Examining the Performance of Six Heuristic Optimisation Techniques in Different Forest Planning Problems

Timo Pukkala and Mikko Kurttila

Pukkala, T. & Kurttila, M. 2005. Examining the performance of six heuristic optimisation techniques in different forest planning problems. Silva Fennica 39(1): 67–80.

The existence of multiple decision-makers and goals, spatial and non-linear forest management objectives and the combinatorial nature of forest planning problems are reasons that support the use of heuristic optimisation algorithms in forest planning instead of the more traditional LP methods. A heuristic is a search algorithm that does not necessarily find the global optimum but it can produce relatively good solutions within reasonable time. The performance of different heuristics may vary depending on the complexity of the planning problem. This study tested six heuristic optimisation techniques in five different, increasingly difficult planning problems. The heuristics were evaluated with respect to the objective function value that the techniques were able to find, and the time they consumed in the optimisation process. The tested optimisation techniques were 1) random ascent (RA), 2) Hero sequential ascent technique (Hero), 3) simulated annealing (SA), 4) a hybrid of SA and Hero (SA+Hero), 5) tabu search (TS) and 6) genetic algorithm (GA). The results, calculated as averages of 100 repeated optimisations, were very similar for all heuristics with respect to the objective function value but the time consumption of the heuristics varied considerably. During the time the slowest techniques (SA or GA) required for convergence, the optimisation could have been repeated about 200 times with the fastest technique (Hero). The SA+Hero and SA techniques found the best solutions for non-spatial planning problems, while GA was the best in the most difficult problems. The results suggest that, especially in spatial planning problems, it is a benefit if the method performs more complicated moves than selecting one of the neighbouring solutions. It may also be beneficial to combine two or more heuristic techniques.

Keywords ecological planning, habitat suitability index (HSI), Hero, genetic algorithms, random search, simulated annealing, tabu search
Authors´ addresses *Pukkala*: University of Joensuu, Faculty of Forestry, P.O. BOX 111, FIN-80101 Joensuu, Finland; *Kurttila*: Finnish Forest Research Institute, Joensuu Research Centre, P.O. Box 68, FIN-80101 Joensuu, Finland
E-mail timo.pukkala@forest.joensuu.fi; mikko.kurttila@metla.fi
Received 12 March 2004 **Revised** 13 August 2004 **Accepted** 27 August 2004

1 Introduction

If forest stands are regarded as being independent of each other in forest management planning, ranking of treatment alternatives can be made at the level of a single forest stand compartment. The management alternative that gives the highest return or minimizes the consumption of resources is selected for every stand. Usually, however, the stand treatments depend on each other. This is the case for instance when the decision makers want an even flow of timber from their forests. In addition to temporal relationships, also spatial dependencies between stand treatments may be considered and explicitly examined in optimisation.

Management objectives in today's forest planning are diverse and complex. Ecologically related objectives emphasize the location of resources in the maintenance of natural diversity. Economic objectives may also stress the locations of resources. For instance, creation of clusters of stands that can be cut during the same period may decrease the costs of logging operations. These kinds of spatial objectives cannot be calculated directly as a sum of stands over the whole planning area, which makes optimisation complicated. In addition, the number of decision makers or persons who want to have a say in the planning process is often more than one, calling for special problem formulations. These changes and the technical developments in the computers' processing capacity have resulted in increased adoption of heuristic optimisation techniques in forest planning calculations.

Heuristic optimisation techniques have been used for instance in harvest scheduling problems including adjacency constraints (e.g. Borges et al. 2002). Approaches based on random search techniques were used first (e.g. O'Hara et al. 1989, Nelson and Finn 1991). Use has also been made of LP in combination with heuristic algorithm. Heuristic rules have been used to process the LP solution and to modify an LP matrix to produce a solution at the next level (e.g. Weintraub et al. 1994, Tarp and Helles 1997).

Towards the end of the 1990's "modern" heuristic techniques (Reeves 1993a) were more frequently applied to forest management problems. Simulated annealing (SA) (e.g. Dahlin and Sallnäs 1993, Lockwood and Moore 1993, Öhman and Eriksson 1998, Öhman 2000), tabu search (TS) (e.g. Bettinger et al. 1997, Boston and Bettinger 1999) and genetic algorithms (GA) (e.g. Lu and Eriksson 2000, Bettinger et al. 2002) have been the most often applied basic techniques. Different modifications of the basic techniques have been examined and presented widely in the field of operational analysis, but also in forestry (e.g. Bettinger et al. 1999, Boston and Bettinger 2002, Falcão and Borges 2001, 2002). Through modified search process and different neighbourhood definitions these modifications aim at increasing the probability for finding close-to-optimal solutions. More details of different heuristic techniques can be found e.g. from Glover and Laguna (1993), Pham and Karaboga (2000) and Bettinger et al. (2002).

Different heuristic techniques use different search procedures and control parameters that guide the search process. Therefore, their ability to solve different forest management planning problems may vary. In planning that requires the use of heuristics, one should be able to design or to select the technique that is best for that particular problem. In this choice, the objective function value and the time needed to find the solution are important criteria.

A few studies have compared different heuristic techniques in forest management planning context (e.g. Boston and Bettinger 1999, Bettinger et al. 2002, Boston and Bettinger 2002, Crowe and Nelson 2002, Falcão and Borges 2002, Nalle et al. 2002, Palahí et al. 2004). Some of them compare a new heuristic to existing techniques (e.g. Falcão and Borges 2002, Nalle et al. 2002, Crowe and Nelson 2002). Bettinger et al. (2002) compared eight heuristic techniques (random search, simulated annealing, great deluge, threshold accepting, tabu search with 1 opt and 2 opt moves, genetic algorithm and a hybrid tabu search / genetic algorithm process). These techniques were applied to three different and increasingly difficult wildlife planning problems, two of the problems including spatial objectives. The main consideration was given to the objective function value that the techniques could find. Also the time needed to produce a single solution was reported.

A problem in the comparison of heuristic methods is that the results depend on the problem type, method-specific parameters, and the time the method is given to search the decision space. Therefore, different studies report differing results on the performance of heuristic methods. For instance Bettinger et al. (2002) concluded that simulated annealing and tabu search are better than genetic algorithm in spatial planning problems, whereas Palahí et al. (2004) found that genetic algorithm was clearly superior to simulated annealing and tabu search in the spatial version of their planning problem.

The parameters of a heuristic determine how much time it will require to complete the search. Even the simplest techniques such as random search can find good solutions, but at a very high computational cost. Because of the positive correlation between solution time and objective function value, ranking of heuristics is difficult without simultaneous consideration of computing time and objective function value. Out of these two criteria, computing time is more problematic because it greatly depends on the quality of the computer program that is used in the analysis (Hooker 1995). In addition, the total computing time may not alone describe the speed of the method well enough since most of the time may be wasted for non-improving search near the final solution if the stopping parameters are inappropriately set. Therefore, to get a better idea about the performance of a heuristic method, the whole sequence of objective function values during the search process should be examined instead of the mere convergence point. Because the type and size of the problem also affect the performance of heuristic methods, they should be tested with several different planning problems.

This study tested six heuristic optimisation techniques in five different planning problems. The optimisation results were analysed in terms of objective function value, computing time and the temporal development of the objective function value. The distorting effects of differences in implementations (programming skills) were minimised by using the same software for all methods. Only the core of the algorithm was programmed separately for every method, but the data structures, routines for producing random initial solutions, calculation of a move's effect on objective function, etc., were exactly the same with all methods.

The formulations of the planning problems were applications of multi-attribute utility theory. The objective function was an explicitly specified additive utility model. The differences between the planning problems resulted from different objective functions: starting from a simple one-objective planning problem and ending with a spatial multiple ownership problem. All problems are relevant and topical in current Finnish forest planning.

2 Materials and Methods

2.1 The Tested Heuristic Optimisation Techniques

Random Ascent

In random ascent (RA), an initial solution is produced by selecting a random treatment schedule for each forest stand from among the treatment alternatives generated for the stand. Then, a stand and one of its treatment schedules not being in the current solution, are selected randomly. The effect of the suggested change on the objective function value is calculated. If the selected treatment schedule improves the objective function value, it is included in the solution, otherwise not. The search procedure is stopped when the maximum number of trials, as specified by the user, is reached.

Hero

In Hero, maximization of the objective function (additive utility function in this study) consists of two steps (Pukkala and Kangas 1993). First, a random selection of a treatment schedule for each stand produces an initial solution. Second, the technique tests one stand at a time to see whether another treatment schedule would improve the objective function value. The stands and their treatment schedules are inspected sequentially. If increase is detected, the treatment schedule that improves the solution replaces the previous one. When all treatment schedules of all stands are examined in this way, the process is repeated until no schedules can be found that would further improve the solution.

Simulated Annealing

Simulated annealing (SA) is a variant of the descent/ascent techniques of local optimization, and its search process resembles the process of RA. The difference is that SA attempts to avoid getting trapped in local optima by allowing random deteriorations in the objective function value (e.g. Dowsland 1993). The moves that improve the value of the objective function are always accepted. Non-improving moves are accepted with a probability of $p = \exp(-(U_{New} - U_{Old})T_i^{-1})$, where T_i is the current "temperature", and U is objective function value (utility in this study). During the optimisation process, the temperature cools (which imitates the cooling process of melted metal), according to a given cooling schedule. The user can define the cooling schedule. At high temperatures the probability for accepting inferior moves is high (the melted metal moves easily), but as the temperature decreases (the metal solidifies), the probability decreases.

Parameters that the user has to specify when using SA are starting and stopping temperature, cooling schedule and the number of iterations at each temperature. The number of iterations can change during the cooling process, for example increase when the temperature cools.

Simulated Annealing + Hero

A new technique based on a combination of SA and Hero (SA + Hero) was also tested. The idea of cooling and that of accepting inferior solutions were applied in the same way as in SA, whereas the neighbourhood was searched in the same way as in the Hero, i.e. sequentially. All moves that improved the objective function value were accepted. At every temperature, all schedules for all compartments were inspected once and sequentially after which the temperature was changed and the same process was repeated until a stopping temperature was reached. This technique has been applied in the study of Kurttila and Pukkala (2003).

Tabu Search

Search memory in the form of tabu lists are the key characteristics of tabu search (TS). They control the search process for instance by prohibiting the repetition of recent moves. The length of the tabu list defines the number of iterations during which a treatment schedule that participated in a move may not be included in or removed from the solution. As in previous techniques, the search process of TS starts from a random initial solution. Then, several candidate moves, i.e. randomly selected treatment schedules of random compartments, are produced. Among these moves, the best non-tabu move is made. If all moves are in the tabu list, the move that has the shortest tabu tenure is selected. However, an elite move, which is a move that produces the best solution so far, is always accepted. The length of the tabu-list (duration of tabu tenure) can be different for moves that enter the solution and for moves that are removed from the solution (for more details, see e.g. Glover and Laguna 1993).

Genetic Algorithms

Unlike the heuristic optimisation techniques described above, the search process of genetic algorithms (GA) is not based on neighbourhood search. Instead, GA is based on an initial population of solution alternatives, their evaluation and their breeding. The alternative solutions are called parent chromosomes, which are processed by crossing over (combining parts of two or more chromosomes) and by mutation (random change in one or several genes, or compartments). These operations result in a new chromosome (offspring). One of the two parents of a new chromosome is selected with the probability proportional to its ranking. The second parent is chosen randomly with an equal probability for all chromosomes. In the incremental GA technique used in this study, the new chromosome replaces one initial chromosome. The removed chromosome is selected based on its objective function value, the probability of removal being highest for chromosomes that have a low objective function value. The updated group of chromosomes is called generation (for more details, see e.g. Reeves 1993b).

2.2 Test Data

The heuristic optimisation techniques were tested with a compartment level forest data collected from North Karelia, Finland. The Forest Centre of North Karelia surveyed the test forest using visual compartment inventory. The total forest area was 984.8 ha and it was divided into 736 stands during the inventory. The forest area includes 26 non-industrial private forest holdings more than 5 ha in size (the holdings range from 5.1 ha to 395.1 ha). The largest holding was formed artificially by combining the compartments of those holdings for which no forest plan had been ordered from the Forest Centre.

In the beginning of the planning period, the mean growing stock volume was 141.0 m^3/ha. The proportions of pine (*Pinus sylvestris*), spruce (*Picea abies*) and broad-leaved trees were, 48%, 35% and 17%, respectively. The initial age distribution of stand compartments was as follows: younger than 20 years 22%; 20–80 years 57%; and more than 80 years 21%. The current annual increment was estimated at 5.7 m^3/ha.

The inventory data were fed into the Monsu forest planning software (Pukkala 2001). The length of the planning period was 60 years, divided into three 20-year-long sub-periods. The automatic stand treatment simulator of Monsu was used to produce one to eight optional treatment schedules for each stand. The average number of treatment schedules per stand was 3.8. The possibility to regenerate a stand depended on stand age and mean tree diameter. With a high enough stand age or mean diameter, determined according to the official treatment recommendations for Finnish private forests (Luonnonläheinen… 1994), it was possible to simulate regeneration chains (including a suitable final felling and regeneration technique) for the stand. In addition to these earliest possible regeneration cuttings and thinnings, the program also simulated alternatives where a possible cutting was postponed to the next period or later. For mature stands, one of the simulated treatment schedules was the “No treatments” option. However, the tending treatments for young stands were never postponed.

2.3 Objective Variables

Altogether six management objectives, corresponding fairly well to the forest management goals examined by Ihalainen (1992), were included in the objective functions of planning problems: asset value of the forest, measured by soil expectation value (SEV) at the end of the three sub-periods (Objectives 1–3), and corresponding to the economic security goal by Ihalainen (1992); discounted net income (NPV) from timber sales (Objective 4) indicating the “released” return; recreation score at the end of the third sub-period (Objective 5), calculated as the mean area-weighted recreation score of all stands (Pukkala et al. 1988), indicating recreation possibilities; and ecological quality of the landscape (Objective 6), representing the nature values of Ihalainen (1992).

Holding-specific importances of management objectives were generated for the 26 holdings of the case study area using the method developed by Pukkala et al. (2003). This method produced a similar set of rankings of management objectives as has been observed in empirical studies (e.g. Ihalainen 1992).

The management objective 6 aimed at improving the living conditions of flying squirrels. The suitability of a forest stand for flying squirrel was determined with a threshold value of a multiplicative habitat suitability index (HSI) according to the principles of an earlier study (Kurttila et al. 2002). The HSI was computed from the following variables: growing stock volume, proportion of the stand volume made up of spruce and deciduous trees, mean diameter of deciduous trees and the volume of dead standing trees (e.g. Mönkkönen et al. 1997, Hanski 1998, Reunanen et al. 2000). In planning problems 2 and 3, which were non-spatial, the aim was to increase the total area of forest stands having a HSI greater than 0.4 (threshold value used by Kurttila et al. (2002)). Planning problems 4 and 5 were spatial and aimed at creating and clustering habitats by using the proportion of “similar-stand boundary” (both neighbour stands having a HSI > 0.4) of the total boundary length as an objective variable. The use of this objective variable requires that the length of the common boundary of each stand with every neighbouring stand must be known.

GIS routines were used to derive the adjacency information.

2.4 Planning Problem Formulations

Each heuristic optimisation method was used 100 times to solve five different planning problems. In these problems, the six management objectives were targeted in various ways to the whole planning area and individual holdings. The objective functions of the problems were written in the form of an additive utility model:

$$U = \sum_{i=1}^{I} b_i u_i(q_i) \qquad (1)$$

where U is the total utility, I is the number of goals, b_i is the relative importance of management objective i, u_i is a scaled sub-utility function for management objective i, and q_i is the value of objective i. The sub-utility functions transform the absolute value of the variable measured in its own units to a relative sub-utility value. These functions were determined through the smallest possible, target level, and the largest possible value of the objective variable, and the respective priorities. The relative sub-utility values were weighted by the relative importance of the objective variable (b_i) and summed (Pukkala and Kangas 1993).

Problem 1

In the first problem, only one variable was included in the objective function (objective weights generated for the holdings were ignored). The net present value (NPV) of cutting revenues from the whole forest area was the maximized objective. In the optimisation, it suffices to find the best treatment alternative among the simulated ones separately for each compartment.

Problem 2

In the second problem, six management objectives were included in the objective function but the holding-specific variation in the importance of management objectives was ignored. The importances of the six landscape-level management objectives (b_i, i = 1,…,6) were determined as means of the holding-specific weights

$$b_i = \frac{1}{K}\sum_{k=1}^{K} a_{ik}$$

where K is the number of holdings, and a_{ik} is the importance of objective i in holding k). They were as follows: $b_1 = 0.0717$, $b_2 = 0.0717$, $b_3 = 0.0717$ (asset value of the forest at the end of the three sub-periods); $b_4 = 0.3672$ (discounted net income from timber sales); $b_5 = 0.2521$ (recreation score at the end of the third sub-period); and $b_6 = 0.1657$ (ecological objective). Note that the overall global importance of management objectives is nearly the same also in Problems 3 to 5.

The sub-utilities (u_i in Eq. 1) of the recreation score, flying squirrel habitat area and NPV increased linearly from their minimum to their maximum values. For the SEVs at the end of the sub-periods, the minimum quantity gave a sub-utility of zero, the planning area's mean initial SEV produced a sub-utility of 0.9, and the maximum possible quantity of the SEV in the holding gave a sub-utility of one.

Problem 3

Problem 3 consisted of a set of partial holding-level optimisation tasks. The formulation took into account the varying objective weights of individual forest owners. The objective function corresponding to this problem was

$$U = \sum_{k=1}^{K} w_k \sum_{i=1}^{I} a_{ik} u_{ik}(q_{ik}) \qquad (2)$$

where K is the number of forest holdings, w_k is the weight of holding k, I is the number of holding-level management objectives (6 in all holdings), u_{ik} is a scaled sub-utility function of management objective i in holding k, q_{ik} is the value of objective i in holding k, and a_{ik} is the relative importance of management objective i in holding k. The relative importance of each holding is the same ($w_k = 1/K$). The holding-specific sub-utility functions (u_{ik} in Eq. 2) were formed in the same way as in Problem 2.

Problem 4

In Problem 4, the objective function corresponds to Problem 2 but the problem was made more complex by changing the non-spatial ecological objective variable (total habitat area) into a spatial one (proportion of similar-stand-boundary).

Problem 5

Problem 5 was the most complicated, including a spatial planning-area-level objective and holding level objectives in the same problem (for more details, see Kurttila and Pukkala 2003). The applied objective function was:

$$U = w_l b_6 u_6(q_6) + \sum_{k=1}^{K} w_k \sum_{i=1}^{5} a_{ik} u_{ik}(q_{ik})$$

where U is the total utility, w_l is the weight of the landscape level, b_6, u_6 and q_6 are, respectively, the relative local importance, sub-utility function and the amount of the landscape level management objective, K is the number of forest holdings, w_k is the weight of holding k, and a_{ik}, u_{ik} and q_{ik} are, respectively, the relative importance, sub-utility function and value of management objective i in holding k.

The importance of the ecological, planning area-level objective, and the importances of individual holdings' utility functions were defined by transferring the global priority of the flying squirrel habitat objective ($w_k a_{6k}$) from the holding-specific utility functions to the landscape level:

$$w_l b_6 = \sum_{k=1}^{K} w_k a_{6k}$$

The respective importances of the holding levels were set to zero. Thus, holdings where the importance of the ecological objective was larger than zero "transferred" the respective proportion of their importance to the planning-area level. The minimum and maximum values of the planning-area-level boundary proportion's sub-utility function were determined for the whole planning area. The weights and sub-utility formulations of the other holding-level objectives were similar to those in the holding level plans of Problem 3.

2.5 Parameter Values for Heuristic Optimisation Techniques

The parameters that control the search process of a heuristic were inspected one at a time, keeping the other parameter values constant. The value of the parameter was increased or decreased until the change no longer resulted in a clear improvement of the objective function value. The eventual combination of parameter values was assumed to result in a search process that, on one hand, produces a very good objective function value but, on the other hand, can be extremely slow. The idea was to find values that allow the method to use enough time to reach a solution that no longer improves or improves only very slowly. The parameter values were defined using Problem 3. The values obtained for a technique were used in all planning problems.

The following parameter values and settings were used in different heuristic techniques:

RA: The number of trials was 500 000 and the initial solution was produced by selecting randomly one treatment schedule for each stand among the simulated schedules.

Hero: The initial solution was produced randomly i.e. as in RA.

SA: The starting temperature (T_0) was computed from $T_0 = 0.1 \times N^{-1}$, where N is the number of compartments in the planning area. This formula is based on the assumption that the effect of a compartment on the objective function value, which ranges from 0 to 1, is at most N^{-1}. Furthermore, because the various treatment schedules of a compartment are, with several conflicting goals, often nearly equally good, N^{-1} was multiplied by 0.1, the result being a guess for the magnitude of local optima. The cooling multiplier was set to 0.95. In the starting temperature, the number of iterations was $5N$. The number of iterations increased by 5% at each temperature change. The search process was stopped when a freezing temperature of $0.01 \times T_0$ was reached.

SA+Hero: In the beginning of SA+Hero, a random treatment schedule was selected for every stand. This selection was used as the starting point for the SA+Hero search process. The initial and freezing temperatures as well as the cooling schedule were similar as in SA.

TS: In the TS applied in this study, an iteration means the production of a set of candidate moves and updating the solution. The number of candidate moves was 100. The number of iterations was set to 10 000. The initial tabu tenure was 100 iterations for schedules that were removed from the solution and 20 iterations for schedules that were included in the solution.

GA: The initial population consisted of 30 random combinations of treatment schedules for compartments. The number of generations was 50 000. In the production of a new chromosome, the first parent was selected so that chromosomes (solutions) with good objective function values had a higher probability to become selected. The probability was directly proportional to the ranking of the solution (see Reeves 1993b). The other parent chromosome was selected randomly, the probability being independent of the objective function value. Two crossing over points were determined randomly, and the sections between these points were taken from the second parent. The number of mutations in the new chromosome was 5, which means that the treatment schedule of five randomly selected compartments was changed. The new chromosome replaced an old one, the probability of becoming replaced being inversely proportional to the ranking of the chromosome.

3 Results

3.1 Objective Function Value

The optimisation was repeated 100 times with each planning problem and optimisation technique. In this section, the maximum and minimum values, standard deviations and averages of these repetitions are reported. In addition, the time consumption of the techniques is analysed.

In Problem 1, Hero and SA+Hero techniques always found the best and the globally optimal solution (Table 1). Also RA, SA and TS often found a solution where the objective function value was very close to 1, but GA was clearly inferior to other methods.

In Problems 2 and 3 the objective function values found by different techniques were generally very close to each other. SA and SA+Hero found the best values in these problems. The results of TS were very close to the values of these two techniques. Hero had the poorest mean objective function value in Problem 2 and GA in Problem 3.

For Problems 4 and 5 with spatial objectives, GA found the best solutions with respect to the maximum and mean objective function value. The results of SA and SA+Hero were very close to each other. Hero found the poorest results for these problems.

The small standard errors of the 100 repeated optimisations result in very narrow confidence intervals, meaning that very small differences (0.001 or more) in mean objective function values are significant. It may for instance be concluded that GA was significantly poorer than all other methods in the simplest planning problem (no. 1) but better than all other methods in the most difficult problems 4 and 5. In addition, Hero but also RA were significantly inferior to the other methods in the spatial planning problems 4 and 5.

3.2 Time Consumption

The time that was used for the optimisation process differed remarkably between the techniques. Because Hero was always the quickest, the other algorithms were compared to it in Fig. 1 (the time that Hero used =1). Also SA+Hero, RA and TS converged relatively fast. The slowest optimisation techniques, SA and GA, used about 200 times more time than Hero did.

The objective function values, when the fastest and the second fastest optimisation techniques converged, are shown in Fig. 2. When Hero converged, TS had still a relatively low objective function value. GA was still producing and evaluating the initial population of solutions, and had not even started its search. Only the combination of Hero and SA in problem 4 had reached an objective function value equally good as Hero.

At the time point when the second fastest technique (RA) finished, all techniques except GA had already found relatively good solutions, and in some problems TS and SA had even found better solutions than RA.

Fig. 3 shows the development of the objective function value in one optimisation for Problem 3.

Table 1. Mean, maximum and minimum objective function values among the 100 repetitions, and the standard deviations of the objective function values. The best value (the highest maximum and minimum and lowest standard deviation) is in boldface and the worst value is in italics.

	Hero	SA+Hero	RA	TS	SA	GA
Problem 1						
Mean	**1.0000**	**1.0000**	0.9995	0.9996	0.9995	*0.9765*
Max	**1.0000**	**1.0000**	**1.0000**	0.9998	**1.0000**	*0.9793*
Min	**1.0000**	**1.0000**	0.9991	0.9991	0.9991	*0.9737*
Sd	**0.0000**	**0.0000**	0.0005	0.0002	0.0005	*0.0012*
Problem 2						
Mean	*0.7577*	0.7713	0.7660	0.7708	**0.7734**	0.7632
Max	*0.7607*	0.7721	0.7690	0.7725	**0.7751**	0.7648
Min	*0.7551*	0.7701	0.7626	0.7694	**0.7717**	0.7604
Sd	0.0012	**0.0004**	*0.0015*	0.0007	0.0008	0.0009
Problem 3						
Mean	0.8058	**0.8084**	0.8069	0.8076	0.8081	*0.8019*
Max	0.8080	0.8091	0.8086	0.8091	**0.8092**	*0.8033*
Min	0.8032	**0.8076**	0.8052	0.8058	0.8072	*0.8000*
Sd	*0.0009*	**0.0003**	0.0008	0.0006	0.0005	0.0007
Problem 4						
Mean	*0.7512*	0.7731	0.7679	0.7723	0.7744	**0.7790**
Max	*0.7568*	0.7744	0.7725	0.7742	0.7766	**0.7875**
Min	*0.7456*	0.7718	0.7639	0.7699	**0.7723**	0.7702
Sd	0.0022	**0.0004**	0.0017	0.0009	0.0009	*0.0032*
Problem 5						
Mean	*0.7998*	0.8076	0.8065	0.8073	0.8076	**0.8170**
Max	*0.8030*	0.8084	0.8085	0.8090	0.8086	**0.8302**
Min	*0.7948*	**0.8067**	0.8038	0.8047	0.8061	0.8058
Sd	0.0019	**0.0003**	0.0010	0.0010	0.0007	*0.0050*

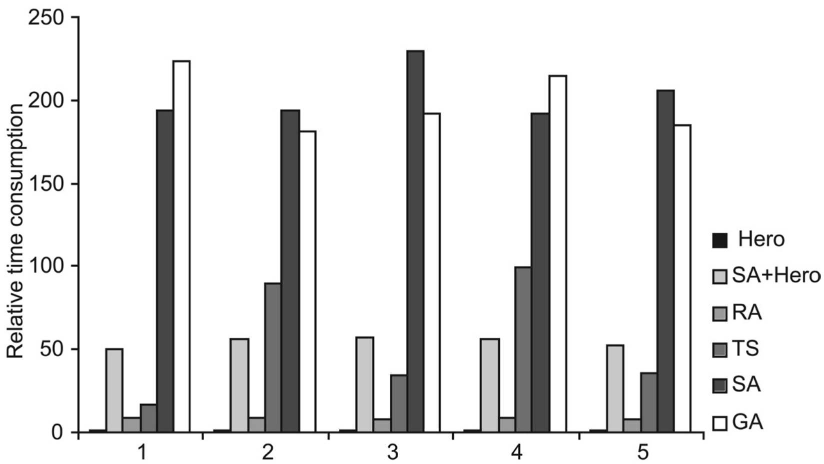

Fig. 1. Duration of the optimisation in five problems in relation to the fastest optimisation technique (Hero = 1).

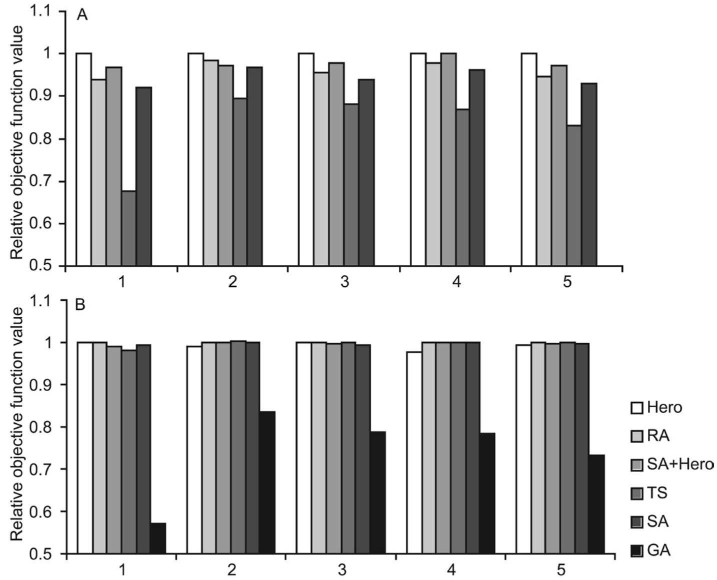

Fig. 2. Relative objective function value of the tested heuristics at the time point when the fastest (A) or second fastest (B) technique converges. In A the relative objective function value of Hero (the fastest technique) is set equal to one, and in B the objective function value of RA (the second fastest technique) is set equal to one.

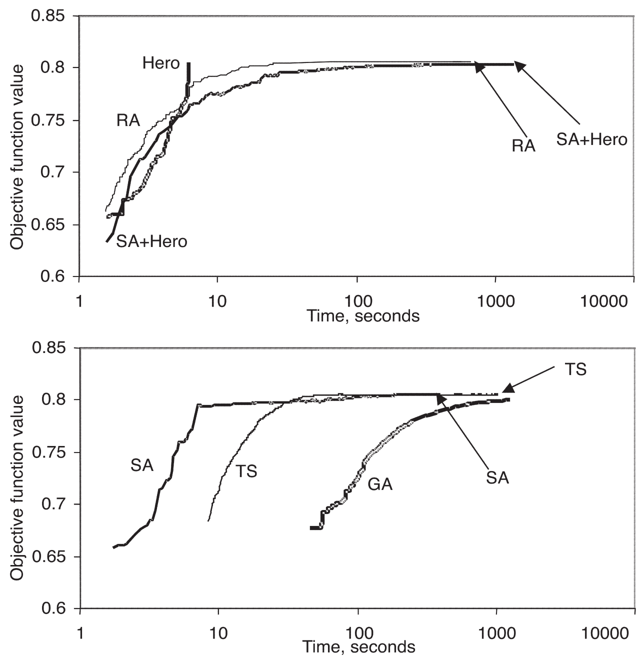

Fig. 3. An example of the development of the objective function value in Problem 3 by different heuristic techniques. The end points for RA, TS, SA, and SA+Hero are shown with an arrow. Note the logarithmic scale of the x-axis.

The quick convergence of the Hero method (less than 10 seconds) and the very slow convergence of GA (more than 1000 seconds) are apparent (note the logarithmic time scale). The figure should be taken as an example rather than universal result, because the shape of the curve depends on method-specific parameter values, as well as on the type and size of the planning problem.

4 Discussion

In this study, six heuristic optimisation techniques were tested in five different forest planning problems, all modelled as utility maximization tasks. The main conclusion that can be drawn from the results is that, with fair parameter values, the methods do not differ much in terms of objective function value. This is not surprising because widely used and accepted techniques were used; the result only indicates that none of the methods is a complete failure. However, some differences in the quality of solutions can be found. Simple techniques such as Hero and RA are good for simple problems only. When the planning problem becomes more complicated, and especially when spatial objectives were included in the problem, the performance of simple techniques (especially Hero) decreases, whereas more complicated techniques improve their performance. According to the results of this study, GA seems especially suitable to solve planning problems with spatial objectives (see Palahí et al. 2004). A result similar as in Palahí et al. (2004) was that GA is not good in very simple problems.

In the study of Bettinger et al. (2002), the techniques were categorized in three classes according to the value of the objective function: very good (simulated annealing, great deluge, threshold accepting, tabu search with 1-opt and 2-opt moves, and a hybrid tabu search / genetic algorithm process), adequate (tabu search with 1 opt moves and genetic algorithm) and less than adequate (random search). The results of our study were not perfectly in line with the results of Bettinger et al. (2002). Our study showed that GA is better than SA and TS in spatial problems, although an opposite result was reported in Bettinger et al. (2002). However, problem formulations, implementation of the technique and parameter settings can greatly affect the performance of different algorithms (Crowe and Nelson 2002), which makes differences between studies understandable. In this study, the good performance of GA in the most difficult problems can result, in addition to the problem and sub-utility formulations, from the fact that it was the only technique where a move could imply more than one change in the solution. The result may imply that it could be beneficial to use more complicated neighbourhood structures in forest planning problems that have spatial objectives (see also Bettinger et al. 2002).

The "true optimum" value is known only for Problem 1. For the other problems it had been possible to try to find out the global optimum using for instance, methods based on extreme value theory (Bettinger et al. 2002). However, the extreme value theory is not reliable (e.g. Reeves 1993a). In addition, relative rather than absolute merits were looked for in this study.

With respect to the total time consumption, the six techniques tested in this study can be divided into three groups: quick techniques (Hero and RA), medium speed techniques (SA+Hero and TS) and slow techniques (SA and GA). The differences in time consumption were more than 100-fold: Hero optimisation could have been repeated 200 times during one GA run. In planning situations where an immediate solution is expected from the planning program, like in interactive optimisation, the use of slow techniques may not be worthwhile if a quicker method is sufficient in the problem at hand. However, since quick techniques are not necessarily good in difficult problems, it may be recommended that interactive heuristic optimisation should be pursued only with simple problems in which quick methods work well.

When observing the temporal development of the objective function value, the picture about the speed of different methods becomes somewhat different. If the time consumption of Hero is taken to mean "quick", then all other methods are slow (except SA+Hero in some problems) because they are still far behind of the objective function value reached by Hero. If the time that RA requires for convergence is regarded to bisect methods to quick and slow ones, then all methods except GA are quick because they have reached good and close-to-final objective function values. If those solutions were regarded acceptable, the search could be stopped. A practical problem is to find such values for the optimisation parameters that will detect when this point is reached.

In the problems of this study, variation in the complexity of the problems was introduced through different objective functions by increasing the number of objectives, interdependencies between objectives and by adding a spatial objective to the objective function. However, compared to e.g. problems with several spatial objectives (e.g minimum and maximum clear-cut area constraints with minimum area constraint for old forest applied in Falcão and Borges (2002), the problem formulations were not computationally extremely complicated. In this study, additional complexity was introduced through combining forest level ecological objectives and holding-specific objectives in the same problem, which is an important topic when new means to improve the biodiversity in Finnish forests are tested.

In this study, the basic versions of the techniques were used. Numerous modifications of the techniques have been presented and tested. These modifications attempt to make the search process more efficient, for instance by guiding the search of alternative solutions to different parts of the solution space. For instance, the 2-opt moves of TS produce candidates where the treatments of two randomly selected compartments are changed (e.g. Bettinger et al. 2002). According to the results of this study, techniques that utilize only simple neighbourhood search during the search process are not so well suited for planning problems with spatial objectives (see also Falcão and Borges 2002), and techniques where several simultaneous changes are made can be successful. In addition, it may be beneficial to consider the stand adjacencies in the selection of 2-opt moves. It also seems that hybrids of different heuristic techniques (and probably also hybrids of LP and heuristics) sometimes work better than one technique alone. Also the hybrid technique of SA+Hero presented in this study performed well; it was fast and produced rather good results.

In the future, the methodological development should concentrate on the specific needs of varying forest planning problems, for example on the stand delineation process and on the use of stands as spatial elements of the planning problem. It looks evident that the use of heuristic optimisation techniques in forest planning calculations will become more common. The most important reasons for this are their simple search process and flexibility, which allows a rather free problem formulation with both additive and multiplicative parts in the objective function (Pukkala 2002), and the possibility to use non-linear sub-utility functions and spatial objective variables. In addition, the search process of heuristic techniques is quite easily understood by decision makers.

References

Bettinger, P., Sessions, J. & Boston, K. 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. Ecological Modeling 94: 111–123.

— , Boston, K. & Sessions, J. 1999. Combinatorial optimization of elk habitat effectiveness and timber harvest volume. Environmental Modeling and Assessment. 4: 143–153.

— , Graetz, D., Boston, K., Sessions, J. & Chung, W. 2002. Eight heuristic planning techniques applied to three increasingly difficult wildlife planning problems. Silva Fennica 36(2): 561–584.

Borges, J.G., Hoganson, H.M. & Falcão, A.O. 2002. Heuristics in multi-objective forest planning. In: Pukkala, T. (Ed.). Multi-objective forest planning. Kluwer Academic Publishers, Dordrecht. p. 119–151.

Boston. K. & Bettinger, P. 1999. An analysis of Monte Carlo integer programming, simulated annealing and tabu search for solving spatial harvest scheduling problems. Forest Science 45: 292–301.

— & Bettinger, P. 2002. Combining tabu search and genetic algorithm heuristic techniques to solve spatial harvest scheduling problems. Forest Science 48(1): 35–46.

Crowe, K. & Nelson, J. 2002. An indirect search algorithm for harvest scheduling under adjacency constraints. Forest Science 49(1): 1–11.

Dahlin, B. & Sallnäs, O. 1993. Harvest scheduling under adjacency constraints – A case study from the Swedish sub-alpine region. Scandinavian Journal of Forest Research 8: 281–290.

Dowsland, K.A. 1993. Simulated annealing. In: Reeves, C.R. (Ed.). Modern heuristic techniques for combinatiorial problems. Blackwell Scientific Publications. p. 20–69.

Falcão, A.O. & Borges, J.G. 2001. Designing an evolution program for solving integer forest management scheduling models: an application in Portugal. Forest Science 47(2): 158–168.

— & Borges, J.G. 2002. Combining random and systematic search heuristic procedures for solving spatially constrained forest management scheduling models. Forest Science 48(3): 608–621.

Glover, F. & Laguna, M. 1993. Tabu search. In: Reeves, C.R. (Ed.). Modern heuristic techniques for combinatiorial problems. Blackwell Scientific Publications. p. 70–150.

Hanski, I. 1998. Home ranges and habitat use in the declining flying squirrel Pteromys volans, in managed forests. Wildlife Biology 4: 33–46.

Hooker, J.N. 1995. Testing heuristics: We have it all wrong. Journal of Heuristics 1: 33–42.

Ihalainen, R. 1992. Yksityismetsänomistuksen rakenne 1990. Metsäntutkimuslaitoksen tiedonantoja – The Finnish Forest Research Institute, Research Notes 405. 41 p. [In Finnish]

Kurttila, M. & Pukkala, T. 2003. Combining holding-level economic goals with spatial landscape-level goals in the planning of multiple ownership forestry. Landscape Ecology 18: 529–541.

— , Pukkala, T. & Loikkanen, J. 2002. The performance of alternative spatial objective types in forest planning calculations: a case for flying squirrel and moose. Forest Ecology and Management 166: 245–260.

Lockwood, C. & Moore, T. 1993. Harvest scheduling with spatial constraints, a simulated annealing approach. Canadian Journal of Forest Research 23: 468–478.

Lu, F. & Eriksson, L.O. 2000. Formation of harvest units with genetic algorithms. Forest Ecology and Management 130: 57–67.

Luonnonläheinen metsänhoito.1994. Metsänhoitosuositukset. Metsäkeskus Tapion julkaisuja 1994:6. 74 p. [In Finnish]

Mönkkönen, M., Reunanen, P., Nikula, A., Inkeriöinen, J. & Forsman, J. 1997. Landscape characteristics associated with the occurrence of the flying squirrel Pteromys volans in old-growth forests of northern Finland. Ecography 20: 634–642.

Nalle, D.J., Arthur, J.L. & Sessions, J. 2002. Designing compact and contiguous reserve networks with a hybrid heuristic algorithm. Forest Science 48(1): 59–68

Nelson, J.D. & Finn, S.T. 1991. The influence of cut-block size and adjacency rules on harvest levels and road networks. Canadian Journal of Forest Research 21: 595–600.

O'Hara, A.J., Faaland, B.H. & Bare, B.B. 1989. Spatially constrained timber harvest scheduling. Canadian Journal of Forest Research 19: 715–724.

Öhman, K. 2000. Creating continuous areas of old forest in long-term forest planning. Canadian Journal of Forest Research 30: 1817–1823.

— & Eriksson, L.O. 1998. The core area concept in forming contiguous areas for long-term forest planning. Canadian Journal of Forest Research 28: 1032–1039.

Palahí, M., Pukkala, T., Pascual, L. & Trasobares, A. 2004. Examining alternative landscape metrics in ecological forest landscape planning: a case for capercaillie in Catalonia. Manuscript.

Pham, D.T. & Karaboga, D. 2000. Intelligent optimisation techniques. Genetic algorithms, tabu search, simulated annealing and neural networks. Springer-Verlag, London. 302 p.

Pukkala, T. 2001. Monsu metsäsuunnitteluohjelma. Ohjelmiston toiminta ja käyttö. Mimeograph at The University of Joensuu, Finland, 75 p. [In Finnish]

— 2002. Measuring non-wood forest outputs in numerical forest planning. A review of Finnish Research. In: Pukkala, T. (ed.). Multi-objective forest planning. Kluwer Academic Publishers. p. 173–207.

— & Kangas, J. 1993. A heuristic optimization method for forest planning and decision making. Scandinavian Journal of Forest Research 8: 560–570.

— , Kellomäki, S. & Mustonen, E. 1988. Prediction of the amenity of a tree stand. Scandinavian Journal of Forest Research 5: 263–275.

— , Ketonen, T., Pykäläinen, J. & Kangas, J. 2003. Predicting the timber supply from private forests – a utility theoretic approach. Forest Policy and Economics 5(3): 285–296.

Reeves, C.R. (Ed.). 1993a. Modern heuristic techniques for combinatorial problems. Blackwell Scientific Publications. 320 p.

— 1993b. Genetic algorithms. In: Reeves, C.R. (ed.). Modern heuristic techniques for combinatorial problems. Blackwell Scientific Publications. p. 151–196.

Reunanen P., Mönkkönen, M. & Nikula, A. 2000. Managing boreal forest landscapes for flying squirrels. Conservation Biology 14: 218–226.

Tarp, P. & Helles, F. 1997. Spatial optimization by simulated annealing and linear programming. Scandinavian Journal of Forest Research 12: 390–402.

Weintraub, A., Barahona, F. & Epstein, R. 1994. A column generation algorithm for solving general forest planning problems with adjacency constraints. Forest Science 40(1): 142–161.

Total of 38 references