

THE USE OF OPTIMISATION METHODS
FOR POLICY DESIGN IN A
SYSTEM DYNAMICS MODEL

R. G. Coyle - University of Bradford
R. K. Appiah - University of Zambia

August, 1982.

Introduction

In a previous paper (1), one of us presented a fairly traditional System Dynamics analysis of a production and raw materials system. It was shown that the system was remarkably unstable, in the sense that the Raw Material Sector of the system responded explosively to an exogenous sine wave simulating a bi-annual seasonality in external orders. It was also shown that bringing to bear the attitude of mind of a control engineer, allied to some simple rules of thumb stemming from that discipline, led to the identification of three alternative control structures, labelled Options I - III, each of which was capable of giving major improvements in the performance of that system. It was suggested that the system dynamics modeller could achieve quite significant improvements by that simple approach, without necessarily knowing very much about control engineering. Much of this ability to effect considerable improvement very easily may, however, derive from the fact that many managed systems are very ill-behaved in the first place, so that almost anything is bound to be an improvement. The real problem for the analyst is, however, to identify and evaluate design alternatives and it is with that problem that the bulk of this paper will deal.

It was readily conceded in that first paper that practitioners of System Dynamics, who seek to analyse the controllability of managed systems, could derive much benefit from a closer acquaintance with control theory. However, one must accept that, as the analyst of managed systems must also know a good deal about management, accounting, economics, etc. there is, in the real world, inadequate time available for a really deep study of control engineering. We therefore sought, in a second paper (2) to bring to bear our respective backgrounds in System Dynamics and

Control Engineering to analyse the ^{plant} source system from the point of view of the control engineer. We showed that the model could be expressed as analogue patch diagrams, and in the form of a control matrix and a plant matrix.

Armed with that formulation of the model we were to suggest a fourth control strategy, Option IV. It was shown that the control matrix formulation was a fruitful means of generating alternative configurations for the design of control policies.

It was demonstrated that System Dynamics models are formally reducible to the more conventional control engineering format, thereby making them more accessible for the control engineer to deploy his considerable analytical methods for the design of control strategies. In particular we proposed that Adaptive Model Following Control was likely to be a fruitful line of approach.

In that paper we were, however, careful to point out the important theoretical and practical differences between system dynamics and control engineering as well as drawing on their similarities, as it had been suggested in reference 1 that system dynamicists ought to do. In particular, we emphasised that the plant and the controller are not independent, as they usually are for the analysis of engineering control problems. The implication is that parts of the plant, and usually the major part of it, are freely changeable by the system dynamicist, a freedom not enjoyed, apart from a major redesign and then only to a limited extent, by the control engineer. Indeed, this interdependence of plant and controller is so pervasive that system dynamicists have habitually made little or no distinction between the two. It is, however, a remarkable coincidence that the distinction has re-emerged in recent work by Coyle and Wolstenholme, (3) in which the 'plant' is represented by flow and behavioural modules, and the 'controller' is

explicitly separated into information and control modules.

Equally, recent work by Coyle and King (4) has related the control rules and the changeable parts of the plant to the concepts of policy and strategy used in the Business Policy literature. If these three lines of research from control engineering, the reformulation of System Dynamics, and the business policy area, can be successfully developed and integrated, then some very stimulating developments towards a more fully satisfactory theory of the behaviour of managed systems become possible.

The present paper carries these converging strands of work forward to address the problem of parameter selection for a given control structure and the comparison of the resultant performance of two competing policy designs.

Control Matrixes for Policy Options

In terms of the notation introduced in the previous paper (2) the original model processed a control matrix of the form

$$G_0 = \begin{bmatrix} 0 & 0 & 1/TABL & 0 & \boxed{1 & 0} & 0 \\ 0 & 0 & 0 & -1/TARMS & \boxed{0 & 1} & 0 \end{bmatrix} \quad (1)$$

The sub-matrix in the small box is better written as

$$\begin{array}{cc} \alpha X_0 & 1 - X_0 \\ (1 - \beta_0) & \beta_0 \end{array}$$

with $\alpha_0 = \beta_0 = 1$

Option III from reference 1 has a control matrix

$$G_3 = \begin{bmatrix} 0 & 0 & 1/TABL - 1/TARMS & \alpha_3 & (1-\alpha_3) & 0 \\ 0 & 0 & 0 & -1/TARMS & (1-\beta_3) & \beta_3 & 0 \end{bmatrix} \quad (2)$$

with $\alpha_3 = \beta_3 = 1$

and Option IV in reference 2 has the matrix

$$G_4 = \begin{bmatrix} 0 & 0 & 1/TABL - 1/TARMS & \alpha_4 & (1-\alpha_4) & 0 \\ 0 & 0 & 1/TABL - 1/TARMS & (1-\beta_4) & \beta_4 & 0 \end{bmatrix} \quad (3)$$

in which $\alpha_4, \beta_4 \neq 0$

All three options (and we do not consider Options I and II from reference 1, for the sake of brevity) also have associated with them 5 other parameters which we may group as follows:-

- a) K and WAP which are respectively the number of weeks of Average Orders which are regarded as forming an acceptable order backlog, and the number of weeks of average production which are to form the desired raw material stock. To a control engineer these are gain elements and reduction of WAP was the basis of Option II, which did indeed lead to increased damping in the system.
- b) TAOR and TAPR which are respectively the average times for Order Rate and Production Rate before the operation of K and WAP. Control engineering practice suggested, in reference 1, that increasing TAPR would stabilise the system which proved to be the case.
- c) DDEL - the delay in receiving new supplies of raw materials.

It is natural to regard all these as parameters of the plant and that was the viewpoint adopted in reference 2. However it is easy to see that although they play no direct part in the control equation which determine the Production Rate and the Raw Material Order Rate, all except DDEL are freely choosable. Even the latter could be changed, by moving to another supplier, and that would be precisely analogous to the control engineer's drastic option of a major redesign of, say, the aircraft.

In terms of the business policy concepts referred to earlier, K and TAOR specify the company's strategy for dealing with its customers by defining the Order Backlog which would be regarded as acceptable in a given set of circumstances ie. for a given Order Rate. Similarly WAP and TAPR reflect the strategy adopted for maintaining internal stacks of raw materials. Those strategies serve as the governing factors on the streams of individual decisions on Production Rate and Raw Material Order Rate, and those decisions are provided by the regulatory, or strategy achievement, policies specified in whichever G matrix the company finds it expedient to use.

The design problem is then to choose a matrix G_0, G_3 or G_4 , and values of its parameter, and perhaps also to choose values for K, WAP, TAPR and TAOR, so as to maximise some measure of performance.

Performance Assessment in S.D. Models

The tradition in System Dynamics for assessing the effects of policy changes has been a visual comparison of plotted output, and, generally, little attention has been paid to numerical comparisons between simulations. Such a practice is not necessarily a bad thing;

many managed systems are so badly controlled in any case that it is relatively easy to make such a large improvement that a numerical Figure of Merit, or Performance Index, would be superfluous. In addition, experienced modellers are acutely aware of the dangers of oversimplifying the complexities of management by forcing them into a numerical index. Work has, however, been done on methods for the formulation of Performance Indices for System Dynamics models (5, 6, 7) and such indices are valuable in those cases, (and this is one such) where one is making comparisons of performance between relatively subtle alternatives.

Developing a Performance Index

In this case there are four factors to be weighed in the index.

- a) The need to prevent sharp variations in production from its previous average.
- b) The corresponding need to avoid variation between the Raw Material Order Rate and its average.
- c) The requirement that Order Backlog should not depart too far from its Desired Value.
- d) A corresponding wish that Raw Material Stocks should be matched to their desired value.

In each case, the penalty function can be expressed as the integral overtime of the squared deviation. The factors are manifestly not of equal importance, and they are assigned weights of 4, 1, 3, and 5 respectively. These reflect a judgement that it is far more importance to keep activity in the factory stable than it is too smooth the pattern of orders to outside suppliers, that it is quite important to keep Raw Material Stock close to target because of the

effect on corporate liquidity, but that it is most important of all to regulate the Order Backlog because of the effect on the customers. These relative weights are easily scaled to make the Index, PI, equal to 100 for the Base Case for the original model, sinusoidal of a regular input. The behaviour in that case was rather awful but only for the Raw Material Order Rate, which we have penalised very lightly.

Using the Index

The purpose of analysis is to suggest alternatives rather than to provide answers. Matrix analysis is a highly systematic way of generating alternatives and, if it can be coupled with an equally systematic way of numerically exploring that 'Alternative Space' ('Policy Space' is a better term), then the path to real system improvement should be open.

Such an approach may well be rather at variance with traditional System Dynamics practice but recent work by Keloharju (8) has provided very powerful tools for applying it. This has involved the complete integration into the DYSMAP simulation package (9) of a hill-climbing optimisation facility thereby rendering automatic a process which is otherwise exceedingly cumbersome, namely interfacing a simulation model with some very advanced optimisation facilities (10).

Armed with those tools, we may now return to the original system, that is before applying any control redesigning options at all. For that version of the system the matrix in G_0 , equation 1, and the original parameter values are $TABL=4$, $TARMS=4$, $\beta_0=1$.

We therefore attempt to minimise the performance index for matrix G_0 alone, that is leaving the 'plant' parameters at their initial values of $K=6$, $WAP=8$, $TAOR=4$, $TAPR=4$.

Initially we constrain $TABL$ and $TARMS$ such that $2 \leq TABL, TARMS \leq 10$ values which were judged to represent the maximum extent to which management would depart from established practice.

The result of the optimisation, which converges after about 70 iterations, is that the PI is reduced from 100 to 8.77 with the optimal values of the parameters determined to be $TARMS=TABL=10$, $\beta_0=0$ and $\beta=0.37$. This is interesting, in that both $TARMS$ and $TABL$ have been forced to their extremes, as has β_0 . The latter suggests that the Raw Material Order policy has in the past been fundamentally incorrect in using $\beta_0=1$. The rationale for that was that the Raw Material Manager took account of Average Production Rate on the plausible grounds that it was his job to see that raw material stock was available to be used up in production. He ignored Average Order Rate on a reasonable basis that orders were what the Production Manager dealt with. The optimiser, in driving β_0 to 0, and therefore setting $(1-\beta_0)$ to 1 is indicated that those apparently sensible attitudes are fundamentally misconceived.

Similarly, the optimiser casts doubt on the Production Manager's attitude that he ought to attempt to correct Backlog errors fairly quickly and keep up with Average Order, which he has implemented by putting $TABL=4$ and $\beta=1$. The optimisation instead indicates a much more relaxed attitude to order backlog, $TABL=10$ and far more attention paid to Average Production than to Average Order Rate. The last is not surprising, given that an index attaches a heavy weight to achieving

stable production patterns but the former is surprising, given that the matching of Order Backlog to its desired value is the most heavily weighted component of the Index.

The reason is the relatively heavy weight attached to maintaining raw material stocks at their desired level a requirement which seems as sensible, but may be as wrong, as the policies discussed above. The ability of the optimisation to cast doubt on its own objective function is a fruitful source of improvement in the difficult area of creating and ordering a good objective function. For the purposes of this paper we shall, however, stay with the one we have.

The real significance of the results for α and β is that the optimiser has taken us far beyond the depth of analysis achieved in reference 1, by the application of simple heuristics. That approach largely accepts the structure of the policy as fixed, as long as that structure seems reasonably sensible, and focusses the main attention on parameter values with relatively limited structural change. The more sweeping search done by the optimiser has produced a more profound result bearing on the nature of the policies to be applied.

When we allow the 'plant', or strategy parameters to enter the optimisation, containing K , WAP , $TAOR$ and $TAPR$ to lie between 2 and 10, a further substantial improvement in the index is achieved. The final value of the index decreases yet again to 1.02 with the optimal parameter values

TARMS=10 $\alpha_0=0.45$ $\alpha_0 = 0.38$
 K=10 $\beta_0=0.034$
 TABL=9.1
 TAPR=10
 WAP=2 $\text{TAOR} = 9.1$

All of the parameters in the first column have been driven to their extreme, except TABL which is near its extreme. In particular, the high value of TAPR and the low value of WAP confirm the control engineering rule of thumb, used in Option II in reference (1) that increasing gain and reducing delay are good ways to increase stability. This is, however, slightly contradicted by the fact that K, which is also a gain element, has risen from its initial value of 6 to its limit of 10. This is the number of weeks that customers will have to wait for their orders, and it has been driven to the value which management feel is the safe limit, bearing in mind that reliability of promised delivery is often more important than the magnitude of the wait. In fact, the high value of K, and the low value of WAP, correspond to the managerial heuristic of making the customers wait as long as one dares, and holding as little stock as possible.

The question to which we now turn is whether G_3 and G_4 can be made to perform equally well, or better.

The first optimisation in this section showed that when the policy matrix G_0 was optimised for a constant contrast strategy, the Performance Index improved from 100 to 8.73, and the second showed that when policy and strategy were optimised together there was a further decrease to 1.02. Finally, therefore, we optimise the strategies alone and find the minimum attainable index to be 12.0 when WAP=2, K=10, TAOR=8.44 and TAPR=10.0. This is very close to the high-backlog-low stock strategy but the main point of these results is that neither strategy nor structure is of dominant importance, and each has to be tuned to the requirements of the other.

Optimisation of Option III

The performance of Option III is already known to be better than the base run on the original model, so that the parameter values in G_3 , equation 2, are set by rule of thumb (or guesswork) to

TABL=TARMS=4 and $\alpha_3=\beta_3=1$

the value of the PI is 8.01

which is better than the minimum for G_0 , when these four parameters are optimised, with the same constraints, the PI reduces to 4.84 with TARMS=10, TABL=9.6, $\alpha_3=0.337$ and $\beta_3=0$. These are practically the same as the values for G_0 confirming the earlier indication that the Production and Raw Materials policies are fundamentally wrong.

For matrix G_3 , joint optimisation of strategy and policy reduces the PI still further to 0.985 with the parameter values TARMS=10, TABL=3.99, K=10, WAP=2, TAOR=8.77, TAPR=10.0, $\alpha_3=0.456$ and $\beta_3=0$.

This is further confirmation of the kind of strategy/policy balance appropriate for this system, but the change in TABL, which has moved very nearly back to its original value of 4 shows how very carefully one has to watch the strategy/policy balance.

In the previous section we also tested strategies alone but that is clearly rather unrealistic, in general, and we make no further attempt to do so.

Optimisation of Option IV

Finally, we examine matrix G_4 , equation 3. With initial values of TARMS=TABL=4 and $\alpha_4=\beta_4=1$ its PI is 15.5, certainly better than the original system, but by no means as good as Option III. If $\alpha_4=\beta_4=0.5$, the PI changes negligibly to 15.3 suggesting that TARMS and TABL are more significant parameters.

This is borne out by the optimisation, which leads, for G_4 optimised alone, to a PI of 6.29 when $TABL=TARMS=10$, $\alpha_4=0.32$ and

$\beta_4=0$.

When strategy and policy are jointly optimised, the PI decreases to 1.57 with all parameters driven to, ~~or close to~~, their extremes, ^{$\beta_4=0.41$} except $\alpha_4=0.34$ and $\alpha_4=0.59$. These last two values are interesting as they imply the Production and Raw Materials Managers really co-ordinating very closely and each giving nearly equal weight to Average Orders and Average Production in their respective decisions. This is rather different from the corresponding results for Option III and the original system in which the Raw Material Manager was told to ignore Average Production and pay attention instead to Average Orders, which was exactly opposite to what he had been doing in the original model, and the Production Manager was recommended to do almost the same thing, and change his behaviour.

The reason is the much closer co-ordination of their activities implied by the presence of TABL and TARMS in both columns 3 and 4 of matrix G_4 . This means that the Production Manager takes into account the Raw Material Stock position as well as the more obvious Order Backlog which was all he considered in the original system. Similarly, the Raw Materials Manager includes Order Backlog in his decision making, as well as the obvious factor of Raw Material Stock. Paradoxically this seems to be too tight a degree of co-ordination, as shown by the fact that the optimal performance of Option IV is worse than that of Option III, and is even somewhat worse than the optimal performance of the original model, as shown in Table 1.

Simulations of the behaviour of each of these cases appear in the following diagrams.

Conclusions from Optimisation Experiments

The simple conclusion to be drawn from Table 1 is that it is feasible to over-control a system, which is what is happening in Option IV. The additional feedback mechanisms create such a tight net of control that the system has too little freedom of manoeuvre.

The second, and perhaps more subtle, conclusion is that very major improvements can be brought about by very small changes, though those changes may run counter to established practice and common sense. The change from 100 to 8.73 in the original system is achieved largely by persuading the two managers to shift their emphasis. The Raw Material Manager should have his ordering not on production, but on customer orders, and the Production Manager should link his production decisions far less heavily to customer orders and far more so to previous production. Naturally, once these results have been obtained, it is easy to advance common-sense arguments for them, but it is far harder to reach them by such arguments in the first place.

The third conclusion is that, if policy and strategy are allowed to enter the analysis, even further improvements are possible, and that these improvements can converge so that different structures are able to give practically identical performances. That does not mean that structure is unimportant as, in this case, one gives slightly worse results than the others.

	Original System	Option III	Option IV
Initial Value of PI	100.00	8.01	15.5
Result from Optimisation of Policy alone	8.73	4.84	6.29
Result from joint optimisation of policy and strategy	1.02	0.98	1.57

Table 1. Summary of Optimisation Experiments:

Alternatives to Matrixes

The matrix approach to policy generation has worked well in this case in that, having generated and analysed a collection of options, we can now be quite confident that there is very little further improvement available for this system. It has been optimised in a far more fundamental sense than simply minimising the Performance Index for any one configuration. The process applied is more accurately called 'simulation through repeated optimisation' than 'optimisation through repeated simulation'.

The matrix approach does, however, have two disadvantages. The first is that it is quite simply not accessible to many otherwise well qualified System Dynamicists. For example, an accountant is well qualified to analyse financial problems, and he can learn the required simulation techniques quite rapidly, but it is rather unlikely that he would have the time or the inclination to learn such matrix algebra.

Secondly, matrixes have a quality of rigidity, which is probably psychological rather than mathematical, which makes it clumsy to write down all the possibilities in a compact form. More seriously, the matrix formulation separates too rigidly the attributes of policy and strategy which have been discussed in this paper.

An alternative is the Policy Option Diagram, the form of which was prepared by Keloharju, through the English name is due to Wolstenholme. The link to policy and strategy explained below is prepared for the first time in this paper.

The Policy Option Diagram, and one cannot resist using the acronym POD, is a tree, the apex being a flow rate in the model, and the branches being all factors which do, or could, form part of the determination of that flow rate. The POD for Production Rate appears in Fig. 1.

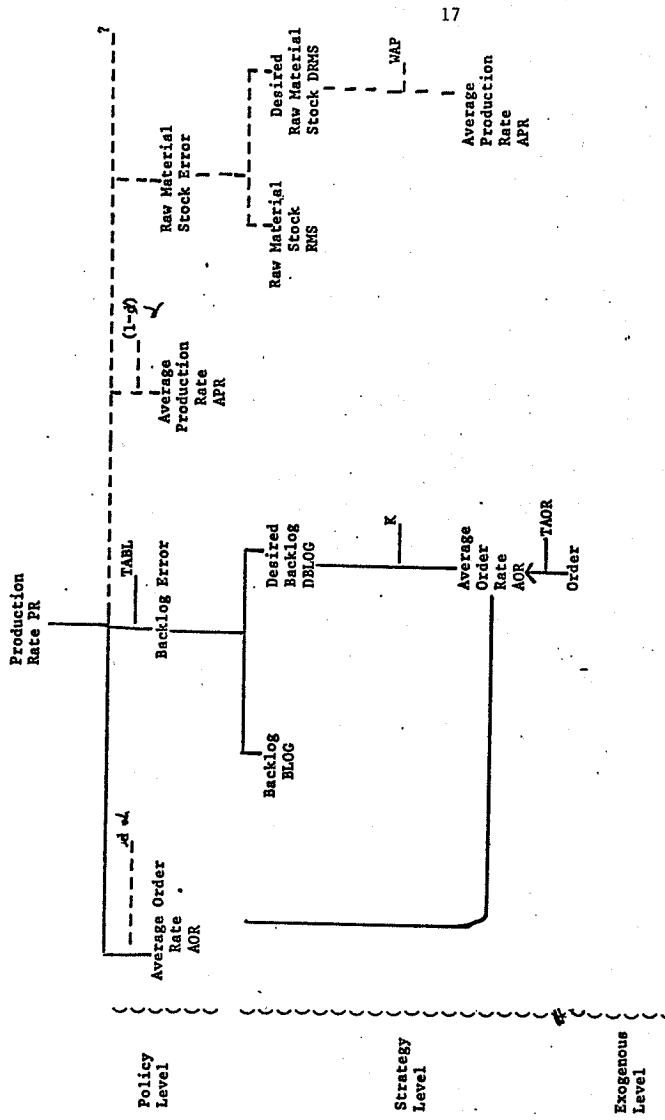


Fig. 1. Policy Option Diagram for Production Rate.

The solid lines at the left of the diagram show the factors taken into account in determining Production Rate in the original system, that is they ^{show} the policy option which management has hitherto chosen to exercise, with what notably had results was demonstrated in reference 1, and indicated by the Performance Index of 100 used in this paper.

The top tier in the POD is labelled the Policy Level because these are the factors used in the daily stream of decision by which the firms attempts to meet its targets for, in this case, Order Backlog. The lower tier is called the Strategy Level because it contains the governing factors which produce that decision stream. The parameters appear at the sides of the branches to which they relate.

The lowest level is named the Exogenous Level because it is here that driving forces enter the system.

The point about the POD is that it should make one think of what might be made to enter into a decision rather than merely concentrating on what does enter it. This is shown by the dotted lines which imply that Raw Material Stock Error might be used to control production (which is essentially Option III). For example, there is no reason why APR should not influence PR as well as, or instead of, AOR. To do so would necessitate a new parameter λ which creates a further dimension in the parameter space.

Any state variable in the model could be used to govern PR and that is implied by the ? at the extreme right hand edge. The only limit is the imagination of the analyst. There is, for example, no law that says that Desired Raw Material Stock should depend only, or at all, on Average Production Rate. The POD is thus at least as effective as control matrixes in generating alternatives, has the advantage of linking policy and strategy, and is a good deal more transparent to many people.

It has the major disadvantage that it does not connect readily to the analytical methods of control engineering, which perhaps reinforces the earlier comment about the pressing need for System Dynamicists to become more knowledgeable about control theory and some of its mathematics.

Conclusion

We make no attempt to generalise the results of this analysis to all production systems, though they do apply to a significant class of such systems. We argue, instead, that the diversity of production environments is so great and the needs of different firms at differing stages of their evolution is so varied, that any generalisation could be highly misleading.

Instead, it is argued that the simulation and optimisation procedures used in this analysis are so easy to learn and to use that it is more cost-effective to deal with each case as one encounters it. A critic could say that that was solving problems ad hoc, though it could equally well be called treating each case on its merits.

References

- Coyle, R.G. 'The Dynamics of a Production and Raw Materials System' IEEE Trans. on Cybernetics Man and Society, September, 1982.
- Appiah, R.K. and R.G. Coyle 'Policy Design in System Dynamics Models - a modern control engineering viewpoint' Submitted to IEEE Trans. on Cybernetics Man and Society.
- Coyle, R.G. and E.F. Wolstenholme *System Dynamics as a Rigorous Procedure for System Description*
J Opt Res Soc To appear
- Coyle, R.G. Equations for Systems
University of Bradford Printer, 1979.
- Coyle, R.G. Modelling the Future of Mining Groups
Trans. Inst. Min. Metal (Sect.A), 90, 1981, pp.A81-88.
- Montaldo, D.H. 'Policy Design for an Operating Mine'
Ph.D Thesis, 1981. unpublished. *University of Bradford*
- Keloharju, R. Relativity Dynamics
Helsinki School of Economics,
Working Paper F33, 1982.
- Cavana, R.Y. and R.G. Coyle DYSMAP Users Manual
University of Bradford Printer, 1982.
- Luostarinen, A. and R.G. Coyle DYSMOD Users Manual
University of Bradford Printer, 1982.
- * Coyle, R.G. and P.D.A. King Towards a Structuralist Theory of Business Policy Issues
Unpublished working paper.

I know this one is in the wrong place but since it has been returned by Admin. Sec, I have left it where it is for the time being.