A forest planning system for solving spatial harvest scheduling problems

Remsoft Inc.
332 Brunswick Street,
Fredericton, NB,
Canada E3B 1H1

1-800-792-9468 or
1-506-450-1511

www.remsoft.com

Finalist

CORS Competition on the Practice of O.R.

CORS National Conference

Windsor, ON

Canada

June 7 - 9, 1999

Karl R. Walters, Graduate Research Assistant,
North Carolina State University,
College of Forest Resources
Raleigh, NC
email: waltek@mindspring.com

Ugo Feunekes, Vice-President
Remsoft Inc.,
332 Brunswick St., Fredericton, NB
Canada E3B 1H1
www.remsoft.com email: ugo@remsoft.com

Andrew Cogswell, Senior Analyst
Remsoft Inc.,
332 Brunswick St., Fredericton, NB
Canada E3B 1H1
www.remsoft.com email: acogswell@remsoft.com

Eric Cox, Forest Planning Project Leader
Champion International Corp.
9485 Regency Square Blvd., Suite #300
Jacksonville, FL 32225
email: coxe@champint.com

Abstract

Spatial harvest scheduling at the forest level is a daunting task. Apart from the sheer amount of data involved and the lingering effect of past activities, the problem is distinguished by contrasting temporal (long- versus near-term) and spatial (forest versus stand) scales and complicated by often disparate planning objectives at each level of planning. Recent changes in forestry practices (e.g., strict limitations on the extent and juxtaposition of openings in the forest) complicate the problem further still and make the actual on-the-ground implementation of a plan tricky at best. It is critical that a problem of this type is tackled in a systematic and organized fashion to ensure high quality solutions that can withstand rigorous evaluation. Remsoft has developed a hierarchical forest planning system that solves spatial harvest scheduling problems ranging from hundreds to hundreds-of-thousands of spatial units in size. The system is composed of a forest modeling system that determines long-term, sustainable activity levels using linear programming and a spatial harvest scheduling system that locates near-term activities in the forest using heuristics. Other than the basic spatial data (available in public domain format), the system is entirely self-contained, capable of producing all necessary data and storing and editing solutions. The merit of the system lies in its broad applicability and ability to be used by any *organizations using it. In this report, we provide a case history of how one forest agency faced with Sustainable Forestry Initiative^{SM} guidelines adopted the system to assist them in their spatial harvest scheduling.*

Introduction

Until very recently, the nature of forest management scheduling on public lands versus private industrial lands has been markedly different. Industrial forest management has concentrated largely on minimizing rotation lengths and costs of harvesting to yield the highest possible economic returns. Public lands management, with its many competing and often conflicting goals of multiple use, has required the consideration of larger sets of constraints on forest management activities. Linear programming based harvest schedules have been favored by both public lands and industrial forest agencies, but the public lands models have doubtless been far larger and more complicated. Moreover, the majority of non-proprietary research into forest modeling techniques has been carried out by the USDA Forest Service or by researchers applying techniques for Forest Service management problems.

In October 1994, the American Forest & Paper Association (AF&PA) drafted a set of principles called the Sustainable Forestry Initiative^{SM} (SFI) to guide forest management activities on the lands of member companies. These guidelines included limits on clear-cut size, minimum buffer widths between openings, minimum time intervals between cuts of adjacent areas (green-up intervals) and other restrictions on forest operations. Not only did these guidelines represent a large change in the operations of most member companies, but compliance with these regulations was made a condition of membership in AF&PA in January 1996. Almost immediately, forest management planning on industrial lands became very focused on spatial planning issues, the scope of which are easily as complex as those of public lands management. To date however, there has been little research effort applied to the creation of tools that can handle the spatial planning problems inherent to SFI. This paper describes the SFI problem in detail, as well as a set of tools that, though not developed specifically for the SFI problem, have been implemented with good success by several AF&PA members, including Champion International Corporation.

Champion International

Champion International is a large forest-products company with land holdings throughout the United States, and significant acreage throughout the Southeast and Texas. Like most other companies in the region, the bulk of these forestlands have been managed as southern pine plantations. These plantations are typically large (several hundred to upwards of several thousand acres) and have been harvested multiple times with rotations ranging from 10 to 30 years, depending on growth rates. Harvest scheduling has typically been stand-level oriented, based on maximization of present net worth and land expectation value. Concentration of harvesting operations to minimize the cost of moving harvesting equipment has been used extensively. The resulting landscape is generally one of overwhelming uniformity, with large tracts of the same species of very similar age.

Champion's Sustainability and Stewardship (S&S) operating guidelines are loosely based on the AF&PA SFI guidelines, and are at least as strict and sometimes more restrictive than those set out in SFI. For example, clearcut harvesting is generally limited to areas less than 240 acres but only where necessary will clearcut areas actually exceed 120 acres. Clearcut harvest areas are considered contemporary if they are established within a fixed number of years of one another. Until better estimates of regeneration response are known, Champion has conservatively set the green-up period at 4 or 5 years in most regions. Clearcutting will not be permitted within 300 feet of a contemporary clearcut harvest area unless a watercourse that requires more restrictive riparian zone management separates the two.

Although it is quite easy to state that forest management planning as an activity is deciding what activities to implement, in what place and at what time, the actual process of making these determinations is far from easy. In years past, locating a harvest activity had little consequence elsewhere in the forest or in future years, except for the fact that the harvested timber was no longer readily available and that regeneration costs would be incurred. In the SFI planning problem, the location of a single harvest block can make large areas of the forest unavailable for harvesting for several years into the future, and so a poor choice in location can severely limit options for the future. In the case of Champion and other companies operating in the southeast, past practices are probably the biggest obstacles to overcome under the new operating guidelines. In general, making harvest blocks as large as possible and keeping green-up delays as short as possible is the best way to mitigate adjacency constraints.

Champion is not unique with its S&S guidelines and other companies in the southeast have adopted similar operating principles. In fact, continued membership in AF&PA is contingent on adoption of SFI principles. The majority of forest companies implementing these new types of operating guidelines face similar difficulties in implementing them: they require some method of determining spatial harvest schedules in a timely manner, and they have to contend with a major shift in harvesting logistics associated with smaller harvest areas dispersed over much wider areas.

Alternative approaches to the SFI planning problem

Much of the literature on spatial planning deals with a model formulation that Murray (1999) calls the Unit Restriction Model (URM). Basically, a preliminary analysis yields integral spatial units such that the area of each one is less than the maximum opening size allowed by regulations. A mathematical model may then be formulated to guarantee that no two adjacent units are simultaneously treated in a management plan and thereby it is certain that the maximum opening restriction is not violated.

Maximize
$Z = \Sigma_i \Sigma_t \alpha_{it} x_{it}$ (1)
Subject to
$\Sigma_t x_{it} \leq 1 \quad \forall i$ (2)
$\Sigma_i \beta_{it} \; x_{it} \geq L_t \quad \forall t$ (3)
$\Sigma_i \beta_{it} \; x_{it} \leq U_t \quad \forall t$ (4)
$x_{it} + x_{jt} \leq 1 \quad \forall i, t, j \in N_i$ (5)
$x_{it} = (0, 1) \quad \forall i, t$ (6)
where
i = index of planning units,
t = index of time periods,
α_{it} = benefit or revenue associated with treating unit i in period t
β_{it} = volume contribution for treating unit i in period t
L_t = lower bound on total volume produced in period t
U_t = upper bound on total volume produced in period t
N_i = set of planning units adjacent to unit i

Decision variables:

$$x_{it} = \begin{cases} 1 & \text{if } \textit{unit } i \text{ is treated in period } t \\ 0 & \text{otherwise.} \end{cases}$$

Formulated as a mixed-integer programming (MIP) model, the URM in theory can be solved exactly by commercial solver software using branch and bound algorithms. However, such solvers are rather limited when it comes to the size of problems that can be solved exactly in a reasonable amount of time. Thus, the rationale for much of the research into URM planning problems has been to reduce the number of adjacency constraints to fit within the capabilities of commercial IP codes. If a given problem can be correctly represented by fewer constraints and good structure, then a correspondingly larger problem should be solvable. However, even the largest problems solved using these methods numbered only a few hundred spatial units.

Murray (1999) defines an alternative spatial harvest scheduling problem to the URM called the Area Restrictive Model (ARM). The ARM is identical to the URM except for how adjacency constraints are handled.

Maximize
$Z = S_i S_t a_{it} x_{it}$ (7)
Subject to
(2) – (4), (6)
$¦_{it}(x) \leq A$ "i, t (8)
where
A = *maximum permissible contiguous area treated*
v_i = *area of unit* i
$f_{it}(x)$ = *recursive function summing all treated neighboring units associated with* x_{it} (*if* $x_{it} = 1$).

In the ARM, spatial units tend to be much smaller than the maximum opening size allowed by regulation. Adjacent units may be harvested simultaneously as long as total opening size of the contiguous units does not exceed the maximum opening size limit. The resulting adjacency constraint function is not only recursive (a spatial unit's neighborhood includes its neighbors and the neighbors of its neighbors), but also nonlinear (the number of adjacencies is not linearly dependent on the number of units within the block, nor on its total size).

The ARM appears to be a more realistic formulation of the SFI management problem on two fronts: 1) it *directly* models maximum opening size restrictions and adjacency relationships that may be based as much on proximity as contiguity, and 2) it does not assume that harvest block configurations be known *a priori*. Many industrial forest planners are facing spatial restrictions for the first time, and because the planning problem has changed so dramatically in the last few years, there is little previous experience to draw on that would indicate whether a given block layout is predisposed to adjacency conflicts. Moreover, results have shown that poor block configuration can contribute to lower objective function values (Jamnick and Walters, 1991; Walters and Feunekes, 1994). Unfortunately, the nonlinear structure of the ARM adjacency constraint function indicates that it is highly unlikely that exact methods for solving the ARM will be developed.

Remsoft s spatial forest planning system

Since about 1990, Remsoft has been conducting original research and development into forest management planning tools, and it currently markets two software packages, Woodstock and Stanley, as an integrated spatial forest planning system. The system is based on the hierarchical planning approach of Jamnick and Walters (1993). First of all, a strategic forest management schedule is developed based on uniform forest strata. The strategic schedule is then allocated to individual stands within the forest to form feasible harvest units that are subsequently scheduled under harvest flow and spatial restrictions. What sets apart this approach from others in the literature are two key assumptions. First, only the strata scheduled for harvest in the initial planning periods of the strategic plan are eligible for allocation in the tactical plan. This greatly reduces the magnitude of the planning problem because fewer periods and fewer strata are considered. It also prevents planners from exploiting other strata in the short-term that are needed for harvest in the future. If such restrictions are not made, short-term feasibility of spatial scheduling may be enhanced, but at the cost of increasingly difficult spatial scheduling in the future.

Second, the approach generally retains the goals and constraints of the long-term forest management plan by preserving the basic timing choices of the strategic harvest schedule. Variations in the timing choices do occur in order to meet spatial restrictions, but in general, stands that are scheduled for harvest early in the strategic planning horizon are similarly scheduled early in the tactical planning horizon. If the strata are quite homogeneous with respect to growth and yield, it can be reasonably assumed that the output flows of the tactical plan will be very similar to the strategic plan if the timing choices in the tactical plan closely approximate the timing choices in the strategic plan.

Woodstock

Woodstock is Remsoft's strategic forest management modeling system. Although it was originally conceived as an inventory projection model, Woodstock now supports binary search, linear programming and Monte Carlo simulation formulations as well (Walters 1996). The key to Woodstock's flexibility in model formulations is its modeling language, which allows forest planners to define forest structure, activities, outputs and reports in a format that best reflects their way of doing things. Rather than attempt to provide a comprehensive list of featured activities and outputs, the Woodstock modeling language provides users with building blocks to explicitly model forest management activities and regimes. Forest types are defined in terms of static attributes (landscape themes) and dynamic attributes (yield components) that completely describe the structure of a forest through time. Activities that are carried out as part of managing the forest can effect change in structure by changing landscape themes, and they can trigger the production of outputs (based on one or more yield components).

Because activities and outputs are separated in the model from control structures that determine the model formulation, converting an inventory projection model to a generalized Model II linear programming (LP) model (Johnson and Scheurman 1977) can be accomplished with a few lines of code: all of the yield, output and activity data that define the essence of the forest model remain unchanged. For an inventory projection model, the Woodstock interpreter executes the model directly; for a LP model, the interpreter generates a LP matrix that may be solved using one of several supported commercial LP solvers. Once an optimal solution is found, the Woodstock interpreter reads in that solution and executes the activities chosen.

Woodstock offers several unique capabilities. The modeling language allows resource planners to describe forest land, activities and outputs in terms specific to their own situation rather than adapting their problem to fit a predefined formal structure with a limited number of silvicultural options. The reporting and graphics features are extensive, including run-time graphs of key outputs and user-defined reports in ASCII text and spreadsheet formats that allow for additional analysis using standard business software. Woodstock allows modelers to compile a Dynamic Link Library (DLL) that calculates yield coefficients in real-time based on inventory information from a geo-referenced database. In addition to minimizing the time and effort required to generate individual yield tables, the DLL allows the same growth and yield functions used in the company inventory system to generate yield estimates that are as current as the GIS database information. As a fully 32-bit Windows program, Woodstock runs natively on NT workstations along with other company planning software. It generates standard MPS files, and directly supports several commercial LP solvers.

Stanley

Stanley (Remsoft 1996) is a tactical planning tool also developed by Remsoft Inc. In most forest management scheduling models, including Woodstock, stands of similar composition and structure are aggregated into a single class, and so management activities in the resulting schedule are assigned to classes rather than individual stands. Stanley attempts to implement a forest management schedule by assigning management activities to appropriate forest stands within the class. Moreover, the assignation of activities is done while recognizing restrictions on opening size and green-up intervals and output flow constraints. This capability is what drew the attention of resource planners at Champion International.

In order to use Stanley, there must be a link between the classes in the strategic scheduling model and the individual stands that comprise the class. Stanley relies on the attribute and spatial information stored in *Shapefiles* (ESRI 1997) to link scheduled forest classes to their component stand polygons. Each polygon in the geographic database includes fields corresponding to the landscape theme attributes used in a Woodstock model. To identify stand polygons that are eligible for a given management activity that was scheduled by Woodstock, Stanley locates those polygons with the correct age and thematic attributes. Next, Stanley searches for adjacent stands that are eligible for a compatible management activity in the same period using the geographic information of the shape file. If it finds one or more compatible stand polygons, the contiguous area defined by the stand polygons is assigned a unit number, indicating that the group of stands should be scheduled for treatment as an operational unit. The key point to remember is that Stanley only considers those forest classes and polygons that are scheduled during initial periods of a Woodstock schedule; forest classes from later periods are not considered (Feunekes and Cogswell 1997).

When choosing among equally good alternatives, Stanley tries to select polygons that contribute to a more compact harvest unit rather than those that yield more amoeboid shapes. To comply with opening size restrictions, Stanley verifies that any potential opening is larger than a user-specified minimum opening size – if the largest opening possible is smaller than the specified minimum, the component stands are flagged as impossible area. Similarly, if the addition of a stand polygon to a contiguous block area will exceed the user-specified maximum opening, Stanley will stop trying to increase the block size and will move on to another potential block. Last, Stanley assigns final harvest periods to each block that closely match the timing choices of the strategic harvest schedule, while maintaining specified opening size and green-up intervals. Stanley does not rely on an optimization algorithm like branch and bound but rather employs several different heuristics to balance flows, control block configuration and minimize timing choice deviations.

Woodstock and Stanley are complementary software tools. During matrix generation, Woodstock generates a *choices file*, a binary representation of the decision variables used in the LP and their objective function coefficients, the basis, and the reduced costs for non-basic variables. Stanley reads this file to help choose appropriate alternatives when it must deviate from the timing choices that were selected in the optimal solution. In turn, Stanley processes the geographic information in the shape files to generate the initial stratum areas file for a Woodstock model. Together, Woodstock and Stanley help to streamline the spatial planning process by automating steps in the modeling process and by maintaining the linkages between strategic and tactical decisions that are necessary for a workable system (Covington et al. 1988, Nelson and Brodie 1990; Jamnick and Walters 1993).

Remsoft's Spatial Planning system is a modified ARM approach. The spatial scheduler assigns harvest prescriptions to individual stands, grouping them where possible, and rearranging them where necessary to avoid adjacency conflicts and opening size violations. Once the prescriptions are assigned, harvest blocks are identified by labeling contiguous areas harvested in the same planning period. Hence, block configuration is a by-product of the solution method, like the ARM. However, the pure ARM strategy is to maximize benefits while simultaneously addressing management objectives and spatial restrictions. In the Remsoft Spatial Planning system, spatial constraints are relaxed yielding the standard LP formulation that has been used for many years. Once solved, the planning horizon is partitioned, and only decision variables from the early portion of the planning horizon are used as input to the spatial scheduler. By working with only a subset of the forest planning problem, it becomes feasible to consider the spatial restrictions directly. Rather than be concerned with competing resources and maximizing outputs, the spatial scheduler simply minimizes deviations from the timing choices of the optimal solution while being sure that spatial restrictions are not violated.

Woodstock/Stanley compared to URM approach

Woodstock/Stanley is not an optimization approach and so satisficing solutions to the spatial planning problem are all that can be expected. Given the nature of the ARM it is unlikely that an optimization approach for solving it exactly will ever be developed (Murray 1999).

However, it is still a valid question to ask how well Woodstock/Stanley performs relative to a URM formulation which can be solved exactly. Since the URM requires a block configuration be established a priori, a valid comparison involves solving a URM MIP based on the same harvest blocks delineated by Stanley. Three considerations are important for this comparison. First, the URM solution is optimal only for the given block configuration; a different block configuration may yield a solution with a higher objective function value. Second, much of the computational effort of the Stanley algorithm is done to generate a feasible block configuration while it assesses spatial restrictions and flow constraints. In the URM MIP formulation, the computational effort is spent only on spatial restrictions and flow constraints because the block configuration is already given. Hence, comparing just solution times is not appropriate.

A hypothetical forest planning problem was developed that represents southern pine plantation management common to forest companies operating throughout the U.S. southeast. The forest was composed primarily of slash and loblolly pine plantations, cypress ponds, and bottomland hardwoods, totaling about 87,000 acres in 13,047 polygons. A strategic forest plan was developed with Woodstock for this forest using 1-year planning periods with a planning horizon of 25 years and an objective function that maximized present net worth subject to even-flow constraints on harvest volume. A tactical plan was developed using Stanley for the first 10 years of the strategic planning horizon. Harvest blocks were constrained to be no smaller than 10 acres and no larger than 120 acres. The green-up interval was set at 5 years and the proximity distance for contemporary harvest blocks was set at 300 feet. Harvest flow fluctuations were constrained to be less than 10% across all periods.

Stanley was run for 15 minutes on a 266 MHz Pentium II Dell Dimension computer, and the best solution was retained. The harvest blocks selected by Stanley were used to develop a MIP model with a maxmin objective function on harvest volume. The solution times and results for each model are presented in Table 1.

Table 1. Execution times and objective function values for Stanley and MIP formulations.

Model	**Execution Time**	**Objective Function Value**
Woodstock–1400 Rows, 8670	44 sec – matrix generation	45,525 cu/yr
columns, 42303 non-zero elements	23 sec – LP solution/con-version	
Stanley – 30,000 ac, 5020 polygons	900 sec	Lowest period volume 34,266 cu/yr, 76.4% of LP volume target, 4.9% flow variation over 10 periods
Total Woodstock/Stanley	967 sec	
MIP (using Ketron MIPIII)	893 sec (stopped after 4 integer solutions found)	Maximum volume: 35,224 cu/yr, 77.4% of volume LP volume target, 0.3% flow variation over 10 periods

The Woodstock model considered the entire forest for a 25 year planning horizon, yet it was able to represent all the timing choices for harvest and the multiple age and forest type combinations in fewer than 9000 decision variables. Because the discrete location of individual forest stands was ignored most forest polygons could be collapsed into a small number of classes. The Stanley run had a planning horizon of only ten years, reducing the number of timing choices and eligible polygons by 60%. However, that still left more than 5,000 polygons to consider for blocking candidates – an unaccountably large number of possible combinations if enumeration was to be used to find an optimal solution. But because the Stanley algorithm knows when these polygons should be cut (from the LP solution), the number of possible combinations is reduced to tractable size since most combinations are not technically feasible.

There is a large reduction in objective function value between the LP relaxation solution and the spatial solutions. In large part, this is due to the spatial arrangement of stands within the forest. There are many tracts of uniform forest types that are several hundred acres in size. Ignoring the 120 acre opening size limit, the LP solution chooses to liquidate entire tracts. However, under SFI operating guidelines, many of these areas cannot be harvested during the 10 year planning horizon without violating the maximum opening size and adjacency constraints. Therefore, a significant portion of the forest must remain unharvested when spatial restrictions are enforced.

The MIP solution yielded a solution within 3.2% of the LP bound, and its 10-year harvest level is about 1% better than the Stanley solution (roughly 352,000 cunits over the planning horizon) and it has better flow fluctuations (see Figure 1.) It is important to remember that the Stanley heuristics attempt to maximize total harvest volume and balance flows across all periods whereas the MIP formulation simply maximizes the lowest harvest level in all periods. The MIP solution left 8 blocks unharvested and presumably some of the blocks in the Stanley solution could have been left unharvested as well. This would have had a small impact on total volume harvested but would have improved the flow fluctuations appreciably. The MIP formulation capitalizes on this tradeoff and so it is not surprising that it outperforms a heuristic. However, the Stanley heuristics are still performing remarkably well.

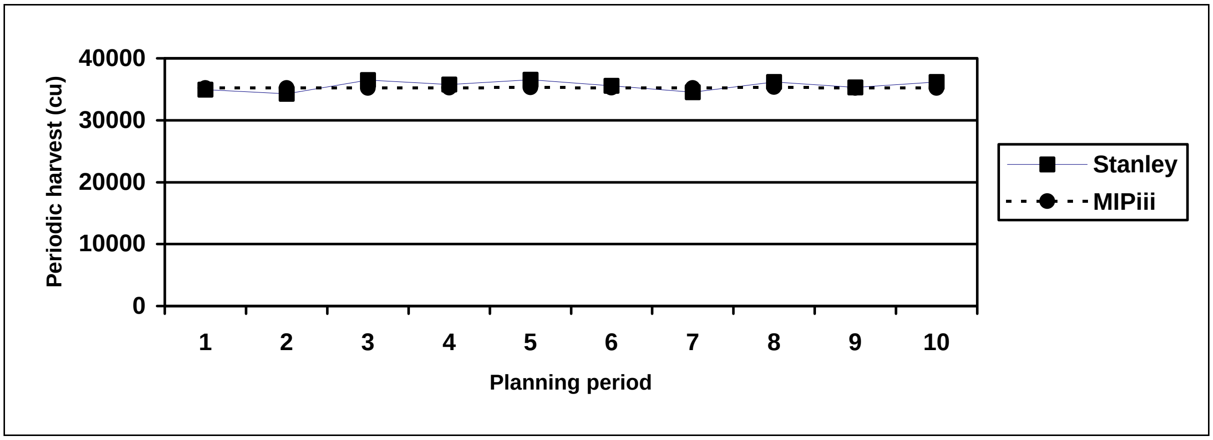

Figure 1. Harvest flow profile of Stanley and MIP spatial harvest schedules.

Champion s experience with Woodstock/Stanley

In 1995, Champion acquired two copies of Woodstock for testing in its Eastern Florida region to determine whether it was suitable for use company-wide. At the time, Stanley was still being developed as a commercial product, and though it was not yet available, the company was aware of its planned capabilities and had access to beta copies of the software. The following year, Champion implemented Woodstock throughout its nine regions and contracted Remsoft to conduct extensive training of regional technical foresters in the use of Woodstock.

Champion has been working at developing a new unified forest information system. A standardized procedure of data flows had to be developed to maintain integrity of the geographic information across the strategic, tactical and operational levels. Implementation of these procedures is ongoing, as is development of new growth and yield models that link directly to Woodstock through DLLs. Stanley has been acquired by most of Champion’s regions and internal Woodstock model development has been ongoing for more than two years. In the long run, Woodstock and Stanley are considered integral parts of Champion’s forest planning efforts.

What has made Stanley well-accepted by its users is the fact that the program quickly generates solutions that, by and large, make intuitive sense. Some users have reported that approximately 60–70% of the harvest blocks selected by the Stanley algorithm are acceptable as-is. To correct the remaining 30–40%, foresters can override choices made by the Stanley algorithms or impose decisions directly, before rerunning the program. Through an iterative process an acceptable operational plan is developed that reflects professional judgement and operational reality: initial algorithmic solutions are incrementally adjusted until a spatially-feasible, operational harvest schedule results. Stanley provides a tremendous leg-up on the problem by giving foresters a head-start on a solution: it is easier to evaluate a solution and make corrections than to start from scratch, which has been the traditional approach to harvest blocking.

One of the strengths of the Woodstock/Stanley approach is that the strong linkages between the GIS database and the planning models prevent planners from hiding or ignoring important information. Analysis will quickly alert a forest planner to small but important details that might otherwise be blurred by the scope of a forest comprised of hundreds of thousands of acres. Such attention to detail is what makes it possible to rigorously comply with SFI guidelines. Woodstock/Stanley provides a good framework for empowering management decisions: planning allows foresters to anticipate potential problems in the future, and by testing alternatives they make sound business decisions that improve the company’s profitability and competitive advantage.

Literature cited

Covington, W. W., Wood, D. B., Young, D. L., Dykstra, D. P. and Garrett, L. D. 1988. TEAMS: A decision support system for multiresource management. J. For. 86(8): 25–33.

Environmental Systems Research Institute. 1997. ESRI Shapefile: A technical description. ESRI White Paper, May 1997. 18pp.

Feunekes, U. and Cogswell, A. 1997. A hierarchical approach to spatial forest planning. IN Proceedings of the 1997 Symposium on Systems Analysis in Forestry. Traverse City MI, May 28-30, 1997. In press.

Jamnick, M. S. and Walters, K. R.1993. Spatial and temporal allocation of stratum–based harvest schedules. Can. J. For. Res. 23: 402 – 413.

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.18.31p.

Murray, A. T. 1999. Spatial restrictions in harvest scheduling. For. Sci. 45(1): 45–52.

Nelson, J. D. and Brodie, J. D.1990. Comparison of a random search algorithm and mixed integer programming for solving area–based forest plans. Can. J. For. Res.20: 934 – 942.

Remsoft Inc.1996. Design and development of a tactical harvest blocking/scheduling tool. Final Report to Canadian Forest Service, Pacific Region. March 1996.

Walters, K. R. 1996. An overview of the Woodstock Forest Modeling System. The Compiler. Fall 1996.