

Spatially constrained timber harvest scheduling

ANTHONY J. O'HARA

Forestry Commission of New South Wales, Sydney 2001, Australia

BRUCE H. FAALAND

School of Business Administration, University of Washington, Seattle, WA 98195, U.S.A.

AND

B. BRUCE BARE

*College of Forest Resources and Center for Quantitative Science in Forestry, Fisheries and Wildlife,
University of Washington, Seattle, WA 98195, U.S.A.*

Received October 19, 1988

Accepted February 6, 1989

O'HARA, A. J., FAALAND, B. H., and BARE, B. B. 1989. Spatially constrained timber harvest scheduling. *Can. J. For. Res.* 19: 715-724.

Multiple-use management of forests often requires imposition of spatial constraints on the selection of units for harvest. To satisfy such constraints, harvest units must be treated as integral units. A biased sampling search technique is used to find integer solutions to operationally sized problems. Solutions found for the sample problems are within 8% of the upper bound of the corresponding linear programming solution and less than 4% below the upper bound on the true optimum as defined by a confidence interval estimator.

O'HARA, A. J., FAALAND, B. H., et BARE, B. B. 1989. Spatially constrained timber harvest scheduling. *Can. J. For. Res.* 19 : 715-724.

L'aménagement polyvalent des forêts exige souvent l'imposition de contraintes spatiales portant sur le choix des unités pour la récolte. Pour satisfaire ces contraintes, les unités de récolte doivent être traitées comme des unités intégrales. Une technique d'échantillonnage par biais a été utilisée pour trouver des solutions intégrées à des problèmes aux dimensions opérationnelles. Les solutions élaborées pour les problèmes échantillonnés se trouvent à moins de 8% de la limite supérieure de la solution correspondant à la programmation linéaire et à moins de 4% sous la limite supérieure de l'optimum vrai tel que défini par un estimateur d'intervalle de confiance.

[Traduit par la revue]

Introduction

Whether a forest is intended for the long-term sustained yield production of timber in the presence of multiple-use considerations or as an industrial wood supply, the spatial relationship of harvest units has a major influence on effective management. Spatial constraints affect the timing of harvest of adjacent units of the forest (e.g., a unit cannot be harvested if an adjacent unit has been harvested in either the same period or in one or more preceding periods).

Explicit recognition of spatial relationships and the consequent constraints on timber harvest scheduling are important for a number of reasons (Bare et al. 1984). First, management for multiple use often requires that managers know the geographic location of specific outputs as well as how much of the output to produce in a given time period. Of equal concern is the scheduling for harvest of adjacent units in the same time period, where the combined area of the two units exceeds the statutory or policy limit on the maximum size of clear-cuts (Hokans 1983). Typical time lags range from 4 years in the redwood forests of California to 30-50 years in such forests as the Shoshone National Forest in Wyoming.

Failure to recognize spatial relationships can also result in environmental problems. Where a buffer strip is to be retained, such as around streams, lakes, recreation areas, roads, or wilderness areas, or where the buffer strip is necessary for regeneration purposes, the spatial relationship between the timber scheduled for harvest and the adjacent stand or feature needs to be recognized.

The spatial relationship between harvest units must be recognized so that specific wildlife habitat objectives, such

as maintenance of adequate degrees of habitat diversity, can be met. For example, Mealey et al. (1982) state that USDA National Forest planners are encouraged to establish wildlife habitat objectives that specify threshold and most desirable levels of specific proportions of forest age-classes that are spatially distributed within the geographic area of the forest, to provide adequate wildlife resources. They further state that while planning efforts to date have been relatively successful in providing adequate quantities of needed habitat, the habitat dispersion necessary for providing cover and edges has not been very successfully met.

Literature review

Although spatial constraints play a major role in forest management planning, literature that discusses the incorporation of spatial constraints into timber harvest scheduling is limited. Concern within the USDA Forest Service, for example, has primarily been with forest-level calculations of the sustainability of harvest levels over multiple rotations from a biological standpoint, rather than with the spatial implications of timber harvesting (Johnson 1981). Such an approach tends to overstate timber harvest capability when additional multiple-use objectives with spatial implications, such as watershed, soil, recreation, visual, and wildlife resources, must be met (Mealey et al. 1982).

The most widely used technique for timber harvest scheduling in the United States is linear programming (LP). One of the earliest models was Timber RAM (resource allocation model (Navon 1971)). Another widely used LP harvest scheduling model, developed for even-aged industrial

forests, was MAXMILLION (Ware and Clutter 1971). No spatial considerations were included in either model.

In an attempt to deal more effectively with site-specific environmental questions, the MUSYC model of Johnson and Jones (1979) was developed. However, this model was basically a more sophisticated timber management model (Iverson and Alston 1986) in which location-specific issues could not be easily addressed. The wholesale revision of MUSYC to FORPLAN created an opportunity for incorporation of spatial constraints, particularly with Version II (Stuart and Johnson 1985), although this requires a specific approach to formulation.

Armell (1986) noted that one of the most frequently asked questions by forest managers concerns how the allocations represented by the standard, stratum-based FORPLAN solution, in which homogeneous forest units are aggregated, can be implemented within a heterogeneous area represented by a given parcel of national forest land. Complications result from the specific placement and management of habitats for wildlife such as spotted owls or pileated woodpeckers, and from the need to consider harvest adjacency constraints. The standard FORPLAN solution does not consider these factors. Approaching the problem by means of the "coordinated allocation choices" option in FORPLAN Version II does not lead to a satisfactory solution to this problem, because resolution of the problem at the harvest-unit level produces a problem of unmanageable size.

Spatial constraints have been explicitly incorporated into LP harvest scheduling models (e.g., Mealey et al. (1982) and Thompson et al. (1973)). However, none of these approaches led to integer solutions. Jones et al. (1986) examined the combined transportation network - harvest scheduling problem with spatial constraints on harvest scheduling, but did not attempt to solve the harvest scheduling portion of the problem to an integer solution. Furthermore, units were individually identified by time period only for the first two decades of a five-decade analysis, and were aggregated within the last three periods. Finally, no harvest flow constraints were imposed.

Recent studies by Jones and Meneghin (1987),¹ Meneghin et al. (1988), and Weintraub et al. (1988)² also address the problem of incorporating spatial integrity into forest planning models. While Jones and Meneghin (1987)¹ and Meneghin et al. (1988) are primarily concerned with the joint transportation - harvest scheduling problem, Weintraub et al. (1988)² are concerned with habitat dispersion. While the former introduces alternative formulations aimed at reducing the number of constraints, the latter focuses on a column-generation technique coupled with a heuristic approach for solving the associated subproblem. Sessions and Sessions (1988) recently developed a heuristic algorithm (SNAP) for solving the combined harvest scheduling - transportation problem in the presence of spatial constraints. Their model schedules for up to three periods and

accepts capacity limitations on the branches of the road network. However, the model is not designed to solve long-term multiple-use forest planning problems.

A problem shared by all LP approaches to solving the spatially constrained timber harvest scheduling problem is that the solutions found are not integral. Most commonly, units are split to meet the spatial constraints. In a mathematical sense, the constraints are met, but in practical terms, field implementation of the solution is not possible unless the solution is integral.

Integer solutions to the problem were found by Hokans (1983, 1984), using a discriminant function calculated from variables used by the manager to define a spatially feasible harvest schedule for a subunit of the forest, in combination with a computerized grid-based geographic data base of stands. The procedure was developed specifically to ensure that policy limits on clear-cut size were not exceeded.

There are few spatially oriented timber harvest scheduling models in general use that employ integer programming (IP) as a solution procedure. Probably the most widely recognized model is that of Kirby et al. (1986). Their integrated resources planning model (IRPM), developed as a mixed-integer programming formulation, is capable of solving modest-sized problems (i.e., less than 50 integer variables and a total of 500 rows and columns), and deals with the joint transportation - multiple-use planning problem. The model has been implemented for a variety of national forests in the western United States at both the forest and subforest levels of resolution. However, as Jones et al. (1986) report, the IP solution procedure must be abandoned in favor of a heuristic algorithm when solving most operationally sized planning problems.

Bare et al. (1984) used IP as a solution procedure for small problems in a research context. Solutions were found to harvest scheduling problems for stands with up to 40 harvest units over five periods, together with spatial constraints that required at least a two-period lag between harvests of adjacent units. Solution time for such a problem was over 3 h on a VAX 11/780.

Bare et al. (1984) also developed a dynamic programming approach to solve problems of similar size. Although their formulation was capable of finding the true optimum for small problems, it required substantial computer storage, which limited its use to problems involving approximately 50 harvest units over five time periods.

Model development

This paper describes a biased sampling search technique for the spatially constrained harvest scheduling problem. In the development of the model, it is presumed that a contiguous planning area has been identified to contain a fixed number of prespecified harvest units. These units may assume any geometrical shape, and typically range in size from less than 6 acres (2.5 ha) to more than 40 acres (16 ha). Furthermore, it is assumed that no harvest unit can exceed a maximum specified size. For example, if a policy requires that the area clear-cut cannot exceed 80 acres (32 ha), no individual harvest unit can exceed this upper limitation.

The integer (binary) decision variables (X_{ij}) defined later may represent clear-cut or partial-cut prescriptions, although in the sample problems discussed later, only the former are considered. However, depending upon the severity of impacts and the resources considered, partial-cut prescrip-

¹J.G. Jones and B. Meneghin. 1987. Some spatial relationships for use in mathematical programming formulations for simultaneously analyzing forest management and transportation alternatives. Mimeograph report. USDA Forest Service, Intermountain Research Station, Ogden, UT.

²A. Weintraub, R. Epstein, and F. Barahona. 1988. Integrating the habitat dispersion problem with timber management planning. Mimeograph report. Industrial Engineering Department, University of Chile, Santiago.

tions are easily incorporated into the model. As no attempt is made to treat road-building activities, it is presumed that all harvest units are available for harvesting each period unless previously harvested or spatially constrained. Thus, the model is not intended to solve the joint transportation - harvest scheduling problem. Instead, its sole purpose is to solve the harvest scheduling problem in the presence of spatial constraints. The following notation is used in defining the model:

$X_{ij} = 1$ if unit j is selected for harvest in period i
0 otherwise

C_{ij} is the volume of harvest unit j in period i

l_i is the lower bound on total harvest volume in period i

u_i is the upper bound on total harvest volume in period i

I is the number of periods in the planning horizon

J is the number of forest units

S_j is the set of all units adjacent to unit j

The basic harvest scheduling problem used in this paper is a model I formulation (Johnson and Scheurman 1977), as follows:

maximize

$$\sum_{i=1}^I \sum_{j=1}^J C_{ij} X_{ij}$$

subject to

$$\sum_{j=1}^J C_{ij} X_{ij} \geq l_i, \quad i = 1, \dots, I \quad \text{lower bound}$$

$$\sum_{j=1}^J C_{ij} X_{ij} \leq u_i, \quad i = 1, \dots, I \quad \text{upper bound}$$

$$\sum_{i=i'}^{i''} (X_{ij} + X_{ij'}) \leq 1, \quad i' = i^*, \dots, I - (i'' - i^*) \text{ spatial constraints}$$

In the spatial constraints, i^* denotes the first and i'' the last period for which a given spatial restriction is in force, $i^{\#}$ denotes the number of periods that the restriction applies, and j' denotes each unit in turn from the set S_j .

Because, as is often the case, a unit can be harvested only once within the planning horizon, the following constraints are added:

$$\sum_{i=1}^I X_{ij} \leq 1, \quad j = 1, \dots, J$$

To illustrate the nature of the spatial constraints in terms of this formulation, the equations for imposing the restriction that adjacent units 1 and 2 cannot be harvested unless there are at least two time periods between harvests are as follows:

$$\begin{aligned} X_{11} + X_{21} + X_{12} + X_{22} &\leq 1 \\ X_{21} + X_{31} + X_{22} + X_{32} &\leq 1 \\ X_{31} + X_{41} + X_{32} + X_{42} &\leq 1 \\ X_{41} + X_{51} + X_{42} + X_{52} &\leq 1 \end{aligned}$$

A wide variety of spatial restrictions can be modelled using this formulation. Restrictions on the total number of con-

tiguous acres (hectares) clear-cut are met in two ways. First, as previously stated, individual harvest units are limited to some prespecified maximum size. Second, a sufficient time lag between harvests in adjacent units is imposed so that openings created by the clear-cutting of a unit can regenerate to the point at which they are no longer considered openings. By applying such restrictions to all adjacent units that, if harvested at the same time, would exceed the statutory or policy limit on total clear-cut size, spatially feasible harvest plans can be developed. Thus, it is possible to allow for adjacent units to be harvested in the same time period, so long as the total area of the opening does not exceed the policy limit on maximum clear-cut size. However, in the sample problems discussed later, harvesting adjacent units (no matter how small or large) is not permitted.

If units selected for harvest are adjacent to buffer strips around streams, lakes, or recreation areas, an appropriate time lag is imposed before the buffer unit is eligible for harvest. Control of visual quality may require a restriction on the number of periods between harvests of adjacent units, or the restriction that no more than a maximum number of units be harvested each period. Forest and wildlife habitat diversity objectives may be met by ascertaining the age-class distribution of the forest and imposing time lags between harvests of adjacent units, so that the age-class distribution is perpetuated or brought closer to a desirable distribution.

Some multiple-use objectives require modification of the basic formulation. If units are not to be harvested at all, they can be excluded from the analysis. Buffer strips are often retained indefinitely, but may be made available for harvest following a prescribed number of periods after harvest of adjacent units.

Although the objective of this formulation is to maximize total harvest volume, there is no difficulty in modifying the algorithm to set the objective function as maximization of net present value. And as demonstrated later, solutions found under this objective differ little from those found under volume maximization.

Spatially constrained resource allocation model (SCRAM)

General algorithm

The enormous number of feasible solutions associated with anything but a very small harvest scheduling model prompted the abandonment of attempts to find the true optimum with the use of either IP or an exhaustive direct-search algorithm. Instead, a heuristic algorithm capable of finding good, feasible solutions for problems of all sizes was developed (O'Hara 1987). It was recognized that the spatially constrained harvest scheduling problem possesses many characteristics similar to those displayed by assembly line balancing problems and that both problems involve difficulties in finding exact solutions (owing to their combinatorial nature). This aided in the development of the heuristic algorithm and led to a solution procedure capable of solving the spatially constrained harvest scheduling problem to an integer (not necessarily optimal) solution.

While many assembly line balancing algorithms rely on characteristics peculiar to a particular problem, Arcus (1966) and Tonge (1965) describe a general Monte Carlo simulation approach (i.e., a random search) to obtain good solutions to several assembly line balancing problems. Their results were further improved by prebiasing the random allocation of tasks to favor tasks considered to be more likely to pro-

duce a better solution. Arcus (1966) demonstrated that very good solutions are possible (and were found) if the number of solutions generated was large.

Random search methods have several advantages, as noted by Karnopp (1963): ease of programming, simplicity and low cost of storage and operation, insensitivity to the type of objective function (especially valuable if the function is discontinuous), efficiency (in both search procedure and setup time), flexibility, and the ability to provide information about the function being searched and to use that information to direct the search. The use of prebiasing is a relatively simple method for favoring better solutions and (or) reducing computation times.

One of the characteristics of the set of feasible solutions to small harvest scheduling problems that suggested that a random search algorithm might be successful in finding good feasible solutions is that there appear to be a reasonable number of solutions that are relatively close to the optimum solution.

The random search algorithm developed for this problem generates combinations of units for harvest over the desired planning horizon, so that both harvest flow and spatial constraints are met. The procedure generates 100 feasible combinations over all periods and selects the best of these as the final solution (the reasons for generating 100 feasible solutions are given later). An estimate of the true optimum is given, along with a $100(1 - e^{-100})\%$ confidence interval for the optimum. The program was written in FORTRAN 77 and includes a random number generator from Wichmann and Hill (1982). A flow chart and listing of the program code is given by O'Hara (1987).

Data requirements include a matrix of volume and value by time period for each unit, and for each unit, a list of the units adjacent to it and the number of periods for which the corresponding spatial restriction is to apply. Finally, minimum and maximum harvest levels for each period are specified. With suitable modification of array size, very large problems can be processed. There are no restrictions on the number of periods for which adjacent units must lag between harvests.

A unit is selected at random for inclusion in the combination of units for period 1 (prebiasing techniques that modify this random selection are discussed later). All units adjacent to the selected unit are flagged for exclusion from consideration, the flag denoting the period at which the unit is again eligible for inclusion in the harvest. A second unit is selected at random from the remaining eligible units. This process continues until the minimum harvest flow constraint is satisfied.

Generation of a feasible combination of units for the second period follows. The first step within the new period is to flag for exclusion any unit with zero volume, and to remove the flag for exclusion from any unit that has met the time lag for a spatial restriction imposed by harvest of an adjacent unit in an earlier period. From this point, generation of a feasible combination of units for the current period continues in the same way as described for the first period.

Processing continues in this manner until a feasible combination of units is generated for each period over the entire horizon. The total volume (value) of the combination is compared with the incumbent solution. The current solution

replaces the incumbent solution if the total volume (value) exceeds the incumbent total volume (value).

After 100 feasible solutions are generated and the best selected as the final solution, a $100(1 - e^{-100})\%$ confidence interval about the optimum is calculated, as described later. The decision to cease processing after the generation of 100 feasible solutions is somewhat arbitrary, but was made on the basis that Arcus (1966) found 100 samples to be a reasonable size, that Zanakis (1977, 1979a, 1979b) found that 100 samples gave reasonable estimators for the confidence interval, and that limits needed to be placed on utilization of computer resources.

Prebiasing techniques

In addition to a random selection of units, an attempt was made to produce better solutions by prebiasing the selection of units to favor units considered more likely to yield good solutions. Three techniques were used.

As the objective function used in the analysis of the random search algorithm is volume maximization, the first prebiasing technique selects units by volume. By replacing volume by value, prebiasing by value can be used if maximization of net present value is the objective. To implement this process, the total volume of all units available for inclusion is determined. The probability of selecting a given unit is set equal to the proportion of total volume contributed by that unit. Hence, units with larger volume have a greater probability of being selected than units with smaller volume.

After some experimentation, prebiasing by volume was applied for the first 3 periods, with no prebiasing used in later periods for problems over 5 periods, and for the first 7 periods for problems over 9 or 10 periods. The rationale for this approach was that in attempts to meet harvest flow constraints, higher volume units are more important early in the planning horizon, when average unit volume is relatively low, than later in the planning horizon, when volume growth enables harvest flow constraints to be met more easily (i.e., with fewer units). Prebiasing over the full planning horizon tends to leave only the smallest volume units in the last period. Because volume growth diminishes as units age, selection amongst these smaller volume units in the last period may yield lower total volumes than if selection is random over the final periods of the planning horizon. Clearly, these empirically derived conclusions only apply to the data set and problem definition employed. However, similar rules of thumb can be easily derived for other data sets.

The second prebiasing technique biases selection in favor of units with the fewest effective adjacent units. A unit is considered to be an effective adjacent unit if selection of the unit for harvest means that the adjacent unit is added to the list of units excluded from harvest (i.e., the adjacent unit is not already on the list). The probability of selection is the inverse of the number of effective adjacent units.

The third prebiasing technique is a combination of the first two techniques. Selection is biased in favor of units with the fewest effective adjacent units, and among the units so chosen, selection is biased in favor of units with larger volume.

Confidence interval estimation

The quality of the solution found by a heuristic random

search algorithm is greatly improved if some estimate is possible of the proximity of the best solution to the true optimum. Golden and Alt (1979) propose a very simple $100(1 - e^{-n})\%$ confidence interval on the true optimum, Z_{opt} . For a minimization problem this is written as

$$Z_1 \geq Z_{opt} \geq Z_1 - \hat{b}$$

where $\hat{b} = Z_{(0.63n)+1} - \hat{a}$, n is the number of solutions in the sample, parentheses denote rounding down to the next integer, and \hat{a} is an estimate of the true minimum (Z_{opt}). If a random sample of iterative solutions to a large minimization or maximization problem is ordered, so that the smallest is set as Z_1 (for a maximization problem, the largest solution is made the smallest by multiplying by -1), then using a procedure described by Cooke (1979), Z_{opt} is estimated as

$$\hat{a} = 2Z_1 - (e - 1) \sum_{i=1}^n Z_i / e^i$$

This estimator was developed from the empirical distribution function for the set of ordered solutions Z_1, \dots, Z_n without making any assumptions about the distribution of the solutions. The estimator applies to any distribution with a location or threshold parameter (Zanakis and Evans 1981).

The confidence interval estimator requires that the random sample of ordered solutions come from a three-parameter Weibull cumulative distribution:

$$F(Z) = 1 - e^{-((Z - a)/b)^c}, \quad Z \geq a$$

with parameters a (location), b (scale), and c (shape) (Zanakis and Evans 1981). Golden and Alt (1979) note that the assumption is derived from a proof by Fisher and Tippet (1928) that for S independent samples, each of size m , from a parent population that is bounded from below by a , the distribution of the smallest value in each sample approaches a Weibull distribution with a as the location parameter as m gets large. The estimator is then derived from the observation that the cumulative distribution function of the sum of the location and scale parameters equals $1 - e^{-1}$, or approximately 0.63. Using this fact, the confidence interval can be derived by a simple manipulation of a probability expression and the cumulative distribution function of the Weibull distribution. The estimator was tested with a variety of travelling salesman problems and found to perform efficiently.

Golden and Alt (1979) recommend refining this confidence interval by using maximum likelihood estimates, but Zanakis and Evans (1981) claim that the added accuracy does not justify the incremental effort and that the estimators given earlier should suffice for most simple heuristic optimization methods.

Data

Data for our tests are drawn primarily from the study by Jones et al. (1986). These data cover 242 units over five 10-year time periods and are from the Twin Rocks section of the Lolo National Forest in western Montana. The harvest units are delineated on the basis of relative uniformity of species, age, and condition, in combination with geographic boundaries, and range in size from 6 to 40 acres (2.5–16 ha) (40 acres is the maximum clear-cut size). Volume and value by unit and period are known, with value recorded in con-

stant 1982 dollars and future prices calculated assuming a growth in real prices of 2% per year.

As no harvest flow constraints were included in the study by Jones et al. (1986), a variation of SCRAM, which ensures that all feasible units are harvested, is used to establish lower bounds on harvest volume under a given set of spatial constraints. With this estimate as a starting point, the lower bound on harvest volume for a period is set at the maximum level for which 100 feasible solutions are found within 15 min of central processing unit (CPU) time on a VAX 8700. Different trials are run in increments of 1000 board ft (1 board ft = 2.360 dm³) to aid in this determination. The upper bound is arbitrarily set 1000 board ft higher for the 242-unit problem. The width of the bounds is of the order of 5–10% of the lower bound. As SCRAM adds units to the harvest in a period only until the lower bound is met, and then adds units to the harvest in the next period, the flow constraints are set no wider than is necessary to avoid, on average, excluding a unit from being harvested because the volume is so large that adding it to the current combination of units will cause the upper harvest bound to be exceeded. For ease of application and interpretation, harvest flow constraints are kept constant over the planning horizon, although SCRAM does not require this.

While there is nothing inherent in SCRAM that requires that a harvest unit can only be harvested once during the planning horizon, this constraint was added to the sample problems discussed later. Similarly, the only silvicultural prescription incorporated into the sample data set involves a final clear-cut. Again, SCRAM permits other silvicultural prescriptions if problem definition dictates this.

Results

In analyses with the SCRAM model, problems covering 25 units over five periods with a three-period lag between harvest of adjacent units are also reported. These were selected because they could always be solved by complete enumeration and, hence, provide a convenient benchmark for comparison.

Results using SCRAM with the Jones et al. (1986) data set are given in Table 1. Interpretation of the percent increase to the upper bound of the confidence interval and to the LP optimum requires some explanation. The upper bound of the confidence interval is a sample-based estimate of the maximum amount by which the true optimum differs from the best solution found, with essentially 100% probability. Because the estimators for the confidence interval are sample based, a solution with a narrower confidence interval is not necessarily better than one with a wider confidence interval. To reinforce confidence in the performance of the algorithm, the percent difference between the best solution found and the LP optimum, which is the absolute upper bound on the solution, is also tabulated. The LP optimum is obtained by excluding the requirement that X_{ij} be integer (binary) variables. Otherwise, the problem formulation remains unchanged.

For the 25-unit problem, prebiasing by volume yields the fastest solution. The best solution, which was actually the optimum, was found by prebiasing by the effective number of adjacent units. This procedure took about 140% longer than the next slowest procedure. The procedure using no prebiasing produced the lowest solution, 2.44% less than

TABLE 1. SCRAM with the data set from Jones et al. (1986)

Data ^a	Prebiasing ^b	Best solution (10 ⁶ bd ft)	Value (\$1000)	% to upper bound ^c		CPU ^d (s)
				1	2	
25/5/3						
1150/1600	Direct search	6 961	2 930	Irrelevant	14.93	124.7
	None	6 795 (-2.44)	2 918	7.71	17.73	288.1
	Volume	6 858 (-1.50)	2 910	8.57	16.65	192.6
	Adjacent	6 961 (0.00)	2 930	10.05	14.93	733.3
	Adjacent/volume	6 890 (-1.03)	2 886	7.22	16.11	397.8
242/5/1						
16 000 / 17 000	None	82 325	34 744	1.66	3.25	89.0
	Volume	82 074	34 511	1.40	3.57	256.4
	Adjacent	82 161	34 571	na	3.46	+ + 16
	Adjacent/volume	81 473	34 462	na	4.33	+ + 14
242/5/3						
8500/9500	None	44 557	18 846	2.41	6.61	428.2
	Volume	45 015	18 899	3.29	5.52	92.2
	Adjacent	45 089	19 097	3.46	5.35	700.0
	Adjacent/volume	44 859	19 038	3.17	5.89	192.2
242/5/V						
10 500 / 11 500	None	54 419	23 245	1.77	5.66	248.4
	Volume	54 601	23 180	2.09	5.31	133.4
	Adjacent	54 460	23 123	na	5.58	+ + 54
	Adjacent/volume	54 554	23 146	na	5.40	+ + 85

^avvv/w/x: vvv is number of units; w is number of periods; x is number of periods for which spatial restriction applies, where V denotes variable from 1 to 4, selected at random before construction of data set. yyyyy/zzzzz: yyyyy is lower bound on harvest flow (Mbf); zzzzz is upper bound on harvest flow (Mbf).

^bDirect search, uses direct search algorithm employing complete enumeration; none, no prebiasing used; volume, prebiasing by volume for first three periods, no prebiasing for later periods; adjacent, prebiasing by effective number of adjacent units for first three periods for one-period spatial restrictions and all periods for other spatial restrictions; adjacent/volume, prebiasing by effective number of adjacent units and then by volume for first three periods for one- and five-period spatial restrictions and all periods for three-period spatial restrictions.

^cPercentage to upper bound has two entries: (1) the percentage difference between the best solution found and the upper bound of the $100(1 - e^{-n})\%$ confidence interval; (2) the percentage difference between the best solution found and the LP optimal solution. For the 25-unit data set, the percent difference between the best solution found and the true optimum is given in parentheses in the best solution data column.

^dCPU s on a VAX 8700. If entered as + + nn, the algorithm failed to generate 100 feasible solutions within 15 min of CPU time. When this occurred, no estimate could be made of the optimum or the confidence interval. Hence, these are tabulated as na (not available). The number nn after + + is the number of solutions found within 900 CPU s.

the true optimum. However, the value of the solution was close to the best found by SCRAM with any other procedure.

The $100(1 - e^{-n})\%$ confidence interval appears to be quite conservative for the small data set, as would be expected for such a high probability. The confidence interval contains the true optimum in all cases, and the upper bound on the confidence interval exceeds the true optimum by at least 5% in all cases.

The performance of SCRAM using each prebiasing technique with the larger data sets varies with the restrictiveness of the spatial restrictions. SCRAM without prebiasing yields the best solution when only a one-period lag between harvests is imposed, and does so more than 280% faster than the next fastest procedure, which is prebiasing by volume. The percent increase for the confidence interval indicates that the solution found is within 1.66% of this upper bound on the true optimum with (essentially) 100% probability. Given the limited spatial restriction, it is not surprising that the procedures that used the effective number of adjacent

units require considerably more computation time than the other procedures.

For the more restrictive data set, in which a three-period lag is imposed, prebiasing by the effective number of adjacent units yields the highest overall solution (no more than 3.46% below the estimated upper bound on the true optimum). Prebiasing by volume is the fastest procedure with this data set and yields the second highest solution. The improvement in performance of prebiasing by effective number of adjacent units and then by volume compared with results with the less restrictive data set is considerable, and indicates that the motivation for the prebiasing technique is well-founded for certain types of problem.

Prebiasing by volume also produces the best solution with the fastest computation time for the data set in which the periods for spatial restrictions are variable. It produces a solution slightly higher than searching with no prebiasing, but approximately 85% faster. The solution is no more than 2.09% below the estimated upper bound on the true opti-

TABLE 2. SCRAM with the data set from Jones et al. (1986) and rotation harvest flow control

Data	Prebiasing	Best solution (10 ⁶ bd ft)	Value (\$1000)	% to upper bound		CPU (s)
				1	2	
25/5/3						
1000/1600	Direct search	6 983	2 924	Irrelevant	14.56	672.0
	None	6 499 (-7.45)	2 788	11.48	23.10	10.2
	Volume	6 538 (-6.81)	2 794	12.63	22.36	8.4
	Adjacent	6 577 (-6.17)	2 795	14.00	21.64	29.1
	Adjacent/volume	6 516 (-7.17)	2 796	11.79	22.77	18.6
242/5/1						
10 000 / 11 000	None	52 554	22 172	2.73	4.65	15.4
	Volume	52 377	22 136	2.28	5.01	23.4
	Adjacent	52 485	22 061	2.48	4.79	71.6
	Adjacent/volume	52 405	22 225	2.49	4.95	64.9
242/5/3						
8000/9000	None	42 184	17 887	2.69	6.68	34.6
	Volume	42 380	17 946	2.97	6.18	29.4
	Adjacent	41 998	17 880	2.21	7.15	135.4
	Adjacent/volume	42 669	17 800	3.65	5.46	71.1
242/5/V						
9 000 / 10 000	None	47 237	19 826	2.42	5.85	15.2
	Volume	47 886	20 322	3.67	4.41	20.5
	Adjacent	47 669	20 195	3.26	4.89	68.6
	Adjacent/volume	47 433	20 014	2.62	5.41	60.4

NOTE: For an explanation of parameters and codes used, see Table 1.

mum. Although neither of the techniques using prebiasing by effective number of adjacent units produces 100 feasible solutions, the best solutions found are of high quality. Of course, only one solution is necessary if it is better than any other, although a confidence interval cannot be estimated for that case.

Solution times are relatively lengthy for the Jones et al. (1986) data set under harvest flow constraints that attempt to harvest close to the maximum available volume over the planning horizon. Such a scenario is not often used in practice, because attempts to meet sustained yield requirements necessitate retention of potential harvest volume over a longer time horizon. To model such a situation, volume and value for the Jones et al. (1986) data set are projected over a nine-period (90-year) rotation. For this planning horizon, harvest flow constraints are set to harvest the maximum available volume, subject to the spatial constraints. The harvest flow constraints so determined are then applied to the original Jones data set. It is expected that these lower (and thus more easily met) harvest flow constraints will considerably improve the solution times. Results of this analysis are given in Table 2.

It is immediately obvious that the supposition that solution times will be dramatically improved for less restrictive harvest flow constraints is valid. The harvest flow constraints for the three-period and variable-period spatial restrictions are reduced by only 5.9 and 5.0%, respectively, but the solution times improve by factors of 3.14 and 6.51%, respectively, for the fastest techniques with these data sets and the original harvest flow constraints (see Table 1).

The other obvious difference between these results and those obtained with the original harvest flow constraints is that differences in solution times for different procedures are less dramatic. There is far less reason to consider a procedure to be at all deficient as a result of a markedly longer solution time. For practical application, this has two important advantages. First, more than one prebiasing technique can be used to obtain different solutions and perhaps increase the likelihood of finding a particularly good solution. Second, a larger number of feasible solutions can be found by a procedure without prohibitive use of computing resources, further increasing the chances of finding a better solution with any one procedure.

Discussion

General comments

SCRAM gives managers the first model capable of finding good, feasible integer solutions to timber harvest scheduling problems in the presence of spatial constraints for large numbers of units over a planning horizon. Before the development of this model, such scheduling was done manually, with a considerable manpower requirement and the ever-present risk of infeasibility.

SCRAM uses data that are available for any harvest scheduling operation. Hence, there are no special data requirements. The model can process schedules with spatial constraints in any combination desired by the user (assuming a feasible solution is possible), limited only by the requirement that a reasonable number of feasible solutions exist. This implies that the harvest flow constraints are not exces-

sively restrictive when considered with the spatial constraints. Moreover, estimates of confidence intervals for the optimum suggest that the best solutions found for large problems are all within 4% of the true optimum, with essentially 100% probability. Even assuming the worst case of the upper bound on the true optimum being the LP optimum, all solutions to large problems with moderate planning horizons are within 8% of this upper bound.

Units selected for harvest by the model are relatively randomly scattered throughout the forest. Such a result may be undesirable if close control over harvests is desired, as evidenced by the geographical confinement of harvesting operations. This scenario can be modelled by only considering for selection those units that are within the defined geographical area. However, this will require definition of harvest flow constraints for that area, and may lead to difficulties in meeting the spatial constraints over the forest as a whole, because the algorithm only considers the spatial relationship within the list of units being considered for selection.

The definition of "period" in the model has been left deliberately vague. Although one period equalled one decade for most data sets in this study, there is no intrinsic requirement in the model for such a period length. The length of a period is thus entirely at the discretion of the user. For example, if the only spatial restriction imposed was the requirement that the harvests of adjacent units be no closer than 5 years apart, the harvest schedule would be modelled with this algorithm by setting the period equal to 5 years and the spatial restriction equal to one period, and then running the model for the desired number of 5-year periods to meet the planning horizon. The only requirements are that the periods be of constant length and that the spatial constraints apply for integer multiples of the period length.

Unless the spatial restrictions are imposed for only one period, unharvested units will always remain at the end of the planning horizon (unharvested units may remain with a one period restriction, depending on how difficult it is to meet the harvest flow constraints). These unharvested units are the units adjacent to units harvested closer to the end of the planning horizon than the length of the spatial constraint, together with units not required to meet the harvest flow constraints. No provision is made for ensuring that sufficient units remain to meet sustained yield requirements beyond the end of the planning horizon. The number of units remaining is obviously variable. Hence, several runs of the model may be necessary to balance the number of units made available for harvest within the planning horizon with the number that must be retained to meet sustained yield volume requirements. Alternatively, the model can be run for a full rotation, although even then, spatial constraints may prevent the harvest of all units within the period of the rotation.

Significance of the harvest flow constraints

Harvest flow constraints have a significant influence on the operation of SCRAM. Recall that the algorithm passes on to selection for the next period as soon as the lower bound on the harvest flow is exceeded for the current period. Thus, it can be argued that the lower bound penalizes the algorithm and leads to dominated solutions, because a better solution may be found by increasing the lower bound and leaving the upper bound unchanged. One approach for problems

with loose lower bounds is to increase the lower bound by 10% increments and run the model for each new set of harvest flow constraints, until either no solutions can be found within what is considered a reasonable amount of computation time, or the lower bound on harvest flow has reached the highest acceptable level.

Prebiasing techniques

Particular types of problems respond to different forms of prebiasing (or no prebiasing at all) in characteristic ways. These results suggest that prebiasing by various methods should be used for certain problems, but random search without prebiasing performs best for other types of problems. Lastly, SCRAM (with or without prebiasing) yields solutions that are essentially equally good for many problems. However, thoughtful selection of which, if any, prebiasing technique should be used, and for what number of periods, can substantially reduce the computation time necessary for finding a good solution.

SCRAM without prebiasing is usually faster than any form of prebiasing when the harvest flow constraints are not very restrictive. Tables 1 and 2 show that for problems in which the spatial restrictions apply for only one period, random search without prebiasing is significantly faster than any prebiasing technique, and also yields the best or second-best solution for either data set (although the latter result may be at least partly due to chance). The reason for this is that imposition of a one-period lag between harvest of adjacent units leaves a large number of units available for selection for harvest in any period. Hence, it is not as important to select the "best" units for inclusion in a particular period in the harvest schedule.

Prebiasing by volume, effective number of adjacent units, or a combination of the two becomes more efficient under more restrictive spatial constraints. Prebiasing by volume yields 100 feasible solutions substantially more quickly than any other technique when a three-period spatial restriction is imposed, and yields the second-best solution for both analyses. Both the other prebiasing techniques work very well under the same constraint, although both require considerably more computation time than prebiasing by volume. The effectiveness of the prebiasing techniques under this constraint is not surprising, because the more restrictive situation favors techniques that facilitate the search for the "best" solution.

For the intermediate situation, in which the spatial constraints are moderately restrictive, no one technique is superior. Random search without prebiasing or with prebiasing by volume produces solutions significantly more quickly than prebiasing by either of the other techniques. However, even in the situation in which prebiasing by effective number of adjacent units produced only 54 solutions in 900 s of CPU time, the best solution found is better than that found by random search without prebiasing in 248.4 s of CPU time. Hence, CPU time alone does not provide a good measure of the quality of a procedure.

In general, prebiasing by volume yields solutions more quickly than any other technique for problems in which the spatial restriction is greater than a one-period lag. In doing this, the best solution found is usually either the best or second-best solution found by any technique. Random search without prebiasing is fastest and best or second best for all problems over planning horizons of five periods with

a one-period spatial constraint. The remaining techniques produce solutions that differ relatively little in quality but usually require considerably greater amounts of computer time.

Confidence interval on the true optimum

Results obtained with the estimator for the $100(1 - e^{-n})\%$ confidence interval suggest that the confidence interval is likely to include the true optimum in many cases. Furthermore, even in the worst case of the upper bound on the SCRAM solution being the LP optimum, the largest difference between a solution found by the random search procedures for large problems and the LP optimum is 7.15%; the usual difference is of the order of 4–6%.

A difference of at most 7.15% (more reasonably, less than 6%) between the best solution found and the LP optimum solution can be considered negligible. First, because the true (unknown) optimum is less than or equal to the LP optimum, we can be confident that a good, feasible solution has been obtained. Secondly, it raises the question of whether additional computing resources should be devoted to computing an even better solution. And thirdly, it must be recognized that when the solution is implemented in an operational setting, other sources of uncertainty will dominate, making it extremely unlikely that model-based timber volumes (values) will be realized. Such a conclusion is reinforced by the knowledge that solutions closer to the LP optimum are likely to be found if the approach described earlier for actual use of SCRAM is followed. For practical purposes, solutions found by SCRAM can be said to be satisfactorily close to the true optimum.

Conclusions

The problem of incorporating spatial constraints into the timber harvest scheduling problem can be expressed in a generalized mathematical formulation. With this formulation, explicit recognition is given to constraints on the spatial relationships between harvesting units. A wide variety of such constraints, drawn from multiple-use objectives, can be successfully represented via this formulation.

SCRAM uses a random search algorithm to find good, feasible solutions to operationally sized problems with a variety of spatial constraints, and is the first model capable of finding good, feasible integer solutions to such timber harvest scheduling problems. Furthermore, SCRAM, unlike IP, always finds feasible solutions. SCRAM also appears to yield solutions that are stable in terms of accuracy as problem size increases. The size of a problem for which a solution can be found has not been determined, but the algorithm requires only modification of array dimensions to deal with much larger problems. No difficulty is expected in producing solutions to problems covering several thousand units, such as may be encountered in practice.

If value rather than volume maximization is desired, the algorithm can be easily modified to meet that objective. Solutions found under this objective differed little from those found under volume maximization for the data sets examined.

In short, SCRAM appears to be immediately applicable to field use for developing harvest schedules for areas in which the road pattern has been built and only temporary logging roads are needed for timber harvest. The harvest schedules generated will meet harvest flow constraints, are high in

volume (value), and are feasible in the field for direct identification of a harvesting pattern that meets spatial restrictions on the harvest of units of the forest.

SCRAM can be used to schedule cultural treatments such as shelterwood harvests, heavy thinnings, and spraying for pest control. The necessary condition for using the model is that the treatment be applied to the unit as a whole at one time. The spatial constraints on such treatments are not as evident as for clear-cut harvesting, but may be relevant in certain situations.

The general approach developed in this study for scheduling harvests over time, subject to specific constraints, can also be applied to the problem of developing sequences of treatments to be applied to a unit over a planning horizon. A series of schedules for each unit can be developed, then selected by a model similar to SCRAM, and examined for feasibility, subject to specific constraints. The spatial relationships of such treatments over time may be of interest to management.

Another important assumption to recall is that roads to service the harvest units are presumed built. For second-growth forests that have already been roaded, this may not be a severe deficiency. However, if the road network has not been built, the harvest schedule will be heavily influenced by the cost of building the roads and the timing of construction. The joint solution of the harvest scheduling problem in the presence of spatial constraints and the road network planning problem is a task of considerable complexity and is an important area for future research in this branch of forest management planning.

Acknowledgment

The senior author acknowledges financial assistance from the Forestry Commission of New South Wales, the Fulbright Foundation, and the College of Forest Resources, University of Washington, Seattle.

- ARCUS, A.L. 1966. COMSOAL: a computer method of sequencing operations for assembly lines. *Int. J. Prod. Res.* 4(4): 259–277.
- ARMEL, N.B. 1986. Area analysis and Version II of FORPLAN. *In* Proceedings of the Workshop on Lessons From Using FORPLAN, Denver, CO, April 29 – May 1. USDA Forest Service, Land Management Planning Systems Section, Washington, DC.
- BARE, B.B., FAALAND, B.H., and GUPTA, I. 1984. Timber harvest scheduling in the presence of spatial constraints. Paper presented at the Joint National Meeting of the Institute of Management Sciences and the Operations Research Society of America, San Francisco, CA, May 14–16.
- COOKE, P. 1979. Statistical inference for bounds of random variables. *Biometrika*, 66: 367–374.
- FISHER, R., and TIPPETT, L. 1928. Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Proc. Cambridge Philos. Soc.* 24: 180–190.
- GOLDEN, B.L., and ALT, F.B. 1979. Interval estimation of a global optimum for large combinatorial problems. *Nav. Res. Logist.* 26: 69–77.
- HOKANS, R.H. 1983. Evaluating spatial feasibility of harvest schedules with simulated stand selection decisions. *J. For.* 81: 601–603, 613.
- . 1984. An artificial intelligence application to timber harvest schedule implementation. *Interfaces*, 14: 77–84.
- IVERSON, D.C., and ALSTON, R.M. 1986. The genesis of FORPLAN: a historical and analytical review of Forest Service planning models. USDA For. Serv. Gen. Tech. Rep. INT-214.

- JOHNSON, K.N. 1981. Timber activity scheduling on the national forests: the second revolution. *In* Forestry predictive models: problems in application. *Edited by* D.C. LeMaster, D.M. Baumgartner, and R.C. Chapman. Washington State University Press, Pullman.
- JOHNSON, K.N., and JONES, D.B. 1979. A user's guide to multiple use sustained yield resource scheduling calculation (MUSYC). USDA Forest Service, Timber Management, Fort Collins, CO.
- JOHNSON, K.N., and SCHEURMAN, H.L. 1977. Techniques for prescribing optimal timber harvest and investment under different objectives: discussion and synthesis. *For. Sci. Monogr.* No. 18.
- JONES, J.G., HYDE, F.C., III, and MEECHAM, M. 1986. Four analytical approaches for integrating land management and transportation planning on forest lands. *USDA For. Serv. Res. Rep. INT-86*.
- KARNOPP, D.C. 1963. Random search procedures for optimization problems. *Automatica*, 1: 111-121.
- KIRBY, M.W., HAGER, W.A., and WONG, P. 1986. Simultaneous planning of wildland management and transportation alternatives. *In* Systems analysis in forestry and forest industries. Vol. 21. *Edited by* M. Kallio, A.E. Andersson, R. Seppala, and A. Morgan. North-Holland/TIMS Studies in Management Science. Elsevier Scientific Publishing Co., New York. pp. 371-387.
- MEALEY, S.P., LISCOMB, J.F., and JOHNSON, K.N. 1982. Solving the habitat dispersion problem in forest planning. *In* Transactions of the 47th North American Wildlife and Natural Resources Conference, Portland, OR, 26-31 March 1982. *Edited by* K. Sabol. Wildlife Management Institute, Washington, DC. pp. 142-153.
- MENEGHIN, B.J., KIRBY, M.W., and JONES, J.G. 1988. An algorithm for writing adjacency constraints efficiently in linear programming models. *In* The 1988 Symposium on Systems Analysis in Forest Resources, Asilomar, CA, March 29 - April 1, 1988. USDA For. Serv. Gen. Tech. Rep. RM-161. pp. 46-53.
- NAVON, D. 1971. Timber RAM. USDA For. Serv. Res. Pap. PSW-70.
- O'HARA, A.J. 1987. Incorporating spatial constraints into the timber harvest scheduling problem. Ph.D. dissertation, University of Washington, Seattle.
- SESSIONS, J., and SESSIONS, J. 1988. Scheduling and network analysis program (SNAP): user's guide. Department of Forest Management, Oregon State University, Corvallis.
- STUART, T.W., and JOHNSON, K.N. 1985. FORPLAN version II: a tool for forest management planning. Paper presented at the Joint National Meeting of the Institute of Management Sciences and the Operations Research Society of America, Boston, MA, April 29 - May 1.
- THOMPSON, E.F., HALTERMAN, B.G., LYONS, T.S., and MILLER, R.L. 1973. Integrating timber and wildlife management planning. *For. Chron.* 49: 247-250.
- TONGE, F.M. 1965. Assembly-line balancing using probabilistic combinations of heuristics. *Manage. Sci.* 11: 727-735.
- WARE, G.O., and CLUTTER, J.L. 1971. A mathematical programming system for management of industrial forests. *For. Sci.* 17: 428-445.
- WICHMANN, B.A., and HILL, I.D. 1982. Algorithm AS 183: an efficient and portable pseudo-random number generator. *Appl. Stat.* 31: 188-190.
- ZANAKIS, S.H. 1977. Computational experience with some non-linear optimization algorithms in deriving maximum likelihood estimates for the three-parameter Weibull distribution. *In* Algorithmic methods in probability. Vol. 7. *Edited by* Marcel F. Neuts. North-Holland/TIMS Studies in Management Sciences. pp. 63-77.
- _____. 1979a. A simulation study of some simple estimators for the three-parameter Weibull distribution. *J. Stat. Comput. Simul.* 9: 101-116.
- _____. 1979b. Extended pattern search with transformations for the three-parameter Weibull MLE problem. *Manage. Sci.* 25: 1149-1161.
- ZANAKIS, S.H., and EVANS, J.R. 1981. Heuristic "optimisation": why, when and how to use it. *Interfaces*, 11: 84-90.