A Minimal Observation Method for Initial Maneuver Determination

Samuel Wishnek
BAE Systems Inc.

ABSTRACT

Maneuver determination is an important capability for maintaining custody of the large and growing space object catalog. Large maneuvers can be particularly difficult to handle as a space object may quickly leave the field of view of an unaware sensor. This paper describes an algorithm for estimating maneuvers from the only a pre-maneuver state estimate and two post-maneuver angles-only observations and analyzes its accuracy and reliability. This algorithm can be used to quickly and efficiently estimate maneuvers and seed accurate initial estimates into higher fidelity algorithms that require more data and computational resources.

1. INTRODUCTION

Maneuvering objects introduce significant additional challenges to the maintenance of object custody and space situational awareness. When an object with a known orbit undergoes an unknown maneuver, the maneuver introduces four new degrees of freedom. These are the time of the maneuver and the three components of the delta-v vector. For an approximately instantaneous maneuver, the only remaining constraint is that the new and old orbits intersect at the point where the object maneuvered. With these constraints and two additional post-maneuver angles-only observations, a unique time and delta-v vector can be found to describe the maneuver and post-maneuver orbit. Inspired by initial orbit determination (IOD) algorithms, this research explains and assesses a new algorithm that uses this minimal amount of data to provide an initial maneuver estimate. Key metrics include the speed, accuracy, and stability of the solution under realistic dynamics and errors.

Much of the existing work on maneuver estimation tends to focus on either comparing pre-maneuver and post-maneuver orbits or running a filter over data that encompasses the maneuver time [1][2]. Both methods assume a relatively large amount of data to perform the estimation and determine a precise solution. An important difference between prior work and the proposed work is that the proposed work can estimate the maneuver with the bare minimum amount of data. Three angles-only observations are required for initial orbit determination when there is no prior information, so the initial maneuver determination (IMD) algorithm necessarily leverages the pre-maneuver state to determine the post-maneuver state. The new algorithm provides a new capability to perform this estimation with too little data for existing algorithms to estimate the maneuver or post-maneuver state.

This approach is inspired by approaches taken by existing algorithms for initial orbit determination. The Gooding algorithm for angles-only initial orbit determination takes three angles-only observations of a space object, assumes range values for the first and third observation, and finds the orbit that fits the first and third position states. The first and third range values are iteratively corrected to minimize the error with the middle observation [3]. The proposed maneuver determination approach uses a similar initial setup by applying trial range values to the latter angles-only observation and then replaces the other range guess with a trial value of the time or eccentric anomaly of the maneuver and then corrects based on the error between the predicted and observed first observation.

The research covers the method of the algorithms and their relative performance in time and accuracy running inputs with realistic dynamics and errors. The accuracy of the solution under measurement error is an essential metric for demonstrating the real-life utility of the algorithms. Analysis of the algorithm includes a build-up to progressively more realistic scenarios. The full analysis shows acceptable performance with realistic initial state error, measurement error, high precision orbital dynamics, and finite maneuvers. This is despite the proposed algorithm using simple two-body dynamics and an assumption of instantaneous maneuvers to optimize the algorithm for speed. Additional results include performance metrics for various maneuver magnitudes, observation geometries, and relative times between observations.

2. METHOD

2.1 Gooding Initial Orbit Determination

The initial maneuver determination algorithm has a central logic inspired by the Gooding IOD algorithm. An understanding of this IOD algorithm is useful for framing how this IMD algorithm achieves results. The Gooding IOD algorithm has become a staple tool for orbit determination thanks to it's broad applicability and reliability [4]. It is found in tools including Orekit and ODTK for angles-only initial orbit determination [5][6]. Compared to other IOD algorithms, the process that it uses to reach a solution is straightforward. It starts with three angles-only measurements of a target in an unknown orbit. Trial values for the first and third range are chosen to generate potential position states for those measurements. A Lambert's problem solver is run using these two positions and the time between them. The Lambert's problem solver generates a full orbit solution and that can be used to predict the expected observation angle to the middle measurement. The error between the expected and observed middle measurement can then be used to adjust the trial range values and iterate until a solution is found [3]. At its core, it uses a numerical multi-variate root finding algorithm to adjust the trial range values and minimize the error to the middle measurement.

2.2 Overview of Method

The initial maneuver determination algorithm uses an approach inspired by the Gooding IOD algorithm. However, there are substantial differences between the two algorithms to account for the different types of available data and desired results from each algorithm. The IMD algorithm uses the time of the maneuver and range to the second measurement as independent variables and the error between an expected first measurement and the observed first measurement as the key dependent variables.

Using the two-body dynamics assumption, the pre-maneuver orbit and the post-maneuver orbit form intersecting ellipses in space with a shared focus. Since a full pre-maneuver position and velocity state is known, assuming the time of the maneuver provides a trial value for the point of intersection between these ellipses. The trial value of the range to the second measurement then provides a second position along the post-maneuver orbit and enough information to use a Lambert's problem solver to find the post-maneuver orbit. A numerical root-finding algorithm can then be applied to calculate corrections to the trial values based on the error between the expected and observed first measurement.

The two unknown measurement ranges are not used as independent variables because it would make the error measurement much more difficult. While there are methods for finding the point and distance of closest approach between two ellipses, that does not account for any potential timing mismatches for the estimated positions along each ellipse. This further complicates the task of aligning the orbits for a feasible transfer. The selection of maneuver time and second measurement range as independent variables avoids this issue by assuming the unknown time value. Any observable error in the trial values shows up as measurable angle error in the first observation.

It should be noted that later sections replace the time independent variable with eccentric anomaly on the pre-maneuver orbit. This substitution improves computational efficiency but does not materially change the logic of the algorithm. Time can be calculated from eccentric anomaly and vice versa. Accordingly, much of the discussion about the algorithm will refer to this first independent variable as time since it is the more intuitive explanation for the process. Section 2.8 discusses this change in detail.

2.3 Initial Conditions

Similar to how the Gooding IOD algorithm can accept an initial guess for the first and third ranges. the IMD algorithm can accept an initial guess for the time of the maneuver and range between the target and observer at the time of the second observation. Gooding showed broadly consistent convergence with arbitrary selection of the initial ranges [3]. Similarly, an arbitrary choice of the initial guess for the maneuver time of halfway between the time of the last known pre-maneuver state and the first post-maneuver observation is recommended for the IMD algorithm. This sets the initial guess halfway between the constraining bounds on the time. The range is arbitrarily set to the distance between the observer position at the time of the second observation and the last known position state of the target. This simply scales the initial guess to reflect the general distances involved in the problem. The reliability of these choices for the initial guess are demonstrated in results section.

An alternative approach is to approximate the post-maneuver orbit as circular which allows for a complete solution to the post-maneuver orbit from only the two angles-only observations. This can then be used to find an initial guess

for the maneuver time based on the intersection of the pre-maneuver and post-maneuver orbit planes. The cross product of the the normal vectors of the two planes will align with the position vector of the target object at the time of the maneuver. One caveat of this approach is that for small maneuvers, small errors in the measurements can potentially push the initial guess beyond the time limits imposed by the known pre-maneuver state and first post-maneuver observation. Care should be taken to catch these cases and enforce realistic limits on the maneuver time if the circular orbit assumption approach is taken.

2.4 Assessing the Trial Values

The trial maneuver time and range for the second observation fully define the post-maneuver orbit. The orbit can be constructed by propagating the initial state to the time of the maneuver and using the range to find the position of the target at the time of the second observation. This defines two full three dimensional position states for the target along with the time between the states. This is the classic Lambert's problem and it can be solved with an implementation of one of the many Lambert's problem solvers available [7][8][9]. Some Lambert solvers can provide a solution after an arbitrary number of orbits if that number is known. The IMD algorithm can inherit this capability if the implemented Lambert solver is capable as well. The Lambert solver will output the velocity of the target object immediately after the maneuver and at the time of the second observation. To correct the trial values, the observed angles from the first post-maneuver observation can be compared to the expected values. The remaining process is identical to the that of the Gooding IOD algorithm.

Propagate the estimated state of the target object from the time of the second observation back to the time of the first to get the expected position of the target at the time of the first observation, $\vec{r}_1$. Define the expected observation vector as $\vec{p}_{1,e} = \vec{r}_1 - \vec{R}_1$. The difference between the normalized expected vector, $\hat{p}_{1,e}$, and the observed pointing vector for the first observation, $\hat{p}_{1,o}$, represents the agreement between the trial values and the observed truth. This is formulated as the vector $\vec{P}$ as expressed in equation 1. Note that $\hat{p}_{1,e}$ is defined as the normalized vector such that $\vec{p}_{1,e} = p_{1,e}\hat{p}_{1,e}$.

$$\vec{P} = (\hat{p}_{1,o} \times \vec{p}_{1,e}) \times \hat{p}_{1,o} = \rho_1 \left(\hat{p}_{1,e} - (\hat{p}_{1,o} \cdot \hat{p}_{1,e})\, \hat{p}_{1,o}\right) \tag{1}$$

The vector valued $\vec{P}$ goes to the zero vector when $\hat{p}_{1,e} = \hat{p}_{1,o}$, and the direction of the vector indicates the direction of the trial value's error. Normalizing $\vec{P}$ and taking the dot product with $\vec{p}_{1,e}$ defines a scalar metric for the error. Equation 2 shows how this can be expressed in terms of the expected and observed pointing vectors. Note that the measurement vector is not normalized and should have magnitude equal to the expected range to the target, ρ_1. We define this metric as F, the cost associated with the trial values.

$$F = \hat{P} \cdot \vec{p}_{1,e} = \rho_1 \sqrt{1 - (\hat{p}_{1,o} \cdot \hat{p}_{1,e})^2} \tag{2}$$

Defining the vector perpendicular to both the observed first measurement, $\hat{p}_{1,o}$, and $\hat{P}$ will be convenient for adjusting trial values as described in the next section. This is the unit vector $\hat{N}$ as described in equation 3. A visualization of these vectors is provided in figure 1

$$\hat{N} = \frac{\hat{p}_{1,o} \times \hat{P}}{\left|\hat{p}_{1,o} \times \hat{P}\right|} \tag{3}$$

2.5 Adjusting the Trial Values

With the cost function F described in the previous section, there are several paths forward for determining the appropriate corrections to the trial values. One easily implemented approach is to apply an optimization algorithm to the function that calculates F. Many pre-built optimizers include useful features such as the ability to define constraints on the independent variables to keep the trial maneuver time and range to the second observation realistic. The basic Nelder-Mead algorithm can easily fill this role. However, a major drawback of this approach is the computational overhead inherit to these generic algorithms. Like IOD, one of the major potential use cases of IMD is its use in computational-resource limited environments. The Newton-Raphson multivariate root finding approach used in the Gooding IOD algorithm may better fit these use cases [3]. Accordingly, the Newton-Raphson method serves as the primary approach for the analysis here. Implementations with a pre-built optimizer are used as points of comparison in the results.

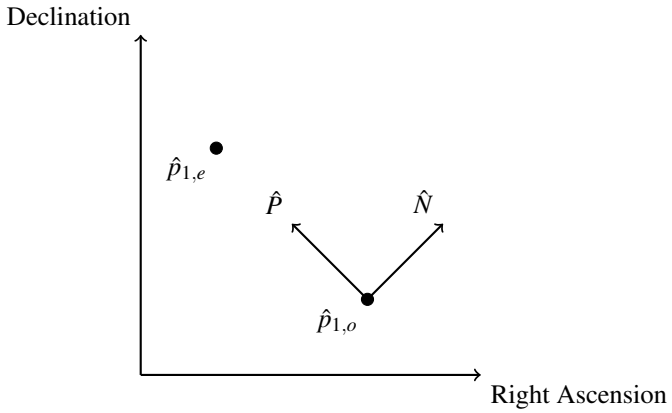

Fig. 1: Visualization of the measurement, trial point, and computed vectors

The $\hat{P}$ and $\hat{N}$ vectors form a basis for the error of the first measurement. Since the axes are defined by the direction of the error, all of the error for the trial values will be aligned with the $\hat{P}$ axis and the component in the $\hat{N}$ direction will be zero. The numerical derivatives of the cost function is calculated by adjusting each of the independent variables by a small amount in both the positive and negative direction and finding the expected target position at the first measurement for each case. That deviated solution is then projected on the $\hat{P}$ and $\hat{N}$. The projection on $\hat{P}$ is the previously defined value F and the projection on $\hat{N}$ is G as shown in equation 4. Since the deviated solution is projected on the original axes, the value of G for small changes in the input variables can be non-zero.

$$G = \hat{N} \cdot \vec{p}_{1,e} \tag{4}$$

Deviating each of the input variables independently, one can construct a two by two Jacobian matrix for the error of the trial values with elements defined in equations 5 through 8. F and G should both go to zero when the expected measurement matches the observed measurement, so the root of this multivariate function is the point where the chosen time and range line up to yield the correct first observation. The function and Jacobian provide everything needed to perform the Newton-Raphson iteration. Equation 9 shows the formula for the update step.

$$\frac{dF}{dt} = \frac{F(t_i + dt, \rho_{2,i}) - F(t_i - dt, \rho_{2,i})}{2dt} \tag{5}$$

$$\frac{dF}{d\rho} = \frac{F(t_i, \rho_{2,i} + d\rho) - F(t_i, \rho_{2,i} - d\rho)}{2d\rho} \tag{6}$$

$$\frac{dG}{dt} = \frac{G(t_i + dt, \rho_{2,i}) - G(t_i - dt, \rho_{2,i})}{2dt} \tag{7}$$

$$\frac{dG}{d\rho} = \frac{G(t_i, \rho_{2,i} + d\rho) - G(t_i, \rho_{2,i} - d\rho)}{2d\rho} \tag{8}$$

$$\begin{bmatrix} t_{m,k+1} \\ \rho_{2,k+1} \end{bmatrix} = \begin{bmatrix} t_{m,k} \\ \rho_{2,k} \end{bmatrix} - \begin{bmatrix} \frac{dF}{dt} & \frac{dF}{d\rho} \\ \frac{dG}{dt} & \frac{dG}{d\rho} \end{bmatrix}^{-1} \begin{bmatrix} F \\ 0 \end{bmatrix} \tag{9}$$

2.6 Determining Convergence

The previous section describe how the algorithm can iteratively correct trial values to reduce the error on the predicted first observation. This can be performed until one of several convergence thresholds is met. The primary method for determining convergence is based the scalar error value, F. This value can be scaled by the estimated orbital radius as the first observation time to remove any unit dependency. The solution can be considered converged if this value falls below some tolerance value. The tolerance should be defined near zero but larger than machine precision. A value

of 1e−14 was used for this implementation. Second, the main loop of the algorithm should have some iteration limit before exiting. The influence of differing iteration limits is discussed in detail in section 3.2.3.

In addition to these criteria, other checks are necessary to catch errors that might arise when the expected and observed measurement angles match. The vector $\vec{P}$ will be zero if $\hat{p}_{1,o} = \hat{p}_1, e$ as the cross product of identical vectors is zero. Attempting to calculate $\hat{P}$ would lead to a division by zero error, so values of $\vec{P}$ near machine precision indicate that the algorithm has successfully found a solution. Similarly, this also occurs in the calculation of $\hat{N}$ and requires a similar check that $\hat{p}_{1,o} \times \hat{P}$ is not near machine precision.

When the algorithm detects successful convergence, the delta-v of the maneuver can be found from the difference in the pre-maneuver and post-maneuver target velocity at the maneuver time and returned as output along with the maneuver time value.

2.7 Multiple Revolution Case

The IMD algorithm is fully capable of handling cases with multiple revolutions. These revolutions can be between the initial state and observations, between the two observations, or a combination of both. However, in this case there may be multiple solutions, and some adjustments to the algorithm will be needed to handle them. The complications arise in the Lambert solver.

Many popular Lambert solver algorithms work fine through multiple orbits, but they need to be provided the number of orbits to arrive at a solution [9]. If the number of orbits for the target are known in advance, that number just needs to be passed along to the Lambert solver algorithm so that it can appropriately account for it in its results. In this case, the IMD algorithm is otherwise unchanged. However, if the number of revolutions is not known, more substantial changes are need. The time between the initial state and observations together with reasonable limits on the minimum semi-major axis of the orbit can provide limits on the maximum number of revolutions. The IMD algorithm can then be run for each revolution count trial value. In most practical cases, this will be a small finite number of test cases. Most of the individual results will either error out as the proposed solution violates physical limits or return very high delta-v values. The subset of remaining results is the set of feasible solutions to the multiple revolution case.

2.8 Adjustments for Improving Speed and Reliability

Kepler's equation, (10), is not analytically solvable [10]. Accordingly, even the most basic two-body propagation of the initial state to the time of the maneuver requires an iterative root-finding or integration approach. While these algorithms can be very fast with only a few iterations, the speed can be improved by substituting the time of the maneuver with the eccentric anomaly of the maneuver. The eccentric anomaly defines the state of the satellite at the maneuver time, and the maneuver time can be found analytically from the eccentric anomaly and initial orbit elements using equation 11. The process of the algorithm is then identical from that point forward. The main drawback of this approach is that the implementation becomes more complicated. The order of times is immediately obvious from the values, but the eccentric anomaly rolls over back to zero with each pass of the periapse. Determining if a new trial value remains within the bounding limits and accounting for the possibility of multiple orbits are both still possible but require more effort to implement. The time of the maneuver must be between the time of the last known state and the first post-maneuver observation. This is trivial to implement when time is the independent variable, but more complex with eccentric anomaly. The check can instead be implemented by initially using the iterative two-body propagation to find the eccentric anomaly of the initial orbit at the bounding time of the first maneuver. This yields the clamping constraint for the eccentric anomaly trial value so that the propagation only needs to occur once at the start rather than in every iteration. This change removes one of the three iterative processes inside the optimization loop of the maneuver determination algorithm. The Lambert solver and propagation to the time of the first observation remain as they do in the Gooding IOD algorithm.

$$M = E - e\sin(E) \tag{10}$$

$$t_{maneuver} = t_{initial} + \sqrt{\frac{a^3}{\mu}}\left(E_{maneuver} - E_{initial} - e\left(\sin(E_{maneuver}) - \sin(E_{initial})\right)\right) \tag{11}$$

The bounds on the independent variables, eccentric anomaly before the maneuver and range of the second observation, provide an additional opportunity to improve both the speed and reliability of the algorithm. Both of these variables

are bound by intuitive limits. The eccentric anomaly cannot be before the last known pre-maneuver state and cannot be so late that the corresponding time is after the first post-maneuver observation. The range is limited to being a positive value and close enough to be feasibly seen by the observer and in the environment where the Earth's gravity dominates its motion. The most basic approach to implementing these bounds is to check if the bounds are violated with each set of new trial values and then enforce the bounds by setting violating values to values just barely inside the limits. However, one of the common failure modes of this algorithm is that a minimum outside the bounds can pull the optimizer to the limits and hold it there until the maximum number of iterations is reached. The result is both wrong and slow since the algorithm must reach its maximum number of iterations. Better performance in both metrics can be achieved by randomly choosing a new trial value within the range rather than keeping the value near the bounds. This allows the algorithm the opportunity to find the true solution rather than getting stuck pressing against the limiting values. The improvement from this change is demonstrated in section 3.2.2.

3. RESULTS

3.1 Performance Against Realistic Deviations and Errors

To keep the computational cost of the IMD algorithm low, it relies on simple two-body orbital dynamics to model the motion of the target satellite. In reality, there are a variety of perturbing forces that shift the experienced dynamics of spacecraft away from the two-body ideal. Higher order components of the Earth's geopotential model, solar radiation pressure, atmospheric drag, and the pull of the Moon and other planets all contribute to the motion of real space objects. The ideal impulsive maneuver modeled in the IMD algorithm is also only an approximation of the finite thrust that a real maneuvering object will experience.

In addition to these differences between reality and the approximations of the model, there is also error in any real angles-only measurement and state solution. Measurement and state uncertainty are unavoidable complications on any set of real observational data. The cumulative impact of these errors and approximations further complicates the task of estimating states and maneuvers. In order to be a useful tool, the IMD algorithm must be robust to both the errors introduced from model simplifications and the inevitable uncertainty in input data. The reliability of the algorithm against varying amounts of error is a key metric for understanding the capabilities and limitations of it. This section will address the effects of these errors and build up to a set of deviations that model realistic orbits and observations.

Monte Carlo tests are used to investigate the response of the IMD algorithm to these errors. Each deviation is applied in series to build up to a realistic model combining all of the expected errors. For these tests, the observers are placed in random positions in low Earth orbit while the target object is generated with a random geosynchronous orbit. The constraints used to generate these orbits are shown in table 1. The test cases are generated with a seeded random number generator to keep each trial as similar as possible. The maneuvers themselves are Gaussian distributed random vectors with each component independently drawn from an unbiased distribution with the stated standard deviation.

Each Monte Carlo trial consists of 1000 randomly generated samples. Results are visualized with two empirical cumulative distribution function (ECDF) plots such as those depicted in figures 2 and 3. The first shows the normalized error between the estimated maneuver and the true maneuver. This is calculated according to equation 12. The second shows the error in the estimated maneuver time in seconds. The vertical axis of an ECDF plot is the percentile of the trial case with the corresponding error. To better show the results, the x-axis for the delta-v error plot is logarithmic while the x-axis for the time error plot is linear.

$$\Delta v \text{ Relative Error} = \frac{|\Delta\vec{v}_{true} - \Delta\vec{v}_{estimate}|}{|\Delta\vec{v}_{true}|} \tag{12}$$

Table 1: Orbital element ranges uses for random orbit generation

Orbit	Semi-Major Axis Range	Eccentricity Range	Inclination Range
LEO	6500 to 8000 km	0 to 0.1	-57.3° to 57.3°
GEO	42160 to 42170 km	0 to 1E-7	-5.73° to 5.73°

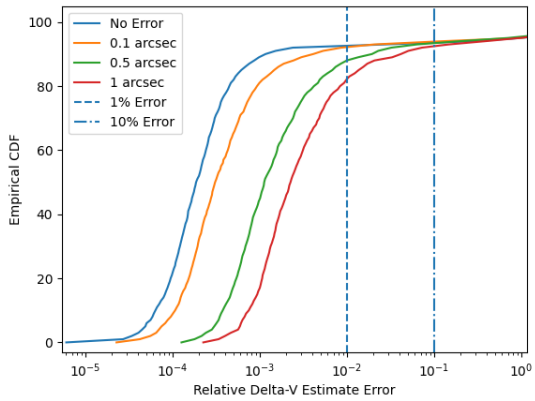

Fig. 2: CDF of normalized maneuver vector estimate errors for varying measurement errors.

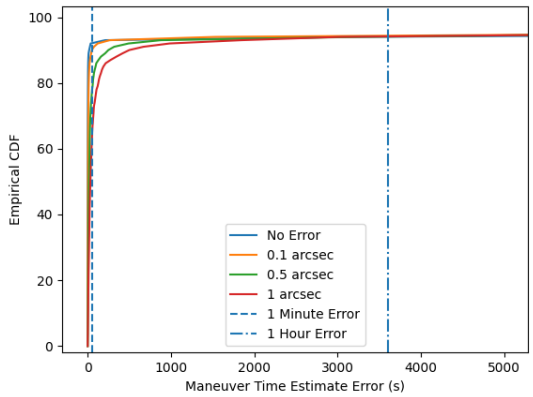

Fig. 3: CDF of maneuver time error for varying measurement errors.

3.1.1 Measurement Error

Measurement error in angles-only measurements is the error in the reading of the perceived angles to the target. This comes from a variety of physical limits on the sensor including the finite pixel size, distortion in optical components, the size of the point spread function on the sensor, and atmospheric distortion for ground-based sensors. For testing purposes, these effects are all modeled together into a symmetric unbiased Gaussian distribution around the truth. The distortion is added by first finding a randomly generated unit vector in the plane perpendicular to the true measurement vector and rotating the true measurement by a random angle around that vector where the angle is selected from a Gaussian distribution with standard deviation equal to the modeled measurement error magnitude.

Measurement error can vary dramatically depending on the hardware used in the observations. The analysis performed here looks at measurement error standard deviations of 1, 0.5, and 0.1 arcseconds as these are a range of values typical to space-based instruments [11]. Figures 2 and 3 show the ECDF of the normalized maneuver vector and maneuver time error for these test cases. The decrease in the accuracy of the estimates with increasing measurement error is clearly visible in both plots. The case with no measurement error has about 94 percent of test cases within ten percent of truth. This drops to about 92.5 percent with a standard deviation of one arcsecond of measurement error. This one arcsecond standard deviation measurement error is carried forward as the default measurement error in subsequent analysis.

3.1.2 Initial State Error

The accuracy of the pre-maneuver state estimate is the other data quality factor that can influence the performance of the IMD algorithm. Any real estimate for the pre-maneuver state will have some amount of error. A useful IMD algorithm must remain accurate for reasonable amounts of error in the pre-maneuver state.

The accuracy of a state estimate can vary depending on the type, quality, and amount of data used to calculate it. The test cases chosen for this analysis use initial position component standard deviations of 10 meters, 100 meters, and 1 kilometer paired with corresponding velocity standard deviations of 1 millimeter per second, 10 millimeters per second, and 100 millimeters per second. These errors are implemented as white Gaussian noise added to the truth state. The empirical CDF plots of the results for each of these cases are provided in figures 4 and 5. This analysis retains 1 arcsecond measurement error as described in the previous section. There is a small increase in maneuver estimate error with increasing state error and a more dramatic jump in the highest error case. Significantly fewer test cases returned results within one percent of truth though the ten percent threshold saw a much smaller impact. The three lower error cases converge on 92.5 percent of results within 10 percent of truth with the outlier one kilometer uncertainty case dropping to 88.8 percent of test cases within 10 percent of truth. The 10 meter position uncertainty and 1 millimeter per second velocity uncertainty are selected as the default for subsequent analysis.

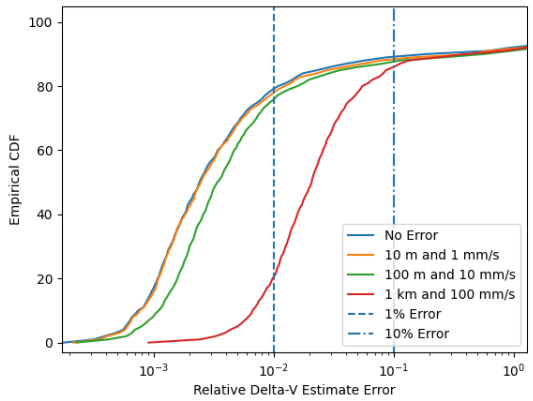

Fig. 4: CDF of normalized maneuver vector estimate errors for varying initial state errors.

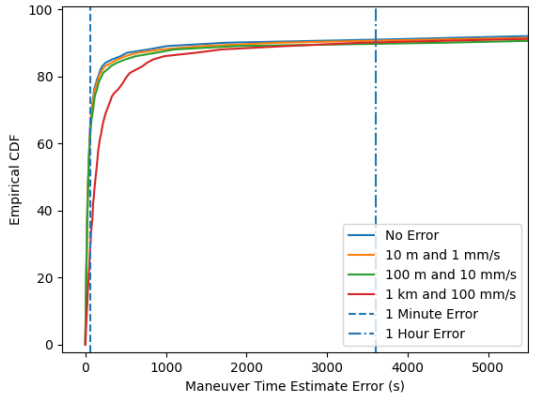

Fig. 5: CDF of maneuver time error for varying initial state errors.

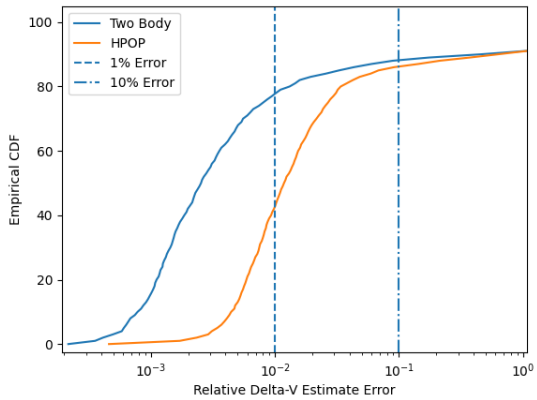

Fig. 6: CDF of normalized maneuver vector estimate errors for two-body and high precision propagation methods in truth modeling.

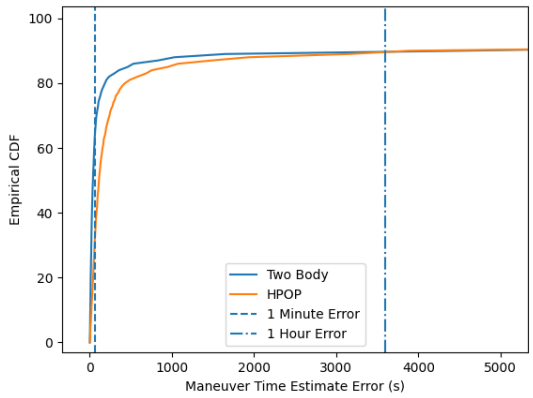

Fig. 7: CDF of maneuver time error for two-body and high precision propagation methods in truth modeling.

3.1.3 Realistic Orbit Deviations

Similar to IOD algorithms, the IMD algorithm uses two-body orbital dynamics to model the motion of the target space object. Real space objects are subject to a variety of perturbing forces including higher order components of the geopotential model, the gravitational pull of other bodies, solar radiation pressure, and atmospheric drag. Accordingly, the IMD algorithm must be stable to this modeling difference to be useful for estimating the maneuvers of real space objects. To test this, the High Precision Orbit Propagator (HPOP) in Ansys's STK software is used to generate truth data. As implemented, this orbit propagator includes a geopotential model up to degree and order 21, solid tides, perturbing point gravity from the Sun and Moon, and solar radiation pressure. The ECDF of the results of the two-body propagated truth and HPOP propagated truth are shown in figures 6 and 7. The change shows a significant reduction in the portion of results within one percent of truth from around 82.4 to 46.1 percent. However, the change at the ten percent threshold is much smaller from around 92.5 percent to 90.4 percent. The error introduced by the differing modeling between truth and the IMD algorithm introduces a layer of noise into the results that deviate estimates away from truth as the best fit maneuver under two-body dynamics does not exactly align with the true maneuver under HPOP dynamics.

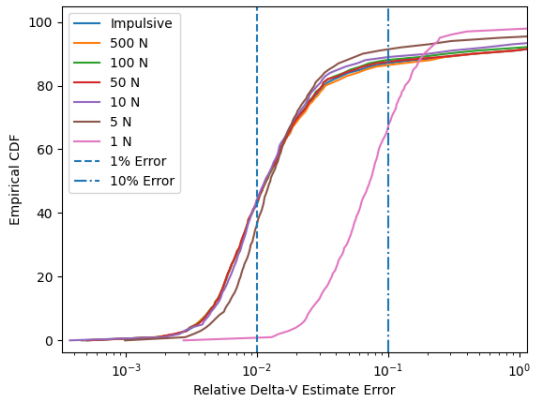

Fig. 8: CDF of normalized maneuver vector estimate errors for varying engine thrust and 1000 kg spacecraft mass.

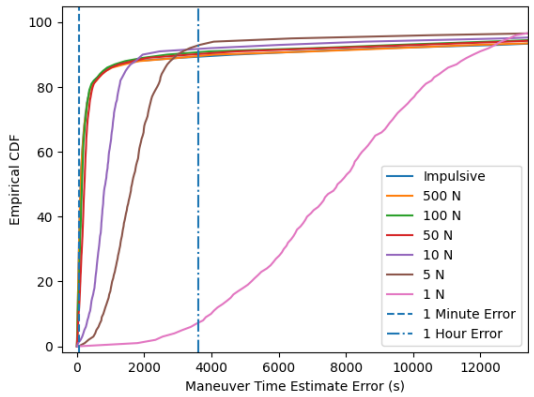

Fig. 9: CDF of maneuver time error for varying engine thrust and 1000 kg spacecraft mass.

3.1.4 Finite Maneuver

The last change implemented to better model realistic satellite motion was to replace the impulsive maneuver with a finite maneuver. This was done by generating test scenarios with STK's "Seed Finite From Impulsive" option to replace the modeled impulsive maneuver with the closest approximation performed by an engine with constant thrust and specific impulse (I_{SP}). The specific impulse of the the engine was set to 300 seconds to model a chemical thruster and the thrust was varied between 1 and 500 newtons to validate the IMD algorithm against test cases both closer and further from the impulsive ideal. The mass of the spacecraft was modeled as 1000 kg, and the accelerations can be calculated by dividing the thrust by the mass. The results for these runs are provided in figures 8 and 9. There is very little change in the delta-v results from impulsive for engines with thrust down to five newtons. With a one newton thrust engine, the difference is more significant with almost none of the trial cases with maneuvers within one percent of truth down from around forty five percent for the bulk of cases and thirty five percent for the five newton engine. The portion within ten percent delta-v error drops to around sixty seven percent for the one newton engine. The change is more visible in the error on the time of the maneuver. Very few of the test cases remain accurate to within an hour and the steady drop in accuracy with decreasing thrust is visible in the decreasing slope of the five and ten newton engines.

These results show the combined impact of realistic deviations on the target satellite and measurements. They build up to a realistic simulation of feasible use cases for the IMD algorithm. While results vary depending on the details of the simulated scenario, higher thrust scenarios were predominantly successful with 80 to 90 percent of cases solving to within 10 percent of truth.

3.2 Algorithm Speed

The IMD algorithm attempts to maximize speed in two important ways. First, by minimizing the amount of data required to estimate a maneuver, an SDA system can find updated state results earlier than a system using traditional methods that filter through the maneuver or require three observations for post-maneuver IOD. Second, the algorithm is designed to be fast to execute and feasible to implement in systems with limited computational resources.

3.2.1 Eccentric Anomaly as an Independent Variable

The Gooding IOD algorithm has two sequential iterative steps inside the main loop of the algorithm. For each trial values of first and third measurement ranges, the algorithm first solves Lambert's problem for the orbit connecting the two position states then propagates one of the solutions to the time of the second measurement. Both steps are implementations of the Newton-Raphson root-finding algorithm to find solutions to these non-analytically solvable

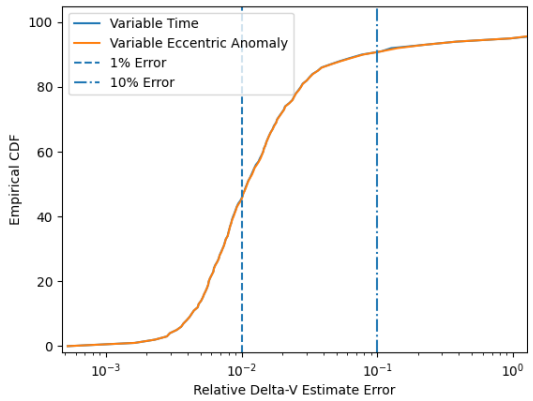

Fig. 10: CDF of normalized maneuver vector estimate errors for time versus eccentric anomaly as a trial value.

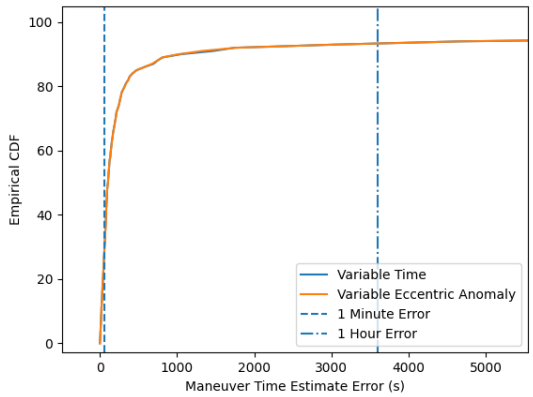

Fig. 11: CDF of maneuver time estimate errors for time versus eccentric anomaly as a trial value.

problems [9][10]. In practice, these two iterative steps contribute to the vast majority of the computational cost of the algorithm.

Implementing the IMD algorithm as initially described with independent variables of time and the second measurement range would require three iterative steps for each trial value of the independent variables. This would be the propagation to the maneuver time, the Lambert solver connecting the maneuver time position and second observation position, and the propagation to the first measurement time. Replacing the maneuver time independent variable with the eccentric anomaly along the pre-maneuver orbit replaces this first iterative step with an analytic one.

To quantify the improvement, a Monte-Carlo test is run using the full set of perturbations described in the last section for both cases. Since computational power varies by machine and implementation details, raw performance is not a clear indicator. Instead, the key results will be the comparative performance between the two. With independent variables of time and the second measurement range, the average time per trial scenario was 11.99 milliseconds. Replacing the maneuver time with the maneuver eccentric anomaly improves this to 8.82 milliseconds. This represents a 26.4 percent improvement for the same test data. Figures 10 and 11 show the relative delta-v error and time error CDFs respectively with no significant difference between the two approaches.This shows that the improvement in computation time comes at no cost to accuracy.

3.2.2 Random Reinitialization

One of the main failure modes of the IMD algorithm occurs when the various unmodeled deviations in the true motion and measurements forms a local minimum that pulls the solution outside of the feasible limits imposed by realistic ranges and the time limits in which the maneuver must have occurred. When the algorithm falls into one of these local minima, it attempts to pass into these non-feasible regions before being corrected by the sanity check on the variables and brought back into the known bounds. This can settle into a loop as it passes back and forth between the impossible values and the corrected values. This can be avoided by replacing a simple clamping limit on the variables with a correction that randomly sets the independent variables to feasible values rather than pushing them back to the nearest feasible state.

Random reinitialization of the independent variables improves both the computation time and reliability of the algorithm. Cases that fall into incorrect out-of-bounds local minima have the opportunity to recover and are no longer stuck until the maximum iteration count is reached. Accordingly, for a one thousand trial Monte-Carlo simulation, the average compute time with clamping limits was 20.9 milliseconds. This improved to 10.1 milliseconds with random reinitialization. This was an improvement of 48.1 percent. The empirical CDFs of the maneuver estimate errors from these runs are shown in figure 12. Random reinitialization improves the portion of cases within ten percent of truth from 83 to 90 percent.

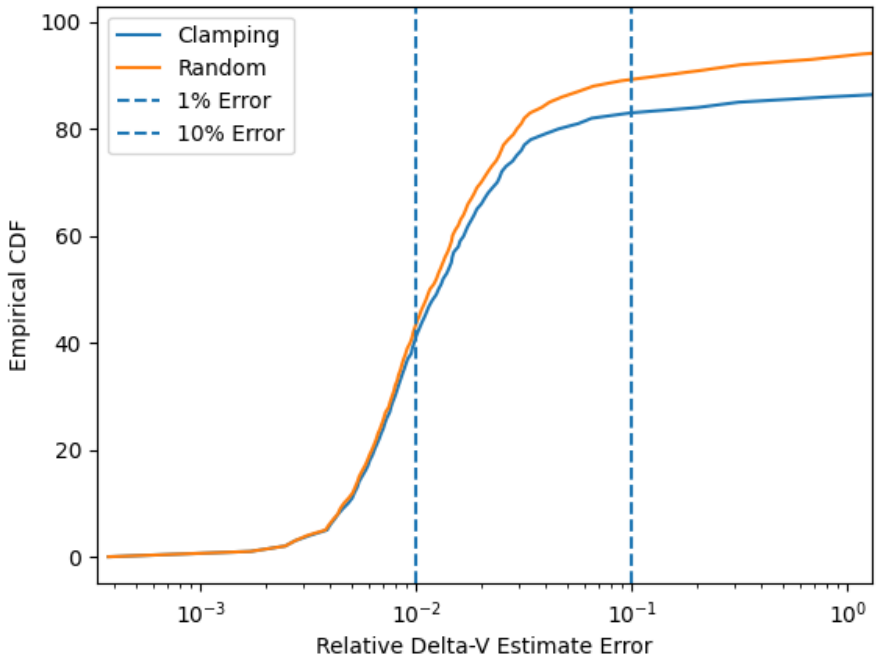

Fig. 12: Relative performance of clamping limits on the independent variables and random reinitialization

3.2.3 Maximum Iteration Limit

With random reinitialization as described in the previous section, one of the main failure modes of the algorithm is closed off. It can no longer get trapped in a cycle attempting to approach a local minimum out of bounds. With this change the algorithm has the potential to benefit from an increased maximum iteration limit as it provides more opportunity to find the true solution. Previously shown test cases used a twenty iteration limit. This value was varied and the results for iteration limits between ten and one hundred are shown in figures 13 and 14. The first plot shows how the delta-v accuracy evolves with changing iteration limit while the second plot shows the average time per run of the algorithm. The dashed line is the linear fit of the data with a slope of 0.0440 milliseconds per additional iteration and an r value of 0.9986 for this small sample size.

The algorithm sees only a slight benefit when increasing the iteration limit above twenty for these test cases, so the maximum remains at twenty for the subsequent analysis described in this paper.

3.3 Performance by Orbit Domain

Previous described examples show results for a geosynchronous target and low-Earth orbiting observers. To assess how different observing geometries can impact the accuracy of results, a series of cases are run with varying target and observer orbital domains. In order to avoid complications and ambiguity from multiple orbit cases, the chosen domains are approximately medium-Earth orbit (MEO) and geosynchronous orbit. The orbits are randomly and uniformly drawn from the ranges specified in table 2. The maneuver time and observation times are randomly and uniformly drawn from the ranges specified in table 3. The IMD algorithm itself is provided the observer positions, randomly perturbed measurements, and randomly perturbed initial state. While the maneuver occurs within an hour of the initial state, this limit is not provided to the algorithm and it must search the full time span between the initial state and the time of the first observation. For these test cases, the magnitude of the maneuver was fixed with a delta-v of ten meters per second and a random direction.

Results for these cases are shown in figures 15 and 16. Labels are formatted with the observer domain first and the target domain second. The performance was best for the cases with the MEO observers looking outward at the GEO targets. The MEO targets performed similar to one another and worse than the GEO targets. This aligns with results

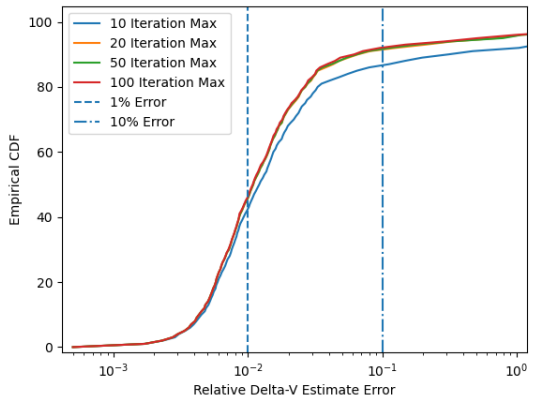

Fig. 13: CDF of normalized maneuver vector estimate errors for varying maximum iteration limit.

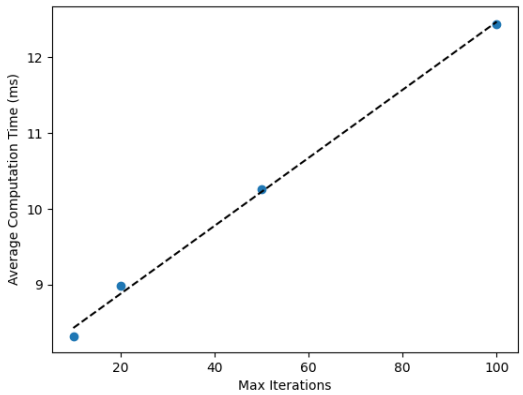

Fig. 14: Average computation time versus maximum iteration limit.

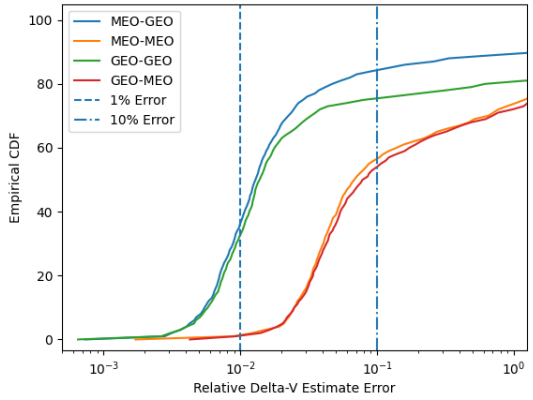

Fig. 15: CDF of normalized maneuver vector estimate errors for varying orbital domains of the observers and targets.

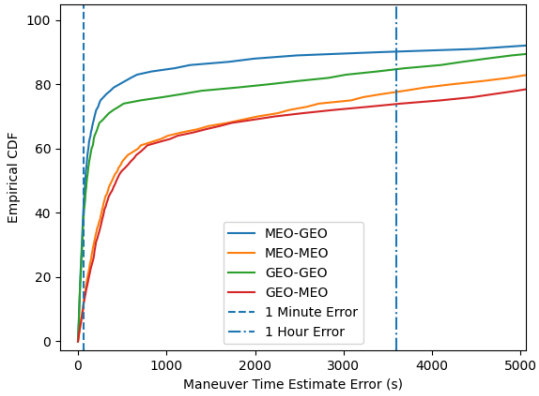

Fig. 16: CDF of maneuver time estimate errors for varying orbital domains of the observers and targets

showing improved performance for outward looking observers rather than inward or co-orbital observers in IOD [4]. While performance is reduced for the MEO targets, the majority of cases are still accurate to within ten percent of truth despite the modeled measurement and state errors as well as the more complex dynamics used for modeling the truth data.

Table 2: Orbital element ranges uses for random orbit generation in domain comparison.

Orbit	Semi-Major Axis Range	Eccentricity Range	Inclination Range
MEO	26500 to 26600 km	0 to 0.1	-57.3° to 57.3°
GEO	42100 to 42200 km	0 to 0.1	-57.3° to 57.3°

Table 3: Time ranges uses for maneuver and observation modeling in domain comparison.

Maneuver Time Range	First Observation Time Range	Second Observation Time Range
0 to 1 hour	3 to 3.5 hours	5 to 5.5 hours

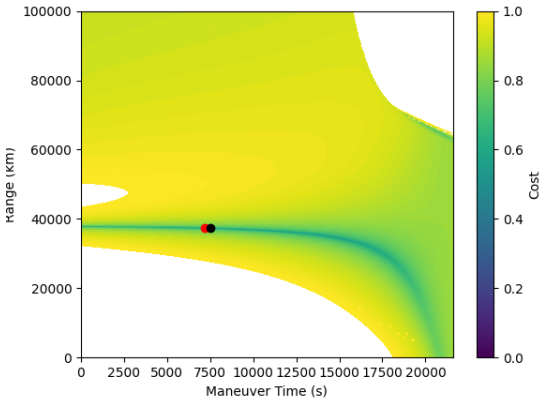

Fig. 17: The log error contour for feasible trial values for a solved maneuver determination case. The red dot indicates the true solution and it sits at the bottom of the contour valley.

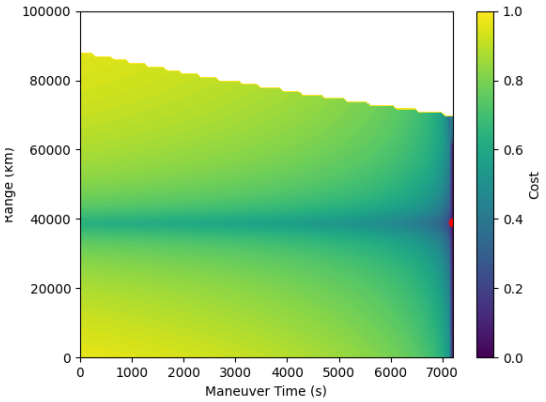

Fig. 18: The log error contour for a case with the first observation at the same time of the maneuver. There is not a unique solution.

3.4 Stressing Cases

While the IMD algorithm has a high success rate in a wide variety of realistic cases, there remain cases where it fails. Understanding these failures can inform when the IMD algorithm is applicable and when another solution should be preferred. Some of these failing cases are due to fundamental limitations in what is computable given the information available.

The investigated cases are the temporal spacing of the maneuver and measurements, the magnitude of the maneuver, and the relative inclination of the observers and the target's post-maneuver orbit. A contour plot of the solution space is used to visualize how these factors impact the observability of the solution. As the algorithm computes a scalar error at each iteration, a trial value's effective cost can be visualized by plotting this scalar error against the trial value maneuver time and range to the second observation. Figure 17 shows the contour for a successfully solved nominal case. The red dot indicating the true solution of the actual maneuver performed without measurement and modeling errors. The black dot indicates the location of the smallest calculated error. A case that is not solvable will have a curve of minima rather than a single point. Error in the input state, measurements, and orbit modeling show up in a shifting of the contours relative to the true solution. The bottom of the contour valley would then not align with the true solution.

3.4.1 Time Spacing

The temporal spacing of the maneuver and measurements stresses the IMD algorithm by increasing the impact of input and modeling errors on the solution. Trivially, a first measurement that occurs at the same time as the maneuver cannot provide any information about the post maneuver orbit. Only the two constraints from the second observation provide any useful information. Accordingly, there would only be two constraints on the post maneuver orbit with four unknowns from the time of the maneuver and three delta-v components.

This case is shown in figure 18. The initial state is two hours before the maneuver and first measurement, so the solution lies on the maximal bounds for possible maneuver time. At this point there is a vertical line of equal minima along the maximal time bound and it is not possible to determine the correct range of the second observation.

While the case of the simultaneous maneuver and observation is trivially indeterminable, close times between the maneuver and observation significantly increase the impact of any errors and can lead to dramatically incorrect solutions. This can also occur if the observations occur very close in time to one another and from similar positions such as the case if the same instrument makes both observations. In this case the orbit has little time to evolve between observations and the small changes can easily be washed out by measurement error.

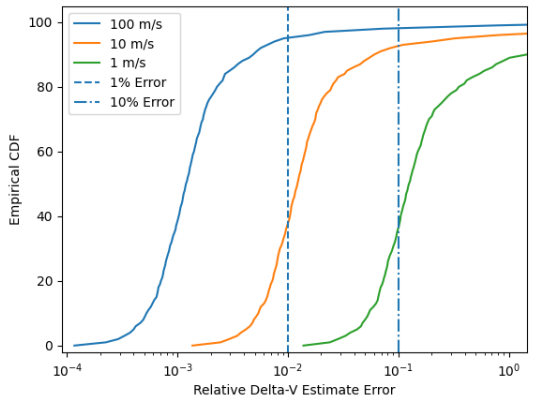

Fig. 19: CDF of normalized maneuver vector estimate errors for varying maneuver magnitude.

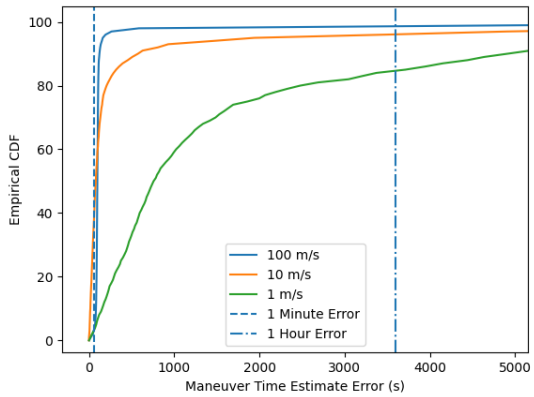

Fig. 20: CDF of maneuver time estimate errors for varying maneuver magnitude.

3.4.2 Maneuver Magnitude

Similar to how brief periods between the maneuver and first measurement can restrict the observability of the new orbit, smaller maneuvers are more difficult to accurately estimate. Once again, this is due to the larger relative impact of errors on the resulting estimate. The decrease in the magnitude of the maneuver increases the relative error for the same absolute error in the estimate. Figures 19 and 20 show the impact of decreasing maneuver magnitude on the cumulative distribution function of relative maneuver errors and absolute time error. As expected, relative errors increase as the maneuver magnitude increases. With a maneuver magnitude of 100 m/s, 98.2 percent of test cases were accurate to within ten percent of the true maneuver. By the point the maneuver has dropped to 1 m/s, only 36.7 percent of cases have relative errors below ten percent.

Figure 21 shows the error contour plot for a case with no maneuver. Rather than a minimal point, there is a zero-valued curve crossing through all times with ranges corresponding to zero magnitude maneuvers. The algorithm will appropriately produce a result of some very small maneuver at a random point along this curve.

3.4.3 Relative Inclination

In angles-only IOD, a target's orbit is not observable if it is coplanar to the observers. Hu et al. showed in their 2020 paper that the beacon-spacecraft OD problem was second-order local weakly observable [12]. As the proof only relied on the relative inclination of the observers and target as well as the angles-only measurement type, it extends without modification to the IMD problem. Accordingly, the maneuver of the target object is not observable if the observations are taken by instruments in the orbital plane of the target after its maneuver. Figure 22 shows the error contour for a case with coplanar target and observations. The minimal curve arcing through the solution space has zero value at all points along its path and the true solution is indistinguishable from any other point along its path. The scenario is similar to the non-coplanar case shown in figure 17. However, in figure 17 the minimal curve trends up as it moves away from the minimum at the solution. By contrast, the coplanar case remains constant all along the curve. This demonstrates how the solution is not observable if the observers are within the post-maneuver orbital plane of the target.

4. CONCLUSION AND FUTURE WORK

A method for initial maneuver estimation using the minimal number of angles-only measurements was described and analyzed in detail. The algorithm's accuracy and reliability under realistic simulated dynamics, maneuvers, and errors was assessed and described. The IMD algorithm is most applicable in cases with large maneuvers and a rapid deviation away from its nominal orbit. In these cases, the amount of correlated data and the time before the object is lost is the

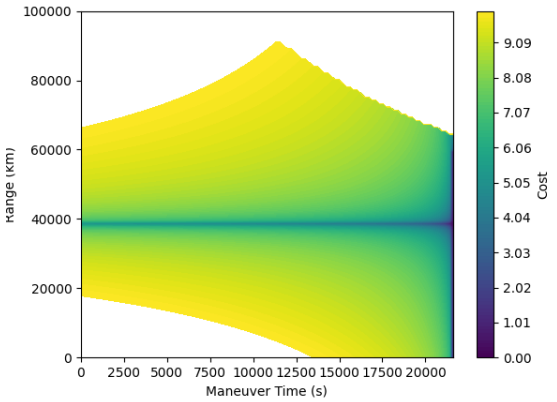

Fig. 21: The log error contour for a case with no maneuver. The minimal curve occupy the final time step as well as the set of points corresponding to no maneuver. While the contour shows a set of solutions, all of them equivalently indicate no maneuver.

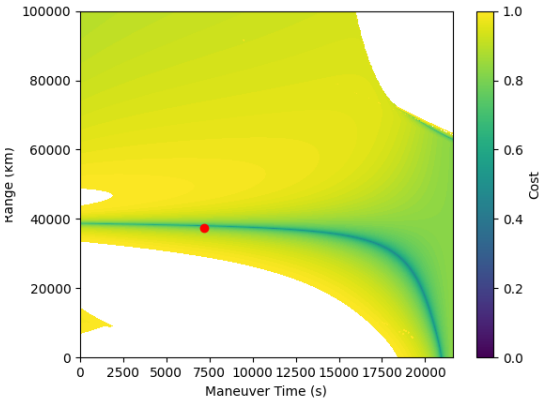

Fig. 22: The log error contour for the case with coplanar observers and post-maneuver target. The minimal curve is equally zero at all points and the algorithm cannot find the true minimum marked by the red dot.

most limited. The IMD algorithm performed best in these cases with 98.2 percent of simulated cases with 100 meters per second of delta-v yielding results accurate to within ten percent of the true maneuver. The accuracy varied based on a wide variety of factors including observation geometry, data errors, maneuver magnitude, and engine thrust. The impact of each of these factors and more were analyzed and discussed to characterize the overall performance of the IMD algorithm in a wide variety of scenarios.

The IMD algorithm is most applicable in cases where there is insufficient data or time to use more robust algorithm that accurately models the true orbital dynamics and a finite maneuver. However, its applicability is not limited to only those cases. Even when more data is available, the IMD algorithm remains useful as a tool for generating accurate initial guesses to help these algorithms more quickly and reliably converge on an accurate solution. Follow on work will include an assessment and quantification of the impact of the IMD algorithm as a tool for seeding higher fidelity solutions.

[1] Xingyu Zhou, Tong Qin, and Linzhi Meng. Maneuvering spacecraft orbit determination using polynomial representation. *Aerospace*, 9(5):257, 2022.

[2] Stijn Lemmens and Holger Krag. Two-line-elements-based maneuver detection methods for satellites in low earth orbit. *Journal of Guidance, Control, and Dynamics*, 37(3):860–868, 2014.

[3] RH Gooding. *A new procedure for orbit determination based on three lines of sight (angles only)*. Defence Research Agency Farnborough, UK, 1993.

[4] David A Vallado. Evaluating gooding angles-only orbit determination of space based space surveillance measurements. *Paper Usr*, 10:1–21, 2010.

[5] ANSYS. Orbit determination toolkit.

[6] CS GROUP. Orekit.

[7] Giulio Avanzini. A simple lambert algorithm. *Journal of guidance, control, and dynamics*, 31(6):1587–1594, 2008.

[8] Nitin Arora and Ryan P Russell. A fast and robust multiple revolution lambert algorithm using a cosine transformation. *Paper AAS*, 13(728):162, 2013.

[9] Dario Izzo. Revisiting lambert's problem. *Celestial Mechanics and Dynamical Astronomy*, 121:1–15, 2015.

[10] F Landis Markley. Kepler equation solver. *Celestial Mechanics and Dynamical Astronomy*, 63:101–111, 1995.

[11] Eric W Deutsch. Empirical uncertainty estimators for astrometry from digital databases. *The Astronomical Journal*, 118(4):1882, 1999.

[12] Yunpeng Hu, Inna Sharf, and Lei Chen. Distributed orbit determination and observability analysis for satellite constellations with angles-only measurements. *Automatica*, 129:109626, 2021.

A. APPENDIX: ALGORITHM PSEUDOCODE

Algorithm 1 Initial Maneuver Determination - Varying time implementation

$t_m \leftarrow \frac{t_1+t_0}{2}$
$\rho_2 \leftarrow \left|\vec{r}_0 - \vec{R}_2\right|$
converged $\leftarrow$ False
tolerance $\leftarrow 1e-14$
itr $\leftarrow 0$
maxItr $\leftarrow 20$
$\varepsilon \leftarrow 1e-6$
while not converged and itr $<$ maxItr **do**
    $\vec{r}_m, \vec{v}_{-m} \leftarrow$ propagate$(\vec{r}_0, \vec{v}_0, t_m - t_0)$ ▷ This propagate step is analytic when using eccentric anomaly.
    $\vec{r}_2 \leftarrow \vec{R}_2 + \rho_2 \hat{p}_2$
    $\vec{v}_{+m}, \vec{v}_2 \leftarrow$ lambert$(\vec{r}_m, \vec{r}_2, t_2 - t_m)$
    $\vec{r}_1, \vec{v}_1 \leftarrow$ propagate$(\vec{r}_m, \vec{v}_{+m}, t_1 - t_m)$
    $\vec{C} \leftarrow \vec{r}_1 - \vec{R}_1$
    $CR \leftarrow \hat{p}_1 \cdot \vec{C}$
    $\vec{u} \leftarrow \hat{p}_1 \times \vec{C}$
    $\hat{P} \leftarrow \frac{\vec{u}\times\hat{p}_1}{|\vec{u}\times\hat{p}_1|}$
    $\hat{N} \leftarrow \frac{\hat{p}_1\times\hat{P}}{\hat{p}_1\times\hat{P}}$
    $F \leftarrow \hat{P} \cdot \vec{C}$
    $\vec{C}_{-x} \leftarrow$ getC$(\vec{r}_0, \vec{v}_0, \vec{R}_1, \vec{R}_2, t_0, t_1, t_2, \hat{p}_1, \hat{p}_2, \rho_2 - \rho_2\varepsilon, t_m)$
    $\vec{C}_{+x} \leftarrow$ getC$(\vec{r}_0, \vec{v}_0, \vec{R}_1, \vec{R}_2, t_0, t_1, t_2, \hat{p}_1, \hat{p}_2, \rho_2 + \rho_2\varepsilon, t_m)$
    $\vec{C}_{-y} \leftarrow$ getC$(\vec{r}_0, \vec{v}_0, \vec{R}_1, \vec{R}_2, t_0, t_1, t_2, \hat{p}_1, \hat{p}_2, \rho_2, t_m - t_m\varepsilon)$
    $\vec{C}_{+y} \leftarrow$ getC$(\vec{r}_0, \vec{v}_0, \vec{R}_1, \vec{R}_2, t_0, t_1, t_2, \hat{p}_1, \hat{p}_2, \rho_2, t_m + t_m\varepsilon)$
    $J \leftarrow \begin{bmatrix} \frac{\hat{P}\cdot\vec{C}_{+x} - \hat{P}\cdot\vec{C}_{-x}}{2\rho_2\varepsilon} & \frac{\hat{P}\cdot\vec{C}_{+y} - \hat{P}\cdot\vec{C}_{-y}}{2t_m\varepsilon} \\ \frac{\hat{N}\cdot\vec{C}_{+x} - \hat{N}\cdot\vec{C}_{-x}}{2\rho_2\varepsilon} & \frac{\hat{N}\cdot\vec{C}_{+y} - \hat{N}\cdot\vec{C}_{-y}}{2t_m\varepsilon} \end{bmatrix}$
    $\rho_2 \leftarrow \rho_2 - J_{2,2}\frac{F}{|J|}$
    $t_m \leftarrow t_m + J_{2,1}\frac{F}{|J|}$
    if $\frac{F}{|\vec{r}_1|} <$tolerance **then**
        converged $\leftarrow$ True
    itr $\leftarrow$ itr+1
$\delta\vec{v} \leftarrow \vec{v}_{+m} - \vec{v}_{-m}$
return $t_m, \delta\vec{v}$

Algorithm 2 getC function

$\vec{r}_m, \vec{v}_{-m} \leftarrow$ propagate$(\vec{r}_0, \vec{v}_0, t_m - t_0)$
$\vec{r}_2 \leftarrow \vec{R}_2 + \rho_2 \hat{p}_2$
$\vec{v}_{+m}, \vec{v}_2 \leftarrow$ lambert$(\vec{r}_m, \vec{r}_2, t_2 - t_m)$
$\vec{r}_1, \vec{v}_1 \leftarrow$ propagate$(\vec{r}_m, \vec{v}_{+m}, t_1 - t_m)$
$\vec{C} \leftarrow \vec{r}_1 - \vec{R}_1$
return $\vec{C}$