Incorporating average and maximum area restrictions in harvest scheduling models

Alan T. Murray, Marcos Goycoolea, and Andrés Weintraub

Abstract: A major goal in natural resource management has long been balancing the multiple uses of forest lands. Timber harvesting remains an important component of natural resource utilization, but must be approached in such a way that recreational use, ecosystem dynamics, species survivability, and other considerations are not sacrificed. One way in which production impacts are mitigated in forest management is by restricting the spatial extent of harvesting activities in developed plans. Through the use of harvest scheduling optimization models, constraints can be structured and imposed to limit local area disturbance associated with harvesting to a stipulated maximum. This represents an approach for regulating impacts in an economically driven management setting. Harvest scheduling research has recognized the challenges in appropriately structuring maximum area restrictions in optimization models, but regulating average disturbance area size may also be desired. This paper develops a model formulation for imposing average and maximum area limits on local impacts in harvest scheduling that can be solved using exact techniques. Application results are presented that highlight the feasibility of this approach. Further, the associated tradeoffs that exist in modeling both average and maximum area restrictions simultaneously are illustrated.

Résumé : Depuis longtemps, un des principaux buts de la gestion des ressources naturelles est de maintenir l'équilibre entre les différentes utilisations de la forêt. Bien que la récolte des arbres demeure une composante importante de l'utilisation des ressources naturelles, elle doit être envisagée de façon telle que les activités récréatives, la dynamique des écosystèmes, la survie des espèces et bien d'autres aspects ne soient pas compromis. Une façon d'atténuer les impacts de la coupe consiste à limiter l'étendue spatiale des activités de récolte lors de la planification. À l'aide de modèles d'optimisation permettant l'ordonnancement des opérations de récolte, des contraintes peuvent être structurées et imposées, afin de limiter la superficie locale des perturbations dues à la récolte à une valeur maximale déterminée. Cette approche représente une façon de contrôler les impacts dans un environnement soumis à des prérogatives de nature économique. La recherche dans le domaine de la planification des opérations forestières reconnaît la difficulté de structurer adéquatement les contraintes de superficies dans les modèles d'optimisation. Mais il peut également être souhaitable de contrôler la superficie moyenne des perturbations. Cet article développe un modèle de planification des opérations pour imposer des valeurs limites moyennes et maximales des impacts locaux et pouvant être résolu par une approche exacte. Les résultats d'une mise à l'essai démontrent la faisabilité de l'approche. De plus, les compromis associés à la modélisation simultanée de contraintes de superficie maximale et moyenne sont illustrés.

[Traduit par la Rédaction]

Introduction

Forests were once considered so vast and plentiful, particularly in the United States, that no level of utilization of this natural resource would have been imagined as having a significant impact on inventories or the environment. This view no longer exists. It is recognized that past forest practices were myopic in the sense that local impacts and long-term sustainability issues were not sufficiently addressed to guarantee a continued resource base capable of supporting consumer demands (i.e., timber extraction, recreational use, and water quality, among others), as well as preserving native flora and fauna. As a result, forest management practices have turned to more extensive and more detailed planning and analysis to ensure the responsible use of our natural resources.

One aspect of current management involves the use of more spatially explicit harvest scheduling optimization models to support detailed operational planning. In doing this, viability, sustainability, and forest health considerations may be represented in applied harvest scheduling models (Barrett et al. 1998). Forest harvest scheduling is generally concerned with maximizing rates of return, balancing annual timber volume flows, ensuring road access to harvest areas, and minimizing the localized impacts of forest disturbance (Kirby et al. 1986). One of the most confounding issues has been this latter concern for limiting the spatial disturbance of harvesting and is the focus of this paper.

Timber harvesting remains a necessary forest activity, but needs to be done in a way that forest health is maintained. A

Received 22 January 2003. Accepted 8 September 2003. Published on the NRC Research Press Web site at http://cjfr.nrc.ca on 19 February 2004.

A.T. Murray.[1] Department of Geography, The Ohio State University, 1036 Derby Hall, 154 North Oval Mall, Columbus, OH 43210, U.S.A.
M. Goycoolea. Department of Industrial Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205, U.S.A.
A. Weintraub. Departamento de Ingenieria Industrial, Universidad de Chile, Casilla 2777, Santiago, Chile.

[1]Corresponding author (email: murray.308@osu.edu).

 doi: 10.1139/X03-217

mathematical approach for regulating local impacts was initially suggested in Thompson et al. (1973), where the extent of harvested areas was restricted. A fundamental challenge continues to be representing and imposing constraints for limiting the spatial impacts of harvesting in optimization models. The intent is to make certain that harvest activity in a contiguous area is restricted from exceeding a specified bound. This is typically referred to as a maximum area restriction. The need and desire to restrict spatial disturbance has become standard practice in the management of both public and private forests (Jones et al. 1991; Barrett et al. 1998; American Forest and Paper Association 2001).

U.S. National Forest management has long required the Forest Service to address local area spatial disturbance (see Thompson et al. 1973; Jones et al. 1991). A recent review and discussion of this topic is provided in Barrett et al. (1998), who note that regulations for clearcut maximums typically range between 40 and 120 acres (1 acre = 0.404 685 ha). Boston and Bettinger (2002) indicate that restrictions are expressly stipulated in Oregon (120 acres), California (40 acres), and Sweden (49 acres). Barrett and Gilless (2000) identify other states and countries having explicit spatial limitations. On a voluntary basis, the Sustainable Forestry Initiative requires participating private forests to be managed in a way that "[t]he average size of clearcut harvest areas shall not exceed 120 acres…" (American Forest and Paper Association 2001, p. 4). Making such standards operational in an optimization model has become a major research challenge (Murray 1999).

Consider the forest region shown in Fig. 1. Depicted are individual management units, each with unique inventory information (e.g., timber species, age class, available acres, volume, value, and slope). Each unit is labeled with a unique identification number (1–9) as well as the total available acres in brackets. Given a stipulated maximum area restriction and the area of each unit, the challenge is to mathematically structure a constraining condition that limits local impacts. If it is assumed that the maximum area restriction is 120 acres, then the units shown in Fig. 1 are obviously much smaller than the maximum restriction. As a result, it is possible to simultaneously harvest a number of neighboring units without violating the maximum area restriction. In fact, it is conceivable that up to five contiguous management units could be harvested in this case without violating the 120-acre maximum. This particular problem is now commonly referred to in the harvest scheduling literature as the area restriction model (ARM).

One question is whether the maximum is in fact a hard constraint. Is the actual limit really 120 acres in this case, or is it possible that some deviation above 120 acres is allowable? After all, the Sustainable Forestry Initiative discusses that the average size must not exceed 120 acres (American Forest and Paper Association 2001). This obviously means that the extent of disturbed areas is not limited to a specific maximum, but rather that when all areas are examined, they average less than or equal to an established size. Lockwood and Moore (1993), Van Deusen (1996), Hochbaum and Pathria (1997), and Murray (1999) discuss the possibility of exceeding such limits in harvest scheduling. All present approaches for relaxing the strict bounds typically imposed in optimization models. Rather than allow violations, Boston and Bettinger (2001) developed a heuristic procedure to account for average area disturbance limits, while at the same time restricting absolute maximums.

This paper presents an exact approach for addressing an average area bound as well as maintaining maximum area restrictions in a harvest scheduling model. The next section reviews previous harvest scheduling work associated with structuring and solving spatially restricted models. This is followed by a section detailing how feasible harvest blocks can be identified and subsequently used in modeling spatial restrictions. The paper then structures an area-based model that explicitly incorporates the notion of maintaining average area. Application results are presented that demonstrate that this formulation can be solved exactly using commercial optimization software for medium to large harvest scheduling problems. Further, the capability for assessing trade-offs associated with varying maximum area limits in relation to an average area restriction is illustrated.

Background

There has been a sustained interest in modeling-based approaches for assisting forest managers in the scheduling of timber harvests. One of the earliest contributions interested in regulating spatial impact explicitly was the work of Thompson et al. (1973). Subsequent work has focused on a range of approaches for imposing spatial restrictions. Murray (1999) identified the two most basic as the unit restriction model (URM) and ARM.

The URM has received considerable attention in past research (Thompson et al. 1973; Jones et al. 1991; Murray 1999). The URM assumes that blocks are structured so that harvesting any two neighboring blocks would violate a maximum area restriction. Block creation along these lines may be readily carried out using commercial geographic information system (GIS) software, as discussed in Barrett (1997). In contrast to the URM, the ARM does not presuppose any spatial structure on forest-reporting boundaries. In fact, it is anticipated that planning units may be much smaller than stipulated harvest area maximums when using the ARM. As a result, it is possible to harvest numerous neighboring planning areas without violating a maximum area restriction. Research devoted to the ARM has substantially increased in recent years. No doubt this is partly attributable to more detailed forest inventory information, accessible using GIS.

The ARM, where one is seeking to optimize economic return subject to maintaining maximum area restrictions, was precisely the concern detailed and solved heuristically in Hokans (1983). Exact solution techniques for the ARM have been somewhat elusive because of the combinatorial complexity of how harvest units could possibly be formed into blocks. That is, the problem was difficult to state mathematically and thought to be computationally prohibitive to solve by exact methods (Murray 1999). As a result, initial efforts focused on heuristic solution techniques for this problem, such as artificial intelligence (Hokans 1983), simulated annealing (Lockwood and Moore 1993), Monte Carlo integer programming (Barrett et al. 1998), tabu search (Clark et al. 2000; Richards and Gunn 2000), and hybrid tabu search –

Fig. 1. Forest units for potential harvest. Each unit is labeled with a unique identification number (1–9) as well as the total available acres in brackets.

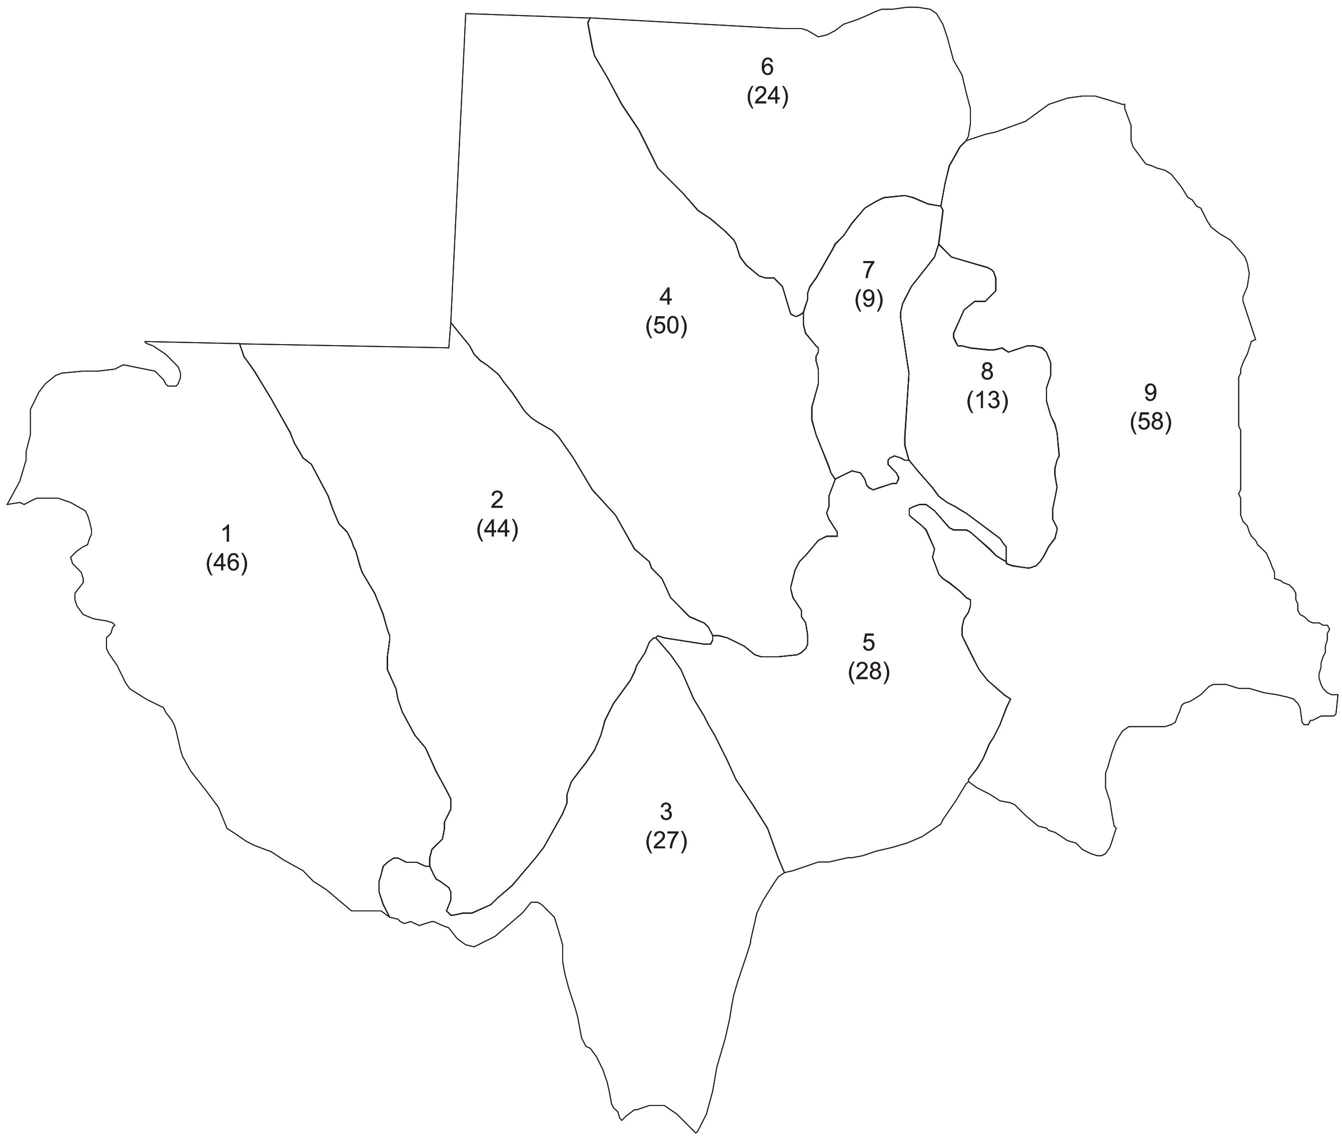

genetic algorithm (Boston and Bettinger 2001). Given increased computational capabilities and finer resolution spatial information, recent work is now being directed at exact solution approaches for the ARM (Barrett and Gilless 2000; McDill et al. 2002; Martins et al. 2001; Goycoolea et al. 2003). The research presented in this paper will build upon the work of Goycoolea et al. (2003) to address the issue of restricting average harvest area size.

Harvesting blocks

One of the major advances for addressing area-based restrictions in the ARM was the recognition in Goycoolea et al. (2003) that potential feasible blocks in harvest scheduling could actually be enumerated a priori, provided that management unit size and the maximum area restriction combine to produce no more than seven or so units in a block. Of course, the concept of forming larger representative blocks out of relatively small reporting units is not a new or novel concept in forestry. Recent work by Barrett (1997) details the historical perspective for doing this. Blocks are typically created by merging smaller units so that harvesting any one block prohibits neighboring blocks from being harvested simultaneously, using the URM. However, it has been found that a priori blocking can introduce biases (Jamnick et al. 1990; Daust and Nelson 1993; Murray and Weintraub 2002). The associated biases are spatial (such blocks are but one realization of possible aggregations of reporting units) and economic (reduced returns and greater operational costs), but have implications for forest productivity and sustainability as well (see Murray and Weintraub 2002).

The approach of Goycoolea et al. (2003) enumerated all feasible blocks in contrast with just one possible realization as would be evaluated in a URM application. Consider the following notation:

- i = index of planning units ($i = 1, 2, \ldots, I$)
- l = index of feasible blocks
- C_l = contiguous set of harvest units satisfying maximum area limit

- a_i = area of planning unit i
- A = maximum permissible contiguous area harvested

Using this notation, it is possible to specify how blocks are identified for use in an area-based model. Here, contiguity is defined as two units sharing a common boundary or point. A block C_l is formally defined as follows: (*i*) $\sum_{i \in C_l} a_i \leq A$ and (*ii*) units in C_l contiguous.

Recall that blocks are composed of individual planning units and are identified in advance.[2] All spatial combinations of units that form feasible blocks are enumerated and subsequently considered in the associated harvest scheduling model. This may be contrasted with the single blocking arrangement utilized in URM approaches (see Daust and Nelson 1993; Murray and Weintraub 2002). The small forest area shown in Fig. 1 will be used to illustrate potential blocks. Again, the unit identification number is given for each of the nine blocks as is its total area in brackets. If 120 acres is the maximum allowable block size, there are 74 unique potential contiguous blocks for this area that maintain the 120 acre restriction. These feasible blocks are enumerated in Table 1. The block sets are as small as one unit (e.g., {1}, {2}, {3},…) and as large as five units (e.g., {3,4,5,7,8}. All sets are less than the maximum area requirement of 120 acres in this case and so they are potential blocks that could be selected for harvesting.

While enumeration is generally considered something to be avoided, it is manageable in this case given unit sizes and the maximum area limit. After all, there are 511 unique combinations of units (many noncontiguous and many exceeding the 120-acre maximum) for Fig. 1, yet there are only 74 feasible blocks. Contributing as well is the fact that potential blocks must be spatially contiguous, which reduces the enumeration search space. With current computing capabilities, we have found that unit areas and maximum area limits resulting in at most eight units in a block are feasible to enumerate. Of course, more will be possible as computing capabilities continue to increase.

Enumeration used in this way is analogous to that detailed in Murray and Church (1996*b*) for constraint lifting in a URM. Further, it is similar to the approach described in McDill et al. (2002) for identifying blocks that violate maximum area restrictions in an ARM. In this latter approach, the associated number of necessary constraints grows rapidly and their constraints lack facet-defining structure, both of which contribute to an inability to solve moderate to large harvest scheduling applications exactly (Murray and Weintraub 2002). This is where the work of Goycoolea et al. (2003) made a major contribution in recognizing that it was possible, under certain conditions, to enumerate potential feasible blocks a priori and structure-associated restrictions.

Table 1. The 74 feasible harvest blocks associated with region shown in Fig. 1.

No. of units				
1	2	3	4	5
{1}	{1,2}	{1,2,3}	{1,3,5,7}	{2,3,5,7,8}
{2}	{1,3}	{1,2,5}	{1,3,5,8}	{3,4,5,7,8}
{3}	{2,3}	{1,3,5}	{2,3,5,7}	
{4}	{2,4}	{2,3,5}	{2,3,5,8}	
{5}	{2,5}	{2,4,6}	{2,4,7,8}	
{6}	{3,5}	{2,4,7}	{2,5,6,7}	
{7}	{4,5}	{2,5,7}	{2,5,7,8}	
{8}	{4,6}	{2,5,8}	{3,4,5,7}	
{9}	{4,7}	{3,4,5}	{3,4,5,8}	
	{5,7}	{3,5,7}	{3,5,6,7}	
	{5,8}	{3,5,8}	{3,5,7,8}	
	{5,9}	{3,5,9}	{4,5,6,7}	
	{6,7}	{4,5,6}	{4,5,6,8}	
	{6,9}	{4,5,7}	{4,5,7,8}	
	{7,8}	{4,5,8}	{4,6,7,8}	
	{7,9}	{4,6,7}	{5,6,7,8}	
	{8,9}	{4,7,8}	{5,6,7,9}	
		{4,7,9}	{5,7,8,9}	
		{5,6,7}	{6,7,8,9}	
		{5,6,9}		
		{5,7,8}		
		{5,7,9}		
		{5,8,9}		
		{6,7,8}		
		{6,7,9}		
		{6,8,9}		
		{7,8,9}		

Restricting spatial impacts

With blocks specified in advance, it is now possible to give an exact formulation of the area-based harvest scheduling model based upon the model structure proposed in Goycoolea et al. (2003). Note that in this research, it is assumed that no unit is available for multiple harvests within the planning horizon and that clear-cutting is the harvesting option. The formulation will utilize the following additional notation:

- t = index of planning periods ($t = 1,2,\ldots,T$)
- b_{it} = net benefit of unit i in period t
- d_{it} = harvesting cost of unit i in period t
- v_{it} = volume of unit i in period t
- U_t = upper bound on total volume harvested in period t
- L_t = lower bound on total volume harvested in period t
- p = green up-exclusion period length
- Ω_l = set of blocks that are incompatible with block l
- N_l = set of blocks that are adjacent to block l
- S_i = set of blocks containing havest unit i
- α_l = total area of block l
- β_{lt} = total net benefit of block l in period t
- δ_{lt} = total volume of block l in period t
- f_t = fixed harvesting equipment costs in period t
- $x_{lt} = \begin{cases} 1 \text{ if block } l \text{ is selected for harvest in period } t \\ 0 \text{ otherwise} \end{cases}$

Assuming that all feasible blocks have been identified, there are associated attributes that may be easily computed. In fact, it is possible to account for both linear and nonlinear relationships among units in a block. The three of interest in

[2] Of course, heuristic solution techniques would not necessarily need blocks to be identified in advance, and typically do not (see Lockwood and Moore 1993; Clark et al. 2000; Richards and Gunn 2000; Boston and Bettinger 2001).

this paper are total area (α_l), total net benefit (β_{lt}), and total volume (δ_{lt}) for each block. A linear specification is used in this research and is given for block l as follows:

$$[1] \qquad \alpha_l = \sum_{i \in C_l} a_i$$

$$[2] \qquad \beta_{lt} = \sum_{i \in C_l} (b_{it} - d_{it}) - f_t$$

$$[3] \qquad \delta_{lt} = \sum_{i \in C_l} v_{it}$$

It is conceivable that nonlinear specifications of these variables could be important. Incorporating nonlinear relationships for block attributes may be readily structured in eqs. 1–3 without impacting the properties of the linear formulation given below.

It is typically the case that many of the specified blocks share a common planning unit(s). That is, $i \subset C_l$ and $i \subset C_j$ when $l \neq j$. In such a case, these two blocks (l and j) are incompatible in the sense that they cannot simultaneously be harvested. A planning unit may be scheduled for one harvest at most, which means that it can necessarily belong only to one harvested block of units. Further, two adjacent blocks l and j, whose combined area exceeds the stipulated maximum, are considered incompatible as well. Incompatibilities may be expressed in the following way:

$$\Omega_l = [j | (C_l \cap C_j \neq \varnothing) \cup j \in N_l]$$

It is worth clarifying that potential feasible blocks are not prohibited in the definition of Ω_l. As an example, consider units 1 and 2 in Fig. 1 with respect to a 120-acre block limit. These units represent viable potential sets C_l, i.e., {1}, {2}, but would be defined as being incompatible in Ω_l, as they are neighbors. However, they may be simultaneously harvested. The reason this is not a problem is that all feasible blocks are enumerated, so there exists a set C_j representing their combination as a unique block, i.e., {1,2}. This means that nothing is lost by restricting them as independent blocks.

An integer programming formulation for the ARM can now be specified using the above notation and definitions. The model below represents an extension of the basic form proposed in Goycoolea et al. (2003).

Area restriction model (ARM)

$$[4] \qquad \text{Maximize} \sum_l \sum_t \beta_{lt} x_{lt}$$

Subject to

$$[5] \qquad \sum_{l \in S_i} \sum_t x_{lt} \leq 1 \quad \forall i$$

$$[6] \qquad \sum_l \delta_{lt} x_{lt} \leq U_t \quad \forall t$$

$$[7] \qquad \sum_l \delta_{lt} x_{lt} \geq L_t \quad \forall t$$

$$[8] \qquad \sum_{t'=t-p}^{t+p} (x_{lt'} + x_{jt'}) \leq 1 \quad \forall l, j \in \Omega_l, t \in (p+1, T-p)$$

$$[9] \qquad x_{lt} = (0,1) \quad \forall l, t$$

Objective 4 of the ARM maximizes total benefit over time. Constraints 5 prohibit a management unit from being harvested more than once. Constraints 6 and 7 establish upper and lower bounds on harvest volume in each planning period, respectively. Constraints 8 restrict incompatible blocks from being simultaneously harvested. Finally, constraints 9 impose integer restrictions on decision variables.

This formulation of the ARM represents a simplified harvest scheduling model. Additional constraints and alternative objective functions can easily be accommodated, though computational requirements could be prohibitive. Typical extensions include addressing road building, revenue expectations, nondeclining yield, etc. (Kirby et al. 1986).

Substantial research continues to focus on regulating spatial disturbance in harvest scheduling. As such, there is also a need for a modeling approach that addresses average harvest area size across a forest region to ensure that it remains at or below a stipulated threshold. As mentioned previously, the Sustainable Forestry Initiative suggests that average harvest area size not exceed 120 acres (American Forest and Paper Association 2001). Imposing a maximum area restriction would guarantee that average size requirements are maintained, but it fails to allow for the greater flexibility implied by the notion of "average". Boston and Bettinger (2001) appear to be the first to explicitly model average harvest area limits combined with absolute maximum area restrictions. In particular, they model average harvest area size of no more than 120 acres (as suggested in American Forest and Paper Association 2001), as well as impose a range of maximum area restrictions (148–220 acres) in the analysis of a forest plantation in Georgia. It is emphasized in Boston and Bettinger (2001) that forestry companies often develop and apply maximums in this range in practice, even though it is not expressly stipulated in the Sustainable Forestry Initiative.

Restricting average area

Addressing the average area issue in the context of the ARM is the primary objective of this paper. There are, in fact, a number of potential ways to approach this. Let $\bar{A}$ be the average area limit. The most straightforward approach for restricting average area would be to incrementally increase A, the maximum area limitation, until $\bar{A}$ is reached, provided that it is acceptable to have individual blocks exceeding $\bar{A}$. The rationale for doing this is that for a given maximum A, the resulting average size of harvested blocks will always be less than A.

Another possibility for addressing average area is to allow selective violations of constraints 8 by introducing a variable to track violations and minimize their occurrence, similar to what was suggested in Hochbaum and Pathria (1997). An example is as follows:

$$[10] \qquad \sum_{t'=t-p}^{t+p} (x_{lt'} + x_{jt'}) - y_{ljt} \leq 1 \quad \forall l, j \in N_l, t \in (p+1, T-p)$$

where y_{ljt} represents an binary decision that equals 1 when the adjacency condition is not maintained. If we add to our objective minimizing the weighted summation of all y_{ljt} variables, the controlled violation of strict maximum area limits is achieved. Note also that constraints 8 imposing no block overlap would still be necessary.

Table 2. Area restriction model (ARM) and Average-ARM results for Butter Creek with an average area constraint ($\bar{A}$) of 80 acres and varying maximum harvest area limits.

		ARM			Average-ARM	
A	No. of feasible blocks	Objective	Average harvested block size (acres)*	Solution time (s)	Objective	Solution time (s)
80	3 141	8 706.66	60.98	0.14	—	0.18
85	3 996	8 962.95	64.81	0.22	—	0.42
90	4 975	9 167.52	67.95	0.25	—	0.36
95	6 274	9 359.21	73.04	0.33	—	0.73
100	7 945	9 582.17	77.56	0.45	—	1.74
105	10 088	9 724.51	83.27	0.52	9 721.57	2.27
110	12 766	9 851.56	86.76	0.68	9 830.12	6.30
115	16 226	10 061.59	87.54	0.91	10 028.06	12.42
120	20 625	10 110.88	95.18	1.52	10 051.69	14.04
125	26 352	10 292.36	97.47	1.69	10 196.80	21.48
130	33 529	10 397.73	100.74	2.93	10 286.26	34.38

*1 acre = 0.404 685 ha.

Finally, an explicit approach for tracking and imposing an average area restriction is the following:

$$[11] \quad \frac{\sum_l \sum_t \alpha_l x_{lt}}{\sum_l \sum_t x_{lt}} \le \bar{A}$$

This may be simplified to produce the following linear constraint:

$$[12] \quad \sum_l \sum_t (\bar{A} - \alpha_l) x_{lt} \ge 0$$

This constraint directly addresses the issue of average area and does so in a way that can be readily evaluated and monitored, in contrast with the discussed alternatives. Further, constraint 12 may be readily included in the above ARM formulation. The Average-ARM (average area restriction model) is the above ARM formulation 4–9 with constraint 12 imposed as well. Operationally then, a property of the Average-ARM is that $\bar{A} \le A$. When this is not the case, i.e., $\bar{A} = A$, the problem reduces to the ARM, and constraint 12 is not necessary.

Formulation strengthening

To solve the Average-ARM, it is important to exploit spatial properties of the problem to produce a tighter formulation that can be more easily solved using exact techniques. Goycoolea et al. (2003) illustrate that the ARM is nothing other than a node-packing type problem. The graph for the ARM is structured using feasible blocks C_l as nodes and arcs defined by incompatibilities Ω_l. This projected graph is then solved as a node-packing problem. Given the relationship between the ARM and the Average-ARM, it is possible to utilize this same graphic representation.

The major opportunity for Average-ARM formulation improvement is that constraints 8 can be strengthened by utilizing higher-ordered cliques as well as introducing other facet-defining constraints (e.g., lifting, maximal cliques, odd cycles, web–antiweb, and K4 reduction). This is not a particularly new concept in either harvest scheduling (see Murray and Church 1996*a*) or node packing (Nemhauser and Woolsey 1988; Murray and Church 1997). Using the sets of block incompatibilities Ω_l, is it possible to identify all necessary cliques to impose in the Average-ARM formulation. Let k denote a particular clique set Φ_k, for which there are K such sets. Constraints 8 may be replaced with the following:

$$[13] \quad \sum_{j \in \Phi_k} \sum_{t'=t-p}^{t+p} x_{jt'} \le 1 \quad \forall k, t \in (p+1, T-p)$$

Clique sets may be identified using the approach suggested in Murray and Church (1996*a*). The result is a structurally superior formulation. All reported findings in this paper make use of higher-ordered clique constraints given in 13 to solve the Average-ARM efficiently.

Harvest scheduling application

Two forest regions located in northern California were used to assess the utilization of average area conditions in harvest scheduling. The first is referred to as Butter Creek and second is El Dorado. In this analysis, only a single time period will be examined to illustrate the properties of the Average-ARM harvest scheduling approach. Butter Creek has 351 harvest units averaging 25 acres in size, and El Dorado has 1351 harvest units averaging 38 acres in size. An average area restriction ($\bar{A}$) of 80 acres was imposed for a range of maximum area sizes (A = 80–130 acres for Butter Creek and A = 80–160 acres for El Dorado) for discussion purposes. For the results presented, single time period problems with no temporal flow requirements were considered. As the spatial nature of average and maximum area restrictions is fundamentally important, the intent is to highlight performance characteristics of the Average-ARM approach. Such characteristics are not necessarily obvious when constraining conditions associated with harvest volume requirements are introduced. Nevertheless, we do note computational experience with multiperiod application instances.

A C++ program was written to structure the associated ARM and Average-ARM applications, subsequently calling CPELX version 7.1, a commercial optimization package, as

Fig. 2. Area restriction model (ARM) solution A = 80 (maximum harvested area of 80 acres; 1 acre = 0.404 685 ha).

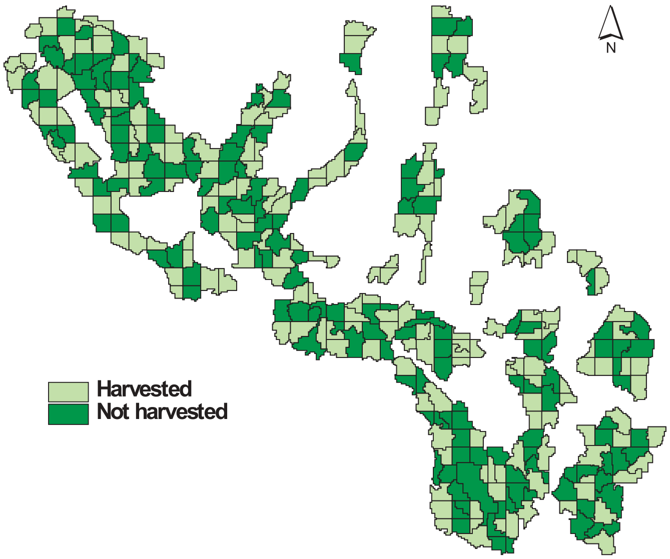

Fig. 3. Average-area restriction model (Average-ARM) solution for average harvest opening of 80 acres and maximum opening of 110 acres (1 acre = 0.404 685 ha).

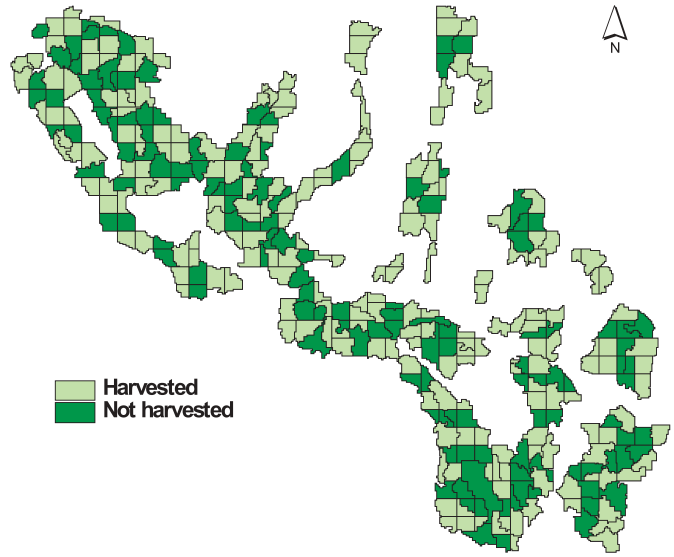

Fig. 4. Average-area restriction model (Average-ARM) solution for average harvest opening of 80 acres and maximum opening of 130 acres (1 acre = 0.404 685 ha).

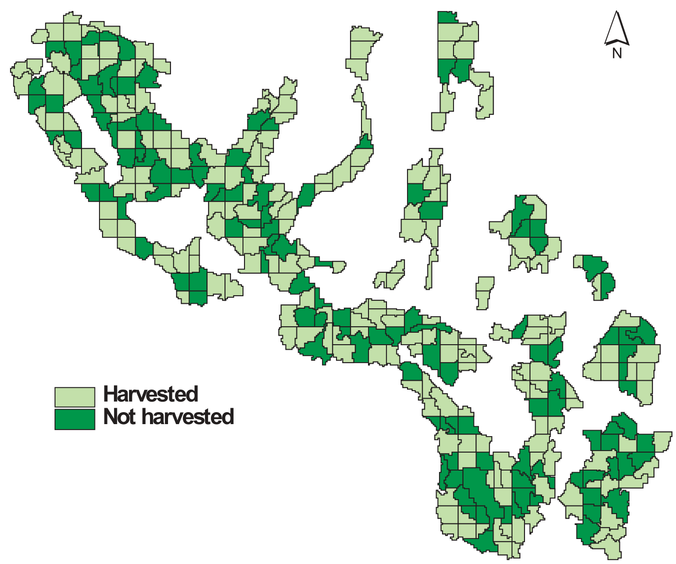

a callable library to solve the associated problems. The analysis was carried out on a Pentium IV-2000 processor personal computer with 2 GB of RAM running Mandrake Linux 8.2. Solution times are reported in CPU seconds. ArcView version 3.2, a commercial GIS package, was used to manage, analyze, and display the associated applications.

As suggested previously, an initial start for addressing average area conditions across a region is to use the ARM and systematically increment the maximum area limit. This analysis is given in Table 2 for the Butter Creek application. Note that processing time to generate the associated optimization problems was less than 9 s (including the time necessary to enumerate all feasible blocks). Figure 2 shows the ARM solution for A = 80. As the permissible maximum harvest block area increases in Table 2, the number of feasible blocks increases (e.g., 3141 blocks for A = 80 compared with 33 529 blocks for A = 130). Further, as the maximum area of an allowable block increases, there is a significant increase in the associated objective value. For example, when A = 80, the ARM objective is 8706.66, whereas for A = 100, the ARM objective is 9582.17. This is over a 10% increase in total return, and both solutions satisfy the 80-acre average harvest area limit. Table 2 suggests that there is a problem using the ARM in the context of average conditions when $A \geq 105$. In these cases, associated solutions are not viable planning alternatives because of the fact that they exceed the average harvest area requirement of 80 acres.

An alternative approach is to explicitly model the average area restriction using the Average-ARM formulation. Findings for the Butter Creek application are also given in Table 2. All Average-ARM solutions summarized in Table 2 satisfy the 80-acre average harvest area requirement, in contrast with the ARM results. In the cases where the average area is less than the maximum, e.g., A = 80–100 in Table 2, objective values are the same for the two models. This makes sense given that the ARM results satisfy the average area requirement up to this point. However, for A > 100, the two models differ substantially. Figures 3 and 4 show Average-ARM solutions for A = 110 and A = 130, respectively. An important observation in the Table 2 results is that it is possible to attain the average area requirement without substantial objective value degradation. For example, using the Average-ARM approach for A = 130 ensures that the average harvest area limit of 80 acres is maintained. This is in contrast with the ARM approach, which results in an average harvested area of 100.74 acres for A = 130. The Average-ARM approach is able to achieve this constraint with only a slight decrease in total return (less than 1%) and computational effort remains reasonable (less than 35 s).

Similar comparative results were found for the El Dorado application. Table 3 gives ARM and Average-ARM results.

Table 3. Average-area restriction model (Average-ARM) results for El Dorado with an average area constraint ($\overline{A}$) of 80 acres and varying maximum harvest area limits.

		ARM			Average-ARM	
A	No. of feasible blocks	Objective	Average harvested block size (acres)*	Solution time (s)	Objective	Solution time (s)
80	5 252	1 511 490	48.49	0.32	—	0.36
85	6 396	1 546 040	51.03	0.44	—	0.63
90	7 754	1 570 870	54.70	0.66	—	0.80
95	9 381	1 585 205	57.48	0.85	—	1.01
100	11 462	1 608 200	61.69	1.12	—	1.60
105	13 937	1 631 070	64.88	1.25	—	1.57
110	17 001	1 655 435	69.41	1.69	—	4.25
115	20 813	1 672 900	71.11	2.09	—	8.05
120	25 540	1 697 695	73.77	2.60	—	13.30
125	31 358	1 709 565	77.38	3.41	1 709 550	21.66
130	38 584	1 720 720	79.19	7.03	1 720 720	42.11
135	47 701	1 744 425	85.01	4.75	1 744 425	63.51
140	59 038	1 763 535	83.77	8.32	1 763 535	109.45
145	73 255	1 777 220	88.14	9.03	1 777 220	175.06
150	90 963	1 787 425	91.00	10.50	1 786 705	302.47
155	113 188	1 792 915	94.32	15.17	1 791 430	455.67
160	141 185	1 804 135	98.16	20.44	1 800 595	710.93

*1 acre = 0.404 685 ha.

The processing time to generate the associated optimization problems was less than 224 s in both cases. For $A > 135$, the average condition of 80 acres is exceeded using the ARM. In contrast, the average area restriction of 80 acres is maintained in all of the Average-ARM solutions. Further, this is achieved without significantly impacting the total return (less than 1% for $A = 160$), similar to the findings for Butter Creek.

Discussion and conclusions

One significant issue is that feasible potential harvesting blocks grow substantially as A increases. The relationship between the number of associated blocks and the maximum allowable harvested area is nonlinear in Tables 2 and 3, so computational limitations could quickly become an issue. However, no limitations were encountered in our applications.

While not reported in this paper, we have also examined the impacts of addressing multiple planning periods for these applications. Adding the temporal dimension to the Average-ARM, along with associated constraints, does increase problem complexity considerably. However, we were able to achieve reasonable results (<2% optimality gap) for three time period instances of El Dorado in less than 18 h of processing time. Again, it should be noted that the number of planning units is 1351 in this case, so this is a fairly large planning problem to begin with.

This paper presented an exact approach for modeling and solving a harvest scheduling problem where addressing average area of disturbance is important. This work builds upon the ARM approach to account explicitly for average harvested area. The application results demonstrated that the model is computationally feasible, but more importantly it was shown that attempting to address average area implicitly using an ARM rather than the Average-ARM could be problematic. It is possible using the Average-ARM to ensure that average area impacted is limited, while allowing larger maximum areas to be disturbed, if this is desired. After all, guidelines such as those outlined in the Sustainable Forestry Initiative regulate only average conditions, not absolute maximums. This is significant because failing to model the spirit of a policy or guideline could result in substantial economic or environmental losses, possibly leading to greater overall natural resource degradation in the long term.

It is important that exact approaches for forest planning problems be developed when possible. One reason is that heuristic solution techniques, like the hybrid tabu search – genetic algorithm developed by Boston and Bettinger (2001), need to be evaluated in terms of their efficiency and solution quality. Without exact approaches, there is no way to assess heuristics. This is an issue because heuristic solution approaches remain essential for large-scale harvest scheduling applications, so there is a need for knowing how well such heuristics perform given application characteristics. Another important reason that exact approaches are necessary is that evaluating policy impacts and interpreting the significance of constraining conditions can only be done with certainty if exact techniques are used. Differences in findings using heuristics may only be a byproduct of suboptimality or local optima as opposed to representation, model used, or policy context. The developed Average-ARM provides a sound basis for examining the impacts of average harvest area size restrictions as well as absolute maximum area bounds.

Acknowledgements

Partial funding for the first author was provided by the National Science Foundation (Geography and Regional Science Program and the Decision, Risk, and Management Science Program) under grant BCS-0114362. Partial funding for this research was also received from FONDECYT under grant 1020784 and Nucleous Milenium Complex Engineering Systems. The authors thank Klaus Barber of the USDA Forest Service for making the utilized forests available for analysis.

References

American Forest and Paper Association. 2001. Sustainable forestry initiative standard [online]. Available from http://www.afandpa.org/ (cited 17 January 2001).

Barrett, T.M. 1997. Voronoi tessellation methods to delineate harvest units for spatial forest planning. Can. J. For. Res. **27**: 903–910.

Barrett, T.M., and Gilless, J.K. 2000. Even-aged restrictions with sub-graph adjacency. Annals of Operations Research, **95**: 159–175.

Barrett, T.M., Gilless, J.K., and Davis, L.S. 1998. Economic and fragmentation effects of clearcut restrictions. For. Sci. **44**: 569–577.

Boston, K., and Bettinger, P. 2001. The economic impact of green-up constraints in the southeastern United States. For. Ecol. Manage. **145**: 191–202.

Boston, K., and Bettinger, P. 2002. Combining tabu search and genetic algorithm heuristic techniques to solve spatial harvest scheduling problems. For. Sci. **48**: 35–58.

Clark, M.M., Meller, R.D., and McDonald, T.P. 2000. A three-stage heuristic for harvest scheduling with access road network development. For. Sci. **46**: 204–218.

Daust, D.K., and Nelson, J.D. 1993. Spatial reduction factors for strata-based harvest schedules. For. Sci. **39**: 152–165.

Goycoolea, M., Murray, A.T., Barahona, F., Epstein, R., and Weintraub, A. 2003. Harvest scheduling subject to maximum area restrictions: exploring exact approaches. Oper. Res. In press.

Hochbaum, D.S., and Pathria, A. 1997. Forest harvesting and minimum cuts: a new approach to handling spatial constraints. For. Sci. **43**: 544–554.

Hokans, R.H. 1983. Evaluating spatial feasibility of harvest schedules with simulated stand-selection decisions. J. For. **81**: 601–603, 613.

Jamnick, M.S., Davis, L.S., and Gilless, J.K. 1990. Influence of land classification systems on timber harvest scheduling models. Can. J. For. Res. **20**: 172–178.

Jones, J.G., Meneghin, B., and Kirby, M. 1991. Formulating adjacency constraints in linear optimization models for scheduling projects in tactical planning. For. Sci. **37**: 1283–1297.

Kirby, M.W., Hager, W.A., and Wong, P. 1986. Simultaneous planning of wildland management and transportation alternatives. TIMS Studies in the Management Sciences, **21**: 371–387.

Lockwood, C., and Moore, T. 1993. Harvest scheduling with spatial constraints: a simulated annealing approach. Can. J. For. Res. **23**: 468–478.

Martins I., Constantino, M., and Borges, J. 2001. A column generation approach for solving a non-temporal forest harvest model with spatial structure constraints. Technical Report Instituto Superior de Agronomía Universidad Técnica de Lisboa.

McDill, M.E., Rabin, S.A., and Braze, J. 2002. Harvest scheduling with area-based adjacency constraints. For. Sci. **48**: 631–642.

Murray, A.T. 1999. Spatial restrictions in harvest scheduling. For. Sci. **45**: 45–52.

Murray, A.T., and Church, R.L. 1996*a*. Analyzing cliques for imposing adjacency restrictions in forest models. For. Sci. **42**: 166–175.

Murray, A.T., and Church, R.L. 1996*b*. Constructing and selecting adjacency constraints. INFOR, **34**: 232–248.

Murray, A.T., and Church, R.L. 1997. Facets for node packing. European Journal of Operational Research, **101**: 598–608.

Murray, A.T., and Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. For. Sci. **48**: 779–789.

Nemhauser, G.L., and Wolsey, L.A. 1988. Integer and combinatorial optimization. John Wiley & Sons Inc., New York.

Richards, E.W., and Gunn, E.A. 2000. A model and tabu search method to optimize stand harvest and road construction schedules. For. Sci. **46**: 188–203.

Thompson, E.F., Halterman, B.G., Lyon, T.S., and Miller, R.L. 1973. Integrating timber and wildlife management planning. For. Chron. **47**: 247–250.

Van Deusen, P. 1996. Habitat and harvest scheduling using Bayesian statistical concepts. Can. J. For. Res. **26**: 1375–1383.