Risk Maps for Conjunction Potential Throughout the Cislunar Domain

Alexander Koenig
Massachusetts Institute of Technology
Phillip M. Cunio
ExoAnalytic Solutions, Inc.
Juliana Chew
Massachusetts Institute of Technology

ABSTRACT

To achieve persistent cislunar space domain awareness (SDA), adequate coverage of and revisit rates for all of cislunar space are required. Not all areas of cislunar space should be weighted equally in terms of this coverage, however: the depth of the coverage should vary depending on parameters such as proximity to space-based infrastructure or potential risk posed by the objects in question. To create such a weighting system, we develop a "risk map" within this paper, which captures the time it would take an object at any position in cislunar space to reach a set of a few destinations of interest where relevant infrastructure currently resides or likely will reside in the near future. This "transit time" is the theoretical minimum time it would take for an object in a given region to intercept any individual target in the set of space-based infrastructure assuming an upper bound on the magnitude of the initial maneuver used to achieve the transfer.

This risk-map-based weighting system can be used to better tune our systems to establish coverage of regions of cislunar space in a manner consistent with the risk profile the objects in those regions pose. That is not to say that revisit rates only exactly corresponding to the transit times are required; plenty of buffer is required to pick up movement in order to acquire and maintain custody, and to not have to perform wide searches after missed observations. However, the risk map assists in evaluating where coverage is most critical, and creates a new characterization of the shape of cislunar space in terms of its transit times.

This paper presents a method for calculating such a risk map under the assumption of Keplerian motion for orbit transfers, discusses a potential implementation considering the complete three-body system, and shows results from a software tool developed to calculate a risk map if given initial inputs.

1. INTRODUCTION AND BACKGROUND

The cislunar regime extends from above the Geosynchronous Earth Orbit (GEO) neighborhood to the 2nd Lagrange point on the far side of Luna, Earth's moon, and is likely in the near future to contain many key transit routes and infrastructure nodes as humans continue to expand operations beyond Low Earth Orbit (LEO) and toward the entire solar system. As such, achieving sufficient situational awareness with frequent enough state updates to support the complex dynamics associated with operations within and affording a viable space traffic management (STM) scheme for these routes and nodes will be a key way to support the burgeoning trillion-dollar space economy, of which cislunar infrastructure, operations, and logistics are expected to be a large part [1].

Key regions of space around the Earth have been understood as ideal locations to emplace infrastructural elements since before the end of World War II [3]. Modern-day reality admits to the many benefits of infrastructure with global scale and reach, and there is no reason to doubt that humans will continue to utilize spaceborne infrastructure at increasing rates in the foreseeable future.

Infrastructure may be understood to include many means of distributing or managing valuable commodities, even intangible commodities, in a way that enables many economic or cultural activities. Infrastructure may be built and operated by governments (as the interstate highway system is) or by corporations (as the early-1900s private railway network in the United States was) or even by individuals (as mountain hiking shelters sometimes are). While

current spaceborne infrastructure primarily deals in information (communications packets, weather data, and pointing/navigation/timing service), future spaceborne infrastructure may well deal in pharmaceuticals, recreational experiences, or exotic/advanced materials. The infrastructural elements to create, store, disseminate, or utilize all of these tangible and intangible products may begin to be emplaced sooner than anyone today expects.

Infrastructure may also be distinguished by a generally large scale (relative to the size or number of commodities handled and customers served), and corresponding needs for monitoring, maintenance, and security. In the case of spaceborne infrastructure, maintenance may be achieved by active operation of spacecraft in various orbits, and security will likely be achieved by means of legal regulations and informal arrangements. Neither of these, however, is feasible at scale without a capable and detailed monitoring regime. Particularly given the tendency of infrastructure to expand alongside human activity and the renewed human exploration focus on the Moon and surrounding space, it is highly probable that infrastructural elements will soon be emplaced at distances up to or even beyond the lunar orbit radius (approximately 400,000 km from the Earth). This creates a cislunar sphere with a volume about 1000 times the volume of the current infrastructural region (cis-GEO space, below the geosynchronous orbit radius of approximately 42,000 km from the Earth's center). Monitoring all of this volume, at ranges of up to 10 times the range of the current communications satellite ring (and about 20 times the range of the GPS constellation), will be a notable challenge.

Accordingly, developing means to comprehend the challenge and address the most critical parts of this volume is of interest. For the purposes of this paper, infrastructure will be the subject of monitoring (that is, we assume it is more important to watch actual space objects which comprise or may affect the infrastructural networks operating in space than it is to watch actual empty space or the distant stars), and infrastructure will be understood as space objects that serve an infrastructural purpose (or those other objects that may interact with them).

Figure 1 shows a conceptual sketch of the types and possible locations of infrastructural elements in the cislunar sphere. Note the concentration at or near Lagrange points — this paper focuses on the moon and its Lagrange points as primary areas for infrastructural expansion, along with the present GEO belt which carries critical infrastructure today.

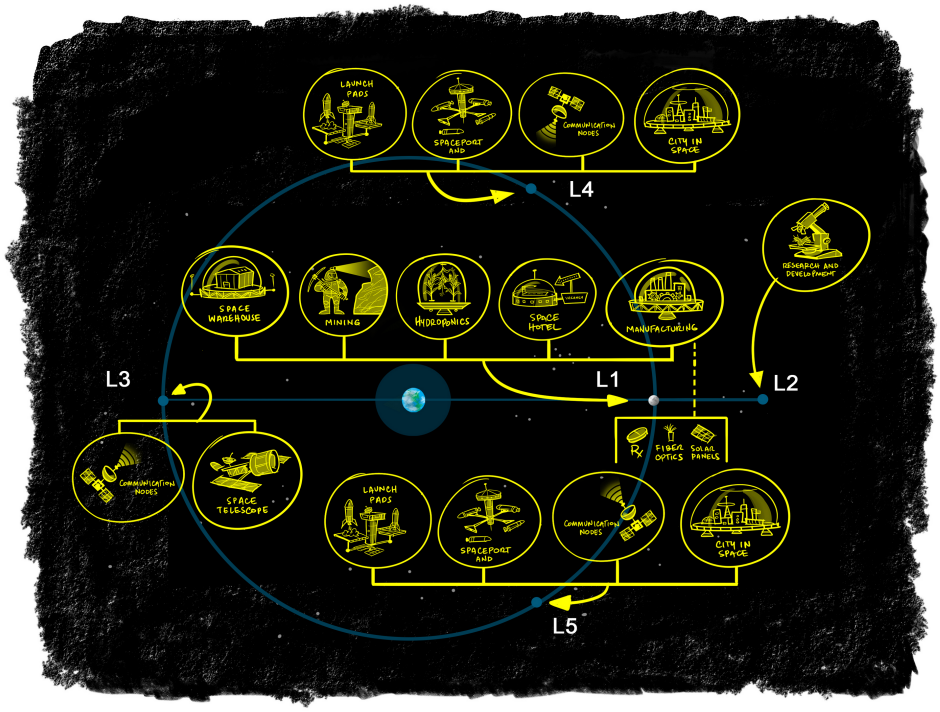

Fig. 1: Potential future infrastructure within cislunar space.[1]

[1]Original artist: www.danhamiltonart.com

Additionally, while infrastructure may be viewed as somewhat static and long-lived, traffic elements (consisting of resupply or repair for infrastructure, scientific expeditions, or humans themselves) will be increasingly likely to share the cislunar volume with infrastructural objects. Space traffic management (STM), the art of coordinating all such traffic to ensure the safety of all, will be a practice tied to infrastructure management and undergirded by the same type of in-depth monitoring capacity.

Monitoring the infrastructure and traffic within this volume is a challenge of allocating observation, processing, and attention resources thoughtfully. A key tool for such allocation is the use of schemes which identify natural regions of more or less likelihood of demanding rapid responses. This paper will develop such a tool, in the form of risk maps for the cislunar volume, which detail and contour levels of transit times from any location in the cislunar volume to the nearest likely infrastructure site as a proxy for the risk. After a discussion of appropriate approach methods, the general means of calculating these risk maps is presented, and the results are shown. Finally, impacts and import of the results are discussed.

2. RISK MAP CALCULATION METHODS

Building a risk map is a matter of iterating over several layers of computational loops. At the highest level, cislunar space over a region of interest is discretized into a grid of starting locations for analysis. For every point within this grid, the fastest time to reach any infrastructural node from that point (the "transit time") is then computed, achieved by iterating over a feasible range of initial velocity vectors and destination pairs and thereafter selecting the minimum time from that set. For each of these orbit and destination pairs, the two-point boundary value problem (TPBVP) is solved with increasing time increments until a transfer duration is found which does not exceed a prespecified ΔV limit (where only the initial transfer maneuver is considered, rather than an additional velocity-matching rendezvous maneuver). The TPBVP solver is at the core of the process, returning a transfer trajectory given an initial location, a destination, and a transfer duration. This overall process is shown in Figure 2.

2.1 Map Discretization

Although conceptually the risk map is continuous, there is no convenient analytic function to calculate the transit time; hence, the region considered (typically a spherical surface or spherical shell) is discretized. The transit time is determined at each point and a continuous map is recovered through nearest-neighbor interpolation. In this paper, we consider risk maps at spherical surfaces at various altitudes, which involves discretizing a 3D spherical shell represented in spherical coordinates. We use an equirectangular map projection to display the individual 2D risk maps, which involves simply assigning the right ascension (in map terms, the longitude) as the horizontal location, or abscissa, and the elevation (in map terms, the latitude) as the vertical location, or ordinate. To have equal spacing of points within the output projection, each risk map is discretized as a rectangular grid of points equally spaced across latitude and longitude. Both two- and three-body system dynamics are symmetric across the horizontal plane, so the risk map is computed for just one half of the system, which then may be reflected across this plane.

For this present study, this analysis of cislunar space has been discretized into 31 altitude slices ranging from $0.25\times$ to $1.5\times$ mean Earth-moon distance, 96 right ascension (longitude) slices ranging from 0° to 360°, and 48 elevation (latitude) slices ranging from 0° to 180°. This discretization comprises in excess of 1.4×10^5 individual points for computation, a computational challenge discussed further in Section 2.6.

2.2 Destination Selection

Each of the discretized points within the map represents a hypothetical starting location for a spacecraft. The transit time at that location is the minimum of the durations it would take the spacecraft to reach each infrastructural node, given a ΔV constraint. These durations are determined individually and the minimum value is taken from the set, thereby considering the worst-case scenario.

The destinations considered in this paper include the GEO belt, the Earth-moon Lagrange points, and the moon itself, as these locations are the most likely to harbor future infrastructure (or already do presently).

In the two-body, inertial case, the Earth-moon system dynamics are rotationally symmetric. Since the infrastructural nodes are at one of two altitudes — either roughly lunar altitude, as for the moon and Lagrange points, or at the altitude of the GEO belt — the risk map can be computed for just one representative destination at each altitude and rotated to create the data for the other destinations. This step is particularly important for the GEO belt, which would otherwise require discretization into numerous distinct locations, significantly impacting runtime.

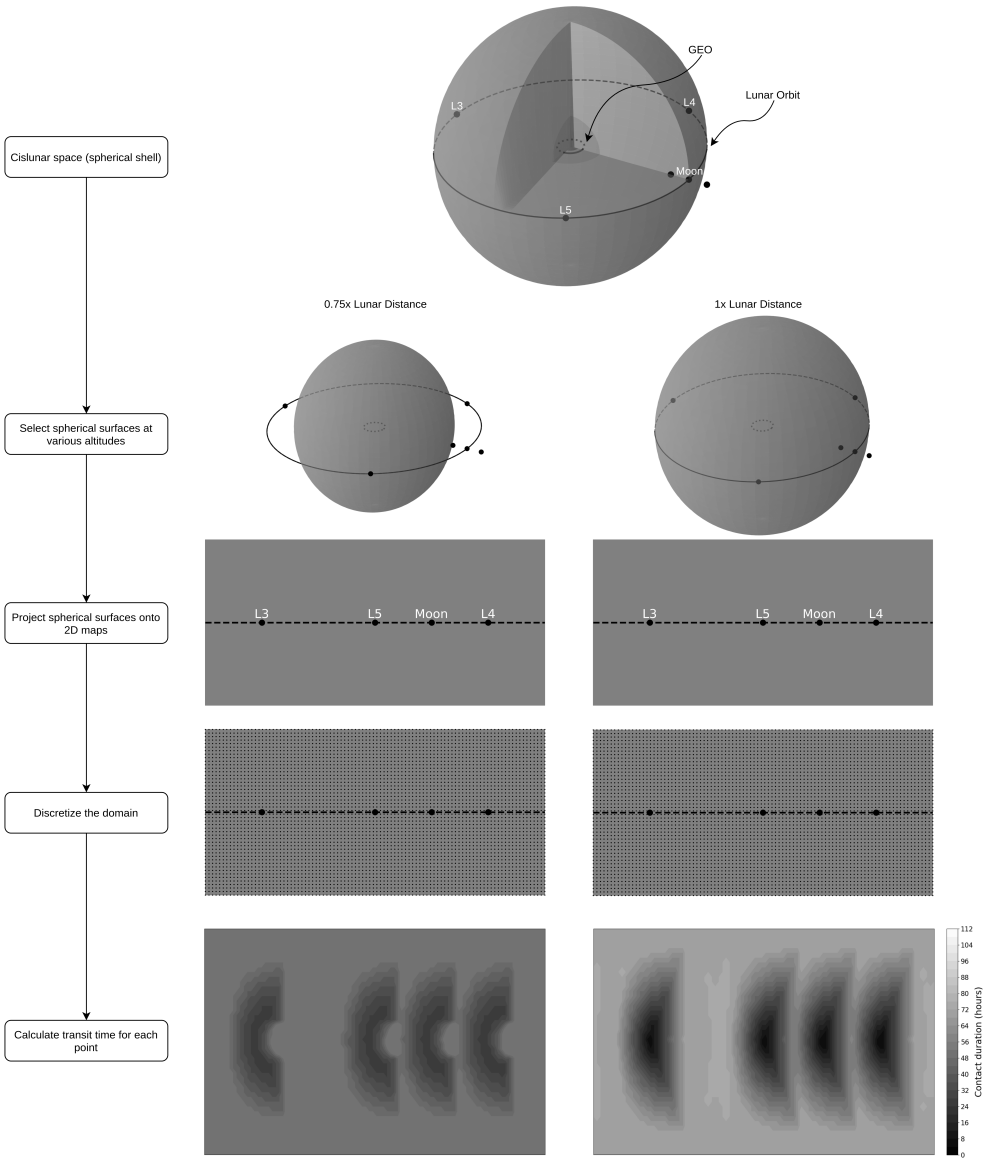

Fig. 2: Risk map creation process.

Here, we also group L1, L2, and the moon into a single destination for the sake of runtime. Realistically, infrastructure is likely to be spread out into periodic orbits about these points rather than grouped exclusively at the points themselves, so keeping these locations defined separately for the sake of accuracy is not a primary concern, although it does affect the output map.

2.3 Initial Orbit Selection

To solve the TPBVP given a ΔV constraint, we have to know not only the starting location but also the initial velocity from which the transfer maneuver is performed. Desiring to remain agnostic to the initial orbit, assumed to be circular yet unknown, we consider a number of initial trajectories passing through the starting location of interest. Some constraints must be imposed on the search, else the maneuvers could be constructed with arbitrarily high velocities and thus would achieve arbitrarily short transit times. For simplicity, only circular initial orbits are considered in this paper. Although considering elliptical starting trajectories may further reduce the ΔV needed to reach a particular destination, it prohibitively impacts the computational runtime required for this effort due to the additional degree of freedom in the search space. To compensate for this limitation, the ΔV constraint is somewhat increased to allow an

"ellipticalization" budget.

2.4 Transfer Duration Determination

Given each starting location and destination pair, the transit time is calculated by iterating over the TPBVP solver with increasing durations until a transfer trajectory not exceeding the ΔV constraint is found. This paper uses 4-hour increments, neatly corresponding to half a typical work shift. For computational efficiency, if a certain threshold duration is exceeded, the loop is ended with no duration found. Since the minimum value is taken from the set of all transfer durations at a particular location, these no-result cases do not make their way up to the risk map as a whole (where they do, the threshold duration is increased and the risk map is recomputed).

2.5 ΔV Selection

To select the ΔV constraint on transfers, we first allow a budget for orbit ellipticalization since only initial circular orbits are considered, as discussed in Section 2.3. Achieving an eccentricity e=0.5 from a circular orbit at lunar distance requires a ΔV of 230 m/s. This quantity is an effective ΔV and does not necessary represent a propulsive maneuver the spacecraft would actually perform, but rather is intended to alleviate the limitations of the circular orbit assumption.

On top of the ellipticalization budget, we consider what a reasonably swift spacecraft would be capable of performing in a single propulsive maneuver. Only chemical propulsion is relevant, as electric propulsion would operate on much longer timescales. With an I_{sp} of 350 s, if the spacecraft were to utilize 20% of its total mass for fuel, it could achieve a ΔV of around 750 m/s. Although this is a considerably large fuel mass budget for just a single propulsive maneuver within the spacecraft's overall lifetime, a cislunar SDA system should be able to robustly handle the worst-case scenarios, and such a maneuver is within the realm of possibility. Note that this ΔV is only the initial transfer; a velocity-matching rendezvous maneuver is not considered.

Adding the two ΔV budgets provides a ΔV limit of $\sim$1 km/s, which will be the scenario analyzed in this paper. To see how the risk map changes with different ΔV limits, we also investigate the 650 m/s and 300 m/s scenarios (approximately $2/3$ and $1/3$ of the original limit, respectively) at a subset of the risk map altitudes.

2.6 TPBVP Solver: Two-Body versus Three-Body Dynamics

The TPBVP solver takes as inputs the transfer duration, initial location, and destination, then outputs the required transfer trajectory. The method must be able to smoothly handle elliptic, parabolic, and hyperbolic trajectories. In the two-body case, this requires solving Lambert's problem, to which the series solution is known analytically. In the three-body case, no analytic solutions exist; popular methods include the multiple shooting method and nonlinear least-squares optimization.

Due to the layered nature of the risk map computations, on the order of 10^7 to 10^8 TPBVP solves are required in order to achieve a reasonable level of fidelity across the entirety of cislunar space; computational efficiency is a key consideration. The number of required solves is brought down by 1-2 orders of magnitude by leveraging planar and rotational symmetries as well as other computational shortcuts. Rotational symmetry is only available in a two-body approach, and furthermore, two-body TPBVP solvers are considerably faster than their three-body counterparts. For these reasons, we downselected to the two-body approach. The Python astrodynamics library Poliastro [5] was used as the Lambert solver, since it can straightforwardly handle elliptic, parabolic, and hyperbolic orbits. While an implementation of a three-body nonlinear least-squares optimization TPBVP solver was tested, described in Section 2.7, the runtime of the solver was such that the complete risk map would take many months to run.

While it precludes certain maneuver optimizations, using a two-body approach, and thereby ignoring lunar gravity, may still be is a reasonably justifiable approximation to build a realistic risk map. On one hand, all minimum-time transfer trajectories involve high velocities and low durations, and thus are less affected by weak perturbations such as lunar gravity. For those trajectories in close proximity to the moon, the additional gravity well could only be expected to increase transit times to all other destinations, given that we are only interested in the minimum duration to all infrastructural nodes. Although the destination list does include regions where third-body effects are prominent (i.e., the moon and its Lagrange points), trajectories that take advantage of these strategic dynamics to reach a particular destination are likely to be precluded in favor of higher energy transfers anyway.

2.7 Three-Body TPBVP Solver

Several algorithmic methods to solve the three-body TPBVP exist, but we downselected to nonlinear least squares optimization for its relatively high computational efficiency, as demonstrated by [2], where such a solver was implemented

specifically for the Earth-Moon system. As per their method, we formulated the restricted three-body TPBVP solver using the recently-developed Theory of Functional Connections (TFC) [4]. This TFC-based solver finds a trajectory that both meets given boundary constraints and best fits the three-body equations of motion by performing nonlinear least squares regression on adaptable free functions. Figure 3 provides an outline of the procedural strategy.

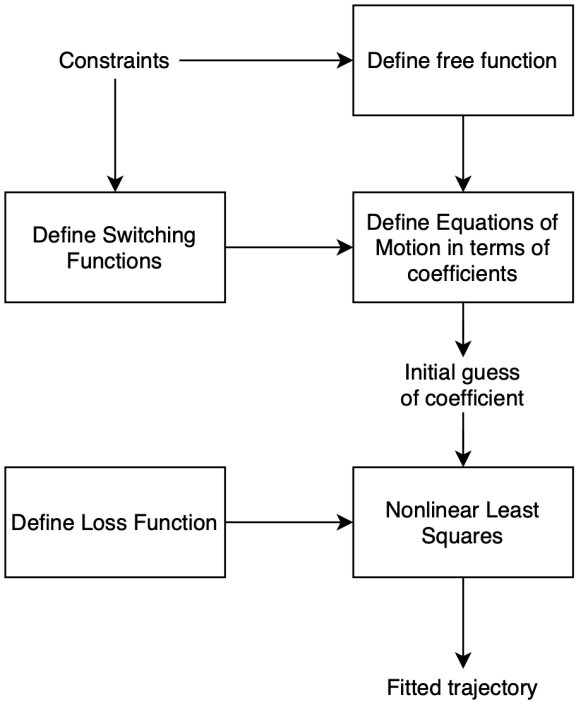

Fig. 3: A diagram of the implementation of a TFC-based three-body TPBVP solver.

First, the space of possible orbit solutions is defined by specifying separate constraints (switching functions) and free functions — the constraints match the position boundary values and are therefore constant, whereas the free functions are adaptable in order to fit the equations of motion. Next, the free functions are specified with a set of basis functions and corresponding coefficients; it is these coefficients which are iteratively refined via nonlinear least-squares optimization to find a satisfactory orbit. For this regression, the loss function for the equations of motion is defined, which computes the generated orbit's consistency with the equations of motion.

The position along each axis is defined as

$$X_i(t, g_i(t)) = g_i(t) + \sum_{j=1}^{k} \phi_j(t)\rho_j(t, g_i(t)), \tag{1}$$

where X_i defines the location of a body on the i^{th} axis, and $g_i(t)$ is the free function for the i^{th} axis. The k equations of $\phi_j(t)$ are switching functions for the k given constraints, which ensure the boundary conditions are met. The ρ_j equations are the projection functionals that project $g_i(t)$ onto the domain of functions that meet the constraints. For more information about these equations, refer to [2].

Each $g_i(t)$ is defined in terms of unique coefficients ξ_i, which are collectively referred to as Ξ; these coefficients are weightings of each vector in the set of basis vectors. When performing nonlinear least squares, the TFC solver iteratively refines Ξ to find a satisfactory orbit.

Nonlinear least squares regression was performed using Python's SciPy package [6], with $L(\Xi)$ as the loss function:

$$L(\Xi) = \Big[\widetilde{F}_1(t_0,\Xi),...,\widetilde{F}_1(t_i,\Xi),...,\widetilde{F}_1(t_f,\Xi), \\ \widetilde{F}_2(t_0,\Xi),...,\widetilde{F}_2(t_i,\Xi),...,\widetilde{F}_2(t_f,\Xi), \\ \widetilde{F}_3(t_0,\Xi),...,\widetilde{F}_3(t_i,\Xi),...,\widetilde{F}_3(t_f,\Xi)\Big] \tag{2}$$

where $\widetilde{F}_i$ is residual from the equation of motion for the i^{th} axis defined in terms of the time t_j and coefficients Ξ.

The original implementation paper [2] achieved runtimes on the order of seconds for each TPBVP solve in the Earth-Moon system. While our implementation achieved similar runtime, this computational speed is much too slow for generating the risk map, and in the end this approach was abandoned. Even if only one TPBVP solve was required per starting location and destination pair, the complete map would require a minimum of two months to generate. (Realistically, due to the time incrementation method, an average of 10 required solves per pair is typical, since the time step is 4 hours and the mean transfer time is roughly 40 hours). Furthermore, the solver tended to produce non-ideal solutions due to Runge's phenomenon at the boundary points, which prohibited accurate ΔV estimation even when run on small subregions of cislunar space. To solve the latter problem of Runge's phenomenon, rather than using high-order polynomials to approximate the transfer trajectories, using quadratic splines would decrease velocity error at the boundaries. One complication is that a spline-based solution increases the number of coefficients required to be optimized, and therefore increases runtime.

While the scope of this work was too limited to permit further investigation into runtime optimization for the TFC-based three-body TPBVP solver, there are a number of possible routes for future work in this area. The initial transfer trajectory guess used in the nonlinear optimizer is an important variable in its runtime; we simply began with a guess of a linear trajectory from the starting location to destination. Using the two-body solution as an initial guess of the transfer trajectory would potentially lessen the number of iterations required in the nonlinear least squares optimizer (although determining the coefficients Ξ that represent the two-body transfer trajectory solution is a separate regression problem on its own). For starting locations nearby an already-run location, the three-body solver output of the already-run location could be used as an initial guess instead. Furthermore, the transfer time search space could be reduced to be within a certain range of the two-body solution, decreasing the overall number of solves required. Additional shortcuts would likely be necessary to achieve a feasible runtime, however.

3. RESULTS

3.1 2D Risk Maps

2D risk maps were computed at 31 equally-spaced altitudes ranging from 0.25× to 1.5× the mean Earth-moon distance. These maps show projections of different altitude slices of cislunar space onto a 2D plane. For brevity, only four of these maps are displayed in Figure 4; the rest are utilized in Section 3.2 to display a 3D risk map.

These maps reveal a few different characteristics of transit times within cislunar space. The time to transit to GEO at any particular altitude is essentially isochronal across the spherical surface at that altitude, and therefore all altitudes have a unique background "floor duration" where the fastest transit time represents heading directly to the GEO belt. At 1× the Earth-moon distance, the floor duration is 72 hours; at 0.5×, the floor duration is 40 hours. Most of the map comprises this floor duration; only in a subset of the map (different regions centered around each of the Lagrange points) is it faster to transfer to infrastructural nodes within cislunar space rather than the GEO belt. These regions are generally taller than they are wide; i.e., a spacecraft can approach cislunar infrastructural nodes faster from out-of-plane locations compared to equally far away in-plane locations.

At all altitudes below 0.6× the Earth-moon distance, it is faster to reach GEO than the moon or any Lagrange point, suggesting this transition altitude is a potential demarcation for the inner boundary of cislunar space, and also one of relevance for designing SDA systems. Below this transition altitude, the SDA system is more GEO-oriented; above this altitude, it is cislunar-oriented. Note, however, that a system desiring to address both infrastructural risks as well as space traffic management must be able to monitor objects transitioning between these regions, and therefore may avoid viewing one region or the other exclusively.

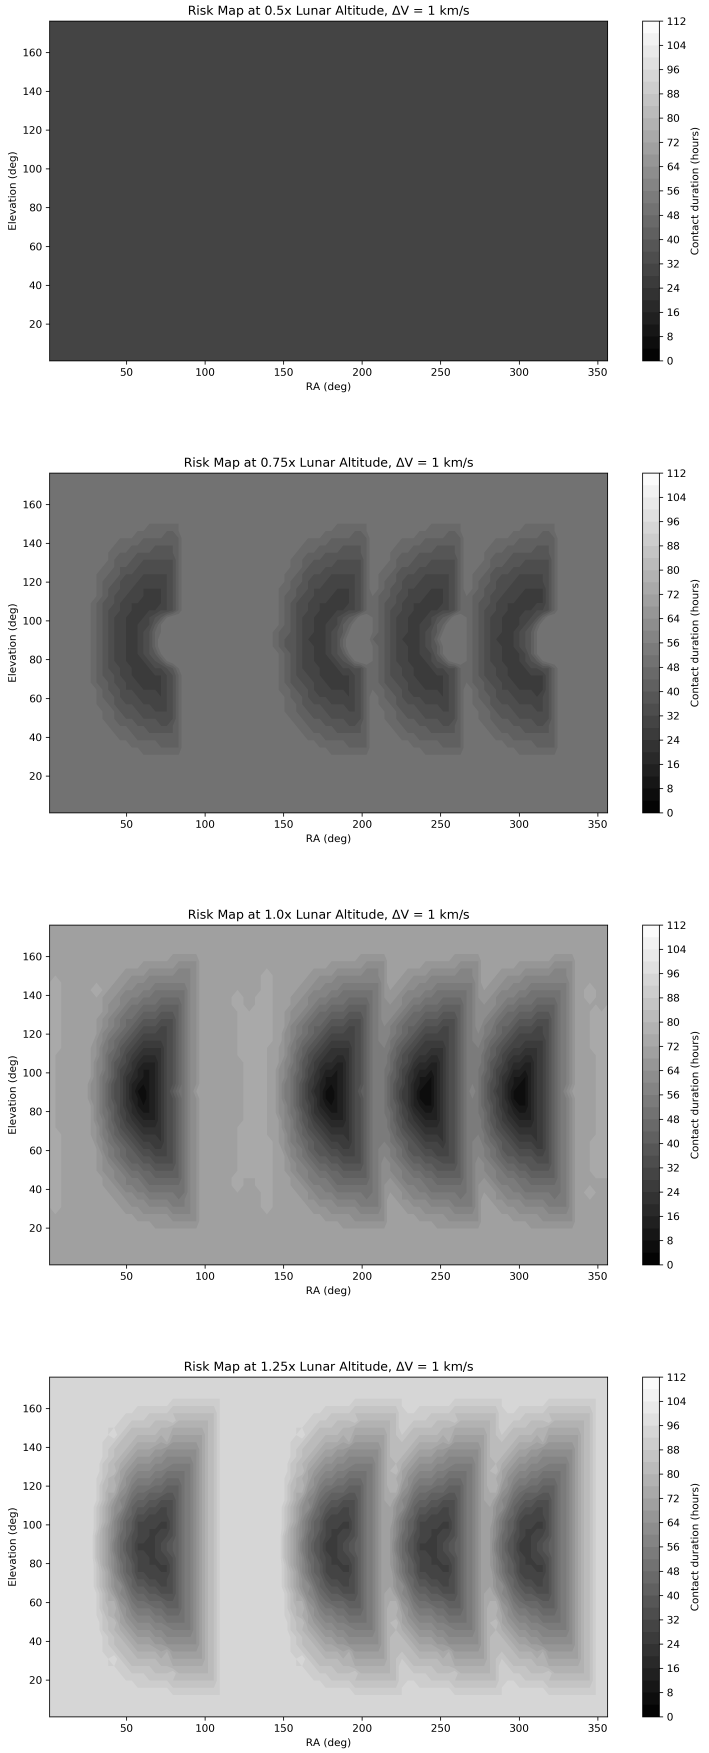

Fig. 4: Risk maps at increasing altitudes for ΔV = 1 km/s. The moon is located at 240° horizontally and 90° vertically. L4 and L5 are offset by ±60° from the moon, and L3 is offset by 180°.

Each of these findings regarding the background the floor duration and GEO-cislunar transition altitude alleviate one of the key challenges for cislunar SDA, which is that the volume in question is on the order of 1000$\times$ greater than the volume encompassed by the GEO region. A considerably smaller volume of cislunar space may be considered if one is primarily interested in sub-regions of potentially rapidly-evolving risk. For example, in the 1$\times$ lunar altitude map, regions of $\leq$ 24 hour transit duration comprise only roughly 0.5 steradians, or less than 5% of the overall area which may otherwise be observed. Section 3.2, which displays results for 3D risk maps, explicitly quantifies the volume of these sub-regions of cislunar space.

Figure 5 displays risk map results for varied ΔV values of 650 m/s and 350 m/s, which are respectively 2/3 and 1/3 of the assumed nominal ΔV of 1 km/s. Note the different time color scale used for the 300 m/s case, which has transit times of up to 300 hours, compared to 112 hours in the 650 m/s case and 72 hours in the 1 km/s case. As expected, the transit times increase in the case of decreased spacecraft maneuverability. The general shape of the regions identified by the risk map remain similar in character, however — namely, there are crescent-shaped regions of potentially rapidly-evolving risk nearby infrastructural nodes, and a broad isochronal background region where the fastest transits are to GEO. (Although at sufficiently low ΔVs it becomes energetically impossible to reach GEO from cislunar space, such a maneuver capability is far from assumed the worst-case scenario).

Note that these latter maps for lower ΔVs are discontinuous in certain regions which seems likely to be an aphysical result arising from discontinuities between elliptic and hyperbolic transfer trajectories in the Lambert solver itself. These discontinuities are not present in the 1 km/s maps.

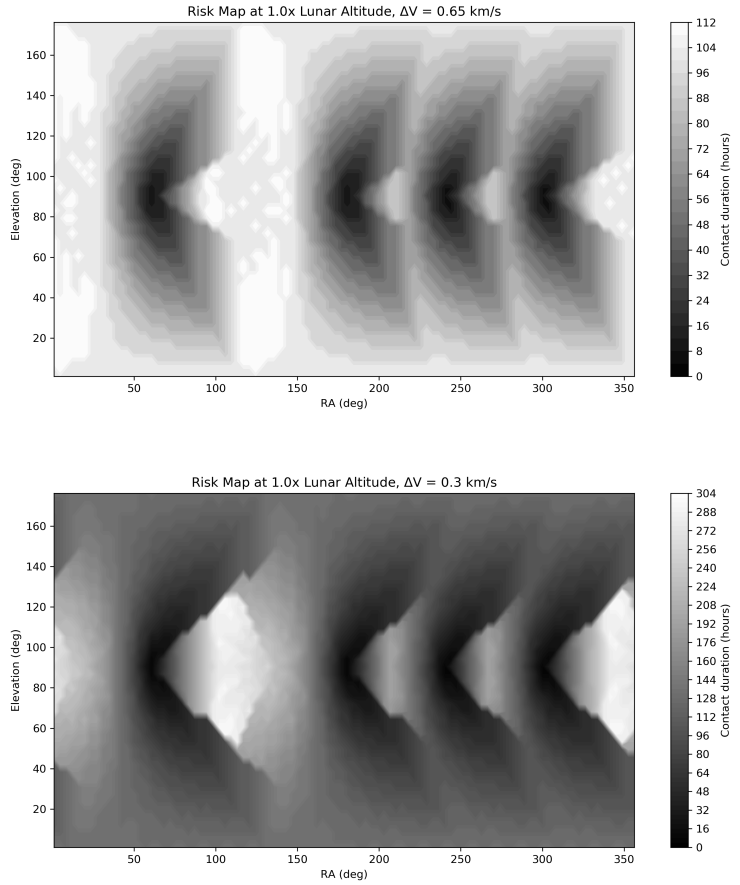

Fig. 5: The risk map at 1$\times$ lunar altitude for different ΔV limits (650 m/s, top, and 300 m/s, bottom).

3.2 3D Risk Maps

The 3D regions of cislunar space where the transit time is less than a certain threshold can also be shown with the risk map tool; Figure 6 displays these regions for transit times of $\leq$12 hours, $\leq$24 hours, and $\leq$36 hours. These visualizations display regions of potentially rapidly-evolving risk nearby infrastructural nodes, and therefore outline where heightened SDA coverage and faster observational revisit rates are most important. All distances are normalized to the Earth-moon distance (i.e., 1 unit = $3.844 \cdot 10^5$ km). The moon has coordinates of (1,0,0) and orbits counterclockwise as viewed from above.

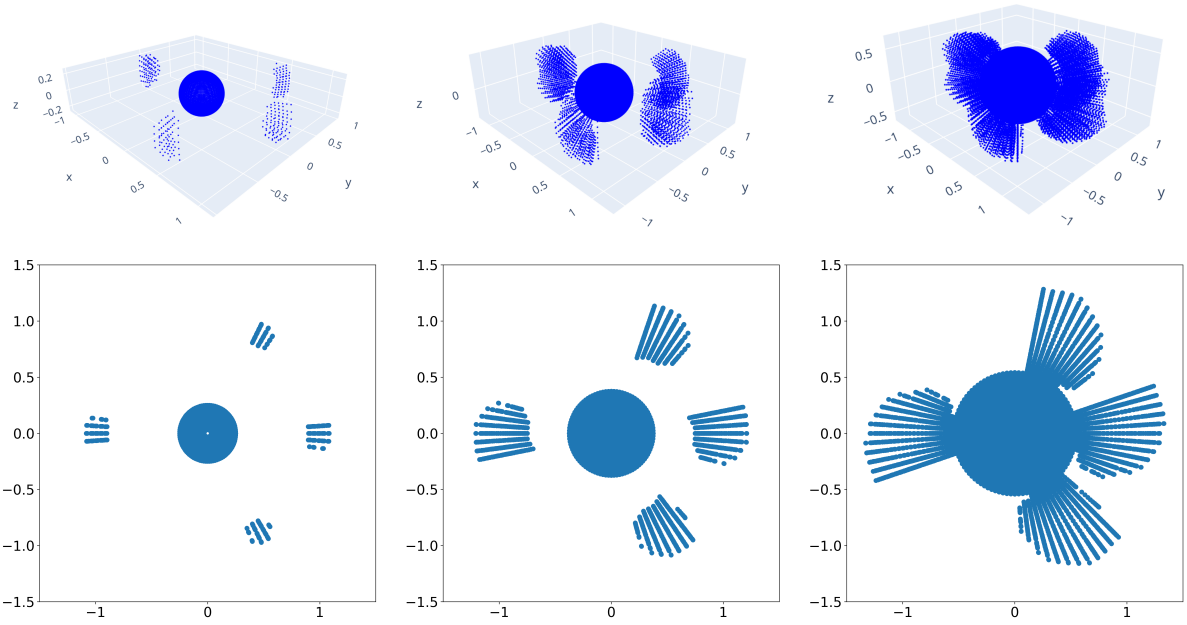

Fig. 6: Regions of transit times $\leq$12 hours (left), $\leq$24 hours (center), and $\leq$36 hours (right). Distances are normalized to the mean Earth-moon distance.

These maps generally have five regions: one essentially spherical region centered on the Earth where heading directly to the GEO belt represents the fastest risk, and four aspherical regions centered around the moon, L3, L4, and L5 where heading to each of those destinations respectively represents the fastest risk. Since this paper used a two-body approach, the four aspherical regions are identical in shape (in reality, three-body effects would somewhat alter the shape of each of these regions). Figure 7 shows the shape of the four aspherical regions and Table 1 quantifies their physical sizes.

Note that these regions are comprised of a grid of dots equidistant in each spherical coordinate axis rather than filled contours, which does not reflect physical reality but rather is a byproduct of the discretization method described in Section 2.1.

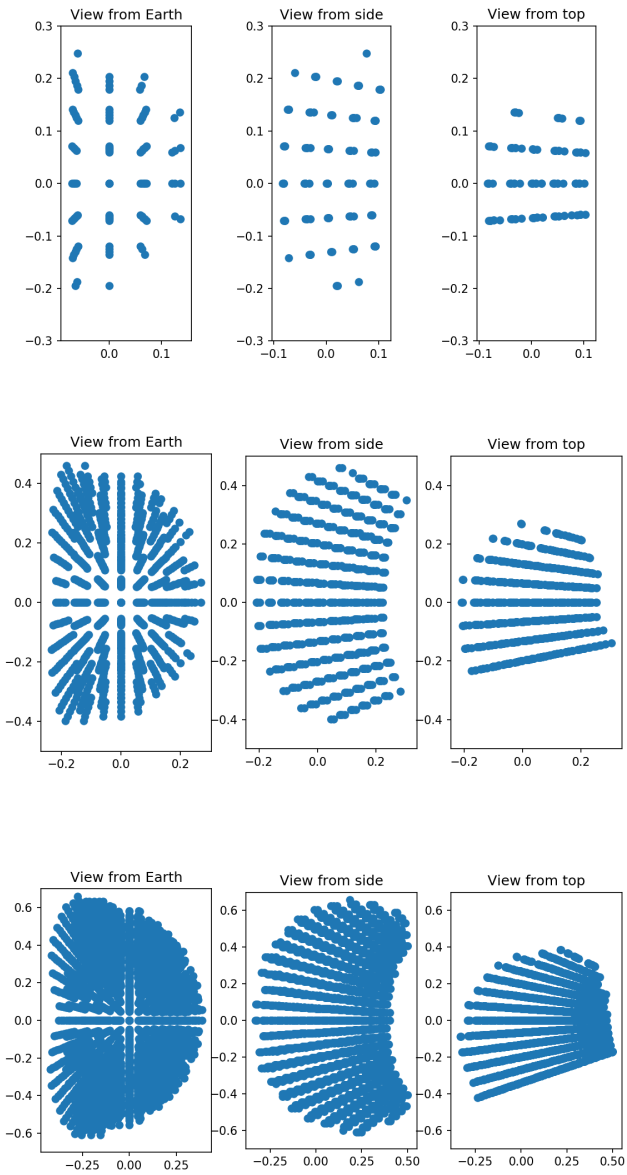

Fig. 7: The Lagrange/Moon regions of $\leq$12 hours transit time (top), $\leq$24 hours (center), and $\leq$36 hours (bottom). The perspective given in the left column is similar to the one given by the 2D risk maps in Section 3.1, though mirrored. The top-down perspective views the Earth-moon plane from above, and the side perspective displays the view from the remaining axis. Distances are normalized to the Earth-moon distance.

Table 1: Approximate physical extent of each Lagrange/Moon region.

Transit time	$\leq$12 hours	$\leq$24 hours	$\leq$36 hours
Height (km)	$1.5 \cdot 10^5$	$3 \cdot 10^5$	$4.5 \cdot 10^5$
Depth/width (km)	$7.5 \cdot 10^4$	$1.5 \cdot 10^5$	$2.9 \cdot 10^5$
Volume (km^3)	$9 \cdot 10^{14}$	$7 \cdot 10^{15}$	$3 \cdot 10^{16}$
Height (Earth-moon distance)	0.4	0.8	1.2
Depth/width (Earth-moon distance)	0.2	0.4	0.75
Volume (Earth-moon distance^3)	0.016	0.128	0.375

These results not only reconfirm that a considerably smaller volume of cislunar space may be considered if one is primarily interested in high-risk sub-regions, but also quantifies their volumes. For example, all regions of $\leq$12 hour transit times comprise a total volume only 60$\times$ larger than the sphere encompassing the GEO belt, compared to the approximately 1000$\times$ larger volume of the overall cislunar sphere.

4. DISCUSSION AND IMPLICATIONS

The results presented point to the importance of an integrated STM regime. Different regulations may apply at different altitudes given the difference in relative transit times, and this indicates that different frequencies of notice filing (and of sensor scanning to provide verification of current state) are also appropriate. These may be roughly analogous to airport or transit zones as understood in modern air traffic control: much as traffic obeys different rules in different volumes for other domains (consider the distinctions among airspace classes under FAA regulations), so might distinct rule regimes be utilized in STM. The FAA distinguishes between airspace classes to apply varying operational requirements appropriate to regions of varying risk; e.g., operations near airports involve stricter operational requirements so as to ensure appropriate risk management, whereas operations away from people, aircraft, and infrastructure are permitted greater freedom. This paper outlines what analogous regions in cislunar space may be, including their shapes, volumes, and relative levels of risk.

Fundamental STM practices are already used daily for many of the elements of spaceborne infrastructure that operate in the geosynchronous or low Earth orbit (GEO and LEO) regimes today, and are being extended and improved as these orbits become more crowded. While these same practices can be deployed in the entire cislunar volume, they will require updating for various types of infrastructure and longer observational ranges involved. (As an example, consider that the International Space Station's orbit, the most traffic-controlled region in space, is about 400 km from the Earth's surface and relies on radars to a notable extent, which is infeasible for cislunar-scale distances.) And much as monitoring infrastructure is critical, monitoring traffic will be as well.

The various contours of the risk maps presented in this paper show both clear, underlying structure and scaling of this structure with parameters such as total transit time and ΔV expended. They thus also suggest natural boundaries for STM regulatory regimes, in that certain regions may be considered more prone to risk, and resources both for monitoring and for updating guidance may be intelligently applied first to high-risk regions. Alternatively, more experimental activities (or more experimental configurations of monitoring and legal liability application) may be positioned in regions with natively lower risk contours.

As an example, a key finding of this paper is the floor altitude, where at a given radius every infrastructure element there or below may be considered threatened for a sufficiently short transit time parameter. For this paper, no transit time value below 12 hours was evaluated, as this was considered a short response time. Future work may evaluate the floor altitudes present when transit times are set much lower.

5. CONCLUSIONS

This paper provided the framework to calculate a cislunar space domain awareness (SDA) coverage weighting by using transit durations as a proxy for risk. The time it takes an unknown object in a particular region to reach an infrastructural node is considered here to be the relative risk inherent to that region in general. This work quantified the intuitive concept that in a situation of limited SDA resources or for long technology rollout timeframes, such

resources can be prioritized towards more important applications closer towards infrastructural nodes. Furthermore, to alleviate some of the core challenges associated with cislunar SDA such as the vast volume compared to regions with existing SDA coverage and corresponding long observational ranges, this weighting system can also be used to determine where and with what density to allocate SDA resources in the long-term.

Future work includes refining the risk map through the implementation of a three-body calculation approach, which was tested in this work but abandoned due to the insufficient computational efficiency of the solver. More work is required to enhance the runtime of the three-body two-point boundary value problem solver before use in this application. Outside of the risk map calculation methods, future work also remains in further understanding how to appropriately implement this concept for direct cislunar SDA and STM applications, such as in using the computed transit times to estimate appropriate observational revisit rates of each region of cislunar space by a cislunar SDA system.

6. ACKNOWLEDGEMENTS

We extend our thanks to several people who aided and advised this work, including Marcus Bever, Brien R. Flewelling, and Chris Ingram of ExoAnalytic Solutions, Inc., as well as Richard Linares of the Massachusetts Institute of Technology.

7. REFERENCES

[1] O'connell remarks to u.s. chamber of commerce. *Office of Space Commerce*, 2019. https://www.space.commerce.gov/oconnell-remarks-to-u-s-chamber-of-commerce/.

[2] A. Almeida, H Johnston, C Leake, and D Mortari. Fast 2-impulse non-keplerian orbit-transfer using the theory of functional connections. *European Physical Journal Plus*, 136(223), 2021.

[3] Arthur C. Clarke. Peacetime uses for v2. *Wireless World Magazine*, 1945. http://lakdiva.org/clarke/1945ww/.

[4] Daniele Mortari. The theory of connections: Connecting points. *Mathematics*, 5(4), 2017.

[5] J et al Rodríguez. *Poliastro: Astrodynamics in Python*. Zenodo, 2015. http://doi.org/10.5281/zenodo.17462.

[6] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.