Forestry Canada Forêts Canada

Dynamic programming applications to stand level optimization

William A. White

Information Report BC-X-305
Pacific Forestry Centre

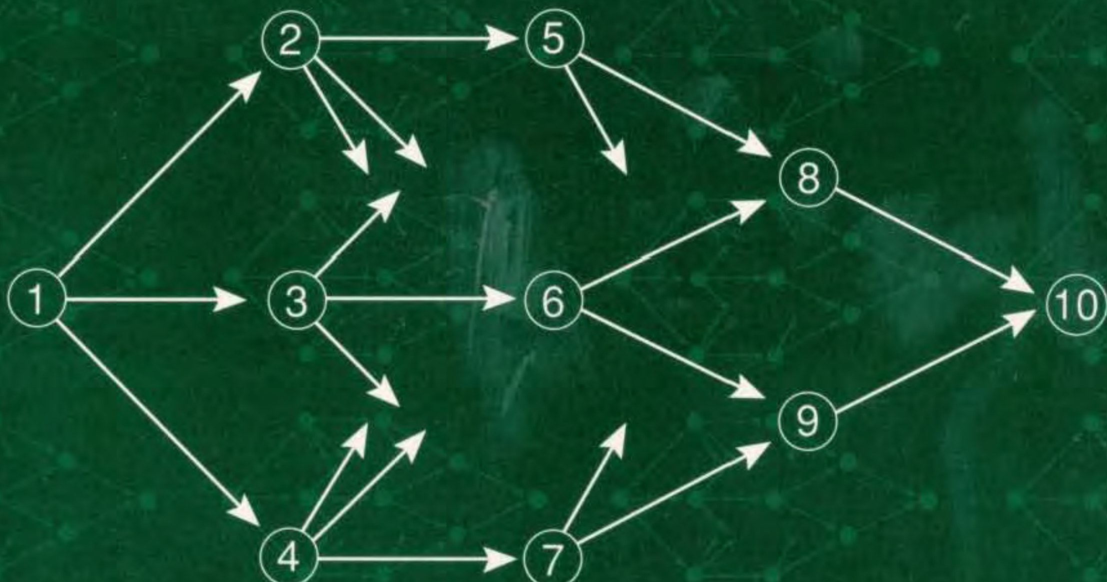

Dynamic programming applications to stand level optimization

William A. White

Forestry Canada
Pacific Forestry Centre

BC-X-305

1989

Forestry Canada
Pacific Forestry Centre
506 West Burnside Road
Victoria, British Columbia
V8Z 1M5
Phone (604) 388-0600

ISSN 0830-0453
ISBN 0-662-16701-5
Cat. No. Fo46-17/305E
Printed in Canada

Additional copies of this publication
are available in limited quantities
at no charge from the
Pacific Forestry Centre

Microfiches of this publication may be purchased from:

MicroMedia Inc.
Place du Portage
165, Hotêl-de-Ville
Hull, Quebec
J3X 3X2

Abstract

The features which characterize dynamic programming problems are reviewed. Optimal forest management problems are then shown to fit into the framework for solution using dynamic programming. A chronological review of the literature of dynamic programming in forestry follows.

Résumé

Le rapport examine d'abord les caractéristiques des questions résolues par programmation dynamique. Il montre ensuite que les problèmes d'aménagement forestier optimal tombent dans la catégorie de ceux auxquels cette méthods convient. Il revoit ensuite, dans l'ordre chronologique, la documentation portant sur la programmation dynamique appliquée en foresterie.

Acknowledgements

I am grateful for reviews of an earlier version of this report by Terry Heaps and Ernie Love of Simon Fraser University.

Contents

Introduction

Forest planning can be carried out at a number of levels: a single stand, multiple stands, a timber supply area, a province or state, or even a nation. Of these, stand level solutions are the basic unit of analysis. An effective plan at this level can be integrated into higher level plans to improve those analyses. In recent years, dynamic programming has emerged as a powerful approach to stand level problems. The purpose of this report is to acquaint the reader with dynamic programming and its application to forestry. This will be accomplished by reviewing the basic features which must characterize a problem which can be solved using dynamic programming. A discussion will follow of the suitability of dynamic programming to stand level optimization problems. This will be followed by a review of application of dynamic programming to forestry problems with emphasis on stand level optimization applications.

Characteristics of dynamic programming problems

Dynamic programming is essentially an optimization approach that simplifies complex problems by transforming them into a sequence of smaller simpler problems (Bradley et al. 1977). However, not all problems can be broken down and simplified. Following the presentation in Hillier and Lieberman (1980) the basic features which characterize dynamic programming problems are discussed below:

1. The problem can be divided into *stages* with a *policy decision* required at each stage. Stages normally represent time periods in a planning horizon, but they can represent anything which divides the problem up into sections with a decision required at each section; for example, legs of a journey might be considered stages of a long journey. The policy decision is the action to be taken at each time period, stopping point, etc. A further characteristic of dynamic programming problems is that the sequence of policy decisions are interrelated.

2. Each stage has a number of *states* related to it. The state is a description of the various possible conditions the system may be in at a given stage. A state is expressed in terms of one or more state variables, such as stock inventory level, a degree of machine wear, or a description of a forest stand in terms of the number of trees, basal area, and stand age. The state description is a vital component of the DP formulation. It must be detailed enough to accurately describe the system being studied but simple enough to limit the number of states at each stage. As the number of states at each stage grows the problem is beset with the "curse of dimensionality" and becomes difficult if not impossible to solve. For instance, the forest could be described by the number of trees, basal area, tree height, stand age, soil characteristics, years since last thinning, and so on, but by doing so each stage would have a very large number of states and the problem could not be solved. Instead, detail is sacrificed to allow a solution to be obtained.

3. Each time a policy decision is taken it transforms the current state into a state associated with the next stage. The new state may be determined by both the policy decision and a probability distribution. For instance, given a level of product inventory (a state variable), the decision (a stage) of how much to produce today (the policy decision) will influence how much inventory (the state) we have at the next stage. A decision to thin a stand at age fifteen will influence the characteristics of that stand at age twenty. The probability distribution recognizes that we do not know with certainty what the stand will look like. Some dynamic programs can be set up as networks and solved as linear programs (Hillier and Lieberman 1980, chapter 10). In Figure 1 the columns refer to the stages while the nodes correspond to states. Policy decisions move us through the network. Contributions to the objective function from the various policy decisions are represented by values assigned to the branches which connect the nodes.

4. The optimal policy for all remaining stages is entirely independent of policies adopted in previous stages.

This feature of dynamic programs is commonly called the principle of optimality. Wagner (1975, page 266) describes this principle as follows: "An optimal policy must have the property that regardless of route taken to enter a particular state, the remaining decisions must constitute an optimal policy for leaving that state." It is this property that allows dynamic programs to be broken up into a series of smaller, simpler problems. We

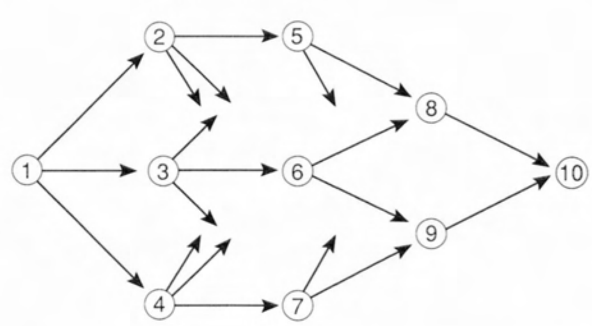

Figure 1. The network form of a dynamic program.

will see in the next section how this property places some important constraints on the formulation of problems involving policies or actions which have a lasting effect, such as forest management operations.

5. The solution procedure must begin by finding the optimal policy for either the first or last stage.

These techniques are called forward induction and backward induction respectively. Backward induction is more commonly used, but in the forestry literature (Brodie et al. 1978) forward induction is more common. The advantages and disadvantages of each method are discussed in the next section.

6. Finally, a recursive optimization procedure can be developed "which builds to a solution of the overall N-stage problem by ... sequentially including one stage at a time and solving one stage problems until the overall optimum has been found" (Bradley et al. 1977, page 462). While the recursive relationship will vary according to the problem, a general formulation suitable to deterministic problems is described below.

A general forward recursive function is:

$$f^*_n(s) = \max_{x_{n+1}} \{ f^*_{n+1}(x_{n+1}) + c_{sx_{n+1}} \}$$

where: $f^*_n(s)$ is the return from being in any state in the set of all feasible states from the first stage of the problem to the present stage (n) and state (s); x_{n+1} is any state in the set of all feasible states from which the current state s can be reached; and $c_{sx_{n+1}}$ is the return associated with going from state x_{n+1} to the current state s. The problem is solved by sequentially choosing the maximizing values of x_{n+1}.

Forestry and dynamic programming

The previous section outlined the basic features of problems that can be solved using dynamic programming. Are forest management problems amenable to solution using dynamic programming?

Many dynamic programming problems are divided into stages by dividing up an extended period of time into equal sections. Stand optimization problems are suited to be divided up into time segments of 1, 5, 10 or any other number of years with an action or policy decision such as thin, fertilize, or harvest required at each stage. The role of the stage interval in a dynamic programming problem in forestry is considered by Kao and Brodie (1979) and to a lesser extent by Roise (1986). Reducing the stage interval from 10 years to 1 year increased the soil expectation value of the stand by about 1% according to Kao and Brodie. Roise increased the soil expectation value by 39% by reducing the stage interval from 15 to 10 years. Since the gains obtained by reducing the interval to less than 10 years are small, this interval is often used in the literature.

As mentioned in the previous section, the state variable should describe the forest in a simple way. A balance must be achieved between describing the forest as accurately as possible and limiting the number of states so that the problem can be solved in a reasonable time. Previous work, described in more detail in the next section, has limited the description of the forest stand to age, number of trees per hectare, basal area per hectare, and sometimes one other variable such as time since last thinning or type of thinning. As well as limiting the number of variables which describe the forest, it is also important to render these variables discrete.

Kao (1980) shows that, if all diameter classes are to be considered for thinning, the difference in the number of trees per hectare between states must be at least 15. Similarly, basal area is described in intervals of, for example, 0.5 m^3/ha. Kao (1980) discusses the impact of various intervals of the number of trees and basal areas. A small interval of about 15 trees provides more accuracy for a given basal area interval, but increases computing time. Increased precision can be obtained by decreasing both the tree interval and the basal area interval. Kao (1980) shows that reducing the basal area interval from a 3.7 m^2/ha to 0.37 m^2/ha increases the soil expectation value by about 5%. The greatest benefits come in reducing the interval from 3.7 m^2/ha to 1.9 m^2/ha. Kao states that any interval between 0.37 m^2/ha to 1.9 m^2/ha should be suitable along with a tree interval of 15 trees. Larger basal area intervals were shown to significantly reduce computing time.

Smaller tree and basal area intervals also reduce the occurrence of artifact effects. Such effects occur when more than one alternative is found to lead to a state. The alternative with the greatest cumulative net worth is selected to represent that state. It is possible, however, that an alternative not selected could contribute more in future periods and ultimately would have been the best choice by a small amount. Kao (1980) and Brodie and Kao (1979) stress that the differences in value will be very small. The errors caused by artifact effects could only be eliminated by using continuous states. They can be minimized by using large tree intervals or by using a combination of small tree intervals and small basal area intervals.

The key factor in determining interval widths is the same as that for determining the number of variables used per state: they must not be so fine as to create an inoperable number of states, or so broad that they do not distinguish between forest stands that should be classed differently.

Policy decisions in forest stand problems involve forest management decisions such as degree of juvenile spacing, level of fertilization, degree of thinning, and, eventually, the decision to harvest and to return to bare land or an initial stocking density. These decisions move the stand from an existing state to a state associated with the next stage. As with the descriptions of states, these decisions are made in intervals. For example,juvenile spacings of 10, 20, and 30% of stems or basal area might be compared.

The principle of optimality places an important constraint on forest stand problems. The effects of an action taken today must be completely described by the state variables. This is exemplified by fertilization. Let us assume that the state variables being used are number of trees and basal area per hectare and the stages are time intervals of 10 years. If in year 30 a stand is fertilized, this action accelerates growth and in year 40 the stand is again described by the number of trees and basal area per hectare. The principle of optimality requires that all the effects of the fertilization be completed in the 10-year period. Another way of stating this is that a stand with N trees per hectare and G basal area per hectare today which has been fertilized in the past must grow from now on in the same way as a stand with the same number of trees and basal area which has not been fertilized. If the time interval between stages is not sufficient to guarantee this condition, then an additional state variable is required to described the stand as fertilized, or the stage interval must be increased.

When a single-tree distance-dependent growth model is used, the configuration of the trees will also affect stand growth. As noted, the principle of optimality requires stands in the same state to grow in the same way. Number of trees and basal area per hectare will not be sufficient descriptors if stand configurations vary. The nonrandom element which will have the greatest effect on where the trees are in the stand, and in particular how close one tree may be to another, is the initial stocking density of the stand. Stands with the same initial density must be simulated repeatedly with the same initial tree locations. This will prevent differences in growth caused by tree location. Where initial density is an important consideration, it must be added to the list of state descriptors.

Forward recursions have been favored in the forestry dynamic programming literature largely because they do not require "a separate pass through the network ... for each candidate rotation" (Brodie et al. 1978, page 517). Forward recursions require only paths of interest to be searched. The advantage of this technique is that it minimizes the solution time and provides optimal paths to rotations shorter than those paths being investigated. It fails, however, to provide optimal regimes for states not on the optimal path; in such cases, the recursion must be restarted at the "nonoptimal" state and the problem resolved. Backward recursions provide optimal regimes for all states considered in the problem. Within a large problem, such as most forestry problems, and where the optimal path is of primary importance, the time savings involved in forward recursions are significant when only one pass through the network is required. If the user is primarily interested in what happens outside the optimal path or if the problem being studied is relatively small, then backward recursions appear to be superior.

Forest stand optimization problems fit the framework for dynamic programming. However, the principle of optimality must be adhered to, and the problem must not be too large to be solved within a reasonable time. It must also be remembered that the growth model used does not provide exact stand growth but only an approximation based on limited information. Characteristics such as species, age, site quality, density, management regime, or tree quality are related directly through equations to stand parameters such as number of trees, average diameter, height, basal area and volume (Mitchell 1980). The solutions will only be as precise as the data.

Alternative methods

The problem of optimal forest management regimes has been (and continues to be) solved by methods other than dynamic programming. Indeed, the early dynamic programming literature (such as Amidon and Akin 1968) used the same problems solved earlier by other means to verify their solutions. Other than dynamic programming, there are three dominant solution methods: marginal analysis, control theory, and comparative simulation

Marginal analysis was one of the earliest approaches to the stand optimization problem. Early approaches include USDA Forest Service (1963), Duerr and Christiansen (1964) and Chappelle and Nelson (1964). Schreuder (1971) criticized the work of Chappelle and Nelson and noted some weaknesses in the marginal analysis approach. In forestry, marginal analysis involves comparing the marginal value growth percent of the stand with the marginal cost of capital for each consecutive time period. The optimal stocking level is then determined. Schreuder states that the main problem with this approach is that it may not provide a global optimum, because it takes into account only one period at a time and thus does not consider the interdependencies of the periods. This approach excludes precommercial thinnings, which are defined as thinnings which cost more than they enhance the value of the stand in the same period, although such thinnings have been shown to be economic over the life of a stand. Schreuder also says that marginal analysis becomes very difficult if not impossible to use when prices and costs vary.

Brodie et al. (1978) provide a brief review of the important contributions to the stand optimization problem by writers working in a control theoretic framework. Naslund (1969) formulated the problem but provided no solution. Anderson (1976) formulated a solution which showed consistency with the traditional Faustman model (Samuelson 1976) and other more general dynamic models but again no actual solutions are shown. Clark and de Pree (1979) formulate the problem as one of linear control and provide numerical results. McDonough and Park (1975) and Dixon and Howitt (1980) are reported in Cawrse et al. (1984) to have obtained local optimum solutions using iterative techniques. As evidenced by the papers noted here, the major shortcoming of this approach is the difficulty in obtaining numeric solutions to problems. Cawrse et al. (1984) attempt to overcome this difficulty by employing a variational solution technique. Their solution, however, would only be effective for simple models of stand growth which use only one equation.

Comparative simulation has been very popular in recent years. Simulators "grow" a stand to a given age, thin the stand, grow the stand again, and so on. A path is followed to some rotation age. Paths are compared, and the path which maximizes a particular characteristic of the stand, such as volume or market value, is considered best. Complete enumeration of every possible path through a network is not feasible, and even approaching that lofty goal is time-consuming and expensive. Randall (1977), Reukema and Bruce (1977), and Sleavin (1983) discuss comparative simulation techniques and provide examples.

The method of solution of stand optimization problems which has shown the greatest promise in recent years is dynamic programming. It has the potential to overcome the difficulties noted in each of the alternative methods discussed above. It is a multiperiod approach; it provides global optimal solutions; variable factor and product prices can be allowed for; and every possible path need not be searched to obtain the solution. In sum, dynamic programming is a useful tool for this problem because it is more efficient than simulation, and simpler and more versatile than marginal analysis and control theory.

Literature review

Hool (1966) provided the first North American[1] application of dynamic programming to a forest production problem. Using data from the Darlington Woods, Indiana Continuous Forest Inventory System, he defines states by groupings of volumes, tree counts, and by whether or not the stand had been thinned. Stages are divided into 2-year intervals over a total time period of 16 years. The choices left to the decision makers are undisturbed growth, thinning, and harvesting; the harvesting option was divided into selection harvesting and clearcutting. The problem is then solved as a Markov chain using the general recurrence relation:

$$V^*_m(i) = \max_k \{\sum_{j=1} P^k_{ij} [r^k_{ij} + V^*_{m-1}(j)]\}$$

where i represents states of which there are n; m represents stages; P^k_{ij} represents the probability of a transition from state i to state j given decision k; r^k_{ij} represents the return associated with the transition from i to j given decision k; $V^*_{m-1}(j)$ represents the optimal accumulated expected returns given that you start from state j in stage m-1; $V^*_m(i)$ represents the optimal expected return of being in state i at stage m.

Hool uses the above to obtain the optimal mangement activity policy for each state and then develops prescriptions by assuming that once the optimal policy decisions have been made that the most likely transitions occur. He goes on to calculate the mean number of stage transitions in which each state is expected to remain unchanged, the mean number of transitions the woodlot in a given state requires in order to go to another state by undisturbed growth, and the probability of a state existing if the stand were allowed to grow undisturbed for a long period of time. It is a Markovian problem since the probability and return values are specified to be independent of time and the probability of transition from state i to state j depends only on state i and not on the history of the system.

While the Hool paper is valuable in demonstrating a forestry application of the Markov chain, it is applied to a very restricted problem. The number of states is small, and time horizon is not only finite (16 years) but very short for a forestry problem. The technique could, however, be generalized and used for larger problems.

Two researchers in the USDA Forest Service provided the first widely circulated application of a deterministic dynamic problem to a forest stand prob-

[1]He cites a Japanese paper that appears to be the first dynamic programming application to forestry.

lem (Amidon and Akin 1968). The authors use dynamic programming to confirm the results obtained through marginal analysis by Chappelle and Nelson (1964). They use a backward recursion and therefore solve the problem for several different rotations. They restrict the decisions to be made on their sawtimber stand to the following options: a light thinning that leaves the stand with 1 MBF more volume of sawtimber at the start of the next stage 5 years hence, or a heavier thinning that leaves the stand with 1 MBF less sawtimber at the start of the next stage. This highly restricted problem is solved by using the following recursive structure:

$$T(x,y) = \max \{D(x,y) + T(x+1, y+1); \\ I(x,y) + T(x+1, y-1)\}$$

where x refers to stand age in 5-year intervals; y refers to sawtimber volume in 1 MBF intervals; $D(x,y)$ is decrement in growing stock value from $x+1$, $y+1$ to x,y; $I(x,y)$ is increment in growing stock value from $x+1$, $y-1$ to x,y; $T(x,y)$ is the optimal total value from (x,y) up to the rotation age; and $I(n,y)$ is the initial condition, i.e., final harvest value for the rotation corresponding to n.

The results confirm the work of Chappelle and Nelson and note the more flexible nature of dynamic programming in handling changes such as increases in product prices. The shortcoming of the work was the highly restricted set of thinning options and the simple description of the stand in terms of sawtimber volume alone.

Gerard Schreuder (1971) recasts a continuous time model as a discrete time dynamic program to obtain an optimal thinning schedule and rotation age. This recasting simplifies the solution of the problem. The stages are time periods and states are described by the total volume per acre. The decision variable is net cut (cut minus growth). The problem is solved by backward recursion and improves upon the Amidon and Akin solution by including the cost of land. This paper also uses an oversimplified description of the stand, but it was valuable in demonstrating the flexibility of the dynamic programming approach compared to other solutions.

Lembersky and Johnson (1975) and Lembersky (1976) use a Markovian decision process to optimize timber production in the face of uncertainties. The uncertainties stem from the growth and natural mortality of the trees in the latter paper, while the former also includes uncertain future prices. Lembersky and Johnson (1975) seek to maximize the financial return from the forest while Lembersky (1976) is interested in maximizing average annual volume. According to Lembersky (1976, page 69): "The two criteria receive separate treatments because there are fundamental differences between Markov Decision Processes (MDP) with average undiscounted rewards (here volumes) and MDP with discounted summed rewards (dollar revenues in Lembersky and Johnson). The differences are not in the underlying basic structure, but result from mathematical properties of aggregated expected rewards."

The properties differ because undiscounted returns tend to infinity over an infinite horizon and therefore must be averaged to provide meaningful comparisons. Discounted returns are almost always finite and therefore comparable.

Lembersky and Johnson expand on the work of Hool (1966) by stretching it to an infinite horizon and including market behavior. The state space includes the state variables of average dbh, number of trees per hectare and relative market price. By not including an indicant of thinning, the authors are assuming that whether the stand reaches its current state naturally or by thinning it will grow the same thereafter. Forty-eight forest states are considered and given five relative market values so the problem consists of 240 states in total. Fifteen management actions are considered ranging from leaving the stand undisturbed to clearcutting and replanting at various densities. Optimal actions are then derived for each state.

The strength of these papers is that they recognize risk in forest production problems; their weakness is that the range of states and management decisions, as in Hool's work, is very limited.

J. Douglas Brodie and his colleagues and students at Oregon State University have made a number of valuable contributions to the literature of dynamic programming and forestry. Brodie et al. (1978) using data on yield curves from McArdle et al. (1961) use a forward recursion to simultaneously determine optimal stocking levels and rotation. The stage intervals are 10 years and the stands are described by volume per acre. A major improvement over previous work is the increased number of states and therefore policy decisions included. The authors consider interval units of 100 cubic feet per acre (7 m^3/ha). At each state all actions which could achieve a lower stocking level are considered using the recursive function described by Kao (1980) as:

$$T(x+1, y) = \max_{y'} \{T(x, y') + P(y', y)\}$$

where y' is any node stocking level in the current stage that can reach stocking level y in the next stage; $P(y', y)$ is the net revenue obtained from the transition from state y' to y; $T(x+1, y)$ is the total value of the optimal schedule up to $(x+1, y)$; $T(x, y')$ is total revenue of the optimal schedule up to the stage being considered.

Since forest growth is continuous and dynamic programming requires states to be described in discrete terms, some stands required a mandatory thinning so as to fit into a state. This is overcome in Brodie and Kao (1979) where network nodes are treated as "neighborhood storage locations" for the exact values which describe the state and optimize the present stage. For example, all stands with between 148 and 185 stems and a basal area of between 10.5 and 11.5 m^2 could be stored at node (185, 11). The values used for subsequent growth would be the actual optimal continuous values and not the neighborhood values. Previous solutions have used the neighborhood values. As well, thinnings are only considered every 10 years. The effects of varying this stage length was discussed in the previous section.

The primary shortcoming of Brodie et al. (1978) was "the model could not explicitly treat the accelerated diameter growth associated with more intensive thinning" (Brodie and Kao 1979, page 665). This is overcome by replacing the yield tables of the previous paper with the biometric model DFIT (Bruce et al. 1977) which simulates the growth of Douglas-fir stands. The incorporation of tree size is important as it affects selling prices and logging costs. In order to incorporate tree size into the dynamic programming solution, the state descriptors are altered. The stand age remains one descriptor but the volume descriptor has been replaced by two descriptors: the number of trees and basal area. The solution is again obtained by using a forward recursion and the algorithm DOPT. DOPT is simply DFIT with an algorithm to solve dynamic programs. This allows for the greatest amount of variability seen in the literature to date. Intervals of 15 trees, 4 square feet per acre (0.89 m^2/ha) of basal area, and stages of 10 years are used in the paper. The optimal value function is defined as the value of the present net worth "path" from regeneration to age t, number of trees N, and basal area G for the stand:

$$f(N,G) = \max_{\{n,g\}} \frac{P_D T - L_D T - C}{(1+i)^t} + f_{(t-10)} \ (n,g)\text{"}$$

(Brodie and Kao 1979, pages 668-669)

where P_D is the selling price of logs of average diameter D at breast height; T is the volume of thinnings [including mortality captured]; L_D represents logging costs given average diameter D; C is a fixed entry cost which is incurred if thinning or harvest take place; i is the interest rate; (n,g) is the set of feasible basal areas and number of trees at age (t-10) from which the current level of N and G can be reached.

The starting condition is:

$$f_{30}(N,G) = R$$

where R represents the value cost of all regeneration and other treatments before age 30[2].

This paper represented the culmination of work in the field of deterministic dynamic programs related to forestry. Kao (1980) expands on this article in his dissertation and considers more management variables than just thinning. Riitters et al. (1982) investigate the optimization of timber production and grazing in ponderosa pine (*Pinus ponderosa* Laws) but "the structure of the dynamic programming algorithm essentially is the same" (page 519). Hann et al. (1983) used the same structure to study initial planting density and precommercial thinning.

DOPT is used again in Riitters et al. (1982) to compare volume versus value maximization. Sleavin (1983) develops a similar model for the growth simulator DFSIM. Martin (1978) and Haight et al. (1984) utilize single-tree simulators as opposed to whole stand simulators such as DFIT and DFSIM but modify the output to permit the computational format to remain essentially the same (Brodie and Haight 1985).

While this research was going on at Oregon State University, a few other papers linking dynamic programming and forestry appeared independently of that main body of research. The most fundamental were those of Chen et al. (1980a, 1980b). The state variable used is basal area and the decisions are thinning by amount of basal area. De Kluyver et al. (1980) determine optimal stand management using a two-stage approach for a forest consisting of many stands. A dynamic program is used to find the most efficient regime for each stand, then a multiple objective linear program is used to determine the optimal policies for the whole forest. The state space is a yield distribution over various products such as veneer, sawlogs, or pulpwood. The problem solved is very restricted and only eight different regimes are investigated. The paper has very limited practical application.

With the exception of the papers that constructed dynamic programs as Markovian decision processes (Hool 1966; Lembersky and Johnson 1975; Lembersky 1976) the questions of risk and uncertainty have not been taken into account in the papers reviewed[3]. The

[2] The authors simulate a naturally regenerated stand to age 30 and make management decisions from that point on.

[3] Risk and uncertainty are used in the Knightian sense. Risk refers to decision making with known probabilities of occurence, while uncertainty refers to complete ignorance of the probabilties.

risk of fire is handled in a Markov process dynamic program by Martell (1980). He shows that the soil expectation value and the optimal rotation fall as the risk of fire increases. Reed and Errico (1985) reach a similar conclusion. A more general review of risk and uncertainty in stand level problems is available in Kao (1982, 1984); these two papers investigate optimal stocking levels and rotation under risk and uncertainty, respectively.

Under deterministic assumptions, when we know the level of growing stock[4], growth is a fixed value. When risk is present, growth becomes a random variable and is characterized by a probability function. In Kao (1982) stocking level is used as the state variable and volume is maximized. There is a probability of being in each state and an expected thinning associated with being in that same state.

If we are in state *A* and desire to be in state *Y'* then the probability of being in state *Y'* is the cumulative probability of having a growing stock greater than that level in the next stage. The stand could then be thinned back to that level and the appropriate amount of thinning can be calculated. It should be stressed that the structure of network nodes and the basic method of solving the problem remain unchanged from the deterministic case. However, more calculations are needed as each node must now carry information about the probability of which node would be reached in the next stage.

The recursion used by Kao to solve this problem is the following:

$$R_{(n+1)}(Y) = \max_{n} \{P_{n+1)}(Y \mid X) P_n(X)[T_{(n+1)}(Y \mid X) + C_n(X) + Y]\}$$

where $R_{(n+1)}(Y)$ is the expected total volume which can be captured from state *Y* (a stocking level) at stage $n+1$; *X* is the decision variable. It is the stocking level we choose to be in at stage *n* to get to state *Y* at stage $n+1$. $P_{n+1}(Y \mid X)$ is the probability of being in state *Y* in stage $n+1$ conditional on being in state *X* in stage *n*; $P_n(X)$ is the probability of being in state *X* in stage *n*; $T_{n+1}(Y \mid X)$ is the expected volume of thinning associated with making transition from state *X* in stage *n* to state *Y* in stage $n+1$; $C_n(X)$ is the total expected thinnings associated with being at state *X* in stage *n*; *Y* is the stocking level in state *Y*.

Note that the recursion is composed of two parts. The first part concerns the probabilities of being in the desired states and is obtained by multiplying the probability of being in state *Y* at stage $n+1$ (conditional on being in state *X*) by the probability of being in state *X* at stage *n*. The second is composed of the expected volume including accumulated thinnings contained in state *Y*. These two parts are multiplied together to yield the expected total volume associated with being in state *Y*. The optimal value for *X* is eventually used to generate values for $P_{n+1}(Y)$ and $C_{n+1}(Y)$.

While the method employed by Kao allows for more flexibility in state description than the earlier Markov decision process approaches (it could even be described as quasi-continuous), it would be much more difficult to employ if financial results were desired. The state space would need to be expanded to include number of trees and a proxy for tree size such as basal area. This would expand the elements of risk and more probabilities would need to be included. The "curse of dimensionality" would soon become a factor.

Kao (1984) focuses on optimal stocking levels and rotation under uncertainty. Kao stresses that pure uncertainty, which he defines (page 922) as "the complete ignorance of the probabilities attached to each outcome" does not really exist since we have at least a vague idea of possible outcomes or at least of impossible outcomes. Kao modifies the definition of uncertainty (page 923) to "no probability is known exactly." Kao estimates growth from a growth function derived as the stand develops. Probabilities of future states can also be estimated from the growth function and thus the problem is transformed from one of uncertainty to one of risk and the recursion equation of Kao (1982) can be used. In this process of adaptive optimization, the optimal stocking levels and rotation are recalculated each time more information about the true growth function becomes available. The results show that under uncertainty expected returns are less than under any of the levels of risk used in Kao (1982). The optimal rotation, however, is 70 years, which is longer than the rotation at the highest level of risk which was 50 years. Closer investigation of the uncertainty solution shows that the level of risk is actually less in Kao's uncertainty case than in the highest risk case.

Conclusion

The structure of forest stand optimization problems has made them an ideal candidate for solution using dynamic programming. In particular, the fact that the solutions require a series of interrelated decisions that makes dynamic programming particularly suitable. The fundamentals have been covered and the field is now open for more sophisticated applications of dynamic programming to forestry. These include the financial implications of risk and uncertainty, integrating

[4]The level of growing stock in cubic metres per hectare in a stand

dynamic programming in more models (TASS, for example) and linking dynamic programming stand models with multiple stand models. Although some work has been done by Williams (1976) and Nazareth (1973), more work in this area is required.

References

Amidon, E.; Akin, G. 1968. Dynamic programming to determine optimal levels of growing stock. For. Sci. 14:287-291.

Anderson, F. 1976. Control theory and the optimal timber rotation. For. Sci. 22: 242-246.

Bradley, S.; Hax, A.; Magnanti, T. 1977. Applied mathematical programming. Addison-Wesley. Reading, Mass.

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

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

Brodie, J.D.; Adams, D.; Kao, C. 1978. Analysis of economic impacts on thinning and rotation for Douglas-fir, using dynamic programming. For. Sci. 24:513-522.

Bruce, D.; Demars, D.; Reukema, D.L. 1977. Douglas-fir managed yield simulator. DFIT: User's Guide. USDA Forest Serv. Tech. Rep. PNW-57, Pac. Northwest Forest and Range Exp. Stn. Portland, Oregon. 26 p.

Cawrse, D.; Betters, D.; Kent, B. 1984. A variational solution technique for determining optimal thinning and rotational schedules. For. Sci. 30:792-802.

Chappelle, D.; Nelson, T. 1964. Estimation of optimal stocking levels and rotation ages of loblolly pine. For. Sci. 10:421-502.

Chen, C.; Rose, D.; Leary, R. 1980a. Derivation of optimal stand density over time — a discrete stage, continuous state dynamic programming solution. For. Sci. 26:217-227.

Chen, C.; Rose, D.; Leary, R. 1980b. How to formulate and solve optimal stand density over time problems for even-aged stands using dynamic programming. U.S. For. Serv. Gen. Tech. Rep. 56, North Central For. Exp. Stn. St. Paul, Minnesota.

Clark, C.W.; De Pree, J.P. 1979. A simple linear model for the optimal exploitation of renewable resources. Appl. Math. Opt. 5:181-186.

de Kluyver, C.; Daillenbach, H.; Whyte, A. 1980. A two-stage, multiple objective mathematical programming approach to optimal thinning and harvesting. For. Sci. 26:674-686.

Dixon, B.L.; Howitt, R.E. 1980. Resource production under uncertainty: A stochastic control approach to timber harvest scheduling. Am. J. Agric. Econ. 62:499-507.

Duerr, W.; Christiansen, N. 1964. Exercises in managerial economics of forestry. Oregon State University, Corvallis. 325 p.

Haight, R.; Brodie, J.D.; Dahms, W. 1984. A dynamic programming algorithm for optimization of lodgepole pine management. For. Sci. 31:321-330.

Hann, D. 1980. Development and evaluation of an even- and uneven-aged ponderosa pine/Arizona fescue stand simulator. U.S. For. Serv. Res. Pap. INT-267. Intermount. For. and Range Exp. Stn., Ogden, Utah.

Hann, D.; Brodie, J.D.; Riitters, K. 1983. Optimum stand prescriptions for ponderosa pine. J. For. 81:595-598.

Hillier, F.; Lieberman, G. 1980. Introduction to operations research, 3rd Edition. Oakland, Holden Day.

Hool, J. 1966. A dynamic programming-Markov chain approach to forest production control. For. Sci. Monogr. 12, 26 p.

Kao, C. 1980. A study of optimal timing and intensity of silvicultural practices - commercial and precommercial thinning, fertilization and regeneration effort. Ph. D. Thesis. Oregon State University. Corvallis.

Kao, C. 1982. Optimal stocking levels and rotation under risk. For. Sci. 28:711-719.

Kao, C. 1984. Optimal stocking levels and rotation under uncertainty. For. Sci. 30: 921-927.

Kao, C.; Brodie, J.D. 1979. Determination of optimal thinning entry interval using dynamic programming. For. Sci. 25:672-674.

Lembersky, M. 1976. Maximum average annual volume for managed stands. For. Sci. 22:69-81.

Lembersky, M.; Johnson, K.N. 1975. Optimal policies for managed stands: An infinite horizon Markov decision process approach. For. Sci. 21:109-122.

Martell, D. 1980. The optimal rotation of a flammable forest stand. Can. J. For. Res. 10:3-34.

Martin, G., Jr. 1978. A dynamic network analysis of silviculture alternatives for red pine. Ph.D. thesis, University of Wisconsin, Madison.

McArdle, R.E.; Meyer, W.H.; Bruce, D. 1961. The yield of Douglas-fir in the Pacific Northwest. U.S. For. Serv. Tech. Bull. 210.

McDonough, J.M.; Park, D.E. Jr. 1975. A discrete maximum principle solution to an optimal control formulation of timberland and management problems. PROSE Res. Pap., PROSE, Inc., Los Ange-

leal processes. New York, Interscience.
Mitchell, K. 1980. Distance dependent individual tree stand models. Pages 100-137 *in* K.M. Brown and F.R. Clarke, eds. Proceedings of the workshop on forecasting forest stand dynamics. School of Forestry, Lakehead University, Thunder Bay, Ontario.
Mitchell, K. 1986. Comparison of McArdle, DFSIM and TASS growth and yield models. Pages 350-359 *in* C. Oliver, D. Hanley and J. Johnson, eds. Douglas-fir: stand management for the future. College of Forest Resources, University of Washington, Seattle.
Naslund, B. 1969. Optimal rotation and thinning. For. Sci. 15:446-451.
Nazareth, L. 1973. A resource allocation model that combines Dantzig-Wolfe decomposition and dynamic programming. M.Sc. Thesis. Coll. of Engineering, University of California, Berkeley.
Randall, R. 1977. Financial consequences of commercial thinning regimes in young-growth Douglas-fir. U.S. For. Serv. Res. Note PNW-292, Pac. Northwest For. and Range Exp. Stn. Portland, Oregon, 8 p.
Reed, W.; Errico, D. 1985. Assessing the long-run yield of a forest stand subject to the risk of fire. Can. J. For. Res. 15:680-687.
Reukema, D.; Bruce, D. 1977. Effects of thinning on yield of Douglas-fir: Concepts and some estimates obtained by simulation. U.S. For. Serv. Tech. Rep. PNW-58, Pac. Northwest For. and Range Exp. Stn. Portland, Oregon, 36 p.
Riitters, K.; Brodie, J.D.; Hann, D. 1982. Dynamic programming for optimization of timber production and grazing in ponderosa pine. For. Sci. 28:517-526.
Roise, J. 1986. A nonlinear programming approach to stand optimization. For. Sci. 32:735-748.
Samuelson, Paul A. 1976. Economics of forestry in an evolving society. Econ. Inq. 14:466-492.
Schreuder, G. 1971. The simultaneous determination of optimal thinning schedule and rotation for an even-aged forest. Forest Sci. 17:333-339.
Sleavin, K. 1983. Searching the response surface of stand simulators under different objectives and constraints...DFSIM as a case study. Pages 106-128 in Proceedings of the national silviculture workshop on economics of silvicultural investments. Eugene, Oregon, May 16-20, 1983. USDA Forest Service, Timber Management Branch, Washington, D.C.
USDA Forest Service. 1963. Timber trends in western Oregon and western Washington. U.S. For. Serv. Res. Pap. PNW-5, Pac. Northwest For. and Range Exp. Stn., Portland, Oregon, 154 p.
Wagner, H. 1975. Principles of operations research, 2nd Edition. Prentice-Hall. Englewood Cliffs, New Jersey.
Williams, D. 1976. Integrating stand and forest models for decision analysis. Ph.D. Thesis. Department of Forestry. University of British Columbia, Vancouver.