

# Autonomous Visual-Inertial Navigation for Networked Satellites

Frederik Markus  
Robotics Institute  
Carnegie Mellon University  
Pittsburgh, PA 15213  
fmarkus@cmu.edu

Zachary Manchester  
Robotics Institute  
Carnegie Mellon University  
Pittsburgh, PA 15213  
zacm@cmu.edu

**Abstract**— Current orbit-determination methods rely heavily on Earth-based infrastructure like the Deep Space Network (DSN), ground-based radars, or Global Navigation Satellite Systems (GNSS). This paper introduces a joint orbit-and-attitude-determination method using one or more low-cost RGB cameras, angular-velocity data from a gyroscope, and ranging data from other satellites. Machine-vision algorithms are used to identify known landmarks in images captured by the spacecraft. These landmark locations are fused with gyro data and satellite-ranging measurements and combined with a spacecraft dynamics model in an Iterated Extended Kalman Filter (IEKF) that estimates the position, velocity, and attitude of the spacecraft, as well as the gyroscope bias. All measurements and computations are performed onboard the spacecraft, without the use of any exogenous Earth-based inputs. We explore the benefits of extending the number of networked satellites and investigate different network topologies. We validate the performance of the approach in simulation using real satellite imagery.


## TABLE OF CONTENTS

|  |   |
|--|---|
| 1. INTRODUCTION.....                         | 1 |
| 2. BACKGROUND AND RELATED WORKS.....         | 2 |
| 3. MULTI-SATELLITE ESTIMATOR DESIGN .....    | 2 |
| 4. IMPACT OF RANGING ON STATE ESTIMATE ..... | 4 |
| 5. LIMITS ON NAVIGATION ACCURACY .....       | 5 |
| 6. NETWORK TOPOLOGY .....                    | 5 |
| 7. EVALUATION.....                           | 5 |
| 8. CONCLUSIONS & OUTLOOK .....               | 7 |
| ACKNOWLEDGEMENTS.....                        | 8 |
| REFERENCES .....                             | 8 |
| BIOGRAPHY .....                              | 9 |

## 1. INTRODUCTION

As more spacecraft venture into deep space, there is a growing need for scalable navigation systems that are less dependent on Earth-based resources like ground-based radar and Global Navigation Satellite Systems (GNSS). Current missions beyond Earth orbit rely on ground stations such as NASA's Deep Space Network (DSN) [1] or ESA's European Space Tracking (ESTRACK) system [2]. These systems, with their large dishes, can only support a limited number of spacecraft and are already oversubscribed. Meanwhile, there are a growing number of missions planning to operate in cislunar space and at Mars, including by private companies like SpaceX [3].

To address the growing need for navigation beyond Earth



**Figure 1.** The spacecraft uses visual bearing measurements to landmarks on Earth’s surface and ranging measurements to other spacecraft to estimate both its orbital state and attitude.

orbit, novel navigation methods have to be developed that can provide consistent, accurate information independently of Earth-based resources. We propose a method that provides a solution for networked satellite systems. Each satellite in the network can autonomously localize itself and can thus be used to bootstrap accurate position estimates for other spacecraft or vehicles on a planetary surface. While our simulations are based on the parameters of the upcoming Argus low-Earth orbit demonstration mission, redeployment the method to new planetary bodies is straightforward — requiring only the retraining of landmark-classifier neural networks with new surface imagery and adjustment of model parameters.

Our contributions include: 1) the development of a joint orbit-and-attitude-determination method using one or more low-cost RGB cameras, angular-velocity data from a gyroscope, and ranging data from multiple networked satellites; 2) an analysis of scaling effects present in these satellite networks; and 3) simulations to validate the theoretical analysis.

The paper proceeds as follows: We first provide a short overview of the history of vision-based satellite navigation and some theoretical background. We then provide a descrip-

tion of the system and our state estimation approach. Network scaling effects, their limits, and the impacts of changing the satellite network topology are then analyzed. Finally, we provide simulations to validate the previous analysis and provide a glimpse of promising applications of this technology.

## 2. BACKGROUND AND RELATED WORKS

### *Vision-Based Satellite Navigation Methods*

Various vision-based approaches have been investigated for precise satellite navigation. For relative navigation to passive client spacecraft, bearing-only methods [4, 5] and range-only methods [6] have both shown promise, as well as detecting and tracking multiple other satellites in visual imagery [7]. Previous work [8] in the field of visual-inertial navigation has aimed to solve the closely related “lost-in-space problem,” where a spacecraft aims to localize itself after initializing in an unknown state, using solely bearing measurements from an RGB camera and an inertial measurement unit (IMU) with a least-squares approach. This low-cost method performed favorably against much more expensive systems, achieving position errors less than 5 km within four orbits 85% of the time in Monte-Carlo simulations [8]. Optical navigation techniques, such as those based on star trackers [9, 10] or planetary appearance analysis have shown consistent performance thanks to their usage of well-studied spatial geometries and classical computer-vision techniques [11–13].

### *Relative Navigation*

Research in the field of relative navigation for satellites has produced a broad array of applicable measurement types. Using carrier-phase differential GNSS signals, millimeter localization accuracy has been achieved [14]. Light detection and ranging (Lidar) methods have been used for proximity operations near non-cooperative objects [15, 16]. Extensive work has also been conducted on using camera-based measurements, both in monocular [17] and stereo [18] setups. Extensive work has been done on camera-based relative bearing measurements as shown in [4, 5, 19, 20]. While this may be theoretically possible with the current satellite framework design, as we have multiple onboard cameras, it is unlikely to work in practice with the same camera hardware, as our chosen cameras have been optimized for detecting surface features on Earth or other planetary bodies. This type of camera would not function well with the lower signal-to-noise ratio and higher dynamic range encountered when imaging space objects [21]. For the purposes of this paper, we will focus on ranging measurements that measure the distance between two satellites because this can easily be implemented using existing satellite radio hardware. The following sections will provide a detailed analysis of the benefits of including ranging measurements in the filter design and showcase key results associated with this setup.

### *Posterior Cramér-Rao Bound (PCRB)*

The Cramér-Rao Bound (CRB) is a useful statistical result that provides a lower bound on the performance of an estimator. In words, it states that the inverse of the Fisher information matrix is a lower bound on the covariance of any unbiased estimator. This bound works well for time-invariant systems where estimated parameters  $\theta$  do not change during the estimation period. For dynamical systems, we can extend the concept to the Posterior Cramér-Rao Bound (PCRB) [22]. We can define  $\mathbf{I}_P$  as representing the prior information and  $\mathbf{I}_M$  representing the posterior after information from a set of observations  $\mathbf{x}_k$  at timestep  $k$  is incorporated by means of the

function  $g(\mathbf{x}_k)$ . Using the algorithm presented in [23],  $\mathbf{I}_{P,k}$  can be defined recursively at timestep  $k$  using  $\mathbf{I}_{k-1}$ :

$$\mathbf{I}_{P,k} = (\mathbf{A}^\top \mathbf{Q}^{-1})^\top (\mathbf{I}_{k-1} + \mathbf{A}^\top \mathbf{Q}^{-1} \mathbf{A})^{-1} (\mathbf{A}^\top \mathbf{Q}^{-1}), \quad (1)$$

where  $\mathbf{A}$  is the discrete-time dynamics Jacobian and  $\mathbf{Q}$  is the process-noise covariance. If prior knowledge of the system states exists, the initial information matrix  $\mathbf{I}_0$  can be initialized with the prior knowledge. Otherwise a zero matrix can be used. Once the measurements are made, the standard Fisher information matrix term can be calculated as:

$$\mathbf{I}_{M,k} = \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H} \quad (2)$$

Combining the two we get the new information matrix:

$$\mathbf{I}_{k+1} = \mathbf{I}_{P,k} + \mathbf{I}_{M,k} \quad (3)$$

This information matrix gets propagated to the next timestep and used in the calculation of the next  $\mathbf{I}_P$ . To now calculate the PCRB, we simply invert the information matrix. This value defines the lower bound on the covariance at every timestep. We can assess the performance of our filter by checking the gap between its actual covariance and the PCRB. For our system with unbiased estimators, the covariance can never be lower than the PCRB. We can state this as:

$$\mathbb{E}([g(\mathbf{x}_k) - \theta][g(\mathbf{x}_k) - \theta]^\top) \geq \mathbf{I}_k \quad (4)$$

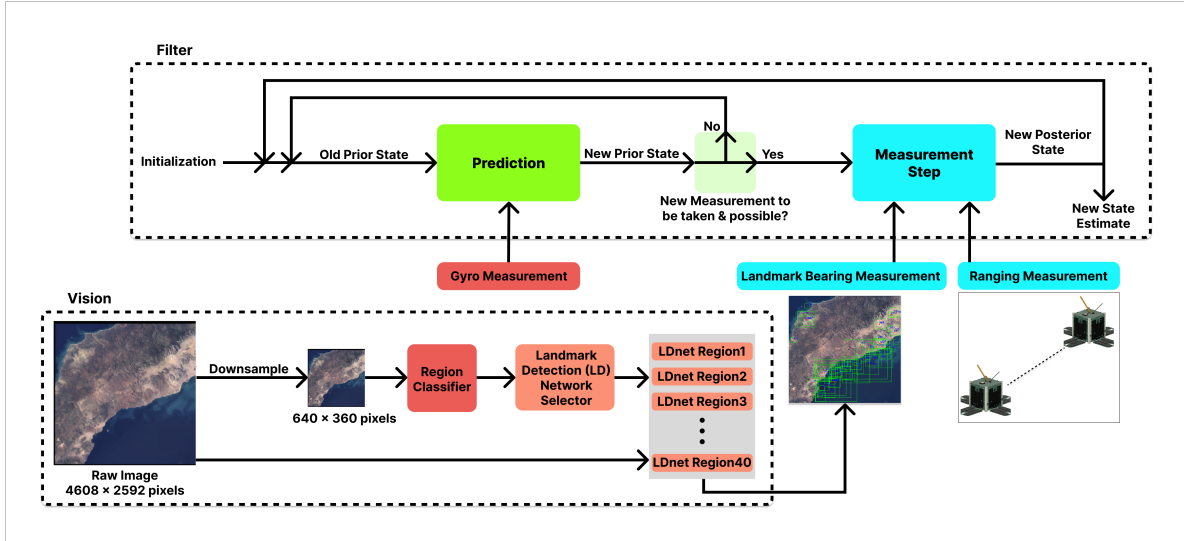
## 3. MULTI-SATELLITE ESTIMATOR DESIGN

In this section we consider the architecture of a multi-satellite system that fuses different measurement types to estimate the full pose of each satellite. The existence of multiple satellites allows us to consider measurements between the individual satellites on top of measurements that each satellite obtains individually. A schematic overview of this extended system is provided in Figure 2.

### *Filter States*

The filter estimates a total of 17 state variables: the spacecraft’s position  $\mathbf{r}$  in an Earth-centered inertial (ECI) frame, the spacecraft’s velocity  $\mathbf{v}$ , also in ECI, the spacecraft’s attitude  $\mathbf{q}$ , represented using a quaternion, the gyroscope’s bias  $\mathbf{w}_b$ , an unmodeled accelerations term  $\mathbf{u}_a$  that aims to estimate dynamics terms that are not otherwise modeled, and a drag scale factor  $d_{est}$  that is used to provide an estimate of the atmospheric drag on the spacecraft. We can express the continuous time system dynamics of the filter as follows:

$$f_c(\mathbf{x}) = \begin{cases} \dot{\mathbf{r}} = \mathbf{v} \\ \dot{\mathbf{v}} = \frac{-rGM_{Earth}}{\|\mathbf{r}\|^3} + \mathbf{u}_a + \mathbf{a}_{drag-est} + \mathbf{a}_{J2} \\ \dot{\mathbf{w}}_b = 0 \\ \dot{\mathbf{q}} = \frac{1}{2} \mathbf{q} \otimes \hat{\mathbf{w}}_{IMU} \\ \dot{\mathbf{u}}_a = 0 \\ \dot{d}_{est} = 0 \end{cases} \quad (5)$$



**Figure 2.** A diagrammatic overview of the system. The filter design is a variation of an Iterated Extended Kalman Filter. The measurements we receive are introduced in the prediction step, for the gyro measurement, and the measurement step for the landmark-bearing and inter-satellite ranging measurements.

It is important to note the difference in acceleration dynamics between the ground-truth simulation and the filter. The filter  $\dot{\mathbf{v}}$  dynamics shown above are simplified. The ground-truth simulation uses:

$$\dot{\mathbf{v}} = \frac{-\mathbf{r}GM_{\text{Earth}}}{\|\mathbf{r}\|^3} + \mathbf{a}_{\text{drag-gt}} + \mathbf{a}_{\text{J2}} + \mathbf{a}_{\text{sun}} + \mathbf{a}_{\text{moon}} \quad (6)$$

where  $\mathbf{a}_{\text{J2}}$  represents the acceleration due to the J2 perturbation [24], and  $\mathbf{a}_{\text{moon}}$  and  $\mathbf{a}_{\text{sun}}$  account for third-body accelerations due to the moon and sun, respectively.

*Dynamic Model Compensation (DMC)*—We utilize an unmodeled acceleration term  $\mathbf{u}_{\mathbf{a}}$  to compensate for differences between the ground-truth dynamics and filter dynamics. Since these differences are time-correlated, simply using zero-mean Gaussian process noise does not suffice. The addition of this  $\mathbf{u}_{\mathbf{a}}$  term is commonly known as Dynamic Model Compensation (DMC) and it has been shown to work for satellites in both lunar [25] and low-Earth orbits [26]. For our purposes, we assume that the unmodeled accelerations follows a stochastic random-walk process. We can express this in discrete time as:

$$\mathbf{u}_{\mathbf{a}}(k+1) = w(k), \quad w(k) \sim \mathcal{N}(0, \sigma_{u_{\mathbf{a}}}) \quad (7)$$

*Drag Estimation*—For the purpose of our ground-truth simulation, we calculate the atmospheric density using the Harris-Priester model [27, 28]. This takes the form:

$$\mathbf{a}_{\text{drag-gt}} = -\frac{C_d L}{2M} \rho_{HP} \mathbf{v} \|\mathbf{v}\| \quad (8)$$

where  $C_d$  represents the drag coefficient of the spacecraft,  $L$  is the surface area of the spacecraft,  $M$  the mass of the spacecraft and  $\rho_{HP}$  is the atmospheric density as calculated by the Harris-Priester model. Due to the multiple orders-of-magnitude range of the drag acceleration that the spacecraft encounters [29], directly estimating the density is difficult.

Instead, we estimate a drag scaling parameter, resulting in the following model for the atmospheric density based on [30],

$$\rho_E = d_{\text{est}} \rho_0 e^{-\frac{\|\mathbf{r}\| - R_E}{H}} \quad (9)$$

where  $\rho_0$  refers to a nominal density value,  $R_E$  refers to the equatorial radius of Earth, and  $H$  is a reference height set, in our case, to 600km, a reasonable altitude for a low Earth orbit spacecraft. Replacing the atmospheric density in the drag equation with this model, we get the drag acceleration used by the filter:

$$\mathbf{a}_{\text{drag-est}} = -\frac{C_d L}{2M} \rho_E \mathbf{v} \|\mathbf{v}\| \quad (10)$$

*Gyro Bias Estimation*—The spacecraft possesses a gyro that is used to measure angular velocity. These measurements are incorporated in the prediction step of the filter. The simulation uses the following gyro model [31]:

$$\mathbf{y} = \begin{bmatrix} 1 + s_1 & m_{12} & m_{13} \\ m_{21} & 1 + s_2 & m_{23} \\ m_{31} & m_{32} & 1 + s_3 \end{bmatrix} \begin{bmatrix} w_x \\ w_y \\ w_z \end{bmatrix} + \mathbf{w}_{\mathbf{b}} + \mathcal{N}(0, \sigma_w) \quad (11)$$

The gyro model takes a ground-truth angular velocity  $\mathbf{w}$  and produces the output  $\mathbf{y}$ . The constants  $s_i$  refer to scaling factors present in the gyro, leading to a “stretching” of the signal. The  $m_{xy}$  terms refer to axial misalignment between axis  $x$  and  $y$ . The gyro also possesses an additive bias  $\mathbf{w}_{\mathbf{b}}$  that we need to consider in order to correctly estimate the attitude state of the spacecraft. The gyro bias is modeled to follow a random walk. The final term in the gyro model applies white noise sampled from a zero-mean Gaussian distribution.

#### Measurement Model

Each satellite has a set of four cameras with which it captures RGB images. We developed a machine-vision pipeline that can extract usable bearing measurements from these images.

These bearing measurements take the form:

$$\tilde{b}_{k,i,j,c} = \frac{r_{k,i} - l_j}{\|r_{k,i} - l_j\|} \mathbf{R}_c \mathbf{S} \quad (12)$$

where  $r_{k,i}$  is the position of satellite  $i$  at timestep  $k$ ,  $l_j$  is the position of landmark  $j$ ,  $\mathbf{R}_c$  is the rotation from the body frame to the camera frame  $c$ , and  $\mathbf{S}$  is the noise rotation matrix, created by sampling an axis-angle vector from a Gaussian, that impacts the bearing measurement.

We can express the range measurement model between satellite  $i$  and satellite  $j$  as:

$$\tilde{r}_k = \|r_{i,k} - r_{j,k}\| + v_k \quad (13)$$

where  $\|r_{i,k} - r_{j,k}\|$  is the norm of the difference in state position between the two satellites at timestep  $k$  and  $v_k$  is Gaussian white noise assumed to be drawn from a normal distribution with zero mean.

For the purpose of our simulation, we also have to ensure that a measurement is only generated when two satellites have a direct line of sight. We achieve this by calculating a ray between the two satellites and checking whether this ray intersects with an oblate-spheroid Earth model at any point. If an intersection exists, we disregard the ranging measurement. The full measurement vector takes the form:

$$\tilde{\mathbf{y}} = \begin{bmatrix} \tilde{b}_1 \\ \vdots \\ \tilde{b}_m \\ \tilde{r}_1 \\ \vdots \\ \tilde{r}_n \end{bmatrix} \in \mathbb{R}^{3m+n} \quad (14)$$

where  $m$  represents the numbers of visible landmarks and  $n$  the number of visible satellites.

The ranging measurement naturally has some amount of noise associated with its measurement. For a current generation CubeSat, previous work [32] has shown that the ranging error expected with this type of system is around one meter. Consequently we can extend our measurement noise covariance  $\mathbf{W}$  with a new diagonal matrix with these entries.

The  $g_{land}$  function described in Algorithm 2 is the landmark bearing function that one satellite uses to estimate bearings to landmarks it observes on the Earth. For notational simplicity we define the rotation matrices  $\mathbf{B}$  as the rotation from body frame to ECI frame and  $\mathbf{E}$  as the rotation from ECI to ECEF frame.

The  $g_{range}$  function described in Algorithm 3 is the ranging estimation function that one satellite uses to estimate the distance between itself and other satellites. The function  $f_{vis-sats}$  is used to extract the identity of the other satellite from the received ranging measurement.

#### 4. IMPACT OF RANGING ON STATE ESTIMATE

The most straight-forward way to fuse data from multiple satellites is to consider them to all be part of a single larger

---

#### Algorithm 1 Measurement Step

---

**Given:**  $\mathbf{P}_{k+1|k} \in \mathbb{R}^{16 \times 16}$ ,  $\mathbf{x}_{k+1|k} \in \mathbb{R}^{16}$ ,  $\text{num\_iter} \in \mathbb{R}$ ,  $(\tilde{\mathbf{b}}, \mathbf{l}) \in (\mathbb{R}^{m \times 3})^2$ ,  $\tilde{\mathbf{r}} \in \mathbb{R}^n$ ,  $\mathbf{t} \in \mathbb{R}$

**for**  $i = 1:\text{num\_iter}$  **do**

**if**  $i == 0$  **then**

$\bar{\mathbf{x}}^i \leftarrow \mathbf{x}_{k+1|k}$

**else**

$\bar{\mathbf{x}}^i \leftarrow \mathbf{x}_{k+1|k+1}^{i-1}$

**end if**

$\mathbf{h}_{land}^i \leftarrow g_{land}(\bar{\mathbf{x}}^i, \mathbf{l}, \mathbf{t})$

$\mathbf{H}_{land}^i \leftarrow \left. \frac{\partial g_{land}}{\partial \mathbf{x}} \right|_{\bar{\mathbf{x}}^i, \mathbf{l}, \mathbf{t}}$

$\mathbf{h}_{range}^i \leftarrow g_{range}(\bar{\mathbf{x}}^i, \tilde{\mathbf{r}})$

$\mathbf{H}_{range}^i \leftarrow \left. \frac{\partial g_{range}}{\partial \mathbf{x}} \right|_{\bar{\mathbf{x}}^i, \tilde{\mathbf{r}}}$

$\mathbf{h}^i \leftarrow \text{concatenate}(\mathbf{h}_{land}^i, \mathbf{h}_{range}^i)$

$\mathbf{H}^i \leftarrow \text{concatenate}(\mathbf{H}_{land}^i, \mathbf{H}_{range}^i)$

$\tilde{\mathbf{y}} \leftarrow \text{concatenate}(\tilde{\mathbf{b}}, \tilde{\mathbf{r}})$

$\mathbf{K}^i \leftarrow \mathbf{P}_{k+1|k} \mathbf{H}^i (\mathbf{H}^i \mathbf{P}_{k+1|k} \mathbf{H}^{i\top} + \mathbf{W})^{-1}$

$\mathbf{x}_{k+1|k+1, q}^i \leftarrow \bar{\mathbf{x}}^i + \mathbf{K}^i (\tilde{\mathbf{y}} - \mathbf{h}^i)$

$\mathbf{x}_{k+1|k+1, q}^i \leftarrow \bar{\mathbf{x}}^i * \mathbf{K}^i (\tilde{\mathbf{y}} - \mathbf{h}^i)$

$\mathbf{P}_{k+1|k+1} \leftarrow (\mathbb{I}_{16} - \mathbf{K}^i \mathbf{H}^i) \mathbf{P}_{k+1|k} (\mathbb{I}_{16} - \mathbf{K}^i \mathbf{H}^i)^\top + \mathbf{K}^i \mathbf{R} \mathbf{K}^{i\top}$

**end for**

**return**  $\mathbf{x}_{k+1|k+1}, \mathbf{P}_{k+1|k+1}$

---



---

#### Algorithm 2 $g_{land}$

---

**Given:**  $\mathbf{x} \in \mathbb{R}^{16}$ ,  $\mathbf{l} \in \mathbb{R}^{m \times 3}$ ,  $\mathbf{t} \in \mathbb{R}$

$\mathbf{B} \leftarrow f_{qtoR}(\mathbf{x}_q)$

$\mathbf{E} \leftarrow f_{ECItoECEF}(\mathbf{t})$

$\mathbf{l}_{ECEF} \leftarrow \mathbf{E} \mathbf{l}$

$\mathbf{r}_{ECEF} \leftarrow \mathbf{E} \mathbf{x}_r$

$\mathbf{c}_{ECEF} \leftarrow f_{cam}(\mathbf{r}_{ECEF}, \mathbf{E} \mathbf{B})$

$\hat{\mathbf{b}}_{ECEF} \leftarrow (\mathbf{l}_{ECEF} - \mathbf{c}_{ECEF}) / \|\mathbf{l}_{ECEF} - \mathbf{c}_{ECEF}\|$

$\hat{\mathbf{b}}_{body} \leftarrow (\mathbf{E} \mathbf{B})^\top * \hat{\mathbf{b}}_{ECEF}$

**return**  $\hat{\mathbf{b}}_{body}$

---

system. This means we stack the state vectors and combine the covariance matrices of all satellites, allowing us to properly consider cross terms between all the states of all the satellites. It is interesting to consider how exactly the information from the ranging measurement flows into the measurement update from satellites that may or may not share a measurement link. For this, we must consider the measurement Jacobian of a multi-satellite network as this is the mechanism by which information between satellites is exchanged. For illustration purposes, we will consider a system with two satellites and a single ranging measurement between them, but in theory this can be extended to any size of satellite network. Consider the measurement vector for this system, assuming that the two satellites can see each other and both see a certain number of landmarks (denoted as  $m_1$  and  $m_2$  respectively):

$$\mathbf{y} = \begin{bmatrix} \tilde{b}_1 \\ \tilde{r}_1 \\ \tilde{b}_2 \\ \tilde{r}_2 \end{bmatrix} \in \mathbb{R}^{3m_1+3m_2+2} \quad (15)$$

---

**Algorithm 3**  $g_{range}$ 

---

```
1: Given:  $\mathbf{x} \in \mathbb{R}^{16}$ ,  $\tilde{\mathbf{r}} \in \mathbb{R}^n$ 
2:  $visible\_sats \leftarrow f_{vis-sats}(\tilde{\mathbf{r}})$ 
3:  $ranges \leftarrow \text{None}$ 
4: for  $sat \leftarrow 1$  to  $visible\_sats$  do
5:    $ranges \leftarrow \text{append}(ranges, \|\mathbf{x}_r - sat_r\|)$ 
6: end for
7: return  $ranges$ 
```

---

$\tilde{r}_1$  and  $\tilde{r}_2$  represent the respective ranging measurements that satellite 1 and satellite 2 take to each other. Since these are both noisy, they will not match precisely. We can now consider the system measurement Jacobian  $\mathbf{H}$  of this system in block form:

$$\begin{matrix} & \text{Sat 1} & \text{Sat 2} \\ \begin{matrix} 1 \\ 2 \\ \vdots \\ m_1 + 1 \\ m_1 + 2 \\ \vdots \\ m_1 + m_2 + 2 \end{matrix} & \begin{pmatrix} \mathbf{H}_b(l_{11}) & \mathbf{0}_{3 \times 16} \\ \mathbf{H}_b(l_{12}) & \mathbf{0}_{3 \times 16} \\ \vdots & \vdots \\ \mathbf{H}_r(\tilde{r}_1) & \mathbf{H}_r(\tilde{r}_2) \\ \mathbf{0}_{3 \times 16} & \mathbf{H}_b(l_{21}) \\ \vdots & \vdots \\ \mathbf{H}_r(\tilde{r}_1) & \mathbf{H}_r(\tilde{r}_2) \end{pmatrix} \end{matrix} \quad (16)$$

The landmark bearing measurement Jacobians, denoted  $\mathbf{H}_b$ , are not coupled between satellites and are identical to what would be found in a single-satellite case. Each of these Jacobian blocks consists of:

$$\mathbf{H}_b(l) = \left[ \begin{array}{ccc|cc} \frac{\partial z_l}{\partial \mathbf{r}} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \frac{\partial z_l}{\partial \mathbf{q}} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{array} \right]_{z_l=l} \quad (17)$$

The matrix block that is responsible for allowing the information from one satellite's measurements to "flow" and impact the other satellite's states are the ranging measurement Jacobians. Each block consists of:

$$\mathbf{H}_r(r) = \left[ \begin{array}{cccccc} \frac{\partial z_s}{\partial \mathbf{r}} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{array} \right]_{z_s=r} \quad (18)$$

Since one individual ranging measurement results in ranging Jacobians  $\mathbf{H}_r$  for both satellites involved in the measurement, the system Jacobian  $\mathbf{H}$  introduces coupling effects between the states and covariances of different satellites. The implication is that any one satellite can benefit from another satellite's bearing measurements, even if those measurements have nothing to do with the first satellite. The entirety of the information contained in these measurements is exchanged solely through the cross terms in the measurement Jacobian. The result is that the state estimate of a satellite can be drastically improved even if it gets no bearing measurements of its own. We can show that the state estimate of these "blind" satellites is improved by the simple fact that they can take a single ranging measurement to another satellite.

We can also conduct a residual analysis to check convergence of the filter. To this end we are interested in analyzing the Normalized Innovation Squared (NIS). Using the definition of the innovation  $\mathbf{v}$  as the difference between the measurement and the estimate, and the innovation covariance  $\mathbf{C}$  that we calculate as part of the measurement step in Algorithm 1, we can define the NIS as:

$$\begin{aligned} \mathbf{v} &= \tilde{\mathbf{y}} - \mathbf{h} \\ \mathbf{C} &= \mathbf{H}\mathbf{P}\mathbf{H}^\top + \mathbf{W} \\ \text{NIS} &= \mathbf{v}^\top \mathbf{C}^{-1} \mathbf{v} \end{aligned} \quad (19)$$

The NIS is a dimensionless measure of the size our innovation relative to the uncertainty that the filter predicts for it. A larger NIS value therefore indicates a more "surprising" measurement, suggesting that the filter has diverged in its system outlook from the ground truth.

## 5. LIMITS ON NAVIGATION ACCURACY

Using the information matrix and Cramér-Rao Bound presented earlier, it is possible to show that the amount of information in a system cannot decrease through additional measurements. This does not necessarily mean that more measurements will guarantee a better state estimate, as the measurements may not have any impact from an information-theoretic point of view. Instead, this means that any additional measurement will not result in a worsened upper bound on the uncertainty present of the system. We can assess the system uncertainty expressed by the covariance matrix against the Posterior Cramér-Rao Bound (PCRB). By checking the gap between the Root Mean Square Error (RMSE) and the PCRB we can assess the performance of our estimation filter. A smaller gap indicates that the filter is operating closer to the optimal state it can achieve, while larger gaps indicate weaker performance. We can use this knowledge to effectively tune the initialization and noise parameters associated with the filter itself. It is straightforward to apply this concept of bounded uncertainty to the presented satellite networks. In this case, additional measurements are introduced by adding additional satellites to the system, which can be used for additional ranging measurements by other satellites. It is then of interest to observe how the satellites' state estimates perform with each new additional measurement.


## 6. NETWORK TOPOLOGY

With multiple satellites, there is also the possibility of utilizing different network topologies. In this paper, we differentiate between two approaches: a mesh network, where every satellite exchanges information with every other satellite, and a star network where one "chief" satellite takes measurements with all other satellites, but those other satellites do not take measurements among each other. Figure 3 illustrates these two network topologies for a set of four satellites.

Different network topologies provide different advantages and disadvantages: A more connected network will provide the ability to gather more information for each satellite and thus improve their own state estimate. However, this comes at the cost of requiring more exchanges of measurements and states between satellites, and can increase power consumption. Especially on small-scale CubeSats, this cost has to be carefully weighed against other vital objectives such as down-linking information to a ground station.

## 7. EVALUATION

The following simulation setup was used to benchmark the system. Besides the given parameters, it must also be noted that all spacecraft were placed in (near) polar circular orbits,



**Figure 3.** Comparison of the different satellite topologies. *Left:* Mesh network with ranging measurements between all satellites. *Right:* Star network with one central “master” satellite collecting measurements to all other surrounding satellites.

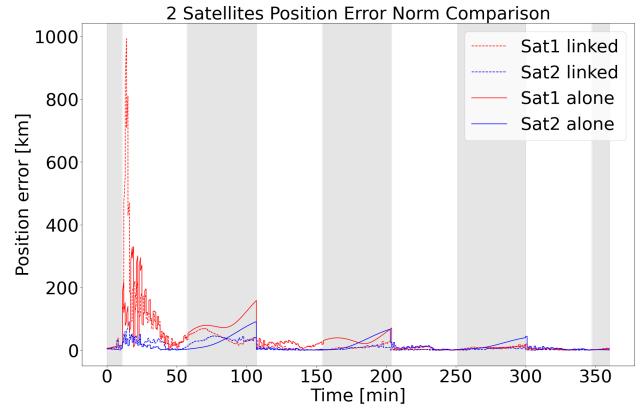
similar to what would be expected for the upcoming Argus CubeSat mission. The satellites were normally distributed around the latitude, longitude pair stated in the parameters below. Furthermore, the angular velocity was parameterized to describe a rotation akin to a wobbling motion.  $\Delta t$  corresponds to the simulation step size and  $t$  to the current step.

| Parameter                         | Value   | Unit                             |
|-----------------------------------|---|----------------------------------|
| Time Step                         | 0.5   | s                                |
| Sim Duration (Multi Sat.)         | 21600   | s                                |
| Starting Epoch                    | 60431   | MJD                              |
| Iterated EKF Steps                | 5   | –                                |
| Monte Carlo Trials                | 5000  | –                                |
| Initial Altitude                  | 600   | km                               |
| Initial Latitude                  | 0.0   | °                                |
| Initial Longitude                 | -73.0   | °                                |
| Measurement Rate                  | 60  | s                                |
| Angular Velocity x                | $\cos\left(\frac{2\pi t}{10/\Delta t}\right)$ | rad s <sup>-1</sup>              |
| Angular Velocity y                | $\sin\left(\frac{2\pi t}{10/\Delta t}\right)$ | rad s <sup>-1</sup>              |
| Angular Velocity z                | $\frac{\pi}{18}$                              | rad s <sup>-1</sup>              |
| Init. position error sample cov.  | 10000   | m <sup>2</sup>                   |
| Init. velocity error sample cov.  | 10  | m <sup>2</sup> s <sup>-2</sup>   |
| Init. att. error sample cov.      | 0.001   | – (DCM)                          |
| Init. gyro bias error sample cov. | $2 \times 10^{-3}$                            | rad <sup>2</sup> s <sup>-2</sup> |
| Init. drag scalar                 | 1   | –                                |


**Table 1.** Simulation Parameters

### Exchanging Ranging Measurements

The benefit of adding a single measurement between two satellites can be seen in Figure 4. In the first situation, shown by the solid lines, the two satellites are not coupled through a ranging measurement, essentially creating two independent, single satellite systems. The observed error shows a significant blowup occurring during the first eclipse, when vision measurements are unavailable, and subsequent reductions in error once those measurements are available again. The second scenario, shown by the dashed lines, illustrates the benefit of a ranging measurement. Starting in the first full orbit, the ranging measurement reduces the observed peak error magnitude significantly. It is notable that when landmark measurements are available, the benefit that the ranging measurement yields diminishes greatly.



**Figure 4.** Comparative plot showing the error norm differences between two satellites operating independently (solid lines) and two satellites coupled by a single ranging measurement between them (dashed lines). Gray shading indicates eclipsed parts of the orbit where landmark bearing measurements are not possible. The estimation filter requires about half an orbit to stabilize after initialization.

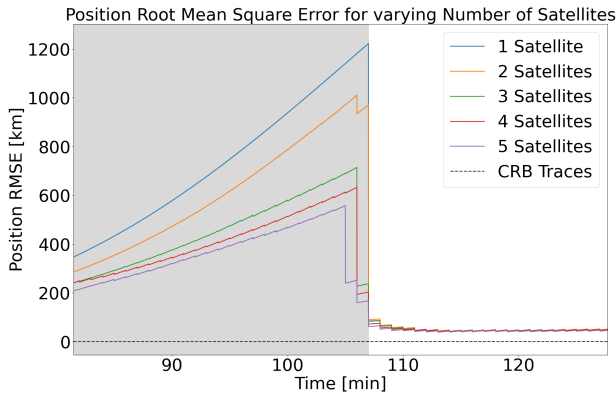


**Figure 5.** Root-mean-square error for fully connected networks containing between 1 and 5 satellites over four orbits. Gray shading indicates eclipsed parts of the orbit where landmark bearing measurements are not possible. Again note the initial error blowup after initialization and subsequent decrease. The more satellites are added, the smaller the error blowup during eclipse.

### Network Scaling Effects

A question that naturally arises after seeing the benefit of adding a single satellite ranging measurement is; what happens when more satellites are added?

Figure 5 illustrates the position residual as a function of the number of satellites. The satellites for this experiment are all arranged in a fully connected network. The plot clearly illustrates that adding more satellites to a network improves the performance of the whole system, as the Root Mean Square Error (RMSE) moves closer to its respective CRB (shown as the black dashed line). Note that the CRBs for all numbers of satellites lie so closely together that differences in value are not discernible on this scale. Figure 6 also illustrates that each additional satellite results in a diminishing marginal



**Figure 6.** A close up of Figure 5. We can see that the gains made by adding a satellite are not linear, as each additional satellite results in a diminishing marginal gain. More satellites also lead to an earlier reduction in error (indicated by earlier drops for higher numbers of satellites).

gain. The performance gain from adding the fifth satellite is not as great as the gain from the fourth and the fourth satellite in turn has less impact than the third. And this is only for the night-section of the orbit. During the daytime, there is little discernible difference at all. This reinforces the point made previously that when landmark bearing measurements are available, the benefit of having an additional ranging measurements is greatly diminished.

#### Mesh vs. Star

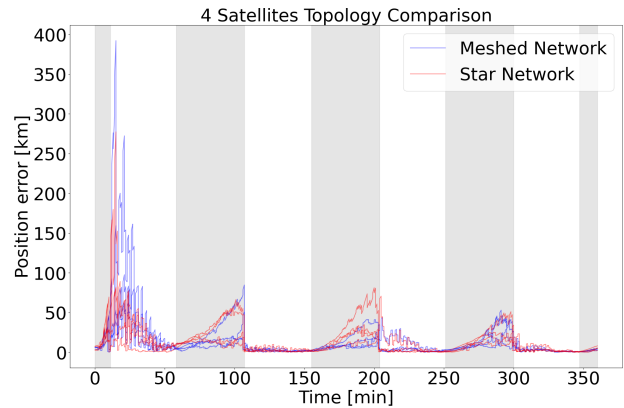
As discussed in Section 6, two natural network arrangements that trade off estimation accuracy with power consumption are the “star” network on one side, where all ranging measurements are made with one central satellite, and the “mesh” network, where all satellites exchange measurements with one another. Figure 7 compares the performance of the two discussed approaches for a set of four satellites. The mesh network performs favorably against the star network during eclipse after the filter states have settled. Note the consistent estimation accuracy advantage that the mesh network has over the star network. However, the plot once more illustrates the point made previously: when more landmark-bearing measurements are available, the estimation performance of the two networks are nearly equivalent.

#### Convergence Analysis via NIS


We can consider the performance of the filter through the lens of the NIS as defined in Section 4. For this purpose we consider the case of two satellites exchanging ranging measurements and taking landmark measurements.

## 8. CONCLUSIONS & OUTLOOK

The primary contribution of this work is the creation of a framework for vision-based navigation for satellite networks. While our figures only present results for position states, it is important to note that our system estimates the full spacecraft pose. We neglected to present error plots for other states, such as attitude, as these are not impacted by the ranging measurements that we are primarily interested in. We have shown that the full pose can be accurately estimated using



**Figure 7.** A comparative plot of a fully connected (“meshed”) network of four satellites, shown in blue, and a centralized (“star”) network, shown in red. Exchanging more measurements reduces the impact of the eclipse blowup, but also comes at a higher power consumption cost which, especially on a constrained system, may be more limiting. Initial error spikes (especially for the meshed network) are indicative of poor initial state estimates being exchanged between satellites.



**Figure 8.** A plot illustrating the dimensionless magnitude of the NIS as a function of time for two satellites, averaged over all Monte Carlo simulations. Note the high spikes when the satellite initially observes the Earth and takes measurements and the subsequent decrease in magnitude as the measurements becomes less “surprising.” As the filter must propagate its state and covariances through the eclipse of the orbits, in gray, small spikes are observed on recommencement of Earth observation. As the orbits progress, these continue to decrease in magnitude. During the eclipse itself, satellites can only exchange measurements between each other which on their own carry little interesting information, thus leading to minimal spikes.

only measurements derived from a realistic gyroscope and a machine-learning-based vision pipeline. We have also analyzed the positive impact on state error of extending the satellite measurements to incorporate measurements received from other satellites in the form of the inter-satellite distance.

Our primary conclusion is that extending the network to more satellites has a benefit on the state error that diminishes as more satellites are included in the network. We have shown that different network topologies bring different trade-offs. One interesting avenue for future work is to consider the network topology through the lens of energy consumption. Since the mesh network has more exchanges and thus a higher power usage, it is hard to justify using it when the star network performs just as well. A conclusion that all multi-satellite evaluations were able to reach is that exchanging measurements with other satellites becomes drastically less important when other types of measurements are also available.

The next step in this project is to test this system in low-Earth orbit on Carnegie Mellon’s upcoming ARGUS CubeSat mission. Future work will also study porting our framework to Lunar and Martian orbits. The vision pipeline can, in principle, be re-trained using existing datasets collected from these bodies. The lack of a significant atmosphere and resulting visual occlusions (such as clouds) on these bodies means that feature matching should work even better than on Earth.

## ACKNOWLEDGEMENTS

The authors would like to thank all those that helped in the course of this research, including the members of CMU’s Spacecraft Design-Build-Fly Laboratory course for the many fruitful discussions that helped refine the concepts in this paper.

## REFERENCES

- [1] D. J. Mudgway, *Uplink-downlink: a history of the nasa deep space network, 1957-1997*. National Aeronautics and Space Administration, 2001, vol. 4227.
- [2] Y. Doat, M. Lanucara, P.-M. Besso, T. Beck, G. Lorenzo, and M. Butkowic, “Esa tracking network—a european asset,” in *2018 SpaceOps Conference*, 2018, p. 2306.
- [3] J. L. Heldmann, M. M. Marinova, D. S. Lim, D. Wilson, P. Carrato, K. Kennedy, A. Esbeck, T. A. Colaprete, R. C. Elphic, J. Captain, K. Zacny, L. Stolov, B. Mellerowicz, J. Palmowski, A. M. Bramson, N. Putzig, G. Morgan, H. Sizemore, and J. Coyan, “Mission architecture using the spacex starship vehicle to enable a sustained human presence on mars,” *New Space*, vol. 10, no. 3, pp. 259–273, 2022, pMID: 36199953. [Online]. Available: <https://doi.org/10.1089/space.2020.0058>
- [4] S. D’Amico, J.-S. Ardaens, G. Gaias, H. Benninghoff, B. Schlepp, and J. L. Jørgensen, “Noncooperative rendezvous using angles-only optical navigation: System design and flight results,” *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 6, pp. 1576–1595, 2013. [Online]. Available: <https://doi.org/10.2514/1.59236>
- [5] J. Sullivan and S. D’Amico, “Nonlinear kalman filtering

- for improved angles-only navigation using relative orbital elements,” *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 9, pp. 2183–2200, 2017. [Online]. Available: <https://doi.org/10.2514/1.G002719>
- [6] J. A. Christian, “Relative navigation using only inter-satellite range measurements,” *Journal of Spacecraft and Rockets*, vol. 54, no. 1, pp. 13–28, 2017.
- [7] J. Kruger, S. S. Hwang, and S. D’Amico, “Starling formation-flying optical experiment: Initial operations and flight results,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.06748>
- [8] K. McCleary, S. Gurumurthy, P. R. Fisch, S. Tayal, Z. Manchester, and B. Lucia, “Vinsat: Solving the lost-in-space problem with visual-inertial navigation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 11 774–11 780.
- [9] S. Casini, A. Cervone, B. Monna, and P. Visser, “Design and testing of star tracker algorithms for autonomous optical line-of-sight deep-space navigation,” *Applied Optics*, vol. 62, no. 22, pp. 5896–5909, 2023.
- [10] G. Carmeli and B. Ben-Moshe, “Ai-based real-time star tracker,” *Electronics*, vol. 12, no. 9, p. 2084, 2023.
- [11] J. A. Christian, “Optical navigation using planet’s centroid and apparent diameter in image,” *Journal of guidance, control, and dynamics*, vol. 38, no. 2, pp. 192–204, 2015.
- [12] —, “Optical navigation using iterative horizon reprojection,” *Journal of Guidance, Control, and Dynamics*, vol. 39, no. 5, pp. 1092–1103, 2016.
- [13] —, “A tutorial on horizon-based optical navigation and attitude determination with space imaging systems,” *IEEE Access*, vol. 9, pp. 19 819–19 853, 2021.
- [14] V. Capuano, A. Harvard, and S.-J. Chung, “On-board cooperative spacecraft relative navigation fusing gnss with vision,” *Progress in Aerospace Sciences*, vol. 128, p. 100761, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0376042121000646>
- [15] J. O. Woods and J. A. Christian, “Lidar-based relative navigation with respect to non-cooperative objects,” *Acta Astronautica*, vol. 126, pp. 298–311, 2016, space Flight Safety. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576515301661>
- [16] A. Rhodes, E. Kim, J. A. Christian, and T. Evans, “Lidar-based relative navigation of non-cooperative objects using point cloud descriptors,” in *AIAA/AAS Astrodynamics Specialist Conference*, 2016, p. 5517.
- [17] S.-H. Kim, H.-L. Choi, and H. Shim, “Mono-vision based satellite relative navigation using active contour method,” *Journal of the Korean Society for Aeronautical & Space Sciences*, vol. 43, no. 10, pp. 902–909, 2015.
- [18] S. Segal, A. Carmi, and P. Gurfil, “Stereovision-based estimation of relative dynamics between noncooperative satellites: Theory and experiments,” *IEEE Transactions on Control Systems Technology*, vol. 22, no. 2, pp. 568–584, 2014.
- [19] V. Pesce, R. Opromolla, S. Sarno, M. Lavagna, and M. Grassi, “Autonomous relative navigation around uncooperative spacecraft based on a single camera,” *Aerospace Science and Technology*, vol. 84, pp. 1070–1080, 2019. [Online]. Avail-

able: <https://www.sciencedirect.com/science/article/pii/S1270963818317346>

- [20] G. Napolano, C. Vela, A. Nocerino, R. Opromolla, and M. Grassi, "A multi-sensor optical relative navigation system for small satellite servicing," *Acta Astronautica*, vol. 207, pp. 167–192, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S009457652300125X>
- [21] F. Sansone, F. Branz, and A. Francesconi, "A relative navigation sensor for cubesats based on led fiducial markers," *Acta Astronautica*, vol. 146, pp. 206–215, 2018.
- [22] P. Tichavsky, "Posterior cramer-rao bound for adaptive harmonic retrieval," *IEEE Transactions on Signal Processing*, vol. 43, no. 5, pp. 1299–1302, 1995.
- [23] P. Tichavsky, C. H. Muravchik, and A. Nehorai, "Posterior cramer-rao bounds for discrete-time nonlinear filtering," *IEEE Transactions on signal processing*, vol. 46, no. 5, pp. 1386–1396, 1998.
- [24] J. KECHICHIAN and T. KELLY, *Analytic solution of perturbed motion in near-circular orbit due to J2, J3 earth zonal harmonics in rotating and inertial Cartesian reference frames*, 1989. [Online]. Available: <https://arc.aiaa.org/doi/abs/10.2514/6.1989-352>
- [25] B. Tapley and B. Schutz, "Estimation of unmodeled forces on a lunar satellite," *Celestial mechanics*, vol. 12, no. 4, pp. 409–424, 1975.
- [26] K. A. Myers and B. D. Tapley, "Dynamical model compensation for near-earth satellite orbit determination," *AIAA Journal*, vol. 13, no. 3, pp. 343–349, 1975.
- [27] D. L. Dowd and B. Tapley, "Density models for the upper atmosphere," *Celestial mechanics*, vol. 20, no. 3, pp. 271–295, 1979.
- [28] I. Harris and W. Priester, "The upper atmosphere in the range from 120 to 800 km," Tech. Rep., 1964.
- [29] W. J. Larson and J. R. Wertz, "Space mission analysis and design," 1999.
- [30] O. Montenbruck, E. Gill, and F. Lutz, "Satellite orbits: models, methods, and applications," *Appl. Mech. Rev.*, vol. 55, no. 2, pp. B27–B28, 2002.
- [31] M. Kirkko-Jaakkola, J. Collin, and J. Takala, "Bias prediction for mems gyroscopes," *IEEE Sensors Journal*, vol. 12, no. 6, pp. 2157–2163, 2012.
- [32] M. Holliday, R. Ticknor, J. Stupl, R. Hunter, P. Fisch, I. S. Sow, J. Willis, and Z. Manchester, "The py4 mission: A low-cost demonstration of cubesat formation-flying technologies," 2024.

## BIOGRAPHY



**Frederik Markus** is a Guidance, Navigation & Control Engineer at the European Space Agency (ESA). He received a MSc in electrical engineering in 2025 and a BSc in electrical engineering in 2022, both from ETH Zürich. His research interests include distributed control and state estimation with applications to aerospace.



**Zac Manchester** is an assistant professor in the Robotics Institute at Carnegie Mellon University and founder of the Robotic Exploration Lab. He received a PhD in aerospace engineering in 2015 and a BS in applied physics in 2009, both from Cornell University. His research interests include control and optimization with application to aerospace and robotic systems with challenging

nonlinear dynamics.