The Use of Genetic Algorithms in Multilayer Mirror Optimization

Shannon Lunt   R. S. Turley

Department of Physics and Astronomy
Brigham Young University, Provo, UT 84602

Abstract

We have applied the genetic algorithm to extreme ultraviolet (XUV) multilayer mirror optimization. We have adapted the genetic algorithm to design optimal bifunctional mirrors for the IMAGE/EXPLORER Mission. Our best design, a 16-layer aperiodic stack of alternating layers of Y_20_3 and Al had a predicted reflectivity of 36% at 304 Å and 0.2% at 584 Å.

1 Introduction

In this section we will discuss the motivation for the design of bifunctional XUV multilayer mirrors, the IMAGE/EXPLORER mission. The mission placed unique requirements on the multilayer mirror we needed to design. We adopted the genetic algorithm (GA) to meet these design challenges.

1.1 IMAGE Mission-XUV and Specifications

We designed a mirror for the XUV section of the IMAGE Mission which will be launched in early 2000 and whose goal is to take images of the earth's magnetosphere (magnetic field lines surrounding the earth contained by the solar wind). Particles from the solar wind compress and confine the magnetic field on that side and stretch it out behind the earth on the other side. In the process of compression, particles are transferred from the solar wind to the magnetosphere. The IMAGE Mission's goal is to study the interaction of these particles in the magnetosphere. The mirror was specified at 14.5 degrees from normal to be highly reflective (> 20%) at 304 Å to see the singly ionized helium lines from the magnetosphere and to be non-reflective (< 0.2%) at 584 Å to cut out the bright neutral helium lines from the earth's atmosphere which would saturate the detector.

We were encouraged that such a design would be possible because 584 Å is a little less than twice as 304 Å. We thought that we could find layer thicknesses giving constructive interference at 304 Å that would interfere destructively at 584 Å. Because the magnitude of the index of refraction of materials in the XUV is close to one, they are not very reflective. As a result, to get high reflectivity, multilayer mirrors are made of many layers of materials. Figure 1 shows reflections from multiple layers can add. Depending on the thicknesses of the materials (and therefore the relative phase of the various reflections), the outgoing waves either add (constructive interference) or cancel each other (destructively interference).

Meeting the project specifications was difficult because most materials are more reflective at 584 Å than at 304 Å. Also, aperiodic stacks, multilayer mirrors with each layer having a different thickness, were found to produce better mirrors for this problem than periodic mirrors did, contrary to most previous expectations [2]. However, using aperiodic stacks increases the size of the solution space. In this case we typically need to select two materials for the stack (from a list of dozens of possibilities) and the thickness of each of typically 16 layers.

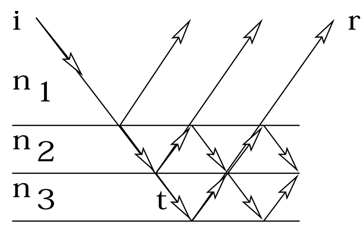

Figure 1: Multiple reflections from multilayers

1.2 A New Application of the GA

Genetic algorithms have been used in optimization problems in fields related to multilayer mirrors and seem to produce better results than alternative methods in problems with discrete variables, discontinuities in the solution, and multiple parameters [1]. However, they have never been used in the XUV region for optimizing multilayer mirrors, especially bifunctional mirrors.In applying the genetic algorithm (GA), many of the parameters in the code had to be determined that are specific to the problem of designing a mirror for the IMAGE Mission. Also, each problem to which the GA is applied requires a unique merit function, a function that tells how good or bad a solution is. Our implementation of the GA was hybrid of the standard GA approach and a local optimizer using the simplex algorithm to reduce computation time. Hybrid approaches have been used effectively by others in different applications of the GA such as using the GA with simulated annealing in calculating the optical constants of materials [3].

1.3 Outline

In Section 2 we will summarize the genetic algorithm and its application to the mirrors for the IMAGE Mission. In Section 3 we present and discuss results obtained for this particular application of the GA. The conclusion in Section 4 discusses when it is useful to use the GA, some rules to apply it, and future work to be done.

2 Genetic Algorithm

2.1 Description

As the name implies, GA's use a similar technique to nature's process for optimization and refinement through the use of DNA and survival of the fittest [1]. Table 1 describes some of the terminology used in the GA. The attributes of each member of the *population* to be optimized are encoded in a DNA-like array within *chromosomes.* For the specific problem described in Section 1.1, the materials and thicknesses in the multilayer were encoded into a *gene*, an array containing the materials and thicknesses in the stack, as shown in Figure 2. Each *allele* in the *gene* was stored in a byte so there were constraints on the thicknesses due to the storage constraints. The initial *population* of mirrors was chosen randomly with the program choosing the two materials to use and each layer thickness, with the number of layers being fixed for each run.

Population	set of trial solutions
Generation	successively created population
Gene	array containing materials and thicknesses
Allele	each material or thickness in the gene
Parent	member of the current generation
Child	member of the next generation
Chromosome	coded form of a trial solution consisting of genes made of alleles

Table 1. Terminology in the GA

Material 1	Material 2	1 A - Thickness 1	2 A - Thickness 2	1 B - Thickness 3	2 B - Thickness 4

Figure 2: Composition of the Chromosomes

Parents are then selected based on the value of their merit function, which contains the information to be optimized. The merit function in this application of the GA included the specifications of the mirror design for the IMAGE Mission. *Children* are then produced by crossover and mutation of the parents' genes. The next *generation is* composed of these *children* and the best *parents* of the current *generation* with the process continuing until the merit function ceases to change significantly. A schematic of the GA is shown in Figure 3.

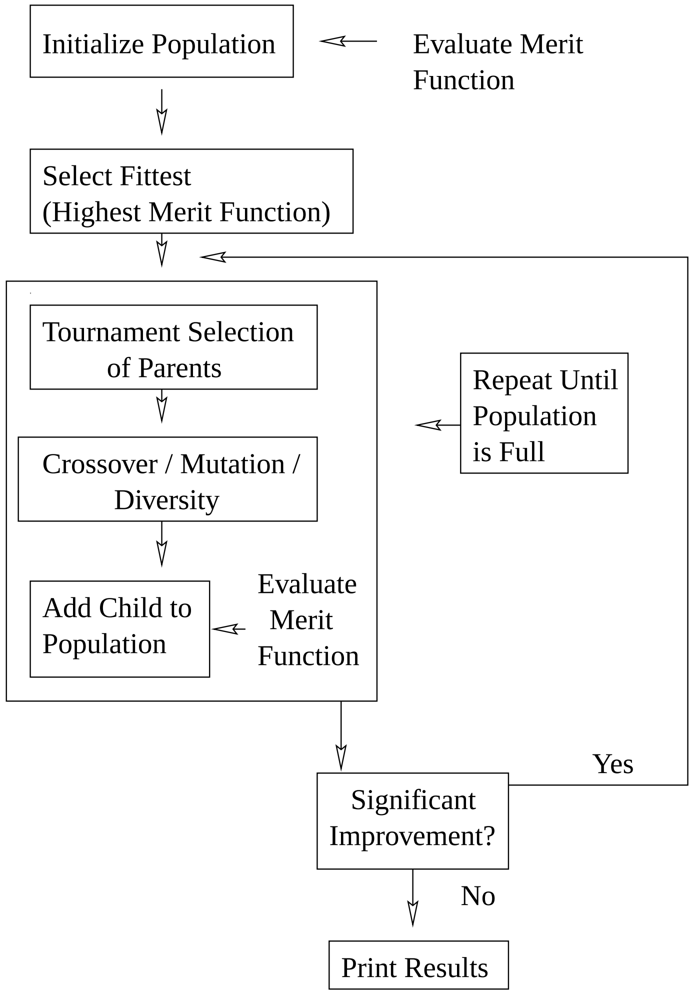

Figure 3: Schematic of the Genetic Algorithm

2.2 Advantages and Disadvantages for this Class of Problems

The GA is a global optimization procedure that overcomes many of the problems associated with local optimization procedures. Although this technique finds global extremes, it usually converges slowly to a solution since it initially fills the population in a hit-and-miss fashion and does not take into account any details of the shape of the merit function surface. However, one does not need to compute gradients for convergence, as in some of the local optimization techniques, so the actual encoding of the problem is quite simple. Also, it is not very dependent on an initial solution, a very useful feature when encountering a new problem for which one has little intuition. In addition, the GA handles discrete and constrained variables much more easily than most local optimizers.

2.3 The GA Applied to the IMAGE Mission

In applying the GA to the design of mirrors for the IMAGE Mission, many parameters had to be chosen and studies done to ensure that the code would find the best solution. Table 2 lists the parameters that must be chosen each time the GA is applied to a new problem.

Crossover Probability
Mutation Probability
Population Size
Replacement Percentage
fraction of population used for hybrid (if applicable)

Table 2: Parameters specific to each application of the GA

2.3.1 Materials and Thicknesses

Before choosing the parameters specific to the problem, information about the problem must be encoded in the GA as described earlier in this section. For the IMAGE Mission mirrors, the encoded information included the thickness and composition of each layer. The code allowed for the possibility of two oxides on top of the stack of either fixed or variable thickness. The two alternating materials of the mirror could be fixed by the user or chosen from a database by the program. The database was a compilation of many common materials and included the optical constants of these materials at 304 Å and 584 Å.

There were also constraints placed on the thicknesses of each layer. The alleles in the gene were each stored in a byte. This created an upper limit on the thicknesses of each layer of 255 Å, which was extended to 275 Å by adding 20 Å to each thickness at the end of each run. This made the lower limit for each thickness to be 20 Å, a constraint imposed by our sputtering process.

2.3.2 Selecting Parents and Reproduction

One of the decisions to be made in applying the GA is the selection of parents. Some possible strategies include:

- *population decimation* where only the members with merit functions above a cutoff value are kept,

- *proportionate selection* where the probability of choosing a parent is based on its fitness and
- *tournament selection* [1]

In the application of the GA described here, several strategies were tried but the tournament selection was chosen as it seemed to work the best at the end. Proportionate selection does not distinguish well between good solutions and slightly better solutions. Thus, as you get near the end of a run and most of the members of the population are good, the program goes through many iterations before it ends while tournament selection allows the program to terminate faster. In tournament selection, two members of the population are chosen randomly and their merit functions are compared. The member with the higher merit function is chosen to be one of the two “parents.” The genes from the parents are combined with the possibility of crossover and mutation to produce two children.

When crossover occurs, as shown in Figure 4, a random byte (allele) in the gene is chosen and at that location, the byte becomes a hybrid of the two parents' genes at that byte. For a discrete variable, the new allele is a combination of the bits of the parents' allele at that point. For a continuous variable, the allele in the children becomes an average of the two parents' alleles. The bytes before the location of the hybrid are then copied from one parent and all of those after the location are copied from the other parent. The degree to which crossover occurs is based on a crossover probability determined in encoding the GA for a specific problem. In mutation, as shown in Figure 5, a random byte is chosen in the gene and is randomly altered. The mutation probability is similar to the crossover probability in that it determine how often, if ever, the mutation is applied in producing children.

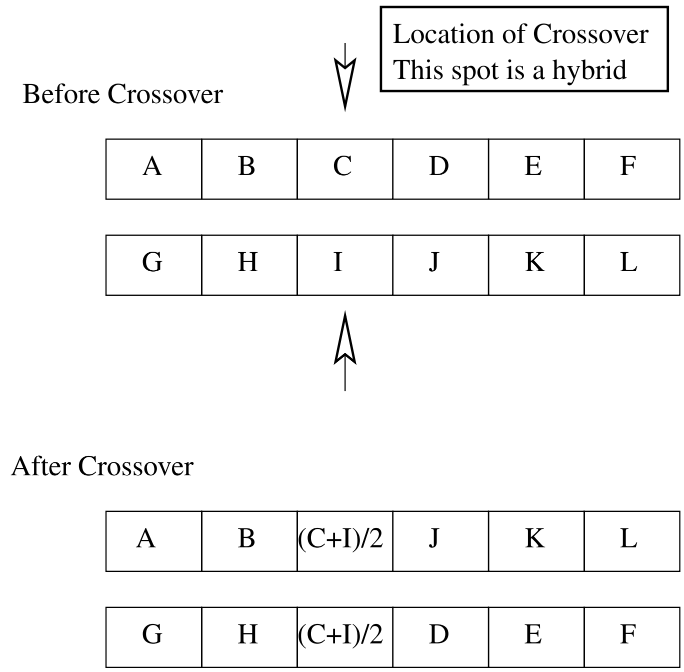

Figure 4: Genetic Algorithm Crossover

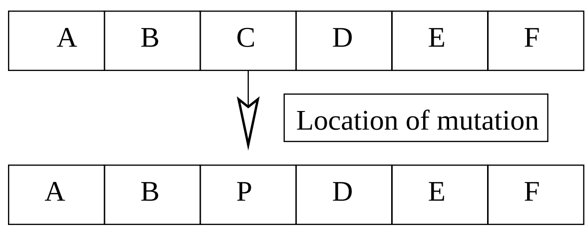

Figure 5: Genetic Algorithm Mutation

This selection of parents and reproduction is repeated until the number of children desired is produced. These children and the best members of the parent population or previous generation then make up the new generation. The percent of the previous population to be kept is one of the parameters to be chosen in each application of the GA (see Table 2).

An attempt was made to include diversity the code to better simulate natural selection. In nature, if animals compete for the same resources, the strongest will survive by being able to meet their needs. Diversity exists in nature because different animals use different resources and do not compete with each other. In the code, multiple occurrences of a set of materials would be like animals competing for the same resources. The first member in a population with certain materials was given a weight of one and all other occurrences of those materials were given a successively lower number as a weight. This weighting was then included in the reflectivity calculations, with the first instance having the highest reflectivity. By decreasing the reflectivity of multiple instances of the materials, the merit function is decreased and the chances of the mirror surviving to the next generation are decreased. This diversity attempts to sample more of the solution space and get out of local minima if the code gets stuck. Unfortunately, we never got this feature working to our satisfaction.

2.3.3 Merit Function

The merit function in this problem compared the reflectivity of each member at 304 Å against its reflectivity at 584 Å to fit the specifications of the mirror design, as explained in Section 1.1. The merit function used follows, where R_{304} and R_{584} are the reflectivities at 304 Å and 584 Å respectively:

$$\frac{R_{304}}{\max(.002, R_{584})} \quad (1)$$

One can see from this relation that stacks with high reflectivity at 304 Å and low reflectivity at 584Å will be favored above other stacks by having a higher merit function and, thus, being more likely to survive in subsequent populations and to produce children. The program seemed to favor minimizing the reflectivity at 584 Å over maximizing the reflectivity at 304 Å. As a result, mirrors were found with very good low reflectivity at the longer wavelength that more than satisfied the requirements but the reflectivity at the shorter wavelength was not very high. It was more important, then to only get the reflectivity at 584 Å down to 0.2% and not any lower. By taking the maximum value between .2% and the reflectivity of the mirror at 584 Å when calculating the merit function, the program was forced to maximize the reflectivity at 304 Å and better mirror designs were found.

2.3.4 Hybrid Used

The GA uses a time-intensive hit-and-miss approach to initialize and alter each generation. To cut down on the execution time of this program and to make sure the entire solution space was sampled, the genetic algorithm (a global optimizer with slow convergence) was combined with a simplex algorithm (a local optimizer with rapid convergence). The GA was used to initialize the population and then the simplex was applied to the thickness of the mirrors to improve a certain fraction of the population. In the simplex algorithm, the thicknesses of a stack are encoded into a simplex geometric shape representing the parameter space which is altered until the optimal solution is found [4]. The simplex was only applied to a small amount of the population (one hundredth) and it was found that this made for more rapid convergence to the optimal solution and allowed the GA to be applied to smaller population sizes without sacrificing performance.

Varying Materials					
Hybrid	Gen	R_{584}	R_{304}	Mats	Time(sec)
yes	14	.2%	34.2%	U/Te	2763.4
no	3	.2%	20.72%	Y_20_3/Al	329.0

Table 3: 16 layers on $Si0_2$ with a population of 6000

The data in Table 3 was obtained with the GA choosing the materials for a mirror made up of 16 layers (8 layer pairs) on $Si0_2$ with a population of 6000. The GA was first run with the hybrid acting on one hundredth of the population and then as the straight GA. The GA with hybrid found a much better solution than the straight GA, as can be seen by comparing the reflectivities at 304 Å. The computation time for the hybrid run on a DEC Alpha 200 4/233 workstation was about nine times greater than for the run without hybrid probably because more generations were produced before a good solution was found. The straight GA run looks like it got stuck in a local maximum very early on in the run and was unable to get out. The hybrid, on the other hand, sampled more of the solution space and so was able to find a better solution.

Fixed Materials					
Hybrid	Gen	R584	R304	Mats	Time(sec)
yes	10	.2%	36.45%	Y203/Al	1967.2
no	41	.2%	36.23%	Y203/Al	4019.0

Table 4: 16 layers on $Si0_2$ with a population of 6000

The data in Table 4 was obtained by running the GA with the materials fixed to see how the two versions of the program compared. This study shows the difference in the computation time between the GA with hybrid and without. From the time of computation and the number of generations taken, one can see that the hybrid was able to converge to a solution much quicker than the regular GA. It looks like each generation of the regular GA takes less computation time, though, which is very surprising. Although the simplex algorithm to a solution faster than the GA does, it requires more iterations for each step so it solution than the GA, as in Table 3.

There are many other local optimizers that could have been used in conjunction with the GA in this application, such as conjugate gradient or BCGS. The simplex is probably the slowest local optimizer known but it is also the most robust and easiest to implement.

3 Results

We discuss the results obtained using the GA for the IMAGE Mission and the mirror design actually used for the IMAGE project in Section 3.1. The values of the parameters used in this application are given in Section 3.2.

3.1 Mirrors Designed for the IM

In searching for the optimal design of the mirror following specifications: hybrid, 16 layers, mutation probability=0.05, crossover probability=0.75, population size=8000, replacement percentage=50%, and with the algorithm choosing the materials. The best design found was Y_20_3/Al with a reflectivity at 304 Å of 36% and a reflectivity at 584 Å of < .02%, as shown in Figures 6 and 7.

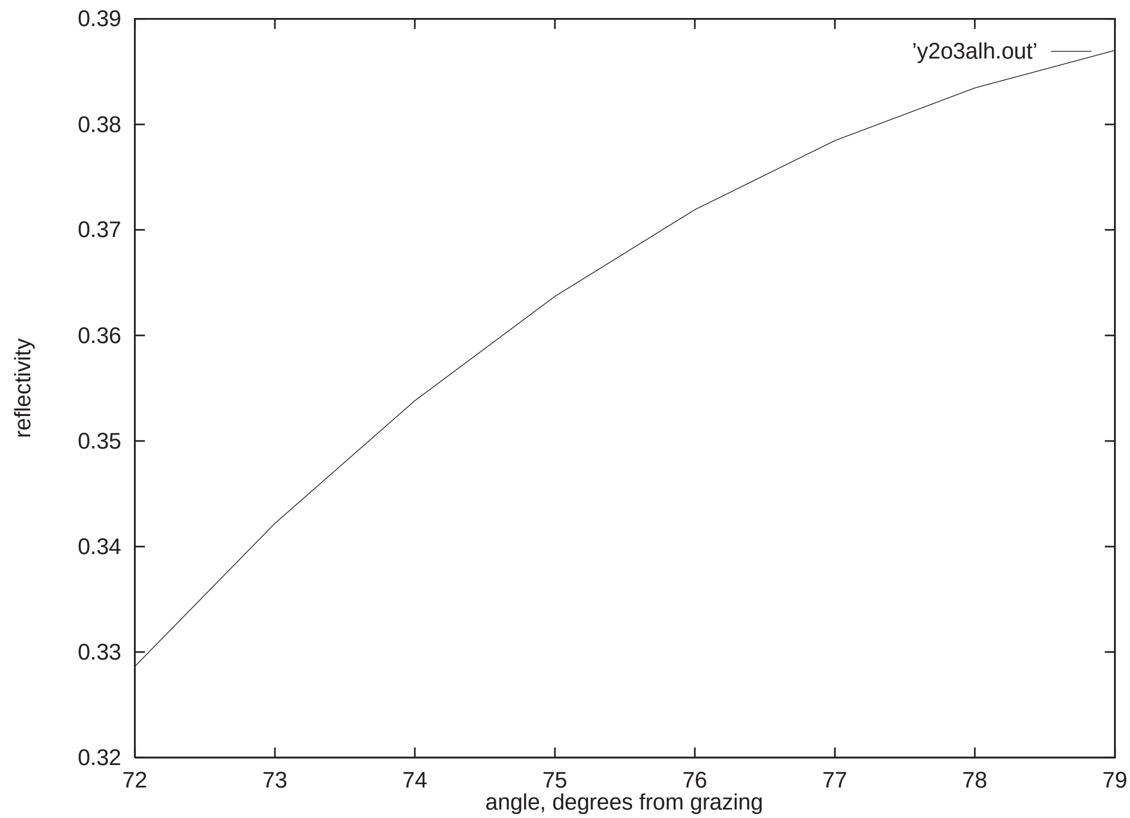

Figure 6: Reflectivity of best Y2O3/Al mirror at 304 Å.

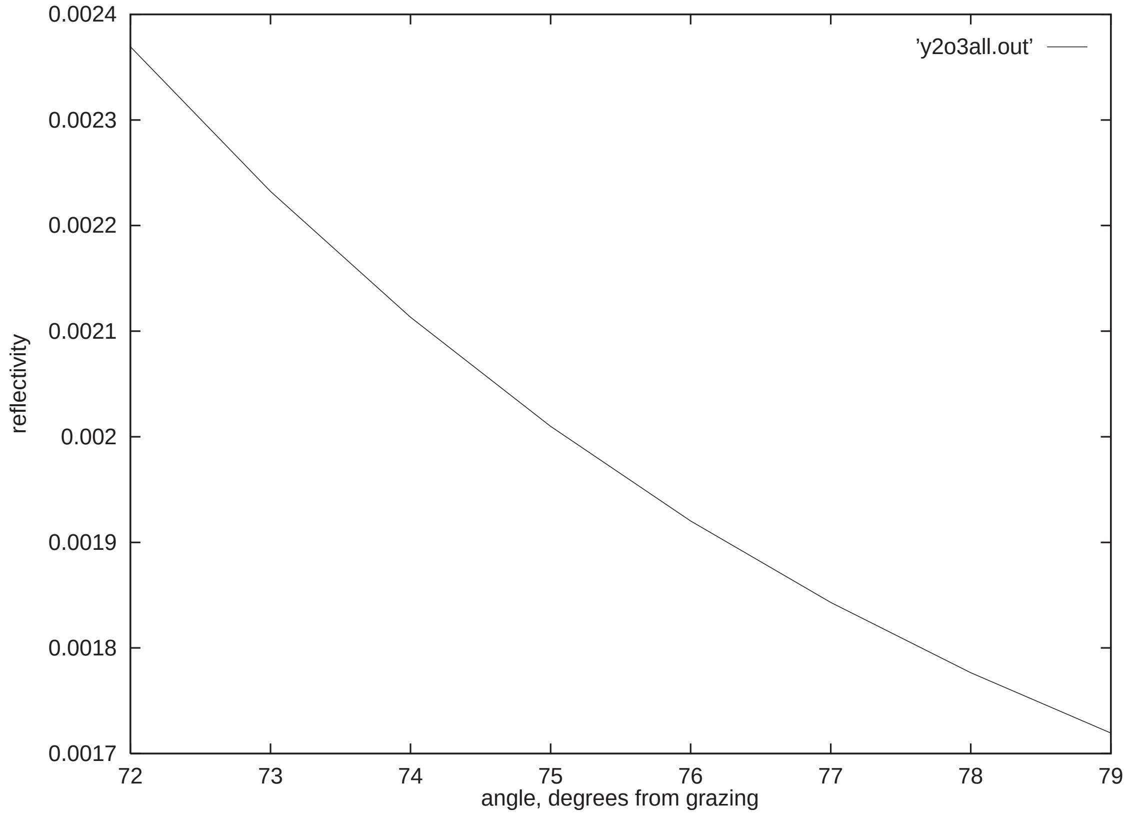

Figure 7: Reflectivity of best Y2O3/Al mirror of 584 Å.

This design was very surprising and non-intuitive because it included an oxide. The presence of oxygen in the form of an oxide on the top of a mirror usually decreases the reflectivity significantly. Also, elements were thought to reflect better than other designs from an experimental view , though, since it already includes an oxide. Most mirrors are made and then the top layer or so oxidizes and the reflectivity decreases because oxygen has a very low index of refraction. The Y_20_3/A1 mirror would not have that problem since the Y is already in an oxide compound

Attempts were made to grow the Y_20_3/Al mirror but many difficulties were encountered. Due to the nature of Y_20_3, a ceramic, the sputtering process used at the time to produce the mirrors would not work. We are currently working on setting up an RF sputtering system to be able to make this mirror and other films.

3.2 Optimal Parameters

For each application of the GA to a different problem, one must find the optimal values for certain parameters used in the code. The main parameters to be chosen are the fraction of each population that is replaced by new children, the crossover probability, and the mutation probability. If a hybrid is used, one must also choose the fraction of each population it is applied to. Each problem also requires its own merit function specific to what is being optimized. Johnson and Rahmat-Samii [1] give a range of values in which the optimal solution for most problems will be found: smaller replacement percentages usually lead to faster convergence, 0.6-0.9 for the crossover probability, and 0.01-0.1 for the mutation probability. For the problem of finding a mirror that is highly reflective at 304 Δ and non-reflective at 584 Δ, the best values for the parameters listed above were found to be: replacement percentage=50%, crossover probability=0.7, mutation probability=0.05, and the hybrid applied to the population divided by 100.

4 Conclusion

4.1 Where the GA is Valuable

In doing optimization, it is often useful to have a good solution. Unfortunately, there is not a the advantages and disadvantages of the local most beneficial.

The GA is useful in optimization problems will look like. It also is very useful when satisfyi has discrete parameters and discontinuities. The random nature of this algorithm allows more of the solution space to be sampled and, if the population is large enough, to find the global extreme and not get stuck in a local extreme. This feature makes the GA inherently time intensive, though. In designing a mirror to meet the specifications of the IMAGE Mission, a hybrid approach worked best-the GA to find solutions all over the parameter space and a simplex algorithm to converge to the best solution. This hybrid significantly reduced the computation time and allowed smaller populations to be used without sacrificing the optimal solutions.

4.2 Rules for Application of the GA

When applying the GA to any problem, there are many parameters to be chosen: the population size, the mutation and crossover probabilities, the amount of each generation to be kept, and the merit function. Also, if time is a concern, a hybrid with another algorithm should be used. If a hybrid is used, one must decide how much of each generation to apply it to. For the problem of finding a mirror that was highly reflective at 304 Å and non-reflective at 584 Å, the following parameters were found to work best:

crossover probability=0.7, mutation probability=0.05, replacement percentage=50%, and the hybrid was applied to one-hundredth of each generation (population/ 100).

4.3 Future Research

More work still needs to be done on this application of the GA before it is complete. The code used in this optimization did not take into account manufacturability or feasibility of a design. There was a feature built in that would weight different materials more than others to represent lower cost or easier manufacturing but this feature was never fully developed or used. Also, the code is designed to optimize a mirror for broadband reflectivity . This can be changed with minimal altering of the code but the merit function may need to be changed as well. We plan to make the Y_20_3/Al mirrors described in Section 3.1 and compare the actual reflectivities with the calculated values.

There were many problems with getting the calculated and measured reflectivities to match. This is probably due to the uncertainty in optical constants in the XUV region. Many sources have vastly different values for the index of refraction so it was difficult to know which to use. Also, for many wavelengths in the XUV, the index of refraction is not even known. We plan to make the Y_20_3/Al mirrors described in Section 3.1 and compare the actual reflectivities with the calculated values. Obviously, this region in the spectrum still has many interesting problems that need to be studied to increase understanding of the XUV region and help the theory come closer to reality.

References

[1] J. Michael Johnson and Yahya Rahmat-Samii, "Genetic Algorithms in Engineering Electromagnetics," *IEEE* Antennas and Propagation Magazine, 39, Aug 1997, pp. 7-21.

[2] Masaki Yamamoto and Takeshi Namioka, "Layer-by-layer design method for soft-x-ray multilayers," Applied Optics, 31, Apr 1992, pp. 1622-1630.

[3] Aleksandra B. Djuriši□, Joven M. Elazax, and Aleksandar D. Raki□, "Simulated-annealing-based genetic algorithm for modeling the optical constants of solids," Applied Optics, 36, 1997, pp. 7097-7103.

[4] William H. Press, et al., Numerical Recipes in C. the Art of Scientific Computing, New York, Cambridge University Press, 1992.

Published in the 1999 J of the Utah Academy. Vol.**76**, pp.61-73.