

LINEAR PROGRAMMING AND RELATED APPROACHES IN FOREST PLANNING*

OSCAR GARCÍA

Forest Research Institute, Ministry of Forestry,
Private Bag 3020, Rotorua, New Zealand

(Received for publication 19 July 1990; revision 5 December 1990)

ABSTRACT

A number of approaches to forest management planning are surveyed and analysed. A unified treatment of the various models and optimisation techniques is attempted, aiming at a better understanding of their essential characteristics and interrelationships. Ways of improving computational efficiency are examined. Structures for the forestry submodel are classified into three basic types, and interpreted in terms of network flows. A utilisation / economics submodel, comprising an objective function and constraints representing supply requirements, yield regulation, terminal conditions, and other aspects of the management problem, completes the formulation. A new criterion for yield regulation, based on irreversible investment concepts, is discussed in detail. The optimisation methods proposed in the literature for solving forest planning problems are evaluated, together with other potentially useful techniques and strategies. Topics covered include the use of standard Linear Programming packages, the prospects of decomposition and partitioning methods, the potential for utilising oldest-first and other cutting priorities, the explicit use of time structure in Dynamic Linear Programming and Optimal Control Theory algorithms, and the Hoganson-Rose procedure and its relationship to Lagrangian relaxation.

Keywords: forest management; timber harvest scheduling; forest estate modelling; Operations Research; mathematical programming; control theory; forest economics.

INTRODUCTION

I will consider the problem of planning the management of forest estates, that is, of aggregates of forest stands at the forest, enterprise, regional, or national levels. The stands are the basic components, areas of forest taken as homogeneous for management purposes. Only even-aged management, where a stand contains trees of the same or similar ages, will be discussed. I will also focus largely on intensively managed production forests. In this context, forest management planning deals mainly with decisions about when and where to harvest and plant, and about what silvicultural regimes to use. A silvicultural regime

* Paper originally presented at IFORS'90—12th Triennial Conference of the International Federation of Operations Research Societies, Athens, Greece, 25–29 June 1990.

specifies stand treatments, including initial spacing, timing and intensity of thinnings (partial cutting), pruning, fertiliser treatment, and application of pesticides (García 1986, 1988).

The most commonly used techniques in forest planning are simulation and Linear Programming (LP). The best-known LP-based forest planning systems are probably Timber RAM (Navon 1971), MAXMILLION (Ware & Clutter 1971), and FORPLAN (Hoekstra *et al.* 1987). Other proposed methods include models with nonlinear objective functions, approximate heuristic techniques, and algorithms using ideas based on Control Theory.

Johnson & Scheurman (1977) reviewed and analysed many of the available approaches. Their paper has become a classic, and their grouping of forest LP models into Model I and Model II forms is widely used. Building upon Johnson & Scheurman's work, I suggest a revised classification that seems useful for comparing a wide variety of existing models, and for guiding future developments. This and other problem formulation issues are presented in the next section.

In the section on Solution Strategies, possible methods for computing optimal solutions are discussed. This section requires the reader to be familiar with Mathematical Programming techniques, and would be of more interest to specialists in forest planning, operations research, or computer science.

Over-all, the emphasis is on a better understanding of the essential characteristics of the various approaches and their interrelationships.

PROBLEM STRUCTURE

It is convenient to analyse separately two components of the problem formulation: a model of the forest sub-system, representing development of the forest and forest management decisions, and a model of the utilisation/economics sub-system, representing the "rest of the world". I will use a somewhat simplistic view of an even-aged forest, avoiding excessive generalisations. The formulation is just sufficiently general to allow a unified treatment of the commonly used approaches.

The Forestry Submodel

This submodel describes the physical aspects of forest management decisions and of the forest response.

The forest (or forest estate) consists of a number of cutting units or stands, areas that are considered homogeneous. At any point in time, a stand is characterised by its age. A stand can be subjected to various sequences of silvicultural treatments (thinning, pruning, fertiliser), resulting in streams of inputs and outputs during its life, and at the time of harvesting. In these models we will assume a unique silvicultural regime for each stand. Alternative regimes are considered as originating different stands.

Optionally, a number of stands may be grouped into a timber class or crop-type. Stands in a crop-type have the same silvicultural treatment, and the same potential stream of inputs and outputs. They may differ in their ages and times of harvesting.

Time is measured in periods of uniform length. Ages are specified as age-classes, with widths equal to the period length.

Model A: a state space view

Here the state of the forest at the start of each period is described by the area in each crop-type and age-class. During a period some of the area in each class may be cut, and the remaining area moves into the next age-class at the beginning of the next period (Fig. 1). The harvested areas are replanted (or regenerated) immediately into the same or different crop-types. Certain linear combinations of the harvested and residual areas constitute inputs and outputs in the utilisation submodel (yields, revenues, costs, use of resources).

The model is not as restrictive as it may appear at first sight. One or more “bare land” crop-types can account for new land of various qualities available for planting, or left unplanted. Any unavoidable regeneration delays can be represented through dummy age-classes. The aggregation of stands into crop-types is optional; each stand could constitute a unique crop-type. Transfers of area between crop-types can be allowed in order to model such things as alternative silvicultural regimes or changes of ownership.

This model can be interpreted in terms of the state space approach of Systems Theory. Any system that evolves in time can be described by a state that characterises the system at some point in time, and a transition function that specifies how the state changes over time (e.g., García 1979, 1988, and references therein). It can be seen that the state of the forest (areas by crop-type and age-class) at the start of period $t+1$ is a linear function of the state

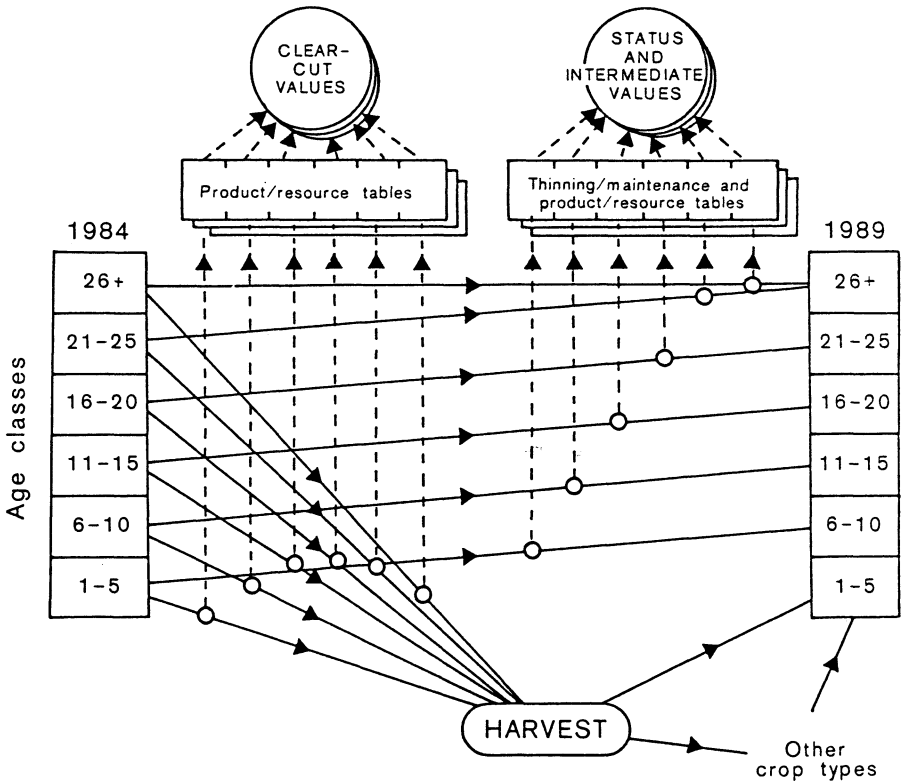


FIG. 1—Basic structure of a forest planning problem (from García 1984)

at t and the actions (areas cut and planted) in period t . This point of view is pursued further in the section on “Time Structure” below.

A visualisation in terms of network flows is useful. We represent the age-classes in Fig. 1 by network nodes, linking the flow of areas across periods. An example with three time periods is given in Fig. 2. In the example the crop-type initially has area only in the second age-class (it might represent a single stand). The doubled lines represent more than one network arc. Flow is conserved at each node, with the sum of the areas going into the node equalling the sum of areas going out. These relationships can be written as LP constraints, one for each node.

Without loss of generality, the age-classes are shown as completely cut after the last period. It will be convenient to add age-classes as needed, instead of having an open-ended oldest class as in Fig. 1. Note that the initial area can be transferred to other crop-types that might represent alternative silvicultural regimes. For optimisation it is sufficient to allow these transfers at the beginning; in simulators these transfers could be allowed in any period (García 1981, 1984).




FIG. 2—Model A network. All arcs are directed from left to right

Model B: an extension of Johnson & Scheurman’s Model II

In Fig. 2 the age-class nodes have at most one incoming arc (except for the initial ones). The flow on this arc can be expressed as the sum of the flows on the outgoing arcs, and this relationship used to eliminate the arc and node from the network (Fig. 3). Applying this simplification repeatedly to Fig. 2 results in Fig. 4. For consistency, the network has been extended to the left with nodes for periods in the past. An alternative interpretation is that this network shows the possible paths between planting and harvesting nodes.

The curved arc joining period i and period j represents the area planted in period i harvested in period j . The straight arcs correspond to planting, except for the double lines in past




FIG. 3—Network simplification step




FIG. 4—Model B network

periods, which are possible inter-crop-type transfers. The areas left standing are not represented explicitly in this network, but can be obtained from the equations in Fig. 3.

The flow conservation constraints in this network are essentially those in Johnson & Scheurman’s Model II, except for the addition of transfers and replanting across crop-types. For each period, including a number of periods in the past equal to the oldest initial age-class, there are two nodes, a “harvesting” node and a “planting” node (Fig. 5). In the Figure, x denotes flows of forest areas, r flows of bare land, and a are initial areas. The indices i and j indicate crop-types, while s and t are time periods. Obviously, not all the arcs shown are relevant for every node, and some nodes may be redundant.

Model C: Johnson & Scheurman’s Model I

In Fig. 2 and 4 an initial hectare of forest can follow many different paths from a “source” to a final “sink” node. Tracing the possible paths we obtain something like Fig. 6. Only a few of the paths that move across crop-types are shown. This corresponds to a chain decomposition of the network flow (Dantzig 1963, Section 19-1, Theorem 2).

In this model the decision variables are the number of hectares sent from each source through each path. There is only one flow constraint for each source. The paths correspond to harvest/planting sequences. As in Models A and B, each of these generates a sequence of inputs and outputs for the utilisation submodel.

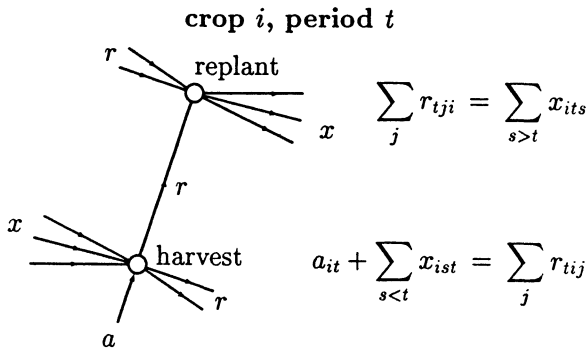


FIG. 5—Building blocks for Model B




FIG. 6—Model C network. Most inter-crop paths are not shown

Discussion

It is clear that Models A, B, and C are essentially equivalent in their ability to model forest planning problems. In particular, Model B does not inherit the drawbacks of Model II compared to Model I often mentioned in the literature. According to Johnson & Scheurman, “Model I preserves intact the hectares that form any beginning management unit (age class) through the entire planning horizon”, while “Model II allows hectares from beginning age classes ... to be broken up and combined with hectares from other age classes ...” (see also Barros & Weintraub 1982; Gunn & Rai 1987). Actually, in general the initial stands in Model I (or C) may be subdivided, with parts of their area allocated to different harvest/planting sequences (Fig. 6). It is true, though, that in Model I different stands do not merge on harvesting, as usually happens with stands (age-classes) within a crop-type in Models II or B. However, the aggregation of stands into crop-types in Model B is optional, so that the identity of the stands can be maintained if so desired. García (1984) showed how to handle thinnings and other treatments in Model B.

Model A is perhaps the most natural, and versions of it have been used by foresters for many decades in manual forest planning procedures, and later in computer simulations (García 1981, 1984, 1986). The state/transition function interpretation is used in the Control Theory and Dynamic LP approaches.

Model B results in a much smaller number of constraints (nodes) and variables (arcs) in LP formulations than Model A. One must be wary, however, of the common misconception that the LP computing effort depends primarily on the number of constraints. Modern large-scale LP codes are very efficient at exploiting sparsity, and their computing time is most sensitive to the number of nonzeros in the LP matrix. The reduced problem size in Model B is achieved by eliminating the residual (standing) area variables, expressing them in terms of the areas harvested (Fig. 3), but this substitution can reduce the sparsity of the utilisation submodel constraints. In addition, an LP formulation based on Model A can include the residual areas as free variables; they are forced to be non-negative by the other constraints (free variables always stay in the basis, and can be handled with much less computational effort than the others). Nevertheless, in most situations the utilisation constraints deal mainly with harvesting variables, and the reduction in problem size is likely to be advantageous.

Model C has been the most commonly used in LP-based forest planning systems, at least until recently. It results in fewer constraints than either A or B, but the LP matrix tends to be dense. Also, it is obvious that the number of variables can be very large (Fig. 6). Often, only a small subset of the feasible harvesting/planting sequences is considered, accepting the possibility of suboptimisation. Eriksson (1983) showed how to generate these sequences as needed in the course of the simplex algorithm, using column generation techniques. Essentially, the column generation steps search for paths in the Model A or B networks, demonstrating the similarities between the three models.

It is not difficult to see how these models can be extended to variable period lengths, at least when the period length and age-class width are successively doubled. Barros & Weintraub (1982) used various period lengths (in a Model I formulation) in order to improve precision in the initial periods, while saving computational effort in the less critical later periods.

The Utilisation Submodel

The "utilisation" submodel represents the environment within which the forest manager operates. It includes the economic and financial aspects, markets and/or processing plants for the timber, and the resources needed for managing the forest, such as labour, logging machinery, and roading.

In a Mathematical Programming approach this submodel consists of an objective function that must be maximised or minimised, and a set of constraints that are added to the forestry submodel constraints.

Objective function

In most instances the objective is the maximisation of a discounted cash flow. Sometimes the objective is to minimise the discounted total cost of supplying some utilisation plant. Occasionally, the planners may prefer to maximise total physical production, i.e., a zero discount rate is used.

Discounting implies a weighted sum of inputs and outputs for each period. The weights are exponentially decreasing functions of time of the form α^t , where t is time and α is a constant discount factor. An interesting alternative has been proposed by Harvey (1986). In

agreement with the intuitive feeling of many foresters, he suggested that exponential discounting undervalues the future too strongly, and argues for weights of the form $(1/t)^{\beta}$.

Yield regulation

In a perfect market, with unit costs and prices not affected by production levels, the optimum would be achieved by optimising independently the management of each individual stand. With an irregular age-distribution in the forest estate, this can result in large fluctuations in the total production over time. Similar fluctuations would appear in the revenues, and in the labour, equipment, and finance requirements for harvesting and silvicultural operations. It is generally accepted that in most practical situations this is undesirable, and various devices are used in forest planning models to smooth-out production levels.

A major reason for smoothing is that costs and revenues, assumed to be constant in the single-stand analysis, are actually affected by the scale of production. There are costs associated with expansion of production/processing capacity, and with maintenance of unused capacity. In addition, there may be constraints arising from such things as contractual commitments (García 1986).

Integrated models: One of the reasons for regulating production may be the need to maintain reasonable levels of wood supply to processing plants. Perhaps the most satisfactory approach is then to model in some detail the processing subsystem and its interactions with the forestry subsystem. Examples of this have been given by Barros & Weintraub (1982) at the enterprise level, and Kallio *et al.* (1981) at the national level. These models tend to be problem specific, and the development effort and information requirements can be considerable. In addition, often the problem details are not sufficiently well defined, as in many indicative planning studies. It may then be appropriate to use the simpler yield regulation techniques discussed below.

Apart from yield regulation, more realistic modelling of some aspects of the utilisation system can improve the quality of the management decisions. For example, it can be relatively easy to add a log allocation and transport model (Manley & Threadgill 1987).

Single-period conditions: A common practice is to include in the model simple upper and/or lower bounds on inputs and outputs for each period. These may be hectares harvested or planted, volume production, net cash revenue, and can represent industry demands or operational limitations (Clutter *et al.* 1983).

Another view assumes that the undesirability of fluctuating production levels is due to nonlinearities in the relationship between revenue and production within each period. Price is then expressed as a function of the quantity produced (a demand function). The result is a separable nonlinear objective, quadratic if the demand function is linear (Johnson & Scheurman 1977). The problem can be linearised using separable LP techniques (Hrubec & Navon 1976), or approached by other methods (Walker 1976; Lyon & Sedjo 1986).

Inter-temporal constraints: Constraints involving pairs of consecutive periods are frequently used. Typically, the absolute or percentage decrease (and/or increase) in the total volume harvested between one period and the next is constrained to be not greater than some given

value. A popular special case is the condition of nondeclining yield, where the volume harvested in one period is constrained to be not less than that harvested in the previous period. Nondeclining yield conditions have been imposed by law in the management of public forests in the United States.

More generally, constraints of this type may be specified for each of a number of products, or even for inputs such as employment levels. The constraints may also be enforced for groups of crop-types, e.g., forests within a region (Manley & Threadgill 1987).

A capacity expansion approach: Although sometimes nonlinearities in single-period net revenues might be significant, more often it is the changes in production from one period to the next that are important. Arguments against large changes include the need to supply existing industry, the maintenance of employment levels, and the lack of resources for coping with large temporary increases in production. Note that difficulties are caused mainly by production reductions, a motivation for the nondeclining yield criterion. Nondeclining yield at all costs, however, can be too strong a constraint, and ad-hoc bounds on period-to-period variations are difficult to justify.

A more satisfactory model can be based on the idea of capacity expansion. I will introduce a very simple model, some aspects of which may be considered somewhat unrealistic. It can be extended and generalised, but the purpose here is to present an alternative only slightly more complex than nondeclining yield.

Let us measure production by the yield of some product, e.g., roundwood volume. In a Model B setting, the yield in period t is

$$y_t = \sum \sum v_{ist} x_{ist} ,$$

where x_{ist} is the area in crop-type i planted in period s and cut in period t , and v_{ist} is the corresponding yield per hectare (usually a function of the age $t - s$). Assume that for each period there is a production capacity z_t that cannot be exceeded:

$$y_t \leq z_t .$$

This capacity may reflect the availability of such things as wood-processing facilities, harvesting machinery and manpower, roads and transportation. In any period the current capacity can be increased, at a cost:

$$z_t = z_{t-1} + w_t, \quad w_t \geq 0.$$

Including the expansion costs in the objective function completes this utilisation submodel:

$$\text{maximise } \sum \alpha^t p_t y_t - c \sum \alpha^t w_t .$$

Here α is the discount factor, the p_t are net marginal revenues, and c is the expansion cost per unit of capacity.

Note that in the LP the z_t can be free variables, unconstrained in sign; the other constraints force them to be non-negative. We could also eliminate these variables, and the last set of constraints, with the substitution

$$z_t = z_0 + \sum_1^t w_s ,$$

where z_0 is the initial capacity.

An alternative derivation, that might provide additional insight and suggest appropriate values for the expansion costs, is presented in Appendix 1. Clark *et al.* (1979) discussed a related model for fisheries, with an infinite horizon, in continuous time, but with only one state variable.

It is easy to see ways of extending the model. The option of reducing capacity is easily added, substituting

$$w_t = w_t^+ - w_t^-,$$

and assigning to the capacity reduction w_t^- a scrap value lower than the capacity expansion cost for w_t^+ . Depreciation, or loss of capacity over time, can also be included (Clark *et al.* 1979). Yields and capacities associated with various products and resources could be considered. All this, however, adds to the requirements for information and detailed knowledge of the utilisation system.

Perhaps the most useful way of using this model would be to use parametric LP, varying the expansion cost c . For $c = 0$, the problem is equivalent to an unconstrained stand-level optimisation. Increasing c would give gradually increasing degrees of smoothing. It seems plausible that as c increases the solution would converge to the nondeclining yield one, although this conjecture has not been proven.

End effects

The use of a finite planning horizon to model what is essentially an infinite horizon problem can cause difficulties. Without any provision for end effects, the optimisation will normally cause the forest to be liquidated in the last period. This generates a distortion that propagates back for several periods near the end, producing solutions significantly different from what would be a long-term optimum.

One way of reducing this effect is to introduce constraints on the ending inventory (Johnson & Scheurman 1977). For example, the total standing volume left at the end might be forced to be no less than the volume in the first period.

Another approach is to assign a value to the area left in each age-class at the end, and to include this in the objective function. A value that has given good results, and that is easy to calculate, is the discounted cash flow over infinite periods with no yield regulation or other constraints. This is equivalent to solving an infinite horizon problem, where the constraints other than area flow conservation apply only to a finite number of initial periods.

With some yield regulation models, such as the capacity expansion model, it is possible to prove that there exists a point in time after which the constraints are not active. The infinite discounted cash flow approach can then result in an infinite horizon optimum. Otherwise, some distortion due to end effects occurs, and the results for a certain number of the final periods should be discarded. Grinold (1983) discussed end effects in the general context of multi-stage problems.

Other constraints

Apart from area flow conservation and yield regulation, other constraints may be needed. These may reflect resource limitations, accessibility, contractual commitments. Some may be enforced simply by omitting arcs in the area flow networks.

SOLUTION STRATEGIES

Brute Force

By far, the most common approach in practice is the straightforward application of a general-purpose LP package. A matrix generator program is used to translate the area flow conservation constraints from a Model B or Model C formulation, and the constraints and objective function from the utilisation submodel, to the format required by the LP system. Problems with thousands of constraints and tens of thousands of variables are routinely solved.

I have already discussed some aspects of the LP solution, and the trade-offs between sparsity and number of constraints. Experience with FOLPI (García 1984) has also shown the importance of redundant variables and of degeneracy, contrary to textbook folklore. It is well known that network constraints result in massive LP degeneracy. García (1984) reported slow convergence with the LP2900 (ICL 1980) and MINOS (Murtagh & Saunders 1978) LP systems due to degeneracy (*see also* Kent *et al.* 1987). Right-hand-side perturbation was tried, but it did not reduce the number of iterations. A modification of MINOS (M.A. Saunders, pers. comm. 1985) improved the situation somewhat. The much better performance of the SCICONIC package (Manley & Threadgill 1987) was found to be due largely to an optional pre-processing step (PRESOLVE) that reduces the number of redundant variables and constraints, and therefore degeneracy. A similar procedure used with MINOS gave comparable results.

In any given problem, many (most?) of the harvest, replanting, and transfer arc flows in the full Model B network are not permissible. Excluding these arcs (variables) causes nodes (constraints) to become redundant, allowing network simplifications. These simplifications may be propagated in several cycles. The first version of FOLPI did take most of these into account when generating the LP constraints, except for some redundant variables arising from empty age-classes and from replanting defaults. A more thorough analysis, together with some fine-tuning of algorithmic parameters (mainly pricing block size), has resulted in order-of-magnitude improvements in solution times over those originally reported. There is still a large proportion of degenerate pivots, however, and some technique such as those proposed by Perold (1980) or Gill *et al.* (1989) might be advantageous. Making use of oldest-age or other priorities (*see below*) could also help.

Another strategy that can improve the efficiency of the LP solution is to substitute inequalities for the node equality constraints of the forestry submodel. In most properly formulated problems, the solution should not change if the flow out of a node is constrained to be less than or equal to the flow into the node, instead of being forced to be equal. A different solution would mean that it is profitable to leave unused hectares behind. Even if this is so, it can be argued that this is usually a feasible alternative, so that the inequalities formulation would be more appropriate. The advantage of relaxing the node equalities, at least for the source nodes, is that often an initial basic feasible solution, zero flows on all arcs, is immediately available, making the LP Phase I unnecessary. This has resulted in savings of around 40% with some LP packages. With SCICONIC there was little or no gain, however, because it uses a special non-basic Phase I procedure that seems to be very efficient. Instead of the all-zeroes flow, various heuristics could be used to generate better initial solutions.

Decomposition and Partitioning

It is clear that without any constraints in the utilisation submodel it would be very easy to solve the optimisation problem. The optimum would be achieved by optimising independently the management of each individual stand. Faustmann showed in 1849 how to do this, in a stationary model with an infinite planning horizon. In Model C the problem is trivial—just assign each initial area to the most profitable harvest/planting sequence. In Models A or B, it is necessary to find a maximum profit path in the network for each crop-type and age-class. Dynamic Programming and other simple and efficient network path algorithms are available for this (actually, most algorithms find all these paths simultaneously).

The forest planning problem can be seen then as a number of easy-to-solve problems (the scheduling of stands in the forestry submodel) linked by complicating constraints (the utilisation submodel). The best-known approach for exploiting this structure is the Dantzig-Wolfe decomposition algorithm, and numerous proposals have appeared advocating its use. Other decomposition methods have been reviewed by Molina (1979). Less well-known, but apparently more promising, are various LP partitioning techniques.

Dantzig-Wolfe decomposition

It will be useful to explain the Dantzig-Wolfe algorithm in a slightly more general setting than usual. Consider an LP problem where we distinguish two sets of constraints:

maximise $c'x$ subject to:

$$A x = b, \quad (1)$$

$$B x = d, \quad (2)$$

$$x \geq 0.$$

The problem would be easy to solve if the constraints in (2) were absent. In forest planning (1) may be the forestry submodel, and (2) the utilisation submodel.

Let X be a matrix where the columns are all the basic feasible solutions of (1). The solution to the full LP must be a convex linear combination of these solutions (assuming, for simplicity, that the solution set is bounded):

$$x^* = X w, \quad \text{with } 1'w = 1, \quad w \geq 0.$$

The problem can then be stated as finding w to

maximise $(c'X) w$ subject to

$$(B X) w = d,$$

$$1'w = 1,$$

$$w \geq 0.$$

This is the Dantzig-Wolfe master. The columns of $B X$ are not all produced in advance, but are generated as needed. At any iteration, the next column $B x$ to enter the master basis is found by solving the subproblem

maximise $(c' - p'B) x$ subject to

$$A x = b,$$

$$x \geq 0.$$

where p are the simplex multipliers from the master.

The subproblem is easy to solve, and the master has fewer constraints than the original problem.

The Dantzig-Wolfe algorithm is usually presented as applicable to situations where A is block-diagonal (e.g., Dantzig 1963), so that the subproblem above decomposes into a number of independent subproblems (as in the forest planning problem). However, it can also be useful when a special structure in A , such as a network structure, makes the subproblem easy to solve (as in the forest planning problem). In this situation the algorithm is sometimes called a column generation procedure (Eriksson 1983).

Despite its undeniable aesthetic appeal, the performance of the decomposition algorithm in practice has not been encouraging (Orchard-Hays 1973; Ho & Loute 1983). It seems that it can be justified for problems that are too large to be solved directly, but can rarely compete with state-of-the art general LP codes otherwise. It is often mentioned that convergence is very slow, owing to the huge number of columns in BX . Perhaps more significant in forest planning is the fact that the number of rows in B can be relatively large, and that, although B is sparse (constraints rarely involve variables in more than two periods), the basis for BX in the master will be dense.

Partitioning

Partitioning methods seek to exploit special structure in the basis when solving the equations of a simplex-based method. The special structure arises from an “easy” subset of constraints, Equation (1). Partitioning methods are often called basis factorisation methods, but they should not be confused with the triangular LU basis factorisation techniques used in modern LP codes. These methods have been developed mainly for two types of problems. In one, the matrix A is of the block diagonal or related forms (block angular, block triangular, staircase) typical of multi-stage problems (e.g., Winkler 1974; Dantzig *et al.* 1981). In the other, the special constraints correspond to some kind of network (e.g., Dantzig 1963; Chen & Saigal 1977; Kennington & Helgason 1980). The constraints (2) are called complicating, coupling, or side constraints. It is possible to deal also with complicating or side variables (Winkler 1974; Glover & Klingman 1981, 1985; McBride 1985). Birge (1985) examined some relationships between partitioning and decomposition methods. The use of partitioning in harvest planning was suggested by García (1984), and attempted by Broad (1990).

At every iteration, the simplex and related methods need to solve systems of linear equations involving a basis and its transpose. In the problem of the previous section, the basis can be partitioned as

$$\begin{bmatrix} A_1 & A_2 \\ B_1 & B_2 \end{bmatrix}$$

where A_1 and A_2 are submatrices of A , and B_1 and B_2 are submatrices of B . The matrices A_1 and B_2 are square. The equations are solved taking advantage of the fact that the systems

$$A_1 x = b \quad \text{and} \quad A_1' y = c$$

are easily solved making use of the special structure of A_1 . The basic idea is to make use of the partitioned inverse, or of a block-triangular factorisation, so that only the inverse or a factorisation of a “working basis”

$$H = B_2 - B_1 A_1^{-1} A_2$$

needs to be stored and manipulated directly. Although developed primarily to solve problems with embedded network or block diagonal components, it is clear that these techniques are applicable whenever A_1 is easily invertible—for example, if A_1 is triangular.

Complementing this data-handling structure, several variations on the simplex strategy may be advantageous. A good unifying review of these has been given by Winkler (1974).

In a Model B forest planning problem, A may correspond to the forestry submodel constraints, a network. This is a particularly simple type of network; it is acyclic (no loops), and uncapacitated (no upper bounds on arc flows) (García 1984; Gunn & Rai 1987; Broad 1990). Adolphson (1988) showed how to streamline the general procedure in precisely this situation.

It may also be possible to include in A some constraints from the utilisation submodel. This is so with some constraints involving harvested volumes, resulting in an A with a generalised network structure (McBride 1985).

For problems in which the number of side constraints is small relative to the number of network constraints, significant (although not spectacular) computational savings over general-purpose LP solution methods have been reported. As with decomposition, loss of the sparsity of the side constraints may be a problem in the forest planning situation. Although both A and B are sparse, the working basis H is likely to be dense. In a sense, both (1) and (2) are "easy" sets of constraints in the absence of the other set. Unless the structure of the side constraints can be exploited, only modest computational gains seem likely. Winkler (1974, p.62) gave some pointers on what could be done, for the block-angular case.

Generalised upper bounds

The generalised upper bounding (GUB) algorithm of Dantzig and Van Slyke (e.g., Murtagh 1981) can be seen as a special case of partitioning. Here the special structure consists of a set of GUB constraints, bounds on sums of disjoint sets of variables. All the forestry submodel constraints in Model C are GUBs. In Model B, either the harvesting or the replanting constraints can be written in this form.

Some LP packages support GUB, making it easy to take advantage of this structure. With at least one LP system, however, it was found that the GUB formulation used more computing time than straight LP (García 1984). Similar results were reported by Kent *et al.* (1987). Glover & Klingman (1981) commented that "... the GUB feature has been dropped from most of the commercial LP codes ...". This relates to the fact that both IBM and CDC dropped GUB because of added difficulty of code maintenance and MIP interfacing (D. Klingman, pers. comm. 1983).

Prioritising: The Oldest-first Principle

It seems obvious to most foresters that, other things being equal, the oldest stands should be cut first. More precisely, cutting from an age-class in some period would not be optimal if there are hectares available in an older age-class in the same crop-type. This principle is used, implicitly or explicitly, in most simulations and in some specialised optimisation procedures (Walker 1976; Lyon & Sedjo 1986; *see* next section). Current LP-based systems, however, evaluate all available classes for possible cutting in the pricing step of each iteration.

It is possible to construct examples where the oldest-first ordering is not optimal, so that the principle is not universally valid for arbitrary planning models. It would be useful to

establish precisely under what conditions the restriction to an oldest-first ordering is optimal. An optimality proof under conditions that cover many practical situations is given in Appendix 2. In special circumstances, prioritising rules that work across crop-types can be obtained (Walker 1976; García 1981).

Restricting the LP bases to those having a structure consistent with a prioritising rule would vastly reduce the number of candidate columns for entering and leaving the basis that need to be considered. Even if the optimal basis has this special structure, however, there is no guarantee that the ordinary simplex method would reach it stepping only through the special bases. An idea of Markowitz for exploiting the triangularity of the optimal basis in certain generalised networks is applicable though (Dantzig 1963, p. 424). He proposed using parametric LP, moving through bases that are optimal for the current value of the parameter. In the forestry planning problem we could parametrise the right-hand-side of the utilisation constraints, or the objective function as in the capacity expansion approach. In both instances, we can start with the parameter value corresponding to the unconstrained solution.

It can be seen that an oldest-first condition implies planar crop-type subnetworks. That is, the Model B network can be drawn so that harvesting arcs do not cross each other (Appendix 2). An intriguing possibility is to somehow exploit this planarity. For example, Berge (1962) showed that on a planar network the maximum flow problem is equivalent to the much-easier shortest-path problem.

Prioritising could also be used with less drastic changes to existing systems, and when the oldest-first principle may not be always optimal. It seems likely that large computational gains could be achieved by using prioritising rules to guide the selection of entering columns in a modified pricing procedure, and/or to resolve ties when pivoting. The effort spent in pricing, the most time-consuming part of the LP iteration, could be reduced, and the impact of degeneracy might possibly diminish. Even something as crude as ordering the LP matrix columns in decreasing order of age has been found to reduce the number of iterations (García 1984).

Time Structure

Forest planning obviously deals with dynamic systems, evolving in time. This time dimension does not play a fundamental role in the optimisation methods discussed above. A number of techniques, though, are based directly on the time structure that is characteristic of the state space perspective.

The model

As already mentioned, the time structure is most apparent in Model A, and in particular in the transition function. The state of the forest in period t can be described by a vector \mathbf{x}_t , where the elements are the current number of hectares in each of the cutting units (crop-type / age-classes). There is a decision or control vector \mathbf{u}_t , specifying the areas harvested from each cutting unit, and areas replanted into various crop types. The state in period $t+1$ is a linear function of the state and control for period t (the transition function):

$$\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \mathbf{u}_t,$$

with initial conditions giving the areas \mathbf{x}_1 for the first period (*see*, for example, Lyon & Sedjo 1983).

The control must satisfy certain linear constraints:

$$C_t x_t + D_t u_t = b_t, \quad u_t \geq 0.$$

These ensure that areas harvested and replanted do not exceed those available in the originating cutting unit. Single-period and two-period utilisation submodel constraints can be included here, using the transition function and including slack or surplus variables in u , if necessary. Other state and control variables can also be added, such as the current capacity and capacity expansion values.

The objective is, for example, the maximisation of the sum of the net discounted revenues for each period. These are linear functions of x_t and u_t (nonlinear functions may be used in some instances). Alternatively, we can define another state variable equal to the accumulated net discounted revenue, and the objective is then to maximise its final value. These two variations are called the Lagrange and the Mayer forms, respectively. A problem with both periodic and terminal terms in the objective is said to be a problem of Bolza. These names originate in the Calculus of Variations, an analytical approach that can be used in some of the simpler situations.

The optimisation of the model just described is sometimes called a dynamic, or multi-stage LP problem. Writing down the constraints in an LP matrix, the so-called staircase structure results. This has submatrices (the A, B, C, and D matrices above) linking pairs of consecutive periods. The LP matrix can also be rearranged into block-angular and other patterns.

Dynamic LP and forestry applications

The main solution approaches, other than straight LP, are based either on Dantzig-Wolfe decomposition, or on partitioning. These have already been explained in general terms. Details of a representative sample of suggested solution methods can be found in several papers in the publication by Dantzig *et al.* (1981). Note that, despite the similarity in name, standard Dynamic Programming methods are useful only when the dimension of the state vector does not exceed 2 or 3.

A number of techniques are inspired by ideas from Optimal Control Theory—specifically, by the (discrete) Optimum Principle of Pontryagin. Optimal Control Theory is an offshoot of the Calculus of Variations and of System Theory, and deals with more general problems than those discussed here. In the linear dynamics case, Control Theory models and methods can be described using well-known Linear Programming concepts, so that the differences from LP are largely a matter of terminology. There has been, though, a valuable cross-fertilisation of ideas, with LP methods used in control systems (Canon *et al.* 1969), and Control Theory methods suggesting LP solution strategies (Propoi & Krivonozhko 1978).

Pontryagin's Optimum Principle specifies necessary conditions for an optimum, but does not by itself provide an algorithm for finding it. The basic idea, in Mathematical Programming terms, is that if the values of the dual variables were known, it would be easy to compute the optimum primal values (u_t and x_t), period by period, starting with the initial x . Conversely, with known primals, the duals (adjoint variables in Control Theory jargon) could be obtained going back in time, starting from an easily determined final value. Calculating the values of both the primal and dual variables using these relationships is what is known as a two-point boundary value problem (initial boundary conditions for the primals, and final boundary

conditions for the duals). This is not an easy problem to solve in general, and a number of iterative algorithms have been used. Propoi & Krivonozhko (1978) proposed a method for the Dynamic LP problem inspired by some of these ideas, combined with LP partitioning.

Walker (1976) developed a dynamic planning model and solution technique named ECHO. ECHO was originally derived from first principles, using ideas from forestry financial maturity and economic marginal analysis. Later, connections with Control Theory were pointed out by McDonough & Park (1975), and with Mathematical Programming by Johnson & Scheurman (1977). The model considers a single crop-type, although Walker (1976) indicated that it can be generalised if a fixed cutting priority is given. Yield regulation is single-period, with a demand curve and possibly nonlinear logging costs. An oldest-first priority is used in harvesting. This results in simplifications that essentially amount to having a one-dimensional dual vector, i.e., one dual variable per period. Walker computed approximately optimal solutions with a procedure that iteratively adjusts harvest levels.

Lyon & Sedjo (1983, 1986) refined Walker's procedure, approaching it from a control-theoretical perspective (*see also* Sedjo & Lyon 1990, Chapter 7). Taking advantage of the one-dimensional dual, they used a variation of the "shooting method" for boundary value problems. Although the details differ, the basic idea of this method is to assume an initial value for the dual, and use this to compute the dual and primal values forward in time. If the final value obtained for the dual is close enough to what it should be, the solution is accepted. Otherwise, the initial dual value is changed, and the process repeated. The adjustments to the initial dual value can be done systematically, using binary search or some other univariate optimisation technique. Following Walker, Lyon & Sedjo used the harvest level instead of the dual variable, making use of a relationship between consecutive harvest levels related to the one between consecutive dual values obtained by Pontryagin's principle.

Kallio *et al.* (1981) developed a forest planning model in a Dynamic LP framework. It is an integrated model intended for planning at the national level. The forestry submodel is relatively simple, but it has a fairly elaborate utilisation component. The authors indicated that the algorithm of Propoi & Krivonozhko (1978) or other specialised Dynamic LP algorithms could be used for the optimisation, although standard LP was used in practice.

Other temporal aspects

Although a state-space interpretation is most natural in the Model A setting, it is also applicable to Model B. The state at any period in Model B can be described by the area planted in each of the previous periods that has not yet been cut. It would be possible, then, to make use of the temporal structure in this model also.

To examine one possibility, consider the two main aspects of the problem that are exploited in Dynamic LP algorithms. One is data-handling, using the block structure of the LP matrices to save storage, and maybe some computations. The other is the fact that variables are directly linked only to variables that are close to them in time. This motivates pivoting strategies that limit the extent of matrix updating necessary, and other computational shortcuts. It appears that this second characteristic of dynamical systems, possibly combined with an oldest-first priority rule, should produce quasi-triangular Model B bases—that is, lower-triangular matrices plus a small number of diagonal rows above the main diagonal. In some simple problems, nondeclining yield with one crop-type, I have found that, with an

appropriate ordering of rows and columns, optimal bases are triangular except for one additional diagonal row. If quasi-triangular bases could be maintained, techniques similar to those of Huard (1981) might produce considerable computational gains.

Another possibility worth exploring relates to the nature of finite-horizon models. Often we are interested in the optimal decisions for one or a few initial periods, but we are forced to plan over a longer horizon to avoid distortions created by end effects. Aronson *et al.* (1985) developed the forward simplex method for this type of problem. The forward algorithm proceeds by first solving the one-period problem, augmenting its solution to obtain the optimal solution for the two-period problem, etc. After several steps the initial part of the solution stabilises, and when the first-period decision is guaranteed to be "good enough", computation can stop. Although the forward simplex method was developed for multi-stage problems, it might be applicable also to Model B.

Returning to the state space formulation, in many problems the coefficients in the transition and constraint equations are independent of t . In other words, the problem is time-invariant, or stationary (not to be confused with stationary solutions). It seems natural to think that this property could be exploited, making it easier to compute solutions, even (or at least) with an infinite horizon. Simple results have been obtained in special cases, e.g., for dynamic networks (Ford & Fulkerson 1958), for transition matrices with a Leontief structure (Dantzig 1959), and for systems with no constraints on the state (Dantzig & Sethi 1981). For more general problems, however, no practical ways of making use of stationarity seem to be available.

Other Strategies

Hoganson-Rose, and Lagrangian relaxation

Hoganson & Rose (1984) developed an algorithm for obtaining approximate solutions to forest planning optimisation problems. They used a Model C formulation, but the approach could be also used with Models A or B. With reference to the equations in the section on decomposition, the basic idea is to remove the complicating constraints (2) from the LP rows, and place them as a penalty term in the objective function. Estimates of the dual variables for (2) are used as penalty weights. The transformed problem then has a form identical to the Dantzig-Wolfe subproblem, and is easily solved. If in this solution the constraints (2) are satisfied to within an acceptable tolerance limit, the computation is finished. Otherwise, the dual estimates are modified, and the procedure is repeated.

The Hoganson-Rose algorithm is essentially equivalent to a Lagrangian relaxation procedure used mainly to generate bounds in integer linear programming algorithms (Fisher 1981, 1985). The alteration of the dual estimates is different, though. It seems likely that the Hoganson-Rose procedure could be streamlined and accelerated by using the subgradient method of adjusting the duals suggested by Fisher.

Even with full convergence to a feasible solution, in general this solution will not be optimal, although it may be "close enough", as pointed out by Hoganson & Rose. This is easily seen from the fact that the number of basic variables in this solution equals the rank of the forestry constraints (1), which is less than that for the fully constrained problem. In particular, this solution cannot "split" the initial areas, because the bases for the unconstrained forestry network are trees rooted at the sink. This may be seen as an advantage, because it

may be undesirable to subdivide stands. Still, this is likely to be different from an integer programming optimal solution.

To ensure an optimal solution to an LP problem by Lagrangian relaxation it is necessary to use a nonlinear penalty function. Gol'shtein (1981) showed how to do this with a quadratic. A similar approach has been used by Gunn & Rai (1987) with forest planning models.

Other potentially useful techniques

Tseng & Bertsekas (1987) have proposed a new non-simplex method for linear programming. This method has given good results for networks (Bertsekas & Tseng 1988), which suggests that it might be advantageous on the simpler types of forest planning problems. Another method that deviates from the simplex method has been studied by Mitra *et al.* (1988). They pointed out similarities with the usage of nonbasic solutions in SCICONIC, a feature that may account for the already-mentioned Phase I efficiency of this LP package.

Two-level linear programming (Balas & Karwan 1984) deals with problems that have some similarities to forest planning. This approach considers systems with two subsystems in a hierarchical planning structure, e.g., forestry and utilisation. The objective function and constraints for a planner at one level are determined, in part, by the other level.

The degree of aggregation (that is, period length, age-class width, and number of crop-types) has a large impact on the cost of solving a forest planning LP problem. Manley & Threadgill (1987) provided some comparative figures. Shetty & Taylor (1987) considered a strategy that starts with small, highly aggregated problems. At each iteration, bounds on the objective value are obtained and, if the solution is not acceptable, the problem is modified to one with a lower level of aggregation.

It was mentioned in the previous section that sometimes it is not desirable to subdivide the initial areas. In general, an optimisation with this constraint would be costly, requiring the use of integer programming. Simple cases, such as nondeclining yield, have a structure similar to the generalised assignment problem (e.g., Fisher *et al.* 1986), for which a number of good algorithms are available. It could be possible to obtain an efficient algorithm along similar lines, which might even be competitive for continuous variables.

I have considered only discrete formulations of forestry planning problems. It is also possible to formulate problems in continuous time, and/or with continuous age distributions. Models of this kind have often been used to obtain analytical results in simple situations (e.g., Johansson & Löfgren 1985). There are instances where a continuous optimisation problem is easier to solve than the corresponding discrete version (Anderson & Nash 1987). In this context, it is interesting to note that there are strong similarities between continuous versions of simple constant or nondeclining yield problems, and the cutting-and-filling problem (Kantorovich 1958; Anderson & Nash 1987). In the cutting-and-filling problem, material is removed from one area, and used to fill a space of a given shape. In the forestry problem, hectares are removed from an age distribution, and allocated trying to fill "gaps" in a production curve. The fact that the optimal material movements in cutting and filling do not cross each other, has a parallel in the oldest-first principle, where the harvesting allocations do not cross (Appendix 2).

DISCUSSION AND CONCLUSIONS

The classification of forest planning models presented here appears to be a reasonable way of comparing the various approaches that have been used. It may also help in guiding future developments. Models A, B, and C are essentially equivalent in their power for describing forest planning problems. They do differ in their usefulness in association with specific solution techniques.

Network concepts have been found very useful, not only for visualising and thinking about models, but also for training and communication with forest managers. It is important to realise that the managers' acceptance of these systems depends on their understanding of the model. Ideally, they should be able to use the system by themselves, without a detailed knowledge of irrelevant mathematical and computational technicalities (García 1984). I have not dealt with these aspects here, but the development of appropriate metaphors and user interfaces is likely to have a larger impact than any new mathematical or computational advances.

For completely general forest management problems, it seems difficult to achieve any large computational advantages with other than general-purpose LP basis-handling techniques. Following Winkler (1974), we can distinguish in an algorithm the data-handling aspects (use of special structure to save in storage and computation) and the strategy aspects (primal, dual, or other strategies, pivoting criteria, pricing approaches). In data-handling, modern LP codes are extremely effective in exploiting sparsity with arbitrary structures. It is in the strategy side where major improvements may be possible. Modified pricing and pivot selection procedures using prioritising rules appear to be a promising approach.

Important gains from specialised algorithms seem likely with more restricted types of problems. Satisfactory procedures are already available for some instances of single-period conditions. Other examples are the nondeclining yield and the expansion cost models. Efficient solution techniques for these could be very useful in combination with interactive simulators.

The most often-mentioned method of exploiting special structure in forestry problems, Dantzig-Wolfe decomposition, appears particularly unattractive. Partitioning methods are more promising, although sparsity would need to be preserved. Dynamic LP methods are largely experimental, but could be useful provided the structure within blocks is utilised. Lagrangian relaxation and related techniques have been successful recently on a number of structured LP and integer problems.

There is little or no experience with most of these methods on real-life forestry problems, and most of the other suggestions here have not appeared before in the forestry literature. In any case, I believe that any real breakthroughs will require combining ideas from several approaches.

REFERENCES

- ANDERSON, E.J.; NASH, P. 1987: "Linear Programming in Infinite-dimensional Spaces—Theory and Applications". Wiley, Chichester.
- ADOLPHSON, D.L. 1988: Solving acyclic shortest path networks with side constraints with an application to capacitated lot sizing problems. Pp.304–14 in Rand, G.K. (Ed.) "Operational Research '87" Elsevier Science Publishers, North-Holland.

- ARONSON, J.E.; MORTON, T.E.; THOMPSON, G.L. 1985: A forward simplex method for staircase linear programs. *Management Science* 31: 664–79.
- BALAS, W.F.; KARWAN, M.H. 1984: Two-level Linear Programming. *Management Science* 30: 1004–20.
- BARROS, O.; WEINTRAUB A. 1982: Planning for a vertically integrated forest industry. *Operations Research* 30: 1168–82.
- BERGE, C. 1962: Teoría de las redes y sus aplicaciones. Continental S.A., México. [Transl. from French edition Dunod, Paris, 1958. English edition: Wiley, New York, 1962].
- BERTSEKAS, D.P.; TSENG P. 1988: Relaxation methods for minimum cost ordinary and generalized network flow problems. *Operations Research* 36: 93–114.
- BIRGE J.R. 1985: A Dantzig-Wolfe decomposition variant equivalent to basis factorization. *Mathematical Programming Study* 24: 43–64.
- BROAD, L.R. 1990: Note on area conservation mechanisms associated with forest management. *New Zealand Journal of Forestry Science* 20: 120–7.
- CANON, M.D.; CULLUM, C.D.; POLAK, E. 1969: “Theory of Optimal Control and Mathematical Programming”. McGraw-Hill, New York.
- CHEN, S.; SAIGAL, R. 1977: A primal algorithm for solving a capacitated network flow problem with additional linear constraints. *Networks* 7: 59–79.
- CLARK, C.W.; CLARKE, F.H.; MUNRO, G.R. 1979: The optimal exploitation of renewable resource stocks: problems of irreversible investment. *Econometrica* 47: 25–47.
- CLUTTER, J.L.; FORTSON, J.C.; PIENAAR, L.V.; BRISTER, G.H.; BAILEY, R.L. 1983: “Timber Management: A Quantitative Approach”. Wiley, New York.
- DANTZIG, G.B. 1959: On the status of multistage linear programming problems. *Management Science* 6: 53–72.
- 1963: “Linear Programming and Extensions”. Princeton University Press, Princeton New Jersey.
- DANTZIG, G.B.; SETHI, S.P. 1981: Linear optimal control problems and generalized linear programs. *Journal of the Operational Research Society* 32: 467–76.
- DANTZIG, G.B.; DEMPSTER, M.A.H.; KALLIO, M.J. (Ed.) 1981: “Large-scale Linear Programming” (2 vol.). International Institute for Applied Systems Analysis, Laxenburg, Austria.
- ERIKSSON, L.O. 1983: Column generation applied to long range forestry planning models. *The Swedish University of Agricultural Sciences, Department of Operational Efficiency, Report No. 155*.
- FISHER, M.L. 1981: The Lagrangian relaxation method for solving integer programming problems. Pp.581–616 in Dantzig, G.B.; Dempster, M.A.H.; Kallio, M.J. (Ed.) “Large-scale Linear Programming”. International Institute for Applied Systems Analysis, Laxenburg, Austria.
- 1985: An applications oriented guide to Lagrangian relaxation. *Interfaces* 15(2): 10–21.
- FISHER, M.L.; JAIKUMAR, R.; Van WASSENHOVE, L.N. 1986: A multiplier adjustment method for the generalized assignment problem. *Management Science* 32: 1095–103.
- FORD, L.R.; FULKERSON, D.R. 1958: Constructing maximal dynamic flows from static flows. *Operations Research* 6: 419–33.
- GARCIA, O. 1979: Modelling stand development with stochastic differential equations. Pp.315–34 in Elliott, D.A. (Comp.) “Mensuration for Management Planning of Exotic Forest Plantations”. *New Zealand Forest Service, Forest Research Institute Symposium No.20*.
- 1981: IFS, an interactive forest simulator for long range planning. *New Zealand Journal of Forestry Science* 11: 8–22.
- 1984: FOLPI, a forestry-oriented Linear Programming interpreter. Pp.293–305 in Nagumo, H. et al. (Ed.) “Proceedings IUFRO Symposium on Forest Management Planning and Managerial Economics”. University of Tokyo.
- 1986: Forest estate modelling (Part 2). Pp.97–9 in Levack, H. (Ed.) “1986 Forestry Handbook”. New Zealand Institute of Foresters (Inc.), Wellington.

- 1988: Growth modelling—A (re)view. *New Zealand Forestry* 33(3): 14–7.
- GILL, P.E.; MURRAY, W.; SAUNDERS, M.A.; WRIGHT, M.H. 1989: A practical anti-cycling procedure for linearly constrained optimization. *Mathematical Programming* 45: 437–74.
- GLOVER, F.; KLINGMAN, D. 1981: The Simplex SON algorithm for LP/embedded network problems. *Mathematical Programming Study* 15: 148–76.
- 1985: Basis exchange characterizations for the Simplex SON algorithm for LP/embedded networks. *Mathematical Programming Study* 24: 141–57.
- GOL'SHTEIN, E.G. 1981: An iterative linear programming algorithm based on the modified Lagrangian. Pp.617–30 in Dantzig, G.B.; Dempster, M.A.H.; Kallio, M.J. (Ed.) "Large-scale Linear Programming". International Institute for Applied Systems Analysis, Laxenburg, Austria.
- GRINOLD, R.C. 1983: Model building techniques for the correction of end effects in multistage convex programs. *Operations Research* 31: 407–31.
- GUNN, E.A.; RAI, A.K. 1987: Modelling and decomposition for planning long-term forest harvesting in an integrated industry structure. *Canadian Journal of Forest Research* 17: 1507–18.
- HARVEY, C.M. 1986: Value functions for infinite-period planning. *Management Science* 32: 1123–39.
- HO, J.K.; LOUTE, E. 1983: Computational experience with advanced implementation of decomposition algorithms for Linear Programming. *Mathematical Programming* 27: 283–90.
- HOEKSTRA, T.W.; DYER, A.A.; Le MASTER, D.C. (Ed.) 1987: FORPLAN: An evaluation of a forest planning tool. *USDA Forest Service. General Technical Report RM-140*.
- HOGANSON, H.M.; ROSE, D.W. 1984: A simulation approach for optimal timber management scheduling. *Forest Science* 30: 220–38.
- HRUBES, R.J.; NAVON, D.I. 1976: Application of Linear Programming to downward sloping demand problems in timber production. *USDA Forest Service, Research Note PSW-315*.
- HUARD, P. 1981: Solving large scale linear programs without structure. Pp.19–54 in Dantzig, G.B.; Dempster, M.A.H.; Kallio, M.J. (Ed.) "Large-scale Linear Programming". International Institute for Applied Systems Analysis, Laxenburg, Austria.
- ICL 1980: "LP2900 Applications Manual". International Computers Ltd, Technical Publication 6955.
- JOHANSSON, P.-O.; LÖFGREN, K.-G. 1985: "The Economics of Forestry and Natural Resources". Blackwell, Oxford.
- JOHNSON, K.N.; SCHEURMAN, H.L. 1977: Techniques for prescribing optimal timber harvest and investment under different objectives—Discussion and Synthesis. *Forest Science Monograph* 18.
- KALLIO, M.; PROPOI, A.; SEPPÄLÄ, R. 1981: A model for the forest sector. Pp.1055–100 in Dantzig, G.B.; Dempster, M.A.H.; Kallio, M.J. (Ed.) "Large-scale Linear Programming". International Institute for Applied Systems Analysis, Laxenburg, Austria.
- KANTOROVICH, L. 1958: On the translocation of masses. *Management Science* 5: 1–4.
- KENNINGTON, J.E.; HELGASON, R.V. 1980: "Algorithms for Network Programming". Wiley, New York.
- KENT, B.M.; KELLY, J.W.; FLOWERS, W.R. Jr 1987: Experience with the solution of USDA Forest Service FORPLAN models. In Dress, P.E.; Field, R.C. (Ed.) "The 1985 Symposium on Systems Analysis in Forest Resources". The Georgia Center for Continuing Education, Athens, Georgia.
- LYON, K.S.; SEDJO, R.A. 1983: An Optimal Control Theory model to estimate the regional long-term supply of timber. *Forest Science* 29: 798–812.
- 1986: Binary-search SPOC: an Optimal Control Theory version of ECHO. *Forest Science* 32: 576–84.
- MANLEY, B.R.; THREADGILL, J.A. 1987: Testing FOLPI, a Linear-Programming-based forest estate model. Pp.275–302 in Kininmonth, J.A. (Ed.) "Proceedings of Conversion Planning Conference". *N.Z. Ministry of Forestry, FRI Bulletin No. 128*.


- McBRIDE, R.D. 1985: Solving embedded generalized network problems. *European Journal of Operations Research* 21: 82–92.
- McDONOUGH, J.M.; PARK, D.E.Jr. 1975: A discrete maximum principle solution to an optimal control formulation of timberland management problems. *PROSE Research Note* 7 [cited by McDonough & Park 1976].
- 1976: Nonlinear optimal control approach to interregional management of timber production and distribution. Pp.198–210 in Meadows, J.; Bare, B.; Ware, K.; Row, C. (Ed.) “System Analysis and Forest Resource Management”. Society of American Foresters, Washington D.C.
- MITRA, G.; TAMIZ, M.; YADEGAR, J. 1988: Experimental investigation of an interior search method within a Simplex framework. *Communications of the ACM* 31: 1474–82.
- MOLINA, F.W. 1979: A survey of resource directive decomposition in Mathematical Programming. *Computing Surveys* 11: 95–104.
- MURTAGH, B.A. 1981: “Advanced Linear Programming: Computation and Practice”. McGraw-Hill, New York.
- MURTAGH, B.A.; SAUNDERS, M.A. 1978: Large scale linearly constrained optimization. *Mathematical Programming* 14: 41–72.
- NAVON, D.I. 1971: Timber RAM—A long-range planning method for commercial timber lands under multiple-use management. *USDA Forest Service, Research Paper PNW-70*.
- ORCHARD-HAYS, W. 1973: Practical problems in LP decomposition and a standardized phase I decomposition as a tool for solving large scale problems. Pp.153–55 in Himmelblau, D.M. (Ed.) “Decomposition of Large-scale Problems”. North-Holland, Amsterdam.
- PEROLD, A.F. 1980: A degeneracy exploiting LU factorization for the Simplex method. *Mathematical Programming* 19: 239–54.
- PROPOI, A.; KRIVONOZHKO, V. 1978: The simplex method for dynamic linear programs. *International Institute for Applied Systems Analysis, RR-78-14*.
- SEDJO, R.A.; LYON, K.S. 1990: The long-term adequacy of world timber supply. Resources for the Future.
- SHETTY, C.M.; TAYLOR, R.W. 1987: Solving large-scale linear programs by aggregation. *Computers and Operations Research* 14: 385–93.
- TSENG, P.; BERTSEKAS, D.P. 1987: Relaxation methods for linear programs. *Mathematics of Operations Research* 12: 569–96.
- WALKER, J.L. 1976: ECHO: solution technique for a nonlinear economic harvest optimization model. Pp.172–88 in Meadows, J.; Bare, B.; Ware, K.; Row, C. (Ed.) “System Analysis and Forest Resource Management”. Society of American Foresters, Washington D.C.
- WARE, G.O.; CLUTTER, J.L. 1971: A mathematical programming system for the management of industrial forests. *Forest Science* 17: 428–45.
- WINKLER, C. 1974: Basis factorization for block-angular linear programs: Unified theory of partitioning and decomposition using the simplex method. *Stanford University, Systems Optimization Laboratory, Technical Report SOL 74-19*.

APPENDIX 1

ANOTHER VIEW OF THE EXPANSION COSTS MODEL

A simple yield regulation approach was presented in the section on Utilisation Submodels "A capacity expansion approach". An alternative motivation for the model follows.

Consider the classical model for a firm (e.g., a sawmill) with a linear cost function:



Assume that there are constant returns to scale in the industry. Then, this model would also apply to the industry as a whole, and the fixed cost would be proportional to the total capacity (think of the industry as the sum of many small mills):

$$F = fz.$$

Assume also that in period t the installed capacity z_t can be increased by any amount w_t , e.g., by adding new mills, but it cannot be reduced (the model could be extended by including scrap values). We have then the same constraints as in the section on "A capacity expansion approach", with the objective of maximising the discounted total profits:

$$\text{maximise } \sum \alpha^t (p_t y_t - f z_t) = \sum \alpha^t p_t y_t - f \sum \alpha^t z_t.$$

The second term can be written as a function of the w_t :

$$\begin{aligned} f \sum \alpha^t z_t &= f z_0 \sum \alpha^t + f \sum \alpha^t \sum_1^t w_s \\ &= f z_0 \alpha(1-\alpha^T)/(1-\alpha) + f/(1-\alpha) \sum (\alpha^t - \alpha^{T+1}) w_t, \end{aligned}$$

where T is the planning horizon. For an infinite horizon, this objective function is the same as the one in the section referred to, except for a constant term. The expansion cost is related to the fixed cost per unit by

$$c = f / (1-\alpha).$$

There is a small discrepancy between the two objectives for finite planning horizons. It is caused by ignoring here the fixed costs incurred after T due to expansion decisions before T .

Another interesting relationship for the expansion or fixed cost is

$$f/p = B/z.$$

This is easily obtained from the diagram. The right-hand side is the break-even utilisation as a fraction of the installed capacity, and it may be easier to estimate than trying to estimate f or c directly.

APPENDIX 2

OPTIMALITY OF THE OLDEST-FIRST RULE

Under some conditions, optimal solutions are consistent with the oldest-first principle. The following result is not the most general possible, but it is easy to understand and covers many practical situations.

Theorem:


In Model B, assume that

- (a) The area constraints can be given as inequalities, as described in the last two paragraphs of the section on “Brute Force”
- (b) The objective and utilisation constraints depend on the harvesting area flows only through the yields $y_{it} = \sum v_{ist} x_{ist}$, and
- (c) The logarithms of the yield functions v_{ist} are concave functions of the age $t-s$.

Then, there exists an optimal basis such that if x_{ist} and $x_{is't'}$ are in the basis and $s < s'$, then $t \leq t'$.

Proof:

Suppose that an optimal basis does not satisfy the rule, that is, it contains the variables $x_{ist'}$ and $x_{is't}$, with $s < s'$ and $t < t'$ (the arcs “cross”, see the solid arcs in the diagram below)



Re-direct a small number δ of hectares from s to t , adjusting the other flows to maintain the same yields at t and t' , and the same supply at s' :

- $x_{ist} = \delta$,
- decrease $x_{is't}$ by $(v_{ist} / v_{is't}) \delta$,
- $x_{is't'} = (v_{ist} / v_{is't'}) \delta$,
- decrease $x_{is't'}$ by $(v_{is't'} / v_{is't'}) (v_{ist} / v_{is't}) \delta$.

Increase δ until one of $x_{is't'}$ or $x_{is't}$ drops out, obtaining a new solution with no crossing arcs. The net change in the availability of hectares at s is

$$[(v_{is't'} v_{ist}) / (v_{ist'} v_{is't}) - 1] \delta.$$

It can be shown that the concavity of $\log v$ implies that this change is non-negative, so that the new solution is feasible.