Available online at www.sciencedirect.com
SCIENCE DIRECT®

Int. J. Production Economics 103 (2006) 742–754

international journal of production economics

www.elsevier.com/locate/ijpe

An empirical comparison of Tabu Search, Simulated Annealing, and Genetic Algorithms for facilities location problems

Marvin A. Arostegui Jr.[a], Sukran N. Kadipasaoglu[b], Basheer M. Khumawala[b,*]

[a]United States Air Force, USA
[b]Bauer College of Business, Decision and Information Sciences Department, University of Houston, 270D Melcher Hall, Houston TX 77204 6027, USA

Received 1 February 2004; accepted 1 August 2005
Available online 15 June 2006

Abstract

Algorithms to solve Facility Location Problems (FLP) optimally suffer from combinatorial explosion and resources required to solve such problems repeatedly as required in practical applications become prohibitive. In these cases heuristic methods are the only viable alternative. We compare the relative performance of Tabu Search (TS), Simulated Annealing (SA) and Genetic Algorithms (GA) on various types of FLP under time-limited, solution-limited, and unrestricted conditions. The results indicate that TS shows very good performance in most cases. The performance of SA and GA are more partial to problem type and the criterion used. Thus, in general we may conclude that TS should be tried first to the extent that it always yields as good or better results and is easy to develop and implement.

Keywords: Facility location; Tabu Search; Genetic Algorithms; Simulated Annealing; Combinatorial optimization

1. Introduction

Decision makers have a natural desire to find optimal solutions to the various operations management problems they face every day. Unfortunately, there is a class of common problems that are extremely difficult to solve optimally. Such problems include assembly line balancing, lot sizing, project scheduling, job and flow-shop scheduling, facilities location and layout. Although algorithms to solve such problems optimally exist, they suffer from combinatorial explosion. These problems belong to the class of problems known as NP-hard, and in most cases, the time and computing resources required to solve such problems repeatedly in practical applications become prohibitive. In these cases heuristic methods are the only viable alternative. Van Laarhoven and Aarts (1987) classify heuristics into "tailored" and "general". While tailored heuristics have a limited applicability to a specific problem, general algorithms define a strategy for obtaining approximate solutions and thus are widely applicable to various forms of combinatorial optimization problems.

The most well known of the general heuristic methods are Tabu Search (TS), Simulated Annealing (SA), and Genetic Algorithms (GA). The popularity of these heuristics has flourished in

*Corresponding author. Tel.: +1 713 743 4721; fax: +1 713 743 4940.
E-mail address: bkhumawala@uh.edu (B.M. Khumawala).

0925-5273/$ - see front matter
doi:10.1016/j.ijpe.2005.08.010

recent years and several published studies can be found in the literature where they outperform the tailored counterparts. However, only a few studies provided comparisons of these three heuristics in depth. In this paper, we compare the relative performance of TS, SA and GA on various Facilities Location Problems (FLP). The choice of FLP is made due to its strategic importance in the design of the supply chain network. Our motivation is to contribute further to the understanding of which of these three heuristics may be more effective under different circumstances.

We assume the reader is familiar with (1) The FLP, its importance and its pertinent literature, (2) The three heuristics, TS, SA and GA and therefore for brevity we abstain from their basic descriptions. For the interested reader, there are several articles and texts that provide a good introduction to TS (Glover, 1989, 1990a, b, 1993), SA (Kirkpatrick et al., 1983; Cerny, 1985), and GA (Goldenberg, 1989).

In the remainder of the paper, we briefly provide the references to the pertinent FLP literature. Then we discuss the results from the few previous studies that have compared the three general heuristics. This is done to establish the need for our research. In Section 3, we present the formulation of three versions of the FLP selected for this research and the details of our implementation of TS, SA and GA. The details of the empirical comparison are described in Section 4 and the results and conclusions are discussed in Section 5.

2. Background

Facilities location decisions have attracted a great deal of attention from researchers in the last couple of decades (Harkness and ReVelle, 2003; Drezner et al., 2002; Canel and Das, 2002; Nozick, 2001; Canel et al., 1996, 2001; Melkote and Daskin, 2001; Giddings et al., 2001; Canel and Khumawala, 1996, 2001; Hinojosa et al., 2000; Tragantalerngsak et al., 2000; Avella et al., 1998; Owne and Daskin, 1998; Volgenant, 1996). Consequently, there is now a variety of approaches for solving these problems. These include 0-1 MIP (Haug 1985; Cohen and Lee 1989), dynamic programming (Pomper, 1976), breakeven analysis (Jucker and Carlson, 1976), quadratic programming (Hodder and Jucker, 1985, Hodder and Dincer, 1986), and fuzzy set theory (Naik and Chakravarthy, 1993). Variations of the FLP are solved by Hakimi (1983), Drezner (1994) and by Drezner et al. (1998, 2002) using various "tailored" heuristic methods. Independent of the methodology used, it is a well-known fact that facilities location decisions significantly influence profitability and productivity. With the current trend in globalization and the associated complexities in supply chain practices, the FLP is now taking on even greater importance (Porter, 1998; Shaver, 1998; Bartmess et al. 1994; Bartmess and Cerney, 1993).

The earliest comparison of TS, SA, and GA is due to Sinclair (1993) who used the quadratic assignment problem as a test domain. Although he also included the Great Deluge and the Record-to-Record Travel techniques, we will limit our discussion to the three heuristics of interest. Sinclair developed the application of these heuristics and their comparisons on the hydraulic turbine runner-balancing problem. Sinclair (1993) developed basic implementations of each heuristic and selected the parameters that resulted in the best performance in terms of solution quality and computation time. Sinclair's study (1993) found that GA performed significantly worse than the other heuristics even while requiring more computation time. Between SA and TS, TS provided better solutions in 28 out of 37 cases. However, the computation times for both of these heuristics were comparable.

Lee and Kim (1996) applied the TS, SA and GA heuristics to the resource-constrained project-scheduling problem. They found that, in general, SA and TS had nearly the same performance, with GA lagging slightly behind, especially as the number of activities is increased. However, they found that GA had the least variability overall.

The results of Sinclair (1993) and Lee and Kim (1996) suggest that GA might not perform as well as TS or SA. Kincaid (1992) compared TS and SA for a location problem of noxious facilities and concluded that TS outperformed SA. However, other papers have found SA to outperform TS for the lot sizing problem (Kuik, et al., 1993) and the flow shop problem (Marett and Wright, 1996). Tian et al. (1996) conducted a study applying SA to the Quadratic Assignment Problem (QAP), more specifically to the facility location problem. Their results suggest that SA is an effective approach. Drezner et al. (2002) compared the effectiveness of five heuristics in solving the p-median problem (a variant of FLP) and reported that SA (in discrete space) combined with an ascent algorithm (in continuous space) gave the most satisfactory results.

More recently, Syarif et al. (2002) and Zhou et al. (2002) have used GA in solving problems involving facilities location and allocation of customers to facilities and reported satisfactory results. Chaudry et al. (2003) applied GA to the relatively large-sized, constrained p-median problem. They reported that GA was able to obtain optimum results for most of the problems, however, in a few cases the computational burden was in excess of an hour. A study by Ghosh (2003), reported that TS when compared to Complete Local Search (CLM) algorithm returned better quality solutions for the un-capacitated facility location problem—although both heuristics gave results very close to optimum. Wang et al. (2003) applied TS to the budget-constrained location problem where opening and closing of facilities were considered simultaneously. They compared TS to greedy interchange and Langrangian relaxation approximation, and reported that solutions generated by TS were the most satisfactory in terms of solution quality. The authors also confirmed the suitability and desirability for utilizing such heuristics as they perform very well for large problems. Corinthal and Captivo (2003), similarly, reported very satisfactory results with TS. Finally, Chamberland (2004) proposed a heuristic based on the TS principle to solve the network subsystem expansion problem and found that TS provides good solutions.

This evidence suggests that for a given problem domain, one of the heuristics may be preferable to others, thus justifying the need for further experimentation in other problem domains such as the FLP. This is in line with the recommendation in Drezner et al. (2002) as the authors proposed the application of meta heuristics such as TS, GA and SA as future research.

3. Problem domains

We use three variations of the FLP to compare the performance of TS, SA, and GA. These three variations are (1) the capacitated FLP (CFLP), (2) the multi-period FLP (MPFLP), and (3) the multi-commodity FLP (MCFLP).

It is important to note that representation selection (matrix vs. vector) as well as the parameter values for TS, SA and GA—described below in detail—are determined based on the results of extensive pilot experiments and testing for the three variations of the FLP. In these experiments, the effect of various parameter settings on solution quality and computation time is assessed for each of the heuristics and parameter values are set accordingly. The details of pilot experiments and the statistical tests are given in Arostegui (1997).The representation selection and parameter values used for the three FLP are shown in Table 1. In this section, we describe the solution procedure of each FLP with the three heuristics.

3.1. *Capacitated facilities location problem (CFLP)*

For CFLP, the problem is to determine which of N capacity-constrained facilities should be used to satisfy the demand for M customers at the lowest sum of fixed and variable costs. The problem is formulated as in Khumawala (1974).

Table 1
Representation selection and parameter values

	CFLP matrix representation	MPFLP matrix representation	MPCFLP vector representation
Tabu search			
LTM	10	5	5
STM	3000	3000	50
TLSIZE	5	7	2
Simulated annealing			
Rate	.90	.90	.90
Accept	2500	2500	50
Factor	2	3	3
Genetic algorithm			
POP	40	30	40
GENS	2000	2000	150
XOVER	.85	.80	.95
MUTATE	.05	.05	.10

For all three heuristics, this problem is represented in the form of an $N \times M$ matrix, where each cell in the matrix represents the demand of a customer j ($j = 1 \ldots M$) served from a facility i ($i = 1 \ldots N$). For TS and SA, a neighbor solution is defined as a solution where a single allocation of customer demand is moved from its original supplier to another supplier with some positive capacity. The size of this neighborhood, typically a very large number, varies depending on the number of positive allocations in the incumbent solution.

For TS only a subset of the total neighborhood is evaluated at each iteration. The subset consists of randomly selecting an existing allocation of demand for each of M customers, and reallocating as much as possible to a different, randomly selected supplier with some positive remaining capacity. Since only one reallocation is made per customer, the resulting neighborhood is of size m. Two parameters are defined: Long-term Memory iterations (LTM) and Short-term Memory iterations (STM). Each LTM consists of a restart with a randomly generated solution and an empty tabu list. For each LTM, a number of the normal TS algorithms are executed. For tabu list size (TLSIZE) we use a random tabu list size as in Taillard (1991). TLSIZE represents the maximum number allowed for the TLSIZE at any time. The values of these parameters are given in Table 1. The TS procedure is terminated when a certain number of iterations have been completed. Further details on termination procedures are given in Section 4.

For SA, the neighborhood structures for SA are the same as those for TS as described above. However, instead of evaluating all or part of an incumbent solution's neighborhood, we only evaluate a randomly selected neighbor. We randomly select a customer. For this customer we randomly select an existing positive allocation of demand to some supplier. We then randomly select a different supplier with some positive capacity and reallocate as much of the demand as possible from the old supplier to the new supplier.

For SA, the cooling schedule consists of an initial temperature, a probability function, and a temperature adjustment function. The initial temperature is selected so that the resulting probability of accepting non-improving solutions is 95%. To find this initial temperature, we take the difference in cost function between the incumbent solution and the first 100 non-improving solutions and get the average cost. Using this average we compute the required temperature parameter to achieve the 95% acceptance rate. Finally, the temperature adjustment function is a simple decay function of the "Rate" parameter. The function is applied at the end of each epoch.

The length of an epoch is controlled by two parameters: Accept and Factor. An epoch consists of evaluating enough solutions such that at least some number of solutions are accepted (improving or non-improving). This number of solutions is controlled by the "Accept" parameter. Since the rate at which the solutions accepted decreases as the temperature parameter is lowered, the total number of solutions evaluated should be limited to prevent lengthy epochs at lower temperatures. The maximum number of solutions that may be evaluated per epoch is a factor of the "Accept" parameter controlled by the "Factor" parameter. Thus, en epoch will end when "Accept" solutions have been accepted or when "Factor $\times$ Accept" solutions have been evaluated, whichever comes first. Values for the "Rate", "Accept" and "Factor" parameters are given in Table 1. The algorithm will terminate whenever the system is deemed frozen. In our implementation, frozen is defined as an acceptance rate of less than 1% during any one epoch. Further details on termination procedures are given in Section 4.

In the GA implementation, the fitness function assigns a "goodness" value to each chromosome or solution at each generation. In our case, the goodness value is represented by the objective function value, which is the total cost of the solution. Using this value, we assign to each chromosome a probability of selection as a parent for next generation. Our probability assignment, taken from Murata et al. (1996). The probability of selection computed as in Murata et al. (1996) is actually used in selecting only the first parent for a crossover operation. The second parent is selected at random with all chromosomes having an equal probability of selection. This selection procedure is described in Holland (1975).

The crossover operation is applied probabilistically to a set of two parents from which two offsprings are generated for the new generation. With probability XOVER, the crossover operator is applied to the parents to generate the first offspring. With complementary probability, the offspring is a copy of the first parent. With probability XOVER, the crossover operator is applied to the parents (in reverse order) to generate the second offspring.

With complementary probability, the offspring is a copy of the second parent. In the implementation of GA a single crossover point operation is used. This crossover point is randomly selected as one of the M customers. The first offspring will receive the columns to the left of the crossover point from the first parent matrix and the remaining columns from the second parent. The operation is the same for the second offspring with the roles of the parents reversed. Unfortunately, there is no guarantee that the resulting offspring solutions derived in this manner will be feasible, thus requiring a repair operation. The repair operation is applied to each row in the matrix (i.e., each potential facility), whenever the assigned allocations exceed the corresponding capacity. The repair operation consists of randomly selecting the allocations and moving them to other randomly selected facilities with available capacity. This is done until the total demand supplied from a facility is less than its capacity, and until the solution is feasible for each facility.

The mutation operator is applied probabilistically to each gene (column) in each chromosome with MUTATE probability. The mutation consists of selecting a random allocation for the customer, and reallocating from the current supplier with some positive capacity. The smaller of the existing allocation or the available capacity in the new supplier is reallocated. The results are always feasible.

The GA is terminated whenever a fixed number of generations (GENS) have elapsed, or if the algorithm converges to a single solution (i.e., all of the chromosomes in a generation represent the same solution). Convergence only tends to happen when the mutation rate MUTATE is small (e.g. 1%). Further details on termination procedures are given in Section 4. Values for the "Rate", "Accept" and "Factor" parameters are given in Table 1.

3.2. *Multiple-period facilities location problem (MPFLP)*

The MPFLP consists of determining which of the N capacity-constraint facilities to be used over a planning horizon of V periods, given the fixed and variable costs in the CFLP; while simultaneously considering the costs of opening and closing facilities in different time periods. This problem is formulated as in Hormozi and Khumawala (1996).

For all three heuristics, the MPFLP problem is again represented by a matrix with rows representing the facilities at N locations and $M \times V$ columns representing the customers in the V time periods. Thus, the matrix representation has a direct connection to the x_{tij}'s in the problem formulation. The first M columns represent the customers in the first period; the second M columns represent the customers in the second period, and so on. The matrix is filled by assigning values to each element such that the sum of each column matches the corresponding customer's demand, and the sum of each row (by groups of M columns) does not exceed the corresponding facility's capacity. A matrix thus filled represents a solution. The y_{ti}'s can be easily determined by examining the values in the matrix to determine which facilities have been assigned some demand in any period.

For TS and SA, a neighbor solution is defined as a change to two of the matrix elements within a column while maintaining the demand/capacity constraints. This change amounts to a reallocation of demand from one facility to another for a single customer in a single period. This is a very large and variable neighborhood since there may be as many reallocations possible as there are positive elements in the matrix, and each reallocation can be by as little as a quantity of one or as much as the value of the source element. Accordingly, for TS we only evaluate a subset of this neighborhood at each iteration. The subset consists of making one reallocation for each column in the matrix for a neighborhood size of nv. The two elements involved in the reallocation are randomly selected, and the quantity of demand reallocated is the minimum of the source element's value or the destination element's remaining capacity. This procedure always results in feasible solutions. For SA, only one neighboring solution is considered at a time.

For the GA, we again implemented a single-point crossover. We select a crossover point $x \in \{1, \ldots, V\}$, and the first offspring consists of the matrix columns from the first parent representing customer demand allocations for periods less than x and the columns from the second parent representing customer demand allocations for periods x and above. For the second offspring, the roles of the parents are reversed. The resulting offsprings are always feasible. The crossover operation is applied probabilistically in the same manner as was done for the CFLP.

3.3. Multiple-commodities facilities location problem (MCFLP)

For MCFLP, M customers have demand for one or more W commodities. The objective of MCFLP is to determine which of N capacity-constrained potential facilities should supply which commodities to the M customers, such that each facility supplies only one kind of commodity while minimizing the overall sum of fixed and variable costs. We used a capacitated version of Warszawski's (1973) formulation. See Arostegui (1997) for details.

MPFLP is best represented by vector representation due to the special constraint that each location may only supply a single commodity. This particular constraint makes it defining neighbor moves impractical for a matrix representation. As vector representation requires solving transportation problems, the solution space is smaller than for a matrix representation. As we did for CFLP and MPFLP, we conducted pilot experiments to get a clear picture of how parameter settings affect solution quality and computation time. The details of these pilot experiments are provided in detail in Arostegui (1997). The solution representation and parameter settings are given in Table 1.

A particular solution is represented by a vector $s = \{y_1, y_2, \ldots, y_N\}$, where $y_i \in \{0, 1, \ldots, W\}$ indicates what commodity, if any, is supplied by facility at location i (a zero value indicates the facility in location i is not in operation). Using this vector solution, the problem is partitioned into W single-commodity facility location problems. The x_{rij}'s are determined by solving the corresponding transportation problems for each single-commodity problem.

For TS and SA, a neighbor solution is defined as a solution where the commodity assignment of one of the vector elements is changed to some other commodity or to a closed status (no commodity assigned). The complete neighborhood for an incumbent solution then is the set of all solutions resulting from changing the commodity assignment of each potential facility from its current assignment to one of the other commodities in turn or to no commodity assignment. For a problem with w commodities and n potential facilities, the neighborhood size is wn. Although this is not a very large neighborhood, having to solve transportation problems for each neighbor is a slow process, especially as the problems grow in size. Therefore, for TS only a portion of this neighborhood is evaluated. At each iteration, we only evaluate the n neighbor solutions resulting from changing the current assignment of each potential facility in turn to some other commodity (or no commodity). The new commodity assignment is selected at random. In some cases, this operation will result in neighbor solutions that are not feasible due to insufficient capacity assigned to a particular commodity to meet all demand. In such cases, the resulting neighbor solution is ignored (i.e., no attempt is made to evaluate n feasible solutions, just n neighbor solutions, feasible or not).

For SA, given an incumbent solution, we randomly select one of the wn neighboring solutions to evaluate next. The neighbor solution is selected by randomly choosing one of the potential facilities and changing its current commodity assignment to some other, randomly selected, commodity (or no commodity). If the resulting neighbor solution is not feasible, another one is tried. Although neighbor solutions that are not feasible are ignored, they do count toward the total number of solutions considered. This is a necessary condition to avoid lengthy searches for feasible solutions in highly constrained problems, especially toward the end of the SA algorithm, when the probability of accepting solutions becomes rather small.

4. Empirical comparison

In this study, specifically, we addressed the following research questions: Is there a dominant heuristic for all three types of facilities location problems? If not, is there a dominant heuristic for each type of facilities location problems? For each of the heuristics, what is the relationship between time allowed for computation and quality of solutions?

To address the research questions above, the performance of the three heuristics was evaluated along three dimensions: time-limited, solutions-limited, and unrestricted. We used these three dimensions for the following reasons: (1) Computational time limitation is a common criterion for evaluating heuristics as has been employed by many researchers. (2) We included the solution-limited dimension because it is possible that the rate of improvement in solution quality, after a while, becomes negligible and therefore may not be worth the additional computational resources. (3) We included the unrestricted dimension to allow the heuristics to fully run to their completion as described in Section 3.

For the time-limited evaluation, all of the three heuristics were allowed a fixed amount of computation time (300 s) at which point the best solution they had achieved was noted. In the solutions-limited evaluation, of the three heuristics, the genetic algorithm evaluates the fewest solutions for the given parameters. Therefore, the TS and SA were limited to 76,040, the number of candidate solutions the GA evaluated for the CFLP and similarly to 56,030 for MPFLP. For MCFLP, the use of a vector representation means the GA is no longer the heuristic with the fewest solutions explored. Accordingly, to determine this limit we looked at the number of solutions evaluated by each heuristic during the unrestricted runs. Then, for each group of four problems representing a specific size, we picked a lower bound. For example, for the first four problems, TS evaluated the fewest solutions. Therefore, we set the solution limits for all the heuristics to this lower bound. In this manner, we established limits for the four problem size groups as 1522, 2016, 3536, or 5740 candidate solutions. This procedure allows comparisons based on the ability of the heuristics to make the most out of the same amount of information.

For the unrestricted dimension, the heuristics were allowed to complete their run according to their parameter settings. The above parameters for each heuristic were established following several pilot studies in which several experiments were conducted with different settings. For details see Arostegui (1997).

For each variation of FLP, a set of 20 test problems were generated using our database of 250 US cities and 1990 census population, for a total of 60 test problems. Each problem is considered ‘‘large’’ in terms of the computational requirements; in general, none of the problems could be solved to optimality within any reasonable amount of time. For each type of FLP, The 20 problems represent various combinations of the problem size, capacity and fixed costs consistent with those used by several researchers, including the original Kuehn and Hamburger data. The problems for each type of FLP are shown in Tables 1–4. Following the method proposed by Dudewicz and Dalal (1975) as fully described in Law and Kelton (1991), we used a $P^* = .90$ (probability of being correct) and $n^* = 20$ (number of replications generated with different random number streams). Subsequently, we used Friedman's F_r test to assess the difference in performance among the three heuristics across all types of FLP. To determine which heuristics are in fact different from each other, we used Wicoxon's signed rank test for paired differences. For brevity, we do not provide the statistical tables

Table 2
CFLP test problems

Problem	Size[a]	Capacity	Fixed cost
CAP301	25 × 200	50,000	7500
CAP302	25 × 200	50,000	17,500
CAP303	25 × 200	100,000	7500
CAP304	25 × 200	100,000	17,500
CAP305	25 × 250	50,000	7500
CAP306	25 × 250	50,000	17,500
CAP307	25 × 250	100,000	7500
CAP308	25 × 250	100,000	17,500
CAP309	50 × 200	25,000	7500
CAP310	50 × 200	25,000	17,500
CAP311	50 × 200	50,000	7500
CAP312	50 × 200	50,000	17,500
CAP313	50 × 250	25,000	7500
CAP314	50 × 250	25,000	17,500
CAP315	50 × 250	50,000	7500
CAP316	50 × 250	50,000	17,500
CAP317	100 × 250	12,500	7500
CAP318	100 × 250	12,500	17,500
CAP319	100 × 250	25,000	7500
CAP320	100 × 250	25,000	17,500

[a]Size = facilities × customers.

Table 3
MP-FLP test problems

Problem	Size[a]	Capacity	Demand	Open	Close
MP301	20 × 15 × 38	2500	Increasing	4937	3702
MP302	20 × 15 × 38	2500	Increasing	2370	1777
MP303	20 × 15 × 38	5000	Increasing	4937	3702
MP304	20 × 15 × 38	5000	Increasing	2370	1777
MP305	25 × 10 × 25	2500	Decreasing	3113	2335
MP306	25 × 10 × 25	2500	Decreasing	1494	1121
MP307	25 × 10 × 25	5000	Decreasing	3113	2335
MP308	25 × 10 × 25	5000	Decreasing	1494	1121
MP309	25 × 15 × 38	5000	Concave	7646	5734
MP310	25 × 15 × 38	5000	Concave	3670	2752
MP311	25 × 15 × 38	10000	Concave	7646	5734
MP312	25 × 15 × 38	10000	Concave	3670	2752
MP313	30 × 10 × 25	2500	Convex	1660	1245
MP314	30 × 10 × 25	2500	Convex	797	598
MP315	30 × 10 × 25	5000	Convex	1660	1245
MP316	30 × 10 × 25	5000	Convex	797	598
MP317	30 × 15 × 38	2500	Seasonal	3777	2832
MP318	30 × 15 × 38	2500	Seasonal	1813	1360
MP319	30 × 15 × 38	5000	Seasonal	3777	2832
MP320	30 × 15 × 38	5000	Seasonal	1813	1360

[a]Size = periods × facilities × customers.

Table 4
MC-FLP test problems

Problem	Size[a]	Capacity	Fixed Costs
MC201	$5 \times 10 \times 50$	120,000	90,000
MC202	$5 \times 10 \times 50$	120,000	180,000
MC203	$5 \times 10 \times 50$	240,000	90,000
MC204	$5 \times 10 \times 50$	240,000	180,000
MC205	$5 \times 30 \times 100$	80,000	90,000
MC206	$5 \times 30 \times 100$	80,000	180,000
MC207	$5 \times 30 \times 100$	160,000	90,000
MC208	$5 \times 30 \times 100$	160,000	180,000
MC209	$7 \times 15 \times 50$	110,000	90,000
MC210	$7 \times 15 \times 50$	110,000	180,000
MC211	$7 \times 15 \times 50$	220,000	90,000
MC212	$7 \times 15 \times 50$	220,000	180,000
MC213	$7 \times 20 \times 100$	160,000	90,000
MC214	$7 \times 20 \times 100$	160,000	180,000
MC215	$7 \times 20 \times 100$	320,000	90,000
MC216	$7 \times 20 \times 100$	320,000	180,000
MC217	$10 \times 30 \times 100$	150,000	90,000
MC218	$10 \times 30 \times 100$	150,000	180,000
MC219	$10 \times 30 \times 100$	300,000	90,000
MC220	$10 \times 30 \times 100$	300,000	180,000

[a]Size = commodities × facilities × customers.

here. For detailed statistical results the interested reader can refer to Arostegui (1997).

5. Results and conclusions

5.1. Time-limited results

For the time-limited evaluation, all of the three heuristics were allowed a maximum time of 300 s and the best solutions from each heuristic were noted. This approach evaluates the efficiency with which the three heuristics reach quality solutions over time. Figs. 1–3 illustrate the performance of the heuristics for CFLP, MPFLP, and MCFLP, respectively, for a sample problem. (Note that the figures show performance over time well past the 300 s limit at which the heuristics were compared.) Fig. 1 clearly shows that for CFLP, TS is the first to reach the lowest-cost solution. SA does not have as steep a descent, but eventually reaches the same quality solutions as TS. The GA has an initially steep descent, but slows down well before reaching solutions as good as the other two heuristics and would take significantly more computation time to reach the same solutions. For CFLP, TS gives best results in terms of rapidly reaching low-cost solutions, followed by SA and GA, respectively.

The performance of TS is again superior to the other heuristics for MPFLP as illustrated by Fig. 2. Statistically, TS shows the best performance in terms of rapidly reaching low-cost solutions, followed by SA and GA, respectively.

For MCFLP, however, as illustrated by Fig. 3, the results are significantly different. GA shows a much faster descent—reduction in cost—than TS, before TS catches up. Statistically, however, the performance of GA and TS are similar. The SA algorithm does not perform well at all for this class of problems.

Overall, given the same amount time, for the parameters used in this experiment, TS gave, in general, the best results. The improvement in the performance of GA for the MCFLP may be attributed to the vector representation used for this type of FLP. The vector representation requires parameters that result in fewer solutions to be evaluated and GA appears to perform well in this environment, when information is limited.

5.2. Solutions-limited results

Figs. 4–6 show the performance of the heuristics over the number of candidate solutions evaluated on the three variations of FLP for a sample problem. (Note that the figures show performance of the heuristics over candidate solutions evaluated well past the limit at which the heuristics were compared.) Statistically, SA performs similar or better, i.e. reaches a similar or better solution, than TS in many cases. However, for all three FLPs, it is evident that the GA is quicker than TS or SA in initially finding low-cost solutions. However, it begins to level off early, and is surpassed by TS and SA.

Statistically, SA is the best performer for the CFLP, followed by TS and GA. For the MPFLP, GA is the clear winner. There is not a significant difference in performance between TS and GA. Similarly, for MCFLP, GA is very efficient. In this case there is no statistically significant difference between the performance of GA and that of TS. SA does not perform well at all for the MCFLP.

5.3. Unrestricted results

When all the heuristics were allowed to finish their run according to their parameters, the statistical tests showed significant differences among

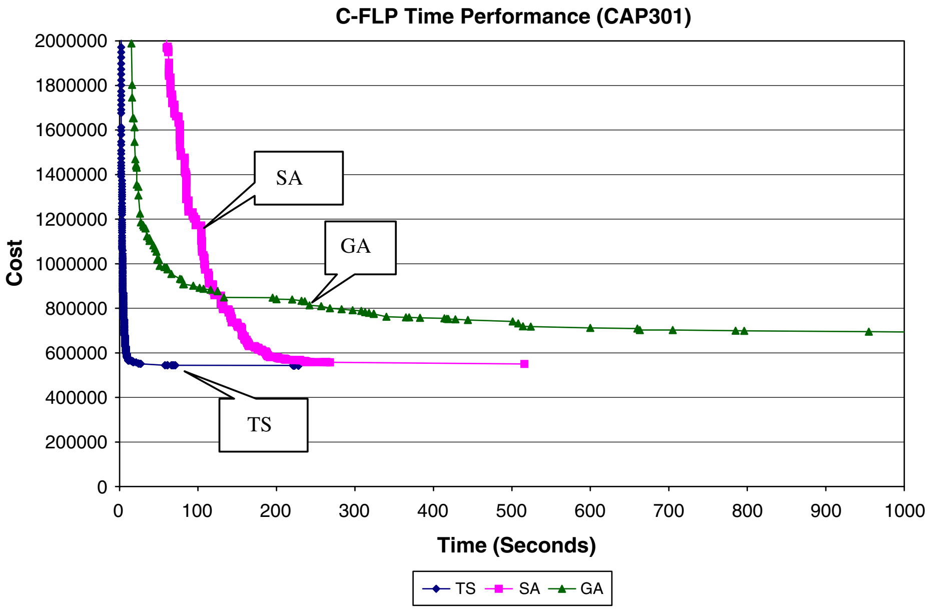

Fig. 1. CFLP—time-limited performance.

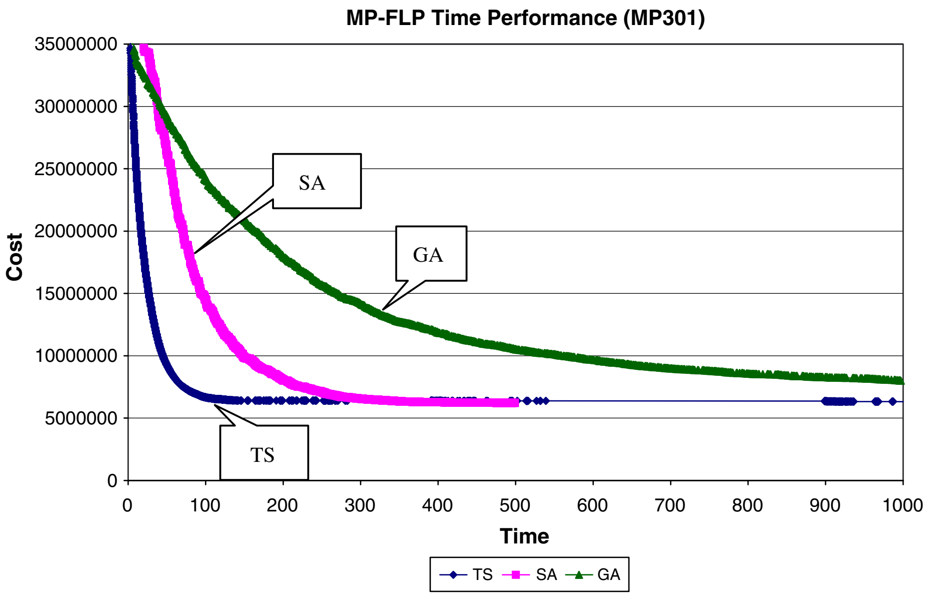

Fig. 2. MPFLP—time-limited performance.

their performance. TS performed best for CFLP and MCFLP while SA performed best for MPFLP. GA performed just as well as TS for the MCFLP. The unrestricted results are not illustrated due to space considerations; however, details are available from Arostegui (1997).

6. Conclusions

In general, we can conclude that the performance of TS, SA and GA for the different types of FLP is situational principally based on the measure of comparability. For the CFLP, TS provides the best

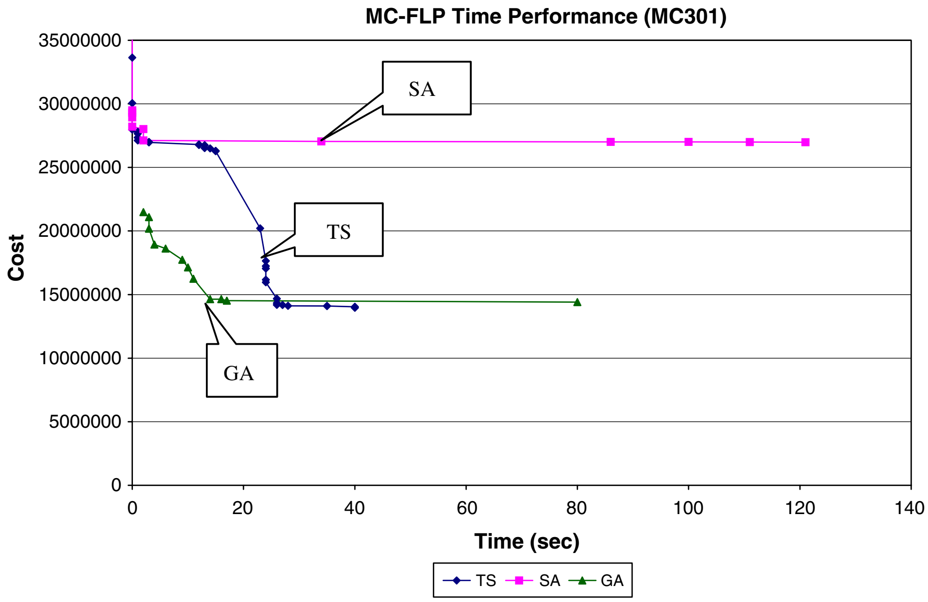

Fig. 3. MCFLP—time limited performance.

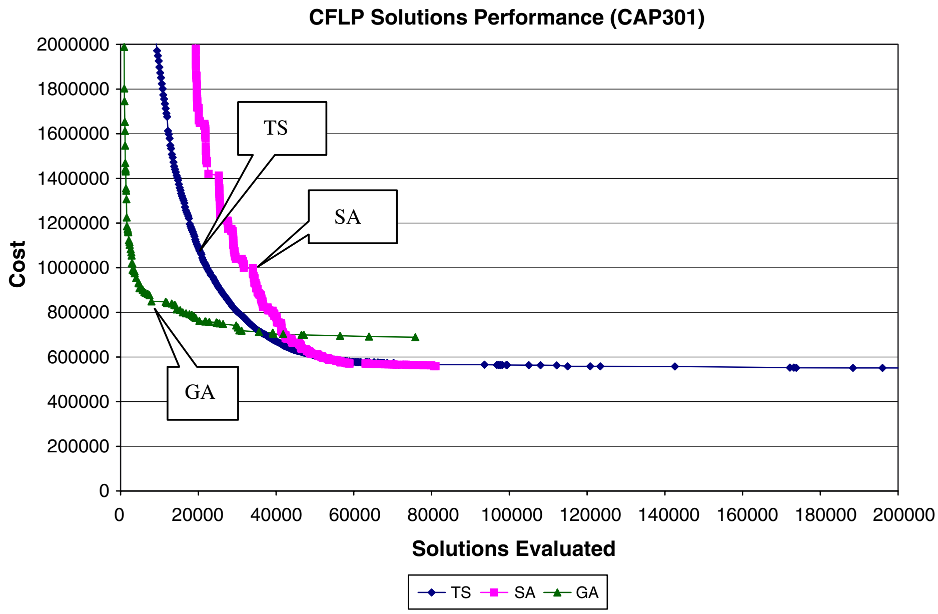

Fig. 4. CFLP—solution-limited performance.

performance under time-limited evaluation, while SA provides the best performance under solution limited evaluation. TS shows superior performance when no restrictions were imposed for the CFLP. For the MPFLP, under time-limited evaluation TS and SA performed similarly (superior to GA), however, under solution-limited evaluation GA was superior. For the unrestricted evaluation case SA performed best with TS a close second for the MPFLP. For the MCFLP, under time-limited evaluation TS and GA give best results. Under solution-limited and unrestricted evaluation TS was the best performer.

Overall, it appears that TS shows very good performance for all types of FLP under the time-limited evaluation. Also, and the longer the solution

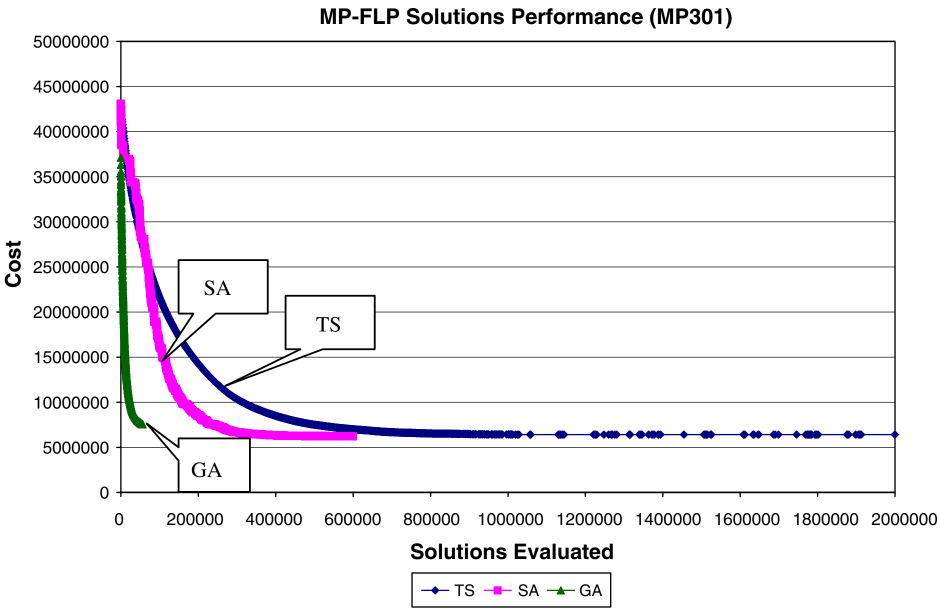

Fig. 5. MPFLP—solution-limited performance.

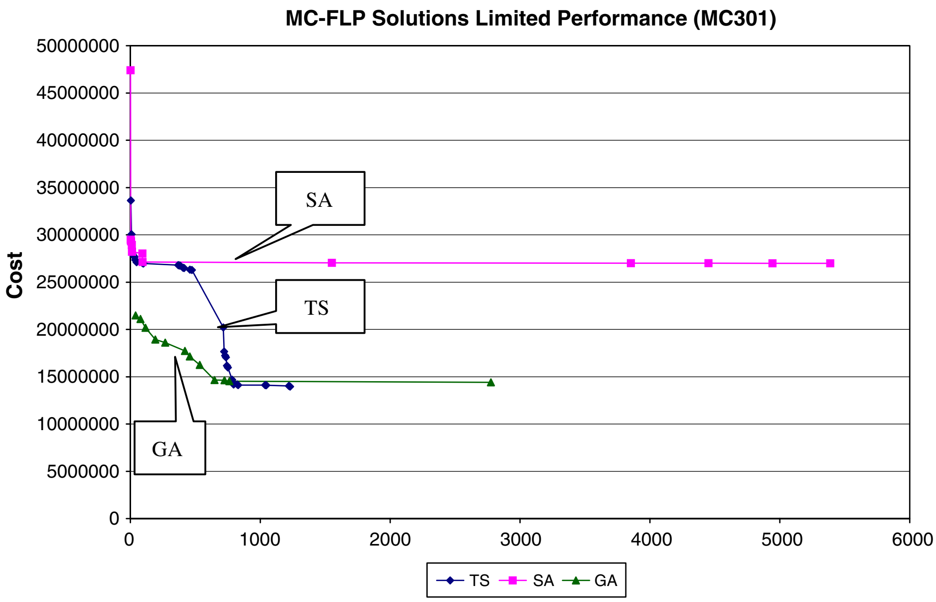

Fig. 6. MCFLP—solution-limited performance.

time, the better the chances that TS will show superior performance. The performance of SA and GA are more partial to problem type. GA can extract more information from fewer solutions than TS and SA. It seems to improve solution quite rapidly. The advantage TS and SA over GA depends on the number of candidate solutions they are allowed to evaluate and the extent of the information they can extract from those solutions for a given problem type. Thus, in general we may conclude that TS should be tried first to the extent that it gives robust results and is easy to develop and implement.

It is important to note that the parameters of the three heuristics used in the study may affect its results and the conclusions. Similarly, the selection of time limit and the number of solutions allowed may affect the results. Therefore, even though using several forms of FLP gives significant breadth to our conclusions, it is important to not over-generalize the results. However, our study has significant contributions to the FL research; it is—to our knowledge—the first to compare TS, GA and SA using more than one problem type. More importantly, it is the first known attempt to apply TS, SA and GA to several classical facility location problems.

References

Arostegui, Jr., M.A., 1997. An empirical comparison of tabu search, simulated annealing and genetic algorithms for facilities locations problems. Doctoral Dissertation, University of Houston.

Avella, P., Benati, S., Cánovas, M.L., Dalby, K., Di Girolamo, D., Dimitrijevic, B., Ghiani, G., Giannikos, I., Guttmann, N., Hultberg, T.H., Fliege, J., Marin, A., Munõz, M.M., Ndiaye, M.M., Nickel, S., Peeters, P., Pérez, B.P., Policastro, S., de Gama, F.A., 1998. Some personal views on the current state and future of location analysis. European Journal of Operations Research 104 (2), 269–287.

Bartmess, A., Cerney, K., 1993. Building competitive advantage through a global network of capabilities. California Management Review 35 (2), 78–103.

Bartmess, A.D., Gray, A.E., Ferdows, K., Cerny, K., Rosenfield, D.B., MacCormack, A., Newman Jr., L.J., Wheelwright, S.C., Kotlier, B., 1994. The plant location puzzle. Harvard Business Review 72 (2), 20–37.

Canel, C., Das, S.R., 2002. Modeling global facility location decisions: integrating marketing and manufacturing decisions. Industrial management and Data Systems 102 (2), 110–118.

Canel, C., Khumawala, B.M., 1996. A mixed-integer programming approach for the international facilities location problem. International Journal of Operations and Production Management 16 (4), 49–68.

Canel, C., Khumawala, B.M., 2001. International facilities location: a heuristic procedure for the dynamic uncapacitated problem. International Journal of Production Research 39 (17), 3975–4000.

Canel, C., Khumawala, B.M., Law, J., 1996. An efficient heuristic procedure for the single-item, discrete lot sizing problem. International Journal of Production Economics 43, 139–148.

Canel, C., Khumawala, B.M., Law, J., Loh, A., 2001. An algorithm for the capacitated, multi-commodity multi-period facility location problem. Computers & Operations Research 28 (5), 411–427.

Cerny, V., 1985. Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm. Journal of Optimization Theory and Applications 45, 41–51.

Chamberland, S., 2004. An efficient heuristic for the expansion problem of cellular wireless networks. Computers & Operations Research 31, 1769–1791.

Chaudry, S.S., He, S., Chaudry, P.E., 2003. Solving a class of facility location problems using genetic algorithms. Expert Systems 20 (2), 86–91.

Cohen, M.A., Lee, H.L., 1989. Resource deployment analysis of global manufacturing and distribution networks. Journal of Manufacturing and Operations Management 2, 81–104.

Corinthal, M.J., Captivo, M.E., 2003. Upper and lower bounds for the single source capacitated location problem. European Journal of Operational Research 151, 333–351.

Drezner, T., 1994. Locating a single new facility among existing, unequally attractive facilities. Journal of Regional Science 34 (2), 237–253.

Drezner, Z., Wesolowsky, G.O., Drezner, T., 1998. On the logit approach to competitive facility location. Journal of Regional Science 38 (2), 313, 15.

Drezner, T., Drezner, Z., Salhi, S., 2002. Solving the multiple competitive facilities location problem. European Journal of Operational Research 142 (1), 138–151.

Dudewicz, E.J., Dalal, S.R., 1975. Allocation of observations in ranking and selection with unequal variances. Sankhya, B 37, 271–281.

Ghosh, D., 2003. Neighborhood search heuristics for the uncapacitated facility location problem. European Journal of Operational Research 150, 150–162.

Giddings, A.P., Bailey, T.G., Moor, J.T., 2001. Optimality analysis of facility location problems using response surface methodology. International Journal of Physical Distribution & Logistics Management, Bradford 31 (1), 38–52.

Glover, F., 1989. Tabu search—part I. ORSA Journal of Computing 1 (3), 190–206.

Glover, F., 1990a. Tabu search—part II. ORSA Journal of Computing 2 (3), 4–32.

Glover, F., 1990b. Tabu search: a tutorial. Interfaces 20 (4), 74–94.

Glover, F., 1993. A user's guide to tabu search. Annals of Operations Research.

Goldenberg, D.E., 1989. Genetic algorithms in search, optimization, and machine learning. Addison Wesley Publishing Company, Inc, Reading, MA.

Hakimi, S.L., 1983. On locating new facilities in a competitive environment. European Journal of Operational Research 12 (1), 29–36.

Harkness, J., ReVelle, C., 2003. Facility location with increasing production costs. European Journal of Operational Research 145 (1), 1–13.

Haug, P., 1985. A multi-period, mixed integer programming model for multinational facility location. Journal of Management 11 (3), 83–92.

Hinojosa, Y., Puerto, J., Fernandez, F.R., 2000. A multiperiod two-echelon multicommodity capacitated plant location problem. European Journal of Operational Research 123 (2), 271–291.

Hodder, J., Dincer, M.C., 1986. A multifactor model for international plant location and financing under uncertainty. Computers & Operations Research 13 (5), 601–609.

Hodder, J.E., Jucker, J.V., 1985. International plant location under price and exchange rate uncertainty. Engineering Costs and Production Economics 9, 225–229.

Holland, J.H., 1975. Adoption in Natural and Artificial Systems. The University of Michigan Press, Ann Arbor.

Hormozi, A.M., Khumawala, B.M., 1996. An improved algorithm for solving a multi-period facility location problem. IIE Transactions 28, 105–114.

Jucker, J.V., Carlson, R., 1976. The simple plant-location problem under uncertainty. Operations Research 24 (6), 1045–1055.

Khumawala, B.M., 1974. An efficient heuristic procedure for the capacitated warehouse location problem. Naval Research Logistics Quarterly 21 (4), 609–623.

Kincaid, R.K., 1992. Good solutions to discrete noxious location problems via metaheuristics. Annals of Operations Research 40, 265–281.

Kirkpatrick, S., Gelatt, C., Vicchi, P., 1983. Optimization by simulated annealing. Science 220, 671–679.

Kuik, R., Salomon, M., Van Wassenhove, L.N., Maes, J., 1993. Linear programming, simulated annealing and tabu search heuristics for lot sizing in bottleneck assembly systems. IIE Transactions 25 (1), 62–72.

Law, A.M., Kelton, W.D., 1991. Simulation Modeling and Analysis, second ed. McGraw-Hill, New York.

Lee, J., Kim, Y., 1996. Search heuristics for resource constrained project scheduling. Journal of the Operational Research Society 47, 678–689.

Marett, R., Wright, M., 1996. A comparison of neighborhood search techniques for multi-objective combinatorial problems. Computers & Operations Research 23 (5), 465–483.

Melkote, S., Daskin, M.S., 2001. Capacitated facility location/network design problems. European Journal of Operational Research 129 (3), 481–495.

Murata, T., Hisao, I., Tanaka, H., 1996. Genetic algorithms for flow shop scheduling problems. Computers and Industrial Engineering 30 (4), 1061–1071.

Naik, B., Chakravarty, A.K., 1993. A fuzzy QFD approach for a global strategy of manufacturing facility. Production and Operations Management Meeting Proceedings, Boston.

Nozick, L.K., 2001. The fixed charge facility location problem with coverage restrictions. Transportation Research Part E 37 (4), 281–297.

Owne, S.H., Daskin, M.S., 1998. Strategic facility location: a review. European Journal of Operations Research 111 (3), 423–447.

Pomper, C., 1976. International Investment Planning: An Integrated Approach. North Holland Publishing Company, Amsterdam.

Porter, M., 1998. Clusters and the new economics of competition. Harvard Business Review 76 (6), 77–90.

Shaver, J.M., 1998. Do foreign-owned and US-owned establishments exhibit the same location pattern in us manufacturing industries. Journal of International Business Studies 29 (3), 469–492.

Sinclair, M., 1993. Comparison of the performance of modern heuristics for combinatorial optimization on real data. Computers & Operations Research 20 (7), 687–695.

Syarif, A., Yun, Y., Gen, M., 2002. Study on multi-stage logistic chain network: a spanning tree-based genetic algorithm approach. Computers & Industrial Engineering 43 (1–2), 299–315.

Taillard, E., 1991. Robust taboo search for the quadratic assignment problem. Parallel computing 17 (4–5), 443–455.

Tian, P., Wang, H., Zhang, D., 1996. Simulated annealing for the quadratic assignment problem. Computers and Industrial Engineering 31 (3/4), 925–928.

Tragantalerngsak, S., Holt, J., Ronnqvist, M., 2000. An exact method for the two-echelon, single-source, capacitated facility location problem. European Journal of Operational Research 123 (3), 473–489.

Van Laarhoven, P.J.M., Aarts, E.H.L., 1987. Simulated Annealing: Theory and Applications. Reidel, Dordrecht.

Volgenant, T., 1996. Facility location: a survey of applications and methods. Interfaces 26 (3), 117–133.

Wang, Q., Batta, R., Bhadury, J., Rump, C.M., 2003. Budget constrained location problem with opening and closing of facilities. Computers & Operations Research 30, 2047–2069.

Warszawski, A., 1973. Multi-dimensional location problems. Operational Research Quarterly 24 (2), 165–179.

Zhou, C., Andrew, D., Yang, J.Y., 2002. Locational determinants of Japanese foreign direct investment in china. Asia Pacific Journal of Management 19 (1), 63–87.