

MODELLING VERY LARGE COMPLEX SYSTEMS USING DISTRIBUTED SIMULATION: A PILOT STUDY IN A HEALTHCARE SETTING

Sally Brailsford¹, Korina Katsaliaki¹, Navonil Mustafee² and Simon Taylor²

¹ School of Management,
University of Southampton
Southampton SO17 1BJ, UK
s.c.brailsford@soton.ac.uk, korina@soton.ac.uk

² Centre for Applied Simulation Modelling
School of Information Systems, Computing and Mathematics
Brunel University
Uxbridge, Middlesex, UB8 3PH, UK.
navonil.mustafee@brunel.ac.uk, simon.taylor.brunel.ac.uk

ABSTRACT:

Modern manufacturing supply chains are hugely complex and like all stochastic systems, can benefit from simulation. Unfortunately supply chain systems often result in massively large and complicated models, which even today's powerful computers cannot run efficiently. This paper presents one possible solution - distributed simulation. This pilot study is implemented in a healthcare setting, the supply chain of blood from donor to recipient.

Keywords: distributed simulation, healthcare, blood transfusion, supply chains

1. INTRODUCTION

The increase in the power of computers in recent years has enabled massive simulation models to be developed, on a scale unthinkable a decade ago. However this has only served to stimulate the demand for modelling larger and larger systems. Moreover, the internet has created the capability to run models in different locations, able to communicate with each other and to exchange data. This has stimulated research into the interoperability of models (even using different modelling software or operating systems) which can be run in different sites.

An area in which such massive models have enormous potential is that of supply chain modelling. At its simplest, a supply chain is the entire process by which a product is manufactured and sold, starting “upstream” with the procurement of raw materials and moving downstream through manufacture, assembly, distribution, sale and support. It has long been realised that the downstream material flow of product is not the only flow in a supply chain, and that there are backwards flows of information which can affect the efficiency of the process.

One classic effect of a delay in the information flow is the so-called bullwhip effect [1], whereby errors in order quantity are successively amplified, leading to huge variations in inventory. Moreover, it is now recognised that the chain functions much better if viewed as a whole, rather than by a “silo” approach where each component in the chain acts independently of its neighbours. Modern supply chain management approaches favour a global, holistic view in which the individual parts share information and trust each other, rather than simply trying to optimise their own local processes. This philosophy has been successfully applied in supply chains where the individual components are different companies, possibly in competition with each other and potentially unwilling to share commercially valuable data.

2. BACKGROUND TO THIS STUDY

This study was carried out in collaboration with the National Blood Service (NBS). The NBS consists of 15 Process, Testing and Issuing (PTI) Centres which together serve 316 hospitals across England and North Wales. Each PTI Centre thus serves around 20 hospitals. We worked particularly with the Southampton PTI Centre.

The NBS collects blood by voluntary donation, mainly from local venues such as church halls or places of employment. The blood is transported back to the nearest PTI Centre where it is tested for ABO and Rhesus grouping and infectious diseases such as HIV. It is then processed into around 120 different products, of which the main three are red blood cells (RBC), platelets and plasma. RBC have a shelf life of 35 days and platelets of 5 days, but plasma can be frozen and stored for up to a year. In this study, we only modelled RBC and platelets, which together comprise 85% of stock and are the chief source

of wastage and shortages. RBC are measured in “units” of 400ml and these form one of the basic entity types in our model. Blood products are stored in the PTI Centre’s blood bank until they are requested by the hospitals served by that Centre. There is also a nationally coordinated scheme for transferring excess stock between Centres.

The ordering system is highly complex. Local practice varies and all hospitals have slightly different ordering policies. Hospitals determine their own optimal stock levels according to their estimates of demand. An order is placed with the local PTI Centre when inventory falls below a predetermined order point, or when rare products not held in stock are requested for particular patients. Different types of order can be placed, each with different associated costs.

Individual doctors are responsible for the quantity of blood products ordered for each patient in the hospital. In theory, doctors order blood according to the Maximum Surgical Blood Ordering Schedule (MSBOS) [2] which specifies how much blood is required for a given operation. The MSBOS is conservative, to allow for cases where extra blood might be needed if complications arise, but many doctors still over-order to be on the safe side. Patients should ideally be given blood of the same type but “mismatching” is possible in emergencies – for example, O-negative blood can be given to anybody.

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, 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 £120.

3. THE SIMULATION MODEL

This system is clearly stochastic since the demand for blood is variable (even for elective surgery) depending on the type of operation and the occurrence of complications requiring extra transfusions. The supply is also variable since it relies on volunteers showing up to donate. Other organisational issues arise since the NBS manages the supply side but the hospitals manage the demand side.

Complex stochastic multi-product, multi echelon perishable inventory problems are intractable by analytic techniques [3]. Thus simulation was chosen to model this system.

The model was built in Simul8 [4]. It is described in detail elsewhere [5] and is very large and complex, requiring extensive data. Nineteen months’ data from the Southampton PTI Centre was provided and analysed 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. The entities in the model are the individual units of RBC and platelets, and the orders (representing the backwards flow of information mentioned above).

The simplest version of the model, including RBC and one medium-volume hospital, takes approximately 4 CPU minutes to run for six simulated months. The runtime rises dramatically when we tried to run the same model for one simulated year. Similarly, we were unable to run a model with more than two medium-size hospitals even in a very powerful computer. The enormous number of number of entities in the system, each of which carries many attributes, increases the computation time exponentially.

The complete model, in the case we are concerned with, should actually comprise 16 hospitals. It is clearly not feasible to run such a model in a single PC, but the use of distributed simulation allows us the possibility of running the full model.

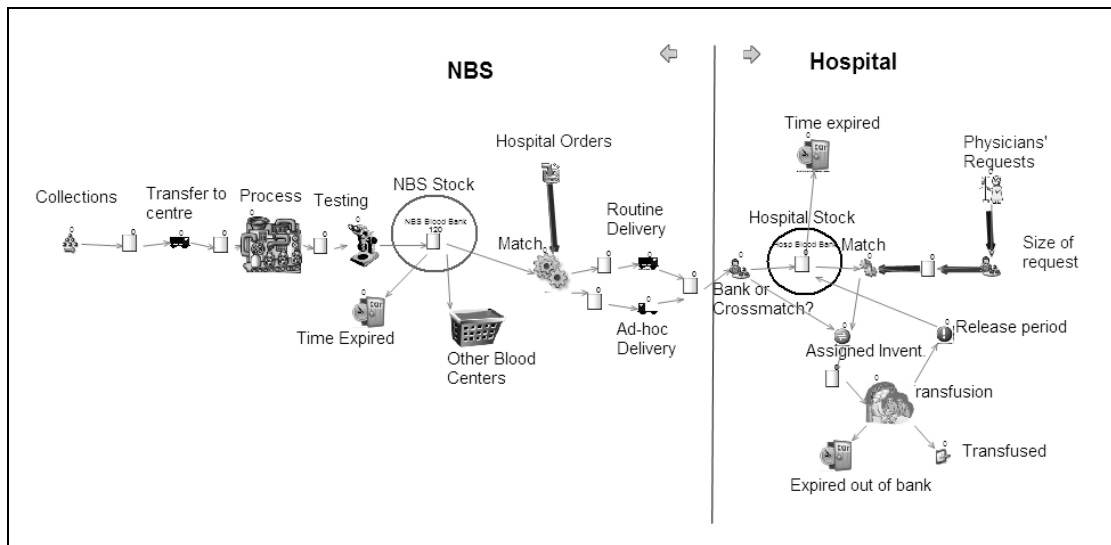


Figure 3.1. Screen shot of a simplified version of the Simul8 model, showing one hospital only

4. TECHNICAL DETAILS

Distributed simulation allows us to execute a single simulation model, perhaps composed of several distinct models, over multiple computers that are connected through a network [6]. A distributed simulation middleware is generally needed to coordinate the advancement of simulation time and for passing messages between the individual models. The Run Time Infrastructure (RTI) is soon becoming the middleware of choice for distributed simulation. This is because it implements the interface specifications which are part of the High Level Architecture (HLA), an IEEE 1516 standard for distributed simulation [7]. We have used the Defense Modeling and Simulation Office implementation of RTI (RTI 1.3-NG) to interoperate two Simul8 models.

The HLA facilitates the creation of distributed simulation (federation) by linking together new or existing simulation models (federates). All communication between federates takes place through the RTI.

The Simul8-HLA federation consists of two Simul8 CSP federates, called NBS and Hospital, and one Manager federate which coordinates the execution of the distributed simulation. Each federate runs on a different computer. The Simul8 federates are connected to the HLA-RTI through the CSP Controller Middleware (figure 4.1). Mustafee and Taylor [8] discuss the approach used to achieve this.

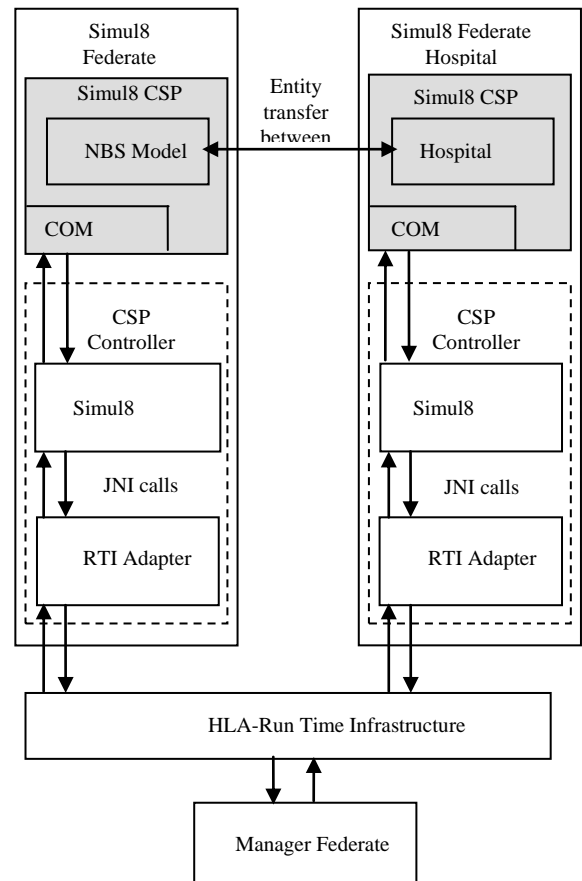


Figure 4.1. Simul8-HLA Federation Architecture

However, in order to accommodate (1) attribute transfer between the models and (2) existing model designs, this approach has been modified in the following ways:

(1) The hospital places orders for blood with the NBS at specific hours of the day. During each such interactions, entities with attribute *blood type* are transferred from the Hospital federate to the NBS federate to model the placing of orders. The incoming orders are collected in a queue “Order H1” (see figure 4.2). The NBS federate then matches its stock of blood (queue “NBS Blood Bank”) with the order received based on blood type (workstation “Match H1”). The matched blood units are then sent back to Hospital federate as entities (with the blood type attribute) to model the transfer of blood from the NBS to the hospital. The queue “From NBS” in federate Hospital receives the incoming blood units from the NBS model. Figure 4.2 shows only the relevant part of the models that are involved with exchange of entities (the actual model is substantially larger!).

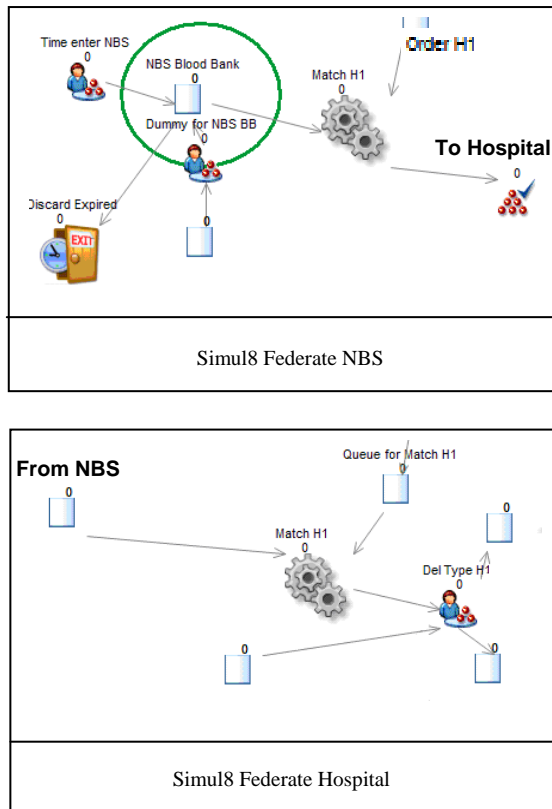


Figure 4.2. Models Simulated by Simul8 Federates

The entity exchange between these federates take the form of HLA interactions, as specified by the emerging standards of SISO’s CSPI-PDG [9]. Since we are concerned with transferring only one attribute, viz. *blood type*, it is possible to model orders being sent from the hospital and blood units being returned by the NBS simply by using a single interaction. This is achieved by constructing a String message with 7 fields and using it to populate an object of class SuppliedParameters, to be passed as an argument

to the method *sendInteraction* of class RTIambassador [10].

Each field provides information on the quantity of blood ordered or received for each of the 8 blood types. This is shown in figure 4.3 below; 0, 2, 1, 4, 0, 5, 1 and 3 units of blood have been ordered / received for blood groups O+, O-, A+, A-, B+, B-, AB+ and AB- respectively. Thus, one HLA interaction carries all the *blood type* attribute information between the Simul8 models.

| | | | | | | | |
|----|----|----|----|----|----|-----|-----|
| 0 | 2 | 1 | 4 | 0 | 5 | 1 | 3 |
| ↑ | ↑ | ↑ | ↑ | ↑ | ↑ | ↑ | ↑ |
| O+ | O- | A+ | A- | B+ | B- | AB+ | AB- |

Figure 4.3. Attribute *blood type* represented as a string

The entity exchange between these federates take the form of HLA interactions. After the interaction is received by the RTI Adapter of CSP Controller Middleware, it invokes the following method defined in Simul8 Adapter.

IntroduceEntity(time, entity)

This introduces an *entity* into Simul8 at the current simulation time + *time* (as required by the CSP). In our Simul8 model the parameter *entity* is a String datatype as discussed above. The Simul8 adapter then uses this and the Simul8 COM interface to introduce all the entities with attribute blood type into the model.

(2) The existing models have been programmed to enable Simul8 to feedback into Microsoft Excel the following: blood required by Hospital federate and blood supplied by NBS as a response to the hospital requirements. Simul8 appends this information into the files at specific times.

The CSP Controller middleware therefore has to accommodate this behaviour. This has resulted in a new **private** method being implemented in Simul8 Adapter that deals with Excel file handling. The implementation of this method is specific to this work. Likewise, implementation of some CSP Controller Middleware interfaces have also changed but their method signature, of course, remains the same.

Work is currently underway to extend this approach to accommodate multiple attributes for each entity that are passed by the method described in this paper.

5. DISCUSSION

The work described in this paper is essentially a feasibility study for adapting a standard, widely used simulation package for use in a distributed simulation model. The potential for such an approach in healthcare simulation modelling is huge. There is an increasing recognition that healthcare systems do not exist in a vacuum and that even seemingly well-defined subsystems such as emergency departments, operating theatres or out-patient clinics have complex interconnectivity with other parts of the overall healthcare system, both within the hospital and outside its walls. This can lead either to the development of enormous models which attempt to capture these relationships, or to oversimplification by ignoring them and making the model boundaries artificially narrow. In this feasibility study we have demonstrated that distributed simulation offers a viable solution to this problem, using low cost off-the-shelf software which is widely available and increasingly used in the NHS.

6. ACKNOWLEDGEMENTS

The authors thank the following people for their time and their help with data and model validation: Crispin Wickenden and Andrew Oliver from the NBS; the employees of the Southampton PTI Centre, particularly Mike Northcott; Tracey Lofting from Southampton General Hospital; Rob Hick and Judith Chapman from the Blood Stocks Management Scheme.

7. REFERENCES

- [1] H.L. Lee, V. Padmanabham and S. Whang (1997). Information Distortion in a Supply Chain: The Bullwhip Effect, *Management Science*, 43, 4: 546-558.
- [2] BCSH Blood Transfusion Task Force (1996). Guidelines for pre-transfusion compatibility procedures in blood transfusion laboratories, *Transfusion Medicine*, 6: 273-283.
- [3] M.A. Cohen and W.P. Pierskalla, (1979a). Target Inventory Levels for a Hospital Blood Bank or a Decentralised Regional Blood Banking System. *Transfusion*, 19,4: 444-454.
- [4] www.simul8.com, accessed 14.10.05.
- [5] K. Katsaliaki and S.C. Brailsford, Using Simulation to Improve the U.K. Blood Supply Chain, submitted to *Journal of the Operational Research Society*, 2005.
- [6] R.M. Fujimoto (2001). Parallel and Distributed Simulation Systems. In: *Proceedings of 33rd Conference on Winter Simulation*, December 09 – 12, Arlington, Virginia. Winter Simulation Conference. IEEE Computer Society, Washington, DC, 147-157.
- [7] IEEE 1516, "IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA)," New York, NY: Institute of Electrical and Electronics Engineers, 2000.
- [8] N. Mustafee and S.J. Taylor, Investigating Distributed Simulation with COTS Simulation Packages: Experiences with Simul8 and the HLA, paper submitted to this conference.

[9] S.J.E. Taylor, S.J. Turner and M.Y.H. Low (2005). The COTS Simulation Interoperability Product Development Group. In: Proceedings of the 2005 European Simulation Interoperability Workshop. Simulation Interoperability Standards Organization, Institute for Simulation and Training, Florida, 05E-SIW-056.

[10] US Department of Defense. High Level Architecture Run-Time Infrastructure RTI 1.3-Next Generation Programmer's Guide. September 1999.

development of standards for distributed simulation in industry.

AUTHOR BIOGRAPHIES

SALLY BRAILSFORD is currently Senior Lecturer in Management Science in the School of Management at the University of Southampton. Her research interests include simulation modelling methodologies, system dynamics, health service research and disease modelling, and the modelling of human behaviour in healthcare systems.

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 PhD student in the same University. Her research area is logistics in hospitals with the use of simulation.

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.

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 since 1996 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