

**DISTRIBUTED SIMULATION WITH COTS SIMULATION PACKAGES:
A CASE STUDY IN HEALTH CARE SUPPLY CHAIN SIMULATION**

Navonil Mustafee
Simon J. E. Taylor

Centre for Applied Simulation Modelling
School of Information Systems, Computing & Maths
Brunel University
Uxbridge, Middlesex, UB8 3PH, UK

Korina Katsaliaki
Sally Brailsford

School of Management
University of Southampton
Southampton, SO17 1BJ, UK

ABSTRACT

The UK National Blood Service (NBS) is a public funded body that is responsible for distributing blood and associated products. A discrete-event simulation of the NBS supply chain in the Southampton area has been built using the commercial off-the-shelf simulation package (CSP) Simul8. This models the relationship in the health care supply chain between the NBS Processing, Testing and Issuing (PTI) facility and its associated hospitals. However, as the number of hospitals increase simulation run time becomes inconveniently large. Using distributed simulation to try to solve this problem, researchers have used techniques informed by SISO's CSPI PDG to create a version of Simul8 compatible with the High Level Architecture (HLA). The NBS supply chain model was subsequently divided into several sub-models, each running in its own copy of Simul8. Experimentation shows that this distributed version performs better than its standalone, conventional counterpart as the number of hospitals increases.

1 INTRODUCTION

The UK National Blood Service (NBS) is a public funded body that is responsible for distributing blood and associated products. The analysis of this health care supply chain is of particular interest as blood donors are in short supply, the shelf-life of blood products is relatively short and blood product ordering policies are potentially complex. Discrete-event simulation is a commonly used modeling technique (Robinson 2005, Law and Kelton 2000, Pidd 1998) that can be used to analyze supply chains (Banks et al. 2002). Practitioners in this area typically use visual interactive modeling environments or *Commercial Off-The-Shelf Simulation Packages* (CSPs), such as AnyLogic, Arena, Simul8, and Witness, that exploit developments in visual interactive modeling to facilitate, for example, model building and experimentation (Swain 2003). The NBS sup-

ply chain simulation described in this paper consists of the Southampton area NBS Processing, Testing and Issuing (PTI) facility and over twenty associated hospitals. The original model was built using the CSP Simul8. Researchers at the University of Southampton discovered that as this model grew in size with the addition of more hospitals, the time taken to perform one simulation run increased to a point that made the use of simulation infeasible (Katsaliaki and Brailsford 2006).

One potential solution to this problem is to use distributed simulation. This is a technique where models are implemented over many computers in a parallel or distributed fashion with the goals of reducing the execution time of a single simulation run, sharing the memory needs of a simulation across several computers and the linking of simulations sited in different locations (Fujimoto 1999; 2003). Some form of distributed simulation middleware is generally needed to coordinate the advancement of simulation time and for passing messages between the individual models, for example, a Run Time Infrastructure (RTI) as specified by the High Level Architecture (HLA) (IEEE1516 2000). The Simulation Interoperability Standards Organization's (SISO) CSP Interoperability Product Development Group (CSPI PDG) <www.cspi-pdg.org> is developing a standard approach to using the HLA to support distributed simulation with CSPs such as Simul8. In an attempt to solve the runtime problems of the NBS supply chain simulation, this paper investigates how a "conventional" model of a supply chain can be distributed over several computers running different copies of a CSP linked together by a HLA RTI. This paper demonstrates that distributed simulation is a useful technique for this type of health care supply chain and, as standards for CSPs are maturing, argues that urgent research is needed in this area to determine if distributed simulation can be widely used to support the analysis of supply chains and other such applications in health care.

This paper is structured as follows. Section 2 gives a background of the NBS supply chain. Section 3 and Section 4 presents the conventional and the distributed approach to simulating the NBS simulation. This is followed by experimentation and presentation of results in Section 5. Section 6 discusses the findings and Section 7 draws the paper to a close.

2 BACKGROUND TO THE UK SUPPLY CHAIN OF BLOOD

The UK NBS is a part of the National Health Service (NHS) Blood and Transplant (NHSBT) organization. NHSBT was formed on 1st October 2005 by the merger of National Blood Authority (NBA) (which manages the NBS, Bio Products Laboratory and the International Blood Group Reference Laboratory) and UK Transplant (NHS Blood and Transplant, 2006). The NBS is responsible for collecting blood through voluntary donations, testing the blood for ABO and Rhesus grouping and infectious diseases such as HIV, processing the blood into around 120 different products (of which the main three are Red Blood Cells, plasma and platelets), storing the stockpile and transferring excess stock between different NBS centers, and finally issuing the different blood products to the hospitals as per their needs (Brailsford et al. 2006). The NBS infrastructure consists of 15 Process, Testing and Issuing (PTI) centers which together serve 316 hospitals across England and North Wales.

Blood products are stored in PTI Centres until they are requested by the hospitals served by that Centre. A hospital places an order for blood products when its inventory falls below a predetermined order point, or when rare products not held in stock are requested for particular patients. Hospitals normally receive their orders daily and the blood remains in the hospital bank until it is cross-matched (tested for compatibility) for a named patient. It is then placed in “assigned inventory” for that patient for a fixed time after the operation. If it is not used, it is returned to “unassigned inventory” and can be cross-matched again for another patient. On average a unit will be cross-matched four times before it is used or outdated. In practice, however, only half of the cross-matched blood is actually transfused. This clearly represents a huge potential for savings since the cost of a single unit of RBC is around £132. Using simulation to model the NBS supply chain can potentially help to decrease shortages and wastage of blood products, reduce collection and delivery costs and increase service levels.

3 THE “CONVENTIONAL” NBS SIMULATION

In this study we modeled the NBS Southampton PTI and include only RBC and platelets, which together comprise 85% of issues and are the chief source of wastage and shortages. The model was originally built using the Simul8

and is described in detail elsewhere (Katsaliaki and Brailsford 2006). The supply chain model is very large and complex, and requires extensive data. Nineteen months’ data from the Southampton PTI Centre was provided and analyzed using the NBS information system PULSE. This gave details of the products supplied to each hospital, by date, time, delivery type, quantity and blood group. Questionnaires were sent to the hospitals supplied by the Southampton centre, and interviews conducted with NBS staff and hospital blood bank managers.

There are two main categories of entities in the model; items and orders. Items are the individual blood units (RBC and platelets) delivered from the NBS Centre to the hospitals in a one-way direction, since returns of products are not allowed. Orders are placed by the hospital blood bank managers to the NBS Centre for blood products, and represent the backwards flow of information. Requests are matched with items according to their characteristics (attributes) as in a Kanban system and delivered as appropriate.

While the model runs, data such as the day and time of placing an order, the requested blood product (RBC or platelets), the amount by blood group, etc. are reported to an Excel file. The model runs in minutes (for a single hospital), and the remaining shelf-life of blood products is counted in minutes. However, the hospitals’ blood bank stock for placing orders to the NBS is checked only every hour.

The smallest version of the model contains the processes of the NBS Centre, from collection of whole blood to delivery of blood products, and the processes within a single medium-volume hospital. The model captures physicians’ requests for blood and the processes whereby the hospital blood bank checks its stock levels and places orders. Figure 1 shows a simplified illustration of this model. For multiple hospitals, Figure 2 shows an example of the relationships between the NBS supply centre (PTI) and the hospitals it serves. The order entities and item entities are represented in figure 2 as information flow and material flow respectively. They all run in a single copy of the CSP Simul8 and on a single PC.

The time taken to execute the conventional, single computer NBS model rises exponentially as more hospitals are added to the model (see section 6). Thus, it is argued that (a) as the size of the supply chain increases it is beyond the capability of a single computer to simulate it, and (b) distributed simulation offers a potential solution to this (Taylor et al. 2002; Lendermann, Gan, and McGinnis 2001; Gan et al. 2000). It allows us to execute a single simulation model, perhaps composed of several distinct models, over multiple computers that are connected through a network (Fujimoto 2001). In the next section, we consider our alternative to the “conventional” approach by introducing a distributed approach to modelling the NBS supply chain simulation inspired by the work of SISO’s CSPI PDG.

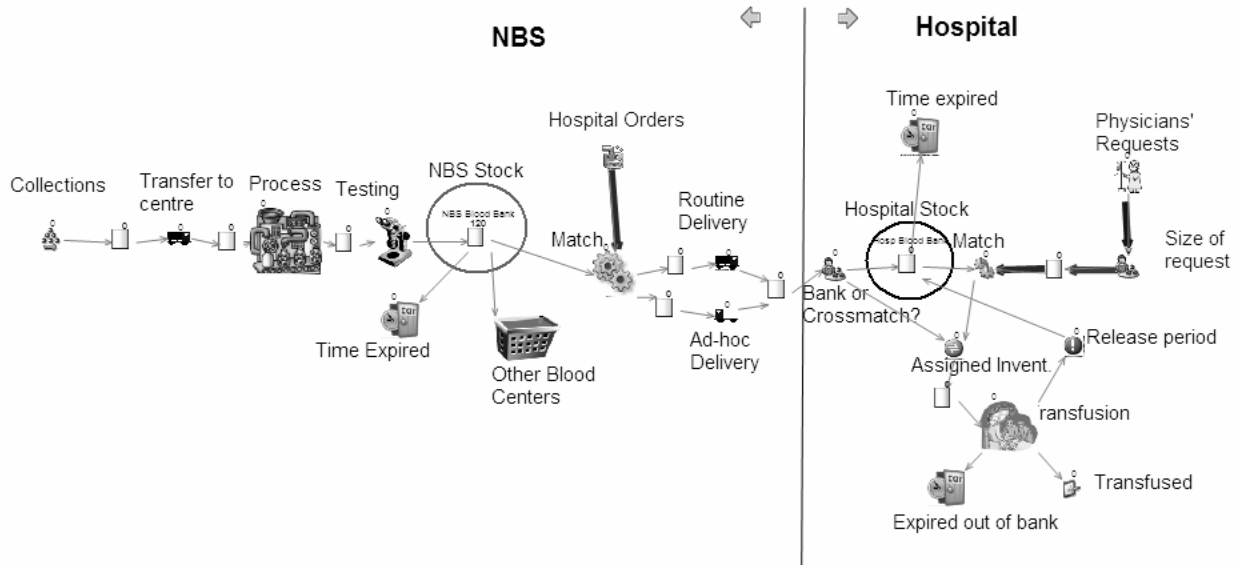


Figure 1: Screenshot of a Simplified Version of the Simul8 Model Showing One Hospital

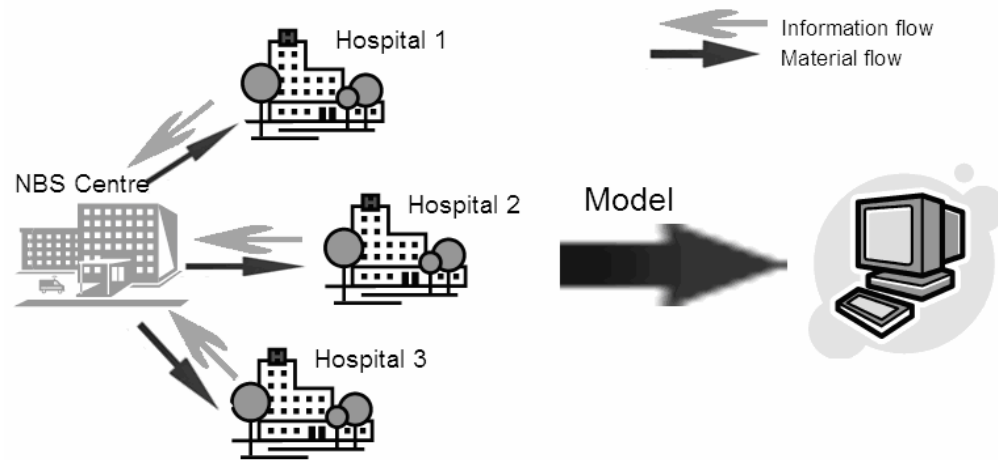


Figure 2: Conventional Simulation Approach

4 THE DISTRIBUTED NBS SIMULATION

As presented elsewhere in WSC 2006, SISO's CSPI PDG is developing a standard approach to using the HLA with CSPs for distributed simulation. (Taylor et al. 2006b). This work has identified what effectively amounts to "templates" for different distributed simulation (interoperability) requirements made by CSPs. These are the so-called Interoperability Reference Models (IRMs). Each IRM is supported by a data exchange specification and an interoperability framework. The distributed NBS simulation fits

the profile of the Type I IRM *Asynchronous Entity Passing* and our approach to developing a distributed version of Simul8 and the NBS model is inspired by this (see Taylor et al. (2006a) for an in-depth discussion of these issues). Following this, to create a distributed simulation of the NBS supply chain we (1) divide up the conventional model into different model elements, (2) use an RTI conforming to the HLA and (3) interface the Simul8 CSP with RTI as suggested by the CSPI PDG Type I interface framework (using interactions).

The model decomposition creates individual models of the Southampton PTI and hospitals (in this case 4 different

hospitals have been built). These model run in separate copies Simul8. Together these form federates that interact by timestamped messages that represent the interaction of one model part with another (e.g., when an entity leaves one part of a model and arrives at another). These are mapped onto HLA interactions. The complete model, constituted of distributed federates, form our NBS supply chain federation. Note that in this work, Simul8, like most of the other CSPs, does not provide inbuilt support for distributed simulation. The modifications made to the Simul8 CSP are described in Mustafee and Taylor (2006).

Figure 3 shows the distributed model of the NBS. Each part model (federate) represents either the NBS supply centre or a hospital running in a separate copy of Simul8 running on a separate computer. A *Manager* federate (not shown), a special program that coordinates the execution of the federation through registration of synchronization points (Kuhl, Weatherly, and Dahmann 1999), runs with the HLA-RTI process on the same computer as the NBS model.

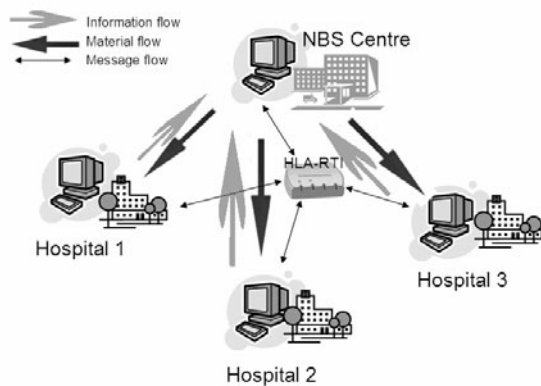


Figure 3: Distributed Simulation Approach

In this investigation, to make our approach as transparent as possible, interaction between the models/Simul8 and the HLA-RTI is via an Excel file. For example, entities representing orders are written into the file by Simul8 during the execution of hospital models. The HLA-RTI, augmented by the CSPI PDG standards, then correctly transfers this information to the NBS model by means of HLA interactions. The incoming orders from each hospital are collected into their corresponding queues in the NBS model and the orders are matched with the available stock of blood. The resulting matched units are written into an Excel spreadsheet in the NBS federate. This information is then sent to the different hospital models in a similar manner. The decision to implement the distributed supply chain

in this manner was motivated by issues of end user transparency and ease of implementation.

5 EXPERIMENTATION AND RESULTS

To investigate the distributed approach against the conventional approach, four scenarios were investigated. These were one NBS supply centre serving one, two, three and four hospitals respectively. The hospitals which were added to the models were all of the same size. For instance, physician requests were around 1000 blood units for each hospital per month, with each hospital diverging by a small percentage ($\leq 6\%$) from the mean. Before experimentation commenced, the outputs for the conventional and distributed models were compared to check that the same results for a year's run was produced. This was done to validate the minor modifications to link Simul8/Excel/HLA-RTI in the distributed model did not artificially increase/decrease the workload. All experiments were conducted on Dell Inspiron laptop computers running Microsoft Windows XP operating system with 1.73GHz processors and 1GB RAM connected through a 100Mbps CISCO switch. The same computer specifications were used to guarantee consistency in runtimes. The results of the execution times for each of the models are based on the average of 5 runs.

Figure 4 shows the execution time in seconds for both conventional and distributed approaches as the NBS simulation progress, month by month. The results show that the conventional model with one hospital took approximately 14 minutes to run for a whole simulated year. The run time rose to 78 minutes when the model ran with two hospitals and to approximately 17.5 hours with three hospitals. The addition of the fourth hospital increases the execution time to 35.8 hours. The distributed model with one NBS supply centre and one hospital ran in approximately 8.5 hours, with two hospitals in 9.8 hours, with three hospitals in 12.7 hours and with four hospitals in 16.5 hours.

6 DISCUSSION

From the results it is apparent that the versions with one or two hospitals are less time consuming to run using the conventional approach. Conversely, when a third and fourth hospital are added then the distributed method bests the runtime of the conventional approach. There also appears to be an exponential escalation of the runtime in the conventional version while increasing the number of hospitals in the model. This is quite a contrast to the substantially smaller and smoother rise in the runtime in the distributed method.

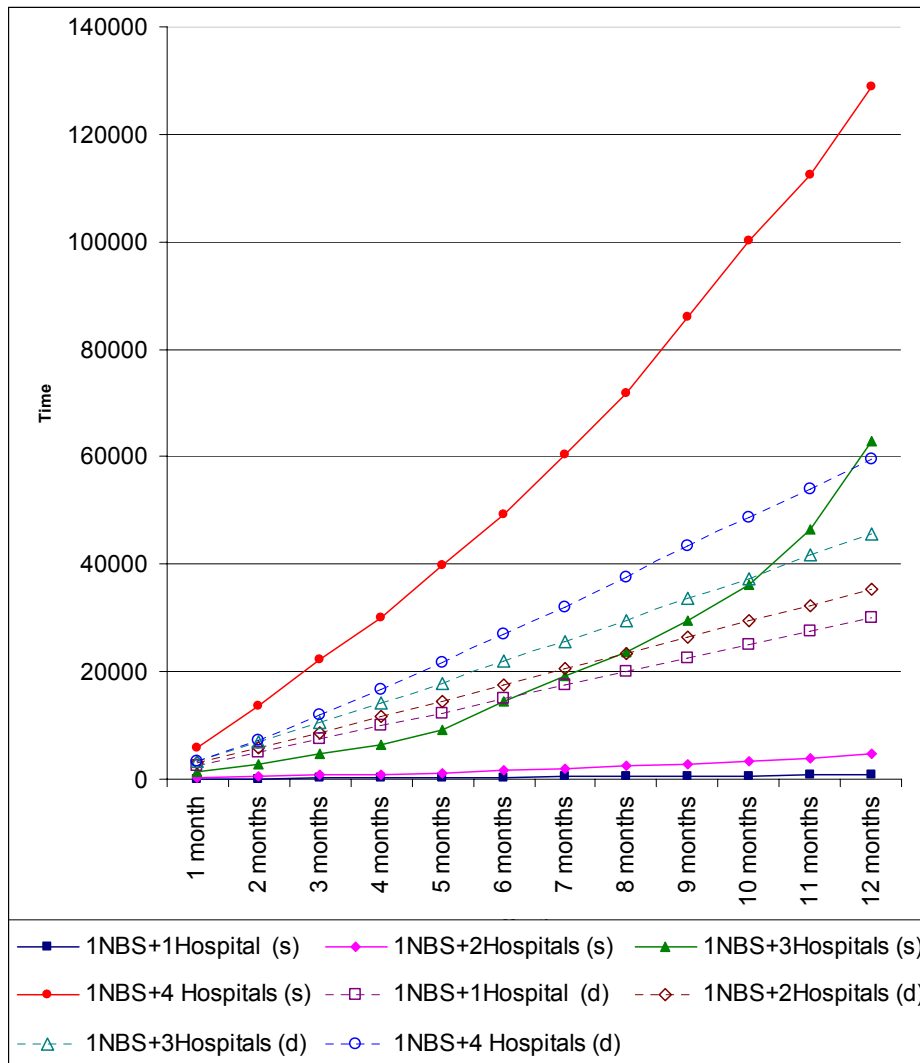


Figure 4: Runtimes of Conventional (s) and Distributed Method (d) for One NBS Centre with One to Four Hospitals

The enormous number of entities in the system, each of which carries many attributes, increases the computation time of the conventional model exponentially even though there is no exponential element in the functions of the model. This exponential increase appears to result from a combination of two different factors. Firstly, the massive amount of information generated by the model cannot be accommodated in the random access memory (RAM) alone, and hence the operating system has to keep swapping information to the hard disk. The part of the hard disk which is kept aside for use as swap space is called virtual memory. It takes much more time for the processor to recall information from the virtual memory when compared to RAM. Thus, as the models get bigger, more information is generated resulting in more swaps between the RAM

and the virtual memory, thereby contributing to an increase of execution time. Secondly, the behavior of the system being modeled is such that all entities (blood units) in the system have a limited shelf life. This behaviour is modelled in the NBS simulation by continually scheduling events that decrease the shelf-life of each entity by the minute. This results in more computations as the number of entities flowing through the system increases. Thus, the increase in runtime appears to be primarily due to a large event list caused by a combination of the volume of entities and the “counting down” of the shelf-life. The large event list in turn causes swapping between RAM and virtual memory, which further causes long runtimes. Our results suggest that the distributed approach allows the processing

and the memory demands made by large event lists to be shared over several computers.

7. CONCLUSIONS

This paper has described an investigation into using conventional and distributed approaches to simulating the supply chain of blood from a National Blood Service Centre to hospitals with the simulation package Simul8. For this scenario at least, when the supply chain grows in size and complexity, distributed simulation appears to offer a viable alternative to conventional simulation by sharing the processing and memory requirements of the simulation across multiple computers.

Future work will involve simulating larger NBS supply chain models and investigating the effect that using different RTIs have on performance. Alternative implementation strategies with the CSP are also intended to be investigated.

ACKNOWLEDGMENTS

The authors would also like to thank Dr. Mark Elder (founder and CEO of Simul8 Corporation) for providing the Simul8 licences and his generous on-going support.

REFERENCES

- Banks, J., S. Buckle, S. Jain, P. Lendermann, and M. Manivannan. 2002. Panel Session: Opportunities for Simulation in Supply Chain Management. *Proceedings of the 2002 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc02papers/226.pdf [accessed June 30, 2006].
- Brailsford, S., K. Katsaliaki, N. Mustafee, and S. J. E. Taylor. 2006. Modelling Very Large Complex Systems using Distributed Simulation: A Pilot Study in a Healthcare Setting. In *Proceedings of the 2006 Operational Research Society Simulation Workshop (SW06)*, Leamington Spa, UK. March 28-29, 2006. pp. 257-262.
- Fujimoto, R. M. 1999. Parallel and Distributed Simulation. *Proceedings of the 1999 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc99papers/016.PDF [accessed June 30, 2006].
- Fujimoto, R. M. 2001. Parallel and Distributed Simulation Systems. *Proceedings of the 2001 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc01papers/018.PDF [accessed June 30, 2006].
- Fujimoto, R. M. 2003. Distributed Simulation Systems. *Proceedings of the 2003 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc03papers/015.pdf [accessed June 30, 2006].
- Gan, B. P., L. Li, S. Jain, S. J. Turner, C. Wentong, and J. H. Wen. 2000. Distributed Supply Chain Simulation Across Enterprise Boundaries. *Proceedings of the 2000 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc00papers/168.PDF [accessed June 30, 2006].
- IEEE 1516. 2000. IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA). New York, NY: Institute of Electrical and Electronics Engineers.
- Katsaliaki, K., and S. B. Brailsford. Using Simulation to Improve the U.K. Blood Supply Chain. *Journal of the Operational Research Society* [to appear].
- Kuhl, F., R. Weatherly, and J. Dahmann. 1999. Creating Computer Simulation Systems, An Introduction to the High Level Architecture. Prentice Hall PTR.
- Law, A.M., and W. D. Kelton. 2000. Simulation Modeling and Analysis, 3rd ed. McGraw-Hill: New York.
- Lendermann, P., B. P. Gan, and L. F. McGinnis. 2001. Distributed Simulation with Incorporated APS Procedures for High-Fidelity Supply Chain Optimization. In *Proceedings of the 2001 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via www.informs-cs.org/wsc01papers/154.PDF [accessed June 30, 2006].
- Mustafee, N., and Taylor, S. J. E. 2006. Investigating Distributed Simulation with COTS Simulation Packages: Experiences with Simul8 and the HLA. In *Proceedings of the 2006 Operational Research Society Simulation Workshop (SW06)*, Leamington Spa, UK. March 28-29, 2006. pp. 33-42.
- NHS Blood and Transplant. © NHS Blood and Transplant 2006. Available via www.nhsbt.nhs.uk [accessed June 4, 2006].
- Pidd, M. 1998. Computer Simulation in Management Science, 4th ed. Wiley: Chichester, UK.
- Robinson, S. 2005. Discrete-event Simulation: from the pioneers to the present, what next? *Journal of the Operational Research Society*, 56(6), pp. 619-629.
- Swain, J. J. 2003. Simulation Reloaded: Sixth biennial survey of discrete-event software tools. *OR/MS Today*, 30(4), pp. 46-57.
- Taylor, S. J. E., R. Sudra, T. Janahan, G. Tan, and J. Ladbrook. 2002. GRIDS-SCS: An Infrastructure for Distributed Supply Chain Simulation, *Simulation*, 78(5), pp. 312-320.

Taylor, S. J. E., X. Wang, S. J. Turner, and M.Y.H. Low. 2006a. Integrating Heterogeneous Distributed COTS Discrete-Event Simulation Packages: An Emerging Standards-based Approach. *IEEE Transactions on Systems, Man and Cybernetics: Part A*, 36(1), pp. 109-122.

Taylor, S. J. E., S. J. Turner, M. Y. H. Low, J. Ladbrook, and X. Wang. 2006b. The COTS Simulation Package Interoperability Product Development Group (SISO CSPI PDG). *Proceedings of the 2006 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

simulation modeling methodologies, system dynamics, health service research and disease modeling, and the modeling of human behavior in healthcare systems. Her e-mail address is s.c.brailsford@soton.ac.uk.

AUTHOR BIOGRAPHIES

NAVONIL MUSTAFEE is a research student at the Centre for Applied Simulation Modelling (CASM) in the School of Information Systems, Computing and Mathematics, Brunel University, Uxbridge, United Kingdom. His research interests are in grid computing and parallel and distributed simulation. His e-mail address is navonil.mustafee@brunel.ac.uk and his Web address is <http://people.brunel.ac.uk/~cspgnmm>.

SIMON J E TAYLOR is the co-founding Editor-in-Chief of the UK Operational Research Society's (ORS) *Journal of Simulation* and the Simulation Workshop series. He has served as the Chair of the ORS Simulation Study Group between 1996 to 2006 and was appointed Chair of ACM's Special Interest Group on Simulation (SIGSIM) in 2005. He is also the Founder and Chair of the COTS Simulation Package Interoperability Product Development Group (CSPI-PDG) under the Simulation Interoperability Standards Organization. He is a Senior Lecturer in the Centre for Applied Simulation Modelling in the School of Information Systems, Computing and Mathematics at Brunel University and a visiting Associate Professor at Nanyang Technological University. His recent work has focused on the development of standards for distributed simulation in industry. His e-mail address is simon.taylor@brunel.ac.uk.

KORINA KATSALIAKI graduated from the University of Macedonia, Greece in 2002 after attending the 4 year degree in Business Administration. She acquired her Msc in Management Sciences from the University of Southampton and is currently a Ph.D. student in the same University. Her research area is logistics in hospitals with the use of simulation. Her e-mail address is korina@soton.ac.uk.

SALLY BRAILSFORD is currently Senior Lecturer in Management Science in the School of Management at the University of Southampton. Her research interests include