

A Greedy Genetic Algorithm for the
Quadratic Assignment Problem

by
Ravindra K. Ahuja
James B. Orlin
Ashish Tiwari

WP #3826-95

June 1995

A Greedy Genetic Algorithm for the Quadratic Assignment Problem

Ravindra K. Ahuja
Dept. of Industrial & Management Engg.
Indian Institute of Technology
Kanpur 208016 INDIA
email: ahuja@iitk.ernet.in

James B. Orlin
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
email: jorlin@mit.edu

Ashish Tiwari
Department of Computer Science
State University of New York
Stony Brook, NY 11794-4400

A Greedy Genetic Algorithm for the Quadratic Assignment Problem

Ravindra K. Ahuja¹, James B. Orlin², and Ashish Tiwari³

ABSTRACT

The Quadratic Assignment Problem (QAP) is one of the classical combinatorial optimization problems and is known for its diverse applications. In this paper, we suggest a genetic algorithm for the QAP and report its computational behavior. The genetic algorithm incorporates many greedy principles in its design and, hence, is called the *greedy genetic algorithm*. The ideas we incorporate in the greedy genetic algorithm include (i) generating the initial population using a randomized construction heuristic; (ii) new crossover schemes; (iii) a special purpose immigration scheme that promotes diversity; (iv) periodic local optimization of a subset of the population; (v) tournamenting among different populations; and (vi) an overall design that attempts to strike a balance between diversity and a bias towards fitter individuals. We test our algorithm on all the benchmark instances of QAPLIB, a well-known library of QAP instances. Out of the 132 total instances in QAPLIB of varied sizes, the greedy genetic algorithm obtained the best known solution for 103 instances, and for the remaining instances (except one) found solutions within 1% of the best known solutions. Based on our computational testing, we believe that the greedy genetic algorithm is the best heuristic algorithm for dense QAP developed to date in terms of the quality of the solution.

¹ Dept. of Industrial & Management Engg., Indian Institute of Technology, Kanpur 208 016, INDIA.

² Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

³ Department of Computer Science, State University of New York, Stony Brook, NY 11794.

1. INTRODUCTION

The Quadratic Assignment Problem (QAP) is one of the classical combinatorial optimization problems and is widely regarded as one of the most difficult problem in this class. Given a set $N = \{1, 2, \dots, n\}$, and $n \times n$ matrices $F = \{f_{ij}\}$, $D = \{d_{ij}\}$, and $C = \{c_{ij}\}$, the QAP is to find a permutation ϕ of the set N which minimizes

$$z = \sum_{i=1}^n \sum_{j=1}^n f_{ij} d_{\phi(i)\phi(j)} + \sum_{i=1}^n c_{i\phi(i)}.$$

As an application of the QAP, consider the following campus planning problem. On a campus, new facilities are to be erected and the objective is to minimize the total walking distances for students and staff. Suppose there are n available sites and n facilities to locate. Let d_{kl} denote the walking distance between the two sites k and l where the new facilities will be erected. Further, let f_{ij} denote the number of people per week who travel between the facilities i and j . Then, the decision problem is to assign facilities to sites so that the walking distance of people is minimized. Each assignment can be mathematically described by a permutation ϕ of $N = \{1, 2, \dots, n\}$ such that $\phi(i) = k$ means that the facility i is assigned to site k . The product $f_{ij} d_{\phi(i)\phi(j)}$ describes the weekly walking distance of people who travel between facilities i and j . Consequently, the problem of minimizing the total walking distance reduces to identifying a ϕ that minimizes the function z defined above. This is an application of the QAP with each $c_{ik} = 0$. In this application, we have assumed that the cost of erecting a facility does not depend upon the site. In case it does, we will denote by c_{ik} the cost of erecting facility i at site k , and these costs will also play a role in determining the optimal assignment of facilities to sites.

Additional applications of QAP include (i) the allocation of plants to candidate locations; (ii) layout of plants; (iii) backboard wiring problem; (iv) design of control panels and typewriter keyboards; (v) balancing turbine runners; (vi) ordering of interrelated data on a magnetic tape; (vii) processor-to-processor assignment in a distributed processing environment; (viii) placement problem in VLSI design; (ix) analyzing chemical reactions for organic compounds; and (x) ranking of archaeological data. The details and references for these and additional applications can be found in Burkard [1991], Malucelli [1993], and Pardalos, Rendl and Wolcoviez [1991].

On account of its diverse applications and the intrinsic difficulty of the problem, the QAP has been investigated extensively by the research community. The QAP has been proved to be an \mathcal{NP} -complete problem, and a variety of exact and heuristic algorithms have been proposed. Exact algorithms for QAP include approaches based on (i) dynamic programming (Christofides and Benavent [1989]); (ii) cutting planes (Bazaraa and Sherali [1980]); and (iii) branch and bound

(Lawler [1963], Pardalos and Crouse [1989]). Among these, the branch and bound algorithms obtain the best solution, but are unable to solve problems of size larger than $n = 20$.

Since many applications of QAP give rise to problems of size greater than 20, there is a special need for good heuristics for QAP that can solve large size problems. Known heuristics for QAP can be classified into the following categories: (i) construction methods (Armour and Buffa [1963], Buffa, Armour and Vollmann [1962]); (ii) limited enumeration methods (West [1983], Burkard and Bonniger [1983]); (iii) local improvement methods (Li, Pardalos and Resende [1991]); (iv) simulated annealing methods (Wilhelm and Ward [1987]); (v) tabu search methods (Skorin-Kapov [1990], Taillard [1991]); and (vi) genetic algorithms (Fleurent and Ferland [1993], Bean [1994]). Among these, (i) the tabu search method due to Skorin-Kapov [1990] and (ii) the (randomized) local improvement method due to Li, Pardalos and Resende [1994] and Resende, Pardalos and Li [1994], which they named as *GRASP* (Greedy Randomized Adaptive Search Procedure), are the two most accurate heuristic algorithms to solve the QAP.

In the past, genetic algorithms have been applied to a variety of combinatorial optimization problems. In this paper, we suggest a genetic algorithm for the QAP which incorporates many ideas based on greedy principles. We call this algorithm the *greedy genetic algorithm*. The ideas we incorporate include (i) generating the initial population using a good (randomized) construction heuristic; (ii) new crossover schemes; (iii) special-purpose immigrations that promote diversity; (iv) periodic local optimization of a subset of the population; (v) tournamenting among different populations; and (vi) developing an overall design that attempts to strike a right balance between diversification and a bias towards fitter individuals. We test each of these ideas separately to assess its impact on the algorithm performance, and also compare our final algorithm on all benchmark instances in QAPLIB, a well-known library of QAP instances compiled by Burkard, Karisch and Rendl [1994].

We find the greedy genetic algorithm to be a very robust algorithm for QAP; out of the 132 total instances in QAPLIB of varied sizes, it obtained the best known solution for 103 instances, and for the remaining instances (except one) found solutions within 1% of the best known solutions. We also compared the greedy genetic algorithm with GRASP on a selected set of instances from QAPLIB and found it to outperform GRASP. Based on our computational testing, we believe that the greedy genetic algorithm is the best heuristic algorithm for dense QAP developed to date in terms of the quality of the solution. Our algorithm is also highly parallelizable. We note that the performance of our greedy genetic algorithm depends upon the various parameter settings and using a carefully selected mixture of various greedy principles, which we obtained through extensive computational testing. We anticipate that extensive testing and experimentation would be valuable if one were to apply the ideas presented here to solve related combinatorial problems.

2. Greedy Genetic Algorithm for QAP

Genetic algorithms represent a powerful and robust approach for developing heuristics for large-scale combinatorial optimization problems. The motivation underlying genetic algorithms can be expressed as follows: Evolution has been remarkably successful at developing complex and well-adapted species through relatively simple evolutionary mechanisms. A natural question is the following: What ideas can we adapt from our understanding of evolution theory so as to solve problems in other domains? This fundamental question has many different answers because of the richness of evolutionary phenomenon. Holland [1975] provided the first answer to this question by developing genetic algorithms.

Genetic algorithms (GA's) imitate the process of evolution on an optimization problem. Each feasible solution of a problem is treated as an individual whose fitness is governed by the corresponding objective function value. A GA maintains a population of chromosomes on which the concept of the survival of the fittest (among string structures) is applied. There is a structured yet randomized information exchange between two individuals (*crossover* operator) to give rise to better individuals. Diversity is also added to the population by randomly changing some genes (*mutation* operator) or bringing in new individuals (*immigration* operator). A GA repeatedly applies these processes until the population converges.

GA's can be implemented in a variety of ways. The excellent books by Goldberg [1989] and Davis [1991] describe many possible variants of GA's. We also refer to these books for various GA notation such as chromosomes, genes, etc. In our research, we have used the GA with "steady-state reproduction" instead of the traditional "generational reproduction" because researchers have found it to be faster in practice (see, for example, Davis [1991]). Our GA has the following high-level description:

```

algorithm genetic;
begin
    obtain initial population;
    repeat
        select two individuals  $I_1$  and  $I_2$  in the population;
        apply the crossover operator on  $I_1$  and  $I_2$  to produce a child  $I_3$ ;
        choose two individuals out of  $I_1$ ,  $I_2$  and  $I_3$ ;
        occasionally perform immigration;
    until the population converges;
end;

```

This is the simplest description of our GA. Each execution of the repeat loop is called a *trial*. We will now describe how various steps of the GA are implemented when specialized for the QAP. We have incorporated several heuristics based on greedy principles to improve the performance of the GA. We will also describe the details of these heuristics.

Coding Scheme

A coding scheme maps feasible solutions of the problem to strings. The effectiveness of the crossover operator depends greatly on the coding scheme used. The coding should be such that the crossover operator preserves high performing partial arrangements (schemata) of strings, and minor changes in the chromosome translate into minor changes in the corresponding solution. For the QAP, a natural encoding is the permutation of n numbers in the set $N = \{1, 2, \dots, n\}$, where the j th number in the permutation denotes the facility located on site j . In our implementation, we have used this encoding scheme.

Initial Population Generation

The performance of a GA is often sensitive to the quality of its initial population. The "goodness" of the initial population depends both on the average fitness (that is, the objective function value) of individuals in the population and the diversity in the population. Losing on either count tends to produce a poor GA. By having an initial population with better fitness values, we typically get better final individuals. Further, high diversity in the population inhibits early convergence to a locally optimal solution.

For the QAP, several heuristics can be used to generate the initial population. In our implementation, we have used the initial solutions produced by GRASP (see Li, Pardalos, and Resende [1991]). GRASP consists of two phases: a construction phase and an improvement phase. The construction phase uses a randomized greedy function to assign facilities to locations one-by-one, at each step minimizing the total cost with respect to the assignments already made. The improvement phase uses a neighborhood search technique to improve the solution obtained by the first phase iteratively. Each application of GRASP yields a (possibly) different solution because of the randomization used in the construction phase. In our implementation of the GA, we have used the initial population comprising of solutions produced by the construction phase of GRASP. We didn't apply the improvement phase because it would have taken substantially more time and computational tests suggested that it would not have led to better overall results.

Selection

The selection criteria is used to select the two parents to apply the crossover operator. The appropriateness of a selection criteria for a GA depends upon the other GA operators chosen. In the literature, a typical selection criteria gives a higher priority to fitter individuals since this leads to a faster convergence of the GA. Nevertheless, we obtained better overall results when the selection is not biased towards fitter individuals. In the GA, we selected both parents randomly giving equal probability of selection to each individual in the population.

Crossover

The crossover scheme is widely acknowledged as critical to the success of GA. The burden of "exploration" and "exploitation" falls upon it simultaneously. The crossover scheme should be capable of producing a new feasible solution (i.e., a child) by combining good characteristics of both parents. Preferably, the child should be considerably different from each parent. We observed that when the initial population is generated using a good heuristic method, then standard crossover schemes, such as the order crossover given in Davis [1991], generally fail to produce individuals better than both the parents. As the order crossover scheme did not give good results, we tested two more sophisticated crossover schemes, called the (i) path crossover scheme, and (ii) optimized crossover scheme. We shall now describe these schemes in more detail. We remark that these crossover schemes are not specific to the QAP, and would apply to any problem with chromosomes represented by permutations.

Path Crossover

Let I_1 and I_2 denote two individuals. Applying a crossover scheme to these two parents produces a child, which we denote by I_3 . Let the chromosomes of I_1 , I_2 , and I_3 be represented by $a_1 - a_2 - \dots - a_n$, $b_1 - b_2 - \dots - b_n$, and $c_1 - c_2 - \dots - c_n$, respectively. In any crossover scheme that we consider, the child inherits any genes common to both the parents; that is, if $a_k = b_k$ for some k , then $c_k = a_k = b_k$. In the two types of path crossovers described next, we define a path $p_1 - p_2 - \dots - p_r$ where (i) each node in the path corresponds to an individual; (ii) the first node p_1 corresponds to I_1 , the last node p_r corresponds to I_2 , and, (iii) for each $1 \leq k \leq r-1$, p_{k+1} is obtained from p_k by performing a primitive transformation. We considered two types of primitive transformations: *swap* and *insert*. If we use a swap transformation, we get the *swap path crossover* scheme. If we use an insert transformation, we get the *insert path crossover* scheme. The path crossover was originally developed by Glover [1994], who referred to it as "path relinking." We use the alternative term "path crossover," because we think that it is more consistent with the usual terminology used in genetic algorithms.

In the swap path crossover scheme, we start at some position of the chromosomes, which is determined using a random process, and examine the chromosomes corresponding to I_1 and I_2 from left to right in a cyclic fashion. If the alleles at the position (or gene) being looked at are the same, we move to the next position; otherwise, we perform a swap of the alleles of two genes in I_1 or in I_2 , whichever gives the fitter solution, so that the alleles at the position being looked at become alike. We repeat this process until all genes have been considered. All chromosomes obtained using this process are valid children of I_1 and I_2 ; out of them we let I_3 denote the fittest child.

We illustrate the swap path crossover in Figure 1(a). Suppose we start at the first position. In parent 1, facility 5 is located at the first site and in parent 2 facility is located at this site. There are two ways in which the two parents can move closer to one-another; by swapping the sites of the facilities 2 and 5 in parent 1 or in parent 2. We compute the objective function values of these two solutions and we perform the swap for which the corresponding solution has a lower cost; the resulting solution is the first child. We then consider the alleles in the two resulting solutions starting at the second position and obtain possibly several children; the best among these is I_3 . In this example, we start at the first position; in general, we start at a position selected randomly and examine subsequent positions from left to right in a wrap-around fashion until all positions have been considered.

The insert transformation is similar to the swap transformation; the only difference is that to have the same alleles at a given position in the two chromosomes, we insert an allele at a new place thereby shifting the other alleles to the right in a way so that common alleles do not move. We illustrate the insert path crossover in Figure 1(b). For parent 1, the insert transformation consists of removing facility 2 from site 2 and moving to site 1 and facility 5 shifts to site 2. For parent 2, this transformation causes bigger changes. It moves facility 5 to site 1, facility 2 moves to site 2, facility 1 moves to site 5 (because facilities at sites 3 and 4 are fixed), and facility 6 moves to site 6. It is easy to see that the insert transformation causes changes which are radically different than the swaps.

Parent 1:	5 - 2 - 3 - 4 - 1 - 7 - 6	Swap in Parent 1:	2 - 5 - 3 - 4 - 1 - 7 - 6
Parent 2:	2 - 1 - 3 - 4 - 6 - 5 - 7	Swap in Parent 2:	5 - 1 - 3 - 4 - 6 - 2 - 7
(a)			
Parent 1:	5 - 2 - 3 - 4 - 1 - 7 - 6	Insert in Parent 1:	2 - 5 - 3 - 4 - 1 - 7 - 6
Parent 2:	2 - 1 - 3 - 4 - 6 - 5 - 7	Insert in Parent 2:	5 - 2 - 3 - 4 - 1 - 6 - 7
(b)			

Figure 1. Illustrating swap and insert path crossovers. The common genes are shown in bold type.

Use of the swap or insert path crossover scheme yields one child I_3 . One possible scheme is to keep the best two individuals among the two parents (I_1 and I_2) and the one child (I_3). We found that this greedy scheme leads to premature convergence. If the parents are very good, the child is usually less fit than both parents, and the crossover operation does not change the population. Rather, we required the child to replace one of the parents. If the child is fitter than both the parents, then it replaces the parent with which it is more similar; otherwise, the child replaces

the worse parent. This crossover scheme ensures that the children not only inherit useful information from their parents but also search new regions of the solution space. Further, our rule for replacement of a parent by the child ensures that the best individual in the population is never replaced. Such an approach is called as *elitist* approach in the literature.

Optimized Crossover

In the optimized crossover scheme, the genes of the parents are selected so as to optimize the fitness of the child. Use of optimized crossover in GA's has yielded encouraging results for the independent set problem (Aggarwal, Orlin and Tai [1994]).

Using the optimized crossover scheme, the two parents produce a child, which is called the *optimized child*. We will now explain how to determine an optimized child for QAP. Let $I_1 = a_1 - a_2 - \dots - a_n$ and $I_2 = b_1 - b_2 - \dots - b_n$ be two parents. An optimized child of I_1 and I_2 is an individual $I_3 = c_1 - c_2 - \dots - c_n$ (i) which satisfies $c_i = a_i$ or b_i for all $1 \leq i \leq n$, and (ii) has the smallest objective function value among individuals satisfying (i). To understand how to obtain the optimized child, a network formulation will be helpful. Each solution of a QAP can be represented by a perfect matching in a complete undirected bipartite graph $G = (U \cup V, U \times V)$ with $U = \{u_1, u_2, \dots, u_n\}$ representing facilities and $V = \{v_1, v_2, \dots, v_n\}$ representing sites. We can represent an individual I by an arc set E in G as follows: $(u_i, v_j) \in E$ if and only if facility i is located at site j . Now, consider the union of two arc sets E_1 and E_2 corresponding to two parents I_1 and I_2 . An optimized child of I_1 and I_2 is a perfect matching in $G' = (U \cup V, E_1 \cup E_2)$ with the least objective function value. See, for example, Figure 2. $E_1 \cup E_2$ is a collection of (i) singleton arcs (corresponding to common alleles), and (ii) node-disjoint cycles (corresponding to different alleles). Suppose that there are k cycles. In this case, there are exactly 2^k perfect matchings in G' , because once we select any arc in any node-disjoint cycle for inclusion in the matching, we must select every other alternate arc. Hence arcs in any cycle can be selected in exactly two ways. Since there are k cycles, there will be exactly 2^k children of I_1 and I_2 ; for each such child, we determine the objective function value, and the child with the least objective function value is the optimized child. This child replaces one of the parents using the same strategy as in the path crossover scheme. One might wonder whether we can identify an optimized child in polynomial time. This is not likely because it can be shown (proof omitted) that determining an optimized child is also an \mathcal{NP} -hard problem. Nevertheless, the exhaustive enumeration method may be efficient in practice because k may be very small. For most problems in our test bed, the number of cycles is less than 5.




Figure 2. Illustrating the optimized crossover. There are two cycles and four candidates for the optimized child.

In our investigations, we used all three crossover schemes outlined above. In preliminary testing, we obtained best results from the insert path crossover scheme. Subsequently, we performed detailed computational testing using the insert path crossover scheme.

Mutation/Immigration

Mutation refers to the process of increasing diversity in the population by introducing random variations in the members of the population. Researchers have used immigration in place of mutation to achieve the same purpose (see, for example, Davis [1991]). Immigration refers to the process of replacing poor members of the current population by bringing in new individuals. For our implementation of GA, we choose immigration. We also used a variant of mutation, which we discuss below in the subsection on local optimization.

We could use the same method to generate immigrants which we used to generate the initial population. We found that this did not introduce new genetic material to the population and the GA converged prematurely, possibly because the method did not introduce sufficient new genetic material. Therefore, we use a different scheme to generate immigrants that produces individuals significantly different than the ones already in the population and hence forces the algorithm to search newer regions of solution space.

Let us define an $n \times n$ matrix $seen[i, j]$, whose $(i, j)^{th}$ position stores the number of individuals in the population (present as well as all past members) in which facility j has been assigned to site i . A lower value of $seen[i, j]$ implies that fewer individuals in the population had facility i assigned to site j . To promote diversity, we encourage this assignment in new individuals. We

generate a new individual for immigration using the following method. We first select a permutation of n numbers obtained using a random process; this gives us the order in which we consider the sites one-by-one. To each site under consideration, say site k , we assign the unassigned facility l for which $\text{seen}[k, l]$ is minimum. This assignment promotes individuals from the under-explored solution space, and though the process does not explicitly use randomness, it typically yields new individuals.

We experimented with three rates of immigration: 10% (that is, one immigration after every 10 trials), 20%, and variable immigration rate. In the variable immigration rate, we start with the immigration rate of 10% and increase this rate by 2% after every 200 trials. We also introduced immigration when we found that the diversity in the population is below an accepted threshold level.

Local Optimization

Commonly used approaches for solving \mathcal{NP} -hard problems are local optimization, simulated annealing, and tabu search; all three approaches rely on a neighborhood structure. For the QAP problem, the 2-exchange neighborhood (that is, obtained by swapping two numbers in the permutation) has been extensively investigated and found to be a fairly powerful heuristic. Local optimization based upon the 2-exchange neighborhood can be used to improve the quality of any population. We observed that periodic local optimization of a subset of the population improves the performance of the GA, and hence we incorporated it into the algorithm. After every 200 trials, we locally optimize 20% of the population. We store individuals using arrays: after 200 trials we locally optimize the first 20%, after 400 trials we locally optimize the next 20%, and so on. When we reach the end of the array, we return to the beginning of the array.

Tournamenting

In principle, one can apply a GA many times starting with different populations, and choose the best individual obtained among all the runs. As an alternative, we consider the following approach, one that takes the final population of two different runs, keeps 50% of the individuals in the two runs, and applies the GA again with this mixed population as the initial population. The basic idea behind this scheme is that it allows the best alleles of two different runs to be combined through crossover operator to create fitter individuals. In a sense it permits the GA to search regions that are "in between" the regions to which the two populations converged. Generalizing this idea further for multiple runs gives us a GA with "tournamenting". Figure 3 describes the tournamenting scheme with four teams; one can easily generalize this approach to a larger number of teams.

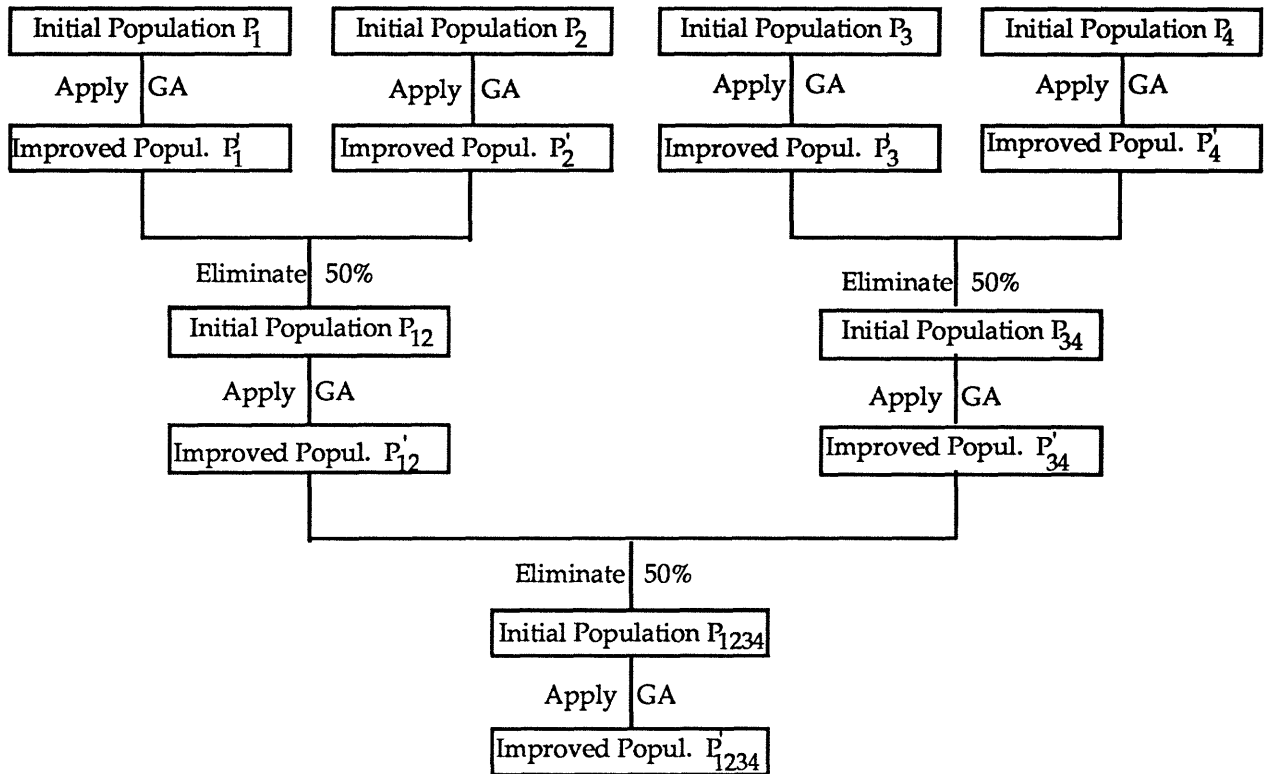


Figure 3. GA with tournamenting with four teams.

When we use tournamenting, we terminate the execution of a GA before the population converges, since we found that allowing the GA to run until the population converges leads to poorer solutions, possibly because of the lack of diversity in the final populations. We found that after 1,000 trials of our GA, the population quality is good and there is sufficient diversity. We choose to terminate the GA after 1,000 trials.

We considered several options for eliminating 50% of the two populations, say P'_1 and P'_2 in order to obtain the initial population P_{12} for the next round, including the following two options: (i) select the best 50% of the union of the two populations; and (ii) the union of the best 50% of the populations P'_1 and P'_2 . Both these schemes have the drawback that multiple copies of the fitter individuals might dominate the population. We obtained a better mix by performing a one-to-one competition between individuals of the two populations; we compare the k th individual of P'_1 with the k th individual of P'_2 and the fitter one moves to the next round.

3. COMPUTATIONAL RESULTS

In this section, we present the computational results of the greedy GA when applied to all the 132 instances of the QAPLIB compiled by Burkard, Karisch and Rendl [1994]. We present results that give insight into the behavior of our algorithm for QAP, and also give insight into the various

greedy principles we have incorporated into our algorithm. We also compare our algorithm with GRASP, and the random keys GA due to Bean [1994].

Computational Time and Accuracy. We present in Table 1 the running times and the accuracy of the greedy GA applied to all the instances in QAPLIB. Problem sizes in QAPLIB vary from 10 to 100, and all problems are dense. The middle numeric digits in the problem identification state the problem size. We present results for the following three versions of the algorithm.

Version 1. The initial population is generated using Phase 1 of GRASP, all individuals in the population are equally likely to be selected for crossover; insert path crossover is used; immigration rate is set to 10%; after every 200 trials, 20% of the population is locally optimized; and tournamenting with four teams is used.

Version 2. Same as Version 1 except that the immigration rate is variable. The algorithm starts with the immigration rate of 10%, and after every 200 trials the immigration rate is increased by 2%.

Version 3. Same as Version 2 except that the tournamenting with eight teams is used.

The objective function value of the best known solution is given in Burkard, Karisch and Rendl [1994]. We will give only the percent deviation from the best known solution, in the columns marked GA-1, GA-2, and GA-3. We give the computational times only for version 1; the time taken by version 2 is comparable to the time taken by version 1, and version 3 takes time approximately twice the times taken by version 1.

We find that the second version (in which immigration rate is variable) has a somewhat better overall performance than the version in which the immigration rate is constant. For small size problems, all versions obtain the best known solutions. Only for the larger size problems are the algorithmic differences apparent. The greedy GA with tournamenting with eight teams has a very robust performance. We find that out of 132 problem instances, the greedy GA obtains the best known solution for 103 problem instances; and out of the remaining 29 problems only one problem has deviation more than one percent from the best known solution. Moreover, we obtained these results by applying the algorithm only once. We anticipate that the results will improve if we apply the greedy GA several times starting with different seeds for the random number generator.

We now present results contrasting the various schemes employed by us in the greedy GA. We select Version 1 of the greedy GA (without tournamenting) as the benchmark algorithm and investigate how its performance changes as we modify its attributes (such as, the selection scheme, or the crossover). For the sake of brevity, we do not present results for modifications made in two or more schemes simultaneously; we tried some of these modifications, and found them to be by and

Table 1. Computational results of several versions of greedy genetic algorithms.

	Problem Name	GA-1	Time	GA-2	GA-3	Problem Name	GA-1	Time	GA-2	GA-3
1	bur26a.dat	0.00	117.3	0.00	0.00	esc16h.dat	0.00	24.0	0.00	0.00
2	bur26b.dat	0.00	112.7	0.00	0.00	esc16i.dat	0.00	25.8	0.00	0.00
3	bur26c.dat	0.00	113.5	0.00	0.00	esc16j.dat	0.00	201.1	0.00	0.00
4	bur26d.dat	0.00	106.7	0.00	0.00	esc32a.dat	3.08	190.9	0.00	0.00
5	bur26e.dat	0.00	109.1	0.00	0.00	esc32b.dat	0.00	200.1	0.00	0.00
6	bur26f.dat	0.00	102.2	0.00	0.00	esc32c.dat	0.00	194.6	0.00	0.00
7	bur26g.dat	0.00	97.1	0.00	0.00	esc32d.dat	0.00	176.3	0.00	0.00
8	bur26h.dat	0.00	101.9	0.00	0.00	esc32e.dat	0.00	184.9	0.00	0.00
9	car10ga.dat	0.00	3.9	0.00	0.00	esc32f.dat	0.00	184.4	0.00	0.00
10	car10gb.dat	0.00	3.9	0.00	0.00	esc32g.dat	0.00	185.5	0.00	0.00
11	car10gc.dat	0.00	4.3	0.00	0.00	esc32h.dat	0.00	174.5	0.00	0.00
12	car10gd.dat	0.00	4.2	0.00	0.00	esc64a.dat	0.00	1315.4	0.00	0.00
13	car10ge.dat	0.00	4.0	0.00	0.00	kra30a.dat	1.34	150.7	1.57	0.00
14	car10gf.dat	0.00	4.4	0.00	0.00	kra30b.dat	0.18	165.3	0.08	0.00
15	car10gg.dat	0.00	4.3	0.00	0.00	lipa10a.dat	0.00	4.8	0.00	0.00
16	car10gh.dat	0.00	4.0	0.00	0.00	lipa10b.dat	0.00	4.5	0.00	0.00
17	car10gi.dat	0.00	3.9	0.00	0.00	lipa20a.dat	0.00	37.4	0.00	0.00
18	car10gj.dat	0.00	4.1	0.00	0.00	lipa20b.dat	0.00	37.2	0.00	0.00
19	car10gk.dat	0.00	3.8	0.00	0.00	lipa30a.dat	0.00	172.3	0.00	0.00
20	car10gl.dat	0.00	4.9	0.00	0.00	lipa30b.dat	0.00	168.6	0.00	0.00
21	car10gm.dat	0.00	4.2	0.00	0.00	lipa40a.dat	0.00	510.9	1.10	0.96
22	car10gn.dat	0.00	4.3	0.00	0.00	lipa40b.dat	0.00	513.2	0.00	0.00
23	car10go.dat	0.00	3.9	0.00	0.00	lipa50a.dat	0.95	743.1	0.93	0.95
24	car10pa.dat	0.00	4.6	0.00	0.00	lipa50b.dat	0.00	754.7	0.00	0.00
25	car10pb.dat	0.00	4.3	0.00	0.00	lipa60a.dat	0.90	1528.5	0.82	0.77
26	car10pc.dat	0.00	4.3	0.00	0.00	lipa60b.dat	0.00	1523.7	0.00	0.00
27	car10pd.dat	0.00	4.4	0.00	0.00	lipa70a.dat	0.72	3074.2	0.73	0.71
28	car10pe.dat	0.00	3.9	0.00	0.00	lipa70b.dat	0.00	3061.6	0.00	0.00
29	car10pf.dat	0.00	4.5	0.00	0.00	lipa80a.dat	0.65	4759.5	0.67	0.61
30	car10pg.dat	0.00	4.5	0.00	0.00	lipa80b.dat	19.69	4749.4	0.00	0.00
31	car10ph.dat	0.00	4.2	0.00	0.00	lipa90a.dat	0.59	6179.2	0.62	0.58
32	car10pi.dat	0.00	4.0	0.00	0.00	lipa90b.dat	0.00	6159.7	0.00	0.00
33	car10pj.dat	0.00	4.6	0.00	0.00	nug05.dat	0.00	1.2	0.00	0.00
34	car10pk.dat	0.00	4.4	0.00	0.00	nug06.dat	0.00	1.8	0.00	0.00
35	car10pl.dat	0.00	4.6	0.00	0.00	nug07.dat	0.00	2.4	0.00	0.00
36	car10pm.dat	0.00	4.4	0.00	0.00	nug08.dat	0.00	3.4	0.00	0.00
37	car10pn.dat	0.00	4.1	0.00	0.00	nug12.dat	0.00	9.5	0.00	0.00
38	car10po.dat	0.00	3.8	0.00	0.00	nug15.dat	0.00	20.7	0.00	0.00
39	chr12a.dat	0.00	9.8	0.00	0.00	nug20.dat	0.00	48.9	0.00	0.00
40	chr12b.dat	0.00	9.2	0.00	0.00	nug30.dat	0.07	177.1	0.07	0.07
41	chr12c.dat	0.00	10.1	0.00	0.00	rou10.dat	0.00	5.1	0.00	0.00
42	chr15a.dat	0.00	20.3	0.00	0.40	rou12.dat	0.00	9.8	0.00	0.00
43	chr15b.dat	0.00	20.9	0.00	0.00	rou15.dat	0.00	17.3	0.00	0.00
44	chr15c.dat	0.00	22.0	4.59	0.00	rou20.dat	0.08	37.6	0.20	0.16
45	chr18a.dat	0.00	39.5	0.18	0.40	scr10.dat	0.00	5.3	0.00	0.00
46	chr18b.dat	0.00	39.4	0.00	0.00	scr12.dat	0.00	9.4	0.00	0.00
47	chr20a.dat	0.18	47.3	7.21	0.00	scr15.dat	0.00	17.6	0.00	0.00
48	chr20b.dat	7.40	48.2	3.48	5.13	scr20.dat	0.03	39.8	0.03	0.00
49	chr20c.dat	4.72	48.9	0.00	0.00	sko42.dat	0.23	503.1	0.00	0.25
50	chr22a.dat	0.62	72.9	0.62	0.75	sko49.dat	0.27	629.1	0.27	0.21
51	chr22b.dat	1.19	76.1	2.87	0.00	sko56.dat	0.08	1488.0	0.07	0.02
52	chr25a.dat	10.54	96.8	12.86	0.00	sko64.dat	0.38	1894.1	0.17	0.22
53	els19.dat	0.00	40.3	0.00	0.00	sko72.dat	0.44	2539.0	0.27	0.29
54	esc08a.dat	0.00	3.5	0.00	0.00	sko81.dat	0.23	5482.1	0.40	0.20
55	esc08b.dat	0.00	3.7	0.00	0.00	sko90.dat	0.43	6348.9	0.33	0.27
56	esc08c.dat	0.00	3.8	0.00	0.00	sko100a.dat	0.19	8304.1	0.30	0.21
57	esc08d.dat	0.00	3.6	0.00	0.00	sko100b.dat	0.48	7364.7	0.24	0.14
58	esc08e.dat	0.00	3.6	0.00	0.00	sko100c.dat	0.01	10157.1	0.27	0.20
59	esc08f.dat	0.00	3.7	0.00	0.00	sko100d.dat	0.35	10151.1	0.38	0.17
60	esc16a.dat	0.00	23.7	0.00	0.00	sko100e.dat	0.23	10563.5	0.24	0.24
61	esc16b.dat	0.00	24.1	0.00	0.00	sko100f.dat	0.19	10739.5	0.32	0.29
62	esc16c.dat	0.00	26.7	0.00	0.00	ste36a.dat	1.47	354.8	1.39	0.27
63	esc16d.dat	0.00	26.6	0.00	0.00	tho30.dat	0.31	197.8	0.39	0.00
64	esc16e.dat	0.00	23.4	0.00	0.00	tho40.dat	0.33	479.1	0.20	0.32
65	esc16f.dat	0.00	23.0	0.00	0.00	wil50.dat	0.09	1057.6	0.10	0.07
66	esc16g.dat	0.00	24.9	0.00	0.00	wil100.dat	0.18	10271.9	0.21	0.20

large consistent with modifications in single scheme. We investigated the algorithmic modifications using the following 11 benchmark instances selected from QAPLIB:

Prob. No.	1	2	3	4	5	6	7	8	9	10	11
Prob. Id.	kra30b	lipa50a	lipa60a	lipa70a	lipa80a	sko49	sko90	sko100a	tho30	wil50	wil100

Initial Population. We compare GA's using three methods of generating the initial population: (i) random population; (ii) using the first phase of GRASP; and (iii) using both phases of GRASP. Figure 4 gives the % deviation of the best individual for the 11 benchmark instances. It appears that using Phase 1 of GRASP gives the best overall results. Using both phases gives the worst results, possibly because the initial population is not sufficiently diverse.




Figure 4. Comparison of various ways to generate the initial population.

Selection. In Figure 5, we compare two selection strategies: when all individuals are equally likely to be selected for crossover, and when fitter individuals have greater likelihood for selection (the most fit individual is twice as likely to be selected as the least fit individual). For most of the benchmark instances, unbiased selection is as good or better than biased selection.

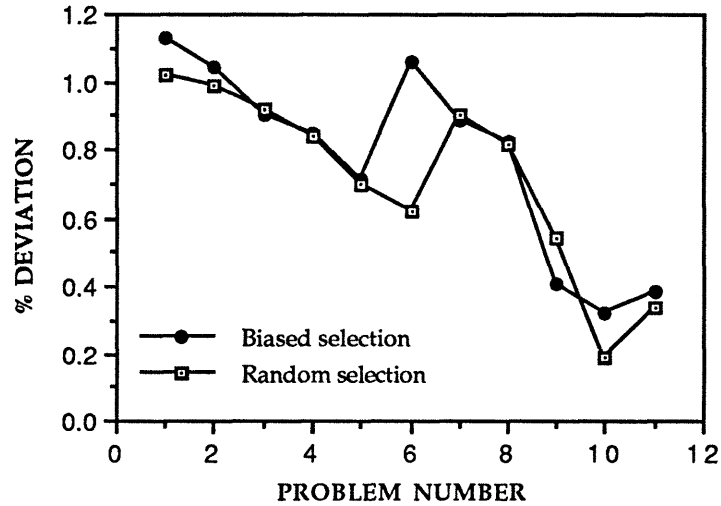




Figure 5. Comparison of biased selection with random selection.

Crossover. In Figure 6(a), we give the % deviation for the best individual obtained by GA using three different crossover schemes: optimized crossover, swap path crossover, and insert path crossover. Though the results are mixed, the insert path crossover gives modestly better results for these benchmark instances. The insert path crossover is found to be more effective than the swap path crossover possibly because the local optimization done periodically also performs swaps, and insert path crossover considers changes that cannot be captured by swaps alone. Surprisingly, the performance of the GA with optimized crossover is worse among all the versions. One possible explanation is the following: in the optimized crossover for any two parents, the optimum child is typically one of the two parents, and so this approach does not add sufficient diversity to the population. Figure 6(b) gives the best objectives function value as a function of the number of trials for the benchmark instance sko100a.dat. After 200 trials, GA with optimized crossover does not improve the objective function value, through the GA with insert crossover occasionally does.



(a)



(b)

Figure 6. Comparison of various crossover schemes.

Immigration. The performance of a GA can depend upon the method used to generate immigrants. Figure 7 compares two schemes of generating the immigrants - the first scheme generates immigrants using the same method as used to generate the initial population and second scheme using a different method that promotes diversity (see Section 2).

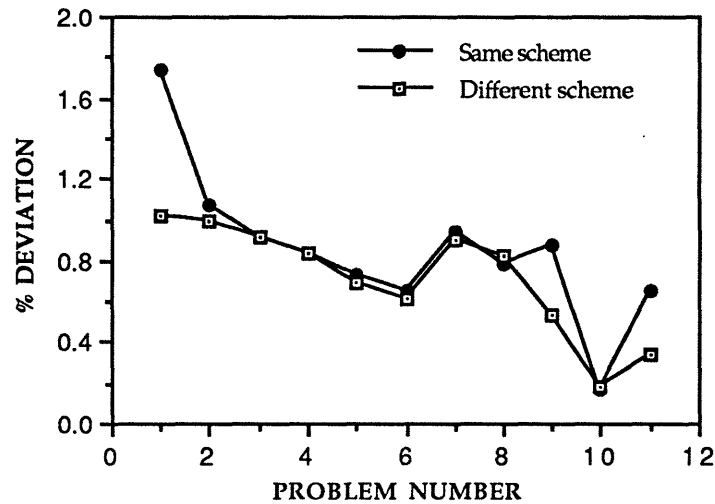

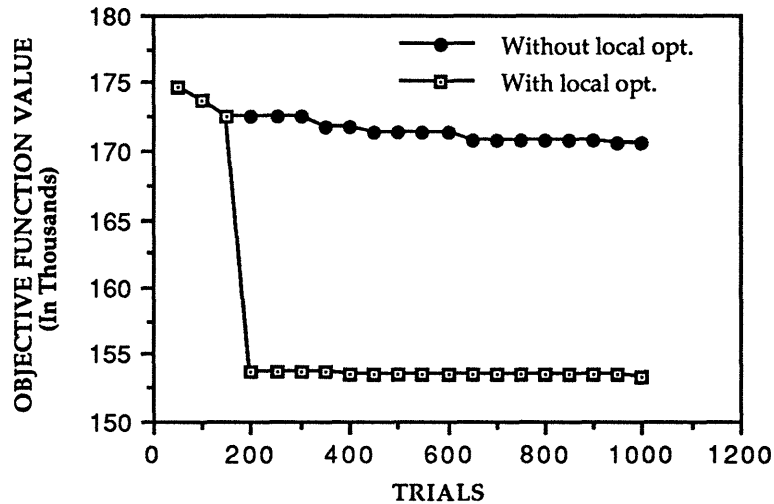


Figure 7. Comparing two ways to generate the immigrants.

Local Optimization. In Figure 8(a), we assess the effect of periodically optimizing a part of the population. Evidently, the GA in which local optimization is performed is much more robust than the GA in which local optimization is not used. To assess the role of local optimization further, we take the benchmark instance sko100 and plot in Figures 8(b) the best objective function values as a function of the number of trials. We see that the local optimization plays an important role in improving the quality of the individuals in the population. In fact, the local optimization contributes more to the quality of the solution than any of the other features that we analysed.




(a)



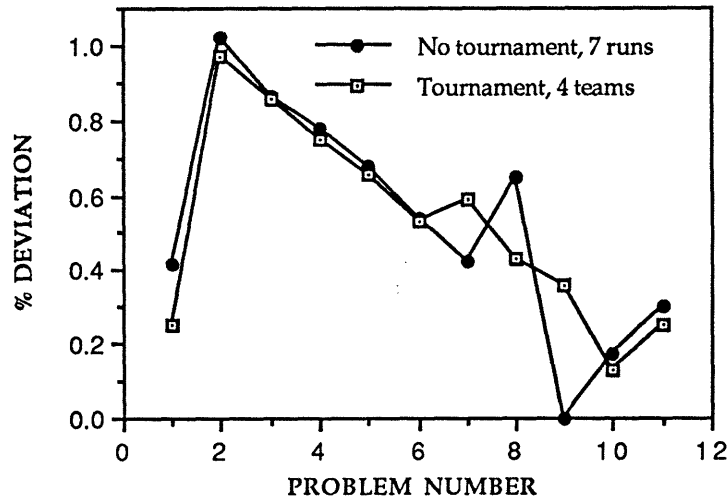
(b)

Figure 8. Effect of local optimization on the performance of GA.

Tournamenting. We now give some graphs to show the effectiveness of tournamenting. In Figure 9(a), we compare the percent deviations of the best objective function value obtained using the GA without tournamenting, and GA with tournamenting (4 or 8 teams). We find that tournamenting improves the performance of the GA substantially over a single run of GA. In the GA with tournamenting with 4 teams, GA is applied 7 times 4 of which generate the 4 teams. As a fairer comparison, we compare the GA with tournamenting with 4 teams with the best individual among 7 runs of the GA without tournamenting in Figure 9(b). Here the GA with tournamenting modestly outperforms the GA without tournamenting. We conclude that tournamenting slightly helps to improve the quality of the individuals. We also investigated how the quality of the best solution improves as tournaments are played. For this, we take a benchmark instance and show in Figure 10 the objective function values of the best individuals at different rounds of tournament. We find that different rounds monotonically improve the objective function value.



(a)



(b)

Figure 9. Effect of tournamenting on the performance of GA.

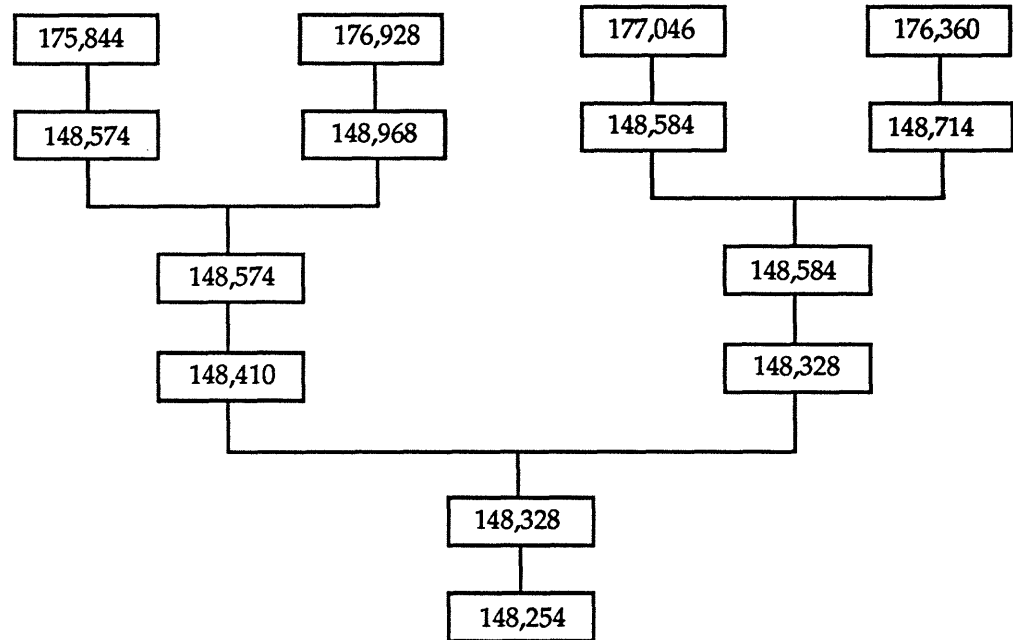


Figure 10. Illustrating tournamenting on a benchmark instance.

Comparison with Random keys Algorithm. Bean [1994] gives a general purpose GA for an optimization problems with a permutation encoding, called the *random keys GA*. This approach applies to QAP as well. We developed a computer program for this algorithm and compared it to the greedy GA on the benchmark instances. Figure 11 presents these results. The greedy GA is substantially superior to the random keys GA in terms of the quality of the solution obtained. We note that the random key GA is a generic approach and was not fine-tuned for the QAP. In contrast, the greedy GA incorporates a number of features that exploit the structure of the QAP.

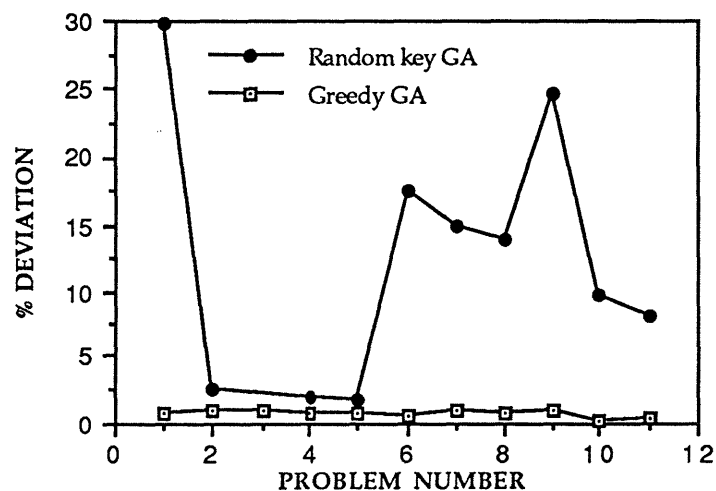


Figure 11. Comparing greedy GA with random key GA.

Comparison with GRASP. Li, Pardalos and Resende [1994] developed GRASP and presented extensive computational results that suggested that GRASP is the most accurate method for solving QAP for medium to large size problems. GRASP is an iterative method and is applied a large number of times with (possibly) different starting solutions. The accuracy of the best solution improves as more iterations are performed. In a similar fashion, the greedy GA can be applied several times and tournamenting with a larger number of teams can be used to improve the performance of the algorithm. To obtain a fair comparison between GRASP and GA, we compared the accuracy of the two algorithms while running GRASP for the same time as that taken by the greedy GA. We apply GA with tournamenting (4 teams) once and note the time taken, say x seconds. We then run GRASP for x seconds and compare the best solution obtained with the GA solution. Figure 12 gives this comparison on the benchmark instances. We could not do the comparison on all the 11 benchmark problem instances because 4 of these instances are nonsymmetric and GRASP can be applied to symmetric instances only. Nevertheless, the comparison using 7 benchmark instances suggests that the greedy genetic algorithm is superior to GRASP. To compare these two algorithms from a different perspective, we take two benchmark instances, will50 and tho40, and plot the best objective function value as a function of time; we give these plots in Figures 13(a) and (b). These plots confirm our previous finding that the greedy GA converges more quickly to the optimum solution than does GRASP. We plotted these curves for several other benchmark instances too, and found them to be similar to the ones given in Figure 13.


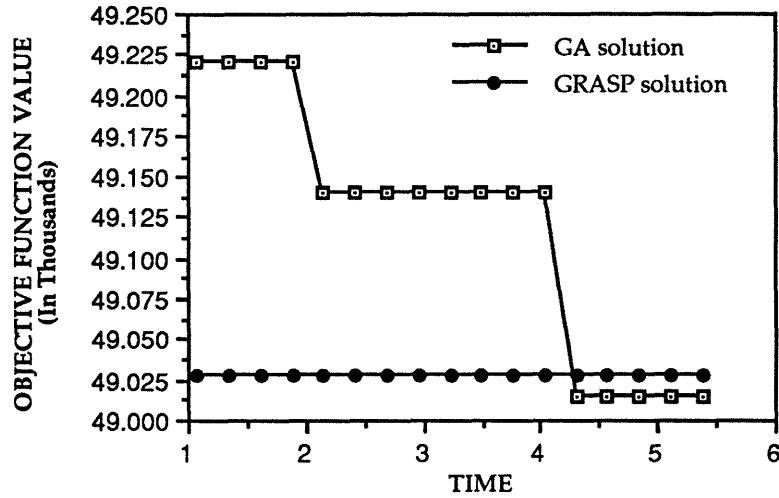
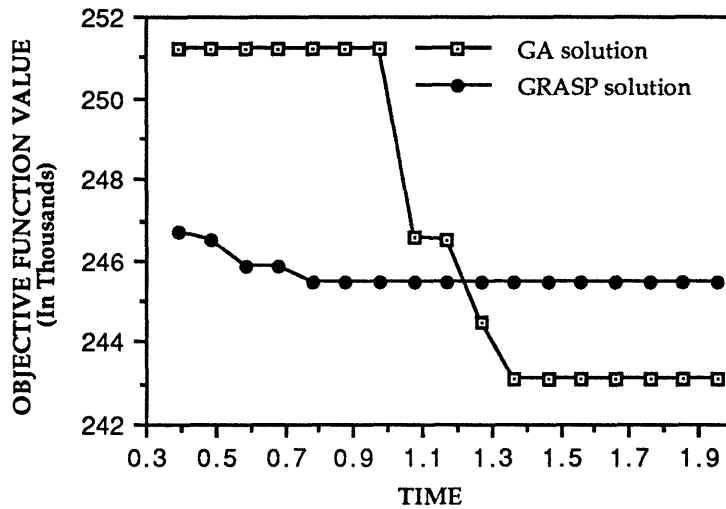


Figure 12. Comparison of greedy GA with GRASP.



(a)



(b)

Figure 13. Solutions of will50 and tho40 by GRASP and GA on the same time scale.

4. SUMMARY AND CONCLUSIONS

In this paper, we developed a new genetic algorithm for QAP and presented computational results of the algorithm on a large set of standard problems in the literature. We call our algorithm the *greedy genetic algorithm* since it incorporates many greedy principles. The principal findings of our paper are the following:

1. The quality of the initial population affects the performance of the genetic algorithm. Randomizing a (greedy) construction heuristic method is one reasonable method to generate the initial population.
2. Path crossover performed better than optimized crossover, possibly because it keeps the population sufficiently diverse.
3. The algorithm had better performance when the method used to generate immigrants was different than the method used to generate the initial population.
4. Periodically optimizing a subset of the population improves the overall performance of the genetic algorithm substantially.
5. Tournamenting was modestly helpful in obtaining high quality individuals.

We find that success of a greedy genetic algorithm depends upon striking a right balance between greediness and maintaining diversity of the population. Greediness improves the quality of the individuals substantially, but overdoing it is detrimental to the overall performance of the genetic algorithm. We obtained the right balance through extensive computational testing. We believe that the greedy principles obtained in this paper might lead to improved genetic algorithms for other combinatorial optimization problems as well. It is plausible that using these ideas one could develop genetic algorithms for several hard combinatorial optimization problems that are better than or comparable to algorithms based on other heuristic optimization techniques.

ACKNOWLEDGMENTS

We thank Bezalel Gavish for suggesting us to investigate genetic algorithms for QAP. We also thank Fred Glover, Bettina Klinz, Abraham Punnen, Skorin-Kapov, Panos Pardalos and Mauricio Resende for their valuable help in this research by providing their feedback/papers/programs. This research was partially supported by ONR Grant N00014-94-1-0099 and a grant from UPS Foundation.

REFERENCES

- AGGARWAL, C.C., J.B. ORLIN, AND R.P. TAI. 1994. Optimized crossover for the independent set problem. Research Report, Operations Research Center, MIT, Cambridge, MA.
- ARMOUR, G.C., AND E.S. BUFFA. 1963. Heuristic algorithm and simulation approach to relative location of facilities. *Management Science* 9, 294-309.
- BAZARAA, M.S., AND M.D. SHERALI. 1980. Benders' partitioning scheme applied to a new formulation of the quadratic assignment problem. *Naval Research Logistics Quarterly* 27, 29-41.

- BEAN, J. C. 1994. Genetic algorithms and random keys for sequencing and optimization. *ORSA Journal on Computing* 6, 154-160.
- BUFFA, E. S., G. C. ARMOUR, AND T.E. VOLLMANN. 1962. Allocating facilities with CRAFT. *Harvard Business Review* 42, 136-158.
- BURKARD, R.E. 1991. Location with spatial interactions: The quadratic assignment problem. *Discrete Location Theory* (Eds. P.B. Mirchandani and R.L. Francis), John Wiley.
- BURKARD, R.E., AND T. BONNIGER. 1983. A heuristic for quadratic boolean programs with applications to quadratic assignment problems. *European Journal of Operations Research* 13, 374-386.
- BURKARD, R.E., S.E. KARISCH, AND F. RENDL. 1994. QAPLIB - A quadratic assignment program library. Technical Report, Institut für Mathematik, Technische Universität, Graz, Austria.
- CHRISTOFIDES, N., AND E. BENAVENT. 1989. An exact algorithm for the quadratic assignment problem. *Operations Research* 37, 760-768,
- DAVIS, L. 1991. *Handbook of Genetic Algorithms*. Van Nostrand, New York.
- FLEURENT, C., AND J. A. FERLAND. 1994. Genetic hybrids for the quadratic assignment problem. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*.
- GLOVER, F. 1994. Genetic algorithms and scatter search: Unsuspected potential. *Statistics and Computing* 4, 131-140.
- GOLDBERG, D.E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Inc.
- HOLLAND, J.H. 1975. *Adaptations in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.
- LAWLER, E. L. 1963. The quadratic assignment problem. *Management Science* 9, 586-599.
- LI, Y., P.M. PARDALOS, AND M.G.C. RESENDE. 1994. A greedy randomized adaptive search procedure for the quadratic assignment problem. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*.
- MALUCELLI, F. 1993. *Quadratic assignment Problems: Solution Methods and Applications*. Unpublished Doctoral Dissertation, Dipartimento di Informatica, Università di Pisa, Italy.
- PARDALOS, P.M., AND J. CROUSE. 1989. A parallel algorithm for the quadratic assignment problem. *Proceedings of the Supercomputing 1989 Conference*, ACM Press, pp. 351-360.
- RESENDE, M.G.C., P.M. PARDALOS, AND Y. LI. 1994. Fortran subroutines for approximate solution of dense quadratic assignment problems using GRASP. To appear in *ACM Transactions on Mathematical Software*.
- SKORIN-KAPOV, J. 1990. Tabu search applied to the quadratic assignment problem. *ORSA Journal on Computing* 2, 33-45.
- TAILLARD, E. 1991. Robust tabu search for the quadratic assignment problem. *Parallel Computing* 17, 443-455.

WEST, D.H. 1983. Algorithm 608 : Approximate solution of the quadratic assignment problem. *ACM Transactions on Mathematical Software* 9, 461-466.

WILHELM, M.R., AND T.L. WARD. 1987. Solving quadratic assignment problems by simulated annealing. *IEEE Transactions* 19, 107-119.