# Leveraging Low-Altitude Constellations in Ad-Hoc Lunar PNT System for Distributed Spacecraft Autonomy

Yeji Kim[1]
*Yonsei University, Seodaemun-gu, Seoul, 03722, South Korea*

Brian C. Kempa[2], Caleb Adams[3], and Jeremy D. Frank[4]
*NASA Ames Research Center, Moffett Field, California, 94035, USA*

Richard J. Levinson[5]

*KBR Wyle Services, LLC, Moffett Field, California, 94035, USA*

**This paper investigates leveraging a low-altitude Lunar Position, Navigation, and Timing (LPNT) constellation comprising 100 small satellites at an altitude of 100 km to provide reliable, ad-hoc navigation services for future lunar exploration missions. Autonomous operations of this large-scale swarm rely on two-way Inter-Satellite Link (ISL) measurements, utilizing pseudoranges and relative velocities among satellites. We describe four localization methodologies of the swarm assets utilizing Centralized Extended Kalman Filter (CEKF) and Decentralized Extended Kalman Filter (DEKF) algorithms. We evaluate their localization performance to quantify the positioning accuracy for both satellite swarm assets and lunar surface users, computational efficiency, and robustness to communication failures. We numerically assess four candidate localization methodologies (Redundant CEKF, Leader-Follower CEKF, Distributed CEKF, and DEKF) to address challenges posed by limited ISL measurements and computational resource constraints inherent in small satellite platforms. Simulation results indicate that Redundant CEKF can achieve the highest position accuracy with a median satellite position error of 371.2 m and a lunar surface user localization error of 16.69 m, although it incurs significant computational and bandwidth requirements. Leader-Follower and Distributed CEKF approaches demonstrate balanced trade-offs, achieving slightly degraded position accuracies of 419.6 m median satellite position error but offer more efficient computational and network resource utilization. The DEKF method shows the lowest computational demands but exhibits higher sensitivity to communication failures, with a median satellite position error of 541.7 m. This research explores critical trade-offs among positioning accuracy, computational complexity, and communication overhead, providing valuable insights into the operation and scalability of distributed spacecraft autonomy for future lunar navigation missions.**

[1] Master's student, Department of Astronomy, AIAA Student Member.
[2] Research Engineer, Intelligent Systems Division.
[3] Project Manager, Intelligent Systems Division.
[4] Planning & Scheduling Group Lead, Intelligent Systems Division.
[5] Computer Scientist, Intelligent Systems Division.

## I. Nomenclature

| | | |
|---|---|---|
| $x$ | = | Estimated state vector of a satellite |
| $P$ | = | Estimated state covariance matrix |
| $\bar{x}$ | = | Predicted state from the dynamic model |
| $\bar{P}$ | = | Predicted covariance from the dynamic model |
| $\hat{x}$ | = | Updated state after measurement update |
| $\hat{P}$ | = | Updated covariance after measurement update |
| $u$ | = | Estimated position of the lunar ground user |
| $H$ | = | Measurement sensitivity matrix |
| $R$ | = | Measurement noise covariance matrix |
| $K$ | = | Kalman gain matrix |
| $z_i$ | = | Inter-satellite measurement |
| $z_a$ | = | Measurement from anchor node |
| $f$ | = | Frobenius norm of post-measurement covariance for a satellite triplet |
| $G_H$ | = | Broadcast network topology for state and covariance exchange |
| $G_M$ | = | Measurement network topology for selected ISL measurements |
| *BPLAN* | = | Broadcast network planner |
| *MPLAN* | = | Measurement network planner |
| *HP* | = | Hamiltonian path used for $G_H$ construction |
| $M_S$ | = | Median position error of a satellite |
| $M_G$ | = | Median position error of a ground user |
| $\sigma_S$ | = | Standard deviation of satellite position error |
| $\sigma_G$ | = | Standard deviation of user position error |
| *PDOP* | = | Position dilution of precision |
| *Avail* | = | Availability of navigation service at the Lunar surface |

## II. Introduction

Driven by the growing involvement of public and private exploration partnerships, traditional low Earth orbit missions are shifting to beyond geosynchronous orbit [1]. These upcoming missions aim to foster a sustainable and innovative exploration program, in collaboration with commercial and international partners, to facilitate human expansion throughout the solar system and return new knowledge and opportunities to Earth [2].The Lunar Position, Navigation, and Timing (PNT) technology plays a key role in providing reliable and robust navigation services on the Moon's surface and the South pole region, which is the focal area for upcoming scientific and human exploration missions. To support this new era of lunar operations, including small satellites launched through NASA's Commercial Lunar Payload Services (CLPS) program, the Lunar PNT architectures must be adaptable to smaller platforms like CubeSats. As part of this trend, there are increasing efforts to utilize science missions in Lunar orbit to form non-dedicated, ad-hoc PNT networks.

Two traditional approaches, the Deep Space Network (DSN) and weak signal Global Positioning System (GPS), are established as deep-space navigation technologies for missions beyond the geosynchronous orbit. Beginning in 1958, the DSN was developed to communicate with the Explorer 1 spacecraft based on the use of radiometric tracking in spacecraft navigation [3]. The DSN can provide nearly unfettered coverage to spacecraft beyond low-Earth orbit (LEO), however, increased space mission volume has created concerns about its scalability for future lunar missions [4]. For cislunar applications, DSN can offer positioning accuracy of 100 m (3σ) using three or more geometrically diverse Earth stations [5]. However, its dependence on Earth visibility limits continuous operation. In addition, its limited localization performance and the need for high-power onboard antennas make DSN unsuitable for future lunar missions that demand continuous tracking and precise positioning. To satisfy the increasing requirements of DSN in Lunar applications, spacecraft are also required to improve their onboard antenna power and efficiency of the transmission. However, there is an important aggregate cost trade, in perspective of size, weight, and power (SWaP), between adding capabilities to every spacecraft and adding to a capacity on the ground that serves multiple spacecraft [6, 7].

A weak GPS signal can provide PNT service while the user spacecraft is bound to the Moon, leveraging a single, steerable high gain antenna with the relatively narrow beam which includes all the sources in its field of view [8]. However, at lunar distances, the higher the altitude the receiver is above the GPS constellations, the poorer and the weaker are the relative geometry and the received signal powers, respectively, leading to a significant navigation accuracy reduction [9]. The transmitted power becomes weaker with increasing distance from the Earth as well as

signals tracked from one of the side lobes of the GPS antenna pattern. Additionally, the number of visible satellites and relative geometric condition of the GPS satellites at very high altitude drops dramatically and reduces the navigation solution accuracy. Thus, the weak GPS system is also not an ideal way for lunar PNT service, especially considering the sparse geometry and reduced signal power at high altitudes.

An alternative approach is the design of dedicated lunar navigation constellations, analogous to Earth Global Navigation Satellite System (GNSS), which aim to provide global, continuous 24/7 coverage. For example, the Lunar Communications Relay and Navigation Systems (LCRNS) by NASA and Lunar Navigation Satellite Systems (LNSS) by JAXA are designed to serve as dedicated Position, Navigation, and Timing (PNT) systems for the Moon. However, developing such systems for the lunar environment poses unique challenges, including limited payload capacity of small spacecraft, few lunar ground stations, and lower investment levels compared to GNSS on Earth [10].

To overcome these constraints, NASA is pursuing low-cost lunar navigation platforms by leveraging small-sat assets originally intended for science missions. These satellites can be repurposed to form a low-cost, autonomous, ad-hoc, and on-demand mission-centric Lunar PNT swarm capable of providing PNT services to the future lunar missions [11]. As many planned missions will operate at low altitudes from 30 km to 100 km for mapping and observation, the low-altitude orbital constellations can offer region-specific PNT services. However, these low-altitude orbits are typically unstable and require frequent maintenance or are limited to short durations, often less than 90 days. Moreover, at an altitude of 100 km, satellites offer only limited visibility windows above the horizon from the perspective of ground users, thereby limiting the continuity of PNT service availability.

To enable flexible support under these constraints, we explore the use of a non-dedicated, ad-hoc lunar navigation constellation that facilitates on-demand PNT services. A preliminary study of ad-hoc Lunar PNT system was conducted using 21 spacecraft in 6,500 km altitude frozen orbits to test its feasibility and a basic performance of orbital asset localization among ad-hoc Lunar constellations in small satellites format [11]. These swarm assets are designed for autonomous operations with minimal Earth interaction, reducing dependency on bandwidth and ground resources. The design in [11] demonstrated the feasibility of a decentralized PNT approach, employing a Centralized Extended Kalman Filter (CEKF) and Distributed Extended Kalman Filter (DEKF) algorithms which minimizes onboard operating costs. A follow-up study in [13] introduced communication constraints and a matching algorithm to optimize ISL measurements [13]. Each spacecraft was limited to 2 communications antennas, forcing the selection of measurements and scheduling spacecraft activities to perform the measurements. A matching algorithm is implemented to select the best measurements and schedule position estimation updates. The localization performance of LPNT system is then investigated with increasing levels of network degradation for swarm assets considering the impact of intermittent and permanent communication failure, to demonstrate the robustness and fidelity of proposed ad-hoc PNT service [14]. The study confirmed that the frozen-orbit lunar constellations are highly robust and resilient to communication failures.

In this study, we focus on leveraging low-altitude small-sat science missions, at an altitude of 100km, to support ad-hoc PNT services. We propose a large-scale constellation comprising of 100 satellites, where communications are neither pervasive nor free in the small-sat platform. However, the upcoming low-cost Lunar missions are not able to fully support network resources this large-scale PNT system would require due to their limited communication between systems, computational power, and memory, which are distributed across the systems [15]. Given the limited lunar ground infrastructure and tight SWaP constraints, decentralized, autonomous architectures are necessary [10]. In response to these operational challenges, we explore four candidate localization approaches with distinguishable estimation strategies, based on a centralized extended Kalman filtering (CEKF) and a distributed extended Kalman filter (DEKF) to state estimates, for minimal onboard operating costs. Each satellite performs two-way ranging with mutually visible neighbors, using inter-satellite link (ISL) measurements, including pseudoranges and relative velocity between visible satellites, as sensor inputs [13]. To assess the feasibility of this low-altitude LPNT system in delivering reliable navigation services, we conduct a numerical evaluation of localization accuracy for both the swarm assets and a representative lunar surface user. We also analyze the computational characteristics of each approach, highlighting differences in measurement selection strategies, scheduling mechanisms, and their suitability for constrained onboard environments. Furthermore, the performance of decentralized filtering architectures is examined under the two-ISL antenna constraint, with a focus on how this limitation is addressed through the filter design and matching algorithms. The study also evaluates the system's robustness to network degradation, analyzing the effects of communication failures on the localization performance of both the satellite constellation and the surface user.

The remainder of this paper is organized as follows. Section 3 introduces the proposed low-altitude LPNT constellation comprising 100 small satellites. Section 4 presents the four candidate localization approaches based on CEKF and DEKF algorithms. Section 5 discusses the numerical evaluation results, including localization accuracy, computational performance, and robustness under failure scenarios. Finally, Section 6 concludes the study and outlines directions for future work.

## III. Low-Altitude Lunar Constellation

In preliminary ad-hoc LPNT studies, the non-dedicated swarm assets of the constellation are comprised of 21 small satellites in frozen orbits at 5,500-6,500 km altitude, evenly spaced around 3 circular, 40° inclination orbital planes. The previous studies on frozen orbit constellations mainly focus on investigating the feasibility of low-cost Lunar PNT constellations as a part of NASA's Distributed Spacecraft Autonomy (DSA) project [16], as well as the applicability of DSA technologies to provide PNT services on and around the Moon in the absence of a dedicated constellation [13]. The reliability of the proposed low-cost LPNT system has been further validated in [14], by assessing the potential communication failures, demonstrating the robustness and resilience of frozen-orbit constellations.

Despite the advantages of frozen orbits such as minimal fuel requirements for orbit maintenance and favorable coverage for PNT services at high altitudes, they are not well suited to the orbital profiles of upcoming lunar science missions, which typically require lower altitudes for high-resolution mapping and surface observation. As shown in Table 1, many past and proposed scientific Lunar missions will be mainly in low altitude around 100 km, and are primarily focused on observations near the lunar poles. These planned small-satellite assets present an opportunity to form a non-dedicated, on-demand PNT network, minimizing the need for new infrastructure while leveraging existing spacecraft. This concept can significantly reduce mission design costs by maximizing the use of science-driven platforms to deliver PNT services to lunar surface users [13].

**Table. 1 Proposed Lunar small-sat science missions**

| Mission | Status | Altitude (km) | Inclination (deg) |
|---|---|---|---|
| LunaH-map [17] | Completed | 12 | 90 |
| Lunar Reconnaissance Orbiter (LRO) [18] | Completed | $50-200$ | 90 |
| Lunar Flashlight [19] | Failed | $12.6-52.4$ | 90 |
| Lunar Atmosphere and Dust Environment Explorer (LADEE) [20] | Completed | $20-250$ | 157 |
| Danuri (KPLO) [21] | Ongoing | 100 | 90 |
| Lunar Ice Cube [22] | Failed | 100 | 90 |
| Kaguya (SELENE) [23] | Completed | 100 | 90 |
| Lunar Trailblazer [24] | Failed | 100 | 90 |

We investigate a new low-altitude LPNT scenario using a large-scale constellation of 100 satellites, designed to reflect the orbit characteristics of upcoming science missions. This system consists of 98 satellites in low-altitude polar orbits at 100 km, evenly distributed across seven circular planes, and two additional satellites in frozen orbits at 5,500 km, as shown in Figure 1. The constellation configuration and orbital parameters are summarized in Table 2.

Unlike frozen-orbit constellations, 100-node LPNT swarm operating at an altitude of 100 km depends on spacecraft primarily intended for science objectives. These orbits are inherently unstable and require frequent maintenance, making them suitable only for short-duration missions, typically less than 90 days. This limited time and resources would require satellites to schedule communication activities to deliver the best possible quality of PNT service to Lunar users. In addition, the availability of LPNT service is constrained by each satellite's limited visibility window from the lunar surface, as individual satellites remain within line-of-sight for only a few minutes. This restricted visibility poses challenges to ensuring continuous navigation support, especially for users at the lunar surface. Consequently, PNT service availability is geographically limited, with the proposed constellation primarily supporting operations in the lunar South Pole region, where missions like Polar Resources Ice Mining Experiment 1 (PRIME-1) and Volatiles Investigating Polar Exploration Rover (VIPER) under NASA's Artemis program are focused [23]. These missions aim to characterize water ice resources, critical for future exploration and in-situ resource utilization planning. It is also important to note that the orbits of real science missions may differ from the reference constellation used in

this study. Each mission carries unique science-driven constraints, resulting in variations in resource availability, coverage, and operational flexibility, particularly when operating within the limitations of small-satellite platforms.

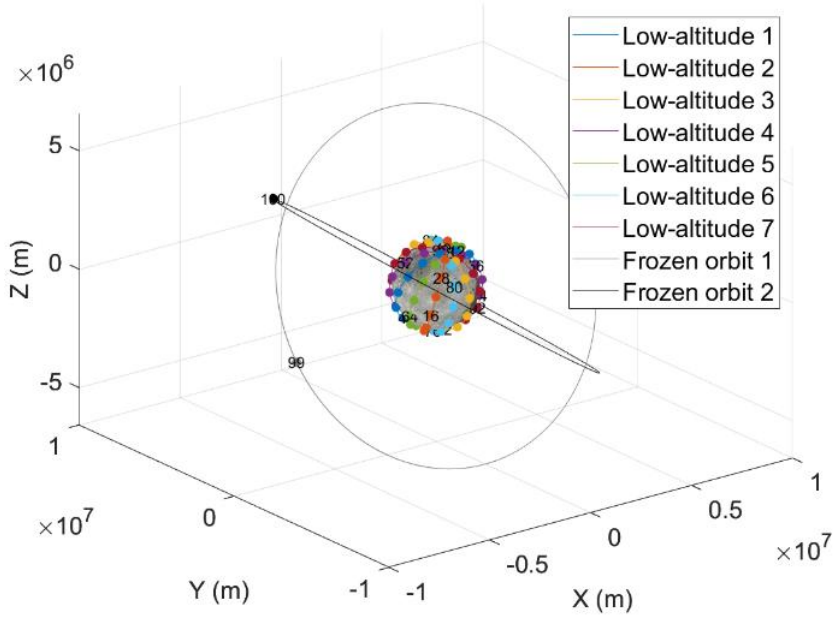



**Fig. 1 Low-altitude Lunar constellation**

**Table. 2 Low-altitude constellations of Lunar PNT system**

| Orbit type | Altitude | Inclination | Planes | Assets | Coverage |
|---|---|---|---|---|---|
| Low altitude | 100 km | 90° | 7 | 98 | South Pole |
| High altitude | 5,500 km | 39.71° | 2 | 2 | |

## IV. Four Candidate Localization Approaches

The proposed Lunar constellation estimates the states of the swarm assets using two-way ISL measurements from mutually visible asset nodes ($z_i$), which include pseudoranges and relative velocities, as well as reference observations from anchor nodes ($z_a$) in centralized and decentralized extended Kalman filtering approaches. These ISL measurements are obtained using only two directional antennas per satellite, while anchor node measurements come from external receivers that do not interfere with other functionality of the spacecraft. In the current LPNT system, two ground-based anchor nodes, one at each pole, and one asset node on the Lunar Gateway space station in Near Rectilinear Halo Orbit (NRHO) [24] are incorporated into the estimation process with an assumption of no position error. These anchor nodes serve as ground-truth references, based on independent positioning system such as DSN or weak-GPS, and are essential for resolving absolute positioning, since swarm-only ISL data cannot fully determine satellite states without an external reference.

To manage limited ISL resources from the two-antenna constraint, this study employs two types of network topologies for data exchange among distributed spacecraft: a broadcast network topology ($G_H$) and a measurement network topology ($G_M$), each serving distinct purposes and requiring different setup techniques to support global state exchange and localized measurement selection, respectively. The broadcast topology exchanges each satellite's state ($x$) and covariance ($P$) globally across the swarm. The broadcast topology is constructed by solving a Hamiltonian path using a network planner (*BPLAN*). After establishing $G_H$, each satellite $i$ computes the Frobenius norm ($f_{ijk}$) of the post-measurement covariance, assuming $i$ obtains measurements from visible asset nodes $j$ and $k$. The $f_{ijk}$ values are globally shared over $G_H$, allowing each node $i$ to construct the measurement topology $G_M$ by solving a multi-

matching problem using a measurement planner (*MPLAN*), as described in the study [13]. The multi-matching solver determines the best set of ISL measurement pairs to address the limitation of using only two directional antennas in the decentralized processing algorithm.

To address the challenges of managing network systems under the constrained communication resources of the proposed lunar constellation, this study introduces four candidate swarm localization approaches based on CEKF and DEKF algorithms These approaches are designed to support autonomous operation of large-scale satellite swarms, while managing both network complexity and computational load. For swarm localization, we implement the initial CEKF and DEKF models described in [11] for state estimation of LPNT swarm assets, along with the matching algorithm proposed in [13], which selects the best pairs of visible neighbors in each satellite for ISL measurement generation. Building on these foundations, we coordinate the scheduling of two network systems, $G_H$ and $G_M$, to facilitate the filter update cycle and manage data exchange across the swarm in communication-constrained environment.

### A. Redundant CEKF

The CEKF algorithm consists of "prediction" and "update" steps for onboard state estimation. During the filter prediction step, the filter propagates the state ($\bar{x}$) and covariance ($\bar{P}$) using a two-body dynamic model. In the update step, the state ($\hat{x}$) and covariance ($\hat{P}$) are corrected using available ISL measurements. The CEKF processes $x$ and $P$ from all swarm assets simultaneously in a single processor, which is a common architecture for current autonomous satellite systems [25-29]. By leveraging the complete global state and covariance information of 100 satellites, the CEKF provides highly accurate state estimates for precise swarm localization. Its rapid convergence and reliable performance in nonlinear environments allow it to effectively manage complex inter-satellite dynamics.

In the *redundant* CEKF approach, each satellite maintains a full global copy of the filter, independently estimating the $x$ and $P$ for 100 satellites. This design ensures consistent filter behavior across the swarm, as each node runs an identical algorithm with access to the same visibility and measurement information. As long as communication links remain intact, all satellites generate the same measurement network by solving the identical matching problem. This redundancy minimizes network variability and reduces the risk of inconsistent updates. However, maintaining global filter state onboard each node imposes memory and computational demands, which challenges small-satellite capabilities on the Moon [15]. The full algorithmic flow of the *redundant* CEKF approach is outlined in Algorithm 1.

**Approach 1:** Redundant CEKF

**Given**: $x_0,\ P_0,\ R,\ Q$
**While** $T \geq 1$ **do**
  **Case** T — ◆Step 1: Network Topology Setup and Data Exchange
    **Switch** 1:
      Compute $f$ locally
      $G_M = MPLAN(f)$
      Slew antennas to $G_M$
    **Switch** 2:
      Join $G_M$
      Perform $z_i$
    **Switch** 3:
      $G_H = BPLAN(x)$
      Slew antennas to $G_H$
    **Switch** 4:
      Join $G_H$
    **Switch** 5:
      Global broadcast of $x, P, z_a, z_i$
  **return**
  **if** $T \leq 5$ **then** — ◆Step 2: Filter Update
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_a)$
    $T \leftarrow T + 1$
  **else**
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_i, z_a)$
    $T \leftarrow 0$
    **return** $\hat{x}_T, \hat{P}_T$
  **end if**
  $[\bar{x}_{T+1}, \bar{P}_{T+1}] = Prediction(\hat{x}_T, \hat{P}_T)$ — ◆Step 3: Filter Prediction
**return**

In this approach, the configuration of the measurement network topology is preceded by the broadcast network setup, which leverages the Frobenius norm ($f$) values computed locally by each satellite. Satellite visibility and the corresponding broadcast topology can be precomputed offline based on orbital dynamics or, alternatively, generated in real time through an onboard Hamiltonian path solver. While the current simulation framework does not explicitly incorporate a Hamiltonian solver within the data exchange process, it implements an optimized connection scheme that is functionally equivalent to the solver's output. Consequently, the simulation represents the performance expected from real-time onboard reconstruction. The update cycle occurs every 5 minutes, with a one-minute time interval, $T$, where each satellite updates its anchor node measurement individually. The broadcast and measurement network setup requires up to 2 minutes for slewing and joining the network, given an antenna slew rate of 3° per second [29]. An additional minute is allocated for data transmission to exchange $x$, $P$, $z_a$, and $z_i$ between two visible neighbor satellites, thus, the minimum frequency for filter updates is 5 minutes. Once the network system is established and swarm communication is enabled, each node performs a full CEKF update using its onboard processor. The predicted states and covariances are then propagated forward from $T$ to $T + 1$ and used in the next cycle.

### B. Leader-Follower CEKF

To alleviate the high computational and memory burden imposed by the *redundant* CEKF approach, the *Leader-Follower* CEKF approach introduces a hierarchical structure within the swarm. In this scheme, a single satellite is designated as the 'leader' node, which is responsible for executing the full centralized filter. The remaining satellites operate as 'follower' nodes, receiving and applying the leader's estimation results without conducting any independent filtering operations.

During each update cycle, the leader satellite performs the full CEKF algorithm, which includes the prediction of global state and covariance using a two-body orbital model, the integration of anchor node measurements, the construction of the broadcast network topology, and the execution of the multi-matching algorithm to form the measurement network. Based on these network structures and ISL measurements, the leader computes updated state estimates ($x^+$) and covariances ($P^+$) for all nodes in the swarm. Once the filtering process is complete, the leader broadcasts the updated state and covariance information to all follower nodes via the broadcast network. Each follower stores the received estimates and uses them to initialize its own prediction step for the next cycle. This architecture significantly reduces the onboard processing and memory requirements for the majority of the swarm, as only one satellite performs the computationally intensive update step. The remaining nodes are relieved of storing and manipulating large state vectors and covariance matrices, which can otherwise exceed the available resources on small-satellite platforms. However, the *Leader-Follower* architecture introduces a critical single point of failure. If the leader satellite experiences a communication loss, hardware fault, or isolation from the rest of the swarm satellites, the entire filtering process can be disrupted. To mitigate this risk, the leader role may be dynamically reassigned based on swarm topology or fault detection logic. For example, if a leader fails to broadcast within a predefined time window, a backup node can initiate a re-election process based on link stability, remaining power, or centrality within the network graph.

In nominal operation, the *Leader-Follower* CEKF follows the same scheduling framework as the *redundant* CEKF. Each cycle begins with the reception of anchor measurements, followed by antenna slewing and network setup for $G_H$ and $G_M$. However, this approach requires a global broadcast of $G_M$ to all nodes, which introduces an additional 3 minutes of network establishment and data transmission time compared to the *redundant* CEKF. As a result, the minimum feasible update frequency in the *Leader-Follower* CEKF architecture becomes 8 minutes. Once the network is established, the leader performs a centralized update and disseminates the results to all followers. The estimation fidelity of the centralized filter is preserved, assuming the leader maintains full visibility of required measurements and the swarm remains well-connected.

One of the key benefits of the *Leader-Follower* approach is that it balances estimation performance with hardware efficiency. It achieves centralized filter accuracy while imposing minimal resource requirements on most satellites. This makes the architecture well-suited for SWaP-constrained platforms, where allocating memory and processor cycles to a full EKF may not be feasible. On the other hand, the approach's dependence on a single node for filter execution introduces inherent vulnerabilities. In addition to the failure risks, this approach places a higher bandwidth demand on the broadcast network compared to the *redundant* CEKF. After the leader performs the filter update, the updated$x$ and $P$ for all satellites must be shared via $G_H$. Moreover, follower satellites cannot generate Frobenius norm metrics ($f_{ijk}$) or solve the $MPLAN$ or multi-matching problem, as they do not maintain the full state and covariance set. Consequently, the system becomes increasingly dependent on the leader's communication reliability and the timely broadcast of the network topology. This dependency increases network variability and amplifies the communication burden, particularly under dynamic or degraded link conditions Robust implementation of the *Leader-*

*Follower* CEKF requires careful design of failover strategies, leader election protocols, and communication prioritization to ensure that state updates consistently reach all follower nodes. Without these measures, the architecture may suffer from partial estimation failures or delayed updates that degrade overall performance.

In summary, the *Leader-Follower* CEKF approach offers a resource-efficient alternative to the *redundant* CEKF, preserving the estimation accuracy of centralized filtering while reducing the system-wide processing load. However, this efficiency comes at the cost of reduced redundancy and increased dependence on a single asset, requiring additional design considerations for fault tolerance and communication resilience in large-scale lunar constellation. The update cycle is illustrated in Algorithm 2, operating on a fixed 8-minute schedule.

**Approach 2:** Leader-Follower CEKF (Leader)

**Given**: $x_0,\ P_0,\ R,\ Q$
**While** $T \geq 1$ **do**
  **Case** T
    **Switch** 1:
      Compute $f$ on the Leader
      $G_M = MPLAN(f)$
      $G_H = BPLAN(x)$
      Slew antennas to $G_H$
    **Switch** 2:
      Join $G_H$
    **Switch** 3:
      Global broadcast $G_M$
    **Switch** 4:
      Slew antennas to $G_M$
    **Switch** 5:
      Join $G_M$
      Perform $z_i$
    **Switch** 6:
      Slew antennas to $G_H$
    **Switch** 7:
      Join $G_H$
    **Switch** 8:
      Global broadcast of $x, P, z_a, z_i$
  **return**
  **if** $T \leq 8$ **then** ◆Step 2: Filter Update
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_a)$
    $T \leftarrow T + 1$
  **else**
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_i, z_a)$
    $T \leftarrow 0$
    **return** $\hat{x}_T, \hat{P}_T$
  **end if**
  $[\bar{x}_{T+1}, \bar{P}_{T+1}] = Prediction(\hat{x}_T, \hat{P}_T)$ ◆Step 3: Filter Prediction
**return**

## C. Distributed CEKF

The *distributed* CEKF is designed to reduce the communication and computational load of LPNT swarm localization while preserving key advantages of centralized estimation. Unlike the *redundant* or *Leader-Follower* CEKF methods, the *distributed* CEKF assigns filtering tasks among multiple satellites, spreading the computational load across the swarm. In this decentralized architecture, each satellite estimates only its own state, selectively incorporating shared state and covariance information from neighboring satellites. This approach removes the need for each node to store and process the global state vector and covariance matrix, thereby significantly reducing computational demands. Each satellite executes a local CEKF update using its own state, ISL measurements with neighboring satellites, and anchor node measurements when available. Before performing the local update, each satellite requires knowledge of neighbors' filtered states and covariances. This information is exchanged via the additional broadcast network setup and treated as fixed input data, allowing satellites to account for their neighbors' estimation uncertainties without explicitly computing cross-correlations. The selection of ISL measurements for each

node's local update is performed using a global multi-matching algorithm. Each satellite computes and broadcasts the Frobenius norm of its post-update covariance. Using these values, the matching planner ($MPLAN$) generates the measurement network, ensuring compliance with the two-antenna constraint. After matching, satellites clearly know their measurement responsibilities and data exchange requirements.

The detailed filtering process is illustrated in Algorithm 3. Each filter update cycle begins with anchor node measurements, followed by the state prediction step. Next, $G_H$ is established, during which satellites exchange their filtered states and covariances. Subsequently, the swarm initiates the multi-matching procedure by exchanging computed Frobenius norms through $G_H$. Based on these values, satellites individually execute the $MPLAN$ algorithm onboard to construct the measurement network, identifying optimal ISL measurement pairs. Each satellite then performs its local CEKF update, restricted to measurements involving itself. Similar to the *redundant* CEKF, this distributed method employs centralized filtering logic onboard every satellite. However, the distributed computation of $f_{ijk}$ necessitates the additional establishment of $G_H$ within the 8-minute update cycle, introducing approximately 3 extra minutes for antenna slewing, network joining, and data exchange. Despite this added complexity, the distributed approach significantly mitigates computational demands by addressing the scaling challenges of centralized multi-matching algorithms.

Unlike Algorithms 1 and 2, Algorithm 3 performs $MPLAN$ immediately before antenna slewing, constructing $G_M$ based on the globally exchanged Frobenius norms. Although each satellite independently solves $MPLAN$, the shared information ensures identical matching solutions across all satellites. By distributing computational tasks, complexity scales approximately as $O(n^2)$ rather than the centralized $O(n^3)$, substantially reducing the processing burden and enhancing scalability. Nevertheless, this computational advantage must be balanced against the additional complexity and communication burden of establishing $G_H$, introducing trade-offs in operational complexity and reliability for practical lunar constellations.

**Approach 3:** Distributed CEKF

**Given**: $x_0,\ P_0,\ R,\ Q$
**While** $T \geq 1$ **do**
  **Case** T — ◆Step 1: Network Topology Setup and Data Exchange
    **Switch** 1:
      Compute $f$ globally
      $G_H = BPLAN(x)$
      Slew antennas to $G_H$
    **Switch** 2:
      Join $G_H$
    **Switch** 3:
      Global broadcast $f$
    **Switch** 4:
      $G_M = MPLAN(f)$
      Slew antennas to $G_M$
    **Switch** 5:
      Join $G_M$
      Perform $z_i$
    **Switch** 6:
      Slew antennas to $G_H$
    **Switch** 7:
      Join $G_H$
    **Switch** 8:
      Global broadcast of $x, P, z_a, z_i$
  **return**
  **if** $T \leq 8$ **then** — ◆Step 2: Filter Update
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_a)$
    $T \leftarrow T + 1$
  **else**
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_i, z_a)$
    $T \leftarrow 0$
    **return** $\hat{x}_T, \hat{P}_T$
  **end if**
  $[\bar{x}_{T+1}, \bar{P}_{T+1}] = Prediction(\hat{x}_T, \hat{P}_T)$ — ◆Step 3: Filter Prediction
**return**

### D. DEKF

The DEKF approach fully decentralizes the filtering process, further minimizing computational and communication burdens compared to CEKF-based methods. In contrast to the *redundant*, *Leader-Follower*, or *distributed* CEKF approaches, DEKF eliminates the need for global state exchange, centralized multi-matching algorithms, and their associated network topologies. Each satellite independently estimates its own state using local ISL measurements and anchor node observations, simplifying onboard processing and dramatically reducing network overhead.

The preliminary design of the DEKF is outlined in [11], where it is assumed that each satellite incorporates information from all visible neighbors into its local EKF update. The method distributes computation across the swarm, reducing the overall burden while preserving estimation accuracy in orbital ephemeris and clock offsets comparable to centralized methods [12]. Notably, the DEKF does not require the exchange of raw ISL or anchor node measurements ($z_{ij}, z_a$), as it is exclusively based on local state updates. Once states and covariances are shared, each satellite computes the measurement sensitivity matrix ($H$), Kalman gain ($K$), and subsequently $f$ using globally consistent $x$ and $P$ values. Satellites then solve the $MPLAN$ algorithm independently to determine identical ISL measurement plans, ensuring coordinated measurement selection while respecting the two-antenna constraint.

This process is described in Algorithm 4, which operates on a 6-minute cycle: 3 minutes for broadcast network setup and state exchange, 2 minutes for measurement network configuration, and 1 minute for Frobenius norm exchange. Since $f_{ijk}$ is exchanged immediately after $x$ and $P$, the broadcast network is reused without re-establishment, saving up to 2 minutes compared to Algorithms 2 and 3, which require additional network setup for $f_{ijk}$ exchange. The DEKF greatly reduces communication complexity and system demands. By removing globally coordinated multi-matching and centralized $MPLAN$ execution, it avoids broadcasting of full $G_M$ each cycle. This yields lower per-node computational complexity approximately $O(n)$ as each satellite processes only local measurements. Without reliance on centralized filtering, the system can also tolerate communication losses and node failures. Each satellite operates autonomously, therefore, local disruptions minimally affect overall performance. However, without global state awareness or coordinated scheduling, satellites may form redundant or incomplete ISL pairings, leading to degraded localization accuracy.

**Approach 4**: DEKF

**Given**: $x_0,\ P_0,\ R,\ Q$
**While** $T \geq 1$ **do**
  **Case** T — ◆Step 1: Network Topology Setup and Data Exchange
    **Switch** 1:
      $G_H = BPLAN(x)$
      Slew antennas to $G_H$
    **Switch** 2:
      Join $G_H$
    **Switch** 3:
      Global broadcast $x$ and $P$
      Compute $H, K, f$ globally
    **Switch** 4:
      Global broadcast $f$
      $G_M = MPLAN(f)$
    **Switch** 5:
      Slew antennas to $G_M$
    **Switch** 6:
      Join $G_M$
      Perform $z_i$
  **return**
  **if** $T \leq 6$ **then** — ◆Step 2: Filter Update
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_a)$
    $T \leftarrow T + 1$
  **else**
    $[\hat{x}_T, \hat{P}_T] = Update(\bar{x}_T, \bar{P}_T, z_i, z_a)$
    $T \leftarrow 0$
    **return** $\hat{x}_T, \hat{P}_T$
  **end if**
  $[\bar{x}_{T+1}, \bar{P}_{T+1}] = Prediction(\hat{x}_T, \hat{P}_T)$ — ◆Step 3: Filter Prediction
**return**

In summary, DEKF provides a scalable, lightweight, and robust solution for resource-limited lunar swarms. It minimizes computational and communication loads while supporting autonomous localization. Still, its reduced accuracy and lack of coordinated estimation may limit suitability for precision navigation. In such cases, CEKF-based methods may be preferred, while DEKF remains attractive for failure-tolerant, bandwidth-constrained missions.

## V. Numerical Assessments

A numerical evaluation of the low-altitude LPNT system is performed using a MATLAB-based simulation of a 100-satellite constellation. This study investigates the localization performance of low-altitude swarm, focusing on both on-orbit asset positioning and Lunar surface user localization, primarily targeting the South Pole region. For user evaluation, a single ground user located at 84.9°S, 137.5°E near the South Pole is considered. In addition to localization performance, we examine the computational efficiency of the four proposed filtering methods in managing a large-scale Lunar constellation under constrained communication resources. Furthermore, we analyze the impact of limited two-way ISL measurement availability on constellation-wide estimation, exploring the scalability of NASA's DSA technologies under increasing network complexity. Finally, we assess the resilience of the 100-node swarm to network degradation, evaluating how various levels of communication failure affect the feasibility of an ad-hoc LPNT system. This study provides reference results for the low-altitude LPNT concept, validating its technical feasibility, algorithmic methodologies, and overall performance, while also evaluating the practical fidelity and robustness of the system.

### A. Simulation Setup

Figure 2 illustrates the overall simulation architecture of the ad-hoc LPNT system. The General Mission Analysis Tool (GMAT) is used to generate the ground-truth ephemerides of all 100 satellites, incorporating initial conditions such as orbital parameters, spacecraft mass, and perturbation effects. GMAT outputs are expressed in the Moon International Celestial Reference Frame (ICRF) inertial coordinate system, and these ephemerides serve as the reference for localization analysis. Pseudorange and relative velocity measurements are synthesized according to the formulations in [11], covering both asset-to-asset and asset-to-anchor links. Gaussian noise is added to each measurement using satellite positions derived from GMAT. To minimize signal reflection and interference from the Moon's irregular terrain [30], a 10° elevation mask is applied to all asset-to-anchor links. Localization of the satellite swarm is performed once every minute over a 7-day simulation period using all four candidate estimation methods. A simultaneous user localization task is also executed every minute, employing a weighted least-squares estimator that integrates the uncertainty of each visible PNT asset. A convergence threshold of 0.01 m is applied for user position estimation.

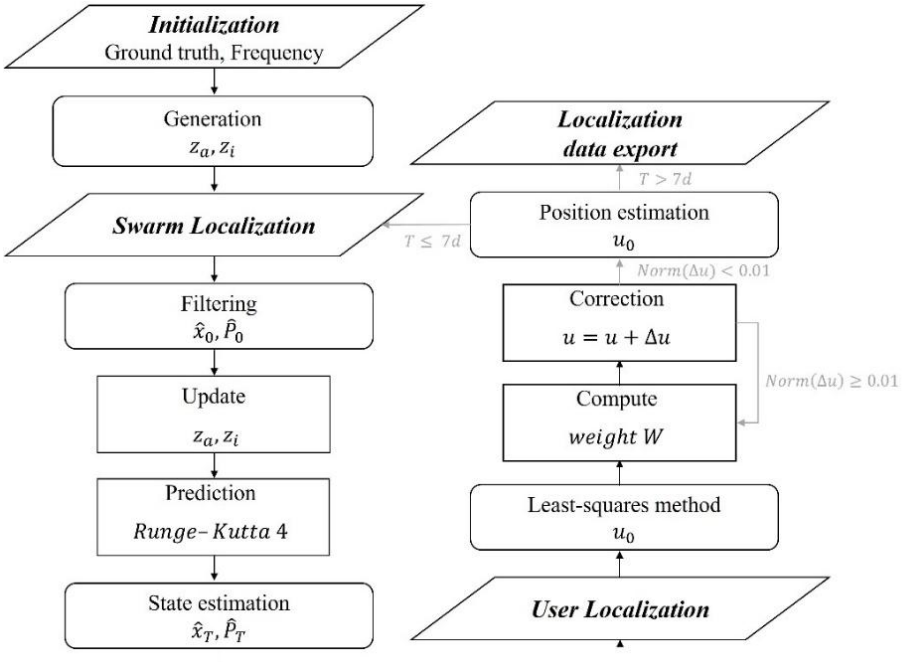



**Fig. 2 Ad-hoc Lunar PNT simulation flowchart**

ISL measurements are assumed to be two-way line-of-sight (LOS), enabling the resolution of relative clock biases by averaging round-trip timing, which can resolve relative clock bias by averaging the travel time between nodes to within tens of nanoseconds. Measurement noise is modeled as zero-mean Gaussian, accounting for position uncertainty, clock bias, and timing jitter. Pseudorange and pseudorange rate accuracies are set to 10 m and 0.1 m/s, respectively. Jitter noise is excluded due to its negligible contribution relative to other noise sources.

Ground anchor nodes are assumed to have no clock errors or positional uncertainty, reducing the covariance of pseudorange measurements between asset and anchor nodes by half compared to asset-to-asset measurements. ISL measurements are modeled in the Lunar-centered inertial frame when generating both asset-to-asset and asset-to-anchor observations. While signal path loss is a common concern in near-Earth GNSS environments, the error model used here focuses on signal attenuation from internal and external interference, since Lunar constellations are unaffected by atmospheric delay due to the Moon's near-vacuum atmosphere, which contains fewer than $10^6$ molecules per cubic centimeter. In the user measurement model, both clock bias and measurement noise terms are included. A weight matrix is applied to enhance user positioning accuracy by adjusting the influence of each asset's measurement according to its associated uncertainty. A recursive least-squares update is performed until convergence of the correction term is achieved, ultimately estimating the position of the South Pole user.

### B. Localization Performance

Previous studies on ad-hoc LPNT systems [10, 12, 13] have evaluated the localization performance of frozen-orbit constellations with a 10-minute filter update cycle, reporting a median swarm position error ($M_S$) of 110.4 m and a median user position error ($M_G$) of 7.783 m. While this study adopts the same process noise model as in previous work, accounting for unmodeled dynamic acceleration with an empirically tuned value of 0.01, we assume a more conservative measurement covariance matrix, setting the pseudorange uncertainty to 200,000 m and the pseudorange rate uncertainty to 20 m/s. This conservative approach accounts for the limited anchor visibility in the low-altitude constellation, which poses additional challenges for maintaining accurate swarm localization.

We analyze the localization performance of both the swarm assets and a representative Lunar surface user, assuming full availability of satellite links and measurements. Because user positioning accuracy strongly depends on the fidelity of swarm localization, maintaining precise asset-level estimates is critical. By evaluating the localization accuracy of the low-altitude swarm, we aim to demonstrate the scalability of DSA technologies under large-constellation operational conditions. Four candidate localization approaches are numerically assessed, each based on either centralized or distributed Kalman filtering and characterized by distinct update cycles: 5 minutes for Approach 1, 8 minutes for Approaches 2 and 3, and 6 minutes for Approach 4, as shown in Figure 3.

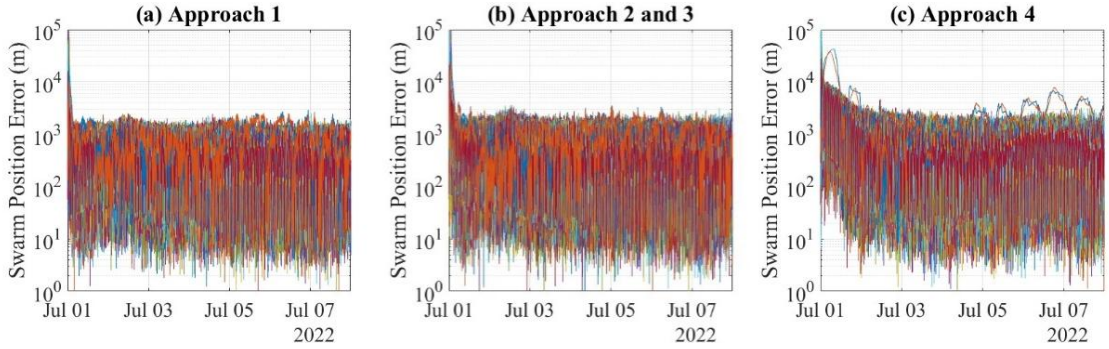



**Fig. 3 Localization performance of low-altitude constellation with four localization approaches**

Approach 1 proves to be the most effective localization method, leveraging centralized data processing and the highest filter update frequency, resulting in $M_S$ of 371.2 m. Approaches 2 and 3 demonstrate slightly reduced performance, with $M_S$ of 419.6 m, primarily due to the extended filter update interval caused by the additional broadcast of $G_M$. Approach 4 yields the lowest localization accuracy, with an $M_S$ of 541.7 m. In this case, localization performance slightly varies with orbital configuration, where high-altitude satellites exhibit larger estimation errors than their low-altitude counterparts. This outcome highlights a key limitation of the DEKF algorithm, which uses a reduced covariance matrix and cannot fully capture inter-satellite correlations driven by differing orbital dynamics.

Figure 4 presents the user localization performance of Approach 1, depicting the PNT service availability across the entire lunar surface over a 7-day simulation. An availability value of 0 indicates no service availability during the

simulation period, while a value of 1 represents continuous availability. As illustrated, the low-altitude constellation provides effective navigation support primarily for users located near the lunar South Pole. Although the frozen-orbit constellation is also designed primarily for South Pole coverage, its service area is even more constrained due to the limited ground view at an altitude of 100 km. Our analysis includes position error and the number of satellites in view at the South Pole, both of which vary significantly due to the limited ground visibility of low-altitude constellation. Over a 7-day period, Approach 1 achieves $M_G$ of 16.69 m, with a median availability of 0.7312. Similarly, Approaches 2 and 3 yield $M_G$ values of 16.88, with LPNT availability at 0.7311. Approach 4 performs the worst among the four approaches, yielding a $M_G$ of 17.16 and availability of 0.7275, consistent with its reduced estimation performance.

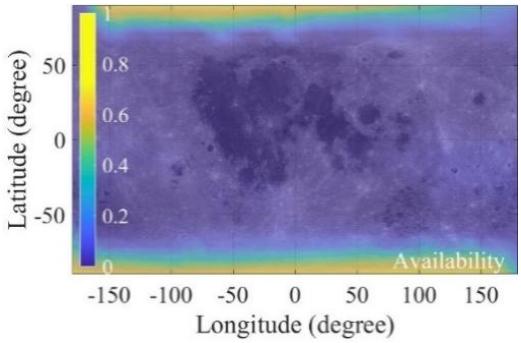



**Fig. 4 User localization performance of Approach 1**

### C. Computational Performance Analysis

Each of the four candidate localization approaches presents distinct trade-offs in computational efficiency and network performance. To rigorously assess these trade-offs, we evaluate key factors, including communication topology dynamics, bandwidth utilization, computational complexity and simulation time, and memory requirements associated with onboard data storage. We systematically compare these computational performance metrics across the four localization approaches, providing a comprehensive evaluation of their respective advantages and limitations. This investigation establishes a foundational framework for evaluating the performance of each localization approach, enabling upcoming space operations to select the most suitable method based on specific mission requirements and operational constraints. These comparisons are especially critical for future swarm missions that require autonomous operation in communication-constrained or failure-prone environments.

Communication topology dynamics refer to the frequency of network topology variations, quantifying how often the network structure must adapt to support different functions such as broadcasting satellite information or establishing ISL measurement pairs. The network topology requirements for each localization algorithm are systematically assessed by quantifying the distinct configurations necessary for swarm operations. Approaches 1 and 4 employ only two topologies: one dedicated to broadcasting and another for measurements. In contrast, Approaches 2 and 3 necessitate an additional broadcast topology at the beginning of each cycle to enable the exchange of $G_M$ and $f$, thereby introducing further network operational complexity. This added requirement may reduce available communication time for other mission-critical functions such as command, telemetry, or payload data downlink.

Bandwidth consumption is analyzed by assessing the message traffic generated within satellite communication, assuming the use of high-bandwidth directional radios, such as S-band, for data exchange. To quantify bandwidth requirements, we consider both an intermediate node and a terminal node at the network boundary. In a broadcast network consisting of 100 nodes, intermediate node 50 represents a midpoint, where nodes 1–49 communicate with nodes 51–100, yielding a total of 2,450 messages. In the worst-case scenario, node 50 must simultaneously relay this maximum volume of traffic. However, this represents a highly conservative estimate, as message propagation does not necessarily converge at node 50 instantaneously. A terminal node with a single direct connection to another node in the swarm is expected to receive transmissions from all other nodes, amounting to 99 messages. Again, assuming simultaneous message arrival constitutes an overestimation of the actual bandwidth demand. This comparison can provide a useful upper bound for designing communication hardware and protocol buffering requirements. In the computation of bandwidth consumption, it is also necessary to account for variations in Frobenius norm traffic based on satellite visibility in low-altitude and high-altitude constellations. A high-altitude configuration provides a substantially greater number of satellites in view, 65 nodes, compared to 30 in a low-altitude configuration. The number of messages handled by the intermediate node is divided into 2,352 and 98 for low- and high-altitude configurations, respectively, while terminal nodes exhibit message counts of 97 and 2 under analogous conditions. This sensitivity to satellite visibility highlights the need for dynamic scheduling strategies that adapt to orbital

geometry. A comprehensive summary of bandwidth utilization across the broadcast network topology is provided in Table 3.

**Table. 3 Bandwidth utilization across the broadcast network topology**

| Type | $G_H$ | Data | Size | # | *Intermediate* BW [Bps] | *Terminal* BW [Bps] |
|---|---|---|---|---|---|---|
| Approach 1 (Redundant CEKF) | 1 | Pairs | 1 byte [unit8] | 4 [Matching pair] | 163.3 | 6.6 |
| | | Pseudorange | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Pseudorange rate | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Anchor Pairs | 1 byte [unit8] | 2 [anchors] | 81.67 | 3.3 |
| | | Anchor Pseudorange | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | Anchor Pseudorange rate | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | State | 8 byte [double] | 6 [components] | 1960 | 79.2 |
| | | Covariance | 8 byte [double] | 36 [components] | 11760 | 475.2 |
| | | | | Total | 16578.3 | 669.9 |
| Approach 2 (Leader-Follower CEKF) | 1 | Matching solution | 1 byte [unit8] | 4 [Matching pair] | 163.3 | 6.6 |
| | 2 | Pairs | 1 byte [unit8] | 4 [Matching pair] | 163.3 | 6.6 |
| | | Pseudorange | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Pseudorange rate | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Anchor Pairs | 1 byte [unit8] | 2 [anchors] | 81.67 | 3.3 |
| | | Anchor Pseudorange | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | Anchor Pseudorange rate | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | State | 8 byte [double] | 6 [components] | 1960 | 79.2 |
| | | Covariance | 8 byte [double] | 36 [components] | 11760 | 475.2 |
| | | | | Total | 16741.7 | 676.5 |
| Approach 3 (Distributed CEKF) | 1 | Frobenius Pairs in low-altitude | 1 byte [unit8] | 1800 [Frobenius pair] | 70560 | 2910 |
| | | Frobenius norm in low-altitude | 8 byte [double] | 900 [Frobenius norm] | 282240 | 11640 |
| | | Frobenius Pairs in high-altitude | 1 byte [unit8] | 8450 [Frobenius pair] | 13801.7 | 281.7 |
| | | Frobenius norm in high-altitude | 8 byte [double] | 4225 [Frobenius norm] | 55206.7 | 1126.7 |
| | 2 | Pairs | 1 byte [unit8] | 4 [Matching pair] | 163.3 | 6.6 |
| | | Pseudorange | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Pseudorange rate | 8 byte [double] | 2 [antennas] | 653.3 | 26.4 |
| | | Anchor Pairs | 1 byte [unit8] | 2 [anchors] | 81.67 | 3.3 |
| | | Anchor Pseudorange | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | Anchor Pseudorange rate | 8 byte [double] | 2 [anchors] | 653.3 | 26.4 |
| | | State | 8 byte [double] | 6 [components] | 1960 | 79.2 |
| | | Covariance | 8 byte [double] | 36 [components] | 11760 | 475.2 |
| | | | | Total | 438386.7 | 16628.2 |
| Approach 4 (DEKF) | 1 | State | 8 byte [double] | 6 [components] | 1960 | 79.2 |
| | | Covariance | 8 byte [double] | 36 [components] | 11760 | 475.2 |
| | | Frobenius Pairs in low-altitude | 1 byte [unit8] | 1800 [Frobenius pair] | 70560 | 2910 |
| | | Frobenius norm in low-altitude | 8 byte [double] | 900 [Frobenius norm] | 282240 | 11640 |
| | | Frobenius Pairs in high-altitude | 1 byte [unit8] | 8450 [Frobenius pair] | 13801.7 | 281.7 |
| | | Frobenius norm in high-altitude | 8 byte [double] | 4225 [Frobenius norm] | 55206.7 | 1126.7 |
| | | | | Total | 435528.4 | 16512.8 |

Computational complexity and simulation time of each approach are systematically evaluated using a MATLAB-based simulator, accounting for the time required for filter updates, state prediction, and multi-matching problem-solving, as summarized in Table 4. Here, $n$ denotes the number of satellite nodes available in the LPNT constellation, and $m$ represents the number of total measurements processed per timestep. In Approach 2, only the leader node performs the full $MPLAN$ computation, while follower nodes bypass this step, resulting in significantly reduced computational load for the majority of the swarm. Approach 4 applies an increased measurement covariance matrix technique, allowing each satellite to estimate its own state while accounting for the full uncertainty associated with its predicted measurements [11]. This methodology results in increased computational overhead, leading to longer processing times in Approach 4, surpassing Approach 3 based on the CEKF method. Despite the higher computational burden in MPLAN, the Update and Prediction steps in Approach 4 remain faster than those of all other approaches. These distinctions are essential when selecting algorithms for onboard flight processors with strict timing budgets.

**Table. 4 Computational Complexity and Simulation Time**

| | Computational Complexity | | | Simulation Time | | |
|---|---|---|---|---|---|---|
| *Functions* | *MPLAN* | *Update* | *Prediction* | *MPLAN* | *Update* | *Prediction* |
| Approach 1 (Redundant CEKF) | $O(n^3)$ | $O(n^4m^4)$ | $O(n^2)$ | 0.2288312 s | 0.02089370 s | 0.01113146 s |
| Approach 2 (Leader-Follower CEKF) | $O(n^3)$ | $O(n^4m^4)$ | $O(n^2)$ | 0.2288312 s | 0.02089370 s | 0.01113146 s |
| Approach 3 (Distributed CEKF) | $O(n^2)$ | $O(n^4m^4)$ | $O(n^2)$ | 0.002288037 s | 0.02089370 s | 0.01113146 s |
| Approach 4 (DEKF) | $O(2n^2)$ | $O(m^4)$ | $O(1)$ | 0.02149823 s | 7.356813e-5 s | 0.01112692 s |

Onboard memory constraints are also analyzed to assess the data storage requirements for both individual satellites and the entire swarm. Across the four localization approaches, four distinct types of data are essential to support their respective estimation algorithms: global broadcast network topology, states and covariance, $f$, and measurements. The global broadcast network topology, representing satellite connectivity, remains constant at 100 bytes across all approaches, with Approach 2 incorporating the leader's role in computation. The storage requirements for states and covariances, and measurements, are identical in Approaches 1, 2, and 3: 4,800 bytes for six state components, 28,800 bytes for 36 covariance components, and 3,200 bytes for up to four measurements, including ISL and anchor node observations. In contrast, Approach 4 significantly reduces storage requirements, utilizing only 48 bytes for six state components, 288 bytes for 36 covariance components, and 32 bytes for measurements, as it processes only a single satellite's information within the filter. Additionally, the storage demand for f in Approaches 1 and 2 amounts to 96,650 bytes, corresponding to Frobenius norm sizes of 30 for 98 satellites in a low-altitude and 65 for 2 satellites in a high-altitude constellation. In contrast, Approaches 3 and 4 require only a single satellite's Frobenius norm, with a size of 65, leading to a total memory footprint of 4,225 bytes. Minimizing the memory footprint of $f$ directly reduces contention over onboard buffers during time-sensitive updates. Table 5 presents a comprehensive evaluation of the computational performance of the four localization approaches, with red indicating the worst-performing case and blue denoting the best.

**Table. 5 Evaluation of the computational performance**

| | Approach 1 (Redundant CEKF) | Approach 2 (Leader-Follower CEKF) | Approach 3 (Distributed CEKF) | Approach 4 (DEKF) |
|---|---|---|---|---|
| Communication Topologies | 2 | 3 | 3 | 2 |
| Minimum Update Cycle | 5 minutes | 8 minutes | 8 minutes | 6 minutes |
| Bandwidth (Intermediate) | 16.58 KBps | 16.74 KBps | 438.4 KBps | 435.5 KBps |
| Bandwidth (Terminal) | 0.6699 KBps | 0.6765 KBps | 16.63 KBps | 165.1 KBps |
| Localization Complexity | $n^3 + n^4m^4$ | $n^3 + n^4m^4$ | $2n^2 + n^4m^4$ | $n^2 + m^4$ |
| Simulation Time | 0.2608564 s | (Leader) 0.2608564 s<br>(Follower) 0.0320206 s | 0.03431320 s | 0.03269872 s |
| Memory | 134.0 Kbyte | (Leader) 134.0 Kbyte<br>(Follower) 0.472 Kbyte | 41.53 Kbyte | 4.697 Kbyte |
| Onboard information | $G_H$, $x_i$, $P_i$, $f_i$, $Z_i$<br>$(All;\ i = 1 - 100)$ | $G_H$, $x_i$, $P_i$, $f_i$, $Z_i$<br>$(Leader; i = 1 - 100)$ | $G_H$, $x_i$, $P_i$, $f_i$, $Z_i$<br>$(All;\ i = 1 - 100,\ l = 1)$ | $G_H$, $x_l$, $P_l$, $f_l$, $Z_l$<br>$(All;\ i = 1 - 100,\ l = 1)$ |

From the perspective of communication topology dynamics and the minimum frequency of filter updates, Approaches 1 and 4 demonstrate superior efficiency. Approach 1, in particular, exhibits superior bandwidth performance by eliminating the need to transmit matching solutions or $f$ values, in contrast to Approach 3, which incurs the greatest bandwidth overhead. From a computational standpoint, Approach 4 demonstrates the lowest simulation time, as each node independently computes its own $f$ value. This operation requires evaluating only up to $n^2$ edge pairs in the worst-case scenario, similar to Approach 3. On the other hand, Approaches 1 and 2 require significantly higher computations, reaching $n^3$ complexity in the worst case, as each node must evaluate $n^2$ edge pairs.

Regarding memory usage, Approaches 1, 2, and 3 retain the full CEKF matrices while scaling with $n$ for intermediate $f$ computations, while Approach 4 is based on reduced matrices in the DEKF algorithm. This analysis provides valuable insight into the trade-offs across network dynamics, bandwidth, processing time, and storage efficiency, offering a practical framework for selecting the optimal approach under mission-specific constraints.

### D. Failure Test

We evaluate the impact of communication failures on a low-altitude Lunar PNT constellation, considering the practical constraint that each satellite is equipped with only two directional antennas. This hardware limitation introduces vulnerabilities to network degradation, which can significantly impair the performance of distributed state estimation. To systematically characterize these vulnerabilities, we define four distinct failure scenarios: (1) failure in delivering the matching solution ($G_M$), (2) failure in performing ISL measurements ($z_i$), (3) failure in exchanging Frobenius norm values ($f_{ijk}$), and (4) failure in communicating the state and covariance ($x, P$). This section investigates the fault sensitivity and resilience of all four localization approaches through simulation experiments. Each failure scenario is evaluated by analyzing the resulting degradation in localization accuracy. The performance degradation is quantified as the cumulative percentage increase in localization error, computed by comparing 7-day simulation results under each failure mode, using only operational nodes with those from an idealized baseline in which all assets maintain full communication reliability. The results provide insight into how architectural choices, such as centralized versus decentralized processing, and *redundant* versus *leader-follower* coordination, affect system robustness under realistic fault conditions. Among the four failure types, Approach 1 is most impacted by ISL measurement and state data transmission failures, which disrupt the centralized update mechanism. In Approach 2, which employs a leader-follower architecture, performance is sensitive to failures in transmitting the matching solution, measurements, and global state updates, as it depends on the leader's ability to disseminate information. Approaches 3 and 4, which utilize distributed filtering, are particularly vulnerable to disruptions in $f_{ijk}$ exchange, since this information is essential for solving the multi-matching problem. Interruptions in $f_{ijk}$ exchange can prevent proper measurement pairing, leading to degraded network connectivity and compromised estimation quality.

To evaluate these effects in simulation, we group the four failure types into two representative scenarios: (1) Frobenius norm and state exchange failures, and (2) matching and measurement exchange failures. Scenario (1) is modeled by removing a complete satellite's state and covariance information from the network prior to the matching process, simulating a persistent failure such as a total node outage or sustained communication dropout. From the start of the simulation, these failures are assumed to persist throughout the full 7-day duration, reflecting their critical role in maintaining inter-satellite correlation. Scenario (2), by contrast, represents shorter-duration communication disruptions in the ISL measurement generation. The relevant measurement observations are removed for 24 to 48 hours prior to the measurement update step, where it can isolate the impact of degraded measurement availability without affecting filter initialization or long-term orbital memory. In each scenario, a fixed set of seven satellites, one from each of the seven orbital planes in the LPNT constellation, is selected to undergo the simulated failures. This ensures consistent cross-plane representation and captures the impact of failures distributed across the network's geometric structure. Through this setup, we aim to analyze not only the fault tolerance of each localization approach, but also the extent to which architecture and redundancy influence resilience to dynamic network disruptions. The simulation outcomes for the two fault scenarios are illustrated in Figures 5 and 6, highlighting the impact of each failure type on localization performance.

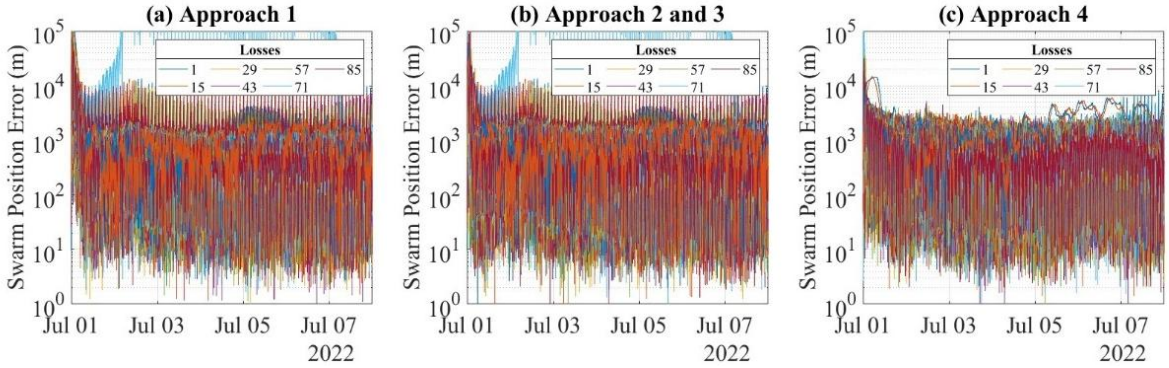



**Fig. 5 Localization performance of low-altitude constellation under *f* and state exchange failures**

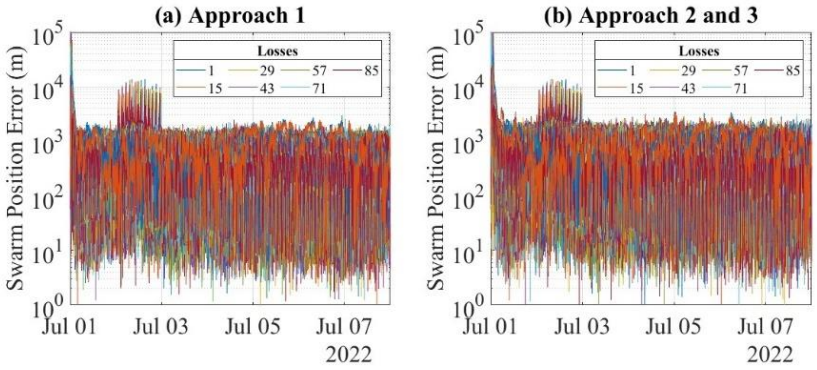



**Fig. 6 Localization performance of low-altitude constellation under matching and measurement failures**

In Approach 1, the $M_s$ increases (slightly) from 371.2 m to 372.7 m under matching and measurement exchange failures, and further rises to 387.8 m when Frobenius norm and state data are lost. Approaches 2 and 3 exhibit greater performance degradation, with $M_s$ values rising from 419.6 m to 421.5 m and 440.1 m, respectively, indicating higher sensitivity to failures in matching, measurement, and state dissemination. Approach 4, which operates independently of matching solutions but remains highly dependent on state and $f$ exchanges, suffers the most substantial impact, yielding an increase in $M_s$ from 541.7 m to 572.3 m. To evaluate the impact of this degradation on user-level accuracy, user positioning error is also assessed for Approach 4 as a representative worst-case. Despite the increase in swarm localization error of 5.5%, the corresponding user localization error $M_G$ increases only modestly, from 17.16 m to 17.39 m (1.3%), indicating that user-level performance remains largely resilient to moderate swarm-level degradation in this architecture.

Table 6 summarizes the degradation in localization performance across all four approaches, quantified as the percentage increase in $M_s$ relative to the idealized communication case. To ensure that the analysis reflects only the impact on the operational network, localization results from nodes undergoing failure are excluded during the failure periods. Specifically, for Scenario (1) which represents persistent Frobenius norm and state exchange failures, the seven affected satellite nodes are entirely removed from the performance evaluation over the full 7-day simulation period. In contrast, for Scenario (2) which represents transient matching and measurement exchange failures for one day, the performance of the identical seven nodes is excluded only during their respective 24- to 48-hour failure intervals within the 7-day window. Each of the four localization approaches employs a different number of required communication topologies, 2, 3, 3, and 2, respectively. Accordingly, a single entire network component failure was introduced for each required topology per approach to represent a consistent failure condition. It is important to note that the degradation percentages are significantly higher in Scenario (1), not only due to the more severe impact of structural data loss on filter operation rendering the failing nodes not available, but also because the failure persists over the full duration in the simulation. In contrast, Scenario (2) has a more limited effect, as the impacted nodes retain valid state information outside of the short-term failure window for one day while failing nodes can still continue contributing to the estimation process during the failure.

**Table. 6 Low-altitude constellations of Lunar PNT system**

| Scenario | Failure Type | Approach 1 (Redundant CEKF) | Approach 2 (Leader-Follower CEKF) | Approach 3 (Distributed CEKF) | Approach 4 (DEKF) |
|---|---|---|---|---|---|
| (1) | $f_{ijk}$ | 0 % | 0 % | 4.88 % | 5.65 % |
| | $x$ and $P$ | 4.47 % | 4.88 % | 4.88 % | 5.65 % |
| (2) | $G_M$ | 0 % | 0.453 % | 0 % | 0 % |
| | $z_i$ | 0.404 % | 0.453 % | 0.453 % | 0 % |

The fault tolerance analysis reveals that Approach 4 exhibits the lowest robustness in the failure scenario (1), primarily due to its susceptibility to state and covariance exchange in the DEKF algorithm. Without these data, DEKF nodes cannot reliably update their own estimates or participate in the global matching process. In contrast, Approach 1 demonstrates the highest resilience to both failure scenarios (1) and (2), due to its use of redundant filtering across

all nodes and simplified, static network topology. This approach benefits from inherent redundancy, where each node independently maintains a full global copy of the state, reducing the impact of any single-point failure. Approaches 2 and 3 are comparatively more vulnerable, not only because of their dependence on shared measurement and state updates, but also due to the frequency with which the global communication topology must be re-established. These frequent topology reconstructions increase the likelihood of encountering data loss or timing mismatches during operation. Moreover, the leader-follower design of Approach 2 introduces a single point of failure, which may compromise the entire network if the leader is isolated or fails to transmit updates. However, in the current simulation setup, the leader node in Approach 2 was deliberately excluded from the set of failing satellites to preserve the functionality of the centralized update process. If the leader were to experience a communication failure in practice, the impact would be catastrophic, as no other node is capable of independently executing the required centralized filter update. This failure analysis highlights the critical role of architectural redundancy and communication simplification in enhancing the robustness of lunar swarm localization systems. The results demonstrate that Approach 1 offers the highest fault tolerance among the four, exhibiting the lowest chance of losing network connectivity and the least degradation in localization performance. In contrast, the remaining approaches show varying degrees of sensitivity to communication failures, depending on their operational conditions. Approaches 2 and 3 show slightly better resilience to failure scenario (1) compared to Approach 4, benefiting from reduced degradation performance. However, they are susceptible to performance degradation under failure scenario (2), unlike Approach 4 which can also benefit from the minimal computational load. These variations in failure response emphasize the importance of selecting architecture types based on mission-specific operational constraints. For long-duration autonomous missions with limited access to ground support, such trade-offs must be carefully considered when balancing system efficiency against resilience.

## VI. Conclusion

This study presents a comprehensive evaluation of a low-altitude LPNT system utilizing a 100-node swarm of small satellites operating at 100 km altitude. By leveraging both CEKF and DEKF algorithms, we explored four candidate localization architectures (*Redundant* CEKF, *Leader-Follower* CEKF, *Distributed* CEKF, and DEKF) to address the technical and operational challenges of decentralized spacecraft autonomy under severe communication and resource constraints. Through high-fidelity simulation, the *redundant* CEKF approach demonstrated the highest localization accuracy, achieving a median satellite position error of 371.2 m and a surface user error of 16.69 m, at the cost of substantial computational and bandwidth requirements. The *Leader-Follower* and *distributed* CEKF models offered more efficient alternatives, preserving estimation fidelity while reducing onboard processing demands. The DEKF approach, although the most resource-efficient, showed higher sensitivity to communication degradation, underscoring the trade-off between robustness and computational simplicity. This work also quantified the computational complexity, memory usage, and communication overhead of each approach, offering a practical framework for selecting appropriate estimation architectures based on mission-specific constraints. Furthermore, fault tolerance analysis revealed the importance of architectural redundancy and network simplification in sustaining reliable performance under real-world communication failures. The results establish the feasibility of leveraging existing science missions to deliver region-specific, ad-hoc LPNT services, particularly for high-priority exploration zones like the lunar South Pole. Future research should explore adaptive matching strategies, autonomous network configuration, and onboard fault detection and recovery mechanisms to enhance the scalability and reliability of decentralized LPNT systems. In addition, the integration of optical crosslink technologies, joint navigation-payload scheduling, and hybrid constellation planning could significantly extend the operational lifespan of small-satellite lunar missions. By addressing the interrelated challenges of accuracy, efficiency, and robustness, this study contributes to the foundational design of decentralized large-scale lunar navigation systems. The findings serve as a guiding basis for the implementation of low-cost, autonomous LPNT services and long-term exploration efforts on the Moon.

## Acknowledgments

The Distributed Spacecraft Autonomy (DSA) project is funded by the NASA Space Technology Mission Directorate's Game Changing Development and Small Spacecraft Technology programs. This study was also supported by Hyundai Motors Chung Mong-Koo Foundation.

## References

[1] Zucherman, A. P., Braun, B. M., and Sims, E. M., "Navigating the Policy Compliance Roadmap for Small Satellites," *Journal of Space Safety Engineering*, Vol. 9, No. 4, 2022, pp. 582–599. doi: 10.1016/j.jsse.2022.07.009.

[2] Gill, T., "NASA's Lunar Orbital Platform-Gateway," *Proceedings of the 45th Space Congress – The Next Great Steps*, Cape Canaveral, FL, Feb. 2018.
[3] Mudgway, D. J., *Uplink-Downlink: A History of the Deep Space Network 1957–1997*, National Aeronautics and Space Administration, Washington, DC, 2001.
[4] Steffes, S. R., Barton, G. H., Bhatt, S. A., Fritz, M. P., King, E. T., and Woffinden, D. C., "Deep Space Autonomous Navigation Options for Future NASA Crewed Missions," *Proceedings of the AAS Guidance, Navigation, and Control Conference*, Breckenridge, CO, Feb. 2018, AAS Paper 18-083.
[5] Collicott, B. C., and Woffinden, D., "Lunar Navigation Performance Using the Deep Space Network and Alternate Solutions to Support Precision Landing," *AIAA SciTech 2021 Forum*, Jan. 2021, AIAA Paper 2021-0375. doi: 10.2514/6.2021-0375.
[6] Cesarone, R. J., Abraham, D. S., and Deutsch, L. J., "Prospects for a Next-Generation Deep-Space Network," *Proceedings of the IEEE*, Vol. 95, No. 10, 2007, pp. 1902–1915. doi: 10.1109/JPROC.2007.905043.
[7] Kim, Y., Kim, P., Ryu, H.-G., and Park, S.-Y., "Performance Analysis of a Pointing, Acquisition, and Tracking System for the VISION Laser Crosslink Mission," *Proceedings of the 38th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, Aug. 2024. [Online]. Available: https://digitalcommons.usu.edu/smallsat/2024/all2024/16/
[8] Witternigg, N., Obertaxer, G., Schönhuber, M., Palmerini, G. B., Rodriguez, F., Capponi, L., Soualle, F., and Floch, J.-J., "Weak GNSS Signal Navigation for Lunar Exploration Missions," *Proceedings of the 28th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2015)*, Tampa, FL, Sept. 2015, pp. 3928–3944.
[9] Capuano, V., Basile, F., Botteron, C., and Farine, P.-A., "GNSS-Based Orbital Filter for Earth-Moon Transfer Orbits," *Journal of Navigation*, Vol. 69, No. 4, 2016, pp. 745–764. doi: 10.1017/S0373463315000843.
[10] Bhamidipati, S., Mina, T., and Gao, G., "Design Considerations of a Lunar Navigation Satellite System with Time-Transfer from Earth-GPS," *Proceedings of the 34th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2021)*, St. Louis, MO, Sept. 2021, pp. 950–965. doi: 10.33012/2021.18021.
[11] Hagenau, B., Peters, B., Burton, R., Hashemi, K., and Cramer, N., "Introducing the Lunar Autonomous PNT System (LAPS) Simulator," *IEEE Aerospace Conference*, Big Sky, MT, Mar. 2021, pp. 1–11. doi: 10.1109/AERO50100.2021.9438473.
[12] Wen, Y., Zhu, J., Gong, Y., Wang, Q., and He, X., "Distributed Orbit Determination for Global Navigation Satellite System with Inter-Satellite Link," *Sensors*, Vol. 19, No. 5, 2019, pp. 1031. doi: 10.3390/s19051031.
[13] Frank, J., Levinson, R., Hillsberg, E., Cramer, N., and Burton, R., "Distributed Scheduling of Position Estimation Updates in Ad-Hoc Lunar Constellations," *AAAI Spring Symposium Series*, Palo Alto, CA, 2022.
[14] Kim, Y., Kempa, B., Adams, C., Levinson, R., and Frank, J., "Localization of Ad-Hoc Lunar Constellations in Communication Failure Modes for Distributed Spacecraft Autonomy," *Proceedings of the 2025 International Technical Meeting of The Institute of Navigation (ION ITM 2025)*, Long Beach, CA, Jan. 2025. doi: 10.33012/2025.20026.
[15] Pedroso, L., and Batista, P., "Distributed Decentralized EKF for Very Large-Scale Networks with Application to Satellite Mega-Constellations Navigation," *Control Engineering Practice*, Vol. 135, 2023, Art. 105509. doi: 10.1016/j.conengprac.2023.105509.
[16] Adams, C., Kempa, B., Vaughan, W., and Cramer, N., "Development of a High-Performance, Heterogeneous, Scalable Test-Bed for Distributed Spacecraft," *IEEE Aerospace Conference*, Big Sky, MT, Mar. 2023, pp. 1–8. doi: 10.1109/AERO55745.2023.10115695.
[17] Hardgrove, C., Lazbin, I., Heffern, L., and Prettyman, T., "Lunar Polar Hydrogen Mapper (LunaH-Map) CubeSat Mission Launch, Early Operations and Lunar Flyby Neutron Data Collection," *54th Lunar and Planetary Science Conference (LPSC)*, The Woodlands, TX, Mar. 2023, Abstract 2960.
[18] Steigerwald, W., "NASA's LRO Moves Closer to the Lunar Surface," NASA Press Release, May 2015. *(Formal publication not available; NASA news release.)*
[19] Cheek, N., Gonzalez, C., Adell, P., Baker, J., Ryan, C., Statham, S., Lightsey, E. G., Smith, C. R., Awald, C., and Ready, J., "Systems Integration and Test of the Lunar Flashlight Spacecraft," *Proceedings of the 36th Annual Small Satellite Conference*, Logan, UT, Aug. 2022, Paper SSC22-II-06.
[20] Hine, B., "LADEE PDS Mission Description," NASA Ames Research Center, Document No. DES-12.LADEE.LPMD, 2014.
[21] Jeon, M.-J., Cho, Y.-H., Kim, E., Kim, D.-G., Song, Y.-J., Hong, S., Bae, J., Bang, J., Yim, J. R., and Kim, D.-K., "Korea Pathfinder Lunar Orbiter (KPLO) Operation: From Design to Initial Results," *Journal of Astronomy and Space Sciences*, Vol. 41, No. 1, 2024, pp. 43–60. doi: 10.5140/JASS.2024.41.1.43.
[22] Folta, D. C., Bosanac, N., Cox, A., and Howell, K. C., "The Lunar IceCube Mission Design: Construction of Feasible Transfer Trajectories with a Constrained Departure," *Proceedings of the 26th AAS/AIAA Space Flight Mechanics Meeting*, Napa, CA, Feb. 2016, AAS Paper 16-285.
[23] Kato, M., Sasaki, S., Takizawa, Y., and the Kaguya Project Team, "The Kaguya Mission Overview," *Space Science Reviews*, Vol. 154, Nos. 1–4, 2010, pp. 3–19. doi: 10.1007/s11214-010-9678-3.
[24] Ehlmann, B. L., Klima, R. L., Seybold, C. C., Klesh, A. T., Au, M. H., Bender, H. A., Bennett, C. L., Blaney, D. L., Bowles, N., Calcutt, S., ..., and Zareh, S., "NASA's Lunar Trailblazer Mission: A Pioneering Small Satellite for Lunar Water and Lunar Geology," *IEEE Aerospace Conference*, Big Sky, MT, Mar. 2022. doi: 10.1109/AERO53065.2022.9843663.

[25] Kassas, Z. M., Morales, J. A., and Khalife, J. A., "New-Age Satellite-Based Navigation—STAN: Simultaneous Tracking and Navigation with LEO Satellite Signals," *Inside GNSS*, Vol. 14, No. 4, 2019, pp. 56–65.
[26] Ren, X., Yang, Y., Zhu, J., and Xu, T., "Comparing Satellite Orbit Determination by Batch Processing and Extended Kalman Filtering Using Inter-Satellite Link Measurements of the Next-Generation BeiDou Satellites," *GPS Solutions*, Vol. 23, No. 1, 2019, pp. 25. doi: 10.1007/s10291-018-0816-9.
[27] Dave, P. K., "Autonomous Navigation of Distributed Spacecraft Using Intersatellite Laser Communications," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2020.
[28] Grenfell, P., "Systems Performance Analysis for Autonomous Spacecraft Navigation Within Satellite Constellations Using Intersatellite Optical Communications Links," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2024.
[29] Sin, E., Arcak, M., Nag, S., Ravindra, V., Li, A., and Levinson, R., "Attitude Trajectory Optimization for Agile Satellites in Autonomous Remote Sensing Constellations," *AIAA SciTech Forum*, Jan. 2021, AIAA Paper 2021-1110. doi: 10.2514/6.2021-1110.
[30] Santra, S., Paet, L. B., Staudinger, E., and Yoshida, K., "Radio Propagation Modelling for Coordination of Lunar Micro-Rovers," *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Oct. 2020, pp. 1–8.