AAS 07-348

DESIGN SPACE PRUNING TECHNIQUES FOR LOW-THRUST, MULTIPLE ASTEROID RENDEZVOUS TRAJECTORY DESIGN

Kristina Alemany[*], Robert D. Braun[†]

In 2006, the 2nd Global Trajectory Optimization Competition (GTOC2) posed a "Grand Asteroid Tour" trajectory optimization problem, where participants were required to find the best possible low-thrust trajectory that would rendezvous with one asteroid from each of four defined groups. As a first step, most teams employed some form of design space pruning, in order to reduce the overall number of possible asteroid combinations. Because of the large size of the problem, teams were not able to determine if their pruning technique had successfully eliminated only bad solutions from the design space. Therefore, a small subset of the GTOC2 problem was analyzed, and several design space pruning techniques were applied to determine their effectiveness. The results indicate that the pruning techniques chosen by the participants likely eliminated good solutions from the design space, because they either did not accurately represent the low-thrust problem or could not be considered independently without the effect of other factors.

INTRODUCTION

In 2006, the 2nd Global Trajectory Optimization Competition (GTOC2)[1] posed a trajectory optimization problem of a "Grand Asteroid Tour." Participants were required to design the best possible trajectory, using electric propulsion, that would rendezvous with one asteroid from each of four defined groups. The given objective function rewarded trajectories with low propellant consumption and low total flight time. The candidate asteroids totaled almost 1000, resulting in over 41 billion possible asteroid combinations. Furthermore, launch date, launch v_∞, times of flight, and stay time at each asteroid were free design variables. The large number of possible asteroid combinations prohibited each and every one from being examined, and the multi-modal nature of the design space with respect to the other design variables prohibited a simple gradient-based optimizer from being used for a single asteroid combination. In addition to the large size of the global optimization problem, each local trajectory optimization required determining the best thrust profile to minimize propellant consumption. In order to make the problem more manageable, all of the participating teams first employed some form of pruning step in order to eliminate what they believed to be the worst solutions from the design space. This step included removing both asteroids and asteroid combinations, as well as portions of the launch date and flight time domain for particular asteroid combinations. Although the teams were confident in their local trajectory optimization capabilities, most teams cited one of their major weaknesses to be their chosen pruning technique, believing that they actually eliminated some of the best solutions from the design space. Because the true optimum can not be determined for such a large problem, it is not possible to determine if this conclusion is true.

[*] Ph.D. Candidate, Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, kristina_alemany@ae.gatech.edu
[†] David and Andrew Lewis Professor of Space Technology, Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, robert.braun@ae.gatech.edu

Some of these pruning techniques, however, can be evaluated on a subset of the GTOC2 problem, to determine their efficacy at quickly eliminating bad solutions while keeping the best asteroids and asteroid combinations in the design space. A set of 22 asteroids was chosen from the full GTOC2 set, the resulting problem was discretized in terms of launch date and times of flight, and all feasible solutions were found for the discretized design space. Using these results, a number of the pruning techniques used by the GTOC2 teams were applied to the sub-problem, and evaluated based on the number of good solutions that would have been eliminated from the design space, if any. These pruning techniques fall into several groups: ephemeris-based metrics that eliminate specific asteroid combinations, phase-free approximations that also eliminate asteroid combinations, and metrics that take phasing into account, to eliminate areas of the time domain.

PROBLEM STATEMENT

In order to evaluate the effectiveness of various pruning techniques on a particular problem, the optimal solution must be known. Because this work was motivated by the asteroid tour problem posed in GTOC2, a subset of that problem was solved, which serves as the basis for evaluating the pruning techniques. The original GTOC2 problem asked participants to determine a low-thrust trajectory that maximizes the ratio of final spacecraft mass to total time of flight, while rendezvousing with one asteroid in each of four different groups: Group 1 is made up of the Jupiter Former Comets and contains 96 asteroids, Group 2 is made up of C- or M-class asteroids and contains 176 asteroids, Group 3 is made up of S-class main belt asteroids, and contains 300 asteroids, and Group 4 is made up of Aten asteroids and contains 338 asteroids. Additionally, launch from Earth must occur between 2015 and 2035, with a total flight time not to exceed 20 years. A minimum stay time of 90 days was required at each asteroid.

The problem solved in this work contains 6 Group 1 asteroids, 8 Group 2 and 3 asteroids (combined due to their similar semi-major axis values), and 8 Group 4 asteroids. Figure 1 plots these asteroids, as a function of their semi-major axis, eccentricity, and inclination. As can be seen, these asteroids lie in three distinct groups, at increasing distance from the Earth. These asteroids were chosen as a representative sample, in terms of their orbital elements, from the full set of GTOC2 asteroids. In order to simplify the problem further, the objective function was chosen simply to be the final mass of the spacecraft, and the following constraints were placed on the flight times: Earth to Group 4 $\leq$ 600 days, Group 4 to Group 2/3 $\leq$ 1800 days, and Group 2/3 to Group 1 $\leq$ 1200 days. These constraints assume that the asteroids will be visited in order of increasing semi-major axis. (This validity of this assumption will be addressed later in the paper.) Lastly, the launch window was shortened to fall between 2015 and 2025, inclusive, and the stay time at each asteroid was fixed at 90 days. While flight time no longer directly appears in the objective function, it is dealt with in the chosen constraints. The other assumptions laid out in GTOC2 were not changed in the sub-problem. Launch from Earth is constrained by a hyperbolic excess velocity (v_∞) of up to 3.5 km/s with no constraint on direction. The spacecraft has a fixed initial mass of 1500 kg, which does not change with launch v_∞, and a minimum final mass of 500 kg. The propulsion is modeled to have a constant specific impulse of 4000 s and a thrust level of 0.1 N, and can be turned on and off at will.

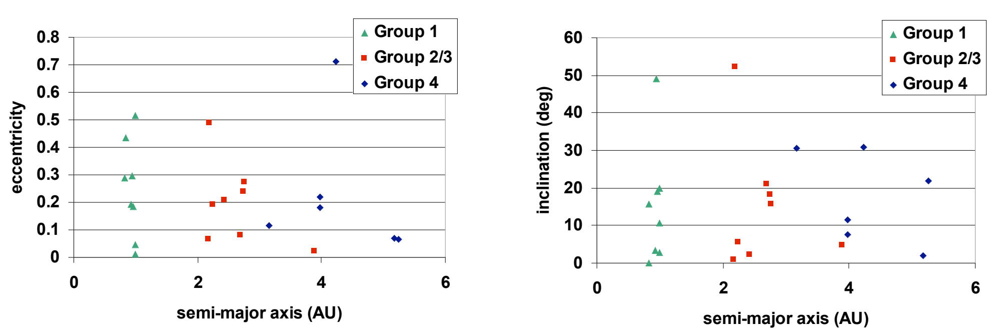

Figure 1 Set of asteroids for GTOC2 sub-problem.

This problem was solved using MALTO, a low-thrust trajectory optimization tool developed at JPL, based on a direct method by Sims and Flanagan[2]. In MALTO, the trajectory is divided into legs that begin and end at control nodes. On each leg is a match point, and the trajectory is propagated forwards from the previous control node and backwards from the subsequent control node to the match point. Each leg is also subdivided into numerous segments containing an impulsive ΔV at the middle of each segment, in order to approximate a continuous thrust problem. The resulting constrained, nonlinear optimization problem is solved within MALTO using SNOPT, which was developed at the University of California San Diego. When MALTO was developed, it was verified against SEPTOP for a number of different trajectories types, including a flyby of the asteroid Vesta with a Mars gravity assist, a rendezvous with the comet Tempel 1, and a flyby of Pluto with two Venus gravity assists and one Jupiter gravity assist[3].

In order to solve the GTOC2 sub-problem, the design space was discretized in terms of launch date and times of flight, and each leg of the trajectory was analyzed separately. The launch date from Earth was discretized in 30 day steps, and the time of flight to the first asteroid (Group 4) was discretized in 100 day steps up to 600 days. MALTO was used for each case to determine the thrust profile and launch v_∞ that maximizes the final mass at the arrival asteroid, based on a 1500 kg initial spacecraft mass. The time of flight for the 2[nd] leg was also discretized in 100 day increments, up to 1800 days. For each feasible leg 1 trajectory (final mass greater than 500 kg), the corresponding leg 2 trajectory was calculated to each of the Group 2/3 asteroids, for each of the discretized times of flight. Finally, the set of leg 3 trajectories was calculated in a similar fashion, starting from all of the feasible leg 2 trajectories. This approach allows not only the best asteroid combination to be determined, but the entire set of feasible solutions, ranked by final mass. In mission design, a number of other factors must be considered in choosing a “best solution” in addition to final mass and flight time. Furthermore, several of the asteroid combinations yield very similar values of final mass, making them essentially indistinguishable within the uncertainty introduced by a discretized solution. Also, in order to evaluate the candidate pruning techniques, it is beneficial to have the full set of trajectory solutions. In pruning the initial design space, the goal is not only to keep the optimum solution in the design space, but also to keep the entire set of best solutions.

Assuming the asteroids are visited in order of increasing semi-major axis (Earth – Group 4 – Group 2/3 – Group 1), the resulting set of feasible solutions contains only 41 of the possible 384 asteroid combinations. This set of solutions contains 4 Group 1 asteroids, 5 Group 2/3 asteroids, and 4 Group 4 asteroids (although certainly not every permutation of these 13 asteroids). Table 1 lists the asteroids that appear in the feasible combinations, along with their pertinent orbital elements.

Table 1: Orbital elements of asteroids appearing in feasible combinations.

Asteroid Name	Group #	semi-major axis (AU)	eccentricity	inclination (deg)	longitude of the asc. node (deg)	argument of periapsis (deg)
“2006 QQ56”	4	0.987	0.047	2.83	163.33	332.96
“2002 AA29”	4	0.994	0.013	10.74	106.47	100.61
“2004 FH”	4	0.818	0.289	0.021	296.18	31.32
Apophis	4	0.922	0.191	3.33	204.46	126.40
Geisha	2/3	2.24	0.193	5.66	78.34	299.88
Aquitania	2/3	2.74	0.237	18.13	128.31	157.68
Medusa	2/3	2.17	0.065	0.937	159.65	251.13
Hertha	2/3	2.43	0.207	2.31	343.90	340.04
Daphne	2/3	2.77	0.272	15.76	178.16	46.22
Kostinsky	1	3.99	0.220	7.64	257.11	163.00
Caltech	1	3.16	0.114	30.69	84.61	294.92
Pandarus	1	5.17	0.068	1.85	179.86	37.74
Potomac	1	3.98	0.181	11.40	137.51	332.82

The best discretized solution, plotted in Figure 2, departs Earth on March 1, 2015 with a launch v_∞ of 2.59 km/s. The total flight time from Earth departure to the final asteroid rendezvous is 3580 days, which includes the two 90-day stay times at each intermediate asteroid, and the arrival mass is 903.27 kg. This best solution visits the following asteroids in order: "2006 QQ56" – Medusa – Kostinsky.

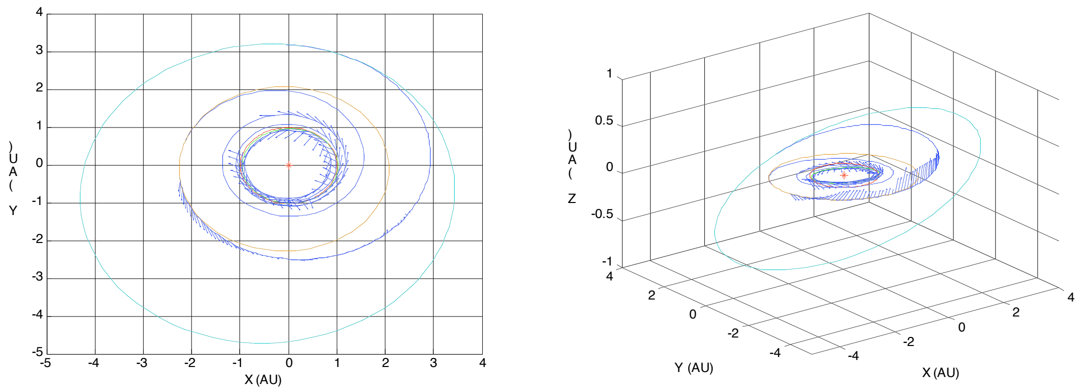

Figure 2: Best asteroid trajectory.

Table 2 then lists the 10 best asteroid combinations, ordered in terms of final mass. It is interesting to note that although flight time does not appear explicitly in the objective function, several of the best trajectory solutions still have flight times less than the constraint value. Although the final mass for one particular leg may be larger for longer flight times, the shorter flight time produces a better overall solution because of phasing considerations with subsequent asteroids.

Table 2: Ten best asteroid combinations based on final mass.

Earth Dep. Date	Ast. 1	Ast. 2	Ast. 3	Leg 1 TOF (days)	Leg 2 TOF (days)	Leg 3 TOF (days)	Mf (kg)
03/01/2015	"2006 QQ56"	Medusa	Kostinsky	600	1600	1200	903.27
03/29/2021	Apophis	Hertha	Pandarus	300	1800	1200	855.02
01/01/2015	"2002 AA29"	Medusa	Kostinsky	600	1700	1200	831.42
09/11/2018	"2006 QQ56"	Geisha	Kostinsky	600	1700	1200	826.20
08/28/2015	"2006 QQ56"	Geisha	Caltech	600	1700	1200	811.74
03/01/2015	"2004 FH"	Medusa	Kostinsky	500	1800	1200	807.01
02/17/2023	Apophis	Geisha	Kostinsky	600	1800	1200	773.23
01/30/2015	Apophis	Medusa	Kostinsky	600	1800	1200	765.22
06/29/2015	"2002 AA29"	Geisha	Caltech	600	1800	1200	759.99
06/23/2016	"2006 QQ56"	Hertha	Kostinsky	600	1800	1200	758.79

EVALUATION OF GTOC2 PRUNING TECHNIQUES

A number of different design space pruning techniques were employed by the GTOC2 participants in an attempt to quickly eliminate bad solutions from the design space. The fastest of the pruning techniques employed was ephemeris-based, using semi-major axis, inclination, and longitude of the ascending node. Many teams also made use of approximations to low-thrust trajectories, including two-impulse Lambert solutions with either single or multiple revolutions. Finally, phasing was taken into consideration using a number of different screening methods. Because the full GTOC2 problem was so large, however, the participants could not definitively conclude whether the techniques they employed were successful or if they inadvertently eliminated the best solutions from the design space. Therefore, some of the more widely used pruning techniques were applied to the sub-problem described above in order to evaluate their effectiveness.

For many of the pruning techniques considered, the maximum final mass for each asteroid combination is plotted against the metric in question. For the leg 1 trajectories, the initial mass for all cases is 1500 kg, making all of the final masses directly comparable. The initial masses for the leg 2 and leg 3 trajectories, however, are based on the corresponding final masses from the previous leg. Therefore, the final masses for these legs differ both because of the mass fraction for each particular trajectory but also because of the initial mass of that trajectory. Therefore, for each of the pruning methods examined, the effect on the mass ratio of leg 2 and leg 3 was examined, to make sure the observed trends are not solely due to different values of initial mass. Additionally, for leg 3, all of the possible asteroid combinations were examined over the date range for an initial mass of 1500 kg, in order to warrant a more fair comparison between cases. Although this additional data is not presented here, it was found to be consistent with the mass data presented below.

Ephemeris-based Pruning Techniques

The one pruning method that almost all participants employed was to visit the asteroids in order of increasing semi-major axis. Therefore, for the original problem, the chosen order was: Earth – Group 4 – Group 2/3 – Group 2/3 – Group 1 (Group 2 and Group 3 overlap significantly in terms of semi-major axis). Although it makes intuitive sense to visit the asteroids in order of either increasing or decreasing semi-major axis, it is not immediately apparent that they should be visited specifically in order of closest to furthest from Earth. Because time of flight appeared in the objective function for the original GTOC2 problem, however, it was imperative to visit the asteroids in order of increasing semi-major axis to reduce the overall flight time of the mission. By adding this restriction to the original GTOC2 problem, participants were able to immediately reduce the number of total asteroid combinations from over 41 billion to 3.4 billion. Therefore, the same restriction was applied to the GTOC2 sub-problem, reducing the number of asteroid combinations from 2304 to 384.

Most teams also screened out asteroids and asteroid combinations based on their inclination, based on the conjecture that large inclination changes require significant amounts of propellant, as is the case for impulsive orbit transfers. Figure 3 plots the maximum final mass for each asteroid combination as a function of the absolute value of the inclination change between the starting and ending body (no differences were found in the results if a distinction was made between positive and negative inclination changes). As will be true for all similar plots presented, only asteroid combinations that were actually analyzed are plotted. For example, four of the eight Group 4 asteroids yielded no feasible leg 1 trajectories, and were therefore not considered in analyzing subsequent leg 2 and leg 3 trajectories. For leg 1, because there were only eight possible combinations, additional Group 4 asteroids were randomly selected and analyzed in order to add more data points. Furthermore, any asteroid combinations that resulted in a maximum final mass less than 500 kg were deemed infeasible and appear as 0 kg in the plots.

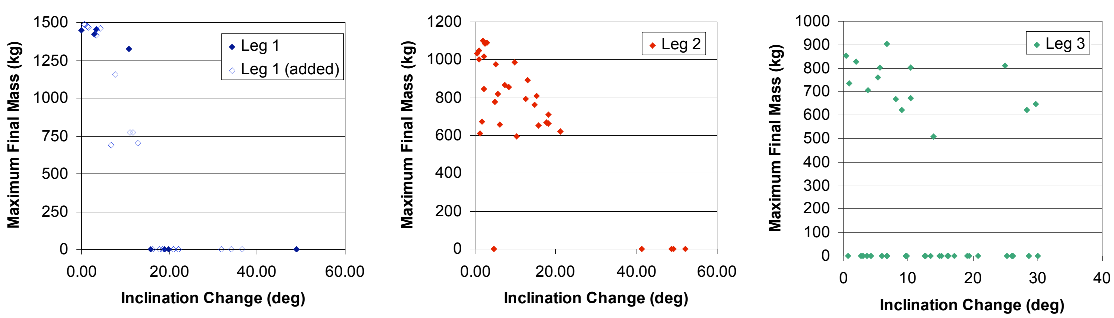

Figure 3: Maximum final mass for each trajectory leg as a function of inclination change.

For leg 1 and leg 2, there is a perceptible correlation between maximum final mass and inclination change. For the asteroids considered, there is also a maximum value of inclination change above which there are no feasible solutions for the date range considered. Therefore, for either of these legs, inclination change could certainly be used to prune certain asteroid combinations from the design space. The challenge, which will be addressed later in the paper, is to determine the inclination change above which asteroid combinations will be removed for the full GTOC2 problem. The correlation for leg 3, however, is not as apparent. One should note that the number of infeasible trajectories is largely due to the piecewise approach to solving the sub-problem. The initial mass for many of the leg 3 cases is only slightly greater than 500 kg. This results in numerous infeasible trajectories. There are feasible solutions, however, for large values of inclination change between the Group 2/3 asteroids and the Group 1 asteroids. In fact, the 5th best asteroid combination has a leg 3 inclination change of 25°. If inclination-based pruning were to be used to eliminate Group 1 asteroids, it is likely that some of the best solutions would have been eliminated from the design space. If all of the trajectories are plotted together, however, as in Figure 4, it is more likely that the feasible leg 3 asteroid combinations would not have been eliminated from the design space. Of course, this would still depend on what maximum allowable inclination is chosen for this pruning step.

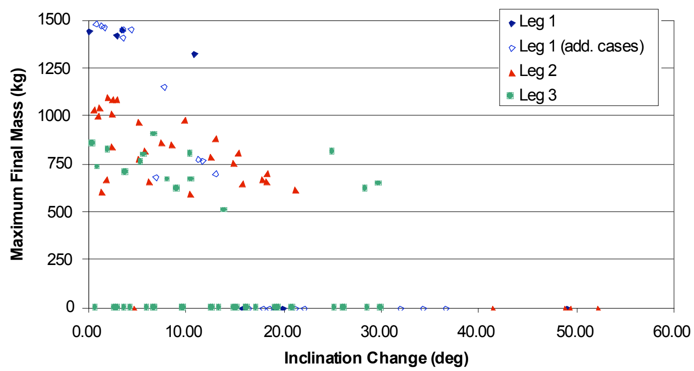

Figure 4: Maximum final mass for all trajectory legs as a function of inclination change.

Some of the GTOC2 teams attempted to prune the design space by combining inclination change with change in the longitude of the ascending node between two asteroids. The premise behind this pruning method is that if the change in the longitude of the ascending node is small, the inclined orbits are more closely aligned, therefore resulting in less propellant required. In this paper, these two variables were combined as follows:

$$J = W_i \cdot abs\left(\frac{\Delta i}{(\Delta i)_{\max}}\right) + W_\Omega \cdot abs\left(\frac{\Delta\Omega}{(\Delta\Omega)_{\max}}\right) \tag{1}$$

Figure 5 plots the maximum final mass for each asteroid combination analyzed, as a function of the above weighted combination. Two weighting scenarios are examined. First (on the left) is a 50/50 split between inclination and ascending node. Second (on the right) is a 75/25 split between inclination and ascending node. As can be seen, the correlation is not nearly as pronounced as when considering inclination change alone. As seen above, three leg 3 asteroid combinations that appear among the final feasible solutions have particularly high inclination changes (25.03°, 28.38°, and 29.75°). The corresponding ascending node change for each of these combinations is 6.27°, 100.71°, and -75.04°, respectively. Although the combination with the smallest value of ascending node change yields the highest final mass of this set (811.74 kg), the other two larger values still yield feasible solutions.

While most participating teams used inclination and ascending node as pruning techniques, a handful of other teams considered several other ephemeris-based methods, none of which proved to be reliable when applied to the sub-problem. One method was to choose Group 1 asteroids that have low energies – therefore, asteroids with the smallest values of semi-major axis. In the sub-problem, however, one of the Group 1 asteroids (Pandarus, a = 5.17 AU), appeared numerous times in the set of feasible trajectories, including in the 2nd best overall trajectory. Another method screened out asteroid combinations that had the largest distances between the first asteroid's apoapsis and the second asteroid's periapsis. For the sub-problem, there was little to no correlation between this distance and the final mass of that particular asteroid pair.

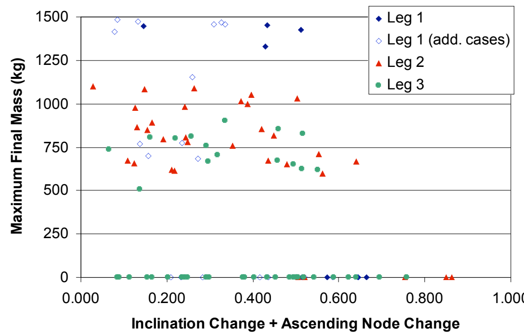

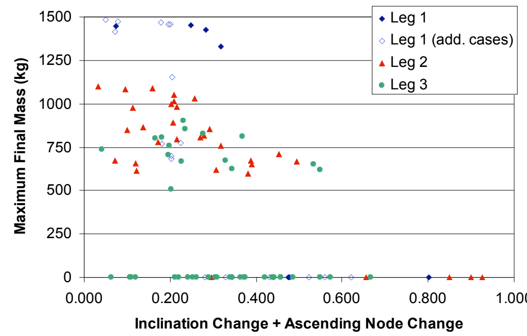

Figure 5: Maximum final mass as a function of a weighted combination of inclination change and ascending node change (50/50 weighting on left, 75/25 weighting on right).

Phase-Free, High-Thrust Approximations

In addition to ephemeris-based pruning, most of the teams used Lambert high-thrust solutions as an approximation to the low-thrust problem. Before considering phasing for particular asteroid combinations, a number of teams looked at phase-free, optimal transfers in order to further reduce the number of asteroid combinations remaining in the design space. The intent behind using phase-free two impulse transfers is to quickly determine the most reachable asteroids. Of course, there is no guarantee that the optimal asteroid configuration for a given asteroid pairing will occur during the date range given in the problem. In this paper, the method developed by Shen and Tsiotras is used, which extends Battin's method to calculate multiple-revolution Lambert solutions[4].

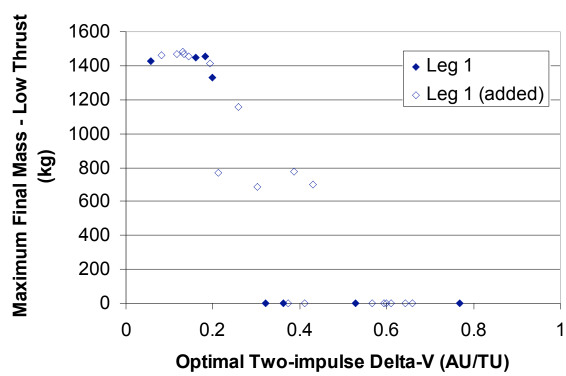

Figure 6: Maximum final mass as a function of optimal two-impulse delta-v for leg 1 asteroids.

Figure 6 plots the optimal, two-impulse transfer for all of the leg 1 asteroids, with N = 1 revolution. The time of flight is set at 600 days, to represent the maximum allowable flight time for that leg. As can be seen, there is a definite correlation between increasing delta-v and decreasing final mass for the low-thrust trajectory solution. Figure 7 then plots the optimal two-impulse transfer for all of the leg 2 asteroids. Although leg 2 had the longest transfer time and the greatest number of revolutions, the two-impulse approximation is still strongly correlated to the final low-thrust mass for that leg. For leg 2, a time of flight of 1800 days was initially used, resulting in a maximum of one revolution for most of the asteroid pairs (left plot). In order to find a solution more representative of the low-thrust solutions, with N = 2 revolutions, the time of flight was increased to 3600 days (right plot). One GTOC2 team that used Lambert arcs between asteroids added additional time to allow for "spiraling." As can be seen from the plot, however, increasing the number of revolutions actually decreases the correlation between the low thrust final mass and the high-thrust delta-v.

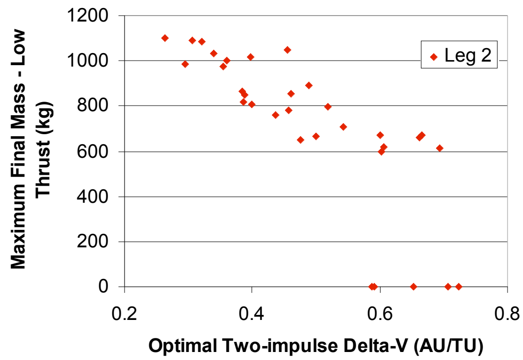

Figure 7: Maximum final mass as a function of optimal two-impulse delta-v for leg 2 asteroid pairings. N = 1, TOF = 1800 days (left). N = 2, TOF = 3600 days (right).

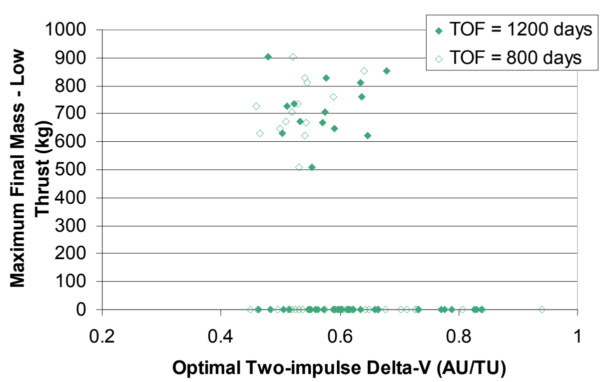

Figure 8: Maximum final mass as a function of optimal two-impulse delta-v for leg 3.

Figure 8 plots the optimal two-impulse transfer for leg 3, with N = 0, and for two times of flight. Although there is an apparent delta-v value above which there are no feasible trajectories, below that delta-v value, there appears to be no correlation between the low-thrust final mass of each asteroid pairing and its corresponding two-impulse delta-v. Just as with the ephemeris-based pruning methods, the phase-free high-thrust approximation seems fail when applied to the final leg.

Pruning Techniques Based on Phasing Considerations

Once the participating teams reduced the number of candidate asteroids using the above techniques, many of the teams then reduced the design space further by taking phasing into consideration. Almost all of the teams used two-impulse Lambert solutions in an attempt to take asteroid phasing into account. While the optimal, phase-free, two-impulse approximation appeared to be an adequate screening technique for eliminating asteroid combinations, it does not consistently represent the low-thrust trajectory over a range of launch dates and times of flight. For example, consider the leg 1 trajectory from Earth to Apophis, departing Earth on February 22, 2022, with a time of flight of 600 days. Both the low-thrust and high-thrust solutions are plotted in Figure 9. The Earth is plotted in green, Apophis in red, and the spacecraft trajectory in blue. The high-thrust solution represents an N=1 Lambert solution. Because of the large inclination of the high-thrust transfer orbit, and therefore the large delta-v required for this transfer, this launch date would likely be eliminated from the design space. The low-thrust solution, however, yields a leg 1 final mass of 1421.45 kg, which is among the best solutions over the range of launch dates.

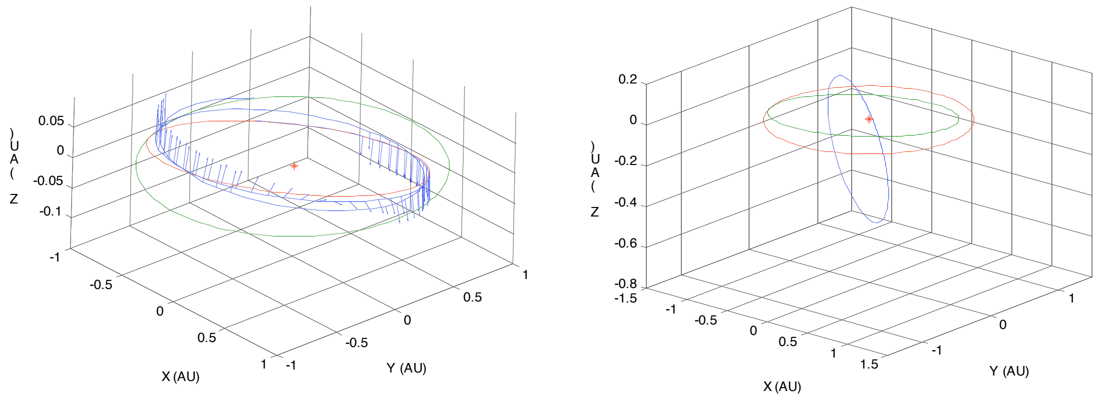

Figure 9: Low-thrust (left) and high-thrust (right) solutions for Earth-Apophis, departing on 2/22/22 with a 600-day time of flight.

For each of the candidate asteroid combinations in the sub-problem, a two-impulse solution was computed over the range of departure dates in consideration, and was plotted against the low-thrust solutions for the same date range. Because of the large amount of data generated, these are not presented here. In general, when the high-thrust final mass was plotted as a function of departure date for each asteroid combination, the resulting shape matched well when compared to the corresponding low-thrust solution. Because of situations like the one presented above, however, this two-impulse approach to pruning can not be consistently relied upon to preserve areas of the design space that yield the best solutions.

Another approach used to address phasing was to determine when the Group 1 asteroids were at their perihelion, noting that it is efficient to rendezvous with this last asteroid just after its perihelion passage. The previous asteroids and corresponding departure dates were then chosen such that the spacecraft would in fact arrive at the last asteroid just after perihelion. Figure 10 plots all of the feasible trajectories generated, along with the corresponding true anomaly of the last asteroid as a function of arrival date. Although the largest values of final mass do not occur directly after perihelion passage for all four Group 1 asteroids, they certainly all lie in the vicinity of perihelion. Therefore, this could be used as an effective pruning method to limit the date range examined.

Almost all of the participating teams attempted to address asteroid phasing, either using the two techniques explained here, or using other more qualitative methods. Based on the team presentations of their GTOC2 results, how to properly account for phasing in the pruning process seemed to elude most teams. At the same time, most teams identified the ability to account for phasing as one of the most important pruning steps.

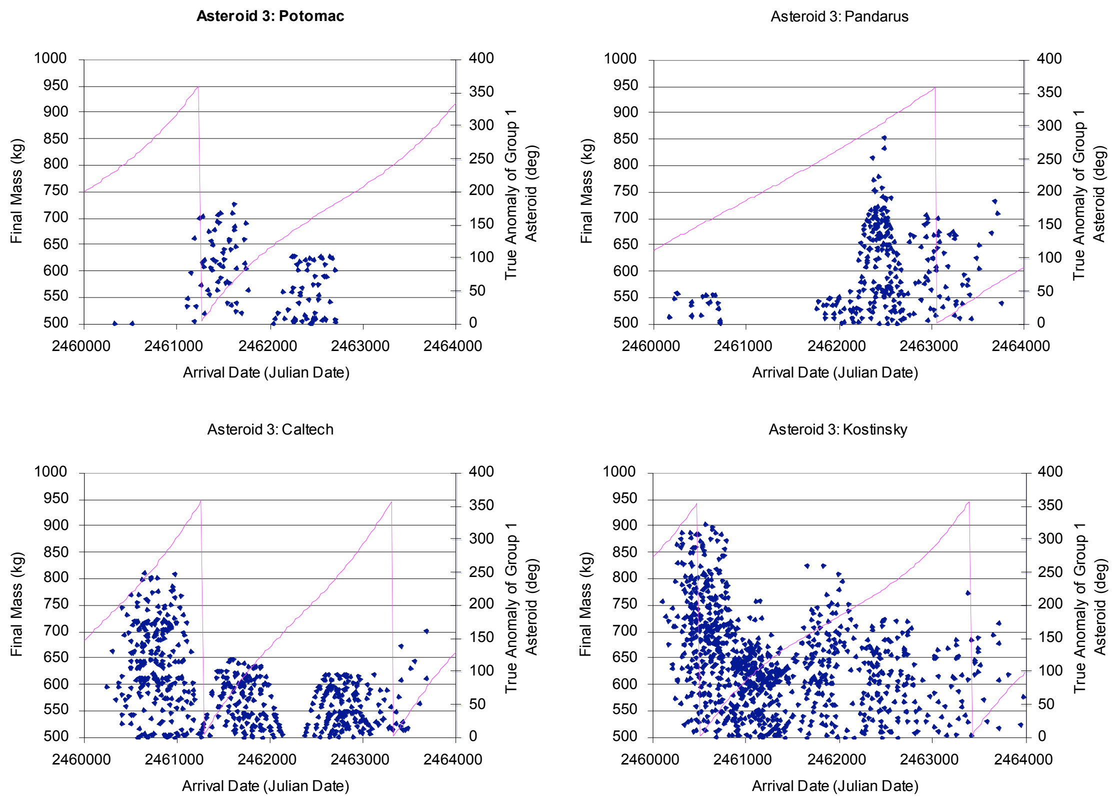

Figure 10: Final mass and true anomaly of arrival asteroid for all feasible trajectories.

SUMMARY OF PRUNING TECHNIQUES

At the conclusion of GTOC2, most teams identified their chosen pruning technique(s) as the biggest question mark in their solution process. Many teams believed that the method they chose actually eliminated some of the best solutions, although due to the size of the problem, they were unable to verify this statement. Therefore, the most popular methods used were applied to a sub-problem of GTOC2 in order to test the effectiveness of these techniques at keeping the best solutions in the design space. Based on these results, it appears likely that many of the teams did, in fact, eliminate some of the best solutions during their pruning step. Although the pruning rules chosen generally worked for most of the asteroid combinations or departure dates considered, there always seemed to be an exception to the rule. It is because of these exceptions that good designs were likely eliminated during the pruning process.

Two general types of metrics were used by the GTOC2 participants: ones that did not take phasing into account and simply eliminated asteroid combinations, and those that did consider phasing in order to shrink the design space in terms of Earth launch date, asteroid departure times, and/or times of flight. In terms of the metrics that did not consider phasing, inclination change and optimal, phase-free, two-impulse delta-V were best suited for screening the asteroid combinations for the sub-problem. Both metrics, however, appeared to yield better results for the leg 1 and leg 2 trajectories. It is likely that feasible trajectories would have been eliminated from the design space if these metrics were applied to the leg 3 trajectories. In terms of considering phasing, the two-impulse Lambert solutions approximated the low-thrust final masses fairly well for most asteroid combinations and date ranges. Several exceptions were found, however, which would have also likely eliminated good solutions from the design space. As a general rule, the last asteroid was intercepted near perihelion (true anomaly ranging between 300° and 70°) for the best trajectories, although other feasible solutions rendezvoused with the last asteroid at other points in the trajectory. Because just the best set of solutions is sought after though, this rule could effectively be used to limit the range of departure dates that would have to be considered.

Another major challenge identified by the teams was how to choose the limit beyond which asteroid combinations would be eliminated for each metric. For example, how does one determine the inclination change above which asteroids are screened out of the design space? Or, how does one choose the range of true anomalies to consider for arrival at the last asteroid? One possibility that presented itself when analyzing the sub-problem results is to use a probabilistic approach to the pruning process. A distribution can be fit to each metric, and then this distribution can be used to eliminate the worst asteroid combinations. Eliminating asteroid combinations by fitting a distribution to each metric is not a foolproof method, however, as will be seen in the subsequent example. The decision must still be made at what percentile to begin eliminating asteroids, although this is a more intuitive decision than simply choosing an arbitrary inclination value.

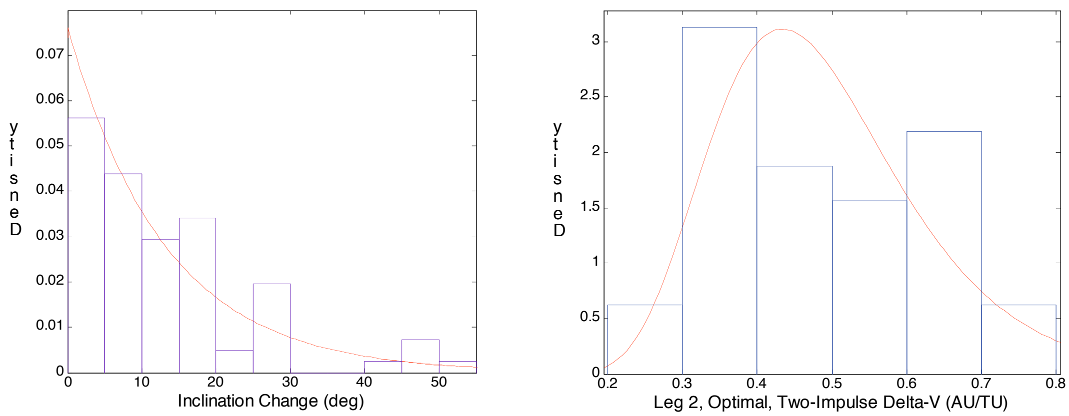

Figure 11: Probability density function of inclination change and optimal, phase-free, two-impulse delta-v for all asteroid combinations.

Figure 11 (left) plots the probability mass function for inclination change for all of the asteroid combinations analyzed (leg 1, 2, and 3). An exponential distribution can be fit to this data, as is also shown in the figure. If the top 10^{th} percentile were to be eliminated from this data set, all asteroid combinations with inclination changes greater than 30.2° would be removed from the design space. If the inclinations were to be considered separately for each leg, then inclination changes greater than 32.7°, 30.1°, and 29.5° would be eliminated for leg 1, leg 2, and leg 3, respectively. When all the inclinations are considered jointly, the 30.2° limit eliminates 90 asteroid combinations, none of which appear in the final set of feasible solutions. By taking the leg-by-leg approach, 92 asteroid combinations would be eliminated. One of the additional leg 3 combinations that would be eliminated, however, appears four times in the final set of feasible solutions. The largest final mass of this set is only 648.7 kg, ranking it 26^{th} out of 41 combinations, so although a feasible solution is being eliminated, it is not among the best solutions. Also plotted in Figure 11 is the distribution for the leg 2, optimal, phase-free, two-impulse transfer (N=1, time of flight = 1800 days). For this metric, each leg must be considered separately, since the delta-v values cannot be compared from leg to leg. For brevity, only the leg 2 example is shown. For the delta-v values, a lognormal distribution is more suited to fitting the data. If the top 10% of cases are again removed, all leg 2 asteroid combinations resulting in an optimal two-impulse delta-v greater than 0.68 AU/TU are eliminated. While this only eliminates three leg 2 combinations, none of these appear in the set of feasible solutions.

Another issue is that many of these metrics were considered individually during the pruning process. As a general rule, large inclination changes result in infeasible or bad solutions, as seen above. For the leg 3 trajectory, however, several asteroid combinations with large inclination changes yielded feasible solutions. In fact, the 5^{th} best overall solution has a leg 3 inclination change of 25°. While this asteroid combination has a small change in the ascending node (6.27°), the particular trajectory that yields the largest final mass also intercepts the last asteroid near its descending node. Therefore, there are a number of other factors that may not have been considered by the GTOC2 teams that contribute to the goodness of a particular asteroid combination. A large inclination change can yield a good solution because the change in ascending node is small, the change in semi-major axis is small, and the arrival asteroid is intercepted near the node. A final asteroid with a large semi-major axis can yield a good solution, as is the case with the Group 1 asteroid Pandarus, because the inclination change is small and the final asteroid is intercepted near periapsis.

Based on the application of many of the GTOC2 pruning techniques to the sub-problem, several recommendations can be made to improve the pruning process. As was done in solving the sub-problem, a leg-by-leg approach appears to be appropriate for the pruning phase as well. This type of approach has been used as a pruning technique for the global optimization of ballistic (high-thrust), multiple gravity assist trajectories[5]. First, feasible trajectories are identified over the launch date range to the first gravity assist planet. Because most of the search space is actually infeasible, a reduced number of trajectories can be analyzed for the subsequent leg, and so on. Similarly, pruning metrics can be used to first eliminate asteroids from the first leg of the trajectory. If just 2 of the 8 asteroids can be eliminated, this immediately reduces the number of possible combinations from 384 to 288, and the number of leg 2 combinations from 64 to 48. This process can then be continued for each leg, greatly reducing the total number of combinations that need to be analyzed in each step. Eliminating portions of the design space by fitting a distribution to the data provides a more objective method than simply choosing an upper bound based on visual inspection or perceived intuition. Certain metrics have been identified as well suited to pruning the design space, while others employed by many of the GTOC2 teams were not successfully applied to the sub-problem. Additional pruning metrics should also be taken into consideration and combined into an overall cost function, to better account for their interdependencies. While many teams attempted to consider inclination change along with ascending node change, there are even more factors that must be considered before definitively eliminating asteroid combinations from the design space.

CONCLUSION

After testing numerous design space pruning methods on a subset of the GTOC2 problem, it appears that teams were correct in their conclusion that they inadvertently eliminated some of the best solutions from the design space. It is also apparent that the pruning step is not as straightforward as many of the teams had hoped it would be, although it is still a necessary step when faced with a problem containing 41 billion possible combinations. Simply applying several metrics individually to the design space is likely to eliminate some of the best solutions, since the metrics being considered are highly coupled with regards to required propellant mass. As explained above, a large inclination change, for example, does not guarantee a bad solution, when other factors such as change in the ascending node or arrival conditions are favorable.

Although the original purpose of this work was to evaluate the effectiveness of the pruning methods employed by the GTOC2 participants, analyzing the sub-problem results led to other potential pruning techniques and methods not used by these teams. These results need to be analyzed further in order to determine if these other techniques could be applied effectively. Future work would use the sub-problem results to develop a more rigorous pruning methodology that incorporates a number of different metrics. This method would then have to be verified against another small problem with a known solution before being applied to the full GTOC2 problem.

REFERENCES

1. "Global Trajectory Optimisation Competition," European Space Agency – Advanced Concepts Team Mission Analysis, http://www.esa.int/gsp/ACT/mad/op/GTOC/index.htm [accessed 8 August 2007].
2. Sims, J.A., Finlayson, P.A., Rinderle, E.A., Vavrina, M.A., Kowalkowski, T.D., "Implementation of a Low-Thrust Trajectory Optimization Algorithm for Preliminary Design", AIAA 2006-6746, AIAA/AAS Astrodynamics Specialist Conference, 21-24 August 2006, Keystone, CO.

3. Sims, J.A., Flanagan, S.N., "Preliminary Design of Low-Thrust Interplanetary Mission," AAS 99-338, AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska, 16-19 August 1999.
4. Shen, H., Tsiotras, P., "Using Battin's Method to Obtain Multiple-Revolution Lambert's Solutions," AAS 03-568, *Advances in the Astronautical Sciences*, Vol. 116, Part II, pp. 1067-1084.
5. Becerra, V.M., Myatt, D.R., Nasuto, S.J., Bishop, J.M., Izzo D., "An Efficient Pruning Technique for the Global Optimization of Multiple Gravity Assist Trajectories," International Workshop on Global Optimization, Almeria, Spain, 18-22 September, 2005.