USING DYNAMIC PROGRAMMING TO EXPLORE HARDWOOD SILVICULTURAL REGIMES

Matthew H. Pelkki[1]

Abstract—Forest managers face a difficult problem when required to estimate future physical and economic returns from timber stands when management strategies can vary by number and timing of thinnings, thinning type, thinning intensity and final harvest age. Field-based testing of these variables provides the most reliable information on stand and individual tree response to these practices, but is limited by the time and cost of these studies. Simulation studies based on growth models which derive their tree and stand response data on field observations are one method of exploring the response curve in greater detail and at a much lower cost in time and effort. Dynamic programming is an ideal method for exploring both stand and individual-tree response to new economic conditions and rotation schemes. Dynamic programming can be used to investigate possible responses to different thinning strategies, intensities, and timing of thinnings. It can also provide insights into management responses to changing markets and prices for wood products. However, simulation approaches are hindered by the weakness of long term (40+ years) growth projections and the ability of the computer systems to accurately model real-world constraints. In this paper, a dynamic programming (DP) based computer optimization technique is used and compared with the results of DP-based studies to previously published field studies of yellow-poplar. The similarities found lend credence to the biological reasonableness of the simulations and the ability of the simulation to provide detailed sensitivity analysis demonstrates its strengths in identifying new research questions for field-based studies.

INTRODUCTION — TIMBER MANAGEMENT

Timber managers face a myriad of questions in planning a stand from regeneration to final harvest. Two of the most basic questions center around rotation length and stand density. These two questions are appropriate whether the objective is fiber or financial maximization. To answer these questions, foresters must pull out the “crystal ball” and make predictions, using some form of growth and yield model to estimate future stand characteristics (total volume, diameter distribution, species composition). If financial maximization is desired, then foresters must also attempt to predict future economic conditions such as the cost of capital, inflation, future stumpage prices, and the costs of harvesting. These data are important for making decisions about sustainable forest-wide management decisions as well as good stand-level decisions. For the forest cover types in the central United States, this usually involves planning and prediction of forest and financial conditions on a time frame from 25 to 100 years.

Two approaches have been applied in an effort to provide forest managers with some guidelines for decision making about rotation length and stand density. The first and oldest technique is field testing. Representative sites are selected and management alternatives are actually implemented and the physical results are measured and reported. The second approach has grown in use since the advent of powerful, low-cost computers. Computer simulation of a forest stand is based upon the use of growth and yield models and this is coupled with economics to provide estimations of future outputs.

The following discussion will compare these two approaches and then provide a detailed description and one example of an implementation of dynamic programming which represents one technique for computer simulation of forest stands. One overall assumption will be that the problem under consideration is an even-aged stand, although all techniques are equally applicable to uneven-aged stand structures.

TWO APPROACHES — COMPUTER SIMULATION AND FIELD TESTING

Necessary Resources for Each Method

Actual field testing of silvicultural regimes requires access to representative sites and the ability to actually implement and then observe the results of management actions. Considerable financial, personnel and equipment resources are needed. This method also requires time for the stand to respond to the management actions and for repeated monitoring of stand parameters. It is assumed that environmental conditions (site, climate) will be typical of the region to which the study applies.

Computer simulation, on the other hand, requires good data on initial stand conditions as a starting point for the simulation. These initial conditions may be a selection of actual representative sites or a hypothetical composite site that is considered typical of the region. Growth and yield models that are responsive to the range of actions and stand conditions simulated are necessary. For example, the TWIGS individual-tree growth simulator for the Central United States (Miner and others 1988) was developed from a data base that poorly represented yellow-poplar, Eastern redcedar, and tupelo/gum. Therefore, this growth simulation would not be an optimal choice for computer simulation of these species in the Central United States. Also, whole stand models such as

[1] Associate Professor of Forest Management and Economics, University of Kentucky, Department of Forestry, Lexington, KY 40546-0073.

Citation for proceedings: Stringer, Jeffrey W.; Loftis, David L., eds. 1999. Proceedings, 12th central hardwood forest conference; 1999 February 28-March 1-2; Lexington, KY. Gen. Tech. Rep. SRS-24. Asheville, NC: U.S. Department of Agriculture, Forest Service, Southern Research Station. 293 p. [Peer-reviewed paper].

yield tables are unable to accurately predict how a stand might grow if the diameter distribution is altered from thinnings from above or below. Predicting future tree quality is also very difficult with whole stand models. Furthermore, computer simulations will need estimates of future financial conditions. Data or additional models that predict stumpage prices, the cost of capital (interest rates), and the cost of harvesting operations must be included.

Strengths and Weaknesses of Each Approach

Table 1 is a list of some major strengths and weaknesses of each approach. Field-based studies have several advantages. They do not rely upon predictions of growth and yield or economic conditions. It is difficult to argue with observed results when properly replicated. However, it should be noted that even observed economic results, when applied in planning future rotations, assume that the future will be the same as the past. The alternatives for management that are chosen can be put into practice with relative ease. Due to the great length of time required and the costs related to acquiring the sites and managing them, field-based studies are expensive. Sensitivity analysis is also difficult and type and number of alternatives is usually kept very low so that adequate replication can be achieved.

Computer simulation, on the other hand, can test a great number of alternatives and combinations. Due to the speed and power of modern personal computers, simulations can test millions of combinations in hours rather than decades and at a far lower cost. However, long-term projections with growth and yield models may be substantially in error (Rauscher and others 1997) and at best might be considered only relatively correct. Another subtle, but important problem with computer simulation is "artifacts." An artifact represents a difference in a result that is not likely to be observed in the real world. For example, Beck and Della-Bianca developed yield equations for unthinned (1970) and thinned (1975) yellow-poplar stands. The two equations are disjoint at any single given age. When these two models are used together in a computer simulation, when a stand is thinned, it immediately shows a slight increase in volume at the same age! This directs the simulation to thin at an early age to capture this "free" volume which, in reality, does not exist. Another form of an artifact may occur from rounding of values in a simulation. For example, it may be possible to test the sensitivity of thinning in basal area increments of one square foot or less, but would the results be practicable in the real world?

The remainder of this paper presents the basics of one computer simulation technique, dynamic programming (DP). A non-mathematical description of DP will be presented, followed by a mathematical formulation and then a description of a DP-based computer program (NESTER). The results of published research on field-based studies of yellow-poplar will be compared to some NESTER results.

OPTIMIZATION USING DYNAMIC PROGRAMMING

Dynamic programming is a computational method that takes large sequential problems and breaks them down into solvable, related subproblems which can be linked together to achieve an optimal solution to the entire problem. Let's think about a 20-year-old even-aged hardwood stand. Let us further assume that management operations will occur only once every 10 years, and that the alternatives available are, to do nothing, thin 20 percent or 40 percent of the volume, or clear fell the entire stand (regeneration harvest). Figure 1 shows a diagram with a node or state representing our 20-year-old stand. The state is described by the total cubic foot volume per acre, in this case 700 ft^3/acre. Over the 10 years, this stand will grow 700 ft^3/acre and will have a total of 1400 ft^3/acre. If we do nothing, the stand remains at 1400 ft^3/acre. If we thin, the stand volume will be reduced to either 1120 ft^3/acre (20 percent thinning) or 840 ft^3/acre (40 percent thinning). If we clear fell the stand it has no residual volume. From one state in the initial stage (age 20), we now have four possible states in the first stage. Figure 2 shows a partial diagram of the progression into the second stage. Note that two paths compete for the state described by 1600 ft^3/acre. As each stage continues, we will have more states in the decision space and as that happens, it is more likely that two or more paths will

Table 1—Strengths (+) and weaknesses (–) of field-based studies and computer simulations

Field-based	Computer Simulation
+ Long term growth and yield valid + Economic data is observed, not predicted + Alternatives selected are relatively easy to put into practice – Costly to implement – Difficult to test large numbers of alternatives and combinations/intensities – Long time for study completion	– Long term problems with growth and yield models – Long term economic assumptions can be invalid – Model may reflect artifacts based on structure of model + Inexpensive to test + Excellent sensitivity analysis + Fast results

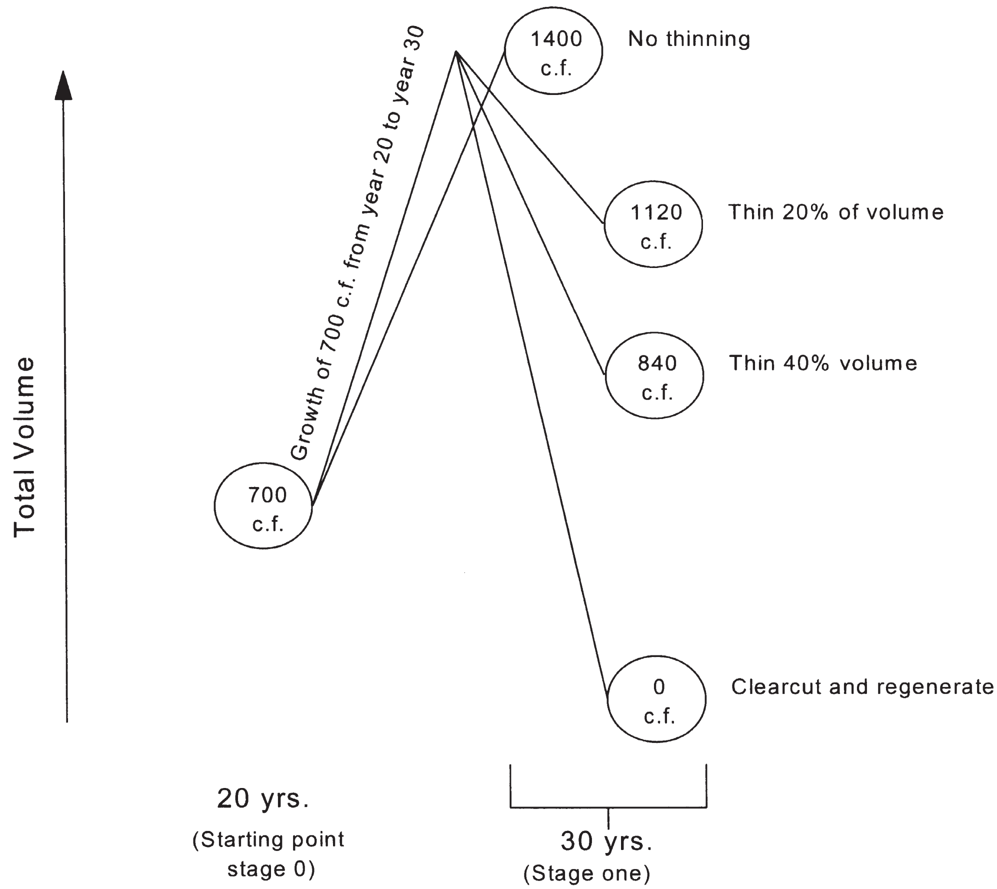

Figure 1—Initial condition and stage one states for forward-recursive dynamic programming network.

compete for a given state. The underlying assumption in dynamic programming is that any two or more paths that compete for the same state represent an identical forest condition with the same future value growth potential. If we assume that it doesn't matter how we achieve 1600 ft^3/acre in a 30-year-old stand, this assumption holds true. This is called the principle of optimality (Dykstra 1984). It is also called the memoryless principle and it means that an optimal path through the network of all possible states can be solved one stage at a time.

More formally, the subproblems are often referred to as stages, and within each stage are multiple states. A recursive relationship links the stages together. This recursive relationship contains a method for choosing optimal states within a stage and for linking optimal stage policies together into an optimal problem solution. To do so, certain information must be stored at each state. This includes 1) a link to a state in an adjacent stage, 2) the objective function value for reaching this particular state, and finally, 3) information necessary to describe the state and make a decision regarding options for moving it to the next adjacent state and determining its value relative to the objective function (Figure 3). The formal mathematics for dynamic programming are provided in the next section.

Dynamic Programming Formulated Mathematically

The objective function shown in equation one relates the value of N management decisions, $f_N^*(Y_N)$, which yield a stand described by Y_N (a regeneration harvest in even-aged stands). A management decision is defined by the variable T_n.

$$f_N(Y_N) = \sum_{n=0}^{N} r_n(T_n) \tag{1}$$

Equation two relates the current state of a stand in stage n to a state in the next stage n+1 ($Y_n \rightarrow Y_{n+1}$) by taking the current stand condition (Yn) and adding the next current growth ($G_{n+1}(Y_n)$) and subtracting out any intermediate harvests (T_{n+1}).

$$Y_n + G_{n+1}(Y_n) - T_{n+1} = Y_{n+1} \quad (n = 0, 1, 2, \ldots, N-1) \tag{2}$$

Equation three simply relates the volume of a stand after it has grown (X_n) to its final volume in stage n (Y_n) by showing the difference as the amount of volume harvested (T_n).

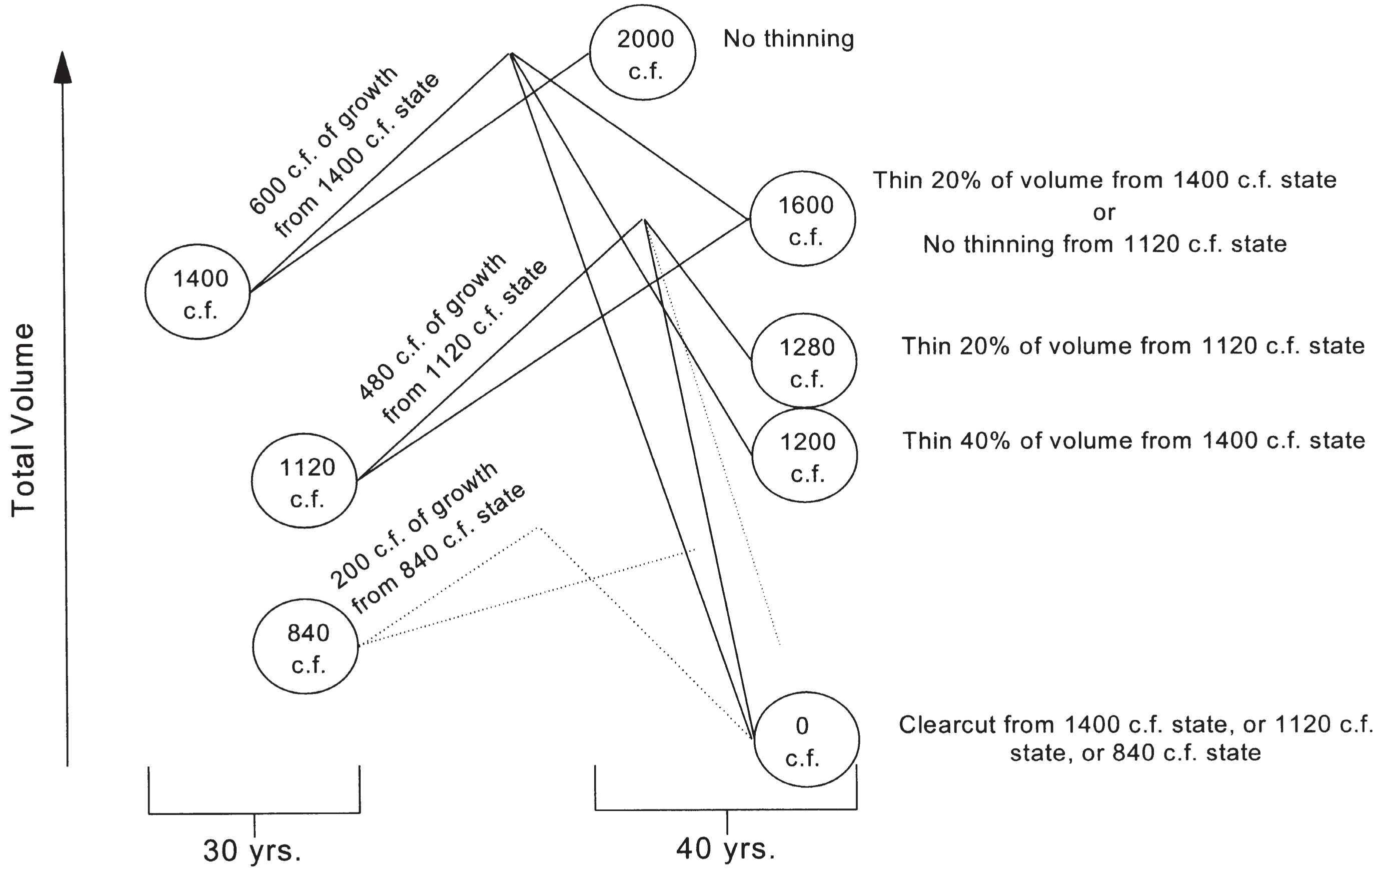

Figure 2—Stages one and two of a dynamic programming network.

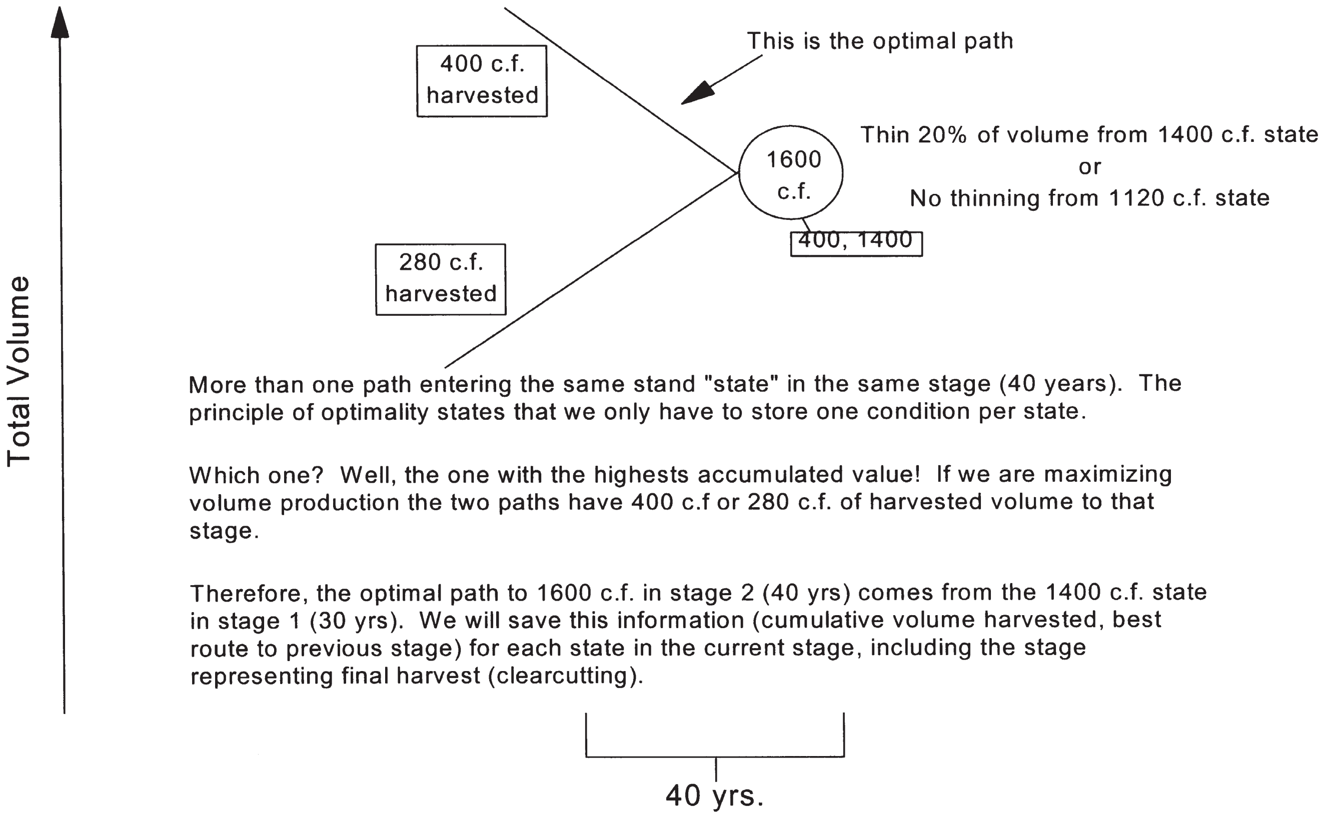

Figure 3—Competition for nodes in the DP network and tracing the optimal path.

$$X_n - T_n = Y_n \qquad (n=0,1,2,...,N) \tag{3}$$

Equation four assures that the final stand condition is a clearcut. If uneven-aged management is applied, then a different ending condition, one with a residual volume, is defined.

$$X_N - T_N = 0 \tag{4}$$

Equation five is the recursive relationship that links each stage together. The function $r_n(X_n,T_n)$ is the return generated at stage n from decision T_n taken on stand described by X_n. The function r can maximize any quantifiable criteria but is usually a maximum volume or value function.

$$f^*_n(Y_n) = \max_{(Y_{n-1},T_n)} [r_n(X_n,T_n) + f^*_{n-1}(Y_{n-1})] \tag{5}$$

The variables are defined again for the reader:

$f^*_N(Y_N)$ = objective function value of N management decisions yielding a stand described by Y_N. In the final stage, N, Y_N is a clear cut stand.

r_nT_n = return generated at stage n by decision T_n.

X_n = state vector describing stand at stage n after it grows from state Y_{n-1}.

Y_n = state vector describing stand at stage n.

T_n = management decision taken at stage n.

$G_{n+1}(Y_n)$ = growth of stand at stage n to stage n+1. Along with T_n, this constitutes the transformation function.

$r_n(X_n,T_n)$ = return generated at stage n from decision T_n taken on stand described by X_n.

Some Comments About Dynamic Programming

A single variable such as total volume is usually insufficient to adequately describe a stand. Two stands, both having a total volume of 1000 ft^3/acre may have very different average diameters and thus different future growth and value potential. Therefore, two or more state variables are used to define a stand condition within a stage. However, this leads to another problem known as the curse of dimensionality (Arthaud 1986). State variables used to describe forest stands are usually continuous. If the variables cubic volume per acre and number of trees per acre are used to define states within a stage and it is assumed that the maximum volume is 5000 ft^3/acre and the maximum number of trees is 500/acre, then the DP program must be able to represent 2.5 million states per stage (assuming that all volumes and trees per acre are rounded to the nearest integer). If yet another state variable is added, the memory requirements are likely to exceed the primary memory capacity of most University mainframe computers. To effectively use continuous variables and to reduce the size of the solution network, the idea of state neighborhoods (Brodie and Kao 1979) was introduced. For example, state neighborhoods of ± 5 trees per acre and ± 5 ft^3/acre would reduce the two state memory requirements from 2.5 million to 100,000 states per stage. Figure 2 demonstrates an example of state neighborhoods. If the neighborhood interval was 100 ft^3/acre, then the paths leading to 1280 ft^3/acre and 1200 ft^3/acre would compete for the same node in the network. The key in using state neighborhoods is making sure that your intervals are not so large as to violate the principle of optimality (Pelkki 1997).

AN EXAMPLE OF DP AND FIELD TESTING - YELLOW POPLAR MANAGEMENT

Description of NESTER

Nester (NEighborhood STate EvaluatoR) is a forward recursive dynamic programming computer program for Windows-based personal computers (Pelkki 1997). It utilizes the GROW subroutine (Brand 1981) from the TWIGS individual-tree growth projection system (Miner and others 1988). NESTER allows the user to choose any combination of six state variables (cubic foot volume per acre, number of trees per acre, average diameter, basal area per acre, number of thinnings, and average tree grade per acre). Stage intervals can be as short as one year or as long as 30 years. As NESTER utilizes an individual-tree growth model, stumpage prices must be provided for each species group included in the stand. NESTER can be used to model mixed species stands, however, at the present time, there are no thinning algorithms that select by species. NESTER does project tree grade using the method presented by Yaussy (1993). Also, because of the individual-tree growth model, NESTER can simulate many types of thinning operations, including mechanical thinning, thinning from above, thinning from below, and thinnings based on tree quality (improvement thinning and high-grading). NESTER can simulate constant, real price increases, or stochastic price changes independently for sawtimber and other roundwood size classes.

NESTER Studies on Yellow Poplar

Two studies have been published to date using NESTER to study yellow-poplar (Pelkki and Arthaud 1998, Pelkki 1999). One study focused upon changing markets and prices and their effects on yellow-poplar management, the other examined changing thinning strategies and the impact on financial returns from yellow- poplar. The results of these studies indicate the strength of dynamic programming for sensitivity analysis and the overall findings of these papers will be compared to published results of field studies reported in Beck and Della-Bianca (1981). The complete results of those studies are not included here. This paper merely highlights some of the new information obtained through computer simulations and identifies some areas where additional field-based work is needed.

Input Data for NESTER

Initial stand conditions were derived from published diameter distributions (McGee and Della-Bianca 1967, Knoebel and others 1986) for 20-year-old, average-stocked, yellow-poplar stands with site indices of 90, 110, and 130 ft at 50 years. Potential grade distributions (Hanks 1976) representing low, average, and high stem quality classes were defined using U.S. Forest Service Forest Inventory and Analysis data for yellow-poplar in the region. Thus, nine different combinations of site index and stem quality distribution served as starting points for the DP simulations.

Stumpage prices were obtained from regional price reports and from Timber Mart South (1994). Capital costs were generally set at 4 percent. Sensitivity to real price increases and higher and lower costs of capital were investigated.

NESTER runs were completed using two states, number of trees per acre and cubic foot volume per acre, with the state neighborhoods of 10 trees per acre and 10 ft^3 per acre, respectively. The stage interval was set at 2 years, but additional runs were investigated using different stage intervals. All thinning options but mechanical thinning were simulated (thinning from above, thinning from below, thinning from above and below, high-grading, and improvement thinning) at intensities ranging from 10 percent to 50 percent of the initial basal area in 5 percent increments. Thus, from each state, 47 options were simulated. The state network represented 46^n+1 states (where n = stage). The initial stage was age 20, and each stage increased by two year increments. Therefore, stage 11 (age = 42) represented $46^{11}+1$ possible ways to manage a stand from age 20 to age 42.

Comparing Results of NESTER-Based Studies to Previous Work

When thinning yellow-poplar stands, Beck and Della Bianca (1981) note that cultural work in sapling and pole-sized stands is very costly and there are few markets for this material. When thinnings are done, Beck and Della-Bianca (1981) recommend thinnings from below to concentrate the value on larger, high-value stems. They also noted that intermediate yellow-poplar trees would respond to thinning from above. In a recently completed DP study, Pelkki (1999) found that improvement thinnings, while earning negative to very low initial returns (-$5.1 to $31.5 per acre) were commonly part of the optimal financial stand regime. These improvement thinnings removed first large trees of low quality, then very small trees, and finally low to higher grade factory sawtimber grade trees until a basal area target was reached. In this manner, the trees with the most valuable future growth potential were retained.

In stands with high-quality stem distributions, the DP-based studies (Pelkki and Arthaud 1998, Pelkki 1999) included thinnings that removed 40-50 percent of the basal area. Furthermore, some of the optimal financial regimes included 4-6 thinnings (not all at the 40-50 percent intensity level). Beck and Della-Bianca (1981) report considerable leeway in manipulating yellow-poplar stocking levels to achieve diameter growth and quality goals without sacrificing volume production. However, the frequency and intensity of the DP-based harvest may initiate advance regeneration which would increase the final harvest and site regeneration costs.

In studies reported in Beck and Della-Bianca (1981), first thinnings are recommended as early as 15-20 years and can be repeated every 5-15 years over the rotation. The DP-based studies (Pelkki and Arthaud 1998, Pelkki 1999) have an initial stand state that is 20 years of age and so cannot simulate thinnings before the age of 22 years (1st stage). However, in the DP-studies, most initial thinnings occurred between the ages of 22 and 28 years. Multiple thinnings in the DP-based studies ranged from 2-12 years apart with most being 4-10 years apart.

In regimes with multiple thinnings, Beck and Della-Bianca (1981) recommend that later thinnings be lighter because basal area growth response in older stands is lower. The DP-studies (Pelkki and Arthaud 1998, Pelkki 1999) had heavier later thinnings, possibly due to the financial, rather than volume objectives of the simulation. Both Beck and Della-Bianca (1981) and the DP-based studies (Pelkki and Arthaud 1998, Pelkki 1999) found that intense thinnings shortened the rotation.

While not explicitly discussed in Beck and Della-Bianca (1981), the yield tables suggest an optimal rotation of 50-70 years for fiber production. The DP-based studies (Pelkki and Arthaud 1998, Pelkki 1999) found maximum financial returns with rotations as short as 32 years or as long as 66 years. Factors contributing to shorter rotations were a good market for pulpwood, OSB, and other fiber-based products, high interest rates, and higher site indices. Factors contributing to longer rotation lengths were high stem quality distributions, lower interest rates, and higher sawtimber prices (or an absence of sub-sawtimber markets). These factors favored longer, sawtimber-focused rotations.

For yellow-poplar management, the DP-based studies uncover three major issues that need additional research in a field-based setting. First and foremost, can improvement thinnings, based on potential tree grading bring the economic returns projected by the computer simulations? Additional studies have shown (Pelkki and Ringe 1998) that the earlier a valid potential tree grade can be applied, the greater the economic returns. Secondly, are harvest costs in poletimber stands prohibitive? The DP-simulations assume a fixed entry cost on all thinnings and a 15 percent thinning penalty on all stumpage prices (to reflect the added cost of thinning over a final harvest). They also charge harvest costs for pre-merchantable stems at the rate of 10 percent of the merchantable price (if the stems were of merchantable size). Studies by Kluender and others (1996) suggest that harvest cost of small diameter stems are greater than proportional to volume in southern pine stands. With the increase in markets for hardwood pulpwood throughout the central Appalachian region, thinning and site preparation cost studies would appear necessary. Finally, the DP-based simulations suggest that multiple thinning regimes with 4-6 entries prior to final harvest may optimize financial returns in some yellow-poplar stands. Some field-based tests of such regimes would answer questions related to problems in stand integrity and advance regeneration treatment costs.

SUMMARY - WHY BOTH APPROACHES ARE NEEDED

Unquestionably, both field-based and simulation-based approaches are necessary for forest researchers in the future. Field-based studies provide on-the-ground confirmation of new practices and working methods for field application. They cannot, however, respond quickly to new markets or changing economic conditions. While computer simulations can perform thorough sensitivity analyses to many variables, they lack the real world operational constraints that are often too complex to model and may even be unanticipated by the researcher. Computer simulations are best left to exploratory research with their results confirmed by field-based studies which can then lead to changes in actual management practices.

REFERENCES

Arthaud, G.J. 1986. Economic comparisons of thinning from above and below in loblolly pine plantations using dynamic programming. Blacksburg, VA: Virginia Polytechnic Institute and State University. 139 p. M.S. thesis.

Beck, D.E.; Della-Bianca, L. 1970. Yield of unthinned yellow-poplar. Res. Pap. SE-58. U.S. Department of Agriculture, Forest Service. 20 p.

Beck, D.E.; Della-Bianca, L. 1975. Board-foot growth and yield of yellow-poplar after thinning. Res. Pap. SE-123. U.S. Department of Agriculture, Forest Service. 20 p.

Beck, D.E.; Della-Bianca, L. 1981. Yellow-poplar: characteristics and management. Agric. Handb. 583. U.S. Department of Agriculture, Forest Service. 91 p.

Brand, G.J. 1981. “GROW” - A computer subroutine that projects the growth of trees in Lake States’ forests. Res. Pap. NC-207. U.S. Department of Agriculture, Forest Service.

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

Dykstra, D.P. 1984. Mathematical programming for natural resource management. McGraw-Hill. 318 p.

Hanks, L.F. 1976. Hardwood tree grades for factory lumber. Res. Pap. NE-333. U.S. Department of Agriculture, Forest Service. 81 p.

Kluender, R.; Lortz, D.; McCoy, W. [and others]. 1996. Harvesting time, costs, and productivity in Arkansas. In: Caulfield, J.; Bullard, S., eds. Proceedings of 1995 southern forestry economics workshop; April 17-19; New Orleans, LA: 260-274.

Knoebel, B.R.; Burkhart, H.E.; Beck, D.E. 1986. A growth and yield model for thinned stands of yellow-poplar. For. Sci. Monograph 27. 62 p.

McGee, C.E.; Della-Bianca, L. 1967. Diameter Distributions in natural yellow-poplar stands. Res. Pap. U.S. Department of Agriculture, Forest Service. 7 p.

Miner, C.L.; Walters, N.R.; Belli, M.L. 1988. A guide to the TWIGS program for the North Central United States. Gen. Tech. Rep. NC-125. U.S. Department of Agriculture, Forest Service.

Pelkki, M.H. [In press]. Impacts of thinning strategies on financial returns from yellow-poplar stands. Southern Journal of Applied Forestry.

Pelkki, M.H. 1997. The effects of neighborhood storage size on dynamic programming solutions. Forest Science. 43(3): 387-395.

Pelkki, M.H.; Arthaud, G.J. 1998. Evaluating changes in yellow-poplar management in response to new low-end markets. Northern Journal of Applied Forestry. 15(1): 37-42.

Rauscher, H.M.; Young, M.J.; Webb, C.D. 1997. Testing the accuracy of growth and yield models for southern hardwood forests. Unpublished manuscript. Asheville, NC: U.S. Department of Agriculture, Forest Service, Southern Research Station, Bent Creek Research and Demonstration Forest.

Timber Mart South. 1994. First and second quarter reports. Athens, GA: University of Georgia. Vol. 19.

Yaussy, D. 1993. NE-TWIGS update: incorporating tree quality. The Compiler. Vol. 11(1): 8-12.