

A Three-Stage Heuristic for Harvest Scheduling with Access Road Network Development

Mark M. Clark, Russell D. Meller, and Timothy P. McDonald

ABSTRACT. In this article we present a new model for the scheduling of forest harvesting with spatial and temporal constraints. Our approach is unique in that we incorporate access road network development into the harvest scheduling selection process. Due to the difficulty of solving the problem optimally, we develop a heuristic that consists of a solution construction stage and two solution improvement stages. We call our approach INROADS and compare it to three other approaches by employing hypothetical example problems with 225 stands (or cut blocks) over a three-period planning horizon. Thirteen example forests that vary in terms of stand value and spatial dispersion are used to evaluate our heuristic, which outperforms the other approaches tested. *For. Sci.* 46(2):204–218.

Additional Key Words: Forest planning, forest roads, spatial constraints, area-restricted planning, spanning trees.

The harvest scheduling problem is a difficult combinatorial optimization problem. The forest is typically divided into stands and the main question is to determine which stands to harvest during which period with which road network. Spatial and temporal constraints further increase the difficulty of the problem.

Access roading has been a very critical component in solving the harvest scheduling problem. Obviously, the lack of access to a stand can prohibit harvesting, but excessive road costs can be prohibitive as well. The most common means for modeling access roading has been with a fixed or prepositioned road network. Generally, the network is modeled such that each stand in the forest has limited access. As we illustrate later, this potentially restrictive method of accessing stands could be cost-prohibitive for harvesting some stands. With the increased pressures to reduce roading and roading costs, road networks should be flexible by providing multiple roading alternatives that may offer a less expensive route.

The forest harvesting problem has traditionally been hierarchical with two main levels. The higher level is longterm and tends to consider what volumes should be cut during

which period. The lower level is the shorter in length and seeks to determine which stands to harvest during which period that will satisfy the long-term plan. In addition to these two main levels, there is planning that takes place at the stand level. For example, questions concerning the number and placement of landings are answered so that the costs of roading and skidding within the stand is minimized. In practice, hauling roads that are built at the stand level are often reused for access to another stand. We develop a model that integrates the access roads and the roads within the stand in the event that doing so could potentially reduce road costs.

The most commonly used methodology to solve the harvest scheduling problem is mixed-integer programming (MIP). However, for large problems, even if roads are not integrated, MIP cannot be used due to runtime considerations. Since the main goal of a modeling approach is to be able to incorporate all of the necessary aspects of the planning problem (Murray and Church 1995)—and we believe roading decisions are necessary aspects of the harvest scheduling problem—we develop an integrated road approach. Since MIP cannot be used effectively on our problem, even for small instances, we develop a heuristic for it.

Mark M. Clark, Department of Industrial and Systems Engineering, Auburn University, AL 36849-5346—Phone: (334) 844-1424; Fax: (334) 844-1381; E-mail: clarkma@eng.auburn.edu. Russell D. Meller and Timothy P. McDonald, USDA Forest Service, Southern Research Station, Devall Drive, Auburn, AL 36849.

Acknowledgments: This research was supported by the USDA Forest Service under USDA Contracts 19-95-062 and 19-96-067. The authors would like to thank the associate editor and the referees for their very thoughtful and constructive comments.

In this article, we present a new model for the harvest scheduling problem. To illustrate the application of the model, we develop a heuristic to solve it. Our approach is unique in two respects. First, we integrate roading and stand selection by employing a node-based approach, rather than prepositioned roads, which is commonly used in other approaches. The node-based approach allows flexibility in roading, which in turn, reduces road costs. Our roading approach also integrates the access roading between stands with the roading within stands. Secondly, we employ a three-stage heuristic for solving the harvest scheduling problem. Our three-stage heuristic uses a greedy approach in the construction stage and pairwise interchange in the second two stages.

After presenting our literature review, we describe the problem in detail, and then we present comparison approaches. We evaluate the approaches next, and then offer our conclusion.

Literature Review

The harvest scheduling problem can be broken into two broad categories (Nelson et al. 1991). The first category is the long-range, strategic plan that considers several rotations, and the second deals with finding a spatially feasible solution for a shorter planning horizon. The strategic plan has historically been modeled as a linear programming problem (Iverson and Alston 1986). The strategic models lack the spatial resolution needed to determine the development schedules of stands and roads (Nelson and Finn 1991). Thus, mixed-integer programs that address both stands and road networks can be formulated. However, only small problems can be solved within a reasonable amount of time because of the integer restrictions (Kirby et al. 1980, Nelson and Brodie 1990, Jones et al. 1991). Therefore, heuristic-based approaches have been employed.

The initial heuristic-based approach for the harvest scheduling problem employed random search algorithms. Bullard (1985) was one of the first to use random search and showed that the results produced by Monte Carlo Integer Programming (MCIP) were within 1% of the optimal nonlinear programming results. They used a randomized search to solve the single-stand, thinning and final-harvest-timing problem, where multiple stands are selected for harvesting.

O'Hara et al. (1989) developed a randomized search heuristic to maximize volume harvested. The heuristic had the capability to pre-bias stands for selection based on volume and adjacency. Due to its volume-maximizing objective function, harvesting costs and road costs were not considered. However, they did employ even-flow volume constraints and spatial constraints. O'Hara applied the algorithm to a forest with 242 units, and the results indicated that the heuristic approach ranged from 6% to 7% of the optimal LP solution.

Nelson and Brodie (1990) developed a randomized search heuristic that sought to maximize net present value minus the road construction costs in the objective function. The algorithm resembles a pyramid approach by generating a large number of solutions initially, and keeping only the good

solutions to expand and build on for subsequent periods. Nelson and Brodie considered an area with 45 harvest units and 52 prespecified road links. After finding the optimal solution using an MIP approach, they found that their heuristic approach found a solution with an objective function within 3% of the optimal MIP solution and several other solutions within 10% of the optimal. Although their heuristic specifies the cost of building roads in the objective function, the road network is fixed or prepositioned such that most of the stands, especially those furthest from the main road, are accessed by only one road.

Clements et al. (1990) modified Nelson and Brodie's (1990) heuristic by extending the planning horizon and changing the objective function to a volume-maximizing function to observe the effects of different block sizes and levels of adjacency constraints on harvestable volume. Although their results indicate a reduction in harvestable volume when the exclusionary periods increase, it is unclear the effect maximum opening sizes have on the harvestable volume.

Another search algorithm is CRYSTAL, which was developed by Walters (1991). Stands are pre-biased for selection based on seven different criteria, such as area, perimeter, and stand type. Once an initial stand is chosen, it is referred to as a seed stand. Adjacent neighbors to the seed stand are then examined to determine if any of them are eligible to be harvested in that period. If so, the seed and the neighboring stand are aggregated into a potential harvest block (Jamnick and Walters 1993). As each neighbor is added to the harvest block, other stands that become neighbors are also considered for selection. This process continues until the annual allowable cut (AAC) is met or the maximum clearcut size is exceeded. If the clearcut size is reached first, a new seed stand is selected.

CRYSTAL's performance was compared to that of Clements et al. (1990) on the basis of volume harvested. The test forest had 277 harvest units that averaged 45 ha. The heuristic developed by Clements et al. yielded an average of 3.7% higher harvest volume than CRYSTAL.

Nelson and Finn (1991) adopted a volume-maximizing model and added a road network to observe the effects of reduced cut-block size and increased exclusion period on volume harvested and road networks. Their results indicated that as cut-block size decreases and exclusionary periods increase, volume harvested decreases. Also, their results show that with smaller blocks, the road construction costs increased due to the additional roads required to access eligible blocks in the early periods in order to meet the volume requirements.

Another heuristic approach to the harvest scheduling problem employs the simulated annealing (SA) algorithm that was presented by Kirkpatrick (1983). Lockwood and Moore (1993) also applied SA by using a series of penalty costs to reflect low volume, violations of the adjacency constraint, and the maximum opening constraint. Although Lockwood and Moore do not consider an access road network, they do solve the much more difficult area-restricted problem, which we will discuss in more detail later. Dahlin

and Sallnas (1993) use SA for a problem where an access road network is considered. Both Lockwood and Moore (1993) and Dahlin and Sallnas (1993) show that SA is a viable approach to large problems.

Murray and Church (1995) compare SA to two other heuristic solution approaches to the harvest scheduling problem. Pairwise interchange, SA, and Tabu search are applied to two actual problems. The results indicate that all three methods are viable solution techniques, but the Tabu search performed better than the other two approaches.

Most of the heuristics that have been developed have not considered roads. However, the design and implementation of forest road networks have been the focus of some researchers for decades. Most of the early work was accomplished by employing a manual search for roading alternatives. Since then the use of digital terrain models have been employed to help the planner find good roading alternatives considering the topography of the forest (Liu and Sessions 1993). More recently, Dean (1997) called this difficult problem the Multiple Target Access Problem (MTAP) and developed a heuristic called Branch Evaluation to solve it. The Branch Evaluation heuristic was compared to optimal solutions of small problems and proved to be effective. Later, Murray (1998) developed a mathematical model of the MTAP so that the roading problem for a stand can be generalized and solved optimally.

In summary, it has been shown that MIP can be used for solving the harvest scheduling problem; however, it is limited to only smaller problems due to the excessive number of decision variables. Thus, several heuristic approaches have been developed for larger problems and have generated near-optimal solutions. The heuristics have been developed for problems with and without road networks. Hopefully those that consider road networks in the solution methodology will lead to more accurate results. The models that have included road networks have considered prepositioned road-networks, which has advantages and disadvantages. The advantage to having prepositioned access roads is that the roads in the final network are likely to be in locations that are feasible. The disadvantage is that the flexibility in accessing stands is limited. On the other hand, the benefit to not having prepositioned roads is that there is inherently more flexibility; however, some road locations might not be feasible with a general road network approach. The recent work in the automation of road networks by Dean (1997) and Murray (1998) should help in finding solutions that are both flexible and feasible.

In this article we develop a model that integrates access road development with stand selection for the harvest scheduling problem. To solve the model, we develop a heuristic for it. In our heuristic we do not use a prepositioned road network and hope to capitalize on this increased flexibility in developing the access road network.

Problem Description

The harvest scheduling problem has two forms: the area-restricted problem (ARP) and the unit-restricted problem (URP) (Murray 1999). The difference in the two forms is in the typical stand sizes and the size of the allowed maximum

opening. Stands in the URP are typically large enough to reach the maximum opening size individually. However, the ARP is more difficult than the URP because multiple contiguous stands (or clusters of stands) are usually required to reach the maximum opening size. Our problem is similar to that of Lockwood and Moore (1993) in that we solve the ARP version of the problem, but with an equal-area grid-based representation of the forest.

Although the harvest scheduling problem considers spatial constraints, it generally does not address constructing an access road network. However, the MIP formulation developed by Nelson and Brodie (1990), and later enhanced by Murray and Church (1995), addresses roading decisions. It is this enhanced version of the harvest scheduling problem that we solve. However, we approach the problem of access roads differently than the previous MIP, and consider an ARP, as opposed to a URP.

Since roads play such an important role in the forest harvesting problem, we decided to increase the options available in constructing the road network in our problem. Thus, we generalize the prepositioned road approach of previous work and develop a road network based on a node-based system as shown in Figure 1(a).

In Figure 1(a), there is a node located at the center of each stand and a node at each possible access point from an existing road (located at the top of the forest in this case). Based on which stands are selected, the nodes in Figure 1(a) must be connected to form an access road network. More detailed networks may be developed utilizing more than one node per stand.

Our approach to roading is based on a very well-known network problem, the minimum spanning tree (MST) problem (Sedgewick 1988), which is similar to the also well-known shortest path problem. Both problems involve choosing a set of links that have the shortest total length, given a connected, but undirected network. Each link in the network represents some positive length or cost. For the shortest path problem, the chosen links must provide a path between the origin and the destination. The minimum spanning tree problem is different in that there is no specified destination. Instead, the required property is that the chosen links must provide a path between each pair of nodes in the network.

Consider the 16 stands that are in Figure 1(a). Assume that we are to harvest the two bottom stands in the lefthand corner. (Recall that there is an existing road to access these stands located along the top boundary of the 16 stands.) If we had not known that we would harvest these two stands simultaneously, we may have considered the prepositioned

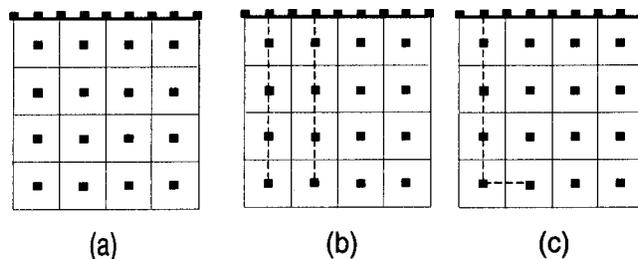


Figure 1. Node-based approach.

road network similar to Figure 1(b). However, if we did not consider any specific access road locations *a priori*, we could use a minimum spanning tree approach to access the two stands. Assuming that the existing road is a node and the other landing locations are also nodes, we could connect the nodes as in Figure 1(c). Depending on road costs, the solution in Figure 1(c) could be much better than the solution in Figure 1(b).

Problem Formulation

We develop a MIP for the harvest scheduling problem that considers an integrated road approach as outlined above (see box below). The notation used in our formulation follows.

- a_{it} discounted revenue generated from harvesting stand i in period t
- c_{jit} discounted cost to build a road segment from stand j (or existing road) to stand i during period t (assume $c_{jit} > 0 \quad \forall j, i, t$)
- A maximum opening area (maximum contiguously connected area allowed for harvest in one period)
- e exclusionary period width (a stand can only be harvested once in every $2e+1$ periods)
- v_i area of stand i
- S set of all stands
- N set of all existing roads (nodes) at the beginning of the planning horizon
- N' $S \cup N$ = set of all stands and set of all existing roads (nodes) at the beginning of the planning horizon
- T planning horizon
- H_t harvest target (area) in period t
- x_{it} 1, if stand i is harvested in period t , and 0 otherwise

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

$N(t)$ set of all harvested stands and set of all existing roads (nodes) at the beginning of period t

r_{jit} 1, if road segment is built from stand j (or existing road) to stand i during period t , and 0 otherwise

The formulation seeks to maximize total net revenue minus road costs subject to spatial and temporal constraints. The net revenue is the stumpage value minus the harvesting costs. The harvesting costs include the felling, skidding, landings, and the interior road costs (roads within a stand). Therefore, before any access road calculations are made, the gross revenue, the harvesting costs, and the net revenue are known. The other component of the objective function is the access roading costs. The access road consists only of the roading required to be built from an existing road in the forest to the entry point of the stand in question.

The first constraint of the model, (1), enforces the exclusionary period on harvesting a stand as defined above. The second constraint, (2), ensures that the harvest target is met for each period during the planning horizon. [Note that (2) is written as an equality and may have to be adjusted for a particular problem.] The third constraint, (3), is required as a result of our minimum spanning tree approach to roading decisions. This constraint is necessary because one may only access a stand to be harvested from a previously harvested stand or an existing road in the forest. Since any previously harvested stand must have been connected in the spanning tree, the result will also be a spanning tree. Our fourth constraint, (4), limits the maximum opening area to less than or equal to a prespecified size. The definition of adjacency in our model is when two stands share a common boundary as

$$\begin{aligned}
 \max \quad & \sum_{i \in S} \sum_{t \in T} a_{it} x_{it} - \sum_{j \in N'} \sum_{i \in S} \sum_{t \in T} c_{jit} r_{jit} \\
 \text{s.t.} \quad & \sum_{\ell=t-e}^{t+e} x_{i\ell} \leq 1 \quad \forall i \in S, t \in [e+1, T-e] \quad (1) \\
 & \sum_{i \in S} v_i x_{it} = H_t \quad \forall t \quad (2) \\
 & \sum_{j \in N(t)} r_{jit} \geq x_{it} \quad \forall i \in S, t \in T \quad (3) \\
 & \sum_{\ell \in P_{it}} v_{\ell} x_{\ell t} \leq A \quad \forall i \in S, t \in T \quad (4) \\
 & x_{it} = [0, 1], r_{jit} = [0, 1] \quad (5)
 \end{aligned}$$

opposed to a single point. Our last constraint, (5), specifies that the decision variables for stand selection and road segments are binary. Constraints that restrict the harvesting of stands near recently harvested stands can also be incorporated into the model.

The above formulation was presented to illustrate the problem, but not as a recommended approach to solve it. Note that the set P_{it} must be defined after the x_{it} variables have been set for period t (see Murray 1999 for a complete discussion). Furthermore, our more general consideration of the road network via the MST comes at a price. The set $N(t)$ can only be defined after the r_{jit} variables have been set for period $t - 1$. Also, the minimal spanning tree(s) cannot be solved by the minimal spanning tree algorithm since it involves staging and the tree(s) utilizes more nodes and arcs than the number of stands being harvested would predict. Thus, attempting to solve this model would pose even more of a challenge than the size of the integer program would suggest. Due to these complications, we chose to develop a heuristic for this problem rather than attempt to simplify and solve our formulation.

Integrated Road Networking and Stand Selection: INROADS

We employ a grid-based approach to represent our forest. In general, a grid refers to two sets of equally spaced parallel lines where the two sets are orthogonal to each other. The result is a two-dimensional plane of equal-sized squares. Every cell on this grid represents a stand or cut block. This representation of the forest is similar to that used by Snyder and ReVelle (1996). Furthermore, if we place a node at the intersection of the gridlines, and subsequently superimpose these nodes over the forest, x and y coordinates can be established to represent physical locations in the forest. At the operational level, we use these nodes to identify the locations of access roading and the boundaries of stands. As we will show later, the nodes can also be used at the stand level to identify landing and interior roading locations if there is an adequate number of nodes per stand.

There are various levels of detail that we could employ to build our road network. First we could use very little detail and only have nodes, say, at the centroid of each stand once the stand boundaries are identified (as was the case in Figure 1). Alternatively, we could have many nodes for each stand so that interior roads and landing locations could be represented. For the purposes of integrating the roads at the operational level with the roads at the stand level, we use a fairly high number of nodes for each stand. We assume all stands are square-shaped and equal in size. Since we will build our road network from the nodes in our model, this implies our roads are orthogonal and that rectilinear distances are appropriate. We further assume the terrain is flat, with no obstructions, to simplify the access road cost model.

Integrating Access Roads and Roads within the Stand

In order to stay abreast of where the access roads are in the forest, nodes are flagged as either active or inactive. An active node means that it represents the endpoint of a road segment.

We assume that the link between two active nodes represents a road segment. Likewise, an inactive node means that it does not act as a road segment endpoint. A requirement for the forest is that there is at least one active node (existing road) in the forest at the beginning of the first period. If there is in fact only one active node, then we have a Steiner tree, as Murray (1998) describes. Otherwise, we can solve the road network problem using multiple trees if necessary.

For our problem, each stand is represented by equal-sized, square-shaped stands. In isolation, each stand has an optimal harvest pattern, which considers three operational cost components: landing, roading, and skidding costs. In other words, based on the magnitude of these costs and the parameters of the stand, such as volume and size, we can determine the optimal spacing of roads, landings, and skidding regions within a bounded stand using the model developed by Clark et al. (1999). Thus, the harvest pattern specifies the number and locations of landings and the interior roading pattern that minimizes total harvesting costs. Based on the optimal harvest pattern, a single stand is subdivided by a series of vertical and horizontal lines. Consider Figure 2, where there are seven example subdivisions of a square-shaped stand, any of which could be optimal depending on the stand parameters and harvesting costs.

In order to accommodate all of the harvest patterns in Figure 2, each stand consists of 49 nodes. These 49 nodes are located at the intersections and endpoints of the vertical and horizontal lines of the harvest patterns. If these seven patterns are superimposed on one another, the stand would look similar to the illustration labeled "master" in Figure 2. (Note that if we attempted to represent the "master" node system in our previous MIP formulation we would be required to modify it substantially, which would ultimately make the MIP formulation even more difficult to represent and solve.)

Once the optimal harvest pattern is known for a specific stand, we can then calculate its net revenue since we can determine the landing locations and the interior roading. We incorporate the landing locations into our minimum spanning tree approach to develop the access road network. Consider the 16 stands shown in Figure 1(a). Now suppose that we have determined the optimal harvest patterns for the same two stands to be harvested (bottom lefthand corner) and determined the landing locations, as shown in Figure 3(a). We can employ the minimum spanning tree approach and develop a road network that will access each of the landing locations as

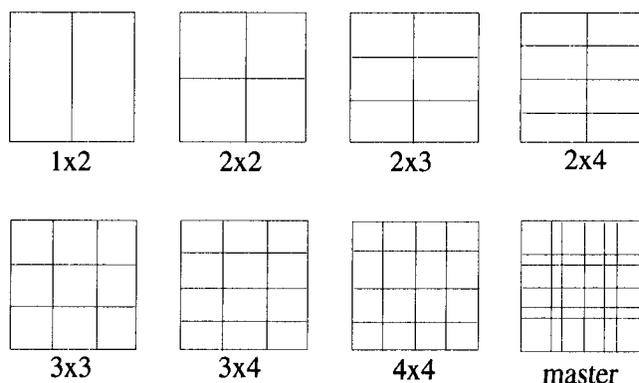


Figure 2. Subdivisions of a square-shaped stand.

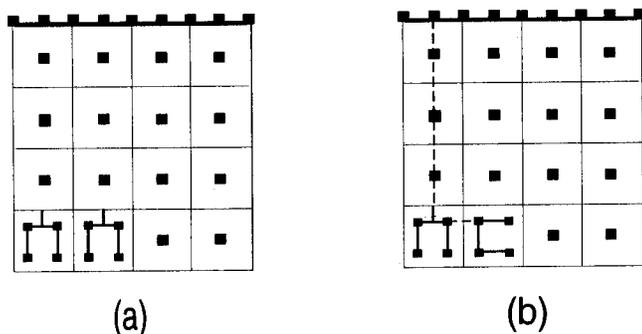


Figure 3. Minimum spanning tree approach.

shown in Figure 3(b). Note how the harvest pattern can be rotated to minimize road costs.

In summary, the nodes in our approach can represent stand boundaries, landing locations, interior roading, and access roading. This type of approach provides flexibility in roading alternatives, which promotes good roading decisions.

Access Road Cost Determination

The access road cost for a specific stand is found by calculating the distance from that stand to an existing road. Roads are represented by active nodes on the master node list. We use four search routines to search for the active node nearest the entry node of a stand. The approach for identifying for the nearest active node is systematic; i.e., we start searching near the stand and systematically move away from the stand in an effort to locate the nearest active node (Clark 1998). Once an active node is found, the distance from the active node to the stand in question can be calculated and the road cost can be assigned to this stand. If there are multiple active nodes, an evaluation of these nodes is made, and only the active node that corresponds to the lowest road cost is retained. After these costs are found, the profit for each stand can be calculated.

The INROADS Heuristic

INROADS is a three-stage heuristic where the first stage is a construction stage and the second and third stages are improvement stages. For multiperiod problems, each stage is performed once for each period. A flow chart for the multiperiod problem is shown in Figure 4.

Stage 1: Directed Search

Stage 1 of INROADS is the construction stage. Its purpose is to find a good initial solution. We refer to Stage 1 as the Directed Search because we employ a greedy approach to the stand selection process for the initial solution. After we have selected the stands for harvest with this greedy approach, we attempt to reduce the access roading cost by finding the minimum spanning tree solution for the selected stands.

The Stage 1 greedy approach simply selects stands, one by one, that have the highest profit values (net revenue—access road costs) among the eligible stands. A stand that is not harvested and does not violate the maximum opening constraint is considered eligible. The heuristic is sequential and considers the access roads that were required for stands previously selected. That is, after selecting a stand for harvest, access roads are temporarily mapped to

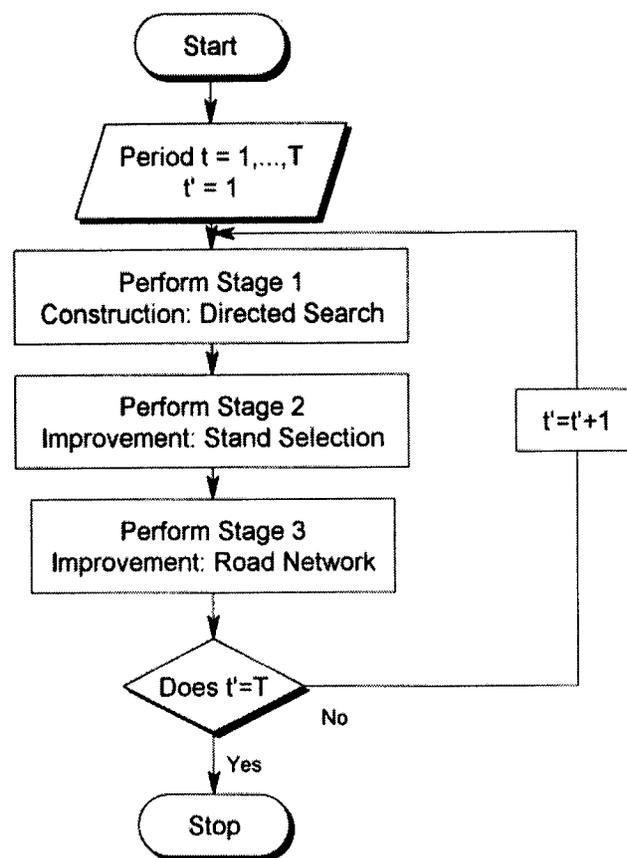


Figure 4. Main flowchart for INROADS.

that stand. When selecting the next stand for harvest, profit values for each remaining eligible stand is recalculated taking into consideration the access roads that were mapped to previously selected stands. For example, assume initially that I stands are to be harvested but none have been chosen. The heuristic calculates the profit (net revenue—road costs) for each stand in the forest. The stand with the highest profit value is chosen for harvest. The nodes that represent the endpoints of the road segments required to access and harvest the selected stand are marked active. The selected stand is flagged for harvest and the process starts over, looking for the next best stand. The profit values for eligible stands are recalculated, taking into consideration the new active nodes. Of the remaining eligible stands, the stand with the highest profit value is selected and marked for harvest. The nodes required to access this stand are made active. The process is repeated until I stands have been selected. Once the I stands have been selected, a total profit value (sum of the profits for the I stands) is calculated. A flowchart of Stage 1 is shown in Figure 5.

Before completing Stage 1, the potential benefit of lower access road costs is explored. In order to realize the potential road cost savings, the selection sequence of the I stands is altered, and the roads that were constructed earlier in the stage are erased. Given the I selected stands, we find the minimum spanning tree for these stands. This step is accomplished by the subroutine called Solution Evaluation, which is shown in the flow chart in Figure 6. The sum of the profit

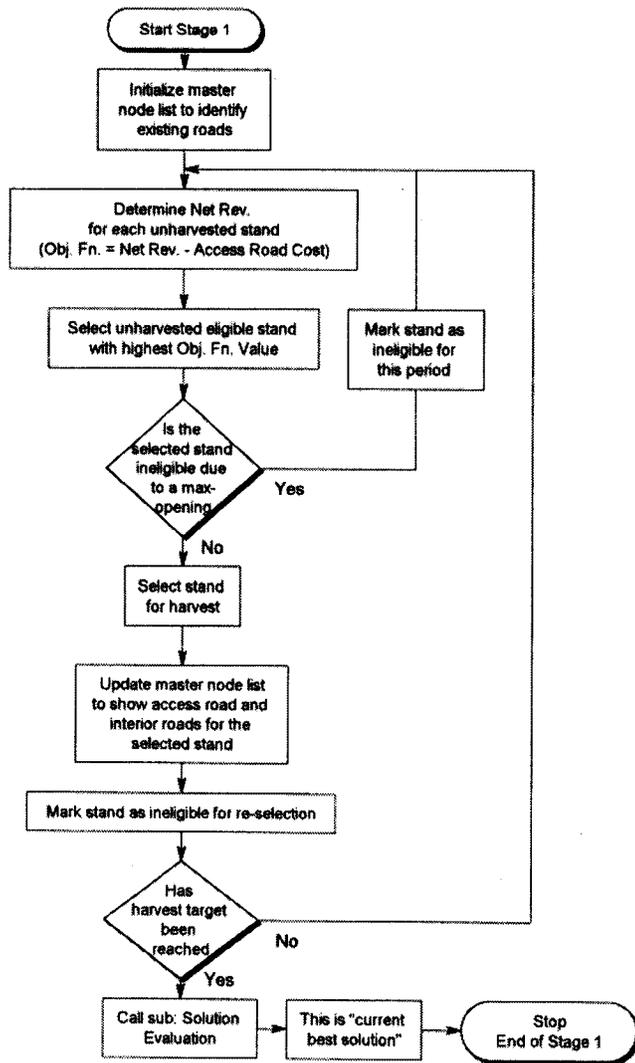


Figure 5. Flowchart for Stage 1 of INROADS.

values from each selected stand that is calculated in the solution evaluation step is now the current best solution.

Stage 2: Stand Selection Improvement

Once the construction stage of Stage 1 is complete and the I stands are marked for harvest, Stage 2 seeks to improve the total profit value through pairwise exchange of stands. Essentially, the I stands are taken out of the solution one at a time and replaced one at a time with an eligible stand. The total profit value is recalculated using the minimum spanning tree approach (the subroutine Solution Evaluation method shown in Figure 5) and compared to the current best objective function value. If a better solution value is found, the "old" current best solution is replaced with the "new" current best value. This pairwise interchange process continues until every stand has been removed and replaced by every other nonharvested stand. If an improvement is found in the first iteration, then a second iteration of pairwise interchanges is performed. This iterative process continues until no further pairwise interchange improvements are found. At the end of Stage 2, the stands that are in the current best solution represent the final stands selected for harvest. The flow chart for Stage 2 is shown in Figure 7.

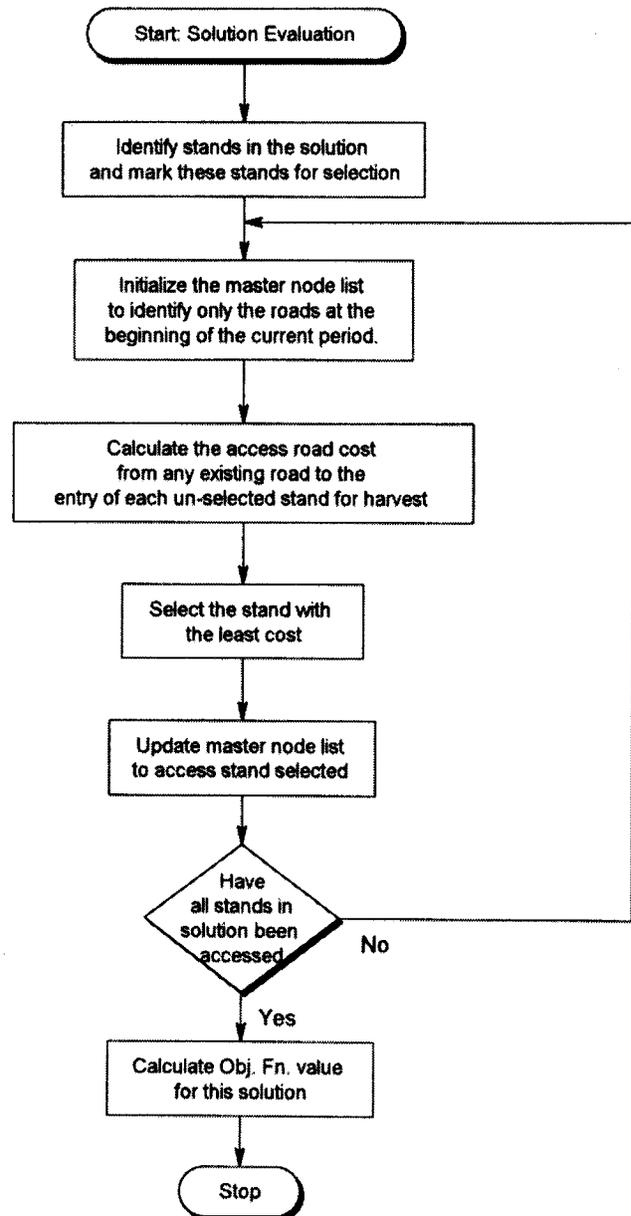


Figure 6. Solution evaluation flowchart.

If all J stands that are not in the solution are included in the pairwise interchange, then $(I \times J)$ exchanges would be considered per iteration. Thus, Stage 2 can be very time consuming, especially for larger problems. The number of pairwise interchanges can be reduced by eliminating stands that have poor profit values, thus decreasing the value of J . We reduce the total number of interchanges by flagging the top $I + n$ stands in Stage 1, and then perform the pairwise interchanges with only the additional n stands flagged in Stage 1. This is denoted as "Interchange n " in our heuristic.

Stage 3: Access Road Cost Improvement

After Stage 2 is complete, the third and final stage of the procedure is performed. This stage is an improvement stage as well. This stage seeks to improve the total profit value by reducing the road costs required to access the final stands selected for harvest in Stage 2. That is, we know which stands are to be harvested, but we do not know the sequence that the stands will be harvested during that specific period. Once the inter-

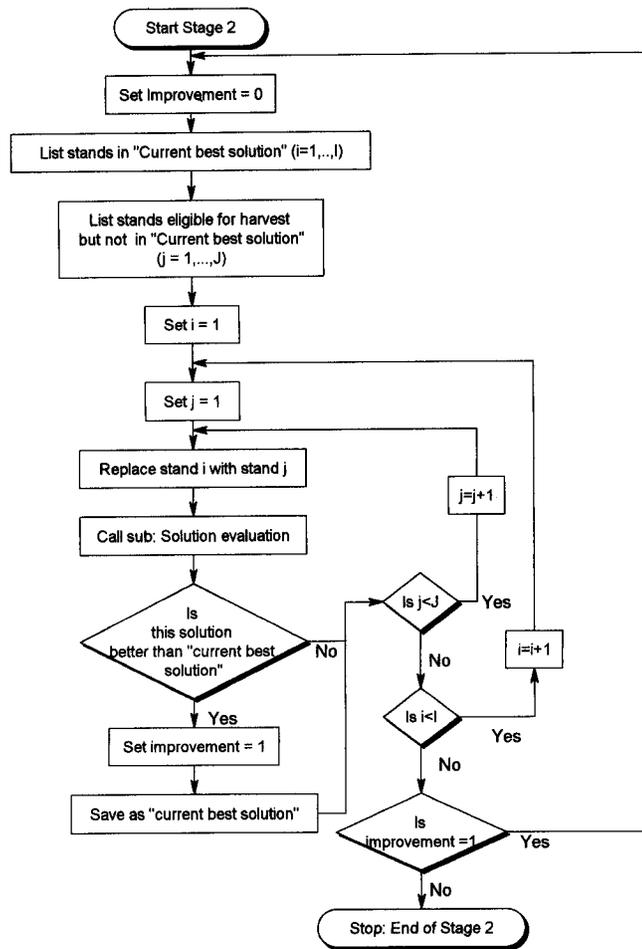


Figure 7. Flowchart for Stage 2 of INROADS.

change procedure was completed in Stage 2 and the selected stands were identified, the stands entered the solution based on the lowest access road costs. This was done in an effort to reduce road cost by harvesting stands closest to the existing roads first. Since the access road network uses orthogonal roads, plus the fact that new roads can be built from existing roads, the sequence that the stands enter the solution could have an impact on the final cost of the road network. Therefore, the purpose of Stage 3 is to evaluate other possible road networks in an effort to further reduce road cost.

In Stage 3, we employ pairwise interchange again, but this time the sequence that the stands enter the solution is interchanged. A description of Stage 3 is shown in the flow chart in Figure 8.

Comparison Approaches

In order to determine the effectiveness of INROADS, four other heuristics have been developed for comparison purposes: (1) a progressive approach, (2) a seed-stand/adjacency approach, (3) a prepositioned road approach, and (4) the Directed Search approach (which is a simplified version of INROADS).

The results from heuristic approaches are usually also compared to the results of an optimal solution. The difficulty with this is that it is not easy to solve even small instances of

our problem optimally. Secondly, as Murray and Church (1995) explain, one of the main goals of heuristics is to be able to incorporate all the necessary aspects of the planning problem that might not be possible to represent in an LP or other optimization approaches, as well as being capable of providing good feasible solutions. In the case of this problem, the complexity of the mathematical model makes it nearly impossible to solve for our integrated road approach. Therefore, we make no attempt at making a comparison to an optimal solution.

The stand selection criterion for the progressive approach is road costs. The secondary criterion is profit. Thus, the stand with the lowest road cost is selected first for harvest. If there are multiple stands that are tied for having the lowest road cost, then net revenue is used to break the tie. In the absence of maximum opening constraints, the heuristic continues to select stands in this fashion until the annual allowable cut is reached. When maximum opening constraints are present, adjacent stands become ineligible once the maximum opening is reached. Thus, stands with the lowest access road cost that are eligible will be selected for harvest.

The second comparison approach is the seed-stand/adjacency approach. The underlying principle of seed-stand/adjacency approach is to select a stand for harvest based on one of several different criteria and then select the stands adjacent to this seed stand until either the maximum opening or the annual allowable cut (AAC) has been reached. There can be multiple criteria for selecting the seed stand such as value, volume, or size. Our criterion for selecting seed stands is net profit. After the seed stand has been selected, adjacent stands are selected for harvest if they are eligible. Depending on the eligibility of the stands and the maximum opening size, the harvested opening could take on many different shapes.

The prepositioned road approach is the third comparison approach. Representing roads in a prepositioned fashion has been the traditional means in solving the forest harvesting problem. One of the early applications of prepositioned roads was its use in mixed integer programming (MIP) models. In an effort to maintain the spirit of this approach to access roading, a similar application is made with a grid layout of stands with prepositioned roads. However, we employ a heuristic-based stand selection procedure in place of an MIP. Using an MIP would have likely generated better results than a heuristic-based approach; however, the problem complexity made this prohibitive.

In the prepositioned road approach, we select stands in a greedy fashion, where profit is the selection criterion. A two-step process will be employed, just as in Stage 1 of our three-stage heuristic. The only difference in this approach and Stage 1 of our three-stage heuristic is that the prepositioned road approach will be constrained to using only the prepositioned roads.

The fourth and final comparison approach is the Directed Search approach, which is identical to Stage 1 (the construction stage) of INROADS. The main purpose for evaluating the results of the construction phase with that of INROADS is to determine the impact of the improvements made in Stages 2 and 3.

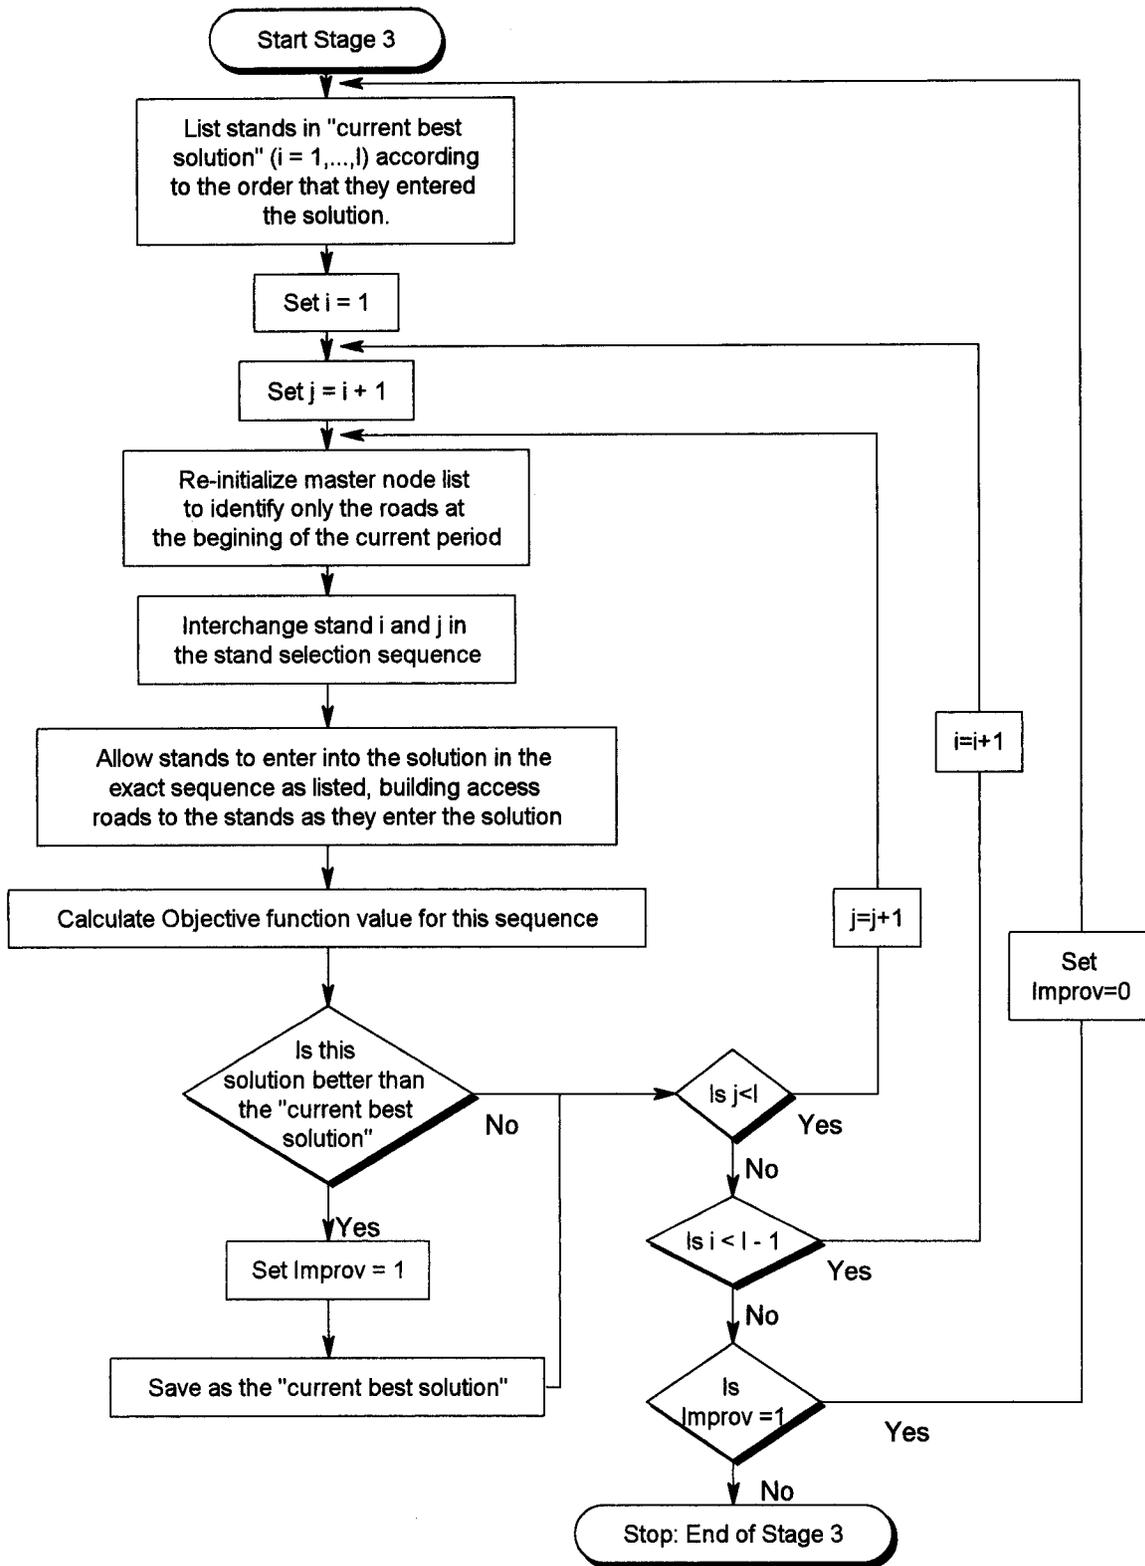


Figure 8. Flowchart for Stage 3 of INROADS.

Example Problems

Thirteen example problems were developed in order to evaluate the heuristics. The 13 examples varied in two respects. The first is in the value of each stand and the second is the dispersion of stand value throughout the forest. Thus, each example can be described by these two factors.

In addition to comparing INROADS to the heuristics from the previous section, INROADS will be further analyzed by three levels of "interchange." The level of interchange refers to the number of candidates considered for pairwise interchange during Stage 2 of the heuristic. In general, as the number of candidates increases, the performance increases as well. However, as the number of candidates increases, the

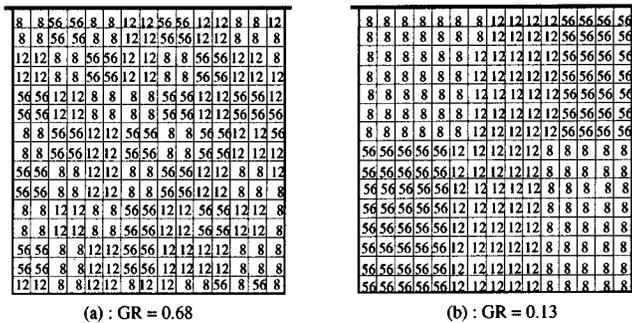


Figure 9. Structure of example forest (70/30).

time required to evaluate the interchanges increases. Thus, there is a “cost-benefit” decision to be made for the number of interchange candidates.

The example forest is a 3642.25 ha (9,000 ac) hypothetical, square-shaped forest. There are 225 square-shaped stands, where each stand represents 16.187 ha (40 ac). There is one existing road that runs horizontally along the northern boundary of the forest. The cost parameters that were used to construct these example problems are listed in the Appendix.

In an effort to simplify the value of timber across 225 (15 × 15) stands in the forest, a single scalar value was assigned to each stand. This scalar is a constant that represents the value of timber per unit area within a stand. Careful attention was given to the assignment of these scalar values. The intent was to represent two extreme cases with three intermediate cases. One extreme case was for all of the stands to be uniform in value. The other extreme case was for a majority of the forest value to be concentrated in a relatively small number of the stands.

To characterize the stand value distribution in the five cases, we use a well-known concept known as Pareto’s Principle (Nahmias 1997). Pareto’s Principle is sometimes referred to as the 80/20 rule, as in, for example, 80% of a nation’s wealth is possessed by 20% of the population. Many other situations, including the stand value distribution of a forest, can be described by using this phenomenon.

We begin with the uniform case. If the forest is uniform in value then we can say that 50% of the value of the entire forest is represented by 50% of the stands in the forest (or any other equal percentage). For simplicity, we represent this case as 50/50. The next four cases get progressively concentrated with respect to value. They are 55/45, 60/40, 65/45, and 70/30 where, for example, 70/30 indicates that 70% of the value of the forest is represented by 30% of the stands. Table 1 shows the quantity and value of the stands in each of these five cases. Two examples of the 70/30 case are presented in Figure 9 with the scalar values from Table 1. The scalar values represent the value per unit area of the unharvested

stand. For example, there are 225 stands, 78 of the 225 have a value of \$800 per unit area (acres in these examples), 79 stands at \$1200 per unit area, and 68 stands with \$5600 per unit area.

As Figure 9 suggests, in addition to the stand value distribution, a way to characterize the physical placement of the stands in the forest is also needed. The physical placement of the stands in the forest is characterized by spatial dispersion. Another way of describing the spatial dispersion is to determine the amount of clustering of similar valued stands, which is sometimes called “spatial autocorrelation” (Griffith and Amrhein 1991). The amount of clustering of similar valued stands is important in evaluating the different approaches for stand selection. If all of the higher valued stands are located in one large contiguous area, then it would be fairly straightforward where harvesting should be conducted to maximize profits. However, as the high valued stands are dispersed, the decisions for selecting stands becomes more difficult due to the cost of access roads.

In the case of the uniform forest, physical placement has no real effect on the dispersion of value since all of the stands in the forest are equivalent. However, in the other four cases, spatial autocorrelation plays an important role. We employ the well-known Geary Ratio (Griffith and Amrhein 1991) as a means to quantify the spatial autocorrelation. Geary Ratios with values near zero indicate that similar valued stands tend to cluster together (positive autocorrelation), ratios near 1 indicate that the stands are randomly scattered (zero correlation), and ratios near 2 indicate that stands with dissimilar values tend to cluster together (negative correlation). In this article, we define adjacent stands as two stands sharing the same border. Stands that join at a single point, such as a corner, are not considered adjacent.

In our example problems, we have three categories of spatial autocorrelation: “low,” “medium,” and “high,” which are shown in Table 2 for the 13 example problems [all 13 example forests are presented in Clark (1998)]. Note how the Geary Ratio decreases as we move from “low” to “medium” to “high” spatial autocorrelation. Also note that in the 50/50 case, the autocorrelation is undefined since all the stand values are equivalent. Note that the Geary Ratio is equal to 0.68 (“medium”) for the forest in Figure 9(a), but the Geary Ratio is equal to 0.13 (“high”) for the 70/30 forest in Figure 9(b).

In each example, the results were compared between INROADS and the other approaches by calculating the net present value of each solution. Each problem had a three-period planning horizon, where each period spanned 10 yr. Thirty stands were chosen for harvest in each period. Spatial constraints were included in the problems, with a maximum

Table 1. Stand value distribution (quantity and value per unit area) for five cases used in example problems.

| 50/50 Case | | 55/45 Case | | 60/40 Case | | 65/35 Case | | 70/30 Case | |
|------------|--------|------------|--------|------------|--------|------------|--------|------------|--------|
| Quantity | Scalar |
| 225 | 24 | 37 | 14 | 45 | 8 | 76 | 12 | 78 | 8 |
| | | 42 | 20 | 45 | 16 | 70 | 14 | 79 | 12 |
| | | 44 | 24 | 45 | 24 | 43 | 43 | 68 | 56 |
| | | 49 | 28 | 45 | 32 | 36 | 47 | | |
| | | 53 | 30 | 45 | 40 | | | | |

Table 2. Spatial dispersion (Geary Ratio) for 13 example problems.

| | 50/50 Case | 55/45 Case | 60/40 Case | 65/35 Case | 70/30 Case |
|----------------------------------|------------|------------|------------|------------|------------|
| No spatial autocorrelation | Und. | — | — | — | — |
| "Low" spatial autocorrelation | — | 1.104 | 1.007 | 1.096 | 1.391 |
| "Medium" spatial autocorrelation | — | 0.396 | 0.793 | 0.631 | 0.680 |
| "High" spatial autocorrelation | — | 0.259 | 0.146 | 0.329 | 0.130 |

opening area of five stands and the restriction of harvesting each stand no more than once during the planning horizon. A constant growth factor of 6% per year was employed. In order to calculate the present value for each problem, the time value of money was 8% per year. Revenues and costs for harvested stands were realized at the beginning of the period in which they were harvested.

We use several different terms to describe the results. In order to avoid confusion in terminology, we define *discounted net revenue* (DNR) as the discounted value of the net revenue (gross revenue – operational costs). Thus, the DNR does not consider any access road costs. *Profit* is defined as the net revenue minus the access road costs for a specific period, *discounted profit values* are the profit values discounted back to time zero, and *net present value* is the sum of the discounted profit values over all periods. The objective function value is equivalent to the net present value.

Example Results

In order to demonstrate the output of the different approaches, two period 1 solutions for the 70/30 case are shown in Figure 10. This example has "medium" dispersion, with a Geary Ratio of 0.680. A graphics driver was embedded in the code so that the results can be seen graphically. Figure 10(a) and 10(b) illustrate the prepositioned roads solution and the INROADS (Interchange 5) solution, respectively. As can be seen, the two solutions are different, in that the prepositioned roads are located such that each selected stand can be accessed by main roads that stem from the existing road, rather than a minimum spanning tree approach that INROADS utilizes. The objective function values of the individual example problems are shown in Figure 11. There were 13 problems and 7 solution procedures. Besides detailing the output from the tests, one can also see how the objective function value increases as the stand value becomes more spatially concentrated (i.e., from 50/50 to 70/30).

Aggregated Results

Since there are significant differences in the objective function values between the data sets, comparing aggregated

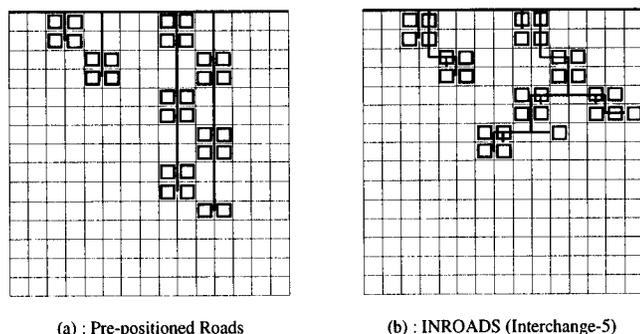


Figure 10. Period 1 solutions for two approaches.

results can be rendered meaningless if care is not taken in how that is done. The aggregated results we present were carefully considered. For each data set, we noted the maximum objective function value. For each approach, we then computed what percent each achieved of this maximum. We present the average and standard deviation of this data for the 13 example problems in Table 3.

The aggregated results show that the progressive approach performed the worst. This is somewhat intuitive since its primary stand selection criteria is very simplistic, focusing on minimizing access road cost. The seed stand/adjacency performed slightly better since it considers stand value to a greater extent. The prepositioned road approach was next best. This approach is the first of the comparison approaches to consider all unharvested stands at every stage. The Directed Search approach was next. The Directed Search and the prepositioned road approach have the same stand selection procedure; however, the approaches to access roading are different. The results from the Directed Search approach, where access roading is more flexible, are better than the prepositioned road approach by approximately 3%, which averages approximately \$250M for the 13 problems.

We only utilized one prepositioned road network for these problems. Note that the performance of the prepositioned road approach would improve if additional road networks were utilized. In essence, the general road approach of Directed Search and INROADS is equivalent to trying all possible prepositioned road networks. [A more complete description of how the prepositioned road approach is formulated for these example problems can be found in Clark (1998)].

As expected, INROADS with Interchange 10 and INROADS with Interchange 20 outperformed all other heuristics, with INROADS with Interchange 20 slightly outperforming all other INROADS approaches tested.

After discussing the runtimes of the heuristics, we will discuss the results of the INROADS heuristics in more detail.

Comparison of Runtimes

All heuristic approaches were coded in C. The example problems were run on a personal computer with a Pentium 120 MHz processor. The fastest approach was the prepositioned road approach, which had an average runtime of 54 sec. The main reason that this approach is so fast is simply that it only considers a very limited road network problem—that is, the roads are prepositioned and each stand could only be accessed by one road. All of the other approaches had multiple means of egress. Therefore, more nodes had to be considered.

INROADS was the slowest approach, with an average runtime of 30,777 sec when Interchange 20 was used. The Directed Search approach, which is the approach most closely

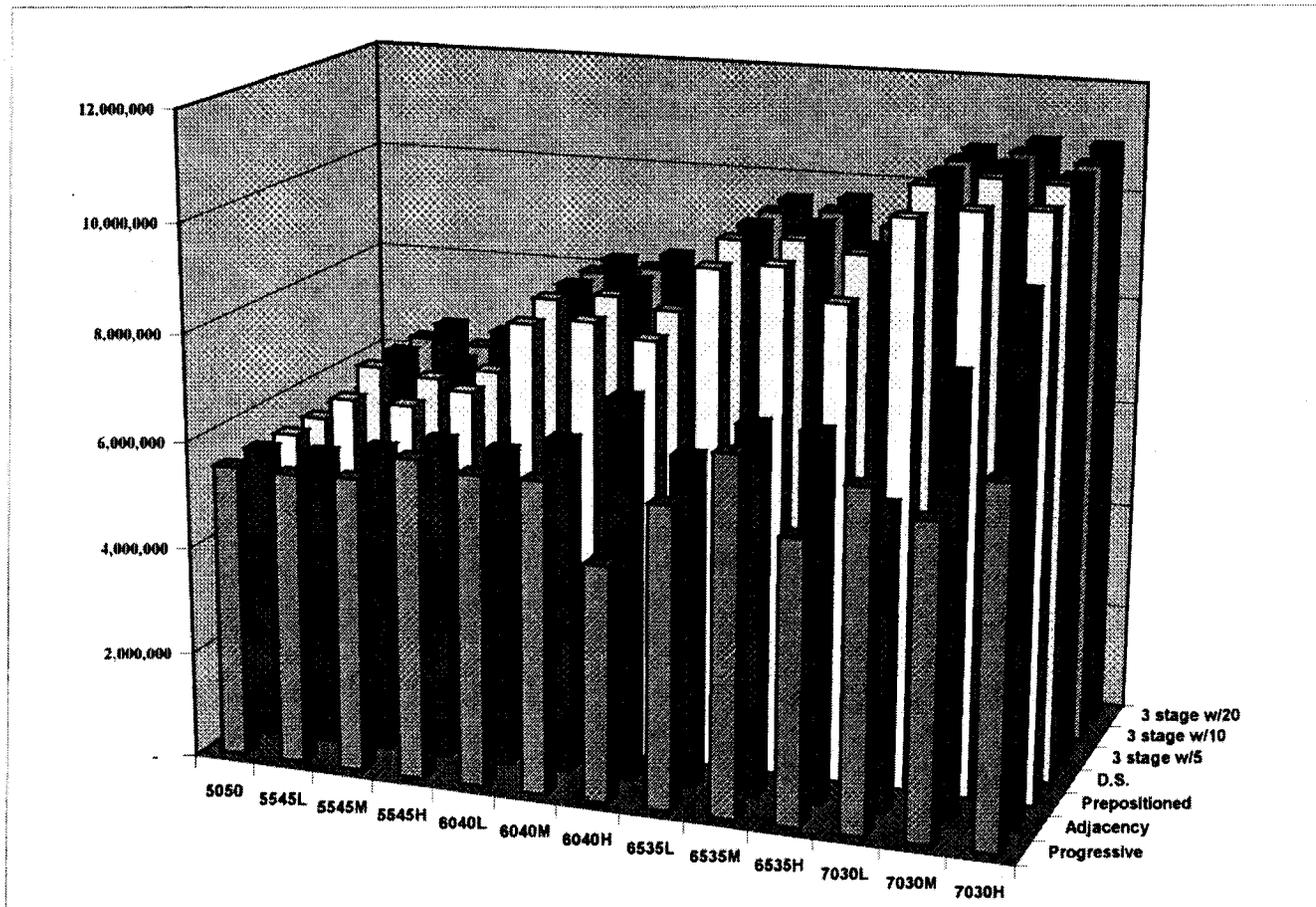


Figure 11. Results of 13 example problems for seven approaches.

related to INROADS, had an average runtime of 422 sec. Thus, the increased profit achieved via a more effective heuristic must be balanced with an increased runtime. The average runtime results (in seconds) are shown for the other approaches in Table 4. The runtime results, coupled with the performance results earlier, indicate that employing Directed Search during the preliminary stages of a study, and employing INROADS toward the end, would be the most productive approach.

We also examined the runtime per iteration. Doing so showed that as the level of interchange increases, the runtime per iteration increases as well. Clearly, this is due to the number of additional combinations in Stage 2. Also, as the periods progress, the runtime per period decreases. This decrease occurs for two reasons. First, it decreases since there are more roads (active nodes) in the forest; thus, it is easier to identify the closest active node when developing the access road network. Second, as the periods progress, there are

Table 3. Aggregated results over 13 example problems for seven approaches.

| Approach | Ave % max obtained | SD % max obtained |
|--------------------------|--------------------|-------------------|
| Progressive | 69.51 | 15.123 |
| Seed stand/adjacency | 78.08 | 12.666 |
| Prepositioned roads | 96.68 | 1.691 |
| Directed search | 99.63 | 0.403 |
| INROADS (Interchange-5) | 99.60 | 0.611 |
| INROADS (Interchange-10) | 99.75 | 0.550 |
| INROADS (Interchange-20) | 99.99 | 0.039 |

fewer eligible stands in the forest since we do not harvest stands more than one time. Thus, the runtime per iteration decreases as the periods progress. The conclusion to be drawn here is that larger levels of interchange will likely have a more significant effect on runtime than an increase in the number of periods.

Discussion of Results

In this discussion we will compare the Directed Search to that of the prepositioned road approach since they use different approaches to access roads, but identical stand selection procedures. Afterwards, we will compare the Directed Search approach to INROADS since they use different stand selection techniques, but have the same access road approach.

The only difference in the Directed Search and the prepositioned road approach is that the latter is constrained in its construction of access roading. Therefore, we can make a clear comparison of the two access road approaches. The results of the 13 example problems are shown in Tables 5 and 6. Directed Search always obtained

Table 4. Average approach runtime results (in seconds).

| | Runtimes |
|--------------------------|----------|
| Progressive | 1,100 |
| Seed stand/adjacency | 135 |
| Prepositioned roads | 54 |
| Directed search | 422 |
| INROADS (Interchange-5) | 5,777 |
| INROADS (Interchange-10) | 12,238 |
| INROADS (Interchange-20) | 30,777 |

Table 5. Directed search results (\$) for 13 example problems.

| | Period 1 | Period 2 | Period 3 | Total |
|--------|-----------|-----------|-----------|------------|
| 50/50 | 2,182,494 | 1,817,676 | 1,518,040 | 5,518,210 |
| 55/45L | 2,637,416 | 2,167,294 | 1,815,163 | 6,619,873 |
| 55/45M | 2,514,423 | 2,104,573 | 1,876,046 | 6,495,042 |
| 55/45H | 2,694,819 | 2,248,613 | 1,779,632 | 6,723,064 |
| 60/40L | 3,489,450 | 2,728,444 | 1,993,876 | 8,211,770 |
| 60/40M | 3,636,630 | 2,684,290 | 2,033,196 | 8,354,116 |
| 60/40H | 3,358,964 | 2,788,336 | 2,018,998 | 8,166,298 |
| 65/35L | 4,105,966 | 3,407,090 | 2,049,994 | 9,563,050 |
| 65/35M | 4,128,887 | 3,423,590 | 2,087,473 | 9,639,950 |
| 65/35H | 4,095,138 | 3,450,862 | 1,913,476 | 9,459,476 |
| 70/30L | 5,027,406 | 4,233,253 | 1,497,132 | 10,757,791 |
| 70/30M | 5,164,026 | 4,274,809 | 1,511,437 | 10,950,272 |
| 70/30H | 5,090,106 | 4,211,016 | 1,598,600 | 10,899,722 |
| | | | Average | 8,566,049 |

a better solution than the prepositioned road approach, with an average difference in objective function values of \$260,263.

We analyzed these results further and found that the major difference between the results of these two approaches is access road cost. On average, the prepositioned road approach had \$180,883 more in access roading cost than the Directed Search approach. This indicates that not only did the roading cost more, as we have shown, but the road costs had a negative impact on the profit of the stands that were selected by an average of \$79,380. By observing the period-by-period results of the discounted net revenue results, the majority of the difference is in Period 1.

The aggregate results indicate that the three-stage heuristic, INROADS, is only slightly better than that of our one-stage heuristic, Directed Search. It is interesting to note that the Directed Search results are better than the INROADS results with Interchange 5. One possible reason for this is due to the sequential fashion in which the multiperiod problem is solved. It is sequential in that a solution is found for the first period and then afterwards, the second and third periods are solved.

Of the 13 examples, INROADS with Interchange 5 performed better than the Directed Search in 9 of 13 cases. However, on average, the Directed Search performed better. The four examples where the Directed Search performed better than INROADS were in cases with "high" spatial autocorrelation (low Geary Ratios). In all four of these cases, INROADS performed better in the first two periods, but much worse in the last period.

INROADS was developed to solve the problem sequentially by period rather than across the planning horizon. Thus, it seeks to select the best stands regardless of the second-best stand's value. Clearly, INROADS performed better than the Directed Search by selecting more of the highest valued stands early (i.e., in the first two periods). Conversely, the Directed Search selected more of the lesser valued stands early, but still did better over three periods. This result is an indication of the shortcoming of INROADS's approach in solving the problems sequentially rather than across the entire planning horizon.

Another reason for INROADS's performance might be due to the problem constraints. That is, most or all of the stands may

Table 6. Prepositioned road results (\$) for 13 example problems.

| | Period 1 | Period 2 | Period 3 | Total |
|--------|-----------|-----------|-----------|------------|
| 50/50 | 2,168,634 | 1,816,759 | 1,519,173 | 5,504,567 |
| 55/45L | 2,436,561 | 2,052,207 | 1,788,864 | 6,277,632 |
| 55/45M | 2,440,057 | 2,101,481 | 1,727,336 | 6,268,874 |
| 55/45H | 2,623,467 | 2,227,530 | 1,785,718 | 6,636,715 |
| 60/40L | 3,198,406 | 2,634,942 | 2,174,664 | 8,008,012 |
| 60/40M | 3,439,841 | 2,597,050 | 2,100,556 | 8,137,447 |
| 60/40H | 3,003,061 | 2,793,839 | 2,087,192 | 7,884,092 |
| 65/35L | 3,941,179 | 3,260,270 | 2,089,959 | 9,291,408 |
| 65/35M | 4,031,599 | 3,304,958 | 2,080,709 | 9,417,266 |
| 65/35H | 3,774,948 | 3,279,902 | 1,780,810 | 8,835,660 |
| 70/30L | 4,860,426 | 4,085,364 | 1,475,038 | 10,420,828 |
| 70/30M | 5,063,706 | 4,152,586 | 1,396,231 | 10,612,523 |
| 70/30H | 4,906,626 | 4,161,516 | 1,612,055 | 10,680,197 |
| | | | Average | 8,305,786 |

be ineligible for interchange in later periods. The reason that this may affect INROADS more than the Directed Search is due to the spatial dispersion of the remaining stands in later periods. Since the Directed Search approach is more "greedy," the remaining stands are more dispersed, which will lessen the impact of spatial constraints later in the solution process.

A similar observation may be made in the cases of "medium" spatial autocorrelation for INROADS with the Interchange 5 and the Directed Search. The third period results are better for INROADS in the lower concentrated valued stands (55/45 and 60/40), but in the higher concentrated valued stands (65/35 and 70/30), the Directed Search performs better. As we move from the 55/45 case to the 70/30 case, the Directed Search approach performs better by a wider margin.

INROADS with Interchange 10 performed better, on average, than the Directed Search and INROADS with Interchange 5. The one case that the Directed Search did better was the 60/40 case with a "high" dispersion factor. Again, the results are that INROADS performed better than the Directed Search in the first two periods and much worse than the Directed Search in the last period.

The results for INROADS with Interchange 20 are better than all other example results (see Table 7) except for 55/45H, which again indicates suboptimizing by period can lead to a worse solution. Overall, these results suggests that as the level of interchange increases beyond 10, INROADS is very likely to perform better than the Directed Search procedure.

Table 7. Three-stage heuristic (Interchange 20) results (\$) for 13 example problems.

| | Period 1 | Period 2 | Period 3 | Total |
|--------|-----------|-----------|-----------|------------|
| 50/50 | 2,185,134 | 1,820,425 | 1,519,032 | 5,524,591 |
| 55/45L | 2,637,416 | 2,170,044 | 1,827,954 | 6,635,414 |
| 55/45M | 2,525,643 | 2,187,805 | 1,863,069 | 6,576,517 |
| 55/45H | 2,701,794 | 2,257,551 | 1,754,055 | 6,713,400 |
| 60/40L | 3,489,450 | 2,732,417 | 2,000,249 | 8,222,116 |
| 60/40M | 3,637,950 | 2,716,528 | 2,026,026 | 8,380,504 |
| 60/40H | 3,607,590 | 2,749,833 | 1,908,050 | 8,265,473 |
| 65/35L | 4,126,426 | 3,423,362 | 2,058,589 | 9,608,377 |
| 65/35M | 4,189,964 | 3,402,199 | 2,095,121 | 9,687,284 |
| 65/35H | 4,096,458 | 3,454,834 | 1,914,184 | 9,465,476 |
| 70/30L | 5,032,686 | 4,227,447 | 1,506,621 | 10,766,754 |
| 70/30M | 5,171,946 | 4,284,892 | 1,504,780 | 10,961,618 |
| 70/30H | 5,109,906 | 4,325,225 | 1,504,073 | 10,939,204 |
| | | | Average | 8,595,902 |

As the number of interchange candidates increases, the number of iterations also increases. For example, in the Interchange 5 case there were 91 Stage 2 iterations for the 13 example problems. In the Interchange 10 and Interchange 20 cases there were 115 and 127 Stage 2 iterations, respectively. This is intuitive since there are more solution combinations when the number of interchange candidates increases.

It seems as though the Directed Search approach might be the most attractive heuristic among the different approaches that were tested (remember that the Directed Search is equivalent to INROADS with Interchange 0). The results of the Directed Search, on its own, performs well for the problems considered, in terms of average objective function value, coupled with a moderate runtime. This further emphasizes our earlier comment that there is a "cost/benefit" tradeoff to be considered when utilizing the second stage of INROADS, since INROADS is likely to obtain a more beneficial result, but at the cost of increased runtime.

Conclusions

In this article we presented a new model for the harvest scheduling problem. Our approach is unique in two respects. First, our approach integrates roading decisions into stand selection by using a minimum spanning tree approach, rather than prepositioned roads, which is commonly used in other approaches. The node-based minimum spanning tree approach allows flexibility in roading, which in turn, reduces road costs. The shortcoming of our roading approach is that it assumes that an access road can be built anywhere in the forest, which may or may not be a feasible location. We also integrate the roading at the operational and stand levels. Roading is a significant component within and between stands. Thus, by using the roads at the stand level to act as access roads to other stands allowed the two types of roading to be integrated. The second contribution is in the heuristic we developed, INROADS. INROADS is unique due to the three-stage heuristic we developed. INROADS uses a greedy approach in the construction stage and pairwise interchange in the second two stages.

In order to evaluate the effectiveness of INROADS for the harvest scheduling problem, we developed three comparison heuristics, which were based on the work of others. Our intent was not to duplicate the previous work, but to construct approaches that maintained the spirit of the heuristics. This allowed us to make some clear comparisons to INROADS. Thirteen example problems were developed in order to make the comparison. The example problems were developed by changing the individual stand values and the spatial dispersion of the timber volumes, while holding the overall value of the forest constant. INROADS performed better, with respect to the objective function value, than the other three heuristics in all 13 examples. However, the other heuristics performed better in terms of runtime. Based on our results, the recommended strategy appears to be using INROADS with only

the construction stage, which we called Directed Search, in the beginning of a study when numerous runs are usually required, and adding the improvement stages in the latter phases of a study.

In addition, we compared the performance of INROADS to the Directed Search heuristic. We limited the number of pairwise interchange candidates during Stage 2 of INROADS to three levels (namely 5, 10, and 20). We found that as the number of interchange candidates increased, so did the objective function value with one interesting exception. In examples with "high" dispersion factors, INROADS did better than the Directed Search in the first two periods, but not as well in the third period. This indicates that future research is needed for procedures that do not solve the problem period by period, but rather look at all periods simultaneously.

Overall, INROADS performed better than the other heuristics on these test problems. However, there exist avenues for future research. In addition to solving all periods of the problem simultaneously, research considering alternate stand selection search techniques should be considered as in (Lockwood and Moore 1993) and (Murray and Church 1995), where interchange, SA, and Tabu searches have been examined. Integrating these searches with road network development will likely improve the results of INROADS. Still, other roading techniques that consider the inherent difference in terrain, such as in Dean (1997), should be explored. Variations in maximum opening sizes and exclusionary periods, along with multiple prescriptions, should also be explored.

Literature Cited

- BULLARD, H.S., H.D. SHERALI, AND W.D. KLEMPERER. 1985. Estimating optimal thinning and rotation for mixed-species timber stands using a random search algorithm. *For. Sci.* 31:303-315.
- CLARK, M.M., R.D. MELLER, T.P. McDONALD, AND C.C. TING. 1999. A new harvest operation cost model to evaluate forest harvest alternatives. *Annals of Operations Research*, forthcoming.
- CLEMENTS, S.E., P.L. DALLAIN, AND M.S. JAMNICK. 1990. An operational spatially constrained harvest scheduling model. *Can. J. For. Res.* 20:1438-1447.
- DAHLIN, B., AND O. SALLNAS. 1993. Harvest scheduling under adjacency constraints—a case study from the Swedish sub-alpine region. *Scan. J. For. Res.* 8:281-290.
- DEAN, D.J. 1997. Finding optimal routes for networks of harvest site access roads using gis-based techniques. *Can. J. For. Res.* 27:11-22.
- CLARK, M.M. 1998. The forest harvesting problem: Integrating operational and tactical planning. Ph.D. Diss., Auburn University.
- GRIFFITH, D.A., AND C.G. AMRHEIM. 1991. *Statistical analysis for geographers*. Prentice-Hall, Englewood Cliffs, NJ. 613 p.
- IVERSON, D.C., AND R.M. ALSTON. 1986. The genesis of FORPLAN: A historical and analytical review of forest service planning models. USDA For. Serv. Tech. Rep. INT-214.
- JAMNICK, M.S., AND K.R. WALTERS. 1993. Spatial and temporal allocation of stratum-based harvest schedules. *Can. J. For. Res.* 23:402-413.
- JONES, J.G., M.L. MEACHAM, A. WEINTRAUB, AND A. MAGENDZO. 1991. A heuristic process for solving large-scale, mixed-integer mathematical programming models for site specific timber harvest and transportation problems. P. 31-34 in *Proc. of the 1991 Symp. on Systems analysis in forest resources*. USDA For. Serv., Southeastern For. Exp. Sta. Asheville, NC.

- KIRBY, M.W., P. WONG, W.A. HAGER, AND M.E. HUDDLESTON. 1980. A guide to the integrated resource planning model. USDA For. Serv., Berkeley, CA. 211 p.
- KIRKPATRICK, S., C.D. GELATT, AND M.P. VECCHI. 1983. Optimizing by simulated annealing. *Science* 220:671-680.
- LIU, K., AND J. SESSIONS. 1993. Preliminary planning of road systems using digital terrain models. *J. For. Eng.* 4:27-32.
- LOCKWOOD, C., AND T. MOORE. 1993. Harvest scheduling with spatial constraints: A simulated annealing approach. *Can. J. For. Res.* 23:468-478.
- MURRAY, A.T. 1998. Route planning for harvest site access. *Can. J. for. Res.* 28:1084-1087.
- MURRAY, A.T. 1999. Spatial restrictions in harvest scheduling. *For. Sci.* 45(1):45-52.
- Murray, A.T. and R. Church. 1995. Heuristic solution approaches to operational forest planning problems. *OR Spektrum.* 17:193-203.
- NAHMIAS, S. 1997. *Production and operations management.* Ed. 3. Irwin, Chicago, IL. 858 p.
- NELSON, J., AND J.D. BRODIE. 1990. Comparison of a random search algorithm and mixed integer programming for solving area-based forest plans. *Can. J. For. Res.* 20:934-942.
- NELSON, J., J.D. BRODIE, AND J. SESSIONS. 1991. Integrating short-term, area-based logging plans with long-term harvest schedules. *For. Sci.* 37:101-122.
- NELSON, J.D., AND S.T. FINN. 1991. The influence of cut-block size and adjacency rules on harvest levels and road networks. *Can. J. For. Res.* 21:595-600.
- O'HARA, A.J., B.H. FAALAND, AND B.B. BARE. 1989. Spatially constrained timber harvest scheduling. *Can. J. For. Res.* 19:715-724.
- SEDGEWICK, R. 1988. *Algorithms.* Addison Wesley, New York. 583 p.
- SNYDER, S., AND C. REVELLE. 1996. The grid packing problem: selecting a harvesting pattern in an area with forbidden regions. *For. Sci.* 42(1):27-34.
- WALTERS, K.R. 1991. Spatial and temporal allocation of strata-based harvest schedules. M.Sc. thesis, Univ. of New Brunswick, Fredericton.

APPENDIX: Example Stand Parameters

The following parameters were used throughout to calculate the optimal harvest patterns for each stand in the forest.

| | | |
|----------|--------------------------------|------------------------|
| <i>A</i> | Total area of each stand | 161,869 m ² |
| <i>c</i> | Volume capacity of the skidder | 5.0 m ³ |
| <i>s</i> | Variable skidding costs | \$.0134/m |
| <i>x</i> | Fixed skidding cost per turn | \$2 |
| <i>r</i> | Road cost per unit distance | \$6.56/m |
| <i>l</i> | Fixed cost per landing | \$300 |