

NEAR-OPTIMAL ANTENNA PLACEMENT USING GENETIC SEARCH*

Lara S. Crawford[†] and Victor H. L. Cheng[‡]
Optimal Synthesis Inc.
Los Altos, CA

Rich Burns[§]
Air Force Research Laboratory
Kirtland Air Force Base, NM

Shiang Liu^{**}
Aerospace Corporation
Los Angeles, CA

Abstract

The optimal placement of ground stations and antennas in the Air Force Satellite Control Network (AFSCN) to support the operation of a large set of satellites is a very difficult problem involving many complex, interacting constraints and requirements. This paper explores the application of genetic search techniques to antenna-placement optimization. The performance index used in the genetic search optimization was based on a candidate configuration's cost. Requirements on the number of satellite contacts that could be scheduled with the configuration were imposed as constraints.

Three different genetic search formulations of the antenna placement optimization problem were designed using this set of performance index and constraints. The first two genetic search formulations assumed that antenna placement was restricted to a pre-existing set of candidate ground site locations. The third genetic search formulation allowed for ground stations to be placed at any location on the Earth, though locations not inside predefined "feasible" regions were eliminated. The three genetic search formulations all found cost-effective solutions to the placement problem while maintaining serviceability.

1. Introduction

The Air Force Satellite Control Network (AFSCN) provides support for a large number of communications, surveillance, navigation, environmental, and research satellites. The AFSCN consists of a number of ground sites and antennas, along with support equipment and mission control centers. Economic considerations have recently motivated consolidation of the satellite support network resources. This research addresses the problem of placing the antennas and ground stations in a near-

optimal manner to minimize cost while ensuring adequate support for a given set of satellites.

The current configuration of the eight main tracking ground stations in the AFSCN is shown in Figure 1. There are sixteen tracking antennas with varying capabilities at these sites.^{1,2,3} The main control nodes for these "common-user" stations are at Onizuka Air Force Base in California and Schriever Air Force Base in Colorado. In addition, the AFSCN contains many special-purpose (dedicated) antennas, such as those for the Global Positioning System (GPS) and the Defense Meteorological Satellite Program (DMSP).




Figure 1. Locations of the eight primary tracking sites in the AFSCN. There are sixteen telemetry, tracking, and commanding (TT&C) antennas at these locations. The stations shown, from west to east on the map, are: Hawaii Tracking Station, Onizuka Air Force Base (control node), Vandenberg Tracking Station, Schriever Air Force Base (control node) / Colorado Tracking Station, Thule Tracking Station, New Hampshire Station, Telemetry and Command Station, Diego Garcia Station, and Guam Tracking Station.

The AFSCN must maintain support for Department of Defense satellites, NASA space vehicles, and some allied nation satellites. This commitment requires over 400 satellite contacts every 24 hours, for the purposes of tracking, satellite health and orbit maintenance, and data transmission. Requests for contact time are submitted in advance to Air Force personnel, who schedule the contacts for each antenna.⁴ Though the scheduling is currently done manually, efforts are being made to automate this process. The contact schedule takes into account the windows of visibility of each satellite with respect to each antenna, the capabilities of each antenna and ground station, and the contact

* This research was sponsored by the Air Force Research Laboratory, Space Vehicles Directorate, under contract number F29601-99-C-0087.

[†] Research Scientist

[‡] Vice President; AIAA Associate Fellow

[§] Research Engineer, Space Vehicles Directorate; AIAA Member

^{**} Project Engineer

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2000		2. REPORT TYPE		3. DATES COVERED -	
4. TITLE AND SUBTITLE Near-Optimal Antenna Placement Using Genetic Search				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFRL,Kirtland AFB,NM,87117				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 2. Antenna placement optimization concept.

requests. Since the placement of the remote tracking stations determine the satellite windows of visibility, the ground station and antenna configuration is critical to the successful performance of the satellite control network.

The ground stations and antennas are also a large component of the overall cost of maintaining the satellite control network. The problem, then, is to use the fewest possible ground stations and antennas while still maintaining the quality of the satellite support. The many complex constraints arising from scheduling and geopolitical requirements make this optimization extremely difficult. The quality of a station/antenna configuration is a non-analytic, non-convex function of the antenna placement and types. This type of discontinuous, non-smooth performance index, however, is not a barrier to genetic search methods. Genetic search methods provide a way of automating and giving direction to trial-and-error searches, and thus are not gradient-based. These methods are therefore an ideal approach for the antenna/ground station placement problem. Several genetic search formulations of this problem are presented here.

Section 2 discusses the characterization of the antenna placement problem. Section 3 presents the genetic search placement algorithms. Results using these algorithms are described in Section 4. Finally, Section 5 gives some conclusions and discussion of future research directions.

2. Placement Problem Characterization

2.1 Optimization Concept

Figure 2 illustrates the different processes involved in the optimization for the antenna placement problem. The optimization program creates candidate solutions, which are evaluated against the performance index and constraints. In order to produce useful antenna placements that improve satellite service and are cost-effective, all relevant variables and constraints must be

considered when designing the optimization procedure. These variables include the current network configuration, network performance requirements, cost, and allowable candidate antenna and station locations. Since one of the key performance factors in the antenna placement problem lies in the ability to support all requested satellite communications in a satisfactory manner, the processes to determine the satellite orbits, visibility windows, and contact schedules are explicitly shown in Figure 2. In formulating the optimization problem, there is a certain degree of freedom in whether to use the contact schedule in defining the performance index or in defining constraints.

2.2 Scenario Description

The scenario used in this research included 110 satellites of various types, each with daily contact requirements. This set of satellites is summarized in Table 1. The scheduling requirements were defined in terms of the average number of contacts required per day, the minimum and maximum pass duration, and the minimum and maximum time required for prepass preparation. Ephemeris data for the satellites was generated using Analytical Graphics, Inc.'s Satellite Tool Kit® (STK).⁵ Six days of ephemeris data was generated and saved at one-minute intervals.

Information about the current ground site and antenna configuration was also considered as part of the scenario. All antenna types were treated as identical for this work. Ownership of ground sites and special-purpose dedication of antennas (such as for GPS) were also ignored. In addition, all antenna locations were treated as if they were at sea level. No range or azimuth limits were used, and an elevation limit of zero degrees was used in all directions for all stations regardless of terrain.

Performance measure calculation necessarily involved the development of an algorithm to calculate the visibility windows for the satellites. In addition, to enable inclusion of performance measures requiring




Figure 3. Feasible regions for ground station locations. Existing ground stations are shown with “x”es.

explicit contact scheduling, a schedule-packing algorithm was developed. For ease in prototyping, the antenna placement optimization algorithms were developed within the MATLAB[®] environment. The computationally intensive scheduling and visibility window algorithms were implemented in C, however.

Table 1. Scenario satellite summary.

Satellite Class	Orbit Types	# of Satellites
Environmental	Circular Sun-Synch & Sun-Earth Libration Point (L1)	8
Communications	GEO	32
Navigation (GPS)	12-hr Circular	27
Surveillance	LEO, GEO, and Molniya	25
R&D	LEO	6
Others (Pseudo)	GEO	12

3. Placement Algorithms

3.1 Optimization Problem Formulation

There are many possible ways to formulate the antenna placement problem as an optimization problem. For this research the problem was posed as:

Minimize: Cost

Subject to: Satisfaction of all contact requirements

The contact scheduling requirements were specified as:

1. The percent of contacts satisfied must be greater than 99.5%.
2. The percent of satellites with all contacts satisfied must be greater than 99%. This value corresponds to all but one of the 110 satellites having all contacts satisfied.

In defining the cost of a given antenna configuration, geographical considerations as well as the number of antennas were taken into account. For the purposes of demonstration, 26 polygonal regions were defined as enclosing “feasible” ground station locations. Each of

these regions was assigned a cost value between 1.0 and 1.7. The regions are shown in Figure 3. The total cost of a candidate antenna/ground site configuration was defined as the sum of the station location cost (determined by the region in which the ground site is located) for each antenna. For most of the formulations discussed here, the performance index to be maximized was simply defined as the negative of the configuration cost.

The following sections discuss the genetic search techniques and representations used for carrying out the antenna placement optimization with these cost and scheduling specifications.

3.2 Genetic Search Techniques

Genetic search techniques (see, for example, references 6, 7, and 8) are a type of directed trial and error, inspired by the biological processes of natural selection and Darwinian evolution. Genetic search methods have an advantage over traditional optimization methods in problems with many complex, discontinuous constraints such as the antenna placement problem. Because genetic methods do not rely on any performance derivative information, discontinuous and non-smooth performance indices and constraints can be easily accommodated. In addition, genetic search does not have a problem with local minima, as do many other methods. Because genetic methods are not analytic, the performance index can be arbitrarily complex, as long as it can be calculated. The main drawback of the genetic approach is that it can be computationally wasteful. In the antenna placement problem, however, computational speed is not a major consideration, as the decision to close a ground station or build new ones generally would require careful, lengthy study to justify. Genetic search methods are thus ideal for the antenna placement problem.

A flow chart describing the genetic search process is shown in Figure 4. In a genetic search formulation,

a candidate problem solution is formulated as a set of one or more *chromosomes*. An initial population of these candidate solutions that spans the search space is defined. Each member set of chromosomes (candidate solution) in the initial population is evaluated with a performance measure, or *fitness*. Then a search loop is executed repeatedly. During the search loop, new members are created from parent members in the population through genetic operations such as *mutation* or *crossover*. The parent members are generally selected probabilistically in a manner dependent on their fitness values. The newly-generated members are then evaluated. This evaluation, as for the initial population, is based on a performance index; here, that index was the negative of the cost, as described above. New members that violate constraints can be assigned poor fitness values or can simply be deleted. At each iteration through the search loop, members with poor fitness may be deleted from the population to maintain the population size.

The implementation of genetic search algorithms in this work was simplified considerably by the use of a proprietary software package.^{9,10} This software, which runs within MATLAB®, provides an integrated environment for facilitating genetic search design.




Figure 4. Genetic search flow chart.

3.3 Formulation Using Candidate Sites

The antenna placement problem was first formulated using a predefined set of candidate sites at which antennas could be placed. Implementation and evaluation of this type of formulation provided insight for the development of the more general formulation allowing ground sites to be located anywhere, described in Section 3.4. Using only candidate sites, two chromosome representations were developed: the baseline representation and the ground-up representation. Both of these representations used the set of fifteen candidate sites shown in Table 2, selected based on an earlier Air Force study. This list of candidate sites included the eight Common User

Element (CUE) sites of the AFSCN, which form the core antenna resources of the network. The CUE sites are shown above in Figure 1.

Table 2. List of candidate ground sites for baseline and ground-up formulations.

1.TTS	6.HTS	11.Kwajalein
2.FTS (Fairbanks)	7.VTS	12.Argentia
3.TCS	8.CTS	13.Ascension
4.DGS	9.NHS	14.Onizuka
5.GTS	10.Det. Alfa	15.Camp Parks
	Prospect Harbor, ME	

Baseline Representation

In the baseline representation, a configuration of antennas was represented with respect to a current, or baseline, configuration through “antenna-add” and “antenna-delete” operations. A set of 29 antennas located at the 15 candidate sites was used as the base case.

For this representation, a chromosome was defined as a string of integers between 1 and 30, with a number n between 1 and 15 indicating an “antenna-add” operation at the corresponding ground site number n , and a number n between 16 and 30 indicating an “antenna-delete” operation at ground site $n-15$. Thus, the string of ordered add and delete operations given in the diagram:

add8	del6	add11	del7	del4	add1
------	------	-------	------	------	------

would be represented by the chromosome string “8 21 11 22 19 1.” If a chromosome were to call for deletion of an antenna from a station with no antennas, that chromosome would be considered to have violated the problem constraints. In this case, or if either of the scheduling constraints defined above were violated, the member was assigned a fitness of $-\infty$. With this chromosome representation, the crossover operation was used to create new offspring. The parent chromosomes were chosen using rank-based selection with a selection ratio of 5 (the best member of the population was five times as likely to be chosen as the worst member).

Ground-Up Representation

The second chromosome formulation using candidate sites incorporated a “ground-up” representation of antenna configurations; no base configuration was required. In this representation, the chromosome was a string of numbers directly representing an antenna assignment to a location. Thus, the chromosome “5 9 10 3 5 7 1” corresponded to a configuration with one antenna each at locations 1, 3, 7, 9, and 10, and two at location 5. The crossover operation was also used with this representation, and the parents were again chosen using rank-based

selection with a selection ratio of 5. The scheduling constraints were implemented in the same way as for the baseline representation.

3.4 Formulation Using All Possible Sites

A more general formulation of the antenna placement problem for genetic search was also developed in this research. This formulation allowed for ground stations at arbitrary locations on the Earth.

Chromosome Representation

The chromosome representation developed to support this all-possible-sites formulation used a single chromosome made up of 30 units of 7 symbols each, for a total length of 210 symbols. Within each unit of 7 symbols, the first three were hexadecimal digits representing the latitude of a ground station, the second three were hexadecimal digits representing the longitude of a ground station, and the last was a decimal digit representing the number of antennas at the specified ground station. This choice of representation thus resulted in a precision of 0.044° in latitude and 0.088° for longitude. Since this chromosome representation required the chromosome to always have a length of 210 symbols, with symbols at different locations having distinctly different meanings (part of a hexadecimal latitude or longitude or a decimal number representing a number of antennas), length-preserving genetic operations were essential in the genetic search.

This third chromosome formulation allowed a ground station to be placed anywhere on the globe, so some means of restricting station locations to acceptable sites was needed. Such a means was provided by the feasible regions defined above in Figure 3. Only locations inside these regions were deemed acceptable; during the genetic optimizations, candidate ground station locations not inside these polygons were eliminated. Thus, every chromosome represented a valid configuration, but did not necessarily contain a full list of 30 viable station locations (even including locations with zero antennas).

Basic Search

As in the candidate-sites formulations, the performance measure (fitness) for this formulation was simply the negative of the cost, where the cost was obtained by summing the regional station location cost for each antenna. The two scheduling constraints described above were again imposed. If either constraint was violated, the fitness was assigned a value of $-\infty$.

Because the chromosomes were so long, a modified crossover operation was used to promote faster exchange of genetic material between chromosomes. In each generation, two parent chromosomes were chosen using rank-based selection with a selection ratio of 3. A fairly low selection ratio

was chosen to promote genetic exchange throughout the population. A random number (between 1 and 10) of sequential crossover operations was then performed, starting with these parents, and only the final resulting pair were retained as offspring.

Search Variation

As a variation, a second genetic search was performed using the same chromosome representation. In this variation, two sets of genetic operations were carried out in each generation. In the first set of operations, two parent chromosomes were selected and put through a random number of crossovers, as described above. For the second set of genetic operations, two new parent chromosomes were selected. Each of these chromosomes was expanded from its hexadecimal/decimal representation into a binary representation, and binary mutation was performed on the results, with a bit flip probability of 0.004 (0.4%). The resulting binary string was converted back into the hexadecimal/decimal representation used in this formulation. The binary mutation served as a mechanism for finer adjustment of the chromosomes. As it was more useful to do fine adjustment on the better chromosomes in the population, the parents for the binary mutation operations were chosen using rank-based selection with a selection ratio of 5, a steeper probability slope than the ratio of 3 used for the crossover operation.

The fitness measure was also modified in this search variation. In the previous search, the scheduling requirements were used as constraints. For this modified search, these requirements were instead incorporated into the performance measure. The fitness thus became the negative of the cost, as before, minus penalty terms if either the percent contacts achieved or the percent satellites with all contacts achieved were below the values given above. Defining a as the percent of contacts achieved, and b as the percent of satellites with all contacts achieved, the penalty terms were:

1. 30 if either $a < 0.995$ or $b < 0.99$
2. $25(0.995 - a)$ if $a < 0.995$
3. $25(0.99 - b)$ if $b < 0.99$

These penalty terms were additive; thus, if both $a < 0.995$ and $b < 0.99$, the total penalty would be $30 + 25(0.995 - a) + 25(0.99 - b)$. As a comparison, the magnitude of the cost of the best configuration achieved by this variation (described in Section 4) was 12.45. By incorporating the scheduling requirements as penalties instead of hard constraints, a wider variety of chromosomes may be retained in the population, promoting faster population evolution.

Table 3. Number of antennas at each site in the best configuration at every 2000 generations of the baseline search, compared to the base case. In generations in which several members of the population share the best fitness, only one member is shown.

Site	Base	Gen 0	Gen 2000	Gen 4000	Gen 6000	Gen 8000	Gen 10000	Gen 12000	Gen 14000	Gen 16000	Gen 18000	Gen 20000
TTS	4	3	2	1	2	1	1	1	1	2	2	2
FTS	4	4	1	1	1	1	1	1	1	1	1	1
TCS	2	1	2	2	1	1	1	1	1			
DGS	1	1		1	1	1	1	1	1			
GTS	2	1	2	2	2	2	2	2	2	2	2	2
HTS	2	2	2	2	1	1	1	1	1	2	2	2
VTs	3	2				1	1	1	1			
CTS	1	1		1	1					1	1	1
NHS	3	2	3	2	2	2	2	2	2	2	2	2
Det. Alfa (ME)	1	1		1								
Kwajalein	2	1										
Argentia	1	1				1	1	1	1	1	1	1
Ascension	1		1		1	1	1	1	1	1	1	1
Onizuka	1	1	1	1	1	1	1	1	1	1	1	1
Camp Parks	1		1									
Number of Antennas	29	21	15	14	13	13	13	13	13	13	13	13
Number of Sites	15	13	9	10	10	11	11	11	11	9	9	9
Cost	35.05	24.80	18.20	16.60	16.40	15.75	15.75	15.75	15.75	15.65	15.65	15.65
% Contacts Satisfied	1.0000	0.9996	0.9993	0.9996	0.9989	0.9993	0.9993	0.9993	0.9993	0.9993	0.9993	0.9993
% Satellites with All Contacts Satisfied	1.0000	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909

4. Placement Results

4.1 Formulation Using Candidate Sites

Baseline Representation

The initial population used for the baseline search was selected largely randomly (by hand), but such that all possible add and delete operations were represented. The best member in the initial population met the constraints using 21 antennas, with a fitness of -24.8 . From this starting point, the genetic search was run for 20,000 generations, during which the population size was kept at 500 by deleting the worst members when the population became too large. After the 20,000 generations, several solutions using only 13 antennas had been found. The best of these configurations had a cost of 15.65, or a fitness of -15.65 . A comparison of the best configuration at every 2000 generations with the base configuration is shown in Table 3.

A map comparing the best configuration in the final evolved population with the baseline configuration is shown in Figure 5. The number on the upper left in each box indicates the number of antennas at the location in the baseline configuration, and the number on the lower right indicates the number in the final configuration obtained by genetic search. This configuration actually appeared eight times in the final population, represented by four distinct chromosomes.

Ground-Up Representation

The initial population for the ground-up search was again selected largely randomly, while ensuring that all fifteen available locations were represented. The best

member in the initial population met the scheduling constraints using 15 antennas, with a cost of 17.8, or a corresponding fitness of -17.8 . From there, the search was run for 10,800 generations, with the population size kept at 500 by deleting the worst members when the population became too large. The ground-up search was somewhat more efficient than the baseline search, as every new chromosome generated in the ground-up search represented a valid antenna configuration, while with the baseline configuration, a new chromosome sometimes included an attempt to delete a nonexistent antenna, and thus was rendered invalid. The search with the baseline representation therefore wasted a large number of offspring, which, on the other hand, never had to be evaluated through scheduling. The ground-up search thus took computationally longer per generation, on average, but was more efficient in terms of number of generations.

After the 10,800 generations, several solutions using only 13 antennas had been found. Of these, four chromosomes in the population shared the best fitness value of -15.45 . Three of these chromosomes represented distinct configurations. A comparison of the best configuration at every 1000 generations with the base configuration is shown in Table 4. The three configurations in the final population with the fitness of -15.45 are shown in Figure 6. The number on the upper left in each box represents the number of antennas at that location in the base configuration, and the other three figures indicate the number of antennas at the site in the three best chromosomes in the final population.

Table 5 gives a comparison of the final configuration from the baseline search, the three final

Table 4. Number of antennas at each site in the best configuration at every 1000 generations with the ground-up search, compared to the base case. In generations in which several members of the population share the best fitness, only one member is shown.

Site	Base	Gen 0	Gen 1000	Gen 2000	Gen 3000	Gen 4000	Gen 5000	Gen 6000	Gen 7000	Gen 8000	Gen 9000	Gen 10000
TTS	4	1	1	1	2	2	1	1	1	1	1	1
FTS	4	1	1	2	2	2	2	2	2	2	2	1
TCS	2	1	1		1	1	1	1				1
DGS	1	1	1	1	1	1	1	1	1	1	1	1
GTS	2	1	1	2	1	1	2	2	1	1	1	1
HTS	2	1	1	2	1	1	1	1	1	1	1	2
VTS	3	1	1									
CTS	1	1	1		1	1	1	1				
NHS	3	1	1		1	1	1	1		2	2	1
Det. Alfa (ME)	1	1	1	1	1	1	1	1	1			1
Kwajalein	2	1	2		1	1	1	1				1
Argentina	1	1	1	2	1	1	1	1	3	2	2	2
Ascension	1	1	1	2					2	2	2	
Onizuka	1	1		1					1	1	1	1
Camp Parks	1	1										
Number of Antennas	29	15	14	14	13	13	13	13	13	13	13	13
Number of Sites	15	15	13	9	11	11	11	11	9	9	9	11
Cost	35.05	17.80	17.25	16.90	16.10	16.10	15.80	15.80	15.55	15.50	15.50	15.45
% Contacts Satisfied	1.0000	0.9993	0.9986	0.9989	0.9996	0.9996	0.9993	0.9993	0.9989	0.9989	0.9989	1.0000
% Satellites with All Contacts Satisfied	1.0000	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909	1.0000

configurations from the ground-up search, and the base configuration. On the whole, the ground-up and the baseline searches performed similarly, though the ground-up search did achieve slightly better fitness values. The differences in the number of ground stations used and in the scheduling percentages achieved, though interesting, are not relevant for evaluating the search algorithms, as they were not incorporated into the performance index (fitness).

4.2 Formulation Using All Feasible Sites

Basic Search

The initial population for the all-feasible-sites search consisted of two sets of members. The first set consisted of 30 chromosomes, each of which contained, somewhere in the chromosome, one of 30 current ground station locations. That location was assigned the number of antennas it currently possesses. All of the other 29 locations represented by the chromosome were randomly generated. The first chromosome contained an existing location in the first unit of the chromosome (the first 7 symbols), the second chromosome contained an existing location in the second unit, and so on. Thus, each of the 30 units in the chromosome representation contained a different existing location somewhere in the initial population. As an example, the first chromosome in this part of the initial population was:

B 7 4 3 5 7 1 E 7 F 1 7 2 0 B A 1 3 F A 1 A 9 A 8 E 4 3 B D
C 2 1 1 1 3 0 C 0 7 5 5 F 0 B 1 C F 1 A C A 6 6 7 4 7 5 8 8 0
B 2 A F 7 A 1 B 1 D C C F D D 5 9 1 A 3 1 F 4 8 6 C C A 5 1
0 9 3 D C 4 4 D D 0 C 8 4 0 1 6 F 2 8 5 4 D 1 2 E 8 A 1 E A 9
7 F A 1 E B 8 0 C 9 5 E C A 1 9 8 4 E A 3 3 2 5 5 9 3 7 8 8 4
1 A 7 1 F A D 5 D 5 9 A B E 4 2 5 C 4 5 3 3 9 3 A A 3 E 3 4
5 8 4 2 D 1 A 1 A A D D 0 3 A E C 2 B 5 6 C 9 5 4 9 1

The first seven symbols of this chromosome, ‘B 7 4 3 5 7 1’, represent one antenna located at Colorado Tracking Station.

The second set of members of the initial population was included to ensure that, somewhere in the population, every possible value was present at every location within the chromosome. This set contained sixteen chromosomes composed of a single repeated digit (0 through F). Only digits between 0 and 5 were included for the number of antennas, however; for other digits a zero was substituted.

From this initial population, the genetic search was run for 32,800 generations. The population size was kept at 2000 by deleting the worst members if the population got too large. The evolution of the population is shown in Table 6. For the first several thousand generations, none of the members of the population met the scheduling constraints. After the 32,800 generations, the best member of the population represented a configuration with 11 antennas at 4 ground stations, with a cost of 12.45. This configuration, shown in Figure 7, allowed 99.57% of contacts to be scheduled, with 99.09% (all but one) of the satellites having all contacts scheduled.

Table 5. Comparison of the results from the baseline and ground-up genetic search formulations. The final generation of the ground-up search contained three configurations with the same fitness.

Site	Base	Baseline gen 20000	Ground-up gen 10800		
TTS	4	2	1	1	1
FTS	4	1	1	2	2
TCS	2		1	1	1
DGS	1		1	1	1
GTS	2	2	1	1	1
HTS	2	2	2	1	1
VTs	3				
CTS	1	1			1
NHS	3	2	1	1	1
Det. Alfa (ME)	1		1	1	1
Kwajalein	2		1	1	1
Argentina	1	1	2	2	2
Ascension	1	1			
Onizuka	1	1	1	1	
Camp Parks	1				
Number of Antennas	29	13	13	13	13
Number of Sites	15	9	11	11	11
Cost	35.05	15.65	15.45	15.45	15.45
% Contacts Satisfied	1.0000	0.9993	1.0000	0.9993	0.9993
% Satellites with All Contacts Satisfied	1.0000	0.9909	1.0000	0.9909	0.9909

Table 6. Best fitness in the population after every 2000 generations, using the basic all-feasible-sites search.

Generations	Best Fitness
0	-∞
2000	-∞
4000	-∞
6000	-22.65
8000	-22.65
10000	-17.85
12000	-17.85
14000	-17.65
16000	-14.80
18000	-14.45
20000	-14.45
22000	-13.55
24000	-13.55
26000	-12.50
28000	-12.50
30000	-12.45
32000	-12.45

Search Variation

The variation on the all-feasible-sites search, described in Section 3.4, was initiated using the population that had been obtained after 1000 generations of the previous search. The fitness of this population was recalculated using the modified fitness measure. From that starting point, the modified search was run for 9900 generations (comparable to 19,800

generations of the previous search, as twice as many new chromosomes were being created in each generation). Again, the population size was kept at 2000. The evolution of the population during this modified search is shown in Table 7. Note that by 1000 generations, the best member in the population met the scheduling constraints. After the 9900 generations with the new search, the best fitness in the population was again -12.45. One typical configuration with this fitness had 11 antennas and 5 ground sites. This configuration is shown in Figure 8. Though the best fitness obtained in this search was identical to that found in the previous search using all possible sites, this configuration achieved slightly better scheduling performance, with 100% of contacts scheduled. This difference, however, was not captured in the fitness measure, and was therefore not selected for during the search.

The modified search appeared to converge somewhat faster than the original search, even taking into account that each generation of the variant search was equivalent to two of the original. It is difficult to say whether either the reformulation of the constraints as penalty functions or the addition of binary mutation caused this effect, or whether it was simply a random occurrence. On the final fitness value achieved, however, the search modifications seemed to have no material effect.




Figure 5. Map showing the locations of the antennas in the base configuration (left in boxes) and in the configuration obtained after 20,000 generations of genetic search using the baseline formulation (right in boxes).




Figure 6. Map showing the locations of the antennas in the base configuration (upper left in boxes) and in the three best configurations obtained after 10,800 generations of genetic search with the ground-up formulation.




Figure 7. Ground station configuration after 32,800 generations using the all-feasible-sites formulation. The number of antennas at each ground site is indicated next to the $\times$ marking the site location.




Figure 8. Sample ground station configuration after 9900 generations using the modified search with the all-feasible-sites formulation. The number of antennas at each ground site is indicated next to the $\times$ marking the site location.




Figure 9. Ground site locations appearing in the 219 members of the final population (of the modified search) with fitness -12.45 . The same antenna assignments to the five sites were retained in all 219 members.

Table 7. Best fitness in the population after every 1000 generations, using the variant all-feasible-sites search. The generation 0 population was obtained by recalculating the fitness of the 1000-generation population from the previous search.

Generations	Best Fitness
0	-46.45313
1000	-23.95
2000	-16.20
3000	-16.20
4000	-15.60
5000	-14.30
6000	-14.25
7000	-12.45
8000	-12.45
9000	-12.45

The final population in the modified genetic search contained 219 members with the same fitness value. These 219 members exhibited some slightly different antenna configurations, some of which showed slightly different scheduling effectiveness. Figure 9 shows all the ground station locations included in these top 219 members. The previous genetic search, without the binary operation, showed only one solution with the fitness of -12.45 in the final population, though this solution was repeated in the population 114 times. The binary mutation did therefore seem to introduce some fine-tuning capability into the search mechanism, even though the fitness, as defined here, was not affected.

The results from all four genetic searches described in this paper are compared with each other in Table 8. The two candidate-sites optimizations, the baseline and the ground-up representations, were able to automatically find solutions with low cost and high scheduling performance. The all-feasible-sites formulation was able to significantly improve on these results; by judiciously positioning the ground sites, it

was able to find solutions requiring only 11 antennas. These results demonstrate that the genetic search technique can be a quite successful automatic means of performing antenna placement optimization.

5. Conclusions and Future Work

The research reported here demonstrates the power of genetic search techniques in specifying near-optimal antenna placement solutions for AFSCN operations. Because genetic search methods can handle any constraint or performance measure, as long as it can be calculated, they work very well for problems like near-optimal antenna placement, which features extremely complex performance requirements and constraints.

During this research, three different chromosome representations were developed and evaluated. In two of these formulations, the antennas were restricted to a predefined set of candidate ground sites. In the third formulation, however, ground stations could be located anywhere inside “feasible” geographical regions. The genetic searches optimized the antenna/ground site configuration on the basis of minimizing the cost of the configuration subject to constraints imposed by scheduling requirements. All of the genetic search formulations were successful in obtaining cost-effective configurations that still were able to satisfy the scheduling requirements. The success of the third formulation is particularly promising, since that representation embodied a very general placement scenario, without any restriction to existing locations.

More complete performance-measure definitions and more sophisticated optimization formulations remain for future work. For example, a more complicated cost structure will likely need to be developed in order to reflect more accurately the true configuration costs. In particular, a cost associated with ground sites (in terms of both number and location) will need to be included along with the per-antenna cost

Table 8. Comparison of placement results.

	base case	baseline search	ground-up search	all sites search	all sites variation
Number of sites	15	9	11	4	5
Number of antennas	29	13	13	11	11
Number of generations	--	20,000	10,800	32,800	9,900
Number of offspring created	--	40,000	21,600	65,600	41,600
% contacts scheduled	1	0.9993	1	0.9957	1
% satellites fully scheduled	1	0.9909	1	0.9909	1
Configuration cost	35.05	15.65	15.45	12.45	12.45

used here. Also, the true cost probably does not increase linearly with the number of antennas at a site; for example, maintenance for two antennas will be less than twice that for one.

The most useful formulation for meeting the goals of the AFSCN resource consolidation analysis will likely be a combination of the formulations developed in this research. Such a formulation would include costs for closing and opening ground sites, so that existing sites factor into the optimization, but would also enable searching for new sites.

References

- [1] *Air Force Satellite Control Network*, <http://www.safaq.hq.af.mil/aqsl/afscn/>, February 2000.
- [2] *50th Space Wing Fact Sheets*, http://www.schriever.af.mil/fact_sheets/index.htm, Schriever Air Force Base, February 2000.
- [3] *Detachment 2, 22nd Space Operations Squadron*, <http://www.diego.af.mil/reef/default.htm>, Diego Garcia Station, February 2000.
- [4] Abbott, R. J., Campbell, M. C., Gallini, T. E., Moran, P. J., and Schulenburg, D.J., "Automated Scheduling in the Satellite Control Network," Aerospace Report No. TOR-96(1571)-1, 30 January 1996.
- [5] "STK User's Manual," Analytical Graphics, Inc., Malvern, PA, 1998.
- [6] D.E. Goldberg, *Genetic Algorithms*, Reading, MA: Addison-Wesley, 1989.
- [7] J.H. Holland, *Adaptation in Natural and Artificial Systems*, Cambridge, MA: The MIT Press, 1993.
- [8] J.R. Koza, *Genetic Programming*, Cambridge, MA: The MIT Press, 1992.
- [9] *Genetic Search Toolbox™ User's Manual*, Optimal Synthesis Inc., 1998-99.
- [10] *Genetic Search Toolbox™ Application Notes*, Optimal Synthesis Inc., 1998-99.