

TOWARDS THE HOLY GRAIL: COMBINING SYSTEM DYNAMICS AND DISCRETE-EVENT SIMULATION IN HEALTHCARE

Sally C. Brailsford
Shivam M. Desai
Joe Viana

School of Management
University of Southampton
Southampton SO17 1BJ, UNITED KINGDOM

ABSTRACT

The idea of combining discrete-event simulation and system dynamics has been a topic of debate in the operations research community for over a decade. Many authors have considered the potential benefits of such an approach from a methodological or practical standpoint. However, despite numerous examples of models with both discrete and continuous parameters in the computer science and engineering literature, nobody in the OR field has yet succeeded in developing a genuinely hybrid approach which truly integrates the philosophical approach and technical merits of both DES and SD in a single model. In this paper we consider some of the reasons for this and describe two practical healthcare examples of combined DES/SD models, which nevertheless fall short of the “holy grail” which has been so widely discussed in the literature over the past decade.

1 INTRODUCTION

The idea of combining discrete-event simulation and system dynamics has been a topic of debate for over a decade, prompted by a recognition that the division between strategic, tactical and operational decision-making was becoming increasingly blurred. During the 1990’s there was a growing realization that because of the speed of communications and the complexity of global organizations, decisions made at one end of the operations-strategy spectrum could rapidly impact on outcomes at the other end. For example, the operational decision for a business to reduce staffing levels in its customer services call centre could impact upon the quality of service perceived by customers, which could ultimately result in lower sales, a decline in the share price and the need for a change in market positioning. It is often difficult to draw clearly-defined boundaries round any part of a large system or organization and say that the resulting subsystem can be studied in isolation. This is particularly (although not uniquely) true in healthcare organizations, where “everything affects everything else”.

Operations Research (OR) has historically been focused at the operational end of the spectrum, but during the 1990’s new OR techniques such as Strategic Options Development and Analysis (Eden, 1989) were developed and successfully used for strategic decision-making (Dyson and O’Brien, 1998). Focusing in on simulation, these two aspects are exemplified by discrete event simulation (DES) and system dynamics (SD). DES is essentially an operational tool, designed for optimization of system performance at a very detailed level. Typically, DES is used for modeling queuing systems where stochastic variability is important. On the other hand SD is a more strategic tool, used at a much higher level for understanding overall system behavior.

Over the past decade there has been considerable interest in which approach should be used and when (Brailsford and Hilton 2001; Brailsford, Churilov and Liew 2003; Morecroft and Robinson 2006). Others

have compared the differences in model-building approaches by users of DES and SD (Tako and Robinson, 2009). However the concept of genuinely combining both approaches in the field of OR remains unrealized. This is not the case in other disciplines, such as engineering or computer science, but there the concept and nature of simulation are slightly different. For example, hybrid simulation executives designed to handle both continuous and discrete model parameters can be found from the 1990's in the computer science literature (Monsef et al: 1995, Music and Matko: 1999) and the continuous/discrete software Modelica (www.modelica.org) was developed for mechanical engineering applications. However, we argue that although these tools do produce models which have both continuous and discrete variables, they do not really capture the spirit and flavor of SD and DES as understood by OR modelers.

Chahal and Eldabi (2008) identify three modes in which DES and SD can be combined in practice. The simplest is the "hierarchical" mode in which there are actually two distinct models which simply pass data from one to the other. Rabelo (2003) and Venkateswaran et al (2005) use this approach to develop architectures for combining DES and SD in one conceptual model in the context of manufacturing enterprise. They develop a united architecture to govern the run-time infrastructure but use two separate software packages, Arena for DES and Powersim for SD, to demonstrate the methodology. The second mode is the "process environment" where there are still two distinct models, but the DES model actually sits inside the SD model and models a small section of the system, which then interacts cyclically with the wider SD environment. This is the approach used in the two healthcare examples discussed in this paper. Finally, in the genuine "integrated" mode, there is one single hybrid model with no clear distinction between the discrete and continuous parts. We argue that this "nirvana state" has not yet been achieved.

Several OR software packages, such as Witness www.lanner.co.uk, provide the facility to model continuous, as well as discrete event phenomena. They can therefore be adapted to provide the underlying structures of SD models. Martin and Raffo (2001), for example, developed a hybrid simulation by modifying ExtendTM, a package designed for both discrete event and continuous simulations. AnyLogic www.xjtek.com/AnyLogic is another package which can provide both DES and SD models (and also agent-based models) with nearly all the main features available in individual software packages such as Vensim, Powersim and iThink (SD tools) and Arena, Simul8 and ProModel (DES tools). However despite these efforts, these packages remain essentially either a DES environment with some continuous features, or a continuous environment with some discrete or stochastic features.

2 THE HOLY GRAIL

Despite this interest, and the early literature, we shall assert in this paper that the modeling community in 2010 is not really much closer to producing a genuinely hybrid modelling methodology that combines the characteristic features of both DES and SD, than it was back in 2000. Merely including both continuous and discrete variables in the same model is only half the story. The "holy grail" is a methodology which combines the benefits and virtues of each approach, allowing a truly holistic systems view yet at the same time capturing the detailed individual variability within parts of that system which is the strength of DES models. SD has a qualitative aspect which enables insights to be obtained very rapidly without the need for a massive data collection exercise. Quantitative SD models are fast to run and can be used in an interactive gaming mode with stakeholders in order to gain understanding of non-intuitive system behavior. Yet SD ignores stochasticity and can lead to misleading conclusions in systems where there is uncertainty or a need to model individual variability. A truly integrated approach would be advantageous because at a macro level, it could describe the movement of individual entities as a homogeneous flow, which would be fast and data-efficient, whereas at a micro level, where there were detailed interactions that affected the overall behaviour of the system, it would be possible to incorporate individual characteristics. The problem is therefore, not to develop the software, which is clearly now available, but to develop both a conceptual philosophy and a practical methodology for combining SD and DES in a real context.

There are also technical issues for a genuinely hybrid model, first identified in this context by Ruth Davies (2001, private communication), in moving backwards and forwards between continuous and dis-

crete models. We shall not discuss these issues here in any detail, but they raise some very interesting research questions. How much information loss can be tolerated in switching from DES to SD and back again? How can essential information be retained? These technical questions do not appear to have been addressed in much detail in the mainstream OR literature, where most papers have been at more of a paradigm or conceptual level. It would seem to be relatively easy to merge a group of individual entities into a continuous mass, but how much of the “population profile” would one need to retain and how would one do it? And coming back the other way, when one starts with a continuous population and wants to split this up into individual entities (which seems much harder to do with any degree of rigor), one would need to define algorithms which transformed the original population (the last time it was discrete) into a new population, based on what had happened to it during the time it was continuous. Could these difficulties be avoided by not actually “passing” entities from one part of the model to another?

In this paper we shall address the following issues. Starting with a brief technical and philosophical overview of the two approaches, we shall consider when and why a combined model would be needed in practice, what benefits such a model could provide, and which parts of the system would need to be modeled in sufficient detail to warrant using a DES model. How can experimentation be performed effectively with a combined model, given the totally different experimental approaches of SD and DES? We then consider two case studies, both in a health or social care context, and discuss these questions. We shall conclude by discussing how far we have come along the road to the “holy grail”.

3 SD AND DES: THE TECHNICAL AND PHILOSOPHICAL DIFFERENCES

3.1 System dynamics

The basic principle underlying system dynamics is that the structure of a system determines its behaviour over time (Forrester 1961; Sterman 2000). In other words, the way that the separate components of any system relate to and affect each other determines the emergent behavior of the system as a whole. Such emergent behavior can be counterintuitive, and it is only by analysis of the component subsystems that the reasons for this unexpected behavior can be understood. SD has two distinct aspects, one qualitative and one quantitative. The qualitative aspect involves the construction of *causal loop* diagrams. Through discussions with problem owners and other stakeholders, the relationships between identified system elements are graphically depicted by a system of arcs and nodes, where the polarity of an arc indicates the direction of influence, positive or negative. The aim is to identify feedback loops, which can be of two kinds: *balancing loops* which retain a steady-state, or *vicious circles* leading to uncontrolled growth. The understanding and insights that this approach can bring are very useful. However the overall net effect of all the feedback loops in a very complex system cannot be determined merely by inspecting the diagram. To do this it is necessary to quantify the variables, and this is not always straightforward if some variables (e.g. “happiness”) are qualitative. Quantitative SD modeling requires the use of stock-flow diagrams. These models are best conceptualized as a system of water tanks connected by pipes. A domestic central heating system is not a bad analogy! Water flows from tank to tank and the rate of flow is governed by taps or valves on the pipes. The “water” which flows around such a system is a continuous quantity. Mathematically, stock-flow SD models are a discretization of a set of ordinary differential equations representing the rates of change of the level of each stock; these ODEs are solved numerically using a discrete time-step dt . Clearly, these models are deterministic. For a fuller description see Brailsford (2008).

3.2 Discrete event simulation

DES is a widely established simulation approach which needs little introduction to the WSC readership. Essentially, it is a stochastic modelling approach ideally suited to queuing network systems, where state changes occur at discrete points of time and individuals (*entities*) move stochastically through a system of queues and “services” or activities whose durations are governed by probability distributions. It is a high-

ly flexible approach in which almost anything can be coded; models can be incredibly detailed and most DES software has a graphical interface which allows the user to see the system operating on the screen, almost like watching a movie. Because of stochasticity, multiple replications of simulation runs are required to obtain statistically significant results.

3.3 Brief comparison of SD and DES

For detailed comparisons of the two approaches, and in particular excellent descriptions of the different world-views of DES and SD modelers, see Lane (2000) and Morecroft and Robinson (2006). Here we briefly outline some of the key differences.

SD is essentially a deterministic approach, and purists might argue that it is not strictly true to call it simulation. It does not consider individual “entities” or handle variability very effectively, despite efforts by software vendors to introduce probability distributions. In comparison with DES, SD models are extremely quick to run as obviously they do not require multiple iterations. SD undeniably lacks the total flexibility of DES, which can use virtually any probability distribution function, or empirical data, to model state dwelling times. Traditionally SD has been used at a higher, more aggregated and strategic level than DES. Forrester (1961) believed strongly that SD models were “learning laboratories” and were definitely not optimization tools. The data requirements of an SD model are generally much less than for DES. Validation of SD models is a contentious issue given their qualitative nature; it is not possible to apply the same battery of statistical validation tools as to a DES model, but other methods have been developed (Rodrigues and Bowers, 1996). While SD models are useful for clarifying the complexities of organizational behavior, their simplified representation of systems, the necessity to aggregate entities, and the use of average flow rates are some of their significant limitations. In contrast, DES models are much more flexible, capturing interactions between entities and detailed characteristics of the system being modeled. However, a major disadvantage of DES is that the data requirements, and the need to maintain a next-events list and to carry out long runs or multiple replications to get reliable results, mean that the models are relatively time-consuming to develop and run.

Philosophically, there is a profound difference between the two approaches which over the years has led to the development of two entirely separate research communities, with their own journals and conferences (Lane, 2000). DES has been universally recognized as part of the OR “toolbox” since the 1960’s whereas SD, although it has an equally long pedigree, has really only become widely accepted as a mainstream OR technique in the past 10-15 years (Brailsford, 2008). Basically, SD modelers and DES modelers just see the world in a different way; this affects the way they conceptualize and model problems, which is far more subtle than merely the distinction between discrete and continuous variables. Very crudely put, an SD modeler sees the world as a holistic synthesis of system elements which are dynamically connected, and takes a “helicopter view” of the world, whereas DES modelers take their microscope to the world and look at the system in detail, paying attention to the variability between individual components. To say that SD is strategic and top-down and DES is operational and bottom-up is too crude an approximation, but it does give a flavor of one of the key conceptual differences between them.

4 WHY COMBINE DES AND SD?

We have already hinted at why it would be useful to combine these two approaches. It is often very difficult to define the boundaries of a model in a system which appears self-contained but is actually very connected into the wider environment. Such systems appear all the time in healthcare. For example consider some of the classic application areas of DES, operating rooms (ORs), Intensive Care Units (ICUs) or Emergency Departments (EDs). There are countless examples of DES models for all of these in the literature. Yet to begin with, the OR is inextricably linked with the ICU – if there are no beds in the ICU because of an unexpected number of emergency admissions, then elective surgeries will be canceled. And both also inextricably linked with the ED, because that’s where the emergency admissions come from.

The ED is inextricably linked with almost every other department in the hospital – the wards, X-ray and other diagnostic services, the pharmacy, pathology, portering and transportation services, availability of specialist physicians, etc. And it gets worse, because the ED is the front door into the hospital from the outside world and so we have links to the ambulance system, primary care, public health, social care, unemployment, the environment ... and so it goes on. Of course the modeler can decide pragmatically to draw a line and say “beyond this I will not go” but this could mean that the model excludes some key elements or drivers of system performance.

A devotee of DES would say that it is possible to include all these aspects in a single model, and of course this is true, it is technically possible. But why would you want to do it? Such a model would be massive, cumbersome, slow to run and data intensive. We are only really interested in the performance of the ED; we don't want to model the whole hospital in detail, let alone the whole health economy. All we want to do (as in *any* model) is to be sure we have captured the important features of the real system, i.e. those key aspects of the “outside world” that impact on the system of interest. Therefore, a broad-brush model which captures these critical features and within which our detailed ED model sits would suit our purpose perfectly. It is those very characteristics of SD which would fit this purpose; the whole-system, big picture approach which slices through all this detail and produces an elegant structure which illuminates the key relationships and captures those external factors which we *cannot* exclude from our ED model.

5 CASE STUDY 1: CHLAMYDIA INFECTION

Chlamydia trachomatis is the most widely prevalent sexually transmitted infection (STI) in the UK. The Department of Health estimates that the chlamydia prevalence rate is around 10% in people aged between 16 and 24 (Department of Health, 2010). Between 1999 and 2008 the number of annual cases of chlamydia more than doubled (National Audit Office, 2009). This is of particular concern because although chlamydia itself is easily treated by a short course of antibiotics, 10-30% of infected women go on to develop pelvic inflammatory disease, which can lead to ectopic pregnancy and infertility. A significant proportion of cases, particularly amongst women, are asymptomatic and so are liable to remain undetected. Population screening for chlamydia infection would identify these asymptomatic cases, but blanket screening of the entire 16-24 year old population is prohibitively expensive and would need to be repeated regularly, since reinfection often occurs. Therefore targeted screening of high-risk people is required.

In 2003 the UK Department of Health set up the National Chlamydia Screening Programme (NCSP). Portsmouth, where this study is based, was one of two NCSP pilot sites. It is a large naval port in the south of England and also has a big student population. Since then the NCSP has expanded the number of places offering chlamydia screening. Between 2008 and 2009 an estimated 16% of young people aged 15-24 in England were tested for chlamydia at a community setting and by 2009 an estimated 1.5 million chlamydia tests had been performed under the programme (National Audit Office, 2009)

The main aim of the modeling study was to provide a planning toolkit for healthcare professionals at the Department of Genito-Urinary Medicine (GUM) in St Mary's Hospital in Portsmouth to inform their decision making. There were a number of other aspects of the study, involving geomapping of the local prevalence data in order to better target screening campaigns, and also a questionnaire survey to develop an understanding of young people's risk-taking behaviour and the impact of education. However the key issue was to estimate future local incidence and prevalence, and then evaluate different clinic designs and staffing policies to improve the clinic performance. Overall, the target was to reduce the level of all STIs in the Portsmouth area.

An SD model was developed in Vensim to represent disease transmission in the community. This was a classical SIR (susceptible-infected-recovered) model, similar to that of Evenden et al (2003), but with two risk categories (high and low), a mixing function between these categories, and a behavior change function which moves people from high risk to low risk, dependent on the success of publicity campaigns. This model generates cases of chlamydia, a proportion of which (depending on the screening

rates) become the demand for the GUM clinic. The model uses data from the clinical literature and from earlier work by the Department of Health (Townsend and Turner 2000), as well as local data from the NCSP.

The DES model of the clinic, developed in Simul8, then processes the demand based on its configuration, i.e. the numbers of doctors, nursing and reception staff, treatment rooms, capacity of waiting areas, and so on. The clinic model is a classical DES: a screenshot is shown below in Figure 1. The model is very detailed and produces many performance measures, such as throughput, waiting times, number of patients not seen at the end of the day, and so on. It uses data collected in the clinic, both by direct observation and from historical records. The main output (the number of successfully treated patients) is then fed back into the SD model, enabling the incorporation of a more realistic and accurate calculation of the number of patients treated than could be obtained from an SD model alone. The interface between these models was automated, and runs in Excel, since both the DES software (Simul8) and the SD software (Vensim) can communicate with Excel using Visual Basic for Applications.

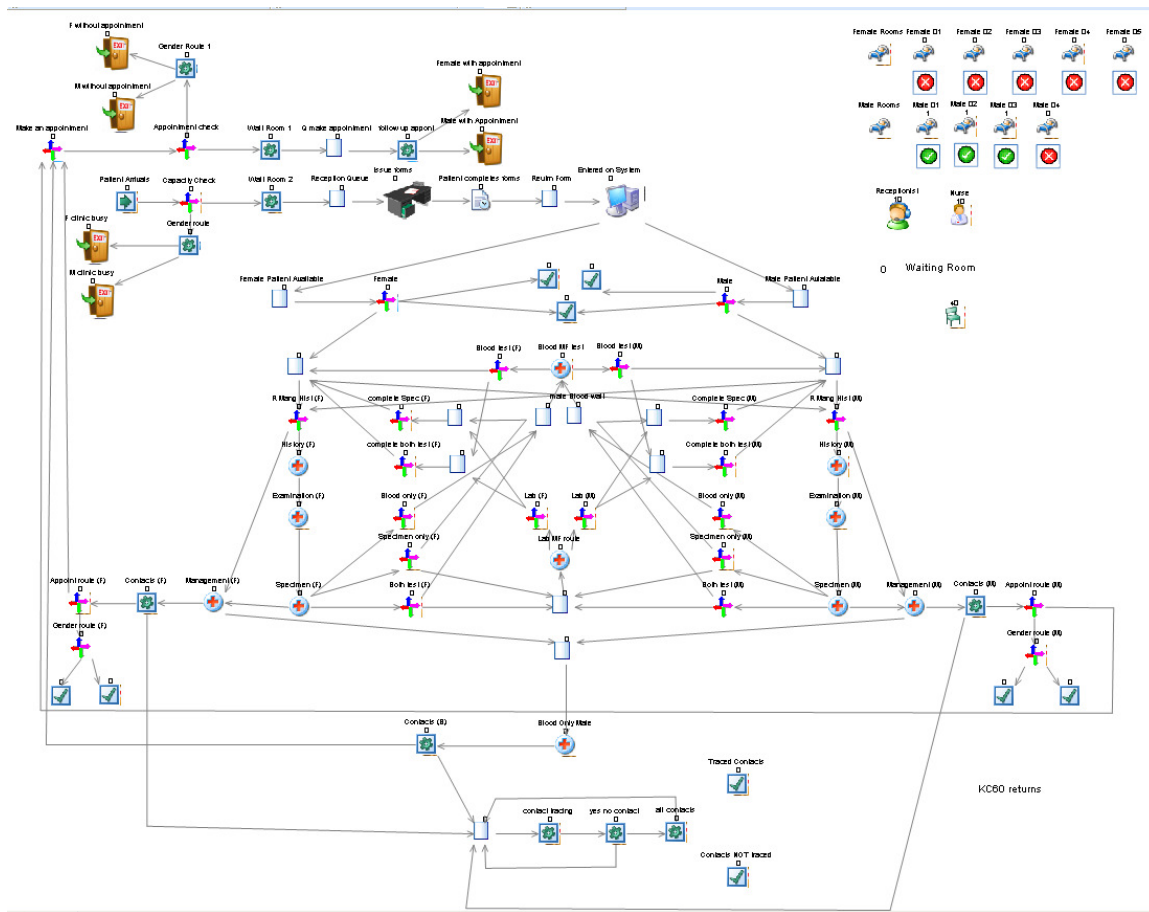


Figure 1: Screenshot of the GUM clinic model

Clearly, the “chlamydia planning toolkit” is not a single hybrid model but consists of two separate models which automatically pass data from one to the other through an Excel interface. The clinic DES model runs for one month with 20 replications and takes under 2 minutes to run on a laptop with an Intel® Core™ i3 CPU 2.13 GHz with 4.00 gigabytes of RAM. The SD model runs for 24 months in monthly intervals. The SD model takes less than one second to run. The models stop after 24 months.

The benefit of including feedback via the SD model is as follows. The medical director of the GUM clinic was worried that the clinic would not be able to cope with the influx of new patients which could arise as a result of a sudden increase in the screening rate, and that long waits and poor performance could actually discourage people from attending. In other words, the NCSP could be counter-productive unless the GUM clinic could deal with the likely demand. Therefore “clinic performance” (in terms of the mean and standard deviation of the waiting time, and the numbers of patients not seen at the end of the day) are included as inputs in the SD model, and have an impact on the screening and detection rate. This in turn obviously impacts on the prevalence and incidence of chlamydia in the community, which ultimately feeds back full cycle in the form of increased arrivals at the clinic.

6 CASE STUDY 2: SOCIAL CARE IN HAMPSHIRE

The second case study is a collaborative research project between the University of Southampton and Hampshire County Council (HCC). Hampshire is a large county on the south coast of England which is attractive for retirement and therefore has a higher than the national average proportion of older people. In common with the rest of the developed world, the population is ageing and the demographic profile of Hampshire is changing. HCC is responsible for providing social support services to all vulnerable people in the community. The problem is that the number of older people requiring care or support is likely to increase in future, as more people survive into their 90's. The aim of the study was to help the local authority better plan services and support for older people. In particular, HCC wanted to assess the impact of a new contact centre (a telephone call centre) which they hoped would speed up access to services and also, make it easier for HCC to “right-size” the services offered to clients, thereby reducing costs.

As in the chlamydia study, two separate models were developed. System dynamics was used at a population level to model long-term dynamic demographic change and thus better predict future demand for long-term care in Hampshire, and DES was used to develop a detailed model for the contact centre. The aim of the SD model was to understand changes in demand for long term care in Hampshire amongst those aged 65 years and over. The DES model addressed operational issues around the organization and staffing of the contact centre, in order to meet the changing needs of the population it serves over the next twenty years.

Long-term care provision is provided from a variety of sources of which the local authority is only one. The SD model, developed in Vensim, calculates how many older people are likely to turn to HCC for help, as well as the private sector, the voluntary sector, and the informal sector. In addition the model also calculates how many people are likely to need care or support but will not receive anything, because their needs are not sufficiently high to meet the eligibility criteria. The data for the model was obtained from a range of sources including the Office for National Statistics, the National Household Survey and HCC themselves. The model showed that HCC needs to work with other stakeholders in order to meet the increase in future demand for services. By supporting some client groups now, the number requiring expensive and intensive support later on their life can be reduced. The SD model replicates the whole system for long-term care and projects it forward to the year 2026. Older people in Hampshire are characterized by age, gender, disability and chronic ill health, using a disability scale compatible with the current eligibility criteria. Older people are categorized by service receipt if they are likely to receive some sort of service.

One of the key outputs of the SD model is the number of older people who are likely to contact the contact centre. This information is fed into the detailed DES model, which was developed in Simul8. In this model the activities of the contact centre are replicated and the performance of the contact centre measured over time. This model calculates the first service received by a new client. This can be changed to model the service received after eight weeks of making contact with the local authority, as many clients' service requirements are likely to change over this period of time. This is also the length of time a client would receive free urgent unplanned crisis care. The data for this model was obtained partly from the contact centre's historical records, but also by direct observation and talking to the staff. One of the

very interesting issues raised was that although the call centre software could provide a wide range of performance data, it could not give us some key parameters for the model, for example call durations. These had to be collected by direct observation. This is typical within public sector organizations: masses of data inappropriate for modeling are collected for performance monitoring whereas data useful for modeling are simply not kept.

Once again, this is two separate models, but in this case, they do not automatically pass data from one to the other, but are manually linked. The contact centre model runs for one month with forty replications and takes around four minutes to run on a laptop with an Intel® Core™ i3 CPU 2.13 GHz with 4.00 gigabytes of RAM. The SD model has been configured to run for a maximum of eighteen years and takes less than one second to run. The hybrid framework is run for the years 2010, 2015 and 2020.

Two key performance indicators are reported by the contact centre model: the percentage of calls answered, and the percentage of calls answered within twenty seconds. The contact centre management endeavors to meet the following two requirements: 95% of calls are answered, and 80% of calls are answered within twenty seconds. If either of these performance indicators is not met a further experiment is carried out where a staff intervention takes place. There are two tiers of call handlers; tier one staff are agents (less experienced) and tier two staff are advisors (more experienced but more expensive). These types of staff have different capabilities, i.e. they can deal with different types of call and take different times to deal with calls. Various combinations of staff are tried in our experiments, increasing the staffing levels until the performance indicators are met.

The benefit of including feedback via the SD model is as follows. It is quite likely that a proportion of people who abandon their calls might attempt to call again at a later point. Most people will probably call back the same day or later that week. However some callers might not bother calling again until much later on, by which time their condition may have deteriorated and their needs become greater. Therefore if the call centre was not staffed sufficiently, this could lead to greater numbers of people with unmet needs in the community, which ultimately would lead to greater demand for more expensive services. Obviously the exact numbers of this type of callback is unknown, so various experiments are performed under different assumptions about the effects of this phenomenon.

An example of this intervention is shown in Figure 2 below, for 2020, with 10% feedback. This means that 10% of abandoned calls lead to a second call one month later in which the needs status of the caller is increased by one severity category compared with the original call. In order to meet the performance standards, two additional advisors were required in January and February. With two additional advisors the percentage of calls fell below the 95% threshold in March. A third additional advisor is required.

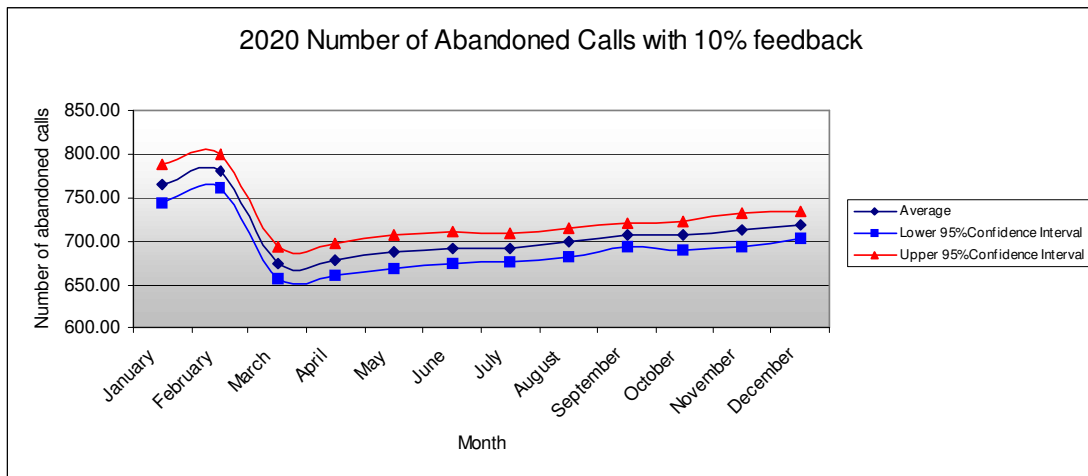


Figure 2: 2020 Number of Abandoned Calls with 10% feedback

7 CONCLUSION

Could we have done either of the case studies in either DES or SD alone? The answer is probably yes, but it would have been a tortuous process: a case of hammering in a screw. In both models we have a very large population in the community (tens of thousands of people) and we are using SD to model a broad process in this community which has a timescale of months (or years, in the case of the contact centre model). In the case of the chlamydia SIR model, SD is ideally suited to modeling the infection process and the feedback between the pools of susceptible and infected people. We could potentially have used DES for these, but it would have increased the runtimes enormously, for doubtful benefit. It is less clear that we could have used SD for the call center and clinic models: these are both highly detailed, stochastic systems where individual variability is key, and the performance measures relate to resource utilization and queuing times, both characteristics for which SD is less well suited. Neither of these examples is a true hybrid: they are both two distinct models with a cycling structure, in which a detailed DES model representing a small part of the overall system is embedded in a “big picture” SD model. This required quite a lot of programming, in the case of the chlamydia study, or actual manual intervention to import and export data and run the models, in the case of the social care model. However it is clear that in both cases there is a benefit from combining these approaches which is greater than the sum of its parts.

There is a school of thought which says that it doesn't matter what modeling approach is used, it is the *process* which counts, and that getting people together round a table to talk about the problem situation is the really important thing. The argument is that in this case the model is of secondary importance and acts only as a catalyst to motivate discussion, or as a boundary object (Carlile 2002) providing a common ground for different stakeholders to interact. While we have sympathy with this view, we would argue that we do believe the approach chosen *does* matter, as it focuses attention on certain aspects of the real-world system which are quite different, according to the approach adopted. Indeed this is one of the hypotheses being tested in a collaborative research project currently under way involving colleagues from Imperial College, London (Bayer et al, 2010).

So, are we any nearer the holy grail than we were 10 years ago? The answer is yes, up to a point. A lot of DES software now has “SD-like” features – and vice versa. However, it is clear that we are still a long way from genuinely combining these approaches, and indeed it may not actually be possible to do so from a philosophical standpoint. However it is undeniably true that we have learned a lot along the road, as the increasing body of literature in this field demonstrates.

REFERENCES

- Bayer S., Bolt T., Brailsford S.C. and Kapsali M. 2010, to appear. The social role of simulation models. *Proceedings of the 2010 International System Dynamics Conference*, Seoul, Korea.
- Brailsford S.C. 2008. System Dynamics: What's in it for Healthcare Simulation Modelers. In *Proceedings of the 2008 Winter Simulation Conference*, eds. S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler, 1478-1483. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Brailsford S.C. and Hilton N.A. 2001. “A comparison of discrete event simulation and system dynamics for modelling healthcare systems”, *Proceedings from ORAHS 2000*, Glasgow, Scotland, ed. J. Riley, 2001, pp 18-39.
- Brailsford, S., Churilov, L. and Liew, S. 2003. Treating ailing emergency departments with simulation: An integrated perspective. In *Proceedings of Western Multiconference on Health Sciences Simulation*, Florida, 2002, ed J. Anderson. Society for Modelling & Simulation International (SCS), San Diego, CA..
- Carlile P. 2002. A pragmatic view of knowledge and boundaries: boundary objects in new product development. *Organization Science* 13: 442-455.
- Chahal, K. and Eldabi, T. 2008. Applicability of hybrid simulation to different modes of governance in UK healthcare. In *Proceedings of the 2008 Winter Simulation Conference*, ed. S. J. Mason, R. R.

- Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler, 1469-1477. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Department of Health, <<http://www.chlamydia-screening.nhs.uk/>>, [accessed 06.10.10].
- Davies, R.M. 2001. Private communication.
- Dyson R.G. and O'Brien F.A. 1998. Strategic Development: Methods and Models. Wiley, Chichester, UK.
- Eden C.L. 1989. Using cognitive mapping for strategic options development and analysis (SODA). In J. Rosenhead (ed.), Rational Analysis for a Problematic World. Wiley, Chichester, UK.
- Evenden, D., Harper, P.R., Brailsford, S.C., Harindra, V. 2006. Improving the cost-effectiveness of Chlamydia screening with targeted screening strategies, *Journal of the Operational Research Society*, 57:1400-1412.
- Forrester J.W. 1961. Industrial Dynamics. MIT Press, Cambridge, MA, reprinted by Productivity Press (1994) and now available from Pegasus Communications, Waltham, MA, USA.
- Lane D.C. 2000. You just don't understand me: Modes of failure and success in the discourse between system dynamics and discrete event simulation. LSE OR Dept Working Paper LSEOR 00-34, London School of Economics and Political Science..
- Martin R. and Raffo D. 2001. Application of a hybrid process simulation model to a software development project, *The Journal of Systems and Software*, 59:237-246.
- Monsef Y., Hallam N. and Khodja R. 1995. NEPT-MIX: An object-oriented mixed simulation language generator. *Proceedings of the First LAAS International Conference on Computer Simulation*, Beirut , pp 43-47.
- Morecroft, J. and Robinson, S. 2006. Comparing discrete event simulation and system dynamics: modelling a fishery. *Proceedings of the 2006 OR Society Simulation Workshop SW06*, 137-148.
- Music G. and Matko D. 1999 Combined simulation for process control: extension of a general purpose simulation tool. *Computers in Industry* 38:79-92.
- National Audit Office, 2009. Young people's sexual health: the National Chlamydia Screening Programme, <<http://www.chlamydia-screening.nhs.uk/ps/publications/reports.html>>, [accessed 06.10.10]
- National Chlamydia Screening Programme Steering Group. 2006. New Frontiers: Annual Report of the National Chlamydia Screening Programme in England 2005/06, London, HPA
- Rabelo L., Helal M., Jones A., Min J., Son Y.J. and Deshmukh A. 2003. A hybrid approach to manufacturing enterprise simulation, In *Proceedings of 2003 Winter Simulation Conference*, eds. S. E. Chick, P. J. Sanchez, D. M. Ferrin, D. J. Morrice, 817-825. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Robinson S. 2005. Discrete-event simulation: from the pioneers to the present, what next?. *Journal of the Operational Research Society*, 56:619-629.
- Rodrigues, A. and Bowers, J. 1996. System dynamics in project management: A comparative analysis with traditional methods. *System Dynamics Review* 12:121-139.
- Sterman J.D. 2000. Business Dynamics: Systems Thinking and Modeling for a Complex World, Irwin McGraw-Hill, Boston, MA.
- Tako A.A. and Robinson S. 2009. Comparing discrete-event simulation and system dynamics: users' perceptions. *Journal of the Operational Research Society*, 60:296-312.
- Townsend, J.R.P. and Turner, H.S. 2000. Analysing the effectiveness of Chlamydia screening, *Journal of the Operational Research Society*, 51:812-82.
- Venkateswaran J. and Son Y.J. 2005. Hybrid system dynamic-discrete event simulation-based architecture for hierarchical production planning, *International Journal of Production Research*, 43:4397-4429.

AUTHOR BIOGRAPHIES

SALLY C. BRAILSFORD is Professor of Management Science at the University of Southampton, UK. She received a BSc in Mathematics from the University of London, and MSc and PhD in Operational Research from the University of Southampton. Her research interests include simulation modeling methodologies, system dynamics, health service research and disease modeling, and the modeling of human behavior in healthcare systems. She is chair of the European Working Group on OR Applied to Health Services (ORAHs) and is on the editorial boards of Health Care Management Science, the Journal of Modeling in Management, the Journal of Simulation and the Flexible Services & Manufacturing Journal. Her email address is: s.c.brailsford@soton.ac.uk.

SHIVAM DESAI is a final-year PhD student in the School of Management at the University of Southampton, UK. His research is sponsored jointly by the UK Engineering and Physical Sciences Research Council and the Economic and Social Research Council. His advisors are Sally Brailsford, Professor Paul Harper from the University of Cardiff, UK and Professor Maria Evandrou from the Gerontology division in the School of Social Sciences at Southampton. Shivam obtained a BSc in Politics and an MSc in Management Science at Southampton. His email address is msd102@soton.ac.uk.

JOE VIANA is also a final-year PhD student in the School of Management at the University of Southampton. His research is sponsored by the UK Engineering and Physical Sciences Research Council. His advisors are Sally Brailsford and Professor Paul Harper from the University of Cardiff, UK, with clinical advice from Dr Veerakathy Harindra from St Mary's Hospital, Portsmouth, UK. Joe obtained a BSc in Sport and Health Sciences and an MSc in Management Science at Southampton. His email address is jv104@soton.ac.uk.