Received: DD MM YYYY | Revised: DD MM YYYY | Accepted: DD MM YYYY
DOI: 10.1109/xxx.xx.xxxx

ORIGINAL PAPER

Time-Correlated Ephemeris Errors and Relativistic Effects for Single-Satellite Doppler Localization of Lunar Rovers

Kaila M. Y. Coimbra | Grace Gao

[1]Department of Aeronautics and Astronautics, Stanford University

Correspondence

*Grace Gao, Department of Aeronautics and Astronautics, Stanford University
Stanford, CA, USA 94305
Email: gracegao@stanford.edu

Abstract

As near-term lunar missions move toward sustained surface operations, rovers must achieve absolute localization well before dedicated navigation constellations become available. This work investigates a minimal-infrastructure scenario in which the rover relies on only its onboard odometry and the Doppler shift of the downlink signal from a single relay satellite without a dedicated navigation payload. We focus on time-correlated satellite ephemeris errors and deterministic relativistic clock rate effects, which are two error sources that dominate in the single-satellite regime yet remain under-treated in prior lunar navigation studies. In simulation, rover odometry is modeled as a noisy incremental displacement, both rover and satellite clock bias and drift are propagated using stochastic dynamics, and broadcast ephemeris errors are captured using a Singer process with periodic orbit determination and time synchronization (ODTS) resets. The measurement simulation and estimator incorporate light-time delay and a consistent application of proper time and coordinate time. To process these time-varying effects under the limited observability of a single satellite, we adopt a tightly coupled Extended Kalman Filter (EKF) augmented with a first-order Gauss-Markov bias state and apply a Rauch-Tung-Striebel (RTS) backward smoothing pass to refine the estimates offline. The framework is validated through 100 Monte Carlo simulations over a 24-hour window using parameters drawn from ESA Lunar Pathfinder and NASA Endurance specifications. The full filter reduces the final mean 3D position error by roughly 19 m (36.7%) compared to the odometry-only baseline. An ablation study shows that omitting the relativistic correction in the filter increases the position error by a factor of 2.88. The RTS smoother further reduces error growth due to the time-correlated ephemeris errors, allowing for more accurate geotagging of samples collected along the rover path.

Keywords

Lunar navigation, time-correlated ephemeris errors, relativistic effects, Doppler shift

1 | INTRODUCTION

In recent years, there has been renewed interest in developing a sustained presence on the Moon. A growing set of lunar mission architectures from space agencies and commercial providers are targeting rover operations supported by dedicated or hosted relay satellites (Fernando et al., 2025; Keane et al., 2022). In the early stages of lunar exploration, prior to full LunaNet constellations (Giordano et al., 2023), absolute rover localization with minimal infrastructure will be critical. In this work, we investigate the localization performance of a single rover operating in the Moon's southern region (an area of high scientific interest) using a single lunar satellite without a dedicated navigation payload. The rover is assumed to rely on only the Doppler shift of the satellite downlink signal, along with its onboard odometry measurements, to estimate its position over time.

1.1 | Related Works and Limitations

This work focuses on improving the simulation fidelity and state estimation for single-satellite Doppler-based localization, with an emphasis on satellite ephemeris uncertainty and relativistic effects. Pöhlmann et al. (2025) have incorporated temporally correlated error models in lunar navigation and argue that temporal correlations are frequently neglected in existing works. Their evaluation emphasizes hybrid architectures with more measurement diversity, including multiple satellites and a lunar reference station. Motivated by this gap, we focus on the single-satellite case, where observability is limited and time correlation in the errors can dominate estimator behavior.

Recent work has also highlighted the need for well-defined lunar time systems, such as Lunar Coordinate Time (TCL), and for consistent conversions between Earth-based coordinate times and lunar time scales. Kopeikin and Kaplan (2024) and Bourgoin et al. (2026) provide formulations for these transformations. Iiyama and Gao (2026) incorporate relativistic modeling when performing orbit determination and time synchronization (ODTS) of a lunar satellite using weak GNSS signals. However, relativistic modeling is still commonly neglected in rover localization studies, even though deterministic relativistic rate terms can couple clock and orbital states in estimation. We address this limitation in the literature by incorporating relativistic corrections.

In prior work by the authors, we investigated rover localization under the same architecture, but with simplified noise models and an anchored batch estimation approach (Coimbra & Gao, 2025b; Coimbra et al., 2025). In this work, we improve the fidelity of the simulated error sources and introduce a tightly coupled filtering and smoothing framework that can process these time-varying effects.

1.2 | Proposed Work

In prior work, we added white noise to rover velocity to model odometry noise, modeled clock effects using a simplified constant bias, and simulated ephemeris uncertainty using white noise perturbations on satellite position and velocity. We also assumed an Earth-based time reference and did not model TCL conversions or the corresponding deterministic relativistic rate terms. With higher fidelity modeling, these assumptions can lead to estimator inconsistency and can mask time-correlated effects that are important in low-observable, single-satellite regimes.

In this work, we refine the simulation fidelity by modeling rover odometry as a noisy incremental displacement measurement. The rover and satellite clocks are modeled using 2-state stochastic dynamics that propagate clock bias and drift. Satellite ephemeris errors are modeled as a time-correlated process with periodic ODTS updates. We also incorporate relativistic clock rate effects and light-time delay in the measurement simulation and estimation framework. Given these time-varying error sources, we move from an anchored batch approach to a tightly coupled Extended Kalman Filter (EKF) with Rauch-Tung-Striebel (RTS) backward smoothing as a post-processing step.

Figure 1 organizes the simulation architecture as a set of high-level modules. The physics module includes rover and satellite dynamics, clock dynamics, and spacetime effects, including relativistic effects and light-time delay. The measurement module includes the measurements available in this limited infrastructure setting: rover odometry, satellite broadcast ephemeris, and observed Doppler measurements. The Doppler measurements differ from an idealized model by accounting for spacetime effects, thermal noise, and clock bias and drift. In the EKF module, the prediction and update blocks emphasize the design choices specific to this scenario. The EKF runs in real-time, whereas the RTS smoother is applied offline in a backward pass to refine the state estimates. Since each module contains stochastic elements, we evaluate the architecture through Monte Carlo simulation.

1.3 | Contributions

The new and innovative aspects of this work are detailed below. This work is adapted from a paper submitted to the Institute of Navigation GNSS+ 2026 conference (Coimbra & Gao, 2026).

- **Effect of time-correlated ephemeris errors for a single satellite on rover localization:** To the best of the authors' knowledge, this work is the first to quantify how time-correlated ephemeris errors degrade single-satellite Doppler-only rover localization performance and to demonstrate mitigation using a tightly coupled EKF with a correlated bias state and RTS smoothing.
- **Relativistic effects on lunar satellite-rover interaction:** To the best of the authors' knowledge, this work is the first to incorporate a consistent proper time and coordinate time treatment for both measurement simulation and estimation for rover localization using a single satellite and to quantify the localization impact when these effects are omitted.

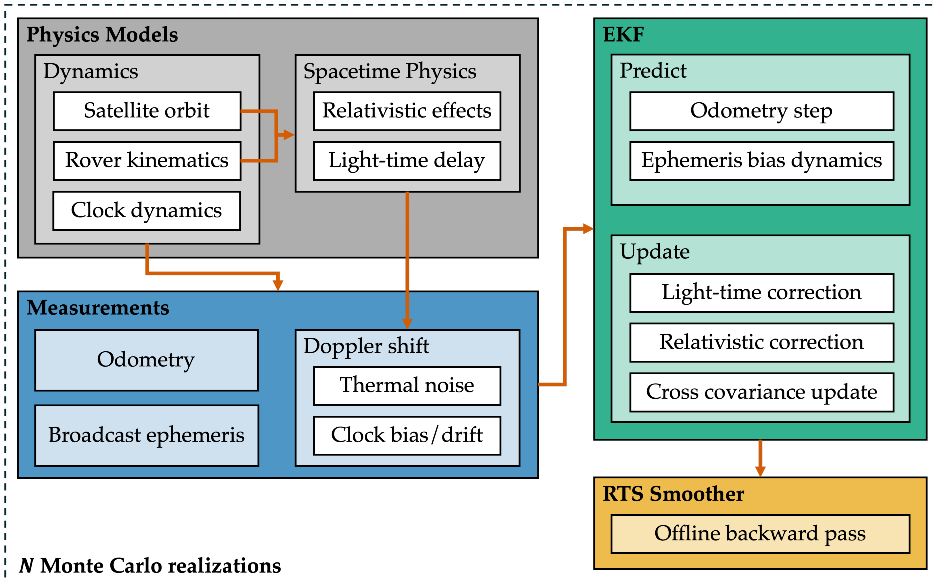

FIGURE 1 Diagram of the simulation architecture.

1.4 | Paper Organization

The paper first discusses how the dynamics, physics, and measurements are modeled in Section 2. In Section 3, the design of the state estimation architecture is described. Then, in Section 4, we define the values and references that are used to test our simulation pipeline for the mission scenario. Section 5 shows what ephemeris and relativistic correction errors arise from the values used in the simulation set-up, as well as the localization performance results from an ablation study on the EKF's design. Finally, we show the key result of the paper, which is the localization performance using the full EKF model and a backward pass through RTS smoothing.

2 | MODELING

A primary contribution of this work is improved modeling fidelity in contrast to prior works. Specifically, we improve the modeling of clock dynamics, spacetime physics, odometry, and ephemeris errors. For modeling aspects that are not primary contributions of this work, we briefly introduce the main concepts and refer the reader to prior studies for additional details.

2.1 | Ground-Truth Dynamics

The ground-truth dynamics include the rover's kinematics, the satellite's orbit propagation, and the clock dynamics for both assets, all of which produce the trajectories and timing references used throughout the rest of the simulation.

2.1.1 | Rover and Satellite Dynamics

The rover kinematic state is parametrized by $\left[\phi, \lambda, \theta, s\right]^\top$, where (ϕ, λ) are the latitude-longitude coordinates, θ is the heading in the East-North-Up (ENU) frame, and s is the rover's speed. In this work, we model the Moon as a smooth sphere, as multipath and other surface effects are outside the scope of this study. Rover-surface interactions, such as wheel slip, are modeled as noise in the odometry measurements. Thus, the rover kinematic state is propagated using spherical kinematics with control inputs of heading rate and linear acceleration. The state is transformed from the local ENU frame to the Moon Principal Axis (PA) frame for consistency between rover and satellite frames. The propagation and transformation equations can be found in Section III.6 a) in Coimbra and Gao (2025b).

The satellite dynamics state $\left[\mathbf{r}_{\boldsymbol{s}}, \mathbf{v}_{\boldsymbol{s}}\right]^{\top}$ is numerically propagated in the Moon-centered inertial (MCI) frame and later converted to the Moon PA frame by the LuPNT simulator (Casadesus Vila et al., 2025). The simulator propagates the satellite state under an N-body dynamical model, with the lunar gravitational field represented by a 100×100 spherical harmonic expansion and the Earth and Sun modeled as point masses. Atmospheric drag and solar radiation pressure are neglected.

2.1.2 | Clock Stochastic Model

Previously, the clock dynamics were modeled as a zero-mean white Gaussian noise with constant variance that was determined by the power spectral density (PSD) coefficients of a specified clock (Coimbra et al., 2025). In this work, the satellite and rover clocks independently follow a 2-state stochastic model that propagates clock bias and drift. We represent the clock state $\mathbf{x}_{\mathrm{clk}}$ for both assets as follows.

$$\mathbf{x}_{\mathrm{clk}} = \begin{bmatrix} \delta t \\ \dot{\delta t} \end{bmatrix} \quad , \tag{1}$$

where δt is clock bias with units of seconds and $\dot{\delta t}$ is clock drift with units of s/s. The clock state is initialized to $\mathbf{0}$.

The clock state is propagated using a transition matrix $\boldsymbol{\Phi}_{\mathrm{clk}}$ and a noise term $\mathbf{w}_{\mathrm{clk}}$ that has a process noise covariance $\mathbf{C}_{\mathrm{clk}}$. The clock process noise $\mathbf{C}_{\mathrm{clk}}$ is parameterized by diffusion coefficients and the time step interval Δt. The diffusion coefficients q_1 and q_2 correspond to the frequency white noise and frequency random walk, respectively (Long & Stacey, 2025).

$$\mathbf{x}_{\mathrm{clk},k+1} = \boldsymbol{\Phi}_{\mathrm{clk}}(\Delta t_k)\mathbf{x}_{\mathrm{clk},k} + \mathbf{w}_{\mathrm{clk},k} \quad , \quad \mathbf{w}_{\mathrm{clk},k} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{\mathrm{clk}}(\Delta t_k)) \tag{2}$$

$$\boldsymbol{\Phi}_{\mathrm{clk}}(\Delta t) = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \tag{3}$$

$$\mathbf{C}_{\mathrm{clk}}(\Delta t) = \begin{bmatrix} q_1 \Delta t + q_2 \frac{\Delta t^3}{3} & q_2 \frac{\Delta t^2}{2} \\ q_2 \frac{\Delta t^2}{2} & q_2 \Delta t \end{bmatrix} \tag{4}$$

2.2 | Spacetime Physics Modeling

Compared with our prior work (Coimbra & Gao, 2025b; Coimbra et al., 2025), this study incorporates relativistic effects and light-time delay into the ground-truth simulation and applies the corresponding corrections within the filter. Although relativistic effects are often neglected in lunar navigation analyses, recent work has established mathematical frameworks for modeling these effects in the lunar regime, making their incorporation into lunar navigation studies more tractable (Bourgoin et al., 2026; Kopeikin & Kaplan, 2024).

2.2.1 | Reference Frames and Coordinate Times

In 2024, the International Astronomical Union (IAU) formally established the lunar celestial reference system (LCRS) and the associated lunar coordinate time (TCL) as a standard framework for lunar applications (International Astronomical Union, 2024). LCRS and TCL are the lunar analogs of the geocentric celestial reference system (GCRS) and geocentric coordinate time (TCG) used for Earth-based applications. Accordingly, the unit of TCL is consistent with the SI second (International Astronomical Union, 2024). As mentioned in Section 2.1.1, satellite dynamics are propagated in the MCI frame. This is a Newtonian approximation to the LCRS: the origin and nominal axis orientation match LCRS, but the equations of motion omit the post-Newtonian metric corrections and use International Atomic Time (TAI) rather than TCL as the time argument. At the accuracy relevant for meter-level Doppler-based navigation over a 24-hour simulation window, the effects of using a true realization of LCRS are negligible compared to the dominant error sources in the measurement model. We confirm this claim by evaluating the leading first post-Newtonian (1PN) acceleration corrections on the lunar satellite (Sośnica et al., 2021), which include the Schwarzschild, Lense-Thirring, and de Sitter effects. Integrated over 24 hours, these terms produce satellite position perturbations of order 1 cm or less and a conservative range rate error below 8 μm/s. The separate time-scale effect from using TAI rather than TCL produces a secular range-rate-equivalent drift below 0.06 mm/s. These errors are well below the modeled measurement errors, including the ODTS ephemeris velocity residual of order 1 mm/s and the Doppler thermal noise floor of order 10^1–10^2 mm/s. We therefore find that a 1PN LCRS realization is unnecessary for the meter-level Doppler navigation analysis considered in this study.

TAI is also the natural time argument for the simulation since the underlying SPICE planetary ephemerides are tagged in Earth-based time scales, so re-tagging every ephemeris query into TCL would shift the same conversion elsewhere in the pipeline without changing the underlying physics. Relativistic clock-rate effects for the satellite and rover are instead modeled separately in the clock dynamics and enter the Doppler measurement model directly. Recent work has derived the mathematical relationships between geocentric and barycentric coordinate times and TCL, as well as the relationship between the proper time of a lunar clock and TCL (Bourgoin et al., 2026; Kopeikin & Kaplan, 2024).

Proper time is the time measured by an idealized clock along its own worldline, which is the path that the clock follows through four-dimensional spacetime. By contrast, coordinate time is the time parameter associated with a specific reference system and is thus frame-dependent but not path-dependent. In navigation applications, the observables are tied to the proper times of individual assets. However, because each clock's rate depends on its position and velocity, the proper times of different clocks must be related to a common coordinate time for consistent comparison. In this study, we therefore compute the rate difference between the proper time of each asset and TCL.

2.2.2 | Relativistic Effects

The dominant relativistic effects are incorporated into the clock and measurement simulation. The proper times of the rover and satellite clocks are computed from each clock's position in the lunar gravitational potential and the special-relativistic time dilation associated with its velocity relative to the inertial frame. Shapiro delay is a gravitational signal propagation delay that becomes significant when the signal path passes close to a massive body, particularly the Sun (Shapiro, 1964). Because the rover-to-satellite link considered in this scenario is much shorter than an Earth-to-Moon link and does not pass near the Sun, we neglect Shapiro delay in our formulation. Kopeikin and Kaplan (2024) provide the transformation between the proper time of a lunar clock τ_L and $s :=$ TCL as follows.

$$\tau_L = s - \frac{1}{c^2}\int_{s_0}^{s}\left(\frac{\dot{r}_L}{2} + U_L(r_L) + V_L^{\mathrm{tide}}\right) ds \quad , \tag{5}$$

where $\dot{r}_L$ is the velocity of the lunar clock in the LCRS, $U_L(r_L)$ is the lunar gravitational potential at the clock's distance from the lunar center of mass r_L, and V_L^{tide} is the tidal gravitational potential of external bodies.

In our simulation, the rover traverses less than 12 km. Over this distance, the variations in lunar gravitational potential due to neglected higher-degree gravity terms produce clock rate perturbations that are orders of magnitude below the Doppler measurement noise floor. For the satellite clock-rate correction, neglected higher-degree terms are further attenuated with altitude, making their contribution even smaller. It is therefore permissible to use a lunar mass-monopole approximation for the rover and satellite clock-rate correction in this scenario: $U_L(r_L) \simeq \mu_L/r_L$, where μ_L is the Moon's gravitational parameter.

According to Kopeikin and Kaplan (2024), the tidal potential contribution to the secular drift of a lunar clock is 8.7×10^{-8} μs/day from the Sun and 1.5×10^{-5} μs/day from the Earth. Given that a stationary clock on the lunar surface experiences a secular drift of 2.7 μs/day, we consider the lunar gravitational potential term to be dominant and neglect the tidal potential term in this work. Thus, the rate difference $\dot{\delta t}_{\mathrm{rel}}$ between the proper time of an asset in the lunar regime and TCL can be simplified as follows.

$$\dot{\delta t}_{\mathrm{rel}} = \frac{d(\tau_L - s)}{ds} = -\frac{1}{c^2}\left(\frac{\dot{r}_L}{2} + \frac{\mu_L}{r_L}\right) \tag{6}$$

A fully consistent simulation would also require converting the clock state from TCL to TAI to maintain consistency with Earth operations and with the fact that the satellite orbit propagation is performed in TAI. However, the TCL-TAI conversion is identical for the satellite and rover clocks, since this common time-scale transformation does not depend on the asset-specific clock trajectory once the clock has been referred to TCL (Kopeikin & Kaplan, 2024). The truth model retains this common-mode term for physical realism, but it cancels in the Doppler measurement model when the satellite and rover clock drifts are differenced, so the details of the TCL-TAI conversion have no effect on the EKF residuals and are omitted from this work.

Since the Doppler measurement is sensitive to clock drift rather than clock bias, we retain the relativistic correction only in the drift component. As shown in Section 2.1.2, the clock bias state is still propagated because the 2-state clock dynamics couples bias and drift through the process covariance $\mathbf{C}_{\mathrm{clk}}$. However, we do not model a separate deterministic relativistic correction to clock bias since clock bias does not enter the measurement model. Under this approximation, the clock drift of each asset is represented as follows, where $\dot{\delta t}_{\mathrm{stoch}}$ refers to the stochastic part of the clock drift $\dot{\delta t}$ as defined in Equation 1.

$$\dot{\delta t} = \dot{\delta t}_{\mathrm{stoch}} + \dot{\delta t}_{\mathrm{rel}} \tag{7}$$

2.2.3 | Light-Time Delay

When the rover receives a Doppler measurement, the reception time t_{rx} is not equal to the transmission time t_{tx}. Therefore, to evaluate the Doppler measurement accurately, we must estimate the satellite position at the time of transmission, $\mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{tx}})$, using the information available at the time of reception t_{rx}. We define the light-time delay as the one-way propagation time from satellite transmission to rover reception, i.e., $\tau_{\mathrm{lt}} = t_{\mathrm{rx}} - t_{\mathrm{tx}}$. We approximate the light-time delay τ_{lt} as the instantaneous range at reception divided by the speed of light c.

$$\tau_{\mathrm{lt}} = \frac{\|\mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{tx}}) - \mathbf{r}_{\boldsymbol{r}}(t_{\mathrm{rx}})\|}{c} \simeq \frac{\|\mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{rx}}) - \mathbf{r}_{\boldsymbol{r}}(t_{\mathrm{rx}})\|}{c} \tag{8}$$

The satellite position at transmission time is then computed using a first-order light-time correction. Although $\mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{tx}})$ could instead be solved iteratively (which is the approach used for ground-truth modeling), we find the first-order correction reproduces the iterative transmission-time position to within 20 μm, which is far below the meter-level navigation accuracy desired. This correction uses the inertial time derivative of the satellite position in Moon PA coordinates. The derivative is expressed as the sum of the satellite velocity $\mathbf{v}_{\boldsymbol{s}}$ in the Moon PA frame and the contribution from the rotation of the Moon PA frame, where $\boldsymbol{\omega}_L$ is the Moon's angular rotation. We further assume that the light-time delay τ_{lt} is sufficiently small such that $\mathbf{v}_{\boldsymbol{s}}(t_{\mathrm{tx}}) \simeq \mathbf{v}_{\boldsymbol{s}}(t_{\mathrm{rx}})$. Under these assumptions, the satellite position at transmission time is computed as follows.

$$\mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{tx}}) = \mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{rx}}) - \tau_{\mathrm{lt}} \cdot (\mathbf{v}_{\boldsymbol{s}}(t_{\mathrm{rx}}) + \boldsymbol{\omega}_L \times \mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{rx}})) \tag{9}$$

In later sections, we use the superscript * to denote states corrected for light-time delay to the first order, such as $\mathbf{r}^*_{\boldsymbol{s}}(t) := \mathbf{r}_{\boldsymbol{s}}(t_{\mathrm{tx}})$.

2.3 | Measurement Modeling

The minimal infrastructure scenario described in this paper is unique in that the setting is confined to very few available measurements. In this section, we describe the models used to generate the measurements that the rover uses for localization: odometry, broadcast ephemeris from the satellite, and the observed Doppler shift from the communication signal of the satellite.

2.3.1 | Odometry

For rover motion sensing, we assume onboard sensors such as inertial sensors, gyroscopes, LiDAR, and wheel encoders provide knowledge of incremental rover motion. We aggregate these effects into an odometry measurement model in which we add white noise to the incremental distance traversed between consecutive states. The noise standard deviation is proportional to the true delta distance traveled over the step Δd_{true}, where α_{odo} is the proportionality constant.

$$\Delta d_{\mathrm{meas}} = \Delta d_{\mathrm{true}} + \varepsilon_{\mathrm{odo}} \quad , \quad \varepsilon_{\mathrm{odo}} \sim \mathcal{N}(0, \sigma^2_{\mathrm{odo}}) \tag{10}$$

$$\sigma_{\mathrm{odo}} = \alpha_{\mathrm{odo}} \cdot \Delta d_{\mathrm{true}} \tag{11}$$

2.3.2 | Broadcast Ephemeris

In addition to odometry, the rover receives broadcast ephemeris from the satellite. The broadcast satellite state available to the rover is assumed to be imperfect. The broadcast satellite position $\tilde{\mathbf{r}}_{\boldsymbol{s}}$ and velocity $\tilde{\mathbf{v}}_{\boldsymbol{s}}$ are modeled as the sum of the true satellite state, propagated according to the dynamics described in Section 2.1.1, and error terms in position $\boldsymbol{\delta r_s}$ and velocity $\boldsymbol{\delta v_s}$.

$$\tilde{\mathbf{r}}_{\boldsymbol{s}} = \mathbf{r}_{\boldsymbol{s}} + \boldsymbol{\delta r_s} \tag{12}$$

$$\tilde{\mathbf{v}}_{\boldsymbol{s}} = \mathbf{v}_{\boldsymbol{s}} + \boldsymbol{\delta v_s} \tag{13}$$

These ephemeris error states are modeled using a time-correlated stochastic process. Pöhlmann et al. (2025) noted that most prior lunar PNT studies model satellite-related measurement errors as white Gaussian noise and introduced temporally correlated models for pseudorange and pseudorange rate biases. Since the measurement architecture in this work is based on one-way Doppler, the same modeling principle is adopted, but implemented in Cartesian ephemeris error states rather than directly in the measurement bias space.

Specifically, we apply a first-order Gauss-Markov (FOGM) process to the satellite acceleration error because ephemeris errors arise from unmodeled forces that act through acceleration. This formulation is commonly referred to as a Singer model, where acceleration is represented as an exponentially correlated random process (Singer, 1970). The ephemeris error state $\mathbf{x}_{\text{eph}}$ consists of the satellite position error $\boldsymbol{\delta r_s}$, velocity error $\boldsymbol{\delta v_s}$, and acceleration error $\boldsymbol{\delta a_s}$ in all three Cartesian axes.

$$\mathbf{x}_{\text{eph}} = \begin{bmatrix} \boldsymbol{\delta r_s} & \boldsymbol{\delta v_s} & \boldsymbol{\delta a_s} \end{bmatrix}^\top \tag{14}$$

Similar to the clock stochastic model, we propagate the ephemeris error state $\mathbf{x}_{\text{eph}}$ using a discrete-time state transition $\mathbf{\Phi}_{\text{eph}}$ and process noise $\mathbf{w}_{\text{eph}}$ defined by the Singer model to generate the ground truth.

$$\mathbf{x}_{\text{eph},k+1} = \mathbf{\Phi}_{\text{eph}}(\Delta t_k)\mathbf{x}_{\text{eph},k} + \mathbf{w}_{\text{eph},k} \quad , \quad \mathbf{w}_{\text{eph},k} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{\text{eph}}(\Delta t_k)) \tag{15}$$

The state transition $\mathbf{\Phi}_{\text{eph}}$ is a function of the time step Δt and the correlation time constant τ_{eph}.

$$\mathbf{\Phi}_{\text{eph}}(\Delta t) = \begin{bmatrix} 1 & \Delta t & \tau_{\text{eph}}^2 \left(e^{-\Delta t/\tau_{\text{eph}}} + \Delta t/\tau_{\text{eph}} - 1\right) \\ 0 & 1 & \tau_{\text{eph}} \left(1 - e^{-\Delta t/\tau_{\text{eph}}}\right) \\ 0 & 0 & e^{-\Delta t/\tau_{\text{eph}}} \end{bmatrix} \tag{16}$$

Since the chosen parameters result in small $\Delta t/\tau_{\text{eph}}$, we use a Taylor series approximation to compute the process noise covariance $\mathbf{C}_{\text{eph}}$. The covariance $\mathbf{C}_{\text{eph}}$ is used to define the process noise $\mathbf{w}_{\text{eph}}$ as shown in Equation 15 and is also a function of the steady-state acceleration variance $\sigma_{e,a}^2$.

$$\mathbf{C}_{\text{eph}}(\Delta t) \simeq \lim_{\Delta t/\tau_{\text{eph}} \to 0} \mathbf{C}_{\text{eph}} = \frac{2\sigma_{e,a}^2}{\tau_{\text{eph}}} \begin{bmatrix} \Delta t^5/20 & \Delta t^4/8 & \Delta t^3/6 \\ \Delta t^4/8 & \Delta t^3/3 & \Delta t^2/2 \\ \Delta t^3/6 & \Delta t^2/2 & \Delta t \end{bmatrix} \tag{17}$$

Under this model, the expected ephemeris error growth is unbounded in the absence of ODTS corrections. We therefore assume the satellite receives periodic ODTS updates from an Earth ground station or other lunar asset, and we model these corrections as resets at a fixed interval. Rather than resetting each error state exactly to zero, we sample each axis independently using reset variances σ_{reset}^2. No time derivative constraints are imposed at the reset instant, but the subsequent FOGM propagation restores dynamical consistency. We assume that the rover has exact knowledge of when these ODTS resets take place.

2.3.3 | Observed Doppler Shift

The third and final measurement that the rover receives is the observed Doppler shift. In practice, we formulate the measurement model in terms of pseudorange rate because Doppler shift is simply a known scalar multiple of pseudorange rate for a given carrier frequency. Using the ground-truth dynamics and spacetime physics models described in Sections 2.1 and 2.2, we synthesize these components into the pseudorange rate form of the Doppler measurement model at reception time t, as shown in Equation 18.

In the measurement model, all satellite states are evaluated at transmission time using light-time correction, and all rover states are evaluated at reception time. The rover and satellite clock drifts, $\dot{\delta t}_r$ and $\dot{\delta t}_s$, each include both stochastic and deterministic relativistic components, as described in Equation 7.

$$\tilde{\dot{\rho}}(t) = (\mathbf{v}_s^* - \mathbf{v}_r) \cdot \frac{\mathbf{r}_s^* - \mathbf{r}_r}{\|\mathbf{r}_s^* - \mathbf{r}_r\|} + c\left(\dot{\delta t}_r - \dot{\delta t}_s\right) + \varepsilon_{\dot{\rho}} \quad , \quad \varepsilon_{\dot{\rho}} \sim \mathcal{N}(0, \sigma_t^2) \tag{18}$$

The additive noise term $\varepsilon_{\dot{\rho}}$ has variance σ_t^2, which represents the thermal noise associated with carrier tracking. The thermal noise variance σ_t^2 is determined from the link budget through the carrier-to-noise density ratio C/N_0. Details of the thermal noise and link budget formulation can be found in Coimbra and Gao (2025a).

3 | STATE ESTIMATION

Prior works evaluated the single-satellite scenario using a weighted batch filter to localize the rover (Coimbra & Gao, 2025b). However, that formulation assumed near-perfect knowledge of the rover's motion and represented clock drift as a single constant. This work adopts a richer error model including FOGM-driven ephemeris errors with ODTS resets, stochastic clock drift, and step-correlated odometry noise. The EKF naturally accommodates these time-correlated processes by augmenting the state to track their evolution and propagating covariance forward at each step. We also investigate the post-processing improvements from RTS backward smoothing, which like a batch filter, refines each epoch's estimate using the entire measurement history.

3.1 | Extended Kalman Filter

The single-satellite measurement geometry leads to a weakly observable estimation problem. This setting motivates a low-dimensional state representation and careful treatment of the dominant error sources. In this section, we present the selected state representation and the corresponding EKF predict and update equations. The resulting formulation captures the dominant measurement effects without unnecessarily increasing the estimator dimension.

3.1.1 | State

In prior work, the rover state estimate included the rover's 3D position and a clock drift term (Coimbra & Gao, 2025b). In this work, to account for biases induced by time-correlated ephemeris errors, we augment the rover state to include a scalar bias term δt_{eph}. The EKF treats the satellite ephemeris as a known input, so unmodeled satellite ephemeris errors enter the Doppler observable as an effective time-correlated measurement bias. Modeling the full satellite error state would unnecessarily increase the estimator dimension in a setting where measurements from a single satellite already provide weak geometric observability. Instead, we model the dominant effect using a single bias state with FOGM dynamics.

Thus, the 5D estimated state $\mathbf{x}$ is composed of the rover's 3D position $\mathbf{r}_{\boldsymbol{r}}$, a relative clock drift term $d_{\text{clk}} := c(\dot{\delta t}_r - \dot{\delta t}_s)$, and the scalar bias term δt_{eph} as follows.

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_{\boldsymbol{s}} \\ d_{\text{clk}} \\ \delta t_{\text{eph}} \end{bmatrix}^{\top} \in \mathbb{R}^5 \tag{19}$$

Under the ablation study where we investigate the effects of not using any satellite-based measurements, the relative clock drift term simplifies to be only the rover's clock drift, $d_{\text{clk}} := c\dot{\delta t}_r$.

3.1.2 | Predict

The predicted state estimate is computed based on the dynamics of each term. Odometry is used as a control input to propagate the rover position, so the rover position is updated using the measured delta distance Δd_{meas} and the rover heading $\mathbf{u}_{\text{hdg}}$, which is recomputed from the current estimated position at every predict step. The relative clock drift d_{clk} is a stochastic process, so its next estimate is its previous estimate. We previously discussed that the ephemeris error is encapsulated by a slowly varying, time-correlated error, or a bias. Carpenter and D'Souza (2025) state that the best practice for bias modeling for sequential navigation filters, such as the EKF, is to use a FOGM process. The FOGM equations are used for the ephemeris bias term δt_{eph} in Equations 20 and 24.

$$\hat{\mathbf{x}}_{k+1|k} = \begin{cases} \hat{\mathbf{r}}_{\boldsymbol{s},k|k} + \Delta d_{\text{meas},k} \mathbf{u}_{\text{hdg},k} \\ \hat{d}_{\text{clk},k|k} \\ e^{-\Delta t/\tau_{\text{eph}}} \hat{\delta t}_{\text{eph},k|k} \end{cases} \tag{20}$$

The predicted state error covariance $\mathbf{P}_{k+1|k}$ is computed using standard EKF equations. The state transition $\mathbf{F}$ is the Jacobian of the state presented in Equation 20. The process noise covariance $\mathbf{Q}$ is composed of the process noise from the 3D rover position $\mathbf{Q}_{\text{pos},k}$, the relative clock drift $Q_{\text{clk},k}$, and the ephemeris bias term $Q_{\text{eph},k}$.

$$\mathbf{P}_{k+1|k} = \mathbf{F}_k \mathbf{P}_{k|k} \mathbf{F}_k^{\top} + \mathbf{Q}_k \tag{21}$$

$$\mathbf{F}_k = \operatorname{diag}\left(\mathbf{1}_4, e^{-\Delta t/\tau_{\text{eph}}}\right) \tag{22}$$

$$\mathbf{Q}_k = \operatorname{blkdiag}\left(\mathbf{Q}_{\text{pos},k}, Q_{\text{clk},k}, Q_{\text{eph},k}\right) \tag{23}$$

Because Doppler measurements from a single satellite provide limited geometric observability, the filter can become overconfident in directions that are weakly constrained by the measurement. To prevent inconsistency, we inject additional process noise in the cross-track direction into $\mathbf{Q}_{\text{pos},k}$ as a conservative model for unmodeled lateral effects, such as heading uncertainty and surface slip, and to maintain realistic covariance growth in weakly observable directions. To define $Q_{\text{clk},k}$, the drift process noise variance rate for a single asset is $\dot{Q} = c^2 q_2$ (Long & Stacey, 2025). Since the clock variances are independent, we can add the contributions from the rover and satellite as shown in σ_{clk}. The process noise for the ephemeris bias $Q_{\text{eph},k}$ follows the FOGM dynamics equations.

We manually tune the cross-track noise multiplier α_{crs} and the steady-state variance in the ephemeris bias process $\sigma^2_{\text{eph,ss}}$.

$$\begin{aligned}
\mathbf{Q}_{\text{pos},k} &= \sigma^2_{\text{odo},k}(\mathbf{u}_{\text{hdg},k}\mathbf{u}^\top_{\text{hdg},k}) + \sigma^2_{\text{crs},k}(\mathbf{u}_{\text{crs},k}\mathbf{u}^\top_{\text{crs},k}) \quad &, \quad \sigma_{\text{crs},k} = \alpha_{\text{crs}}\sigma_{\text{odo},k} \\
Q_{\text{clk},k} &= \sigma^2_{\text{clk}}\Delta t \quad &, \quad \sigma_{\text{clk}} = c\sqrt{q_{2,\text{r}} + q_{2,\text{s}}} \\
Q_{\text{eph},k} &= \sigma^2_{\text{eph,ss}}\left(1 - e^{-2\Delta t/\tau_{\text{eph}}}\right)
\end{aligned} \tag{24}$$

As mentioned previously, we assume the rover has knowledge of ODTS update times. Thus, when an ODTS reset occurs, the final index that corresponds to the bias term δt_{eph} in the predicted state estimate $\hat{\mathbf{x}}_{k+1|k}$ in Equation 20 and in the state transition $\mathbf{F}_k$ in Equation 22 are set to 0. We reduce the bias covariance at the ODTS reset instant to a constant variance $\sigma^2_{b,\text{reset}}$ to reflect improved ephemeris knowledge after correction.

3.1.3 | Update

For the odometry-only ablation study, we run the EKF using only the state propagation defined by the predict step. When applying Doppler measurements, we enable the EKF to also perform the update step using Doppler shift measurements. At each epoch, the filter receives the broadcast ephemeris, which is interpreted as the noisy satellite state $(\tilde{\mathbf{r}}_{\boldsymbol{s}}, \tilde{\mathbf{v}}_{\boldsymbol{s}})$ and the pseudorange rate observables $\tilde{\dot{\rho}}$.

The update step mirrors the chain of operations for ground truth generation. Prior to using the standard EKF update equations, the spacetime physics corrections are applied. First, the noisy satellite states are light-time corrected as shown in Equation 9 based on the filter's noisy knowledge of the satellite states and the estimated rover state. Next, we require an estimate of the rover velocity. We neglect to include the rover velocity in our state $\mathbf{x}$ in order to prevent the state from becoming too high-dimensional with respect to the observability. Instead, we estimate the rover velocity using the odometry measurements.

$$\hat{\mathbf{v}}_{\boldsymbol{r},k} = \frac{\Delta d_{\text{meas},k}}{\Delta t}\mathbf{u}_{\text{hdg},k} \tag{25}$$

The predicted pseudorange rate measurement, with first-order light-time corrections, is as follows.

$$\hat{\dot{\rho}}_k(t) = (\tilde{\mathbf{v}}^*_{\boldsymbol{s}} - \hat{\mathbf{v}}_{\boldsymbol{r},k}) \cdot \frac{\tilde{\mathbf{r}}^*_{\boldsymbol{s}} - \hat{\mathbf{r}}_{\boldsymbol{r},k}}{\|\tilde{\mathbf{r}}^*_{\boldsymbol{s}} - \hat{\mathbf{r}}_{\boldsymbol{r},k}\|} + \hat{\dot{d}}_{\text{clk},k|k-1} + \hat{\dot{\delta t}}_{\text{eph},k|k-1} \tag{26}$$

The deterministic relativistic correction η_{rel} is applied when computing the measurement residual $\mathbf{y}_k$. The rate difference $\dot{\delta t}_{\text{rel}}$ is defined in Equation 6 and is calculated using the filter's knowledge of each asset's states.

$$\mathbf{y}_k = \tilde{\dot{\rho}} - (\hat{\dot{\rho}}_k + c\,\eta_{\text{rel},k}) \tag{27}$$

$$\eta_{\text{rel},k} := \dot{\delta t}_{\text{rel}}(\hat{\mathbf{r}}_{\boldsymbol{r},k}, \hat{\mathbf{v}}_{\boldsymbol{r},k}) - \dot{\delta t}_{\text{rel}}(\tilde{\mathbf{r}}^*_{\boldsymbol{s}}, \tilde{\mathbf{v}}^*_{\boldsymbol{s}}) \tag{28}$$

In Equation 25, we introduced the odometry process noise into the measurement noise when we defined the rover velocity estimate $\hat{\mathbf{v}}_{\boldsymbol{r},k}$ in terms of the measured delta distance Δd_{meas}. Thus, the covariance update must consider correlation between the process and measurement noise in this system. Using Grewal and Andrews (2001), let us define the cross-covariance matrix $\mathbf{C}_k$ as shown in Equation 29. The process noise $\mathbf{w}_k$ here refers only to the component that is cross-correlated with the measurement noise, which is the odometry noise ε_{odo} (defined in Equation 10) projected in the heading direction. The process noise $\mathbf{w}_k$ is defined as the residual that reconciles the filter's deterministic propagation with the ground truth motion, which results in the negative sign. Likewise, the measurement noise $\mathbf{v}_k$ contains the cross-correlated contribution from the odometry-derived velocity estimate $\hat{\mathbf{v}}_{\boldsymbol{r},k}$, defined in Equation 25. The Doppler measurement Jacobian with respect to rover velocity is the negative line-of-sight unit vector from rover to satellite $-\mathbf{e}^\top_{\boldsymbol{r}}$. To extract the component of the velocity error that affects the scalar measurement, we project the velocity vector onto the measurement's gradient with respect to velocity.

$$\mathbf{C}_k := \mathbb{E}[\mathbf{w}_k\mathbf{v}^\top_k] \tag{29}$$

$$\mathbf{w}_k = -\varepsilon_{\text{odo},k}\mathbf{u}_{\text{hdg},k} \tag{30}$$

$$\mathbf{v}_k = \frac{\varepsilon_{\text{odo},k}}{\Delta t}\left(\mathbf{e}_{\boldsymbol{r}} \cdot \mathbf{u}_{\text{hdg},k}\right) \tag{31}$$

For cross-correlated noise, the measurement update equations are as follows (Grewal & Andrews, 2001).

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \mathbf{y}_k \tag{32}$$

$$\mathbf{K}_k = \left(\mathbf{P}_{k|k-1}\mathbf{H}_k^\top + \mathbf{C}_k\right)\left(\mathbf{H}_k^\top \mathbf{P}_{k|k-1}\mathbf{H}_k^\top + R_k + \mathbf{H}_k\mathbf{C}_k + \mathbf{C}_k^\top \mathbf{H}_k^\top\right)^{-1} \tag{33}$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k\left(\mathbf{H}_k\mathbf{P}_{k|k-1} + \mathbf{C}_k^\top\right) \quad , \tag{34}$$

where $\mathbf{K}$ is the Kalman gain, $\mathbf{H}$ is the measurement Jacobian, and R is the measurement noise covariance. Note that Equation 34 reduces to the Joseph form when there is no cross-correlation (i.e., $\mathbf{C}_k = \mathbf{0}$). This becomes more obvious in the expanded form of Equation 34, as shown below.

$$\mathbf{L}_k := \mathbf{I} - \mathbf{K}_k\mathbf{H}_k \tag{35}$$

$$\mathbf{P}_{k|k} = \mathbf{L}_k\mathbf{P}_{k|k-1}\mathbf{L}_k^\top + \mathbf{K}_k R_k \mathbf{K}_k^\top - \mathbf{L}_k\mathbf{C}_k\mathbf{K}_k^\top - \mathbf{K}_k\mathbf{C}_k^\top\mathbf{L}_k^\top \tag{36}$$

The measurement noise covariance R is defined as the sum of the measurement variance, which is the thermal noise σ_t defined in Equation 18, and the line-of-sight projected variance from the rover's velocity uncertainty, similar to the formulation of the measurement noise $\mathbf{v}_k$ in Equation 31.

$$R_k = \sigma_t^2 + \left(\frac{\sigma_{\text{odo},k}}{\Delta t}\left|\mathbf{e}_{\boldsymbol{r}} \cdot \mathbf{u}_{\text{hdg},k}\right|\right)^2 \tag{37}$$

3.2 | RTS Backward Smoother

The forward pass of the EKF processes the measurements in real-time. As a post-processing step, we apply a single pass of the RTS backward smoothing (Rauch et al., 1965) to leverage future measurements and refine prior state estimates. The state estimates are computed through a backward pass from $k = N-1$ to $k = 0$ as follows:

$$\mathbf{A}_k := \mathbf{P}_{k|k}\mathbf{F}_k^\top \mathbf{P}_{k+1|k}^{-1} \tag{38}$$

$$\hat{\mathbf{x}}_{k|N} = \hat{\mathbf{x}}_{k|k} + \mathbf{A}_k\left(\hat{\mathbf{x}}_{k+1|N} - \hat{\mathbf{x}}_{k+1|k}\right) \tag{39}$$

$$\mathbf{P}_{k|N} = \mathbf{P}_{k|k} + \mathbf{A}_k\left(\mathbf{P}_{k+1|N} - \mathbf{P}_{k+1|k}\right)\mathbf{A}_k^\top \tag{40}$$

In practice, this smoother can be applied to refine and improve the geotagging accuracy for samples collected along the rover trajectory.

4 | SIMULATION SET-UP

Using the physics models and state estimation framework described in Sections 2 and 3, we use parameters from real mission specifications (or make assumptions if necessary) to validate our filter design for the single-satellite scenario. Specifically, we describe the rover and satellite dynamics as well as the measurement parameters below.

4.1 | Satellite and Rover

To enable direct comparisons, we employ a similar simulation setup and measurement model to prior work. We model the satellite using ESA SSTL Lunar Pathfinder specifications (SSTL, 2022). The satellite follows an elliptical lunar frozen orbit (ELFO) with an orbital period of 10.84 hours. Given that a detailed operations timeline for the satellite is not yet released, we initialize the satellite's mean anomaly to $49°$ at 00:00:00 UTC on 2030 October 1 so that the satellite just enters the visibility region with respect to a user positioned at the Artemis Basecamp ($87.0°$ S, $50.5°$ E). We employ Lunar Pathfinder antenna parameters and assume the satellite antenna points in the nadir direction.

The rover traverses north at a constant 0.5 km/hr, which is an average traverse speed cited for long-range lunar rovers (Keane et al., 2022). For rover signal availability, we apply an elevation mask of 5 degrees and a carrier-to-noise density ratio threshold of 30 dB-Hz, consistent with Nardin et al. (2023) and Melman et al. (2022). This yields an occultation duration of roughly 3 hours per 10.84-hour orbit, when the rover is located at the Artemis Basecamp. We model the rover high-gain omnidirectional antenna using NASA Endurance rover specifications (Keane et al., 2022). We assume the rover initial error is 5 m standard deviation per-axis in position and 1 mm/s standard deviation per-axis in velocity, motivated by human-in-the-loop map matching or stationary position refinement (Coimbra et al., 2025).

As defined in Section 2.1.2, the satellite and rover clock dynamics are defined using a 2-state stochastic model. We assume that the satellite is equipped with a Spectratime Rubidium Atomic Frequency Standard (RAFS), and the rover is equipped with a less accurate chip-scale atomic clock (CSAC). The diffusion coefficients q_1 and q_2 for both clocks are listed in Table 1.

TABLE 1
Satellite (RAFS) and Rover (CSAC) clock parameters (Ashman et al., 2024; Long & Stacey, 2025)

Clock	q_1	q_2
RAFS	3.70×10^{-24}	1.87×10^{-33}
CSAC	5.8×10^{-22}	1.5×10^{-29}

4.2 | Measurement Parameters

For odometry, we apply a noise standard deviation that is equal to 50% of the true delta distance traveled in the ground-truth measurements. Thus, using notation outlined in Section 2.3.1 and 3.1.2, $\alpha_{\text{odo}} = 0.5$ and the manually-tuned cross-track noise multiplier α_{crs} is set to 0.1.

As corroborated by Grenier et al. (2022), achieving accurate broadcast ephemerides in real time from a lunar navigation system is still an active research area. Grenier et al. (2022) state that 1σ orbit predictions for lunar satellites vary from 5 to 50 m. Since the requirements for lunar satellite ephemeris are yet to be standardized across interoperable partners, we apply a conservative steady-state acceleration error standard deviation. The time-correlated nature of the ephemeris error model results in monotonically increasing errors, necessitating ODTS resets. Table 2 outlines the ephemeris error parameters as well as the values used for simulating ODTS resets. For the bias-augmented filter, the standard deviation for the steady-state bias $\sigma_{\text{eph,ss}}$ was set to 5.05 mm/s to equal the empirical root mean square (RMS) of the line-of-sight projection of the true range-rate ephemeris error. The bias, which follows FOGM dynamics as explained in Section 3.1.2, has a correlation time of 2 hours and a reset bias of 1 mm/s, which is equivalent to the truth-side ephemeris error model as shown in Table 2.

TABLE 2
Time parameters and standard deviations (STDs) used to simulate time-correlated ephemeris errors and ODTS resets

Parameter	**Value**
Correlation time	2 hours
Steady-state acceleration error STD per axis	10^{-4} mm/s^2
Reset time interval	6 hours
Reset position error STD per axis	1 m
Reset velocity error STD per axis	1 mm/s
Reset acceleration error STD per axis	10^{-5} mm/s^2

The odometry, Doppler measurements, and filter updates were applied at a 1 Hz frequency. Due to the stochasticity of several error sources, we validate the filter across 100 Monte Carlo realizations over a 24-hour simulation window.

5 | RESULTS

In this section, we first show how the ground truth is characterized in the simulation and highlight how observability can drastically change depending on the rover's location on the Moon. Then, we discuss the results of the ablation study, which quantifies the localization performance impact of making certain modeling assumptions in the filter, such as neglecting relativistic effects. Finally, we show the filter's best performance using both forward and backward passes.

5.1 | Ground Truth Characterization

As mentioned in Section 2.3.2, the rover receives noisy broadcast ephemeris from the satellite. Using the ephemeris parameters from Table 2, we find that the conservative steady-state acceleration error STD results in a position error of up to 80 m, as shown in Figure 2. The velocity error reaches around 4 mm/s at the end of each 6-hour ODTS reset interval.

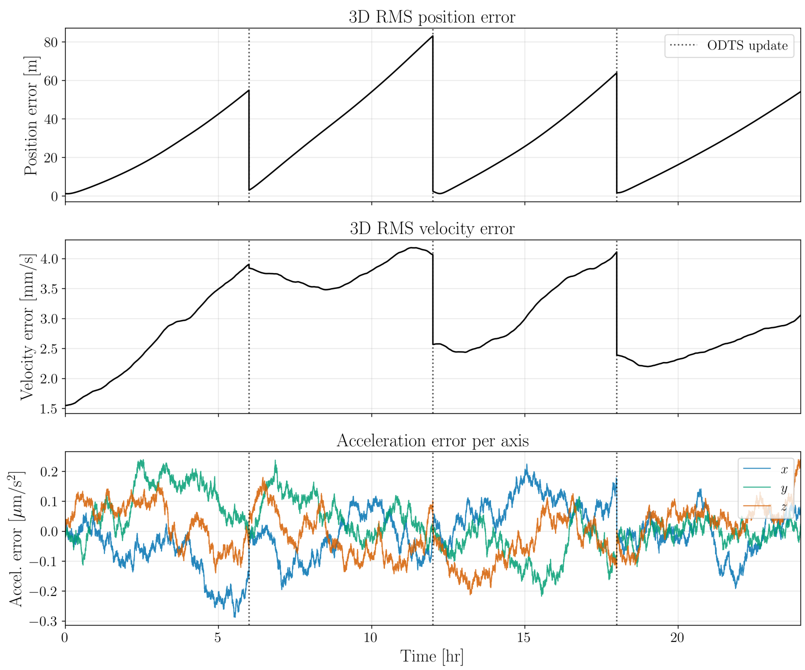

FIGURE 2 3D RMS ephemeris errors in [top] position and [middle] velocity. [Bottom] Per-axis acceleration errors showing the applied steady-state acceleration perturbations. ODTS resets are indicated by the vertical dotted line.

We also investigate the effects of local gravitational and kinematic time dilation on the rate offsets of the satellite and rover clocks. The rate offset for the rover appears as a constant because the time dilation due to gravity dominates the kinematic effects. Although the rover is traversing at 0.5 km/hr, this speed has a negligible effect and the Moon's rotation contributes a kinematic time dilation effect on the order of 10^{-17} s/s. For the satellite, both gravitational and kinematic effects are on the order of 10^{-11} s/s. Figure 3 illustrates the rate offsets of the satellite and rover clocks. To correct for such relativistic effects, we apply a differential local rate in the filter that is equivalent to the difference between the satellite and rover clock rate offsets.

5.2 | Ablation Study

The purpose of the ablation study is to quantify the contribution of filter-side design choices to the localization performance. The rover-satellite parameters and the truth-side physics are identically modeled in all ablation cases. The only modification in each ablation case is how the EKF is modeled. Table 3 denotes which filter design choices are removed for each case. When the bias state is disabled, the EKF state reduces to four dimensions, the 3D rover position and the relative clock drift. When relativistic corrections are disabled, the EKF update step no longer applies the correction in Equation 27. When the light-time delay

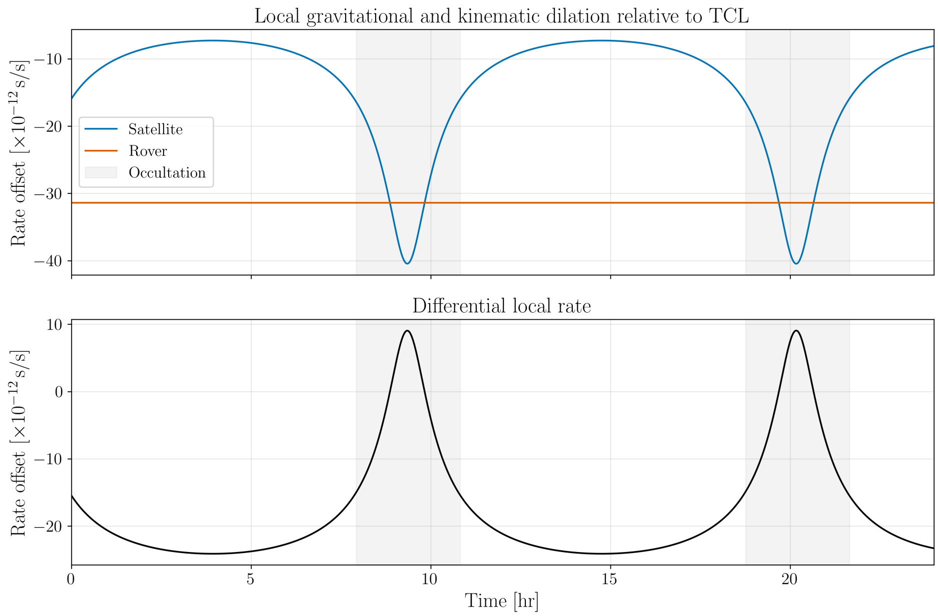

FIGURE 3 [Top] Rate offsets of the satellite and rover clocks based on local gravitational and kinematic time dilation relative to TCL. [Bottom] Differential local rate applied to the filter to correct for relativistic effects.

corrections are disabled, the filter directly applies the noisy satellite states $(\tilde{\mathbf{r}}_s, \tilde{\mathbf{v}}_s)$ in Equation 26. When Doppler measurements are disabled, the EKF is only relying on odometry and is thus only running the predict step. The fourth dimension in the state, the relative clock drift term, reduces to only estimate the rover clock drift.

TABLE 3
Enabled filter design choices and final mean position error for each ablation case. The ablation case labels correspond to the legend in Figure 4.

Ablation case	Doppler	Bias state	Relativistic	Light-time	Final mean position error [m]
Full model	✓	✓	✓	✓	32.9
-[Bias]	✓		✓	✓	34.8
-[Bias, Rel]	✓			✓	94.7
-[Bias, LT]	✓		✓		36.9
-[Doppler]					52.0

In the ablation cases, we disable the bias state when relativistic or light-time corrections are removed. Otherwise, the bias state absorbs part of the line-of-sight projection of the omitted physics, in addition to the intended ephemeris errors, thereby masking the effect of physics mismodeling on the position estimate. Disabling the bias state exposes these errors directly in the positioning performance. When we remove the bias term from the filter, we inflate the measurement noise covariance R, as would be done to account for errors that are not directly estimated by the filter. As a note, inflating the measurement noise does little to mitigate errors that are time-correlated, such as the ephemeris error model used in the study.

The results of the ablation study, shown in Figure 4, reveal two key conclusions: (1) Doppler measurements provide limited position improvement during the first ODTS interval because the time-correlated ephemeris errors dominate within each 6-hour interval, leaving short-term positioning to be driven primarily by odometry. Comparing the “Full model” to the “-[Bias]” case, we find that the bias state is necessary in absorbing these ephemeris errors until each ODTS reset. The fully-modeled filter outperforms the odometry-only case ("-[Doppler]") by taking advantage of the ODTS resets. (2) Omitting relativistic corrections causes substantially larger position errors than omitting light-time corrections. This result suggests that relativistic modeling is essential for achieving meter-level positioning in the lunar environment, while light-time corrections have a comparatively minor effect for the ELFO geometry considered in this study.

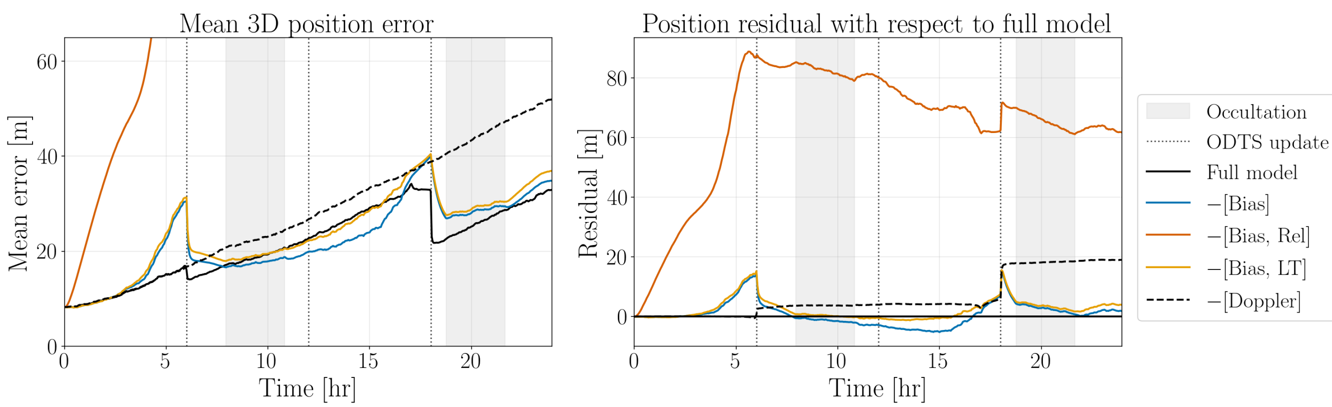

FIGURE 4 [Left] Mean 3D position error across 100 Monte Carlo runs for each ablation case using only the forward EKF pass. [Right] Mean 3D position error residual with respect to the full model. The legend applies to both subfigures.

5.3 | Tightly Coupled EKF with RTS Smoothing

Figure 5 shows the real-time forward pass of the EKF (equivalent to the "Full model" line in Figure 4) and a single backward pass using RTS smoothing. Despite the low observability of the single-satellite Doppler system and the multiple time-varying error sources modeled for higher fidelity validation, we find that the tightly coupled EKF with RTS smoothing is able to extract useful information from minimal measurements and maintain reasonable localization accuracy over time. The backward pass reduces the position error growth following the third ODTS reset, significantly reducing the accumulated error over the first 18 hours. Although the final mean position error remains the same as the forward pass, the backward pass is able to mitigate the errors due to ephemeris errors and take better advantage of the ODTS resets for all previous state estimates. Thus, for a mission scenario where the rover must accurately geotag samples, RTS smoothing proves to be an effective method to further refine the position estimates at sample collection sites during post-processing.

6 | CONCLUSION

In this work, we developed a lunar physics simulator and a sequential state estimator to evaluate rover localization from the Doppler shift of a single relay satellite. To evaluate the estimator at high fidelity, we implemented two effects that prior lunar PNT studies have largely treated as secondary: temporal correlations in ephemeris errors and clock rate offsets due to gravitational and kinematic time dilation. We find that the relativistic correction is in fact a substantial contributor to position error, since omitting it nearly triples the final position error over a 24-hour window. The time-correlated ephemeris error has a more regime-dependent effect, since large errors can force the filter to rely on odometry between ODTS resets and limit how much Doppler contributes within each interval. Across 100 Monte Carlo realizations over a 24-hour window, the filter proposed in this work achieved a final mean 3D position error of 32.9 m, compared to 52.0 m for the odometry-only baseline, which is a 36.7% improvement. The RTS backward pass further leveraged future ODTS resets to consistently reduce error growth as well as uncertainty within the first 18 hours of the trajectory. Future work includes extending the framework to hybrid architectures with multiple relay satellites or a lunar reference station, refining the ephemeris and ODTS parameters against operational performance estimates from planned lunar relay missions, and incorporating lunar terrain into the filter to constrain both the rover's surface position and the satellite signal availability patterns that the rover observes along its traverse.

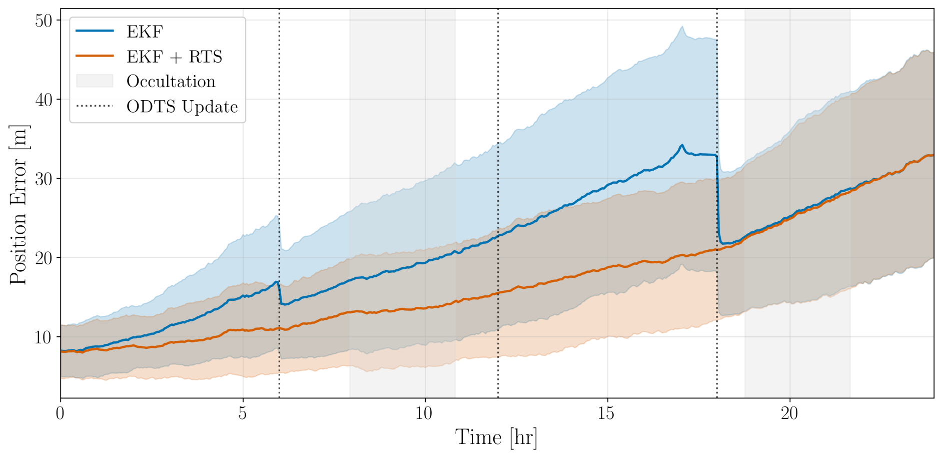

FIGURE 5 Position error over time using an online forward EKF pass and an offline backward RTS pass.

ACKNOWLEDGMENTS

This research is supported by the Stanford Knight-Hennessy Scholars fellowship and the National Science Foundation (NSF) Graduate Research Fellowship No. DGE-2146755. We would like to thank Matthias Kura for reviewing the manuscript and the other members of the Stanford Navigation and Autonomous Vehicles Laboratory for their helpful discussions and support.

REFERENCES

Ashman, B. W., Winternitz, L. B., Stacey, N. I., Long, A. C., Schmidt, M. C., Ryden, G. A., Liounis, A. J., Price, S. R., Bamford, W. A., Hur-Diaz, S. H., et al. (2024). Autonomous navigation of a lunar relay using GNSS and other measurements. *Proceedings of the 37th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2024)*, 3156–3173.

Bourgoin, A., Defraigne, P., & Meynadier, F. (2026). Lunar reference timescale. *Metrologia*, *63*(1), 015003.

Carpenter, J. R., & D'Souza, C. N. (2025). *Navigation filter best practices* (tech. rep.). NASA.

Casadesus Vila, G., Iiyama, K., & Gao, G. (2025). LuPNT: An Open-Source Simulator for Lunar Communications, Positioning, Navigation, and Timing. *2025 IEEE Aerospace Conference*, 1–18.

Coimbra, K. M. Y., Cortinovis, M., Mina, T., & Gao, G. (2025). Single-satellite lunar navigation via Doppler shift observables for the NASA Endurance mission. *NAVIGATION: Journal of the Institute of Navigation*, *72*(3).

Coimbra, K. M. Y., & Gao, G. (2025a). Comparative Analysis and Design of a Dual-Satellite System for Lunar Rover Localization. *2025 IEEE Aerospace Conference*, 1–11.

Coimbra, K. M. Y., & Gao, G. (2025b). Single-Satellite Doppler-Based Localization for Lunar Rovers in Motion. *Proceedings of the 38th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2025)*, 677–688.

Coimbra, K. M. Y., & Gao, G. (2026). Time-Correlated Ephemeris Errors and Relativistic Effects for Single-Satellite Doppler Localization of Lunar Rovers [Abstract accepted]. *Proceedings of the 39th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2026)*.

Fernando, B., Neal, C., Kiraly, J., Fernandez, B., Patterson, R., Gyalay, S., & Lemelin, M. (2025). The Case for Continuing VIPER: A Critical Milestone on the Journey Back to the Moon. *The Planetary Science Journal*, *6*(12), 289.

Giordano, P., Swinden, R., Gramling, C., Crenshaw, J., & Ventura-Traveset, J. (2023). LunaNet Position, Navigation, and Timing Services and Signals, Enabling the Future of Lunar Exploration. *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, 3577–3588. https://doi.org/10.33012/2023.19345

Grenier, A., Giordano, P., Bucci, L., Cropp, A., Zoccarato, P., Swinden, R., & Ventura-Traveset, J. (2022). Positioning and velocity performance levels for a lunar lander using a dedicated lunar communication and navigation system. *NAVIGATION: Journal of the Institute of Navigation*, *69*(2).

Grewal, M. S., & Andrews, A. P. (2001). *Kalman Filtering: Theory and Practice Using MATLAB* (2nd). John Wiley & Sons.

Iiyama, K., & Gao, G. (2026). GNSS-based Lunar Orbit and Clock Estimation With Stochastic Cloning UD Filter. *arXiv preprint arXiv:2601.16393*.

International Astronomical Union. (2024). Resolution II: to establish a standard Lunar Celestial Reference System (LCRS) and Lunar Coordinate Time (TCL) [Official IAU resolution]. Retrieved April 20, 2026, from https://www.iau.org/IAU/Iau/Publications/List-of-Resolutions.aspx

Keane, J. T., Tikoo, S. M., & Elliot, J. (2022). Endurance: Lunar South Pole-Aitken Basin Traverse and Sample Return Rover. *NASA Jet Propulsion Laboratory, Technical Report*. https://tinyurl.com/2p88fx4f

Kopeikin, S. M., & Kaplan, G. H. (2024). Lunar time in general relativity. *Physical Review D*, *110*(8), 084047.

Long, A. C., & Stacey, N. (2025). Goddard Enhanced Onboard Navigation System (GEONS) 3.1 Mathematical Specifications. *NASA Goddard Space Flight Center Technical Report, NASA/TP-20250006484*.

Melman, F. T., Zoccarato, P., Orgel, C., Swinden, R., Giordano, P., & Ventura-Traveset, J. (2022). LCNS Positioning of a Lunar Surface Rover Using a DEM-Based Altitude Constraint. *Remote Sensing*, *14*(16), 3942. https://doi.org/10.3390/rs14163942

Nardin, A., Minetto, A., Guzzi, S., Dovis, F., Konitzer, L., & Parker, J. J. K. (2023). Snapshot Tracking of GNSS Signals in Space: A Case Study at Lunar Distances. *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, 3267–3281. https://doi.org/10.33012/2023.19174

Pöhlmann, R., Staudinger, E., & Seco-Granados, G. (2025). Towards Hybrid Lunar PNT: Error Models, Lower Bounds and Algorithms. *arXiv preprint arXiv:2508.10699*.

Rauch, H. E., Tung, F., & Striebel, C. T. (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA journal*, *3*(8), 1445–1450.

Shapiro, I. I. (1964). Fourth test of general relativity. *Physical Review Letters*, *13*(26), 789.

Singer, R. A. (1970). Estimating optimal tracking filter performance for manned maneuvering targets. *IEEE Transactions on Aerospace and electronic systems*, (4), 473–483.

Sośnica, K., Bury, G., Zajdel, R., Kazmierski, K., Ventura-Traveset, J., Prieto-Cerdeira, R., & Mendes, L. (2021). General relativistic effects acting on the orbits of Galileo satellites. *Celestial Mechanics and Dynamical Astronomy*, *133*(4), 14.

SSTL. (2022). Lunar Pathfinder: Data relay satellite in orbit around the Moon. *SSTL, Service Guide (V4)*. https://www.sstl.co.uk/getmedia/9efe3cc5-6b5b-4fd1-bbd7-443bc2d21848/Lunar-Pathfinder-Service-Guide-V004.pdf