

Minimizing the Cost of Availability of Coverage from a Constellation of Satellites:
Evaluation of Optimization Methods

Clifford Kelley and Maged Dessouky*

Daniel J. Epstein Department of Industrial and Systems Engineering
University of Southern California
Los Angeles, CA
90089-0193

maged@usc.edu

phone: 213-740-4891 fax: 213-740-1120

* Corresponding Author

Abstract

Satellite systems to provide terrestrial surface coverage are currently optimized by independently optimizing parts of the system or by locally starting at a “baseline” configuration and making small changes. Even when using a “baseline” the various components such as constellation configuration, satellite design, and acquisition have generally been optimized separately. A model was developed to minimize the life-cycle cost of the system as a whole with the appropriate linkages between the major parts of the system in both cost and performance. The model consists of a cost model, and a performance model. The cost model includes the cost of purchasing and maintaining an inventory of satellites and launch vehicles, launch operations, and orbit operations. The performance model evaluates the availability over geographic grid averaging areas over a given period of time. The performance of local optimization techniques (Simplex and Complex) was evaluated against global optimization techniques (Simulated Annealing and Genetic Algorithms) for providing availability of coverage. The global optimization techniques performed much better with Genetic Algorithms performing slightly better than Simulated Annealing. The quality of the solutions found was benchmarked against a trade study and local optimization methods, and our proposed global systems approach provided significant cost savings over the existing techniques.

1.0 Introduction

In recent years satellite constellations (a group of satellites working together to provide a service) have been proposed and deployed. They provide communication, navigation, surveillance, and other services. As with any system, the designers are challenged to provide these services at the lowest possible cost. As noted in Blanchard and Fabrycky [1998], many systems have been planned, designed, produced and operated with little concern for economic issues. Often, when cost is considered, it is concentrated on acquisition costs which are only the tip of the iceberg. Blanchard and Fabrycky [1998] also noted that the largest impact on the cost of a system is in the planning and conceptual stage. Because of this, it is critical that a cost and performance model be used during this part of the life of a program. In this paper, we develop a complete life-cycle model for minimizing the cost of availability of coverage from a constellation of satellites.

Satellite systems to provide coverage are currently optimized by independently optimizing parts of the system or by locally starting at a “baseline” configuration and making small changes. Even when using a “baseline” the various components such as constellation configuration, satellite design and acquisition have generally been optimized separately. As illustrated in Wertz and Larson [1996], the satellite design life has been optimized to provide the lowest cost per year for satellite and launch vehicle production. Constellations have been optimized (see also Ballard [1980], Draim [1985, 1986], Lang [1987], Walker [1970,1971,1982, 1984], and Wertz [2001]) by choosing the smallest constellation at the lowest altitude that can provide the required coverage (without failures). Satellite and launch vehicle acquisition are optimized separately based on an estimate of launch demand (based on design life and constellation size).

Although it is recognized that when satellites fail coverage suffers, this service availability has not been treated explicitly in the literature. It is usually stated that to maintain coverage with outages spares will be used. The goals of this research were to develop a model to minimize the life-cycle cost of a satellite system in a global sense with the acquisition, replenishment, and operations costs closely linked to a performance model, which determines availability of service. Further objectives were to determine the best optimization method to use, and to determine how

much cost savings could be achieved over current practices.

2.0 Problem Description

The service examined in this research was the availability of coverage. Coverage is the existence of a clear line of sight from a satellite to a user of the service. In order to determine service availability, it must be clearly defined. A service area or volume must be defined, which specifies where the service will be provided (on the surface of the earth, up to a specified altitude, etc.) along with a criterion. The criterion may be that the availability is an average over the entire service volume, or for every location in the service volume, or there may be a number of defined areas within the service area. For many satellite based systems, grid points across common latitudes are averaged. When the ground track does not repeat for a large number of orbits, the coverage is averaged over the latitude. Since life-cycle cost is used to define the cost of the system, the availability of service is also defined over the life of the system. A set of time points is chosen to determine the service availability for a set of grid averaging areas. Generally these are chosen to span one orbit (if all of the satellite orbits have a common period) or the least common multiple of all of the satellite orbital periods. The time step is chosen to assure that service to a grid point is not skipped which is directly related to the relative angular movement of the satellite and grid point spacing.

The state of the constellation is modeled as a 2 dimensional Markov chain. The number of satellites in the nominal constellation that have failed and the number of spare satellites that have been used defines each state. The rate at which satellites fail moves the chain into higher failure states with lower service availability. The launch rate (which is assumed to be constant for all of the states) moves the chain back into the lower states with higher service availability.

The cost model computes the life-cycle cost of the system and must take into account all of the factors that determine performance. To avoid solutions that make no sense, it is critical that every part of the model that affects availability also has a cost associated with it. A graphical representation of the problem is shown in Figure 1, which includes the supply chain of obtaining

the satellites and launch vehicles, storing them, launching them, operating them to provide service, and disposing of the satellites when they fail. Since coverage is so fundamental, it is also the most studied. Availability of coverage though has not been formally treated in the literature. For many systems, although coverage is necessary, it is not a sufficient condition for service.

Insert Figure 1

Figure 2 illustrates the visibility of the satellites to the surface of the earth. This case is the lowest altitude and smallest number of satellites (in circular orbits) required to achieve single coverage above the horizon worldwide. At this instant in time, everywhere on the earth has at least one satellite in view. A large area of the earth has two satellites in view and some small areas have three satellites in view. Figure 3 shows the same time and same constellation but with one satellite failed. This illustrates how drastically coverage may be affected when even a single satellite fails.

Insert Figure 2

Insert Figure 3

The rest of the paper is organized as follows. The complete life-cycle cost model and the availability of coverage is given in section 3.0. Section 4.0 describes the solution techniques along with the results of the analysis of the techniques. We benchmark our global system approach versus the current practice of locally solving the problem in piecemeal in section 5.0. Conclusions and directions for further research are provided in section 6.0.

3.0 Model Development

The life-cycle model for minimizing the cost of availability of coverage from a constellation of satellites is complex and requires non-linear optimization. The complete mathematical

description of the model is provided in Kelley [2003]. In this section, we provide an overview of the components of the model.

Independent variables:

The independent variables associated with the satellites are the reliability beta of the Weibull reliability distribution, the ratio of design life variance to design life, and the satellite order lead time which includes the development cycle. The independent booster variables are the order cost and lead time. Since there is assumed to be no development cycle for the booster, this is only the production time. The only independent variable associated with the launch is the probability of launch success. The independent variables associated with the service are the geographic grid points, time points and re-phasing delta velocity. The independent variables for the total system are the interest rate, life cycle time, and the required service availability.

Decision variables:

Most of these variables are associated with the satellite constellation such as the orbit altitude, inclination, number of satellites, number of spares, number of planes and phasing of the satellites in each plane. The manufacturing decision variables are satellite design life and reliability. The purchasing variables are the satellite and booster inventory availability. The launch portion determines the mean launch time between the decision to launch and the time the satellite is placed in its assigned slot and starts to provide service.

Computed parameters:

The satellite failure rate and variance are computed based on the combination of a normal distribution for the satellite design life and a Weibull distribution for random failures. The optimal re-order point and order quantity are computed based on the input parameters of order cost, interest rate and the decision variable satellite availability using an inventory model based on Silver [1998]. The booster inventory computations are similar and use the launch vehicle availability from inventory to compute the optimal re-order point and order quantity. The mean re-phase time is computed assuming a fixed delta orbital velocity and is a ratio to the orbit radius. The last variables computed are the total life-cycle cost and service availability.

Computed functions:

These are the probability of occurrence of each constellation state along with the service availability for each of these states. A Markov chain is used to model the failure states of the constellation where the transitions represent the launch of satellites, the failure of satellites and the re-phasing of satellites. The procurement model affects the launch rate because of the probability that both satellites and launch vehicles are available to support a launch. For each of these failure states, the availability of service must be computed for each point in the user service volume. In many cases, the computation time can be greatly reduced due to the limited number of satellites in the constellation that are able to provide service to a user at any given time.

3.1 Life-cycle Cost

The total life-cycle cost C_{total} consists of replenishment, launch, satellite, launch vehicle, and operations costs over the life of the system T_L .

$$C_{total} = (\text{Replenishment} + \text{Launch} + \text{Satellite} + \text{Launch Vehicle} + \text{Operations})T_L$$

The replenishment cost is the cost/month of placing the orders and purchasing the satellites and launch vehicles. Replenishment is linked to the performance model by the fact that if the satellite or launch vehicle is not available from stock, the launch rate will be decreased. Since the satellite and launch vehicle costs are a function of the service and the orbit, only the cost of capital and ordering costs are considered as replenishment costs.

The launch cost is the cost/month of the personnel and facilities required to prepare and launch the satellites and launch vehicles. It links to the performance model by increasing or decreasing the launch rate. At the expense of running more shifts and overtime, the launch rate can be increased to increase the probability of being in a state with fewer satellite failures, which will in turn increase the availability of the service.

The satellite cost is the cost/month for the satellites. The cost of the satellites are a function of the relationship of the number of satellites, the service being provided, the altitude and their failure rate. The failure rate directly affects the state of the constellation and drives the cost of the satellite by requiring more subsystem redundancy, larger solar arrays etc. Since the satellites are assumed to not be an off the shelf design whenever an order is placed the development cost is included.

The operations cost is the cost/month of the infrastructure required to maintain the constellation of satellites. There is a fixed cost for facilities such as office space, antenna sites to communicate with the satellites, and a variable cost of people on hand to monitor and troubleshoot the satellites. Even with a large degree of automation it has been estimated that it takes from 2 to 5 persons per satellite to maintain a constellation.

The launch vehicle cost assumes that unlike when the satellite is ordered the launch vehicle is an off the shelf design so the development cost is not included. The cost is primarily determined by the mass of the satellites carried into orbit, the orbit altitude and inclination of the orbit.

The satellite development, production, and launch vehicle production costs were based on NAFCOM (NASA Air Force COst Model). The rest of the cost parameters were based on cost trends and other data from Wertz and Larson [1999] and Koelle [2000]. While the general trends of the model are based on common practices in government and industry, the model was written to be general enough that an analyst could enter his own CERs (Cost Estimating Relationship) and MERs (Mass Estimating Relationship for estimating the mass of a satellite to determine launch vehicle costs) parameters.

3.2 Availability Model

The availability model consists of two parts. The first part determines the availability as a function of the number of failed satellites. The second part determines the probability the

constellation will have a given number of failed satellites. For coverage a histogram of number of satellites in view is computed for every grid averaging area. A statistical model is then used to compute the availability for every combination of failed satellites.

To determine the probability that a given number of satellites are failed, a two dimensional Markov chain model was embedded into the overall optimization model. Figure 4 shows an example of a 4 satellite constellation in 2 planes with 2 spares. A simple Markov chain for constellation states was described by Durand [1990] and Phlong [1994]. Material on queuing theory and its applications is available from Kleinrock [1975, 1976] and Kao [1997].

Insert Figure 4

The state probabilities of the “nominal” constellation are determined by the position along the horizontal axis. These satellites are used in the computation of service availability. If the required service availability can be achieved with satellites in this “nominal” constellation, the number of satellites that can fail and still achieve the required availability are considered “excess”. The vertical axis is the state probabilities for spare satellites which are not counted for availability. This implies that the “spares” are not positioned to provide service and thus do not contribute to the performance of the constellation. Spares contribute by increasing the probability the constellation has fewer satellites failed because they can be quickly moved to replace a failed satellite. By using a Markov chain, we imply that when the constellation reaches the designed size the launches stop. When spares are available, the preference is to use them. That is, the constellation design, launch, and replenishment strategy allows the constellation to be larger on occasion than needed for basic (no satellites failed) service in order to have satellites available to handle outages. An additional assumption made in this model is that the satellites are not moved once they are placed into a particular orbital “slot” and spares are moved only once during their lifetime.

As seen in figure 4 although many transitions are conceivable some are not desirable. For example when in state 11 through 14, a satellite could be launched. But, since none of the spares

have been used, it is much faster to re-phase a spare. This same philosophy is taken for states 6 through 9. If a spare is available in that plane, then it will be re-phased. Otherwise, a launch will occur. In this example, the probability is 50% for each of these transitions because we have 2 planes and 2 spares. The other strategy can be seen in states 5 and 10. Launches to place spares only occur when the nominal constellation is full.

Another way to reduce cost is to launch more than one satellite per launch vehicle. The cost of launching multiple satellites on a larger launch vehicle is less than the cost of single launches on a smaller launch vehicle. The model assumes that multiple launched satellites cannot be placed into different planes. If this is not the case, it could be accounted for by increasing the launch rate. As shown in figure 5 this complicates the Markov chain and changes some of the launch strategies. Figure 5 uses the same constellation as in figure 4 which is a 4 satellite constellation in 2 planes with 2 spares. In this case when we launch nominal satellites we skip a state since 2 satellites are on board. This dual launch also allows us to transition from state 9 to 15 which launches a spare and a nominal satellite into the same plane. Again, similar to figure 4 there are no launch transitions from state 11 through 14. Because a dual launch would place 2 spares into the same plane from state 5 it is excluded. For this particular constellation, this strategy activates the spares only after failures have occurred.

Insert Figure 5

4.0 Solution Techniques

Because both the cost and availability functions are non-linear with continuous and integer variables, robust solution techniques are required to solve the model. As shown in Kelley [2003] the problem can have multiple local minimums and therefore global optimization techniques are expected to perform better than local techniques. Two general global optimization techniques were used: Simulated Annealing (Aarts [1989]) and Genetic Algorithms (Chambers [2001], Mitchell [1999], Goldberg [1989]). Eight variations of simulated annealing and eight variations of genetic algorithms were evaluated. As with any heuristic method there is no guarantee that they will find the best global optimum.

The simulated annealing method had three user-specified settings. Each setting had two options. The first setting determined the choice of using either the classical or wormhole method. The classical simulated annealing method takes a “seed” solution and makes small changes to it. If the new solution meets the required availability and is less expensive, it is accepted. If it does not meet the required availability or is more expensive, the solution may be accepted depending on the quality of the solution and the current “temperature” level. At the start of the simulation, the temperature is set to a high value, and virtually any solution is accepted. As the simulation progresses, the temperature is reduced and increasingly only better solutions are accepted. If the number of solutions tried is large enough and the cooling is slow enough, the model will find a good solution. Alternatively, the wormhole method only accepts better solutions but creates solutions in a large area of the solution space around the “seed” solution. Thus it can take a “wormhole” path to a new solution without moving through the points in between. As the simulation progresses, the area of the solution space of possible moves is contracted. If the number of solutions tried is large enough and the contraction rate of the wormhole area is slow enough, the model will find a good solution.

The second setting determines the availability computation. It has two options. It can be computed either as a deterministic or stochastic value. Although the stochastic availability can be computed much faster, the cooling rate or contraction rate must be slower due to the uncertainty in the result.

The third setting sets the solution space to be either partitioned or non-partitioned. A partitioned solution space is used when the number of variables in the solution is allowed to change. Constellations can be designed to have uniform or non-uniform spacing. Uniformly spaced constellations have 5 degrees of freedom, orbital altitude, inclination, number of satellites, number of planes (each plane equal), and plane-to-plane phasing (satellite-satellite spacing equal). The orbital eccentricity and argument of perigee are set to zero. When any of these constraints is relaxed, the number of degrees of freedom can be as large as six times the number of satellites. Since this analysis used uniform constellation, the partitioned option was not used.

The genetic algorithm method had two settings. One setting had four options while the other had two. Genetic algorithms optimize by mimicking biological organism evolution. An initial population of solutions is created from a “seed” solution. Each member is assigned a “fitness” which is based on the cost of the solution when it meets the availability constraint, and a penalty cost is added when it does not meet the constraint. The solutions can “mate” and produce an offspring that is similar to its “parents”. In addition random “mutations” are introduced to avoid the population from being captured by local optima. With a large enough population and many generations, the solutions will tend to cluster around a good solution. The first setting selects the method used to decide which solutions can “mate” and produce an offspring. The choices are Roulette, Elitist, Tournament and Uniform. The Roulette method randomly chooses the two parents in proportion to their fitness. The elitist method allows only a subset of the population (the elite) to mate. The uniform method is the opposite of elitist in that all members have an equal chance to “mate”.

Similar to simulated annealing, the second setting for the genetic algorithm determines the availability computation. Although the stochastic availability can be computed much faster the number of generations used must be larger due to the uncertainty in the result.

Two local optimization techniques were also used: the Nelder-Mead simplex method described in Press [1997] and Bazaraa [1993], and the Complex method described by Gupta [1974]. In the simplex method a number of closely spaced separate solutions (one more than the dimension of the problem) are used to find the best direction for the “simplex” to move in search of a better solution. The complex method uses at least two more solutions than the dimension of the problem. The solutions initially span the solution space and gradually converge on the solution centroid by successively moving the worst solution toward the current centroid.

To determine the best solution procedure, the cost of providing coverage to every location on the earth of 1 satellite greater than or equal to 99% of the time with a grazing angle of 0.0 degrees above the local horizon was minimized. The problem has a total of 12 dimensions. The

constellation was constrained to a uniform configuration, and the satellites were constrained to all be at the same altitude and inclination. While satellites at different altitudes and inclinations could be used, the earth's gravitational anomalies would cause the constellation configuration to be unstable, and a large amount of fuel would be required to maintain it.

The dimensions are:

1. Number of satellites (integer)
2. Number of planes (integer, common divisor of number of satellites)
3. Phasing between planes (integer, 0 to number of planes-1)
4. Orbit Semi-major axis (km)
5. Inclination (deg)
6. Satellite design life (months)
7. Satellite reliability (probability)
8. Probability of satellite available for launch from inventory
9. Probability of launch vehicle available for launch from inventory.
10. Number of spare satellites on orbit (integer)
11. Number of satellites launched on each launch vehicle (integer)
12. Mean time to launch a satellite(s) (months)

A total of 30 runs with different random number seeds were run for each method.

4.1 Results

Tables I and II summarize the results. Since the constellations were constrained to uniform configurations, the problem did not require partitioning so it was not one of the chosen simulated annealing methods. Table I compares the performance of each method. Since the computation of service availability and cost dominate the time to run the optimization the number of function evaluations indicates the relative amount of computer time to arrive at an answer. As expected, for a problem with multiple local minima the Simplex method while very fast, performed poorly,

with the best solution almost 30% higher than the global methods. The Complex method did better than the Simplex method (and faster than Simulated Annealing and Genetic Algorithms). While searching over more of the trade space many of the 30 runs would get “stuck” in local solutions and did not show a consistent trend of decreasing cost with the number of evaluations. Although the stochastic performance computations appeared interesting, the extra overhead of more evaluations to compensate for the uncertainty of the performance appeared to outweigh any computation time advantages it had. Both Simulated Annealing and Genetic Algorithms performed well. All of the Simulated Annealing techniques showed a consistent trend of improving the solution with a larger number of function evaluations while some of the Genetic Algorithm techniques got “stuck” in higher cost solutions. In general, when the Genetic Algorithm converged, it consistently appeared to converge somewhat faster than the Simulated Annealing techniques and also was able to find better solutions. While one of the features of the simulation was to make a number of optimization techniques accessible to the analyst, overall the Genetic Algorithms are recommended.

Table II lists the resulting decision variables of the best solutions found. The best every grid point is taken from the optimization comparison and was generated by the Genetic Algorithm Elitist deterministic availability method. Using this same method the problem was also optimized for service provided at every latitude. The solutions have a number of similarities and a number of differences. As expected service provided at every latitude is less expensive than at every grid point since service is averaged across common latitudes. The two solutions are similar in their inventory policy, satellite design parameters and inclination. The major differences are in their replenishment policies and a corresponding difference in their constellations. The every grid point solution used spares and reduced cost with multiple satellite launches and a small number of planes (2). The every latitude solution was at the other extreme with no spares, a single satellite launch and the same number of planes as satellites (5). In addition the every latitude solution used a higher orbit.

5.0 Benchmarking

In order to determine what advantages this method provides, it was benchmarked against current design practices. In current practice two general methods are used, the separate trade study method and the local optimization method.

In the separate trade study method, it is assumed that the parts of the system are orthogonal and thus may be optimized separately. In the coverage problem it is assumed that the satellite costs, constellation costs, and replenishment costs may be optimized independently and when combined produce the lowest cost solution. In the local optimization method, solutions from the trade study method or expert opinions of possible combinations of options are combined and then used as the “seed” for a local optimization method.

The separate trade study method used constellations from 5 to 9 satellites ranging from 2 to 6 planes. This range was chosen based on providing combinations that would range from using excess satellites, spares or combinations thereof. The mean launch time, reliability, design life and inventory policies were treated as given based on other analysis.

The benchmark solution results are presented in table III. The resulting constellations spanned the range from 5 satellites in 5 planes to 6 satellites in 2 planes and 7 satellites in 7 planes. In each case, these constellations are very similar to the Walker starting points. Even though the constellations have similar configurations to the global optimization results their inclinations and orbital altitudes are significantly different. This is likely due to the inclusion of availability of service. Since these constellations were designed for 100% coverage with no satellite failures one would not expect to get the same answer. The fact that the local optimization methods didn't find answers as good as the global methods illustrates that many local optima exist.

The benchmarking methods might provide starting points close enough to the global optimum if the problem were intuitive and the number of dimensions is small. The resulting global solutions even though they were uniform configurations were different in both inclination and altitude than what a simple optimization for altitude and single satellite coverage would suggest. This indicates that when outages are considered a different inclination and altitude has a significant, unanticipated effect on the availability of coverage and the constellation design is not orthogonal

to the rest of the system. As shown in tables II and III the globally optimized solutions provided life-cycle cost reductions of 64 to 235 million dollars or approximately 5% to 17% over the benchmark methods.

6.0 Conclusions and Directions for Further Research

A model that tightly integrates the computation of cost and availability of coverage was developed for optimizing a system consisting of a constellation of satellites. The problem is non-linear and can have multiple local minimums. A number of local and global optimization methods were tried and compared. A number of variations of both simulated annealing and genetic algorithms as well as simplex and complex methods were evaluated. Both Simulated Annealing and Genetic Algorithms consistently converged to good answers with Genetic Algorithms appearing to converge slightly faster and producing slightly better answers.

Further areas for research include more sophisticated Markov chain models for the constellation state; adding the cost of the user equipment and expanding the optimization to a variable number of dimensions. When non-uniform constellations are considered the number of dimensions of the problem varies. The structure of the Genetic Algorithms would make the optimization very easy to implement where Simulated Annealing requires complex partitioning methods. This is equivalent to a number of species competing for the best solution. In addition, Genetic Algorithms would allow proposed solutions to “seed” the initial population and compete directly with other potential solutions.

This model was developed in response to a government initiative to use cost as an independent variable or CAIV. It would also be useful to provide the option to run the optimization backwards (i.e., to constrain the life-cycle cost and maximize availability). Often when developing a new system the budgeting process forces a design to cost (DTC) methodology. Funding systems on a year to year basis forces the designers to minimize the development costs (generally creating a system with higher life cycle cost). Additionally, studies to determine under what conditions excess satellite are preferred over spares could help analysts reduce the

number of dimensions of the problem.

Acknowledgements

This paper is derived from portions of the Ph.D. dissertation of Kelley [2003]. We would like to thank Dr. Elliot Axelband, Dr. George Freidman, Dr. F. Stan Settles, and Dr. Satish Bukkapatnam of the University of Southern California for their helpful comments in guiding the research.

6.0 References

E. Aarts, and J. Korts, *Simulated annealing and Boltzman machines a stochastic approach to combinatorial optimization and neural computing*, John Wiley & Sons, New York, 1989.

A. H. Ballard, Rosette constellations of earth satellites, *IEEE Transactions on Aerospace and Electronic Systems* 16 (1980), 656-673.

M. S. Bazaraa, H.D. Sherali, and C. M. Shetty, *Nonlinear programming: theory and algorithms*, John Wiley & Sons, Inc, New York, 1993.

D. C. Beste, Design of satellite constellations for optimal continuous coverage, *IEEE Transactions on Aerospace and Electronic Systems* 14 (1978), 466-473.

B. S. Blanchard, and W. J. Fabrycky, *Systems engineering and analysis*, Prentice Hall, New Jersey, 1998.

L. Chambers, *The practical handbook of genetic algorithms applications*, Chapman & Hall/CRC, Boca Raton, Florida, 2001.

J. E. Draim, Three and four-satellite continuous-coverage constellations, *Journal Guidance*, 8 (1985), 725-730.

J. E. Draim, A common period four-satellite continuous global coverage constellation, *1986 AIAA/AAS Astrodynamics Conference Proceedings*, 1986, pp. 112-120.

J-M., Durand, and A. Caseau, GPS availability, part II evaluation of state probabilities for 21 satellite and 24 satellite constellations, *Navigation: Journal of the Institute of Navigation*, 37 (1990), 285-296.

D. E. Goldberg, *Genetic algorithms in search, optimization & machine learning*, Addison-Wesley, New York, 1989.

S. K. Gupta, Optimal stabilization policy under price controls: an application of optimal control techniques to economic policy analysis, Ph.D. Dissertation at Washington University, Sever Institute of Technology, St. Louis, MO., 1974.

E. Kao, *An introduction to stochastic processes*, Duxbury Press, Belmont, California, 1997.

C. W. Kelley, "Minimizing the cost of availability of services from a constellation of satellites, Ph.D. dissertation at the University of Southern California, Los Angeles, CA, 2003.

L. Kleinrock, *Queueing systems volume I: theory*, John Wiley & Sons, New York, 1975.

L. Kleinrock, *Queueing systems volume II: computer applications*, John Wiley & Sons, New York, 1976.

D. E. Koelle, Handbook of cost engineering for space transportation systems with TRANSCOST 7.0 statistical-analytical model for cost estimation and economical optimization of launch vehicles", Report No. TCS-TR-168/00, TCS – TransCostSystems, Liebigweg, Germany, 2000.

T. J. Lang, Symmetric circular orbit satellite constellations for continuous global coverage, *AAS* 87 (1987), 1111-1132.

M. Mitchell, *An introduction to genetic algorithms*, MIT Press, Massachusetts Institute of Technology, Boston, MA, 1999.

W. S. Phlong, and B. D. Elrod, Availability characteristics of GPS and augmentation alternatives, *Navigation: Journal of the Institute of Navigation* 40 (1994), 409-428.

W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical recipes in C*, Cambridge University Press, Melbourne, 1997.

E. A. Silver, D. F. Pyke, R. Peterson, *Inventory management and production planning and scheduling*, 3rd edition, John Wiley & Sons, New York, 1998.

J. G. Walker, Circular orbit patterns providing continuous whole earth coverage, RAE Technical Report 70211, 1970.

J. G. Walker, Some circular orbit patterns providing continuous whole earth coverage, *Journal of the British Interplanetary Society*, 24 (1971), 369-384.

J. G. Walker, Coverage predictions and selection criteria for satellite constellations, Royal Aircraft Establishment, Technical Report 82116, 1982.

J. G. Walker, Satellite constellations, *Journal of the British Interplanetary Society*, 37 (1984), 559-572.

J. R. Wertz, and W. J. Larson, *Space mission analysis and design*, 2nd Edition, Microcosm Press, Torrance, CA, Kluwer Academic Publishers, Dordrecht, the Netherlands, 1996.

J. R. Wertz, and W. J., *Reducing space mission cost*, Microcosm Press, El Segundo, CA, Kluwer Academic Publishers, Dordrecht, the Netherlands, 1999.

J. R. Wertz, *Mission geometry; orbit and constellation design and management*, Microcosm Press, El Segundo, CA, Kluwer Academic Publishers, Dordrecht, the Netherlands, 2001.

Biographical Sketch

Dr. Clifford Kelley holds a B.S. in Physics, a B.S. in Electrical Engineering, an M.S. in Systems Engineering and recently completed his Ph. D. in Industrial and Systems Engineering at the University of Southern California. He started his career at Boeing working on computer modeling of missile guidance systems, high energy lasers and synthetic aperture radar. Later he worked for Rockwell International on the B-1B program in the Infrared and Radar countermeasures area, the communications and navigation system and the computer system. He currently works for Boeing as a Senior Systems Engineer on the GPS IIA, GPS IIF and GPS III programs.

Dr. Maged M. Dessouky is an Associate Professor of Industrial and Systems Engineering, University of Southern California. Prior to joining USC, Dr. Dessouky was employed at Hewlett-Packard (Systems Analyst), Bellcore (Member of Technical Staff), and Pritsker Corporation (Senior Systems Analyst). Dr. Dessouky has an in-depth theoretical and practical understanding of models and heuristic methods for high technology system optimization. He is currently an investigator on several transit-related projects, including the NSF-sponsored Real-time Scheduling of Demand Responsive Systems Project. He received his Ph.D. in Industrial Engineering from University of California, Berkeley, and M.S. and B.S. degrees from Purdue University. He is the Area Editor of Planning and Scheduling for the *International Journal of Computers & Industrial Engineering*.

Table I. Computational Results of the Different Solution Techniques

Optimization Method			Number of Function Evaluations to Reach 7.5% of Best	Average Number of Function Evaluations	Best Cost Found \$ (10 ⁶)
Simulated Annealing					
Classic	Deterministic	No Partition	3842	13235.8	1415
Classic	Deterministic	Partitioned	n/a		
Classic	Stochastic	No Partition	>36000	16116.8	1410
Classic	Stochastic	Partitioned	n/a		
Wormhole	Deterministic	No Partition	5047	10422.8	1424
Wormhole	Deterministic	Partitioned	n/a		
Wormhole	Stochastic	No Partition	11000	9784.9	1415
Wormhole	Stochastic	Partitioned	n/a		
Genetic Algorithms					
Roulette	Deterministic		3027	26684.4	1400
Elitist	Deterministic	25%	**	32501.4	1396
Tournament	Deterministic	5 in group	**	14424.9	1403
Uniform	Deterministic		2810	16107.9	1418
Roulette	Stochastic		2840	28686.8	1417
Elitist	Stochastic	25%	**	31661.3	1396
Tournament	Stochastic	5 in group	**	13686.1	1397
Uniform	Stochastic		**	15238.5	1417
Nelder-Mead Simplex			--	1194.3	1795
Complex			**	5859.1	1424

-- did not reach 7.5% of Best

** multiple solutions (no consistent trend)

Table II. Coverage Global Optimization Results

		Every Grid Point Best	Every Latitude Best
Constellation	Nsats	6	5
	Nplanes	2	5
	Nphase	0	3
	Inclination	62.6 deg	62.1 deg
	Sma	26086 km	36531 km
Replenishment	N spares	4	0
	N sats/launch	3	1
	Mean launch time	6.9 mo	4.5 mo
Satellite Design	Reliability	59%	62%
	Design Life	159 mo	169 mo
Inventory Policy	Psv	96%	94%
	Plv	87%	88%
Cost \$ (10 ⁶)		1396	1376

Table III. Coverage Benchmark Results

		Every Grid Point		Every Latitude	
		Local Optimization	Separate Optimization	Local Optimization	Separate Optimization
Constellation	Nsats	7	6	6	5
	Nplanes	7	2	2	5
	Nphase	2	0	0	1
	Inclination	61.6 deg	52.2 deg	55.6 deg	43.7 deg
	Sma	26276 km	16129 km	16261 km	17923 km
Replenishment	N spares	6	6	5	10
	N sats/launch	2	4	3	3
	Mean launch	5.4 mo	3.0 mo	9.2 mo	3.0 mo
Satellite Design	Reliability	61%	65%	63%	65%
	Design Life	159 mo	160 mo	159 mo	160 mo
Inventory Policy	Psv	95%	95%	97%	95%
	Plv	93%	95%	88%	95%
Cost \$ (10 ⁶)		1523	1629	1440	1611

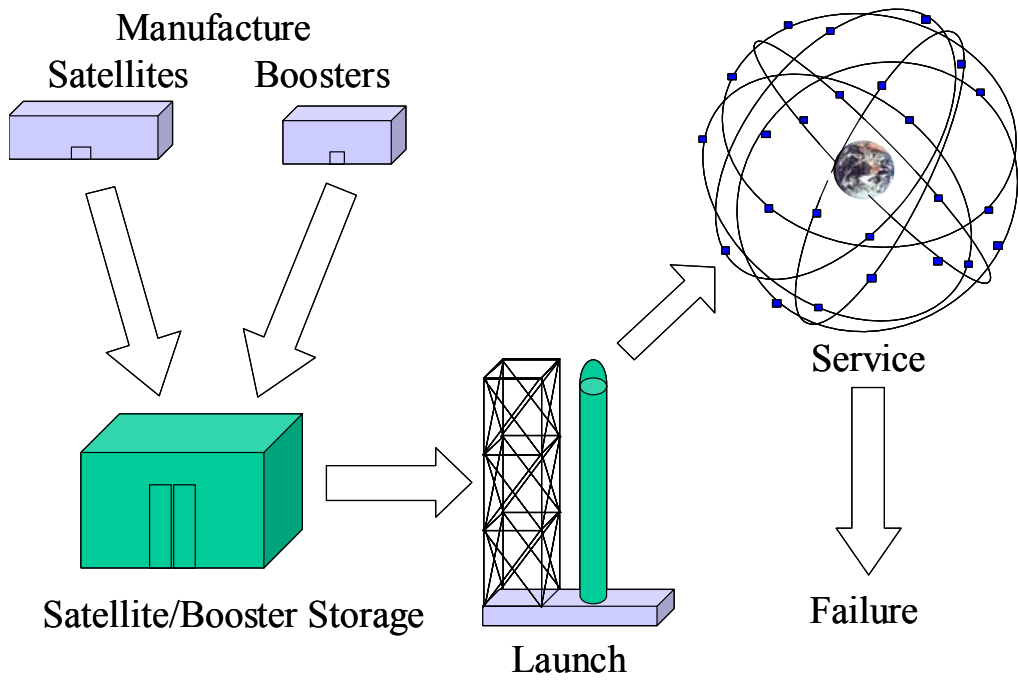


Figure 1
Schematic of Providing Service from a Constellation of Satellites

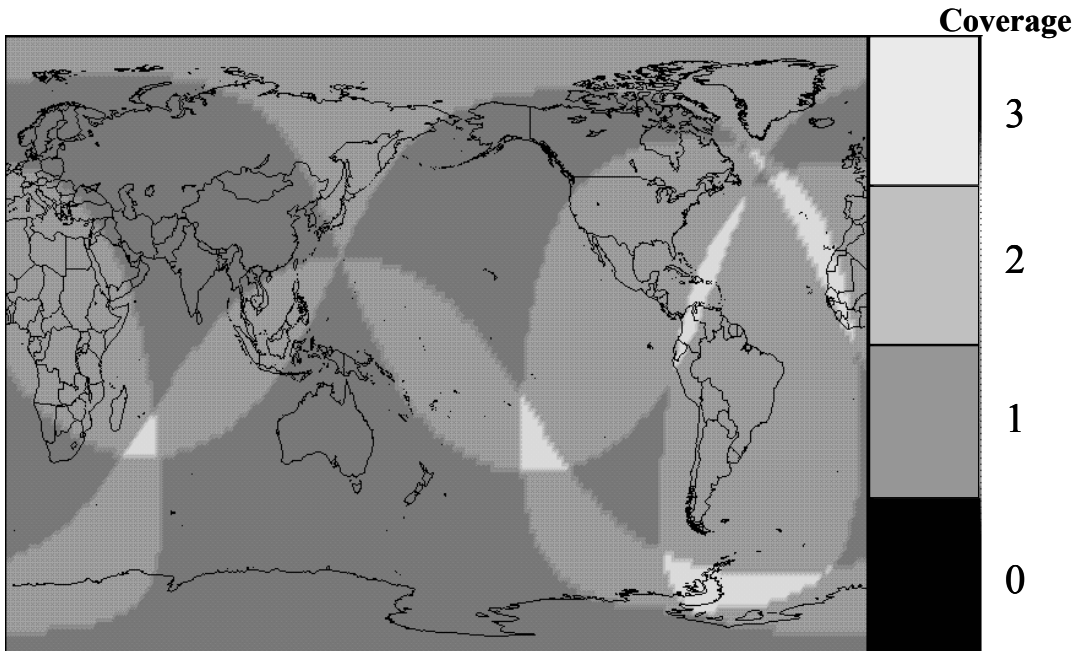


Figure 2
Coverage Plot for 5/5/1 No Satellites Failed




Figure 3
Coverage plot for 5/5/1 with 1 satellite failed

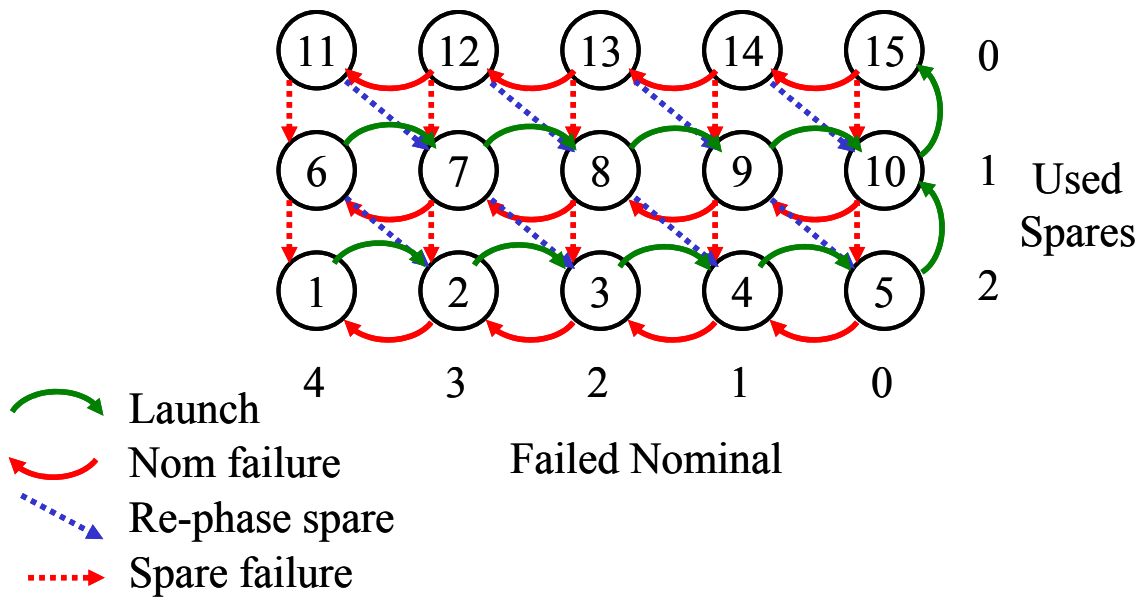


Figure 4
Single Launch Markov Chain




Figure 5
Dual Launch Markov Chain