NEAR EAST UNIVERSITY

Faculty of Engineering

Department of Computer Engineering

Genetic Algorithm Based Optimization

Graduation Project
COM- 400

Student: Yahya Safarini

Supervisor: Assoc. Prof. Dr. Rahib Abiyev

Nicosia - 2003

ACKNOWLEDGEMENTS

ABSTRACT

By increasing complexity of processes, it has become very difficult to control them on the base of traditional methods. In such condition it is necessary to use modern methods for solving these problems. One of such method is global optimization algorithm based on mechanics of natural selection and natural genetics, which is called Genetic Algorithms. In this project the application problems of genetic algorithms for optimization problems, its specific characters and structures are given. The basic genetic operation: Selections, reproduction, crossover and mutation operations are widely described the affectivity of genetic algorithms for optimization problem solving is shown. After the representation of optimizations problem, structural optimization and the finding of optimal solution of quadratic equation are given.

The practical application for selection, reproduction, crossover, and mutation operation are shown. The functional implementation of GA based optimization in MATLAB programming language is considered. Also the multi-modal optimization problem, some methods for global optimization and the application of Niching method for multi-modal optimization are discussed.

TABLE OF CONTENTS

INTRODUCTION

This is an introduction to genetic algorithm methods for optimization. Genetic Algorithms were formally introduced in the United States in the 1970s by John Holland at University of Michigan. The continuing price/performance improvements of computational systems has made them attractive for some types of optimization. In Particular, genetic algorithms work very well on mixed (continuous and discrete), Combinatorial problems. They are less susceptible to getting 'stuck' at local optima than gradient search methods. But they tend to be computationally expensive.

To use a genetic algorithm, you must represent a solution to your problem as a genome (or chromosome). The genetic algorithm then creates a population of solutions and applies genetic operators such as mutation and crossover to evolve the solutions in order to find the best one(s).

This presentation outlines some of the basics of genetic algorithms. The three most important aspects of using genetic algorithms are: (1) definition of the objective function, (2) definition and implementation of the genetic representation, and (3) definition and implementation of the genetic operators. Once these three have been defined, the generic genetic algorithm should work fairly well. Beyond that you can try many different variations to improve performance, find multiple optima (species – if they exist), or parallelize the algorithms.

The genetic algorithm uses stochastic processes, but the result is distinctly non random (better than random).

GENETIC Algorithms are used for a number of different application areas. An example of this would be multidimensional OPTIMIZATION problems in which the character string of the CHROMOSOME can be used to encode the values for the different parameters being optimized.

In practice, therefore, we can implement this genetic model of computation by having arrays of bits or characters to represent the Chromosomes. Simple bit manipulation operations allow the implementation of CROSSOVER, MUTATION and other operations. Although a substantial amount of research has been performed on variable- length strings and other structures, the majority of work with GENETIC Algorithms is focused on fixed-length character strings. We should focus on both this aspect of fixed-length ness and the need to encode the representation of the solution

being sought as a character string, since these are crucial aspects that distinguish GENETIC PROGRAMMING, which does not have a fixed length representation and there is typically no encoding of the problem.

When the GENETIC ALGORITHM is implemented it is usually done in a manner that involves the following cycle: Evaluate the FITNESS of all of the Individuals in the POPULATION. Create a new population by performing operations such as CROSSOVER, fitness-proportionate REPRODUCTION and MUTATION on the individuals whose fitness has just been measured. Discard the old population and iterate using the new population.

One iteration of this loop is referred to as a GENERATION. There is no theoretical reason for this as an implementation model. Indeed, we do not see this punctuated behavior in Populations in nature as a whole, but it is a convenient implementation model.

The first GENERATION (generation 0) of this process operates on a POPULATION of randomly generated Individuals. From there on, the genetic operations, in concert with the FITNESS measure, operate to improve the population.

CHAPTER ONE

WHAT ARE GENETIC ALGORITHMS (GAs)?

1.1 What Are Genetic Algorithms (GAs)?

Genetic Algorithms (GAs) are adaptive heuristic search algorithm based on the evolutionary ideas of natural selection and genetics. As such they represent an intelligent exploitation of a random search used to solve optimization problems. Although randomised, GAs are by no means random, instead they exploit historical information to direct the search into the region of better performance within the search space. The basic techniques of the GAs are designed to simulate processes in natural systems necessary for evolution, specially those follow the principles first laid down by Charles Darwin of "survival of the fittest.". Since in nature, competition among individuals for scanty resources results in the fittest individuals dominating over the weaker ones.

1.2 Defining Genetic Algorithms

What exactly do we mean by the term Genetic Algorithms Goldberg (1989) defines it as:

Genetic algorithms are search algorithms based on the mechanics of natural selection and natural genetics.

Bauer (1993) gives a similar definition in his book:

Genetic algorithms are software, procedures modeled after genetics and evolution.

GAs exploits the idea of the survival of the fittest and an interbreeding population to create a novel and innovative search strategy. A population of strings, representing solutions to a specified problem, is maintained by the GA. The GA then iteratively creates new populations from the old by ranking the strings and interbreeding the fittest to create new strings, which are (hopefully) closer to the optimum solution to the problem at hand. So in each generation, the GA creates a set of strings from the bits and pieces of the previous strings, occasionally adding random new data to keep the population from stagnating. The end result is a search strategy that is tailored for vast, complex, multimodal search spaces.   GAs are a form of randomized search, in that the way in which strings are choisen and combined is a stochastic process. This is a radically different approach to the problem solving methods used by more traditional

algorithms, which tend to be more deterministic in nature, such as the gradient methods used to find minima in graph theory.

The idea of survival of the fittest is of great importance to genetic algorithms. GAs use what is termed as a fitness function in order to select the fittest string that will be used to create new, and conceivably better, populations of strings. The fitness function takes a string and assigns a relative fitness value to the string. The method by which it does this and the nature of the fitness value does not matter. The only thing that the fitness function must do is to rank the strings in some way by producing the fitness value. These values are then used to select the fittest strings. The concept of a fitness function is, in fact, a particular instance of a more general AI concept, the objective function.

1.3 Genetic Algorithms: A Natural Perspective

The population can be simply viewed as a collection of interacting creatures. As each generation of creatures comes and goes, the weaker ones tend to die away without producing children, while the stronger mate, combining attributes of both parents, to produce new, and perhaps unique children to continue the cycle. Occasionally, a mutation creeps into one of the creatures, diversifying the population even more. Remember that in nature, a diverse population within a species tends to allow the species to adapt to it's environment with more ease. The same holds true for genetic algorithms.

1.4 The Iteration Loop of a Basic Genetic Algorithm
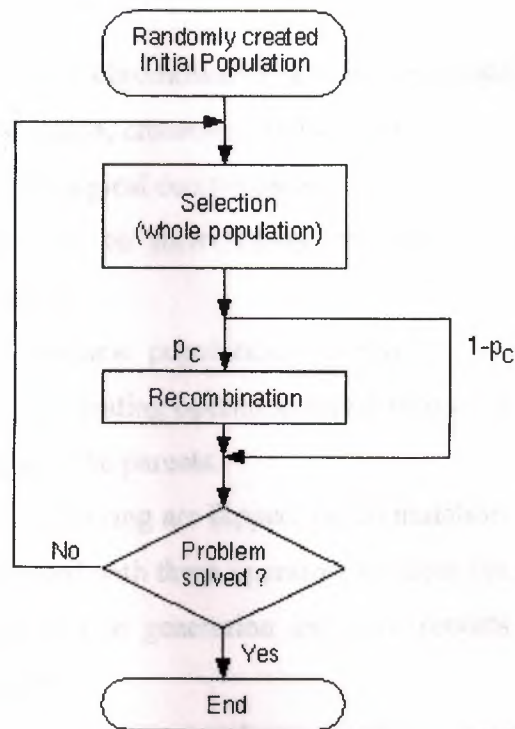
Figure 1.1.

1.5 Biology

Genetic algorithms are used in search and optimization, such as finding the maximum of a function over some domain space.

1. In contrast to deterministic methods like hill climbing or brute force complete enumeration, genetic algorithms use randomization.

2. Points in the domain space of the search, usually real numbers over some range, are encoded as bit strings, called chromosomes.

3. Each bit position in the string is called a gene.

4. Chromosomes may also be composed over some other alphabet than $\{0,1\}$, such as integers or real numbers, particularly if the search domain is multidimensional.

5. GAs are called ``blind'' because they have no knowledge of the problem.

1.6 An Initial Population of Random Bit Strings is Generated

1. The members of this initial population are each evaluated for their fitness or goodness in solving the problem.

5

2. If the problem is to maximize a function f(x) over some range [a,b] of real numbers and if f(x) is nonnegative over the range, then f(x) can be used as the fitness of the bit string encoding the value *x*.

From the initial population of chromosomes, a new population is generated using three genetic operators: reproduction, crossover, and mutation.

1. These are modeled on their biological counterparts.

2. With probabilities proportional to their fitness, members of the population are selected for the new population.

3. Pairs of chromosomes in the new population are chosen at random to exchange genetic material, their bits, in a mating operation called crossover. This produces two new chromosomes that replace the parents.

4. Randomly chosen bits in the offspring are flipped, called mutation.

The new population generated with these operators replaces the old population.

1. The algorithm has performed one generation and then repeats for some specified number of additional generations.

2. The population evolves, containing more and more highly fit chromosomes.

3. When the convergence criterion is reached, such as no significant further increase in the average fitness of the population, the best chromosome produced is decoded into the search space point it represents.

Genetic algorithms work in many situations because of some hand waving called The Schema Theorem.

``Short, low-order, above-average fitness schemata receive exponentially increasing trials in subsequent generations."

1.7 Genetic Algorithm Overview

GOLD optimises the fitness score by using a genetic algorithm.

1. A population of potential solutions (i.e. possible docked orientations of the ligand) is set up at random. Each member of the population is encoded as a chromosome, which contains information about the mapping of ligand H-bond atoms on (complementary) protein H-bond atoms, mapping of hydrophobic points on the ligand onto protein hydrophobic points, and the conformation around flexible ligand bonds and protein OH groups.

2. Each chromosome is assigned a fitness score based on its predicted binding affinity and the chromosomes within the population are ranked according to fitness.

3. The population of chromosomes is iteratively optimised. At each step, a point mutation may occur in a chromosome, or two chromosomes may mate to give a child. The selection of parent chromosomes is biased towards fitter members of the population, i.e. chromosomes corresponding to ligand dockings with good fitness scores.

1.7.1 A Number of Parameters Control The Precise Operation of The Genetic Algorithm, viz.

1. Population size.
2. Selection pressure.
3. Number of operations.
4. Number of islands.
5. Niche size.
6. Operator weights: migrate, mutate, crossover.
7. Annealing parameters: van der Waals, hydrogen bonding.

1.7.1.1 Population Size

1. The genetic algorithm maintains a set of possible solutions to the problem. Each possible solution is known as a chromosome and the set of solutions is termed a population.

2. The variable Population Size (or popsize) is the number of chromosomes in the population. If n_islands is greater than one (i.e. the genetic algorithm is split over two or more islands), pop size is the population on each island.

1.7.1.2 Selection Pressure

1. Each of the genetic operations (crossover, migration, mutation) takes information from parent chromosomes and assembles this information in child chromosomes. The child chromosomes then replace the worst members of the population.

2. The selection of parent chromosomes is biased towards those of high fitness, i.e. a fit chromosome is more likely to be a parent than an unfit one.

3. The selection pressure is defined as the ratio between the probability that the most fit member of the population is selected as a parent to the probability

that an average member is selected as a parent. Too high a selection pressure will result in the population converging too early.

4. For the GOLD docking algorithm, a selection pressure of 1.1 seems appropriate, although 1.125 may be better for library screening since the aim is for faster convergence

1.7.1.3 Number of Operations

1. The genetic algorithm starts off with a random population (each value in every chromosome is set to a random number). Genetic operations (crossover, migration, mutation are then applied iteratively to the population. The parameter Number of Operations (or maxops) is the number of operators that are applied over the course of a GA run.

2. It is the key parameter in determining how long a GOLD run will take.

1.7.1.4 Number of Islands

1. Rather than maintaining a single population, the genetic algorithm can maintain a number of populations that are arranged as a ring of islands. Specifically, the algorithm maintains n_islands populations, each of size popsize.

2. Individuals can migrate between adjacent islands using the migration operator.

3. The effect of n_islands on the efficiency of the genetic algorithm is uncertain.

1.7.1.5 Niche Size

1. Niching is a common technique used in genetic algorithms to preserve diversity within the population.

2. In GOLD, two individuals share the same niche if the rmsd between the coordinates of their donor and acceptor atoms is less than 1.0 Å.

3. When adding a new individual to the population, a count is made of the number of individuals in the population that inhabit the same niche as the new chromosome. If there are more than NicheSize individuals in the niche, then the new individual replaces the worst member of the niche rather than the worst member of the total population.

8

1.7.1.6 Operator Weights: Migrate, Mutate, Crossover

1. The operator weights are the parameters Mutate, Migrate and Crossover (or Pt_cross).

2. They govern the relative frequencies of the three types of operations that can occur during a genetic optimization: point mutation of the chromosome, migration of a population member from one island to another, and crossover (sexual mating) of two chromosomes.

3. Each time the genetic algorithm selects an operator, it does so at random. Any bias in this choice is determined by the operator weights. For example, if Mutate is 40 and Crossover is 10 then, on average, four mutations will be applied for every crossover.

4. The migrate weight should be zero if there is only one island, otherwise migration should occur about 5% of the time.

1.7.1.7 Annealing Parameters: van der Waals, hydrogen bonding

1. The annealing parameters, van der Waals *and* Hydrogen Bonding, allow poor hydrogen bonds to occur at the beginning of a genetic algorithm run, in the expectation that they will evolve to better solutions.

2. At the start of a GOLD run, external van der Waals (vdw) energies are cut off when E_{ij} > van der Waals * k_{ij}, where k_{ij} is the depth of the vdw well between atoms i and j. At the end of the run, the cut-off value is FINISH VDW LINEAR CUTOFF.

3. This allows a few bad bumps to be tolerated at the beginning of the run.

4. Similarlythe parameters Hydrogen Bonding and FINAL_VIRTUAL_PT_MATCH_MAX are used to set starting and finishing values of max_distance (the distance between donor hydrogen and fitting point must be less than max_distance for the bond to count towards the fitness score). This allows poor hydrogen bonds to occur at the beginning of a GA run.

5. Both the vdw and H-bond annealing must be gradual and the population allowed plenty of time to adapt to changes in the fitness function.

1.8 Genetic Operations

In order to improve the current population, genetic algorithms commonly use three different genetic operations. These are reproduction, crossover, and mutation. Both

reproduction and crossover can be viewed as operations that force the population to converge. They do this by promoting genetic qualities that are already present in the population. Conversely, mutation promotes diversity within the population. In general, reproduction is a fitness preserving operation, crossover attempts to use current positive attributes to enhance fitness, and mutation introduces new qualities in an attempt to increase fitness.

1.8.1 Reproduction

The reproduction genetic operation is an asexual operation. Reproduction involves making an exact copy of an individual from the current population into the next generation. The selection of which individuals will be copied into the next generation is done probabilistically based upon relative fitness. Suppose that a gene g exists such that $\forall h \; F(g) \geq F(h)$. Then the reproduction operation will be performed on gene h with a probability $\dfrac{F(h)}{F(g)}$. This selection method ensures all valid genes (i.e. genes that actually solve the problem) have a probability of being chosen for the reproduction operation since $F(h) > 0$ for all genes h that represent a valid solution to the problem that is being solved.

1.8.2 Crossover

The crossover operation is the most important genetic operation. This operation is used to create new individuals by combining the qualities of 2 or more genes (See Fig. 1.2).
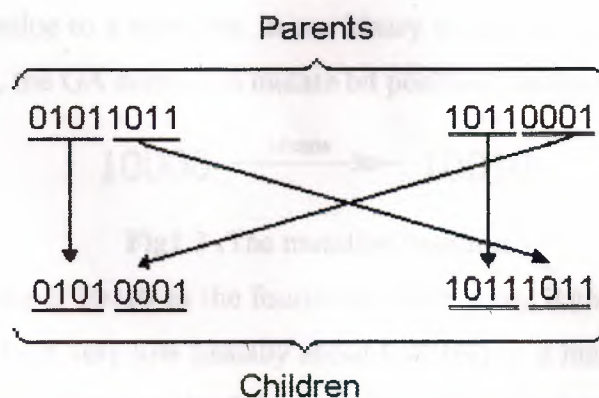
Parents

01011011   10110001

0101 0001   10111011

Children

Figure 1.2. The crossover operation

10

A decision must be made as to which individuals are to be involved in the crossover operation. The method that was used to make this decision in the genetic algorithms under consideration is a form of Boltzmann tournament selection. A Boltzmann tournament proceeds as follows: two genes g and h are selected at random from the current population and are entered into a tournament. If $F(g) > F(h)$ then g wins the tournament, otherwise h wins the tournament. The winner will advance to compete in another tournament with a randomly chosen individual. For the implementation, there were 3 tournaments performed in order to choose each parent. The most general statement of a Boltzmann tournament requires the selection of h to satisfy $|F(g) - F(h)| \geq \phi$, for some ϕ [Mafoud91]. For the implementations, a value of $\phi = 0$ was used.

1.8.3 Mutation

Selection and crossover alone can obviously generate a staggering amount of differing strings. However, depending on the initial population chosen, there may not be enough variety of strings to ensure the GA sees the entire problem space. Or the GA may find itself converging on strings that are not quite close to the optimum it seeks due to a bad initial population.

Some of these problems are overcome by introducing a mutation operator into the GA. The GA has a mutation probability, m, which dictates the frequency at which mutation occurs. Mutation can be performed either during selection or crossover (though crossover is more usual). For each string element in each string in the mating pool, the GA checks to see if it should perform a mutation. If it should, it randomly changes the element value to a new one. In our binary strings, 1s are changed to 0s and 0s to 1s. For example, the GA decides to mutate bit position 4 in the string 10000:

$$10000 \xrightarrow{\text{Mutate}} 10010$$

Fig1.3 .The mutation operator

The resulting string is 10010 as the fourth bit in the string is flipped. The mutation probability should be kept very low (usually about 0.001%) as a high mutation rate will destroy fit strings and degenerate the GA algorithm into a random walk, with all the associated problems.

But mutation will help prevent the population from stagnating, adding "fresh blood", as it were, to a population. Remember that much of the power of a GA comes

from the fact that it contains a rich set of strings of great diversity. Mutation helps to maintain that diversity throughout the GA's iterations.

1.9 Four Differences Separate Genetic Algorithms from More Conventional Optimization Techniques:

1. Direct manipulation of a coding:

Genetic algorithms manipulate decision or control variable representations at a string level to exploit similarities among high-performance strings. Other methods usually deal with functions and their control variables directly.

GA`s deal with parameters of finite length, which are coded using a finite alphabet, rather than directly manipulating the parameters themselves. This means that the search is unconstrained neither by the continuity of the function under investigation, nor the existence of a derivative function. Moreover, by exploring similarities in coding, GAs can deal effectively with a broader class of functions than can many other procedures (see Building Block Hypothesis).

Evaluation of the performance of candidate solutions is found using objective, payoff information. While this makes the search domain transparent to the algorithm and frees it from the constraint of having to use auxiliary or derivative information, it also means that there is an upper bound to its performance potential.

2. Search from a population, not a single point:

In this way, GAs finds safety in numbers. By maintaining a population of well-adapted sample points, the probability of reaching a false peak is reduced. The search starts from a population of many points, rather than starting from just one point. This parallelism means that the search will not become trapped on a local maxima - especially if a measure of diversity-maintenance is incorporated into the algorithm, for then, one candidate may become trapped on a local maxima, but the need to maintain diverity in the search population means that other candidates will therefore avoid that particular area of the search space.

3. Search via sampling, a blind search:

GAs achieves much of their breadth by ignoring information except that concerning payoff. Other methods rely heavily on such information, and in problems where the necessary information is not available or difficult to obtain, these other techniques break down. GAs remain general by exploiting information available in any search problem. GAs process similarities in the underlying coding together with

information ranking the structures according to their survival capability in the current environment. By exploiting such widely-available information, GAs may be applied to virtually any problem.

4. Search using stochastic operators, not deterministic rules:

The transition rules used by genetic algorithms are probabilistic, not deterministic. A distinction, however, exists between the randomised operators of GAs and other methods that are simple random walks. GAs use random choice to guide a highly exploitative search.

CHAPTER TWO

OPTIMIZATION PROBLEM

2.1 Definition for Optimization

A series of operations that can be performed periodically to keep a computer in optimum shape. Optimization is done by running a maintenance check, scanning for viruses, and defragmenting the hard disk. Norton Utilities is one program used for optimizing.

2.2 The Optimization Problem

The gradient-based optimization algorithms most often used with structural equation models (Levenberg-Marquardt, Newton-Raphson, quasi-Newton) are inadequate because they too often fail to find the global maximum of the likelihood function. The discrepancy function is not globally convex. Multiple, local minima and saddle points often exist, so that there is no guarantee that gradient-based methods will converge to the global maximum. Indeed, saddle points and other complexities in the curvature of the likelihood function can make it difficult for gradient-based optimization methods to find any maximum at all. Such difficulties are intrinsic to linear structure models for two reasons. First, the LISREL likelihood is not globally concave. Second, linear structure models' identification conditions do not require and do not guarantee that the model (as a function of the data) will determine a unique set of parameter values outside a neighborhood of the true values. The derivatives of the likelihood function with respect to the parameters are not well defined outside of the neighborhood of the solution. Therefore, outside of the neighborhood of the solution, derivative based methods often have little or no information upon which to advance to the global maximum.

Bootstrap methodology accentuates optimization difficulties, because the bootstrap resampling distribution draws from the entire distribution of the parameter estimates. Even if optimization in the original sample is not problematic, one can expect to encounter difficulties in a significant number of bootstrap resamples. Even if the model being estimated is correctly specified, problematic resamples contain crucial information about the tails of the distribution of the parameter estimates. Indeed, what

bootstrap methods primarily do is make corrections for skewness-for asymmetry between the tails of the distribution of each parameter estimate--that is ignored by normal-theory confidence limit estimates. Tossing out the tail information basically defeats the purpose of using the bootstrap to improve estimated confidence intervals. In general, any procedure of replacing problematic resamples with new resampling draws until optimization is easy must fail, as making such replacements would induce incomplete coverage of the parameter estimates' sampling distribution and therefore incorrect inferences. Ichikawa and Konishi (1995) make this mistake.

Because the nonexistence of good MLEs in bootstrap resamples is evidence of misspecification and because the occurrence of failures affects the coverage of the bootstrap confidence intervals, it is crucial to use an optimization method that finds the global minimum of the discrepancy function if one exists. In order to overcome the problems of local minima and nonconvergence from poor starting values, GENBLIS combines a gradient-based method with an evolutionary programming (EP) algorithm. Our EP algorithm uses a collection of random and homotopy search operators that combine members of a population of candidate solutions to produce a population that on average better fits the current data. Nix and Vose (1992; Vose 1993) prove that genetic algorithms are asymptotically correct, in the sense that the probability of converging to the best possible population of candidate solutions goes to one as the population size increases to infinity. Because they have a similar Markov chain structure, EP algorithms of the kind we use are asymptotically correct in the same sense. For a linear structure model and a data set for which a good MLE (global minimum) exists, the best possible population is the one in which all but a small fraction of the candidate solutions have that value. A fraction of the population will have different values because the algorithm must include certain random variations in order to have effective global search properties. The probability of not finding a good MLE when one exists can be made arbitrarily small by increasing the population size used in the algorithm.

The EP is very good at finding a neighborhood of the global minimum in which the discrepancy function is convex. But the search operators, which do not use derivatives, are quite slow at getting from an arbitrary point in that neighborhood to the global minimum value. We add the Broyden-Fletcher-Goldfarb-Shanno (BFGS) quasi-Newton optimizer as an operator to do the final hill-climbing. We developed and implemented a general form of this EP-BFGS algorithm in a C program called Genetic Optimization

15

Using Derivatives (GENOUD). GENBLIS is a version of GENOUD specifically tuned to estimate linear structure models.

In our experience the program finds the global minimum solution for the LISREL estimation problem in all cases where the most widely used software fails, except where extensive examination suggests that a solution does not exist.

2.3 Genetic Algorithm Optimization

Genetic algorithm optimization is a theoretical improvement over the traditional hill-climb optimization technique that has been employed by TRANSYT-7F for many years. The genetic algorithm has the ability to avoid becoming trapped in a "local optimum" solution, and is mathematically best qualified to locate the "global optimum" solution.

Releases 9 features genetic algorithm optimization of offsets and yields points, using either TRANSYT-7F or CORSIM as the simulation engine. Phasing sequence optimization was introduced in release 9.4 (January 2002), and requires TRANSYT-7F as the simulation engine. Genetic algorithm optimization of cycle length and splits was introduced in release 9.6 (September 2002), and also requires TRANSYT-7F as the simulation engine.
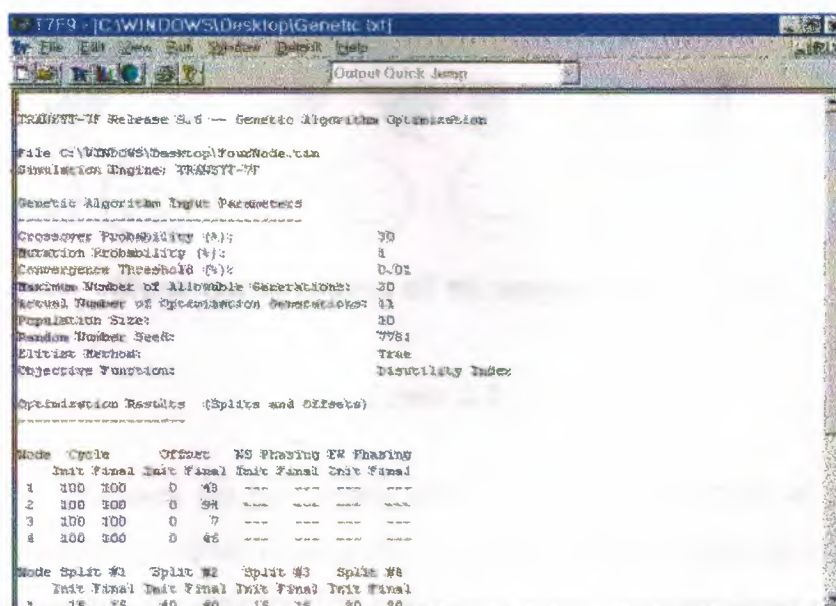

Figure 2.1.

For years, signal timing designers have known that they could often come up with a better control plan by making minor modifications to the so-called "optimal result"

recommended by a computer program. This is a byproduct of the hill-climb optimization process, where most timing plan candidates are not examined, in an effort to save time. Unfortunately, the global optimum solution is often skipped over during the hill-climb optimization process.

Fortunately, with genetic algorithm optimization, the user may have a much more difficult time in coming up with a better solution than the computer program. The genetic algorithm does not examine every single timing plan candidate either, but is a random guided search, capable of intelligently tracking down the global optimum solution. As with the human race, the weakest candidates are eliminated from the gene pool, and each successive generation of individuals contains stronger and stronger characteristics. It's survival of the fittest, and the unique processes of crossover and mutation conspire to keep the species as strong as possible.

Genetic Optimization

Out of the gene pool   **Crossover**

New generation of stronger candidates
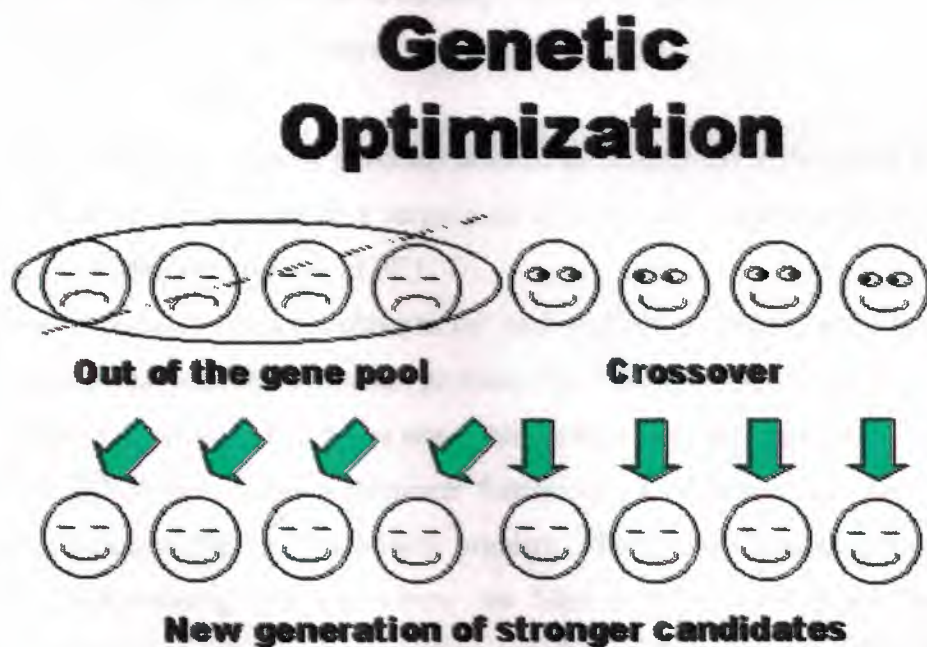
Figure 2.2.

Although it produces the best timing plans, a potential drawback of the genetic algorithm is increased program running times on the computer when optimizing large networks. Fortunately the drawback of increased program running times continues to be minimized by the ever-increasing processing speeds of today's computers. In addition, the TRANSYT-7F electronic Help system offers practical suggestions for reducing the genetic algorithm running times associated with large traffic networks.

2.4 CONTINUOUS OPTIMIZATION

2.4.1 Constrained Optimization

CO is an applications module written in the GAUSS programming language. It solves the Nonlinear Programming problem, subject to general constraints on the parameters - linear or nonlinear, equality or inequality, using the Sequential Quadratic Programming method in combination with several descent methods selectable by the user - Newton-Raphson, quasi-Newton (BFGS and DFP), and scaled quasi-Newton. There are also several selectable line search methods. A Trust Region method is also available which prevents saddle point solutions. Gradients can be user-provided or numerically calculated.

CO is fast and can handle large, time-consuming problems because it takes advantage of the speed and number-crunching capabilities of GAUSS. It is thus ideal for large scale Monte Carlo or bootstrap simulations.

Example

A Markowitz mean/variance portfolio allocation analysis on a thousand or more securities would be an example of a large-scale problem CO could handle (about 20 minutes on a 133 Mhz Pentium-based PC).

CO also contains a special technique for semi-definite problems, and thus it will solve the Markowitz portfolio allocation problem for a thousand stocks even when the covariance matrix is computed on fewer observations than there are securities.

Because CO handles general nonlinear functions and constraints; it can solve a more general problem than the Markowitz problem. The efficient frontier is essentially a quadratic programming problem where the Markowitz Mean/Variance portfolio allocation model is solved for a range of expected portfolio returns, which are then plotted against the portfolio risk measured as the standard deviation:

$$
\begin{aligned}
\min \quad & w_k'' \, \Sigma \, w_k \\
\text{s.t.} \quad & w_k'' \, \mu = r_{k} \\
& w_k 1 = 1 \\
& 0 \le w_k \le 1
\end{aligned}
$$

Where 1 is a conformable vector of ones, and where Σ is the observed covariance matrix of the returns of a portfolio of securities, and μ are their observed means.

18

This model is solved for

r_{ik}, $k = 1, \dots,$

And the efficient frontier is the plot of r_k on the vertical axis against

$$\sigma_k = \sqrt{w_k' \Sigma w_k}$$

On the horizontal axis the portfolio weights in w_k describe the optimum distribution of portfolio resources across the securities given the amount of risk to return one considers reasonable.

Because of CO's ability to handle nonlinear constraints, more elaborate models may be considered. For example, this model frequently concentrates the allocation into a minority of the securities. To spread out the allocation one could solve the problem subject to a maximum variance for the weights, i.e., subject to

$$w_k' \, w_k \leq \varphi$$

Where, φ is a constant setting a ceiling on the sums of squares of the weights.

Table 2.1.

Security	Mean	Std Dev	Correlation Matrix					
T-Bills	10.67	.94	1,					
Bonds	10.54	11.26	.097,	1,				
TSE	12.76	19.21	-.039,	.231,	1,			
S&P	13.67	13.67	.159,	.237,	.672,	1,		
EAFE	17.73	17.73	-.035,	.211,	.391,	.454,	1,	
Synthetic	13.68	12.19	-.024,	.247,	.424,	.432,	.941,	1

This data was taken from Harry S. Marmer and F.K. Louis Ng, "Mean-Semi variance Analysis of Option-Based Strategies: A Total Asset Mix Perspective", Financial Analysts Journal, May-June 1993.

An unconstrained analysis produced the results below:

r_k σ_k^2 w_k

Table2.2.

10.75	0.94	0.99	0.00	0.00	0.00	0.01	0.00
10.81	0.97	0.98	0.00	0.00	0.00	0.02	0.00
10.87	1.02	0.97	0.00	0.00	0.00	0.03	0.00
10.92	1.09	0.96	0.00	0.00	0.00	0.04	0.00
10.98	1.17	0.96	0.00	0.00	0.00	0.04	0.00
11.04	1.26	0.95	0.00	0.00	0.00	0.05	0.00
11.10	1.37	0.94	0.00	0.00	0.00	0.06	0.00
11.16	1.48	0.93	0.00	0.00	0.00	0.07	0.00
11.21	1.59	0.92	0.00	0.00	0.00	0.08	0.00
11.27	1.71	0.91	0.00	0.00	0.00	0.08	0.00

It can be observed that the optimal portfolio weights are highly concentrated in T-bills.

Now let us constrain w´w to be less than, say, .8. We then get:

r_k^* σ_k^2 w_k^*

Table2.3.

10.75	1.38	0.89	0.08	0.01	0.00	0.00	0.02
10.81	1.33	0.89	0.06	0.01	0.01	0.00	0.04
10.87	1.37	0.89	0.03	0.01	0.02	0.01	0.04
10.92	1.36	0.89	0.03	0.00	0.02	0.01	0.04
10.98	1.41	0.89	0.03	0.00	0.02	0.02	0.03
11.04	1.46	0.89	0.03	0.00	0.02	0.04	0.02
11.10	1.52	0.89	0.03	0.00	0.02	0.05	0.01
11.16	1.59	0.89	0.02	0.00	0.02	0.06	0.00
11.21	1.67	0.89	0.02	0.00	0.02	0.07	0.00
11.27	1.76	0.89	0.01	0.00	0.02	0.08	0.00

The constraint does indeed spread out the weights across the categories; in particular stocks seem to receive more emphasis.
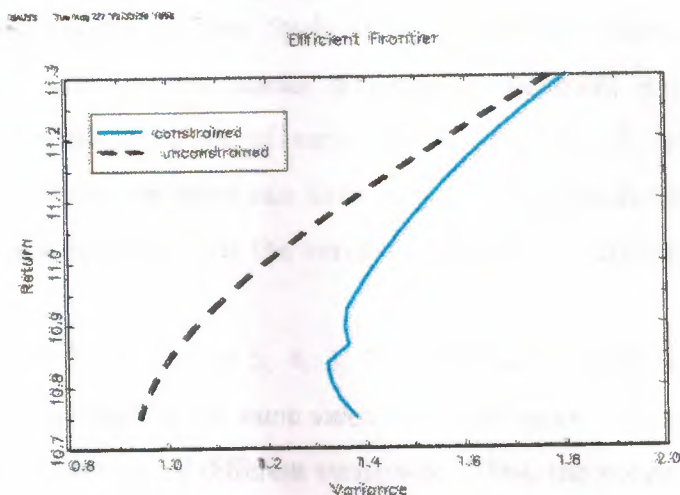

Figure 2.3. Efficient portfolio for these analyses

We see there that the constrained portfolio is riskier everywhere than the unconstrained portfolio given a particular portfolio return.

In summary, CO is well-suited for a variety of financial applications from the ordinary to the highly sophisticated, and the speed of GAUSS makes large and time-consuming problems feasible.

CO comes as source code and requires the GAUSS programming language software. It is available for Windows 95, OS/2, DOS, and major UNIX platforms.

Available for both PC and UNIX system versions of GAUSS and GAUSS Light, CO is an advanced GAUSS Application.

GAUSS Applications are modules written in GAUSS for performing specific modeling and analysis tasks. They are designed to minimize or eliminate the need for user programming while maintaining flexibility for non-standard problems.

CO requires GAUSS version 3.2.19 or higher (3.2.15+ for DOS version).

GAUSS program source code is included.

CO is available for DOS, OS/2, Windows NT, Windows 95, and UNIX versions of GAUSS.

2.4.1.1 Constraint Programming (CP) Problems

The term constraint programming comes from artificial intelligence research, where there are many problems that require assignment of symbolic values (such as positions on a chessboard) to variables that satisfy certain constraints. The symbolic

values come from a finite set of possibilities, and these possibilities can be numbered with integers.

Constraint programming defines "higher-level" constraints that apply to integer variables. The most common and useful higher-level constraint is the **all-different constraint,** which applies to a set of variables, say x1, x2, x3, x4 and x5. This constraint assumes that the variables can have only a finite number of possible values (say 1 through 5), and specifies that the variables must be all different at the optimal solution.

Values such as 1, 2, 3, 4, 5 or 5, 4, 3, 2, 1 for the variables would satisfy this constraint, but any assignment of the same value to two or more different variables (e.g. 1, 2, 3, 1, 4) would violate the all different constraint. Thus, the assignment must be an **ordering or permutation** of the integers 1 through 5.

A classic example of a constraint-programming problem is the **traveling salesman problem**: A salesman plans to visit N cities and must drive varying distances between them. In what order should he/she visit the cities to minimize the total distance traveled, while visiting each city exactly once?

Constraint programming problems have all the advantages and disadvantages of mixed-integer programming problems, and the extra requirements such as "all different" generally make such problems even harder to solve. All of Frontline's solvers support the all-different constraint, but you must bear in mind the implications for solution time if you use such constraints.

2.4.1.2   Solving MIP and CP Problems

The "classic" method for solving MIP and CP problems is called **Branch and Bound**. This method begins by finding the optimal solution to the "relaxation" of the problem without the integer constraints (via standard linear or nonlinear optimization methods). If in this solution, the decision variables with integer constraints have integer values, then no further work is required. If one or more integer variables have non-integral solutions, the Branch and Bound method chooses one such variable and "branches," creating two new sub problems where the value of that variable is more tightly constrained. These sub problems are solved and the process is repeated, until a solution that satisfies all of the integer constraints is found.

Alternative methods, such as genetic and evolutionary algorithms, randomly generate candidate solutions that satisfy the integer constraints. Such initial solutions

are usually far from optimal, but these methods then transform existing solutions into new candidate solutions, through methods such as **integer- or permutation-preserving mutation and crossover**, that continue to satisfy the integer constraints, but may have better objective values. This process is repeated until a sufficiently "good solution" is found. Generally, these methods are not able to "prove optimality" of the solution.

2.4.1.3 Smooth Nonlinear Optimization (NLP) Problems

A smooth nonlinear programming (NLP) or nonlinear optimization problem is one in which the objective or at least one of the constraints is a **smooth nonlinear function** of the decision variables. An example of a smooth nonlinear function is:

$$2 x_1^2 + x_2^3 + \log x_3$$

Where x_1, x_2 and x_3 are decision variables. A quadratic programming (QP) problem is a special case of a smooth nonlinear optimization problem, but it is usually solved by specialized, more efficient methods.

Nonlinear functions, unlike linear functions, may involve variables that are raised to a power or multiplied or divided by other variables. They may also use transcendental functions such as exp, log, sine and cosine.

NLP problems and their solution methods require nonlinear functions that are **continuous**, and (usually) further require functions that are **smooth** -- which means that **derivatives** of these functions with respect to each decision variable, i.e. the function **gradients**, are continuous.

A continuous function has no "breaks" in its graph. The Excel function =IF $(C_1>10,D_1, 2*D_1)$ is discontinuous if C_1 is a decision variable, because its value "jumps" from D_1 to $2*D_1$. The Excel function =ABS (C_1) is continuous, but nonsmooth -- its graph is an unbroken "V" shape, but its derivative is discontinuous, since it jumps from -1 to +1 at $C_1=0$.

NLP problems are intrinsically **more difficult to solve** than LP and QP problems. Because they may have multiple feasible regions and multiple locally optimal points within such regions, there is no known way to determine with certainty that an NLP problem is infeasible, that the objective function is unbounded, or that an optimal solution is the "global optimum" across all feasible r

2.4.2 Unconstrained Optimization

The unconstrained optimization problem is central to the development of optimization software. Constrained optimization algorithms are often extensions of unconstrained algorithms, while nonlinear least squares and nonlinear equation algorithms tend to be specializations. In the unconstrained optimization problem, we seek a local minimizes of a real-valued function, f(x), where x is a vector of *n* real variables. In other words, we seek a vector, x*, such that f(x*) <= f(x) for all x close to x*.

Global optimization algorithms try to find an x* that minimizes f over all possible vectors x. This is a much harder problem to solve. We do not discuss it here because, at present, no efficient algorithm is known for performing this task. For many applications, local minima are good enough, particularly when the user can draw on his/her own experience and provide a good starting point for the algorithm.

Newton's method gives rise to a wide and important class of algorithms that require computation of the gradient vector $\nabla f(x) = \begin{pmatrix} \partial_1 f(x) \\ \vdots \\ \partial_n f(x) \end{pmatrix}$,

And the Hessian matrix, $\nabla^2 f(x) = \left(\partial_j \partial_i (x) \right)$

Although the computation or approximation of the Hessian can be a time-consuming operation, there are many problems for which this computation is justified. We describe algorithms in which the user supplies the Hessian explicitly before moving on to a discussion of algorithms that don't require the Hessian.

Newton's method forms a quadratic model of the objective function around the current iterates x_k, *the model function is defined:*

$$q_k(\partial) = f(x_k) + \nabla f(x_k)^T \partial + \frac{1}{2} \partial^T \nabla^2 f(x_k) \partial .$$

In the basic Newton method, the next iterate is obtained from *the minimizes of* q_k. When the Hessian matrix, $\nabla^2 f(x_k)$, is positive definite, the quadratic model has a unique minimizes that can be obtained by solving the symmetric $n \times n$ linear system:

$\nabla^2 f(x_k) \partial_k = -\nabla f(x_k)$. The next iterate is then $x_k + 1 = x_k + \partial_k$

Convergence is guaranteed if the starting point is sufficiently close to a local minimizes x* at which the Hessian is positive definite. Moreover, the rate of

convergence is quadratic, that is, $\|x_k + 1 - x^*\| \le \beta \|x_k - x^*\|^2$, for some positive constant β.

In most circumstances, however, the basic Newton method has to be modified to achieve convergence.

Versions of Newton's method are implemented in the following software packages:

BTN, GAUSS, IMSL, LANCELOT, NAG, OPTIMA, PORT 3, PROC NLP, TENMIN, TN, TNPACK, UNCMIN, and VE08.

The NEOS Server also has an unconstrained minimization facility to solve these problems remotely over the Internet.

These codes obtain convergence when the starting point is not close to a minimizes by using either a line-search or a trust-region approach.

The line-search variant modifies the search direction to obtain another a downhill, or descent direction for f. It then tries different step lengths along this direction until it finds a step that not only decreases f, but also achieves at least a small fraction of this direction's potential.

The trust-region variant uses the original quadratic model function, but they constrain the new iterate to stay in a local neighborhood of the current iterate. To find the step, then, we have to minimize the quadratic subject to staying in this neighborhood, which is generally ellipsoidal in shape.

Line-search and trust-region techniques are suitable if the number of variables n is not too large, because the cost per iteration is of order n^3. Codes for problems with a large number of variables tend to use truncated Newton methods, which usually settle for an approximate minimizes of the quadratic model.

So far, we have assumed that the Hessian matrix is available, but the algorithms are unchanged if the Hessian matrix is replaced by a reasonable approximation. Two kinds of methods use approximate Hessians in place of the real thing:

The first possibility is to use difference approximations to the exact Hessian. We exploit the fact that each column of the Hessian can be approximated by taking the difference between two instances of the gradient vector evaluated at two nearby points. For sparse Hessians, we can often approximate many columns of the Hessian with a single gradient evaluation by choosing the evaluation points judiciously.

Quasi-Newton Methods build up an approximation to the Hessian by keeping track of the gradient differences along each step taken by the algorithm. Various conditions

are imposed on the approximate Hessian. For example, its behavior along the step just taken is forced to mimic the behavior of the exact Hessian, and it is usually kept positive definite.

Finally, we mention two other approaches for unconstrained problems that are not so closely related to Newton's method:

Nonlinear conjugate gradient methods are motivated by the success of the linear conjugate gradient method in minimizing quadratic functions with positive definite Hessians. They use search directions that combine the negative gradient direction with another direction, chosen so that the search will take place along a direction not previously explored by the algorithm. At least, this property holds for the quadratic case, for which the minimizes is found exactly within just n iterations. For nonlinear problems, performance is problematic, but these methods do have the advantage that they require only gradient evaluations and do not use much storage.

The nonlinear Simplex method (not to be confused with the simplex method for linear programming) requires neither gradient nor Hessian evaluations. Instead, it performs a pattern search based only on function values. Because it makes little use of information about f, it typically requires a great many iterations to find a solution that is even in the ballpark. It can be useful when f is no smooth or when derivatives are impossible to find, but it is unfortunately often used when one of the algorithms above would be more appropriate.

2.4.2.1 Systems of Nonlinear Equations

Systems of nonlinear equations arise as constraints in optimization problems, but also arise, for example, when differential and integral equations are discredited. In solving a system of nonlinear equations, we seek a vector such that f(x)=0 where x is an n-dimensional of n variables. Most algorithms in this section are closely related to algorithms for unconstrained optimization and nonlinear least squares. Indeed, algorithms for systems of nonlinear equations usually proceed by seeking a local minimizes to the problem

$\min \left\{ \|f(x)\| : x \in R^n \right\}$ For some norm $\|.\|$, usually the 2-norm. This strategy is reasonable, since any solution of the nonlinear equations is a global solution of the minimization problem.

Of linear equations $\quad f'(x_k)\partial_k = -f(x_k), \quad (1.1)$

Newton's method, modified and enhanced, forms the basis for most of the software used to solve systems of nonlinear equations. Given an iterate, Newton's method computes f(x) and its Jacobian matrix, finds a step by solving the system and then sets $x_k + 1 = x_k + \partial_k$.

Most of the computational cost of Newton's method is associated with two operations: evaluation of the function and the Jacobian matrix, and the solution of the linear system (1.1). Since the Jacobian is $f'(x) = (\partial_1 f(x), ..., \partial_n f(x))$,

The computation of the ith column requires the partial derivative of f with respect to the ith variable, while the solution of the linear system (1.1) requires order n^3 operations when the Jacobian is dense.

Convergence of Newton's method is guaranteed if the starting is sufficiently close to the solution and the Jacobian at the solution is nonsingular. Under these conditions the rate of convergence is quadratic; that is, $\|x_k + 1 - x^*\| \leq \beta \|x_k - x^*\|^2$, for some positive constant β. This rapid local convergence is the main advantage of Newton's method. The disadvantages include the need to calculate the Jacobian matrix and the lack of guaranteed global convergence; that is, convergence from remote starting points.

The following software attempts to overcome these two disadvantages of Newton's method by allowing approximations to be used in place of the exact Jacobian matrix and by using two basic strategies-trust region and line search-to improve global convergence behavior:

GAUSS , IMSL , LANCELOT , MATLAB , MINPACK-1 , NAG(FORTRAN) , NAG(C) , NITSOL , and OPTIMA .

- Trust Region and Line-search Methods.
- Truncated Newton Method.
- Broyden's Method.
- Tensor Methods.
- Homotopy Methods.

2.4.2.2 Nonlinear Least Squares

The nonlinear least squares problem has the general form

$\min\{r(x) : x \in IR^n\}$ Where r is the function defined by $r(x) = \frac{1}{2}\|f(x)\|_2^2$ for some vector-valued function f that maps R^n to R^m.

27

Least squares problems often arise in data-fitting applications. Suppose that some physical or economic process is modeled by a nonlinear function ϕ that depends on a parameter vector x and time t. If b_i is the actual output of the system at time t_i, then the residual $\phi(x, t_i) - b_i$ measures the discrepancy between the predicted and observed outputs of the system at time t_i. A reasonable estimate for the parameter x may be obtained by defining the ith component of f by $f_i(x) = \phi(x, t_i) - b_i$,

And solving the least squares problem with this definition of f.

From an algorithmic point of view, the feature that distinguishes least squares problems from the general unconstrained optimization problem is the structure of the Hessian matrix of r. The Jacobian matrix of f, $f'(x) = (\partial_1 f(x), ..., \partial_n f(x))$, a be used to express the gradient of r since $\nabla r(x) = f'(x)^T f(x)$. similarly, $f'(x)$ is part of the Hessian matrix $\nabla^2 r(x)$ since

$$\nabla^2 r(x) = f'(x)^T f'(x) + \sum_{i=1}^{m} f_i(x) \nabla^2 f_i(x).$$

To calculate the gradient of r, we need to calculate the Jacobian matrix $f'(x)$. Having done so, we know the first term in the Hessian matrix $\nabla^2 r(x)$ without doing any further evaluations. Nonlinear least squares algorithms exploit this structure.

In many practical circumstances, the first term $f'(x)^T f'(x)$ in $\nabla^2 r(x)$ is more important

than the second term, most notably when the residuals $f_i(x)$ are small at the solution. Specifically, we say that a problem has small residuals if, for all x near a solution, the quantities $|f_i(x)| \|\nabla^2 f_i(x)\|$, $i = 1, 2, ..., n$ are small relative to the smallest eigenvalue of $f'(x)^T f'(x)$.

- Gauss-Newton Method
- Levenberg-Marquardt Method
- Hybrid Methods
- Large Scale Methods
- Techniques for solving L.S. problems with constraints
- Notes and References

2.5 Global Optimization (GO)

A **globally optimal** solution is one where there are no other feasible solutions with better objective function values. A **locally optimal** solution is one where there are no other feasible solutions "in the vicinity" with better objective function values. You can picture this as a point at the top of a "peak" or at the bottom of a "valley" which may be formed by the objective function and/or the constraints -- but there may be a higher peak or a deeper valley far away from the current point.

In certain types of problems, a locally optimal solution is also globally optimal. These include LP problems; QP problems where the objective is positive definite (if minimizing; negative definite if maximizing); and NLP problems where the objective is a convex function (if minimizing; concave if maximizing) and the constraints form a convex set. But most nonlinear problems are likely to have multiple locally optimal solutions.

Global optimization seeks to find the globally optimal solution. GO problems are intrinsically **very difficult to solve**; based on both theoretical analysis and practical experience, you should expect the time required to solve a GO problem to increase rapidly -- perhaps exponentially -- with the number of variables and constraints.

2.5.1 Complexity of the Global Optimization Problem

The global optimization problem is indeed hard. Rinnooy Kan and Timmer [16] claim that the global optimization problem is unsolvable in a finite number of steps. Their argument is as follows:

For any continuously differentiable function f, any point x_* and any neighborhood B of x_*, there exists a function f' such that $f + f'$ is continuously differentiable, $f + f'$ equals f for all points outside B and the global minimum of $f + f'$ is x_*. $((f + f')$ is an of f.) Thus, for any point x_*, one cannot guarantee that it is not the global minimum without evaluating the function in at least one point in every neighborhood B of x_*. As B can be chosen arbitrarily small, it follows that any method designed to solve the global optimization problem would require an unbounded number of steps.

The indentation argument is certainly valid if one wishes to guarantee that an exact point, x_*, is a global minimizer. Indeed, should the exact global minimizer be an irrational number, it is obviously impossible, in a finite number of steps, to numerically

represent this solution. However, one *can*, in a finite amount of time, guarantee that $f(x_*) = f_*$ is within

2.5.2 Solving GO Problems

Multistart methods are a popular way to seek globally optimal solutions with the aid of a "classical" smooth nonlinear solver (that by itself finds only locally optimal solutions). The basic idea behind these methods is to automatically start the nonlinear Solver from randomly selected starting points, reaching different locally optimal solutions, then select the best of these as the proposed globally optimal solution. Multistart methods have a limited guarantee that (given certain assumptions about the problem) they will **"converge in probability"** to a globally optimal solution. This means that as the number of runs of the nonlinear Solver increases, the probability that the globally optimal solution has been found also increases towards 100%.

Where Multistart methods rely on random sampling of starting points, **Continuous Branch and Bound** methods are designed to systematically subdivide the feasible region into successively smaller subregions, and find locally optimal solutions in each subregion. The best of the locally optimally solutions is proposed as the globally optimal solution. Continuous Branch and Bound methods have a theoretical guarantee of convergence to the globally optimal solution, but this guarantee usually cannot be realized in a reasonable amount of computing time, for problems of more than a small number of variables. Hence many Continuous Branch and Bound methods also use some kind of random or statistical sampling to improve performance.

Genetic Algorithms, Tabu Search and Scatter Search are designed to find "good" solutions to nonsmooth optimization problems, but they can also be applied to smooth nonlinear problems to seek a globally optimal solution. They are often effective at finding better solutions than a "classic" smooth nonlinear solver alone, but they usually take much more computing time, and they offer no guarantees of convergence, or tests for having reached the globally optimal solution.

2.6 Nonsmooth Optimization (NSP)

The **most difficult type of optimization problem to solve** is a nonsmooth problem (NSP). Such a problem may not only have multiple feasible regions and multiple locally optimal points within each region -- because some of the functions are non-smooth or even discontinuous, derivative or gradient information generally cannot be used to

determine the direction in which the function is increasing (or decreasing). In other words, the situation at one possible solution gives very little information about where to look for a better solution.

In all but the simplest problems, it is impractical to exhaustively enumerate all of the possible solutions and pick the best one, even on a fast computer. Hence, most methods rely on some sort of random sampling of possible solutions. Such methods are **nondeterministic** or **stochastic** -- they may yield different solutions on different runs, even when started from the same point on the same model, depending on which points are randomly sampled.

2.6.1 Solving NSP Problems

Genetic or Evolutionary Algorithms offer one way to find "good" solutions to nonsmooth optimization problems. (In a genetic algorithm the problem is encoded in a series of bit strings that are manipulated by the algorithm; in an "evolutionary algorithm," the decision variables and problem functions are used directly. Most commercial Solver products are based on evolutionary algorithms.)

These algorithms maintain a **population** of candidate solutions, rather than a single best solution so far. From existing candidate solutions, they generate new solutions through either random **mutation** of single points or **crossover** or recombination of two or more existing points. The population is then subject to **selection** that tends to eliminate the worst candidate solutions and keep the best ones. This process is repeated, generating better and better solutions; however, there is no way for these methods to determine that a given solution is truly optimal.

Tabu Search and Scatter Search offers another approach to find "good" solutions to nonsmooth optimization problems. These algorithms also maintain a **population** of candidate solutions, rather than a single best solution so far, and they generate new solutions from old ones. However, they rely less on random selection and more on **deterministic** methods. Tabu search uses **memory** of past search results to guide the direction and intensity of future searches. These methods generate successively better solutions, but as with genetic and evolutionary algorithms, there is no way for these methods to determine that a given solution is truly optimal.

2.7 Multi-Objective Optimization for Highway Management Programming

Highway infrastructure is a major national investment, and a well-managed highway network forms an integral element of a sustainable economy. An ideal management program for a highway network is one that would maintain all highway sections at a sufficiently high level of service and structural condition, but requires only a reasonably low budget and use of resources, without creating any significant adverse impact on the environment, safe traffic operations, and social and community activities. Unfortunately, many of these are conflicting requirements. For instance, more resources and higher budgets may be needed if the highways are to be maintained at a high state of operability. But this could lead to pavement activities causing longer traffic delays, increased pollution, more disruption of social activities, and inconvenience to the community. Therefore, the decision processes involved in highway management activities requires a multi-objective consideration that addresses the competing requirements of different objectives.

Practically all the pavement management-programming tools in use currently are based on single-objective optimization. In single-objective analysis, the requirements which are not incorporated into the objective function are imposed as constraints in the formulation. This can be viewed as an interference of the optimization process, which artificially sets limits to selected problem parameters. As a result, the solutions obtained from single-objective analysis are sub-optimal with respect to one's derived from multi-objective formulations.

A genetic-algorithm (GA) based formulation for multi-objective programming of highway management activities has been developed at the Centre for Transportation Research. Genetic algorithms, which are a robust search technique formulated on the principles of natural selection and natural genetics, are employed to generate and identify better solutions until convergence is reached. The selection of good solutions is based on the so-called Pareto based fitness evaluation procedure by comparing the relative strength of the generated solutions with respect to each of the adopted objectives.

An important aspect of multi-objective GA analysis is the definition of "fitness" of a solution. The "fitness" of a solution directly influences the probability of the solution being selected for reproduction to generate new offspring solutions. To overcome the

afore-mentioned problem, a rank-based fitness assignment scheme is used where all non-dominated solutions in a given pool of solutions are assigned a rank of 1. Solutions that are dominated by one solution will have a rank of 2, those dominated by two solutions are given a rank of 3, and so on.

In the evaluation of a pool of solutions, one can identify a curve (for the case of two-objective problems) or a surface (for the case three- or higher multi-objective problems), which are composed of all non-dominated solutions. This curve or is known as the Pareto frontier. The GA optimization process seeks to generate new solutions that would give an improved frontier that dominates the existing frontier. This process continues until, ideally, a set of globally non-dominated solutions is found. These globally non-dominated solutions, called the Pareto optimal set, define the Pareto optimal frontier.

The general procedure of GA operations and offspring generation in multi-objective programming is similar to that for single-objective programming. The main difference lies with the evaluation of "fitness" of each solution, which is the driving criterion of the search mechanism of genetic algorithms. An important consideration of the optimization process is to produce representative solutions that are spread more or less evenly along the Pareto frontier. This can be achieved by using an appropriate reproduction scheme to generate offspring solutions and to form a new pool of parent solutions.

The proposed procedure is illustrated by means of a highway maintenance optimization problem at network level subject to five forms of resources and operation constraints. The five constraints were production requirements, budget, manpower availability, equipment availability, and rehabilitation schedule. The problem considered four highway classes, four pavement defects and three levels of maintenance-need urgency. Its objective was to select an optimal set of maintenance activities for an analysis period of 45 working days.

The decision variables were the respective amounts of maintenance work, measured in workdays, assigned to each maintenance treatment type. There were 48 treatment types referring to maintenance repairs arising from 4 distress forms of 3 maintenance-need urgency levels on 4 highway classes. The coded string structure of GA representation thus consisted of 48 cells. Each cell could assume an integer value of workdays from 0 to 45.

In this example, the following four objective functions were considered: (a) minimization of the total maintenance cost; (b) maximization of overall network pavement condition; (c) minimization of total manpower requirement, and (d) maximization of work production in total workday units. For the purpose of illustrating the flexibility of GAs in handling multi-objective problems, two different analyses were performed. They consisted of a two-objective and a three-objective problem. The objectives of the former were minimization of total maintenance cost and maximization of maintenance work production, while those of the latter were minimization of total maintenance cost, maximization of maintenance work production, and maximization of overall network condition.

Figure 2.4 shows the convergence pattern of the Pareto frontier for the case with two- objective problem. The convergence trend was apparent with the Pareto front moving in the direction towards lower maintenance costs and higher maintenance work production. Figure2.5 presents the Pareto optimal frontiers of both the two-objective and three-objective optimization solutions in the same two-dimensional plot of (maintenance cost) vs. (work production) for comparison. It can be seen that although the two Pareto frontiers are very close to each other, the Pareto frontier of the two-objective solutions appear to dominate the frontier of the three-objective solutions in terms of maintenance cost and work production. This was clearly the effect of introducing a third objective, i.e. maximizing pavement condition, in the three-objective problem, in compromising the two initial objectives.
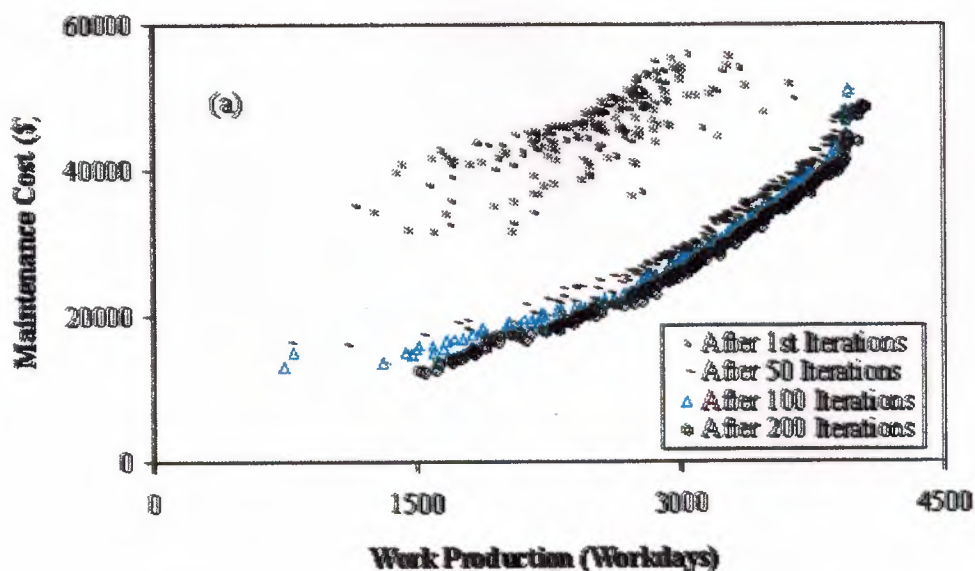
Figure 2.4. Convergence of Pareto frontier in GA Solutions

34
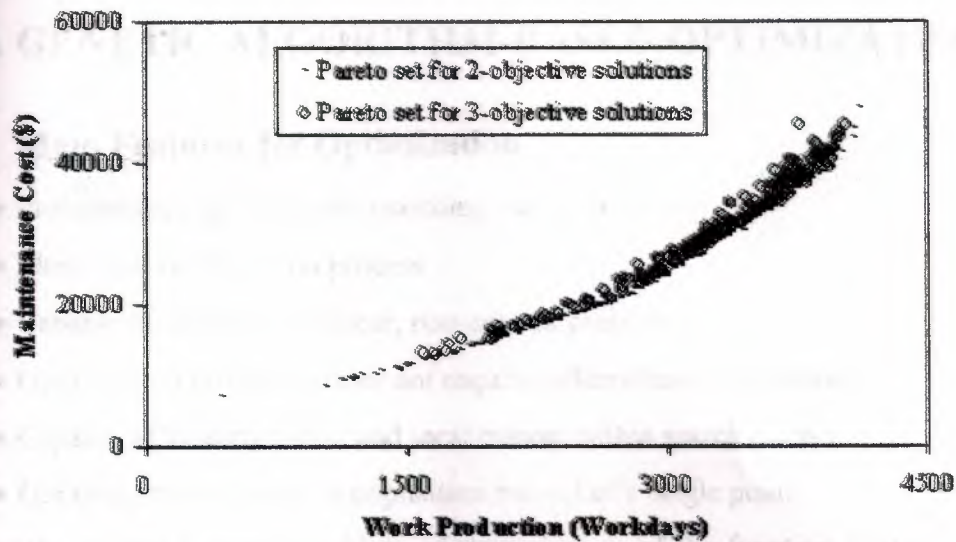
Figure 2.5. Comparison of Pareto frontiers for two- and three-objective solutions.

The robust search characteristics and multiple-solution handling ability of GAs are well suited for multi-objective optimization analysis of multi-objective highway infrastructure management and planning. The concepts of Pareto frontier and rank-based fitness evaluation, and a method of selecting optimal solution from the Pareto frontier were found to be effective in solving the problems. The proposed algorithm was able to produce a set of optimal solutions, which were well spread on the Pareto frontier.

CHAPTER THREE

A GENETIC ALGORITHM-BASED OPTIMIZATION

3.1 Main Features for Optimization

- Probabilistic algorithms for searching and optimization
- Mimic natural evolution process
- Capable of handling nonlinear, non-convex problem
- Optimization procedure does not require differentiation operation
- Capable of locating global and local optima within search domain
- Optimizations is based on population instead of a single point.

Let's consider a simple problem of maximization of the function $F(x)=x^2$, where $x \in [0,31]$ (see Fig).


Figure 3.1.

In order to use GA we should first code variables in to bit strings as any integer number between 0 to 31 may be represented in binary number 5 symbols between $(00000)=0$ and $(11111)=31$ the length of chromosomes will be five.

Lets take 4 chromosomes with random set of genes as an initial population as:

01101

11000

01000

10011

36

Then define their fitness inserting appropriate real values into the function as:

$f(01000)_2 = f(8=64$

Thus the fitness of all individuals in calculated then accordance with the formula

$P^i_5 = f_i/\Sigma^I_{j=1}F^i_5,$ $i = 1..4$

Survival probability for each individual is calculated where as cumulatives

$P^i_{cum} = \Sigma^I_{j=1} P^i_5$ $i = 1..4$

Let's enter all the calculated values in table

Table3.1

Initial Population	Their integer values	$F(\ )=x^2$	P_5	P_{cum}	Num after Selection
01101	13	169	0.14	0.14	1
11000	24	576	0.49	0.63	2
01000	8	64	0.06	0.69	0
10011	19	361	0.31	1	1

Average   293   0.25   1

Maximum   576   0.49   2

Sum   1170   1   4

For the selection process we generate 4 random numbers from the range [0,1]. Suppose that we generated 0.1; 0.25; 0.5 and 0.8.

Comparing these values with cumulative probabilities, we obtaining the following

$0.1 < P^i_{cum}$

$P^I_{cum} < 0.25 < P^2_{cum}$

$P^1_{cum} < 0.5 < P^2_{cum}$

$P^3_{cum} < 0.8 < P^4_{cum}$

Looking at the right sides of these inequalities, one can easily see, that the first and the fourth chromosomes have passed the selection, each four taking a place in the new generation, the second chromosome-the most highly fitted has got 2 copies, while the

third one did not survive at all. These indices are written in the right column of table 3.1.

Then crossover operation is applied. If the probability of crossover Pc=1 is given, it means that, 4.1=4 chromosomes will participate in crossover process.

Let's choose them at random. Suppose that the first and the second strings mate from crossover point 4, and the third with the forth-from crossover point 2.

$$0110 \mid 1$$
$$1100 \mid 0$$
$$11 \mid 000$$
$$00 \mid 011$$

Table3.2

Population after selection	Crossover	New Population	Value of X	$F(x)=X^2$
01101	4	01100	12	144
11000	4	11001	25	625
11000	2	11011	27	729
10011	2	10000	16	256

Average	439
Maximum	729
Sum	1754

Having compared both of the tables we see, for ourselves, that the population fitness is improved and we have come close to the solution. The next recombination operator, mutation, is performed on a bit-by-bit basis. If Pm=0.05 is given, it means that only one of twenty bits in population will be changed: 20· 0.05=1. Suppose, that the third bit of the fourth string undergoes mutation. I.e. x_4=10100. Repeating these operations in a finite number of generations we will get the chromosome (11111) corresponding to problem optimal solution. It is necessary to mention, that GA is especially effective for multi-extreme problems in the global solution search process. For example, if the junction is of the type shown. It is rather difficult to find its global maximum by means of traditional methods. Suppose that the junction is defined as [1]

38

$$f(x)=x\cdot \sin(10\pi\cdot x)+1.$$

The problem is to find x from the range [-1,2], which maximizes the function f, i.e., to find x_0 such that

$$f(x_0)\geq f(x), \text{ for all } x\in[-1,2].$$

It is relatively easy to analyze the junction f. The zeros of the first derivative f' should be determined

$$F(x)=\sin(10\pi\cdot x)+10\pi\cdot x\cdot \cos(10\pi\cdot x)=0$$

The formula is equivalent to

$$\tan(10\pi\cdot x)=(10\pi\cdot x)$$

It is clear that the above equation has an infinite number of solutions,

$$x_i=\frac{2i-1}{20}+\xi_i, \qquad \text{for } i=1, 2, \ldots,$$

$$x_0 = 0,$$

$$x_i=\frac{2i+1}{20}+\xi_i, \qquad \text{for } i= 1, -2, \ldots,$$

Where terms ξ_1 represent decreasing sequences of real numbers (for i= 1, 2, ..., and i=-1, -2, ...) approaching zero. Note also that the function f reaches its local maxima for x, it i is an odd integer, and its local minima for x, if i is an even integer. Since the domain of the problem is $x\in[-1,2]$, the function reaches its maximum for

$$X_{19} = \frac{37}{20}+\xi_{19} = 1.85+\xi_{19},$$

Where $f(x_{19})$ is slightly larger than

$$F(1.85)=1.85 \cdot \sin(18\pi+\pi/2)=1.0=2.85.$$

Assume that we wish to construct a genetic algorithm to solve the above problem, I.e., to maximize the function f. Let us discuss the major components of such a genetic algorithm in turn. [4]

3.1.1 Representation

We use a binary vector as a chromosome to represent real values of the variable x. The length of the vector depends on the required precision, which, in this example, is six places after the decimal point. The domain of the variable x has length 3; the precision requirement implies that the range [-1,2] should be divided into at least $3 \cdot 1000000$ equal size ranges. This means that 22 bits are required as a binary vector (chromosome):

$$2097152=2^{31} < 3000000 \leq 2^{33} = 4194304$$

The mapping from a binary $(b_{21}\ b_{20}\ b_0)$ string into a real number x from the range [-1,2] is straightforward and is completed in two steps:

Convert the binary string $(b_{21}\ b_{20}\ b_0)$ from the base 2 to base 10:

$$(<b_{21}\ b_{20}\ b_0>)_2 = \left(\sum_{i=0}^{21} b_i \cdot 2^i \right)_{10} = x$$

Find a corresponding real number x:

$$x = -10 + x \cdot \frac{3}{2^{22}-1},$$

Where $-1,0$ is the left boundary of the domain and 3 is the length of the domain.

For example, a chromosome

(1000101110110101000111)

Represents the number 0.637197, since

x'=$(1000101110110101000111)_3 = 2288967$ and

$$x=1.0+2288967 \cdot \frac{3}{4194303} = 0.637197.$$

Of course, the chromosomes

(0000000000000000000000) and (1111111111111111111111)

Represent boundaries of the domain-1.0 and 2.0, respectively.

Initial population. The initialization process is very simple; we create a population of chromosomes, where each chromosome is a binary vector of 22 bits. All 22 bits for each chromosome are initialized randomly.

Evaluation function. Evaluation function eval for binary vectors v is equivalent to the function f:

$$Eval\ (v) = f(x),$$

Where the chromosome v represents the real value x.

As noted earlier, the evaluation junction plays the role of the environment, rating potential solutions in terms of their fitness. For example, three chromosomes:

$$V_1 = (1000101110110101000111)$$
$$V_2 = (0000001110000000010000)$$
$$V_3 = (1110000000111111000101)$$

Correspond to values x=0.637197, x=0.958973 and x_3=1.627888, respectively.

Consequently, the evaluation function would rate them as follows:

$$Eval\ (v_1) = f(x_1) = 1.586345$$
$$Eval\ (v_2) = f(x_2) = 0.078878$$
$$Eval\ (v_3) = f(x_3) = 2.250650$$

Clearly, the chromosome v_3 is the best of the three chromosomes, since its evaluation returns the highest value.

During the reproduction phase of the genetic algorithm we would use two classical genetic operators; mutation and crossover.

As mentioned earlier, mutation alters one or more genes (positions in a chromosome) with a probability equal to the mutation rate. Assume that the fifth gene from the v_3 chromosome was selected for a mutation. Since the fifth gene in this chromosome is 0, it would be flipped into 1. So the chromosome v_3 after this mutation would be

$$V_3 = (1110100000111111000101)$$

This chromosome represents the value x_3=1.721638 and $f(x_3)$=-0.082257. This means that this particular mutation resulted in a significant decrease of the value of the chromosome v_3 . On the other hand, if the 10th gene was selected for mutation in the chromosome v_3 them

$$V_3 = (1110000001111111000101)$$

The corresponding value x_3=1.630818 and $f(x_3)$=2.343555, an improve-ment over the original value of $f(x_3)$=2.250650.

Let us illustrate the crossover operator on chromosomes v_2 and v_3. Assume that the crossover point was (randomly) selected after the 5th gene:

$$V_2=(00000 \mid 0111000000010000)$$
$$V_3=(11100 \mid 0000011111000101)$$

The two resulting offspring are

$$V'_2=(00000 \mid 0000011111000101)$$
$$V'_3=(11100 \mid 0111000000010000)$$

These offspring evaluate to

$$F(v'_2)=f(-0.998113)=0.940865.$$
$$F(v'_3)=f(1.666028)=2.459245$$

Note that the second offspring has a better evaluation than both of its parents.

Parameters. For this particular problem we have used the following parameters population size ps=50, probability of crossover p_c =0.25, probability of mutation p_m =0.01. The following section presents some experimental results for such a genetic system.

Experimental results. We provide the generation number for which we noted an improvement in the evaluation function together with the value of the function. The best chromosome after 150 generations was

$$V_{max}=(1111001101000100000101),$$

Which corresponds to a value x_{max}=1.850773

AS exprected, $x_{max}=1.85+\xi$ and $f(x_{max})$ is slightly larger than 2.85.

$\{\} \rightarrow \{f(x)\} \ x_1^t \ [1] \quad |\in\lessgtr \quad \cdot \quad \xi \quad \pi$

Table 3.3.

Generation Number	Fitness Function
1	1.441942
5	2.250003
8	2.250283
9	2.250284
10	2.250363
12	2.328077
36	2.344251
40	2.345087
51	2.738930
99	2.849246
137	2.850217
145	2.850227

3.2 Applications

Genetic algorithms have been very successful at solving many types of problems.

1. Maximizing discontinuous, multimodal, multidimensional functions.

2. They have also been used on discrete problems, including the traveling salesperson, bin packing, and job-shop scheduling.

Here are two example problems, one discrete and one continuous, that are more realistic.

1. For an arbitrary expression in n Boolean variables, find an assignment of true and false values to the Boolean variables in the expression that make the expression true, if such an assignment exists.

2. This is called the Boolean satisfiability problem (SAT).

3. Each assignment of true and false values to the variables can be encoded in a bit string of length n, 0 for false and 1 for true.

4. The fitness of such a chromosome is how ``close" to being true the assignment makes the Boolean expression, such as how many clauses in the expression are individually made true by the assignment.

43

5. Some (in Postscript) SATs are very hard for GAs.

2. Maximize the function:

$f(x_1, x_2, x_3) = 21.5 + x_1 \sin(4 \text{ pi } x_1) + x_2 \sin(20 \text{ pi } x_2) + x_3$

for x_1 in [-3.0,12.1], x_2 in [4.1,5.8], and x_3 in [0.0,1.0].

1. This function is highly multimodal.

2. A chromosome can be encoded in three real numbers, one for each of x_1, x_2, and x_3, over their respective ranges.

3. The fitness of a chromosome is the value of $f(x_1,x_2,x_3)$ for the x_1, x_2, and x_3 encoded in the chromosome.

3.2.1 Difficulties

Hitchhiking.

1. This may cause a GA to perform poorly.

2. Once a schema attains high fitness in a population, then the members of the population containing the schema will proliferate.

3. Bit values in positions not in the schema will *hitchhike* along with the schema as it propagates and will delay the discovery of high-fitness schemas with different values in those other bit positions.

4. For example, consider these 8 schemas.

    11111111...
    ....11111111......................................
    ....11111111.......................................
    ....11111111..............................
    ....11111111......................
    ....11111111................
    ....11111111........
   11111111

5. Define a fitness function $f(x)$ where x is a 64-bit string to be 8 times the number of these schemas it contains.

6. So f(all 64 bits 1) = 64, f(all 64 bits 0) = 0.

7. A GA performs poorly relative to hill-climbing on this because of hitchhiking.

3.2.2 Deception

A problem is called deceptive if the average fitness of schemata not in the global optimum is greater than the average fitness of those that are.

For example, chromosome length = 3, $f(x)$ = fitness of x

$f(111)=1.0$

$f(011)=f(101)=f(110)=0.05$

$f(001)=f(010)=f(100)=0.95$

$f(000)=0.20$

$f(*0*)=f(**0)=f(0**)=(.2+.95+.95+.05)/4=0.5375$

$f(1**)=f(*1*)=f(**1)=(1+.05+.05+.95)/4=0.5125$

$f(00*)=f(*0*)=f(*00)=(.2+.95)/2=.575$

$f(11*)=f(1*1)=f(*11)=(1+.05)/2=0.525$

3.3 The Neighborhood Constraint Method:
A Genetic Algorithm-Based Multiobjective Optimization Technique
3.3.1. Overview

Multiobjective genetic algorithms (MOGAs), including VEGA, Pareto optimal ranking, and niched-Pareto, allow the tradeoffs among design objectives to be estimated in a single optimization model run. This paper presents a new MOGA technique called the neighborhood constraint method (NCM) that uses a combination of a neighborhood selection technique and location-dependent constraints. NCM is demonstrated for a complex, real-world problem, where it is compared with an integer programming (IP) procedure, an iterative procedure using a single-objective GA, and implementations of a Pareto and hybrid niched-Pareto MOGA technique. Preliminary results suggest that NCM is computationally more efficient than the SO approach and may perform better that the other MOGA techniques at promoting and maintaining population diversity over the entire tradeoff curve, using relatively small population sizes. These results show that further evaluation and comparisons of NCM are warranted.

Engineering problems that are faced in everyday decision-making are often complex and multiobjective (MO), requiring solutions to be judged on their ability to satisfy competing design goals. For most multiobjective problems, there exists a set of nondominated, or Pareto, solutions that represent the optimal tradeoff relationship among objectives. For any solution in the Pareto set, any one-design objective cannot be improved without sacrificing another.

The shape of the Pareto set can provide useful information to a decision maker. For example, in a pollution abatement problem, the Pareto set could identify the increasing costs associated with increasing levels of pollutant control. A decision-maker might use this relationship to determine reasonable control levels.

The purpose of an MO analysis is often to estimate the shape of an interesting portion of the Pareto set and to identify solutions along that portion. Several mathematical programming procedures, including the constraint method and the weighting method (Figure 3.2), are available for this purpose. These generating techniques estimate the noninferior set through iterative runs of an optimization model.

Figure 3.2. Classical MO Generating Techniques.

(A) All objectives except one are converted into constraints. The remaining objective is maximized. The constraints are incremented and the model is run iteratively. (B) All objectives are weighted and considered in the objective function simultaneously. A different set of weights is used in each run.

Alternatively, preference-based MO techniques, such as goal programming, may require only a single optimization model run to identify the solution that best meets prespecified "ideal" levels of each objective. Both classes of classical MO procedures are discussed in detail by Cohon (1978).

MO procedures have been used successfully within traditional mathematical optimization procedures, such as linear programming, integer programming and nonlinear programming. Many realistic problems, however, involve complex, nonlinear relationships that do not fit readily into one of these traditional frameworks.

Recently, many such problems have been solved by employing non-gradient-based, probabilistic search techniques such as genetic algorithms (GAs) and simulated annealing (SA). Both GA and SA may be used in the classical MO procedures. However, the inherent parallel structure of a GA provides some important advantages in MO analysis: with the appropriate schemes and operators, the population of solutions carried in a GA can be directed to converge along the Pareto set in a single run. Several multiobjective GA (MOGA) techniques that take advantage of this behavior are summarized in the following section.

47

The objective of this paper is to present a new method for MOGA, called the neighborhood constraint method (NCM). NCM differs from previous MOGA approaches by using a unique combination of a neighborhood selection technique and location-dependent constraints. A description of NCM and its application to an illustrative air quality management problem are provided. Performance of NCM is compared with those of several other MOGA techniques, as well as with the performance of a single-objective GA (SOGA) and an integer programming approach using the classical MO constraint method.

3.3.2 Literature Review: GAs in Mo Analysis

In a review of multiobjective optimization techniques, Fonseca and Fleming (1995) classify MOGA techniques into four major categories: plain aggregation methods, population-based non-Pareto methods, Pareto-based approaches, and niche induction schemes. Plain aggregation methods use GAs in classical MO procedures like those depicted in Figure 3.2. These methodologies do not take advantage of the GA population-based search to generate the Pareto set in a single run, but instead require much iteration.

The most well known population-based, non-Pareto method is the Vector Evaluated Genetic Algorithm, or VEGA (Schaffer, 1991). In the VEGA scheme, the population is divided into groups in each generation, and each group is evaluated with only a single objective. Individuals from different groups are allowed to mate with the intent that offspring will perform well with respect to the objectives of both their parents. While this approach and several of its variations have been successful in locating the elbow and extreme points of the Pareto set, it is recognized that exploration of the entire Pareto set is difficult.

In an intriguing non-Pareto technique, Hajela and Lin (1992) used a combination of restrictive mating, objective weighting, and fitness sharing. Their variable weighting approach treats the weights on objectives as decision variables that are encoded into the GA chromosome for each individual. Combined with fitness sharing, this approach allows individuals to evolve to represent a large number of different weightings. Fleming and Pashkevich (1985) noted that any technique that aggregates objectives linearly in the objective function, as in VEGA and Hajela and Lin's techniques, is not capable of exploring concave regions of the Pareto set.

Pareto-based approaches were proposed by Goldberg (in 1989) and have become a major focus of MOGA research. These techniques explicitly make use of the definition of Pareto optimality. In Pareto schemes, the fitness of an individual is defined in terms of rank. For each generation of the GA, the non-dominated set of solutions in the population is given the top rank and then removed from that population. This step is repeated, where in each iteration the non-dominated set of solutions remaining in the population is given the next rank, until all solutions in the population are assigned a rank. Rankings are then typically used in a tournament selection to encourage the exploration in the direction of nondominated individuals. Successful applications of this procedure have been reported by Cieniawski et al. (1995) and Ritzel et al. (1994). Liepins et al. (1988) found the Pareto optimal ranking scheme to be superior to VEGA. An important advantage of Pareto approaches is that they are able to identify solutions in concave areas of the Pareto set. Pareto schemes typically suffer from population drift, where the population migrates to a small portion of the Pareto set (Goldberg and Segrest, 1987). To assure a uniform sampling of the Pareto set, recent applications of Pareto-based MOGA by Fonseca and Fleming (1993), Horn et al. (1994), and Srinivas and Deb (1994) have incorporated niching schemes.

In both the Pareto and niched-Pareto methods, there may be a likelihood of mating between individuals that are far apart in the nondominated set. These solutions likely will be very different in decision space, and as a result, their children may not meet either objective satisfactorily (Goldberg, 1989). The niched-Parteto scheme by Fonseca and Fleming (1993) and the non-Pareto scheme by Hajela and Lin (1992) deal with this problem by restricting mating to individuals in similar regions in the objective space.

3.3.3 Neighborhood Constraint Method

Overview

The Neighborhood Constraint Method (NCM) presented here is a MOGA technique that:

- Samples effectively from along the entire Pareto set, including the endpoints
- Examines both convex and concave regions of the Pareto set, and
- Provides a computationally-efficient approach

While MOGA techniques by Horn et al. (1994) and Fonesca and Fleming (1993) also address these issues, NCM was developed concurrently and uses a very different approach.

NCM falls into the class of non-Pareto MOGA techniques because it does not make explicit use of Pareto dominance in fitness evaluation. NCM is inspired by the natural phenomena of adaptation to environmental conditions that are location ally, or geographically, dependent (e.g. the gradual darkening of skin color seen in human populations as one travels south from northern Europe to the equator). Adaptation, in this case, arises either from the preferential selection of traits that perform well with respect to the local environment, or through migration of traits to environments where they are better suited.

NCM includes the following major operations: population indexing, location-dependent constraints, and restrictive mating using a neighborhood selection scheme. These operations are described below.

3.3.3.1 Population Indexing

Each individual is given an index to define its location in the population. Indices are assigned in the following manner: for problems with two objectives, individuals are simply indexed from 1 to n, where n is the population size. For three objectives, the population is indexed in a matrix form so that each individual is given a set of x-y ordinal coordinates. The dimension of the matrix is m-by-n, where the product of m-by-n is the population size. Examples for two and three objectives are shown in Figure3.3.

**Two Objective Case**   **Three Objective Case**

Figure 3.3. Indices used in the two and three objective cases. For a problem with K objective, there are K-1 index numbers used to describe each individual.

For problems of four or more objectives, individuals are indexed in a similar fashion, with the number of index values set equal to the number of objectives minus 1.

3.3.3.2 Location-Dependent Constraints

NCM handles multiple objectives in a very similar manner to the classical MO constraint method. For a K-objective problem, K-1 objectives are converted into constraints while the remaining objective is optimized. What makes this method unique is that the values of the constrained objectives are gradually varied at each index location in the population.

Consider the two-objective problem with objectives A and B and a population size of n. The individual with index 1 would be constrained to satisfy objective B at a minimum target level T_{Bmin}. The individual with index n would instead be constrained to a maximum target level T_{Bmax}. The constraint for any individual i in the population is generalized by:

$$T_{Bi} = T_{Bmin} + \frac{(i-1)*(T_{Bmax} - T_{Bmin})}{(n-1)} \quad (1)$$

Conversion of these targets to constraints in the GA is done through the use of penalties. Although the form of the penalty function may be problem specific, it generally can be incorporated in the fitness function as:

$$f = S(A) - w * g(S(B)) \quad (2)$$

$$g(S(B)) = \frac{(T_{Bi} - S(B))}{(T_{Bmax} - T_{Bmin})}$$

$$\text{for } S(B) < T_{Bi} \quad (3)$$

$$g(S(B)) = 0 \quad \text{for } S(B) \geq T_{Bi} \quad (4)$$

Where: f is the fitness of an individual, S (A) and S (B) are the levels to which objectives A and B, respectively, are met, and w is a weighting factor. w is typically set to a very high value to encourage the constraint g, in Equation 2, to be strictly met. Equations 3 and 4 make the constraint on objective B to be one-sided, i.e., there is no penalty for exceeding the constraint.

3.3.3.3 Restrictive Mating: Neighborhood Selection

Neighborhood selection is a restrictive mating scheme discussed in the context of single objective; parallel GAs by DeJong and Sarma (1995). Neighborhood selection schemes restrict the mating of individuals to those within a specified "neighborhood," or vicinity, of each other. A number of different selection techniques, including roulette wheel and tournament selection, can be modified to promote neighborhood selection.

The implementation of neighborhood selection in NCM uses the notions of a selection radius, r, and a reference index location, L. The parameter r is fixed to be some fraction of n, the population size. The parameter L is an index that represents the location of an individual with respect to all other individuals in the population.

The likelihood of an individual being selected for mating is dependent on the proximity of that individual to the location L. An individual is considered for selection only if the distance between that individual and location L is within the selection radius, r. For instance, in a two-objective case, the distance between a reference location at L=4 and an individual at index 8 is 4. If r is greater than 4, the individual is considered for selection. Likewise, in the three-objective case, the distance between L at [4,2] and an individual at [3,7] is $((4-3)^2 + (2-7)^2)^{0.5}$, or 5.1. If r<5.1, then the individual will not be considered for selection.

The neighborhood selection process is described below for a two-objective problem. Assume that the population size, n, is 100; the selection radius, r, is defined as 0.05 * n, i.e., 5; and a tournament selection scheme is carried out. The selection process begins by setting the reference location, L, to 1. Individuals at index locations up to r (=5) locations away from L(=1), i.e., from 1 to 6, will be considered as candidates for selection. Two of the six individuals in this range are selected using tournament selection. The winners of the selection undergo crossover, and the resulting individuals are placed into a new population at locations L and L+1, i.e., at locations 1 and 2. L is then incremented by two, and the selection and crossover processes are repeated. These steps are repeated until the new population has the same number of individuals as the previous population. This implementation of neighborhood selection is depicted in Figure 3.4.
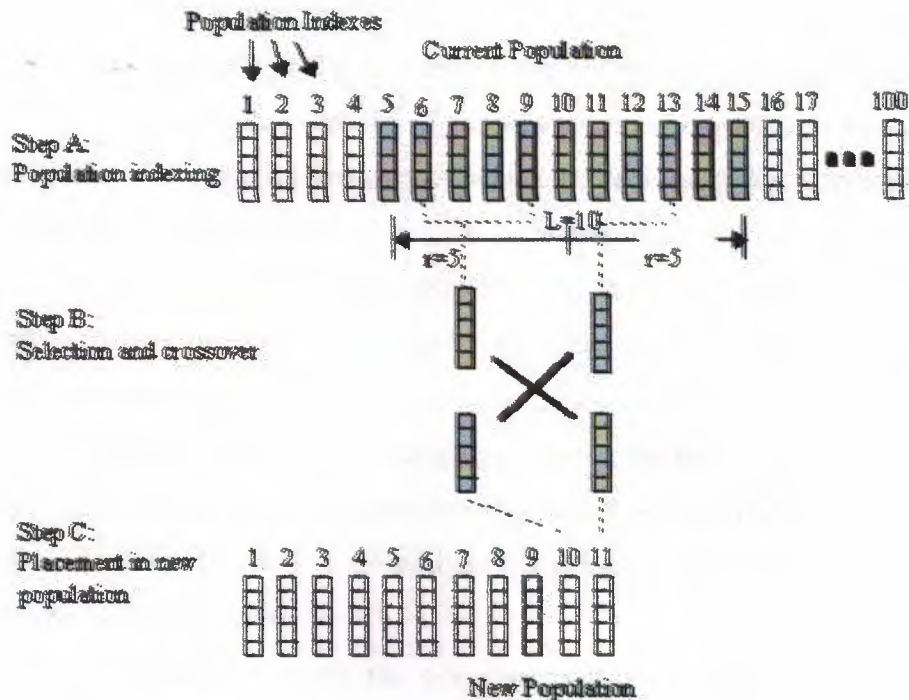
Figure 3.4. Neighborhood Selection.

Step A: individuals are identified that fall within r places of L. Step B: two individuals are selected from this set and undergo crossover. Step C: the resulting individuals are placed into the next population at locations L and L+1.

3.3.3.4 Other Special Operations

In addition to the NCM-specific operations described above, other special operations may be used to improve performance for various types of problems. Several potential modifications are listed and described below.

- **Population sorting** - Because the initial population is often generated randomly and is not necessarily in any order, individuals will likely not be in a location where they perform best. NCM performance may be improved by sorting the initial population, placing individuals closer to their ideal locations.

- **Dynamic scaling of T_{Bmax} and T_{Bmin}** - It may be difficult to estimate the initial values of the minimum and maximum objective targets, e.g., T_{Bmax} and T_{Bmin}, for some problems. When this occurs, T_{Bmax} and T_{Bmin} may be dynamically changed in each generation to reflect the maximum and minimum values in the algorithm to that point.

- **Locational elitism** - To speed convergence, an elitism scheme compares the individual at each location in the population with the individual in that location in the

53

previous generation. If the individual in the previous generation had a better fitness value than that of the current generation, it replaces the current individual.

- **Population distribution** - Using Equation 1, individuals will be spaced evenly along objective B. Alternative distributions of individuals along objective B may be useful in concentrating the search toward areas in which the decision maker is interested.

- **Allow multiple individuals at a location** - Instead of having one set of constraint values per individual, more than one individual can share the same set of constraints. This may provide a more robust search by allocating more GA resources to each of the objective targets.

- **Evaluation of local fitness** - In the selection process, the fitness of individuals within the selection radius can be re-evaluated to determine how well they would perform at location L. Because the children of selected individuals will be placed at L and L+1, this promotes gene migration.

- **Selection of individuals outside the selection radius** - It may be advantageous to select occasionally individuals from outside of the selection radius, which is expected to promote faster gene migration across the population. A factor can be introduced to allow sampling outside the selection radius at a pre-specified rate.

- **Dynamic selection radius** - Our trials to date have used a selection radius that is equal to a fixed percentage of the population size. An alternative would be to allow the selection radius to change dynamically. For instance, the radius would start large and reduce as the population converges to the Pareto set.

3.3.4 Application

NCM was applied to a real-world, air quality management application with two conflicting objectives: minimizing the cost of controlling air pollutant emissions and maximizing the amount of emissions reduction.

This problem included approximately 300 air pollutant sources that varied in emissions amounts from 1 to nearly 8000 tons per year. Each source had from four to seven mutually-exclusive control options, including a no-control option. The average number of control options at a source was approximately five.

The choice of control techniques was discrete, i.e., when a control was chosen, it was applied to the entire emissions emanating from a source. Therefore, this problem requires a combinatorial search through approximately 5^{300} potential control strategies.

Associated with each control technique are emissions removal efficiency and a cost-per-ton of emissions removed. The mathematical representation of the problem is as follows:

$$\text{minimize } C = \sum_j \sum_t A_j [c_{j,t}, u_j] \qquad (9)$$

$$\text{s.t. } \frac{\sum_j [u_j * (e_{j,t} * c_{j,t})]}{\sum_j u_j} \geq T \qquad (10)$$

$$c_{j,t} \in \{0,1\} \quad \forall j,t \qquad (11)$$

$$\sum_t c_{j,t} = 1 \quad \forall j \qquad (12)$$

Where: C is the total cost; A_j is the cost at source j; $c_{j,t}$ is the binary variable representing whether control option t is used at source j; u_j is the emissions at source j; $e_{j,t}$ is the removal efficiency of control t at source j, and T is the target level of emissions reduction. The problem is described in more detail by Laughlin (1995).

This problem can be modeled and solved as an integer programming (IP) formulation, allowing the comparison of the GA results with the actual Pareto otimal solutions obtained using the IP. Also, this problem is closely related to another formulation that we have been examining in our research, in which a highly nonlinear air quality model is used in the evaluation function to predict air quality impacts. This alternative GA formulation seeks to minimize costs subject to an ambient air quality constraint. The complexity of the atmospheric transport and chemistry makes the use of more classical optimization methodologies intractable.

Also, because of the duration of an atmospheric model run, the efficiency and quick convergence of a MOGA technique are important, providing the impetus for the work presented in this paper.

3.3.4.1 Methodology

Four MO methods were tested and compared for the case study problem. Each method is described briefly.

1. Integer Programming

First, an integer-programming (IP) model was set up and solved to determine the true Pareto optimal costs for achieving emissions reductions of 5%, 10%, 15%, 20%, 25%, 30%, 35%, 40%, 45%, 50%, and 55%. The IP model contained approximately 1300 binary decision variables.

2. Single-Objective GA

Next, a single-objective GA (SOGA) was run iteratively, one run for each level of reduction, i.e. 5%, 10%, etc. The SOGA, as well as the MOGA techniques tested in this paper, used real-coded genes. Coding of the decision variables is discussed by Laughlin (1995).

Other parameters used in the SOGA runs included: a mutation rate of 0.1%, a uniform crossover rate of 35%, tournament selection, and elitism to preserve the best individual in each generation. 30% of genes that underwent crossover were subjected to a linear recombination. These parameter values were found to perform well through trials of various parameterizations. Population sizes of 50 and 150 individuals were tested, each for 15 different random number seeds. The runs were terminated when no improvement in fitness was observed in 20 consecutive generations.

3. The Neighborhood Constraint Method

NCM was then used to generate the tradeoff curve in a single run. The NCM-specific parameters used in this problem are listed in Table 3.4.

Table 3.4. Parameters and options used in the neighborhood constraint method.

Parameter	Implementation
Decision variable encoding:	real-valued
Mutation rate:	0.1%
Selection method:	tournament selection (within selection radius)
Uniform crossover rate:	15%
Selection radius:	4% of population size
Selection outside radius:	2%
Elitism:	Locational
Fitness evaluation:	Dynamic T_{bias} and T_{mes}. Local fitness evaluation
Population sorting:	After evaluation of the initial population

For NCM, only 10% of the genes that underwent crossover were linearly recombined. While this value may appear low, trials suggested that a low value allowed for better gene migration across the population with less chance of disruption.

Once again, runs were made for population sizes of 50 and 150 individuals for 15 different random number seeds. Termination occurred after 20 generations with no improvement in the average cost of individuals in the population.

Pareto Ranking and Hybrid Niched-Pareto MOGAs

Pareto ranking and a hybrid niched-Pareto MOGA technique were implemented and tested for the same population sizes and random seeds. The Pareto ranking scheme was based on the procedure described by Goldberg (1989).

The hybrid technique was a combination of two techniques described in the literature, and made use of:

- **Objective sharing** - Combined with tournament selection, this technique encouraged exploration of the tradeoff curve by promoting selection of individuals in relatively nonpopulated areas. The objective sharing implementation was based on the procedure described by Horn et al. (1994).

- **Restrictive mating** - This technique, adapted from Fonseca and Fleming (1993), used a parameter, mating, to encourage the mating of individuals that are similar in objective space.

Tournament selection was used in this hybrid method. A number of individuals equal to 10% (arbitrarily chosen) of the population size was sampled. If there was a tie in the ranking of the best individuals in the sample, continuously- updated objective sharing was used to identify the winner, as described by Horn et al. (1994). Another sampling of roughly 10% of the population was taken to find a mate. After the sampling was completed, the new set of individuals was evaluated to determine whether the individuals fell within an objective-space distance, mating, from the winner of the first tournament. The fitness of individuals that fell within this distance was increased to encourage selection. If there was a tie, objective sharing was carried out in the second sampled set to identify the mate.

4. Results

Table 3.5 summarizes the preliminary results of these trials. Costs obtained from each of the techniques, in millions of dollars per year, are provided for various target levels of reduction. Since the controls being considered are discrete, it may not be possible to meet a target reduction level exactly. Therefore, the cost displayed is for the strategy that most cost-effectively reduces greater than or equal to the current target, but less than the next higher target level.

Table 3.5. Results of the IP, SOGA, and MOGA runs. Population sizes of 50 and 150 were tested for 15 random number seeds.

Redux	Performance	Multiobjective Technique for Various Population Sizes									
		IP (1)	SOGA		NCM		Pareto (2)		Hybrid (3)	
		n/a	50	150	50	150	50	150	50	150
5%	Best	0.53	0.57	0.86	0.77	0.62	n/a	n/a	n/a	0.97
	Mean	n/a	0.93	0.95	1.18	0.72	n/a	n/a	n/a	1.46
	Stdev	n/a	0.25	0.06	0.33	0.07	n/a	n/a	n/a	0.43
10%	Best	1.26	1.33	1.41	1.54	1.31	n/a	n/a	2.32	1.81
	Mean	n/a	1.69	1.57	1.93	1.40	n/a	n/a	2.32	2.23
	Stdev	n/a	0.29	0.12	0.28	0.09	n/a	n/a	0.00	0.23
15%	Best	1.97	2.10	2.16	2.25	2.10	7.47	7.47	2.82	2.45
	Mean	n/a	2.65	2.39	2.84	2.20	7.47	7.47	3.79	3.40
	Stdev	n/a	0.51	0.12	0.41	0.08	0.00	0.00	1.35	0.63
20%	Best	2.73	3.67	3.06	3.12	2.91	n/a	n/a	3.63	3.42
	Mean	n/a	4.34	3.36	3.79	3.05	n/a	n/a	6.45	4.36
	Stdev	n/a	0.51	0.21	0.38	0.09	n/a	n/a	1.44	0.60
25%	Best	3.47	4.57	3.90	4.17	3.66	n/a	n/a	6.51	4.56
	Mean	n/a	5.85	4.11	4.86	3.77	n/a	n/a	7.52	5.78
	Stdev	n/a	0.90	0.26	0.55	0.07	n/a	n/a	0.87	0.81
30%	Best	5.47	6.17	6.17	6.42	6.24	n/a	n/a	7.99	6.74
	Mean	n/a	6.47	6.34	7.16	6.36	n/a	n/a	9.12	8.01
	Stdev	n/a	0.44	0.14	0.53	0.09	n/a	n/a	0.97	0.94
35%	Best	8.74	9.76	9.78	9.59	9.51	13.67	13.71	8.17	11.67
	Mean	n/a	9.78	9.78	11.08	10.17	13.95	13.97	8.53	12.78
	Stdev	n/a	0.00	0.00	2.79	0.40	0.16	0.16	0.40	0.80
40%	Best	13.60	15.05	18.36	17.39	17.64	15.12	14.97	17.10	16.90
	Mean	n/a	20.74	21.66	23.22	22.19	15.24	15.05	18.67	17.62
	Stdev	n/a	3.62	2.16	2.38	1.98	0.14	0.05	1.16	0.85
45%	Best	29.90	31.91	36.57	36.35	33.00	31.39	31.14	33.00	32.68
	Mean	n/a	34.72	36.70	39.11	36.45	31.45	31.17	34.34	33.20
	Stdev	n/a	2.57	0.10	2.82	1.05	0.05	0.02	1.30	0.42
50%	Best	44.25	47.89	47.61	53.35	49.20	47.74	47.61	49.50	48.66
	Mean	n/a	49.19	47.74	61.58	50.19	48.01	47.83	50.73	49.21
	Stdev	n/a	1.35	0.11	6.14	1.02	0.18	0.15	0.91	0.38
Approx. Total Evaluations		n/a	129000	330000	46500	198000	58000	200000	55000	200000

Notes:

(1) Population size, mean, and standard deviations are not applicable for the IP; results are Pareto optimal

(2) No nondominated solutions were found for ranges 5%, 10%, 20% and 25%. For 15%, the solution was one of the seeds used in the initial population.

(3) In roughly half the hybrid runs, no new nondominated solutions in the 0 to 35% reduction range were identified

Results are reported for population sizes of 50 and 150. Also compared are the numbers of fitness evaluations required by each method.

Figure 3.5 shows the nondominated sets and the final populations obtained through the NCM, the Pareto method, and the hybrid method, and compares them with IP the solutions.
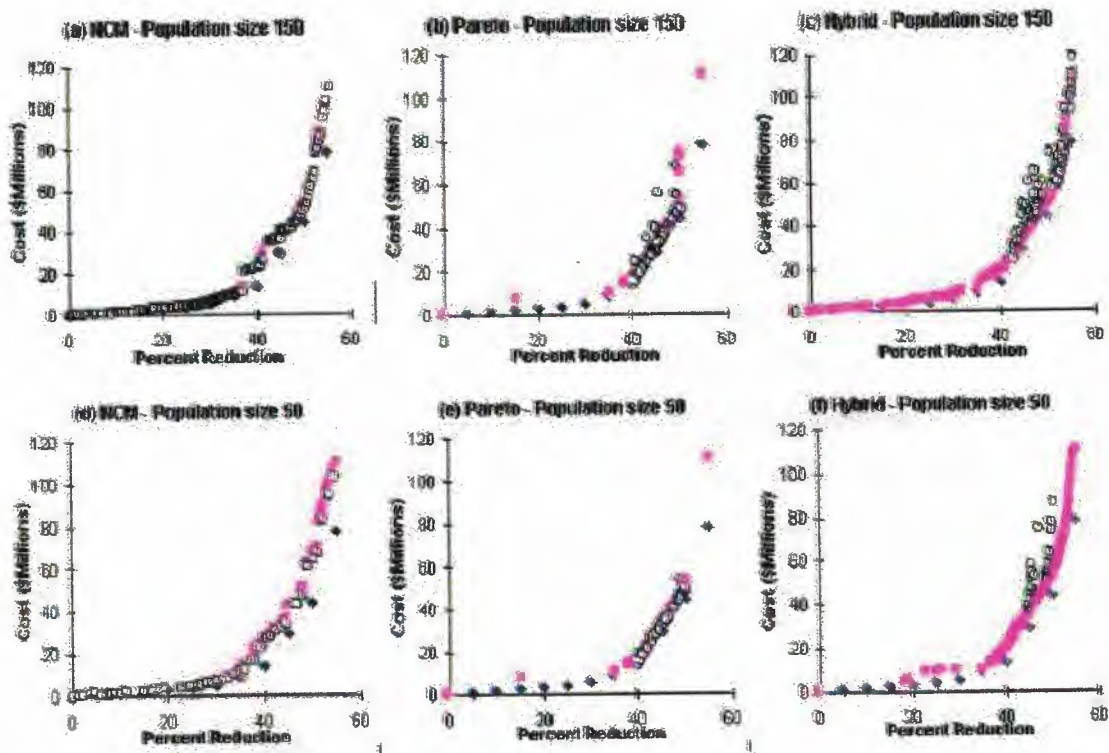

Figure 3.5. Final population and nondominated sets identified throughout a typical run, for each MOGA rechnique. Indeviduals in the final population are represented as light-colored circles. IP-generated Pareto solutions are black diamonds. The nondominated set is represented as gray squares

5. Discussion

Several observations may be made from the results. These include:

- None of the GA methods identified the true optimal points found by the IP. This was not unexpected because of the stochastic nature of GA search and the difficulty of the problem. However, the SOGA and MOGA techniques, with the exception of the Pareto method, did approximate the Pareto set to a degree that may be acceptable to a decision-maker.

- All three MOGA techniques required considerably fewer fitness evaluations and computational time than did the iterative SOGA technique. In about 1/3 of the duration of a SOGA run, which only identified 10 nondominated solutions, the MOGA techniques were able to identify several hundred to thousands of nondominated solutions.

- The Pareto method found better solutions near the elbow region of the curve, compared to the hybrid, NCM, and SOGA. The results suggest that the Pareto technique was superior at the elbow because it tended to direct more search effort in that region than did the NCM and hybrid.

- Of the MOGA runs, the hybrid technique, with a population size of 150, identified the best-nondominated solutions (Fig.3.5 c,f), providing good coverage of the entire curve. However, the hybrid technique was more dependent on the random number seed than NCM, and in several runs failed to identify solutions in the lower portion of the tradeoff curve. Although the hybrid technique was able to find good nondominated solutions, it was unable to maintain them throughout the search. The final population tended to converge toward the elbow. NCM was able to cover the noninferior set well, and also maintained the nondominated solutions throughout the search (Fig.3.5 a, d). This behavior may be desirable because it reduces the book-keeping required to maintain records of solutions along the set. The Pareto technique did poorly in covering the noninferior set and maintaining the nondominated solutions.

- NCM exhibited better coverage of the Pareto set than the hybrid technique for runs where the population was 50.

3.3.5 Summary

This paper presents a new multiobjective genetic algorithm technique, the neighborhood constraint method (NCM). NCM uses a restrictive mating scheme and location-dependent constraints to promote and maintain diversity in the GA population.

The application of the NCM technique to a realistic problem was demonstrated and its performance was compared with those of an integer programming procedure, an iterative procedure using a single-objective GA optimization, and a Pareto and a hybrid niched-Pareto multiobjective GA (MOGA) techniques.

Preliminary results show that NCM performs better than the SOGA and the other MOGA techniques. NCM appears to provide good coverage of the Pareto set, and is capable of promoting and maintaining population diversity throughout the GA run with a relatively small population size. It does appear, however, that the Pareto-based techniques may be better at identifying regions near the elbow of the Pareto set.

These results suggest that NCM is a viable MOGA technique, and point to the need for a more comprehensive investigation of NCM to better identify the effects of the restrictive mating and the location-dependent constraints. A determination of the best selection approach (i.e., roulette wheel, tournament, etc.), selection radius, and recombination and crossover rates to use with NCM would be beneficial. NCM should also undergo a more rigorous comparison with the other MOGA techniques for a set of well-known problems.

Another potential use of NCM that deserves attention is in parallel GAs, since the methodology adapts well to the decentralized selection schemes discussed by De Jong and Sarma (1995).

CONCLUSION

A couple of conclusions from building block theory are of importance to note. Strings with very fit schemata of short length will have a high likelihood of being selected to create the next population, and thus pass on those schemata to strings in the new population. It has been shown that schemata of this form increase in number from one population to the next in an exponential fashion. In other words, n^3 useful schemata are processed per generation, and the majority of these have small orders and lengths associated with them. These schemata are what give a GA the power to efficiently search through a problem space. This n^3 feature is so important to GAs that it has been given a special name, implicit parallelism.

If the conception of a computer algorithms being based on the evolutionary of organism is surprising, the extensiveness with which this algorithms is applied in so many areas is no less than astonishing. These applications, be they commercial, educational and scientific, are increasingly dependent on this algorithms, the Genetic Algorithms. Its usefulness and gracefulness of solving problems has made it the more favorite choice among the traditional methods, namely gradient search, random search and others. GAs are very helpful when the developer does not have precise domain expertise, because GAs possess the ability to explore and learn from their domain.

In this project, the use of operators of GAs in optimization of engineering and commerce are considered. We believe that, by these interesting examples, one could grasp the idea of GAs with greater ease. The different optimization problems are described. The application of GA to solve optimization problem are given the selection procedure model parameters by using GA operators are represented. Also different problems solution, by using GA, is given.

In future, the developments of variants of GAs to tailor for some very specific tasks will be interesting. This might defy the very principle of GAs that it is ignorant of the problem domain when used to solve problem. But we would realize that this practice could make GAs even more powerful

REFERENCES

[1] *Water Resources Research*, Cieniawski, S. E., Eheart, J. W., and Ranjithan, S. (1995).

[2] *Genetic Algorithms in Engineering and Computer Science* , edited by G.Winter, J.Periaux & M.Galan, published by JOHN WILEY & SON Ltd. in 1995.

[3] [Louis 1993] *Genetic Algorithms as a Computational Tool for Design*, by Sushil J. Louis, in August 1993

[4] *Algorithms and Complexity*, by Herbert S.Wilf, in 1986 published by Prentice-Hall, Inc.

Obtained: Central Library of Imperial College (3 Computing 5.25 WIL)

[5] *Proceedings of the Sixth International Conference on Genetic Algorithms*, De Jong, K. A. and Jayshree Sarma (1995).

[6] *Recombination Variability and Evolution: algorithms of estimation and population-genetic models*, by A.B.Korol, I.A.Preygel & S.I.Preygel, in 1994 published by Chapman & Hall.

Obtained : Central Library of Imperial College (4 Life Sciences 575.116.12 KOR)

[7] *Learning Robot Behaviours using Genetic Algorithms*, by ALAN C.Schultz. Navy Center for Applied Research in Artificial Intellignece.

[8] *Genetic Algorithms for Order Dependent Processes applied to Robot Path-Planning*, by Yuval Davidor, in April 1989 published by Imperial College

[9] *Genetic Algorithms in Business and Their Supportive Role in Decision Making*, by Tom Bodnovich, in 16 November 1995, published by College of Business Administration Kent State

[10] *Proceedings of the Fifth International Conference On Genetic Algorithms*, Fonseca, C. M., Fleming, P. J. (1993).

[11] **http://www.cs.qub.ac.uk/~M.Sullivan/ga/ga3.html**

[12] **http://www.doc.ic.ac.uk/~nd/surprise_96/journal/vol4/tcw2**

[13] **http://ltswww.epfl.ch/pub_files/brigger/thesis.html**

[14] **http://www-fp.mcs.anl.gov/otc/Guide/**

[15] **http://www.genetic-programming.com/creation.gif**