INCORPORATING SPATIALLY EXPLICIT OBJECTIVES INTO FOREST MANAGEMENT PLANNING

by

KEVIN CROWE

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
in
THE FACULTY OF GRADUATE STUDIES

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA
October 2004.

ABSTRACT

The increased incorporation of spatially explicit objectives into forest management planning has arisen from a concern over the ecological consequences of landscape-scale disturbance patterns through harvesting. Given the complexity of the ecosystems of forested landscapes, and our incomplete understanding of them, forest managers now commonly design plans to conserve landscape-scale biodiversity through the emulation of natural disturbance patterns. Harvest-scheduling is thereby constrained to imitate not only the aspatial age-class and cover-type distributions of a forest under natural disturbance, but also the patch-sizes and shapes of disturbed and undisturbed forest.

Forest management planning typically requires the use of optimization models because the most efficient allocation of scarce resources is a central economic objective. Incorporating spatially explicit objectives into such models requires that decision variables be binary. This is because, in a spatially explicit plan, a forest stand must be either harvested or not. Hence, integer programming models are needed, and such models are notoriously difficult to solve computationally.

The objective of this research has been to make significant advances in formulating, solving, and understanding three difficult forest planning problems involving spatial objectives.

The first is a tactical planning problem where the objective is to maximize the net present value of a harvest-schedule, subject to the spatial constraint that stands may not aggregate to form harvest-openings greater than a maximum area. In Chapter II, two integer programming models were formulated and solved, using the branch and bound algorithm. It was found that: a) the number of decision variables, and b) the number of opening constraints, ultimately restricts this method from applicability to larger problem instances. Given these limitations, a metaheuristic algorithm, simulated annealing was evaluated in Chapter II. Using the branch and bound algorithm's solutions as upper bounds, the quality of solutions found by the metaheuristic was evaluated. The mean objective function value was within 5% of the optima. Problems instances ranged in size from 1,269 to 36,270 binary decision variables.

The second problem, treated in Chapter IV, concerns the efficient allocation of cutting rights among competing mills within the same management unit. A mixed integer goal programming model was formulated and applied to the Kootenay Lake Timber Supply Area of British Columbia. It was concluded that the model is can a useful tool by which to interactively explore British Columbia's appurtenance policy.

The third problem, treated in Chapter V, is a strategic planning problem: determining the optimal, sustainable rate of harvest while selecting spatially explicit old growth reserves. A mixed integer programming model was formulated and tested on three forests. It was concluded that the formulation appears to be integer-friendly, having solved problems instances containing up to 91,000 decision variables.

The general conclusion of this work is that integer programming is a powerful paradigm by which to incorporate the complexities of spatially explicit objectives within the pragmatic constraints of forest management planning.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

AKNOWLEDGEMENTS

I would like to acknowledge the generous, prudent, and invaluable guidance provided to me by my academic supervisor, Dr. John Nelson. I would also like to acknowledge and thank Dr. Emina Krcmar and Dr. Thomas Maness for serving on my committee. I also wish to thank: Dr. Maurice Queyranne, for having taught me a rigourous course on integer programming; and Tim Shannon, for preparing many GIS data-sets for me over the years.

Chapter I

Introduction

1.1 Introduction to Area of Research

The discipline under which to categorize this dissertation is Forest Resources Management. This specialized discipline exists because of the role forests play in defining the ecological, economic, and social environment of mankind. Ecologically, the world's 40 million km^2 of forests provide habitat to at least 80% of the worlds remaining terrestrial biodiversity (Brooks 1993); and they function as a major carbon sink to regulate the global climate. Economically, almost 1.3 billion m^3 of roundwood is harvested annually from world's forests (UNECE 2003). In Canada alone, 200 million m^3 per year are harvested, generating over $68 billion in sales and providing direct or indirect employment for 1 in 16 Canadians (Natural Resources Canada 2004). Socially, there has been a gradual movement, beginning with grass-roots organizations, but ascending to national and international levels of influence, to proclaim that the ecological benefits of the forest must be conserved. The model of 'spaceship earth' now informs social values; and since forests are a major component of this model, increased public scrutiny and government regulation in the management of forest resources has occurred and will continue. The great significance of conserving the flow of benefits from our forests underlines the need to plan our forest management more thoughtfully, and this requires research into new methods of planning.

The topic of this dissertation is clearly expressed in its title, *Incorporating Spatially Explicit Objectives into Forest Management Planning*. A direct approach to introducing this area of research can therefore proceed by decomposing the title into a subset of broad questions to be addressed; namely:

- What is forest management planning?
- What are spatially explicit objectives? Why are they important?
- How are they incorporated into forest management planning?

1.1.1 What is forest management planning?

The purpose of forest management planning is to support decisions on the assignment of management prescriptions to the stands or land-units of a forest over time (Davis *et al.* 2001). These decisions are not made in isolation from one another. Instead, the set of prescriptions is regarded as a whole, *i.e.*, it is evaluated on its estimated potential to satisfy the aggregate of social, economic, and ecological objectives assigned to the forest as a whole over multiple scales of space and time. The forest management planning problem is characterized by considerable complexity, involves varying degrees of uncertainty, and supports decisions often carrying great economic, ecological, and social consequences. These three characteristics—complexity, uncertainty, and consequence-- warrant a brief expansion in this introduction.

The complexity of the forest management planning problem has several causes. *First*, there are multiple, often conflicting objectives. For example, an economic objective of maximizing a firm's profitability may conflict with an ecological objective of preserving large, contiguous areas of old-growth. Similarly, the social objective of securing an even-flow of timber from the forest may conflict with the economic objective of harvesting a forest as quickly as possible. *Second*, the spatio-temporal scale of planning adds to this complexity. Since trees mature relatively slowly, and forests are composed of many stands, forecasts are commonly made for thousands of stands over hundreds of years. Consequently, the number of decision-variables, representing potential management treatments over time, can be great in number. *Third*, the planning process is interdisciplinary; *i.e.*, various specialists are needed to define the quantitative indicators by which the satisfaction of objectives is measured. It is a complex undertaking both to define an analytically common ground on which specialists may communicate and to select indicators relevant to planning objectives and the availability of data. *Finally*, depending on the regulatory environment, several decision-makers are needed; and this entails that value tradeoffs be articulated. Ultimately, this is a political exercise, but the process can be informed and directed through quantified estimates of how much these values conflict.

The second major characteristic of forest management planning is uncertainty. This stems from many sources; examples are: 1) the accuracy of data representing current forest conditions; 2) the validity of the growth and yield models; 3) the unpredictability of natural disturbances; 4) the fluctuation of market-values for wood products; 5) the changes in values of decision-makers; and 6) the unknown consequences of climate change. Clearly, uncertainty must be formally considered by decision-makers.

Finally, the consequences of forest management planning are of considerable importance. Decisions on how much to cut, where, and when, carry tremendous economic, social, and ecological consequences. Howard (1988) argues that decisions of great consequence, complexity, and uncertainty warrant the development of decision support models within the context of a formal decision-making framework. I will now describe the role of these models within forest management planning.

There is a general framework within which decision support models fit, regardless of the genre of problem. This is illustrated in Figure 1.

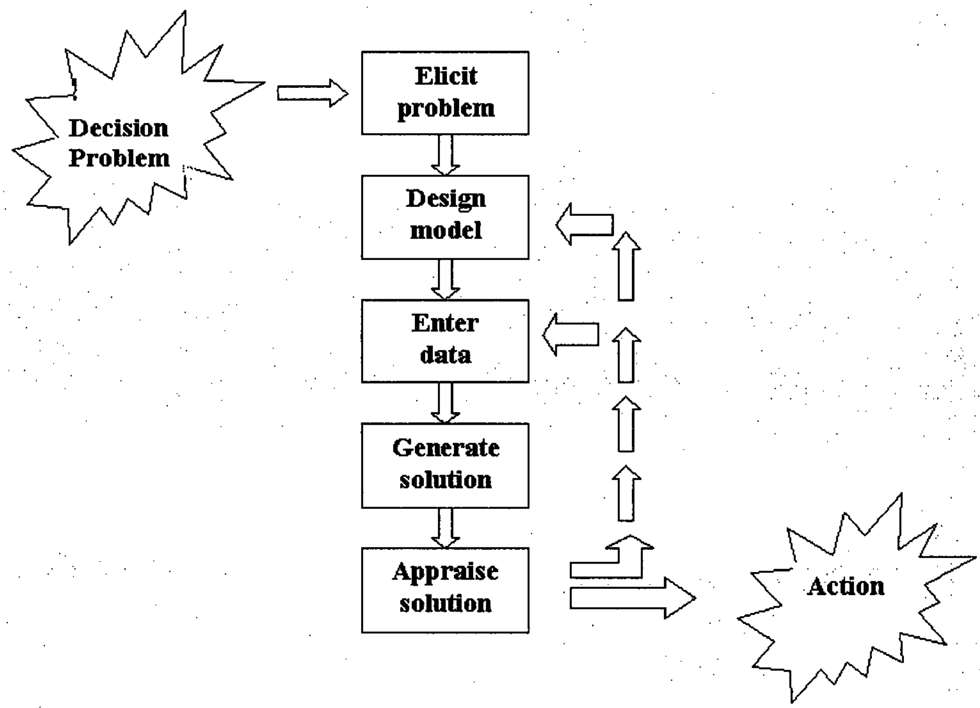

Figure 1.1: Framework within which decision support models operate.

It is noteworthy that the ultimate purpose of the decision-making framework resides in action—not knowledge. Bunnell (1989) usefully distinguishes two types of models: 1) those which are *primarily* used to extend human understanding about the world; 2) those which are *primarily* used to provide information in support of decisions. Decision support models clearly fall into the latter category. Their product is information, not knowledge. Nevertheless, in the design of such models and the appraisal of their solutions, it is not uncommon to gain valuable insights into the fundamental properties of the managed system. Such insights can be of great value to managers.

In Figure 1, the iterative process of appraising a solution and redesigning the model is often referred to as model validation; *i.e.*, testing the hypothesis that the model is a sufficiently precise representation of the essential features of the real system, and that the solutions obtained from the model are valid for the real problem (Hillier and Lieberman 2001). In forest management planning, with its long planning horizons and consequent uncertainty, the model validation process is quite difficult. Therefore, the solutions obtained from these models are accepted as valid for the real world, in very limiting ways.

To clarify and organize the degrees of uncertainty in forest management planning, a hierarchical planning framework of three levels, each with different planning models, is commonly used (Gunn 1991, Martel *et al.* 1998, Davis *et al.* 2001). At the highest level, there are strategic, long-term plans (*i.e.*, greater than one rotation), characterized by great uncertainty and aimed at answering a few general questions; *e.g.*, given our ecological, economic, and social objectives, what is a sustainable rate of harvest for a given forest? The second level in this hierarchy is tactical, which involves mid-term planning (less than one rotation), characterized by less uncertainty, and addressing more particular questions; *e.g.*, given a rate of harvest, and a set of ecological and social constraints, what is the most profitable set of stands to access and harvest? The final level is operational, involving short-term plans with relatively less uncertainty. Operational plans schedule machines and labour to build the roads and harvest the stands laid out in the tactical plan. Uncertainty is addressed in the hierarchical planning framework through what Gunn (1991) calls the *rolling planning horizon*; *i.e.*, plans are renewed well before their planning horizon has passed. The overall result is that a forest management plan is not a static template, but a dynamic, ongoing, adaptive process.

1.1.2 What are spatially explicit planning objectives? Why are they important?

Spatial relationships have always been important in forest management planning. Traditional sustained yield management, for example, had been concerned with finding the most economically efficient way to access and time the harvest of a mosaic of stands dispersed across a forest; and of assigning the right logs harvested to the rights mills. But the evolution of sustained yield management into ecosystem management has introduced new types of spatially explicit objectives and placed great importance upon their achievement—even when these objectives are in acute conflict with traditional economic ones. Hence, it is increasingly important that decision support tools be designed to satisfy both ecological and economic objectives. The underlying assumptions of ecosystem management, and the necessity for some of its objectives to be spatial in nature are now reviewed.

Ecosystem management is both a practice and philosophy aimed at maintaining or enhancing the integrity of an ecosystem while providing resources, products, or non-consumptive values for humans (Gordon 1994). As a philosophy, it envisions a deliberated and managed coexistence of man and nature (Davis *et al.* 2001).

Ecosystem management entails managing the forest at multiple scales. This stems from the very nature of ecology: understanding scale is central to understanding ecology (Levin 1992). The Scientific Panel for Sustainable Forest Practices in Clayoquot Sound (1995) stated that 'planning at a variety of spatial and temporal scales is critical at all stages of forest ecosystem management.' Predicting and planning at a variety of scales requires a conceptual framework; and for forested landscapes, this is provided by ecological hierarchy theory (Eng 1998). The basic premise of hierarchy theory, when applied to landscape ecology, is that hierarchically organized systems can be divided or decomposed into discrete functional components operating at different scales (Urban *et al.* 1987). An example of a hierarchical structure of a forest is, moving from lowest to highest: gap => stand =>watershed =>landscape. Events that occur at a certain level have a characteristic frequency and spatial scale. In general, the higher levels are usually comprised of larger units and operate more slowly than the lower levels (Eng 1998). It is also important to note that the *higher levels tend to constrain the lower levels* (Urban *et al.* 1987). Hence, altering substantial portions of our forested landscape, through harvesting, impacts all scales.

Given the ecological consequence of landscape scale disturbance through harvesting, planning for spatial pattern in forest management is important because it has ecological consequences on *population dynamics* and *ecosystem processes* (Saunders et al. 1991, Forman 1997, Eng 1998, Haila 1999, Spies and Turner 1999). There are many indices of spatial patterns, but most fall within one of three categories (O'Neill *et al.* 1988): 1) those which measure the dimensions of individual elements or patches; 2) those which measure landscape composition (e.g., abundance of patches); and 3) those which measure the spatial arrangement of patches.

Notwithstanding the importance of spatial patterns, there is some difficulty in defining *which patterns* are most relevant to the objectives of ecosystem management. I note two reasons for this difficulty. *First*, it not yet fully known, for all processes and all organisms in a forest, the conditions under which spatial heterogeneity is and is not important (Spies and Turner 1999). Landscape ecology is, after all, a relatively new undertaking, and the complexity of forest ecosystems ensures that it will be many years before science has caught up with the ambitious objectives of ecosystem management. *Second*, spatial forest patterns are meaningful from the perspective of particular organisms (Bunnel 1999, Spies and Turner 1999); *e.g.*, what may appear as fragmented habitat from a human perspective may also appear as continuous from the perspective of another species. In other words, even if it were possible to know all of the relevant spatial patterns for all organisms in a forest, it would be exceedingly difficult to consider and satisfy the perspectives of all species simultaneously.

As a result of these two difficulties, the natural disturbance model of sustainable forest management has been developed (Hunter 1993). By emulating the results of natural processes as closely as possible, it is hypothesized that management through emulation of natural disturbance will minimize the negative impacts of forest harvesting on the biodiversity of forest ecosystems. This underlying assumption is made explicit by the coarse and fine filter analogy first proposed by the Nature Conservancy (1982); *i.e.*, the maintenance of the distribution of habitat types that would occur under natural conditions will satisfy the habitat requirements for most species while recognizing that more specific management prescriptions (fine filter) are needed for species of special concern.

The emulation of natural disturbance patterns has become a model for many forest managers in both Canada (e.g., Canadian Institute of Forestry 2003, Ontario Ministry of Natural Resources 2001, British Columbia Ministry of Forests 2001) and the United States (Hunter 1990). The general approach has been to use historical fire history data and simulation models to estimate the natural, historical, forested landscape pattern against which to compare the landscape patterns of a managed forest. There are many such simulation models currently in use (see Mladenoff and Baker 1999). Harvest scheduling is thereby constrained to imitate not only the aspatial age-class and cover-type distribution of a forest under natural disturbance, but also the patch-sizes of disturbed and undisturbed (i.e., old growth) forest. While idealized landscapes are used as long-term goals, one of the more acute challenges concerns working with the pattern of what has been inherited from both nature and previous management (Bettinger and Sessions 2003).

1.1.3 How are spatial objectives incorporated into forest management planning?

The incorporation of spatially explicit objectives into forest management planning typically (but not exclusively) requires the use of optimization models because the most efficient allocation of scarce resources is a central economic objective in forest management planning (Davis *et al.* 2001). In such optimization models, the decision variables needed to model spatial attributes, such as the size and shape of harvest openings or reserve patches, must be constrained to being binary, not real numbers. This is because a stand is either harvested or it is not-- it cannot be fractionally harvested (Bettinger and Chung 2004). Hence, integer programming models are required and such models are notoriously difficult to solve computationally;

i.e., the computing time needed to generate solutions typically increases exponentially in relation to the number of decision variables (Wolsey 1998, Williams 1999).

For many years, the principal solution technique used in solving integer programming models was the linear programming-based branch and bound, introduced by Land and Doig (1960). Unfortunately, the available software and computers restricted the application of many integer programming models to impractically small problem instances; and since many real-world problems are large, disillusionment with the integer programming paradigm was common (Bixby *et al.* 2004). Over the last 15 years, however, research has progressed along two paths in response to the computational challenges posed by integer programming models (Reeves 1993, Michaeliwicz and Fogel 2000). Along one path, a backlog of theory on integer programming was translated into numerous algorithmic refinements of the branch and bound approach (see Bixby *et al.* 2004). As a result, increasingly large problems can now be solved optimally using commercial solvers. Along another path, an explosion of research into metaheuristic algorithms has occurred (see Osman and Laporte 1996 for bibliography). Metaheuristics are general integer programming techniques not dedicated to the solution of a particular problem; but rather, are designed to be flexible enough to handle many different problems. Metaheuristics have rapidly demonstrated their usefulness and efficiency in solving many large and difficult integer programming problems (Glover 1986, Herz and Widmer 2003). The major shortcoming of metaheuristic algorithms, however, is that they are approximations of optimal solutions, and neither guarantee optimality nor provide any indication of how close their solutions are to being optimal (Reeves 1993).

Hence, insofar as this dissertation is concerned with the incorporation of spatially explicit objectives into forest management planning, it requires attentiveness to the computational challenges posed by integer programming models in general, and the appropriate use of approximation *versus* optimization algorithms in particular. My direction on the appropriate use of these algorithms follows Wolsey's (1998) common-sense advice on this matter; *viz.*, that if at all possible, an *optimal* solution to an integer programming problem should be computed; and that barring this possibility, an approximation to the optimal solution should be found, on the condition that an estimation of the proximity of this solution to the optimal is also formed.

1.2 Problems Treated in this Dissertation

The objective of my research into forest management planning under spatial objectives has been to make a significant advance into formulating, solving, and understanding three difficult planning problems. These problems are:

1. Generating optimal and approximately optimal solutions to tactical planning problem of simultaneously scheduling an optimal flow of timber and optimally aggregating forest stands into harvest-blocks with a limitation on block area;
2. Allocating cutting rights, in the form of discrete areas of a forest management unit, among competing firms with the objective of optimizing economic

objectives—within the context of sustaining landscape-scale ecological values ; and

3. Generating optimal solutions to the problem of simultaneously scheduling the harvest of timber and selecting spatially explicit old-growth reserves.

My treatment of these problems is presented in Chapters II to V of this dissertation, which I now describe.

1.2.1 Description of Chapter II

In this chapter, I present two formulations of the integer programming problem of scheduling an optimal flow of timber subject to limiting the size of the harvest-blocks. These formulations are tested on 30 tactical planning problems using the branch and bound algorithm.

There have been two approaches to modeling this problem: the unit-restricted model and the area-restricted model (Murray 1999). In the unit-restricted model, the boundaries of all potential harvest-openings are predefined. In the area-restricted model, these boundaries are *not* predefined; instead, polygons may be aggregated to form cut-blocks, of a limited size, during the search for an optimal harvest-schedule. In the context of forest management planning, the area-restricted model holds greater appeal than the unit-restricted model because the definition of cut-block boundaries occurs in the context of the search for an optimal flow of timber; it can therefore generate solutions with a higher net present value than those generated using an unit-restricted model. In fact, Murray and Weintraub (2002) rightly observe that the objective function value of optimal solutions to the unit-restricted model can function as lower bounds for the same problem instances when treated with the area-restricted model.

In my analysis of these formulations, particular attention was paid to exploring the effects of problem size, initial forest age-class distribution, and the ratio of polygon area to the area of maximum allowable harvest-opening. The results and discussion of this research both explore and interpret the limitations of this approach, and also provide a set of benchmarks by which to evaluate the effectiveness of metaheuristic algorithms in solving the area-restricted harvest-scheduling model.

1.2.2 Description of Chapter III

In this chapter, I use the benchmarks established in Chapter II to evaluate the effectiveness of a metaheuristic algorithm, simulated annealing, in solving the area-restricted harvest scheduling model. As noted earlier, the major shortcoming of metaheuristic algorithms is that they neither guarantee optimality nor provide any indication of how close their solutions are to being optimal (Reeves 1993).

The objective of this research is to provide an empirical worst-case analysis of the effectiveness of a metaheuristic approach to solving a variety of problem instances; *i.e.*, to gain some idea on how well it performs in general, and in what circumstances it does relatively well or relatively badly. I describe this as a worst-case analysis because the implementation of the simulated annealing algorithm used, although shown to be effective (Bettinger *et al.* 2000), is but one of many possible implementations. That is, the extraordinary flexibility of the metaheuristic approach prevents me from extending judgement beyond its explored potential.

In the context of research on this problem, such analysis, though fundamental, is novel. The area-restricted model was initially introduced by Lockwood and Moore (1993) as a problem to be solved using simulated annealing. Since its introduction, the applications of various metaheuristic algorithms to the area-restricted model have been extensively researched (Bettinger *et al.* 1997, Ohman and Eriksson 1998, Van Deusen 1999, Liu *et al.* 2000, Richards and Gunn 2000, Van Deusen 2001, Sessions and Bettinger 2001, Baskent and Jordan 2002, Richards and Gunn 2003, Caro *et al.* 2003), but these implementations have not been fully evaluated through comparison with optimal solutions. In Chapter III, the comparison of results from applying the simulated annealing algorithm to the benchmarks established in Chapter II indicate that this algorithm can find excellent solutions to the area-restricted model; but that it tends to perform less well on larger problem instances than on smaller ones.

1.2.3 Description of Chapter IV

In this chapter, I present an approach to solving the problem of allocating cutting rights among competing firms with the objective of sustaining landscape-scale ecological values and optimizing the allocation of standing timber. This problem, although not so extensively studied as the harvest-scheduling problem, is nonetheless highly relevant to forest management planning, especially in British Columbia; for it incorporates and allows for evaluation of the recently proposed policy of allocating cutting rights to firms that do not own wood-processing facilities.

My treatment of this problem involves two steps: first, a forest management unit is divided into a set of landscape units, and for each landscape unit, a spatially explicit harvest-schedule is generated, subject to spatially explicit landscape-scale ecological objectives. Second, each landcape unit, and its forecast flow of timber, is assigned to one of many local licensees, using a mixed integer goal programming model. The optimal assignment of a landscape unit to a licensee is determined by its ability to satisfy multiple conflicting objectives for volume, species, log size, hauling cost, and seasonal access. The targets for theses objectives are derived from the demands of each licensee's processing facilities. It is an extension of the classic forestry problem of assigning the "right log to the right mill" (Pearse and Sydneysmith 1966)) constrained and frustrated by the requirement that these logs be assigned as discrete landscape units. The model is solved using the branch and bound algorithm, and is designed to facilitate the interactive exploration of decision makers' preferences.

The constraint that standing timber be assigned to mills in the form of discrete zones of standing timber is a consequence of the Province of British Columbia's long-standing appurtenance policy (Pearse 2001); *viz.*, that allocation of cutting rights be restricted to firms with wood-processing facilities. This policy is currently undergoing re-evaluation, and the option of gradually introducing competitive auctioning of standing timber to logging firms without wood-processing facilities is under consideration. The mixed integer goal programming model described in Chapter IV incorporates this alternative by allowing the assignment of landscape units to logging firms, which in turn redistribute the scheduled flow of logs from these landscape units, to multiple licensees based on optimal satisfaction each licensee's multiple objectives. This refinement of the mixed integer goal programming model allows not only for evaluation of the increased efficiency with which logs can be redistributed when the appurtenance policy is relaxed, but also for the selection of those zones best suited for competitive auction.

1.2.4 Description of Chapter V

In this chapter, I present a mixed integer programming formulation of a model designed to support the strategic-level planning challenge of selecting, or in some cases recruiting, patches of old growth reserves while optimizing the traditional economic objectives of harvest-scheduling. In the context of forest management planning, this model is relevant to many forest managers who are now caught in transition from the sustained yield paradigm to the ecosystem management paradigm, and are managing traditionally regulated forests without sufficiently representative areas of contiguous old ecosystems.

The model is designed to generate optimal solutions using the branch and bound algorithm. Binary decision variables are required to represent each potential reserve, and the model's usefulness hinges upon its ability to solve large problems with a large number of potential reserves. The formulation is tested on several forests, ranging in area from 16,000 to 800,000 ha. Optimal solutions are found for each problem instance tested, including one forest requiring 91,000 binary decision variables.

In the context of other research on this problem, Chapter V is relevant because optimal solutions to this problem have thus far been limited to extremely small problem instances (Hof and Joyce 1992, Hof and Joyce 1993, Rebain and McDill 2003). Given its applicability to large problems with many binary decision variables, the formulation appears to be integer-friendly.

1.3 Format of Dissertation

The format of this dissertation is manuscript-based; *i.e.*, Chapters II to V are structured as journal articles, with separate introductions, methods, and conclusions. Chapter II has been published in the Canadian Journal of Forest Research (Crowe *et al.* 2003). Chapter IV's preliminary results were presented at the Canadian Operational Research Society's annual conference in June 2003. Chapters III, IV, and V are presented in draft form, and will each be submitted for publication as journal articles.

Literature Cited

Baskent, E.Z., and Jordan, E.A. 2002. Forest landscape management using simulated annealing. Forest Ecology and Management 165:29-45.

Bettinger, P., and Chung, W. 2004. The key literature of, and trends in, forest level management planning in North America, 1950-2001. The International Forestry Review 6:40-50.

Bettinger, P., and Sessions, J. 2003. Spatial forest planning: to adopt, or not to adopt? Journal of Forestry 101(2)24-29.

Bettinger, P., Sessions, J., and Boston, K. 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. Ecological Modeling 94: 111–123.

Bettinger P., Graetz, D., Boston, K., Sessions, J., and Chung, W. 2000. Eight metaheuristic planning techniques applied to three increasingly difficult wildlife planning problems. Silva Fennica 36(2):561–584.

Bixby, R.E., Fenelon, M., Gu, Z., and Rothberg, E. 2004. MIP: theory and practice—closing the gap. http://www.caam.rice.edu/~bixby/default. (accessed 12 Jan/04)

British Columbia Ministry of Forests. 1995. Biodiversity Guidebook. Government of British Columbia, Victoria.

Brooks, D.J. 1993. U.S. forests in a global context. U.S. Dept. of Agriculture, Forest Service. Fort Collins, CO.

Bunnell, F.L. 1989. Alchemy and uncertainty: what good are models? USDA Forest Service. Pacific Northwest Research Station. General Technical Report PNW-GTR-232, Portland, OR. 27 pp.

Bunnell, F.L. 1999. What habitat is an island? Pages 1-31 in J.A. Rochelle, L.A. Lehmann, and J. Wisniewski (Eds.). Forest fragmentation: wildlife and management implications. Kononklijke Brill, Leiden, Netherlands.

Canadian Institute of Forestry/Institut forestier du Canada. 2003 Position Statement:Emulation of Natural Disturbance Patterns as a Model for Forest Management. http://www.cif-ifc.org (accessed 6 Feb./04).

Caro, F., Constantino, M., Martins, I., and Weintraub, A. 2003. A 2-opt tabu search procedure for the multiperiod forest harvesting problem with adjacency, greenup, old growth, and even flow constraints. Forest Science 49:738-751.

Crowe, K., Nelson, J., and Boyland, M. 2003. Solving the area-restricted harvest-scheduling problem using the branch and bound algorithm. Canadian Journal of Forest Research 33:1804-1814.

Davis, L.S., K.N. Johnson, P.S. Bettinger, and T.E. Howard. 2001. Forest Management: Fourth Edition. McGraw Hill, New York. 804 pp.

Eng, M. 1998. Spatial patterns in forested landscapes.P. 42-75 In Conservation Biology Principles for Forested Landscapes, Voller, J. and S. Harrison (eds). UBC Press, Vancouver, B.C.

Forman, R.T. 1997. Land Mosaics: The Ecology of Landscapes and Regions. Cambridge Univ. Press, Cambridge, U.K. 632 pp.

Gauthier, S., A. Leduc, and Y. Bergeron. 1996. Forest dynamics modeling under natural fire cycles: a tool to define natural mosaic diversity for forest management. Environmental Monitoring and Assessment 39:pp. 417–434.

Glover, F. 1986. Future paths for integer programming and links to artificial intelligence. Computers & Operations Research 13: 433–549.

Gordon, J. 1994. From vision to policy: a role for foresters. Journal of Forestry 96, 2:16-19.

Gunn, E. A. 1991. Some aspects of hierarchical forest planning in forest management. Pp. 54-62 In Proceddings of the 1991 Symposium on Systems Analysis in forest Resources. Charleston, South Carolina.

Haila, Y., I. K. Hanski, J. Niemel¨a, P. Punttila, S. Raivio, and H. Tukia. 1994. Forestry and the boreal fauna: matching management with natural dynamics. Annales Zoologici Fennici 31: 187–202.

Hertz, A., and Widmer, W. 2003. Guidelines for the use of metaheuristics in combinatorial optimization. European Journal of Operational Research 151:247-252.

Hillier, F.S. and G.L. Lieberman. 2001. Introduction to Operations Research: Seventh Edition. McGraw Hill, New York. 1214 pp.

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

Hof, J.G., and Joyce, L.A. 1993. A mixed integer linear programming approach for spatially optimizing wildlife and timber in managed forest ecosystems. Forest Science 39: 816–834.

Howard, R.A. 1988. Decision analysis: practice and promise. Management Science 14:679-694.

Hunter, M.L. 1990. Wildlife, forests and forestry. Prentice-Hall Inc., Englewood Cliffs, New Jersey, USA.

Hunter, M. L. 1993. Natural fire regimes as spatial models for managing boreal forests. Biological Conservation 65:pp. 115–120.

Land, A.H., and Doig, A.G. 1960. An automatic method for solving discrete programming problems. Econometrica 28:497-520.

Levin, S.A. 1992. The problem of pattern and scale in ecology. Ecology 73:1943-1967.

Liu, G., Nelson, J., and Wardman, C. 2000. A target-oriented approach to forest ecosystem design — changing the rules of forest planning. Ecological Modeling 127:269–281.

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

Martell, D.L., Gunn, E.A., and Weintraub, A. 1998. Forest management challenges for operational researchers. European Journal of Operational Research 104:1-17.

Michalewicz, Z., and Fogel, D.B. 2000. How to Solve It: Modern Heuristics. Springer-Verlag. New York, 467 pp.

Mladenoff, D.J. and Baker, W.L. 1999. Spatial modeling of forest landscape change: approaches and challenges. Cambridge U. Press, Cambridge. 350 pp.

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

Murray, A.T., and Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. Forest Science 48: 779-789.

Natural Resources Canada. 2004. Canada's natural resources: now and for the future. http://www.nrcan-rncan.gc.ca/cfs-scf/index_e.html. (accessed April 2/04).

Ohman, K., and Eriksson, L.O. 1998. The core area concept informing contiguous areas for long-term forest planning. Canadian Journal of Forest Research 28: 1817–1823.

O'Neill, R.V., J.R. Krummel, and R.H. Gardner. 1988. Indices of landscape pattern. Landscape Ecology 1:153-162.

Ontario Ministry of Natural Resources. . 2001. Forest management guide for natural disturbance pattern emulation, Version 3.1. Queen's Printer for Ontario, Toronto, 29 pp.

Osman, I.H. and Laporte, G. 1996. Metaheuristics: a bibliography. In: G. Laporte and I.H. Osman, eds., Metaheuristics in Combinatorial Optimization, Annals of Operations Research 63, Baltzer, pp. 513-623.

Pearse, P.H., and S. Sydneysmith.1966. Method of allocating logs among several utilization processes. Forest Products Journal 16(9):89-98.

Pearse, P.H. 2001. Ready for change: crisis and opportunity in the coast forest industry. A Report to the Minister of Forests on British Columbia's Coastal Forest Industry. Vancouver, B.C.

Richards, E.W. and Gunn, E.A. 2000. A model and tabu search method to optimize polygon harvest and road construction schedules. Forest Science 46:188-203.

Richards, E., and Gunn, E.A. 2003. Tabu search design for difficult forest management optimization problems. Canadian Journal of Forest Research 33: 1126–1133.

Rebain, S. and McDill, M. 2003. A mixed integer formulation of the minimum patch size problem. Forest Science 49: 608-618.

Reeves, C.R. (Editor). 1993. Modern metaheuristic techniques for combinatorial problems. Blackwell Scientific Publications, Oxford, U.K.

Saunders, D.A., R.J. Hobbs, and C.R. Margules. 1991. Biological consequences of ecosystem fragmentation: a review. Conservation Biology 5:18-32.

Scientific Panel for Sustainable Forest Practices in Clayoquot Sound. 1995. Sustainable ecosystem management in Clayoquot Sound: planning and practices. Report 5. Victoria, B.C.

Sessions, J., and Bettinger, P. 2001. Hierarchical planning: pathway to the future? In Proceedings of the First International Precision Forestry Symposium, 17–20 June 2001, Seattle, Wash. College of Forest Resources,University of Washington, Seattle, Wash. pp. 185–190.

Spies, T., and Turner, M. 1999. Dynamic Forest Mosaics. *In:* Maintaining Biodiversity in Forest Ecosystems (Editor: M.L. Hunter). Cambridge U. Press, Cambridge.

Turner, M.G. 1989. The effect of pattern on process. Annual Review of Ecology and Systematics 20:170-197.

The Nature Conservancy. 1982. Natural heritage program operations manual. The Nature Conservancy, Arlington, VA.

Urban, D.L., R.V. O'Neill, and H.H. Shuggart. 1987. Landscape Ecology. Bioscience 37:119-121.

UNECE . 2003. United Nations Economic Commission for Europe: Forest Products Statistics 1998-2002. United Nations Timber Branch, Geneva.

Van Deusen, P.C. 1999. Multiple solution harvest scheduling. Silva Fennica 33(3): 207–216.

Van Deusen, P.C. 2001. Scheduling spatial arrangement and harvest simultaneously. Silva Fennica 35(1): 85–92.

Williams, H.P. 1999. Model building in mathematical programming: fourth edition. John Wiley and Sons Ltd., New York. 354 pp.

Wolsey, L.A. 1998. Integer programming. John Wiley and Sons Ltd., New York. 264pp.

Chapter II

Solving the Area-Restricted Harvest-Scheduling Model

using the

Branch and Bound Algorithm

Introduction

Regulations specifying allowable harvest patterns now commonly limit the size of harvest-openings and restrict harvest activities on adjacent polygons for a fixed period. These spatial constraints complicate the problem of finding optimal harvest schedules because, to model such a problem, the decision variables must be integer. The harvest-scheduling problem with spatial constraints is therefore an integer programming problem; and these problems are, in general, more difficult to solve than similar sized problems with continuous decision variables (Williams 1999). The challenge of developing new formulations or algorithms that are more effective at solving the harvest-scheduling problem with spatial constraints has attracted many researchers during the last decade.

There have been two broad approaches to modeling the harvest-scheduling problem with adjacency constraints (Murray 1999): 1) the unit restricted model (URM) and, 2) the area restricted model (ARM). In the URM, the boundaries of all potential harvest-opening are predefined; *i.e.*, the boundary of each polygon equals the boundary of each potential cut-block (i.e., contiguous area of harvested forest). In the ARM, the boundaries of all potential cut-blocks are not predefined. Instead, polygons may be aggregated to form cut-blocks during the search for an optimal solution. The limit of this aggregation is defined by the maximum allowable opening size. It has been observed (Walters *et al.* 1999, Richards and Gunn 2000) that one advantage of the ARM is that the block configuration emerges in the context of an optimal flow of timber; and that the predefined cut-blocks of the URM may underestimate the potential harvest flow through sub-optimal cut-blocks. Moreover, research indicates that poor block configuration can contribute to lower objective function values (Jamnick and Walters, 1991).

In this paper, two formulations of the ARM as binary integer programming problems are presented. The formulations are tested on tactical planning problems in 6 different forests with 3 different initial age-class distributions. Exact solutions are computed using the branch and bound algorithm. Our objectives are first, to develop and verify formulations which can produce optimal solutions on smaller problems within reasonable computing periods; and second, to discover and analyze any difficulties which might arise in applying these formulations to larger problems.

The outline of this paper is as follows: first, a review of the literature on the harvest-scheduling problem with spatial constraints is presented. Second, our two formulations of opening size constraints are described and illustrated with an example. Third, descriptions of the tactical planning problem and data sets are presented, followed by a presentation of results. Next, we discuss the results of this research with an emphasis on examining possible shortcomings arising from the application of these formulations to larger problems. Finally, we offer our conclusions and suggestions for further research.

Literature Review

Prior to the introduction of adjacency constraints, the harvest-scheduling problem had been modeled using linear programming (LP). A host of LP models was developed; *e.g.* MAXMILLION (Ware and Clutter 1971), Timber RAM (Navon 1971), FORPLAN (Stuart and Johnson 1985) and MELA (Siitonen 1993). The problem shared by all LP approaches is that, since the decision variables are continuous, discrete allocation of cut-blocks is not possible.

Research into solving the URM using branch and bound has taken two directions. First, methods were devised to decrease the number of adjacency constraints through constraint aggregation schemes (Meneghin *et al.* 1988; Torres-Rojo and Brodie 1990; Yoshimoto and Brodie 1994), although this can lead to a loss of efficiency in solving some problems (as discussed by Torres and Brodie 1990). Recent improvements in commercial mixed integer programming (MIP) solvers have made the reduction of constraints less relevant, since many solvers now accept an unlimited number of constraints. The second direction taken was reformulating adjacency constraints to improve the efficiency of the branch and bound search (Yoshimoto and Brodie 1994; Murray and Church 1995b; 1996b; Snyder and Revelle 1996; 1997; McDill and Braze 2000). Beyond improved formulations, McDill and Braze (2000) provided strong evidence that, in addition to the number of decision variables, the initial age-class structure of a forest consistently affects the difficulty of solving problems with branch and bound. Branch and bound is not the only method used to solve the URM optimally: dynamic programming (Hoganson and Borges 1998) and column generation (Weintraub *et al.* 1994) have also been used.

Increased computing speeds and improved commercial solvers have enabled researchers to solve larger problems using the branch and bound algorithm. Recently, McDill and Braze (2001) experimented with different optimality tolerance parameters in the branch and bound algorithm and showed that problems up to 2,500 cut-blocks over three periods can be solved when the acceptable gap between the best integer solution and the LP upper bound is widened.

The initial restriction of exact methods to smaller problems spurred interest in heuristic methods. Solutions formed through heuristic algorithms are not necessarily optimal, but they can provide "good" solutions to larger problems more quickly than exact methods. Monte Carlo integer programming (MCIP) has been used to solve the URM (O'Hara *et al.* 1989; Nelson and Brodie 1990; Clements *et al.* 1990; Boston and Bettinger 1999). When using a heuristic, it is reasonable to ask how far the solution is from the optimal (Wolsey 1998). The MCIP solutions of Nelson and Brodie (1990) and of Boston and Bettinger (1999) were within 10% of the optimum.

Over the last ten years, researchers have improved the effectiveness of heuristic algorithms when applied to difficult integer programming problems. Three metaheuristics have emerged as flexible search strategies for solving such problems: simulated annealing, tabu search, and genetic algorithms. Simulated annealing (Dahlin and Salnas 1993; Murray and Church 1995), tabu search (Murray and Church 1995; Brumelle *et al.* 1998; Boston and Bettinger 1999), and genetic algorithms (Mullen and Butler 1997) have each been applied to the URM and were consistently found to compute solutions closer to the optimum than MCIP. Barrett et al. (1998) also used heuristics to investigate the trade-offs related to opening size problems.

It was in the context of applying metaheuristics to the spatially constrained harvest-scheduling problem that the ARM was conceived and explored. Lockwood and Moore (1993) used simulated annealing on a large application of the ARM (27,548 polygons scheduled over 12 periods). In their objective function, penalty costs were used for violations of opening size or adjacency regulations. Walters *et al.* (1999) developed a heuristic, reliant on an LP solution, to solve the ARM. Richards and Gunn (2000) used tabu search, and Clark *et al.* (2000) used a three-stage heuristic to solve the ARM and simultaneously schedule road construction. Recently, there has been research performed on solving the ARM using the branch and bound algorithm. McDill et al. (2002) developed an algorithm for identifying a set of maximum opening size constraints for solving the ARM using the branch and bound algorithm. They tested this method on forests comprised of 50 and 80 polygons and limited their opening constraints to a maximum of four polygons. Our intention in this paper is to build on this initial research by examining the problems confronted when formulating and solving models of larger problems with more complex opening size constraints.

Methods

Our goal is to find a method to identify a set of opening constraints that ensure that no opening exceeds the maximum opening size. We are also concerned with computational efficiency, and for this reason, two formulations of opening size constraints are tested. These formulations are: 1) maximum opening constraints, and 2) appended clique constraints. Each formulation is described separately.

Maximum Opening Constraints

The set of opening constraints for the ARM is a list of unique inequalities that prevents openings from exceeding a specified size, for every possible combination of adjacent polygons. For each inequality, we identify the minimal set of adjacent polygons that, when simultaneously harvested, creates an opening violation. Next, we form linear inequalities from these opening violations. Since each equation is the minimal set of polygons that violates the opening limit, each inequality requires that only one polygon be removed before it becomes feasible. Therefore, we do not need area coefficients for each decision variable, and the right-hand side of the inequality equals the sum of the decision variables minus 1. For example, consider the four neighbouring polygons in equation [1], where the maximum opening size is 40 ha and each polygon is 12 ha. We let x_i equal 1 if polygon i is cut, and zero otherwise:

[1] $$x_1 + x_2 + x_3 + x_4 <= 3$$

If any 3 polygons in equation [1] are selected for harvest, the maximum opening size remains feasible at 36 ha.

Two types of redundancies are eliminated from the constraint set. First, different orderings of polygons, such as in equation [2] compared to equation [1], do not impact the solution, and are deleted from the constraint set.

[2] $$x_4 + x_3 + x_2 + x_1 <= 3$$

Second, if any subset of the polygons also creates a constraint, the larger constraint is eliminated. This can occur, for example, if polygons x_1, x_2, and x_3 are 2 ha each, and polygon x_4 is 39 ha. A block containing polygons x_1, x_2, and x_3 is feasible, but with the addition of x_4 becomes infeasible. However, just x_1 and x_4 alone break the 40 ha opening limit, and therefore are included in equation [3].

[3] $$x_1 + x_4 <= 2$$

Including equation [1] as a constraint when equation [3] also exists, is superfluous (i.e. ineffective), and therefore it is removed from the constraint set.

Method for Computing the Set of Maximum Opening Constraints

The algorithm used to compute a set of maximum opening constraints is described as follows:

1. Create a set comprised of all polygons in the database.
2. Select from the set a polygon, and determine all its adjacent neighbours.
3. Add a single neighbour to the polygon to create a contiguous pair.
4. If the total area of the pair exceeds the opening size limit, add a new constraint to the constraint set.
5. If the total area of the pair does not exceed the opening size, add it to the set of feasible pairs.
6. Return to step 3 until each neighbour has been selected.
7. Return to step 2 until every polygon has been selected.
8. Weed the constraint set of all superfluous and redundant constraints.

This creates the weeded constraint set for all possible combinations of pairs of polygons, and the set of all possible combinations of feasible polygon pairs that do not break the opening size limit. The eight steps above are now repeated, except the initial polygon set in step 1 is replaced with the set of feasible polygon pairs (produced in step 5), and the neighbours in step 2 are all adjacent neighbours of both polygons. Starting at step 1 with pairs, the process results in a set of feasible polygon triplets, as well as the new set of constraints of infeasible triplets to be appended to the set of infeasible pairs. The entire process is repeated (quadruplets, quintuplets, etc.) until no feasible polygons groupings emerge. The algorithm exhaustively searches all possible additions of single polygons that are neighbours to a feasible polygon cluster.

The algorithm can be thought of as a breadth-first search of trees with branches added for each single neighbouring polygon. In this respect, it differs from the depth-first approach of McDill et al. (2002). Each successive level in the tree increases the number of polygons in a cluster by one. Efficiencies were added to the algorithm by branch bounding, redundancy checking at every level, and using hash values to eliminate redundancies. Branches were bound at levels where they violate constraints, greatly reducing the search area. Branches were also bound that were found to be redundant to other branches (i.e. equations [1] and [2]). Checking for redundancies normally involves comparing the polygons in each cluster to the polygons within every other cluster in a time consuming, exhaustive search. This was avoided by implementing the constraint set as a hash map, and using combinations of the polygon identification numbers as the hash key. This caused clusters with redundant polygon compositions to have identical hash keys. Because the hash map does not accept duplicate values with identical hash keys, redundant constraints were not retained.

Appended Cliques

The second formulation entails appending a set of clique constraints to the set of maximum opening constraints. A clique is defined as a set of mutually adjacent polygons (Murray and Church 1996). Cliques have been used to formulate adjacency constraints by researchers solving the URM as a binary integer program (Meneghin *et al.* 1988; Murray and Church 1996; McDill and Braze 2000). In solving the URM, the idea is that a single clique of more than two polygons (a higher order clique) can eliminate more than one pairwise adjacency constraint. For example, the set of pairwise adjacency constraints applicable to the set of polygons in Figure 2.1 comprises 12 inequalities listed in Table 2.1.

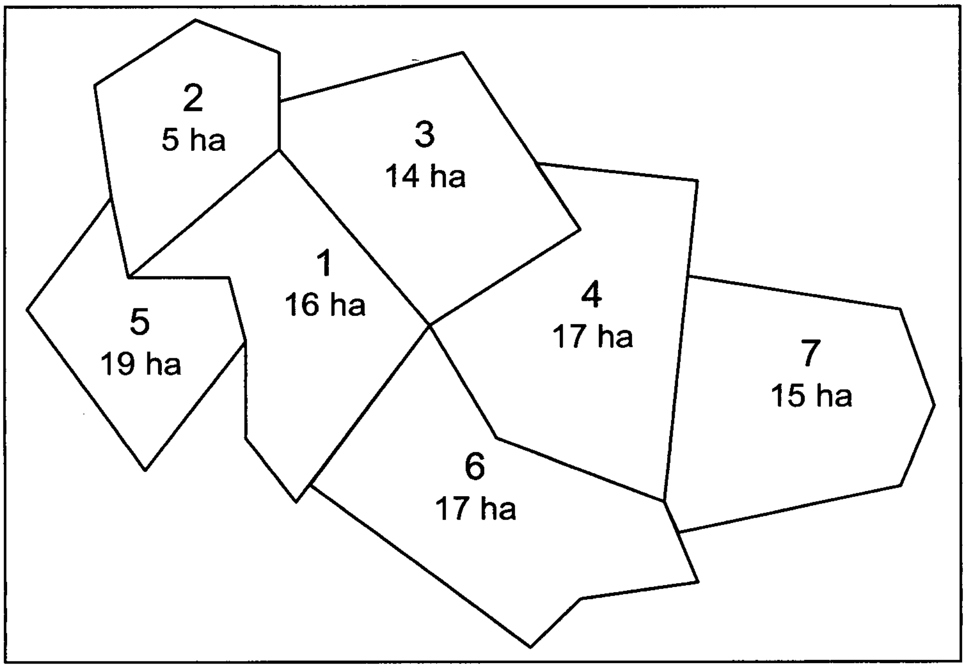

Figure 2.1. Sample forest of seven polygons.

Table 2.1: Twelve URM pairwise inequalities applicable to polygons in Figure 2.1.

Constraint No.	Constraint
1	$x_1 + x_2 \le 1$
2	$x_1 + x_3 \le 1$
3	$x_1 + x_4 \le 1$
4	$x_1 + x_5 \le 1$
5	$x_1 + x_6 \le 1$
6	$x_2 + x_3 \le 1$
7	$x_2 + x_5 \le 1$
8	$x_3 + x_4 \le 1$
9	$x_3 + x_6 \le 1$
10	$x_4 + x_6 \le 1$
11	$x_4 + x_7 \le 1$
12	$x_6 + x_7 \le 1$

The cliques constraints applicable to the set of polygons in Figure 2.1 are:

[4] $x_1 + x_2 + x_5 \leq 1$

[5] $x_1 + x_2 + x_3 \leq 1$,

[6] $x_1 + x_3 + x_4 + x_6 \leq 1$

[7] $x_4 + x_6 + x_7 \leq 1$

Note that clique constraint [6], for example, enforces the same adjacency constraints as the pairwise inequalities numbered 2, 3, 5, 8, 9, and 10 in Table 2.1. Hence, it can replace these pairwise constraints. In addition, researchers have noted that clique constraints are stronger inequalities than the pairwise constraints (Murray and Church 1996) and therefore tend to improve computational efficiency. Both Murray and Church (1996) and McDill and Braze (2000) have concluded that the use of clique constraints with pairwise constraints is, on average, the best formulation of adjacency constraints presently used to solve the URM.

The use of clique constraints in solving the ARM is based on the same idea: one finds higher order cliques and adds them to the minimal opening constraints; but, the definition of higher order cliques differs from that used in the URM. In the context of the ARM, higher order cliques must: 1) comprise at least 3 mutually adjacent polygons; and 2) contain an area greater than the maximum opening size. To illustrate this, observe that Figure 2.1 contains a clique comprising polygons 1, 3, 4, and 6. Letting a_i equal the area of polygon i, the inequality representing this clique constraints is:

[8] $a_1x_1 + a_3x_3 + a_4x_4 + a_6x_6 \leq 40$

Unlike minimal opening constraints, the area coefficient must be used with ARM clique constraints. For example, the above clique inequality, without the area coefficients is:

[9] $x_1 + x_3 + x_4 + x_6 \leq 3$

This inequality is not valid because it allows the simultaneous harvest of polygons 1, 4, and, 6, which violates the maximum opening size of 40 ha.

All clique inequalities were added to the set of maximum opening constraints and incidental redundancies were not removed. Removal of duplicate cliques was left to the preprocessor of the MIP solver. The merit of this formulation requires empirical testing, so if solution times are consistently decreased by the addition of cliques to the minimal opening constraints, then we can conclude that they may be a useful addition to our ARM formulation.

Problem Definition and Model

Formulations of the ARM are tested by scheduling harvest activities over three 20-year periods in a forest with an 80-year rotation period. This is classified as a tactical planning problem because the planning horizon is less than one rotation (Martell *et al* 1998). The objective is to maximize net present value of the harvest activities, where value is measured at \$100 *per* m^3, discounted from the middle of each planning period at 4%. Harvest activities are restricted to clearcutting or doing nothing. The constraints in this problem are:

1) a limit on inter-period harvest volume fluctuations of plus or minus 10%;
2) a limit on opening size of 40 ha, with a 1-period restriction on harvesting of all adjacent polygons (adjacency is defined as sharing a common boundary node);

3) a minimum harvest age of 80-years;

4) an ending-age constraint to ensure that the average age of the forest at the end of the planning horizon be at least 40 years (*i.e.*, the average age of a forest regulated on an 80-year rotation). The formulation of the ending-age constraint is from McDill and Braze (2000).

The notation and mathematical formulation of the model are presented below:

x_{it} = 1 if polygon i is harvested in period t = 1, 2, or 3; 0 otherwise

[Note, if polygon i has not been cut for the entire planning horizon, then x_{it} =1 when t = 0. This is done to implement the ending age constraint].

v_{it} = volume of polygon i in period t (m^3).

r_{it} = discounted net revenue from harvesting polygon i in period t ($).

a_i = area of polygon i (ha)

h_t = total volume harvested in period t (m^3)

e_{it} = ending age of polygon i if harvested in period t (ending age is measured in year 60)

I = number polygons in the forest

P = set of maximum openings in the forest

f_p = number of polygons in the maximum opening constraint p, where $p \in P$

Omax = maximum opening size (ha)

C = set of cliques, where each has at least 3 polygons and a total area > Omax

S_c = set of polygons in clique c, where $c \in C$

Objective Function: maximize net present value

$$[10] \quad \text{Maximize} \quad \sum_{i=1}^{I} \sum_{t=1}^{3} r_{it} x_{it}$$

Subject to:

Each polygon may be assigned not more than one prescription over the planning horizon

[11] $$\sum_{t=0}^{3} x_{it} = 1 \qquad \forall\ i = 1,\ldots,I$$

Accounting variables defining the total volume harvested in each period

[12] $$\sum_{i=1}^{I} v_{it}\, x_{it} - H_t = 0 \qquad \forall\ t = 1,\ldots,3$$

Maximum opening size constraints for each period

[13] $$\sum_{i \in p} x_{it} \leq f_p - 1 \qquad \forall\ p \in P\,;\ t = 1,\ldots,3.$$

Inter-period harvest volumes may fluctuate at most by 10%

[14] $$-.9H_t + H_{t+1} \geq 0 \qquad t = 1,\ldots,2$$

[15] $$-1.1Ht + H_{t+1} \leq 0 \qquad t = 1,\ldots,2$$

Average polygon age of total forest area must be at least 40 years at the end of planning horizon (year 60).

[16] $$\sum_{i=1}^{I} \sum_{t=0}^{3} (e_{it} - 40) a_i x_{it} \geq 0$$

The decision variables are binary

[17] $$x_{it} \in \{0,1\} \qquad \forall\ i = 1,\ldots,I;\ t = 0,\ldots,3$$

Except for the addition of clique constraints, the ARM with clique constraints is identical to the above ARM. The mathematical formulation of the clique constraints is:

$$[18] \qquad \sum_{i \in S_c} a_i x_{it} \leq Omax \qquad \forall\ c \in C;\ t = 1,\ldots,3$$

Description of Data

Spatial data from six forests in British Columbia were used in this study. Polygons in the three smaller forests, Gavin, Hardwicke, and Naka were manually drawn by forest engineers, while the polygons of Stafford and Kootenay were formed by GIS-overlays. The forest of Kootenay I is a subset of polygons extracted from the larger forest, Kootenay II. Polygons ranged in size from 0.06 to 43 ha (a few polygons marginally greater than the 40 ha opening limit are allowed to be harvested by themselves). Table 2.2 summarizes the spatial attributes of these forests, including the size of polygons and the number of polygon adjacencies.

Table 2.2. Spatial attributes of the six forests.

Forest	Total Area (ha)	No. of Polygons	Polygon Area (ha) (minimum, maximum, average)	Polygon Adjacencies (minimum[1], maximum, average)
Gavin	6,193	346	1.1, 43.0, 17.9	1, 16, 5.8
Hardwicke	6,948	423	3.0, 42.0 16.4	2, 10, 6.1
Naka	10,934	785	0.06, 43.0, 13.9	1, 14, 5.2
Stafford	10,421	1,008	2.0, 20.0, 10.3	0, 12, 4.3
Kootenay I	38,441	3,256	6.0, 38.0, 11.8	0, 15, 5.6
Kootenay II	71,257	6,093	6.0, 38.0, 11.7	0, 15, 5.2

1. A few isolated polygons in Kootenay I, Kootenay II and Stafford have no neighbours.

Table 2.3: Three age distributions randomly assigned to each forest.

Eligibility in Period 1	Range of Uniform Distribution of Ages
50%	10-150 years
75%	40-200 years
100%	80-200 years

We were also interested in testing different age-class distributions for each forest. McDill and Braze (2000) found a strong correlation between the initial age-class distribution of a forest and difficulty of solving the URM using a branch and bound algorithm; *viz*., that problems with a high percentage of old-growth forest are, in general, more difficult to solve than others. For this reason, we assigned three age-class distributions to each forest to test the formulations on each. The age-class distributions are based on the percent of polygons eligible for harvest in the first period. The three distributions tested are: 100%, 75%, and 50%. Table 2.3 summarizes the uniform age distributions assigned to each forest using a random number generator.

In addition to the effect of different age class distributions, we also wished to explore the effects of different maximum opening-size limits. Increasing the opening size limit has two effects. First, it allows for more possible combinations of polygons to aggregate into a feasible opening, thereby increasing the number of possible solutions. We wished to assess the degree to which this increases the computing time needed to find an optimal solution. Second, since increasing the opening size limit also requires more maximum opening size constraints, we wanted to evaluate the rate at which the number of opening size constraints increases as the allowable opening size increases. For this reason, we tested 5 different opening size limits on the Stafford forest: 20, 30, 40, 50, and 60 ha opening size limits respectively.

For reasons of practicality, we limited our analysis of the clique formulations to the smaller forests, and set the maximum run-time to four hours. The larger forests in Kootenay I and II were not solved with the clique formulation, and a limit of 24 hours was set as the maximum run-time. In each problem instance the optimality tolerance (the percent difference between the lower best integer program bound and the upper LP bound) was set at 0.5%.

The models were written using MPL® modelling software (Maximal Software Inc., Arlington, VA) and all runs were executed using the CPLEX® 7.5 MIP solver (ILOG Inc., Mountain View, CA) on a Pentium® III 1.0 GHz central processing unit, with a Windows NT® 4.0 operating system and 1.5 gigabytes of RAM. In the CPLEX® solver, all default parameters were used except that the optimality tolerance was widened to 0.5% and the MIP emphasis was changed from optimality to feasibility. The program for computing the opening constraints was written in JAVA and executed on a Pentium 4, 2 GHz central processing unit with a Linux operating system.

Results

The results from computing the opening constraints and cliques are presented in Table 2.4. We first observe that in the Stafford forest, an exponential increase in the number of opening constraints occurs as the maximum opening size increases (Figure 2.2), while only a linear increase occurs in the maximum number of polygons per constraint (Table 2.4).

Table 2.4. Results from computing maximum opening constraints and cliques for each forest.

Forest	Minimal opening constraints per period	Maximum No. of polygons per constraint	Cliques per period	Computing time for min. opening constraints
Gavin	1,925	6	459	4 sec
Hardwicke	4,4495	6	761	25 sec
Naka	21,272	10	472	30 min
Stafford-20 ha	2,183	6	1,047	5 sec
Stafford-30 ha	5,474	7	745	45 sec
Stafford-40 ha	16,049	9	251	9 min
Stafford-50 ha	50,110	10	16	2 hrs
Stafford-60 ha	168,885	13	2	100 hrs
Kootenay I	97,360	6	N/A	5 hrs
Kootenay II	156,563	7	N/A	8 hrs

Figure 2.2. Exponential increase in the number of opening size constraints per period relative to the maximum opening size in the Stafford forest.

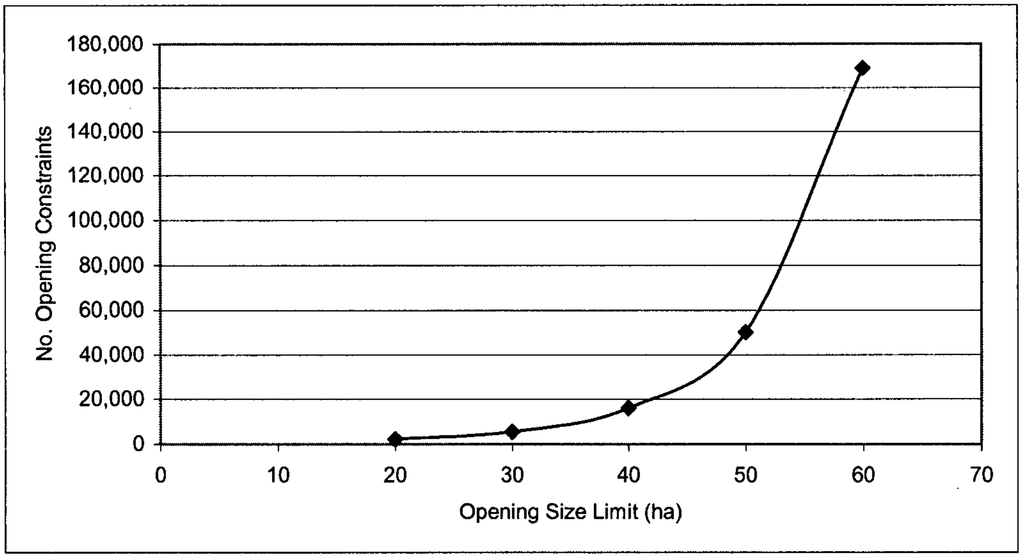

As the number of polygons within an opening increases, the number of adjacent neighbours rapidly expands and the possible combinations of polygons that create unique opening constraints explodes. Another factor contributing to the number of constraints is the polygon line work. For example, Naka has more constraints, but fewer polygons than Stafford-40. This can be attributed to a few, small, linear-shaped polygons in Naka that have many adjacent neighbours. Overall, Naka has more very small polygons and more polygon adjacencies than Stafford-40 (Table 2.2). Small, irregular shaped polygons become non-trivial problems during constraint generation, so every effort should be made to eliminate them *a priori* from the database (e.g. merge them into neighbour polygons).

The computing times needed to calculate the sets of opening constraints are also shown in Table 2.4. The time to process constraints for Stafford increased at an exponential rate relative to the opening size; however, relative to the number of constraints, it increased at a constant rate. However, there are differences between forests, as shown by Stafford-60 and Kootenay II. These forests have a similar number of constraints (> 150,000), but processing time differs by an order of magnitude. This is caused by the high number of polygons per constraint in Stafford-60 (13) compared to Kootenay II (7). The large number of polygons per constraint causes more superfluous constraints that need to be weeded out, and this weeding process can take up to 20 times longer than simply generating the constraint set that includes the superfluous constraints.

Table 2.4 also lists the number of cliques per period per forest. One obvious trend, illustrated in the Stafford forest, is that as the opening size increases, the number of cliques decreases (e.g., Stafford-20 has 1,047 cliques, while Stafford-60 has only 2). It becomes harder to find cliques that by themselves create violations for the larger openings. This demonstrates an obvious limitation inherent in clique constraints for the ARM.

The resulting LP and IP objective functions, relative gaps, and computing times for each problem instance solved using the two formulations are presented in Table 2.5.

Table 2.5. LP solutions, integer solutions, relative gaps, and computing times for each problem instance using the two ARM formulations. (Note 50%, 75%, and 100% refer to the number of polygons eligible for harvest in the first period. "With cliques" refers to the formulation with clique constraints appended; and "without cliques" refers to the formulation without clique constraints).

Forest		50% with cliques	50% without cliques	75% with cliques	75% without cliques	100% with cliques	100% without cliques
Gavin	LP	613,261	613,298	818,072	818,072	900,845	900,860
	IP	607,712	607,924	813,816	813,361	896,085	896,271
	Gap	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
	Time	3.7 min.	3.9 min.	53.4 min.	35.2 min.	1.7 hr.	42.3 min.
Hardwicke	LP	618,152	618,152	905,372	905,399	1,032,840	1,032,854
	IP	614,929	614,524	899,272	900,116	1,017,604	1,017,008
	Gap	Optimal	Optimal	0.60%	0.515	1.48%	1.52%
	Time	58 sec.	2.3 min.	4 hrs.	4 hrs.	4 hrs.	4 hrs.
Naka	LP	2,447,163	2,447,489	3,096,629	3,097,514	3,722,808	3,722,838
	IP	2,436,547	2,436,552	3,077,219	3,077,230	3,694,829	3,694,812
	Gap	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
	Time	5.4 min.	5.2 min.	16.4 min.	14.1 min.	45.3 min.	41.5 min.
Stafford-20	LP	2,206,051	2,206,051	2,844,800	2,845,313	3,335,002	3,336,209
	IP	2,202,451	2,205,899	2,844,771	2,844,751	3,334,967	3,290,517
	Gap	Optimal	Optimal	Optimal	Optimal	1.1%	1.38%
	Time	1.1 min.	47 sec.	1.7 min. 4	45 sec.	4 hrs.	4 hrs.
Stafford-30	LP	2,385,794	2,385,796	3,085,943	3,085,943	3,620,131	3,620,131
	IP	2,380,009	2,378,215	3,077,917	3,073,567	3,601,615	3,601,315
	Gap	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
	Time	1.8 min.	1.6 min.	2.5 min.	3.3 min.	3.3 min.	5.0 min.
Stafford-40	LP	2,437,829	2,437,829	3,131,460	3,131,462	3,667,689	3,667,689
	IP	2,423,872	2,424,135	3,127,728	3,127,943	3,661,992	3,662,742
	Gap	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
	Time	8.9 min.	5.9 min.	6.0 min.	5.2 min.	6.6 min.	5.1 min.
Stafford-50	LP	2,456,458	2,456,458	3,143,926	3,143,926	3,682,274	3,682,274
	IP	2,452,818	2,452,855	3,142,367	3,142,717	3,681,035	3,681,035
	Gap	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
	Time	16.2 min.	15.5 min.	19.8 min.	19.4 min.	34.9 min.	34.6 min.

Forest		50% with cliques	50% without cliques	75% with cliques	75% without cliques	100% with cliques	100% without cliques
Kootenay I	LP	N/A	8,890,022	N/A	12,460,000	N/A	13,432,000
	IP	N/A	8,852,097	N/A	11,178,270	N/A	13,302,383
	Gap	N/A	Optimal	N/A	Optimal	N/A	Optimal
	Time	N/A	2 hrs.	N/A	5.6 hrs.	N/A	7.5 hrs
Kootenay II	LP	N/A	16,426,000	N/A	20,791,000	N/A	24,916,000
	IP	N/A	16,423,000	N/A	No solution	N/A	24,915,000
	Gap	N/A	Optimal	N/A	None	N/A	0.86%
	Time	N/A	6.3 hrs.	N/A	24 hrs.	N/A	17.2 hrs.

Note that the results for Stafford-60 were not included in Table 2.5 because no solutions were produced within the 4-hour limit placed on solving the smaller problems. The optimization software parses the data read, i.e., it builds the matrix, prior to the solution processing. The Stafford-60 forest was still being parsed after 4 hours; hence, no LP solution was calculated within the time limit. The parsing time for the Kootenay II forest was 5.5 hours.

Each solution was entered into a GIS and analyzed to confirm feasibility. For example, Figure 2.3 illustrates an optimal solution for the Gavin forest.

We have several observations regarding the results in Table 2.5. First, the LP solutions produced by the formulation with appended cliques are lower in value than the solutions produced by the formulation without cliques in 11 of the 21 problem instances. This indicates that the addition of clique constraints are somewhat inconsistent in "tightening" the formulation, *i.e.*, reducing the feasible region of the LP problem nearer to that of the integer problem. Our second observation concerns the ability of the clique constraints to decrease computing time needed to find an optimal solution. In very few

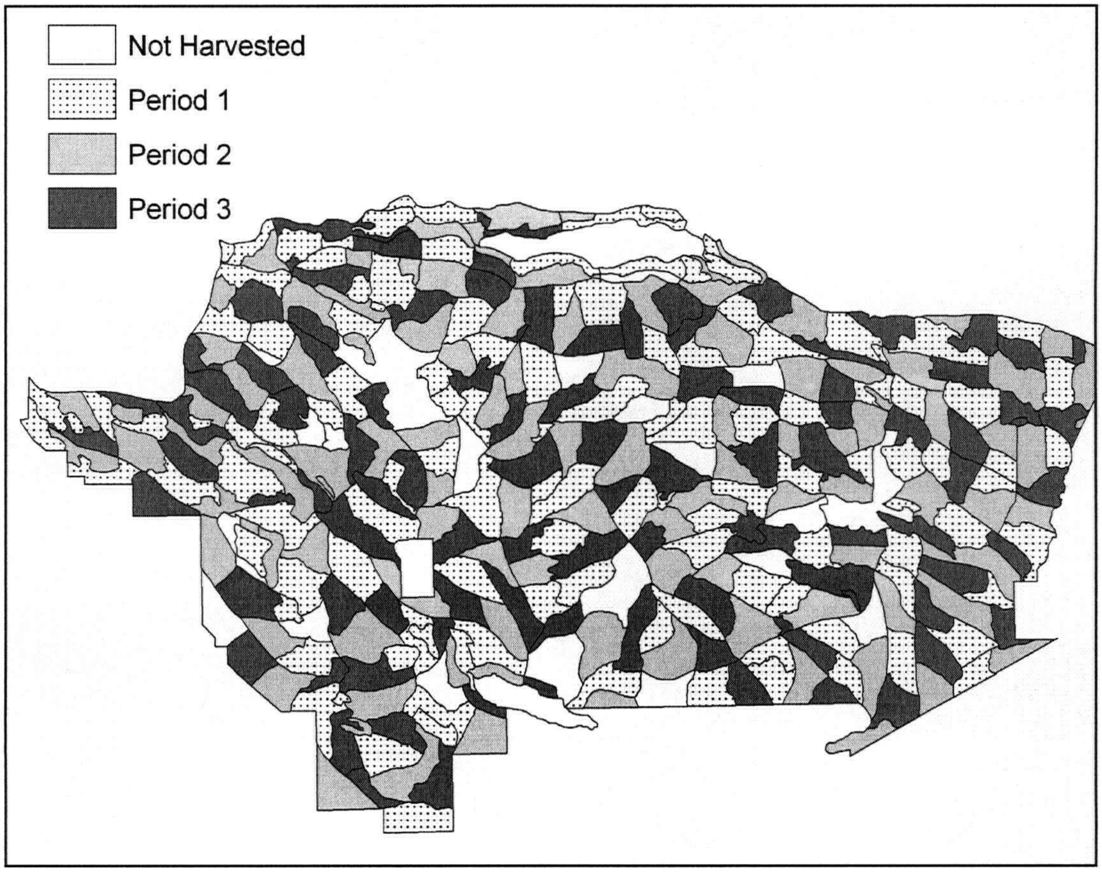

Figure 2.3. An optimal harvest schedule of the Gavin forest mapped for 3 periods.

problem instances did this occur, and in others, the cliques actually hindered the search, despite the fact that a given problem instance with clique constraints began with lower LP bounds than those without clique constraints (e.g., Gavin with 100% eligible). Why did the addition of cliques adversely affect the branch and bound search in some instances? Our reasoning is that different formulations can influence the selection of variables in the branch and bound algorithm. When branching from a node, there is a rule by which the branching variable is selected. In CPLEX®, there are various rules from which to choose and we selected the default parameter "automatic" along with the MIP emphasis on feasibility. For reasons of commercial competition, the CPLEX® manual does not disclose the algorithms underlying these parameters. We can only conclude that, in some cases, the addition of cliques to the formulation 'misled' this process into selecting less promising variables on which to branch. This indicates that more research is needed into selecting a set of parameters to guide the branch and bound algorithm that will better exploit the structure of the ARM.

Our third observation from Table 2.5 is that for each forest, as one moves from problems with younger to older initial age-class distributions, the solution quality (in terms of relative gaps and computing times) tends to deteriorate. This echoes observations made by McDill and Braze (2000) in their work on the URM; namely, that problems with older forests are, in general, more difficult to solve than problems with younger forests. The reason is that these older polygons make very similar contributions to the objective function, hence MIP solvers are unable to rule out as many branches in the branch and bound tree, thus slowing down the process. There are of, course, exceptions to this general rule. One interesting exception is Kootenay II, with 75% of the polygons initially eligible for harvest. A feasible integer solution for this problem could not be found after 24 hours of computing time; and yet, the same forest, with 100% of the polygons initially eligible yielded a near-optimal solution in 17 hours.

The results in Table 2.5 also illustrate, for the Stafford Forest, the effect of relaxed opening constraints upon the net present value for different age-classes. These trends are shown graphically in Figure 2.4 and demonstrate diminishing returns as the opening size increases beyond 30 ha. The greatest gain in net present value is achieved by modest relaxation of the 20 ha opening limit, and the magnitude of change is similar for all age-classes.

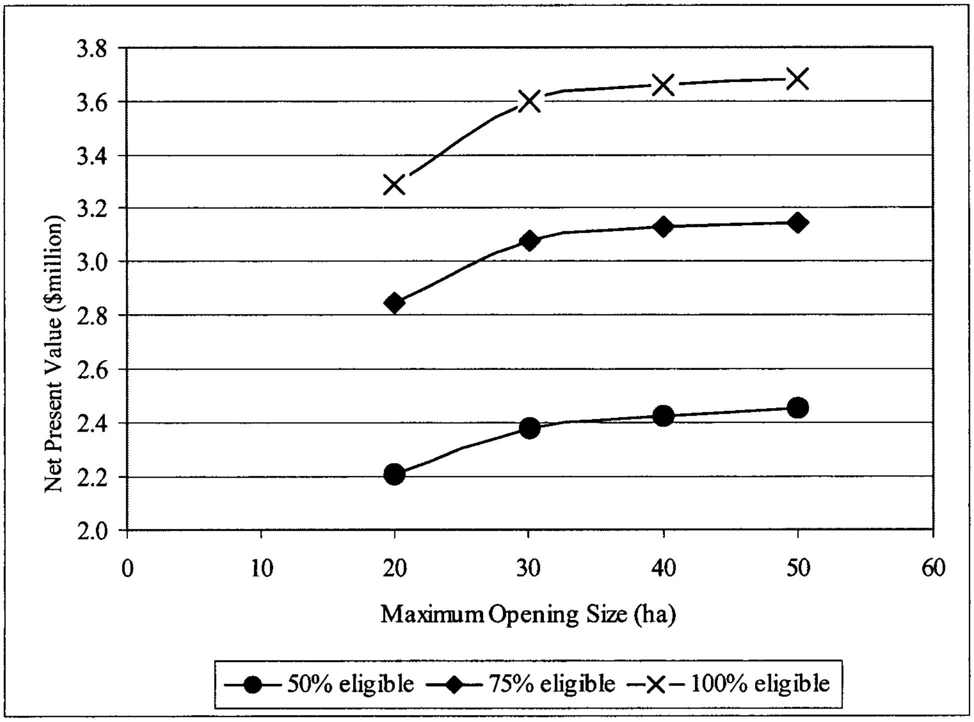

Figure 2.4. Net present value versus maximum opening-size for the three initial age-classes (% polygons eligible for harvest) of the Stafford forest.

Finally, we observe that the number of polygons in a given forest does not always indicate the difficulty of solving it. For example, some problem instances of the Hardwicke and Stafford-20 forests, (423 and 1008 polygons, respectively) proved to be difficult to solve optimally. Initial integer solutions were found within several minutes, but significant, further improvements were not found within the 4-hour limit. The prolonged computing times of Kootenay I and II (at least 3000 polygons), however, indicate that longer computing times can generally be expected when the number of polygons is great. With Kootenay II, in particular, the limits of practical computing effort seem to have been reached.

Discussion

The results indicate that small and medium sized problems of the ARM can be solved optimally (or near-optimally) using the branch and bound algorithm within reasonable periods of computing time. There are however, two questions that must be addressed concerning possible shortcomings of this method.

The first concerns the influence of the polygon size relative to the opening size. We observed, in the case of Stafford, that when this ratio is decreased, the set of opening constraints and the computing time required to calculate it, both increased exponentially. The task of computing opening constraints for Stafford very quickly became impractical, and this is a serious limitation of our method. It might be argued that a more efficient algorithm for preventing the inclusion of duplicate and superfluous constraints can be developed, and that this would help address this shortcoming. There is little doubt that improvements can be made in the coding of our algorithm; but even if such efficiencies were made, there remains the problem of the time-consuming task of parsing hundreds of thousands of constraints. As noted earlier, it took 5 and one half hours to parse the input data for Kootenay II. The solver's preprocessor can be used to remove redundant constraints, but only after the model has been parsed. Therefore, even if our algorithm for computing a set of non-redundant opening constraints were improved, there would remain the problem of the excessive time needed to parse potentially millions of constraints.

A second limitation of this method was illustrated by the computing time needed to solve Kootenay II. Sensitivity analysis of the Kootenay II problem would become highly impractical given solution times in excess of 24 hours. The difficulty in computing solutions to the Kootenay II problems appears to have arisen from the number of decision variables. This is the classic difficulty faced by many integer programming problems; and typical attempts to overcome these problems involve experimenting with better formulations or developing heuristic algorithms.

There is perhaps, one approach to overcome the first and second limitations just discussed. This strategy involves aggregating many small, similar polygons into fewer, larger polygons prior to solving the ARM. This process would both increase the polygon size relative to the opening size and decrease the number of decision variables. Of cardinal importance in this aggregation procedure are the criteria by which polygons are classed as similar; for it is possible that poor aggregation criteria may lead to solutions with significantly lower objective function values than solutions without pre-aggregated polygons. Pre-aggregation of polygons could ultimately lead to developing blocks so large that only a URM and not an ARM would be a suitable model for the harvest-scheduling problem. Such pre-aggregation would utterly defeat the purpose of of using an ARM instead of a URM; for, as Murray and Weintraub (2002) rightly observe, the objective function value of an optimal solution of a URM is always a lower bound for the same problem instance solved solved as an ARM. No doubt, a certain amount of experimentation with aggregation criteria would be needed.

Conclusion

In this research we have shown how to formulate the area restricted harvest-scheduling problem for exact solutions using the branch and bound algorithm. We have also shown that small and medium-sized problems can be solved optimally (or near-optimally) within reasonable time periods. The addition of clique constraints to the original formulation inconsistently provided lower LP bounds, and was of little help, on average, to the branch and bound algorithm in finding solutions more efficiently. As distinct from prior work on this problem (McDill et al. 2002), we have explored the computational limitations of solving the ARM using the branch and bound algorithm by solving larger problems with more complex opening size constraints. We conclude that a) the number of decision variables, and b) the number of opening constraints ultimately limit the applicability of this method to larger problems.

Our results point to several areas in need of further research. First, it would be helpful to have a comprehensive examination of the influence of various branch and bound parameters on the efficiency of solving the ARM. Second, it would also be interesting to see how much further this approach can be taken through the strategy of aggregating smaller polygons prior to solution by the ARM. Finally, it would be useful to evaluate the results of metaheuristic algorithms used to solve the ARM by comparing these results to the optimal solutions found in this paper.

Literature Cited

Barrett, T.M., Gilless, J.K., and L.S. Davis. 1998. Economic and fragmentation effects of clearcut restrictions. For. Sci. 44(4):569-577.

Boston, K., and Bettinger, P. 1999. An analysis of Monte Carlo integer programming, simulated annealing, and tabu search heuristics for solving spatial harvest-scheduling problems. For. Sci. 45(2): 292-301.

Brumelle, S., Granot, D., Halme, M., and Vertinsky, I. 1998. A tabu search algorithm for finding a good forest harvest schedule satisfying green-up constraints. Eur. J. of Operational Res. 106:408-424.

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(2):204-218.

Clements, S.E., Dallain, P.L., and Jamnick, M.S. 1990. An operational, spatially constrained harvest scheduling model. Can. J. For. Res. 20: 1438-1447.

Dahlin, B., and Sallnas, O. 1993. Harvest-scheduling under adjacency constraints – a case study from the Swedish sub-alpine region. Scan. J. For. Res. 8:281-290.

Hoganson, H.M., and Borges, J.G. 1998. Using dynamic programming and overlapping subproblems to address adjacency in large harvest-scheduling problems. For. Sci. 44:526-538.

Jamnick, M.S., and Walters, K.R. 1991. Spatial and temporal allocation of stratum-based harvest schedules. Can. J. For. Res. 23: 402-413.

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

Martell, D.L., Gunn, E.A., and Weintraub, A. 1998. Forest management challenges for operational researchers. Eur. J. Operational Res. 104:1-17.

McDill, M.E., and Braze, J. 2000. Comparing adjacency constraint formulations for randomly generated forest planning problems with four age-class distributions, For. Sci. 46(3): 423-436.

McDill, M.E., and Braze, J. 2001. Using the branch and bound algorithm to solve forest planning problems with adjacency constraints. For. Sci. 47(3):403-418.

McDill, M.E., Rebain, S.E. and J. Braze. 2002. Harvest scheduling with area-based adjacency constraints. For. Sci. 48(4):631-642.

Meneghin, B.J., Kirby, M.W., and Jones, J.G. 1988. An algorithm for writing adjacency constraints efficiently in linear programming models. Pp. 46-53 in The 1988 Symp. on Systems Analysis in Forest Resources. USDA For. Serv. Gen. Tech. Rep. RM-161.

Mullen, D., and Butler, R. 1997. The design of a genetic algorithm based on a spatially constrained timber harvest-scheduling model. P. 57-65 in Proc. of the Seventh Symp. On Systems Analysis in Forest Resources. USDA For. Serv. Gen. Tech. Rep. NC-205.

Murray, A.T., and Church, R.L. 1995a. Heuristic solution approaches to operational forest planning problems. OR Spektrum 17:193-203.

Murray, A.T., and Church, R.L. 1995b. Measuring the efficacy of adjacency constraint structure in forest planning models. Can. J. For. Res. 25:1416-1424.

Murray, A.T., and Church, R.L. 1996. Analyzing cliques for imposing adjacency restrictions in forest models. For. Sci. 42(2):715-724.

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

Murray, A.T. and Synder, S. 2000. Introduction to spatial modeling in forest management and resources planning. For. Sci 46(2):147-156.

Murray, A.T., and Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. Forest Science 48:779-789.

Navon, D. 1971. Timber RAM. USDA For. Serv., Pacific Southwest Forest Range Experiment Station, Res. Pap. PSW-70. 48 pp.

Nelson, J., 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.

O' Hara, A.J., Faaland, B.H., and Bare, B.B. 1989. Spatially constrained timber harvest-scheduling. Can J. For. Res. 19: 715 – 724.

Richards, E.W. and Gunn, E.A. 2000. A model and tabu search method to optimize polygon harvest and road construction schedules. For. Sci. 46(2):188-203.

Siitonen, M. 1993. Experiences in the use of forest management planning models. Tiivistelmä: Kokemuksia mallien käytöstä metsätalouden suunnittelussa. Silva Fennica 27(2):167-178.

Synder, S. and ReVelle, C. 1996. Temporal and spatial harvesting of irregular systems of parcels. Can. J. For. Res. 26:1079-1088.

Synder, S. and ReVelle, C. 1997. Dynamic selection of harvests with adjacency restrictions: The SHARe model. For. Sci. 43:213-222.

Stuart, T.W., and Johnson, K.N. 1985. FORPLAN version II: a tool for forest management planning. Paper presented at the Joint National Meeting of the

Institute of Management Sciences and the Operations Research Society of America, Boston, M.A. April 29 – May 1. 96 pp.

Torres-Rojo, J.M., and Brodie, J.D. 1990. Adjacency constraints in harvest-scheduling: an aggregation heuristic. Can. J. For. Res. 20: 978-986.

Walters, K.R., Feunekes, H., Cogswell, A., and Cox, E. 1999. A forest planning system for solving spatial harvest-scheduling problems. Canadian Operations Research Society National Conference, June 7-9 1999. Windsor, Ontario. 8 pp. Available from http://www.remsoft.com [updated March 2003, cited April 2002].

Ware, G.O., and Clutter, J.L. 1971. A mathematical programming system for management of industrial forests. For. Sci. 17: 428-445.

Weintraub, A., Baharona, F., and Epstein, R. 1994. A column generation algorithm for solving general forest planning problems with adjacency constraints. For. Sci. 40(1):142-161.

Williams, H.P. 1999. Model building in mathematical programming: fourth edition. John Wiley and Sons Ltd., New York. 354 pp.

Wolsey, L.A. 1998. Integer programming. John Wiley and Sons Ltd., New York. 264pp.

Yoshimoto, A., and Brodie, J.D. 1994. Comparative analysis of algorithms to generate adjacency constraints. Can. J. For. Res. 24: 1277-1288.

Chapter III

An Evaluation of Applying the Simulated Annealing Algorithm

to the

Area-Restricted Harvest Scheduling Model

using

Optimal Benchmarks

Introduction and Literature Review

With the advent of ecosystem management, it is increasingly common for the size of harvest-openings to be regulated and for harvest delays to be placed on all stands adjacent to these openings. For example, adoption of the American Pulp and Paper Association's *Sustainable Forestry Initiative* (2001) by more than 90% of the forest companies in the U.S., entails that the average clearcut size, on both private and public land, not exceed 48 ha (Boston and Bettinger 2001). Clearly, such restrictions have become an operational reality; but the challenges they pose to modeling and solving the harvest-scheduling problem have not been fully addressed.

One such challenge is the evaluation of neighbourhood-search metaheuristic algorithms, now widely used for solving the harvest scheduling problem with constraints on opening sizes (e.g., O'Hara *et al.* 1989, Nelson and Brodie 1990, Lockwood and Moore 1992, Dahlin and Sallnas 1993, Murray and Church 1995, Bettinger *et al.* 1997, Ohman and Eriksson 1998, Brumelle *et al* 1998, Boston and Bettinger 1999, Van Deusen 1999, Bettinger *et al.* 1999 , Liu *et al.* 2000, Richards and Gunn 2000, Clark *et al.* 2000, Bettinger *et al.* 2000, Van Deusen 2001, Sessions and Bettinger 2001, Baskent and Jordan 2002, Boston and Bettinger 2002, Crowe and Nelson 2003, Richards and Gunn 2003, Caro *et al.* 2003). Metaheuristic algorithms have been extensively used because the harvest scheduling problem with opening-size constraints is a binary integer programming problem; *i.e.*, for a solution to be spatially explicit, a given stand must be either harvested or not harvested in a given period. Most large integer programming problems are notoriously difficult to solve using exact algorithms (Wolsey 1998). Partial enumeration by exact algorithms often have a slow convergence rate, and only small or medium-sized problems have thus far been solved to optimality (Crowe *et al.* 2003). Since many practical harvest-scheduling problems are large, metaheuristics have been used instead of exact algorithms.

One problem with metaheuristic algorithms is that they produce *approximately optimal* solutions; *i.e.*, they neither guarantee optimality nor provide any indication of how close their solutions are to being optimal (Reeves 1993). Clearly, it is important to form some estimate of how close the objective function of a metaheuristic solution is to optimality for a given type of problem (Wolsey 1998).

Reeves (1993) identifies three methods by which heuristic performance can be evaluated: 1) analytical methods, 2) statistical inference, and 3) empirical testing. Using the analytical method, it is possible to analyze the operations of some heuristics such that their *worst case or average performance* on a problem can be proven. For example, it has been proven that a particular heuristic for the traveling salesman problem will always produce a solution not more than 50% longer than the optimal (Johnson and Papadimitriou 1985). Such proofs, though mathematically challenging, are of limited practical use unless the performance bounds can be made tight. Also, local search methods (on which the metaheuristics of simulated annealing and tabu search are based), because of the random elements in their operation, have been shown to have no performance guarantee for the traveling salesman problem (Reeves 1993). Using the analytic method, it also possible to obtain a bound for a particular problem instance using some form of *relaxation* of the problem; e.g., a relaxation of integer constraints. Again, the usefulness of this approach to evaluating heuristic performance depends on how tight the gap is between the value of the bound and the heuristic solution.

Statistical inference has also been used to estimate the performance of a heuristic. This has been developed by Golden and Alt (1979) and is based on the statistical theory of extreme values (Fisher and Tippett 1928). Golden and Alt argue that each time we apply a heuristic to a minimization problem, we implicitly sample a large number, *m*, of possible solutions among which we find the minimum objective function value, v_i. As m approaches infinity, the distribution of v_i approaches the Weibull distribution. Given *n* independent solutions obtained in this way, it is possible to find a point estimate for the overall minimum and a confidence interval. Golden and Alt (1979) strongly supported the validity of this approach through extensive testing on the traveling salesman problem. Notwithstanding their efforts, according to Reeves (1993) there are few other reported applications of this approach in the literature. Among the few, Boston and Bettinger (1999) tested this approach on the 0/1 harvest scheduling problem and concluded that extreme value statistics provide unreliable estimates of the optimal objective function value and that the quality of the estimate is strongly dependent on the quality of solutions generated by the heuristic procedure.

The empirical approach to testing a heuristic involves comparing its performance with that of existing techniques on a set of benchmarks. Since benchmarks represent only a small fraction of the possible population of instances, they should be sufficiently representative of real problems. Testing a heuristic across a range of problem instances facilitates evaluation of how well the heuristic performs in general, and under what conditions it performs relatively well or poorly. Factors commonly discussed are the influence of problem size, solution quality, solution variance, and computational costs (Reeves 1993).

In research on the integer harvest-scheduling problem, some evaluation of metaheuristic methods has occurred. For example, Nelson and Brodie (1990) used Monte Carlo integer programming to solve a three-period problem with 291 binary decision variables to within 10% of the known optimal. Murray and Church (1995) compared tabu search, simulated annealing, and hill climbing on the same data set used by Nelson and Brodie: the simulated annealing metaheuristic produced solutions averaging within 8% of the objective function of the optimal, and tabu search within 6.3%. On a larger data set, comprised of 1,293 binary decision variables, Murray and Church (1995) found that simulated annealing averaged within 11.8% and tabu search within 4.4% of the known optima. Weintraub *et al.* (1995) interfaced metaheuristic decision rules with a continuous LP solver to produce an integer solution averaging within 6.4% of the known optimal for a transportation and scheduling problem comprising 191 integer decision variables. Boston and Bettinger (1999) compared tabu search and simulated annealing, on four problem instances ranging is size from 3,000 to 5,000 decision variables. The simulated annealing algorithm produced solutions within 3.4, 1.9, 0.2, and 3.5 percent of the known optima of the four problems, while the tabu search produced solutions within 6.3, 2.8, 0.0, and 4.3 percent, respectively.

In all of the above comparisons, the type of harvest-scheduling model evaluated was the unit-restricted model (URM). In the URM, the boundaries of all potential harvest-openings are predefined; *i.e.*, the boundary of each polygon in a problem instance equals the boundary of each potential cut-block (i.e., contiguous area of harvested forest).

There exists another model of the harvest-scheduling problem with opening size constraints, referred to as the area-restricted model (ARM). This model has *not* been fully evaluated. In the ARM, the boundaries of all potential cut-blocks are not predefined. Instead, polygons may be aggregated to form cut-blocks during the search for an optimal solution. The limit of this aggregation is defined by the maximum allowable opening area. It is important that an evaluation of metaheuristic algorithms used to solve the ARM be made, for two reasons: 1) the ARM is, arguably, a more suitable model of the harvest-scheduling problem than the URM; and 2) the ARM is potentially more difficult to solve. I discuss each reason in detail.

First, the ARM has several advantages over the URM. For example, it has been observed (Walters *et al.* 1999, Richards and Gunn 2000) that in the ARM, the configuration of cut-blocks emerges in the context of an optimal flow of timber; and that the predefined cut-blocks of the URM may underestimate the potential harvest flow through sub-optimal cut-blocks. Moreover, research indicates that poor block configuration can contribute to lower objective function values (Jamnick and Walters, 1991). Murray and Weintraub (2002), using several methods by which to pre-aggregate polygons into cut-blocks, demonstrate that the ARM consistently produces superior objective function values to the URM.

Second, the ARM may be more difficult to solve than the URM. My reasoning for this is that the number of feasible solutions to a given problem instance can be much greater when modeled as an ARM than as a URM. The degree to which this difficulty increases depends on the difference in *area* between the polygons and the maximum opening size restriction. For example, given an opening size restriction of 40 ha and an average polygon size of 5 ha, in one case, and an average polygon size of 15 ha in another case then, *ceteris paribus*, the number of feasible combinations in the former case will be much greater than in the latter. The URM avoids this type of expansion of the solution space. Murray and Weintraub (2002) for example, rightly observe that the URM can provide a lower bound on problems modeled as an ARM.

As mentioned above, in the case of the ARM, metaheuristic algorithms have not been fully evaluated. This is because formulations have only recently emerged for exact optimal solutions to the ARM (McDill and Braze 2002, Crowe *et al.* 2003, Caro *et al.* 2003). McDill and Braze (2002) solved problem instances up to 80 polygons over three periods, on a simulated forest designed to allow not more three polygons to aggregate within an allowed opening. Caro *et al.* (2003) evaluated a tabu search metaheuristic on six instances of a forest of 20 stands over 2 periods. Their metaheuristic solutions were within 1% of the optimal; but they note that the instances were too small to evaluate the true merit of their metaheuristic. Caro *et al.* (2003) also used a case study of 574 stands over seven periods and found that their tabu search algorithm produced solutions within 8% of an LP upper bound. Crowe *et al.* (2003) evaluated exact formulations of the ARM on five forests ranging in size from 346 to 6,093 polygons scheduled over 3 periods. These provide solid benchmarks by which to evaluate metaheuristic solutions to the ARM.

The objective of this paper is to use an empirical approach to evaluate the ability of a simple implementation of a metaheuristic algorithm, simulated annealing, to solve a variety of instances of the harvest scheduling problem modeled as an ARM.

Each of these instances has been solved optimally using the formulation methods discussed in Crowe *et al.* (2003), and therefore optimal benchmarks are to be used in the evaluation. By testing a metaheuristic across a range of instances, the objective is to gain some idea on how well it performs in general, and in what circumstances it does relatively well or relatively badly. By using a simple implementation of the simulated annealing algorithm, the intention is to provide an empirical worst-case analysis of the potential of metaheuristic search algorithms in solving the ARM. While simulated annealing has been shown to be quite effective relative to other metaheuristic algorithms in solving the ARM (Bettinger et al. 2000), the extraordinary flexibility of the metaheuristic approach to problem solving (Herz and Widmer 2003) makes it inappropriate for us to speculate on its unexplored potential.

The outline of this paper is as follows: first, I present the formulation of the area-restricted harvest-scheduling problem. Second, I describe the benchmark problem instances to be used in this study. Third, I describe the simulated annealing algorithm used to solve these problem instances. Fourth, I present results, comparing the objective functions of solutions produced using simulated annealing with those of optimal solutions. Finally, I discuss these results and conclude with suggestions on further research.

Problem Definition and Model

Formulations of the ARM are tested by scheduling harvest activities over three 20-year periods in a forest with an 80-year rotation period. The objective is to maximize net present value of the harvest activities, where value is measured at \$100 *per* m^3, discounted from the middle of each planning period at 4%. Harvest activities are restricted to clearcutting or doing nothing. The constraints in this problem are:

1) a limit on inter-period harvest volume fluctuations of plus or minus 10%;
2) a limit on opening size of 40 ha, with a 1-period restriction on harvesting of all adjacent polygons (adjacency is defined as sharing a common boundary node);
3) a minimum harvest age of 80 years;
4) an ending-age constraint to ensure that the average age of the forest at the end of the planning horizon be at least 40 years (*i.e.*, the average age of a forest regulated on an 80-year rotation). The formulation of the ending-age constraint is from McDill and Braze (2000).

The notation and mathematical formulation of the model are presented below:

x_{it} = *1 if polygon i is harvested in period t = 1, 2, or 3; 0 otherwise*

[Note, if polygon i has not been cut for the entire planning horizon, then x_{it} *=1 when t = 0. This is done to implement the ending age constraint].*

v_{it} = *volume of polygon i in period t (m^3).*

r_{it} = *discounted net revenue from harvesting polygon i in period t (\$).*

a_i = *area of polygon i (ha)*

h_t = *total volume harvested in period t (m^3)*

e_{it} = *ending age of polygon i if harvested in period t (ending age is measured in year 60)*

I = *set of all polygons in the forest*

O_{it} = *set of stands, including stand i, that are harvested during period t, and are included in the same opening as stand i (i.e., contiguously connected to stand i)*

Omax = maximum opening area for harvest-blocks (ha)

Objective Function: maximize net present revenue.

$$[1] \quad \text{Maximize} \quad \sum_{i=1}^{I} \sum_{t=1}^{3} r_{it} x_{it}$$

Subject to:

Each polygon may be assigned not more than one prescription over the planning horizon.

$$[2] \quad \sum_{t=0}^{3} x_{it} \leq 1 \qquad \forall\, i = 1,\ldots,I$$

Accounting variables defining the total volume harvested in each period.

$$[3] \quad \sum_{i=1}^{I} v_{it}\, x_{it} - H_t = 0 \qquad \forall\, t = 1,\ldots,3$$

Limit on total area of polygons aggregating into harvest blocks for each period.

$$[4] \quad \sum_{i \in O_{it}} x_{it} a_i \leq Omax \qquad \forall\, i \in I;\, t = 1,\ldots,3$$

Inter-period harvest volumes may fluctuate at most by 10%.

$$[5] \quad -.9H_t + H_{t+1} \geq 0 \qquad t = 1,2$$

$$[6] \quad -1.1Ht + H_{t+1} \leq 0 \qquad t = 1,2$$

Average polygon age of total forest area must be at least 40 years at the end of planning horizon (year 60).

[7] $$\sum_{i=1}^{I} \sum_{t=0}^{3} (e_{it} - 40) a_i x_{it} \geq 0$$

The decision variables are binary

[8] $$x_{it} \in \{0,1\} \qquad \forall i = 1,\dots,I;\ t = 0,\dots,3$$

Description of Benchmarks

Spatial data from seven forests in British Columbia were used in this study. Table 3.1 summarizes the spatial attributes of these forests, including the size of polygons and the mean number of adjacencies per polygon. The number of opening constraints per period refers to the number opening constraints used to control opening size in calculating the optimal solution using the branch and bound algorithm. Polygons in the smaller forests, Gavin, Hardwicke, and Naka were manually drawn by forest engineers, while the polygons of Stafford and Kootenay were formed by GIS-overlays. The forest of Kootenay I is a subset of polygons extracted from the larger forest, Kootenay II, which is similarly a subset of Kootenay III. Polygons ranged in size from 0.06 to 43 ha (a few polygons marginally greater than the 40 ha opening limit are allowed to be harvested by themselves). In addition, the number of opening constraints per period is listed in this table: this represents the number of linear constraints needed to prevent harvest-openings from exceeding 40 ha in the models solved using the branch and bound algorithm.

Table 3.1: Spatial attributes of the seven benchmark forests.

Forest	Total Area (ha)	No. of Polygons	Polygon Area (ha) (minimum, maximum, average)	Opening Constraints per period	Polygon Adjacencies (minimum[1], maximum, average)
Gavin	6,193	346	1.1, 43.0, 17.9	1,925	1, 16, 5.8
Hardwicke	6,948	423	3.0, 42.0 16.4	44,495	2, 10, 6.1
Naka	10,934	785	0.06, 43.0, 13.9	21,272	1, 14, 5.2
Stafford	10,421	1,008	2.0, 20.0, 10.3	16,049	0, 12, 4.3
Kootenay I	38,441	3,256	6.0, 38.0, 11.8	50,110	0, 15, 5.6
Kootenay II	71,257	6,093	6.0, 38.0, 11.7	97,360	0, 15, 5.2
Kootenay III	118,535	12,090	6.0, 39.5, 9.8	156,563	0, 15, 5.2

1. A few isolated polygons in Kootenay I, II, and III, and Stafford have no neighbours.

I was also interested in testing different age-class distributions for each forest. McDill and Braze (2000) found a strong correlation between the initial age-class distribution of a forest and the difficulty of solving the URM using a branch and bound algorithm; *viz.*, that problems with a higher proportion of old-growth forest area are, in general, more difficult to solve than others. For this reason, I assigned three age-class distributions to each forest to test the formulations on each. The age-class distributions are based on the percent of polygons eligible, by age, for harvest in the first period. The three distributions tested are: 100%, 75%, and 50%. Table 3.2 summarizes the uniform age distributions assigned to each forest using a random number generator.

Table 3.2: Three age distributions randomly assigned to each forest.

Percent of Polygons Eligible for harvest in Period 1	Range of Uniform Distribution of Ages
50%	10-150 years
75%	40-200 years
100%	80-200 years

In addition to the effect of different age class distributions, I also wished to explore the effects of different maximum opening-size limits. Increasing the opening size limit allows for more possible combinations of polygons to aggregate into a feasible opening, thereby increasing the number of possible solutions. For this reason, I tested 4 different opening size limits on the Stafford forest: 20, 30, 40, and 50 ha opening size limits, respectively.

The methods used to compute optimal solutions for these forests are described fully in Crowe *et al.* (2003).

Description of the Simulated Annealing Algorithm

The simulated annealing algorithm I used to solve the harvest-scheduling problem is similar to that designed and illustrated by Boston and Bettinger(1999), except for the choice of parameters: *viz.*, the initial temperature, the reduction factor, the definition of a neighbourhood, and the number of iterations at a given temperature. To illustrate the context of these parameter choices, I first present a concise outline (Wolsey, 1998) of the simulated annealing metaheuristic.

1. Get an initial solution, S.
2. Get an initial temperature,**T**, and a reduction factor, **r** with $0 < r < 1$.
3. While T is not frozen, do the following:
   3.1 Perform the following loop **nrep** times:
      3.1.1 Pick a random **neighbour** S' of S
      3.1.2 If S' is feasible:
         3.1.2.1 Let delta = f(S') – f(S).
         3.1.2.2 If delta >= 0, set S = S'
         3.1.2.3 If delta < 0, set S = S' with probability $e^{-\text{delta}/T}$
   3.2 Set T = rT
4. Return best solution.

I defined a *neighbour* of S to be any solution arising from the following operation:

i) Randomly select any binary decision variable, x_{ij}, from the current solution, S (where $x_{ij} = 1$ if polygon i is harvested in period j, 0 otherwise).

ii) If $x_{ij} = 0$, let $x_{ij} = 1$.

iii) Else if $x_{ij} = 1$, let $x_{ij} = 0$.

iv) If this permuted solution is feasible, then it is a neighbour, S'.

Note that only feasible solutions can be accepted in this algorithm. The manner of defining a neighbour of S was based on the idea that small neighbourhoods are preferable to large complex ones (Reeves 1993). I should note the choice of a neighbourhood structure can itself be the subject of some experimentation, and that Bettinger, Boston and Sessions (1999) have initiated such research for a tactical planning problem. The choice of an initial feasible solution followed naturally from this choice of a neighbourhood: I set all decision variables equal to zero.

The parameters requiring experimentation were r, T, and nrep. An experiment was designed to find the best values for these parameters within a reasonable period of time, and it was complicated by the fact that the annealing algorithm is randomized, and therefore results of a single run may not be typical. To simplify the execution of the experiment, I set r = 0.99 as a constant. This choice was supported by the literature on simulated annealing which indicates that most reported successes use values between 0.8 and 0.99 with a bias to the higher end of the range (Reeves 1993).

The value of an initial experimental temperature, T, was determined by setting r = 1.0, iteratively running the algorithm and raising T, until all feasible solution changes were accepted with at least 95% probability. The intention here was to find a value for T allowing an almost free exchange of neighbouring solutions so that the final solution would be independent of earlier solutions.

Finally, the value of nrep was chosen. This value governs the number of repetitions at each temperature, and the theory of simulated annealing suggests that the temperature should converge gradually to a value of zero. If computation times were of no consequence, then nrep should vary exponentially with problem size. I wished this to be a practical test of the algorithm, and were therefore forced to limit nrep. Since the number of iterations at each temperature is related to the size of the neighbourhood (Reeves 1993), I began with a value of nrep equal the number of binary decision variables in each problem instance. I then iteratively doubled the value of *nrep* and re-ran the algorithm until convergence at a sufficiently cool temperature occurred. The stopping criterion was the completion of a fixed number of iterations, N. In this way, the cooling rate can be regarded as a constant relation between the number of iterations and the value of nrep; i.e., if one wishes to run the simulated annealing algorithm for a longer period of time, one must multiply both nrep and N by the same factor.

The harvest-scheduling model using the simulated annealing algorithm was encoded in the C programming language, using Microsoft Visual C++ 6.0. Optimal solutions were computed using CPLEX 8.1. All models were solved using a 2.4 GHz Pentium 4 CPU, with 1.5 gigabytes of RAM, and a Windows XP operating system.

Results

The solution qualities and variances for all problem instances solved using the simulated annealing algorithm are presented in Table 3.3. In Table 3.3, *mean solution quality* refers to the mean objective function value of the solutions produced by simulated annealing divided by the optimal objective function value; *gap variance* refers to the standard deviation of those solutions divided by the optimal objective function value. I use this measure because it describes not only how widely values are dispersed around the mean metaheuristic objective function value, but also because it describes this value in terms of the variance in the *gap* between the metaheuristic and optimal solution values. Each mean requiring three minutes of computing time is based on a sample of thirty solutions; while the means requiring thirty minutes are based on a sample of fifteen solutions. Three-hour runs were executed for any problem yielding a mean objective function value of less than 90% of the optimal on a thirty-minute run. These are presented in Table 3.4. The simulated annealing algorithm evaluated, on average, 136 million solution permutations *per* minute. Results from the Kootenay III forest are not included in Table 3.3 for the instance where 100% of the stands are initially eligible. This is because no optimal solution was found-- the branch and bound solution tree required more memory than was available.

Table 3.3: Solution quality and gap variance from applying simulated annealing to benchmark instances.

Forest	% of stands eligible for harvest	Mean Solution Quality after 3 minutes	Gap variance after 3 minutes	Mean Solution Quality after 30 minutes	Gap Variance after 30 minutes
Gavin	50%	94.2%	0.68%	94.5%	0.70%
	75%	95.4%	0.73%	95.5%	0.20%
	100%	94.9%	0.63%	95.0%	0.82%
Hardwicke	50%	98.0%	0.51%	98.0%	0.55%
	75%	98.3%	0.42%	98.4%	0.58%
	100%	98.6%	0.72%	98.8%	0.42%
Naka	50%	89.3%	0.11%	91.4%	0.49%
	75%	90.5%	0.32%	92.5%	0.32%
	100%	91.8%	0.29%	93.4%	0.34%
Stafford 20 ha	50%	96.5%	0.35%	97.1%	0.26%
	75%	95.2%	0.60%	96.1%	0.61%
	100%	97.1%	0.37%	97.7%	0.36%
Stafford 30 ha	50%	98.4%	0.16%	98.4%	0.18%
	75%	98.5%	0.07%	98.5%	0.09%
	100%	98.7%	0.04%	98.7%	0.05%
Stafford 40 ha	50%	98.8%	0.05%	98.8%	0.06%
	75%	98.4%	0.04%	98.5%	0.04%
	100%	98.6%	0.02%	98.6%	0.04%
Stafford 50 ha	50%	99.1%	0.05%	99.1%	0.06%
	75%	98.4%	0.03%	98.4%	0.03%
	100%	98.6%	0.08%	98.7%	0.05%
Kootenay I	50%	90.7%	0.32%	91.0%	0.13%
	75%	92.3%	2.06%	92.5%	1.07%
	100%	93.8%	1.26%	93.9%	0.98%
Kootenay II	50%	87.1%	1.14%	87.1%	1.46%
	75%	87.3%	2.56%	87.4%	2.31%
	100%	88.5%	1.72%	88.6%	1.57%
Kootenay III	50%	88.5%	0.27%	88.5%	0.17%
	75%	90.81%	1.20%	90.8%	0.78%

Table 3.4: Solution quality and gap variance results from 3-hour runs using the simulated annealing algorithm

Forest	% Stands Harvestable in Period 1	Mean Solution Quality after 3 Hours	Gap Variance after 3 hours
Kootenay II	50%	87.6%	1.29%
	75%	87.9%	1.31%
	100%	88.8%	1.57%
Kootenay III	50%	88.9%	0.17%
	75%	91.1%	0.78%

The objective in testing the simulated annealing algorithm was to observe how well it performs in general, and under what circumstances it will do relatively poorly. This objective will guide analysis of the results.

First, *in general*, the mean solution produced by the simulated annealing algorithm, for all 29 instances, using thirty minutes of computing time, was 94.97% of the optima, and the mean gap variance (as defined above) was +/- 0.51%.

Second, in *particular*, the effects of attributes of particular problem instances which interested us were:

1. the number of decision variables;
2. the percent of stands initially eligible for harvest; and
3. the ratio between the mean polygon area and the maximum opening area.

I examine these separately.

The relation between the number of decision variables and the solution quality, for each of the thirty-minute runs are presented in Figure 3.1.

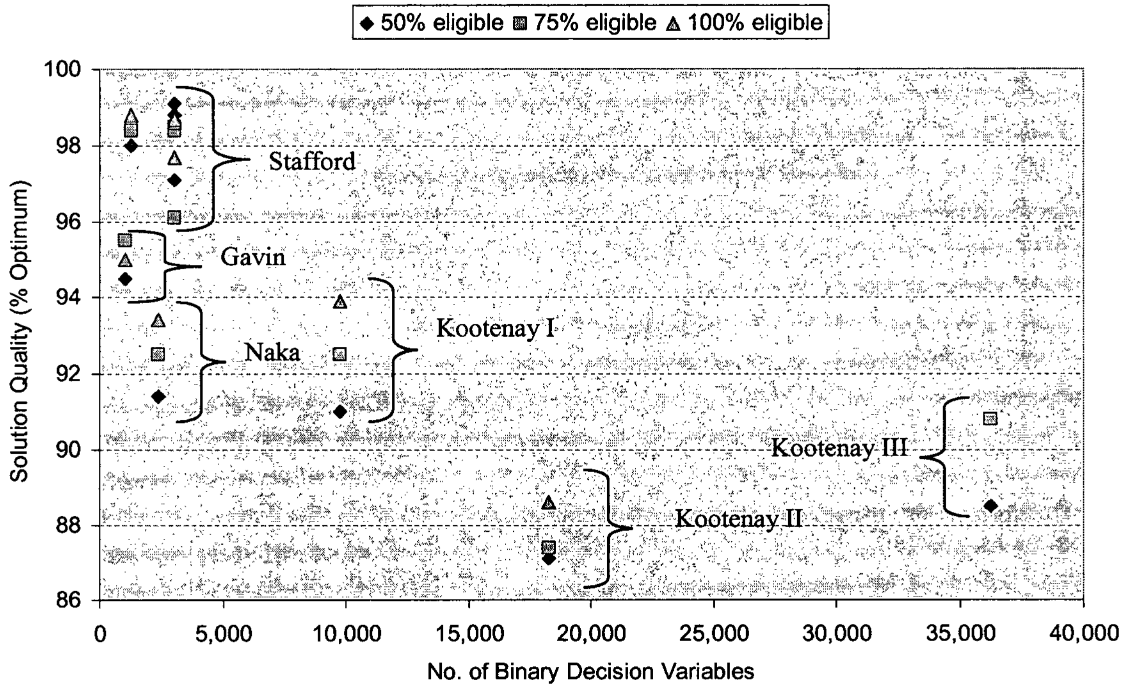

Figure 3.1: Relation between mean solution quality and number of binary decision variables (Note: number of decision variables equals number of polygons times three periods).

Figure 3.1 indicates that there is a downward trend between problem size and the quality of solutions produced by the simulated annealing algorithm. The problem instances which most weaken this trend are from the Naka and Kootenay II forests, comprised of 2,355 and 18,279 binary decision variables respectively. Each of these forests yields solutions relatively inferior to those of larger forests. Clearly there is an attribute, or attributes, other than problem-size also influencing the algorithm's search for better solutions. I considered whether the relatively poor solution qualities of Naka and Kootenay II might have arisen from either: a) the chance that the particular random assignment of ages might have made the problem instance difficult to optimize using neighbourhood search; or, b) the particular spatial arrangement of the polygons in these forests. Looking at Figure 3.1, and noting that variance in solution quality *between* the seven forests is greater than the variance *within* the three different age-classes randomly assigned to each forest, I was inclined to pursue b) as a possible venue for explanation. Unfortunately, the spatial attributes of the forests, presented in Table 3.1, reveal no attribute by which to differentiate both Naka and Kootenay II from the other forests.

The influence of the number of binary decision variables is more pronounced, however, when I compare the progress made on solution quality when moving from the three-minute to thirty-minute runs. Here the progress made on the smaller Naka problems produced a mean improvement of 1.90%, while on the much larger forests of Kootenay II and III, mean progress was 0.20% and 0.07%, respectively. In fact, the solutions to the Kootenay II and III forests benefited very little from the additional 3 hours of computing time (Table 3.4), improving upon the 30-minute solution quality by a mean of 0.47% and 0.35%, respectively.

Apart from the effect of the number of decision variables, I was also interested in the effect of the *mean polygon size* upon the solution quality. In the Introduction of this paper, I reasoned that the smaller the mean polygon area relative to the maximum opening area, the more feasible combinations there would be, the larger the solution space would be, and therefore the more difficult it would be to find near-optimal solutions. The results in Table 3.3 do not support this speculation. For example, in the Stafford forest, the solution quality improves from a 20 ha opening to a 50 ha opening. Figure 3.2 indicates that there is no apparent trend between solution quality and the ratio of mean polygon area to maximum opening area for all of the instances solved.

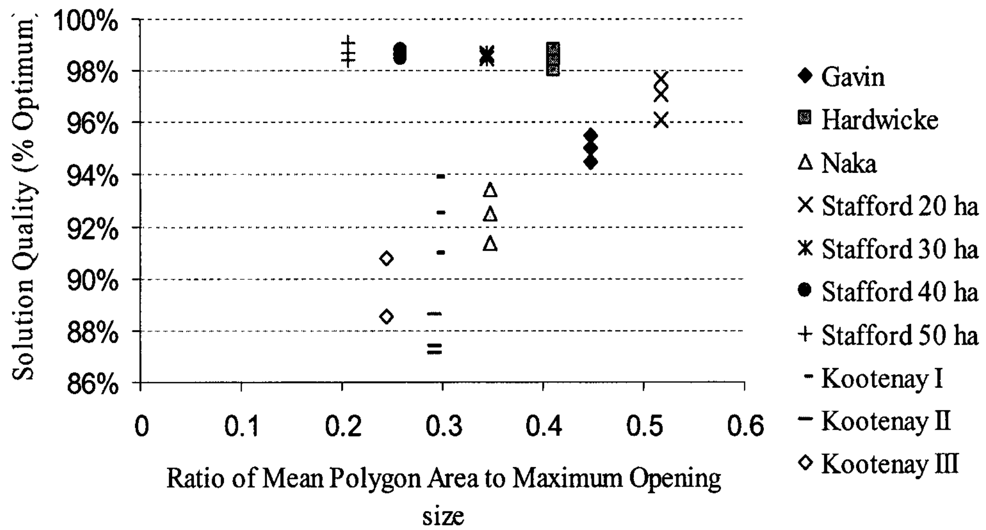

Figure 3.2: Relation between qualities of solutions produced by simulated annealing algorithm and the ratio of mean polygon area to maximum opening area. Results illustrated are from the 30-minute runs, with different solutions for each of the three initial age-class distributions per forest.

Finally, I was interested in the effect of initial age-class distribution on solution quality. The results in Table 3.3 reveal that six of the seven forests yielded their worst mean solution qualities to the youngest forest (50% initially eligible) while 5 of the 7 forests yielded their best mean solution qualities to the oldest forests. This indicates that, problems with fewer eligible stands are more difficult to optimize for a metaheuristic than problems where more stands are eligible. These findings are contrary to those for the branch and bound algorithm, where older forests are more difficult to optimize.

Discussion

The results uncover several trends worthy of discussion. The first, and perhaps most important trend, is that problem size does not appear acutely to affect the ability of the simulated annealing algorithm to find near-optimal solutions. A weak trend was observed between larger problem instances and poorer solution qualities; but the decline in quality was not steep. The simulated annealing algorithm produced excellent results, on average within 5% of the optima over the range of instances tested. Of course, the number and size of the instances solved in this research cannot allow us to generalize this trend with certainty; but the results can provide some confidence to practitioners currently using neighbourhood-search metaheuristics to produce efficient tactical-level plans in forest management. Not all neighbourhood-search metaheuristics have been shown to cope equally well in finding near-optimal solutions as problem size increases (see Johnson and McGeoch 2002 for heuristic algorithms used on increasingly larger instances of the traveling salesman problem).

A second interesting result provided by this research, especially from the the Stafford forest, is that the ratio of the mean polygon area to the maximum opening area did not influence the quality of the best solution found by the metaheuristic. As noted earlier, the smaller this ratio is, the greater is the number of feasible solutions (other things being equal). I am therefore obliged to ask: Why did problems with a smaller ratio *not* yield solutions of lower quality, given the expansion of the solution space? My answer to this is complemented by reflection on another question raised by this research: *viz*., why is it that forests with 100% of the stands initially eligible for harvest yielded solutions of higher quality than forests with only 50% of the stands initially eligible? This question complements the first because, in both cases, problem instances with relatively more feasible solutions yield solutions which are of equal or higher quality than instances with relatively fewer feasible solutions. Why?

The number of feasible solutions relative to the number of decision variables influences solution quality because it influences the neighbourhood search. For example, let S be a set of feasible solutions to a particular problem, and N(s) be the neighbourhood of a solution, s. In neighbourhood search, N(s) is defined as the set of solutions which can be obtained from s by performing a simple permutation operation on s. But not all such permutations upon s produce a solution within the set N(s), because some of these permutations produce infeasible solutions. Hence, if there are, for example, more feasible solutions in problem A than in problem B, then *on average*, each neighbourhood of each solution in A will have a greater number of members than each neighbourhood of each solution in B. This can have two major effects on the neighbourhood-search in problems A versus B: 1) more time is lost in problem B than in problem A by producing infeasible solutions through a permutation operation; and, more importantly 2) the ability of the algorithm to *diversify the search*, i.e., enter new regions of the solution space is, hampered when neighbourhoods are smaller. Hertz and Widmer (2003), for example, have also observed the relative ineffectiveness of neighbourhood-search in highly constrained problems where permutation-operations rarely produce a feasible solution. To promote diversity in the search, they suggest relaxing some constraints and adding penalties to the objective function.

A final point of discussion requires us to interpret the *general* results: a mean solution quality of almost 95%, generally with very little variance. What does this mean for the future of research on metaheuristic applications to the ARM? Clearly, there is little room for improvement, given that this is a worst case analysis. Closing the gap on the final 5% between metaheuristic versus optimal solution quality may present itself as an interesting challenge to researchers; but what might be the practical merit of this? Efficient solutions produced by a symbolic model rarely translate into equivalent results in operational realities. Hence, for the tactical harvest scheduling problem, the minor increases in NPV which might arise by improvements in metaheuristic planning algorithms may never materialize, given uncertainties in field data, growth and yield data, and estimates of harvested log grades and values.

The results of this research therefore point in one direction for future relevant research on applying metaheuristic algorithms to the ARM: evaluate algorithms using much larger problems instances. There are several reasons for this. First, since there was a trend observed between problem size and solution quality, it would be useful to explore this further. Second, the results revealed that the simulated annealing algorithm, when applied to the larger problems (Kootenay II and III), made very little improvement to solution quality between three-minute versus three-hour runs. This indicates that, on very large problems, it can be much more challenging to explore truly different regions of the search space effectively. Improvements in search diversification strategies should therefore be evaluated in the context of larger problem instances. Finally, although the optimal benchmarks by which metaheuristic solutions to larger instances of the ARM are to be evaluated may not be computationally feasible, the research of Crowe *et al.* (2003) and Caro et al. (2003) indicates that LP relaxations of the ARM provide reasonable estimates of the upper bounds.

Conclusions

The objective of this paper was to apply the simulated annealing algorithm to a variety of instances of the area-restricted harvest-scheduling model, and to evaluate its approximately optimal solutions by comparison with optimal benchmarks. Of the 29 instances solved, the average deviation from the optima was only slightly more than 5%. Attributes of the problem, such as the number of decision variables, maximum opening size, and initial age-class distribution were examined for their effect on the metaheuristic's ability to produce good solutions. A weak downward trend was observed on the relationship between solution quality and problem size.

This research is significant because it constitutes the first evaluation of the ability of a metaheuristic algorithm to produce good quality solutions for the ARM. Given that the application of the simulated annealing algorithm required no innovation and was relatively simple to implement, the results constitute a worst-case analysis of the potential of the metaheuristic approach to this problem. The excellent results produced by this worst-case case analysis, coupled with weak downward trend on the relation between problem size and solution quality, suggest that future research on applying metaheuristics to the ARM use relatively large problem instances.

Literature Cited

American Forest and Paper Association. 2001. Sustainable Forestry Initiative (SFI)^{SM} Standard. http://www.afandpa.org/forestry/sfi_frame.html (accessed 6/21/01).

Baskent, E.Z., and Jordan, E.A. 2002. Forest landscape management using simulated annealing. Forest Ecology and Management 165:29-45.

Bettinger, P., Sessions, J., and Boston, K. 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. Ecological Modeling 94: 111–123.

Bettinger, P., Boston, K., and Sessions, J. 1999. Intensifying a metaheuristic forest harvest scheduling search procedure with 2-opt decision choices. Canadian Journal of Forest Research 29: 1784–1792.

Bettinger P., Graetz, D., Boston, K., Sessions, J., and Chung, W. 2000. Eight metaheuristic planning techniques applied to three increasingly difficult wildlife planning problems. Silva Fennica 36(2):561–584.

Boston, K., and Bettinger, P. 1999. An analysis of Monte Carlo integer programming, simulated annealing, and tabu search metaheuristics for solving spatial harvest-scheduling problems. Forest Science 45(2): 292-301.

Boston, K., and Bettinger, P. 2001. Development of spatially feasible forest plans: a comparison of two modeling approaches. Silva Fennica 35(4): 425–435.

Boston, K. and Bettinger, P. 2002. Combining tabu search and genetic algorithm metaheuristic techniques to solve spatial harvest scheduling problems. Forest Science 48:35-46.

Brumelle, S., Granot, D., Halme, M., and Vertinsky, I. 1998. A tabu search algorithm for finding a good forest harvest schedule satisfying green-up constraints. European Journal of Operational Research 106:408-424.

Caro, F., Constantino, M., Martins, I., and Weintraub, A. 2003. A 2-opt tabu search procedure for the multiperiod forest harvesting problem with adjacency, greenup, old growth, and even flow constraints. Forest Science 49(5):738-751.

Clark, M.M., Meller, R.D., and McDonald, T.P. 2000. A three-stage metaheuristic for harvest-scheduling with access road network development. Forest Science 46:204-218.

Dahlin, B., and Sallnas, O. 1993. Harvest-scheduling under adjacency constraints – a case study from the Swedish sub-alpine region. Scandinavian Journal of Forest Research 8:281-290.

Crowe, K., and Nelson, J. 2003. An indirect search algorithm for harvest-scheduling under adjacency constraints. Forest Science 49(1):1-11.

Crowe, K., Nelson, J., and Boyland, M. 2003. Solving the area-restricted harvest-scheduling problem using the branch and bound algorithm. Canadian Journal of Forest Research 33:1804-1814.

Fisher, R. and Tippett, L. 1928. Limiting forms of the frequency distribution of the largest or smallest member of a sample. Proceedings of the Cambridge Philological Society 24:180-190.

Golden, B.L., and Alt, F.B. 1979. Interval estimation of a global optimum for large combinatorial problems. Naval Research Logistics Quarterly 26:69-77.

Hertz, A., and Widmer, W. 2003. Guidelines for the use of metaheuristics in combinatorial optimization. Eurpoean Journal of Operational Research 151:247-252.

Johnson, D. , and McGeoch, L. 2002. Experimental analysis of metaheuristics for the symmetric traveling salesman problem. *In*: The traveling salesman problem and its variations. (Eds. Gutin and Punnen), Kluwer Academic Publishers.

Johnson, D.S., and Papadimitriou, C.H. 1985. Performance guarantees for heuristics. Interfaces (2):145-180.

Jamnick, M.S., and Walters, K.R. 1991. Spatial and temporal allocation of stratum-based harvest schedules. Canadian Journal of Forest Research 23: 402-413.

Liu, G., Nelson, J., and Wardman, C. 2000. A target-oriented approach to forest ecosystem design — changing the rules of forest planning. Ecological Modeling 127:269–281.

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

Martell, D.L., Gunn, E.A., and Weintraub, A. 1998. Forest management challenges for operational researchers. European Journal of Operational Research 104:1-17.

McDill, M.E., and Braze, J. 2000. Comparing adjacency constraint formulations for randomly generated forest planning problems with four age-class distributions, Forest Science 46: 423-436.

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

Murray, A.T., and Church, R.L. 1995. Metaheuristic solution approaches to operational forest planning problems. OR Spektrum 17:193-203.

Murray, A.T., and Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. Forest Science 48: 779-789

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

O' Hara, A.J., Faaland, B.H., and Bare, B.B. 1989. Spatially constrained timber harvest-scheduling. Canadian Journal of Forest Research 19: 715 – 724.

Ohman, K., and Eriksson, L.O. 1998. The core area concept informing contiguous areas for long-term forest planning. Canadian Journal of Forest Research 28: 1817–1823.

Reeves, C.R. (Editor). 1993. Modern metaheuristic techniques for combinatorial problems. Blackwell Scientific Publications, Oxford, U.K.

Richards, E.W. and Gunn, E.A. 2000. A model and tabu search method to optimize polygon harvest and road construction schedules. Forest Science 46(2):188-203.

Richards, E., and Gunn, E.A. 2003. Tabu search design for difficult forest management optimization problems. Canadian Journal of Forest Research 33: 1126–1133

Sessions, J., and Bettinger, P. 2001. Hierarchical planning: pathway to the future? In Proceedings of the First International Precision Forestry Symposium, 17–20 June 2001, Seattle, Wash. College of Forest Resources,University of

Washington, Seattle, Wash. pp. 185–190.

Van Deusen, P.C. 1999. Multiple solution harvest scheduling. Silva Fennica 33: 207–216.

Van Deusen, P.C. 2001. Scheduling spatial arrangement and harvest simultaneously. Silva Fennica 35: 85–92.

Walters, K.R., Feunekes, H., Cogswell, A., and Cox, E. 1999. A forest planning system for solving spatial harvest-scheduling problems. Canadian Operations Research Society National Conference, June 7-9 1999. Windsor, Ontario. 8 pp. Available from http://www.remsoft.com [accessed December 2003].

Weintraub, A, Jones, G., Meacham, M., Magendzo A., and Malchuk, D., 1995. Metaheuristic procedures for solving mixed-integer harvest-scheduling transportation planning models. Canadian Journal of Forest Research 25:1618–1626.

Williams, H.P. 1999. Model building in mathematical programming: fourth edition. John Wiley and Sons Ltd., New York. 354 pp.

Wolsey, L.A. 1998. Integer programming. John Wiley and Sons Ltd., New York. 264pp.

Chapter IV

Integrating the Strategic Allocation of Cutting Rights with Spatially Explicit Timber Supply Planning

Introduction

In British Columbia and elsewhere, publicly owned forest is often managed under volume-based tenure systems. Strategic management plans are made for each forest, and the allocation of cutting rights to forest companies is then made. In some cases, an entire forest management unit is allocated to one licensee; but in other cases, the allocation is made to several competing licensees. The problem of allocating cutting rights among competing licensees entails assigning large discrete units of forest land to these licensees, typically for periods of fifteen to twenty-five years. The units of forest land, known as chart areas, are large because it is economically infeasible for firms to maintain roads and operations widely dispersed across an entire management unit. The assignment of chart areas is a problem with several complicating factors, which will now be described.

First, it has been the government's historic policy to allocate cutting rights to licensees on the condition that they own wood-processing facilities. This is referred to as the appurtenance policy (Pearse 2001). This policy renders the allocation of cutting rights problematic because not all mills process all types of logs with equal economic efficiency. In the last twenty years, many new specialty mills have arisen; and these mills, given certain types of logs, can make more valuable wood products than traditional volume based mills. Hence, solving the classic forestry problem of allocating the "right log to the right mill" is frustrated by the discrete allocation of chart areas to licensees. Failure to allocate logs optimally, in the long run, results in reduced net returns on timber, in the form of both profits to the private sector and stumpage to the government. Clearly, solving the problem of assigning chart areas to licensees necessitates an exploration not only of the costs of this appurtenance policy with regard to inefficient allocation of standing timber; but also, a practicable method by which the appurtenance policy can be relaxed in order to increase the efficient allocation of standing timber. The government of British Columbia is currently considering a policy of relaxing the appurtenance policy through allocation of some chart areas to timber sales.

The second complicating factor in the assignment of chart areas to licensees arises from the introduction of landscape-scale, spatially explicit, harvesting constraints (e.g., targets for seral patch size distributions, old-growth reserves, preservation of special habitat conditions) applied to the management unit as a whole. The imposition of these constraints could potentially cut across the boundaries of chart areas, thus affecting the timber supply of one licensee more than another. These constraints are also important for certification schemes (e.g., Canadian Standards Association), where a certification plan must be made for an entire management unit, regardless of the competing interests of its multiple licensees. Hence, there is a need for a method by which the imposition of spatial constraints may combine harmoniously with the assignment of chart areas to licensees.

The third complicating factor in the assignment of chart areas to licensees is that objectives to be satisfied are, for each licensee, multiple and conflicting. For example, the assignment must not only satisfy each licensee's volume targets over time, but also its targets for certain species and log sizes. In addition, each licensee seeks to minimise the haul distance arising from the assignment, and to secure enough standing timber accessible for harvesting in the winter months. Since the supply of timber is scarce, these conflicting objectives also exist *between* licensees who must compete against one another to satisfy their objectives.

The problem of assigning chart areas to licensees is illustrated in Figure 4.1. Here a forest management unit is divided into 14 chart areas, each of which must be assigned to one of 3 licensees or to timber sales, through which logs may be redistributed to licensees. In this problem, we cannot assume equivalence between a licensee and a mill; for, a licensee may own more than one mill. If a licensee has more than one mill, the licensee's targets are the sum of its mill's targets. The assignment is based on satisfying each licensee's multiple timber and operational objectives over several decades. In addition, all scheduled harvesting must occur such that landscape-scale ecological objectives are satisfied.

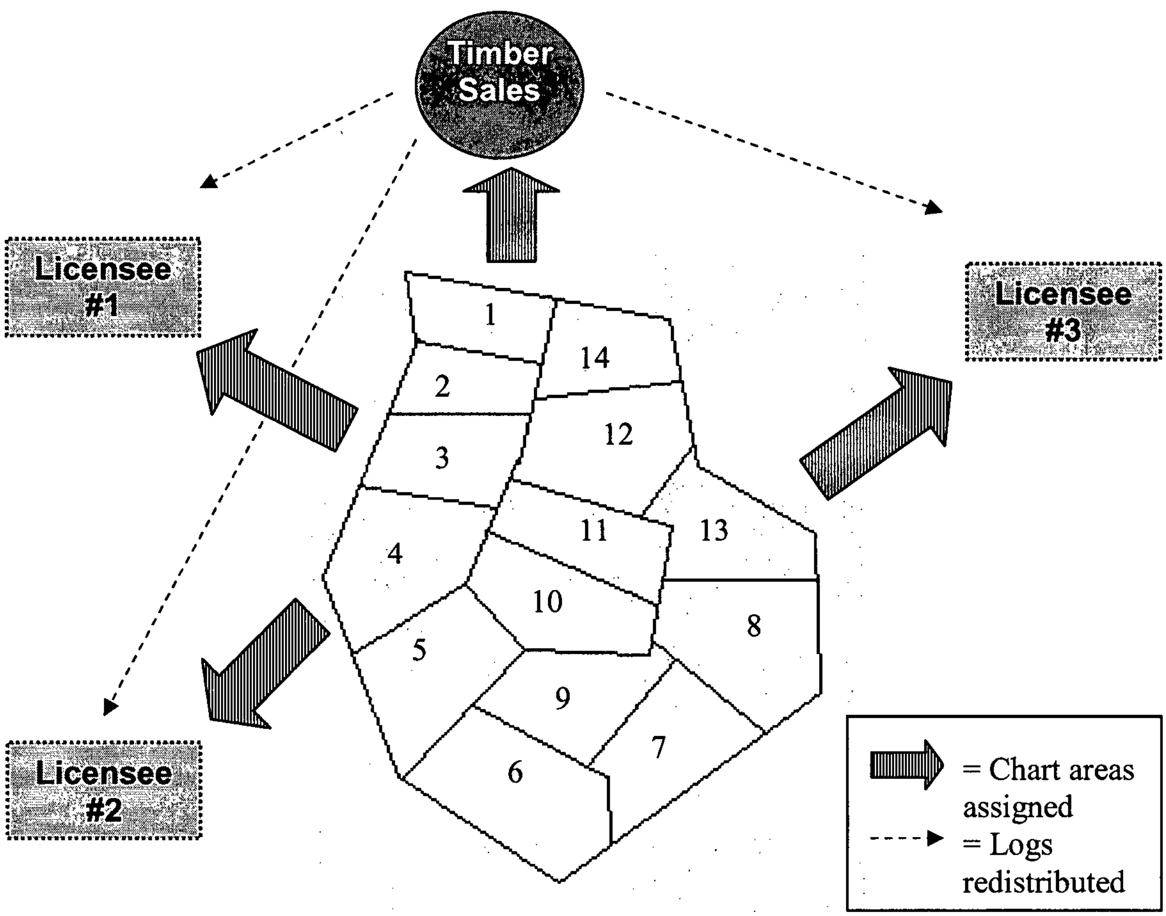

Figure 4.1: A forest management unit divided into 14 chart areas, each of which must be discretely assigned once to either of 3 licensees or timber sales.

The complexity and economic consequences of this problem warrant the design of a decision support model. It is my intention in this paper, to design, test, and evaluate a model for the problem of assigning chart areas to licensees and timber sales. First, a brief literature review on the strategic allocation of standing timber is presented. This will be followed by a description and formulation of the model of assigning chart areas to licensees and to timber sales. Next, the model will be tested on a case study, the Kootenay Timber Supply Area, in British Columbia. Results and discussion will follow, and I will conclude with a balanced evaluation of this model and suggest future research.

Literature Review

The allocation of cutting rights among competing mills may be regarded as an extension of the log allocation problem. Pearse and Sydneysmith (1966) define this problem succinctly: "Given a certain heterogeneous supply of logs in a particular production period, how should this raw material be allocated among the available utilization facilities?" The allocation of cutting rights is similar in its objective, but instead of allocating logs in particular periods, one must allocate heterogeneous groups of standing timber, i.e., chart areas, with the intent to produce certain logs in certain periods.

In the literature, the log allocation problem has been extended from allocating logs to allocating standing timber. Walker and Preiss (1988) developed a model to schedule timber harvesting and delivery activities over five years for a firm with multiple mills in Ontario. The objective was to minimise delivered wood costs while satisfying all mill demands. Interestingly, to avoid dispersal of the harvest, similar stands were aggregated into large blocks (averaging 1,000 ha) and these were treated as binary decision variables, with their harvested contents distributed to various mills. Wrightman and Jordan (1990) modeled the problem of satisfying the demands of competing mills within one timber supply area in New Brunswick. Their objective was to minimise total transportation costs while satisfying each mill's demand and controlling the differences in transportation costs for competing mills. It is noteworthy that, in New Brunswick, competing mills were not constrained to harvest in given chart areas. Burger and Jamnick (1995) designed a linear programming model to procure and distribute logs from multiple sources to multiple mills for the woodlands division of a single firm in Nova Scotia. Mill requirements, product revenues, harvest, transportation, and wood purchasing costs were all considered in the model. In the early 1990's, the first attempts to link the long-term spatial strategic plans with log allocation problem were made. Nelson and Howard (1991) developed a three-stage heuristic for allocating spatially and temporally feasible timber harvesting rights among competing firms in British Columbia. They generated a spatially explicit harvest schedule for a sample problem and then assigned chart areas to competing mills. The heuristic assignment algorithm was designed to satisfy the multiple objectives of each mill (total volume targets, seasonal volume targets, and transportation costs) as closely and equitably as possible. Colberg (1996) integrated strategic timber supply planning with a log procurement system using a set of interdependent models in a hierarchical planning framework for a firm operating in Georgia and Alabama. He sought to integrate three interdependent objectives in one planning framework: 1) to maximize returns from company-owned or controlled timberlands; 2) minimise the cost of wood procurement systems; and 3) manage wood products operations as an integral part of the firm's fibre supply system.

Model Description

The three complicating factors in the assignment of cutting rights to licensees, as described above, are: 1) the imposition of landscape-scale spatial constraints on harvesting across the boundaries of chart areas; 2) the multiple, conflicting objectives within and between licensees; and 3) the problem of evaluating the appurtenance policy, and planning for its relaxation through timber sales. I now describe how each of these, in turn, has been incorporated into the structure and model of the problem.

First, the simplest way to prevent, over the long term, the inequitable or disruptive distribution of landscape-scale spatial constraints across the boundaries of chart areas, is to define a set of landscape-units for the management unit, and let these become the chart areas. Landscape-units are delineated on the basis of readily identifiable physiographic or geographic features such as a single large watershed or a series of smaller watersheds. They may also be established to reflect dominant resource use patterns or administrative boundaries. A harvest schedule can then be produced for each landscape-unit, subject to even-flow, spatial, and other harvesting constraints. This is what I propose doing for solving the problem of allocating cutting rights among competing licensees. The problem then becomes one of assigning landscape-units, with their forecast flow of harvestable timber, to competing licensees.

The second problem is that of conflicting multiple objectives. The assignment of chart areas to licensees should address the multiple conflicting objectives of each licensee (e.g., volume, species, log size, seasonal access, and haul distance objectives). When considering multiple objectives, the problem is to find an efficient solution, instead of an optimal solution. A solution is efficient if the achieved level of any one objective cannot be improved without worsening the achieved level of any other objective (Romero 1991). Assuming each objective, Z_j, is to be maximized, the multi-objective problem can be written as follows:

[1] Maximize $Z_1(X),\ Z_2(X),\ \dots,\ Z_p(X)$

Subject to

[2] $g_i(X) = b_i \quad i = 1, 2, \dots, m$

Where: $Z_j(X), j= 1, \dots, p$, are objective functions, X is the vector of decision variables, and $g_i(X)$ are the problem constraints.

In order to generate a solution, the initial problem is generally replaced by a single objective problem. There are several techniques for this: a) the constrained technique; b) the weighting technique; and c) the goal programming technique.

In the constrained technique, only one objective is maximized subject to lower limits on the other objectives. The problem facing the analyst is that the values of the lower limits are unknown and have to be specified.

The weighting technique entails assigning a relative weight to convert the objective vector to a scalar, Z, which is the weighted sum of the separate objective functions. These weights can be varied over reasonable ranges and the problem facing the analyst is to specify their values.

The goal programming method is devised for problems where targets have been assigned for all objectives and the decision-maker is interested in minimizing the nonachievement of the corresponding goals. To illustrate this model, let each objective function be expressed in general terms as $Z_j(X)$ and the value of the goal associated with each objective j from the set of objectives J, be b_j. A variable is needed by which the deviation between $Z_j(X)$ *and* b_j, either positive, dj+, or negative, dj-, is represented. The goal programming model can therefore be expressed in the following form:

Minimize

[3] $$\sum_{j \in J} d_j^+ + d_j^-$$

Subject to

[4] $$Z_j(X) + d_j^+ - d_j^- = 0 \qquad \textit{for each } j$$

Of course, depending on the decision-maker, the objective function may or may not include both positive and negative deviations for each target. For example, if the decision-maker stipulates only that $Z_j(X)$ be greater than or equal to a given goal, b_j, and that negative deviations are of no consequence, then the associated d_j^- would not enter the objective function. In designing the objective function of a goal programming model, some thought must be given to the desired control over each objective; i.e., whether positive or negative deviations ought to be minimized, or both. The usefulness of the goal programming approach hinges upon the clarity with which the targets, b_j, can be defined and the insights provided by trade-off analysis.

In modeling the multi-objective problem of assigning chart areas to licensees and timber sales, the goal programming technique was selected for several reasons. First, given a clearly defined allowable annual cut for the management unit as a whole, the volume targets for each licensee can also be clearly defined. From this volume target, other targets can be defined; for example, given information on each licensee's marketing and production strategy, targets for species type, seasonal access, and log sizes can also be defined in terms of cubic metres *per annum*. The targets on haul distance can also be estimated with reference to the allowable annual cut (AAC), by measuring these in units of $m^3 \cdot$ km *per annum*. Another reason goal programming is suitable to this multi-objective problem is that, controlling *both* positive and negative deviations from these targets is an important element of this problem. Since every cubic metre of wood in the management unit is assigned, therefore if one licensee exceeds its volume target, then another licensee will necessarily fall short of its volume target. Hence, the ability to control and understand both positive and negative deviations from volume targets makes the goal programming technique desirable for this multi-objective problem.

The most common variants of goal programming are weighted and lexicographic (Tamiz *et al.* 1998). Weighted goal programming minimises a weighted sum of unwanted deviations from the decision-maker's set of targets across a set of objectives. All goals are considered simultaneously. Lexicographic goal programming minimises a ranked vector of unwanted deviations from a set of targets for a number of objectives where different goals are grouped into different levels of priority. Lexicographic goal programming occurs when there exists a natural or desired ordering amongst the goals. Goals in the higher priority levels are satisfied as closely as possible and only then are goals in the lower priority levels considered; i.e., a sequential minimization of priority levels with no degradation in the value of higher priority levels. For the problem of assigning chart areas to licensees, I have chosen to use a weighted goal programming approach because there exists no clear hierarchy among the multiple objectives each licensees desires satisfied by the assignment of chart areas.

In addition to an optimization scheme, solving a multi-objective problem also requires a method by which a decision-maker's preferences are articulated. Hwang and Masud (1979) classify solution techniques for multi-objective problems according to the timing of the requirement for information on preferences. That is, techniques vary depending on whether articulation of preferences occurs: 1) prior to the optimization; 2) in sequence with the optimization; or 3) after the optimization. For the problem of assigning chart areas to licensees, I suggest that the method of articulating preferences in sequence with the optimization is most suitable, for the following reasons. First, it would be too difficult to reach any consensus on a prior articulation of preferences, given that numerous stakeholders are involved as decision-makers. Second, a posterior articulation of preferences requires that a great number of efficient solutions be generated, and this set may be too large for the decision-makers to analyze effectively. Finally, an iterative interaction between the decision-makers and the computer program does not require preference information, which is often difficult for the decision-maker to articulate (Evans 1984). Instead, selective adjustment or readjustment of weights can be made at each iteration. There exists a vast literature on multi-criteria decision-making techniques which can be used, and it is beyond the scope of this paper to review them. For an updated review of interactive goal programming methods, the reader is referred to Jones and Tamiz (1995), Tamiz *et al.* (1998), and Lee and Olson (2000).

The third complicating factor in the problem of assigning chart areas to licensees is that of evaluating the effects of the appurtenance policy, and of providing practicable alternatives to its strict application. The method used in this model is that of allowing for a gradual relaxation of the appurtenance policy by introducing a decision variable representing timber sales. The timber sales variable is assigned a portion of the allowable annual cut as its sole goal, and is therefore eligible to receive chart-areas. The timber sale variable functions to *redistribute* its harvested timber to the other licensees in such a manner that the weighted sum of their deviations from their respective targets is minimised. In effect, the gradual relaxation of the appurtenance policy is accomplished by assigning a greater percentage of the management unit's allowable annual cut to timber sales. In British Columbia, for example, BC Timber Sales (TS) is a recent program allowing for the sale of standing timber to firms that do not possess wood processing facilities. The key to understanding the potential usefulness of timber sales is that it allows for the redistribution of logs from one chart area to multiple licensees, thus allowing greater flexibility in achieving the ultimate objective of assigning the right log to the right mill. One of the central objectives of this research is to provide a decision support tool by which: a) the effects of this redistribution can be evaluated, and b) the selection of chart areas most suited for timber sales can be made.

In summary, the explicit objectives of this model are to minimize the sum of deviations from targets for total volume (m^3/year), volume by species (m^3/year), haul distance (km· m3/year), seasonal volume (m^3/year), and volume by log size class (m^3/year). Specific annual targets by licensee are presented later in Table 4.1.

Formulation of Model

The notation and formulation for the mixed integer goal programming model of assigning chart areas to licensees and timber sales (TS) are presented below.

Let:

r, R = *the index and set of targets types measured in m^3; e.g., m^3 of total volume, of species, of log sizes, or of seasonal wood.*

m, M = *the index and set of licensees.*

z, Z = *the index and set of chart areas.*

t, T = *the index and set of time periods*

x_{mz} = *1 if chart area z is assigned to licensee m, 0 otherwise.*

y_z = *1 if chart area z is assigned to TS, 0 otherwise*

a_{mzt} = *1 in period t if chart area z is assigned to licensee m; 0 otherwise*

b_{zt} = *1 in period t if chart area z is assigned to TS; 0 otherwise*

D_{mz} = *haul distance between licensee m and chart area z (km)*

DV_{mzt} = *haul distance between licensee m and chart area z times volume harvest from chart area z in period t (km m^3)*

d_{mt} = *haul distance incurred by each licensee in each period ($m^3 \bullet$ km)*

d_{mt}^{+} = *positive deviation from target for haul distance incurred by each licensee in each period ($m^3 \bullet$ km)*

d_{mt}^{-} = *negative deviation from target for haul distance incurred by each licensee in each period ($m^3 \bullet$ km)*

N_v_{rmt} = *percent normalization constant for deviation from a target r of licensee m, in period, t*

N_d_{mt} = *percent normalization constant for deviation from haul distance target of licensee m, in period, t*

N_ts_t = *percent normalization constant for deviation from volume target of TS in period t*

Tar_d_{mt} = *haul distance desired by each licensee, m, in each period, t (km m^3)*

Tar_v_{rmt} = *volume desired by licensee m in period t for target r (m^3)*

$Tar\_v\_ts_t$ = *volume target for TS in period t (m^3)*

V_{rzt} = *volume of target r harvestable from chart area, z, in period, t*

v_{rmt} = *volume of target r assigned to licensee m in period t (m^3)*

v_{rmt}^{+} = *positive deviation incurred by licensee m in period t from target r (m^3)*

v_{rmt}^{-} = *negative deviation incurred by licensee m in period t from target r (m^3)*

$v\_ts_t^{\ +}$	= *positive deviation incurred by TS in period t from volume target* (m^3)
$v\_ts_t^{\ -}$	= *negative deviation incurred by TS in period t from volume target* (m^3)
v_ts_{rt}	= *volume of target r assigned to TS in period t* (m^3)
v_tran_{rmzt}	= *volume of target r transferred through TS to licensee m from chart area z in period t*
$W_{rmt}^{\ +}$	= *penalty weight for positive deviation from target r of licensee m, in period, t*
$W_{rmt}^{\ -}$	= *penalty weight for negative deviation from target r of licensee m, in period, t*
$W\_ts_t^{\ +}$	= *penalty weight for positive deviation from volume target r of TS, in period, t*
$W\_ts_t^{\ -}$	= *penalty weight for negative deviation from volume target r of TS, in period, t*
$Wd_{mt}^{\ +}$	= *penalty weight for positive deviation from haul distance target of licensee m, in period, t*
$Wd_{mt}^{\ -}$	= *penalty weight for negative deviation from haul distance target of licensee m, in period, t*

Objective function:

Minimise total weighted percent deviations from all volume and haul distance targets

$$[5] \quad \sum_{r\in R}\sum_{m\in M}\sum_{t\in T} v_{rmt}^{\ +}\ W_{rmt}^{\ +}\ N\_v_{rmt} \ + \ \sum_{r\in R}\sum_{m\in M}\sum_{t\in T} v_{rmt}^{\ -}\ W_{rmt}^{\ -}\ N\_v_{rmt} \ +$$

$$\sum_{m\in M}\sum_{t\in T} d_{mt}^{\ +}\ Wd_{mt}^{\ +}\ N\_d_{mt} \ + \ \sum_{m\in M}\sum_{t\in T} d_{mt}^{\ -}\ Wd_{mt}^{\ -}\ N\_d_{mt} \ +$$

$$\sum_{t\in T} v\_ts_t^{\ +}\ W\_ts_t^{\ +}\ N\_ts_t \ + \ \sum_{t\in T} v\_ts_t^{\ -}\ W\_ts_t^{\ -}\ N\_ts_t$$

Subject to:

Ensure that each chart area is assigned only once, to either a licensee or timber sale (TS).

[6] $$\sum_{m \in M} x_{mz} + y_z = 1 \qquad \forall z$$

Trigger assignment of chart areas across all periods

[7] $$x_{mz} = a_{mzt} \qquad \forall m,z,t$$

[8] $$y_z = b_{zt} \qquad \forall z,t$$

Sum volume of each target, r, assigned to timber sale in each period, t

[9] $$v\_ts_{rt} = \sum_{z \in Z} b_{zt} V_{rzt} \qquad \forall r,t$$

Define the volumes transferred to licensees through assignment of chart areas to timber sale, for each target type, r, in each period, t, from each chart area, z.

[10] $$\sum_{m \in M} v\_tran_{rmzt} = b_{zt} V_{rzt} \qquad \forall r,z,t$$

Define the total volume assigned to a licensee as the sum of the volume assigned directly through chart areas *and* transferred through timber sale

[11] $$v_{rmt} = \sum_{z \in Z} a_{mzt} V_{rzt} + \sum_{z \in Z} v\_tran_{rmzt} \qquad \forall r,m,t$$

Define deviations from each licensee's volume target, r, in each period, t

[12] $$v_{rmt} - v_{rmt}^{+} + v_{rmt}^{-} = Tar\_v_{rmt} \qquad \forall r,m,t$$

Define deviations from timber sales volume target, in each period, t

[13] $$v\_ts_{rt} - v\_ts_{t}^{+} + v\_ts_{t}^{-} = Tar\_v\_ts_{t} \qquad \forall t, r=1$$

Define haul distance *per* licensee *per* period (where v_tran_{rmzt} *refers only to volume transferred, i.e., target r=1)*

[14] $$d_{mt} = \sum_{z \in Z} a_{mzt} DV_{mzt} + \sum_{r=1} \sum_{z \in Z} v\_tran_{rmzt} D_{mz} \qquad \forall m,t$$

Define deviations from haul distance targets

[15] $$d_{mt} - d_{mt}^{+} + d_{mt}^{-} = Tar\_d_{mt} \qquad \forall\, m,t$$

Define binary decision variables

[16] $$x_{mz} \in \{0, 1\} \qquad \forall m,z$$

[17] $$y_z \in \{0, 1\} \qquad \forall z$$

The objective [5] is to minimise the weighted sum of all percent deviations, both positive and negative, from each target. Normalisation techniques are used to overcome incommensurability. This occurs when deviational variables, measured in different units, are summed directly. The direct summation creates a bias toward the objectives with a larger magnitude, causing misleading results. In this model, I use the percentage normalization technique, where the normalization constant (N_{rmz}) equals the one divided by the target value (Tar_{rmz}); *i.e.*, $N_{rmz} = 1/Tar_{rmz}$.. This ensures that all deviations are measured on a percentage scale.

The objective function therefore minimizes the total weighted percent deviations from all volume and haul distance targets. This is subject to the constraint [6] that each chart area be assigned to either a licensee or timber sale. Equation [7] lets the binary decision variable, x_{mz}, trigger the accounting variable a_{mzt} across all periods. Similarly, equation [8] lets the decision variable, y_z, trigger the accounting variable b_{zt} across all periods. This minimizes the number of binary decision variables needed by the model. Equation [9] defines the volume of each target type, *r*, assigned to timber sale for each period. Equation [10] defines volume of each target type, *r*, transferred from chart area, *z*, through timber sale, to licensee *m*. Equation [11] defines total volume of target type, r, both assigned and transferred to each licensee in each period. Equation [12] defines positive and negative deviations from volume targets, for each target type, r. Equation [13] defines positive and negative deviations from volume targets for timber sale. Equation [14] defines volume haul distance incurred by each licensee ($m^3 \cdot$ km). Haul distance incurred by each licensee, in each period, is measured in units, $m^3 \cdot$ km, to better reflect actual haul distance to be incurred from the allocation. Equation [15] defines the positive and negative deviations from haul distance targets. My approach to penalizing deviations from haul distance targets is, therefore, based on the fact that not all licensees have equal annual harvests. Equations [16] and [17] define the binary decision variables in this model.

Case Study Description

The model is applied to a hypothetical allocation problem in the Kootenay Lake Timber Supply Area (TSA) in the southern interior of British Columbia. This forest is comprised of 1.1 million ha and supplies approximately 700,000 m^3 per year to seven local licensees, each with one mill. Historically, the allowable AAC of this forest far exceeded the actual harvest; but, as Figure 4.2 illustrates, recent declines in the AAC have decreased this gap. This implies that there is less opportunity for licensees to "pick and choose" the stands within their assigned chart areas in order to satisfy the demands of their mills. Consequently, the efficient assignment of chart areas is more urgently needed than in the past.

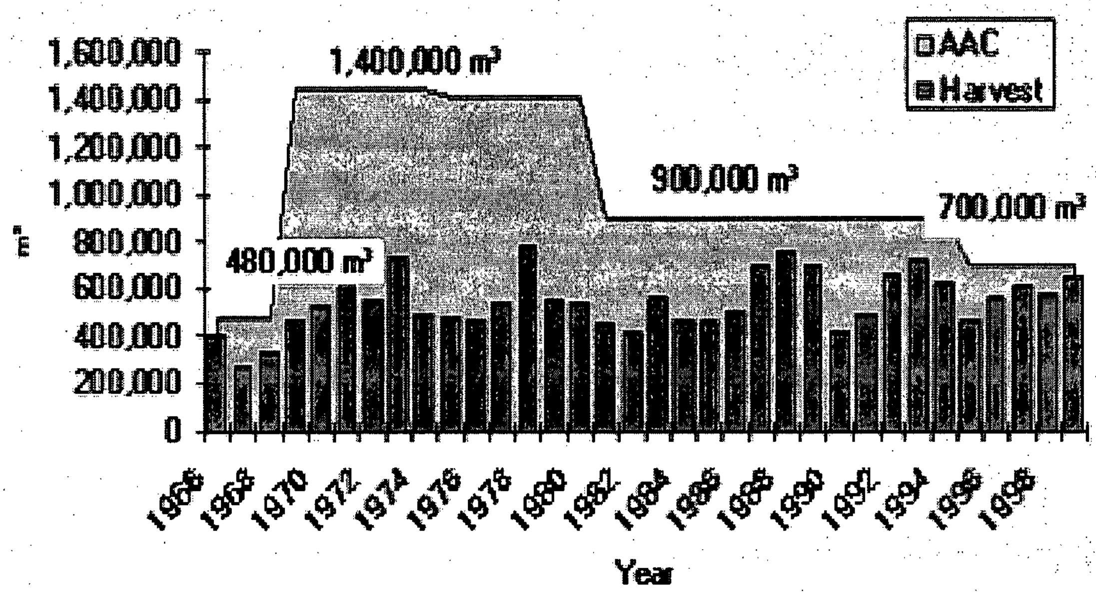

Figure 4.2: Historic allowable cuts and actual harvest levels of Kootenay Lake TSA (source: Kootenay Lake Forest District 2003)

As noted above, the first step in preparing input data for the model is to divide the forest into landscape-units. In the case of the Kootenay Lake TSA, the forest is divided into 26 landscape-units, averaging 42,000 ha (Figure 4.3).

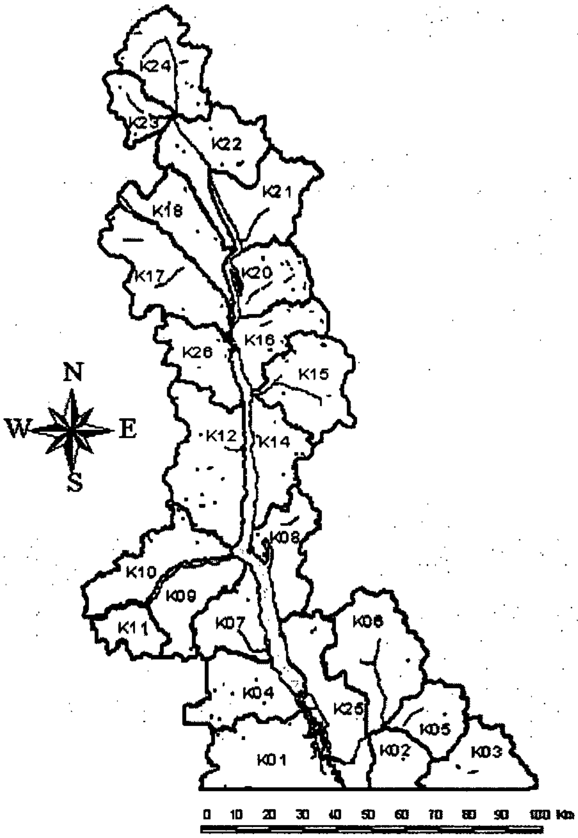

Figure 4.3: Twenty-six Landscape-units of Kootenay Lake TSA (source: Kootenay Lake Forest District 2003).

I used the landscape-units delineated by the British Columbia Ministry of Forests. Next, a harvest schedule, subject to landscape scale spatial constraints, is produced for each landscape-unit using Forest Planning Studio software (Nelson 2001). Each of the 26 schedules has a 200-year horizon, with 5-year periods. The objective is to maximize total volume harvested subject to: 1) constraints on maximum inter-period harvest fluctuations of +/- 10%; 2) preservation of special habitat conditions; and 3) achievement of targets for seral–stage distributions intended to emulate the effects of natural disturbance. These seral constraints comply with the Ministry of Forests *Biodiversity Guidebook* (B.C. Ministry of Forests 1995). For example, in each landscape-unit, ten percent of the forest area is targeted to be in old growth.

Given the 26 harvest-schedules, the problem is then one of assigning the sustainable flow of volume, estimated log-sizes classes, species, and seasonal wood supply to each licensee. Table 4.1 contains the target values used in the problem instance for period one. The assignment problem is for eight periods.

Table 4.1: Annual target values for case study in period one (All values are in m^3, except those for haul distance, which are in m^3km).

Target Type	Licensee. 1	Licensee. 2	Licensee. 3	Licensee. 4	Licensee. 5	Licensee 6	Licensee 7	Target Total	Total Available
Volume	92,645	163,125	111,297	78,190	136,626	92,231	14,179	688,293	688,293
Seasonal	23,161	40,781	27,824	19,548	34,157	23,058	3,545	172,073	49,331
Species 1*	46,322	0	27,824	19,548	0	23,058	7,089	123,841	155,783
Species 2	46,322	0	27,824	19,548	0	23,058	7,089	123,841	152,179
Species 3	0	81,563	27,824	19,548	68,313	23,058	0	220,305	191,378
Species 4	0	81,563	27,824	19,548	68,313	23,058	0	220,305	186,474
Log size class* 1	30,573	81,563	0	0	68,313	23,058	3,545	207,051	65,351
Log size class 2	30,573	81,563	55,648	0	68,313	23,058	7,089	266,244	287,179
Log size class 3	31,500	0	55,648	0	0	23,058	3,545	113,751	189,545
Log size class 4	0	0	0	39,095	0	23,058	0	62,153	99,634
Log size class 5	0	0	0	39,095	0	0	0	39,095	46,583
haul distance	3,713,848	5,578,173	7,536,729	4,072,679	5,770,963	3,436,615	796,170	30,905,177	n/a

* Species 1 = Lodgepole pine, species 2 = Douglas-fir, species 3 = Engelmann spruce; species 4 = Western red cedar.

** Log sizes are derived from the age and yield curve of each stand, where: **class 1** $\geq$ 600 m^3/ha; **class 2** $\geq$ 500 and < 600 m^3/ha; **class 3** $\geq$ 400 and < 500 m^3/ha; **class 4** $\geq$ 300 and < 400 m^3/ha; and class **5** $\geq$ 200 and < 300 m^3/ha.

Determining the target values for this problem is an important step, and care was taken to avoid targets that needlessly result in inefficient solutions. The standard goal programming formulation can produce inefficient solutions if the target values are set too pessimistically (Tamiz *et al.* 1998). In the case of volume targets, each licensee's target is a fraction of the AAC, and the sum of all volume targets equals the AAC. Setting the right-hand side values for haul distance was somewhat different; for there is not a fixed supply of haul distance on which to base this estimate, as there is in the case of volume targets. Haul distance targets were estimated by: i) calculating the *average* distance from each licensee's mill(s) to each chart area; ii) dividing this mean distance by two; and, iii) multiplying this distance value (km) by the periodic volume target (m^3). The average distance was divided two to avoid generating inefficient solutions by setting the targets too pessimistically. My approach to setting target values for log size and species was somewhat different. Here, I recognized that it is more realistic to have the demand for some species and some log sizes exceed supply and in other case, to have supply exceed demand. This is because some types of logs are in higher demand than others. This is evident in the first period demands listed in Table 4.1. Finally, the seasonal targets represent three months of operations (i.e., ¼ of the annual volume targets) reflecting availability at lower elevations.

Although schedules were produced for a 200-year planning horizon, the assignment problem is for the first eight five-year periods. Should decision-makers think that a planning horizon of forty years for the assignment of chart areas is too long, given the expected life of a mill, then weights of zero in the objective function can be assigned to any periodic objectives deemed irrelevant. By default, all weights were assigned a value of 1 in the base scenarios.

Results

The first step taken in the strategic allocation of cutting rights was to produce harvest schedules for each of the landscape-units. As Roucke and Nelson (1995) observe, the flow of wood, when scheduled from a whole management unit, can be greater than the sum of its parts, when scheduled separately. This is because not all landscape-units have equal age-class distributions. I evaluated the cost of scheduling landscape-units separately by comparing the sum of volumes from all landscapes *versus* the volume forecast from the management unit scheduled as a whole: the sum of volumes from the sum of all landscapes was 98.4% of that forecast from the management unit as a whole. This value is consistent with the findings of Roucke and Nelson (1995).

Next, I used the branch and bound algorithm of CPLEX version 7.5 to solve problem instances for mixed integer goal programming model. The number of variables in this problem was 17,864, of which 161 were binary. The solution times were between 30 s. and 5 min. using a 1.0 GHz Pentium III central processing unit. Since the interactive approach to multiple criteria decision-making necessitates speedy responses from the computer (Evans 1984), these computing times indicate that the model can lend itself to the interactive approach.

As mentioned above, the model was designed to evaluate the effects of relaxing the appurtenance policy. This was done by computing solutions using different volume targets for the timber sales based on a percent of the AAC. I compared the effects of redistributing 0, 20, 40, 60, 80, and 100 percent of the AAC through the timber sales. In the scenario where 0% of the AAC is to be redistributed through timber sale, I used a high penalty value to ensure that no deviations from the targeted volume of zero occurred. This is because I wished to ensure a proper comparison on the effectiveness of no timber sale *versus* incremental use of the timber sale. For all other scenarios, and all other targets, a penalty weight of 1 was used for each target deviation. This is because I wished to establish a base-set of scenarios, by which to evaluate the general trends resulting from increased volume in timber sales.

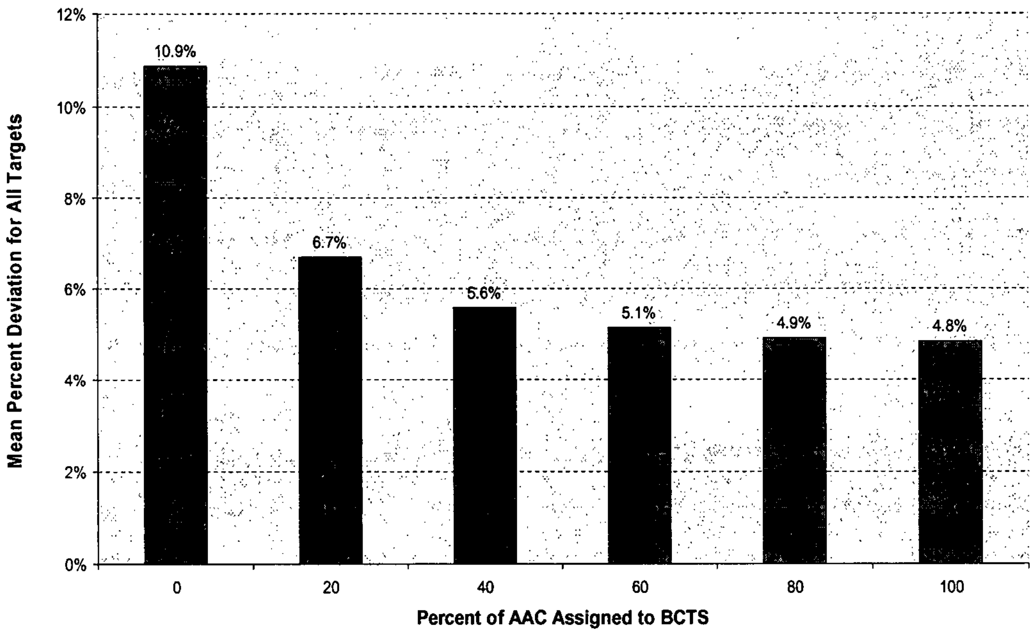

Figure 4.4: Changes in the mean percent deviation from targets relative to changes in the percent of total volume redistributed through BC timber sales.

Figure 4.4 indicates the degree to which redistribution of timber through timber sale can increase the *overall* efficiency with which timber allocation objectives can be met: a decrease in mean of all target deviations from 10.9% to 4.8%. The absolute values of these percent deviations are not so meaningful here as the general trend; for, the absolute values are very much a function of the right-hand side targets. The general trend, though, is much more informative, at least for a base-case scenario. Observing this trend, it is noteworthy that the mean satisfaction of objectives does not improve greatly after 40% of the AAC is set as the target volume for redistribution through timber sales.

Having observed the general trend, I turn now to *particular* objectives. Figure 4.5 illustrates the effect of timber sale on the achievement of volume objectives.

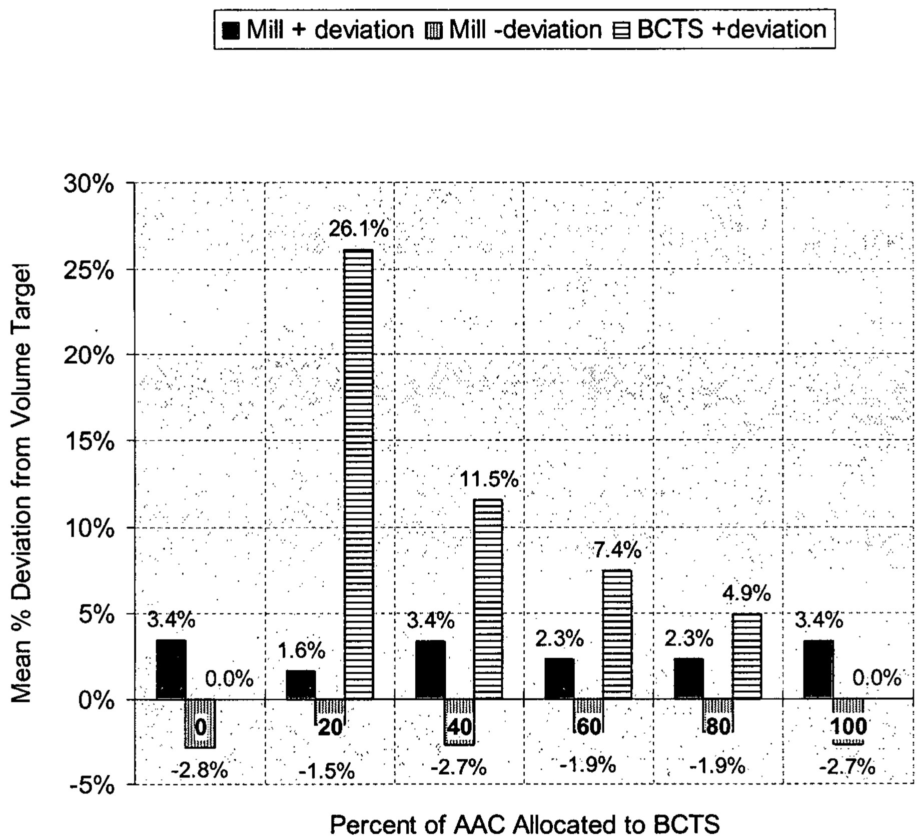

Figure 4.5: Mean percent deviations from licensee and BC timber sales volume targets relative to different proportions of the AAC assigned to timber sales.

Figure 4.5 illustrates two interesting trends with respect to deviations from volume targets. The first is that the proportion of AAC allocated to timber sale has little effect on the mean deviation from the volume targets of the 7 licensees. This is in part because the initial mean deviations are not great, i.e., when 0% of the AAC is redistributed through timber sale the mean deviations are +3.4% and -2.8% respectively. This indicates that, for the given problem instance, it is not difficult to allocate chart areas to licensees such that volume targets can be nearly met. The second interesting trend is that the positive deviations from volume targets for timber sale are relatively high when a small percentage of the AAC is assigned to timber sale. For example, when only 20% of the AAC is set as a target for timber sale, the optimal solution exceeds this target by 26.1%. This indicates that the penalty incurred for such a deviation is compensated for by the greater satisfaction of other objectives in this problem. In other words, in these scenarios, there is a strong tendency for optimal solutions to exceed the volume targets of timber sale in order to exploit the flexibility which timber sales offer through redistribution.

Turning to other targets, I observe the effects of timber sale on the satisfaction of haul distance targets in Figure 4.6: improvements are considerable, beginning with a mean deviation (per licensee per period) of +44% and ending with +27%. Once again, the absolute values of the deviations are not so important as the trend; *viz.*, that redistribution through timber sale can decrease haul distance by an average of 40% per licensee per period.

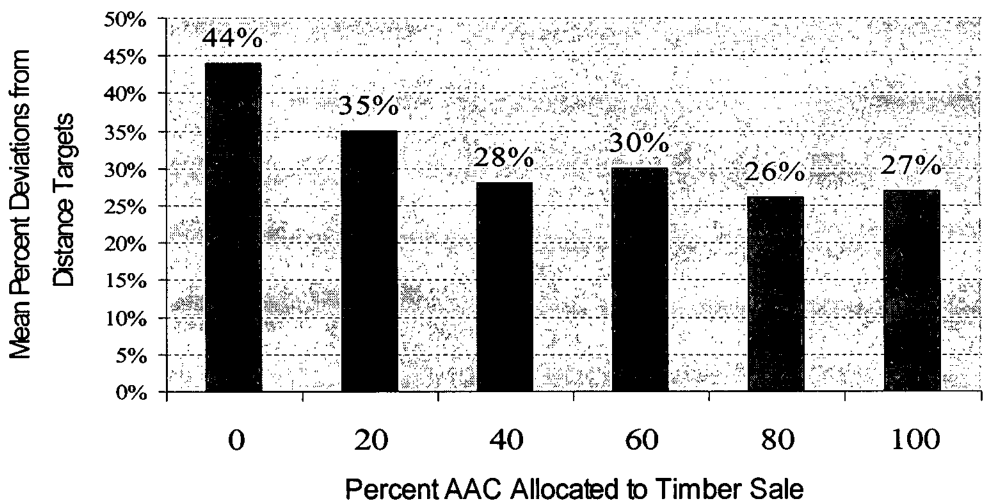

Figure 4.6: Deviations from haul distance targets relative percent of AAC allocated to timber sale.

Positive deviations from targets for haul distance incurred by each licensee, in each period, are measured by the variable d_{mt}^{+}, which is in units, $m^3 \cdot$ km to better reflect *actual* haul distance incurred from the solution. It was necessary to measure haul distance in these units because the volumes of wood redistributed to particular licensees through timber sale are measured in m^3, and are not known in advance, as the volumes from discrete chart areas are. Hence, the satisfaction of haul distance targets, in this model, is necessarily linked to the satisfaction of volume targets; and interpretation of results should be informed by this. This discrepancy between the effect of timber sale on volume (Figure 4.5) *versus* haul distance (Figure 4.6) indicates that achievement of haul distance targets, in this model, is not wholly determined by the achievement of volume targets. Another interesting trend in Figure 4.6 is that little improvement is made after 40% of the AAC is assigned to timber sale. This is consistent with the trend observed for all targets, shown in Figure 4.4.

The relation between the achievement of log size targets and the percent of volume redistributed through timber sale is illustrated in Figure 4.7.

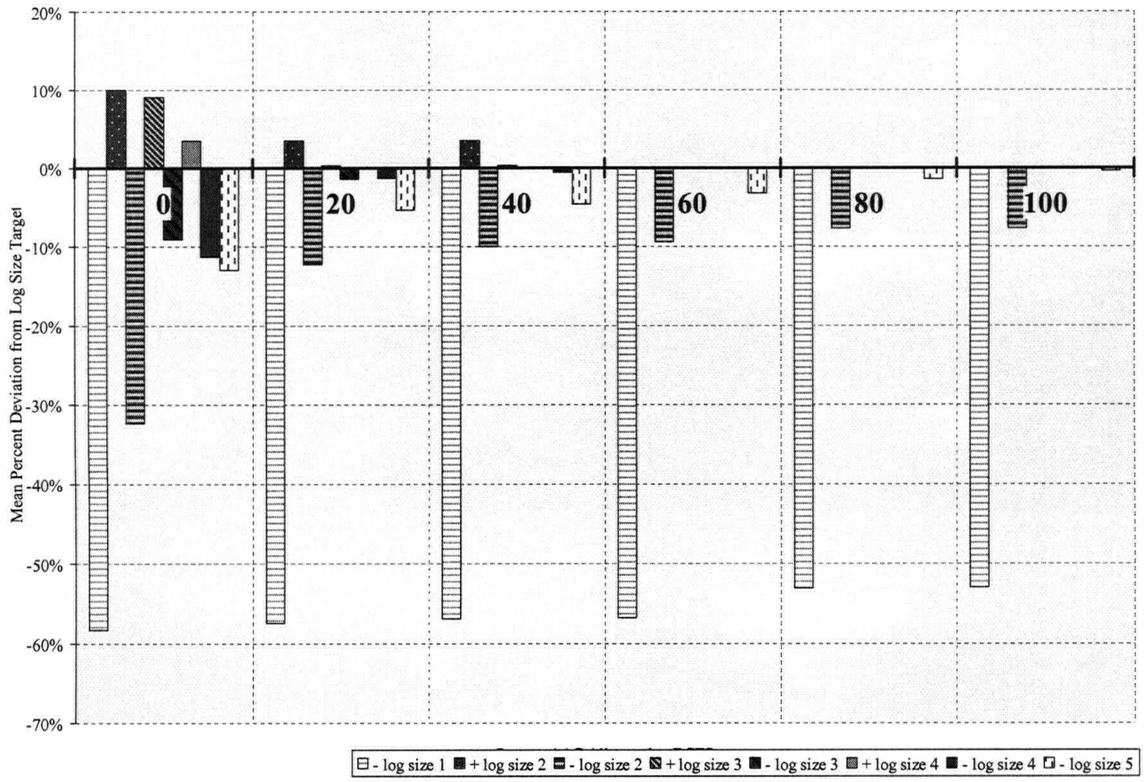

Figure 4.7: The achievement of log size targets versus percent of volume redistributed through timber sale.

The results are similar to those of other targets discussed; i.e., little improvement is made after 40% of the volume is redistributed through timber sale. There is one difference, though, in that log class 1, containing logs with the largest diameter, shows little improvement, even when 100% of the harvested volume is redistributed through timber sale. This indicates that demand, represented by the RHS-targets, greatly exceeds the supply. In this situation, the licensees with targets for log class 1 could use this model to explore various compromise scenarios. This can be done either through altering target values or through altering penalty weights assigned to deviations from this target in a given period. Working on compromise solutions is a useful part of planning, and this model lends itself to this.

Finally, Figure 4.8 illustrates the relation between the achievement of species targets and volume redistributed through timber sale. Little improvement is made after 60% of the volume is assigned to timber sale. Note that, a necessity for comprise also occurs with species 3, where optimal redistribution through timber sale yields mean negative deviations of -14.4% per licensee per period.

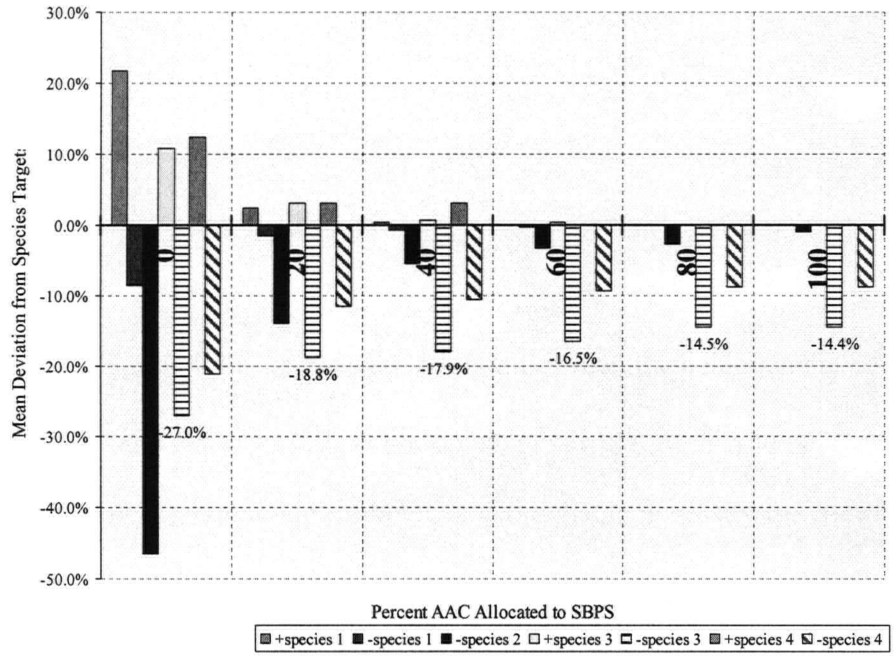

Figure 4.8: Relation between achievement of species targets and volume redistributed through timber sale.

Discussion

The results from the application of this model indicate that timber sale can play a great role in redistributing timber to better satisfy the multiple conflicting objectives of competing licensees. I now address questions on how effectively this model and its application have dealt with the three complicating factors in this problem, discussed earlier; *viz.*, 1) the landscape-scale ecological objectives, 2) the multiple, conflicting objectives of the licensees; and 3) identifying potential effects of timber sale.

The first category of discussion concerns the ability of the model to incorporate *landscape-scale ecological objectives*. On this topic, I must consider the effects of the boundaries of the landscape-units upon the solutions produced by this model. As noted above, the boundaries of the landscape-units are delineated on the basis of readily identifiable physiographic or geographic features (such as watersheds), and with no regard to the problem solved by this model. It might be contended that a different set of landscape-unit boundaries could yield a different, and perhaps more satisfactory, set of solutions to this problem; and that, insofar as this model fails to exploit the flexibility offered through the exploration of different landscape-unit boundaries, it falls short of providing the best possible solutions to this problem. I have two replies to this statement.

First, this model *can* incorporate different landscape-unit boundaries. This would simply entail execution on a different data-set; *i.e.*, a different set of landscape-units with scheduled flows of timber. Second, in designing the goal programming model, I was reluctant to incorporate into it the problem of directly redefining landscape-unit boundaries. In delineating natural landscape-unit boundaries, priority must be given to the physical features of the environment, such as topography and soil type, rather than to vegetative features, such as stand type and age. The former features are, as it were, ontologically prior to the latter (Seymour and Hunter 1999) and therefore constrain them. The allocation problem, insofar as it is concerned primarily with vegetative features, such as stand type and age, is solved in total disregard of these physical features of the environment. It might have been possible to constrain this from occurring, but I thought this might introduce an unnecessary complexity into the model. Given the comparatively simpler, yet powerful, flexibility offered by timber sale, it seemed unnecessary to incorporate into this model the many principles of natural landscape-unit design (see Forman 1997).

The second area of discussion concerns the ability of the model to address the *multiple conflicting objectives* of the licensees. One obvious point of discussion concerns the practicality of solving a problem given so many objectives. In the case study of the Kootenay Lake TSA, for example, there are objectives for volume, seasonal wood, haul distance, four species, and five log sizes for each of the seven licensees in eight periods, in addition to eight periodic volume targets for timber sale. The total number of objectives therefore equals 680, and this could be regarded as an unwieldy number for which to determine a set of penalty weights.

In reply, it should be noted that the many objectives could be approached, initially, by applying aggregates of weights. For example, the eight periods can all be assigned the same weight for a given objective, thereby reducing the number of objectives to 84. Similarly, in first exploring trade-offs between objectives, the weights for all licensees can be aggregated, thus reducing the number of weights to 12. In short, one can reduce the number of weights to a manageable number in the initial stages of using the model—to explore both the general trends in the trade-offs between the multiple objectives and the role of timber sale. Hence, the great number of objectives in this model does not exclude it from being a useful tool.

A final question to address concerns the relation between the formulation of the model and the role of timber sale. It might be thought that the significant role of timber sale is, to some extent, predestined by the model's formulation; and therefore, that its ability to reduce target-deviations is entirely predictable. In short, it might be thought that this model *labours upon the obvious*. To some extent, this is a valid objection; but it overlooks two things: first, that the model is able to quantify the effects of timber sale, and therefore, of the appurtenance policy; and second, that the model is able to *locate,* within the management unit, where implementation of timber sale could occur. These features make it useful in both policy evaluation and planning. For example, the provincial government of British Columbia, it has been argued (Pearse 2001), should re-evaluate its current approach to allocating timber, not only to increase net returns, but to re-engage in free-trade with the United States. In the National Forests of the United States, logging rights are sold by using a competitive auction. Timber sale variables in this model could therefore be used: 1) by policy analysts to evaluate the flexibility offered by using a competitive auction at the management-unit scale; and 2) by planners choosing which portions of the forest to put up for competitive auction. The introduction of timber sale variables into this model therefore greatly strengthens the analytical potential and relevance of this model.

Conclusion

The objective of this research was to develop a decision support model for the problem of integrating the allocation of cutting rights with spatially explicit timber supply planning. Key elements of this problem were: 1) that allocation and scheduling of timber satisfy landscape-scale ecological objectives; 2) that the allocation problem is one of multiple conflicting objectives within and between licensees; and 3) that the costs of the appurtenance policy, in terms of the inefficient allocations entailed by its application, be quantifiable. From this perspective, I conclude that the mixed integer goal programming model formulated and tested in this research is satisfactory. Analysis and discussion of its hypothetical application to the Kootenay Lake TSA revealed: 1) that the model can be solved relatively quickly; 2) that it can quantify general trends in the increased efficiency with which allocation objectives are satisfied, given a relaxation of the appurtenance policy; and 3) that it identify contiguous areas of the forest most suited, not only for allocation to licensees, but to timber sales. The timely relevance of this model is that it can be used to help evaluate and plan for changes to the appurtenance policy in British Columbia currently demanded by our American trading partners.

Literature Cited

Burger, D.H., and Jamnick, M.S. 1995. Using linear programming to make wood procurement and distribution decisions. Forestry Chronicle 71:89-96.

British Columbia Ministry of Forests. 1995. Biodiversity Guidebook. Government of British Columbia, Victoria.

Charnes, A., and Cooper, W.W. 1961. Management Models and Industrial Applications of Linear Programming. John Wiley and Sons, New York.

Colberg, R. E., 1996. Hierarchical Planning in the Forest Products Industry. *In* Proceedings of a Workshop on Hierarchical Approaches to Forest Management in

Public and Private Organizations. Toronto., Canada, May 25-29, 1992 / edited by David L. Martell, L.S. Davis, and Andrés Weintraub. Petawawa National Forestry Institute, 164 pp.

Evans, G.W. 1984. An overview of techniques for solving multi-objective mathematical programs. Management Science 30:1268-1282.

Forman, R.T. 1997. Land Mosaics: The Ecology of Landscapes and Regions. Cambridge Univ. Press, Cambridge, U.K. 632 pp.

Hwang, C.L. and Masud, A.S.M. 1979. Multiple objective decision-making-- methods and applications: a state of the art survey. Springer Verlag, Berlin.

Ignizio, J.P., 1976. Goal Programming and Extensions. Lexington Books, Lexington, MA.

Jones, D.F., Tamiz, M., 1995. Expanding the flexibility of goal programming via preference modelling techniques. Omega. The International Journal of Management Science 23: 41–48.

Kootenay Lake Forest District. 2003. Planning and land information. http://www.for.gov.bc.ca/dkl/planning/ (accessed Jan., 2003).

Lee, S.M., 1972. Goal Programming for Decision Analysis. Auerbach Publishers, Philadelphia.

Lee, S.M., Olson, D., 2000. Goal programming. In: Gal, T., Stewart, T.J., Hanne, T. (Eds.), Multicriteria Decision Making: Advances in MCDM Models, Algorithms, Theory, and Applications. Kluwer Academic Publishers, Boston (Chapter 8).

Nelson, J.D., 2001, ATLAS-FPS V5.0, Windows Edition, Forest Operations Group, Faculty of Forestry, University of British Columbia.

Nelson, J.D., and Howard, A.F. 1991. A three-stage heuristic procedure for allocating spatially and temporally feasible cutting rights on public lands. Canadian Journal of Forest Research 21:762-768.

Pearse, P.H., and Sydneysmith, S.1966. Method of allocating logs among several utilization processes. Forest Products Journal 16(9):89-98.

Pearse, P.H. 2001. Ready for change: crisis and opportunity in the coast forest industry. A Report to the Minister of Forests on British Columbia's Coastal Forest Industry. Vancouver, B.C.

Romero, C., 1991. Handbook of Critical Issues in Goal Programming. Pergamon Press, Oxford, U.K.

Roucke, K. and Nelson, J. 1995. Timber supply and economic impacts associated with sustained yield unit size. Forestry Chronicle 71: 647-656.

Seymour, R.S., and Hunter, M.L. 1999. Principles of ecological forestry. In pp. 22-64, M.L. Hunter (Editor) Maintaining biodiversity in forest ecosystems. Cambridge University Press. Cambridge. 697 pp.

Tamiz, M., Jones, D.F., Romero, C., 1998. Goal programming for decision-making: An overview of the current state-of-the-art. European Journal of Operational Research 111, 569–581.

Walker, H.D., and Preiss, S.W. 1988. Operational Planning using mixed integer programming. Forestry Chronicle. 64: 485-488.

Wrightman, R., and Jordan, G. 1990. Harvest distribution planning in New Brunswick. *In*

Proceedings of the Canadian Pulp and Paper Association, Woodlands Section, 71 Annual Meeting, 20-21 Mar., Montreal. Canadian Pulp and Paper Association, Woodlands Section, Montreal. Pp. E13-E17.

Chapter V

Adding Spatially Explicit Reserve Constraints

to a

Linear Programming Model

of the

Strategic Harvest Scheduling Problem

Introduction

In the management of forests over the last decade, an increased emphasis on the ecological objective to conserve biodiversity has led forest managers to adopt the *ecosystem management* paradigm. The central axiom of this paradigm is that manipulation of a forest ecosystem should work within the limits established by natural disturbance patterns (Seymour and Hunter 1999). One important consequence of this is that harvest levels are now determined through an approach that attempts to maintain, or in some cases recreate, a natural landscape age-class structure. This entails that a proportion of the forest be spared from harvesting at the financial rotation-age and age into an old growth seral stage. Even in forests where intervals of stand replacing fire-disturbance are relatively short, the emulation of natural disturbance requires that some portion of the forest progress into old growth; the quasi-random spatial pattern of such disturbances results in some stands burning repeatedly on short cycles while others escape for long periods (Van Wagner 1978).

Planning a schedule of harvests to satisfy these new age-class objectives poses little challenge to managers with linear programming models. The more serious planning challenge involves satisfying a related objective; namely, that the planned old growth reserves not be fragmented. Planning for this objective is both important and computationally difficult. It is important because harvest activities tend often to create landscapes where the old forest is fragmented (Gustafson and Crow 1994, Franklin and Forman 1987); and small, isolated patches of old growth do not provide adequate habitat for species specialized to live in their interiors, *i.e.*, unaffected by edge-effects

(Harris 1984, Temple 1985, Murcia 1995). It is computationally difficult to plan for the size and shape of old growth patches for the following reasons. First, such planning typically requires the use of optimization models because the most efficient allocation of scarce resources is a central economic objective in forest management planning; second, in such optimization models, the decision variables needed to model spatial attributes, such as size and shape, must be binary. This is because a stand is either in a reserve or it is not-- it cannot be fractionally reserved. Finally, since binary decision variables are needed to model the harvest scheduling problem with old growth reserves, traditional linear programming models do not suffice. Integer programming models are needed, and such models are notoriously difficult to solve computationally; *i.e.*, the computing time needed to generate solutions typically increases exponentially in relation to the number of decision variables (Wolsey 1998, Williams 1999). As a result, many integer programming problems in general (Reeves 1993), and the harvest-scheduling problem with spatially explicit old growth reserves in particular (e.g.,Hof and Joyce 1993, Rebain and McDill 2003), have been solved optimally on only small, impractical problem instances. Realistically large, strategic harvest-scheduling problems with spatially explicit objectives are now commonly solved using approximation, or, heuristic algorithms (Sessions and Bettinger 2001, Nelson 2003), which neither guarantee optimality nor provide any indication of how close the solution is to optimality (Reeves 1993).

The objective of this paper is to present the formulation and evaluation of a strategic harvest scheduling model which can: 1) simultaneously schedule harvests and allocate spatially explicit old growth reserves; 2) yield exact optimal solutions through application of the branch and bound algorithm; and 3) be applied to large, realistic problem instances requiring many thousands of binary decision variables.

The organization of this paper is as follows: first, the literature on related modeling research is reviewed; second, the formulation of the model is presented and clarified through a small, worked example; third, the algorithm used to enumerate the old growth reserve constraints is described; fourth, the problem instances upon which the model is tested, ranging in size from 17,000 to 800,000 ha are described, followed by a presentation of results. Finally, I discuss strengths and weaknesses of this approach, focusing on the flexibility of this model in meeting planning needs and the computational challenges associated with increased numbers of potential reserves.

Literature Review

Strategic forest management planning problems have traditionally been modeled and solved through linear programming (LP) methods (Davis *et al.* 2001, Martell *et al.* 1998), and there is a wealth of knowledge and expertise invested in many LP planning tools; e.g., FORPLAN (Johnson *et al.* 1986), Woodstock (Remsoft 2003), MELA (Siitonen 1995) and FMPP (Jonsson et al. 1993). As mentioned above, the limitation of LP methods is that the selection of spatially explicit reserves must occur outside of the LP model *per se.*

The inability of LP models to produce spatially explicit harvest schedules spurred much research in applying integer programming methods. Much of the initial research was aimed at producing solutions where the area of harvest-openings was constrained through adjacency constraints in tactical planning models. Both exact algorithms were used (Meneghin et al. 1988; Torres-Rojo and Brodie 1990;Yoshimoto and Brodie 1994;Murray and Church 1996; McDill and Braze 2000) and heuristic algorithms (O'Hara *et al.* 1989, Nelson and Brodie 1990, Clements *et al.* 1990, Lockwood and Moore 1992, Murray and Church 1993, Dahlin and Sallnas 1993, Weintraub *et al.* 1995, Bettinger et al. 1997, Ohman and Eriksson 1998, Brumelle *et al* 1998, Bettinger et al. 1999, Van Deusen 1999, Boston and Bettinger 1999 , Liu et al. 2000, Richards and Gunn 2000, Clark et al. 2000,Van Deusen 2001, Sessions and Bettinger 2001, Baskent and Jordan 2002, Boston and Bettinger 2001, Crowe and Nelson 2003, Richards and Gunn 2003, Caro *et al.* 2003). Heuristic algorithms proved capable of solving larger, more realistic problems, and their application was extended to solving multiple rotation strategic planning problems with spatially explicit objectives (e.g., Liu *et al.* 2000, Sessions and Bettinger 2001, Baskent and Jordan 2002, Crowe and Nelson 2003).

Of particular relevance to this research is that heuristic algorithms were also used to solve problems modeled to schedule harvests and to allocate reserves simultaneously. Bettinger *et al.* (1997) used tabu search to schedule timber harvests subject to spatially explicit habitat goals. These goals required that thermal or cover patches (ranging from 3 to 17 ha) be preserved and that forage areas be clustered around them. Ohman and Eriksson (1998) used a simulated annealing algorithm to solve a model designed to maximise the net present value of harvesting and to preserve patches of old growth forest. To reduce the fragmentation of the preserved old growth forest, they maximised the total "core area", *i.e.*, the area of old growth forest not influenced by edge-effects. Introducing a core area objective into the model led to the selection of old growth patches that were both large in size and round in shape. No explicit control of the spatial lay-out of the harvest-blocks was included. Ohman (2000) also applied the model to a forest of 755 stands over 20 periods. She compared the objective function from an integer programming model with that of an aspatial, linear programming model and concluded that the cost of attaining the spatial patterns appeared to be low. Van Deusen (2001) used a simulated annealing algorithm to schedule harvests and simultaneously create buffer strips of uncut forest around ponds and to promote habitat connectivity.

With advances in computing power, the problem of simultaneously scheduling harvests and allocating old growth reserves has been solved on large problems using heuristic algorithms. For example, Liu *et al.* (2000) used a simulated annealing algorithm to simultaneously harvest timber and satisfy old growth patch size targets on a forest of 80,000 ha. Their problem contained 720,000 binary decision variables. Sessions and Bettinger (2001) used a simulated annealing heuristic to achieve similar objectives on a 38,000 ha forest and their problem required 600,000 binary decision variables.

There has been much less research on simultaneously harvesting and allocating old growth patches using exact optimization algorithms, and the problem instances solved have either been very small or restricted to one period. Hof and Joyce (1992) initiated research on this problem by designing a multi-objective non-linear program to maximise three weighted objectives: old growth area, habitat for edge dependent species, and timber volume. They enumerated potential areas of protected old growth as circles where the choice variables were the size and location of these circles. Their exploratory efforts were applicable to very small problems. Hof and Joyce (1993) also designed a mixed integer programming model to maximise population of edge dependent species, old growth dependent species, and timber output. Optimal solutions were generated on a 25-cell planning area for one period. The authors note that the single time period reflects the permanence of the management options. Ohman (2001) formulated a mixed integer programming model to cluster harvest activities and reserves. The model was solved using the branch and bound algorithm and tested on a problem instance of 10,000 pixels, but it did not schedule harvest activities periodically. Rebain and McDill (2003) formulated an integer programming model of a tactical planning problem to schedule harvest units with constraints on harvest openings and constraints on harvesting a subset of old growth patches. The set of potential old growth reserves was enumerated without control over patch-shape. The formulation required many binary decision variables and was slow to solve. It was tested on a forest of 50 stands and scheduled for 3 periods, and required 178 hours of computing time, using a 1.2 GHz cpu, to generate an optimal solution. Martins et al. (2003) developed a column generation approach to solve forest planning problems with constraints on the clearcut size and on the total area of old growth patches with a minimum size requirement. Their approach was tested on a forest of 574 stands for one period.

This literature review would not be complete without mentioning research conducted on a different but related problem: the *reserve selection problem*. In this problem, there is no direct scheduling of harvests; instead the objective is to select, from a set of potential reserves containing diverse conservation elements, a subset that maximises the diversity of biological representation. There are two general approaches to modeling this problem (Camm *et al.* 1996). In the first, the objective is to choose the minimal set of reserves containing all conservation elements, (typically threatened species or habitats) at least once. Underhill (1994) notes that this is a *set covering problem.* In the second, more realistic approach, the objective is to maximise the number of conservation elements when there is a limit on the number of reserves that may be chosen. This is the *maximal covering problem* (Church and ReVelle 1974, Church *et al.* 1996). The maximal covering problem has been refined to control patch shape (Williams and ReVelle 1998) and applied to large, realistic problem instances by Fischer and Church (2003).

Summary of Literature Review and Problem Definition

The preceding literature review illustrates the following key points which help further clarify and define the problem modeled in this research. *First*, that in defining the boundaries of old-growth reserves, attention should be given to shape; i.e., that a reserve round in shape increases the relative 'core-area' of old growth habitat, and that some criterion to measure and produce a round shape should be used in modeling the problem. The *second* element is that a diversity of desired *reserve types* may exist, reflecting the diversity of forest ecosystems in a management unit. A model of this problem should therefore incorporate reserve types. *Finally*, the literature illustrates that in solving models designed to simultaneously harvest timber and select reserve patches, metaheuristic algorithms have been used with greater success than exact algorithms. The computational challenges confronting exact approaches therefore require a formulation which minimizes the complexity of the search for an optimal solution. To minimize this complexity, I suggest that the *tactical problem* of controlling opening size need not be incorporated into a *strategic model.* The work of both Rebain and McDill (2003) and Martins *et al.* (2003) indicate that scheduling harvests while a) controlling opening size, and b) selecting old growth patches, restricts the model to solving small problems. Control over opening sizes requires that each harvestable polygon, for each period, must be represented in the model by a binary decision variable. By avoiding the control over opening sizes, the model in this research is therefore able to avoid the complexity of using this type of binary decision variable.

In this research therefore, the tactical problem of controlling opening sizes will not be included the design of the strategic model. Instead, questions addressed in this model are the following: a) what is the optimal sustainable rate of harvest from a given management unit; and, b) where ought reserves be located in order to minimize the impact on timber flow? These are both critical questions in strategic planning. By excluding the tactical problem of opening size constraints, it is hoped that the model will be applicable to larger problem instances. Of course, a model designed to address such questions has already been presented by Ohman and Eriksson (1998); but what distinguishes my research here is that I use an exact approach, instead of a metaheuristic.

Methods

The model presented here is based on a Model I formulation of the strategic harvest-scheduling problem, as defined by Johnson and Scheurman (1977). The innovation in this research is that a set of spatially explicit reserve constraints is appended to this Model I.

This involves two steps: *first*, I enumerate all feasible reserves, for each reserve class, using an enumeration algorithm. A *feasible reserve* is a patch of forest that meets certain age, size, and shape criteria (described below). This model therefore allows for a multitude of reserve classes which could be based on types of forest ecosystems, or patch sizes, or other planning criteria. In this research, the classes are based on different old growth patch sizes. The *second* step involves translating this set of *feasible reserve patches* into a set of N possible constraints, of which only K of these constraints must hold. Part of the optimization process is to choose the combination of K constraints that permits the Model I objective function to reach its optimal value.

It should be noted that no explicit control of the spatial lay-out of the harvest blocks is included in this model. These are typically addressed in tactical planning models. Since this is a strategic planning model, it addresses long-term planning issues which involve not only estimating the sustainable yield of a forest, but also the spatial allocation of its reserves.

Formulation of Model

The formulation of the model is described through the following example. Given a forest of twelve polygons, each comprised of 10 ha (Figure 5.1), and set of *feasible reserves* (meeting size and shape criteria, listed in Table 5.1), the problem is twofold: 1) formulate a set of potential reserve constraints; and 2) incorporate these constraints into a harvest scheduling model.

1	2	3	4
5	6	7	8
9	10	11	12

Figure 5.1: Sample forest of twelve 10-ha polygons

Table 5.1: Set of reserves in reserve-class K, i.e., where size is $\geq$ 40 ha and $<$ 50 ha, and the shape criterion restricts feasibility to only to square openings, i.e., excluding rectangular openings.

Reserve Id	Set of polygons within feasible reserve
1	{1,2,5,6}
2	{2,3,6,7}
3	{3,4,7,8}
4	{5,6,9,10}
5	{6,7,10,11}
6	{7,8,11,12}

To formulate the reserve constraints, let:

x_{ij} = the area of polygon i harvested under timing choice j (this is the classic Model I decision variable)

y_k = 1 if potential reserve k is *not* selected to be an actual reserve, 0 if it is selected.

$y_k \in \{0,1\})$

M = a number of an arbitrarily high value (e.g., > that the total area of the forest).

j, J = the index and set of timing choices.

k, K = the index and set of feasible reserves.

I_k = the set of polygons in feasible reserve k

Translation of the reserves into constraints requires two linear inequalities. First, an inequality [1] by which a harvested or unharvested polygon triggers y_k to be 1 if potential reserve k is *not selected* to be a reserve, and 0 otherwise.

[1] $$\sum_{i \in I_k} \sum_{j \in J} x_{ij} - M y_k \leq 0 \qquad \forall\, k \in K$$

From the example illustrated in Figure 5.1 and Table 5.1, the inequalities (ignoring timing choice) would be:

[2] $$x_1 + x_2 + x_5 + x_6 - M y_1 \leq 0$$

[3] $$x_2 + x_3 + x_6 + x_7 - M y_2 \leq 0$$

[4] $$x_3 + x_4 + x_7 + x_8 - M y_3 \leq 0$$

[5] $$x_5 + x_6 + x_9 + x_{10} - M y_4 \leq 0$$

[6] $$x_6 + x_7 + x_{10} + x_{11} - M y_5 \leq 0$$

[7] $x_7 + x_8 + x_{11} + x_{12} - \; M\, y_6 \leq 0$

Second, an inequality is needed [8] by which the minimal number of reserves, in each reserve class, K, is implemented.

[8] $$\sum_{k \in K} y_k \leq P_K - D_K \qquad \forall\, K \in W$$

Where P_K = the number of potential reserves in reserve class K; D_k = the number of desired reserves in reserve class K; and W is the set of all reserve classes. As noted above, this model allows for a multitude of reserve classes, and the set, W, is comprised of different size-classes of old growth patches. Also note that, P_k must be > D_k, otherwise there would be no selection problem.

From the illustrated example, assuming two of the potential six reserves were desired, the inequality would be:

[9] $y_1 + y_2 + y_3 + y_4 + y_5 + y_6 \leq 4$

There are several observations worth noting about this formulation as it currently stands.

First, there is no aggregation of polygons into strata. In this model, each polygon is represented by a decision variable, and this is needed to trigger the spatially explicit reserves. The aggregation of polygons into strata is a traditional approach to solving LP models of harvest scheduling and is intended to decrease the number of decision variables; but the need to reduce the number of decision variables is no longer of great importance, given the advances in computing hardware and LP-solvers.

Second, the number of binary decision variables in each problem equals the number of potential reserves. This is important because the number of binary decision variables greatly influences the computing time required to reach an optimal solution in many integer programming problems (Wolsey 1998). In other words, a problem instance with many thousands of potential reserves may prove to be computationally intractable. It is, of course, difficult to know *a priori*, how computationally challenging a large problem instance is to solve; and a major objective in this research is to determine whether or not this model can be applied to large, realistic problem instances with many thousands of potential reserves.

Third, the formulation, as currently expressed, does not prevent the overlap of reserves. For example, the two reserves to be selected from the sample forest (Figure 5.1) could be 1 and 2, which both share polygons 2 and 6. Depending on the *type* of reserves selected this may or may not be desirable. On the one hand, if one were selecting a set of reserves to maximise the number of species represented by the selection of a fixed number of areas for a reserve system, then one may wish to choose the minimal set of reserves-- and thus overlapping would be desirable. In the reserve selection problem (Camm *et al.* 1996), a *set covering model* is used where overlapping of reserves results from maximizing species representation on the minimal area of land. On the other hand, if one has area targets to meet—e.g., 10% of the management unit must be reserved for old growth—then the formulation, as it currently stands, is problematic because the sum of all areas reserved may not equal the sum of all reserve-areas; *i.e.*, one could count twice those areas which overlap. To prevent this, three area-accounting constraints are needed.

For the first constraint, let r_i = a non-binary variable representing each polygon, i, within a feasible reserve. The value of r_i is triggered by y_k .

[10] $$y_k = r_i \qquad \forall \quad i \in I_k \text{ and } k \in K$$

Although r_i functions as a binary decision variable (assuming the value of 0 or 1), because it is triggered by y_k, it need not, and therefore is not, explicitly recognized in the model as such. It is recognized as a real number. This is because I do not wish to expand the size of the branch and bound tree unnecessarily.

The second constraint needed for area-accounting of reserves is to place a limit on the total harvestable area of the forest, *i.e.*, control the area not placed in a potential reserve for each reserve class, W.

[11] $$\sum_{i \in C_K} r_i A_i \leq R_K - F_K \qquad \forall K \in W$$

Where A_i = area of polygon I; C_K= *the set of polygons in potential reserve-class K;* and R_K= the total area of potentially reservable forest in reserve class K; and F_k = the desired area of forest in reserve class K to be reserved. Once again, $R_K > F_K$, otherwise there is no selection problem. It is noteworthy that this constraint will account for only the largest reserve class; *i.e.*, smaller reserves nested partially or entirely within larger reserves can still satisfy the area constraint [11]. I include this constraint within the model because of the importance of meeting objectives for the largest reserve targets. To ensure that the total old growth reserve area is correct, an additional constraint is needed.

$$[12] \qquad \sum_{i \in Q} r_i A_i \leq PR - \alpha T$$

Where Q = the set of polygons within *all* potential reserves; *PR* = the total area of forest within the set of all potential reserves; T = the total area of the forest; and α = the percent of total area of the forest desired to be within a spatially explicit reserve. Constraints [11] and [12] together will ensure: 1) that the desired area of forest reserved *within the largest* reserve size-class will be satisfied; and 2) that the *total area* of forest is placed within a spatially explicit reserve achieves a desired target.

The complete mathematical formulation of the mixed integer programming model for strategic harvest scheduling problem with spatially explicit reserves is presented below. Let:

$i,,I$ = the index and set of polygons.

j,J = the index and set of timing choices.

k,K = the index and set of feasible reserves of a given reserve-class

t,T = the index and set of planning periods.

I_k = the set of polygons in feasible reserve k

x_{ij} *= the harvested area of polygon i under timing choice j*

y_k *= 0 if reserve k is selected for implementation, 1 otherwise*

r_i *= a non-binary trigger variable for each polygon, i, used in accounting the total area placed in reserve*

h_t = *total volume harvested in period t*

a = *percent area of forest area to be placed in a reserve*

A_i = *area of polygon i*

C_K = *the set of polygons in potential reserve-class, K*

D_K = *the desired number of reserves, in reserve-class K, to be implemented.*

F_K = *desired area of forest to be reserved in reserve-class, K (ha)*

$LTSY$ = *the long term sustained yield of the forest (m^3 per period)*

M = *an arbitrarily high number; e.g., a value greater than the total number of hectares in the forest*

P_K = *the number of potential reserves in reserve-class, K.*

PR = *total area of forest within all classes of potential reserves (ha)*

Q = the set of polygons within *all* potential reserves

R_K = *total area of forest in reserve-class, K (ha)*

G = *total area of forest (ha)*

V_{ijt} = *harvest volume per hectare of polygon i under timing choice j in period t*

W = *the set of all reserve-classes*

Maximise total volume harvested over all periods

[13] $$Max. \sum_{t \in T} h_t$$

Subject to:

Accounting variable for periodic volume, h_t

$$[14] \quad h_t = \sum_{i \in I} \sum_{j \in J} x_{ij} V_{ijt} \qquad \forall\ t = 1,\dots,T$$

Limit area harvested for each polygon

$$[15] \quad \sum_{j \in J} x_{ij} \leq A_i \qquad \forall\ i = 1,\dots,I$$

Inter-period harvest fluctuations limited to +/- 20%

$$[16] \quad -.8h_t + h_{t+1} \geq 0 \qquad t = 1,\dots,|T|-1$$

$$[17] \quad -1.2h_t + h_{t+1} \leq 0 \qquad t = 1,\dots,|T|-1$$

Harvest in last period, n, not to exceed long term sustainable yield

$$[18] \quad h_n - LTSY \leq 0$$

For each reserve, let any harvested polygon within a potential reserve, k, trigger the binary variable, y_k*, to be 1, and unharvested polygons trigger* y_k *to be zero*

$$[19] \quad \sum_{j \in J} \sum_{i \in I_k} x_{ij} - M y_k \leq 0 \qquad \forall\ k \in K \text{ and } \forall\ K \in W$$

The minimal number of reserves, in each reserve class, *K*, must be implemented for all reserve classes, W

[20] $$\sum_{k \in K} y_k \leq P_K - D_K \qquad \forall\ K \in W$$

Let each reserve variable, $y_{k,}$ trigger each polygon variable, $r_{i..}$

[21] $$y_k = r_i \qquad \forall i \in I_k \text{ and } \forall k \in K$$

For each reserve class, K, ensure that the total reserve-area target is satisfied.

[22] $$\sum_{i \in C_K} r_i A_i \leq R_K - F_K \qquad \forall K \in W$$

Ensure that the sum of all polygon areas placed in a reserve satisfies the target for total area of forest to be reserved.

[23] $$\sum_{i \in Q} r_i A_i \leq PR - a\ G$$

y_k is a binary variable

[24] $$y_k \in \{0,1\}$$

Non-negativity constraints

[25] $$x_{ij} \geq 0 \qquad \forall\ i, j.$$

[26] $$r_i \geq 0 \qquad \forall\ i$$

Enumeration Algorithm for the Set of Potential Reserves

To provide a set of potential reserve constraints, all feasible reserves must first be enumerated. The reserve enumeration algorithm produces an exhaustive list of feasible reserve blocks, *i.e.*, aggregates of contiguous polygons in which: 1) the total area is within a desired range; 2) the age of each polygon is greater than or equal to the age of old growth; and 3) the shape is relatively round.

The criteria by which patch-shape can be measured are many (Forman 1997). In this research, I wished to reduce edge-effects through compactness and used a circularity ratio (Unwin 1981) as the criterion by which compactness of reserve shape is measured. This circularity ratio equals the area of the patch divided by the area of the smallest circle enclosing the patch.

The enumeration algorithm executes as follows:

1. From the set of all polygons comprising the forest, define the subset of polygons meeting the age-class criterion to be subset A.
2. From subset A, copy all polygons which also satisfy both size and shape criteria to the *feasible list*. This is the list of feasible reserves.
3. For each polygon in subset A, form all possible couplings with adjacent polygons also in subset A. The set of all possible couplings is subset B.
4. From subset B, copy all couples satisfying both shape and size criteria to the *feasible list*. From subset B, remove all couples which exceed the area limit.
5. For each couple in subset B, form all possible triplings with polygons which are both adjacent to the couple and within subset A. The set of all possible triplings is subset C.
6. From subset C, copy all triplings satisfying both area and shape criteria to the *feasible list*. From subset C, remove all triplets which exceed the area limit.
7. Continue expanding the *feasible list* of old growth patches in this manner until all possible combinations have been enumerated. Eventually no candidate blocks

are created because there are either no further neighbours of blocks to be combined, or all blocks created are over the maximum size limit.

8. Once the *feasible list* can increase no more, remove all duplicate sets from this list (e.g., the reserve {1,2,3} is a duplicate of {3,2,1}).

Application of Model

The model was tested on three different data-sets representing managed forests of different sizes in British Columbia—Stafford, Kootenay, and Arrow (see Table 5.2). I followed the British Columbia Ministry of Forests *Biodiversity Guidebook* (B.C. Ministry of Forests 1995) on setting the old growth patch-targets; *viz*, that 10% of the total area be placed in old growth reserves, distributed in patch-size between 40 and 200 ha. I sought an equal distribution of reserved area in patches of three different sizes: 40-80 ha, 81-140 ha, and 141-200 ha.

Table 5.2: Area and polygon numbers of forests used in testing the MIP model.

Forest	Area (ha)	# Polygons	Mean Polygon Area (ha)
Stafford	16,874	1,233	13.7
Kootenay	71,245	6,093	11.7
Arrow	799,211	34,054	23.5

In testing the model, two questions were of special interest to us:

1) How will the number of feasible patches increase in relation to: a) the percent of forest area in old growth; and, b) the circularity ratio by which patches are selected as feasible? and, *2) Is this model capable of solving problems with many thousand potential reserves?*

To pursue these questions, I altered the age-class distributions of these forests by controlling the percent of area in old growth stands. That is, a fixed percentage of forest area was randomly assigned an age greater than or equal to the old growth age. In all instances, the ages randomly assigned were between 225 and 500 years. The remaining polygons of the forest were randomly assigned an age between 0 and 224.

A third question which interested us in testing the model was: how would the solutions from the model compare with the solutions produced by a traditional, aspatial method of allocating old growth area, using LP? In other words, what is the *cost* of imposing spatial objectives upon the old growth reserves? To address this question, I compare the results from the MIP model with those of an LP model. The LP model is the same as the model described in equations 13-18 above, with the following constraint added:

Let the total harvested area of old growth polygons not exceed a fixed percentage of the total area of the forest.

[27] $$\sum_{i \in R} \sum_{j \in J} x_{ij} \leq PR - a\,T$$

The general parameters for the harvest scheduling problem were: a 200 year planning horizon of twenty 10-year periods. Each polygon had a rotation-age of eighty years and there were eight timing choices for the Model I decision variables. Inter-period fluctuations of volume harvested were constrained to within +/- 20%.

The models were built with the MPL Modeling System and solved with CPLEX version 8.1, on a 2.4 GHz Pentium 4 central processing unit with 1.5 gigabytes of RAM. Selection of branch and bound parameters can greatly affect the efficiency of the search (Wolsey 1998). The CPLEX software allowed us several alternatives, and those chosen, after informal experimentation, are presented in Table 5.3.

Table 5.3: MIP strategy options selected in CPLEX.

MIP Strategy Options in CPLEX	Option Selected
Node Selection: the rule for selecting the next node to process when backtracking	Best Estimate
Variable Selection: the rule for selecting the branching variable at the node which has been selected for branching	Automatic
MIP Probe: determines the amount of variable probing to be performed on a problem	Probing level 3 (maximum)
Branch Direction: decides which branch, the up branch or the down branch, should be taken first at each node	Algorithm select
MIP Emphasis: strategy used to inform automatic variable selection and branch direction (options are optimality, feasibility, best bound, or balanced).	Best bound

Results

The resulting number of feasible old growth reserve patches, produced by the enumeration algorithm, based on different old growth percentages and circle ratios, are presented in Table 5.4. These values indicate the rate at which binary decision variables increase in each problem instance, and are therefore important indicators of computational difficulty.

Table 5.4. Resulting number of eligible reserve-blocks for each forest based on different old growth percentages and circle ratios.

Forest	Circle Ratio	% Old growth 15%	20%	25%	30%	35%	40%	45%
Stafford	0.40	0	0	0	1	2	2	4
	0.35	1	1	1	6	10	15	65
	0.30	3	9	11	27	39	81	1,108
	0.25	6	22	38	99	202	396	11,256
	0.20	14	50	177	407	1,059	1,816	65,465
	0.15	23	112	452	875	1,604	3,023	111,098
	0.10	24	128	525	981	1,756	3,482	121,810
	0.00	24	128	525	981	1,756	3,534	121,925
Kootenay	0.40	0	1	1	3	4	na *	na
	0.35	0	5	18	34	53	na	na
	0.30	3	27	91	167	293	na	na
	0.25	22	133	519	1,104	2,576	na	na
	0.20	55	492	2,228	14,578	69,902	na	na
	0.15	77	1,087	7,539	66,017	295,000	na	na
	0.10	102	1,331	12,034	163,137	1,152,424	na	na
	0.00	113	1,380	12,592	191,013	1,499,425	na	na
Arrow	0.40	29	34	55	84	108	157	na
	0.35	109	148	241	405	552	869	na
	0.30	313	538	950	1,674	2,450	4,571	na
	0.25	740	1,427	2,810	5,334	8,781	19,860	na
	0.20	1,461	3,031	6,825	14,171	34,873	91,461	na
	0.15	2,238	5,038	12,728	28,994	70,115	209,935	na
	0.10	2,790	6,622	18,317	45,141	121,143	392,072	na
	0.00	3,032	7,168	20,838	65,184	192,202	818,737	na

*Note: the feasible reserves for some forests with higher old growth percentages were not enumerated because the number of reserves increased at such a high rate as to imply computational infeasibility for the given MIP formulation. These instances are represented by na

In Table 5.4, I observe an exponential increase in the number reserve-blocks as the percent of old growth in the forest was increased. The effect of relaxing the circle ratio also increased the number of feasible reserve-blocks. In addition, the mean polygon size affected the number of feasible reserve-blocks; e.g., Kootenay, with the smallest mean polygon area, has the greatest rate of increase in feasible reserves as old growth percentages increase and circle ratios are relaxed.

In Table 5.5, model statistics and solution information are presented for a subset of the problem instances described in Table 5.4. The criterion for selecting instances to solve was based on problem size; *i.e.*, that they be large enough to allow for evaluation of the computational practicability of this model. Note that an estimate of problem size could not be formed prior to running the enumeration algorithm. All instances solved had a circularity ratio of 0.2, and were solved optimally. The optimality tolerance was defined to be within 0.5% of the LP upper bound.

Table 5.5: Model statistics and solution information for instances solved.

Forest	% Old growth	Constraints	Variables	Binary Variables	Parsing Time	Solution Time	% LP* Solution
Stafford	40	15,003	11,945	1,816	35 sec.	49 min.	90.8
Stafford	45	811,546	75,671	65,465	10.4 hr	56 min.	94.5
Kootenay	30	177,215	64,216	14,578	41 min.	6 min.	97.6
Kootenay	35	934,434	119,903	69,902	13.6 hr.	58 min.	97.8
Arrow	35	248,776	314,205	34,873	5.2 hr.	68 min.	98.7
Arrow	40	719,737	372,702	91,461	20 hr.	126 min.	96.8

*this refers to the LP model described in equations 13-18 and 23.

My first observations concern the computational practicability of this model. Table 5.5 indicates that the times needed for computing optimal solutions is consistently brief, given the size of some instances. For example, the greatest solution time was 126 minutes needed to solve an instance with 91,461 binary decision variables. This indicates that the model may be integer-friendly, and therefore useful to planners with large, realistic problems. In addition to solution time, Table 5.5 also presents parsing times. This is the time needed for the MPL software to read through the model and data files and store them in the form of a matrix. In the larger problems, the parsing time was quite long; *e.g.*, for the largest problem, parsing required 20 hours. The parsing is of practical concern, but fortunately it increases linearly with the problem size; therefore even larger problems may not be out of reach.

In Table 5.5, I also observe that the objective function values of the spatially explicit models are very close to those of the linear programming model. In other words, the cost of imposing a spatial arrangement on the retained old growth was relatively small, on average yielding a 4% lower objective function value. These results are consistent with those of Ohman and Eriksson(1998).

Discussion

In this discussion, I pursue two paths of inquiry concerning the practicality of this model. The first addresses the computational limits of this model, given the rate at which reserves, and therefore binary decision variables, can increase. Second, I reveal the flexibility implicit in the formulation of this model, focusing on how this model may be used to solve more complex planning problems with spatially explicit reserve constraints.

First, the usefulness of this model may be limited by the exponential rate at which the number of feasible reserves can increase relative to: a) the relaxation of the circle ratio criteria for feasible reserves; b) the percentage of the forest area in old growth; and c) a decreasing mean polygon-area relative to the patch-size area. In Table 5.4, for example, I observe one problem instance requiring 1.5 million binary decision variables. In light of this, what qualified statement can be made on the potential usefulness of this model?

I address points a, b, and c, separately. *First*, the relaxation of the circle ratio criteria, by which reserves are deemed feasible, in practice ought to proceed to a limit—which should be greater than zero. For example, relaxing the circle ratio to a value of zero conflicts with one of the main objectives of planning for spatially explicit old growth; *viz.* to minimize area of old growth forest not affected by edge. Hence, insofar as this model may require higher circle ratios to limit the number of binary decision variables, this limitation is not entirely regrettable; for, higher circle ratios more effectively satisfy the ecological objective of preserving core areas of old growth.

Second, the percentage of forest area in an initial old growth state can greatly affect the number of feasible reserves. For example, in Table 5.4, in instances where the initial old growth area exceeds 30 to 40 percent of the total forest area, the number of feasible reserves is so great in number that some problems may be intractable. As a practical planning concern, however, this may not be relevant; for it is rare in managed forests today to have this percentage of old growth (FAO 2001).

Third, a relatively small mean polygon-area relative to the patch-area of the reserves can greatly increase the number of feasible reserves. This is true, and it can present a serious obstacle to solving problems. Instances with extremely small old growth polygons relative to the patch-area would require aggregation of smaller, similar, and contiguous polygons into fewer, larger polygons. Further, this aggregation must precede execution of the enumeration algorithm in order to reduce the number of feasible reserves and resultant binary decision variables. Of course, any time spatial data sets are aggregated to a coarser resolution (i.e., a smaller number of larger spatial units), one must ask what problems this might produce. This is referred to as the modifiable area unit problem, or MAUP (Openshaw 1981) and it produces two general effects. First, the *scale effect*, occurs when values are averaged over the process of aggregation, and variability in the dataset is lost; therefore values of statistics computed at the different resolutions change. In the case of the MIP model, given the potential necessity to aggregate similar old growth polygons, the variability lost through aggregation to a coarser scale may not be great. It would, for example, be much greater if I proposed aggregating *all* polygons to a coarser scale of resolution. The second effect, referred to as the *zoning effect*, arises from the particular method by which zones at a coarser scale of resolution are delineated; for different delineation methods can produce different values or statistics within each zone. In the case of the MIP model, the method by which new boundaries of old growth patches are delineated, could influence and limit the flexibility with which the optimization algorithm selects certain old growth stands as actual reserves in the solution. To minimize this loss of flexibility, aggregation, if necessary, should be minimized.

My second line of inquiry concerns the flexibility of this model in solving more complex and practical planning problems. For example, the model, as described above, allocates reserves statically, rather than dynamically. Dynamic allocation may be a relevant planning objective because, as the forest ages within the strategic planning horizon, some stands may *enter into* an old growth condition while other stands may decay and *pass out of* the old growth condition. Closer investigation of the model indicates that it can incorporate analysis of the dynamic nature of old growth reserves, without any changes to the formulation.

There are two ways by which dynamic selection of old growth patches can be analyzed using this model. First, there is the problem of planning for those stands which *enter into* the old growth condition at some point within the planning horizon. This is a pressing problem for many management units with a deficit of old growth. The transition from the sustained yield paradigm to the ecosystem management paradigm has left many managers with a deficit of representative old growth. The problem facing such managers is how best to recruit patches of old growth while also meeting the economic objectives of the forest. In this model, recruitment of old growth would be examined by changing, in the enumeration algorithm, the age at which a forest polygon is eligible to become part of feasible old growth patch. For example, suppose a management unit currently had only 3% of its area in feasible old growth patches and the objective were to reach 10%. By incrementally lowering the feasible age of old growth in the enumeration algorithm, and incorporating these younger patches into the MIP model, one could examine trade-offs between rate at which these targets can be met and the resultant cost in terms of timber-supply. From Table 5.4, I see that relaxation of the age-criteria can proceed until 30 to 45 percent of the forest is eligible, at which point the number of reserves can become computationally impractical. Such economic trade-off analysis is particularly well suited to this model because, unlike models using heuristic algorithms, it produces optimal rather than *approximately optimal* solutions.

Second, there is the problem of planning for stands which *pass out of* old growth condition. This problem can be addressed in the model through the use of different *classes* of reserves. In the enumeration algorithm, just as one can form different classes based on size, or on species, so too can one do this with age. As a result, one can reserve a uniform distribution of old growth age-classes to regulate the rate at which the area of old growth reserves break up over time. Execution of this can be synchronized with the recruitment method described above. Therefore, the dynamic nature of old growth patches, while not explicitly incorporated into the formulation of this model, can nevertheless be thoughtfully addressed through controlling eligibility parameters in the enumeration algorithm.

Conclusions

In this paper I have presented the formulation and evaluation of a strategic harvest scheduling model which can: 1) simultaneously schedule harvests and allocate spatially explicit old growth reserves; 2) yield exact optimal solutions through application of the branch and bound algorithm; and 3) be applied to large, realistic problem instances requiring many thousands of binary decision variables.

In testing the formulation, it was shown that the model may be integer friendly, having solved a problem instance with over 91,000 binary decision variables on a forest of 700,000 ha. Although the analysis was of static recruitment of old growth reserves, the model can facilitate analysis of dynamic recruitment of old growth. The ability of this model to solve large, realistic problem instances is relevant to the problem faced by forest managers currently facing a deficit of old growth and seeking an optimal strategic recruitment method. It was also found that the cost of imposing a spatial arrangement on the retained old growth was relatively small. The mixed integer programming model on average yielded a 4% lower objective function value than that of the relaxed LP.

Future research on this modeling approach can take several interesting directions. First, this model can be used to perform trade-off analysis on the rates at which spatially explicit old growth targets can be met versus the costs in terms of allowable annual cuts. Such analysis would be informative on forests where representative old growth is highly fragmented; and it would be quite innovative because, for the first time, this analysis would use optimal solutions on large problems. A second direction of future research is to apply this approach to a Model II formulation of the harvest-scheduling problem (Johnson and Scheurman 1977). This would be a rather challenging formulation because, in the Model II, individual hectares of forest cannot be tracked. Nevertheless, it could be a worthy direction of research because the Model II formulation allows for the incorporation of fire probability statements (Davis *et al.* 2001); and an optimization model capable of combining spatially explicit old growth recruiting strategies with distributions of fire probabilities may yield valuable insights into the effectiveness of recruiting strategies.

Literature Cited

Baskent, E.Z., and Jordan, E.A. 2002. Forest landscape management using simulated annealing. Forest Ecology and Management 165:29-45.

Bettinger, P., Sessions, J., and Boston, K. 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. Ecological Modeling 94:111–123.

Bettinger, P., Boston, K., and Sessions, J. 1999. Intensifying a heuristic forest harvest-scheduling search procedure with 2-opt decision choices. Canadian Journal of Forest Research 29: 1784-1792.

Boston, K., and Bettinger, P. 1999. An analysis of Monte Carlo integer programming, simulated annealing, and tabu search heuristics for solving spatial harvest-scheduling problems. Forest Science 45: 292-301.

Boston, K. and Bettinger, P. 2001. Development of spatially feasible forest plans: a comparison of two modeling approaches. Silva Fennica 35(4): 425–435.

British Columbia Ministry of Forests. 1995. Biodiversity Guidebook. Government of British Columbia, Victoria.

Brumelle, S., Granot, D., Halme, M., and Vertinsky, I. 1998. A tabu search algorithm for finding a good forest harvest schedule satisfying green-up constraints. European Journal of Operational Research, 106:408-424.

Camm, J.D., Polansky, S., Solow, A., and Csuti, B. 1996. A note on optimal algorithms for reserve site selection. Biological Conservation 78: 353-355.

Caro, F., Constantino, M., Martins, I., and Weintraub, A. 2003. A 2-opt tabu search procedure for the multiperiod forest harvesting problem with adjacency, greenup, old growth, and even flow constraints. Forest Science 49:738-751.

Church, R.L. and ReVelle, C. 1974. The maximal coverage problem. Papers of the Regional Science Association 32:101-18.

Church, R.L., Stoms, D.M., & Davis, F.W. 1996. Reserve selection as a maximal covering location problem. Biological Conservation 76:105-112.

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

Clements, S.E., Dallain, P.L., and Jamnick, M.S. 1990. An operational, spatially constrained harvest-scheduling model. Canadian Journal of Forest Research 20:1438-1447.

Crowe, K., and Nelson, J. 2003. An indirect search algorithm for harvest-scheduling under adjacency constraints. Forest Science 49:1-11.

Dahlin, B., and Sallnas, O. 1993. Harvest-scheduling under Adjacency Constraints – A Case Study from the Swedish Sub-alpine Region. Scandinavian Journal of Forest Research 8:281-290.

Davis, L.S., Johnson, K.N., Bettinger, P.S., and Howard, T.E. 2001. Forest management: to sustain ecological, economic, and social values. McGraw Hill. New York. 804 pp.

FAO (2001). Global Forest Resources Assessment 2000. FAO Forestry Paper 140. Rome, Food and Agriculture Organization. http://www.fao.org/forestry/fo/fra/ (accessed 02/01/04)

Fischer, D.T., and Church, R.L. 2003. Clustering and compactness in reserve site selection: an extension of biodiversity management area selection model. Forest Science 49:555-65.

Forman, R.T. 1997. Land Mosaics: The Ecology of Landscapes and Regions.

Cambridge Univ. Press, Cambridge, U.K. 632 pp.

Franklin, J.F., and Forman, R.T. 1987. Creating landscape patterns by forest cutting: ecological consequences and principles. Landscape Ecology 1: 5–18

Gustafson, E.J., and Crow, T R. 1994. Modeling the effects of forest harvesting on landscape structure and the spatial distribution of cowbird brood parasitism. Landscape Ecology 9: 237–248.

Harris, L.D. 1984. The fragmented forest: island biogeography theory and the preservation of biotic diversity. University of Chicago Press, Chicago, Ill.

Hof, J.G., and Joyce, L.A. 1992. Spatial optimization for wildlife and timber in managed forest ecosystems. Forest Science 38(3): 489–508.

Hof, J.G., and Joyce, L.A. 1993. A mixed integer linear programming approach for spatially optimizing wildlife and timber in managed forest ecosystems. Forest Science 39: 816–834.

Johnson, K.N., and Scheurman, H.L. 1977. Techniques for prescribing optimal timber harvest and investment under different objectives—discussion and synthesis. For. Sci. Monogr. No. 18. 31 p.

Johnson, K.N., Stuart, T.W., and Crim, S.A. 1986. FORPLAN version 2: an overview. USDA Forest Service, Land Management Planning System Section, Washington, D.C.

Jonsson, B., Jacobsson, J., and Kallur, H. 1993. The Forest Management Planning Package: theory and application. Stud.For.Suec. 189.

Liu, G., Nelson, J., and Wardman, C. 2000. A target-oriented approach to forest ecosystem design — changing the rules of forest planning. Ecological Modeling 127:269–281.

Lockwood, C., and Moore, T. 1992. Harvest-scheduling with spatial constraints: a simulated annealing approach. Canadian Journal of Forest Research 23:468-478.

Martell, D.L., Gunn, E.A., and Weintraub, A. 1998. Forest management challenges for operational researchers. European Journal of Operational Research 104:1-17.

Martins, I., Constantino, M., and Borges, J.G. 2003. A column generation approach for solving a non-temporal forest harvest model with spatial structure constraints. European Journal of Operational Research. In press.

McDill, M.E., and Braze, J. 2000. Comparing adjacency constraint formulations for randomly generated forest planning problems with four age-class distributions, Forest Science 46: 423-436.

Meneghin, B.J., Kirby, M.W., and Jones, J.G. 1988. An algorithm for writing adjacency constraints efficiently in linear programming models. In The 1988 Symposium on Systems Analysis in Forest Resources, Asilomar, CA, March 29 – April 1, 1988. USDA For. Serv. Gen. Tech. Rep. RM-161. pp. 46-53.

Murcia, C. 1995. Edge-effects in fragmented forests: implications for conservation. Trends in Ecology and Evolution 10:58-62.

Murray, A.T., and Church, R.L., 1993. Heuristic solution approaches to operational forest planning problems. Published in 1995 in OR Spektrum 17:193-203.

Murray, A.T., and Church, R.L. 1996. Analyzing cliques for imposing adjacency restrictions in forest models. Forest Science 42:715-724.

Murray, A.T., and Weintraub, A. 2002. Scale and unit specification influences in harvest scheduling with maximum area restrictions. Forest Science 48:779-789.

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

Nelson, J. 2003. Forest-level models and challenges for their successful application. Canadian Journal of Forest Research 33:422-429.

O' Hara, A.J., Faaland B.H., and Bare, B.B. 1989. Spatially constrained timber harvest scheduling. Canadian Journal of Forest Research 19:715 – 724.

Ohman, K., and Eriksson, L.O. 1998. The core area concept in forming contiguous areas for long-term forest planning. Canadian Journal of Forest Research 28:1817–1823.

Ohman, K. 2000. Creating continuous areas of old forest in long-term forest planning. Canadian Journal of Forest Research 30: 1032–1039.

Ohman, K. 2001. Forest planning with consideration to spatial relationships. Doctoral Thesis. Swedish University of Agricultural Sciences, Umea. 132 pp.

Openshaw, S., and Taylor, P. 1981. The modifiable area unit problem. In Quantitative Geography: A British View (Editor: N. Wrigley). Routledge and Regan Paul, London.

Rebain, S. and McDill, M. 2003. A mixed integer formulation of the minimum patch size problem. Forest Science 49:608-618.

Reeves, C.R. (Editor). 1993. Modern heuristic techniques for combinatorial problems. Blackwell Scientific Publications, Oxford, U.K.

Remsoft. 2003. Intelligent software for the environment. www.remsoft.com (accessed 05/06/03).

Richards, E.W. and Gunn, E.A. 2000. A model and tabu search method to optimize polygon harvest and road construction schedules. Forest Science 46:188-203.

Richards, E., and Gunn, E.A. 2003. Tabu search design for difficult forest management optimization problems. Canadian Journal of Forest Research 33: 1126–1133

Sessions, J., and Bettinger, P. 2001. Hierarchical planning: pathway to the future? *In* Proceedings of the First International Precision Forestry Symposium, 17–20 June 2001, Seattle, Wash. College of Forest Resources, University of Washington, Seattle, Wash. pp. 185–190.

Seymour , R.S., and Hunter, M.L. 1999. Principles of ecological forestry. In pp. 22-64, M.L. Hunter (Editor) Maintaining biodiversity in forest ecosystems. Cambridge University Press. Cambridge. 697 pp.

Siitonen, M. 1995. The MELA system as a forestry modeling framework. Lesnictvi For. 41: 173–178.

Temple, S.A. 1985. Predicting impacts of habitat fragmentation on forest birds: a comparison of two models. In Wildlife 2000: Modeling habitat relationships of terrestrial vertebrates. Edited by J. Verner, M.L. Morrisson, and C.J. Ralph. University of Wisconsin Press, Madison. pp. 301–304.

Torres-Rojo, J.M., and Brodie, J.D. 1990. Adjacency constraints in harvest scheduling: an aggregation heuristic. Canadian Journal of Forest Research 20: 978-986.

Underhill, L.G. 1994. Optimal and suboptimal reserve selection algorithms. Biological Conservation 70:85-7.

Unwin, D.J. 1981. Introductory Spatial Analysis. Methuen, London.

Williams, J. and ReVelle, C. 1998. Reserve assemblage of critical areas: a zero-one programming approach. European Journal of Operational Research:104:497-509.

Van Deusen, P.C. 1999. Multiple solution harvest scheduling. Silva Fennica 33(3):207–216.

Van Deusen, P.C. 2001. Scheduling spatial arrangement and harvest simultaneously. Silva Fennica 35(1): 85–92.

Van Wagner, C.E. 1978. Age-class distribution and the forest fire cycle. Canadian Journal of Forest Research 8: 220-227.

Weintraub, A., Jones, G., Meacham, M., Magendzo, A., and Malchuk, D. 1995. Heuristic procedures for solving mixed-integer harvest-scheduling-transportation planning models. Canadian Journal of Forest Research 25: 1618-1626.

Williams, H.P. 1999. Model building in mathematical programming: fourth edition. John Wiley and Sons Ltd., New York. 354 pp.

Wolsey, L.A. 1998. Integer programming. John Wiley and Sons Ltd., New York. 264pp.

Yoshimoto, A., and Brodie, J.D. 1994. Comparative analysis of algorithms to generate adjacency constraints. Canadian Journal of Forest Research 24: 1277-1288.

Chapter VI

Summary and Conclusion

In this dissertation, I have presented advances in decision support models for three important forest management planning problems with spatial constraints. In Chapter II, using the branch and bound algorithm, I explored the limitations of my formulations for the area-restricted model and found that: a) the number of decision variables, and b) the number of opening constraints, ultimately restrict this method from applicability to larger problems instances.

These results directed me to the objective of Chapter III: a worst-case analysis of the strengths and weaknesses of the metaheuristic approach to solving the area-restricted model. I concluded that problem size does not appear acutely to affect the ability of the simulated annealing algorithm to find near-optimal solutions. A weak trend was observed between larger problem instances and poorer solution qualities; but the decline in quality was not steep. Another interesting result, of relevance to understanding the relation between neighbourhood search and the peculiar structure of the area-restricted model, was that the ratio of the mean polygon area to the maximum opening area did not influence the quality of the best solution found by the metaheuristic. Also, as the percent of old-growth area in the forest increased, it was found that relative solution qualities improved for the simulated annealing algorithm, but that the search was more time-consuming for the branch and bound algorithm.

In Chapter IV, my formulation of a mixed integer goal programming model was described and applied to the Kootenay Lake Timber Supply Area. My conclusions were: a) the model is applicable to multiple criteria decision-making; b) that, given a relaxation of the appurtenance policy, it can be used to quantify the increased efficiency with which allocation objectives can be met; and c) it can be used to identify contiguous areas of the forest best suited not only for allocation to mills, but also for timber sales. One of the merits of this model is its relevance to the direction of tenure policy re-evaluation in British Columbia.

In Chapter V, I presented and tested my formulation for a mixed integer programming model to schedule harvests and select old growth reserves. My conclusions were: a) that the formulation appears to be integer friendly, having solved problem instances with 91,000 binary decision variables; and b) that a relatively small mean polygon-area relative to the patch-area of the reserves can greatly increase the number of feasible reserves, and therefore present an obstacle to solving certain problems (This obstacle, I also reasoned, may be surmountable through a careful pre-aggregation of smaller old-growth polygons into larger ones); and c) that the *recruitment* of old growth patches can be achieved indirectly through controlling eligibility parameters in the algorithm used to enumerate patches.

Having summarized my work, I conclude with a general observation on my treatment of the topic of this dissertation, *Incorporating Spatially Explicit Objectives into Forest Management Planning*. For reasons stated in the Introduction, I have chosen to incorporate spatially explicit objectives into forest management planning through integer program models; and in each of the preceding chapters, I have demonstrated a rigourous interest in both optimization and large problems. This, of course, arises from the applied nature of my discipline: forest management is concerned with selecting the most efficient allocation of scarce resources, and planning models ought to be solvable for real-world problem instances. But I would not like this concern with optimization and problem size to overshadow another concern implicit throughout my work. For, optimization and the ability to solve large problems are, as it were, constraints within which this other concern is feasible; *viz.*, the ability to strengthen the relationship between the model and the reality of the problem modeled through an added layer of complexity. In each of the three problems treated in this paper, there was a concern to incorporate this added layer of complexity, which was spatial in nature. For example:

- ❖ In the area-restricted model, one can not only find the optimal sequence of non-adjacent cut-blocks to harvest, *but also* delineate the optimal boundaries of those cut-blocks.
- ❖ In modeling the allocation of cutting rights, one can not only assign chart areas to mills, *but also* to assign chart areas to timber sales based on their ability to redistribute logs to those mills.
- ❖ In the strategic harvest-scheduling problem, one can not only select an optimal sequence of stand-type areas to harvest or set aside as old growth reserves, *but also* delineate the optimal boundaries of those reserves.

In this dissertation, therefore, I have also shown that integer programming is a powerful paradigm by which to incorporate the *complexities* of spatially explicit objectives within the pragmatic constraints of forest management planning.