FOREST MONITORING AND HARVESTING POLICIES

John L. Casti
University of Arizona, Tuscon, Arizona, USA
and
International Institute for Applied Systems Analysis, Laxenburg, Austria

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
Laxenburg, Austria

Research Reports, which record research conducted at IIASA, are independently reviewed before publication. However, the views and opinions they express are not necessarily those of the Institute or the National Member Organizations that support it.

FOREWORD

One of the main themes of the Adaptive Resource Policies Project is to study the problem of identification of aspects of natural resource systems useful for policy design and management. In this article, John Casti investigates the problem of the information needed to construct appropriate feedback policies for forest harvesting. Both the deterministic and the stochastic versions of state identification are considered, as well as the problem of how often measurements of the forest need to be taken in order to determine its state. The article concludes with an appendix discussing some general system-theoretic aspects of measurement and control.

CARL J. WALTERS
Leader
Adaptive Resource Policies Project

Forest Monitoring and Harvesting Policies*

J. Casti

University of Arizona
Tucson, Arizona
and
International Institute for Applied Systems Analysis
Laxenburg, Austria

ABSTRACT

This paper considers the problem of what information must be measured in a forest management model in order to generate optimal feedback harvesting policies. Together with this question, the paper also addresses the issue of how LP-based models can be embedded within a dynamic programming framework, so that feedback rather than open-loop decisions can be determined.

The appendix to the paper presents a more general system-theoretic framework within which forest management may be studied. Issues of surprise, response to unknown disturbances, and robustness of policies are examined and a program for systematically investigating such management questions is outlined.

I. INTRODUCTION

Current forest management policies are often based upon large-scale mathematical models of the linear or nonlinear programming type. While there is considerable merit in structuring the complexities of the decisionmaking process in such a fashion, there are nontrivial difficulties as well, difficulties which are *not* related to the issues of parameter accuracy, linearity versus nonlinearity, measurement data, etc., but which stem from the basic system-theoretic structure of the model. In particular, we cite the problem of monitoring and feedback decisionmaking. Mathematical programming models are inherently unsuitable for situations in which the initial assumptions regarding the relationship between variables, values of parameters, effect of

---

19

decisions, and validity of objectives are subject to modification based upon the actual behavior of the system as the process unfolds. For such updating we need a model which naturally incorporates a feedback mechanism coupled with a monitoring scheme to provide information about the system's behavior to the decisionmaker. In this report, we shall discuss various approaches to the monitoring and control problem which enable one to retain the computational advantages of the mathematical programming formulation without sacrificing the adaptive advantages of feedback control.

As prototypical examples of the two model classes which we shall consider, let us briefly review the so-called Models I and II of Johnson and Scheurman [13] and the LQG model of Dixon and Howitt [11]. These two modeling efforts are of particular interest for us for at least two reasons. First of all, a number of forests are currently being managed according to the schema of Model I or II, so it is clearly not just a textbook issue to address the methodological questions raised by such models. And secondly, the comparison of LQG with Models I and II brings the basic system-theoretic issues surrounding open-loop versus feedback decisionmaking into sharp focus within the context of specific forest models.

Model I has the mathematical form

$$\max \sum_{i=1}^{U} \sum_{q=1}^{R_l} D_{lq} x_{lq},$$

under the area constraint

$$\sum_{q=1}^{R_l} x_{lq} = A_l, \qquad l = 1, \ldots, U,$$

where

  x_{lq} = number of hectares of management unit l assigned to regeneration harvest sequence q,
  A_l = number of hectares in management unit l,
  U = number of management units (i.e., number of age classes present in period 1),
  R_l = number of possible regeneration harvest sequences over the planning horizon for management unit l,
  D_{lq} = discounted net value per hectare of management unit l over the planning horizon, if assigned to regeneration harvest sequence q.

It is further assumed that $x_{lq} \geq 0$. It should be noted here that a particular regeneration harvest sequence has its own built-in constraints involving the

time that must elapse between the harvesting of a certain unit and when it can be reharvested. Since it is not important for our subsequent discussion, we shall not elaborate upon the details of Model I here, but refer to [13] for full details and examples.

In contrast to Model I, where it is assumed that hectares that form any management unit (age class) at the beginning are preserved intact throughout the process, Model II allows hectares from beginning age classes to be broken up and combined with hectares from other age classes when only part of the beginning age class is regeneration harvested in a given period. Thus, in Model II hectares are placed in a new management unit each time they are regeneration harvested.

Mathematically, Model II takes a form very similar to that of Model I, with the only real difference being the introduction of additional variables to account for hectares regeneration harvested in a given period and left as inventory at the end of the process.

For our mathematical needs in this paper, the important point to observe about both Models I and II is that they are of a standard linear programming form with equality constraints, i.e., abstractly both Models I and II can be cast in the form

$$\max \sum_{i=1}^{N} c_i x_i,$$

subject to the constraints

$$\sum_{j=1}^{N} a_{ij} x_j = b_i,$$

$$x_i \geq 0, \qquad i = 1, 2, \ldots, M.$$

Assuming that the problem data (the numbers a_{ij}, b_i, c_i) are known, determination of the maximizing $\{x_i\}$ is a straightforward computation utilizing any of a number efficient LP packages.

A well-known difficulty associated with such a setup is that the optimal policy (i.e., the sequence of optimal x_i-values) is inherently "open-loop." In other words, if the problem data are inaccurate or change unexpectedly during the course of the process, there is no way to modify the values of the x_i to account for such changes. The basic assumption is that the data are given once and for all, *in advance*, and the associated sequence $\{x_i^*\}$ is computed which is optimal for those given data (and generally no others). Thus, if we change our mind about any part of the problem as circumstances or knowledge change, in principle we must start all over again with the new data set and recompute the new optimal decision sequence.

Another difficulty with the standard LP setup is that it is rather awkward to account for stochastic variation in the problem data. Since the decision rule is determined throughout the evolution of the process by the initial specification of the data, there is no natural mechanism whereby the decisions can be modified to take advantage of the way the data happen to "come out" as the process unfolds.

To offset the above analytic disadvantages of the mathematical programming formulation, we have the advantage of computational tractability. Problems involving thousands of variables can be rather routinely dealt with by current algorithms. So, in situations of low uncertainty about the basic problem relations and with accurate data, it is reasonable to expect that the mathematical programming models will yield very satisfactory results for forest harvest planning.

Now let us consider the LQG forest model of Dixon and Howitt [11], which is specifically designed to allow for feedback decisions in the face of uncertainty. The general form of an LQG optimization process is as follows. We wish to minimize (over u) the expected value of the quadratic form

$$J = \sum_{t=0}^{T-1} \left[(x(t), Q(t)x(t)) + (u(t), R(t)u(t)) \right] + (x(T), Mx(T)),$$

where the matrices $Q(t) \geq 0$, $R(t) > 0$, and $M \geq 0$. The n-vector x and m-vector u are related through the linear dynamics

$$x(t+1) = F(t)x(t) + G(t)u(t) + v(t),$$

where $F(\cdot)$ and $G(\cdot)$ are matrix functions of appropriate sizes and $v(\cdot)$ is a zero-mean gaussian random vector with known covariance. In addition, we usually cannot observe the state $x(t)$ directly, but instead observe

$$z(t) = H(t)x(t) + w(t),$$

where $H(\cdot)$ is a matrix of size $p \times n$ and $w(\cdot)$ is a zero-mean gaussian observation noise with known convariance. Further, it is assumed that the noise processes $v(\cdot)$ and $w(\cdot)$ are uncorrelated.

The solution of the LQG problem is well known and is given in many books (see [14], [6]). Basically, the solution proceeds in two steps:

(1) Obtain the unbiased least-squares estimate of the state using the observations z. This estimate—call it $\hat{x}(t)$—is obtained from the celebrated Kalman filter.

(2) Pretend $\hat{x}(t)$ is the *true* state and apply standard *deterministic* linear-quadratic control theory to obtain the optimal control $u^*(t) = -K(t)\hat{x}(t)$, where $K(t)$ is obtained from the solution of a matrix Riccati equation.

The entire procedure is depicted in Figure 1. The details of the computation can be found in the books cited above. The important point to note about the above procedure for our purposes is the need to estimate the state of the system before the controlling action can be computed. Fortunately, in the LQG setting the Kalman filter provides exactly the needed estimate. In more general settings the situation is far more complicated and no such separation of estimation and control is possible.

In the model of Dixon and Howitt [11], the state variables $x(t)$ are taken to be the volume and basal area, and the decision vector $u(t)$ is the timber removal actions. The matrices $F(\cdot)$ involve various aspects of the growth of both basal area and volume from one period to the next, while the matrices $G(\cdot)$ relate the removal of timber volume and basal area to decisions involving either thinning or regeneration harvesting. Finally, the stochastic vectors $v(\cdot)$ represent the effect of omitted variables such as weather, mortality, and insect infestation. In the quadratic objective function, the matrices $Q(\cdot)$, $R(\cdot)$, and M represent the costs of having the system state and controls deviate from some prescribed desired levels. Finally, the observation system is such that the matrices $H(\cdot)$ specify which elements of the state can be directly observed, while the noise vector $w(\cdot)$ accounts for uncertainty in the observing process.

It is important to note that the optimal controlling action (harvesting or thinning) is determined as a (linear) function of the current state of the system, i.e., $u^*(t) = -K(t)x(t)$. Thus, in order to calculate the control we must first measure the state $x(t)$. However, the observation process produces only the element $z(t) = H(t)x(t) + w(t)$. Thus, a central feature of the monitoring problem is the construction of a good estimate of $x(t)$, given $z(t)$. For the LQG situation, this estimate is provided by the Kalman filter. In the
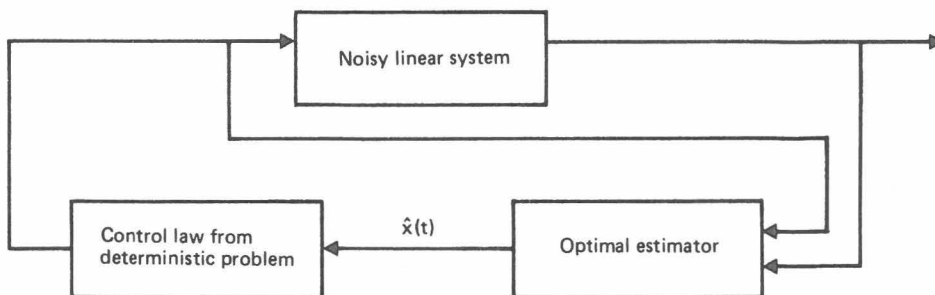

FIG. 1.   Schematic diagram for solution of LQG problem.

absence of noise or in situations where the assumptions of the Kalman filter are not valid (the usual case in real problems), other types of state estimators or observers must be used to obtain $x(t)$. We shall present some of the more effective approaches below.

Before closing this section, it is worthwhile to point out that the deterministic version of the LQG problem [where $v(t)$ and $w(t)$ are identically zero] can easily be formulated in mathematical programming terms and considered as a quadratic programming problem [5]. In that event, a more or less direct comparison between Models I and II and the LQG model could be made. But such a comparison would miss the main point, namely, the difference between making decisions as a function of how the process actually unfolds, which enables one to see how the noise factors $v(t)$ and $w(t)$ actually are realized, and having to assume the realizations of $v(t)$ and $w(t)$ *in advance* for the entire duration of the process. Since it is an easily established result in optimal-control theory that, for a given problem, the optimal feedback control law is always at least as good as the optimal open-loop law, it is immediately clear that if computational constraints were not an issue, we would always prefer the feedback law. However, since the National Forests are constrained to utilize a model similar to Model I or II in their planning, the question arises whether or not a feedback decisionmaking scheme can be superimposed upon the LP model in order to more adequately account for ignorance, uncertainty, and/or unexpected developments during the course of the process. The consideration of this question forms the basis for the remainder of this paper.

II. BASIC MONITORING PROBLEM

As indicated above, if we lived in a perfect world in which there were no uncertainties, surprises, or unexpected developments, the LP planning model exemplified by Models I and II would be the perfect vehicle for forestry management. Decisions could be made with complete confidence, the plan would always be precisely on target, and there would be no need to worry about events such as forest fires, drought, disease, insect defoliation, and so forth, as these factors would already be perfectly accounted for in the model. Obviously, such a world is sheer fiction, and the uncertainties of the real forest environment force us to continually monitor the forest and to occasionally revise the preprogrammed harvesting decisions when reality begins to depart too greatly from our model.

Faced with the real-world uncertainties surrounding forest management, we are immediately led to the following *monitoring problem*: what is it we need to know about the forest in order to decide when the current plan needs *revision*, and, if a revision is called for, how should it be carried out in order

to bring the overall forest plan back on target? Clearly, this is a very vague, general statement and must be translated into more formal mathematical terms before system and control methodology can hope to shed light on the situation.

For the sake of definiteness, let us assume that our forest model is of the LQG type. Furthermore, we assume that a forest plan $\bar{u}(t)$ is given which, in the absence of uncertainty, generates a corresponding system trajectory $\bar{x}(t)$. We are concerned about deviations of the real system from the programmed state $\bar{x}(t)$ and changes in the plan needed to return the system to the desired level $\bar{x}(t)$. Introducing the variables $x(t) = x^*(t) - \bar{x}(t)$, $u(t) = u^*(t) - \bar{u}(t)$, where $x^*(t)$, $u^*(t)$ represent the actual decisions implemented and the corresponding trajectory, the LQG problem is to minimize

$$\sum_{t=0}^{T-1} \left[(x(t), Q(t)x(t)) + (u(t), R(t)u(t)) \right] + (x(T), Mx(T)),$$

subject to

$$x(t+1) = F(t)x(t) + G(t)u(t), \qquad x(0) = x_0.$$

Here the matrices Q, R, F, G, M are as described in the Dixon-Howitt model discussed earlier. Clearly, the entries of $Q(\cdot)$ represent the weights we attach to keeping the various components of $x^*(t)$ near $\bar{x}(t)$, while those of $R(\cdot)$ are the weights attached to deviations of the plan from its preprogrammed levels. In most forest environments, it can be expected that the costs in R are greater than those represented by Q. Under the assumption that $x(t)$ can be measured *exactly*, the foregoing linear-quadratic problem has the known solution

$$u(t) = -K(t)x(t),$$

where

$$K(t) = \left[R(t) + G'(t)P(t+1)G(t) \right]^{-1} G'(t)P(t+1)F(t),$$

with $P(t)$ satisfying the Riccati difference equation

$$P(t) = Q(t) + F'(t)P(t+1)F(t) - F'(t)P(t+1)G(t)$$

$$\times \left[G'(t)P(t+1)G(t) + R^{-1}(t)G'(t)P(t+1)F(t) \right]$$

$$P(T) = M.$$

1. Monitoring Problem (Deterministic Version)

Since, in general we cannot measure $x(t)$ directly, a more precise statement of the monitoring problem is as follows: Given the measurements $y(t) = H(t)x(t)$ and the decisions $0 \leqslant t \leqslant t^*$, is it possible to uniquely reconstruct the system state $x(t^*)$?

The solution to the monitoring problem is given by the following basic result from linear system theory. Assume that the matrix pair $(F(\cdot), H(\cdot))$ is *completely constructible*, i.e., the matrix

$$M(s, t^*) = \int_s^{t^*} \Phi_F'(\sigma, t^*) H'(\sigma) H(\sigma) \Phi_F(\sigma, t^*) \, d\sigma$$

has rank n for all $s \leqslant t^*$ (here we shift over to continuous time to avoid notational complexity). Let $L(t)$ be any matrix function such that $F(t) - L(t)H(t)$ is a stable matrix function, i.e., the equation

$$\dot{w}(t) = [F(t) - L(t)H(t)] w(t), \qquad w(0) = w_0,$$

is such that $w(t) \to 0$ as $t \to \infty$. Note also that $\Phi_F(\cdot, \cdot)$ is the fundamental matrix for the function $F(\cdot)$, i.e., $\partial \Phi / \partial \sigma = F(\sigma)\Phi$, $\Phi(t^*, t^*) = I$. Then the system given by

$$\frac{d\hat{x}}{dt} = F(t)\hat{x}(t) + L(t)[y(t) - H(t)\hat{x}(t)] + G(t)u(t),$$

$$\hat{y}(t) = \hat{x}(t)$$

is an asymptotic state estimator of $x(t)$. That is, $x(t) - \hat{x}(t) \to 0$ as $t \to \infty$, with the rate of approach depending upon the rate at which $w(t) \to 0$.

Thus, after a sufficient period of time [determined by the arbitrarily selected matrix $L(t)$], $\hat{x}(t)$ is as close as desired to the unknown state $x(t)$. The point of the constructibility requirement on $(F(\cdot), H(\cdot))$ is to guarantee that a matrix function $L(t)$ exists with the desired stability properties. If $L(t)$ is selected so that $\hat{x}(t) \to x(t)$ rapidly, we can expect that use of the estimated state $\hat{x}(t)$ for computing the feedback law $u(t)$ will provide results close to those obtained if we knew the exact state $x(t)$. Of course, $\hat{x}(t)$ is not exactly equal to $x(t)$, so that the control actually applied, $\hat{u}(t) = -K(t)\hat{x}(t)$, is not equal to the theoretical optimal control $u(t) = -K(t)x(t)$. However, it can be shown that if the original system is completely controllable, then the error made in applying the wrong control does not become magnified in later stages, but rather the "wrong" control $\hat{u}(t)$ asymptotically approaches the "correct" control $u(t)$.

The preceding result also provides a basis for the solution of an important variant of the monitoring problem which arises when the measurement process $y(t)$ [i.e., the matrices $H(\cdot)$] is not prescribed in advance, but is to be selected by the forest manager. Since the matrix function $H(\cdot)$ determines exactly what components (or linear combinations of components) of the forest can be measured, the system designer would like to have $H(\cdot)$ be as simple as possible, consistent with the ability to use the measurements generated to reconstruct $x(t)$. From the preceding development, it is clear that the simplest such $H(\cdot)$ would satisfy the following requirements.

(1) the pair $(F(\cdot), H(\cdot))$ is constructible [needed to guarantee the existence of a matrix $L(\cdot)$ generating a stable state estimator];

(2) the matrix $H(\cdot)$ has as few rows as possible consistent with the constructibility requirements [the number of rows of $H(\cdot)$ determines how many system outputs are actually measured].

Clearly, the generation of efficient and simple monitoring schemes will depend upon precisely which combination of components of $x(t)$ can be measured directly. The simplest case would be if the sum $\sum_{i=1}^{n} \alpha_i x_i(t)$ could be measured for some set of $\{\alpha_i\}$. Then we would have only a single output, and it can be shown that almost every nonzero choice of coefficients $\{\alpha_i\}$ would satisfy the constructibility requirement above. Basically, the foregoing discussion answers the question of *what* variables we should measure in order to generate a suitable feedback controller.

2. *Monitoring Problem (Stochastic Version)*

The above development assumed that the system output $y(t)$ was measured without error. Furthermore, it was tacitly assumed that the system dynamics were known exactly. If we relax these assumptions to the extent of letting uncorrelated additive gaussian noise processes corrupt the observation and account for uncertainty in the dynamics, then we are faced with the stochastic control problem outlined earlier and depicted in Figure 1. Let us now consider this situation in a bit more detail.

Let the system dynamics be given by

$$\frac{dx}{dt} = Fx(t) + Gu(t) + v(t)$$

with the observations

$$y(t) = Hx(t) + w(t).$$

We assume that the noise processes $v(\cdot)$ and $w(\cdot)$ are uncorrelated, zero-mean,

white gaussian processes with known covariances

$$E[v(t)v'(\tau)] = \hat{Q}(t)\delta(t - \tau),$$

$$E[w(t)w'(\tau)] = \hat{R}(t)\delta(t - \tau).$$

If we want to minimize the expected value of the quadratic criterion

$$\int_0^T [(x, Qx) + (u, Ru)]\, dt, \qquad Q \geqslant 0, \quad R > 0,$$

the problem is to express the optimal control $u^*(t)$ as a function of the observations $y(s)$, $0 \leqslant s \leqslant t$.

The well-known solution of the problem is to estimate the system state by computing a minimum-variance estimator $\hat{x}(t)$ according to the equation

$$\frac{d\hat{x}}{dt} = F(t)\hat{x} + G(t)u + P(t)H'(t)R^{-1}(t)[y(t) - H(t)\hat{x}],$$

$$\hat{x}(t_0) = 0,$$

where $P(\cdot)$ satisfies

$$\frac{dP}{dt} = F(t)P + PF'(t) - PH'(t)\hat{R}^{-1}(t)H(t)P + \hat{Q}(t),$$

$$P(0) = P_0 = E[x_0 x_0'].$$

The control $u^*(t)$ is then given by $u^*(t) = -K(t)\hat{x}(t)$, where $K(t)$ is generated exactly as in the deterministic case. This is an application of the so-called "separation theorem" of LQG theory, enabling us to separate the problem into two phases: estimation of $\hat{x}$ and control by standard deterministic means.

For forestry management and monitoring, the interesting problem here is the generation of the estimate $\hat{x}(t)$. The information needed to implement the foregoing scheme is the covariance matrices $\hat{Q}$, $\hat{R}$, and P_0. A discussion of how to determine these matrices under a specified observation process is given in [11]. The matrix $P(t)$ represents the covariance of the *error* made in estimating $x(t)$ by $\hat{x}(t)$ and, as a consequence, we see that the magnitude of the entries in $P(t)$ tells us how accurately $\hat{x}(t)$ approximates $x(t)$. Since $\hat{Q}$ and $\hat{R}$ determine $P(t)$, we have a definite prescription for determining how

accurate our measurements of both the system and its output must be in order to achieve a given estimator accuracy.

From a measurement-design point of view, the problem is composed of two parts: (1) what observations to take and (2) how accurate the observations need to be. The first part has exactly the same answer as in the deterministic case. Namely, choose the observation matrix $H(t)$ as simple as possible (minimal number of rows), consistent with the constructibility requirement of the pair (F, H). The accuracy question is answered by determining how small the covariances P_0, $\hat{Q}$, and $\hat{R}$ must be in order to make $P(t)$ small. From the equation for P we see that there is clearly some tradeoff in this regard, since a large uncertainty in the dynamics could be at least partially compensated for by very accurate measurements of the output. Detailed treatment of these and many other issues of sensitivity and error estimates may be found in any of a number of standard works on stochastic control and Kalman filtering (see [1], [4]).

III. MONITORING, FEEDBACK, AND OPEN-LOOP MODELS

The discussion of the preceding sections makes it evident that there is a basic practical conflict in forestry modeling: On the one hand, the LP-type models such as Models I and II have the capability of dealing with large numbers of variables in a small amount of computing time provided that the variables and constraints satisfy certain linear relationships. However, such models are inherently incapable of adapting to changes in the problem relationships, parameter values, or constraints, and generally must be entirely rerun when such changes occur. On the other hand, feedback-control structures such as the Dixon-Howitt LQG model are inherently well adapted to responding to problem changes, but generally require substantial computing resources unless the linear-quadratic-gaussian structure is present as in the Dixon-Howitt situation. In the case of general nonlinear dynamics, to obtain a feedback control law one would have to employ a dynamic-programming-type scheme with its well-known attendant dimensionality difficulties. In this section, we shall explore various avenues for combining the two approaches to take advantage of the computational virtues of the open-loop models without sacrificing the adaptive capability of the feedback schemes.

To begin with we consider deterministic models of the LP type. To facilitate the discussion, we shall take the canonical LP model as our prototype. Thus, our basic model is

$$\max_{\{x\}} (c, x),$$

subject to the constraints

$$Ax \leqslant b,$$

$$x \geqslant 0,$$

where x and c are n-dimensional vectors, b is an m-vector, A is an $m \times n$ matrix and $(\cdot, \cdot)$ denotes the vector inner product. Both Model I and II are of the above form with suitable choice of A, b, and c.

Similarly, as our prototypical feedback model we wish to maximize (over the $\{u_i\}$)

$$J = \sum_{i=1}^{N} g_i(x_i, u_i),$$

where x and u are connected by the relation

$$x_{i+1} = f_i(x_i, u_i), \qquad x_0 = \alpha,$$

$x_i \in R^n$, $u_i \in R^m$. If the functions $g_i(\cdot, \cdot)$ and $f_i(\cdot, \cdot)$ are linear, then it is an easy exercise to see that this problem can be formulated as a canonical LP problem by introducing the new $(N+1)(n+m)$-dimensional vector

$$z = \left(x_0^1, x_0^2, \ldots, x_0^n, x_1^1, \ldots, x_1^N, \ldots, x_N^n, u_1^1, \ldots, u_N^m \right)$$

into the LP format, and choosing A, b, and c appropriately. (*Remark*: It may also be necessary to augment z if x_i and/or u_i can take on negative values. We omit this technicality here for ease of exposition.)

In the usual situation, the function g_i and/or f_i are nonlinear and the above problem is solved by some variant of dynamic programming. If we introduce the new function

$\phi_N(\alpha) =$ minimum value of J when there are N stages remaining, the system is in state α, and an optimal policy is used,

then an easy application of the principle of optimality [15] yields the recurrence relation

$$\phi_N(\alpha) = \min_{u_N} \left[g(\alpha, u_N) + \phi_{N-1}(f(\alpha, u_N)) \right],$$

$$\phi_1(\alpha) = \min_{u_1} \left[g(\alpha, u_1) \right].$$

(Here we assume the functions f and g are independent of the stage. If not, the above reference should be consulted for the necessary modifications to the dynamic-programming algorithm). For now, the important feature of the above DP scheme is that the optimal decision u_i^* is determined sequentially as a function of the current state, i.e., $u_i^* = u_i^*(\alpha)$. The computational problems arise from the need to know the function $\phi_{N-1}(\cdot)$ for all values of its argument in order to calculate $\phi_N(\cdot)$. If the state vector is quantized to, say, k levels in each component, then we must store k^n numbers in high-speed memory. Since it may commonly occur that $k \cong 10^2$ and $n \cong 10\text{--}20$, we see that the amount of storage required for a straightforward application of DP can be astronomical. This is a manifestation of Bellman's "curse of dimensionality." Before discussing ways to lift this curse, in Table 1 we briefly compare LP and DP with regard to practical implementation in a forest management environment.

From Table 1 we conclude that LP and DP are complementary in the sense that the strengths of each are the weaknesses of the other. In the procedures sketched below, we indicate various means for combining the two methods to provide a solution procedure which is superior to each taken separately.

1. Reformulation of LP Model as DP Model

We have outlined above how the prototypical control problem can be considered as an LP problem if the objective function and dynamics are linear. Now let us see how to formulate the canonical LP model as a dynamic program.

TABLE 1

LP VERSUS DP

Linear programming (Models I and II)	Dynamic programming (Dixon and Howitt)
Can handle high-dimensional state	Restricted to low-dimensional state
Linear objective and linear side con-straints only	General nonlinear functions can be used
Open-loop solution	Closed-loop (feedback) solution
Stochastic features difficult to incorporate	Easy to include stochastic and adaptive aspects
Very efficient algorithms available for large problems	Substantial computational requirements for moderate-size problems

Consider the objective function

$$J = \sum_{i=1}^{n} c_i x_i$$

and constraints $Ax \leq b$, $x \geq 0$. Let

$\phi_n(b_1, b_2, \ldots, b_m)$ = value of J when there are n stages remaining, the constraint resource available is $b = (b_1, \ldots, b_m)$, and an optimal policy is used.

In [15] it is shown that ϕ_n satisfies the relation

$$\phi_n(b_1, \ldots, b_m) = \max_{0 \leq x_n \leq z_n} \left[c_n x_n + \phi_{n-1}(b_1 - a_{1n} x_n, \ldots, b_m - a_{mn} x_n) \right],$$

where $z_n = \min_i (b_i / a_{in})$. The initial function is

$$\phi_1(b_1, \ldots, b_m) = \max_{0 \leq x_1 \leq z_1} \left[c_1 x_1 \right].$$

The above formulation shows quite clearly that the resource constraints b_i play the role of state variables in the DP version of the LP problem. The decision variables are then determined as a function of the amount of resource remaining to be allocated. The computational tractability of the foregoing scheme is determined by the number of resource constraints m. If m is small, say ≤ 10, we can expect a more or less routine solution. For m large, the curse of dimensionality arises and dimensionality-reduction procedures must be employed.

In the forestry Model I, the number m equals the number of age classes which contain hectares in the initial period. In Model II, m equals the number of periods before zero in which the oldest age class present was regenerated plus the total number of periods. While in both these cases m is likely to be larger than is comfortable for routine computation, it should be noted that the side constraint $Ax \leq b$ is strict equality, i.e., $Ax = b$ in both Models I and II. Thus, the additional structure this fact provides, coupled with the fact that we know $\phi(b)$ must be a linear function of b, may be enough to reduce the computation to digestible proportions. If so, then the above formulation will yield a solution $x^* = x^*(b)$, i.e., the optimal harvest plan as a function of the number of hectares of each age class present.

2. *LP/DP Successive Approximations*
The most natural procedure for generating a feedback control from the LP model is to use the solution to the LP model as the initial approximation in

some sort of successive-approximation scheme. We have already seen that the standard control problem can be formulated as a mathematical-programming problem by partitioning the variables in the programming problem into two classes: states and controls at each time period. Thus, when we solve the programming problem we generate a nominal control law $u^0(t)$ and a corresponding nominal state trajectory $x^0(t)$. The question is how to use $u^0(t)$ and $x^0(t)$ to generate a feedback control $u(x(t))$.

Since the one class of high-dimensional control problems for which it is possible to routinely calculate feedback controls is the LQ case, the most sensible approach is to linearize the system dynamics and quadraticize the costs, using the nominal state and control x^0 and u^0. Thus, if the actual problem is to minimize

$$\sum_{i=1}^{N} g(x_i, u_i)$$

with

$$x_{i+1} = f(x_i, \dot{u}_i), \qquad x_0 = x_0,$$

then we calculate the quantities

$$F(i) = \left(\frac{\partial f}{\partial x}\right)_{x=x^0, u=u^0} \qquad G(i) = \left(\frac{\partial f}{\partial u}\right)_{x=x^0, u=u^0},$$

$$Q(i) = \left(\frac{\partial^2 f}{\partial x^2}\right)_{x=x^0, u=u^0}, \qquad S(i) = \left(\frac{\partial^2 g}{\partial x \partial u}\right)_{x=x^0, u=u^0},$$

$$R(i) = \left(\frac{\partial^2 g}{\partial u^2}\right)_{x=x^0, u=u^0}$$

and solve the LQ problem

$$\min \sum_{i=1}^{N} \left[(z_i, Q(i)z_i) + 2(z_i, S(i)v_i) + (v_i, R(i)v_i) \right],$$

$$z_{i+1} = F(i)z_i + G(i)v_i, \qquad z_0 = z_0,$$

where $z_i = x_i - x_i^0$, $v_i = u_i - u_i^0$. That is, the states and controls of the LQ problem represent the *deviations* (or corrections) from the nominal state and control generated by the mathematical-programming solution. Obviously, once the controls $\{v_i\}$ and states $\{z_i\}$ are computed, the process can be

repeated with the v_i and z_i taking the roles of u^0 and x^0. In this way, we could hope to converge to the optimal *feedback* control $u_i^*(x_i)$ and corresponding trajectory.

As in all successive-approximation schemes, the utility of the method hinges critically upon the initial approximation. In the forestry situation, however, this should be a lesser problem, since the outputs from either Model I or II should provide a reasonably good approximation to the true solution. Thereafter, the computed feedback controls $\{v_i(z)\}$ represent the *deviations from the original plan* needed in order to account for new information that becomes available. This new information is represented by the quantity $\{z_i\}$, which measures the deviation of the system state from its planned level. Furthermore, the cost matrices $\{Q(i), S(i), R(i)\}$ account for the importance which one attaches to deviations in various aspects of either the plan or the state. So, by introduction of a suitable measure (or norm) into the control space, we could use the quantity $\| v_i \|$ as an indicator as to whether or not it is necessary to revise the plan. If this number is above a certain threshold, then the plan would be revised and we would begin the process over again with a new nominal u^0 and x^0. If not, then minor changes to the current plan would be made as dictated by the control law $v_i(z_i)$.

3. *LP within DP*

Often it may not be feasible to employ the LQ successive-approximation scheme outlined above due to discontinuities of one sort or another in the cost and/or dynamics. However, even in these cases the standard dynamic-programming algorithm will produce the desired feedback control if the dimensionality difficulties referred to earlier do not arise. In other cases, there may be sufficient linear structure present to effectively employ an LP algorithm within the overall dynamic-programming framework in the following manner.

Assume that the problem we want to solve is that of minimizing

$$\sum_{i=1}^{n} g(x_i, u_i),$$

$$x_{i+1} = f(x_i, u_i), \qquad x_0 = \alpha.$$

We know that the principle of optimality yields the recurrence relations

$$\phi_n(\alpha) = \min_{u_n} \left[g(\alpha, u_n) + \phi_{n-1}(f(\alpha, u_n)) \right],$$

$$\phi_1(\alpha) = \min_{u_1} \left[g(\alpha, u_1) \right].$$

So, in the calculation of $\phi_n(\alpha)$, for fixed α, we must search over all admissible values of u_n and compare the quantities $g(\alpha, u_n) + \phi_{n-1}(f(\alpha, u_n))$. If the functions $g(x, u)$ and $f(x, u)$ are linear in u, for fixed x, then the above minimization is a simple LP problem over the variables comprising the vector u. Thus, instead of doing a brute-force search over the components of u, we can employ the vastly more efficient simplex-based LP algorithms.

Fortunately, in some cases of interest g and f do possess the requisite linearity. In fact, it is often this linearity which makes it impossible to use the LQ approach sketched earlier. If g and f are nonlinear, then we would employ a multiple linear regression scheme for each fixed α in order to replace f and g by their best linear approximation. Since such regression packages are readily available, there is no major obstacle to carrying out such a computation, and, as such operations involve only linear algebraic equations, the computing burden is not too onerous as long as it is not necessary to redo the regression too frequently. This will be dependent upon the degree of nonlinearity of the functions f and g. In the forestry problems f and g are only mildly nonlinear, in general, so the above procedure may be expected to meet with some success.

4. *Stochastic Effects and Adaptive Policies*

The methods outlined above have been developed under the assumption that no uncertainty exists in the system objective, dynamics, or constraints. If this is not the case, then a qualitatively different set of circumstances arises under which some of the preceding ideas either do not apply or require substantial modification.

Consider the case in which the performance criterion is known, but the system dynamics contain a stochastic component, i.e., the evolution of the system takes place according to the dynamics

$$x_{i+1} = f(x_i, u_i, r_i), \qquad x(0) = x_0,$$

where r_i is a random variable subject to a *known* probability density function (pdf) $p(r_i)$. For simplicity, assume all $\{r_i\}$ obey the same probability law. In this case, the DP successive-approximation procedure is of no obvious value, since the system trajectory is not *uniquely* specified by the choice of control. Thus, the state-control pair (x^i, u^i) needed to generate the next approximate control cannot be calculated in a well-defined way. On the other hand, the LP-within-DP idea carries over to this setting without changes other than minor modifications to the DP recurrence relation to account for the stochastic vector r_i.

Another way in which stochastic factors can enter the forest model is exhibited in Dixon and Howitt's work, where an originally deterministic

nonlinear model is replaced by a stochastic linear model, with the random variables introduced to account for the loss of information due to the approximation. Such an approach leads to the stochastic LQG control problem discussed above. It should be clear that randomness of this sort is of a quite different character than that entering into the original dynamics themselves. Here, the randomness is injected into the model as an attempt to compensate for an approximation process and has nothing directly to do with one's ignorance of the system. Thus, such randomness is a consequence of the methodological approach and must be carefully distinguished from that which arises from basic measurement and conceptual uncertainties about the dynamics and/or objectives of the system.

With additional computational complications, whatever procedures are used to deal with stochastic control when the pdf of r is known, can be "souped up" to handle the so-called adaptive control situation in which the pdf of r is unknown and must be inferred as the process unfolds through a combination of probing to learn about the system and control of its behavior.

TABLE 2

COMPARISONS OF OPEN-LOOP → FEEDBACK METHODS

Method	Strengths	Weaknesses
Reformulate LP as DP	Can deal with stochastic/ adaptive situation with respect to parameter values and/or constraints. Provides feedback law as a function of resources available, i.e., solves a family of forest problems.	Computational burden quite high, especially if the forest contains many age classes.
LP/DP successive approximation	Good, quick algorithms available.	Needs good initial approximation. Requires sufficient degree of *nonlinearity* in costs and dynamics. Difficult to handle stochastic/adaptive situations.
LP within DP	Reduces DP computing time substantially through efficient search.	Most effective only if system dynamics are "almost linear."

Since it would take us too far afield to go into the details of these matters here, we refer to [3, 16] for a complete discussion with many worked examples.

In Table 2 we summarize the strengths and weaknesses of the various approaches sketched above for generating feedback policies from open-loop models.

IV. MONITORING THEORY AND FOREST MANAGEMENT

The preceding discussion has centered upon the problem of reformulating standard forest-management models in closed-loop form, in order to admit the possibility of feedback decisionmaking. In this section, we take up another aspect of the monitoring problem having to do with the stability of a given policy under variations in the problem environment—in other words, under what circumstances we can adhere to a given plan without serious degradation of overall system performance, and what aspects of system behavior provide the signal to indicate that a change of plan is required. In addition, we shall also examine the question of how to control the system when the measurements are interrupted.

If we are using either Model I or II to determine our management policy, then it is possible to make use of more or less standard routines to analyze the stability of the computed plan with respect to changes in either the criterion function or the resource constraints. Given the LP model

$$\max_{x \geq 0} (c, x), \qquad Ax \leq b,$$

the type of questions that arise are of the following sort:

(1) How much can the vector c be changed before the optimal solution (plan) will no longer be optimal?

(2) For a given change in c, how do we proceed to calculate a new optimal plan?

(3) How much change in the resource constraint b can be allowed before the computed plan is no longer optimal?

(4) If a given change in resources b results in the plan no longer being optimal, how can we determine the new optimal plan from the previous one?

(5) How can we handle the addition of a new variable x_{n+1} without having to restart the entire solution procedure?

(6) How can the insertion of a new constraint be incorporated into the system?

All of these questions form the basis of more or less standard sensitivity analysis in mathematical programming, and for the most part their resolution is part of the output of modern LP computer packages. For this reason, we shall not waste space in describing the results here, referring to standard references [10, 12] for details. We only note that it is reasonable to suspect that in the forest models we would be somewhat more likely to encounter situations in which the cost function changes rather than the resources available, although the latter situation may arise from shifting political or economic factors.

As a technical aside, it is worthwhile to note that the dynamic-programming formulation of the LP problem explicitly exhibits the optimal cost as a function of the "natural" state variables, the components of the resource vector b. Thus, it is reasonable to think about deriving an expression for how the costs change with changes in b, i.e., the quantity

$$\frac{\partial \phi_n}{\partial b_i}(b), \quad i = 1, 2, \ldots, m.$$

Explicit computation of this quantity for various choices of i then tells us what components of b must be "monitored" in order to calculate the costs of projected changes in the plan.

The preceding sensitivity questions also suggest an intriguing type of "inverse" problem, namely, if a given plan is suggested, to characterize all cost vectors c for which that plan is optimal. This amounts to picking one vertex of the convex polytope formed by the constraints and finding all separating hyperplanes containing this vertex. This problem is particularly interesting for forest planning, since there is usually considerable uncertainty in what the actual cost vector c should be. Thus, for a given projected plan one could compute upper and lower bounds for all components of c which would be optimal for that plan. In this case, a change of plan would be warranted only if it could be shown that some component of the cost was outside the range allowable for that plan. Hence, the monitoring problem in this case would revolve about accurate assessments of the cost vector c, rather than the constraint vector b.

If our basic model is of the closed-loop type, then we have a quite different array of mathematical machinery with which to study how the plan (control) changes if various aspects of the system change. For example, in the LQ case we have already seen that the optimal closed-loop plan is given by

$$u^*(t) = -K(t)x(t),$$

where

$$K(t) = [R(t) + G'(t)P(t+1)G(t)]^{-1}G'(t)P(t+1)F(t),$$

with

$$P(t) = Q(t) + F'(t)P(t+1)F(t) - F'(t)P(t+1)G(t)$$
$$\times \left[G'(t)P(t+1)G(t) + R^{-1}(t)G'(t)P(t+1)F(t) \right],$$

$$P(T) = M.$$

The above expressions make it quite evident how to account for any changes in the system costs (Q, R, M) or dynamics (F, G). We simply put in the appropriate values for the current time period and compute the quantity K. Furthermore, the above expressions can be used to derive explicit analytic representations for the sensitivity functions $\partial K / \partial F_{ij}, \partial K / \partial G_{ij}$, etc. [2].

A considerably more general type of situation arises when we allow the possibility of generating a feedback control $u(t, x)$ for the system

$$\dot{x} = f(x, u, \lambda), \qquad x(0) = x_0,$$

in order to insure that the trajectory which minimizes

$$\int_0^T g(x, u, \lambda) \, dt$$

is as insensitive as possible to changes in the parameter vector λ. More specifically, assume that the optimal state trajectory is $x_0(t)$ when the parameters are at some nominal value λ_0. Associated with the states $x_0(t)$ there is an optimal *open-loop* control $u_0(t)$. Now assume the parameter values change to some new, unknown value λ. Then we would like to find a closed-loop control $u(x, t)$ such that $x(t)$ remains "close" to $x_0(t)$ for all variations λ in the parameters.

For this problem, it can be shown that the minimum-sensitivity feedback compensator is similar in structure to classical input compensators and that the closed-loop system error is related to the transfer function for measurement noise, which is the limiting factor for classical sensitivity-reduction techniques. Under the assumption of small parameter variations and quadratic sensitivity cost terms, the general problem reduces to one for which an explicit solution can be obtained for a linear feedback gain function. The conditions developed for the minimum sensitivity gain effectively extend the LQ theory to include unknown disturbances in the system parameters. The details of these results, along with many examples and comparisons with dynamic parameter estimators (Kalman filters or Luenberger observers), are given in [9].

All of the preceding development has been based on the assumption that at each moment *some* measurement could be made of the system output. The problem has been to decide what measurement to take and/or how to process the measurement. Now we want to consider briefly the situation in which at any given moment, there is a definite probability that no measurement at all can be taken. In this case, we must make the "best" current decision on the basis of the information available, information which may or may not include a current measured output.

Since there may be several successive stages for which the state of the system is unknown, we specify the state of the system by three quantities— m, n, and c. Here n is the number of stages remaining in the process, m is the number of stages in the past at which the controller last had perfect information, and c is the physical state of the system at that time. For simplicity, assume we wish to minimize the expected value of some function ψ of the terminal state of the process, i.e.,

$$J = \mathop{E}_{\{r(i)\}} \left[\psi(x_N) \right];$$

the system dynamics are

$$x(k+1) = f[x(k), u(k), r(k)], \qquad x(0) = c;$$

and the $\{r(i)\}$ constitute a set of independent random vectors with common distribution function. Furthermore, assume that at any stage there is a fixed probability p that we shall not be able to observe the actual physical state of the system.

Let

$z_m(c) =$ the random variable designating the observed physical state of the system after m stages have passed without observation of the system, given that the state m stages back was c.

Note that $z_m(c)$ will depend upon the random vectors $r(0), r(1), \ldots, r(m-1)$ and the control vectors $u(0), u(1), \ldots, u(m-1)$.

We note that with n stages remaining, with probability $1-p$ we may be able to observe the physical state, thereby obtaining a physical state denoted by the quantity $z_{m+1}(c)$, conditioned on the random elements $r(0), \ldots, r(m+1)$; then the number of remaining stages becomes $n-1$, and the number of stages in the past at which we last had perfect information becomes 0. If, on the other hand, we cannot observe the system, then the number of remaining stages still becomes $n-1$, m is transformed to $m+1$, and the last perfect observation is unchanged. Let $I(m, n, c)$ denote the value of the criterion

when m, n, and c are as above and we use an optimal policy. Then using the principle of optimality and forward dynamic programming, we obtain

$$I(m,n,c) = \min_u \Big\{ p[I(m-1,n-1,c)]$$

$$+ (1-p)\, \underset{r}{E}\big[I(0,n-1,z_{m+1}(c))\big]\Big\}.$$

The minimizing u is a function of m, n, and c, i.e., $u_{\min} = u(m,n,c)$.

To make use of the equation for I we must determine the conditional expected value of $I(0,n-1,z_{m+1}(c))$, given the information that for $m+1$ successive stages the actual state of the system could not be observed and that the last known state was c. At the first stage, the control vector $u(0,m+n,c)$ (even though $m=0$ here, it is convenient to adopt this notation for a moment) was used, resulting in the new state

$$x(1) = f[c, u(0,m+n,c), r(0)].$$

At this stage, with no information on the actual state, the control $u(1, m+n-1, c)$ was used. This led to the stochastic state

$$x(2) = f[x(1), u(1,m+n-1,c), r(1)].$$

We see that the random variable $z_{m+1}(c)$ depends upon the control vectors $u(k, m+n-k, c)$ for $k = 0, 1, \ldots, m$. Thus, the equation for I determines functions $I(m,n,c)$, $u(k, m+1-k, c)$ for $0 \leqslant m+n+1 \leqslant N$, $k = 0, 1, \ldots, m+1+n$. It is an implicit system with no straightforward way of determining the functions one at a time. For computational purposes, a method of successive approximations would most likely be used.

"Interrupted" control processes of the foregoing type are quite likely to occur in many forestry management environments when budgetary cuts, climatic disturbances, political changes, and other factors cause the measurement system to fail over some period of time. These ideas are also of obvious relevance to the situation in which decisions must be made on the basis of incomplete data about the past, i.e., when the data have gaps. Unfortunately, the analytic and computational study of interrupted control processes has not progressed beyond the functional equation given above for I. In view of the substantial advances in computing power in recent years, it would seem appropriate to reexamine the practical feasibility of using the equation for I in a real decisionmaking environment.

Finally, we observe that in some instances it is more important to insure that the system trajectory $x(t)$ remains within some admissible or desired

region of state space, than it is to optimize some economic criterion. In other words, we may be willing to sacrifice economic efficiency for an enhanced ability of the system to respond to unknown (and perhaps unknowable) disturbances in its dynamics. Such a criterion of system performance is outside the bounds of optimal-control theory and requires the development of a theory of system *resilience*. Some work along these lines, together with the connections between resilience and catastrophe theory, is reported in [7].

In a related direction, we may be interested in determining all, or at least one, policy which will transfer the forest from its current state into some desired region of state space at the end of some prescribed period. This is a problem of *reachability* and has been fairly extensively studied, especially for linear systems. For instance, if the system dynamics are given by

$$x(t+1) = Fx(t) + Gu(t), \qquad x(0) = 0,$$

with $x \in R^n$, $u \in R^m$, then it can be shown that a given state x^* is reachable after at most n time periods if and only if x^* is contained in the vector space generated by the set

$$\{g_1, g_2, \ldots, g_m, Fg_1, Fg_2, \ldots, Fg_m, F^2g_1, \ldots, F^2g_m, \ldots, F^{n-1}g_1, \ldots, F^{n-1}g_m\},$$

where g_i is the ith column of G. If the dynamics are nonlinear, the situation is considerably more complicated, depending upon the exact type of nonlinearity present. Full details of these matters may be found in [6, 8].

V.   CONCLUSIONS AND RECOMMENDATIONS

Our main conclusion from investigation of the forestry management models currently in use is an almost self-evident one: the models contain no capability to adapt to either unanticipated changes in the internal sociopolitical structure or surprises in the operating physical environment. To truly account for such phenomena, of course, one would need a well-developed nonprobabilistic "theory of surprises," as well as a reasonably complete theory of "anticipatory" systems. Neither of these system-theoretic desiderata is, as yet, a reality, so we must conclude that the only short-term remedy that can be offered for forestry management is to include some type of feedback loop in the system model. The uncertainty present in all aspects of the forest system makes it quite evident that the open-loop Models I and II and their successors will be inherently incapable of accommodating themselves to the

changes in the operating environment of the forest manager. Some type of feedback model measuring critical system variables and adjusting the open-loop policy (plan) must be injected into the forest-planning process if there is to be any hope of making the system responsive to unpredictable outside changes.

Operationally, there are several paths that such a feedback modeling process could follow, depending upon the nature of the uncertainty faced by a particular forest, the resources available to cope with the uncertainty (i.e., money, manpower, computer power, time, measuring equipment, etc.), the accuracy needed in the final result, and so forth. Consequently, we have sketched several alternative routes above but can recommend no one approach which is uniformly most powerful. Each situation must be examined before an appropriate recommendation can be made. The principal consolation here is that we definitely feel that all of the procedures mentioned are computationally feasible, given the current and projected state of computer resources.

By way of a future research program based upon the foregoing ideas, we feel that the most important problem is to actually start working with some *specific* forest model and to examine all of the preceding questions within the context of the given system. There is probably very little of lasting value that can be accomplished by examination of an "abstract" forest. The joy, as always, is in the details, and the details can only come from a real forest-management situation.

On the other hand, a valid case *could* be made for instituting a research program on the topic of surprise and anticipation. Such a program would obviously have implications far beyond the narrow confines of forest management. However, even such an abstract system study needs a basis in reality upon which to focus the research, and we feel that the forestry area would be particularly promising, since there is a substantial body of data available on past surprises, historical management reactions, and the results of trying to make decisions on the basis of the projected future of the system.

APPENDIX. SOME GENERAL SYSTEM-THEORETIC ISSUES IN FOREST MODELING

1. Observables

In a fairly general sense, we can say that the basic forestry monitoring problem reduces to the following: given a system state space X and a collection of observables $f_i: X \to R$, what can we know about the true state of the system by measuring the real numbers given by the observables?

Let us assume, for a moment, that we have only a single observable f. Then it is clear that if f is one-to-one, given a measured value $r \in R$, a *unique*

state $x(r)$ can be inferred. This is a general version of our notion of observability described earlier in this report. Most of the time, though, the observable f is not one-to-one, which means that several physically distinct states *appear* to be the same from the standpoint of the observer (modeler). In other words, we can define a relation F on the state set X by saying that x_1 is F-related to x_2 if $f(x_1) = f(x_2)$. It is trivial to verify that F is an *equivalence* relation, thereby enabling one to partition the state space X into equivalence *classes* of states. A crucial question now arises for the decisionmaking process: are states which are equivalent under the observable f also equivalent with respect to the purposes of control? This question is tantamount to defining a second equivalence relation on X —call it C—determined by saying that x_1 is C-related to x_2 if the result of using a feedback control law $u(x)$ is invariant with respect to the substitution $x_1 \to x_2$. Such a relation also implicitly defines a new observable which we can term g. Note that this does not require $u(x_1) = u(x_2)$, although such equality is clearly sufficient. In fact, the answer to the question is provided by the notion of the *linkage* of two observables.

We have the following definition of linkage between observables f and g.

DEFINITION A1.   Let X be a set of states, and let $\pi_f: X \to X/F$, $\pi_g: X \to X/C$ be the natural mapping of states to their respective equivalence classes under f and g. Given a class $[x]_f \in X/F$, consider the collection of C-classes which intersect $[x]_f$, i.e., the set

$$Z = \pi_g \circ \pi_f^{-1} [x]_f.$$

We say that

(1) g is *totally linked* to f at $[x]_f$ if $Z = $ (a single C-class).
(2) g is *partially linked* to f at $[x]_f$ if $Z = $ (more than one C-class) $\subset X/C$.
(3) g is *unlinked* to f at $[x]_f$ if $Z = X/C$.

Intuitively, the idea of linkage tells us to what degree the observable g splits the classes of X/F, i.e., to what degree g can distinguish between states of X which appear the same under f. Thus, the more unlinked g is to f at $[x]_f$, the greater the extent to which g can "split" the class $[x]_f$.

Now we can see that the answer to the above control question is quite clear. The information provided by the observable f is sufficient for feedback control if and only if the observable g is totally linked to f at each F-class $[x]_f$. Basically what this means is that g is a *function* of f in the precise sense in which the term "function" is used in mathematics. The problem of determining when f and g are totally linked can abstractly be settled by the following result.

THEOREM.   *g is totally linked to f if and only if F refines C.*

Here we use the notion that *F refines C* if $(x_1, x_2) \in F \Rightarrow (x_1, x_2) \in C$. Thus, *g* is totally linked to *f* if and only if any two states which are equivalent under *F* are also equivalent under *C*. When stated in these terms it becomes evident that total linkage of *g* to *f* is, in general, not to be expected, which leads to the question of whether or not it is possible to devise some monitoring scheme which will provide suitable information for proper control.

Within the above framework, there exist only three possibilities for generating the proper linkage between observables for feedback control: (1) change the control relation *C*, which induces a change of the observable *g*; (2) change the state set *X*; (3) change the monitoring scheme *f* by introducing new observables $\{f_i\}$ and/or changing the nature of the single observable *f* (e.g., by new instrumentation to measure more accurately). Which of these approaches should be followed in any particular forestry situation depends upon many factors surrounding the individual forest, the measurement systems available, the desires of the decisionmakers, and so on and forms the basis for an important research program in forestry modeling and management.

2. *Bifurcations and Surprises*

The above abstract formalism also enables us to investigate the general question of bifurcation, i.e., when one description (model) of a forest essentially departs (bifurcates) from another. For the sake of exposition, let us consider the case when we have two descriptions $f: X \to R$, $g: X \to R$ of a given forest (here we are *defining* a description to be an observable). Let the associated equivalence relations on the set *X* be denoted by R_f and R_g, respectively. To speak of the *f*-description of the forest bifurcating from the *g*-description intuitively means that two elements $[x]_f$, $[x']_f$ which are "close" in X/R_f are not close in X/R_g. To formalize this notion we make

DEFINITION A2.   If $[x]_f$ is such that given $\varepsilon > 0$, there is a $\delta > 0$ such that

$$\left| [x]_f - [x']_f \right| < \delta \quad \Rightarrow \quad \left| [x]_g - [x']_g \right| < \varepsilon,$$

then we call $[x]_f$ a *stable* point of X/R_f with respect to X/R_g. The complement of the set of stable points is called the *bifurcation set* of X/R_f relative to X/R_g.

It is an easy exercise to verify that the stable points of X/R_f constitute an open set, while the bifurcation set is closed. Also, note that the concept of bifurcation relates to a *pair* of descriptions of the system and can be changed by changing either description.

As illustration of the above ideas, consider the case of describing the root structure of the set of all cubic equations $x^3 + ax + b = 0$. Here we can take

$$X = \{x^3 + ax + b = 0 : a, b \text{ real numbers}\}.$$

For one description we take

$$f \colon X \to R \times R,$$

$$x^3 + ax + b \mapsto (a, b),$$

while for a second description we can take the number of real roots, i.e.,

$$g \colon X \to \{1, 2, 3\},$$

$$x^3 + ax + b \mapsto \{\text{number of real roots}\}.$$

It is clear that the bifurcation points of the f-description from the g-description consists of the set $B = \{(a, b) : 4a^3 + 27b^2 = 0\}$, since it is precisely on this cusp curve in R^2 that the root structure of X changes character. In other words, in the euclidean topology on R^2, two equations which are "close" in the f-description will also be "close" in the discrete topology on the set $\{1,2,3\}$ except on the set B. This example will be recognized as the famous *cusp* catastrophe of the Thom theory and shows how our bifurcation definition subsumes the usual catastrophe-theory setup.

The foregoing development provides a basis for a discussion of the question of when one model can be said to "improve" upon another. The point of using an alternate description of a system is to gain new information. Thus, a criterion for the equivalence of f and g is that the bifurcation sets in X/R_f, X/R_g induced by the pairs (g, f) and (f, g), respectively, shall be empty. In general, given the descriptions f and g, the stable points of X/R_f and X/R_g represent elements of X for which the two descriptions coincide (i.e., yield the same information). On the other hand, on the bifurcation points the two descriptions differ and we gain information about the corresponding elements of X by employing both descriptions.

To summarize the above discussion, we have the following

IMPROVEMENT THEOREM. *The description f is an improvement upon the description g if and only if every point of X/R_f is stable relative to the topology imposed upon X/R_f by g, while the bifurcation set in X/R_g arising from the topology imposed by f is not empty.*

The concept of a bifurcation also enables us to speak more precisely about the occurrence of "surprises" in a system's behavior. Intuitively, a surprise occurs whenever reality departs significantly from our concepts of reality. In the above context, we may regard our concept of reality as being specified by some model f, while reality is characterized by another description g. By the above development, a departure of f from g occurs on the bifurcation set of points where states are close in the g-description but far apart in the f-description. As soon as we specify how far apart is far enough to constitute a surprise, we have a reasonably clear-cut approach to the general issue of surprises. It is of special interest to note that this is a nonprobabilistic attack on the question of surprise and does not rely upon any statistical analysis of past behavior in order to be effectively employed.

3. *General Systems Research for Forestry Models*

Within the above formalism, a considerable number of questions involving forestry management immediately suggest themselves:

(1) What is the basic state space and what are the observables for the principal forest models (e.g., Models I and II, Dixon and Howitt's model, etc.)?

(2) What are the bifurcation sets associated with the pairwise comparison of these models?

(3) How can we meaningfully combine various models to obtain improvements upon each model taken separately?

(4) What is the intrinsic information content in any forest model?

(5) Are there "surprises" inherent in the models which the manager should be alert for?

(6) If potential surprises are present, can we devise management schemes to avoid them (or perhaps induce them to happen)?

These are clearly not easy questions and require a substantial research effort to get a handle on. However, even partial answers would provide a much more solid understanding of the foundations of forest modeling, which in turn can be expected to lead to simpler, more effective management policies.

REFERENCES

1   K. Åström, *Introduction to Stochastic Control Theory*, Academic, New York, 1970.
2   S. Barnett and C. Storey, *Matrix Methods in Stability Theory*, Nelson, London, 1970.
3   R. Bellman, *Adaptive Control Processes*, Princeton U.P., Princeton, 1961.

4 R. Bucy and P. Joseph, *Filtering for Stochastic Processes with Applications to Guidance*, Interscience, New York, 1968.

5 M. Canon, J. Cullum, and C. Polak, *Optimal Control and Mathematical Programming*, McGraw-Hill, New York, 1969.

6 J. Casti, *Dynamical Systems and their Application: Linear Theory*, Academic Press, New York, 1977.

7 J. Casti, *Connectivity, Complexity and Catastrophe in Large-Scale Systems*, Wiley, New York, 1979.

8 J. Casti, Recent developments and future perspectives in nonlinear system theory, *SIAM Rev.*, 24(1982).

9 R. Crane and A. Stubberud, Closed-loop formulations of optimal control problems for minimum sensitivity, in *Control and Dynamic Systems* (C. Leondes, Ed.), Academic, New York, 1973.

10 G. Dantzig, *Linear Programming and Extensions*, Princeton U.P., Princeton, 1963.

11 B. Dixon and R. Howitt, *Uncertainty and the Intertemporal Forest Management of Natural Resources: An Empirical Application to the Stanislaus National Forest*, Giannini Foundation Monograph No. 38, Univ. of California, Berkeley, 1979.

12 G. Hadley, *Linear Programming*, Addison-Wesley, Reading, Mass., 1962.

13 K. Johnson and H. Scheurman, Techniques for prescribing optimal timber harvest and investment under different objectives—discussion and synthesis, *Forest Sci.*, Monograph 18, 1977.

14 R. Kalman, P. Falb, and M. Arbib, *Topics in Mathematical System Theory*, McGraw-Hill, New York, 1968.

15 R. Larson and J. Casti, *Principles of Dynamic Programming*, Vol. 1, Marcel Dekker, New York, 1977.

16 R. Larson and J. Casti, *Principles of Dynamic Programming*, Vol. 2, Marcel Dekker, New York, 1982.