RAMBLING RALF LOOKS AT BUDWORM OPTIMIZATION

Carl J. Walters

August 1975 WP-75-95

RAMBLING RALF LOOKS AT BUDWORM OPTIMIZATION

Carl J. Walters

When work on the budworm started at IIASA last year, considerable emphasis was placed on the sad fact that myopic management has led to explosive outbreak conditions. It seems that the government has been spraying hell out of each outbreak area as it appears, with little thought for large scale spatial consequences. In a great leap sideways, the IIASA groups managed to formalize this myopic viewpoint with site optimization by dynamic programming; it should come as no surprise that

(1) the formal myopic solutions closely resemble actual practice (managers are not <u>that</u> stupid), and

(2) for New Brunswick as a whole, the myopic optimization still gives poor results (trotting blindly toward a brick wall is not very different from running at it full speed). There have been some attempts to temper the optimization by inserting different objective functions and multiarea constraints on total cutting and spraying, but the results are not very encouraging. Some totally different approach to the optimization is clearly needed, if the study is to avoid going all the way back to brute force searching methods involving gaming simulations.

My interest in this note is to suggest a possible optimization procedure based on the notion that any optimal solution will involve some closed loop "control law" to specify actions as a function of system state (or output observations extracted from the state). If the form or shape or equation of this control law is specified in advance, then only the control law parameters (rather than all possible input system states) need be varied in searching for optimal solutions. To make this notion clear, consider our salmon optimization studies. When we use dynamic programming for the salmon, we try to find optimal harvest rates for many possible stock sizes (states):

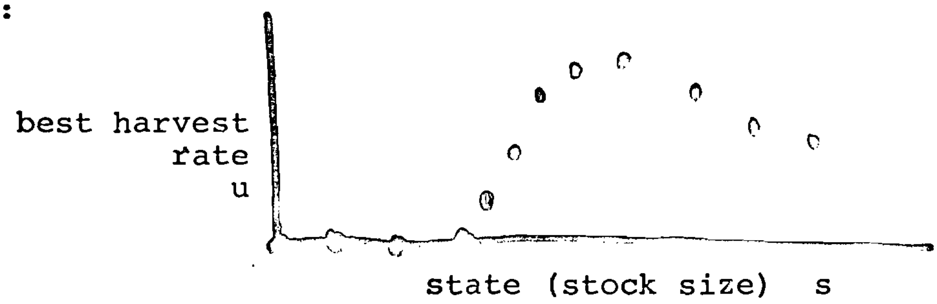

A control law is then found by interpolating between the test states. The alternative would be to assume some functional form for the control law, e.g.

$$u = a + bs + cs^2 + ds^3$$

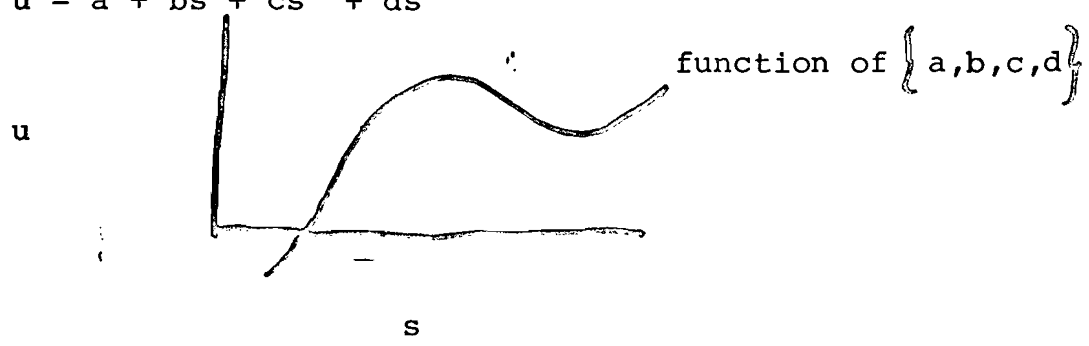

Then we would use some nonlinear programming or response surface search method to find best values for $\{a, b, c, d\}$. That is, in the simplest procedure we would:

(1) run several simulations with a particular combination (a_o, b_o, c_o, d_o), using different random input sequences, and compute the average total multi-year utility or value obtained from the system.

(2) Then vary one parameter to get (a', b_o, c_o, d_o) and repeat step (1); doing this for each parameter gives a gradient

$$\frac{\partial \text{ (total utility)}}{\partial \text{ (a, b, c, and d)}}$$

(3) use the gradient to jump to a new starting point (a_1, b_1, c_1, d_1) and repeat steps (1) and (2).

If the control law has few parameters, we might even do a systematic or grid examination of all possible parameter combinations.

The really critical trick in this alternative optimization approach is to find a reasonable functional form for the control law. The problem is exactly analogous to the modelling problem of how best to represent a complex system in terms of reasonably simple functional relationships. We must identify the following transformations and simplifications:

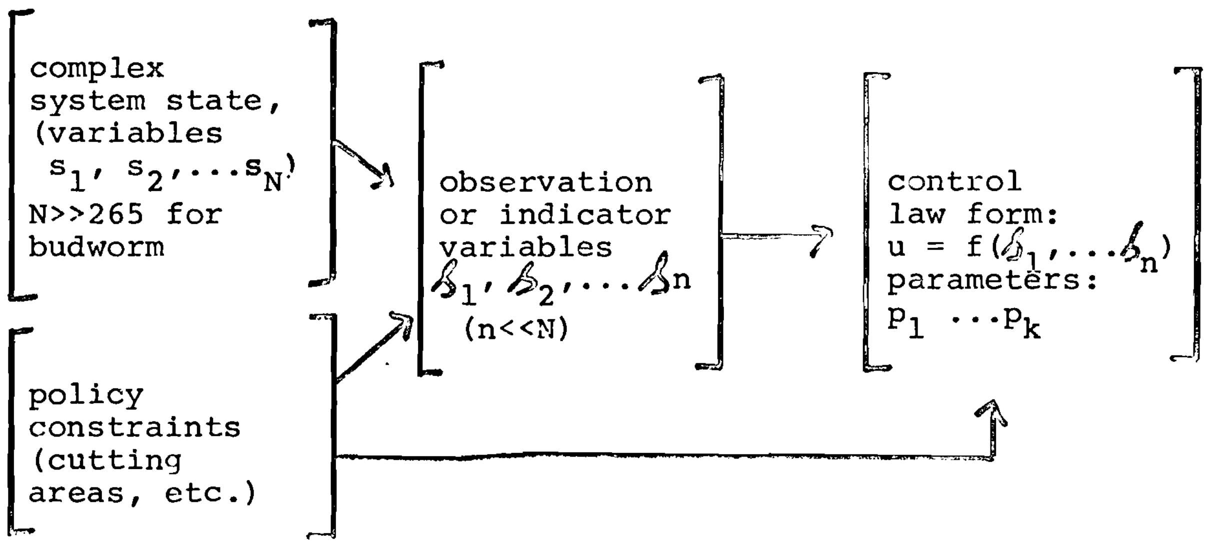

In the remainder of this note, I will try to suggest one possible form for the control law. In arriving at this form, I have tried to take into account the hierarchical nature of the decision problem as well as the need to base control actions on simple indicators and measurements. I also assume that the optimization will be based on some objective function that places a premium on transferring temporal variability into spatial variability as quickly as possible (no attempt is made here to develop such objective or utility functions -- Fiering and Clark are well on the way to that goal).

Making lots of little pieces into a few big pieces

As a first step in developing a simplified control law, I think it is essential to recognize that decisions must take place on at least three levels:

I. Basic Resource Allocation: a variety of private and public decisions combine to set basic limits on:

$$\begin{bmatrix} C_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} \text{forest area cut per year} \\ \text{forest area sprayed per year} \end{bmatrix}$$

These limits are not fully controllable by public policy decisions, and any major change in either of them may lead to political and economic contraints on further actions (options foreclosure idea).

II. Spatial Allocation: within the limits set by C_t and δ_t, the budworm optimization problem becomes more well defined: (1) how should the 265 sites be ranked so as to allocate C_t and δ_t most effectively? and

(2) should the full resources available even be used? A series of public versus private issues also arise at this level: if it appears that the best cutting sites are not economically optimal (location problems, etc.), should government subsidy or direct control be used to force the redistribution?

III. Implementation tactics: Given a spatial allocation, how should forest industry resources be mobilized and scheduled, what pesticides should be used and how should they be applied, how can wildlife and fisheries damage be prevented, etc., etc.?

Let us assume now that the level I decisions have been fixed as a time stream of C_t, s_t values (to be realistic, our recommendations would not likely have much effect at level I in any case).

Gaming procedures could help identify such time streams:

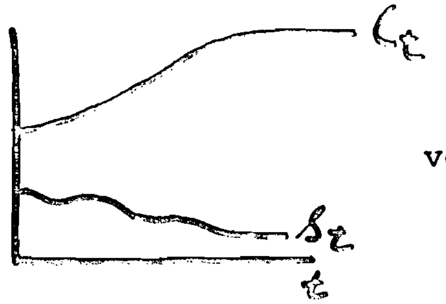

versus

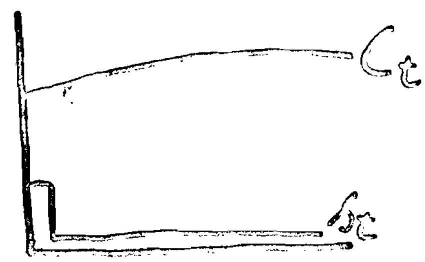

Note that the economic, political, and social issues involved in establishing (predicting, attempting to plan or implement, etc.) these trends are only in part dependent on budworm questions -- and it would be deceptive to pretend otherwise by somehow attempting to prescribe them in relation to some budworm objective function.

Next let us imagine some giant control law to specify whether or not to cut or spray each of the 265 sites as a function of the state (forest budworm) of that site and of all other sites. Though such a control law could never be computed, it would have some basic properties that we might be able to approximate with much simpler functions:

(1) it would implicitly be assigning ranks to each site such that only the top ranking sites would receive control actions in any year;

(2) though there might be separate implicit ranking systems for cutting versus for spraying, these ranking systems could be combined to give a single rank index R_j for each area j, where this index would order the areas such that the top-ranking ones would be cut (j_1, j_2,...j_{c_t}) and the next lower ones would be sprayed ($j_{c_{t+1}}$, $j_{c_{t+2}}$, ...$j_{c_t+s_t}$).

(3) The rank index for any site would be most sensitive to system state in the site and in a domain of adjacent areas near enough to provide or receive dispersing budworms.

If we are willing to believe that the full control law could be rewritten as or expressed in terms of such a ranking system, it seems reasonable to search for approximate control laws *in terms of approximations to the complex function that assigns the rankings R_j*. We could, for example, try an approximation that ignores all surrounding sites and ranks each site in terms of its profitability for logging. This approximate control law would then specify cutting for the c_t most profitable sites and spraying for the s_t next most profitable sites.

The simplest ranking function that would take some account of adjacent areas is the linear approximation:

$$R_j = c_1 (V_j - C_j) + c_2 (E_j + EIN_j)$$

where c_1 and c_2 are ranking coefficients

V_j = the gross logging value of site j

C_j = the logging cost for site j

E_j = the budworm egg density on site j

EIN_j = the expected number of eggs that will be dispersed into area j this year.

The coefficients c_1 and c_2 can be thought of as the unknown parameters of the closed loop control function. We can then by repeated simulations find the total utility (over time and space) to be expected from any assignment of c_1 and c_2:

FOR A UTILITY FUNCTION PLACING HIGH VALUE ON DIVERSITY:

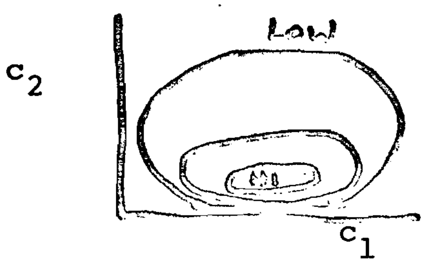

FOR A UTILITY FUNCTION PLACING HIGH VALUE ON PROFITS:

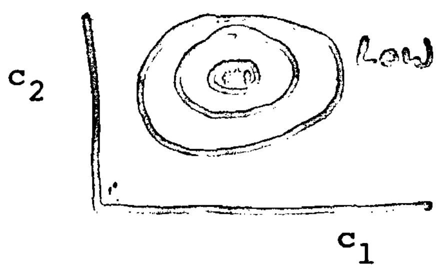

(If a high c_2 value is used, large areas that contribute eggs to one another will tend to be sprayed or cut together which might be good for profits but would tend to reduce diversity.)

An important point is that the best choice of a ranking function depends very much on the objective function -- in general the shape of any closed loop control function depends on what is being optimized. One way to get around this problem is to include all sorts of observations or indicators in the ranking function and hope that the coefficients associated with these will be made small by the optimization when appropriate (as the c_2 coefficient in the example above). We could, for example, extend the linear approximation for R_j to include terms like

$$R_j = \text{above} + c_3\bar{a}_d + d_4\bar{D}_d + c_5 PF_j$$

where

$\bar{a}_d$ = mean tree age in the domain of adjacent sites around area j

$\bar{D}_d$ = forest diversity in adjacent sites; $\bar{D} = \sum_k (a_k - \bar{a}_d)^2$

PF_j = proportion of fir in area j.

Also, the ranking function can be made nonlinear ($+c_6\bar{a}_d^2 + c_7E_j^2 +$ + etc...), so that the optimization may select some coefficients so as to "favor" particular indicators only when they are at relatively low or high values.

The ranking problem can be much simplified if it is assumed that locations for forest cutting are not within the domain of public control. Then one ranking model may be used to simulate private cutting decisions (presumably chosing sites j to maximize $V_j - C_j$), while simpler ranking functions might then be appropriate for prescription of spraying policies.

It might well turn out that there is no reasonably simple ranking function whose coefficients can be optimized to reasonably approximate the full control law, at least for some objective functions.

But the method at least allows enormous reductions in computational effort as compared to other optimization approaches, and it offers an opportunity for formal analysis of alternative indicator systems that would be essential for implementing any policies for allocating resources over many spatial sites.