AAS 19-874

CLOUD COMPUTING METHODS FOR NEAR RECTILINEAR HALO ORBIT TRAJECTORY DESIGN

Sean M. Phillips*, Diane C. Davis†, and Daniel J. Sweeney‡

Complicated mission design problems require innovative computational solutions. As spacecraft depart from a proposed Gateway in a Near Rectilinear Halo Orbit (NRHO), recontact analysis is required to avoid risk of collision and ensure safe operations. Escape dynamics from NRHOs are governed by multiple gravitational bodies, yielding a trajectory design space that is exhaustively large. This paper summarizes the recontact analysis for departure from the NRHO and describes how the Deep Space Trajectory Explorer (DSTE) trajectory design software incorporates high performance cloud computing to compute and visualize the orbit design space.

INTRODUCTION

Recent focus on exploration missions to cislunar space has kindled accelerated interest in multibody orbit solutions. Trajectory analysis in the presence of multiple gravity fields is complex, and innovative computational tools are needed to simplify complicated design spaces, to generate large quantities of data quickly, and to visualize the output for user accessibility. The Gateway mission is a prime example. The Gateway[1] is proposed as a human outpost in deep space. The current baseline orbit for the Gateway is a Near Rectilinear Halo Orbit (NRHO) near the Moon.[2] The NRHO exists in a regime that experiences the gravitational effects of the Earth and the Moon simultaneously, complicating orbit analysis. The mission design process benefits greatly from updated computational tools for multibody missions like the Gateway.

As an example, consider the problem of assessing the risk of collision in an NRHO. As a staging location to missions to the lunar surface and beyond the Earth-Moon system, the Gateway will experience spacecraft and other objects regularly arriving and departing. Departing objects potentially include spent logistics modules, visiting crew vehicles, debris objects, wastewater particles, and cubesats. Each departure is governed by the dynamics of the Gateway orbit and the surrounding dynamical environment. Over time, any unmaintained object in such an orbit eventually departs due to the small instabilities associated with the NRHOs. A separation maneuver speeds the departure from the NRHO, but the effects of the maneuver on the spacecraft behavior depend on the location, magnitude, and direction of the burn. Escape dynamics from the NRHO with regard to these maneuver options open up an enormous potential trajectory design space where subtle changes in input can produce dramatically large changes in the results. Any departing object must avoid recontacting the Gateway as it leaves the lunar vicinity, and a recontact analysis thus involves a significant number of computations and extensive output data.

To explore the dynamics of this extensive design space, the Deep Space Trajectory Explorer[3] (DSTE) trajectory design software incorporates new High Performance Computing (HPC) services and novel interactive visualizations. This paper details the HPC and cloud infrastructure techniques that are

* Senior Professional Staff II, Johns Hopkins University Applied Physics Lab, 11100 Johns Hopkins Road, Laurel, Maryland 20723 sean.phillips@jhuapl.edu.
† Principal Systems Engineer, a.i. solutions, Inc., 2224 Bay Area Blvd, Houston TX 77058, diane.davis@ai-solutions.com.
‡ Gateway Integrated Spacecraft Performance lead, NASA Johnson Space Center, daniel.j.sweeney@nasa.gov

implemented in the DSTE, applying the new capabilities to analysis of recontact risk with the Gateway in NRHO.

NEAR RECTILINEAR HALO ORBITS

The Gateway is planned to fly in a lunar NRHO as its baseline orbit. The NRHO families of orbits are subsets of the larger halo families, which originate from planar orbits near the L_1 and L_2 libration points; the Earth-Moon L_2 halo family appears in Figure 1. Each halo orbit is perfectly periodic in the Circular Restricted 3-Body Problem (CR3BP) and becomes a quasi-periodic orbit in a higher fidelity ephemeris force model. The NRHOs are defined as those members of the halo family with bounded stability properties;[2] they pass near the Moon at perilune and are nearly polar. Families exist with apolunes located both above the lunar north pole and above the lunar south pole; the Gateway is planned to reside in a southern L_2 NRHO in a 9:2 resonance with the lunar synodic period. The 9:2 NRHO is characterized by a period of about 6.5 days, a perilune radius of about 3,500 km, and an apolune radius of about 71,000 km; it is strongly affected by the gravity of both the Earth and the Moon simultaneously. This NRHO offers extended communications with assets on the south pole of the Moon,[4] as well as low-cost orbit maintenance and attitude control,[5] favorable eclipse avoidance properties,[6] and inexpensive transfers from Earth and to other destinations.[5,7] The NRHO portion of the southern L_2 halo family is highlighted in black in Figure 1, and the 9:2 NRHO appears in blue.

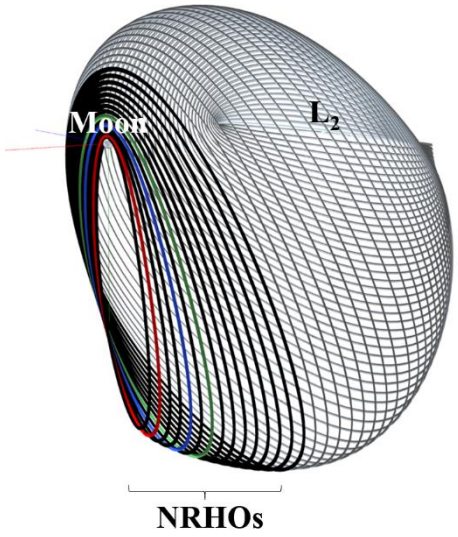

Figure 1. Southern Earth-Moon L_2 Halo family; NRHOs highlighted in black.

DYNAMICAL MODELS AND ASSUMPTIONS

After an object has separated from the Gateway and before it departs the NRHO, the primary gravitational influences are the Earth and Moon, and the CR3BP is an effective approximation for post-separation dynamics.[7] Thus, the current analysis employs the CR3BP to exploit its advantages, such as time independence and symmetries, to assess the risk of recontact for objects departing the Gateway. The CR3BP[8] describes the motion of a massless spacecraft affected by two primary gravitational bodies such as the Earth and the Moon. The model assumes that the two primary bodies are point masses orbiting their center of mass in circular orbits. The spacecraft moves freely under the influence of the two primaries, and its motion is described relative to a rotating reference frame. The CR3BP equations of motion are written

$$\ddot{x} - 2\dot{y} = \frac{\partial U^*}{\partial x}$$
$$\ddot{y} + 2\dot{x} = \frac{\partial U^*}{\partial y} \quad (1)$$
$$\ddot{z} = \frac{\partial U^*}{\partial z}$$

where the pseudopotential U^* is defined as

$$U^* = \frac{1}{2}(x^2 + y^2) + \frac{\mu}{r} + \frac{(1-\mu)}{d} \quad (2)$$

The coordinates *x*, *y*, and *z* are components of spacecraft position in the barycentered rotating frame. The values *d* and *r* represent the distances between the spacecraft and P_1 and P_2 respectively, and $\mu = \frac{m_2}{m_1+m_2}$ is the mass parameter of the system where m_1 and m_2 are the masses of the two primaries. All quantities are nondimensional, where the characteristic length, *l**, is defined as the constant distance between P_1 and P_2, the characteristic mass, *m**, is the combined mass of the two primaries, and characteristic time is $t^* = \left(\frac{l^{*3}}{Gm^*}\right)^{1/2}$, where *G* is the gravitational constant. No closed-form solution exists to the CR3BP equations of motion, but five equilibrium solutions, the libration points, are denoted L_1 through L_5. Stable and unstable periodic orbit families, including the L_2 NRHOs, emerge in the vicinity of the libration points.

When a separation maneuver is applied to an object in the NRHO, maneuver location is parameterized by true anomaly, TA. Maneuver direction is defined in the Velocity-Normal-Binormal (VNB) frame. The VNB directions are defined along the orbit such that the V direction aligns with the rotating velocity, V = $(\dot{x}, \dot{y}, \dot{z})$, the N direction corresponds to the orbit normal, N = $(x, y, z) \times (\dot{x}, \dot{y}, \dot{z})$, and B completes the right handed system. The yaw angle ranges from -180° to 180° in the V-N plane, and the pitch angle ranges from -90° to 90° in the V-B plane.

After the separation maneuver is applied, the object diverges from the Gateway, and it may risk recontacting the Gateway prior to NRHO departure. Recontact risk is defined to occur when an object, after reaching a range from the gateway greater than a certain threshold, returns within the same threshold prior to NRHO departure. In the current study, the threshold is set to 100 km.

To define departure from the NRHO, a momentum integral is employed.[9] The momentum integral, MI, is a line integral of the position vector from the initial time, t_0, to the current time, *t*,

$$\mathrm{MI}(t) = \int_{t_0}^{t} x(\tau)\dot{x}(\tau) + y(\tau)\dot{y}(\tau) + z(\tau)\dot{z}(\tau)d\tau \quad (3)$$

where *x*, *y*, and *z* are components of the position vector relative to the barycenter in the Earth-Moon rotating frame and $\dot{x}, \dot{y},$ and $\dot{z}$ are components of the velocity vector in the same frame. For a perfectly periodic halo orbit in the CR3BP, the MI is also periodic and returns to zero after each period. Over time, as the orbit of a perturbed or unmaintained spacecraft diverges from the NRHO, the MI also diverges, and departure is defined in terms of the divergence of the MI. When the magnitude of the MI crosses a particular threshold, the debris object is considered 'departed' from the NRHO. The specific value of MI threshold depends on the application; in the current study a value of 0.1 is selected.

THE DEEP SPACE TRAJECTORY EXPLORER

When working in a regime in which multiple gravitational bodies must be simultaneously considered, sophisticated design tools offer a significant advantage. The DSTE is JavaFX-based program designed specifically for preliminary multibody trajectory design analysis. The purpose of the DSTE is to provide a visual, interactive workflow that allows the user to identify, compare, and evaluate suitability of trajectories for a given application. The major design goal is to provide tools that allow a user to rapidly traverse broad data spaces; in essence, to help the user find needles in haystacks. To facilitate this design goal, the focus is on two primary implementation factors. The first is performance, notably speed of trajectory integration, map creation, and visualization. The second is user experience, that is, facilitating the creation and manipulation of maps and the selection and visualization of trajectories. Software development for the DSTE is implemented on the latest version of the Java Development Kit (JDK, as of this writing 64-bit 8u161). Using JDK 8 provides two very important implementation benefits: Parallel Streams and the JavaFX GUI toolkit. Parallel Streams are heavily used throughout the application as a means of performing dynamic calculations across extremely large data sets. By using Parallel Streams to perform operations upon a collection of objects

in memory, the DSTE automatically parallelizes key bottleneck computations such that there is a linear scaling with available cores. Primary examples are periapsis Poincaré map generation and periapsis point searches. The latest JavaFX GUI toolkit is used for generation of all visualizations and interfaces within the DSTE. Of particular note is the high performance JavaFX Canvas component that provides a low latency rendering tool for the Poincaré map data. Traditional event-managed plotting toolkits are quickly overwhelmed both in memory and CPU usage given the range and density of this data space. The direct memory rendering mechanism of the JavaFX canvas easily meets the performance needs to handle the data density necessary to achieve the design goals stated above. In the current analysis, cloud capability is added to the DSTE to further improve performance for applications that require the generation of large datasets. The cloud infrastructure is applied to the analysis of recontact risk for objects departing the Gateway spacecraft.

RECONTACT ANALYSIS

Any object or spacecraft separating from the Gateway and departing the NRHO for another destination in space must depart safely, without a risk of colliding with the Gateway. During maneuver planning, assessment must be performed to determine the maneuver parameters that ensure safe departure. The Gateway's recontact analysis problem is unique given the sensitivity of the cislunar orbit regime. The multibody dynamical environment in the NRHO results in a vast dataspace, with outcomes highly dependent on initial conditions. Without previously modeled recontact analysis in this regime, the map generation in DSTE represents the first step towards intuition development for behavior of objects departing the Gateway.

Problem Setup and Constraints

The 9:2 southern L_2 NRHO is baselined as the Gateway reference trajectory. The periodic orbit appears in a CR3BP propagation in the Earth-Moon rotating frame in Figure 2. The 9:2 NRHO is nearly stable in a linear analysis. In a CR3BP propagation, an object in the NRHO remains in orbit for as long as 35 revolutions if no maneuver is applied. In a higher fidelity ephemeris force model, an object in the corresponding NRHO departs within 5-20 revolutions, depending on the perturbations acting on the object. Of course, a maneuver accelerates the departure from the NRHO.[10] For a burn in a given direction, a larger maneuver magnitude leads to faster departure from the NRHO. The burn direction also has a large effect on the time to departure, as does the location of the maneuver, with burns in the velocity and anti-velocity direction at perilune leading to the fastest departures from the NRHO. For an object that separates from the Gateway, the maneuver size, location, and direction determine its post-maneuver behavior. While it remains in the vicinity of the Moon, the object risks recontacting the Gateway structure. Thus, it is necessary to understand the magnitude, direction, and location of separation maneuvers that do and do not lead to safe departure from the NRHO. For the sake of this analysis, these aspects of the problem are designed and constrained as follows.

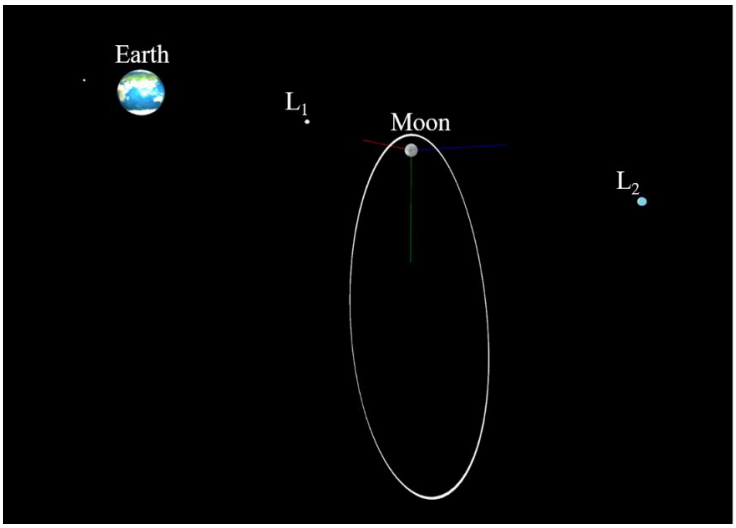

Figure 2. Gateway NRHO in the Earth-Moon rotating frame

Magnitude

Relatively small maneuver magnitudes, up to 20 m/s, are considered for this phase of the recontact analysis. This range represents departure velocities expected for wastewater venting and cubesat deployment (at the lower end) and logistics module or large object disposal (at the higher end). Each burn is considered as a single impulsive maneuver. In the DSTE, the Δv magnitude range and step size are user inputs, as shown in Figure 3.

Figure 3. DSTE user controls for Parameter Space setup

Direction

A sphere of possible burn directions is considered by using a combination of a yaw and pitch parameters, defined relative to the spacecraft velocity vector at any point. For the sake of managing the Monte Carlo space, the combination of yaw and pitch is used to identify each possible maneuver, with each yaw-pitch pair considered as a "bin". Each bin is assigned a maneuver magnitude. The yaw, pitch and maneuver Δv parameters form a triplet that is easily reused at any starting location along an origin trajectory for the basis of a separation maneuver. In the current analysis, the full possible range of yaw (-180° to 180°) and pitch (-90° to 90°) are considered at 1° steps; in the DSTE, the ranges and step sizes are user inputs, as in Figure 3.

Location of the separation maneuver

Departures at any location around the NRHO are available for analysis. To discretize the location state of an origin trajectory, an osculating true anomaly (TA) calculation is used, although the NRHO is not a Keplerian orbit. A depiction of osculating TA around the NRHO in the CR3BP appears in Figure 4 in the Earth-Moon rotating frame, discretized for the visualization in 20° steps. A value of TA = 0° represents perilune and TA = 180° occurs at apolune. In the current analysis, a 1° step size in TA is implemented; this parameterization yields 360 initial locations that, when combined with the maneuver bins, yields a large potential design space, but one that is logically intuitive to manage. Managing groups of yaw, pitch and Δv bins from one separation maneuver location to the next has the added benefit of being conducive to parallel computations.

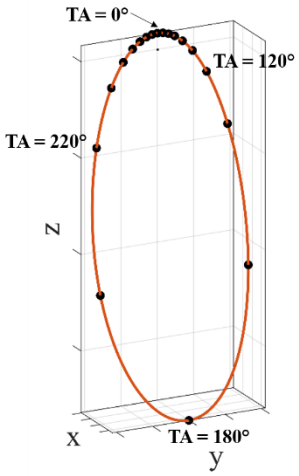

Figure 4. TA marked every 20° along the NRHO

Range checks

Once the separation maneuver is executed, the departing object moves away from the Gateway along a path determined by the maneuver parameters. Some departing trajectories move directly away from the Gateway without reapproaching; however, other maneuvers result in departing paths that repeatedly reencounter the Gateway. For example, consider an object jettisoned from the Gateway with a 1 m/s separation burn in the binormal direction (yaw = 0°, pitch = 90°). The range from the object to the Gateway over time changes depending on the separation maneuver location. A plot of range histories for separation TA values spanning 0° - 360° appears in Figure 5. The colors in the plot correspond to departure TA, as noted in the legend. For example, a separation maneuver at time t = 0 corresponds to a burn at apolune, TA = 180°, and the resulting range over time is represented by a blue curve. Separation at time t = 3.25 days corresponds to a burn at perilune, TA = 0, and the subsequent range to the Gateway over time appears as the first red curve. The final green curve is

associated with a separation maneuver at t = 6.5 days, equivalent to the next apolune. In this way, separation burns around the NRHO are sampled. It is immediately apparent in Figure 5 that some of the separation maneuvers return to within 100 km of the Gateway after their initial increase in range. Between t = 12 days and t = 20 days, close approaches with range under 100 km occur for separation from many regions around the NRHO. Also notable is the fact that some separation maneuvers occurring after apolune, represented by blue curves, result in repeated close approaches with the Gateway for many weeks after the separation. More analysis appears in a previous study.[10]

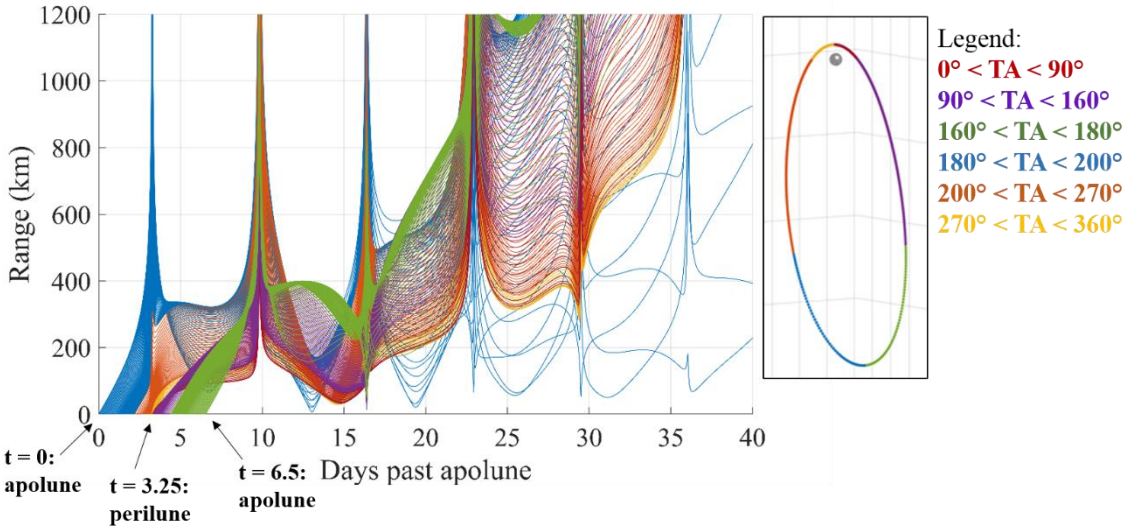

Figure 5. Range to the Gateway over time for objects departing with a 1 m/s separation burn in the binormal direction around the NRHO.

The range history plot in Figure 5 identifies trajectories that risk recontacting the Gateway after separation maneuvers of a single burn magnitude (1 m/s) in a single direction (binormal; yaw = 0°, pitch = 90°). To assess the recontact risk over the full span of magnitude and direction parameters, recontact maps are generated to condense the information contained in the range plots into an easily comprehensible format. At each integration step along each trajectory, the range between the Gateway and the departing object is computed to assess the trajectory for recontact risk. On average, each departing trajectory contains about 600 states that must be compared to the Gateway state before the object either departs the vicinity of the NHRO or impacts the Moon. This large quantity of range checks for each trajectory in the study expands the computational demands significantly.

Taking into account the range and discretization in maneuver location, magnitude, yaw, and pitch, the ranges of the Monte Carlo parameter space, including the total number of maneuvers, are listed in Table 1. All maneuvers are calculated as part of a CR3BP trajectory integration.

Table 1. Parameter Space for Monte Carlo Simulations

Magnitude	20 m/s Δv discretized at 0.5 m/s	40 impulsive magnitudes
Direction (Yaw, Pitch)	360° in yaw ∗180° in pitch	64,800 impulsive bins
Location of maneuver	True Anomaly	360 origin trajectory states
Total CR3BP Trajectories	64,800 bins * 40 magnitudes * 360 states	933,120,000 trajectories
States per departure trajectory	average for departure	600 states per trajectory
Total Close Approach Checks	**Departure Trajectories * Departure States**	**559,872,000,000 checks**

Stopping conditions

After each separation burn, the departing object is propagated forward in time until one of four stopping conditions is met. Two conditions are user-configurable in the DSTE, as in Figure 6. The first is the recontact threshold, or navigation accuracy. This parameter, selected as 100 km in the current analysis, represents the range below which an object is considered a recontact risk. If this range threshold is violated, the maneuver is marked as a recontact and the integration is terminated. The second user-configurable stopping condition is the momentum integral threshold. In the current study, a threshold of 0.1 is selected to represent the point when a separated object has safely departed the vicinity of the NRHO and the risk of recontact is past. If the MI threshold is crossed, the maneuver is marked as a safe departure and the integration is terminated. Finally, if the departing object impacts either the Moon or the Earth, the maneuver is marked as an impact and the propagation is terminated.

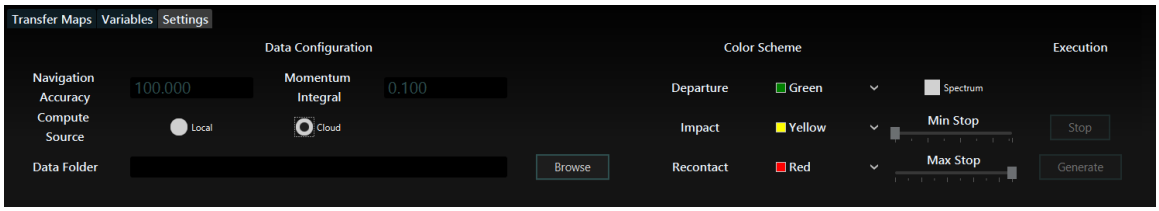

Figure 6. Stopping conditions and map preference user controls in the DSTE.

Each computation represented in Table 1, along with the four stopping conditions, are used to generate a set of recontact maps, which display the results of the separation burns as defined by the maneuver location, Δv, and yaw-pitch direction.

Recontact Maps

The parameters detailed in Table 1 represent an exhaustive design space, and in this analysis, the DSTE is employed to condense this design space into a set of recontact maps.[7] Each map represents departure from a single location along the NRHO at a single maneuver magnitude for the full range of maneuver directions. A sample map representing deployment from the 9:2 NRHO at a true anomaly TA = 318° with a 5 m/s separation burn appears in Figure 7 on the left. The center of the map represents a maneuver in the rotating velocity direction. The horizontal axis represents maneuver yaw with respect to the velocity vector and ranges from -180° to 180°. The vertical axis represents maneuver pitch with respect to the velocity vector and spans 90° to -90°. Note that in the current investigation, the bottom of the map represents a pitch of 90° while the top of the map corresponds to pitch = -90°. The VNB directions are marked on the map as white points. Green regions in the map represent separation burns that lead to departures without risk of Gateway recontact. Red regions correspond to separation burns that result in recontact risk (in this case, defined by range $\leq$ 100 km) prior to NRHO departure. Yellow regions identify separation maneuvers that lead to lunar impact.

Sample departing trajectories appear in Figure 7 on the right. The green trajectories depart the NRHO without recontacting the Gateway, with time to depart ranging from approximately 16 to 60 days depending on maneuver direction. The yellow trajectories in Figure 7 on the right impact the Moon. Impact trajectories arising from the yellow lobe centered at the anti-velocity direction impact the Moon after a flight time of about 24 days after several revolutions below the lunar south pole. One sample impact trajectory in Figure 7 completes two revolutions over the lunar north pole prior to impact. This trajectory arises from the maneuver indicated by one of the individual yellow points marked by a yellow arrow in Figure 7 on the left. It impacts the moon after a 71-day flight. The red orbits originating from the small, oval red areas in the map (marked by a red arrow) risk recontacting the Gateway at the next perilune passage, about 1.5 hours after separation. The red orbits originating from the sine wave-shaped red pattern in the map, marked by a black arrow, risk recontacting the Gateway after approximately 34.5 days, or about five revolutions after separation.

Recontact maps provide a visual method to quickly assess the risk of recontact or lunar impact for an object separating from the Gateway at a given location and with a particular Δv magnitude in any direction. To fully explore the design space, recontact maps are generated for many locations along both the NRHO for a range of values of Δv magnitude.

Figure 7. Recontact map and departing trajectories from the 9:2 NRHO.

DSTE Recontact Integration

The integration process itself is logically broken into three distinct aspects: the initial conditions, the numerical integrator, and the dynamic stopping condition. Each distinct case in the Monte Carlo simulation as defined in Table 1 has a unique initial condition. This initial condition is defined as a maneuver formulated from its yaw, pitch and Δv bin, added to the initial state vector defined by the location of the maneuver. The location of the maneuver originates from the origin trajectory, i.e. the orbit of the spacecraft from which separation occurs and against which each future state is checked for recontact. The DSTE maintains programmatic data structures for representing initial conditions so that they may be integrated in parallel. This parallelization is described in more detail in the subsequent "DSTE Cloud Service" section.

The numerical integrator employs a Runga-Kutta 7-8 order adaptive step size integrator. All Monte Carlo cases use the same integrator regardless of initial condition. For this analysis, the adaptive step size accuracy is selected to be $1x10^{-12}$ nondimensional units, balancing accuracy in the Earth-Moon system with computational performance. The DSTE programmatically manages each independent integration via an integrator object. This object accepts as parameters the initial condition object, described previously, and a dynamic stopping condition, which tells the integrator when to stop integration.

As discussed, the integration process is designed to stop for several logical events of interest, each of which is considered a possible outcome of an object separating from the origin trajectory:

1. Impact with smaller primary (Moon)
2. Impact with larger primary (Earth)
3. NRHO departure (successful departure)
4. Close Approach To Origin Trajectory (recontact)

At each integration step, the departing object's position is checked against each of these events. First, if the position at any point lies within the physical radius of either body, then the fate of the separation trajectory is marked as an impact, and the integration is terminated. Pseudo-code Listing 1 details the computational process for the recontact integration event detection. The logic for impact detection is found on lines 13 and 14 of the listing.

Second, NRHO departure is defined when the momentum integral has passed its specified threshold, a user input in the DSTE. The momentum integral of the departing object is accumulated at each time step and

is used to determine escape from the local system, since a simple crossing in the physical space is insufficient to determine the fate of a departing object in this regime. The momentum integral is accumulated by a simple trapezoid integration. The logic for tracking the momentum integral appears in lines 4 through 10 of Pseudo-code Listing 1. If this condition is met, the fate of the separation trajectory is marked as a successful departure. Dynamic stopping conditions, when framed within the DSTE software, not only provide logical conditions to stop integration, but also track and maintain information between each time step.

1.	For each separation trajectory
2.	Update previous radius
3.	Calculate current radius
4.	Update momentum for departure check
5.	mom = x position * x velocity + y position * y velocity +
	z position * z velocity
6.	Calculate current momentum integral
7.	//Use one trapezoid from the previous step approach
8.	Momentum Integral = Momentum Integral + ((Current Time –
	Previous Time)) / 2 * (Previous Momentum + Current Momentum)
9.	Update previous momentum
10.	Update previous Range difference for recontact check
11.	Calculate new range difference between separation state and the origin state
12.	//Perform Event Checking
13.	If separation state inside radius of either primary or secondary body
14.	Then Impact detected. Stop Integrating.
15.	Else If enough time elapsed to begin performing recontact event checking
16.	If we have gone past the preset number of skip steps
17.	If our Momentum Integral is high enough
18.	Then we have "departed". Stop Integrating
19.	Else If any position vector components fails a distance check
20.	//We have potential impact event.
21.	//Are we moving closer or farther away?
22.	Then If we are also moving closer
23.	Then we have recontact. Stop Integrating.
24.	Event was not detected... keep integrating

Pseudo-code Listing 1. Dynamic Recontact Condition Logic

The final and most critical fate is detected when a separation trajectory makes a close approach to the origin spacecraft. A close approach for this analysis is logically defined as any state vector during integration whose distance between itself and the equivalent origin spacecraft state is less than a specified safety radius and decreasing; the safety radius is user defined in the DSTE. This logic ensures the trajectory is integrated after the initial separation beyond the point when it first passes outside the specified safety radius. Both the departure check and the recontact checks are skipped for a parameterized preset number of integration steps, in this case 6, to minimize absurd logical corner cases that are encountered from the initial separation maneuver. Validating the separation itself via proximity operations is beyond the scope of this analysis. The logic for tracking a close approach appears on lines 15 through 23 of Pseudo-code Listing 1. If the recontact condition is met, the fate of the separation trajectory is marked as a recontact and the integration is terminated. For each separation trajectory, the fate information is collected within the resulting recontact map.

HIGH PERFORMANCE COMPUTING SETUP

The fate of each trajectory is captured and exported as both a data file and a visual recontact map. The DSTE automatically separates map generation using the origin state as the logical demarcation. Logical separation in this manner facilitates broad simulations that can span any possible range of departures from a given origin trajectory. To accelerate the computations, the DSTE automatically executes recontact

integration for each separation transfer combination in parallel using the Java Parallel Streams API and all available cores. For input ranges that are exhaustively large, the DSTE Cloud Service provides a highly scalable recontact integration capability that can be optionally connected to via the user interface. As the computations are completed, the results and fate of each separation transfer are asynchronously transferred to and visualized as a Recontact Map, as in Figure 2.

DSTE Cloud Service

To facilitate highly parallel and deep Monte Carlo scenarios, the DSTE Recontact Integrator and associated APIs are wrapped as Representational State Transfer (REST) end points and re-packaged as a lightweight RESTful Web Service (RWS). The DSTE RWS is deployed to a high performance cloud instance served off the Oracle Cloud Infrastructure (OCI) platform.[11] The shape of the OCI instance is a Bare Metal (BM) compute instance with a total 72 total cores and 512 GB of RAM. The BM shape used for this analysis features 2.3 GHz Intel Xeon E5-2669, at this time the fastest available core speeds for a cloud instance. Selecting OCI BM instances is driven by cost for performance. Further, each BM instance is a dedicated barebones Ubuntu image running directly on a single hardware backplane, ie, no virtual machine (VM) or container. This allows all of the trajectory integrations, which are the computational bottlenecks, to execute and stay resident in RAM within a single instance. The DSTE RWS is served using the Payara Micro application server allowing for fast configuration-less deployment that is flexible to scale horizontally as needed.[12] The software stack deployed to the cloud instance is outlined in Figure 8.

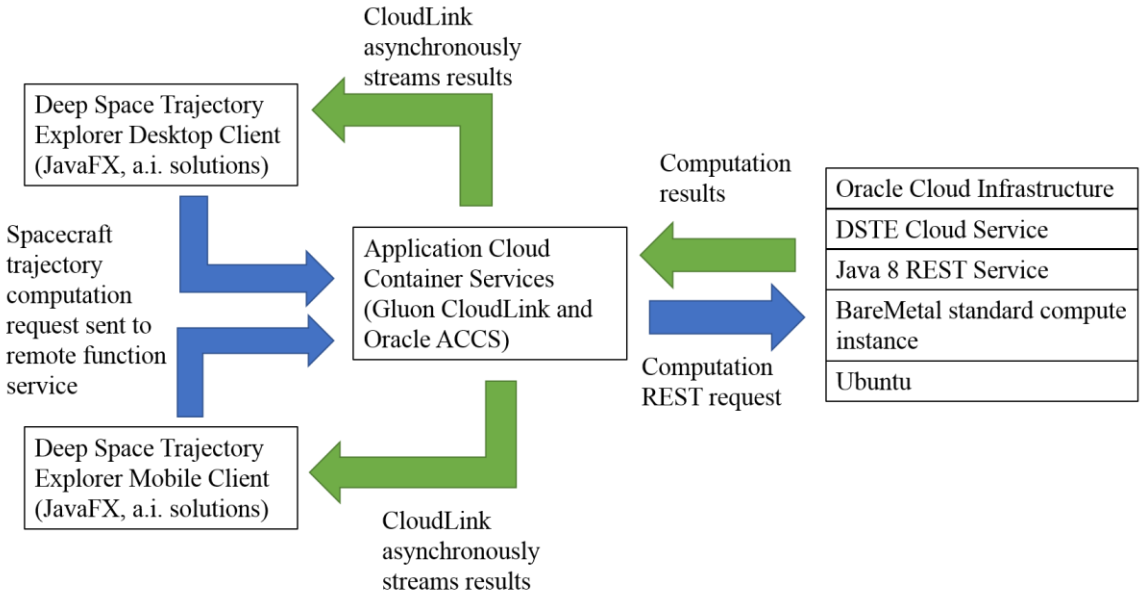

Figure 8. DSTE Cloud Service Architecture used for Recontact Analysis

Computation within the DSTE RWS must be flexible as it is intended to process requests that require more time than is typical for a traditional REST service. The service itself is designed to keep a client connection request alive and stream the results asynchronously back to the client as each integration completes via a JavaScript Object Notation (JSON) exchange. The exchange is facilitated by leveraging Gluon's CloudLink libraries and service. Gluon CloudLink is powerful service API that makes client and cloud deployment simple by providing a "bi-directional, secure, and controlled communication between your mobile apps and your enterprise apps".[13] The Gluon CloudLink service itself is by default provided by Gluon and for this analysis is hosted on Oracle's Application Cloud Container Services. Importantly, the CloudLink libraries map the REST endpoint connection as an asynchronous function which is received by the client

through an observable pattern, described below in the Parallelization Logic section. This is important as unlike typical REST requests, the Recontact Monte Carlo computation takes an extended, and unpredictable, amount of time to complete.

Parallelization Logic

The recontact analysis is logically broken apart at each individual separation maneuver, as the integration for each of these is naturally embarrassingly parallel. The Java Parallel Streams API (at the time of the analysis Java 8u161), is used to automatically allocate each separation integration to a separate worker thread as it becomes available. This process takes the form of a single ArrayList collection containing each initial condition object, each defined and instantiated from the yaw, pitch, Δv and separation location bins described previously. This greatly simplifies the code footprint as any parallelStream() execution automatically creates a thread for each available core (in this case 72). While the thread pool is easily configured for more or less than the default, it was found to be unnecessary largely in part to the selection of the BM instance shape.

Up until this point, standard Java APIs are used. To facilitate the asynchronous connection a combination of calls from Gluon CloudLink and the GlassFish ChunkedOutput API are used. Within the parallelStream(), as each integration is completed, the results are pushed onto a ConcurrentlyLinkedQueue collection. A thread periodically checks the percent complete and polls the queue members at one percent intervals. At each interval the queue is emptied with the completed results written to the ChunkedOutput collection. This automatically sends those results (converted to JSON) back to the client via the REST connection itself. Chunking the output keeps the REST connection alive, as typically connections timeout after 60 seconds. Most importantly, the client receives the results *as they complete*, allowing the DSTE client to process and render them for immediate examination by an analyst. While complete analysis cannot be achieved without the entire result set being returned, the streaming results provide a powerful and immediate look into what a given Monte Carlo parameter set will yield. The analysis of the results is described below in the Recontact Map Analysis section.

Using this design and connection strategy yields a dramatically improved time of computation as compared to running locally on a four core machine. Using a full set of yaw, pitch and Δv bins for a single separation location as a comparison metric, an average time of 23 seconds is achieved, as opposed to 4 minutes on a local four core machine. The 23 seconds includes the time of transmission of the results back to the client. This time will vary, of course, based on the network latency incurred between the client and the cloud instance. The total byte count of JSON averages less than 14 MB, making the total transfer more than acceptable for typical broadband connections. The client reception of the JSON stream is enabled and simplified by using the GluonObservableList API. A GluonObservableList serves as a specialized collection, wrapping a standard Java collection such that any remote REST connection mapped through the Gluon Function Portal will be asynchronously updated. This connection is very convenient as the client code, in this case the JavaFX based DSTE, merely implements a ListChangeListener to be notified and process the results.

RECONTACT MAP ANALYSIS

The output of the recontact integrations are compiled into recontact maps. Recontact maps provide a visual method to quickly assess the risk of recontact or lunar impact for an object separating from the Gateway at a given location and with a particular Δv magnitude in any direction. Complementary analysis then investigates the destination of a departed object: the impact location on the lunar surface, for example, or whether the object escapes to heliocentric space or remains in orbit near the Earth.[14] For the 9:2 NRHO, a set maps corresponding to a Δv magnitude of 1 m/s for 16 separation burn locations around the NRHO appears in Figure 9. Recall that red points on the map correspond to maneuver directions (yaw and pitch) that produce a trajectory that violates the recontact threshold, here set to 100 km. A yellow point designates a burn that leads to lunar impact, and a green point represents a separation burn that results in a safe departure from NRHO, with departure defined by a MI threshold set to 0.1. A Δv magnitude of 1 m/s leads to very limited opportunities for (or risk of) lunar impact. For departure locations between apolune and perilune ($TA \geq 180°$) and for $0 \leq TA \leq 140°$, a few yellow points appear in each map. While the points are few, they form distinct patterns. The time of flight between separation and impact generally falls into one of three groups; about 60 days, about 75 days, and about 100 days. For departure locations between perilune and apolune ($140° < TA$

< 180°), no yellow points appear on the maps; thus, there is no maneuver direction for which a 1 m/s burn leads to lunar impact for such separation locations. Risk of recontact, represented by red regions on the map, occurs for a wide range of maneuver directions, especially for separation after apolune, 180° < TA < 282°. Depending on the location and direction of the separation, the time of flight between separation and recontact can range from a few hours to several weeks.

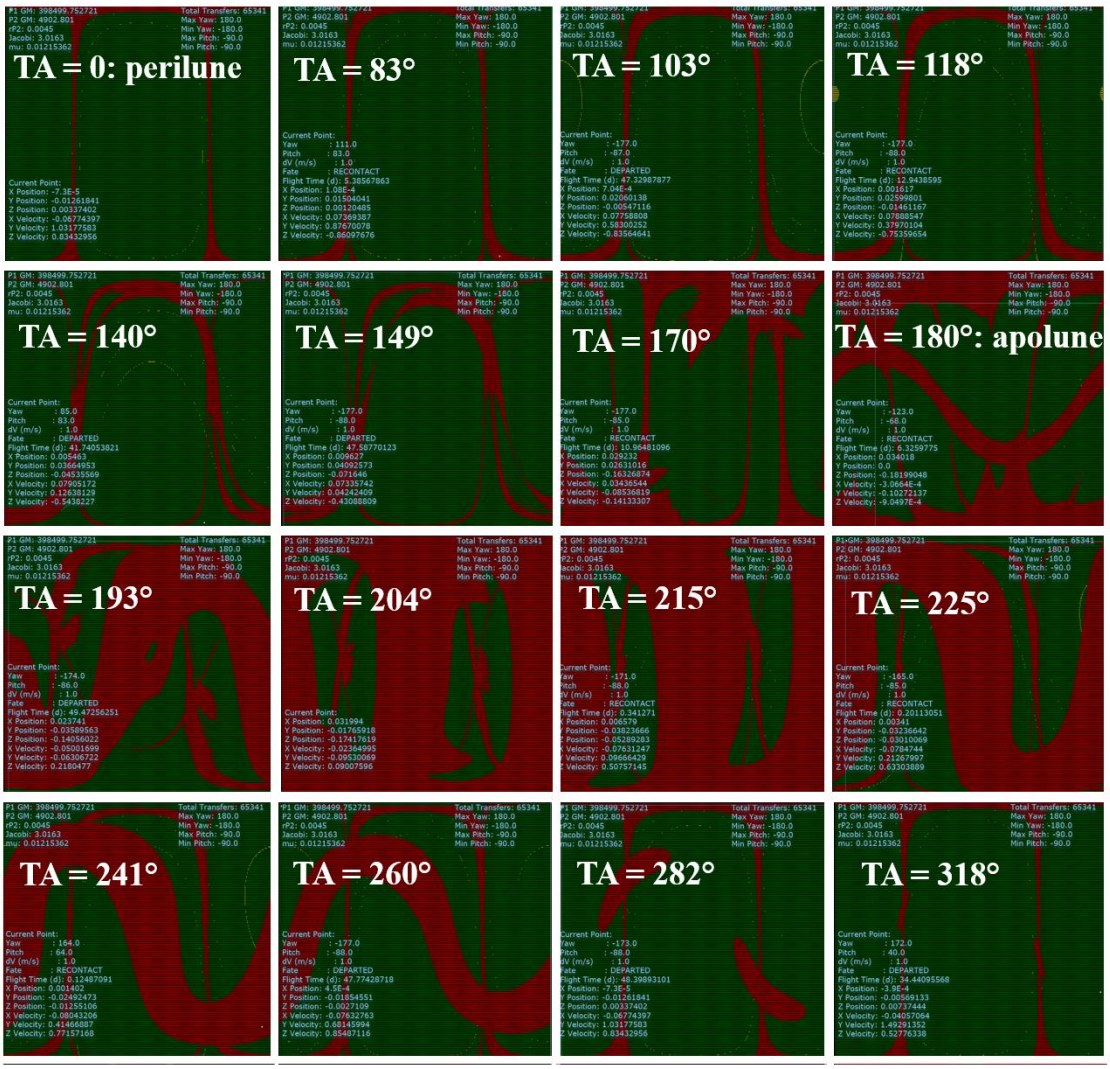

Figure 9. Recontact maps for 1 m/s separation burns from the 9:2 NRHO.

Figure 10. Recontact maps for 3.5 m/s separation burns from the 9:2 NRHO.

When the separation Δv magnitude is increased, the characteristics of the recontact maps change. Figure 11 consists of maps representing departure from the 9:2 NRHO with a 5 m/s separation burn. The larger separation burn reduces the risk of recontact and increases the risk of (or opportunity for) lunar impact. The risk of recontact does not disappear, however. For certain maneuver directions, separation after apolune where $200° \leq TA \leq 280°$ leads to recontact risk at the next perilune with a 5 m/s burn. However, recontacts after long times of flight are rare with a 5 m/s burn. In contrast, maneuver options for lunar impact are more widespread. For example, a range of burn directions centered on the anti-velocity direction at perilune result in lunar impact, as seen by the two yellow half circles in the map corresponding to TA = 0°.

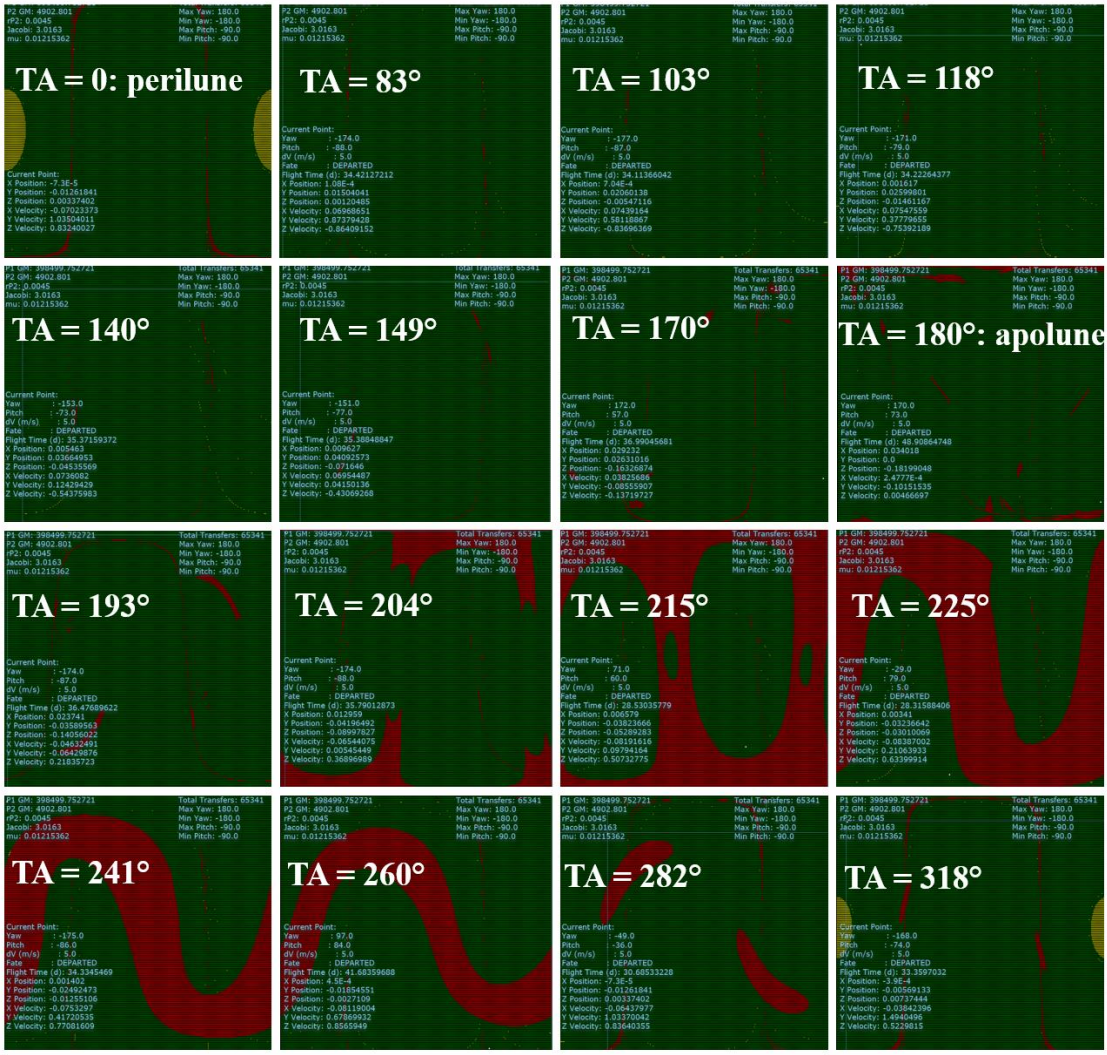

Figure 11. Recontact maps for 5 m/s separation burns from the 9:2 NRHO.

Increasing the Δv magnitude further to 15 m/s again reduces the risk of recontact and opportunities for lunar impact. A set of maps corresponding to 15 m/s separation maneuvers appears in Figure 12. Red regions still exist in the maps; departures between apolune and perilune result in recontact risk at the subsequent perilune for certain maneuver directions. However, recontact is not a risk for departure from locations corresponding to TA < 180°. Yellow regions in the maps, representing maneuver directions that lead to impact, are observed centered on the anti-velocity direction near perilune and the anti-binormal direction between apolune and perilune. Departures from apolune rarely lead to lunar impact regardless of maneuver direction for separation maneuvers up to 15 m/s in magnitude.

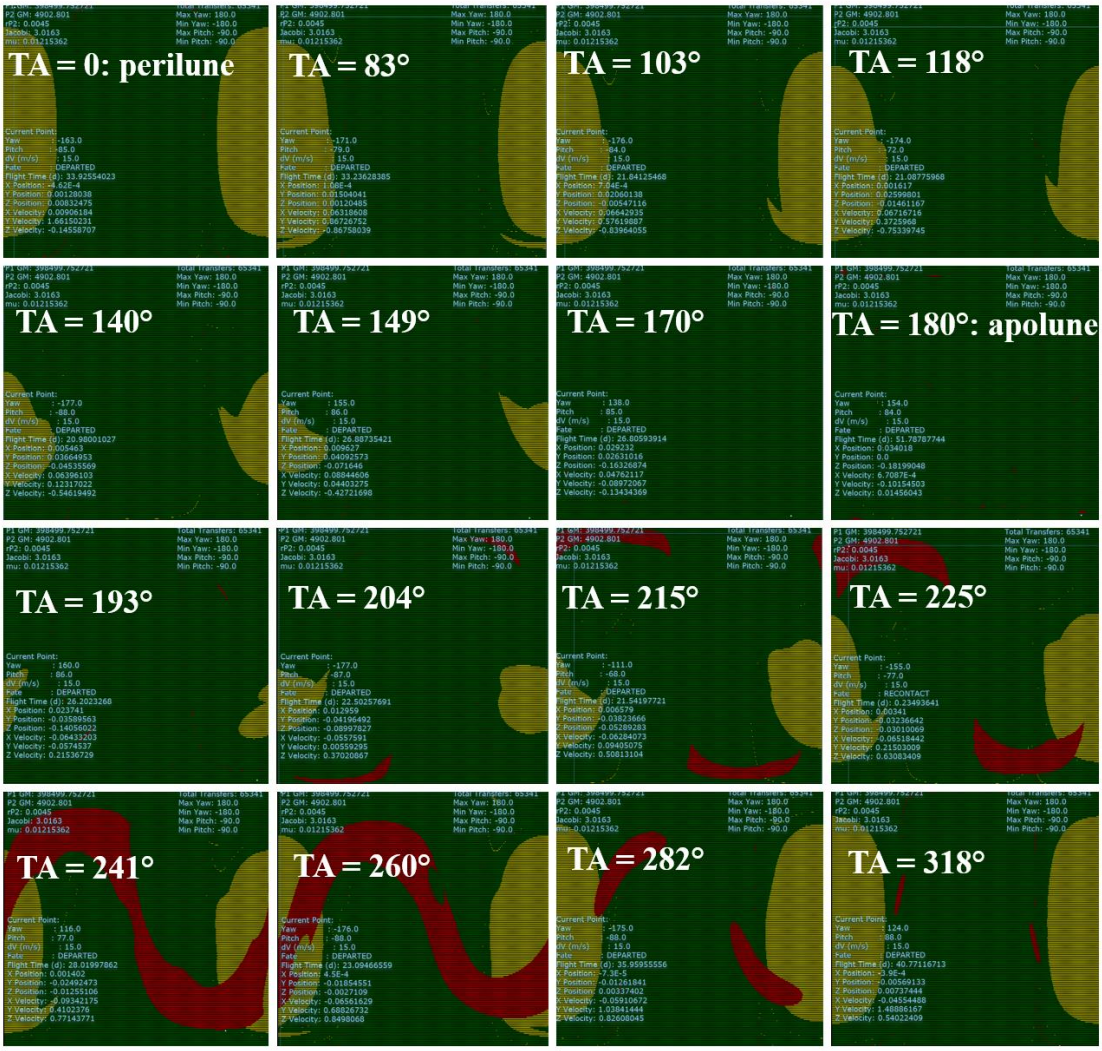

Figure 12. Recontact maps for 15 m/s separation burns from the 9:2 NRHO.

GATEWAY APPLICATION: JETTISON

A concurrent investigation[10] investigates options for jettison of objects and spacecraft from the Gateway; potential examples include spent logistics modules, cubesats, and a used, disposable lunar ascent module. The study identifies a set of separation maneuvers from the Gateway that lead to reliable escape to heliocentric space and another set that results in lunar impact. These examples are now assessed for the risk of recontact prior to departure.

First, the time of flight from Gateway separation to either lunar impact or NRHO departure is assessed. The longer the time of flight, the longer the object presents a potential risk to the Gateway and the longer effects of navigation errors, maneuver execution errors, and other perturbations have to accumulate and affect the ultimate destination of the departing object. Maps colored according to the time of flight appear in Figure 13.

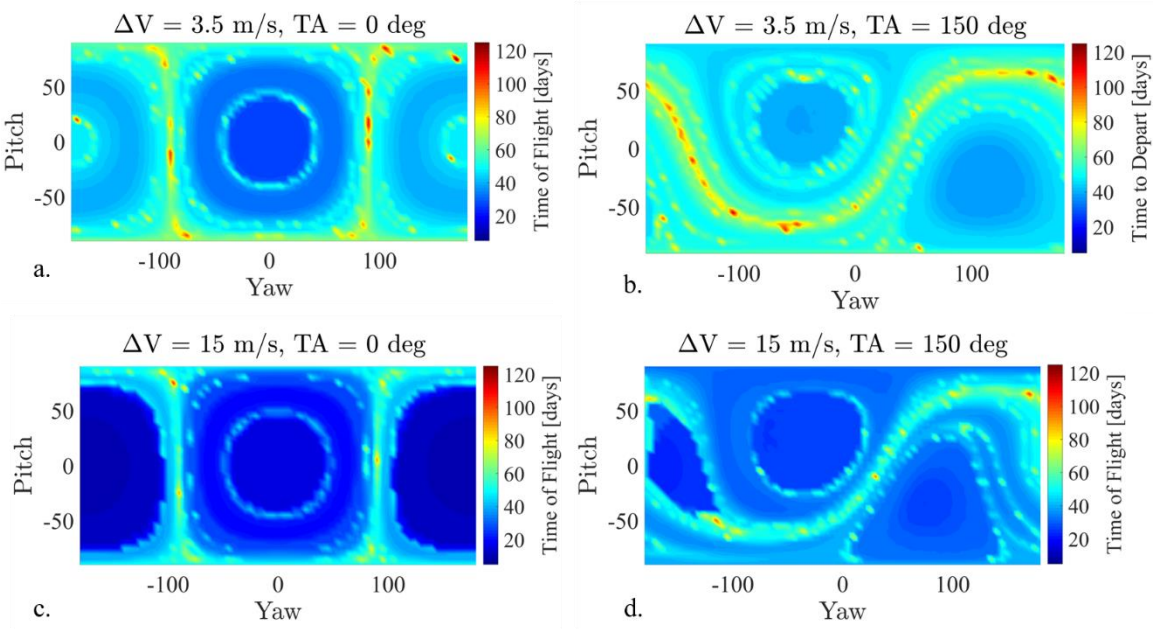

Figure 13. Time of flight to impact or NRHO departure for four separation maneuver sets.

CONCLUDING REMARKS

The assessment of recontact risk in cislunar orbits is a computationally intensive process. HPC methods accelerate the data generation and management. The DSTE employs an asynchronous implementation pattern for performing a HPC execution of a Monte Carlo scenario that leverages cloud resources. This pattern leverages standard Java APIs combined with lightweight open source frameworks available from Glassfish and Gluon. The asynchronous design combined with Gluon function mapping provides a new breed of REST service which can not only leverage cloud scale computation, but provide immediate feedback for analysts. This design can dramatically increase engineering cycles by simplifying the process by which a subject matter expert analyst initiates and processes large Monte Carlo simulations.

The data processing and visualization approaches of the DSTE have proven quite effective when applied to traversing the complicated trajectory data space associated with multibody orbit design. The DSTE is extended from a rich client mission design tool to a powerful RESTful service in a manner that is scalable. Future considerations and work that naturally extend from this include compiling the DSTE service using GraalVM to achieve a significant increase in computational performance.[15] Further, this service could be placed inside a container such that larger requests could be easily be broken apart across multiple cloud instances using a container clustering strategy such as Kubernetes.

ACKNOWLEDGMENTS

The authors would like to thank Brian McCarthy and RJ Power for assistance with plotting and for verifying the velocity transformations within the DSTE. Portions of this work were completed at NASA JSC through contract #NNJ13HA01C.

REFERENCES

[1] Gates, M., M. Barrett, J. Caram, V. Crable, D. Irimies, D. Ludban, D. Manzell, and R. Ticker, "Gateway Power and Propulsion Element Development Status," 69th International Astronautical Congress, Bremen, Germany, October 2018.
[2] Zimovan, E., K. C. Howell, and D. C. Davis, "Near Rectilinear Halo Orbits and Their Application in Cis-Lunar Space," 3rd IAA Conference on Dynamics and Control of Space Systems, Moscow, Russia, May-June 2017.

[3] Davis, D.C., S. M. Phillips, and B.P. McCarthy, "Trajectory Design for Saturnian Ocean Worlds orbiters using multidimensional Poincaré maps," *Acta Astronautica*, November 2017.
[4] D. Grebow, D. Ozimek, K. Howell, and D. Folta, "Multibody Orbit Architectures for Lunar South Pole Coverage," Journal of Spacecraft and Rockets, Vol. 45, Mar. 2008.
[5] Newman, C. P., D. C. Davis, R. J. Whitley, J. R. Guinn, and M. S. Ryne, "Stationkeeping, Orbit Determination, and Attitude Control for Spacecraft in Near Rectilinear Halo Orbit," AAS/AIAA Astrodynamics Specialists Conference, Snowbird, Utah, August 2018.
[6] J. Williams, D. E. Lee, R. L. Whitley, K. A. Bokelmann, D. C. Davis, and C. F. Berry, "Targeting Cislunar Near Rectilinear Halo Orbits for Human Space Exploration," 27th AAS/AIAA Space Flight Mechanics Meeting, Feb. 2017.
[7] Parrish, N., E. Kayser, S. Udupa, J. Parker, B. Cheetham, and D. Davis, "Survey of Ballistic Lunar Transfers to Near Rectilinear Halo Orbit," AAS/AIAA Astrodynamics Specialists Conference, Portland, Maine, August 2019.
[8] Szebehely, Z., *Theory of Orbits: The Restricted Problem of Three Bodies*, Academic Press, New York, 1967.
[9] Boudad, K. K., D. C. Davis, and K. C. Howell, "Disposal Trajectories From Near Rectilinear Halo Orbits," AAS/AIAA Astrodynamics Specialists Conference, Snowbird, Utah, August 2018.
[10] Davis, D.C., K.K. Boudad, and K.C. Howell, "Disposal, Deployment, and Debris in Near Rectilinear Halo Orbits," AAS/AIAA Space Flight Mechanics Meeting, Maui, Hawaii, January 2019.
[11] https://cloud.oracle.com/en_US/cloud-infrastructure, Web, Last Accessed July 2019.
[12] https://www.payara.fish/software/payara-server/payara-micro, Web, Last Accessed July 2019.
[13] https://gluonhq.com/products/cloudlink/, Web, Last Accessed July 2019.
[14] Davis, D.C., K.K. Boudad, R.J. Power, and K.C. Howell, "Heliocentric Escape and Lunar Impact from Near Rectilinear Halo Orbits," AAS/AIAA Astrodynamics Specialists Conference, Portland, Maine, August 2019.
[15] https://www.graalvm.org/ , Web, Last Accessed July 2019.