Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

SYSTEM DYNAMICS: WHAT'S IN IT FOR HEALTHCARE SIMULATION MODELERS

Sally C. Brailsford

School of Management
University of Southampton
Southampton SO17 1BJ, U.K.

ABSTRACT

In the past decade there has been an explosion in the use of system dynamics modeling in healthcare. Despite this, the approach is still far less well known than discrete-event simulation in the mainstream Operations Research community. This paper contains an introduction to system dynamics, illustrated by several examples in the field of healthcare, and discusses some of the possible reasons for the growth in the popularity of this approach for healthcare modeling.

1 INTRODUCTION

Traditionally, in the US, the UK and elsewhere, the academic research communities of system dynamics (SD) and discrete-event simulation (DES) have been fairly separate. They hold their own conferences, have their own journals, and indeed have their own views of the world. In 2000, the Simulation Special Interest Group of the UK OR Society held a joint meeting with the UK Chapter of the System Dynamics Society, entitled "Never the Twain Shall Meet". At this meeting David Lane presented a ground-breaking paper (Lane, 2000) in which he discussed the differences between SD and DES, and posed the question about whether they were "chalk and cheese", or were actually two sides of the same coin. This meeting led to the foundation of a new Special Interest Group called "SD+" whose aim was to bring the SD and DES communities together. It also led to a number of papers exploring the differences between DES and SD, including the well-known work by Robinson and Morecroft (2005, 2006) which compared the use of DES and SD to model the same system, the dynamics of fishing fleets and fishery stocks.

Despite these initiatives, it is still true to say that system dynamics is less well known to OR people than DES. It is highly probable that the number of papers on DES at this conference will greatly exceed the number of SD papers. Nevertheless in some application areas the use of SD is expanding rapidly, and healthcare is one such area. Figure 1 shows the number of hits in a search of the electronic database of literature Web of Science using the search string <"*system dynamics" AND "health**">, from 1980 to 2007. The apparent decline in 2007 is probably an artifact of the lag in adding publications to the electronic databases, rather than an actual drop in the numbers.

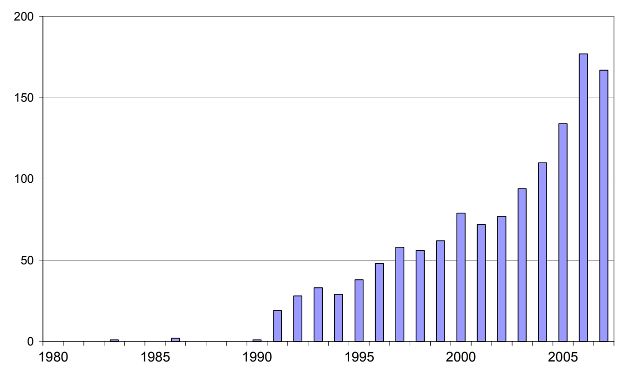

Figure 1: Health SD publications by year, 1980 – 2007

I should declare my own personal interest. Having started out in the early 1990's as a dedicated user of DES, I became interested in SD as a result of the "Never the Twain" meeting, and subsequently became a zealous convert, using SD myself for several modeling studies. Like many other researchers, I have long been interested in the relationship between DES and SD, in particular in the domain of healthcare (Brailsford and Hilton, 2001), but I first used SD in practice in a project to model demand for emergency health care in Nottingham, England (Brailsford et al, 2004).

2 THE PRINCIPLES OF SYSTEM DYNAMICS

System dynamics is an analytical modeling approach whose roots could be said to lie in the general systems theory approach of von Bertalanffy (1968). However the foundations of SD were undoubtedly laid in the 1950's at

MIT by Jay Forrester in his pioneering work on "industrial dynamics" (Forrester 1960, 1961). The fundamental principle of SD is that *structure determines behavior*: 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. Often such emergent behavior can be counterintuitive, and it is only by analysis of the component parts that the reasons for this unexpected behavior can be understood.

2.1 Influence diagrams

SD has two distinct aspects; one qualitative and one quantitative. The qualitative aspect involves the construction of *causal loop* or *influence* diagrams, which depict graphically the way in which the system elements are related. The aim is to enhance understanding of a problem situation through the structure of the system and the relationships present between relevant variables. Through discussions with problem owners and other stakeholders, the identified system elements are represented in the form of an causal loop diagram (CLD), an example of which is shown in Figure 2.

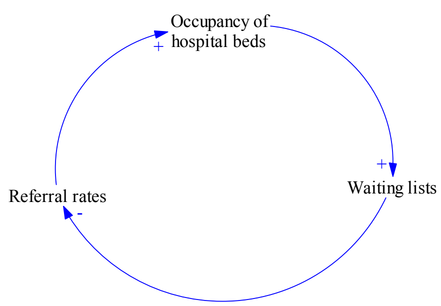

Figure 2: CLD for UK hospital admissions

The system depicted here is the process by which patients in the UK are admitted to hospital. A patient requiring hospital treatment has to be referred to a hospital consultant by their General Practitioner (GP) or family doctor. The patient is then placed on a waiting list and is admitted when a bed becomes available. Long waiting lists are a major political issue in the UK.

The identified elements are connected by arrows. The "+" and "–" signs denote the direction of the influence, but not its magnitude. For example, as referral rates increase the bed occupancy increases, shown by a "+". However, as waiting lists increase, GP referral rates tend to decrease, since long waiting times will lead GPs to seek other alternatives to hospital admission. This effect is denoted by a "–" . In this way complex and informative diagrams can be built up to represent and clarify the system being investigated, providing insights into how the various components interact.

In many cases the qualitative analysis of these diagrams is of value in its own right. The aim of this analysis is to find loops, as in the above example, where elements are connected by a directed cycle of arrows. Balancing loops contain an odd number of "–" signs, whereas reinforcing loops or vicious circles contain an even number of "–" signs. Balancing loops retain the status quo and keep the system in steady-state, whereas in vicious circles the system spirals out of control. Figure 2, a balancing loop, shows that waiting lists have an important role to play in regulating the hospital admissions system and keeping it in steady state. Identifying both types of loop can be very helpful in understanding system behavior.

One of the benefits of qualitative modeling is that it can cast light on the unintended consequences of actions. For example, as we said, long waiting lists are a key issue in the UK's National Health Service and are one of the performance indicators on which hospitals are rated. As waiting lists grow, there is increasing political pressure on hospitals. A proposed solution might be to provide funding for more beds, thus increasing capacity and reducing occupancy. This leads to another balancing loop, shown in the top right-hand side of Figure 3. However there is an unintended effect of "throwing money" at the problem in this way, namely that the inevitable publicity about these extra beds leads to behavioral change in both patients and GPs. GPs are obviously more likely to refer if they think there are new beds available, and patients are more likely to insist on referral. This leads to an increase in referral rates, and thus to a vicious circle, shown in the top half of Figure 3, in which all the links are positive. We see that longer waiting lists lead to greater political pressure, which results in additional extra funding, which increases referral rates, which lead to higher occupancy, which leads to longer waiting lists, etc, etc.

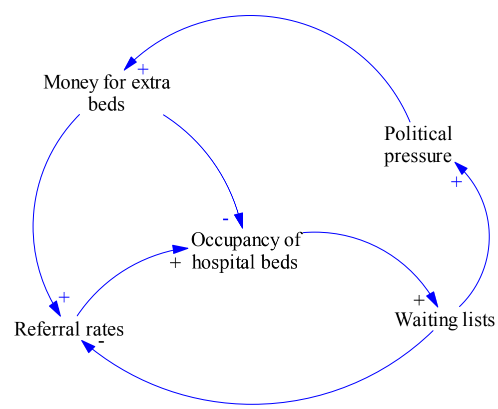

Figure 3: Unintended consequences

Of course in practice the effect of this vicious circle is mitigated by the two balancing loops (the new one and the original one). The overall net effect of these three loops cannot be determined merely by inspecting the diagram. To do this we would need to quantify the variables in Figure 3, and this is not always straightforward, given that variables like "political pressure" are not easily quantified. However the understanding and insights that this approach can bring are very useful, even if no further quantitative modeling is carried out.

2.2 Stock-flow models

For quantitative SD modeling, the CLD has to be converted to a stock-flow diagram. 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 corresponds to the entities in a DES model, only in SD it is a continuous quantity. Thus it may represent money, people, material, product, and so on.

Figure 4 shows the stock-flow diagram constructed from the CLD of Figure 2 in the notation of the software *Vensim* (Ventana Systems, 2008). Here the "water" represents patients. The two "clouds" represent a source and a sink, in other words patients outside the system. The rectangle represents a stock or level (an accumulation), of patients, namely patients occupying beds in hospital. The two arrows (pipes) represent admissions and discharges, and the valves represent the rate of flow along these pipes. We have added a discharge flow (not required in the original CLD) since otherwise patients would never leave hospital.

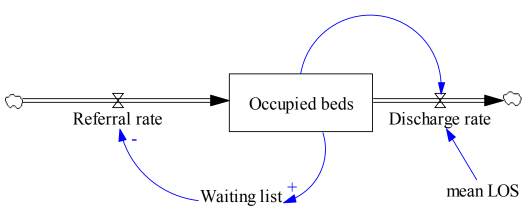

Figure 4: Stock-flow model for hospital admissions

We have retained some of the original notation from the CLD, in that "Waiting List" is now an *auxiliary variable* which influences the referral rate, and is itself influenced by the size of the "Occupied beds" stock. The form of these influences needs to be quantified and the software allows a variety of ways in which this can be done, ranging from analytical mathematical relationships to graphical functions depicting how X varies with Y.

The model now needs to be parameterized, by defining the referral and discharge rates and the initial value of the Occupied beds stock.. The discharge rate could be a simple function of the average length of stay, which is another auxiliary variable Defining the "Waiting list" variable and its interactions is more complex and in practice would require additional auxiliary variables, for example the total number of available beds.

The model can now be run or "simulated", although it should be pointed out that although most SD software does allow limited use of probability distributions, there is generally no variability or stochastic aspect in an SD model. The model shows how variables change over time, allowing their behavior to be monitored and analyzed. Time is handled in these quantitative models by a discretization process where the time-step, *dt*, is usually chosen such that all the rates can be regarded as constant over the period dt. In the above example, denoting by *Occ(t)* the bed occupancy at time *t*, we have the level equation

$$Occ(t+1) = Occ(t) + referral\ rate*dt - discharge\ rate*dt$$

The time spent in each stage can be modeled by the use of delays. The simplest type of delay is the exponential delay. If for example the average length of stay is 4 days and *dt* is 1 day, the discharge rate is equal *to Occupied Beds/4*. Modern SD software has the facility to implement more complex types of delay function, such as pipeline delays and batch delays, which permit non-exponential dwelling times in the various stages to be modeled. However SD undeniably lacks the total flexibility of discrete event simulation, which can use virtually any probability distribution function, or empirical data, to model state dwelling times.

Figure 5 shows an example of output from the above model with highly simplified data (given in the Appendix) to illustrate how the number of occupied beds reaches a steady state after about 2 months, showing the effect of the balancing loop.

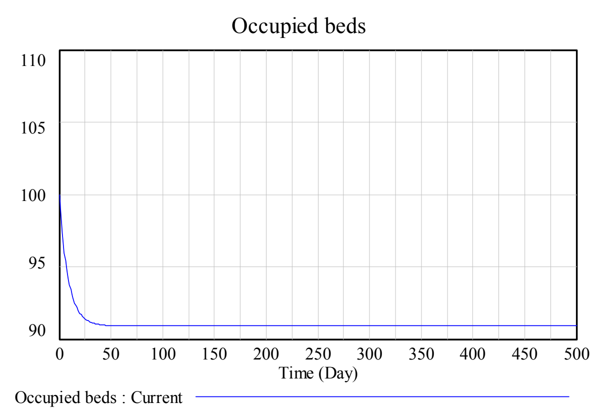

Figure 5: Illustrative Vensim output

2.3 Brief comparison with DES

For detailed comparisons of the two approaches, and in particular excellent descriptions of the different world-views of DES and SD modelers, we refer the reader to Lane (2000) and Robinson and Morecroft (2005, 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, for example the use of devices such as "conveyors" in the software *Stella* (2008) – also known as *ithink* . Time, and material, are continuous variables, although the models are discretized and solved mathematically by difference equations rather than differential equations. SD models are extremely quick to run as obviously they do not require multiple iterations.

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. SD models can be highly complex but they exhibit *dynamic complexity* rather than *detail complexity*.

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).

In recent years there has been a trend towards user-friendly SD software that does not require any knowledge of computer programming. Both *Stella* and *Vensim*, for example, have "drag and drop" interfaces so that the user can select icons for levels, rates, etc, connect them up, and edit their properties; the software then automatically generates the underlying equations and runs the model, collecting and presenting the output. However, there is a much smaller range of SD software available and it would be fair to say that this is not yet as sophisticated in terms of graphics and user interfaces as DES software.

Some DES proponents argue that there is nothing which SD does which DES could not do equally well. Software such as *Anylogic* (<www.xjtek.com> 2008) is capable of using both approaches (although not, strictly speaking, at the same time!) Other researchers are interested in combining the two approaches.

3 SYSTEM DYNAMICS IN PRACTICE: SOCIAL CARE IN THE SOUTH OF ENGLAND

To illustrate the use of SD in practice, we describe a model for the social care system in the county of Hampshire, UK (Desai et al, 2008). The objective of the UK Social Services system is to work with vulnerable people in the community to promote independence, reduce harm and provide support. Many of its functions interface with other agencies, mainly the National Health Service (NHS). The elderly are high consumers of health and social care and so these people may move from Social Services' care to NHS care, and back again, many times. In the UK the social care service is publicly funded, but unlike health care it is means-tested and there are strict eligibility criteria. Therefore, people requesting assistance from Social Services are carefully assessed to determine their actual need and their financial status. There are four levels of need: Critical, Substantial, Moderate and Low. At the time of this study, only clients with Critical or Substantial need were deemed eligible to receive support.

Hampshire Social Services (HSS) serves a population of over 1.2 million people in the south of England. The area is attractive for retirement and so the proportion of people aged 85 and older is higher than the UK national average, and is increasing. All UK Social Services departments are facing significant challenges in balancing increasing demand for older people's services (due to the ageing population) with severe budgetary constraints. To help meet these challenges, HSS needed to project this future demand up to 2011 and consider potential interventions, in the face of major funding cuts.

A system dynamics model of the social care system was developed, incorporating both qualitative and quantitative aspects. A qualitative CLD was developed through a series of interviews with HSS staff, in which the major issues affecting demand for care were identified, along with the external and internal constraints and pressures on HSS staff which affected their ability to deliver care. At the same time, during these interviews a flowchart was developed to depict how clients initially access care, are assessed, and then move through the social care system. Finally, based on a combination of this flowchart and the CLD, a quantitative stock-flow model was developed using the software *Stella* (2008), and a user interface created in Microsoft Excel to facilitate future use of the model by HSS staff. This allowed the model parameters to be changed without an expert knowledge of *Stella* and the model re-run for different assumptions.

For a detailed description of the flow diagram and the stock-flow model which was developed in *Stella*, see Desai et al (2008). A simplified version of the CLD is shown in Figure 6. This contains a number of feedback loops, for example a vicious circle showing that if more applicants for care are rejected initially, they will return later, possibly with greater needs. There is also a balancing loop, not dissimilar to the simple waiting list example in section 2, showing how referral rates serve to regulate the system.

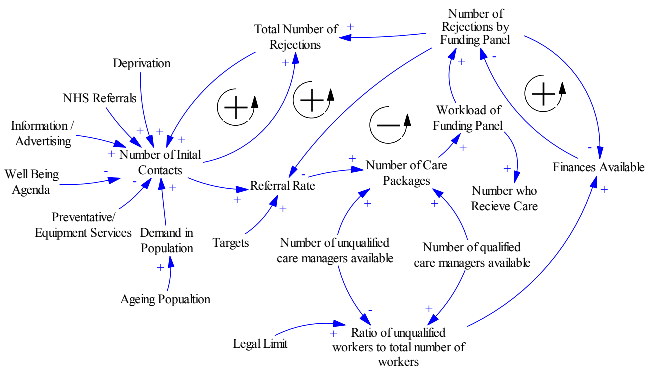

Figure 6: CLD for Social Services care system

The stock-flow model is complex and involves thirteen different categories of care, including daycare, various levels of domiciliary care (care delivered in clients' own homes), various categories of residential care, and mental health services including dementia. Although SD does not model individual characteristics, it does allow groups with given characteristics to be modeled separately via the use of arrays. Through discussions with HSS staff the following three key factors were identified: age, referral source (NHS or elsewhere) and initial level of need. Unfortunately *Stella* only allows for two dimensions of characteristics, so one dimension combined the age group and the referral source, and the other dimension was the initial level of need.

At the time a hot debate within HSS concerned reducing the numbers of clients receiving care, by changing the eligibility criteria for funding. The knock-on effects of this were difficult to anticipate, and widely different views were expressed. Another scenario of interest involved changing the staff mix of care workers; again the effects of this could not be directly anticipated. Both these issues were highly politically sensitive.

The results of the baseline run of the model confirmed that the numbers requiring care will increase over the next five years, particularly among clients aged over 85. The effects of two possible interventions were then explored through changing the model parameters. This demonstrated that providing care only to the highest priority (Critical) clients will indeed reduce the numbers receiving care, and the costs of that care. However, a decrease in the number of "Substantial" clients funded did not lead to the same percentage decrease in the total costs of care. Some of the short-term savings were offset by future increases due to these clients requiring extra care later on, as shown by the CLD in Figure 6. The overall effect was reduced to about 84% of the expected benefit. The second scenario gave approximate values for the savings that may be achieved by varying the proportions of qualified care workers and unqualified care workers.

There were numerous benefits in using SD for this study, rather than DES which might have seemed the more natural choice for a "patient flow, queuing network" type of problem. Firstly HSS were less concerned with developing precise forecasts of future demand and capacity, than with understanding the drivers for that demand and the overall likely effects of various proposed strategies for managing it. The conversation held with HSS staff were therefore at a higher level of abstraction than would have been the case for a DES model. Secondly, the social care system is hugely complex and the temptation in a DES model would have been to develop a massive, detailed model which would have been very time-consuming to build and parameterize. One of the strengths of SD is the way it focuses attention on the system structure, and not the data, which is not always available or accurate.

4 WHY SD IS GOOD FOR YOUR HEALTH

There are a number of reasons why SD is particularly well suited to healthcare problems, and could help to account for the explosion in applications shown in Figure 1. These reasons are illustrated by the case study in section 3. This model is trying to represent a very large, complex system whose boundaries overlap with other organizations. It is rare in a healthcare setting that one is able to draw well-defined boundaries around the system of interest and ignore any interactions with the wider environment. In many healthcare problems there are multiple stakeholders with conflicting objectives and different levels of ownership and power within the system as a whole. The qualitative aspects of SD are very helpful when trying to understand such issues.

Data availability and quality is another issue. While any simulation model can be run with estimated data or "expert opinion", a DES model really does require a substantial amount of detailed data in order to be able to fit distributions and so on. There is only so much which can be usefully gained from sensitivity analysis using limited data. On the other hand, the data requirements of SD are typically very much less, since SD models are usually higher-level and more aggregated. More useful modeling can be done with less data. Some healthcare settings are rich in useful data, but most are not – the key word being "useful".

In a DES model, we often cannot see the wood for the trees – we are so obsessed with the detail that we lose sight of the big picture. In SD, we cannot model this detail, but we can gain understanding of the dynamic complexity of the system. Unlike DES, SD models are not dependent on vast quantities of high-quality data, and so can be used at a more speculative or strategic level, for larger populations and longer time-horizons. A key advantage of SD is that the models generally run very fast (and of course do not require multiple iterations), so can be run interactively in real time with decision-makers.

In April 1999 Geoffrey Coyle and John Morecroft edited a Special Issue of the *Journal of the Operational Research Society* on system dynamics. This issue marked a key point in the acceptance of SD by the mainstream OR community – it predated the Warwick meeting, and brought SD to the attention of modelers who might not otherwise have considered this approach. It contained a paper on European healthcare SD applications by Brian Dangerfield in which he commented that at the time, "*the literature is not vast…*" but "*There is clear potential for system dynamics to be employed in support of health care policy.*" (Dangerfield, 1999). Since then, the use of SD has continued to grow worldwide, although possibly less rapidly in the US, despite MIT being the birthplace of SD. We hope that this paper will help encourage US healthcare modelers to consider using SD for their next project.

APPENDIX

Equations and data for the illustrative SD model of waiting list behavior

dt = 1 day
Initial number of occupied beds = 100
Mean Length of Stay = 10 days
Waiting List = 0.1* Occupied Beds
Referral rate = 10 – (0.1*Waiting List)
Discharge rate = Occupied Beds / Mean LOS

REFERENCES

AnyLogic. 2008. <www.xjtek.com>.

Bertalanffy, L. von .1968. General System Theory: Foundations, Development, Applications. New York: George Braziller.

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.C., V.A. Lattimer, P. Tarnaras and J.A. Turnbull. 2004. Emergency and on-demand health care: Modelling a large complex system, *Journal of the Operational Research Society* 55:34-42.

Dangerfield, B.C.. 1999. System dynamics applications to European health care issues. *Journal of the Operational Research Society* 50:345-353.

Desai, M.S., Penn, M.L. Brailsford, S.C. and Chipulu M. 2008. Modelling of Hampshire Adult Services - gearing up for future demands. *Health Care Management Science*, 11:167-176.

Forrester, J.W. 1960. The impact of feedback control concepts on the Management Sciences. In: Collected Papers of J.W. Forrester, (1975 collection), pp 45-60, Cambridge, MA: Wright-Allen Press.

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..

Morecroft, J. and Robinson, S. 2005. Explaining Puzzling Dynamics: Comparing the Use of System Dynamics and Discrete-Event Simulation. The 23rd International Conference of the System Dynamics Society, July 17-21, 2005, Boston. Conference CD.

Morecroft, J. and Robinson, S. 2006. Comparing Discrete-Event Simulation and System Dynamics: Modelling a Fishery. Proceedings of the Operational Research Society Simulation Workshop 2006 (SW06) (J. Garnett, S. Brailsford, S. Robinson and S. Taylor, eds.). Operational Research Society, Birmingham, UK, pp.137-148.

Rodrigues, A. and Bowers, J. 1996. System dynamics in project management: A comparative analysis with traditional methods. *System Dynamics Review* 12:121-139.

Stella. 2008. High Performance Systems, 145 Lyme Road, Hanover, NH.

Sterman J.D. 2000. Business Dynamics: Systems Thinking and Modeling for a Complex World, Irwin McGraw-Hill, Boston, MA.

Vensim. 2008. <www.ventanasystems.com>.

AUTHOR BIOGRAPHY

SALLY BRAILSFORD is Professor of Management Science at the University of Southampton, UK. She received a BSc in Mathematics from the University of London, and an 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 on the editorial boards of Health Care Management Science, the Journal of Modeling in Management and the Journal of Simulation. She is currently Vice President of the UK OR Society and co-chair of its 50th Annual Conference.
Email: <`s.c.brailsford@soton.ac.uk`>.