

Improving Insight and Understanding by Optimising System Dynamics Models

Eric F. Wolstenholme and Sattar A. Al-Alusi
University of Bradford Management Centre
Emm Lane, Bradford, West Yorkshire, BD9 4JL
United Kingdom

ABSTRACT

This paper will outline the concept of system dynamics optimisation using the DYSMOD software and present a case study of its use to analyse a defence problem. The insights into the problem, which were generated from a conventional system dynamics model and its policy design experiments, will be given. This will be followed by the presentation of results from a set of optimisation experiments, utilising a range of objective functions and structural design parameters. The paper will focus on the value added to the understanding of the problem which resulted from this process. The overall conclusion is that optimisation subsumes conventional sensitivity analysis as well as providing an holistic interpretation of the behaviour of a system dynamics model.

INTRODUCTION

Traditionally, system dynamics has relied on the use of intuition and experience by system owners and analysts to help design policies for improving system behaviour over time. This situation is now changing and much effort is being expended in the development of policy design methods. Basically, two schools of thought are emerging. The first of these concerns the application of control theoretic methods and the second simulation by optimisation (Keloharju (1983)). This approach relies fundamentally on computer software and the software to be described and applied for optimisation in this paper was originally developed as an appendage to DYSMAP (Dynamic Simulation Model Application Programme) (Cavana and Coyle (1982)) and is known as DYSMOD (Dynamic Simulation Model Optimiser and Developer).

The DYSMOD software uses a hill climbing routine to heuristically determine the optimum values for any number of model parameters relative to predefined objective functions or performance measures using a system dynamics model as a starting point. Optimisation in parameter space is achieved by interleaving simulation and optimisation. One iteration of the procedure consists, firstly, of a DYSMAP simulation run, in which the value of the objective function is calculated and secondly in a run of the optimiser to choose parameter values which might improve the objective function. Subsequent iterations consist of rerunning DYSMAP to test out the resultant improvement in the objective function under the

new parameters and further refinement of them by optimisation. Any one experiment with the software might take a 100 or more iterations.

A DEFENCE MODEL

The model to be developed here using optimisation is a defence model (referred to as the armoured advance model (Wolstenholme 1987)). The armoured advance model was constructed initially to examine the effects of alternative formation change strategies by an attacking force (red), under a variety of fire delivery strategies on the part of the defending force (blue).

The red strategies considered were to change formation at a fixed distance of advance or to change formation at a variable distance of advance using a range of different variables on which to base this decision; such as its speed and force size (number of units advancing).

Basically, the purpose of red delaying its formation change point was to facilitate recovery of speed lost (due to blue's fire) by staying in a fast moving but dense battalion formation. Conversely, the purpose of red advancing its formation change point was to conserve its force size by changing to a slow moving but less dense company column formation.

A pseudo parameter SZS (speed/size switch) was defined in the model to allow any combination of weightings to be attached to the speed and size variables in the formation change strategy equation. In the base model this parameter was set to 0.5. When speed was used as a determinant of the formation change point, a speed multiplier function was invoked, which progressively increased the planned distance to the formation change point. When force size was used, a size multiplier was invoked, which progressively decreased the planned distance to the formation change point.

The blue fire strategies were based on criteria involving red's distance of advance, red's speed of advance and red's momentum (speed * numbers). Blue's fire was assumed to reduce both the number of red's units and their speed of advance. However, speed was assumed to be recovered by red when blue fire ceased. Both the speed and attrition effects were modelled as the product of the rate of blue fire (shells/min) and the productivity of fire (attrition rate/shell or speed reduction rate/shell). The latter was assumed to depend on the density of the target [that is, the formation in which red was advancing] and the accuracy of fire, which was made to increase as the distance between the contestants reduced.

The strategy of blue delivering fire on the basis of red's distance of advance, was achieved by defining the percentage of the distance of advance over which fire would take place. This was initially chosen as the

first and last 10% of the distance advanced by red in each formation. The strategy of blue delivering fire on the basis of red's speed and momentum, was achieved by defining upper and lower limits of those variables; fire being delivered whenever the upper limit (or trigger value) was achieved by red and switched fire off when the value of the variable had been driven down to the lower limit. The blue strategies were replicated for two different firing rates (light and heavy).

The results from this model suggested that from red's point of view it was better to deploy a variable formation change strategy and to delay the formation change point. This gave a shorter arrival time at the blue position whilst not sacrificing numbers too much and hence a generally improved arrival momentum (numbers arriving/arrival time). This suggests the emphasis by red should be on staying as long as possible in battalion formation. The rationale for this was thought to be the higher speed achievable in this formation and because speed unlike size is recoverable if blue firing stops.

From blue's point of view the results suggested that it was best to first deliver fire on a criterion of red momentum, secondly to use a criterion of red speed and thirdly to use a criterion of red distance. Here, the aim was to prolong the red advance and reduce as much as possible the numbers arriving and their arrival momentum.

It should be appreciated that the initial results from the model depended on the assumptions made about the strategies used. For example, red's formation change point depended on the value of SZS chosen or on the shape of the multiplier functions. Likewise, blue's fire delivery strategies depended on the percentage of the distance of red's advance over which firing took place [and, indeed, the location of that percentage] and the parameters chosen for the upper and lower limits of speed and momentum at which fire was switched off and on.

The conventional approach to this problem in system dynamics is to perform sensitivity analysis. However, this is only possible to achieve in a limited way for extreme parameter values due to the enormous number of permutations of parameters involved and hence the enormous number of computer runs of the model necessary. Optimisation can help here by seeking out that permutation of parameters which gives the best outcome according to given objective functions.

Optimisation of the Defence Model

The approach in applying optimisation to the armoured advance model was to firstly define relevant objective functions and then to define for each objective function the parameters whose values would be chosen by the optimisation procedure; and the upper and lower limits of the feasible

ranges for these values. In general the guideline followed in defining parameters was that only red's strategy parameters could be involved in experiments using red's objective functions and only blue's strategy parameters could be involved in experiments using blue's objective functions.

The Optimisation of Red Strategies

From red's point of view the main objectives are to maximise it's size and momentum on arrival at the blue position or to minimise it's total advance time. For all three of these objective functions, the relevant parameters to be chosen are those involving the red formation change decisions. These are whether red's speed or size is used as a basis for modifying the distance to company column deployment, and within this, the shape of the speed or size multiplier table used. Freedom was given to the optimiser to choose the value of the speed/size parameter and each of the y-coordinates for the speed and size multiplier functions.

Results from Experiments with Red's Objective Functions

When red's arrival size was maximised it was found that in every case that the optimised arrival size significantly exceeded the arrival size attained in the base model. This was achieved by the optimiser employing size as the sole determinant of the formation change point and adjusting the shape of the size multiplier to produce a very aggressive response in the planned distance to the formation change point, as soon as the actual red size fell below the planned red size. The consequence was that the whole of the red advance was carried out in company columns. In all cases the focus on a long advance in company columns meant that red's total advance time was much slower. Additionally, blue used much ammunition and the efficiency of use was correspondingly low.

This result, of carrying out the whole advance in company columns, represents the maximum extent to which red can protect itself against losses incurred in force size. It might be expected that savings in size resulting from this extreme strategy would at least compensate for the longer advance time and hence not seriously affect the momentum of red's arrival at blue's position. However, this was not the case and the deterioration in momentum can best be explained in terms of the additional exposure time of red to blue's fire, which results from the strategy.

When red's arrival time was minimised it was found in every case that the optimised arrival time was significantly less than that recorded from the base model. This improvement was achieved in the optimisation process by employing speed as the sole determinant of the formation change point and by again adjusting the shape of the speed multiplier to give a much

more aggressive response in the planned distance to the formation change point, as soon as red's actual speed fell below the planned speed. The consequence was that the whole of the red advance was now carried out in battalion formation.

Suprisingly, the arrival time did not suffer too greatly and hence in most cases the arrival momentum achieved by red was as good as, or better, than that from the base model. Blue's use of ammunition was also low and its efficiency of use high. It would, therefore, appear that the gain in speed achieved by red by staying in battalion formation more than compensates for the higher attrition. This result is again perhaps best explained in terms of the low exposure time to blue's fire arising from red maintaining a higher speed formation.

When red's arrival momentum was maximised it was found that in every case the optimised value of red's arrival momentum was significantly better than that recorded in the results from the base model. This result was achieved by red predominantly choosing speed as the criterion for formation change with the shape of both the speed and size multiplier functions chosen for aggressive responses. The results tended much more towards those obtained in minimising red's arrival time rather than those obtained from maximising red's arrival size. Whilst the whole advance was not now in battalion formation the time to company column deployment was extended with little effect on the arrival size.

These results are described in more detail elsewhere (Wolstenholme and Al-Alusi, 1988) and provide substantial confirmation of the previous conclusions from the non-optimised model, that red's best strategy is to stay in battalion formation for as long as possible. The investigation of what might be intuitively considered as trivial runs from two extreme situations (maximising red's arrival size and minimising red's arrival time), was shown to facilitate the formation of a perspective concerning the trade off between the two, which is clearly confirmed in the run involving maximisation of red's arrival momentum. The optimisation approach added to the analysis by clearly highlighting a further effect by which blue's strategies interact with reds. Originally the red desire to stay in battalion formation was seen primarily as being based on increasing it's ability to recover speed, and hence momentum, when blue's firing ceased. Additionally, however, as shown here, it is clear that this strategy further minimises the time over which red is exposed to blue's fire.

The Optimisation of Blue Strategies

From blue's point of view it's main objectives are the opposite to those defined for red. These are, to minimise red's arrival size, maximise red's arrival time and to minimise red's arrival momentum. Additionally, it has

the objective of trying to achieve these objectives, particularly that of minimising red's momentum with the minimum use of ammunition. A further objective function is therefore relevant in blue's case. This involves maximising the average reduction in red momentum per shell fired. For all these objective functions the relevant parameters to be chosen are those involving blue fire delivery. That is (i) the proportion of the red advance over which blue delivers fire (where the latter is carried out on a distance criteria); (ii) the upper and lower limits of red speed at which blue fire is switched on and off (where blue fire is delivered on a speed criterion), and (iii) the upper and lower limits of red momentum at which blue fire is switched on and off (where blue fire is delivered on a momentum criterion). Optimisation experiments were carried out for each of these objective functions and at two levels of fire delivery (light and heavy).

Results from Experiments with Blue's Objective Functions

It was found that the results for the first three of blue's objective functions were identical. It would appear that if red's flexibility concerning its choice of formation change point is removed, then the same effect is achievable by blue via any of these objectives.

Under the blue strategy of delivering fire on a distance criterion, optimisation of red's arrival size, time and momentum was achieved by firing during the whole of the red advance; that is by deploying continuous fire. Under the blue strategy of delivering fire on a speed or momentum criteria, optimisation of red's arrival size, time and arrival momentum was achieved by switching fire on and off at the lowest points within the range defined for red's speed or momentum.

Overall blue achieved the longest time for red's advance and the lowest values of red's arrival speed and momentum when delivering fire on a distance criteria. This was because such a strategy not only allowed continuous fire but also allowed continuity of fire delivery. When fire was delivered on a speed or momentum criteria there were periods when fire was switched off, which facilitated red speed recovery. This result exposes an inherent weakness of these strategies.

The foregoing results generated a very high usage of ammunition. When red's arrival momentum per shell was maximised the ammunition usage was much less. This improvement was achieved by blue only choosing to fire whilst red was in battalion formation. This is an interesting insight which results because of the density of this formation. The productivity of the blue fire directed at red in battalion formation is higher than that directed at red in company column formation. However, because of the selectivity of blue fire, red's arrival size, time and momentum were all improved which is, of course, detrimental to blue. Nevertheless, the

savings in ammunition were phenomenal, which might more than compensate for this deterioration if ammunition was limited.

In addition, when red chose to delay its formation change point the optimisation procedure led to a further intriguing insight. Here, blue chose only to fire on red at the most productive point during the red advance. This was chosen by the optimiser to be at the very end of red's battalion formation advance and this choice is explainable since it is when red is not only in its most dense formation, but also at its closest point to blue (hence blue's accuracy of fire is higher). This result provides a good example of how optimisation facilitates an holistic appreciation and interpretation of results.

Yet another interesting insight, arising when blue delivered fire on a speed or momentum criterion, was that the greatest reduction in momentum per shell was achieved by blue deploying a light rate of fire rather than a heavy rate. This is because not only does light fire save on ammunition, but also it again results in more consistent fire. Heavy fire, by definition of these fire delivery strategies, is more intermittent than light. It quickly drives both red's speed and momentum down to the point at which fire is switched off (thus allowing these attributes to recover) rather than reducing them more gently but for longer periods of time.

CONCLUSIONS

It is concluded that a number of significant additional insights into both red and blue strategies were achieved by optimisation of the armoured advance model. These were:

- red should maintain a high density, high speed formation for as long as possible since this retains the flexibility of speed recovery and minimises the total advance time and hence the period of exposure to blue fire.
- it is important for blue to maintain continuity of fire. This is facilitated by blue firing on a distance rather than a speed or momentum criteria, since the latter results in periods of zero fire when the trigger points, defined in these fire delivery strategies, come into play. This in turn allows red to recover speed which shortens the advance time.
- when ammunition is limited it is more productive for blue to restrict fire to periods when red is advancing in battalion formation and to apply light but consistent fire. For optimum productivity from limited ammunition blue should focus fire over the later stages of the red advance in battalion formation.

the results from the situation where blue is allowed to manipulate red's formation change point, confirm the previous results from the experiments

on red's objective functions. That is, red's speed is a more important variable than red's size to both sides. It is best for red always to extend its battalion formation advance as long as possible. This increases speed and, by reducing the exposure time to blue fire, reduces attrition.

- when ammunition is limited it is to blue's advantage if red does prolong it's advance in battalion formation, since blue can fire for longer in its more productive mode.

The original insights into the armoured advance model were obtained from 12 DYSMAP computer runs (3 blue strategies at two levels of firing against 2 red strategies). The optimisation experiments involved a total of 8400 DYSMAP computer runs (the red experiments repeated the 12 runs for 3 objective functions of 100 iterations each and the blue experiments covered 4 objective functions of 100 iterations each). To achieve the same level of analysis using conventional sensitivity analysis would have required a repeat of the 12 original runs for all permutations of the values of the parameters used in the optimisation runs. This is conservatively estimated at over 10 times the number of optimiser runs. Further, given the degree of automation of the optimisation process, the time saving over conventional sensitivity analysis would be enormous.

REFERENCES

- Cavana, R.Y. and R. G. Coyle (1982) DYSMAP User Manual, University of Bradford.
- Keloharju, R. (1983). Relativity Dynamics. Helsinki School of Economics, Helsinki, Finland.
- Wolstenholme, E. F. (1988). Defence Operational Analysis using System Dynamics, European Journal of Operational Research, vol. 34 no. 1, Feb. 1988, pp. 16-18.
- Wolstenholme, E.F. and A. S. Al-Alusi (1988) System Dynamics and Heuristic Optimisation in Defence Analysis, System Dynamics Review, vol. 3, no. 2, Winter 1987, pp. 102-116.

Acknowledgements

The authors wish to acknowledge that some of the work described has been supported by the Procurement Executive, Ministry of Defence. However, any views expressed are those of the authors and do not necessarily represent those of either this Department or H.M. Government.