EARTH ORBIT RAISE DESIGN FOR THE ARTEMIS MISSION

Gregory J. Whiffen*and Theodore H. Sweetser†

The Artemis mission is an extension of the Themis mission. The Themis mission[1] consisted of five identical spacecraft in varying sized Earth orbits designed to make simultaneous measurements of the Earth's electric and magnetic environment. Themis was designed to observe geomagnetic storms resulting from solar wind's interaction with the Earth's magnetosphere. Themis was meant to answer the age old question of why the Earth's aurora can change rapidly on a global scale. The Themis spacecraft are spin stabilized with 20 meter long electric field booms as well as several shorter magnetometer booms. The goal of the Artemis[2] mission extension is to deliver the field and particle measuring capabilities of two of the Themis spacecraft to the vicinity of the Moon. The Artemis mission required transferring two Earth orbiting Themis spacecraft on to two different low energy trans-lunar trajectories ultimately ending in lunar orbit. This paper describes the processes that resulted in successful orbit raise designs for both spacecraft.

The two spacecraft used in the Artemis mission are designated "P1" and "P2". The trajectory design for the orbit raise phase of both P1 and P2 turned out to be significantly more difficult than anticipated. There were numerous reasons for this. The reasons include the spacecraft maneuvering constraints, the highly constrained ΔV budget, the precision phasing required to reach the proposed low-energy transfers to the Moon, the original proposal assumed significantly different initial states when designing the orbit raise and trans-lunar trajectories, the original proposal approximated the orbit raise with a single impulse when in reality many small burns are required, the fact that optimal design of highly elliptical transfers is numerically difficult, and the fact multiple lunar approaches during a long slow orbit raise are unavoidable creating a very complex design space.

Spacecraft maneuvering constraints that made the design challenging include: the majority of the orbit raise must be done using pulsed thrust limited to the spacecraft spin plane, the pulse size was limited for certain propellant loads to avoid fuel slosh resonances, the thrust levels are relatively low resulting in long burns and potentially large gravity losses, the spacecraft can only thrust in Sun light, spin plane thrusting results in spin axis precession affecting future maneuvers and potentially violating attitude constraints, burn directions are limited to inertially fixed directions, thruster specific impulse and thrust magnitude drop significantly with declining propellant loads, and the ΔV budget.

*Senior Engineer, Mission Design and Navigation Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109

†Senior Engineer, Mission Design and Navigation Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109

Several factors combined in unfortunate ways to further complicate the design. The Earth's shadow covers perigee for most of the orbit raise time period resulting in no thrusting allowed near perigee. This limitation increased gravity losses over those accounted for in the proposal. The propellant load for the inner spacecraft P2 forced a large fraction of the maneuvers to be performed at a lower duty cycle (shorter pulse) due to the propellant load being near a slosh resonance. This further exacerbated gravity losses, forcing more maneuvers, which in turn made it more difficult to reach the low energy transfer interface on time. Another problem was that orbit raise maneuvers happened to cause the P2 spacecraft to precess in a bad direction towards attitudes that violate operational constraints.

Another challenge is associated with maneuvers limited to a plane. The usual intuition of 1 burn allowing targeting of 3 elements and 2 burns separated in time allows for the targeting of 6 elements is not correct. In fact, even 3 separated burns can fail to provide 6 element targeting when all maneuvers are confined to a single plane.

As a result of science planning decisions that could not be anticipated at the time of the Artemis proposal, the actual initial orbits at the start of the orbit raise phase were significantly different than those assumed in the proposal. For example, see Figure 1 which show the dramatic difference between the proposal and actual initial orbit raise orbits for the P2 spacecraft.

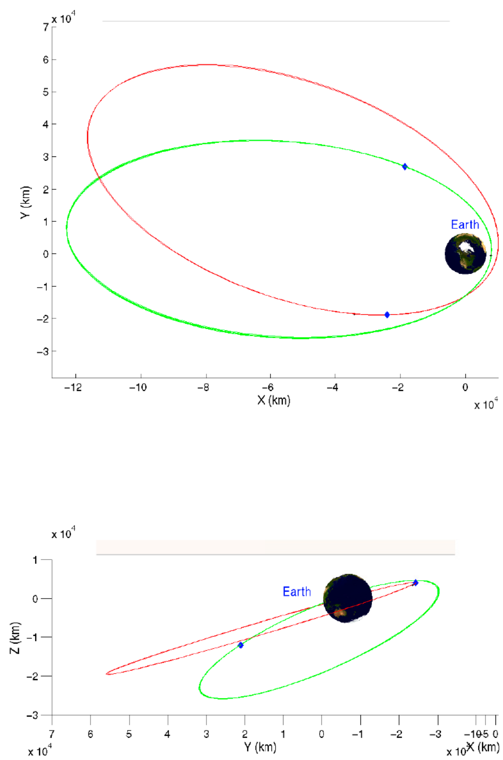

Figure 1. The assumed initial orbit of the "P2" spacecraft at the start of the orbit raise. The initial orbit assumed when Artemis was proposed is green, the actual orbit that began the orbit raise is red. Both a an ecliptic projection and a ecliptic edge on view are provided.

The P1 and P2 orbit raise designs were constructed using Mystic software[3].[4] Mystic was able to accommodate the mission constraints outlined above. However, the complex (and often treacherous) design space resulting from numerous Lunar approaches during the orbit raise phase made simple design strategies impossible. In order to provide some robustness against missed burns, burn passes were allowed only as often every other perigee passage. Sometimes it proved advantageous to separate burns by more than a single periapsis passage to take advantage of or avoid certain lunar interactions. Most perigee burns were divided into two burn arcs, one on either side of the Earth's shadow. The duration and pointing of each burn was fully optimized subject to a propellant minimization objective.

Several different strategies were attempted for the P1 orbit raise design. The strategy that proved successful for the P1 trajectory was to first optimize sets of burns on three alternating perigees to reach an orbital period of 131 hours. From states near this point forward there existed a tremendous number of possible paths involving differing lunar interactions, numbers of Earth revolutions, plane changes, and node changes over the next 140 days of ballistic propagation. It was not at all obvious which of these many paths might be feasible, and then which feasible path is the best to rejoin the low energy transfer. To address this problem, a large number of the ballistic trajectories were used as initial guesses for targeting and optimization - see Figure 2 for a small subset of the candidate trajectories that were used. Different families were organized based on the number of Earth revolutions. A computer cluster was used for this compute intensive process. Trajectories that were found to be feasible or near feasible were then further refined by moving the targeting (or re-joining) of the low-energy transfer to successively later dates.

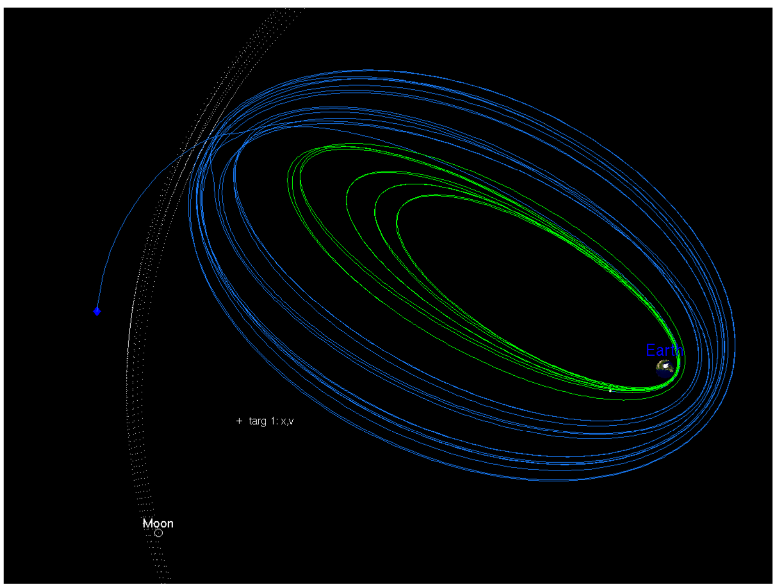

Figure 2. An example of a candidate transfer before optimization and targeting.

The key to reaching the P1 low energy transfer was a pair of lunar flybys separated by only 14 days. For any transfer to remain under the ΔV budget, it was necessary to match these flybys as closely as possible. Exact matching is not necessary. Intuitively, re-joining the low-energy transfer at later times will provide increasing efficiency. It was expected (and was found) that re-joining much beyond the second lunar flyby provided diminishing returns. The final design required only 105.1 [m/s] to reach the lissajous injection state at the end of the deep space low energy transfer.

Figure 3 illustrates the large plane change required during the transfer for P1.

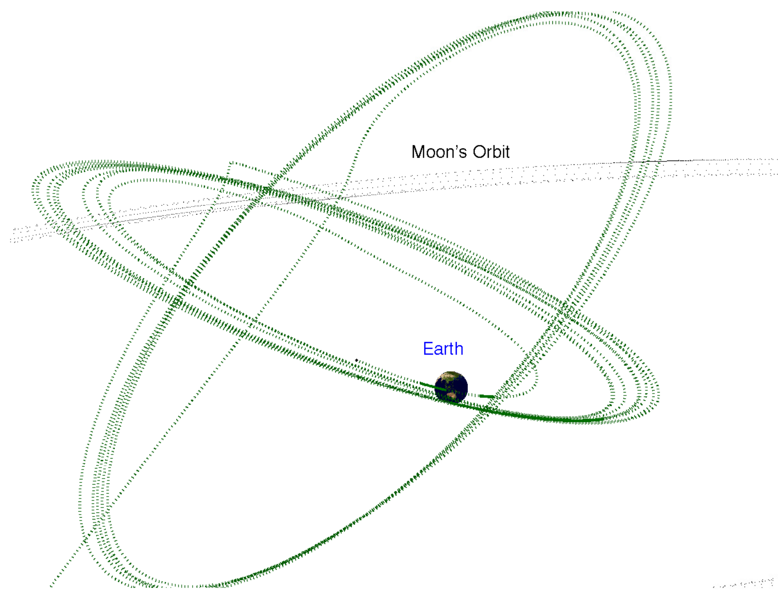

Figure 3. An oblique view of the final orbits of P1 during the orbit raise transfer.

The P2 orbit design was more complex than the P1 design due to the fact P2 begins in a much smaller orbit. A process similar to the P1 design process was used to develop the P2 orbit raise design. Very carful planning of distant lunar approaches was necessary to keep under the allocated ΔV budget. The ΔV budget was more constraining for P2. The P2 orbit raise ended up requiring 42 burns The method used was a branching process. Each orbit raise maneuver was designed to reach several different orbital periods (different period = different "branch"). Subsequent maneuvers reaching longer periods were designed for each branch. The most promising branches were continued while poor performing branches were abandoned. Poor performing branches often lead to situations where lunar interactions reduced the orbit period or required long periods without maneuvers to avoid disadvantageous lunar interactions. High performing branches ended up with advantageous distant lunar interactions early on. Distant lunar interactions were sought that provided maneuver savings as little as 1 meter per second early in the orbit raise. The final few orbit raise maneuvers required very careful planning to maximize the positive influence of the Moon.

ACKNOWLEDGMENT

This research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract to the National Aeronautics and Space Administration.

REFERENCES

[1] V. Angelopoulos, "The Themis Mission," *Space Science Reviews*, No. 141, 2008, pp. 5–34.

[2] V. Angelopoulos, "The Artemis Mission," *Space Science Reviews*, No. Submitted, 2010, pp. ?–?

[3] G. J. Whiffen, "Static/Dynamic for Optimizing a Useful Objective," United States Patent, No. 6,496,741, Filed March 1999. Issued December 2002.

[4] G. J. Whiffen, "Mystic: implementation of the Static Dynamic Optimal Control Algorithm for High-Fidelity, Low-Thrust Trajectory Design," *Proceedings of the AIAA/ASS Astrodynamics Specialists Conference*, Keystone, Colorado, 21-24 August 2006. Paper AIAA 2006-6741.