A HIERARCHICAL APPROACH TO SPATIAL FOREST PLANNING

Ugo Feunekes and Andrew Cogswell[1]

ABSTRACT. —This paper describes a hierarchical approach to spatial forest planning developed by Remsoft that can be used to create harvest blocks that meet maximum opening, green-up, and other constraints. The approach relies heavily on decisions and trade-offs made during long-term strategic analysis to simplify the blocking phase and to ensure that the harvest blocks generated meet long-term objectives and constraints. A series of case studies illustrates the implementation of the system on the ground and highlights key components of the approach.

INTRODUCTION

The practice of forestry is increasingly coming under the regulatory influence of many different organizations, from environmental watchdogs to Federal agencies (e.g., NEPA). Concerns over the impacts of forest operations on the forest landscape have compelled a number of agencies to take initiatives to address these issues. One result of this has been the definition of explicit limits on how forest operations should be implemented and the establishment of procedures to measure compliance with these guidelines. The American Forest & Paper Association, for example, has taken a proactive approach by outlining guidelines for member forest companies to implement sustainable forestry practices (Wallinger 1995), many of which represent a fundamental change in how forest planning and operations are practiced. In addition, many of these guidelines have a significant impact on both present and future harvest volumes.

At the same time, forest management planning, particularly harvest scheduling under spatial constraints, has witnessed notable advances in the techniques and approaches used to solve the problem. Nevertheless, practical, comprehensive procedures and tools that aid forest managers in developing sustainable, spatially feasible management plans at an operational-scale are still rare. This is because the majority of techniques cannot deal with problems on the scale faced by forest managers—typically hundreds of thousands of polygons. Also, issues surrounding the actual implementation of a planning system have not yet been adequately addressed.

This paper describes an approach developed by Remsoft and used by a number of forest companies throughout North America to produce spatially feasible harvest schedules that meet long-term sustainability criteria. The approach has been shown to work on both small case studies as well as operational-sized problems.

BACKGROUND

Spatial planning problems are especially difficult to solve for three reasons. First, most scheduling problems involve large numbers of stands and/or harvest blocks. Sheer size limits the techniques that can be effectively applied to the problem. Second, a long-term look is required to address sustainability—usually several rotations. As planning horizons increase, the decision variables and constraints necessary to represent adjacencies increase exponentially. Finally, spatial allocation and scheduling in the second-growth forest is often dubious because of the uncertainty in regeneration responses. When all of these factors are considered together, it is clear that finding a true optimal solution to an unrestricted problem is virtually impossible. As a result, every spatial planning approach has focused on finding good or near-optimal, feasible solutions, but on simplified problems.

[1] Ugo Feunekes, Vice President, and Andrew Cogswell, Analyst, Remsoft Inc. 620 George Street, Suite 5, Fredericton, NB, Canada E3B 1K3.

A number of different techniques have been employed to solve spatially-constrained harvest scheduling problems. These include various mixed-integer programming formulations (Meneghin et al. 1988; Jones et al. 1991; Weintraub et al. 1994; Yoshimoto et al. 1994), binary search or inventory projection models (Baskent 1990), simulated annealing (Lockwood and Moore 1990), and Monte Carlo integer programming (MCIP) (O'Hara et al. 1989; Clements et al. 1990), to name a few.

The size of the problem that can feasibly be addressed and the high cost of finding solutions limit purely mathematical programming techniques (Weintraub et al. 1995). To be solvable in this manner, the problem must be simplified, usually by limiting the planning horizon (Nelson and Brodie 1990), constraint set, or the number of harvest units. Others have combined optimization techniques and heuristics to find very good, near-optimal solutions to spatial planning problems (Yoshimoto et al. 1994; Clements et al. 1990; Weintraub et al. 1995). However, none of these approaches have been demonstrated to work for problems on an operational scale.

Simulated annealing has several potential advantages over mixed-integer programming, including the ability to model a large number of stands. Given sufficient computational effort, simulated annealing models will theoretically converge to an optimal solution, although the time to convergence may be excessive. Lockwood and Moore (1993) used simulated annealing for a problem with a large number of stands where each stand could only be considered for one treatment over the planning horizon. However, regenerated stands were not considered, thus limiting the process to long-term sustainability issues.

Baskent (1990) used an inventory projection model in combination with an aggregation heuristic to solve large-size spatial scheduling problems over the long-term. While simulation circumvents the problem size limit inherent in purely mathematical programming techniques, the inability to balance multiple product flows is a major drawback. Moreover, any approach that schedules the harvest sequentially is susceptible to future infeasibilities as the number of options is reduced over time (Brodie and Sessions 1991; Remsoft 1996).

Others have tried to solve the spatial scheduling problems in two steps, separating the problem based on the amount of time and spatial detail considered. Nelson et al. (1991) use such a hierarchical approach to develop spatially feasible harvest schedules. First, a stratum-based linear programming (LP) approach is used to solve the long-term scheduling problem subject to forest-wide constraints but without specific spatial detail. Then, in a separate step, the cut is allocated to blocks for the first few periods of the long-term planning horizon and scheduled using a MCIP model. Beyond the problem of manually delineating blocks, the difficulty with this approach is that the only explicit linkage between the strategic and tactical planning models is the allowable harvest. Thus, the tactical solution may be incompatible with more long-term goals.

Jamnick and Walters (1991) also used a hierarchical approach in which the problem was solved in two distinct phases. The first was a strategic planning phase in which long-term (two or more rotations), nonspatial objectives and constraints were evaluated using traditional LP techniques. Stand-level spatial resolution was not considered in this phase. The second phase was a spatial planning phase where harvest activities were scheduled subject to adjacency delay, opening size, and harvest flow constraints over a much shorter time frame (one rotation). What differentiates this approach from others is that the two phases are closely linked by the harvest timings and stand types chosen in the strategic phase (Cogswell 1996). In addition, an automated harvest block generator (Walters 1991) was used, eliminating the need to manually delineate blocks prior to the spatial phase.

REMSOFT'S SPATIAL PLANNING SYSTEM

For a number of years, Remsoft has been refining and adapting the Jamnick and Walters approach, building on its inherent strengths, addressing its weaknesses, and turning the approach into an implementable spatial planning system. Our design criteria included the following: the system had to (1) be flexible to allow it to operate in different jurisdictions; (2) be scalable; that is, it should work on small problems as well as operational-sized problems containing hundreds of thousands of polygons; (3) be theoretically sound in terms of ensuring that long-term outlooks are not ignored; (4) make reasonable demands in terms of hardware, time to solution, and other system requirements; and finally, (5) provide reasonable, near-optimal solutions.

As the system was refined and tested (Walters and Feunekes 1994; Cogswell 1996), one key strength became evident: since most of the problem is solved in the strategic, nonspatial phase, the spatial problem has been sufficiently simplified to allow our design criteria to be met. By using the results of the first phase as the starting point and attempting to implement the strategic harvest schedule, the spatial phase inherits all of the outcomes of the strategic analysis without explicitly recognizing the individual constraints. Short- and long-term trade-offs among silvicultural levels, habitat levels, spatial trade-offs, product flows etc., have already been evaluated and optimized and are accounted for in the harvest schedule.

For example, in a strategic model it may be determined that to produce an optimal flow of a certain product, one should harvest 1,000 acres of 40-year-old pine in year five, followed by an additional 500 acres in year six. It is irrelevant to the strategic model that there are in fact 12 forest stands within the 1,500 acres of scheduled area, or which of the 12 are harvested in year five versus year six. All that is relevant is that 1,000 acres is cut in year five followed by 500 acres the following year. If this occurs, then the objectives and constraints will be satisfied.

In the spatial phase, the problem therefore becomes one of determining which of the 12 stands should be harvested in year five and which should be harvested in year six. At this stage, it is irrelevant why 1,000 acres should be harvested in year five and 500 in year six. Furthermore, if stands in the spatial phase are scheduled to match the optimal timings determined in the strategic phase, all of the constraints from the strategic phase will be met.

The spatial planning system developed by Remsoft includes a forest modeling system; Woodstock (Walters 1993), a harvest block allocation and scheduling tool; Stanley (Remsoft 1996), a tool for forest-stand allocation; and a collection of utilities to generate the required spatial data and view the resulting harvest blocks. Woodstock generates LP matrices using a generalized Model II formulation and produces optimal solutions for the long-term, strategic portion of the harvest scheduling problem. Using the harvest schedule from Woodstock, Stanley allocates forest stands to harvest blocks subject to adjacency, maximum opening size, and harvest flow constraints.

CASE STUDIES

We present three different case studies to demonstrate how the approach is implemented, the types of information that can be generated using the system, approximate times to solutions, and the way in which the strategic model can be used to control the spatial model. In particular, the case studies will illustrate the flexibility of the approach and its ability to address spatial planning problems at multiple scales. In all examples, Woodstock generated optimal harvest schedules, and Stanley produced spatially feasible harvest block patterns. Solutions were generated on a Dell Dimension XPS Pro 180n personal computer (180 MHz Pentium Pro) with 64 Mbytes of memory, running Windows NT. It is worth noting that all of the examples required no more that 16 Mbytes to solve.

Case One

This example illustrates the effect of varying spatial regulations on final blocked solutions. The case study is drawn from a moderate-sized pine forest located in the Southern United States. The tract covers approximately 212,000 acres and is comprised of 31,917 polygons for an average of about 6.6 acres per polygon. The forest is a complex mosaic of large pine plantations, cypress ponds, and natural pine. Many of the older plantations exceeded 1,000 acres, well beyond the 120-acre opening size limit, and the geographic data set had been preprocessed to subdivide the large polygons into a several smaller pieces. Figure 1 illustrates a small section of this forest.

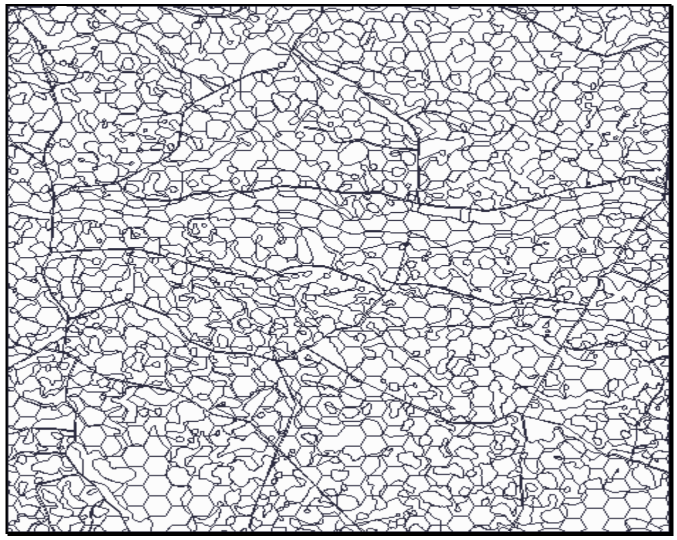

Figure 1. —Forest used in case one. Note the hexagon grid used to divide stands.

In this case, the strategic objective was to maximize present net value subject to even-flow harvest and ending inventory constraints over a 40-year planning horizon. Stands were eligible for thinning and clearcut logging, as well as preharvest cultural treatments such as herbicide application. The LP matrix generated by Woodstock was 3,300 rows by 34,000 columns and 136,000 non zero elements. Woodstock produced the matrix in about 3 minutes and CWHIZ (Ketron 1992) found an optimal solution in about 2 minutes.

While a significant portion of the present net value was generated from commercial thinning, Stanley scheduled only clearcuts. When commercial thins were included in the Stanley runs, loss due to blocking was reduced by almost half (e.g., run 9 scores were increased to almost 91 percent of optimal). However, commercial thins had much less restrictive spatial constraints than clearcuts and were intentionally omitted so as to extend the effects being demonstrated. Values for each of the runs in Table 1 were generated by running Stanley for approximately 5 minutes; within that time Stanley generated approximately 500-600 different layouts, each time retaining the layout that yielded the highest score.

Table 1. —Results of 11 Stanley runs using different minimum and maximum block sizes and adjacency delays

Run	1	2	3	4	5	6	7	8	9	10	11
Min. block (acres)	0	5	10	20	10	10	10	10	10	10	10
Max. block (acres)	-	-	-	-	120	120	120	120	120	500	-
Adjacency delay (years)	0	0	0	0	1	2	3	4	5	5	5
Impossible area (%)	0.0	0.9	2.2	5.1	2.2	2.2	2.2	2.2	2.2	2.2	2.2
Max. block (acres)	1093	1270	867	753	120	120	120	120	120	499	890
Avg. block (acres)	60	78	92	115	53	39	44	52	52	75	90
Score (%)	99.1	98.2	96.6	94.4	95.3	94.5	92.7	89.3	81.3	89.3	89.9

Since most of the harvest scheduling problems have already been solved in the strategic phase, Stanley requires very few input parameters. Those that seem to have the greatest effect on solutions are adjacency delays and the minimum and maximum acceptable opening size (Jamnick and Walters 1991). In this forest, harvest blocks were considered adjacent if they were within 300 feet of one another. The blocks were subject to a 5-year green-up delay, a maximum opening size of 120 acres, and a minimum acceptable block size of 10 acres (Table 1, run 9).

In addition to statistics such as minimum, maximum, and average block size, Stanley calculates a score signifying how close it was able to come to the strategic optimal solution. It was expected that Stanley could not exactly match optimal solutions since Woodstock did not consider stand-level spatial constraints. When spatial constraints were imposed and tightened, we expected the scores from Stanley to decrease accordingly. When constraints were light, Stanley generated solutions that were very close to the strategic optimal. For example, in run 6, Stanley blocked 94.5 percent of the optimal even though 2.2 percent of the area was impossible to block. In other words, Stanley was able to block 94.5 percent of the strategic optimal solution using only 97.8 percent of the eligible forest. As constraints became more limiting, scores were reduced (Table 1).

The results of the Stanley runs show two key effects on the final blocked solutions. First, as the minimum acceptable block size increased, so did the area impossible to block (Table 1, runs 1–4). This effect was attributed to the structure of the forest and not the approach and/or algorithms. The impossible area represents stands or collections of stands that cannot be blocked because of their location in the forest. Often these stands are isolated from others by features such as buffers, rivers, or lakes. When Stanley attempts to create harvest blocks, no configuration exists that will allow these stands to be included.

Second, as the maximum opening size decreased and the adjacency delays increased, scores were reduced. To a large extent, the reduction is a function of the forest (age classes and homogeneity of stand types within the forest, for example). However, as the adjacency delay increased relative to stand rotations, the magnitude of the reductions increased. In this forest, a 5-year adjacency delay was approximately one-quarter of the average plantation rotation, the cost of which was much higher than in a forest with longer rotations. Relaxing the maximum opening size limit resulted in higher scores when the adjacency delay was long (Table 1, run 9 vs. run 10). In this situation, Stanley alleviated adjacency conflicts by aggregating conflicting harvest units to form larger blocks with a single timing choice.

Case Two

Case two illustrates how changing the nonspatial model can affect the resulting block layouts. This forest in Western Canada is 3,191 hectares (ha) in size and contains 3,034 individual polygons. It is dominated by large tracts of even-aged lodgepole pine intersected by many trails, roads, and seismic lines (Fig. 2). For analysis, the forest was classified into a northern zone and southern zone. Regulations dictated that cut blocks not exceed 100 ha and that there be a 20-year green-up delay. The minimal acceptable harvest block size was specified as 10 ha.

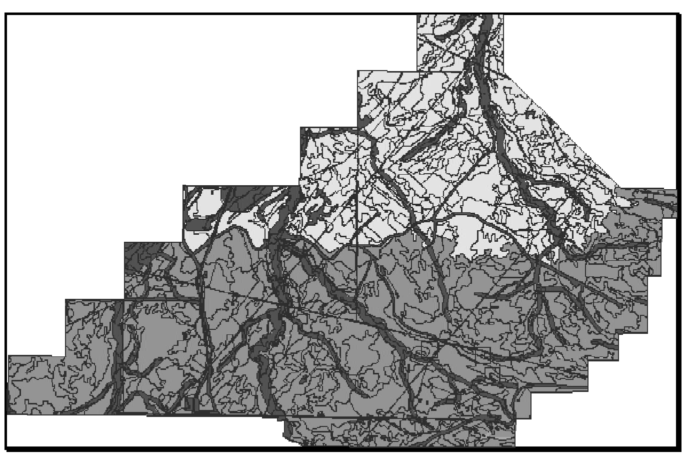

Figure 2. —Forest used in example two. Stands in the north zone are shown in light grey and stands in the south zone are shown in dark grey. Buffers and other excluded areas are shaded black.

The strategic objective was to maximize the first-period coniferous volume subject to an even-flow harvest over a 200-year planning horizon, while maintaining at least 600 ha of mature coniferous forest in every period. Stands were eligible for clear-cut harvesting, after which they could either be naturally regenerated or site-prepared and planted. Two models were generated by Woodstock; harvesting was unrestricted in the first, while in the second, the north zone was inaccessible for the first seven decades. The matrix generated for the unrestricted model contained 2,400 rows by 9,900 columns and 48,000 nonzeros, and the restricted model had 2,200 rows by 8,200 columns and 39,000 non-zeros. Both were generated and solved in under 1 minute. As expected, the restricted model's optimal solution was lower than the first (around 10 percent less).

Both strategic solutions were run through Stanley for 100 years into the future. The only difference between the two runs was the strategic harvest schedule used by Stanley—the spatial constraints were identical in both cases. In the first case, Stanley blocked 95.6 percent of the strategic optimal solution after 5 minutes (over 4,600 alternative layouts evaluated), and blocks were distributed throughout the entire forest (both management zones) in all periods. In the second case, 96.1 percent was blocked after 5 minutes, and no harvest blocks were located in the north zone prior to period seven. It is worth noting here the way that the results were produced. Using the hierarchical approach, we were able to control the general location of blocks by changing the strategic model, which is a relatively simple process, and not the spatial model, which would be a much more complex operation.

Case Three

The final example represents a large operational scale problem representative of the type of problem faced by forest planners today. Case three comes from a large tract of land in Eastern Canada. The forest covers approximately 176,884 ha and is comprised of 88,843 polygons, averaging about 2 ha per polygon. The forest is being intensively managed for softwood pulp and logs. Stands were eligible for several pre-harvest treatments including planting, spacing, and thinning; harvest treatments included clearcut and two-pass logging. The forest is highly fragmented, containing numerous riparian buffers, habitat exclusion zones, and roads. Figure 3 depicts a small section of the forest.

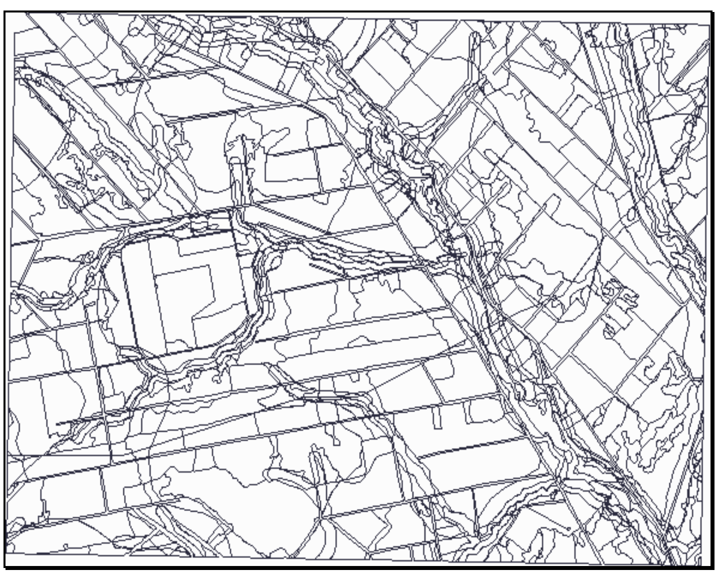

Figure 3. —Forest used in example three. Area shown represents 1 of the 76 maps sheets in the entire forest.

The strategic objective was to maximize spruce-fir harvest volume subject to even-flow harvest and ending inventory constraints over an 80-year planning horizon. Additional constraints limited the production of hardwood volume and the amount of planting and spacing within each planning period. The periodic change in the area within broad forest cover types was also limited to maintain biodiversity levels over the long term. The LP matrix generated by Woodstock was fairly large, approximately 9,800 rows by 71,000 columns and 922,000 nonzeros. Woodstock produced the matrix in about 15 minutes and CWHIZ found an optimal solution after about 35 minutes.

Spatially feasible harvest blocks were created for 25 years into the future, subject to a maximum opening size on clearcuts of 100 ha, a minimum acceptable block size of 5 ha, a 10-year adjacency delay, and maximum harvest flow fluctuation of 10 percent. Two Stanley runs were required to allocate all of the harvest activities to blocks because the spatial restrictions for partial harvest activities were different than those for clearcut harvesting. The first Stanley run (clearcuts) evaluated about 1,000 different block layouts in around 30 minutes while the second run (partial harvest) required about 35 minutes. In both cases, the layout that provided the highest score (nearest to optimal) was kept. When the results from the two runs were combined, Stanley was able to allocate 89.8 percent of the strategic optimal. The final block layout contained 4,000 harvest blocks ranging in size from 5 to 100 ha with an average of 45 ha.

The ability of the approach to deal with problems of this size is well illustrated in this case. To carry out a similar analysis using the more traditional approaches, including manual blocking, would take something in the order of weeks, if not months, to produce a single block layout, rather than the thousands produced in a matter of hours using this approach.

SUMMARY AND CONCLUSIONS

In summary, the system is based on a hierarchical approach introduced by Jamnick and Walters that greatly simplifies the solution of the overall problem. It demonstrates that the spatial complexity of the problem can be increased and solved without increasing the complexity of the spatial algorithm. In effect, the two phases worked together to solve a more complicated problem.

The three case studies presented show that the approach works at various spatial scales on realistic-sized problems and considers long-term trade-offs. Pseudospatial resolution, in which subforest areas are recognized, can easily be accommodated in this planning phase, allowing for the distinction of features such as watersheds, management units, forest districts, etc. An essential factor is that long-term trade-offs and constraints are handled in the strategic phase, and that the software required to allocate stands to specific periods (spatial phase) can therefore be much simpler. Also, since the long- and short-term phases are so closely linked, any changes to the strategic model are directly communicated to the spatial stage.

Work is ongoing to refine the system.

LITERATURE CITED

BASKENT, E.Z. 1990. Spatial wood supply modeling: Concept and practice. M.Sc.F. thesis. Fredericton, NB: University of New Brunswick.

BRODIE, J.D., AND J. SESSIONS. 1991. The evolution of analytical approaches to spatial harvest scheduling. In *Proceedings of the 1991 Symposium on Systems Analysis in Forest Resources*. Gen. Tech. Rep. SE-74. 1991 March 3–7. Charleston, SC. Charleston, SC: USDA, Forest Service, Southeastern Forest Experiment Station.

CLEMENTS, S.E., P.L. DALLAIN, AND M.S. JAMNICK. 1990. An operational, spatially constrained harvest scheduling model. *Canadian Journal of Forest Resources* 20: 1438–47.

COGSWELL, A. 1995. Evaluating the spatial feasibility of future timber harvest allocations using a repeated planning process. B.Sc.F. thesis. Fredericton, NB: University of New Brunswick, Faculty of Forestry and Environmental Management.

JAMNICK, M.S., AND WALTERS, K.R. 1991. Harvest blocking, adjacency constraints and timber harvest volumes. In *Proceedings of the 1991 Symposium on Systems Analysis in Forest Resources*, 255–261. Gen. Tech. Rep. SE-74. 1991 March 3–6. Charleston, SC. Charleston, SC: USDA Forest Service, Southeastern Forest Experiment Station.

JONES, J.G., B.J. MENEGHIN, AND M.W. KIRBY. 1991. Formulating adjacency constraints in linear optimization models for scheduling projects in tactical planning. *For. Sci*. 37: 1283–97.

KETRON MANAGEMENT SCIENCE. 1992. *MPSIII mathematical programming system: CWHIZ linear programming optimizer*. Arlington, VA.

MENEGHIN, B.J., M.W. KIRBY, AND J.G. JONES. 1988. An algorithm for writing adjacency constraints efficiently in linear programming models. In *The 1988 Symposium on Systems Analysis in Forestry Resources*, eds. Kent, B. and L.S. Davis, 46–53. Gen. Tech. Rep. RM-161. 1988 March 29–April 1. Asilomar, CA. USDA Forest Service, Rocky Mountain Forest Range Experiment Station.

MOORE, T.G.E., AND LOCKWOOD, C.G. 1990. *The HSG wood supply model: Description and user's manual.* Info. Rept. PI-X-98. Petawawa Nat. Forestry Inst., Forest Canada.

NELSON, J., AND J.D. BRODIE. 1990. Comparison of a random search algorithm and mixed integer programming for solving area-based forest plans. *Canadian Journal of Forest Resources* 20: 934–42.

NELSON, J., J.D. BRODIE, AND J. SESSIONS. 1991. Integrating short-term, area-based logging plans with long-term harvest schedules. *Forest Science* 37(1**)**: 101–22.

O'HARA, A.J., B.H. FAALAND, AND BARE, B.B. 1989. Spatially constrained harvest scheduling. *Canadian Journal of Forest Resources* 19: 715–24.

REMSOFT INC. 1996. Final report: Design and development of a tactical harvest blocking/scheduling tool. Study funded through the Canada-British Columbia Partnership agreement of Forest Resource Development: FRDA II.

WALLINGER, S. 1995. Sustainable forestry initiative. *Journal of. Forestry*. 93(1): 17–19.

WALTERS, K.R. 1991. Spatial and temporal allocation of strata-based harvest schedules*.* M.Sc.F thesis. Fredericton, NB: University of New Brunswick.

WALTERS, K.R. 1993. Design and development of a generalized forest management modeling system: Woodstock. In *Proceedings of the International Symposium on Systems Analysis and Management Decisions in Forestry*, 190–196. 1993 March 9–12. Valdivia, Chile.

WALTERS, K.R. AND FEUNEKES, U. 1994. A hierarchical approach to spatial planning: a report card. In *Proceedings of the GIS'94 Symposium*, 1994 February.Vancouver, BC.

WEINTRAUB, A., F. BARAHONA, AND R. EPSTEIN. 1994. A column generation algorithm for solving general forest planning problems with adjacency constraints. *Forest Science* 40(1): 142–61.

WEINTRAUB, A., G. JONES, M. MEACHAM, A. MAGENDZO, A. MAGENDZO, AND D. MALCHUK. 1995. Heuristic procedures for solving mixed-integer harvest scheduling **-** transportation planning models. *Canadian Journal of Forest Resources* 25: 1618–26.

YOSHIMOTO, A., J.D. BRODIE, AND J. SESSIONS. 1994. A new heuristic to solve spatially constrained long-term harvest scheduling problems. *Forest Science* 40(3): 365–96.