never isolated. Figure 11 visualizes trends in objects tasked during the isolated case, and the structures present here are almost identical to those presented in Figure 7. As such, the isolated observer continues to visit all objects that it is able to observe, again supporting the robustness of the MCTS methodology.

It is briefly worth revisiting structures in estimation error for this scenario, and a random object is selected for visualization. Figure 12 visualizes estimation error for an object following a L1 Northern Halo orbit with three prominent maneuvers at 6, 14, and 22 days into simulation. The estimator is largely consistent throughout the simulation, with a brief period during which positional errors exceed $3-\sigma$ covariances. Smoothing results were note logged in this scenario, so maneuver estimates are not visualized, but covariances successfully account for maneuvers in each case and grow to approximately 0.2 meters during each maneuver epoch. Because of a diversity of viewing geometries, few structures in uncertainties or estimation error as a function of orbital phase are observed.

Finally, it is useful to consider whether uncertainty projections into measurement space remain small enough that custody is maintained for all objects. Figure 13 visualizes measurement space projections, again for the lunar north pole observer, and while spikes occur around maneuver epochs, uncertainties are largely maintained to several arcseconds $3-\sigma$ in measurement space. Compared to Figure 6, a small increase in projected uncertainties occurs between 3 and 9

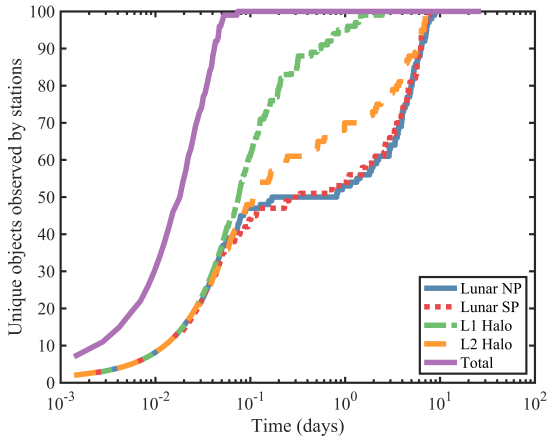

(a) Unique objects tasked by each sensor in cislunar space.

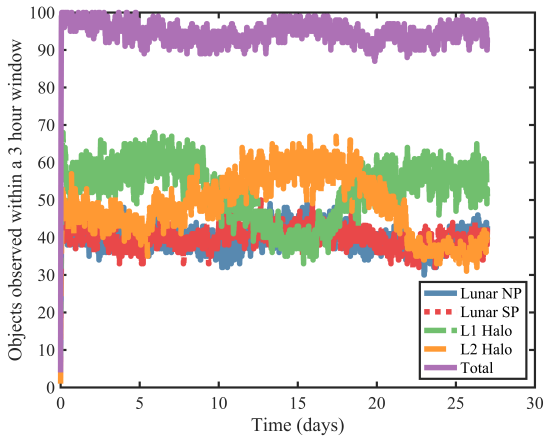

(b) Unique objects tasked by each sensor over sliding three hour windows.

Fig. 11: Tasking data features for the robust communication case.

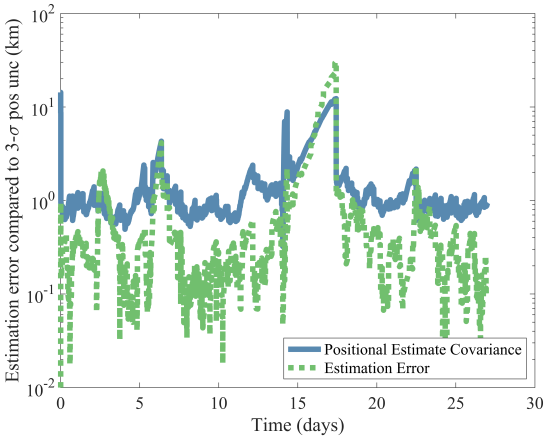

(a) Positional estimation error for a L1 Northern Halo.

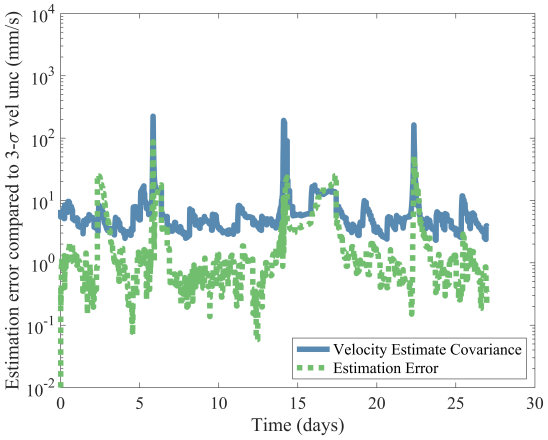

(b) Velocity estimation error for a L1 Northern Halo.

Fig. 12: Estimation error structures with decentralized MCTS tasking and communication failures during observation.

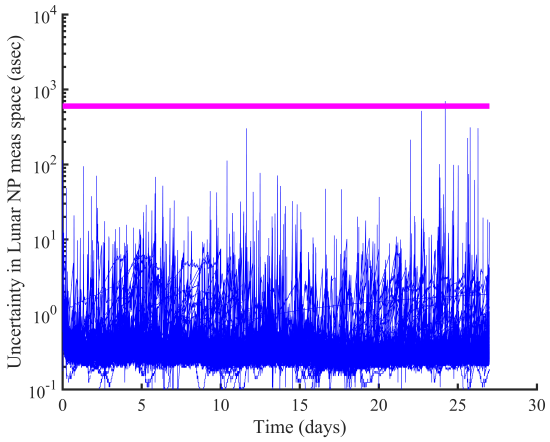

Fig. 13: State uncertainties using decentralized MCTS with communication failures projected into the field of regard of an observer at the lunar north pole.

days as the L2 Northern Halo is isolated, with uncertainties increasing to 10 arcseconds on average during this period.

The results of this scenario establish robustness of decentralized MCTS to failures in communication, and safe operation over large gaps in communication is an ideal feature for space domain awareness in the cislunar regime.

5. OBSERVATORY TASKING

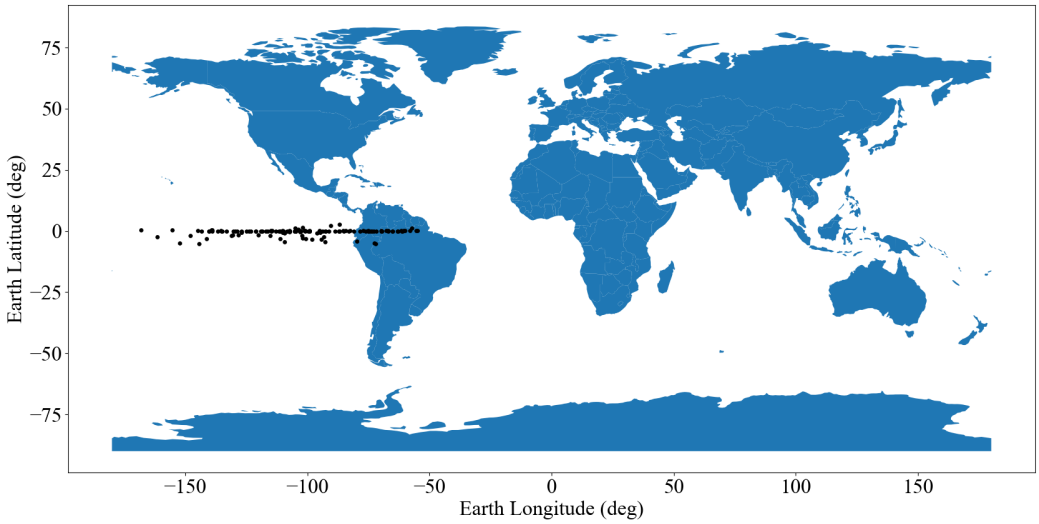

Fig. 14: A geostationary catalog projected onto the surface of the Earth.

The tasking methods presented in this paper are next applied asynchronously with the observatory control, imaging, and processing pipelines outlined in Section III with the overall goal of maintenance of a large catalog of space objects.

An example campaign is performed with decentralized MCTS studying the local geostationary environment. Tracks are maintained for a total of 141 space objects following near-geostationary orbits, with data collected between the dates of May 1st, 2023 and May 9th, 2023. The catalog is visualized as a set ground tracks in Figure 14. As is visible, the catalog consists of a mix of active, publicly known geostationary objects and defunct objects operating in graveyard orbits about the geostationary belt. The catalog is instantiated on May 1st, 2023, using Two-Line Elements and conservative uncertainties.

In Figure 15, we first present catalog uncertainties with tracking using a combination of optical observations and intermittent full state updates sourced from Two-Line Elements. Three data tasking periods are incorporated in the study; the first occurs during the first four hours of observation, the second occurs for four hours at 5 days into the study, and the last occurs at 8 days into the study. Significant reductions in state uncertainties may be noted during the periods of observation, and each period is framed in the Figure, while periods without observation are made opaque. Two-Line Element-based updates are also performed at five noted points in the study, at approximately 1.5, 3.5, 6, 7, and 7.5 days into the test scenario. These updates ensure that state uncertainties don't grow so large that custody is lost, a feature that is especially relevant because of initially conservative $3-\sigma$ uncertainties of 45 kilometers in each position axis and 300 millimeters per second in each velocity axis.

The majority of the catalog is eventually maintained with positional uncertainties on the order of 1 kilometer for the vast majority of the catalog and velocity uncertainties on the order of 100 millimeters per second $3-\sigma$. This is consistent with the results presented in Chapter 5, but because of weather constraints, it was challenging to consistently observe the catalog. This motivates further data collection in Fall 2023, with the goal of nightly observation over at least a one week period. Additionally, the catalog should have been instantiated with much less conservative initial state estimates. It is also worth noting that several objects needed to be reinstantiated because of filter divergence. This was likely a result of small maneuvers that are expected to be quite common for geostationary objects, and a reasonable solution to this issue is the utilization of filters such as the unscented Optimal Control-Based Estimator [30, 9] in a manner that could be integrated with industry standards for propagation such as Orekit. Such a filter would better account for scenarios in which measurement residuals quickly become large as a result of maneuvers. Another potential source of these errors could be uncertainty in the astrometric solutions found using Astrometry.net, as well as inconsistent rates at which astrometric solutions were found for imagery. Because stars are resolved in imagery as small streaks, it is likely that incorporation of a matched filter using the expected profile of stars from mount information could improve pointing estimation.

It is also useful to visualize detection capabilities during periods where observatory tasking is performed. The total number of objects tracked on two nights of observation are visualized in Figures 16 and 17. Note a discrete spike in the second case, where a TLE update was performed midway through the tasking period. In each case, approximately 120 of the 141 space objects were visited. More detailed analysis could be performed considering visits by sensor, since one may expect more objects to be detected by the PANOPTICON sensor, because of a larger field of regard, while the SITH sensor offers more detailed measurement updates.

These results demonstrate the utility of the methods presented in this paper for online and decentralized sensor tasking in a near-optimal manner. Physical results validate the underlying algorithms and motivate use of Monte Carlo Tree Search in more complex sensor tasking scenarios. MCTS shall continue to be applied to the VADeR observatory, considering the long-term impact of autonomous sensor tasking on observatory operations.

6. DISCUSSION AND CONCLUSIONS

This paper presents contributions that greatly augment the autonomous capabilities of optical sensors across a variety of domains, including the VADeR observatory. We first present a methodology for decentralized sensor tasking that scales to many-sensor tasking problems that is demonstrated for a challenging cislunar space object tracking problem. We then outline pipelines developed for safe operation, autonomous imaging, image processing, and data association. These tools are then utilized alongside a Pythonic interface to MCTS libraries to autonomously task the VADeR observatory. Several nights of autonomous observation were performed demonstrating MCTS as a viable algorithm for autonomous space domain awareness in a real scenario. State estimates are maintained for the public catalog of geostationary objects visible from Boulder, Colorado, with further studies planned in the coming months.

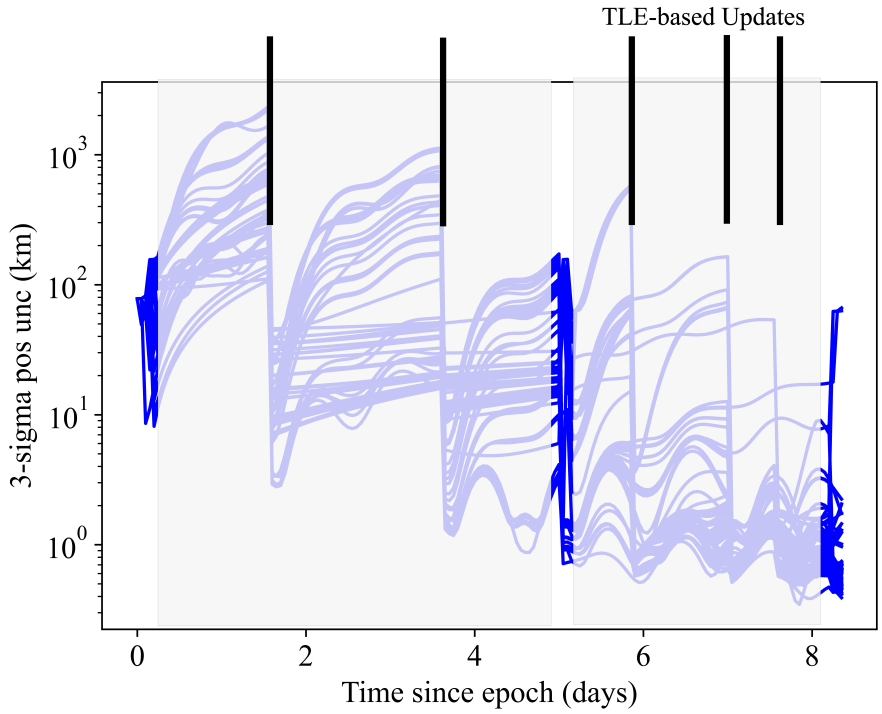

(a) 3-σ positional uncertainties using the VADeR observatory and TLE updates.

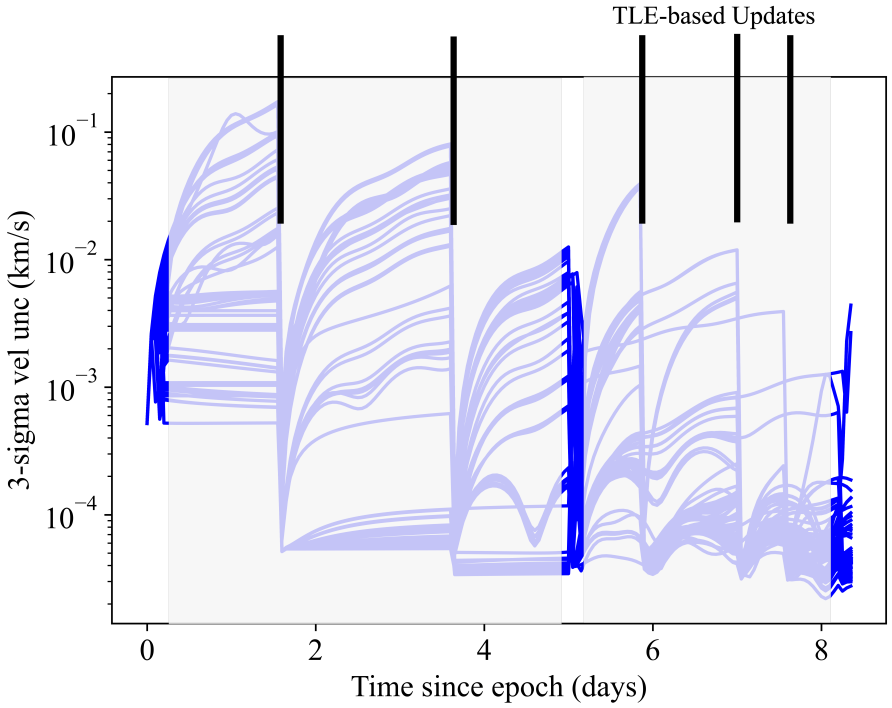

(b) 3-σ velocity uncertainties using the VADeR observatory and TLE updates.

Fig. 15: Catalog uncertainty traces over three nights of observation using the VADeR observatory.

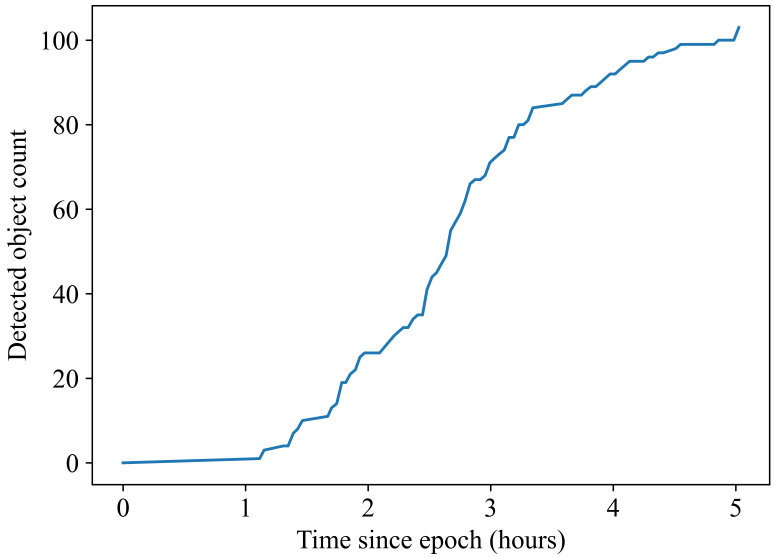

Fig. 16: Unique objects tracked by the observatory on May 1st, 2023.

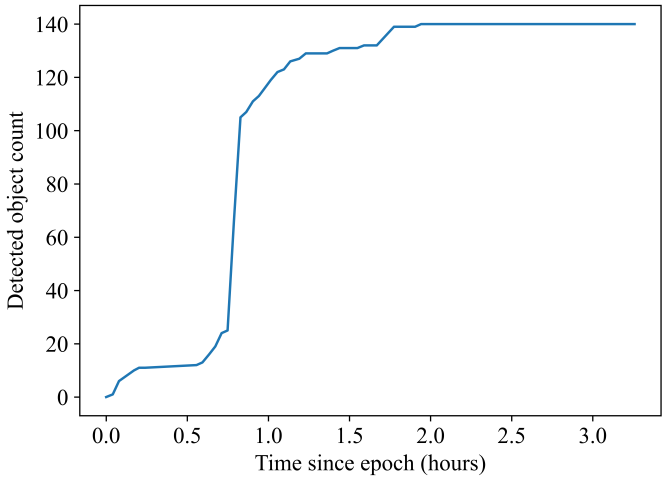

Fig. 17: Unique objects tracked by the observatory on May 6th, 2023.

7. ACKNOWLEDGEMENTS

This research was supported by the Draper Scholars and the NSF Graduate Research Fellows Program Grant Number DGE 1650115. Support for presentation of this work was provided by the Johns Hopkins University Applied Physics Laboratory.

REFERENCES

[1] Keric Hill, Paul Sydney, Randy Cortez, Kris Hamada, Daron Nishimoto, Pacific Defense Solutions, N Holopono St, Kim Luu, and Paul W Schumacher. Dynamic Tasking of Networked Sensors Using Covariance Information. *Advanced for Maui Optical and Space Surveillance Technologies Conference*, (Scenario 1), 2010.

[2] R. Scott Erwin, Paul Albuquerque, Sudharman K. Jayaweera, and Islam Hussein. Dynamic sensor tasking for space situational awareness. *Proceedings of the 2010 American Control Conference, ACC 2010*, pages 1153–1158, 2010.

[3] Patrick S. Williams, David B. Spencer, and Richard S. Erwin. Coupling of estimation and sensor tasking applied to satellite tracking. *Journal of Guidance, Control, and Dynamics*, 36(4):993–1007, 2013.

[4] Richard Linares and Roberto Furfaro. An Autonomous Sensor Tasking Approach for Large Scale Space Object Cataloging. *Advanced Maui Optical and Space Surveillance Technologies Conference*, pages 1–17, 2017.

[5] Peng Mun Siew, Daniel Jang, Thomas G. Roberts, and Richard Linares. *Space-Based Sensor Tasking Using Deep Reinforcement Learning*, volume 69. Springer US, 2022.

[6] Zachary Sunberg and Mykel J. Kochenderfer. Online algorithms for POMDPs with continuous state, action, and observation spaces. In *Twenty-Eighth International Conference on Automated Planning and Scheduling*, 2018.

[7] Michael H. Lim, Claire J. Tomlin, and Zachary N. Sunberg. Sparse tree search optimality guarantees in POMDPs with continuous observation spaces. 2019.

[8] Samuel Fedeler, Marcus Holzinger, and William Whitacre. Sensor tasking in the cislunar regime using Monte Carlo Tree Search. *Advances in Space Research*, (xxxx):1–19, 2022.

[9] Samuel Fedeler, Marcus Holzinger, and William Whitacre. Cislunar Space Object Tracking Considering Maneuver Estimation and Maneuver Utility. *Journal of Guidance, Control, and Dynamics (submitted)*, (xxxx):1–25, 2023.

[10] Samuel Fedeler, Marcus Holzinger, and William Whitacre. Decentralized decision making over random graphs for space domain awareness. *Advances in Space Research (submitted)*, (xxxx):1–25, 2023.

[11] Junling Hu and Michael P. Wellman. Multiagent Reinforcement Learning: Theoretical Framework and an Algorithm. *ICML*, 98:242–250, 1998.

[12] Devavrat Shah, Qiaomin Xie, and Zhi Xu. Non-Asymptotic Analysis of Monte Carlo Tree Search. *Operations Research*, 70(6):3234–3260, 2022.

[13] Samuel J Fedeler, Marcus J. Holzinger, and William Whitacre. Cislunar Space Object Tracking Considering Maneuver Estimation and Maneuver Utility. *Journal of Guidance, Control, and Dynamics (submitted)*, pages 1–25, 2023.

[14] Graeme Best, Oliver M. Cliff, Timothy Patten, Ramgopal R. Mettu, and Robert Fitch. Dec-MCTS: Decentralized planning for multi-robot active perception. *International Journal of Robotics Research*, 38(2-3):316–337, 2019.

[15] Adrien Couetoux, Jean-baptiste Hoock, Nataliya Sokolovska, and Olivier Teytaud. Continuous Upper Confidence Trees. *Learning and Intelligent Optimization*, (section 3):433–445, 2011.

[16] Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6925 LNAI:174–188, 2011.

[17] David M. Beazley. SWIG: An easy to use tool for integrating scripting languages with C and C++. *4th Annual USENIX Tcl/Tk Workshop 1996, TCL/TK 1996*, (July):1–16, 1996.

[18] Peter B. Stetson. Daophot: a Computer Program for Crowded-Field Stellar Photometry. *Publications of the Astronomical Society of the Pacific*, 99:191–22, 1987.

[19] M. Levesque and S. Buteau. Image processing technique for automatic detection of satellite streaks. (February):60, 2007.

[20] E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *Astronomy and Astrophysics Supplement Series*, 117(2):393–404, 1996.

[21] Samuel Wishnek and Marcus J. Holzinger. Astrometry and time-resolved photometry from streaks using calibrated ultra-wide field of view cameras. *AIAA Scitech 2020 Forum*, 1 PartF(January):1–16, 2020.

[22] Dustin Lang, David W. Hogg, Keir Mierle, Michael Blanton, and Sam Roweis. Astrometry.net: Blind astrometric calibration of arbitrary astronomical images. *Astronomical Journal*, 139(5):1782–1800, 2010.

[23] M. F. Skrutskie, R. M. Cutri, R. Stiening, M. D. Weinberg, S. Schneider, J. M. Carpenter, C. Beichman, R. Capps, T. Chester, J. Elias, J. Huchra, J. Liebert, C. Lonsdale, D. G. Monet, S. Price, P. Seitzer, T. Jarrett, J. D. Kirk-

patrick, J. E. Gizis, E. Howard, T. Evans, J. Fowler, L. Fullmer, R. Hurt, R. Light, E. L. Kopan, K. A. Marsh, H. L. McCallon, R. Tam, S. Van Dyk, and S. Wheelock. The Two Micron All Sky Survey (2MASS). *The Astronomical Journal*, 131(2):1163–1183, 2006.

[24] A Vallenari, A G A Brown, and T Prusti. Gaia Data Release 3: Summary of the content and survey properties. *Astronomy & Astrophysics*, pages 1–23, 2022.

[25] Jason Stauch, Travis Bessell, Mark Rutten, Jason Baldwin, Moriba Jah, and Keric Hill. Joint Probabilistic Data Association and Smoothing Applied to Multiple Space Object Tracking. *Journal of Guidance, Control, and Dynamics*, 41(1):1–15, 2017.

[26] Samuel S. Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1 II):5–18, 2004.

[27] Ba Ngu Vo and Wing Kin Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006.

[28] Simon J Julier and Jeffrey K Uhlmann. New extension of the Kalman filter to nonlinear systems. *Proceedings of SPIE*, (7), 1997.

[29] Luc Maisonobe, Véronique Pommier-Maurussane, and Pascal Parraud. Orekit: an Open-source Library for Operational Flight Dynamics Applications. *4th International Conference on Astrodynamics Tools and Techniques*, (February 2021):5, 2010.

[30] Jesse A Greaves and Scheeres. Relative Estimation in the Cislunar Regime using Optical Sensors University of Colorado Boulder Daniel J . Scheeres University of Colorado Boulder. *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, 2021.