Additional Insight into the Performance of a New Heuristic for Solving Spatially Constrained Forest Planning Problems

Jianping Zhu, Pete Bettinger and Rongxia Li

Zhu, J., Bettinger, P. & Li, R. 2007. Additional insight into the performance of a new heuristic for solving spatially constrained forest planning problems. Silva Fennica 41(4): 687–698.

The raindrop method of searching a solution space for feasible and efficient forest management plans has been demonstrated as being useful under a limited set of circumstances, mainly where adjacency restrictions are accommodated using the unit restriction model. We expanded on this work and applied the model (in a modified form) to a problem that had both wood flow and area restriction adjacency constraints, then tested the problem formulation on six hypothetical forests of different sizes and age class distributions. Threshold accepting and tabu search were both applied to the problems as well. The modified raindrop method's performance was best when applied to forests with normal age class distributions. 1-opt tabu search worked best on forests with young age class distributions. Threshold accepting and the raindrop method both performed well on forests with older age class distributions. On average, the raindrop method produced higher quality solutions for most of the problems, and in all but one case where it did not, the solutions generated were not significantly different than the heuristic that located a better solution. The advantage of the raindrop method is that it uses only two parameters and does not require extensive parameterization. The disadvantage is the amount of time it needs to solve problems with area restriction adjacency constraints. We suggest that it may be advantageous to use this heuristic on problems with relatively simple spatial forest planning constraints, and problems that do not involve young initial age class distributions. However, generalization of the performance of the raindrop method to other forest planning problems is problematic, and will require examination by those interested in pursuing this planning methodology. Given that our tests of the raindrop method are limited to a small set of URM and ARM formulations, one should view the combined set of work as additional insight into the potential performance of the method on problems of current interest to the forest planning community.

Keywords forest planning, harvest scheduling, heuristics, raindrop method
Authors' address Warnell School of Forestry and Natural Resources, University of Georgia, Athens, GA 30602, USA **E-mail** pbettinger@warnell.uga.edu
Received 29 November 2006 **Revised** 5 July 2007 **Accepted** 21 September 2007
Available at http://www.metla.fi/silvafennica/full/sf41/sf414687.pdf

1 Introduction

Forest management plans can be developed with traditional mathematical programming techniques, which include linear programming, mixed integer programming, and integer programming (Bevers and Hof 1999, Hof et al. 1994, Hof and Joyce 1992). Forest planning problems are, however, becoming increasingly difficult to solve with these techniques due to the inclusion of adjacency and green-up constraints, which require the use of binary integer decision variables. Modern forest planning involves solving combinatorial optimization problems, and usually as the problem size increases, the complexity of the problem will increase non-linearly. Thus as the problem size increases, solving it may become impractical using traditional mathematical programming methods (Lockwood and Moore 1993), although advances in mathematical programming software continue to mitigate this issue. Murray and Weintraub (2002) illustrate a case where a significant amount of time was required to generate the constraints for a small spatially-constrained planning problem, and although work continues in this area, they suggested in 2002 that it would be unrealistic to solve large problems, or more difficult problems, with exact approaches. As a result, heuristics such as Monte Carlo simulation (Nelson and Brodie 1990), simulated annealing (Dahlin and Sallnäs 1993, Lockwood and Moore 1993, Murray and Church 1995), threshold accepting (Bettinger et al. 2003), tabu search (Bettinger et al. 1997) and genetic algorithms (Glover et al. 1995, Falcão and Borges 2001, Boston and Bettinger 2002) are increasingly being used to address spatially-constrained forest planning problems. The main disadvantage of using heuristics is that they can not guarantee that the global optimum solution to a problem will be located, but they usually can find good solutions to complex planning problems in reasonable amounts of time.

In this study, we take a heuristic, the raindrop method, which was introduced by Bettinger and Zhu (2006), and apply it to a forest management model that maximizes the net present value of a forest plan, and includes wood flow and area restriction adjacency constraints (Murray 1999). The raindrop method has only been tested on one forest management problem, thus the contribution of this research is to determine whether it is useful, as a search process, in broader applications. Bettinger and Zhu (2006) note that "what remains to be seen is whether the model can be applied successfully to the broader range of operations research problems... [including] forest planning problems with other objectives or a more comprehensive suite of constraints."

In Bettinger and Zhu (2006), the raindrop method was shown to be preferable for locating solutions to problems that maximized even-flow of harvest volume and controlled adjacency with the unit restriction model (URM). The difference between this and the area restriction model (ARM) of adjacency is that the URM controls the scheduling of adjacent units (regardless of size) during the green-up period, while the ARM allows adjacent units to be scheduled during the green-up period as long as the total size of the clearcut area does not exceed the maximum clearcut size that is assumed. Thus this second method of modeling adjacency is more complex to both formulate and solve. Our hypothesis is that the raindrop method will be as effective as other standard heuristics in solving problems that include the ARM. In addition, some insight will be gained in the ability to modify and use the heuristic for a more complex problem.

2 Methods

The methods section will first address the problem formulation for the forest planning model. Then the six hypothetical landowners that are modeled will be briefly described, along with the non-spatial economic assumptions pertinent to this research. Finally, the raindrop method will be briefly described along with the modifications necessary to implement it in a forest planning context.

2.1 Forest Planning Model Formulation

The forest planning problem that we investigate attempts to maximize the net present value of timber harvested. The problem formulation is:

Maximize

$$\sum_{t=1}^{T}\left[\sum_{i}^{N}\left(\frac{(V_{it}X_{it}(P-C_{it}))+(VT1_{it}XT1_{it}(P-C_{it}))+(VT2_{it}XT2_{it}(P-C_{it}))}{1.06^{(t-0.5)}}\right)\right] + \sum_{i}^{N}(V_{i20}(P-C_{it}))/1.06^{19.5} \quad (1)$$

Subject to

$$\sum_{t=1}^{T} X_{it} \le 1 \qquad \forall i \quad (2)$$

$$\sum_{t=1}^{T} XT1_{it} \le 1 \qquad \forall i \quad (3)$$

$$\sum_{t=1}^{T} XT2_{it} \le 1 \qquad \forall i, \text{ where } XT1_{it}=1, \text{ otherwise } XT2_{it}=0 \quad (4)$$

$$X_{it} \sum_{j\in N_i \cup S_i} X_{ju}A_j \le \text{MCA} \qquad \forall i,t,u \in t-1 \; to \; t+1, u>0, u\le T \quad (5)$$

$$\sum_{i=1}^{N} V_{i20} - \sum_{i=1}^{N} X_{it}V_{i20} - \sum_{i=1}^{N} XT1_{it}VT1_{i20} - \sum_{i=1}^{N} XT2_{it}VT2_{i20} > 0.9 * \sum_{i=1}^{N} V_{i1} \quad (6)$$

$$AG_{ic} - AG_{it1} > 5 \qquad \forall i \quad (7)$$

$$AG_{ic} - AG_{it2} > 5 \qquad \forall i \quad (8)$$

$$\sum_{i=1}^{N}\left(X_{it}V_{it} + XT1_{it}VT1_{it} + XT2_{it}VT2_{it}\right) > 0.9 * \sum_{i=1}^{N}\sum_{u=1}^{T} X_{iu}V_{i20}/T \qquad \forall t \quad (9)$$

$$\sum_{i=1}^{N}\left(X_{it}V_{it} + XT1_{it}VT1_{it} + XT2_{it}VT2_{it}\right) < 1.1 * \sum_{i=1}^{N}\sum_{u=1}^{T} X_{iu}V_{i20}/T \qquad \forall t \quad (10)$$

$$\sum_{i=1}^{N}\left(X_{it}V_{it} + XT1_{it}VT1_{it} + XT2_{it}VT2_{it}\right) > 0.9 * \sum_{i=1}^{N}\left(X_{it-1}V_{it-1} + XT1_{it-1}VT1_{it-1} + XT2_{it-1}VT2_{it-1}\right) \qquad \forall t \ge 2 \quad (11)$$

$$\sum_{i=1}^{N}\left(X_{it}V_{it} + XT1_{it}VT1_{it} + XT2_{it}VT2_{it}\right) < 1.1 * \sum_{i=1}^{N}\left(X_{it-1}V_{it-1} + XT1_{it-1}VT1_{it-1} + XT2_{it-1}VT2_{it-1}\right) \qquad \forall t \ge 2 \quad (12)$$

where

A_i = area of management unit i

AG_{ic} = clearcut age for management unit i

AG_{it1} = age when first thinning occurs for management unit i

AG_{it2} = age when second thinning occurs for management unit i

C_{it} = logging cost for management unit i harvested in time period t

i, z = management units

P = stumpage price

MCA = maximum clearcut area

N = the total number of management units

N_i = set of all management units adjacent to management unit i

S_i = the set of all management units adjacent to those management units adjacent to management unit i

t, u = planning periods

T = the total number of planning periods in the planning horizon

V_{it} = the available clearcut timber harvest volume from management unit i during time period t

V_{i20} = the unscheduled timber harvest volume from management unit i at the end of period 20, whether or not a harvest had been applied during the period represented by the plan

$VT1_{it}$ = the available first thinning timber harvest volume from management unit i during time period t

$VT2_{it}$ = the available second thinning timber harvest volume from management unit i during time period t

X_{it} = a binary variable, which = 1 if management unit i is clearcut harvested in time period t, and 0 otherwise

$XT1_{it}$ = a binary variable, which = 1 if management unit i is first-thinned in time period t, and 0 otherwise

$XT2_{it}$ = a binary variable, which = 1 if management unit i is second-thinned in time period t, and 0 otherwise

The objective function assesses the difference between value generation and cost of activities (prices versus logging costs) for each activity (clearcut, first thinning, and second thinning) applied to each stand over the time horizon. Decisions are integer in nature, therefore an activity is applied to an entire stand when it is applied. Values associated with harvest activities are discounted from the middle of each planning period, and the value of the ending inventory (that which remains un-cut in the final planning period) is included to fully value the forest enterprise over the 20-year planning horizon. Eq. 2 indicates that each management unit can only be clearcut harvested one time during the planning horizon. Eq. 3 indicates that each management unit can only be "first thinned" one time during the planning horizon. Eq. 4 indicates that each management unit can only be "second thinned" one time during the planning horizon, assuming that it was previously first-thinned. Eq. 5 ensures that the maximum clearcut size will not be violated (assuming the green-up period is 2 yrs). This set of constraints represents a slight modification of the original ARM model provided by Murray (1999), where S_i is a subset of treated stands containing all stands adjacent to neighbors of stand i and stands adjacent to the neighbors of the neighbors, etc. As Murray (1999) suggests, this constraint is a recursive function that senses a sprawling cluster of stands treated within the green-up period, and that the cluster depends on the contiguity of stands and their direct or indirect relationship to stand i. Eq. 6 is an ending volume constraint, where the residual standing volume must be at least 90% of the initial standing volume. Eqs. 7 and 8 ensure that the separation period between thinning and clear cutting activities is at least six years. Eqs. 9 and 10 constrain the volume harvested in each time period to a proportion of the residual, unscheduled, uncut volume. Eqs. 11 and 12 limit the deviation in harvest volume from one period to the next, as a form of harvest stability. This model formation is a model I (Johnson and Scheurman 1977), integer programming problem. The adjacency restrictions are considered the ARM formulation noted in Murray (1999).

2.2 Spatial and Non-Spatial Data

The hypothetical landowners considered in this research are divided into two ownership size groups, small and medium, and the polygons used to represent the management units (Table 1) were derived from an actual land ownership in the southern United States. We assume that the arrangement of the management units is dispersed, and we assigned three age class distributions to each: a young forest, a normal forest, and an older forest (Fig. 1), with the distribution of land by age class similar for both the small- and

Table 1. Characteristics of the hypothetical landowners.

Type	Management units	Area (ha)	Average polygon size (ha)
Small-sized	279	2942	10.5
Medium-sized	477	5821	12.2

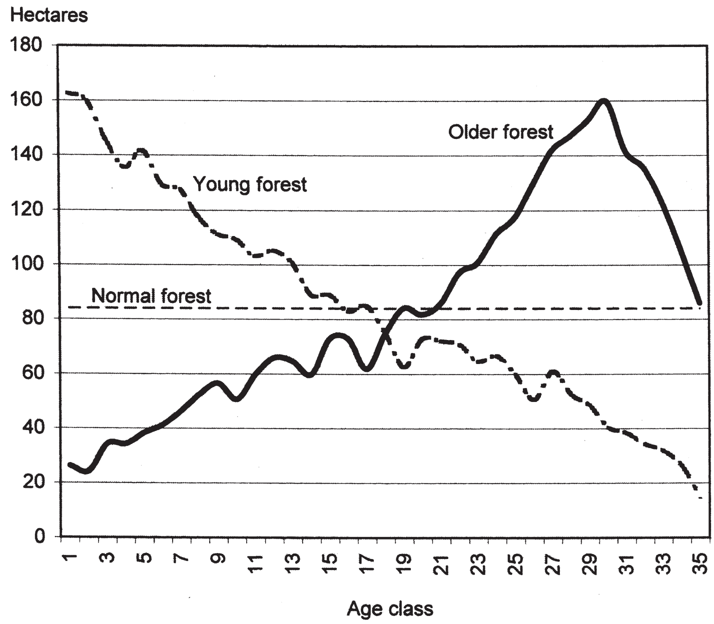

Fig. 1. Initial age class distributions for the small-sized hypothetical landowners examined.

medium-sized landowners. Therefore, a matrix of six hypothetical landowners was available for analysis. The time horizon is 20 years, divided into twenty 1-year time periods. The interest rate assumed is 6 percent. The stumpage prices were obtained from Timber Mart-South (2004), and are \$43.57 per ton for pine sawtimber, \$25.60 per ton for chip-n-saw, and \$6.73 per ton for pine pulpwood. The costs assumed are \$285.89 per ha for mechanical site preparation, \$116.48 for planting, and \$156.66 for a herbaceous weed control treatment. The maximum clearcut size is 97 ha, and the green-up period is assumed to be 2 years.

2.3 Modified Raindrop Method

The raindrop method was introduced by Bettinger and Zhu (2006), where a full description of the process can be found, and is theoretically based on raindrop impact. Like many other heuristics, it is a process that seeks to incrementally improve developing forest plans with iterative changes that are either selected randomly or deterministically. The basic process is that a management unit and a clearcut period are randomly chosen. This choice is forced into the current solution regardless of any potential constraint violations. If there are adjacency constraint violations, they are mitigated in a radiating wave motion. All of the management units contributing to the adjacency constraint violations are added to a list, and the violations are mitigated in order of distance from their centroid to the centroid of the original randomly chosen management unit, with the closest units corrected first. The next best alternative for the affected management unit is inserted into the solution. Any management units that are subsequently affected by this change to the solution are added to the list of affected units. This process continues until all constraint violations have been mitigated. Any mitigation must prevent subsequent adjacency constraint violations with management units which are nearer to the originally randomly chosen unit, thus the impacts radiate outward from the initial choice. Once all infeasibilities have been mitigated, a single iteration of the raindrop method ends.

The main advantage of the raindrop method is that it only uses two parameters, and it has been shown to produce very good solutions for problems that include URM constraints. The two parameters are: 1) the total number of iterations, and 2) the number of iterations that pass before search process reverts to the best solution (stored in memory). Bettinger and Zhu (2006) showed that reverting every 3 or 4 iterations to the best solution produced results that were superior to other heuristics. Simply modeling the raindrop method without reversion does not lead to better solutions. As pointed out by Bettinger and Zhu (2006), one could reasonably assume that as the number of constraints grows, the problem-solving ability of the heuristic may become cumbersome as it attempts to mitigate the infeasibilities that arise. Given our extensive knowledge of the heuristic, we hypothesize that for forest planning problems that include the ARM and wood flow constraints, it may become very difficult to implement this heuristic. This means that if we try to mitigate adjacency violations using a radiating manner, the search for all affected units becomes time-consuming, since when one uses the ARM, one must check the sprawling set of units that describe a clearcut (rather than simply the neighbors that touch each harvest unit). This involves looking backward toward the original randomly chosen management unit more thoroughly than when using the URM constraints (which requires only checking the immediately adjacent units). In addition, changing the timing of harvests to mitigate adjacency constraint violations may result in wood flow constraint violations.

The modified algorithm puts all units that potentially can cause the adjacency constraint violation(s) into a list, then tries to mitigate the adjacency violation(s) by simply setting the units to "no cut" one-by-one for all units in the list. Two options are available:

1) If, after unscheduling a management unit, the adjacency constraints are no longer violated, the next-best harvesting alternative (that does not itself violate constraints backwards, towards the original randomly chosen unit) is assigned to the management unit. This alternative focuses on mitigating any wood flow constraint violations.
2) If, after unscheduling a management unit, the adjacency constraints continue to be violated, the previous status (harvest period) is restored. The logic here is that if unscheduling the harvest does not mitigate the adjacency constraint violation, then the choice of harvest period had limited impact on the constraint violation.

It is possible that after mitigating all of the adjacency constraint violations, one or more of the wood flow constraints could continue to be violated. If this is the case, the process reverts immediately back to the best solution stored in memory. By trial and error, we found that the best parameters for these problems were to run the algorithm for 100,000 iterations on the small-sized problems, and 200,000 iterations on the medium-sized problems. In addition, based on what was learned in Bettinger and Zhu (2006), the search process reverts back to the best solution every 4 iterations.

The modified raindrop method was compared to threshold accepting and 1-opt tabu search. Threshold accepting was initially described by Dueck and Scheuer (1990), and has been applied to forest problems (Bettinger et al. 2002, 2003) with a high level of success. In complex forest planning problems, it has been shown to be as good as simulated annealing and other heuristics, and it is relatively simple to implement, although some parameterization is required. Tabu search was introduced by Glover (1989, 1990), and has also been successfully applied to forestry problems (Bettinger et al. 1997, 1998, 2002). Tabu search with 1-opt moves involves simply changing the status (clearcut period) of one management unit. The change is made deterministically, whereas in threshold accepting the change is made randomly. Each of these heuristics was tested extensively and parameterized for each problem. Zhu (2006) contains a more exhaustive analysis of threshold accepting and tabu search on the problems analyzed here. In order to compare the results, 30 solutions were generated for each of the six hypothetical forest planning problems. Each of the 30 solutions can be considered independent to each other because the initial solution is randomly defined (Bettinger et al. 1998).

3 Results

For two of the three problems that involved small-sized hypothetical landowners, the best forest plans were developed using the modified raindrop method (Table 2). While thirty feasible solutions were generated for each problem and heuristic, tabu search was able to locate the best forest plan that involved the young initial age class distribution. In this case, the best forest plan developed with the modified raindrop method was within 0.4% of the best tabu search forest plan. In each case involving the small-sized hypothetical landowners, the highest average solution value was produced using the modified raindrop method. In addition, the smallest variation was found amongst the plans developed using the modified raindrop method.

When examining the results for the medium-sized hypothetical landowners, the modified raindrop method found the best solution for only one of the three initial age class distributions – the normal age class distribution (Table 2). Tabu search was able to locate the best solution for the young initial age class distribution, and threshold accepting was able to find the best solution for the older initial age class distribution. In case of the young age class distribution, the best solution found using the modified raindrop method was within 1.4% of the best tabu search solution. In the case of the older age class distribution, the best solution found using the modified raindrop method was within 0.1% of the best threshold accepting solution. Tabu search had the highest average solution value in the case of the young age class distribution. In the cases of the normal and older age class distributions, threshold accepting produced the highest average solution values. The variation amongst the solutions was mixed for the different initial age class distributions.

Fisher's least significant difference (LSD) method was applied to the sets of results generated by the three heuristics for each of the six problems. The goal of this analysis was to

Table 2. Quality of solutions generated by three heuristic techniques.

Problem	Best solution	Average solution	Standard deviation
Small-sized forest, normal age class			
Threshold accepting	$14 848 689	$14 632 945	122 556
Tabu search	14 873 804	14 682 830	109 904
Raindrop method	**14 888 572**	14 720 781	107 904
Small-sized forest, young age class			
Threshold accepting	11 404 358	11 209 577	79 117
Tabu search	**11 551 907**	11 376 069	70 280
Raindrop method	11 507 865	11 395 723	66 960
Small-sized forest, older age class			
Threshold accepting	17 902 145	17 534 926	178 391
Tabu search	17 895 195	17 528 975	183 045
Raindrop method	**18 194 232**	17 770 583	167 795
Medium-sized forest, normal age class			
Threshold accepting	30 930 637	30 838 991	58 995
Tabu search	30 950 206	30 714 628	109 727
Raindrop method	**31 083 463**	30 805 706	198 183
Medium-sized forest, young age class			
Threshold accepting	23 045 952	22 888 484	112 586
Tabu search	**23 395 162**	23 217 040	74 480
Raindrop method	23 069 865	22 943 418	78 677
Medium-sized forest, older age class			
Threshold accepting	**36 928 729**	36 792 688	88 160
Tabu search	36 668 720	36 431 426	150 379
Raindrop method	36 888 317	36 737 841	87 160

Table 3. Multiple comparison of the mean solution values for three heuristics when applied to six different forest planning problems. Values represented by different letters are significantly different (analysis of variance F-test $p \leq 0.05$).

Heuristic	Planning problem					
	Small-sized problem			Medium-sized problem		
	Normal	Older	Young	Normal	Older	Young
Tabu search	AB	B	A	B	B	A
Threshold accepting	B	B	B	A	A	C
Raindrop method	A	A	A	A	A	B

determine whether the sample data (30 independent runs of each heuristic) had means that were significantly different at the $p = 0.05$ level. In other words, we used an analysis of variance to test the null hypothesis that the means were equal. Based on the results of the statistical analysis (Table 3), we found that the results generated by the raindrop method were significantly different than the results generated by threshold accepting for the small-sized problem, normal and older forests. The raindrop method results were not significantly different than the tabu search results for the small-sized normal forest problem, even though the raindrop method located the better solution to this problem. For the small-sized problem with the young forest, the raindrop method results were not significantly different than the tabu search results even though tabu search located the better solution. When examining the medium-sized problems, and when considering the older and normal forests, the results generated by the raindrop method were not significantly different than the threshold accepting results, even though threshold accepting located the better solution to one of these two problems. For the medium-sized problem with the young forest, the tabu search results were significantly different than the results generated by the raindrop method.

Each forest plan was developed using a Pentium 4 computer with a 3.0 GHz processor and 1.0 Gb of RAM. The average time to generate a solution for the small-sized forests was 3.0 minutes, 5.6 minutes, and 14.2 minutes for threshold accepting, tabu search, and the modified raindrop method, respectively. The average time to generate a solution for the medium-sized forests was 6.6 minutes, 10.4 minutes, and 25.4 minutes for threshold accepting, tabu search, and the modified raindrop method, respectively. Although anecdotal, we found that developing the programming language for the modified raindrop method (when applied to problems involving the ARM and wood flow constraints) was difficult, and required more time than threshold accepting or tabu search. Even though only two parameters were required, the main issue is in how the modified raindrop method mitigates both ARM and wood flow constraint infeasibilities once they occur.

A final way to evaluate the quality of results generated by the modified raindrop method is to compare them to the results obtained from solving a relaxed linear programming formulation of the same problem. In the case of the relaxed linear programming model, the adjacency constraints were not included. The difference between the best solution generated by the modified raindrop method and the linear programming solution, divided by the linear programming solution, can be used to represent the cost of implementing the green-up and adjacency constraints. We found that the cost of adjacency restrictions ranged from about 1% to about 3.5% of the net present value of each forest plan (Table 4). We found no clear trend that would explain the range of differences from the relaxed linear programming solutions. For example, the modified raindrop method performed very well for the medium-sized hypothetical landowner with either an older or a normal initial age class distribution, but not as well for the small-sized hypothetical landowner with young or

Table 4. Comparison of the results from the raindrop method to results from relaxed linear programming problems.

Problem	Linear programming solution, [a]	Percent difference [b]		
		RD [c]	TS [d]	TA [e]
Small-sized forest, normal age class	$15424070	3.47	3.57	3.73
Small-sized forest, young age class	11897810	3.28	2.91	4.15
Small-sized forest, older age class	18596600	2.16	3.77	3.73
Medium-sized forest, normal age class	31392270	0.98	1.41	1.47
Medium-sized forest, young age class	23689660	2.62	1.24	2.72
Medium-sized forest, older age class	37269520	1.02	1.61	0.91

[a] Relaxed problems, where the adjacency constraints are not included in the problem formulation.
[b] Linear programming solution value – best raindrop method solution value, divided by linear programming solution value
[c] Raindrop method
[d] Tabu search
[e] Threshold accepting

normal initial age class distributions. If one were to surmise that the heuristic did not locate the global optimum solution to the full management problem (even though 30 solutions were generated to each problem), and that the best heuristic solution was within 0.5–1.0% of the global optimum solution to each problem, the cost of the adjacency restrictions would be 0.5–3.0% in net present value of each hypothetical forest. These results are consistent with previously published research on the impact of adjacency constraints for forests in the southeastern U.S. (e.g., Boston and Bettinger 2001).

Solving ARM problems exactly with mixed-integer programming methods has been successful on small problems (McDill et al. 2002, Murray and Weintraub 2002), however, depending on the problem instance, the resulting solution solving process can easily consume the memory of a computer or the software used. In addition, many applications of mixed integer programming assume that a solution is optimal when it is within some pre-defined optimality gap, only searching for solutions with objective function values within some percentage of the best one stored in memory. Raising the optimality gap shortens the computation time, reducing the gap increases computation time as well as memory required. Using a similar sized landscape as our smaller problem, yet with only one planning period, Murray and Weintraub (2002) needed over 60,000 constraints to fully specify the problem, and when solving it with a mixed integer solver, stopped the process when the tolerance gap was about 15%. Based on this guidance, and based on the fact that our problems are larger than those tested in previous research (management units x prescriptions available), we chose to compare our heuristic solutions to results from a relaxed linear programming formulation. If we had attempted to solve exactly a mixed integer problem formulation, we would have found similar mixed integer results using a tolerance of 1 to 3%, far below the tolerance used by Murray and Weintraub (2002).

4 Discussion

Over the past 15 years, researchers have been exploring alternatives to traditional mathematical programming search processes (i.e., linear programming, mixed integer programming) for solving spatial harvest scheduling forest planning problems. One avenue of research has been the development of testing of heuristic search methods. A number of heuristics have shown promise for addressing the development of spatially complex forest management plans. These heuristics can produce near-optimal solutions very quickly (in most cases). However, the main limitation of many heuristics is the extensive testing of parameters which must be performed prior to their use. Bettinger and Zhu (2006) introduced a new heuristic that requires very limited testing, and has shown to be superior to other heuristics on small problems, where the objective is to maximize ever-flow of timber harvest volume subject to URM constraints. When this search process was modified to address larger problems, and problems with more typical southeastern U.S. character (e.g., maximizing net present value, subject to a 97 ha maximum clearcut size modeled using ARM methods), the heuristic was as good as or better than, other standard heuristics in most problems that were examined here.

The raindrop method needed to be modified to accommodate the modeling of ARM adjacency constraints. Since this model uses a spatially-sprawling method to determine all of the sizes of all of the clearcuts, within their respective green-up windows, the preferred method of mitigating the constraint violation was to unschedule the clearcuts of some units until the constraints were again satisfied (i.e., not violated). This modification may have led to suboptimal results, since as Bettinger and Zhu (2006) suggest, the next best choice for an affected unit should be selected. However, to explore other options for affected units (those that are a part of adjacency violation) would require much more processing time, as the ARM adjacency constraints would need to be accessed repeatedly until they were no longer violated. In any event, the modified raindrop method produced as good, or better solution to most problems studied, although the time to generate a solution was much greater than that required by threshold accepting and tabu search.

Tabu search located the best solutions to both the small- and medium-sized hypothetical landowners that were assigned the young initial age class distribution. However, only in the medium-sized forest case were the results significantly different than the raindrop method. The ability of other heuristics to locate good solutions for forests with young age class distributions may therefore be better handled by standard 1-opt tabu search, particularly when larger landscapes are modeled. As a result, the mitigation of both ARM and wood flow constraint violations may be better handled by a deterministic process. Threshold accepting continued to find very good solutions to all problems, and if one were concerned about the time required to generate a solution to a complex forest planning problem, this method may be preferred. The size of the forest for which a spatially-constrained plan is being developed may also influence the performance of the heuristics. As we have shown here, the solutions for the medium-sized forest were closer to the linear programming results (percentage-wise). However, the absolute difference between the relaxed linear programming solution and the best heuristic solutions (regardless of heuristic algorithm) was about the same for both the small- and medium-sized problems. Therefore, the smaller problem seems to have more impact simply because the basis for analysis (the net present value) is inherently smaller.

The clear advantage of the raindrop method lies in solving spatially constrained forest planning problems where the spatial constraints may be relatively simple (i.e., the URM adjacency constraints). If a landowner were interested in acknowledging and incorporating these kinds of spatial constraints, this method may be preferred. In addition, only two parameters are required, and they do not need extensive tests to determine the appropriate values. However, to accommodate ARM and wood flow constraints with the raindrop method requires an extensive and complex set computer code, which is the main disadvantage of using this heuristic.

5 Conclusions

This modified raindrop method still shares the same core idea as the original model introduced in Bettinger and Zhu (2006), where a radiating, spatially-sprawling process is used to find high quality solutions to spatially-constrained forest planning problems. Some modification of the raindrop method is required to accommodate ARM and wood flow constraints, however. Just like the original raindrop method, the main advantage is that it only uses two parameters and it has been shown to generate as good, or higher quality results as compared to threshold accepting or 1-opt tabu search for the six hypothetical landowners examined here. The two disadvantages of this method are 1) the additional programming logic required to accommodate ARM and wood flow constraints, and 2) the additional time required to generate solutions to spatial harvest scheduling problems. One could conclude from this work that it may be advantageous to use this heuristic for problems with relatively simple spatial forest planning constraints, and problems that do not involve young initial age class distributions. Generalization of the performance of the raindrop method to other forest planning problems is problematic, and will require examination by those interested in pursuing this planning methodology. Based on our tests of the method to a limited set of URM and ARM formulations, one should view this set of work as additional insight into the potential performance of the method to problems of current interest to the forest planning community.

References

Bettinger, P. & Zhu, J. 2006. A new heuristic for solving spatially constrained forest planning problems based on mitigation of infeasibilities radiating outward from a forced choice. Silva Fennica 40: 315–333.

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

—, Sessions, J. & Johnson, K.N. 1998. Ensuring the compatibility of aquatic habitat and commodity production goals in eastern Oregon with a Tabu search procedure. Forest Science 44: 96–112.

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

—, Johnson, D.L. & Johnson, K.N. 2003. Spatial forest plan development with ecological and economic goals. Ecological Modelling 169: 215–236.

Bevers, M. & Hof, J. 1999. Spatially optimizing wildlife habitat edge effects in forest management linear and mixed-integer programs. Forest Science 45: 249–258.

Boston, K. & Bettinger, P. 2001. The economic impact of green-up constraints in the southeastern United States. Forest Ecology and Management 145: 191–202.

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

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.

Dueck, G. & Scheuer, T. 1990. Threshold accepting: a general purpose optimization algorithm appearing superior to simulated annealing. Journal of Computational Physics 90: 161–175.

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

Glover, F. 1989. Tabu search – Part I. ORSA Journal on Computing 1(3): 190–206.

— 1990. Tabu search – Part II. ORSA Journal on Computing 2(1): 4–32.

—, Kelly, J.P. & Laguna, M. 1995. Genetic algorithms and tabu search: hybrids for optimization. Computers & Operations Research 22: 111–134.

Hof, J.G. & Joyce, L.A. 1992. Spatial optimization for wildlife and timber in managed forest ecosystems. Forest Science 38: 489–508.

—, Bevers, M., Joyce, L. & Kent, B. 1994. An integer programming approach for spatially and temporally optimizing wildlife populations. Forest Science 40: 177–191.

Johnson, K.N. & Scheurman, H.L. 1977. Techniques for prescribing optimal harvest and investment under different objectives – discussion and synthesis. Forest Science Monograph 18.

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

McDill, M.E., Rebain, S.A. & Braze, J. 2002. Harvest scheduling with area-based adjacency constraints. Forest Science 48: 631–642.

Murray, A.T. 1999. Spatial restrictions in harvest scheduling. Forest Science 45: 45–52.

— & Church, R.L. 1995. Heuristic solution approaches to operational forest planning problems. OR Spektrum 17: 193–203.

— & Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. Forest Science 48: 779–789.

Nelson, J. & Brodie, J.D. 1990. Comparison of a random search algorithm and mixed integer programming for solving area-based forest plans. Canadian Journal of Forest Research 20: 934–942.

Timber Mart-South. 2004. Georgia stumpage prices. Center for Forest Business, University of Georgia, Athens, GA.

Zhu, J. 2006. Spatial harvest scheduling in the southeastern United States: Estimating the impact on landowners of different sizes and spatial configuration of ownership. PhD dissertation. Warnell School of Forestry and Natural Resources, University of Georgia, Athens, GA.

Total of 25 references