(Preprint) AAS XX-XXX

EFFICIENT TRANSFERS SEARCH IN THE THREE BODY PROBLEM LEVERAGING QUASI PERIODIC ORBIT INTERSECTIONS

Matteo Santacesaria[*], Nathaniel Phillip Sailor[†], Andrea Capannolo[‡], and Andrea De Vittori[§]

Given the renewed interest in the cislunar environment and the Lunar Gateway, the necessity of more efficient strategies to reach specific orbits is becoming more and more important. This work leverages manifolds in the Three-Body Problem to retrieve an initial solution for optimizing transfer trajectories to the Gateway. Quasi-periodic tori are generated around initial and target orbits, and analytical curves are derived by interpolation. The intersections between these curves representing the tori correspond to single-impulse transfer opportunities. An optimization algorithm refines the maneuvers, reducing fuel consumption and launch date constraints. This approach enhances mission flexibility and cost-effectiveness for lunar exploration.

INTRODUCTION

With the development of a human outpost orbiting the Moon, the ability to travel efficiently back and forth from our satellite is the key to the sustainability of the project. By this means, lately, new strategies to exploit the n-body dynamics have been explored. Henry and Scheeres[1] introduce a whisker-intersection approach using the stable and unstable manifolds of quasi-periodic tori to generate "natural" transfers between QPOs (Quasi periodic Orbits). Jain and Howell[2] present two complementary frameworks: locally fuel-optimal impulsive maneuvers and two-maneuver interior-type or three-maneuver exterior-type transfers that link Near-Rectilinear Halo Orbits (NRHO) and Distant Retrograde Orbits (DRO) via QPO manifolds. Duering[3] exploits the hyperbolic invariant manifolds of QPOs and periodic orbits to perform transfers among tori in the Earth-Moon system. McCarthy and Howell[4] leverage QPO families and their associated manifolds for cislunar trajectory design and eclipse avoidance. Finally, Jain[5] outlines a systematic methodology that uses the five-dimensional manifolds of QPO families to construct interior and exterior transfers between periodic orbits.

Although these strategies can yield low-cost transfers, they require an optimization process that is highly sensitive to the initial conditions, both in terms of solution quality and computational cost. In fact, all these strategies use disconnected orbits as initial guesses that do not represent actual

[*]M.Sc. Student, Space Engineering, Politecnico di Milano, Italy, (matteo.santacesaria@mail.polimi.it).
[†]M.Sc. Student, School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN, United States, (nsailor@purdue.edu).
[‡]Assistant Professor, School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN, United States, (acapa@purdue.edu).
[§]Assistant Professor, Aerospace Engineering, Politecnico di Milano, Italy, (andrea.devittori@polimi.it).

trajectories. This makes the optimization algorithm converge slower and possibly not to the optimal solution. This work proposes a method to identify feasible transfer orbits and use them as informed initial guesses to guide the optimization process. By narrowing the search space and focusing on promising regions, the approach promises to accelerate convergence and avoid providing solutions with prohibitively high fuel consumption during the design of the finite-burn maneuver. The quasi-periodic tori are generated using the Gomez, Mondelo, Olikara, and Scheeres (GMOS) scheme, and a set of orbits for each torus is propagated to retrieve a grid on which the torus itself can be interpolated via bi-cubic splines. Once analytical curves are recovered for the full state vector, the interpolation points between these two curves are retrieved. In the end, the corresponding impulsive Δv to maneuver to the target torus is computed. These solutions are used as initial conditions for the finite-burn optimization algorithm, searching in the surroundings of the minimum Δv retrieved on the intersection curve. This process, once a starting position and a target object orbiting the Moon are fixed, also allows also for phasing, exploiting quasi-periodic trajectories as parking orbits. Through this approach, a satellite will be able to rendezvous with the desired object with no need for additional active phasing maneuvers.

GRID GENERATION AND TORI INTERPOLATION

In the process described in the introduction, the first step is to identify a periodic orbit in the Moon-Earth Three-Body problem and retrieve a set of quasi-periodic orbits around it. The initial periodic orbit is a 9:2 resonant Near-Rectilinear Halo Orbit (NRHO),[6] from this one another one from the same family has been retrieved, having a period 20% longer. The first one is, as of today, the operative orbit of the Lunar Gateway, while the second one has been selected for the study because is a suitable parking orbit for missions to that station.

Using the GMOS scheme, the initial conditions of a set of quasi-periodic orbits belonging to a family of tori are computed. After selecting a specific torus, these initial states are propagated to generate the corresponding set of orbits, which are then used to construct the interpolating grid for the subsequent analysis. In Table 1 are highlighted the initial conditions for the periodic orbit used to generate the torus, all the values have been non-dimensionalized considering for the Earth-Moon system: $L^* = 384400\ km$, $t^* = 375196\ s$ and $\mu = 0.012150$.

Table 1. Initial conditions for the 9:2 periodic NRHO and the one with 120% the period of the Lunar Gateway.

States	Gateway [nd]	120% Period [nd]
x_0	1.02134	1.043539
y_0	0	$4.878772 \cdot 10^{-3}$
z_0	-0.18162	-0.193594
V_{x0}	0	$5.816054 \cdot 10^{-3}$
V_{y0}	-0.10176	-0.145463
V_{z0}	$9.76561 \cdot 10^{-7}$	$-1.480161 \cdot 10^{-2}$
Period	1.50206	1.802472

After the orbits are propagated under the assumption of Circular Restricted Three-Body Problem (CR3BP), a grid of discretized states is available. However, to achieve an accurate interpolation some post-processing is required: three different techniques to distribute the nodes have been considered, equally spaced in time, space traveled, or velocity variation. In this way, it was possible to analyze the results for a grid with the nodes concentered at the apolune, perilune, or equally spaced. In addition to that, since the quasi-periodic orbits are wrapping around the periodic orbits, a counter-rotation of the nodes was necessary to ensure a regular grid.

Although this process is simple from a logical point of view, its coding implementation requires few steps:

1. A vector of times is created to have the points equally spaced in time, space, or velocity variation.
2. For all the time values retrieved, the initial θ_2 is computed as $\theta_{2i} = \theta_{2\ node} - \omega_2 t_{point}$.
3. An spline interpolation of the initial stroboscopic map is required to be able to start propagating an orbit for any needed value of θ_{2i}, the resulting interpolated map is highlighted in Figure 1 .
4. The initial state is propagated for the desired amount of time and the point is added to the grid.

This ensures that for $t = t_{point}$, $\theta_2 = \theta_{2\ node}$, therefore all the nodes of the same branch have the same θ_2 value. The torus generated and the stroboscopic map interpolated are highlighted in Figure 1.

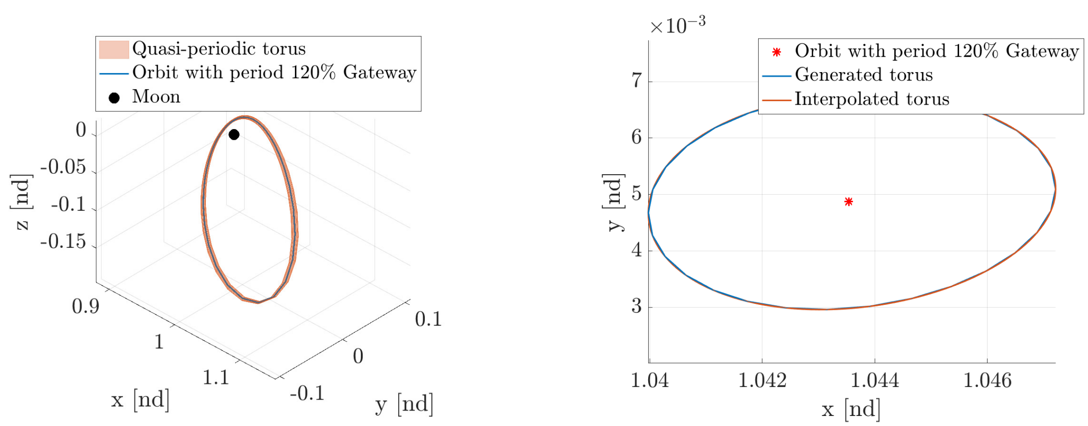

Figure 1. Torus computed around the periodic orbit (left) and interpolated stroboscopic map at $t = t_0$ (right).

Figure 2, 3 and 4 depict the errors resulting from the interpolation with the different grids compared with the propagated torus, the interpolation has been performed using `griddedInterpolant` from Matlab and setting it to use a cubic spline. All grids use 100 nodes for the first frequency (θ_1) and 30 nodes for the second (θ_2), generating a 30 by 100 matrix for each element of the state vector to feed into the interpolation.

The results highlight that the grid based on velocity variation performs slightly better than the one based on the distance traveled, while the one based on time gives less accurate results. The time grid is not capable of representing the torus with good accuracy around the pericenter. The other two grids are working significantly better but the grid based on the distance is less accurate in representing the torus around the pericenter. These results are compliant with what was expected since this grid has a higher concentration of nodes at the pericenter, where the derivative of the state

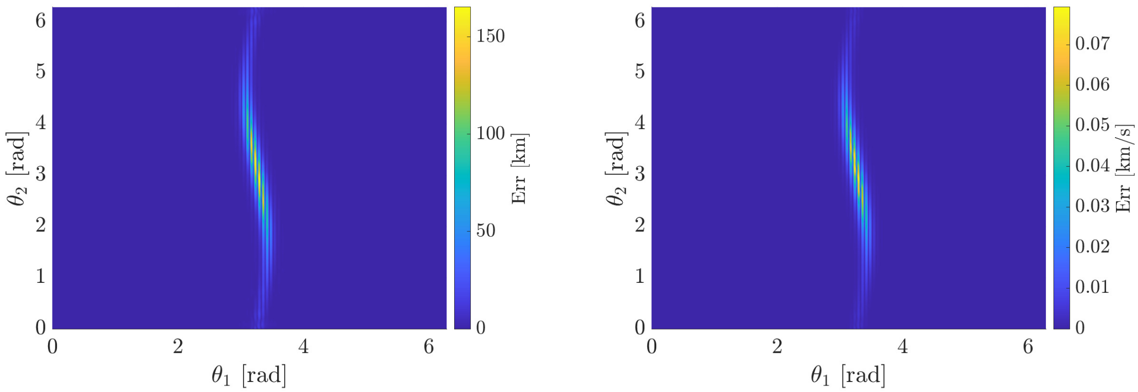

Figure 2. Error in position (left) and velocity (right) interpolating with a grid equally spaced in time, torus with 120% the period of the Gateway.

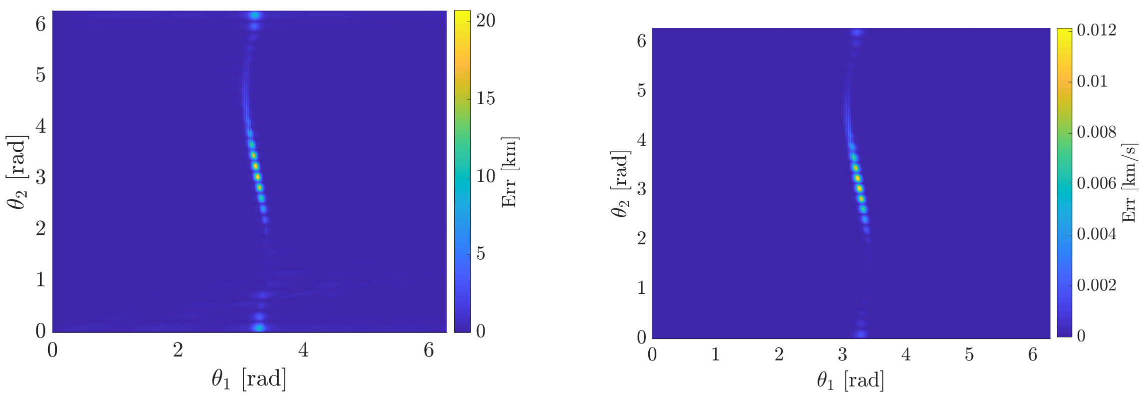

Figure 3. Error in position (left) and velocity (right) interpolating with a grid equally spaced in distance traveled, torus with 120% the period of the Gateway.

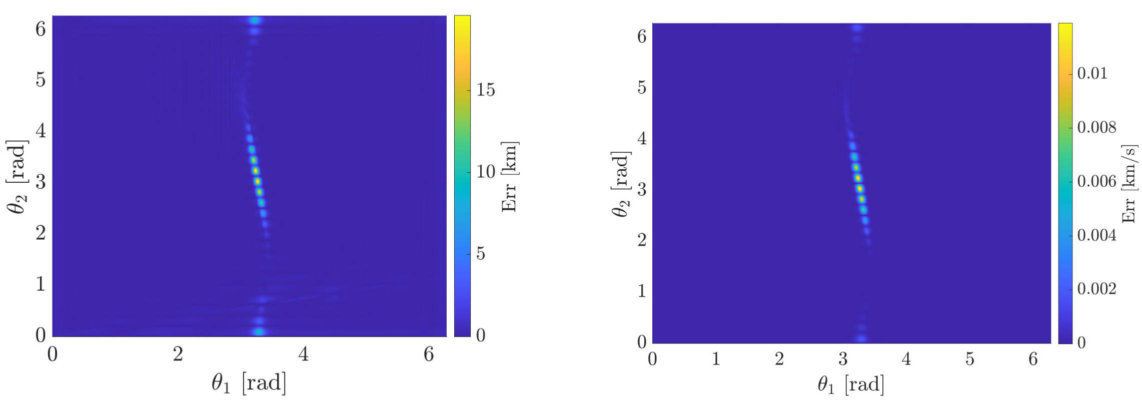

Figure 4. Error in position (left) and velocity (right) interpolating with a grid equally spaced in velocity variation, torus with 120% the period of the Gateway.

vector is larger. For this reason, the grid selected for the interpolation is the velocity-equispaced one. Given the small distance between the two results, it has been decided not to purse grids pushing even more nodes to the perilune, such as building a grid based on acceleration.

Figure 4 shows that the interpolating errors are within acceptable ranges for the torus with $\sim$ 120% the period of the gateway. In fact, given the long distances in place in astrodynamics, and the fact that the interpolation has the only goal of quickly retrieving initial conditions that will be better refined during the design of the maneuver and optimization of the transfer, an error in the order of the km for the distances and of the m/s for the velocities is considered acceptable. The same analysis is carried out for the torus derived from the Gateway's orbit, and after refining the grid by increasing the number of nodes in the stroboscopic map from 30 to 100, acceptable errors are found again. (Figure 5)

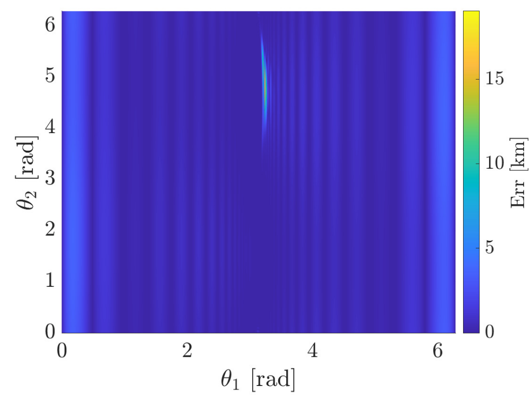

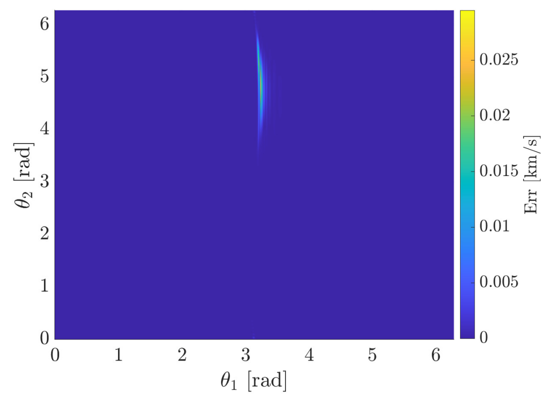

Figure 5. Error in position (left) and velocity (right) interpolating with a grid equally spaced in velocity variation, torus about the orbit of the Gateway.

TORI INTERSECTION POINTS RETRIEVAL

After computing the interpolating function for two tori around two periodic orbits, the intersection points are computed.

The first point is retrieved through an `fsolve` optimization aimed to make the position of the two tori match. The initial conditions for this process are retrieved through a grid search on the points of the grids used to interpolate the tori. Equation 1 shows the match conditions fed into `fsolve`.

$$\begin{cases} r_{x\ torus1} = r_{x\ torus2} \\ r_{y\ torus1} = r_{y\ torus2} \\ r_{z\ torus1} = r_{z\ torus2} \end{cases} \tag{1}$$

The intersection curve between two implicitly defined tori is computed using a predictor-corrector continuation method based on pseudo-arclength parametrization. Continuation along the solution branch is achieved by first determining the tangent direction as the null space of the Jacobian matrix, using the Jacobian computed numerically during the optimization process. A predictor step estimates the subsequent point along the curve as $\mathbf{x}_{k+1}^{(0)} = \mathbf{x}_k + \delta s \cdot \mathbf{v}$, where δs is the pseudo-arclength step size and $\mathbf{v}$ is the tangent vector to the solution curve at the current point. This estimate is then refined through a corrector step by solving an augmented system comprising both the original equations and a pseudo-arclength constraint. The inclusion of this constraint ensures the robustness of the continuation procedure across folds or turning points along the intersection curve.

Figure 6 shows the initial points and resulting intersection curves and the Δv necessary to complete an impulsive maneuver and move to the second torus along with the corresponding angles at which the maneuver is performed on both tori.

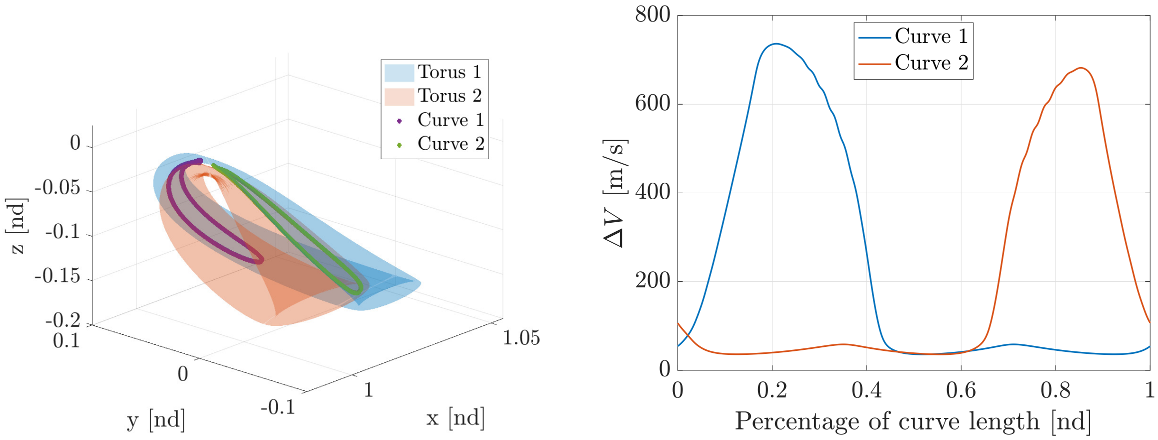

Figure 6. Intersection curve between the tori and ΔV.

From the results, it can be seen that the intersection curve is retrieved as intended, and the ΔV requirement changes depending on the maneuvering position. Looking at the ΔV profile is possible to see a minimum: this value is a well-suited guess for an optimization with the finite-burn model, leading to a more realistic transfer design that withstands the maximum thrust level available on the spacecraft.

As a last step, the transfer opportunities are retrieved setting the initial position of the spacecraft on the larger torus as $\theta_1 = \theta_2 = 0$ at the initial time. Figure 7 shows the results of this analysis.

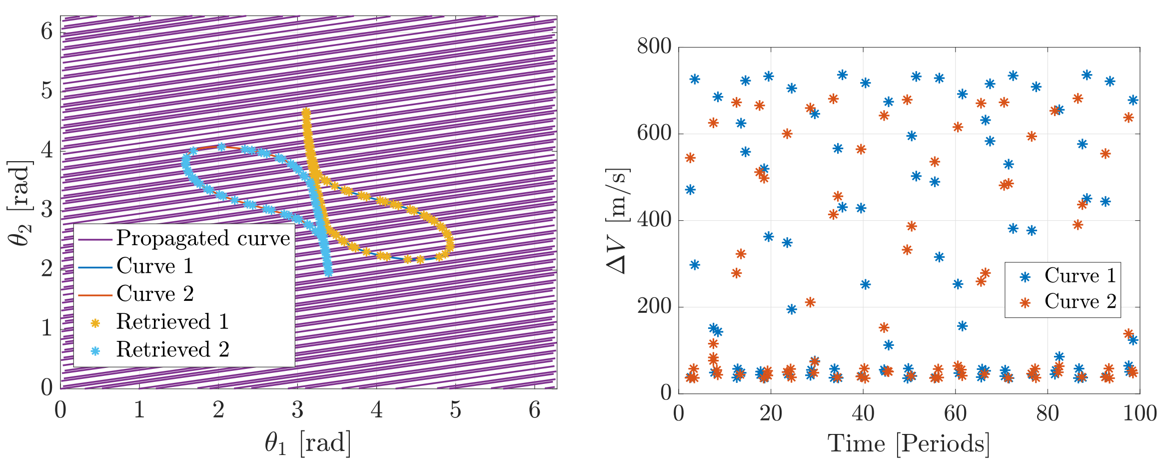

Figure 7. Transfer opportunities over time.

A total of 170 transfer opportunities is retrieved over 100 revolutions on the larger torus, 81 of them being below $100\ m/s$ of required impulse. This mean that, on an average, there is one low-cost transfer opportunity every 9.7 days. This relatively high frequency of encounter with the targeted torus allows great flexibility during the mission design, opening the chance to a larger variety of launch date epochs.

TRANSFER DESIGN

Once all the transfer opportunities are retrieved, the design of the transfers is carried out. The simulation of the complete transfer from a Low Earth Orbit (LEO) to the Gateway is divided into three different steps, obtaining the optimal trajectory for all the legs:

1. Transfer from LEO to the torus with $\sim 120\%$ the period of the Gateway.
2. Transfer from that torus to the one around the orbit of the Gateway.
3. Rendezvous with the station.

This strategy is compared with another one that just targets the orbit of the Gateway. The goal of this analysis is to retrieve the difference in transfer cost and flexibility of the two strategies.

The optimal trajectories are computed using a genetic algorithm (`ga`) for retrieving good initial conditions for all the variables, and a gradient-based algorithm (`fmincon`) to refine the solution and obtain the optimal transfer. This method allows for finding the local minimum within the boundaries provided. In Figure 8 is reported a block diagram of the optimization adopted.

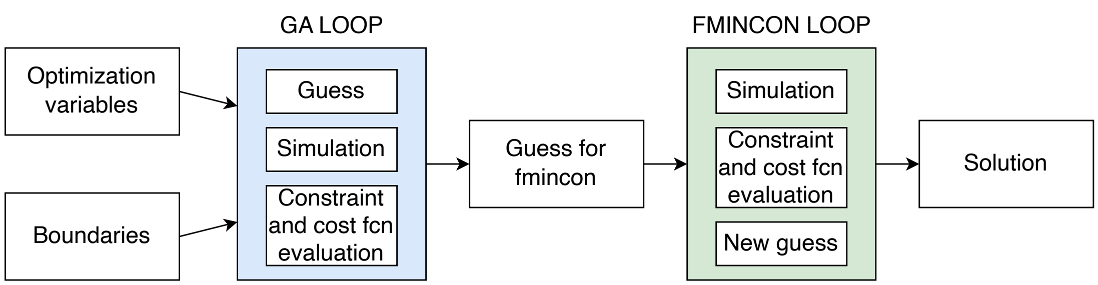

Figure 8. Block diagram of the optimization process.

In terms of workflow, the functions that simulates the trajectory are all structured as in Figure 9, minor variations between the different cases are applied to comply with the specific transfer under investigation. However, all the simulations rely on a two-maneuver strategy to complete the transfers. The simulation takes as input the initial conditions and propagates the trajectory until the starting time of the first burn (t_{1i}), the burn is performed (either impulsive or finite-burn), and the transfer trajectory is propagated until the beginning of the second burn (t_{2i}), also the second burn is performed and the final trajectory is propagated to the apolune for the evaluation of the constraints (if present) and cost function. If the maneuvers are impulsive $t_{1i} = t_{1f}$ and $t_{2i} = t_{2f}$.

The optimization aims at minimizing the total impulse and reads:

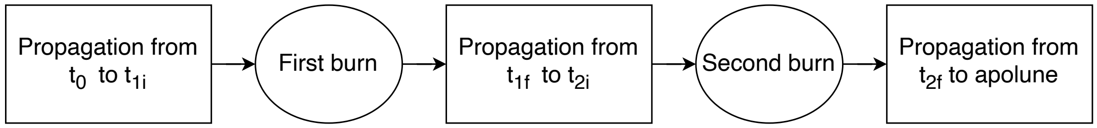

Figure 9. Block diagram of the trajectory simulation process.

$$\begin{aligned}
\min_{\substack{\Delta\mathbf{v}_1,\,\Delta\mathbf{v}_2,\\ t_1,\,t_2}} \quad & J(\Delta\mathbf{v}_1,\, \Delta\mathbf{v}_2) = \|\Delta\mathbf{v}_1\| + \|\Delta\mathbf{v}_2\| \\
\text{S.t.} \quad & \dot{\mathbf{X}}(t) = \mathbf{f}(\mathbf{X}(t), t) \\
& \mathbf{X}(t_0) = \mathbf{X}_0 \\
& t_0 \leq t_1 < t_2 \leq t_f \\
& \mathbf{X}(t_1^+) = \mathbf{X}(t_1^-) + [0,\ 0,\ 0,\ \Delta\mathbf{v}_1]^T \\
& \mathbf{X}(t_2^+) = \mathbf{X}(t_2^-) + [0,\ 0,\ 0,\ \Delta\mathbf{v}_2]^T \\
& \mathbf{X}(t_f) = \mathbf{X}_f
\end{aligned} \tag{2}$$

with $\Delta\mathbf{v_1}$ and $\Delta\mathbf{v_2} \in \mathbb{R}^3$ being the impulses of the maneuvers, $\mathbf{X} \in \mathbb{R}^6$ the state vector, and $\mathbf{f} : \mathbb{R}^6 \to \mathbb{R}^6$ is the function containing the equations of motion of the CR3BP. t_1 and t_2 are the time instants of the first and second burn. The reconstruction of the Δv vectors is done starting from a unit vector in the x direction and performing the rotations around y and z compliant with the σ values retrieved. The mathematical expression is reported in Equation 3.

$$\Delta\mathbf{v}_i = \mathbf{R_z}(\sigma_{iz}) \cdot \mathbf{R_y}(\sigma_{iy}) \cdot [1, 0, 0]^T \Delta v_i \tag{3}$$

with $\mathbf{R_z}$ and $\mathbf{R_y}$ being rotation matrices around the z and y-axis. A constraint is added for the simulation of the trajectory performing the rendezvous with the Gateway: $\theta_{1f}(t) = \theta_{1f}^{GW}(t)$.

Transfer between two tori

Table 2. Boundaries used for the optimization and results.

Variables	Lower Boundaries	Upper Boundaries	Results
$t_{backprop}$ [s]	1	10^5	216
$t_{coasting}$ [s]	1	10^5	571
Δv_1 [m/s]	6	36	21.8
Δv_2 [m/s]	1.8	36	14.5
σ_{1y} [deg]	0	360	35.9
σ_{1z} [deg]	0	360	201.9
σ_{2y} [deg]	0	360	44.9
σ_{2z} [deg]	0	360	217.5

Table 2 shows the optimization variables, the boundaries set and the results obtained for the optimization of the transfer between the tori. $t_{backprop}$ is the how in advance the first burn happens with respect to the intersection point. This specific transfer is also simulated with a finite-burn maneuver retrieving a total mass that differs from the one computed with Tsiolkowsy equation for less than 0.01%, this proves that, in the case simulated, the assumption of impulsive maneuver is accurate.

Table 3. Δv vectors in the rotating frame for the two-impulse transfer between the tori.

$\mathbf{\Delta v_1}$ [m/s]	$\mathbf{\Delta v_2}$ [m/s]
$[-16.4,\ -6.6,\ -12.8]$	$[-8.1,\ -6.2,\ -10.2]$

The Δv vectors are computed through Equation 3 and are shown in Table 3.

Rendezvous with the station

A similar approach is followed to compute the optimal transfer to rendezvous with the station. Both the objects are assumed to be at the apolune at the beginning of the simulation and the maneuvers are considered impulsive. The constraints aim at matching the orbit of the Gateway and also ensure the match of the phase θ_1 at the arrival of the spacecraft on the target orbit. Table 4 shows the data used for the simulation and the results obtained.

Table 4. Boundaries used for the optimization and results.

Variables	Lower Boundaries	Upper Boundaries	Results
t_{prop} [s]	1	$5.6 \cdot 10^5$	$9.8864 \cdot 10^4$
$t_{coasting}$ [s]	1	$5.6 \cdot 10^5$	$2.3815 \cdot 10^5$
Δv_1 [m/s]	1	100	28.4
Δv_2 [m/s]	0.5	100	18.9
σ_{1y} [deg]	0	360	309.1
σ_{1z} [deg]	0	360	81.4
σ_{2y} [deg]	0	360	333.3
σ_{2z} [deg]	0	360	78.7

The Δv vectors are calculated and reported in Table 5.

Table 5. Δv vectors in the rotating frame.

$\mathbf{\Delta v_1}$ [m/s]	$\mathbf{\Delta v_2}$ [m/s]
$[2.7,\ 17.7,\ 22.0]$	$[3.3,\ 16.6,\ 8.5]$

Earth-NRHO transfer

The initial conditions are computed starting from the Keplerian elements of a LEO, this orbit is chosen as it is a typical target orbit for a launch from Cape Canaveral, Florida.[7]

The fitness function receives the spacecraft state in LEO, $\mathbf{X}_{LEO}$, and the state at the target orbit's apolune, $\mathbf{X}_{target}$, then retrieves the impulses necessary to reach the target starting from the initial state. Rather than imposing explicit constraints, each iteration calls `fsolve` to determine the first impulse, $\Delta\mathbf{v}_1$, required to reach the desired apolune position. Once the positional match is achieved, the second impulse, $\Delta\mathbf{v}_2 = \mathbf{v}_{apo} - \mathbf{v}_{target}$, is computed as the difference between the resulting velocity and the target velocity at apolune. In this way, both position and velocity matching are implicitly enforced through the root-finding procedure, even though no constraints are passed directly to the optimizer. Although this approach uses an unconstrained solver, the underlying formulation, seeking two $\Delta\mathbf{v}$ maneuvers that satisfy the full target state, remains unchanged.

Two simulations are performed, one aiming directly for the orbit of the Gateway, and another one aiming for the torus with $\sim 120\%$ period. Table 6 shows the results.

Table 6. Optimization results.

Variables	Gateway	Torus
t_{prop} [s]	2403	2177
t_{coast} [days]	5.2254	5.4406
Δv_1 [m/s]	3179.8	3491.4
Δv_2 [m/s]	861.0	815.5

From the results, it can be noticed that, although the total transfer cost is increasing to get to the torus, the cost of the second maneuver is lower. However, the initial impulse will most likely be given by the launcher. This means that, considering only the maneuvers done by the spacecraft, there is an advantage in targeting the torus instead of the orbit of the gateway for what regards this specific maneuver.

Wrap up

Once all the optimal transfers for the different legs are computed, it is possible to compare the results. The first strategy is traveling directly to the station starting from LEO, the second one is first traveling to the torus originated from the orbit with 120% the period of the Gateway, and later performing the transfers described above to rendezvous with the station.

Table 7. Total Δv and traveling time, direct strategy.

	$\Delta\mathbf{v}$ [m/s]	**Duration** [s]
Total	4040.8	$4.5388 \cdot 10^5$
Spacecraft	861.0	–

Table 7 sums up the results for the first strategy. The spacecraft has to deliver only $861\ m/s$ out of the total because the rest is provided by the launcher to leave the Earth. In this case, the total number of maneuvers is low, but there is a direct correlation between the launch epoch and the arrival on the final orbit. This means that the launch epoch is constrained to the one needed to rendezvous correctly with the Gateway and, if that is not possible, an additional phasing maneuver becomes necessary, increasing the total propellant consumption.

Table 8. Total Δv and traveling time, strategy proposed in this work.

	$\Delta\mathbf{v}$ [m/s]	**Duration** [s]
Total	4382.4	$5.1438 \cdot 10^7$
Spacecraft	891.0	–
$N°1$ (Table 9)	892.0	$\sim 3.2673 \cdot 10^6$

Table 8 sums up the results for the strategy proposed in this work. The first value that stands out is the long transfer time. This is due to the wait time on the first torus before performing the maneuver, being about 577 days. This value is high due to the fact that, while designing the transfer, the minimum-Δv one is selected, regardless from the waiting time. However, looking back at Figure 7 it can be noticed that a large amount of low-cost transfers are available, also during the first few revolutions. In Table 9 are reported the transfer opportunities encountered during the first 10 revolutions and requiring less than $50\ m/s$ for the impulsive transfer. All the transfers shown in the table are close in terms of impulse required with the one previously selected and allow to reach the target, cutting down the transfer time. For example, using transfer number 1 (Table 9) the total transfer time becomes 37.8 days, a value much smaller than the previously retrieved one, and closer to the 5.3 days required by the direct strategy without accounting for the phasing.

Table 9. Low-cost transfer opportunities during the first 10 revolutions.

Transfer $N°$	**Time** [days]	$\mathbf{\Delta v}$ [m/s]
1	19.3664	37.0
2	20.8983	38.5
3	26.6540	36.7
4	60.9250	49.1
5	66.6763	43.9

A more in depth study on the relative motion between the spacecraft on the initial torus and the Gateway on its orbit is performed. The goal is to prove that, after waiting on the first torus for the right amount of time, the spacecraft can rendezvous with the station requiring minimum or no phasing. If this is feasible regardless of the initial position of the Gateway at the arrival of the spacecraft on the torus, allows to decouple the launch date and the encounter epoch.

The analysis is performed considering only the transfers with an impulse lower than $50\ m/s$, and the initial position of the Gateway at the arrival of the spacecraft on the larger torus is computed as: $\theta_{1\ init} = \theta_{1\ transfer} - \omega_{1\ GW} t_{transfer}$. The choice of targeting the θ_1 of the gateway is made to be compliant with the initial condition set in the trajectory optimization code to rendezvous with the station.

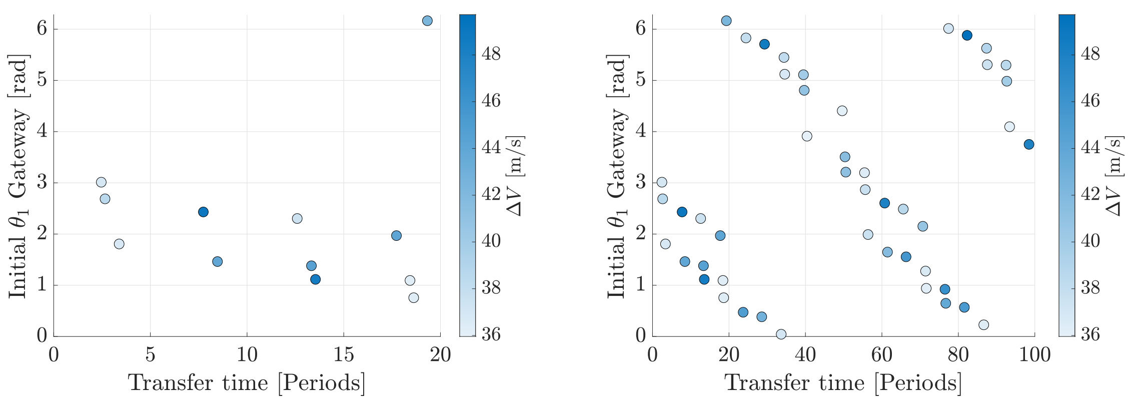

Figure 10. Required initial θ_1 for the Gateway for the different transfer opportunities. On the left a zoom in on the first 20 periods.

Figure 10 shows that every transfer opportunity has a different initial phase of the station required (θ_1) to rendezvous correctly, covering a large variety of values. This shows that, in a large variety of cases, it is possible to complete the transfer with minimum or no phasing. If constraints on the mission duration are present, this method allows to enlarge the launch window by providing different transfer options respecting the constraints. It is however evident that the flexibility grows significantly while looser constraint on the total transfer duration are imposed.

In conclusion, depending on the available launch epochs, position of the Gateway in those epochs and mission constraints given by the nature of the mission itself, the two strategies can be compared, keeping in mind advantages and disadvantages of them for the specific mission under design. It is evident that in manned missions, in which time is critical, the propellant savings are not a driver, and the launch date can be adjusted more freely, having as a priority the safety of the astronauts and the sustainability of the travel for the humans onboard. On the other side, for cargo missions transporting heavy payloads, a small saving in Δv can lead to large savings in propellant consumption, enhancing the cost-effectiveness of the mission.

CONCLUSIONS

This work demonstrates that quasi-periodic torus transfers in the CR3BP:

- Provide reliable initial guesses via low-error interpolation.
- Yield multiple low-Δv intersection opportunities.
- Enable decoupled launch-to-station phasing at a modest Δv penalty.

Future work will extend to continuous-thrust trajectories, higher-fidelity dynamics, and multi-torus chaining for a cislunar transport network.

NOTATION

δs	Pseudo-arclength step size
θ_1	First frequency of the torus
θ_2	Second frequency of the torus
μ	Mass parameter
ω_1	Angular velocity related to the first frequency of the torus
ω_2	Angular velocity related to the second frequency of the torus
L^*	Characteristic length
t^*	Characteristic time
t_0	Initial integration time
t_{point}	Propagation time corresponding to a specific "point" of the grid

REFERENCES

[1] D. B. Henry and D. J. Scheeres, "Quasi-periodic orbit transfer design via whisker intersection sets," *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 10, 2023, pp. 1929–1944.

[2] D. Jain and K. Howell, "Leveraging Quasi-Periodic Orbits to Design Transfers in the Circular Restricted Three-Body Problem," *AIAA SCITECH 2024 Forum*, 2024, p. 0841.

[3] M. Duering, *Station-keeping and orbital transfers in the vicinity of the Moon exploiting quasi-periodic orbit dynamics*. PhD thesis, University of Strathclyde, 2020.

[4] B. P. McCarthy and K. C. Howell, "Leveraging quasi-periodic orbits for trajectory design in cislunar space," *Astrodynamics*, Vol. 5, No. 2, 2021, pp. 139–165.

[5] D. Jain, "Transfer trajectory design strategies informed by quasi-periodic orbits," Master's thesis, Purdue University, 2023.

[6] E. M. Z. Spreen, *Trajectory design and targeting for applications to the exploration program in cislunar space*. PhD thesis, Purdue University, 2021.

[7] E. M. Zimovan, "Characteristics and design strategies for near rectilinear halo orbits within the earth-moon system," Master's thesis, Purdue University, 2017.