AIAA-98-4381

OPTIMAL EARTH ORBITING SATELLITE CONSTELLATIONS VIA A PARETO GENETIC ALGORITHM

William J. Mason*
Victoria Coverstone-Carroll†
John W. Hartmann‡

Department of Aeronautical and Astronautical Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801-2935

Abstract

A Pareto Genetic Algorithm is applied to the design of satellite constellations. The Pareto Genetic Algorithm is a multi-objective optimizer that generates groups, or "families", of optimal solutions. The technique developed in this paper combines the Pareto Genetic Algorithm with Satellite Tool Kit, a commercially available software package that performs satellite performance analyses, to produce a new powerful tool for designing optimal satellite constellations. Several constellations for continuous global coverage are presented.

Introduction

Satellite constellations are becoming more popular choices to achieve mission goals as modern systems are designed for increased capabilities and complex tasks. Traditional design considerations have focused primarily on circular, symmetric configurations.[1,2,3] Such constellation patterns allow for easy analysis and are well understood. In more recent years additional work has been done in non-circular orbit constellations.[4,5] Most of the previous work has focused on providing global coverage and utilized geometric-based techniques; however as the satellite market changes, different criteria emerge as the basis for constellation coverage. Cost, localized coverage, and launch vehicle availability are all issues that must be considered when implementing a constellation. Often times a mission designer would benefit from a group of feasible solutions, with differing costs and performance levels. Such a group of solutions would allow for a more informed design decision, as well as contingency scenarios in case of operational changes, i.e. loss of spacecraft.

This paper addresses the design of optimal satellite constellations. First, an overview of genetic algorithms is presented and is followed by a brief discussion of multi-objective optimization. The genetic algorithm developed to perform the multi-objective optimization is a Pareto Genetic Algorithm termed the Modified Illinois Non-dominated Sorting Genetic Algorithm (MINSGA). The MINSGA's functionality is demonstrated on several test problems. Lastly, the MINSGA is combined with Satellite Tool Kit and used to produce several constellations that yield continuous global coverage.

Genetic Algorithms

Genetic algorithms (GAs) are an optimization technique based on natural selection and Darwinian genetics. These algorithms differ from the more traditional optimization techniques in that they work with a coding of the parameter set, instead of the parameters themselves. This coding is usually done as a binary string, as was done in this study. Genetic algorithms are population based, working with many candidate solutions rather than trying to progress from a single point in the search space. They use only the value of the objective function, termed the "fitness", instead of derivatives or other auxiliary knowledge, and they use probabilistic transition rules instead of deterministic rules. A population of individual solutions is evolved over a series of generational cycles, each undergoing alterations to their respective parameter set, resulting in improvement in their respective fitness value.

The GA used in this paper employs the standard operators of selection, crossover, and mutation to perform the evolutionary search. Selection is the process of choosing two solutions on which crossover of information may occur. The better an individual's fitness (higher for maximization, lower for minimization), the greater the chance that it will be selected from the population for potential crossover. Crossover is used to create new solution strings or children from existing "parent" strings. A string is a

---

* Graduate Research Assistant, Member AIAA
† Assistant Professor, Senior Member AIAA, Member AAS
‡ Graduate Research Assistant, Member AIAA

binary representation of all the independent parameters. Through crossover, beneficial information encoded in the parents has a chance to propagate to future solutions. The role of mutation is mainly that of protecting the population of solutions against irrevocable losses of potentially helpful information that may be contained in individuals with poor fitnesses. Mutation also aids in providing diversity within the population.

"Niching" or "fitness sharing" is a supplemental strategy, which is also used to promote diversity in the solution space. In this procedure, solutions are penalized according to their degree of similarity to others in the population in an attempt to distribute members over large portions of the search space and to avoid premature convergence to a single solution. The population is examined on a per individual basis to determine the number of other members that reside within the neighborhood of the current examinee. The size of the neighborhood is determined by the niching control parameter, σ_{share}. A penalty value is associated with each individual residing within this neighborhood. The magnitude of the penalty is governed by a sharing function whose value decreases by the power α as proximity to the examinee decreases. An α value of 1 corresponds to a linear reduction in magnitude, with the effective penalty approaching zero as individual proximity approaches the neighborhood's boundary. The number of individuals residing within this defined neighborhood and their corresponding penalties are tallied, and the current examinee's fitness is penalized by an amount proportional to that number at the end of each generational cycle. For a more complete discussion of genetic algorithms and niching, see Reference [6].

Pareto Optimization

Pareto optimization can be stated as the problem of determining an optimal solution based on multiple, possibly competing criteria. Reference [7] defines Pareto optimality as;

> The solution to a multi-objective problem is, as a rule, not a particular value, but a set of values of decision variables such that, for each element in this set, none of the objective functions can be further increased without a decrease of some of the remaining objective functions (every such value of a decision variable is referred to as PARETO-OPTIMAL).

Pareto optimal solutions have also been termed non-dominated. This name arises from the fact that no other solution is superior to them in all objectives. In other words, the non-dominated solutions are those solutions that cannot be improved in all objectives simultaneously. A more thorough discussion of Pareto optimality may be found in *Genetic Algorithms*.[6]

Modified Illinois Non-dominated Sorting Genetic Algorithm (MINSGA)

The Pareto GA is based on the concept of non-dominated sorting that was originally proposed by Goldberg and developed by Srinivas and Deb as the Non-dominated Sorting Genetic Algorithm (NSGA).[6,8] The MINSGA replaces the stochastic remainder selection used by Srinivas and Deb with stochastic universal selection. The other major difference is the use of a normalized parameter set in the sharing scheme. The MINSGA was originally developed as a tool to determine optimal low-thrust trajectories.[9,10]

The NSGA, and MINSGA, are not completely new algorithms, rather variations on the fitness evaluation technique used in a traditional GA. The MINSGA uses the concept of non-dominance to sort through a population of possible solutions, assigning each member to a Pareto optimal "front" according to their level of non-dominance. The process consists of two iterated steps. For the first iteration through the sorting routine, each objective is evaluated for every solution. The population is then sorted into dominated and non-dominated solutions. Again, a non-dominated solution is one where no other solution in the population is categorically better in all objectives. The non-dominated solutions are assigned to the first front and given fitnesses equal to the population size. After assignment to a front, niching is applied to these individuals where the fitness is adjusted according to the solutions' proximity to one another in an attempt to evenly distribute individuals across the Pareto front.

After niching is complete, the minimal fitness in the current front is determined. This fitness is then slightly reduced and used as the initial fitness for the next front in the following iteration of the sorting routine. During this iteration and all subsequent ones, individuals previously tagged as non-dominated and assigned to a specific front are excluded from examination and the remaining sub-population is examined to determine the next subgroup of non-dominated solutions. The non-dominated sorting process continues this iterative evaluation of sub-populations until all individuals in the population are assigned to a specific front and given a fitness value.

The benefits of incorporating a Pareto search algorithm in the optimization process are twofold:

i.) Elimination of the problems encountered in classical multi-objective optimization methodologies such as objective conflict.

ii.) Development of a Pareto optimal front of solutions, providing an array of compromise solutions.

When applied to the population based genetic algorithm, these concepts enable automatic generation of Pareto optimal solutions. In the context of satellite constellation design discussed in this study, a Pareto genetic algorithm provides families of solutions illustrating the trades between specific objectives.

Verification of the MINSGA

A set of trial functions comprised from those found in the literature as well as some of original design were used to test the MINSGA. Three functions from the Srinivas and Deb test suite were used along with three new functions that were designed to test some of the more advanced and previously undemonstrated abilities of the modified NSGA. The values for main control parameters were the same as those used in the original NSGA study to allow for meaningful comparison.[8]

Test Case F1

The first test was a simple case to determine if the modified NSGA can identify the Pareto optimal region located between the minima of two parabolas. The problem can be expressed as,

Minimize:
$$f_{11} = x^2 \qquad f_{12} = (x-2)^2 \tag{1}$$

The Pareto optimal region lies in the range, x ∈ [0,2].

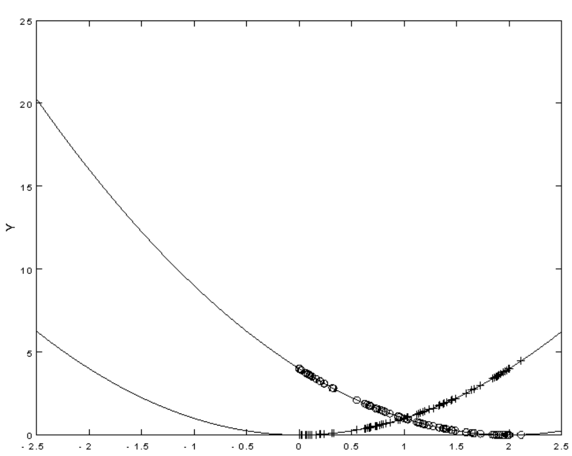

Figure 1: Population at generation 100 for F1.

Figure 1 shows the population plotted for the two functions at generation 100. The population maintained its wide distribution through 500 generations matching the original NSGA's performance.[8] Figure 2 shows the distribution of the population in the objective space. These results show that the MINSGA was able to identify the Pareto optimal region.

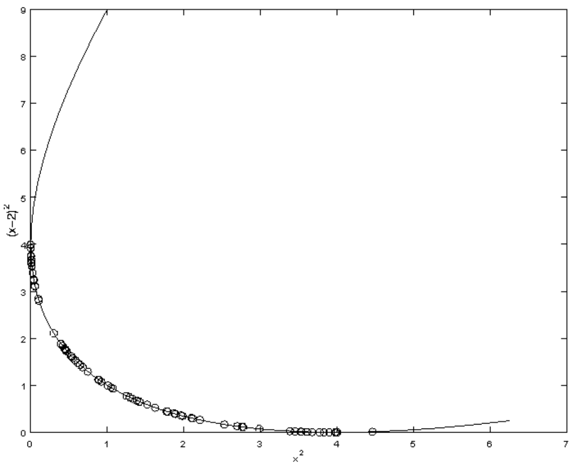

Figure 2: Pareto region at generation 100 for F1.

Test Case F2

The second test case was the same as the first but with a single, linear constraint added. This demonstrated the MINSGA's ability to deal with constraints. The results are not included in this paper but can be found in Reference [11] and compared favorably to those demonstrated in Reference [8].

Test Case F3

The third test case demonstrates that the MINSGA can optimize two objectives in two dimensions with multiple soft constraints. The mathematical representation of the problem is,

Minimize:
$$f_{31} = (x-2)^2 + (y-1)^2 + 2 \qquad f_{32} = 9x - (y-1)^2 \tag{2}$$

Constrained by:
$$x^2 + y^2 \leq 0 \qquad x - 3y + 10 \leq 0 \tag{3}$$

It can be shown analytically that the Pareto optimal region lies along the line $x = -2.5$ in the area above the second, linear constraint and inside the first, circular constraint. As can be seen in Figure 3, the MINSGA correctly identified this region.

Test Case F4

The fourth test case was designed to look at extending the use of a Pareto GA to problems with more than two objectives. To the authors' knowledge, no published results have appeared describing work in this direction, therefore this test was designed to determine the MINSGA's capability to handle such problems. As can be seen in Figure 4, this case consisted of four parabolic functions, each revolved about its respective minimum. The parabolic equations are given in Equation (4).

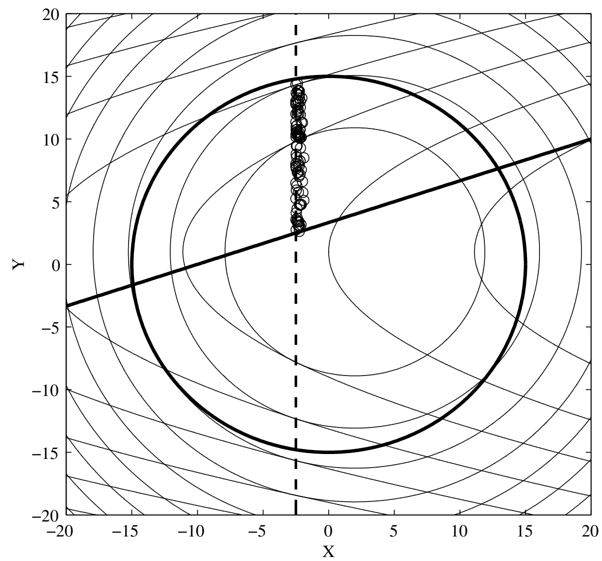

Figure 3: Population at generation 100 for F3.

The Pareto optimal region lies within the square with one of the parabolic minima at each of its corners, i.e. the square defined by the four points (2,2), (-2,2), (-2, -2), and (2, -2).

$$\text{Minimize:}\quad \begin{aligned} f_{41} &= (x-2)^2 + (y-2)^2 \\ f_{42} &= (x+2)^2 + (y\text{-}2)^2 \\ f_{43} &= (x+2)^2 + (y+2)^2 \\ f_{44} &= (x\text{-}2)^2 + (y+2)^2 \end{aligned} \tag{4}$$

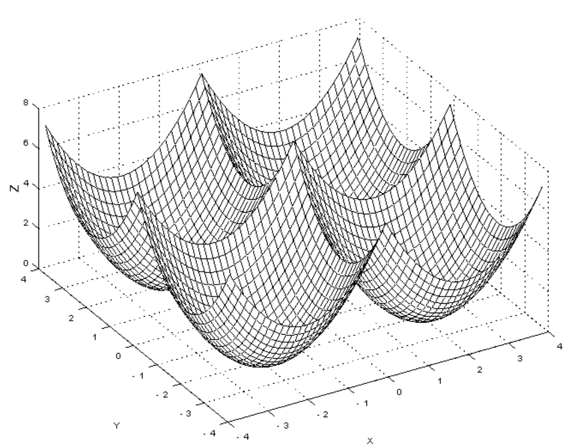

Figure 4: Test case with 4 objectives.

Figure 5 shows that the MINSGA can successfully identify an optimal region for more than two objectives.

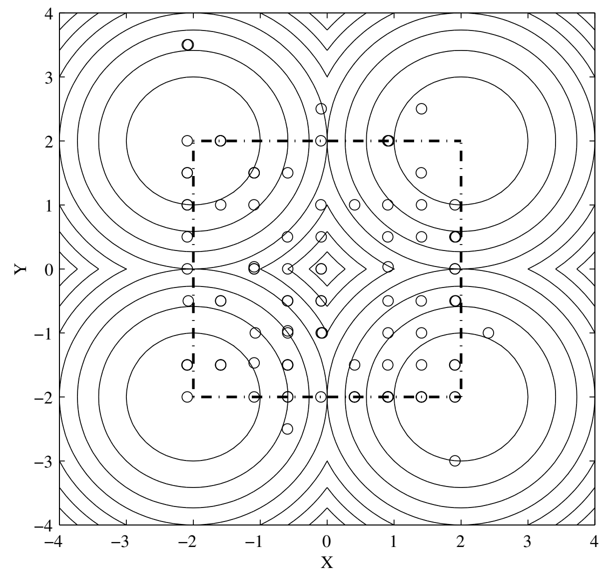

Figure 5: Population at generation 500 for F4.

<u>Test Case F5</u>

The next test case was identical to the previous case F4 but with two additional constraints added to the problem,

$$\text{Minimize F4 with:}\quad \begin{aligned} (x\text{-}1.67)^2 + (y\text{-}.5)^2 &\le 0 \\ y - 0.5x &\le 0 \end{aligned} \tag{5}$$

The Pareto optimal region is the square defined by the four individual function minima, constrained within the ellipse and above the line. Figure 6 shows the locations of the population of solutions on the plane at generation 100. Similar to the previous case, this convergence was reached at the 20th generation and was maintained through generation 500.

Cases F4 and F5 successfully showed the ability of the MINSGA to identify Pareto optimal regions for many objectives both constrained and unconstrained.

<u>Test Case F6</u>

The final test case consisted of two partially out of phase sinusoids.

$$\text{Maximize/Minimize:}\quad \begin{aligned} f_{61} &= 2\left\{\sin\left[2\pi\left(\frac{x}{5}-1\right)\right]+1\right\} \\ f_{62} &= 2\left\{\sin\left[2\pi\left(\frac{x}{5}\right)\right]+1\right\} \end{aligned} \tag{6}$$

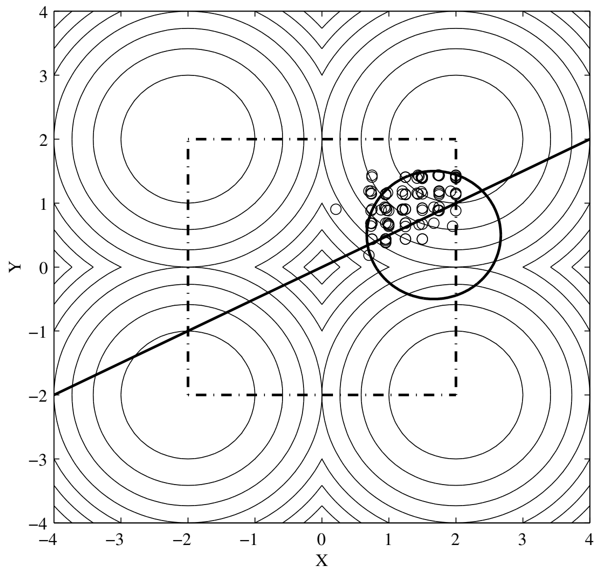

Figure 6: Population at generation 100 for F5.

Four variations are possible and were considered:

a) Minimize both.
b) Minimize f_{61} but maximize f_{62}.
c) Maximize f_{61} but minimize f_{62}.
d) Maximize both.

As can be seen in Figure 7, the MINSGA is able to effectively handle the maximum/minimum problem. Figure 7 displays both objectives as a function of x along with the Pareto optimal region in the objective space. All variations were successfully demonstrated and can be found in Reference [11].

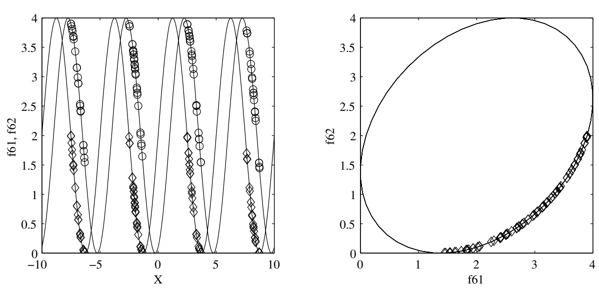

Figure 7: Population at generation 100 for F6.c

The six test cases demonstrate the accuracy and ability of the MINSGA. The results from the first three agree with previous efforts. The last three test cases demonstrate the advanced capabilities of the algorithm.

Constellation Design Using Satellite Tool Kit

Satellite Tool Kit (STK)® is a commercially available software package for satellite performance analyses. It allows the user to create operational scenarios including multiple satellites, targets, ground stations, sensors, and more. STK can provide information regarding coverage characteristics, communications, remote sensing, tracking issues, and link budgets. It also includes many high fidelity models and consideration of higher order effects including aerodynamics models and consideration for Earth geographical data. STK was chosen as the tool to use for satellite orbit propagation and coverage analyses.

MINSGA/STK Synthesis

Using STK/Connect, the communication protocol contained within STK, the MINSGA is able to obtain the objective values for each individual in the GA population. For example, objectives may be based on the constellation's percent global/regional coverage, target revisit time or total number of satellites in the constellation. An intermediary program performs the communications between the MINSGA and STK via a TCP/IP socket connection to the module STK/Connect. This intermediary program decodes the individual’s parameters from MINSGA into a satellite configuration. It then creates the constellation scenario in the format required to initiate STK. A simulation is propagated for a set amount of time to determine the performance of the particular configuration. The results of the simulation are sent from STK back to the intermediary program. The desired figures of merit are calculated or extracted from the results and then the vector-valued objective is returned to MINSGA. The process is repeated for each individual in the GA population. Figure 8 shows the basic operational scheme for the MINSGA/STK synthesis. The overall process that combines MINSGA with STK has been given the acronym CODEC referring to Constellation Optimal Design by Evolutionary Computation.

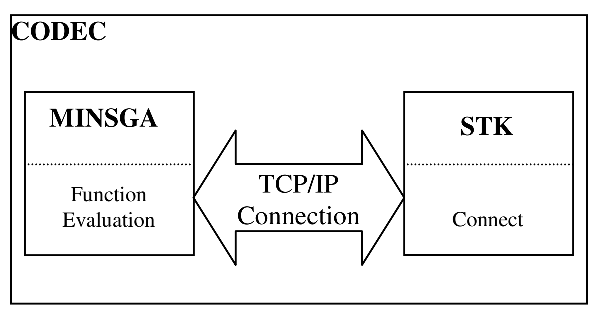

Figure 8: Operational flow of CODEC optimizer.

Verification of the CODEC

Although CODEC was designed with the goal of extending constellation optimization beyond simple circular and symmetric designs, the initial efforts have focused on such cases in an effort to demonstrate the abilities of the algorithm on established results.

The primary test case that has been considered has been that of obtaining maximum continuous coverage with a minimum number of satellites. This also provided the opportunity to measure the algorithm's performance and perform parameter characterization on problems with published results.

Test Case 1: Inclination Characterization

Walker performed his initial work based on the idea of minimizing the maximal distance between the two closest sub-satellite points for any given configuration. He determined that for a five satellite constellation in geo-synchronous orbit the optimal inclination occurs at 43.7°. As a byproduct of minimizing the sub-satellite point distance, he was able to identify configurations that produce continuous global coverage.[1]

The approach taken in this study was to directly target the coverage characteristics as the objective for the optimization. The smallest value of the percent global coverage observed during the scenario period is used as the measure. The first run of CODEC was performed to determine the inclination range for which a Walker 5/5/1 configuration provides continuous global coverage.[1] This is only a single objective optimization, but it provided an opportunity to ensure that the MINSGA/STK coupling was successful.

Niching was performed in hopes of dispersing the population over the entire region of continuous coverage inclinations. In the case of one free parameter, where one wants to spread the optimal population as much as possible, it is easy to determine the niching parameter. Using a normalized sharing scheme the sharing parameter is given by

$$\sigma_{share} = \frac{1}{Population\ Size} \tag{7}$$

The niching parameter is not as easy to determine for more complex problems and search spaces. Table 1 shows the values of the major control parameters for the inclination characterization that was performed. The sharing function parameter value of 2 corresponds to a parabolic sharing function.

A total string length of 8 bits allows 4096 different values for the inclination. In this case where the parameter is the inclination which ranges between 0° and 90°, the string length provides a resolution of approximately 0.02°.

Table 1: Parameters for Test 1.

Control Parameter	Parameter Value
Number of Generations	50
Population Size	50
String Length	8
i range	0.0° - 90.0°
σ_{share}	0.02
α	2

CODEC was able to correctly identify the region between 15° and 90° inclination as providing continuous global coverage. This also demonstrates the advantage of niching which is to produce a range of optimal solutions rather than a single point in the search space.

Test Case 2: Walker Delta Configuration Study

Next, the Walker configuration was included as part of the parameter set. This was an attempt to demonstrate CODEC's ability to handle a true multi-objective optimization. The intended result was to produce a population that illustrated the family of optimal Walker constellations at GEO, thus generating a group of configurations that will allow for a trade between continuous coverage and number of required satellites. The differences in control parameters are displayed in Table 2. Since Equation (7) is no longer adequate, the value of σ_{share} was increased to allow for the additional number of dimensions in the search space. Here P is the total number of planes, S is the number of satellites per plane, and F is the phasing between ascending nodes in adjacent planes.[12] A total string length of 16 bits is a result of assigning 3 bits to the P and F parameters, 2 bits to S, and 8 bits to the inclination.

Table 2: Control parameters for Test 2.

Control Parameter	Parameter Value
P range	1 - 8
S range	1 - 4
F range	0 - (P – 1)
String Length	16
σ_{share}	0.035

The results for the second CODEC run are summarized in Figure 9 and represent the population after 20 generations. The individuals with six satellites survive alongside those with five due to the niching. There is a larger area of the search space where individuals with six satellites can exist, hence they are not penalized as severely by the niching as the configurations with fewer satellites. Results from generation 20 are discussed because subsequent generations were unable to maintain diversity. It is speculated that improper sizing of the niching parameter may be responsible.

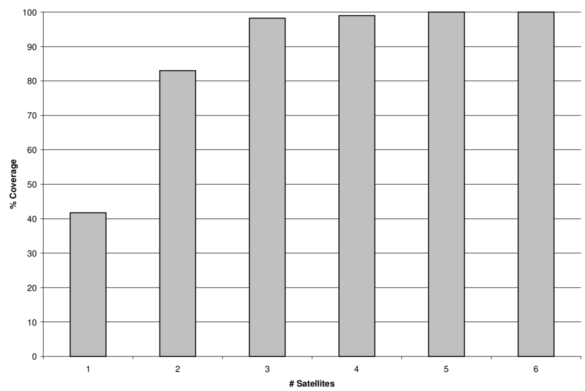

Figure 9: Trade between number of satellites and percent coverage.

In the population at generation 20, there were 42 non-dominated individuals. Of those, 30 had a total of 6 satellites, 4 had 5, 4 had 4, 1 had 3, 2 had 2, and 1 had 1 satellite. Within the individual solutions with 6 satellites, the inclinations ranged from 9.5° to 67°. Figure 10 shows individuals 19 and 35, 2 of the 5 satellite constellations, along with a Walker Delta pattern at his optimal inclination, i_{opt}, of 43.7°. Individual 35 has an inclination of 54° and individual 19 has an inclination of 22°. All three solutions provide continuous global coverage offering a range of options. Individual 35 is represented by the circles and the solid orbit track, individual 19 is shown by crosses with dashed orbit tracks, and the i_{opt} constellation is marked with plus symbols and dotted orbit tracks. Note that for all three constellations, at the instant of this figure, the satellite that has its longitude of ascending node at 0° is also at 0° latitude.

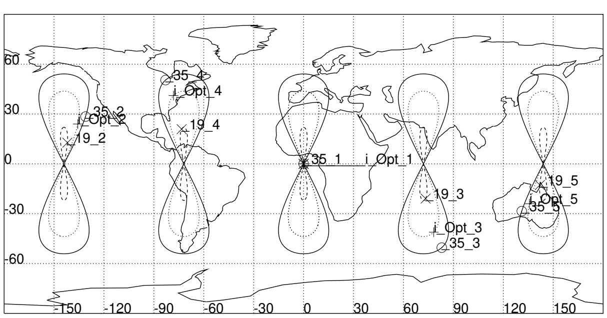

Figure 10: Individuals 19 and 35 Walker Delta

i_{opt}: 5/5/1.

Test Case 3: Open Structure with Fixed Inclination

The third CODEC run addressed continuous global coverage versus number of satellites but without restricting the constellation to a Walker Delta pattern. The right ascension of the ascending node, Ω, and argument of latitude, *u*, were both included as parameters for each of the possible satellites. The inclination of each satellite was set to 43.7°. Table 3 shows the major control parameters that changed from Test 2 for this trial. The total number of parameters was 29, 2 per satellite (Ω, *u*) plus a binary parameter that indicated whether the particular satellite was active in the constellation. At least one satellite was always assumed active.

Table 3: Control parameters for Test 3.

Control Parameter	Parameter Value
String Length	169
Number of parameters	29
Number of Satellites	1 - 10
Ω range	0° - 360°
u range	0° - 360°

The population after 50 generations contained 8 non-dominated solutions. Two of those solutions were identical, containing 5 satellites in a configuration and can be seen in Figure 11. This configuration provides 99% continuous global coverage, with only a small portion of the poles out of view for a short period of time each day. Obviously this is not as good as the known Walker Delta pattern at the same inclination, however this result is obtained with only 50 individuals over 50 generations. The configurations represented in the population begin to evolve Walker like characteristics. The ascending nodes are spread around the equator and the phasing between planes approaches an even distribution.

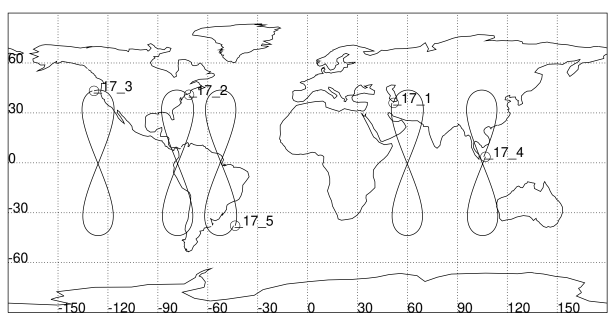

Figure 11: Individual 17 for the *i*=43.7° case.

Test Case 4: Open Structure with Posigrade Inclination

The next trial involved the same scenario as the third, but the inclination of each satellite was freed and allowed to assume any posigrade value, shown in

Table 4. The total string length of 179 bits is a result of assigning 5 bits for inclination, 6 bits for both the argument of latitude and longitude of ascending node, and a single bit representing whether the particular satellite is active. At least one satellite is always active.

Table 4: Control parameters for Test 4.

Control Parameter	Parameter Value
String Length	179
Number of parameters	39
i range	0° - 90°

Figure 12 displays a solution (Individual 7) from this run. This solution provides 99.7% continuous coverage, with short duration coverage gaps near the poles. Given more time to evolve, it is expected that a solution with 100% continuous coverage would appear. The solution was not necessarily expected to mimic a pure Walker Delta. The Walker Delta resides within an optimal plateau. There is a significant neighborhood of the search space surrounding the Walker Delta that also provides 100% global coverage.

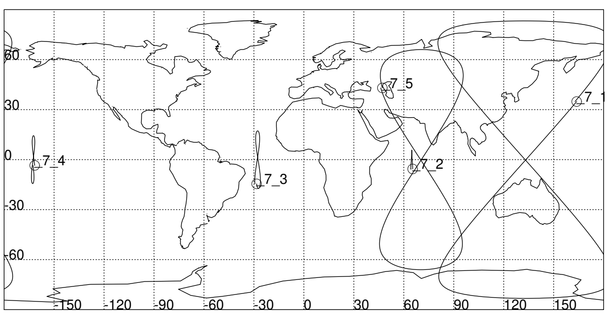

Figure 12: Individual 7 for the posigrade case.

<u>Test Case 5: Open Structure with Open Inclination</u>

The last attempt at the simple global coverage problem extended the case by allowing the inclination to vary between 0° and 180°. Figure 13 shows a typical individual with 5 satellites. This configuration provides 99.2% continuous global coverage. While such a configuration is not desirable when compared with other solutions due to all satellites possessing retrograde inclinations, it does demonstrate the capability of the algorithm to produce creative and novel solutions that may not otherwise be considered. No information about favorable inclination was provided or specified in the algorithm so there was no selection pressure favoring posigrade orbits.

Suggestions for Future Research

There are several areas that require further study. The sizing of niching parameters is not well understood. As was seen in the trial cases, this can be a source of difficulty in obtaining good results. A methodology for determining the niching neighborhood that does not require a priori knowledge of the solutions is sorely needed.

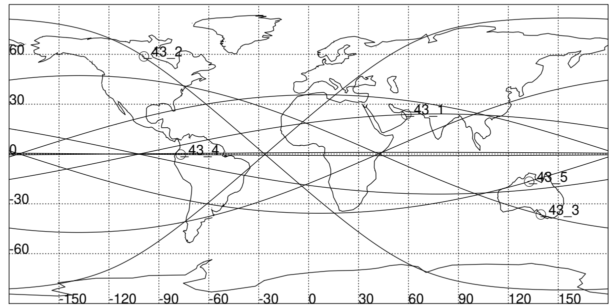

Figure 13: Individual 43 for the open inclination case.

While STK is a useful tool, it is designed for interface with a human operator. The resources and time devoted to graphical calculation and display are wasted in such an application as in this study. A working version of a satellite analysis tool that allows direct access to the computational engine without a graphical user interface should substantially increase speed and performance. Whether STK or some other software tool provides this is unimportant. Any coverage software package could be used in place of STK with proper interfacing.

Some early efforts at more complex objective requirements have looked promising. Problems involving regional and/or multiple coverage, both with and without constraints have shown to be easily within the grasp of CODEC. More advanced problems with link budgets, overall constellation restrictions, inter-satellite links and even moving targets will be investigated.

Conclusion

The results obtained in this study demonstrate the use of a Pareto genetic algorithm to perform satellite constellation design. It has been shown that a group of solutions can be produced demonstrating a trade between objective achievement and cost.

As a secondary point, the benefit of incorporating pre-existing commercial software was demonstrated. Development time was saved and a much richer tool was developed due to the advanced capabilities of an established tool.

Acknowledgments

The authors would like to thank Analytical Graphics, Inc. for their generosity. Their Educational

Alliance program provides Satellite Tool Kit free to universities, thus making this research possible. In particular C. Mulcahy's assistance with STK was essential to the success of this study. The authors would also like to thank D. Goldberg for initially bringing to our attention the NSGA, and for providing many useful leads, including a referral to K. Deb whose advice was critical. Many thanks as well to D. Carroll for developing the original genetic algorithm software used as the foundation for the MINSGA.

[1] Walker, J. G., "Some Circular Orbit Patterns Providing Continuous Whole Earth Coverage", *Journal of the British Interplanetary Society*, Vol. 24, 1971, pp.369-384.

[2] Ballard, A. H., "Rosette Constellations of Earth Satellites," *IEEE Transactions on Aerospace And Electronic Systems*, Vol. AES-16, No. 5, 1980, pp. 656-673.

[3] Rider, L, "Optimized Polar Orbit Constellations for Redundant Earth Coverage," *The Journal of the Astronautical Sciences*, Vol. 33, No. 2, 1985, pp. 147-161.

[4] Draim, J. E., "Three - and Four-Satellite Continuous-Coverage Constellations," *Journal of Guidance, Control and Dynamics*, Vol. 8, No. 6, 1985, pp. 725-730.

[5] Draim, J. E., "A Common-Period Four-Satellite Continuous Global Coverage Constellation," *Journal of Guidance, Control and Dynamics*, Vol. 10, No. 5, 1987, pp. 492-499.

[6] Goldberg, D.E., Genetic Algorithms in Search, Optimization, and Machine Learning, Addison Wesley, 1989.

[7] Principia Cybernetica Web, Web Dictionary of Cybernetics and Systems: OPTIMIZATION. http://pespmc1.vub.ac.be/ASC/OPTIMIZATIO.html, 1998.

[8] Srinivas, N. and Deb, K., "Multi-objective Optimization Using Non-dominated Sorting in Genetic Algorithms", *Evolutionary Computation*, Vol. 2, No. 3, 1995, pp. 221-248.

[9] Hartmann, J. W., Coverstone-Carroll, V., and Williams, S. N., " Optimal Interplanetary Spacecraft Trajectories Via A Pareto Genetic Algorithm", AAS/AIAA Space Flight Mechanics Meeting, Monterey, CA, Feb. 1998, Paper AAS 98-202.

[10] Coverstone-Carroll, V., Hartmann, J. W., Williams, S. N., and Mason, W. M, "New Results in Astrodynamics Using Genetic Algorithms", *Computer Methods in Applied Mechanics and Engineering*, to appear.

[11] Mason, W. M., "Satellite Constellation Design Via Evolutionary Computation", M.S. Thesis, Department of Aeronautical and Astronautical Engineering, University of Illinois at Urbana, Anticipated Submission, Oct. 1998.

[12] Larson, W. J. and Wertz, J. R., Space Mission Analysis and Design, Kluwer Academic Publisher, second edition, 1993.