A dynamic programming algorithm for optimization of uneven-aged forest stands

DAVID J. ANDERSON
Georgia-Pacific Corporation, 900 Southwest Fifth Avenue, Portland, OR 97204-1283, U.S.A.
AND
B. BRUCE BARE[1]
College of Forest Resources, University of Washington, Seattle, WA 98195, U.S.A.
Received October 6, 1993
Accepted March 29, 1994

ANDERSON, D.J., and BARE, B.B. 1994. A dynamic programming algorithm for optimization of uneven-aged forest stands. Can. J. For. Res. **24**: 1758–1765.

A deterministic dynamic programming formulation of the transition uneven-aged stand management problem is presented. Using a previously published northern hardwoods growth model, a forward recursive, discrete, two-state problem that maximizes the net present value of harvested trees at each stage is developed. State variables represent the total number of trees and the total basal area per acre. A neighborhood storage concept previously published is used to reduce the number of states considered at each stage. Two harvest allocation rules are used to assign the harvested basal area to individual diameter classes. Terminal end point conditions and stage to stage sustainability are not required. Results from four base runs of the model are presented and compared with previously published results. Each run produces significantly different optimal paths, with one showing a higher net present value than any previously published. Sensitivity runs illustrate the impact of changes in interest rates, width of neighborhood storage class, and initial conditions. Dynamic programming offers promise for analyzing uneven-aged stand management problems.

ANDERSON, D.J., et BARE, B.B. 1994. A dynamic programming algorithm for optimization of uneven-aged forest stands. Can. J. For. Res. **24** : 1758–1765.

Une formulation en programmation dynamique déterministe est présentée ici pour l'aménagement des peuplements inéquiennes de transition. Le problème est exprimé sous une forme récursive discrète à deux variables, qui maximise la valeur nette actuelle pour les arbres récoltés à chaque étape, en utilisant un modèle de croissance des feuillus nordiques déjà publié. Celles-ci sont le nombre total d'arbres et la surface terrière totale à l'acre. Un concept d'emmagasinage de l'environnement déjà publié est utilisé pour réduire le nombre de possibilités à chaque étape. Deux règles de répartition de la coupe sont utilisées pour déterminer la surface terrière à récolter dans chaque classe de diamètre. L'état final du peuplement et sa stabilité d'une étape à l'autre n'ont pas à être pris en compte. Les résultats d'essais du modèle pour quatre cas de base sont présentés et comparés avec des résultats déjà publiés. Chaque essai donne des solutions très différentes, dont une montrant une valeur nette actuelle plus élevée que celles déjà publiées. Une analyse de sensibilité illustre l'impact de changements dans les taux d'intérêt, la classe de dimension pour l'emmagasinage de l'environnement et l'état initial du peuplement. La programmation dynamique s'annonce prometteuse pour l'analyse de problèmes d'aménagement des peuplements inéquiennes.

[Traduit par la rédaction]

Introduction

Since publication of Adams and Ek's (1974) pioneering work on optimizing decisions in uneven-aged management, forest researchers have developed a variety of optimization techniques in an effort to answer questions posed by forest managers. These include matrix-based approaches, linear and nonlinear programming models, and optimal control formulations utilizing direct search and gradient-based algorithms. Reviews of these works can be found in Gove and Fairweather (1992), Haight (1990*a*, 1990*b*), Haight and Getz (1987), Haight and Monserud (1990*a*, 1990*b*), Bare and Opalach (1987), Buongiorno and Michie (1980), and Michie (1985). While most of these optimization models assume a deterministic decision environment, adaptive strategies for dealing with stochastic events are also reported.

An optimization technique that has been widely used in even-aged stand-level optimization is dynamic programming (Brodie and Haight 1985; Riitters et al. 1982; Haight et al. 1985*a*). However, only two previous reports of its use in an uneven-aged decision environment were found (Hool 1966; Hotvedt and Ward 1990). Although mentioned as a possible technique for optimizing uneven-aged management decisions by Adams and Ek (1974), dynamic programming generally has not been implemented for computational reasons, primarily the large state space that results. Given recent advances in computer technology and continued interest in uneven-aged optimization, this study was initiated to evaluate the feasibility of using dynamic programming to optimize uneven-aged management decisions.

Both static and dynamic problems have been addressed in uneven-aged stand optimization. Static models determine the optimal sustainable steady-state diameter distribution (or level of reserve growing stock), perhaps by species, that perpetuates a regular periodic flow of harvest revenue or volume over time. However, static models do not indicate how (or if) a stand, not in this condition, is to be managed over time to achieve the steady-state condition. Dynamic models search for the optimal harvest trajectory starting with a current stand and show how it should be managed over time to optimize a stated objective. Such models are run with or without either equilibrium or fixed endpoint constraints.

An uneven-aged stand in a steady-state condition is defined as a stand with a diameter class structure such that at every

[1]Author to whom all correspondence should be addressed.

Printed in Canada / Imprimé au Canada

cutting cycle a constant harvest can be removed in perpetuity while maintaining an invariant residual stand structure. Adams (1976) termed one such fixed end point condition investment efficient, if the residual steady-state stand structure produced the maximum land expectation value. Other steady-state stand structures result if the optimization is carried out to maximize volume production or forest rent. While determining the optimal steady-state diameter distribution, residual stand structures may be constrained to take on a geometric progression of trees in successive diameter classes. Such stands are termed "balanced" and follow the constant q-ratio concept of de Liocourt (1989). However, most optimization studies show that these constraints are not necessary to ensure sustainability and may lead to reductions in objective-function performance relative to models without such constraints.

Haight et al. (1985*b*) use an optimal-control formulation to find uneven-aged management regimes that maximize net present value (NPV) of harvests over a 150-year time horizon. Their method solves for both transition and ending stand structures, where the latter may be constrained to converge to either a fixed end point or equilibrium condition. Continuing this comparison of dynamic and static solutions, Haight (1985, 1987) shows that dynamically determined solutions yield NPVs greater than statically determined solutions and that an optimal transition regime does not produce a steady-state stand that maximizes land expectation value.

Hool (1966) was one of the first to adapt dynamic programming to uneven-aged management. His model incorporates Markov chains, thus allowing treatment of the probability of state transformation in the scheduling of stand treatments. Since Hool's study, the bulk of dynamic programming applications in forestry have concentrated on even-aged management. Brodie and Kao (1979), Brodie and Haight (1985), Valsta (1990, 1993), Filius and Dul (1992), Arthuad and Klemperer (1988), and Tait (1986) review some of these works.

Hotvedt and Ward (1990) use dynamic programming to optimize uneven-aged management for loblolly pine (*Pinus taeda* L.) – shortleaf pine (*Pinus echinata* Mill.) stands in the South. They use state variables of total stand residual basal area, sawtimber basal area, and the elapsed time associated with transition from one state to another. Their dynamic model determines the optimal transition harvest strategy that optimizes one of four objective functions. Further, it adopts the concept of a steady-state equilibrium endpoint target and incorporates a constant q-ratio, which previously has been shown to yield lower objective function values (Adams 1976; Haight 1985).

Model formulation

The objective of this paper is to develop a deterministic dynamic programming formulation of the transition uneven-aged problem in the absence of end-point conditions. Dynamic programming is well suited to such problems because they involve a series of sequential, yet independent, decisions (Dreyfus and Law 1977). While the description of the system becomes richer as the number of state variables used increases, the curse of dimensionality (effectuated by a large computational burden) generally limits one to the use of three or four state variables. The principle of optimality (Bellman 1957) and the concept of a recurrence relation are used to select the optimum set of states at each stage. A transformation function is used to link each stage with its successor stage and is a function of the decision taken and the state of the system.[2]

The uneven-aged transition problem is formulated as a forward recursion, discrete, two-state, dynamic-programming problem that maximizes the NPV of harvested trees at each stage. No terminal end point conditions are specified nor is stage to stage sustainability of harvest required. Forward recursion is used because the structure of the network (formed by connecting possible states arising at one stage to those possible at the next) is not known prior to the start of the optimization. State variables represent total trees per acre (TTPA; 1 acre = 0.405 ha) and total basal area per acre (TBAA). This choice was influenced by both the computational burden described above as well as the growth model that uses the same two variables.

An alternative was to use the number of trees per acre in each diameter class. Because this would have required eight state variables and none of the forestry dynamic programming studies reviewed utilized more than four state variables, this was not deemed a viable alternative. However, because the number of trees by diameter class is used by the growth model, a method of allocating the harvested trees back to diameter classes is required. Two harvest allocation methods are used for this purpose and are described below. The length of time between stages was set at 5 years, as this was the time period utilized by the growth model, but any integer multiple of five can be used.

Fundamental to the model is the neighborhood storage concept introduced by Brodie and Kao (1979). A neighborhood storage class divides a continuous state space into discrete regions of a specified width for each state variable. A region can be viewed as a unique point in state space surrounded by a tolerance interval for each state variable. All points within the area enclosed by the tolerance interval are assumed to be equally representative of that region. Employment of the neighborhood storage class concept leads to a reduction in the number of possible states considered at each stage.

Figure 1 portrays a simplified flowchart outlining the basic structure of the UNEVENDP algorithm developed to solve the uneven-aged dynamic programming problem. Fundamentally the algorithm consists of three nested loops. The outer loop controls the number of stages; the middle loop advances through the number of nodes at each stage; and the inner loop applies the selected harvest allocation method at each node. A node represents an exact value of the two state variables. After harvest decisions are determined for a node, the harvest is removed, thus producing a residual stand. This residual stand is then transformed via the growth model into a stand structure that is stored in a temporary file. Once all nodes have been harvested and grown, the temporary file is sorted by NPV and classified into neighborhood storage classes for that stage. Neighborhood classes are then optimized using the dynamic programming recursion. The stand structure possessing the maximum NPV is chosen as the exact stand structure representing each node. Optimal stand structures for each of the neighborhood storage classes become the set of nodes used at the next stage. This process

[2]Dykstra (1984), Kennedy (1986), and Hann and Brodie (1980) are excellent references for readers wishing to further review the basics of dynamic programming.

is repeated until the last stage has been reached. When this occurs, the node at that stage having the largest accumulated recursion function value is declared the optimal end point, and the series of harvest decisions and stand structures are retrieved to form the optimal path.

Incorporated directly within the algorithm and used as the stage transformation function is the growth model developed by Ek (1974) and modified by Adams and Ek (1974). This northern hardwood whole stand diameter class growth model uses 2-in. (1 in. = 2.54 cm) diameter classes starting with the 6-in. class and consists of ingrowth, up-growth, and mortality equations for each diameter class. The growth model has been extensively used in optimization studies of uneven-aged management and is not repeated here (Adams and Ek 1974; Adams 1976; Martin 1982; Bare and Opalach 1988; Haight et al. 1985*b*; Haight 1985).

Dynamic programming model

The variables and the formal model are defined as follows.

Variable	Description	Range
t	Stage index	1, 2, ..., T
j	Node index	1, 2, ..., J
d	Diameter class index	1, 2, ..., N
h	Harvest allocation method index	1, 2, ..., H
$XB_{d,h}(t, j)$	Initial number of trees per acre at stage t, node j, harvest option h, diameter class d	
$H_{d,h}(t, j)$	Number of trees per acre to harvest at stage t, node j, harvest option h, diameter class d	
P_d	Price per tree in diameter class d	
G	Growth model	
V	Values of harvest decision vector	
R	Recurrence relation	

(1) Harvest-decision constraint

$$H_{d,h}(t, j) \le XB_{d,h}(t, j)$$

At stage t, the harvest from diameter class d under harvest allocation method h must be less than or equal to the initial number of trees (XB) in diameter class d. Node j refers to combinations of stand structures as represented by TTPA and TBAA.

(2) Residual stand structure (XR):

$$XR_{d,h}(t, j) = XB_{d,h}(t, j) - H_{d,h}(t, j)$$

The residual number of trees at stage t in diameter class d under harvest allocation method h is equal to the initial number of trees in diameter class d less the harvest from diameter class d under harvest allocation method h.

(3) Start of stage stand structures:

$$XB_{d,h}(t + 1, j) = G[XR_{d,h}(t, j)]$$

The initial number of trees in diameter class d at stage $t + 1$ under harvest allocation method h is equal to the residual number of trees in diameter class d at stage t after being grown for one stage (G).

(4) Value (V) of the harvest-decision vector:

$$V_h(t, j) = \sum_{d=1}^{N} H_{d,h}(t, j)P_d$$

The return from a harvest at stage t is equal to the number of trees harvested from diameter class d under harvest allocation method h times the price per tree (P) in diameter class d.

(5) Net present value of harvest-decision vector (f_h):

$$f_h(t, j) = \frac{V_h(t, j)}{(1 + r)^{(t-1)5}}$$

The NPV of a harvest at stage t under harvest allocation method h equals the value of the harvest divided by the present value factor, where r is the annual real discount rate.

(6) Dynamic-programming recurrence relation ($R_{t,h}$):

$$R(t, j) = \underset{h_k \in H_j}{\text{Max}} \, [R(t - 1, j) + f_h(t, j)]$$

The cumulative NPV at stage t is equal to the maximum of the sum of the cumulative NPV at stage $t - 1$ plus the NPV of the stage t harvest. Maximization takes place over H_j (the set of nodes at stage t − 1 that can be harvested and grown to node j in period t).

Harvest-decision vector determination

Two methods are provided for determining harvest-decision vectors. Both methods specify a target amount of basal area per acre to remove from a stand, with the level given as a percentage of the TBAA present in the stand at the start of a stage. Differences between the two methods lie in how this total basal area is distributed back to each diameter class. Constant between both methods is that all trees in the 20-in. diameter class are always removed, thus a priori limiting the stand to a maximum diameter of 20 in. This was done to avoid extrapolating beyond the data used to fit the growth model.

New to this study is a p-ratio, which synthesizes ideas from even-aged and uneven-aged management. The even-aged analog of the p-ratio is the d/D ratio, which compares the average stand diameter before thinning (d) and average stand diameter after thinning (D). For example, a d/D ratio <1 indicates that, during the thinning, the trees harvested were smaller than the average diameter before thinning occurred. Similarly, a p-patio <1 indicates that the allocation of harvested trees should start from the smallest (6-in.) diameter class and a p-ratio >1 means that the harvest allocation should start from the largest (18-in.) diameter class. The uneven-aged analog of the p-ratio is the q-ratio of de Liocourt (1898). Whereas de Liocourt's q-ratio specifies a constant ratio between the number of trees in successive diameter classes of the residual stand structure, the p-ratio specifies a ratio between trees harvested in successive diameter classes.

Harvest method 1 determines the number of trees to harvest in the starting diameter class by removing the number of trees from that class that corresponds to the total basal area percentage target (i.e., if the total basal area removal target is 10%, then 10% of the trees in the starting diameter class are removed). For the next diameter class, the p-ratio is multiplied times the number removed from the previous class to determine the number of trees to remove. The associated basal area is computed and added to that previously removed. Cutting proceeds successively through each diameter class, summing the total basal area harvested for all diameter classes, including that of the 20-in. diameter class. Cessation of harvest occurs when the target basal area has been exceeded. If the target has not been reached in one pass through the stand, then the process begins again from

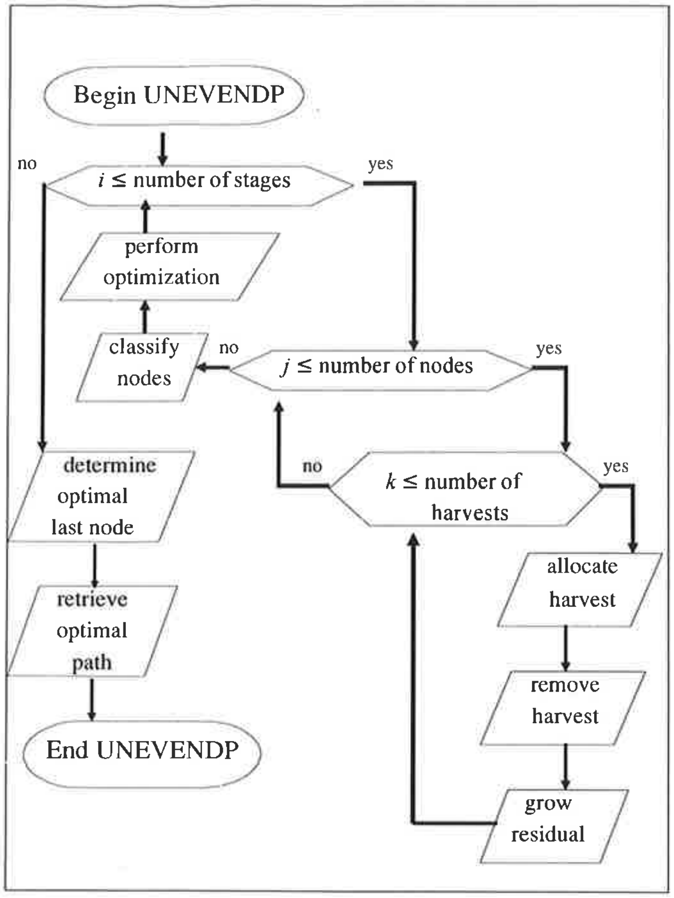

FIG. 1. Flowchart of uneven-aged dynamic programming algorithm.

the starting diameter class using the residual stand structure for the initial estimate.

In method 1, the ratio between trees harvested in successive diameter classes uses the cumulative number of trees removed in the previous diameter class and not the number of trees removed in the current iteration of the harvest method. Thus, the final ratios between harvested trees will not necessarily equal the p-ratio for all diameter classes since the number of trees harvested in each iteration are accumulated. Another factor leading to an unbalanced p-ratio is not completing a harvest allocation iteration due to the total basal area target being exceeded. Also, by calculating the number of trees to harvest using percent basal area removal targets and not the number of trees, it is likely that unbalanced p-ratios will result.

Harvest method 2 uses the p-ratio only to indicate initiation of harvest at the 6- or 18-in. diameter class. A fixed percentage, equal to the specified total harvest basal area percentage, is always harvested from each diameter class. In the UNEVENDP algorithm, a p-ratio <1 and (or) >1 is used as an indicator variable to cover these two cases.

Harvest decision parameters selected for any given dynamic programming solution are combinations of particular values of the p-ratios, total basal area percentage targets, the neighborhood storage class width-tolerance parameters, the discount rate, and the number of stages. Thus, there are potentially an infinite number of harvest decision parameter sets to examine. As shown later in Table 2, to facilitate the examination of a larger number of these possible combinations, the algorithm allows a large number of combinations to be examined at each stage.

TABLE 1. Initial stand structure and price per tree for illustrative optimizations

Diameter class (in.)	Trees/acre	Price/tree ($)
6	69.9	0.04
8	49.3	0.10
10	37.5	0.18
12	29.7	0.26
14	24.1	0.97
16	10.7	2.80
18	2.1	5.04
20	0.3	7.90

NOTE: The stand structure was based on a projection horizon of 150 years, a site index of 60, and a real discount rate of 4%.

Initial input parameters

Diameter classes are defined to be in 2-in. intervals starting from 6 in. to a maximum of 20 in. Initial values for trees per acre and price per tree, for each diameter class are shown in Table 1. The stand structure used as the initial condition for all base runs is the investment efficient stand structure first shown by Adams (1976) and later used by Haight et al. (1985*b*) and Haight (1985). The TBAA for this initial stand is 119.7 ft^2 (1 ft^2 = 0.093m^2). Prices are assumed to be constant over the projection horizon, which consists of 150 years split into thirty 5-year stages. Lastly, a constant 4% real discount rate is used in all runs. Stand structures modelled are assumed to be on high site land that has a site index of 60. By utilizing these parameters, comparisons between the optimal harvest structures produced by the dynamic programming algorithm and previously published studies (i.e., Haight 1985), can be made.

Model experimentation

Four different harvest-decision parameter sets were implemented in this study: two using harvest method 1 and two using harvest method 2. These are shown in Table 2. For any given dynamic programming analysis, multiple sets of p-values and harvest basal area percentage targets may be used. Thus, for harvest method 1 parameter set A, four different p-values and eight different basal area harvest percentage levels are examined. The main difference between decision parameters sets A and B lies in the width of the neighborhood classes used to classify nodes. The only difference between parameter sets C and D is in the definition of the diameter class where harvest may be initiated.

Experimentation results

Shown in Table 2 are the execution times and NPVs of the four base runs. Three things are noteworthy. First, harvest method 1 yields higher objective function values than does harvest method 2, although this gain comes at the expense of increased computational times. Second, parameter set A yields the best solution owing to the greater flexibility allowed in harvest levels examined and in the finer definition of neighborhood storage classes. Lastly, the NPV differences between parameter sets C and D are insignificant even though parameter set C initiates harvest in the 6-inch diameter class and parameter set D in either the 6- or 18-in. classes. To save space, only detailed results for

TABLE 2. Harvest decision parameters for four base runs

	Harvest allocation method 1		Harvest allocation method 2	
Harvest allocation method parameter	Harvest decision parameters A	Harvest decision parameters B	Harvest decision parameters C	Harvest decision parameters D
Interest rate (%)	4	4	4	4
TTPA class width	5	10	10	10
TBAA class width (ft^2)	5	10	10	10
P-values				
P(1)	10.00	5.00	<1	<1
P(2)	5.00	2.50		>1
P(3)	0.20	0.50		
P(4)	0.10	0.25		
Harvest basal area percentage targets				
BA(1)	0	0	0	0
BA(2)	5	5	5	5
BA(3)	10	10	10	10
BA(4)	15	15	15	15
BA(5)	20	20	19	19
BA(6)	30			
BA(7)	40			
BA(8)	50			
Execution time (h:min:s)	3:46:37	0:22:45	0:00:31	0:03:00
NPV/acre ($)	178.83	164.99	156.15	156.66

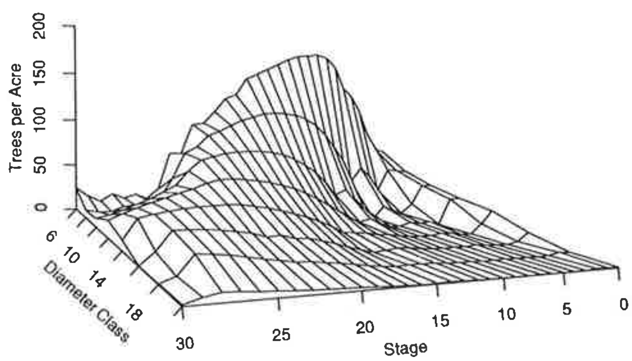

FIG. 2. Start of stage stand structures for harvest method 1 parameter set A.

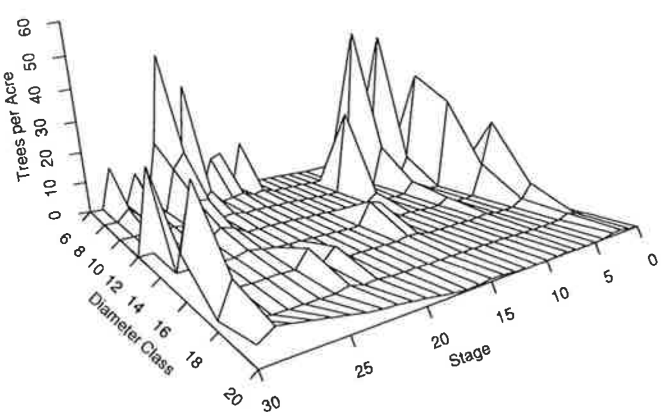

FIG. 3. Harvest control vectors for harvest method 1 parameter set A.

parameter sets A and C will be shown. However, comparisons among all four runs will be drawn as appropriate.[3] As previously discussed, both model formulation and implementation are based on the stand structure that has been harvested at stage t and grown to stage $t + 1$. Thus, these are the stand structures presented in all of the following figures and are referred to as start of stage stand structures ($XB_{d,h}(t, j)$).

Harvest allocation method 1

Figures 2 and 3 show the number of trees per acre by diameter class for the start of stage stand structures and the harvest decision vectors over the 150-year projection horizon for parameter set A. This set of harvest-decision parameters produced the optimal path with the highest NPV ($178.83/acre) of the four harvest decision parameter sets used. The NPV is also greater than previously published results of $171.29/acre (Haight 1985). However, this run took the longest time to execute, largely due to the smaller neighborhood class widths used.[4] A general optimal harvest strategy derived from these results is to (*i*) harvest all large, high value trees within the first 10 years, (*ii*) allow the unused growing space to be occupied by a large number of small ingrowth trees, (*iii*) reduce harvest levels to allow this ingrowth to move into the larger diameter classes, and (*iv*) begin harvesting across all diameter classes later in the projection horizon.

Harvest method 1 parameter set B (not shown), yields the second best NPV among the four base runs. Here, harvest

[3]The full set of results can be found in Anderson (1992).

[4]All computational times are based on an i386/33 processor.

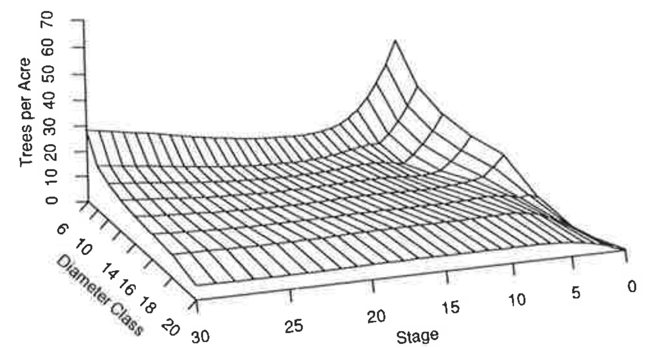

FIG. 4. Start of stage stand structures for harvest method 2 parameter set C.

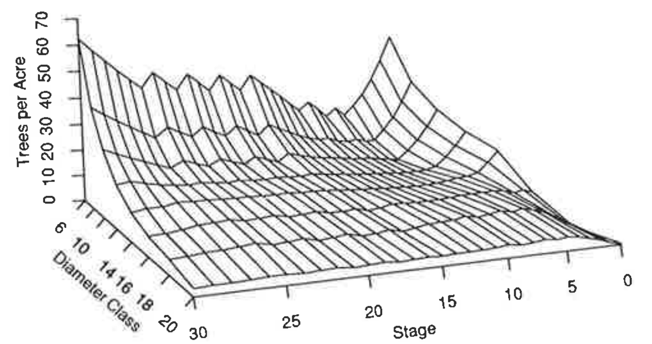

FIG. 6. Start of stage stand structures for harvest method 2 parameter set D.

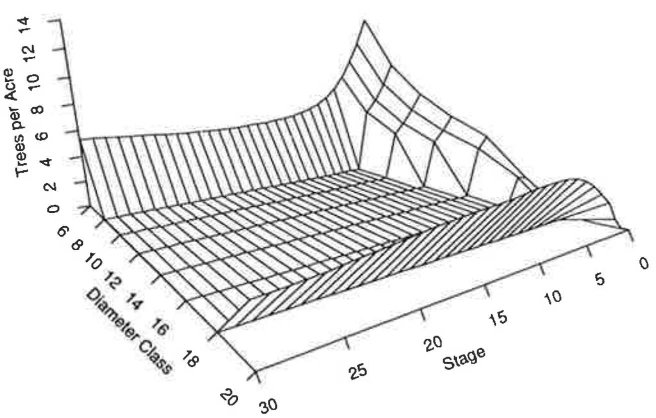

FIG. 5. Harvest control vectors for harvest method 2 parameter set C.

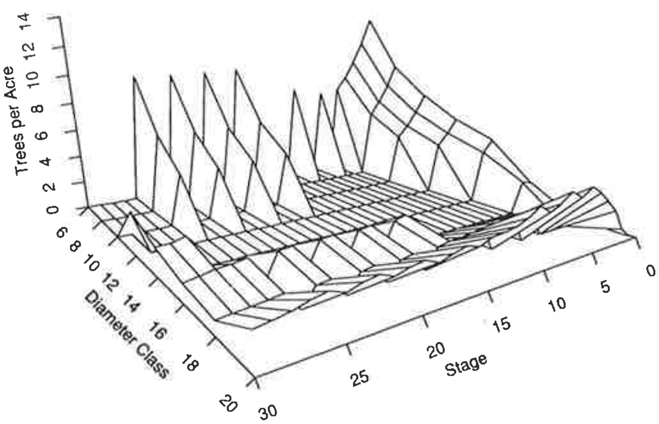

FIG. 7. Harvest control vectors for harvest method 2 parameter set D.

during the first two stages occurs in the smaller diameter classes, followed by a shift into the larger (i.e., 14- to 18-in.) diameter classes until the 20th stage. Then harvest alternates between the upper and lower diameter classes for six stages and finishes by harvesting in the upper diameter classes.

Early rapid drawdown of the original inventory followed by a period during which inventory accumulates characterize both harvest decision parameter sets A and B. Possible long term cycles of heavy harvest followed by a recovery period with lighter harvests could be inferred from the graphs delineating the start of stage stand structure and harvest vector time behavior.

Harvest allocation method 2

Two sets of parameters for harvest allocation method 2 were applied to the start of stage conditions. For harvest decision parameter set C, cutting is constrained to always start in the 6-in. diameter class. For harvest decision parameter set D, harvesting can originate in either the 6- or 18-in. diameter class. Both rules use the same basal area percentage targets yet produce significantly different behaviors. Results for harvest method 2 parameter set C are shown in Figs. 4 and 5. These figures show that over the projection horizon of the run, the start of stage stand structure approaches a one cut steady state after a four-stage transition. The resultant stand structure has fewer trees in the 6- through 14-in. diameter classes and more trees in the 16- through 20-in. diameter classes than the starting stand structure. Thus, stand structure shifts to the higher diameter classes. After the four-stage transition, harvest settles down to only removing trees from the 6- and 20-in. diameter classes. As shown by Haight (1985), this result confirms that an optimal steady-state distribution for this growth model only harvests in the smallest and largest diameter classes.

Parameter set C illustrates the one cut steady state of traditional uneven-aged stand management. It achieves this by removing trees from the smaller, less valuable diameter classes during a 20-year transition period and then maintains this structure by harvests in the 6- and 20-in. diameter classes. However, attaining a steady-state results in less NPV than either of the runs that use harvest allocation method one. Results from harvest method 2 parameter set D shown in Figs. 6 and 7 display a markedly different pattern of start of stage stand structure and harvest behavior over the projection period than harvest method 2 parameter set C. This result is due to the option of initiating harvest in either the 6- or 18-in. diameter classes. During the first four stages, the start of stage stand structure is identical to that of harvest method 2 parameter set C. But, starting at stage 16, a three stage cyclic pattern of harvest emerges. This results in harvests from the 6- to 10-in. and the 16- and 18-inch diameter classes, with the obligatory harvest in the 20-in. diameter class.

Allowing harvest allocation method 2 to start harvesting from either the 6- or 18-in. diameter classes results in a steady-state behavior that breaks from tradition. Whereas harvest method 2 parameter set C exemplifies the traditional steady-state behavior, harvest method 2 parameter set D represents a harvesting concept in which a steady state spans several harvesting periods in a cyclical pattern, such that at each point the stand structures are equal between cycles. Accumulated NPV of harvest method 2 parameter set D is marginally higher than that of harvest method 2 parameter set C, but probably not significant.

Discussion and conclusions

Previous optimization studies using the Adams and Ek (1974) northern hardwood growth model focused on nonlinear programming techniques and optimal control theory approaches for both the static and dynamic formulations of the uneven-aged stand management problem. Recent advances in processing power and storage capacity of microcomputers made the possibility of a dynamic programming approach to the transition uneven-aged optimization feasible to explore. Based on this study, it appears that dynamic programming is a viable methodology to apply to uneven-aged stand management. One atractive feature of dynamic programming is its flexibility in producing high-valued transition regimes as well as steady-state regimes. A number of other questions also can be answered with this approach including the impact of differing harvest specifications and questions concerning the sensitivity of the model to both internal and external parameters.

As part of a sensitivity analysis, three parameters were investigated: (*i*) interest rates, (*ii*) neighborhood storage class widths, and (*iii*) initial stand conditions. Interest rates were varied from 0–8% in 2% increments, with results confirming the generally known result that as interest rates increase the level of the optimal residual inventory decreases.

A concern with using the neighborhood storage concept is that, owing to aggregation, information is lost as class width increases. That is, information related to unique stand structures is lost during the optimization process. To determine the magnitude of differences in NPV, neighborhood storage class widths for TTPA and TBAA were varied by ±5 units from the base run values of 10 units. While the optimal harvest decisions remained unchanged, the NPVs and computational times for the harvest decision vectors did not. For example, decreasing neighborhood storage class widths from 10 to 5 units produced a less than 1% increase in NPV, but a ≥400% increase in execution time. But, increasing neighborhood storage class width from 10 to 15 units reduced NPV by up to 9% and execution time by 83%. Thus, care must be taken when establishing the width of the neighborhood class.

Integral to any growth projection model is specification of the initial states of the resource being modelled: trees per acre in each diameter class for the UNEVENDP model. To explore this, initial stand structures were modified by decreasing the trees per acre in the 6-in. diameter class by one tree and secondly, by decreasing the number of trees in the 20-in. class by 0.1 tree. Performing this modification lead to harvest control vectors and start of stage stand structures that diverged markedly from the base cases. Although harvest decision vectors vary as a result of these changes, the NPVs remained almost constant from the base runs. These results indicate that accuracy in specifying the initial stand structure is essential for reliable long-term optimization projections.

Finally, based on the results presented herein, there appear to be divergent optimal paths that can be followed depending on the harvest method used to allocate total stand harvest values to individual diameter classes. A prior study by Haight (1985) presented an optimal harvest strategy with a NPV of \$171.29/acre. Results presented here for harvest method 1 parameter set A produced a NPV of \$178.83/acre, an increase of \$7.54/acre. This confirms that solutions produced by optimization studies are probably only local optima and not global solutions. Unless better behaved functions are used, this will be the norm in the future as well. It further suggests that there are numerous local optima on the response surface of the Adams and Ek (1974) northern hardwood uneven-aged growth model. Results of previous optimization studies using this model confirm this point. It appears that there are numerous optimal paths that depend on harvest assumptions and other internal and external parameters.

Adams, D.M. 1976. A note on the interdependence of stand structure and best stocking in a selection forest. For. Sci. **22**: 180–184.

Adams, D.M., and Ek, A.R. 1974. Optimizing the management of uneven-age forest stands. Can. J. For. Res. **4**: 274–287.

Anderson, D.J. 1992. A dynamic programming algorithm for optimization of uneven-aged stand management. M.S. thesis, University of Washington, Seattle.

Arthuad, E.L., and Klemperer, W.D. 1988. Optimizing high and low thinnings in loblolly pine with dynamic programming. Can. J. For. Res. **18**: 1118–1122.

Bare, B.B., and Opalach, D. 1987. Optimizing species composition in uneven-aged forest stands. For. Sci. **33**: 958–970.

Bare, B.B., and Opalach, D. 1988. Determining investment-efficient diameter distributions for uneven-aged northern hardwoods. For. Sci. **34**: 243–249.

Bellman, R.E. 1957. Dynamic programming. Princeton University Press, Princeton, N.J.

Brodie, J.D., and Haight, R.G. 1985. Optimization of silvicultural investment for several types of stand projection systems. Can. J. For. Res. **15**: 188–191.

Brodie, J.D., and Kao, C. 1979. Optimizing thinning in Douglas-fir with three-descriptor dynamic programming to account for accelerated diameter growth. For. Sci. **25**: 665–672.

Buongiorno, J., and Michie, B.R. 1980. A matrix of uneven-aged forest management. For. Sci. **26**: 609–625.

de Liocourt, F. 1989. De l'amenagement des sapinieres. Bull. Soc. For. Franche-Comte Belfort.

Dreyfus, S.E., and Law, A.M. 1977. The art and theory of dynamic programming. Academic Press, New York.

Dykstra, D.P. 1984. Mathematical programming for natural resource management. McGraw-Hill, New York.

Ek, A.R. 1974. Nonlinear models for stand table projection in northern hardwood stands. Can. J. For. Res. **4**: 23–27.

Filius, A.M., and Dul, M.T. 1992. Dependence of rotation and thinning regime on economic factors and silvicultural constraints: results of an application of dynamic programming. For. Ecol. Manage. **48**: 345–356.

Gove, J.H., and Fairweather, S. 1992. Optimizing the management of uneven-aged forest stands: a stochastic approach. For. Sci. **38**: 623–640.

Haight, R.G. 1985. A comparison of dynamic and static economic models of uneven-aged stand management. For. Sci. **31**: 957–954.

Haight, R.G. 1987. Evaluating the efficiency of even-aged and uneven-aged stand management. For. Sci. **33**: 116–134.

Haight, R.G. 1990*a*. Economic models for evaluating plantation and uneven-aged forestry. *In* Proceedings of the Southern Forest Economics Workshop on Evaluating Even and All-aged Timber Management Options for Southern Forest Lands. *Edited by* C.A. Hickman. USDA For. Serv., Gen. Tech. Rep. SO-79. pp. 25–34.

Haight, R.G. 1990*b*. Feedback thinning policies for uneven-aged stand management with stochastic prices. For. Sci. **36**: 1015–1031.

Haight, R.G., and Getz, W.M. 1987. Fixed and equilibrium endpoint problems in uneven-aged stand management. For. Sci. **33**: 908–931.

Haight, R.G., and Monserud, R.A. 1990*a*. Optimizing any-aged management of mixed-species stands. I. Performance of a coordinate-search process. Can. J. For. Res. **20**: 15–25.

Haight, R.G., and Monserud, R.A. 1990*b*. Optimizing any-aged management of mixed-species stands. II. Effects of decision criteria. For. Sci. **36**: 125–144.

Haight, R.G., Brodie, J.D., and Dahms, W.G. 1985*a*. A dynamic programming algorithm for optimization of lodgepole pine management. For. Sci. **31**: 321–330.

Haight, R.G., Brodie, J.D., and Adams, D.M. 1985*b*. Optimizing the sequence of diameter distributions and selection harvests for uneven-aged stand management. For. Sci. **31**: 451–462.

Hann, D.W., and Brodie, J.D. 1980. Even-aged management: basic managerial questions and available or potential techniques for answering them. USDA For. Serv., Gen. Tech. Rep. INT-83.

Hool, J.N. 1966. A dynamic programming – Markov chain approach to forest production control. For. Sci. Monogr. 12.

Hotvedt, J.E., and Ward, K.B. 1990. A dynamic programming optimization model for uneven-aged loblolly–shortleaf pine stands in the mid-South. *In* Proceedings of the Southern Forest Economics Workshop on Evaluating Even and All-aged Timber Management Options for Southern Forest Lands. *Edited by* C.A. Hickman. USDA For. Serv., Gen. Tech. Report SO-79. pp. 35–43.

Kennedy, J.O.S. 1986. Dynamic programming applications to agriculture and natural resources. Elsevier Applied Science Publishers, London.

Martin, G.L. 1982. Investment-efficient stocking guides for all-aged northern hardwood forests. College of Agriculture and Life Sciences, University of Wisconsin, Madison. Res. Rep. R3129.

Michie, B.R. 1985. Uneven-aged stand management and the value of forest land. For. Sci. **31**: 116–121.

Riitters. K., Brodie, J.D., and Hann, D.W. 1982. Dynamic programming for optimization of timber production and grazing in ponderosa pine. For. Sci. **28**: 517–526.

Tait, D.E. 1986. A dynamic programming solution of financial rotation ages for coppicing tree species. Can. J. For. Res. **16**: 799–801.

Valsta, L.T. 1990. A comparison of numerical methods for optimizing even aged stand management. Can. J. For. Res. **20**: 961–969.

Valsta, L.T. 1993. Stand management optimization based on growth simulators. Information Systems Group, Finnish Forest Research Institute, Helsinki. Res. Pap. 453.