Comparison of Multi-Objective Genetic Algorithms in Optimizing Q-Law Low-Thrust Orbit Transfers

Seungwon Lee, Paul von Allmen, Wolfgang Fink, Anastassios E. Petropoulos, and Richard J. Terrile
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive M/S 169-315, Pasadena, CA, USA
Seungwon.Lee@jpl.nasa.gov

ABSTRACT
Multi-objective genetic algorithms (MOGA) are used to optimize a low-thrust spacecraft control law for orbit transfers around a central body. A Lyapunov feedback control law called the Q-law is used to create a feasible orbit transfer. Then, the parameters in the Q-law are optimized with MOGAs. The optimization goal is to minimize both the flight time and the consumed propellant mass of the trajectory created by the Q-law, and consequently to find Pareto-optimal trajectories. To improve the qualities of the obtained Pareto-optimal solutions, elitism and a diversity-preservation mechanism are incorporated into MOGA. The MOGA performance with and without the new mechanisms are systematically compared and evaluated with quantitative metrics. The new mechanisms significantly improve the convergence and distribution of the resulting Pareto front for the low-thrust orbit-transfer optimization problem. The new mechanisms also improve the statistical stability and the computational efficiency of the algorithm performance.

Categories and Subject Descriptors
Real-World Applications

Keywords
Multi-objective genetic algorithms, Low-thrust orbit transfer, Q-law, Non-dominated sorting genetic algorithm

1. INTRODUCTION
Many real-world optimization problems involve multiple competing objectives, which give rise to a set of compromising solutions rather than a single optimal solution. Spacecraft trajectory design is such a multi-objective optimization problem. Optimal trajectories are the ones that minimize both flight time and propellant consumption. Reduction of flight-time often competes with propellant saving. The competition leads to a trade-off between the two resources. The reasonable estimation of the trade-off is the key in spacecraft trajectory design.

The trajectories of interest here are orbit transfers around a central body, where the spacecraft uses a low-thrust propulsion system. Low-thrust propulsion is promising for future deep-space missions, because despite its low thrust levels, the momentum transfer per kilogram of propellant is ten to twenty times greater than that for chemical propulsion. In fact, NASA's future space missions Dawn and JIMO will use electric propulsion for inter-planetary cruise and orbital operations. However, the control of low-thrust spacecraft poses a challenging design problem, particularly for orbit transfers around a central body because it involves a large number of revolutions and thrust arcs along these revolutions.

In trajectory optimization, there are two main approaches: One is to develop a thrust control law that guides a trajectory design [2,6-8,10-12], which we call "guided search". The other is to optimize trajectories without a thrust control law [1,5,14,15], which we call "unguided search". The guided search restricts the search space, but simplifies the optimization problem such that it can efficiently estimate optimal trajectories. In contrast, the unguided search explores a wider search space and has a potential to find a better solution than the guided search, but is computationally more demanding.

Here, we take the guided search approach to optimize low-thrust spacecraft orbit transfers. Among the several thrust-control laws available in the literature [2,6-8,10-12], one promising control law is the Q-law [10-12], which is based on Lyapunov feedback control. The Q-law involves a set of control parameters that are left free for the mission designer to manipulate. The Q-law, with default values for the control parameters, provides reasonable estimates of Pareto-optimal solutions, indicating that a suitable Lyapunov function has been found and that optimization of the control parameters should yield near-Pareto-optimal solutions.

Previously [9], we have performed the optimization of the Q-law control parameters with a multi-objective genetic algorithm based on the non-dominated sorting genetic algorithm (NSGA) [13]. We have demonstrated that NSGA efficiently finds optimal Q-law control parameters for a wide variety of orbit transfers [9], and yields Pareto-optimal orbit transfers that are comparable to those obtained with the unguided search [5,14,15]. However, the previous approach leads to Pareto-optimal solutions that are clustered toward the short flight time zone. A wider and uniform spread of Pareto-optimal solutions is desired. Recently, Deb et al. have introduced elitism and diversity preservation mechanisms to NSGA to improve the algorithm performance [3,4]. The revised algorithm, named NSGA-II, has been applied to many theoretical and real-world problems, and is shown to be more efficient than its predecessor [3,4]. Hence, the two new mechanisms are included in our optimization process. The optimization performance with and without the new

GECCO'05, June 25–29, 2005, Washington, DC, USA.

mechanisms are systematically compared and evaluated in order to identify the effect of the new mechanisms.

2. PROBLEM

2.1 Q-law

The Q-law is a Lyapunov feedback control law developed by Petropoulos to provide good initial estimates for optimal orbit transfers between two arbitrary orbits around a central body [10-12]. The Q-law determines when and at what angles to thrust, based on the proximity quotient termed Q. The function Q judiciously quantifies the proximity of the osculating orbit to the target orbit. In the Q-law, the central body is modeled as a point mass, and no perturbing forces are considered.

The Q-law has up to 15 free control parameters, which mission designers can adjust. The control parameters affect the effectivity of thrust usage and the geometry (gradient, maxima, minima, saddle points, etc.) of the proximity quotient Q. A different effectivity of thrust usage leads to a different length or location of a thrust arc, and a different geometry of Q leads to a different thrust angle or shifts thrust-arc location. Hence, the mission designer can acquire a different trajectory for a different set of the Q-law control parameters. For a detailed discussion of the mechanisms of the Q-law, readers are referred to Refs. [10-12].

2.2 Q-law Optimization

For a given set of the control parameters, the output of the Q-law is a series of thrust arcs and thrust angles together with the resulting required flight time and propellant mass. The desired outcome for the mission designer is knowledge of the trade-off between flight time and propellant mass, and the Pareto-optimal trajectory corresponding to each point on the trade-off curve. Therefore, our optimization problem is to minimize the competing objectives of flight time and required propellant mass while varying the Q-law control parameters.

Table I. Initial and target orbit elements, spacecraft thrust, spacecraft specific impulse, spacecraft initial mass, and central body of the orbit transfer considered. The listed orbit elements are semimajor axis (a), eccentricity (e), inclination (i), the argument of periapsis (ω), the longitude of ascending node (Ω), and true anomaly (f). Note that the target values for inclination, argument of periapsis, longitude of ascending node, and true anomaly are left free.

Orbit	a (km)	e	i (deg.)	ω (deg.)	Ω (deg.)	f (deg.)
Initial	24506	0.2	0.6	0.0	0.0	0.0
Target	26500	0.7	Free	Free	Free	Free

Thrust	Specific Impulse (s)	Initial Mass (kg)	Central Body
9.3	3100	300	Earth

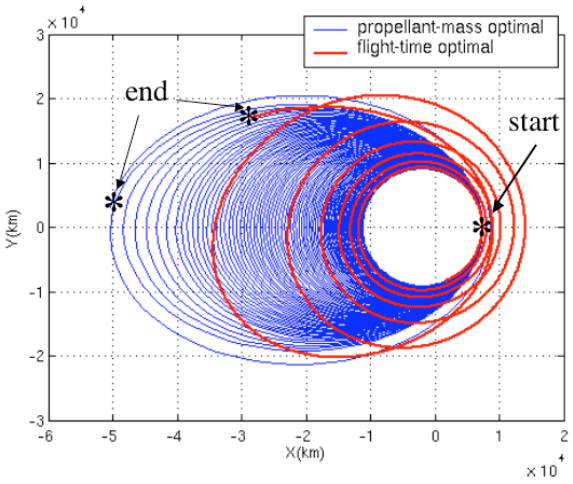

Figure 1. Flight-time optimal and propellant optimal trajectories for the orbit transfer around a central body Low-Thrust Orbit-Transfer. The trajectories end at different locations because the final arguments of periapsis and the final true anomalies are different.

2.3 Low-Thrust Orbit Transfer

The Q-law optimization is applied to an orbit transfer from a low-eccentricity elliptic orbit to a coplanar, high-eccentricity, larger elliptic orbit. The maximum-permitted flight time is 20 days. The details of the orbit transfer are listed in Table I. Typical flight-time-optimal and propellant-optimal trajectories for the orbit transfers are shown in Figure 1. The flight-time-optimal trajectory involves a small number of revolutions because the spacecraft thrusts most of the time. In contrast, the propellant-optimal trajectory involves a large number of revolutions due to long coasting (no-thrust) arcs. Seven Q-law parameters are relevant to this orbit transfer: $\left\{W_e, \eta_{cut}^a, \eta_{cut}^\phi, \varphi_{\min}, m, n, r\right\}$.

3. METHODS

The optimization of the Q-law control parameters is performed with MOGAs. The algorithm parameter setting, algorithm core mechanisms, and performance metrics used for our optimization process are described in the following three sub-sections.

3.1 MOGA Parameters

For all MOGA optimization runs, the following common parameters and operators are used. Each of the 7 Q-law control parameters is represented by a real-valued gene. The population size is set to 100. Parents are selected through binary tournament. Offspring are created with the simulated binary crossover operator and the polynomial mutation operator [3]. The distribution parameters associated with the simulated binary crossover and the polynomial mutation operators are η_c=2 and η_m=100. The crossover probability is 0.8 and the mutation probability per gene is 1/7. After 200 generations, the genetic evolution is terminated. Table II summarizes the genetic algorithm parameters and operators used for our optimization problem.

Table II: Genetic Algorithm Parameter Setting.

Algorithm Parameter	Value
Population Size	100
Generations	200
Number of Parameters	7
Selection Operator	Binary Tournament
Crossover Operator	Simulated Binary with η_c=2
Mutation Operator	Polynomial with η_m=100
Crossover Probability	0.8
Mutation Probability	1/7 per gene

3.2 Core Mechanisms

The core mechanisms used to handle our multi-objective optimization problems are based on NSGA and NSGA-II. NSGA-II [4,13] is an enhanced version of its predecessor NSGA and involves two new mechanisms in comparison to NSGA. Both NSGA and NSGA-II use nondominated sorting to determine preliminary fitness values. The three core mechanisms introduced in NSGA and NSGA-II are briefly summarized below. More details can be found in Refs. [3,4].

3.2.1 Nondominated Sorting

Non-dominated sorting proceeds as follows. First, non-dominated individuals in the current population are identified. The non-dominated individuals are those who are not inferior to any other individuals in the population with respect to every objective. The same fitness value is assigned to all the non-dominated individuals. The individuals are then ignored temporarily, and the rest of the population is processed in the same way to identify a new set of non-dominated individuals. A fitness value that is smaller than the previous one is assigned to all the individuals belonging to the second non-dominated front. This process continues until the whole population is classified into non-dominated fronts with different fitness values.

3.2.2 Elitism

Unlike its predecessors, NSGA-II allows the parents to compete with offspring. In each generation, an offspring population of size N is generated from a parent population of the same size. The two populations are combined, and the first N best-fit individuals from the combined population are chosen to be part of the next generation population. The main purpose of this mechanism is to prevent fit individuals found in earlier generations from being lost easily.

3.2.3 Diversity-Preservation

The original NSGA uses the well-known fitness-sharing approach to preserve the diversity among the Pareto-optimal solutions. Although the fitness-sharing approach is found to maintain diversity in a population, the performance largely depends on its associated parameter. To avoid this sensitivity, Deb et al., have introduced a "crowding distance comparison" approach. First, the crowding distance surrounding a particular solution is measured. The crowding distance is given by the perimeter of the cuboid formed by using the nearest neighbors in the same non-dominated front as the vertices. Second, the crowding distance is used to break a tie when two solutions have the same fitness, i.e. they belong to the same nondominated front. A solution with a higher crowding distance becomes a winner. By preferring the solution with a higher crowding distance, this mechanism encourages population diversity.

3.3 Performance Metrics

The performance assessment of multi-objective optimizers should take at least the following two criteria into account: i) the distance of the obtained solutions to the Pareto front, ii) the distribution of the obtained solutions. Various performance metrics to measure these aspects have been introduced in the literature [3,4,13,16]. Some of the metrics require the knowledge of the "true" Pareto-optimal solutions, which are largely unknown in our optimization problem. Taking this limitation into account, we chose the following three measures: 1) size of the dominated space, 2) coverage of two Pareto fronts, and 3) non-uniformity of the Pareto front. The first two metrics measure the convergence of the Pareto-optimal solutions, while the last metric measures the distribution of the Pareto-optimal solutions.

3.3.1 Size of the dominated space

The size of the dominated space ***S****(A)* is a measure of how much of the objective space is weakly dominated by a given non-dominated set A. As an example, the size of the dominated space is illustrated in Figure 2. Since our optimization involves the minimization of two objectives, a reasonable maximum value for each objective (maxI and maxII) is chosen to determine the size of the dominated space. Higher values of ***S****(A)* indicate better performance.

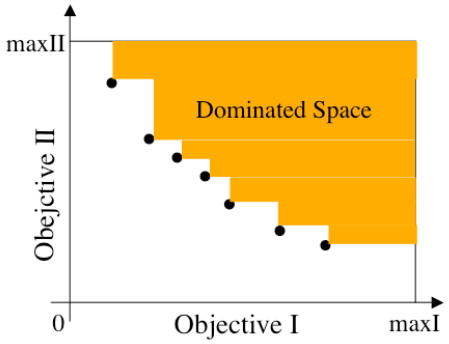

Figure 2. Space dominated (colored in orange) by a given Pareto set when two objectives are minimized.

3.3.2 Coverage of two Pareto fronts

This measure compares two Pareto-optimal sets to each other. When two Pareto-optimal sets *A* and *B* are given, the coverage ***C****(A,B)* of the two Pareto fronts maps the ordered pair (*A,B*) to the interval [0, 1]:

$$C(A,B) = \frac{\left|\{b \in B \mid \exists a \in A : a \succ b\}\right|}{|B|}, \qquad (1)$$

where $|B|$ means the number of solutions in the set *B*, and $a \succ b$ means that solution *a* dominates solution *b*, i.e., the objectives of *a* are both less than those of *b*. Therefore, ***C****(A,B)* gives the fraction of *B* dominated by *A*. For example, ***C****(A,B)*=1

means that all individuals in B are dominated by A. The opposite $\boldsymbol{C}(A,B)$=0 represents the situation that no individual in B is dominated by A. Note that $\boldsymbol{C}(A,B)$ is not necessarily equal to 1-$\boldsymbol{C}(B,A)$. If $\boldsymbol{C}(A,B) > \boldsymbol{C}(B,A)$, this means that the set A has better solutions than the set B.

3.3.3 Non-uniformity of Pareto front

To measure the non-uniformity of the distribution of a Pareto front, we introduce the quantity $\boldsymbol{D}(A)$ given by the distribution of the Euclidian distance (d_i) between two consecutive points (solutions) along the Pareto front:

$$D(A) = \sqrt{\frac{\sum_i \left(d_i/\bar{d} - 1\right)^2}{|A| - 1}}. \qquad (2)$$

This quantity is a standard deviation of the distances normalized by the average distance $\bar{d}$. The Euclidian distance depends on the scaling given to each of the objectives. Here, the propellant mass objective is measured in kilograms, and the flight time objectives in days. When $\boldsymbol{D}(A)$=0, the spacing in the Pareto front is uniform. The higher the value of $\boldsymbol{D}(A)$, the more non-uniform the spacing in the Pareto front. Therefore, a lower value of $\boldsymbol{D}(A)$ is desired.

4. RESULTS

We assess the improvement of MOGA performance due to the elitism and diversity-preservation mechanisms by comparing the following two approaches: 1) MOGA without the new mechanisms (MOGA-I), 2) MOGA with the mechanism (MOGA-II). Both MOGA-I and MOGA-II use non-dominated sorting to assign preliminary fitness values. MOGA-I uses the preliminary fitness values as the final ones, while MOGA-II uses the crowded distance comparison to break a tie when two solutions have the same preliminary fitness value, (i.e., they belong to the same nondominated front). Note that MOGA-I is the algorithm used in our previous studies [9].

Twenty independent runs with different random seeds are performed with MOGA-I and MOGA-II. MOGA-I finds about 35 Pareto-optimal solutions out of total 100 individuals. In contrast, MOGA-II leads to a population of which all the individuals are Pareto-optimal. The mean and standard deviation of the number of Pareto-optimal solutions found by MOGA-I and MOGA-II are listed in Table II. MOGA-II consistently finds more Pareto-optimal solutions than MOGA-I. The quality of the Pareto-optimal solutions obtained with the two algorithms is measured by three qualitative metrics $\boldsymbol{S}(A)$, $\boldsymbol{C}(A,B)$, and $\boldsymbol{D}(A)$.

Table III shows the mean and variance of $\boldsymbol{S}$(A) of MOGA-I and MOGA-II. $\boldsymbol{S}(A)$ is measured in kilogram-days, with maxima in the objectives taken as 40 kg and 20 days, respectively. MOGA-II leads to a higher value $\boldsymbol{S}(A)$, indicating that better solutions are obtained with MOGA-II than with MOGA-I. The difference between the two means is twice as big as the standard deviation of MOGA-I. Furthermore, the standard deviation of MOGA-II is significantly smaller than that of MOGA-I, indicating that the performance of MOGA-II is statistically more stable.

Table IV lists the mean and standard deviation of $\boldsymbol{C}(A,B)$ for MOGA-I and MOGA-II. The value $\boldsymbol{C}$(MOGA-II, MOGA-I)= 0.812 means that 81.2% of the Pareto-optimal solutions obtained with MOGA-I are dominated by the solutions with MOGA-II. Similarly, the value $\boldsymbol{C}$(MOGA-I, MOGA-II)=0.021 represents that only 2.1% of the solutions obtained with MOGA-II are dominated by those with MOGA-I.

The distributions of the solutions obtained with MOGA-I and MOGA-II are evaluated with non-uniformity metric $\boldsymbol{D}(A)$. Table V shows the mean and variance of $\boldsymbol{D}(A)$ of the MOGA-I and MOGA-II runs. A lower value of $\boldsymbol{D}(A)$ means a better/uniform spread of solutions. The comparison of $\boldsymbol{D}(A)$ shows that the spread of the solutions found with MOGA-II is more uniform than with MOGA-I.

All three metric measurements show that MOGA-II clearly outperforms MOGA-I. More specifically, metrics $\boldsymbol{S}(A)$ and $\boldsymbol{C}(A,B)$ measurements show that MOGA-II leads to a better convergence of the Pareto front, and $\boldsymbol{D}(A)$ measurements show that MOGA-II leads to a more uniform distribution of the Pareto front. To illustrate the difference between the Pareto fronts obtained with MOGA-I and MOGA-II, the Pareto fronts obtained with one of the twenty runs of MOGA-I and MOGA-II are plotted in Figure 3.

Table II. Mean and standard deviation of the number of Pareto-optimal solutions obtained with twenty independent runs of MOGA-I and MOGA-II. MOGA-II consistently finds more Pareto-optimal solutions than MOGA-I.

Algorithm	Mean	Standard Deviation
MOGA-I	34.70	3.38
MOGA-II	100	0.00

Table III. Mean and standard deviation of metric $\boldsymbol{S}(A)$, in units of kg*day, of Pareto fronts obtained with twenty independent runs of MOGA-I and MOGA-II. A higher value means a better convergence of the Pareto front.

Algorithm	Mean	Standard Deviation
MOGA-I	399.70	2.50
MOGA-II	405.51	0.35

Table IV. Mean and standard deviation of metric $\boldsymbol{C}(A,B)$ for a pair of Pareto fronts obtained with twenty independent runs of MOGA-I and MOGA-II.

A	B	Mean	Standard Deviation
MOGA-II	MOGA-I	0.812	0.083
MOGA-I	MOGA-II	0.021	0.021

Table V. Mean and standard deviation of the non-uniformity metric $\boldsymbol{D}(A)$ of Pareto fronts obtained with twenty independent runs of MOGA-I and MOGA-II. A lower value means a more uniform spread of the Pareto front.

Algorithm	Mean	Standard Deviation
MOGA-I	0.279	0.071
MOGA-II	0.056	0.007

In addition to the performance at the end of the evolution, the performance of MOGA-I and MOGA-II during the evolution is

also monitored. Figure 4 shows the evolution of metrics ***S****(A)* and ***D***(A) with respect to generation number. MOGA-II shows a faster convergence and a more stable performance than MOGA-I. These results show that MOGA-II not only finds better solutions, but also finds them faster and more reliably. For smooth evolutions like those obtained with MOGA-II, the rates at which the dominated space size metric ***S****(A)* and the non-uniformity metric ***D****(A)* constitute excellent criteria for the optimization termination condition.

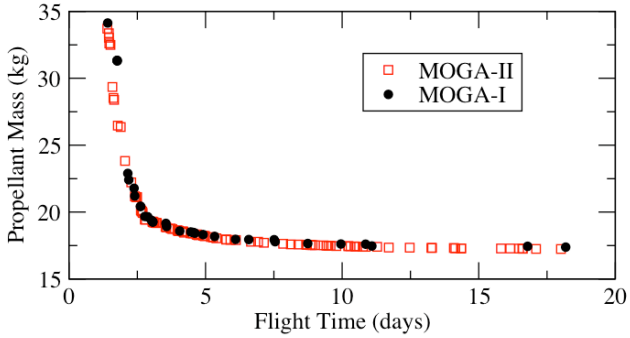

Figure 3. Pareto-optimal solutions obtained with MOGA-I and MOGA-II. The solutions of MOGA-II show a better convergence and distribution of the Pareto front.

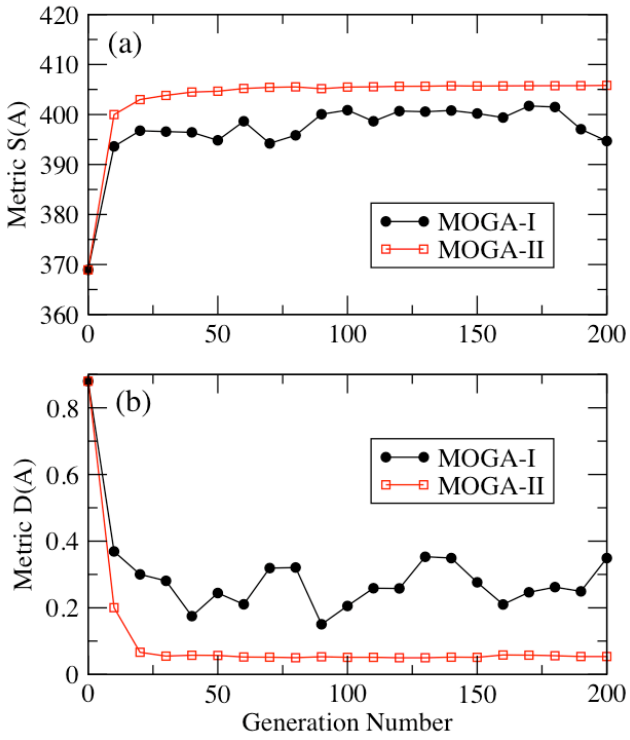

Figure 4. Evolution of metrics *S(A) and D(A)* with respect to generation number for MOGA-I and MOGA-II. MOGA-II shows more efficient performance.

5. CONCLUSIONS

We have applied multi-objective genetic algorithms to the optimization of a low-thrust control law for orbit transfers around the Earth. A Lyapunov feedback control law called the Q-law is used to create an eligible orbit transfer, while the Q-law control parameters are selected with the multi-objective algorithms. Two resources, flight time and propellant mass, are minimized and a trade-off between the two resources is obtained. Elitism and a diversity-preservation mechanism are incorporated into MOGA to improve the optimization performance. The MOGA performance with and without the new mechanisms is systematically compared and evaluated in order to identify the contribution of the new mechanisms. Three quantitative performance metrics were used: 1) size of the dominated space, 2) coverage of two Pareto fronts, and 3) non-uniformity of the Pareto front. The first two metrics measure the convergence of the Pareto front, and the last metric measures the distribution of the Pareto front. We find that the elitism and diversity-preservation mechanisms significantly improve the convergence and spread of the resulting Pareto-optimal solutions. The performance difference between MOGA with and without the new mechanisms is as large as one order of magnitude in terms of the three quantitative metrics. The performance difference also appears in the variation of the metrics with respect to generation number. The new mechanisms lead to a better Pareto front more efficiently and reliably.

6. ACKNOWLEDGMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration. The research was supported by the JPL Research and Technology Development Program.

7. REFERENCES

[1] Betts, J.T., "Survey of Numerical Methods for Trajectory Optimization", *J. Guidance, Control, and Dynamics,* 21(2), 193-207, 1998.

[2] Chang, D. E., Chichka, D. F., Marsden, J. E., "Lyapunov functions for elliptic orbit transfer," *AAS/AIAA Astrodynamics Specialist Conference,* AAS Paper 01-441, July 2001.

[3] Deb, K., "*Multi-Objective Optimization Using Evolutionary Algorithms"* John Wiley & Sons, 2001.

[4] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., "A fast elitist multiobjective genetic algorithm: NSGA-II", *IEEE Transactions on Evolutionary Computation,* 6(2), 182-197, 2002.

[5] Geffroy, S. and Epenoy, R., "Optimal Low-Thrust Transfers with Constraints – Generalization of Averaging Techniques," *Astronautica Acta*, 41, 133-149, 1997.

[6] Gefert, L. P. and Hack, K. J., "Low-Thrust Control Law Development for Transfer from Low Earth Orbits to High Energy Elliptical Parking Orbits," *AAS/AIAA Astrodynamics Specialist Conference,* AAS Paper 99-410, Aug. 1999.

[7] Ilgen, M. R., "Low thrust OTV guidance using Liapunov optimal feedback control techniques," *AAS/AIAA Astrodynamics Specialist Conference,* AAS Paper 93-680, Aug. 1993.

[8] Kluever, C. A., "Simple Guidance Scheme for Low-Thrust Orbit Transfers," *J. Guidance, Control, and Dynamics*, 21(6), 1015-1017, 1998.

[9] Lee, S., von Allmen, P., Fink W., Petropoulos A. E., and Terrile, R. J., "Design and Optimization of Low-Thrust Orbit

Transfers", *IEEE Aerospace Conference Proceedings*, Mar. 2005.

[10] Petropoulos, A. E., "Simple Control Laws for Low-Thrust Orbit Transfers," *AAS/AIAA Astrodynamics Specialist Conference,* AAS Paper 03-630, Aug. 2003.

[11] Petropoulos, A. E., "Low-Thrust Orbit Transfers Using Candidate Lyapunov Functions with a Mechanism for Coasting," *AAS/AIAA Astrodynamics Specialist Conference,* AAS Paper 04-5089, Aug. 2004.

[12] Petropoulos, A.E., "Refinements to the Q-law for Low-Thrust Orbit Transfers", *AAS/AISS Space Flight Mechanics Meeting*, AAS Paper 05-162, Jan. 2005.

[13] Srinivas, N. and Deb, K., "Multiobjective optimization using nondominated sorting in genetic algorithms," *Evolutionary Computation Journal*, 2(3), 221-248, 1994.

[14] Whiffen, G.J. and Sims, J.A., "Application of a Novel Optimal Control Algorithm to Low-Thrust Trajectory Optimization," *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 01-209, 2001.

[15] Whiffen, G.J., "Optimal Low-Thrust Orbit Transfers around a Rotating Non-Spherical Body," *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 04-264, Maui, Hawaii, Feb.8-12, 2004.

[16] Zitzler, E. and Thiele, L., "Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach," *IEEE Transactions on Evolutionary Computation*, 3(4), 257-271, 1999.