

DISTRIBUTED SIMULATION AND INDUSTRY: POTENTIALS AND PITFALLS

Simon J. E. Taylor (Convener)

Centre for Applied Simulation Modelling
Department of Information Systems and Computing
Brunel University
Uxbridge, Middx UB8 3PH, U.K.

Agostino Bruzzone

DIP University of Genoa
via Opera Pia 15
16145 Genova, ITALY

Richard Fujimoto

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280, U.S.A.

Boon Ping Gan

Singapore Institute of Manufacturing Technology
71 Nanyang Drive
638075 SINGAPORE

Steffen Straßburger

DaimlerChrysler
Research & Technology (RIC/EP)
P.O. Box 2360
89013 Ulm, GERMANY

Ray J. Paul ("Discombobulator")

Centre for Applied Simulation Modelling
Department of Information Systems and Computing
Brunel University
Uxbridge, Middx UB8 3PH, U.K.

ABSTRACT

This panel paper presents the views of five researchers and practitioners of *distributed simulation*. Collectively we attempt to address what the implications of distributed simulation are for industry. It is hoped that the views contained herein, and the presentations made by the panelists at the 2002 Winter Simulation Conference will raise awareness and stimulate further discussion on the application of distributed simulation methods and technology in an area that is yet to benefit from the arguable economic benefits that this technique promises.

1 INTRODUCTION

In March 2002, the UK Operational Research Society's Simulation Study Group and the GROUPSIM Network (a forum for innovation in collaborative simulation modelling techniques) held a well attended two-day workshop on innovation in simulation modelling (Eldabi, et al. 2002). Contributions were made from practitioners of simulation in business, defense, manufacturing, and the service industries. It was noted that while several of the defense papers addressed *distributed simulation* as part and parcel of simulation modelling methodology, none of the non-

defense papers made any mention of this technique. Experience has shown that although there are practitioners of distributed simulation in the non-defense industries, there is not a widespread "distribution" of this knowledge across the community. This panel has been convened to investigate why this is the case and each panelist has been approached to share their views on the following basis.

"Distributed simulation, as exemplified by the High Level Architecture (HLA), has had a widespread impact in the defense industry for several years and has arguably led to cost savings through model reuse and interoperability. More recently, there have been several attempts to pass on these cost savings to the industrial modeling and simulation community. This panel has been convened to outline the latest research and commercialization efforts in distributed simulation and to discuss the ultimate potential and pitfalls of this work. Each panelist will present their views on the impact of distributed simulation on commercial modeling and simulation and attempt to answer the question, "Will distributed simulation ever become useful to industry or is it just a pipe dream?"

2 POSITION STATEMENT BY AGOSTINO BRUZZONE

In the last year the industrial production sector, as well as economic-productive activities in general, has been affected by significant changes. The evolution of globalization requires that we quickly review (and update) many current production strategies in order to be competitive and profitable. Today production requires a high level of integration: integration in the market, in the supply chain, in supplier networks, and in finance. For example, companies that in the past hid their corporate data are today requested to share, at least partially, this data in order to be able to access new business areas and partners (one example of this is WILD, the Web Integrated Logistics Designer, a project sponsored by Italian Research National Department involving six partner universities and seven companies with Piaggio Aeroindustries as the industrial main contractor and DIP University of Genoa as Overall General Coordinator).

In this framework production facilities need to be represented at different levels by a large number of specialized entities able to generate an effective interaction (i.e. Progetto ICARO, Integrated Computer Aided Request for an Offer, ricerca sponsorizzata da Nuova Magrini Galileo, Gruppo VaTech). The corresponding decomposition and re-aggregation phases engendered by this multi-level representation needs coordination and the development of a strategically managed supply chain that includes both material requirement support and strategic directions of each organization. This situation becomes very evident in relation to general good mass production. While in the past it was necessary for a production facility to start economic activity, today the investors are offered with a large set of production alternatives in terms of virtual sites based on outsourcing. These practices concentrates economic activity in the high value operations and supports the aggregation of low level activities in high quality specialized entities. This new course moves the production sector into a multi-level network of subjects that actively cooperate in order to reach their common targets, while competing against others. With respect to distributed production realities that already exist, new network architectures are not necessarily based on general common objectives (i.e. production of a new specific goods or service), while they are sharing just sub-level goals that become significant just through supply chain management. It is evident that with this context logistics become the starting point for a general re-engineering of production systems as well as the critical element for supporting competitiveness and profitability of companies. In effect the supply chain management is probably one of the most important challenges of the next ten years of the production sector. Therefore this management is to be considered multi-disciplinary in a wide sense. In effect it includes production systems, as well as micro and macro economic phenomena and eco-

omic-legal aspects, modeling and analysis. This subject has been subjected to traditional studies during last 50 years, however today additional research is needed to integrate such different models in order to provide production system managers with complete methodologies and tools devoted to strategic, tactical and operative planning. Based on a detailed analysis of the national and international production sector new projects are expected to be able to identify a set of indicators devoted to the correlation of the different actors, using neural meta-models, in a supply chain framework modeling their interactions.

The creation of models for each supply chain actor allows the development of a technological demonstrator devoted to supporting managers in the identification of production sites, selection of commercial and productive partners and suppliers, as well as selection of a basis for agreement. This approach allows the possibility of effectively designing the logistics network. The use of Artificial Intelligence Techniques is to be considered fundamental during this phase for the creation of a DSS (decision support system) supporting the decision makers in identifying most convenient solutions (i.e. SGS-AMT, Sistema gestione Scorte, ricerca finanziato al DIP Genova da parte di AMT Genova).

So the final objective is the development of new generation simulators devoted to the quantification of future scenarios and the detailed evaluation of the reliability and robustness of different alternatives. This will lead to the construction of federations of simulators able to navigate complex scenarios from the view point of each supply chain actor, highlighting rules and providing the users-decision makers with the possibility of changing rules, agreements and network structure.

3 POSITION STATEMENT OF RICHARD FUJIMOTO

Interoperable distributed simulations have had a considerable impact in the defense industry. However, to date, this technology has not been widely embraced for non-military applications. Both technical and non-technical issues must be resolved before we will see widespread adoption.

Perhaps most importantly, there must be a clear business case for exploitation of this technology. The "light at the end of the tunnel" must be sufficiently compelling to justify the necessary expenditures in technology and standards development.

Embedded computing offers a rich area of growth for the modeling and simulation industry. An estimated 8.5 billion computers were manufactured in 2001, of which only 200 million, less than 3%, were used in traditional platforms such as personal computers and servers (Tennenhouse 2001). The vast majority of these computers were embedded into equipment, machines, and appliances found in homes, workplaces, automobiles, and even on individuals.

Inexpensive, ubiquitous networking will make distributed systems of embedded computers commonplace. Embedded computing therefore represents a vast area of growth for the modeling and simulation industry.

Simulation can play a major role in these distributed embedded systems. Embedded computers have long been used to create “smart” devices. Simulation can be used to add a predictive capability, enabling devices to anticipate and adapt to future events. *Distributed* simulations constructed by composing simulators of individual elements enable the realization of adaptive systems of such smart devices. These networked devices will sense and interact with the physical world, project future system states, and adapt the system to operate at peak efficiency in dynamically changing environments.

Several technologies must come into play in order to create such embedded distributed simulations. First, embedded simulations must execute on devices with severe power constraints, so power efficient execution will be important. Network bandwidth is limited, and connectivity with other devices may be sporadic and unreliable, especially in mobile settings. Modeling and simulation processes such as data collection, generation of experiments, and output analysis must all be automated since human intervention to manage the simulation is clearly impossible.

Interoperability issues must be addressed along multiple fronts. Simulators must interoperate with data collection and analysis tools, among others. These systems will include components from different manufacturers, raising simulator interoperability issues. Further, even among elements from the same manufacturer, the system must be able to accommodate different versions of components and simulators as new elements are added and old components removed or replaced.

One example illustrating interoperable distributed simulations to manage and control a complex system is described in (Perumalla, et al. 2002). This case study applied distributed simulation methods and interoperability techniques to network simulations for simulation-based on-line control of a military communication network. The on-line simulations model actual networks, including wired ship-board sub-networks connected via satellite links, and wireless mobile devices. The modeled scenario depicted the communication requirements of an amphibious landing where a complex network connects troops ashore and naval vessels. The simulations use a heterogeneous set of tools, including a well-known network simulator ns2 (ns2 2002), and a wireless network simulator called GloMoSim (Zeng, et al. 1998). Though applied to a military network, the same techniques could similarly be applied in non-military settings. While far from the embedded distributed simulations envisioned here, this work does demonstrate the value distributed simulations can bring to manage systems of interoperating components.

4 POSITION STATEMENT OF BOON PING GAN

Promoting a new technology to the industry is a challenging task. To have it adopted, industry must be convinced with the benefits that the technology has to offer. Most importantly, these benefits have to be translated to monetary terms, as adoption of technology has to be justified through the Return of Investment (ROI). Translating the benefits to ROI is a challenge by itself. Involving industry partners to perform a feasibility study in the process of adoption is thus mandated in ensuring the viability of the technology. Also, the implementation of the technology has to be made transparent from the user’s point of view. But is realizing total user transparency possible, and what kind of compromises need to be made?

Distributed simulation technology offers tremendous potentials to industry, particularly in the area of supply chain simulation. Traditionally, a single centralized model that encompasses multiple companies in the supply chain is built when optimization of the chain through simulation is required. But this is not feasible when some partners are not willing to share sensitive/critical data of their business/operation. With the introduction of distributed simulation technology, this is no longer an obstacle as models sit within the company, and data are well encapsulated. Further, models are built by each company that understands its own operation best. Models can be built rapidly and maintained (more) easily as well as individual company is in-charge of their own model. The only issue that the companies need to resolve is the communication (data and objects to be shared) among their models.

Existing models can be reused if the High Level Architecture (HLA) is adopted as the standard for distributed simulation. This makes the technology even more attractive as reusing existing simulation models save tremendous development time and cost. When a model is reused, it does not mean that the model is used as of their existing form. But rather, these existing models are extended/shrunk to meet the objectives of the distributed simulation. Further, HLA also enables models to be built without any restriction on which simulation languages/packages to use as it enables interoperability of simulation languages/packages. Companies can choose the simulation language/package that suits their needs best and one that they are most comfortable with.

Though the technology offers tremendous potentials as discussed, cost savings depend very much on each particular case. Further, to make the distributed simulation concept totally transparent from the user remains a very challenging task. Particularly in the area of time synchronization of simulation federates. Without having the user to understand the basic concept of time synchronization, distributed models built might be highly inefficient. But letting the user to intervene in the performance fine-tuning will make the tech-

nology more difficult to be used. Some form of compromise has to be achieved on this aspect. Reusing of model is not a simple plug-and-play either as models are built for specific simulation objectives. Modification to the existing model is unavoidable. It is also common that companies in a supply chain might want to share different granularity of data with different partner. This issue is not currently addressed by the HLA standard. Lastly, the question of validating distributed simulation models, which involved multiple companies, needs to be addressed.

In short, distributed simulation technology offers tremendous potentials to the industry. Involving industry partners in the development of the technology and making the technology easy to use will be critical to the success of the technology adoption.

5 POSITION STATEMENT OF STEFFEN STRASSBURGER

Interoperability and reusability of simulation models are concerns which are addressed by the IEEE standard HLA. In the military simulation domain HLA is accepted since it is a mandatory standard in many cases. In the civilian simulation community the situation is different (Straßburger 2001): Although the academic scientific community has embraced HLA as the needed standard for distributed simulation, industrial (non-military) clients are only slowly becoming interested in HLA and are still far away from applying it in their day-to-day business.

More severely, many simulation tool developers (e.g., German market leaders Delmia with QUEST and IGRIP, Tecnomatix with eM-Plant and eM-Workplace) have a somewhat reluctant position towards introducing HLA interfaces into their systems. One reason might be that an HLA interface would enable their customers to let their models talk to models developed in tools from a competitor. Thus the customer is no longer bound to use simulation tools of one vendor. To force vendors to integrate HLA interfaces into their systems a strong pressure from large industrial customers is needed.

On the other hand, industrial customers often fail to recognize the value of a standard like HLA. The advantage of HLA as a base technology for the digital factory and virtual manufacturing has not been discovered by many of the potential clients. This is an important problem, because without the demand from the "paying customer" HLA will not gain the deserved impact. For educating potential clients it is therefore necessary to develop practical applications ('killer-applications') which demonstrate the potential of HLA in the daily business. It is not sufficient to have academic examples, since they are often considered as coming from the ivory tower and having no practical relevance.

For the automotive industry HLA could constitute a base technology for the Digital Factory (DF). The general

aim of the DF is that no production facility shall be planned, constructed, and operated without complete prior digital coverage. This aim, of course, requires a large amount of simulation. Since a DF requires many participants to cooperate, distributed simulation can be regarded as one key technology of bringing different independently developed simulation models together. Furthermore distributed simulation can help to reduce time and cost of developing simulations since a simulation can often be constructed by combining existing simulation components.

In summary, the algorithms and technologies in the area of parallel and distributed simulation have reached a very mature level and have led to the development of HLA. Although some standardization issues remain unsolved (Straßburger, et al. 2002) HLA can be considered as state-of-the-art in distributed simulation. Yet, HLA as IEEE standard fails to have the deserved impact in many manufacturing applications and industrial applications in general. One party working on the communication and demonstration of the advantages of HLA is the Daimler-Chrysler Research Center in Ulm, Germany.

6 POSITION STATEMENT OF SIMON J. E. TAYLOR

What is the potential of Distributed Simulation and Industry? There appear to be many diverse applications such as embedded systems, collaborative working, virtual factories, and supply chain simulation. However, for this panel let us consider the simple case of two companies wishing to establish a supply chain. We shall assume that each company has a model of its production processes (although simulation modelling is not a widespread enough technique to always guarantee this) and that each model runs in a COTS simulation package (such as Arena, Simul8, Witness, etc.) How might distributed simulation help the development of this supply chain?

In a supply chain, at some point, a commercial relationship must be formally recognized by collaborating companies. In a simple form this might be an agreement that one partner will supply an amount of material to the other in a particular quantity at a particular frequency. Failure to do this would result in penalty fines being incurred by the supplying company. If distributed simulation allows us to interoperate the two production models (i.e. join them together to simulate the production of the supply chain), then distributed simulation becomes a powerful technique that allows two (or more) companies to explore their relationships and, for example, the implications of changing order portfolios against contractual supply arrangements.

This presentation seems simple enough and suggests the question "Why are we not using distributed simulation?" One possible answer is that the kind of production models that are required for supply chain analysis are not widespread enough to make distributed simulation com-

monplace. This may be true but many multinational companies do have such models and moves are being made by such companies to investigate the application of distributed simulation (as shown by the authors of this paper and by contributions made to this Conference). This shows demand for the technique and some COTS simulation packages even advertise some distributed simulation facilities to make this possible.

Let us consider another possibility why distributed simulation is not widespread in industry. The High Level Architecture (HLA) is considered to be the defacto standard for distributed simulation. It comes in three parts: a set of rules, an Object Model Template (OMT), and a Federate Interface Specification used to interface the federates to each other via a runtime infrastructure (RTI). Elements of the OMT are used to describe the public interface of a federate in terms of object classes (and attributes) and interactions (and parameters). These are published through the federate interface and the RTI is responsible for the timely exchange of this information between the federates. In terms of the COTS simulation packages, a federate is the package and the model that it contains. The OMT is used to specify the data exchanged between the model/packages via the RTI. At the time of writing there is no "standard" approach to the exchange of data between federates. Without this standard, how can we have universal interoperability between our COTS simulation packages? Further, given the relatively simple data exchange requirements of supply chain simulation, is the HLA appropriate at all? In many cases this data exchange is nothing more than a simple timestamped entity with relatively few attributes – a far cry from the massive data exchange requirements of real-time distributed simulation used in the military. I argue that the poor understanding of the needs of industrial simulation modelling practitioners by HLA technologists, and the complexity of the HLA itself, are the major reasons for the almost non-existence of distributed simulation in industry. It might be argued that a simpler alternative is better, one that can be generally understood by practitioners and technologists and can engender understanding between communities. Taylor, et al. (2002) put forward an approach that should be familiar to many from the parallel and distributed simulation community and, it is hoped, conveys the message that sometimes simplicity may be best (especially in terms of validation!)

In conclusion, I consider distributed simulation to hold much potential for industry in the development of new simulation modelling applications and methodologies, and I am sure that simulation modellers will innovate in ways that will surprise! However, the reliance on the relatively complex HLA, the lack of standard, general approaches to the use of this technology for the types of simulation described above, and the lack of cross discipline experts may well be a barrier to the widespread adoption of distributed simulation in industry. In short, if simulation is the hole

and the HLA a peg, then I wonder how big a hammer will we eventually need to make it fit?

7 POSITION STATEMENT OF RAY J. PAUL

Does industry want distributed simulation? What does it want it for? What would it gain from it? The answer is actually not known to be positive for any of these questions.

So, a solution looking for a problem, i.e. distributed simulation is looking for something to do in industry. This is wrong as we don't appear to have a solution. The solution is a fantasy, something that we hope might happen. So, it is a fantasy solution looking for a real problem to solve in industry and not finding one yet.

Why are we wanting to provide fantasy solutions to non-existent problems? Three reasons:

- **Moore's Law.** The rate of technological speed up is exemplified by Moore's Law and creates a belief that we should be bigger and faster as the technology allows us to be. Unfortunately, software ability has always lagged behind technological improvements. So Moore's Law tells us that the rate of disillusionment with software is likely to double every 18 months.
- **The North American Unwritten Research Agreement.** Research funders largely improve their career track by being responsible for larger, rather than smaller budgets. The researchers largely improve their career tracks by attracting large grants and more often. Research funders and researchers are therefore able to offer each other complete satisfaction as long as the issue of feasibility is not discussed seriously. As Paul's Law of the IT Concept Fallacy states, "There exists a belief that any concept can be realized if only enough time and money is spent on the realization attempt." Such a belief is guaranteed to lead to the spending of much time and money."
- **Toys for Boys.** In a male dominated society, toys that are especially bigger and faster or miniaturized or shiny will always take precedence over those that work. Hence I understand that we have a "Hell of a Lousy Architecture" being prejudicially and optimistically promoted and believed in. Hence the contribution to the body of knowledge is known as *Castle Confections Built on Sand*.

I would appeal to all simulationists to heed the words of the great Kiviat who introduced to those attending the 1990 Winter Simulation Conference the concept of SINSFIT – *Simulation is No Substitute for Intelligent Thinking*. I urge readers of this panel paper to go forth and think!

ACKNOWLEDGMENTS

The convener of this panel would like to thank the contributors, John Miller, and Phil Wilsey for their hard work in making this panel possible.

REFERENCES

- Eldabi, T., Robinson, S., Taylor, S.J.E. and Wilcox, P.A. (2002). *Proceedings of the Simulation Study Group Two-Day Workshop*. The Operational Research Society, Birmingham, UK.
- ns2 (2002). The Network Simulator - ns2 homepage, <http://www.isi.edu/nsnam/ns/> [accessed July 31, 2002].
- Perumalla, K., Fujimoto, R., McLean, T., and Riley G. (2002). Experiences Applying Parallel and Interoperable Network Simulation Techniques in On-Line Simulations of Military Networks. In *Proceedings of the 2002 Workshop on Parallel and Distributed Simulation*. IEEE Computer Society. 97-106.
- Straßburger, S. (2001). *Distributed Simulation Based on the High Level Architecture in Civilian Application Domains*. Society for Computer Simulation International, Ghent, Belgium.
- Straßburger, S., A. Hamm, G. Schmidgall, and S. Haasis. (2002). *Using HLA Ownership Management in Distributed Material Flow Simulations*. In *Proceedings of the 2002 European Simulation Interoperability Workshop*. June 2002. London, UK.
- Taylor, S.J.E., Sudra, R., Tan, G., and Ladbrook, J. (2002). GRIDS-SCS: An Infrastructure for Distributed Supply Chain Simulation. *Simulation*.
- Tennenhouse, D. (2001). Opening Address, Intel Developer Forum. San Jose, CA, USA.
- Zeng, X., Bagrodia, R., and Gerla, M. (1998). GloMoSim: A Library for Parallel Simulation of Large-Scale Wireless Networks. In *Proceedings of the 1998 Workshop on Parallel and Distributed Simulation*. 154-161.

AUTHOR BIOGRAPHIES

SIMON J.E. TAYLOR is the Chair of the Simulation Study Group of the UK Operational Research Society and the collaborative simulation modelling forum, the GROUPSIM Network (www.groupsim.com). He is a Senior Lecturer in the Department of Information Systems and Computing and is a member of the Centre for Applied Simulation Modelling, both at Brunel University, UK. He was previously part of the Centre for Parallel Computing at the University of Westminster. He has an undergraduate degree in Industrial Studies (Sheffield Hallam), a M.Sc. in Computing Studies (Sheffield Hallam) and a Ph.D. in Parallel and Distributed Simulation (Leeds Metropolitan). His main research interest is collaborative simulation

modelling. He is also a member of the Purple Theatre Company. His email and web addresses are <simon.taylor@brunel.ac.uk> and <www.brunel.ac.uk/~csstsjt>.

AGOSTINO BRUZZONE began his engineering studies at the Italian Naval Academy with the Faculty of Pisa in 1984. After successfully completing this phase, he transferred to the University of Genoa where he earned his degree in Mechanical Engineering. Since 1991, he has taught "Theories and Techniques of Automatic Control" at the Technical High School of Finale Ligure and has become a member of the industrial simulation work group of Prof. Mosca and Prof. Giribone at the University of Genoa. He has utilized these techniques in the harbor terminal, maritime trade and sailboat racing sectors. He is currently working on a research project involving new modeling designs, AI techniques, and DOE (Design of Experiments); particular attention is being focused on the application of Neural Networks and Fuzzy Logic to industrial plant problems using Simulation and Chaos Theory.

RICHARD FUJIMOTO is a professor in the College of Computing at the Georgia Institute of Technology. He received the Ph.D. degree from the University of California (Berkeley) in Computer Science and Electrical Engineering in 1983, an M.S. degree from Berkeley in 1980, and B.S. degrees from the University of Illinois (Urbana) in 1977 and 1978 in Computer Science and Computer Engineering, respectively. He has been an active researcher in the parallel and distributed simulation community since 1985 and has published numerous papers on this subject. He has given several tutorials on parallel and distributed simulation at leading conferences. He has co-authored a book on parallel processing as well as a second on parallel and distributed simulation systems. He served as the technical lead in defining the time management services for the Department of Defense High Level Architecture (HLA). Fujimoto is an area editor for ACM Transactions on Modeling and Computer Simulation. He also served as chair of the steering committee for the Workshop on Parallel and Distributed Simulation (PADS) from 1990 to 1998.

BOON PING GAN is a Research Fellow with the Manufacturing Planning and Scheduling group at Singapore Institute of Manufacturing Technology (formerly known as Gintic Institute of Manufacturing Technology). He is currently leading a research project that attempts to apply distributed simulation technology for supply chain simulation. He received a Bachelor of Applied Science in Computer Engineering and Master of Applied Science from Nanyang Technological University of Singapore in 1995 and 1998 respectively. His research interests are parallel and distributed simulation, parallel programs scheduling, and applica-

tion of genetic algorithms. His email address is <bpgan@SIMTech.a-star.edu.sg>.

STEFFEN STRASSBURGER is a research assistant at the DaimlerChrysler Research Center in Ulm, Germany. He holds a PhD and a Master's degree in Computer Science from the Otto-von-Guericke University in Magdeburg, Germany. His international experience includes a one-year stay at the University of Wisconsin, Stevens Point and a stay at the Georgia Institute of Technology, Atlanta. He actively participates in several international conferences. His main research interests lie in distributed and web-based simulation and the High Level Architecture.

RAY J. PAUL is a Professor of Simulation Modelling, Director of the Centre for Applied Simulation Modelling, the creator of the Centre for Living Information Systems Thinking, and the Dean of the Faculty of Technology and Information Systems, all at Brunel University, UK. He received a B.Sc. in Mathematics, and an M.Sc. and a Ph.D. in Operational Research from Hull University. He has published widely, in books, journals and conference papers, many in the area of simulation modelling and software development. He has acted as a consultant for a variety of United Kingdom government departments, software companies, and commercial companies in the tobacco and oil industries. He is editor of the Springer-Verlag Practitioner book series. His research interests are in methods of automating the process of modelling, and the general applicability of such methods and their extensions to the wider arena of information systems. He is currently working on aspects of simulation in the social sciences, in particular health management, the new G²R³ modelling technique. His email and web addresses are <ray.paul@brunel.ac.uk> and <www.brunel.ac.uk/~csstrjrp>.