Spatial Forest Planning: Where Did All The Wood Go?

Karl R. Walters
Harvest Scheduling Forester, FTG

&

Andrew Cogswell,
Senior Analyst, Remsoft Inc.

25 November 2002

Abstract

Spatial forest planning has become a hot topic in recent years. Numerous papers in the literature have been published exploring various aspects of the problem, commonly citing significant reductions in achievable harvest volume or present net value due to the imposition of spatial constraints. By and large, the problems associated with spatial planning tend to be driven by economic, social and political requirements. In this paper we examine three different spatial issues and their impacts on management objectives

The basic assumption of all stratum-based harvest schedules is no minimum/maximum block size. Unfortunately, this assumption can severely overestimate the operable land base in regions where forests are heterogeneous and stand size is small relative to economic block size. Strategic models that do not consider spatial operability guidelines tend to severely overestimate harvest volume or present net value, resulting in significant relative shortfalls in the tactical plan (e.g., 23.9% reduction in achievable harvest volume). By applying the spatial operability lock feature of Spatial Woodstock, the strategic harvest volume was reduced by 7.6%. However, the blocked harvest schedule yielded much better results with only an 8.2% shortfall relative to the new strategic volume target.

Harvest block configurations can be limiting under adjacency and green-up restrictions. In areas where there is little flexibility in locating operable harvest blocks, the configuration of blocks can yield significant differences in achievable harvest volume (50.2% shortfall versus 46.2% shortfall under identical conditions).

One of the most onerous aspects of spatial planning is accommodation of green-up intervals, where the harvest of adjacent and proximal harvest blocks must be delayed by a minimum number of years until the current harvest block reaches a desired stand condition. A difference of one year can make a dramatic difference in shortfalls (e.g., 29.6% shortfall with a 4-period green-up interval versus 11.8% shortfall with a 3-period green-up interval). Stand establishment methods that shorten the time to desired stand-condition probably have the largest pay-off.

Keywords: strategic planning, tactical planning, spatial restrictions, constraints, harvest scheduling.

Table of Contents

1 Introduction

Spatial forest planning has become a hot topic in recent years. Numerous papers in the literature have been published exploring various aspects of the problem, including characterizations of the problem structure as well as optimization and heuristic methods of solving these problems thorny planning problems. A central refrain from all these papers is the shortfall in harvest volume/revenue due to spatial constraints:

> *"Present net value of the forest was reduced by 8% relative to the spatially unconstrained harvest schedule".*
>
> *"Volume reductions of up to 11% were observed when spatial restrictions were in force".*
>
> *"Our heuristic came within 5% of the estimated global optimum in our case study."*

A forest planner reading these papers often has to wonder, if these spatial constraints are so costly, how did we get ourselves into this mess? Why is there such variability in the results, and is there really anything I can do to close this "gap" or "shortfall" in expected harvest volumes or revenues? Are there some rules of thumb that might shed some light on the kind of results one can expect for my forest-planning problem?

Over the years, we have had discussions with forest planners and found a few common complaints.

1. Spatially allocating a stratum-based harvest schedule to feasible harvest blocks can be tedious and slow.
2. Everyone assumes that stands are harvested, but in reality we harvest blocks. Stand boundaries usually make poor analogs of final harvest blocks even in regions where stand size approximates economical block size.
3. Scheduling manually delineated harvest blocks often yield volumes/revenues far short of what the non-spatial harvest scheduling model predicted.

By and large, the problems associated with spatial planning tend to be driven by economic, social and political requirements. Harvesting occurs in discrete areas and must be of sufficient magnitude to be economical. In the Pacific Northwest, social demands to minimize clearcut harvesting have resulted in legal restrictions on the size, location and position of harvest areas. And throughout the country, forest industry adopted the Sustainable Forestry Initiative^{SM} that limits the size and scope of harvesting, in a fashion similar to those in force in the Pacific Northwest.

For good or ill, spatial planning has become the norm in virtually all jurisdictions of North America, even though the actual costs of operating under spatial constraints is only just now beginning to be fully appreciated. We have found that adjacency and green-up constraints have become something of a whipping boy, being blamed for shortfalls that really are due to other factors. As we will show, even in the absence of a adjacency/green-up constraints, a fragmented forest may still block poorly and yield lower harvests than predicted in a strategic plan. Some manual intervention by planners helps to improve block layout, but the improvement comes at a cost because it does not jibe with the strategic direction. Moreover, it can be difficult to assess the impacts of manual changes (e.g., "How much does the PNV change by shifting this harvest to period two?").

So really, what IS a spatial shortfall? To call the results of a spatially-explicit tactical model a shortfall is probably a misnomer, since it implies that the problem exists with the tactical side: "If we could plan (tactically) better, we could come closer to the strategic ideal." Research is indicating ways to improve tactical planning and we fully support continuing to seek better methodologies. But we also feel that more can be done on the strategic planning side, along the lines of, "If we could incorporate more of the real spatial issues into the strategic plan, we wouldn't have such unrealistic expectations at the tactical end." One way to achieve this goal is to model spatial relationships directly in the planning model, as is done in HabPlan (NCASI, 2001). In HabPlan, the harvest scheduling algorithm must consider opening size and adjacency of individual polygons explicitly, thereby guaranteeing that the resulting output flows are spatially feasible. In order to achieve this spatial resolution, one must forego optimization in favor of heuristics, and in the process give up a fundamental strength of linear programming: the ability to directly constrain many output flows. General-purpose heuristic algorithms tend to under-perform in highly constrained planning problems and fine-tuning them can greatly improve results for one problem instance, but there is no guarantee that these modifications will work well elsewhere.

In this paper, we examine three different spatial issues, how they can result in significant shortfalls, and what (if anything) can be done to mitigate them. To illustrate these issues, we will apply forest planning tools that have been developed by Remsoft: Spatial Woodstock for strategic planning, and Stanley for tactical planning.

2 Operable Land Base

Long before anyone uttered the words “adjacency” or “green-up”, there was one fundamental spatial constraint employed almost everywhere: the minimum harvest block size. It was never really considered a constraint in harvest scheduling problems - it was simply assumed to be an engineering problem to design a harvest block that minimizes harvest cost and maximizes volume objectives. Determining the allowable cut was the domain of forestry divisions; preparing timber sales or harvest blocks was the domain of procurement or engineering divisions. Once the cut was established, foresters had little input as to how the cut was to be achieved.

The basic assumption of all stratum-based harvest schedules is no minimum/maximum block size. Unfortunately, this assumption can severely overestimate the operable land base in regions such as Atlantic Canada, where stands tend to be small and age classes are heterogeneously dispersed over the landscape. To illustrate this problem, consider a single-species forest composed of 6400 5-ha cells. Each cell is randomly assigned an age ranging from 1 to 40 years, with equal area in each age class (800 ha). Cells with ages greater than or equal to 19 are eligible for harvest.

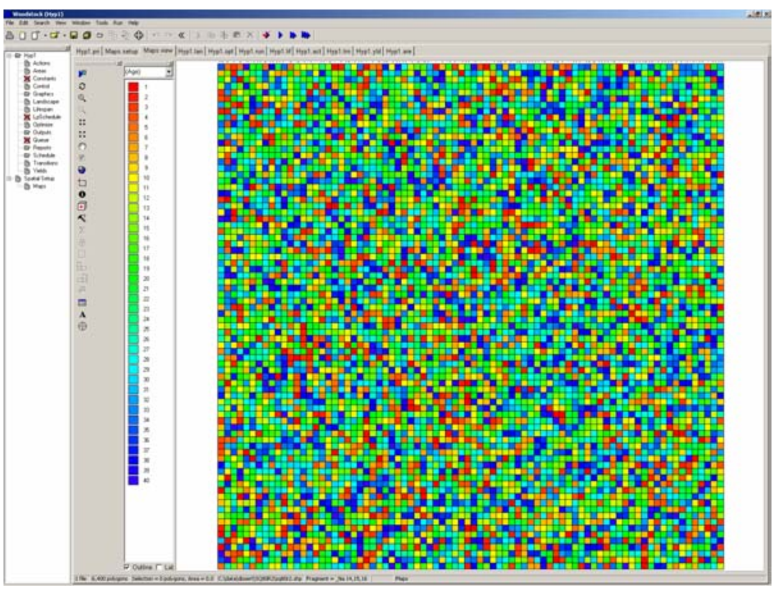

Figure 1. A fully-regulated single-species forest exhibiting scattered age classes.

If we make no assumptions about minimum block size, the maximum allowable cut reported from the stratum-based harvest schedule (Woodstock) is 769,430 m^3/yr. Running this through Stanley, an allocation and scheduling tool, we are able to block out 76.1% of the volume in feasible blocks at least 20 ha in size (4 contiguous cells). Note that this result is in the absence of adjacency or green-up requirements – we are simply trying to harvest using blocks at least 20 ha in size (4 cells).

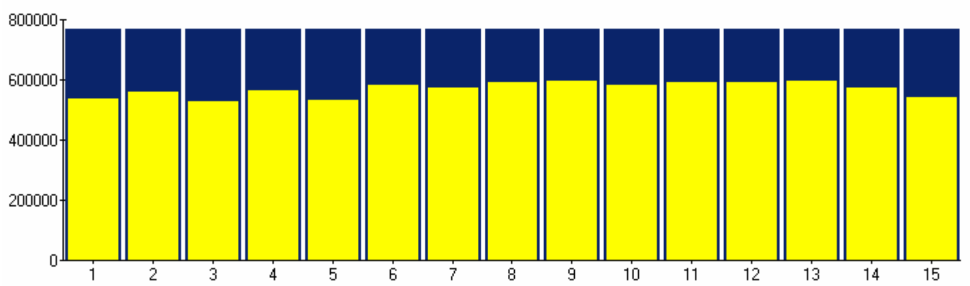

Figure 2. The spatial harvest schedule developed using Stanley, assuming no adjacency or green-up requirements. Over 15 years, we are able to harvest a total of 8,783,043 m^3 (76.1% of the strategic volume estimate for the same interval). Navy bars indicate predicted volumes from Woodstock; yellow bars indicate volumes achieved by Stanley.

Why do we achieve such poor results? Consider what makes a feasible harvest block: at least 20 contiguous ha of area (no corner points), all eligible for harvest in the same time period (age $\geq$ 19). Now, what is the spatial distribution of such potential blocks?

Figure 3. Harvest operability based on minimum block-size requirements. Numerals in cells are stand age; numerals in parentheses are years to first eligibility for harvest.

In Figure 3, consider the middle cell. It is currently eligible for harvest, but none of its neighbors with whom it shares a common boundary are eligible for harvest. However, in the next planning period, the cell below it becomes eligible for harvest, and a feasible harvest block could be configured using the middle cell and the three bottom row cells (yellow cells). Therefore, these four cells should be delayed from harvesting by 1 planning period, even though most of them are biologically mature, so that a feasible harvest block exists. Similarly, by delaying harvest of the magenta cells by 5 planning periods, a feasible harvest block could be configured by combining the magenta and the yellow cells. Finally, when 10 periods of delay have elapsed, another possible block configuration is created combining the yellow, magenta and green cells.

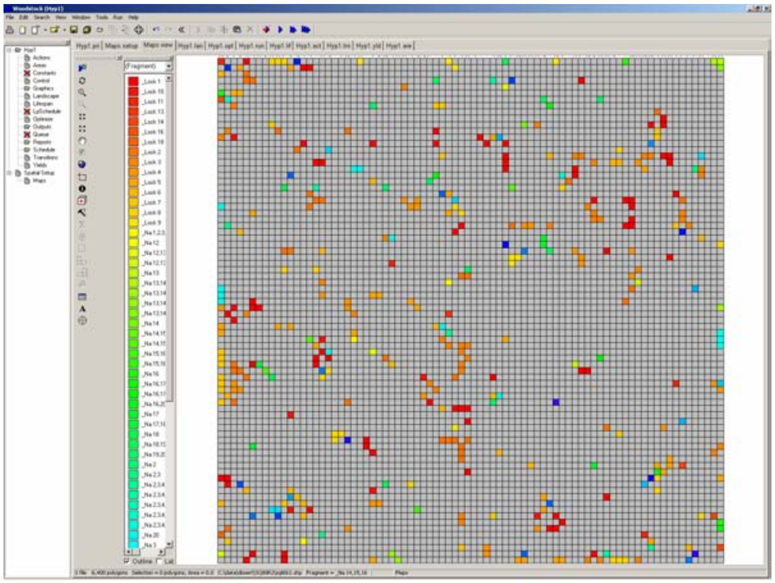

Figure 4. Results of a fragmentation analysis in Spatial Woodstock. Colored cells are those that have had spatial operability locks applied.

Spatial Woodstock can do this *fragmentation analysis* across the landscape and the resulting spatial operability locks can be incorporated into the Woodstock areas file. Any development type class that has been “locked out” will not be eligible for harvest until the lock elapses, even if the stand is otherwise eligible for harvest. The LP then must seek out alternative timing choices to meet the constraints and objective. Since we are clearly constraining the problem, we have to expect that the objective function will decline. As expected, with these spatial operability locks in place, the maximum allowable cut reported from Woodstock is reduced to 711,303/yr, a 7.8% reduction. Since we have reduced the strategic harvest level, we expect that we will be able to block out a higher proportion of the harvest. Not only do we block out a higher proportion of the strategic harvest volume (91.8%), but also the total volume achieved is even higher than what was blocked out by Stanley previously (9,222,884 m^3 versus 8,783,043 m^3 over the 15-yr period).

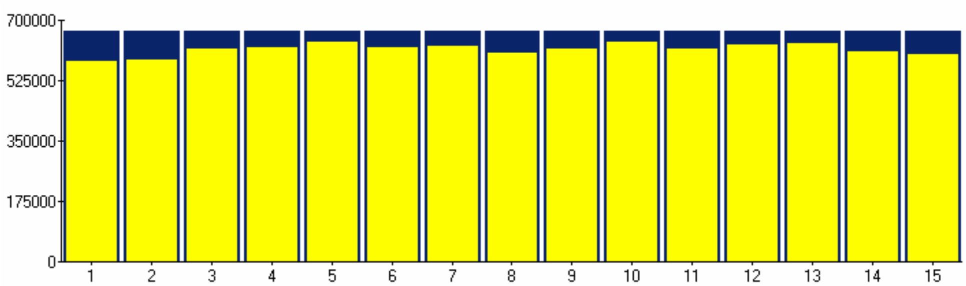

Figure 5. The spatial harvest schedule developed using Stanley, assuming no adjacency or green-up requirements. Over 15 years, we are able to harvest a total of 9,222,884 m^3 (91.8% of the strategic volume estimate for the same interval).

It is reasonable to assume that a significant proportion of the reduction in the strategic volume estimates is due to the minimum block size requirement in the scattered forest because relatively few of them had neighbors of similar age, and therefore few stands would be harvested at an age where MAI is maximized. The spatial operability locks applied after the fragmentation analysis do not guarantee that groups of operable polygons are scheduled simultaneously, but they do help to coordinate the harvest of potential blocks in periods where harvest options are at least present, if not numerous. Research continues into how best to implement these spatial operability locks (i.e., make them a function of MAI or PNV).

3 Block Configuration

Walters & Feunekes (1994) conducted a study comparing a manually developed spatial plan to one developed using computer-based heuristics. Since the heuristics were able to consider far more alternative configurations, they produced superior harvest schedules. Moreover, block configurations themselves were found to have an impact on achieving harvest levels: some block arrangements incur greater adjacency violations than others.

A couple of problems that arise with laying out harvest blocks *a priori* (preblocking) are:

1. A harvest schedule isn't used to help guide the location of blocks. In many cases, it can be difficult to map out candidate stands by treatment and planning period.
2. The block configuration itself may result in higher spatial conflicts than a different configuration.

In the following example, we use the same contrived forest planning problem we used earlier, except that the spatial arrangement of stands is much more ordered. The coloring scheme represents stand age as a gradient from red to dark blue where red is the youngest age class and blue is the oldest. Hence, harvesting will be concentrated in the first 15 periods in the upper right corner where stands are largely dark blue.

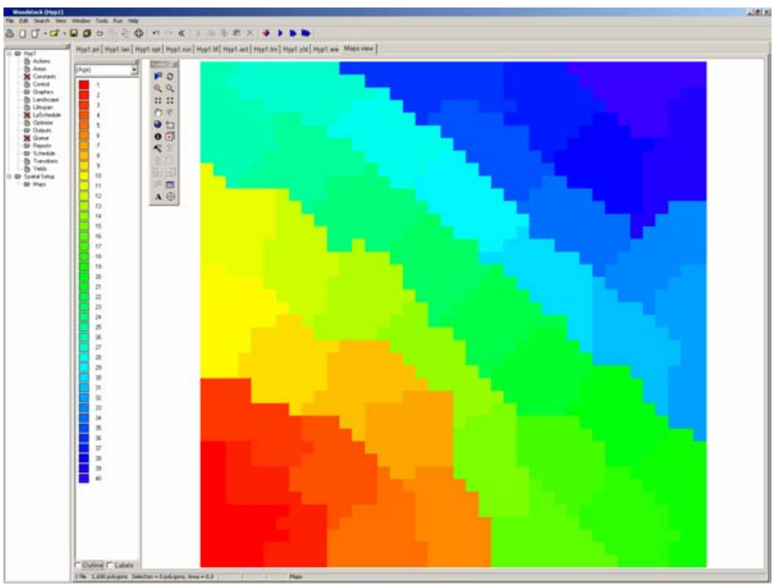

Figure 6. Stand structure of a fully regulated 40-year old forest used in the clustered forest planning problem.

The harvest schedule was formulated with a maximize volume objective, subject to even flow volume constraints and a green-up interval of 3 periods. For the unblocked test, the forest stands were subdivided into square grids of 20 ha. For the blocked test, stands older than 20 years were subdivided into 200 ha blocks. (See Figure 7.) Stanley was used to allocate and schedule the harvest for the first 15 planning periods. In both runs, only blocks 200 ha in size or greater were acceptable.

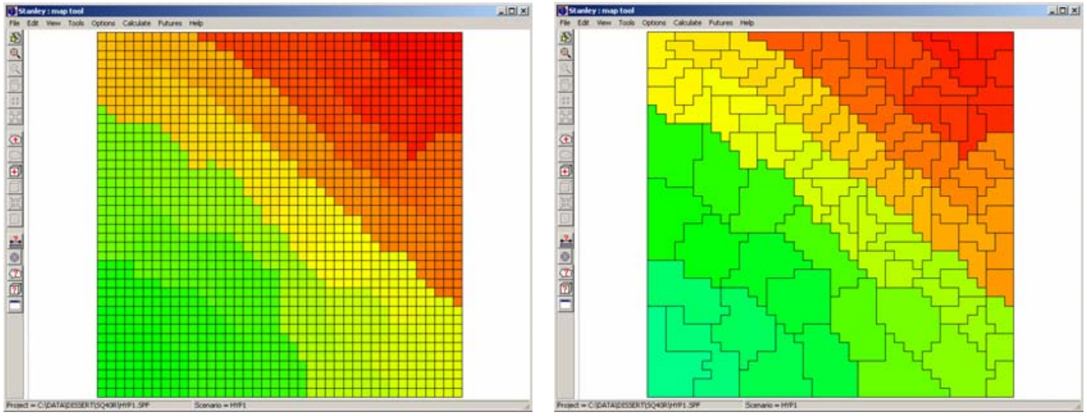

Figure 7. Allocation units used by Stanley to allocate 15-period tactical plan. On the left, Stanley allocates 20-ha cells; on the right, Stanley must allocate 200 ha blocks.

Table 1. Impact of preblocked configuration on harvest achievement. Percent shortfall is given in parentheses.

Adherence to Even Flow	**Unblocked Stanley Solution Achieved volume (yellow) relative to strategic plan (navy)**		**Preblocked Stanley Solution Achieved volume (yellow) relative to strategic plan (navy)**	
	%	**Resulting Flow Profile**	**%**	**Resulting Flow Profile**
10% bound	53.8 (46.2)	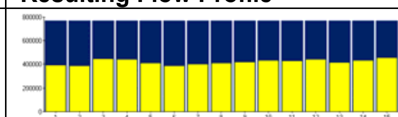	49.8% (50.2)	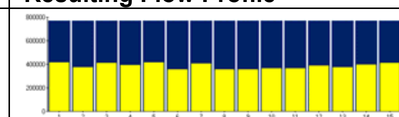
20% bound	63.6 (36.4)	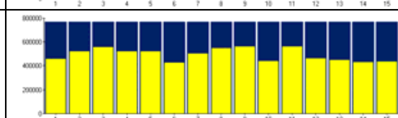	62.9% (37.3)	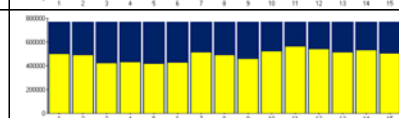
30% bound	70.1 (29.9)	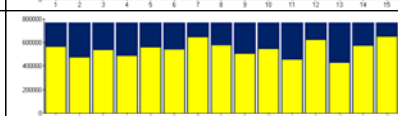	65.8% (34.2)	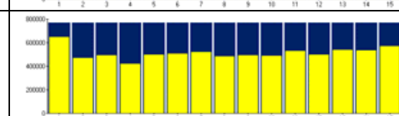
40% bound	73.7 (26.3)	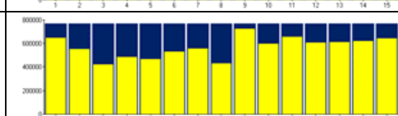	67.7% (32.3)	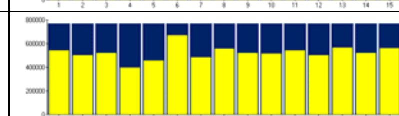
50% bound	74.4 (25.6)	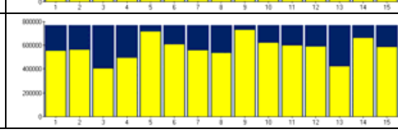	69.9% (30.1)	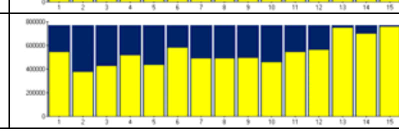

In both cases, there was a significant shortfall in achievable harvest volume when tight bounds were placed on period-to-period volume fluctuations using large harvest blocks. More volume could probably be harvested by allowing some smaller harvest blocks. In mountainous terrain, block configurations may be severely constrained by slope and equipment requirements such that it may be impractical to devise alternative block configurations. The same may be true in lowland or wetland areas where block configurations may be limited by the location of high ground for logging access and decking. However, given the potential impacts due to adjacency constraints, the configuration of harvest blocks should be well thought out to ensure that green-up delays are not costing more in wood volume than is being saved by efficient blocks. Additional volume could probably be achieved by dispensing with even-flow constraints as well. However, if even-flow were not an important goal, why would it be in the strategic model in the first place?

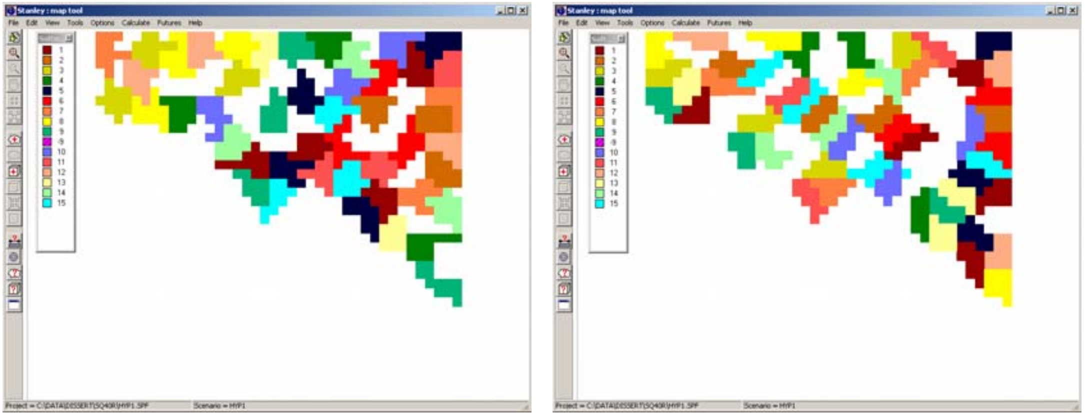

Figure 8. Final block schedules based on Stanley blocks (left) and pre-blocks (right).

In Figure 8, we can see that the pre-blocked configuration required leaving larger gaps between conflicting blocks than did the configuration devised by the Stanley algorithm. Stanley can aggregate polygons into blocks, but cannot subdivide polygons, and therefore must leave entire blocks unharvested if conflicts arise.

4 Green-up Interval Length

Walters and Cox (2001) conducted a study on a forest in the southeastern coastal plain to determine what impact spatial restrictions had on achieving harvest volumes predicted in a strategic harvest schedule. Findings indicated that the most important determinant of shortfalls was the length of the green-up interval; in that study, a volume reduction of about 5% occurred for every additional year of green-up interval. To illustrate the more general case of green-up interval interacting with spatial structure of stands, we revisit the scattered and clustered forest examples. Since the strategic harvest scheduling models are identical, any differences in shortfall observed must be due to the spatial configuration of stands alone.

Using the additional fragmentation analysis provided by Spatial Woodstock, we note an immediate difference between the scattered and the clustered forests: imposing access delays on otherwise impossible-to-block polygons in the scattered forest lowered the even-flow harvest volume by 3.4% relative to the clustered forest (which required no access delays). When these harvest schedules were processed through Stanley using identical spatial parameters[1], we observed that 5-period green-up intervals had much more of an impact on the clustered forest than the scattered forest, yielding less than 50% of the volume predicted by the strategic schedule (see Table 2).

Reducing the green-up interval by 1 period incrementally improves the harvest attainment in both forests, but there is a dramatic improvement in the clustered forest by reducing the green-up interval from 4 periods to 3 (Table 3). This pattern is not observed in the scattered forest. If we compare the mapped solutions for the clustered forest, we see a gradual increase in the area allocated to blocks associated with decreases in green-up until the interval reaches 2 periods. Much more area is allocated at this point, indicating that combinations of timing choices, opening size limits and buffer distances under a 2-period green-up interval allow for dramatically more blocking alternatives.

[1] Proximity distance = 500 m. Minimum block size = 40 ha. Maximum block size = 200 ha. Green-up intervals ranged from 2 to 6 periods. Periodic deviations from strategic timing choices = ±10 periods. Flow variation = +10% relative to minimum.

Table 2. Periodic volume flow and harvest volume achievement for the scattered and clustered forests over 15 periods.

Green-up Interval	**Scattered forest** Annual allowable cut = 742,716 m^3		**Clustered forest** Annual allowable cut = 769,432 m^3	
	%	**Resulting Flow Profile**	**%**	**Resulting Flow Profile**
6 periods	77.2 (22.8)	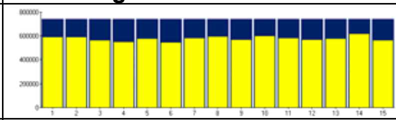	47.8 (52.2)	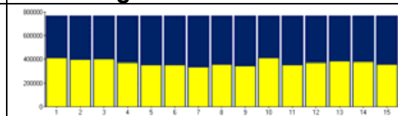
5 periods	82.5 (17.5)	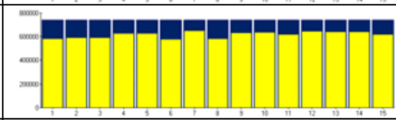	53.6 (46.4)	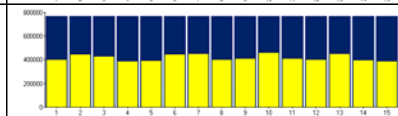
4 periods	90.7 (9.3)	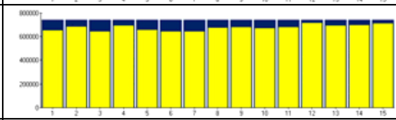	70.4 (29.6)	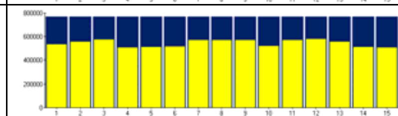
3 periods	93.6 (6.4)	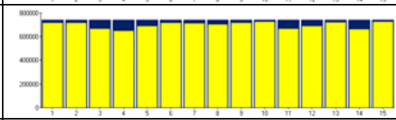	88.2 (11.8)	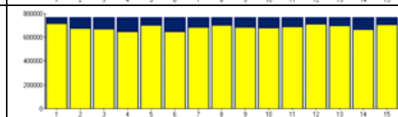
2 periods	94.3 (5.7)	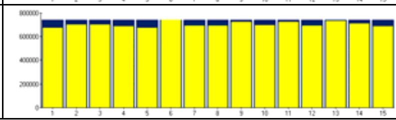	98.2 (1.8)	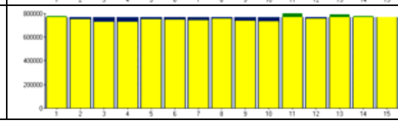

Most of the shortfall in the clustered forest is due to being unable to harvest stands due to adjacency/green-up/maximum opening-size restrictions. Stands are large, contiguous and uniform and in theory should yield nice harvest blocks. However, since the operable development types are clustered in one area, it is difficult to schedule any particular block for harvest without it conflicting with another from a different planning period. The algorithm is then forced to leave unharvested buffers around harvest blocks in order to avoid spatial conflicts, which yields 100% shortfall in harvest on these unharvested areas.

In the scattered forest, much of the area can be harvested because age classes are already dispersed throughout the forest. The difficulty is not adjacency but a lack of contiguity: harvest blocks have to cover a wide range of ages just to meet the minimum block size and therefore very few of them are selected for harvest at their optimum age. The only way to cover a wide range of timing choices for a single development type is to allow timing choice deviations from the strategic solution. If timing choice deviations are not allowed on the scattered forest, Stanley is barely able to schedule any of the harvest.

To bring the strategic and tactical solutions more in line with one another, additional activities or constraints could be incorporated into the strategic harvest schedule. On the clustered forest, intermediate harvests that do not incur adjacency restrictions could produce additional volumes. In the scattered forest, constraints forcing a wider range of timing choices would allow Stanley additional flexibility in blocking stands, without incurring penalties for timing choice deviations. However, the biggest improvement to shortfalls would be had by shortening the length of the green-up interval through more aggressive establishment methods that reduce the time until the stand reaches the desired stand condition (e.g., “free-to-grow”, 5’ height, etc.).

Table 3. Block distributions using green-up intervals ranging from 5 down to 2 periods on the scattered forest and the clustered forest. Colors correspond to harvest period. Grey indicates area scheduled for harvest but left unharvested.

Green-up Interval	Scattered Forest	Clustered Forest
5 periods *Much more area left unharvested (grey cells) in the clustered forest than in the scattered forest.*		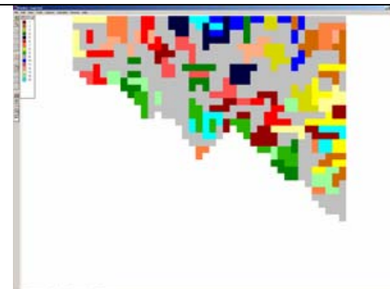
4 periods *Some reduction in area unharvested in both forests.*		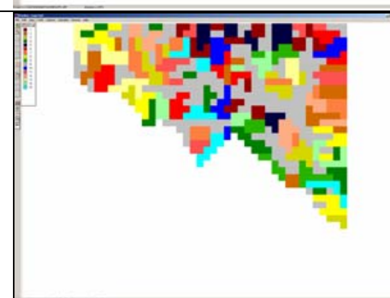
3 periods *Area left unharvested about equal in scattered and clustered forests with dramatic reduction in unharvested areas in clustered forest.*		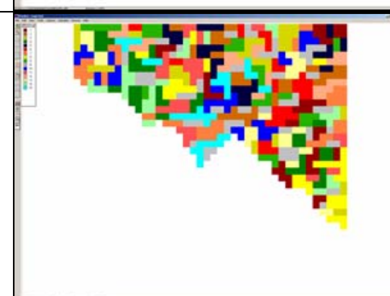
2 periods *Area unharvested about the same in the scattered forest and virtually no unharvested area in the clustered forest.*		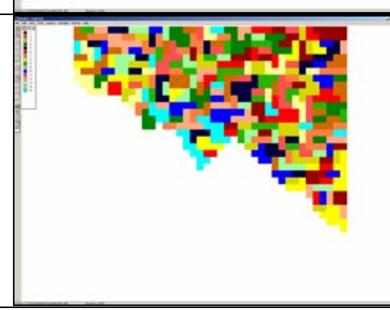

The examples we used represent the worst extremes of spatial distributions and over time we would expect the spatial distributions of both forests to be driven to some middle condition of dispersed but contiguous patches of similar age/stand structure. However, until sufficient time elapses, it is probably safe to assume that forest areas that share characteristics similar to the scattered forest example (New England, Atlantic Canada) will continue to exhibit difficulties in blocking and meeting optimal harvest timing objectives, and that land managers in the U.S. south that have much in common with the clustered forest example will continue to run up against adjacency and green-up violations as they seek to establish harvest blocks.

5 Conclusions

We have shown how three different spatial issues can affect the tactical planning problem of identifying harvest blocks and scheduling them. Overestimating the operable land base can be a significant problem in New England and Atlantic Canada where forests are typically composed of small stands, composed of many different and exhibit many differences in species, age and structure. If a minimally feasible block is 10 ha and the average stand size is less than 4 ha, most harvest blocks will have to be comprised of at least 3 stands. If the age differences among stands are large, one may have to wait well beyond one stand's optimum rotation age until a neighboring stand first become operable. The shortfall from harvesting stands too early or too late can be significant, even without considering financial discounting.

If the strategic harvest schedule chooses these small polygons for harvest, tactical planners may find it difficult to map out any feasible blocks based on the strategic schedule, and they will complain that the strategic plan is useless. Some planners have created potential harvest units in a planning layer, intersected that layer with the forest inventory layer and developed weighted average yield tables for each of the harvest units. The strategic model was then formulated as an area-based model rather than the more conventional stratum-based approach. Although the area-based model was complex to develop, and the shortfalls in harvest level that resulted were large, they felt it was still more cumbersome to spatially schedule their stratum-based schedules.

To address this issue, Spatial Woodstock includes a fragmentation analysis that identifies polygons that should be delayed from harvest until neighboring polygons are also eligible for harvest (spatial operability locks). Remsoft found that the strategic harvest volume on one Crown forest license was reduced by 25% due to the large number of fragmented stands that required spatial operability locks, but the resulting harvest schedule was much easier to block and schedule than the unconstrained model. It is important to remember that the spatial operability locks don't guarantee that groups of operable polygons are scheduled simultaneously, but they do help to coordinate the harvest of tracts in periods in which harvest options are at least present, if not numerous.

In the U.S. southeast, historical fragmentation of industrial forests has not been an issue. Age classes are more typically clustered into large plantations. Yet, a tool that was designed to work with fragmented stands (Stanley) is able to address the problem of adjacency/green-up inherent in SFI planning regulations. By subdividing large plantations (stands) into smaller units, either through GIS intersection of different planning layers, or by applying systematic grids of square or hexagonal polygons, Stanley is able to allocate the strategic harvest schedule and identify harvest blocks that do not result in adjacency conflicts. This tactical solution is already in map-form and can be given to forest planners as a head start to developing a true operating plan. Reports have indicated that 60-70% of a Stanley map solution can be implemented as is, with the remaining 30-40% requiring manual intervention to change boundaries or harvest timing choices (Walters & Cox, 2001).

A few years back, some researchers were investigating how to measure spatial net down factors (Daust & Nelson, 1993). Recognizing that it is difficult to directly incorporate adjacency and green-up constraints into strategic harvest schedules, they hoped to determine reasonable ratios between the theoretically possible and the practically implementable (e.g., say, 85% of a non-spatial harvest level could be blocked out and scheduled). Unfortunately, variations in forest condition, operating conditions and jurisdictions resulted in so many unique circumstances that any general rule-of-thumb was useless. But with the advent of sophisticated planning tools like Spatial Woodstock and Stanley, it is much easier to revisit the strategic model and constrain it based on actual forest conditions so that strategic model results are much closer to realizable goals. We have found that running Woodstock/Stanley in an iterative fashion tends to establish a nice compromise between the promise of an optimal strategic result and a practical, feasible tactical solution.

6 Literature Cited

Daust, D.K., and J.D. Nelson. 1993. Spatial reduction factors for strata–based harvest schedules. *For. Sci. 39(1): 152–165.*

National Council for Air and Stream Improvement, Inc. (NCASI). 2002. HabPlan Documentation: Harvest Scheduling and Landscape Planning Software. 36pp.

Walters, K.R. and E.S. Cox. 2001. An Empirical Evaluation of Spatial Restrictions in Industrial Harvest Scheduling: the SFI Planning Problem. *South. J. Appl. For. 25(2):60-68*

Walters, K.R. and U. Feunekes. 1994. A hierarchical approach to spatial planning: a report card. Proc. GIS '94 Symposium. Vancouver, BC, Canada. Feb. 1994