Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

AN INTEGRATED AND ADAPTIVE DECISION-SUPPORT FRAMEWORK FOR HIGH-TECH MANUFACTURING AND SERVICE NETWORKS

Peter Lendermann
Malcolm Yoke Hean Low
Boon Ping Gan
Nirupam Julka
Lai Peng Chan

D-SIMLAB Programme
Singapore Institute of Manufacturing Technology
71 Nanyang Drive, Singapore 638075
SINGAPORE

Stephen J. Turner
Wentong Cai
Xiaoguang Wang

Parallel and Distributed Computing Centre
Nanyang Technological University
Block N4, Nanyang Avenue, Singapore 639798
SINGAPORE

Loo Hay Lee

Department of Industrial and Systems Engineering
National University of Singapore
1 Engineering Drive 2, Singapore 117576
SINGAPORE

Terence Hung

Software and Computing Programme
Institute of High Performance Computing
1 Science Park Road, Singapore 117528
SINGAPORE

Simon J.E. Taylor

Centre for Applied Simulation Modelling
School of Information Systems, Computing and Maths
Brunel University, Uxbridge, Middx UB8 3PH
UNITED KINGDOM

Leon F. McGinnis

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332-0205
U.S.A.

Stephen Buckley

IBM Thomas J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, New York 10598
U.S.A.

ABSTRACT

This article describes the results of one of the ten pilot programmes under the Integrated Manufacturing and Service Systems (IMSS) initiative pursued by the Agency for Science, Technology and Research (A*STAR) in Singapore. The objective of this particular programme is to investigate how design, analysis, enhancement and implementation of critical business processes in a manufacturing and service network can be realised using one single simulation/application framework. The overall architecture of the framework outlines how commercial simulation packages and web-service based business process application components would have to be connected through a commercial application framework to achieve maximum leverage and re-usability of the applications involved. In the pilot phase of this programme, research issues were also addressed with regard to mechanisms for interoperation between commercial simulation packages, symbiotic interaction between simulation-based decision support components and physical systems, and simulation speed-up through multi-objective optimal computing budget allocation techniques on a grid infrastructure.

1 MOTIVATION

Today, a corporation's operations and business processes are subject to permanent changes. To stay competitive, adaptations and enhancements of manufacturing and service operations and the associated business processes need to take place continuously (Fine 1999). This requires advanced virtual experimentation and techniques, especially in highly capitalised industries where experimentation with the real system would be too disruptive and costly and often is not possible at all. Such techniques also have to take into account the specific characteristics of today's pull-environments in which operational execution plans are the result of a complex translation from frequently-changing customer demand into material quantities to be released into and moved within the manufacturing and logistics systems at pre-specified times (Lendermann et al. 2001). Moreover, in such fast changing business environments it becomes equally difficult to develop, validate and maintain the various models representing the manufacturing and logistics systems.

Where operation of these manufacturing and logistics systems is concerned, advanced methodologies and technologies for planning, scheduling and performance prediction in environments that are subject to high variability and non-linear dynamics with little steady state do exist but mainly for individual sites. However, contemporary high-tech companies operate in global networks that often involve contract manufacturers and third party logistics providers, driving the underlying systems towards mega-networks with large risk and reward, rapid responses and high-speed switching where relationships are formed and dissolved in the context of particular jobs. Complex inter-factory processes, possibly across enterprise boundaries, have become more critical and raised new operational challenges such as how to allocate lots in a wafer fab to customer demand that is placed to a semiconductor assembly and test facility (which typically is geographically separated from the wafer fab), or how to optimise the collaboration between a line maintenance service company for commercial aircraft and a third party logistics provider to minimise the inventory cost for critical spare components. Operation of these complex manufacturing and service networks is still mostly accomplished through myopic, task-centred approaches such as local dispatching, since most of the planning and scheduling problems associated with these networks are very difficult to solve. In addition, when it comes to experimentation with planning and scheduling systems, it turns out that the associated algorithms and procedures are generally modelled in a rather crude manner.

Lastly, adaptations and enhancements of supply network structures and business processes are often discarded at the design table as their implementation involves too much effort and the advantage would be lost by the time the changes are made. One reason is that today most business applications still follow the conventional *integrative* approach: Different subsystems (such as an ERP system and a scheduling application) are connected through point-to-point interfaces; business processes are encapsulated in each application and are constrained by templates; changes can be made only by changing settings and parameters; and each application typically has its proprietary, closed data model with redundancies to other systems. Because of these constraints it is often not possible to actually implement changes derived from virtual experimentation and analysis.

In this setting, our ultimate vision is the creation of an adaptive decision support framework that allows to represent with high fidelity all value-creation processes along a supply chain (not only operational processes such as shop-floor or warehouse operations but also business processes such as planning, order management and scheduling) in a unified business model, improve their performance in a virtual experimentation testbed, and then generate and implement with minimal effort and time the corresponding business application software for critical processes from the same unified model, thereby leveraging existing software application components as much as possible. All critical tactical decisions for supply network management will be enabled through a single IT framework. An immense reduction of the cycle time for business process design-analysis-enhancement-implementation can potentially be achieved.

On the path towards realisation of this vision, we have embarked on a new, dedicated research programme to address some of the major research issues associated with the above-mentioned challenges.

Special attention is given to the specific requirements of two sectors, namely high-tech manufacturing (in our case semiconductor) and service parts supply chain systems (in our case aerospace spare components logistics) since they are strategically important to Singapore's economy and our research team also has ongoing industry projects with prominent companies in these sectors. In their context, the term *service* refers to differentiated services to meet customers' individual needs, superb delivery performance, and rapid reaction and response to requirement changes.

For the resolution of most of the operations management issues arising from these industry projects, a primary focus on simulation techniques with optimisation-based methods playing a supporting role for specific subtasks - rather than the other way around - appears to be the more appropriate paradigm for our effort. Simulation-based techniques are not constrained by analytical simplifications and give *reasonable* solutions (i.e. potentially better than best practice observed in the respective companies) within *reasonable* time, especially after successful resolution of the research issues discussed in this article.

Note that in this research we assume that the work of the simulation modeller is supported by the use of Commercial-off-the-shelf Simulation Packages (CSPs). These include packages such as Arena, AnyLogic, Simul8, etc. (Swain 2003). Our research focuses specifically on the suite of CSPs based on AutoSched AP (Brooks 2005a) and AutoMod (Brooks 2005b) as well as WITNESS (Lanner 2005).

The remainder of this article is outlined as follows: In Section 2 we briefly review previous work and describe the major technology components that will be required for the realisation of our vision. This is followed by an overview of the results that were obtained during the 6-month pilot phase of this research programme in Section 3. An outlook to the major research issues to be addressed in the full-fledged phase of the research programme is given in Section 4. This also includes a discussion of potential collaboration opportunities.

2 TECHNOLOGY COMPONENTS

Distributed, interoperable and reusable decision support system components, comprising a combination of advanced business process application components, discrete event simulation technology and synchronisation/distribution middleware that enables their interoperation for virtual experimentation are required for the realisation of the above-mentioned vision.

Development of enabling distributed simulation technology for supply chain management has been pioneered for example by the D-SIMLAB (Decision Support for Integrated Manufacturing and Logistics in Asset-Intensive Businesses) Research Programme at Singapore Institute of Manufacturing Technology (SIMTech) in collaboration with the Parallel and Distributed Computing Centre at Nanyang Technological University (Gan et al. 2000; Turner et al. 2001) and the Keck Virtual Factory Laboratory at Georgia Institute of Technology (Lendermann et al. 2003), and also by the Manufacturing Engineering Laboratory at the National Institute of Standards and Technology (McLean and Riddick 2000) and the Simulation and Modelling Group at the University of Magdeburg (Strassburger et al. 1999).

We are using the High Level Architecture (HLA) as the mechanism for integrating federates representing supply network operations and decision-making (Kuhl et al. 1999). Currently, the HLA (IEEE standard 1516) is emerging as a standard for Plug & Play of simulation-based decision support components for manufacturing and logistics systems. An international effort led by the Centre for Applied Simulation Modelling at Brunel University, UK, that is driving the development of standards for the interoperation of simulation model components has recently been endorsed by the Simulation Interoperability Standards Organization (SISO) as a standards development group, namely the *Commercial-off-the-shelf Simulation Package Interoperability Product Development Group* (CSPI-PDG), for details refer to (Taylor et al. 2005a, Taylor et al. 2005c).

Terzi and Cavalieri (2004) provide a detailed review of the application of the High Level Architecture (HLA) for manufacturing and supply chain simulation. Specific examples of how HLA-based distributed simulation technology can be applied for decision making in Supply Chain Management are given in (Chong et al. 2004) and (Lendermann et al. 2004).

High-fidelity virtual experimentation testbeds in which planning and scheduling approaches can be investigated and tested will allow experimentation and improvements that otherwise would not be possible, reduce the risk of transition, and ultimately enable faster technology adoption. The most straightforward way to achieve an optimal representation of planning and scheduling processes in a simulation testbed is to integrate the corresponding business process application components with the simulation. However, business process application components by nature are designed to run in real time only but not in simulation time. This can be overcome by making the business process application component itself compliant with discrete event simulation (Lendermann et al. 2002): A replicated copy of the application is incorporated into the simulation model. The same piece of software can be used to run the business in the real world and also to represent the underlying procedures in the simulation world for re-engineering and optimisation of the overall business performance. Synchronisation of real time and simulation time can also be achieved by using the Runtime Infrastructure (RTI) of the High Level Architecture. A manufacturing application of this has been showcased by Julka et al. (2004).

The degree of flexibility as required for the realisation of our ultimate vision can be achieved only through the use of web-service based business process application components. The Service Oriented Architecture (SOA) has been proposed to address this need (Endrei et al. 2004), and web-services are one of the approaches to implement SOA (Gottschalk et al. 2002). Web services allow easy configuration of components to create dynamic applications with flexible data models that can also leverage existing software systems as much as possible (which allows gradual implementation and helps to avoid *big-bang* approaches for going live). This is intended to facilitate a significant reduction of the overall time and effort for the creation of a software application from a business model as well as for reconfiguration and implementation of the application every time the model changes.

A simulation testbed for virtual experimentation that is totally integrated with a business application can be further developed into a symbiotic system that interacts with the business application in a mutually beneficial way

(Dagstuhl 2002). It is highly adaptive in a sense that the simulation system not only performs *what-if* experiments that are used to enhance the physical system, but also accepts and responds to data from the physical system. The physical system benefits from the optimised performance that is obtained from the analysis of simulation experiments. In turn, the simulation system benefits from the continuous supply of the latest input data and the automatic validation of its outputs. To achieve this, software agent technology can play a key role in monitoring, coordination and control of symbiotic simulation systems.

3 RESEARCH ACHIEVEMENTS

The achievements resulting from the 6-month pilot phase of this programme are summarised in the following sections.

3.1 Description and Methodologies

Arising from the above-described vision, the following specific research challenges were addressed during the pilot phase of this programme.

3.1.1 Standardised Representation of Entities across CSP-Based Simulation Models for Reference Model Definition

This involves (i) the development of a generic interface between CSPs and the HLA RTI, making use of modelling standards currently being developed within the CSPI-PDG, (ii) a specific implementation for AutoSched AP, a commonly used CSP for wafer fabs, and (iii) experimentation with different interoperation mechanisms based on inter fab material flow scenarios.

3.1.2 Use of a Computer Grid Infrastructure for Complex Simulation Analysis Tasks

This involves the prototyping of a distributed grid-based virtual infrastructure for concurrent execution of simulation scenarios.

3.1.3 Synergetic Interaction of a Simulation Model with the Underlying Physical System

This involves (i) exploration of the use of a *symbiotic simulation system* to improve the maintenance and adaptation of manufacturing network simulation models and provide decision support to manage changes in the system, and the investigation of how software agent technology can play a key role in monitoring, coordination and control, and (ii) demonstration of feasibility through a case study of a semiconductor backend manufacturing system.

3.1.4 Enhancement of Simulation Analysis Tasks through Optimisation Techniques

This involves (i) investigation of Optimal Computing Budget Allocation techniques to speed up simulation optimisation through a simulation case study based on an ongoing project in the aerospace spare component logistics domain, and (ii) subsequent analysis of two specific problems: determining optimal inventory levels and finding inventory policies for differentiated service levels. Given a set of design alternatives with multiple performance measures, the problem is to find the non-dominated set of designs by running simulation experiments.

3.1.5 Re-Use of Business Applications in a Simulation Environment through a commercially established Application Framework

This comprises development of a framework architecture that integrates simulation model components and web-service based business process application components through the IBM Websphere Business Integration (WBI) suite and definition of a sample process based on an aerospace spare component allocation process.

3.2 Results

The detailed results are summarised in the following sections. More details can be found in the respective references given in each section.

3.2.1 Synchronised Interoperation of Simulation Model Components

Using a defined set of interfaces, CSPI-PDG reference model type I (asynchronous entity transfer) has been implemented on a COTS Simulation Package Emulator (CSPE). The CSPE is designed to investigate the interfaces between CSPs and the HLA RTI and to benchmark alternative interoperability solutions. By adding new features, the CSPE is not only able to support the creation of standalone models (as provided by current CSPs) but also distributed models. Further details, including an evaluation of the CSPE, can be found in (Wang et al. 2005b; Taylor et al. 2005b).

For wafer fab models that previously had been investigated by SIMTech and Chartered Semiconductor Manufacturing using a simple C++ simulator (Lendermann et al. 2004), interoperation has now been achieved for the corresponding AutoSched AP models. The interoperation was implemented using a middleware approach that does not require the modellers to be involved in the technical details of time synchronisation and information exchange among simulation federates (interoperability). Modellers can still build their model in the same way as they have been doing.

The only additional task that they need to do now is to define the situation at which lots are moved from one factory to another. The rest (how information is sent and received, how time is synchronised) is transparent from the modellers. Further details are described in (Gan et al. 2005). Through this experience, we have identified features that are required to realise Plug & Play simulation for other simulation packages such as ProModel (ProModel 2005), WITNESS, and Arena.

Optimistic synchronisation compliant with CSPI-PDG emerging standards has been enabled through a middleware approach. While it is easier to apply conservative synchronisation for the integration of the CSP and the HLA, the conservative approach is heavily dependent on the lookahead value and leads to poor performance with small or zero lookahead. Conversely, the optimistic approach is less constrained and can exploit parallelism in situations where causality errors may occur but in fact seldom occur. To release the modeller from handling the complex rollback procedure, a rollback controller is introduced to perform these tasks on behalf of the simulation model. It provides the possibility to extend the interface for the interoperation of CSPs to include optimistic synchronisation. Further details including the results of performance studies that were conducted are described in (Wang et al. 2005a).

3.2.2 Concurrent execution of virtual experimentation scenarios on a Grid

A generic infrastructure that exploits the use of Grid Computing techniques for both symbiotic simulations and simulation optimisation has been designed. A software prototype that integrates with simulation optimisation uses the Globus Toolkit (Globus 2005) as the core middleware. A thin, generic and configurable client allows users to carry out parameter-sweep type simulations. A server component handles resource discovery and scheduling of jobs on the grid infrastructure. The client can be plugged with an optimisation engine on the user side and can post scenarios based on previous data. The framework is generic enough to support other discrete event simulation applications. Details are described in (Julka et al. 2005).

3.2.3 Symbiotic interaction of a decision support application and a physical system

An interface between the Java JADE agent system (JADE 2005) and a WITNESS simulation model has been developed and concurrent creation and execution of multiple WITNESS models has been enabled. A prototype of a symbiotic simulation system (where the real system is represented by a second simulation model) has been developed. A simple semiconductor backend ontology has been defined and is used for the communication between a monitoring agent and a simulation optimisation agent.

The feasibility of symbiotic simulation optimisation has been demonstrated through a mechanism that optimises outsourcing decisions in a highly loaded semiconductor backend manufacturing system. Experimental results from the prototype show that the symbiotic integration of simulation-based optimisation with the real system can reduce the operational cost of the real system by controlling the degree of outsourcing. A software agent that monitors the real system constantly can trigger simulation-based optimisation and use the optimisation result as control parameters onto the real system automatically. However, the effectiveness of the control depends on how fast the optimisation result is obtained as well as the level of detail used in each simulation experiment. This is described in more detail in (Low et al. 2005).

3.2.4 Enhancement of simulation analysis through optimisation techniques

A multi-objective optimal computing budget allocation (MOCBA) algorithm has been developed to allocate simulation replications to the designs so that the non-dominated set of designs can be found with high confidence at the least expense in terms of simulation replications. This includes a performance index to measure how non-dominated a design is; two types of errors to measure the quality of the selected Pareto set; asymptotic simulation replications allocation rules derived based on the Lagrangean relaxation method; and a sequential procedure to allocate the simulation replications. The MOCBA algorithm has also been applied to two case study problems: Planning aircraft spare parts inventory among airports, and the differentiated service inventory problem. In both the case studies, interfaces between the MOCBA algorithm and the simulation models have been developed, which are generic enough to be applied to any other simulation models. The numerical results show that the MOCBA is able to provide a speedup of close of 5 times for running simulations. Further details can be found in (Lee et al. 2005).

3.2.5 Overall framework architecture

The framework shown in Figure 1 makes use of the IBM Websphere Business Integration (WBI) suite as a backbone. It comprises of a discrete-event simulation component and a middleware to bridge the business/operation logic to the simulation. It capitalises on WBI's capability to integrate and monitor disparate systems easily. Through the use of the WBI Message Broker as the middleware, the framework can be extended to incorporate other software systems and/or components (e.g. Web service components, backend applications, etc.) and eliminates hard-wiring at each point-to-point interface. The target business application should preferably consist of decoupled atomic business processes (e.g. Web service-components with self-contained business logic) with configurable business process flows such that simulation studies can be performed with ease by reconfiguring, adding or removing processes.

The framework not only serves as an integrated decision support tool for businesses, it can also be combined with the WBI Monitor to transform the existing business application into a comprehensive, continuous process improvement framework. In this way, the business application becomes a *living* system and evolves and grows with the business needs. The existing business/operational logic can be re-used for simulation, and new processes explored through simulation can be deployed directly back onto the live system.

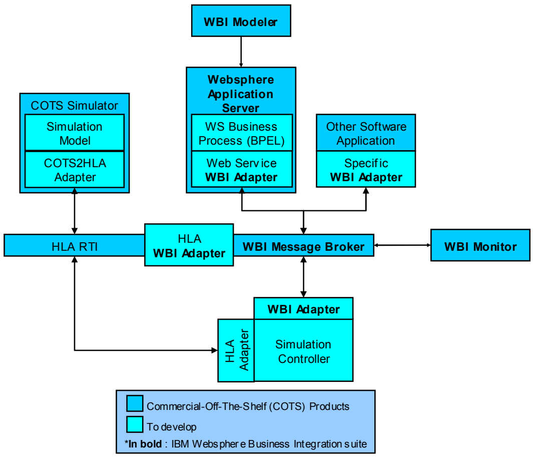

Figure 1: Framework Architecture

The role of monitoring software such as the WBI Monitor in continuous process improvement warrants increased attention at several levels. First, simulation models are based on statistical estimates of live system parameters, such as cycle time and quality. Monitoring software can continuously measure live system parameters and update relevant simulation models accordingly. Second, a business application is typically designed to achieve a specified set of performance targets. Monitoring software can detect shortcomings in live system performance metrics and trigger adaptive improvement activities such as an optimisation algorithm or a symbiotic simulation optimisation as described earlier in this article. Third, optimisation algorithms are typically based on simplifying assumptions about the live system, for reasons of efficiency. Monitoring software can check whether such assumptions continue to hold. Fourth, using business intelligence techniques, monitoring systems can detect new trends such as a change in customer preferences (Kapoor et al. 2005). Trend detection can trigger proactive business application redesign.

3.3 Novelty

The following elements of the achievements during the pilot phase of this programme can be considered novel:

- Plug & Play of simulation model components developed with a COTS simulation package and their efficient interoperation in a standardised manner.
- Application of the Grid Computing technology to execute simulation models based on programming languages such as C++ and Java and seamless integration with optimisation algorithms, in effect making the entire process of simulation experimentation automated, and making the solution scalable to support more complex and time-consuming simulations.
- Use of an agent-based approach for updating, validating and tuning of a simulation system with regard to a real world system based on simulation-based optimisation.
- Enhancement of established Optimal Computing Budget Allocation methods to solve multi-objective optimisation problems.
- A framework architecture enabling maximum reusability of decision making application components in a virtual experimentation environment.

4 FUTURE CHALLENGES

Many important research issues have not yet been addressed during the pilot phase of this programme; however, it is our intention to investigate them in more detail in our future work.

4.1 CSP Interoperation

Where interoperation of COTS simulation packages is concerned, major challenges arise from the fact that different CSPs incorporate different ways to advance time. At the same time, most CSPs do not provide extensions or APIs to access internal states of the simulation needed to support interoperability.

To facilitate this, the modelling process has to be refined to the extent that intervention on the interoperation layer is not required any more, and high level tools need to be developed. To enable execution of realistic, large-scale material and information flow scenarios, more efficient synchronisation mechanisms (e.g., through more advanced optimistic protocols) for the interoperation of CSPs will also be investigated.

As part of the CSPI PDG effort, standards will also be established for interactions involving type II-VI reference models (Synchronous Entity Passing with Bounded Buffers, Shared Resources, Shared Events, Shared Data Structures, Shared Conveyors) as well as for the required data exchange specifications and the corresponding interoperability frameworks.

4.2 Grid Simulation

The present implementation of the grid infrastructure for virtual experimentation can be used only by the optimiser because the optimiser and associated optimisation programs are written in C++ and can be ported to the Linux platform. Cross-platform execution of programs which is essentially required in the case of the symbiotic simulation system (the symbiotic system is developed on Windows and uses CSPs available for Windows only) is still unavailable. Also, gridifying CSPs which are executed through interactive graphical user interfaces will be revisited since human interaction needs to be appropriately represented in the simulation. In future, certain standards need to be developed for vendors to include in the packages for easy integration with a grid infrastructure. Also, appropriate methods for collaborative sharing of distributed data, data/functional decomposition, dynamic resource discovery, decentralised scheduling and task management will be investigated in more detail.

4.3 Symbiotic Simulation

In the area of symbiotic simulation the most important research issues to be addressed can be classified into two areas.

1. Target metrics will be identified based on specific industrial domains, comprising both metrics that measure divergence of a live business process from the simulation model and metrics that define successful improvement of business processes, as well as their interactions with each other. How the respective metrics can be measured and monitored, how frequently to do this, how to filter and analyse the available data, what (additional) data and resources are required, what kind of statistical analysis might be needed, will be investigated as well. The conditions under which improvement in metrics achieved in the simulation analysis may not be observed in the subsequently implemented live process will also be studied. This includes investigation of methods to determine whether such discrepancies are due to insufficient fidelity of the simulation model used for the analysis or changes in the real system occurring before the results of the *what-if* analysis could be implemented. Solutions to how such discrepancies could be overcome will be proposed.
2. Where process improvement methodologies are concerned, we will investigate to what extent a *Continuous Process Improvement* cycle (monitoring, model adaptation and validation, simulation analysis and implementation) needs to be interactive or can be automated. The question how DOE (Design of Experiments) methods and simulation optimisation techniques can help maximise the efficiency of the simulation analysis part will also be addressed.

4.4 Simulation Optimisation

The current MOCBA method has been developed based on the assumption that the set of possible alternative designs is finite and relatively small. When the set of possible designs is infinite or finite but very large, a more sophisticated search procedure (using techniques such as Genetic Algorithms, Simulated Annealing, Nested Partitions, etc.) is required to explore the design space to find the promising design alternatives. This search method needs to be combined with the MOCBA technique or other ranking and selection procedures to find the non-dominated set of designs. In addition, more work needs to be done regarding interfaces and communications between the MOCBA procedure and the simulation packages in a grid computing environment in the real implementation phase.

4.5 Overall Framework Realisation

For the realisation of the overall framework specific issues will be addressed regarding (i) time synchronisation and management as well as state sharing/synchronisation between the simulation and the business process application components, (ii) load sharing and mechanisms to reduce latency between the simulation and the business process application components, (iii) fault handling, notification and recovery between the simulation and the business process application components, and (iv) feasibility of using WBI Monitor for symbiotic simulation systems.

4.6 Standards

Finally, the Plug & Play requirements for the design of today's manufacturing and logistics systems raise the question how standards can play their part to facilitate this: Many initiatives such as MDA™ (Model-Driven Architecture), RM-ODP (Reference Model for Open Distributed Processing), DEVS (Discrete Event Systems Specification), OASIS (Organization for the Advancement of Structured Information Standards) and CSPI-PDG (see Section 2) have been and are being pursued to develop the required *information technology* standards. But standardisation also needs to takes place on the *application* level that would result in archived, re-usable simulation model components that require much less customisation effort.

4.7 Collaboration Opportunities

Where industry collaboration is concerned each of the research issues addressed in this programme - as explained in Section 1 - has been directly related to potential application scenarios arising from industry projects with prominent companies in Singapore:

Interoperation of AutoSched AP wafer fab models is currently under discussion with two semiconductor foundries. A roadmap towards a symbiotic simulation system is being developed with another semiconductor foundry. Integration of business process application components with a simulation package will be applied for the first time in the development of a Virtual Warehouse Management system for aerospace spare component management which will also require a grid infrastructure for concurrent simulation analysis. The grid infrastructure can also be used for simulation-intensive analysis exercises in the semiconductor domain. The MOCBA techniques are applicable for simulation analysis exercises in both the semiconductor and the aerospace domain.

To fully illustrate the potential of this grid computing technology, another future collaborator will be the National Grid Office (NGO) in Singapore. NGO has facilitated in the establishment of a national cyber-infrastructure (NGPP 2005) linking compute resources from the universities and research institutions. The grid-based virtual experimentation prototype could be deployed on the pilot platform to fully appreciate the issues as well as benefits of using an operational grid infrastructure.

ACKNOWLEDGMENTS

This pilot programme was funded by the Agency for Science, Technology and Research (A*STAR) in Singapore under the *Industrial IT – Integrated Manufacturing and Service Systems* initiative.

The authors would also like to thank the following collaborators for their contributions and advice to make this work a success: Associate Professor Ek Peng Chew, Associate Professor I.A. Karimi and Dr Suyan Teng as well as Choon Hwee Koh and Aloysius Chen from the Department of Industrial and Systems Engineering at the National University of Singapore, Surya Leo, Reman Chim and Quoc Khanh Nguyen from the Parallel and Distributed Computing Centre at Nanyang Technological University, Dr Bud Fox and Dr Wei Xiang from the Institute of High Performance Computing in Singapore, and Dr Kong Wei Lye from Singapore Institute of Manufacturing Technology.

REFERENCES

Brooks. 2005a. <http://www.brooks.com/pages/231_autosched_ap.cfm> [accessed March 31, 2005].

Brooks. 2005b. <http://www.brooks.com/pages/237_automod_simulation_software.cfm> [accessed March 31, 2005].

Chong, C.S, P. Lendermann, B.P. Gan, M.B. Duarte, J.W. Fowler and T.E. Callarman. 2004. Analysis of a semiconductor supply chain in a distributed simulation testbed. *Proceedings of the 2004 Winter Simulation Conference*, eds. R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 1902-1909.

Dagstuhl. 2002. <http://www.informatik.uni-rostock.de/~lin/GC/report/PADS.html> [accessed March 31, 2005].

Endrei, M., J. Ang, A. Arsanjani, S. Chua, P. Comte, P. Krogdahl, M. Luo and T. Newling. 2004. *Patterns: Service-oriented architecture and web services*. IBM Red Book.

Fine, C. H. 1999. *Clockspeed : Winning industry control in the age of temporary advantage*. Perseus Books Group.

Gan, B.P., L. Liu, S. Jain, S.J. Turner, W. Cai, and W.J. Hsu. 2000. Distributed supply chain simulation across the enterprise boundaries. *Proceedings of the 2000 Winter Simulation Conference*, eds. J.A. Joines, R.R. Barton, K. Kang, and P.A. Fishwick, 1245-1251.

Gan, B.P., M.Y.H. Low, P. Lendermann, S.J. Turner, X.G. Wang, and S.J.E. Taylor. 2005. Interoperating Autosched AP using the high level architecture. *Proceedings of the 2005 Winter Simulation Conference*, eds. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines.

Globus. 2005. <http://www.globus.org/> [accessed March 31, 2005].

Gottschalk, K. D., S. Graham, H. Kreger and J. Snell. 2002. Introduction to web services architecture. *IBM Systems Journal*. 41(2): 170-177.

JADE. 2005. <http://jade.tilab.com/> [accessed March 31, 2005].

Julka, N., P. Lendermann, C. C. Soon and L. F. M. Liow. 2004. Analysis and enhancement of planning and scheduling applications in a distributed simulation testbed. *Proceedings of the 2004 Winter Simulation Conference*, eds. R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 1799-1806.

Julka, N., P. Lendermann, W.T. Cai, N.Q. Khanh, and T. Hung. 2005. Grid computing for virtual experimentation and analysis. *Proceedings of the 2nd Workshop on Grid Computing & Applications*, May 2005, Singapore.

Kapoor, S., K. Bhattacharya, S. Buckley, P. Chowdhary, M. Ettl, K. Katircioglu, E. Mauch and L. Phillips. 2005. A technical framework for sense-and-respond business management. *IBM Systems Journal*. 44(1): 5-24.

Kuhl, F., R. Weatherly, and J. Dahmann. 1999. *Creating computer simulation systems: An introduction to the high level architecture*. Prentice Hall PTR.

Lanner. 2005. <http://www.lanner.com/products/simulation_suite/witness.php> [accessed March 31, 2005].

Lee, L.H., S.Y. Teng, E.P. Chew, I.A. Karimi, K.W. Lye, and P. Lendermann. 2005. Application of multi-objective simulation-optimization techniques to inventory management problems. *Proceedings of the 2005 Winter Simulation Conference*, eds. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines.

Lendermann, P., B.P. Gan, and L.F. McGinnis. 2001. Distributed simulation with incorporated APS procedures for high-fidelity supply chain pptimization. *Proceedings of the 2001 Winter Simulation Conference*, eds. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer. 1138-1145.

Lendermann, P., N. Julka, L.P. Chan, and B.P. Gan. 2002. Integration of Discrete Event Simulation Technology with Configurable Software Applications. Pending US Patent No. 10/301,788.

Lendermann, P., N. Julka, B.P. Gan, D. Chen, L.F. McGinnis, and J.P. McGinnis. 2003. Distributed supply chain simulation as a decision support tool for the semiconductor industry. *SIMULATION: Transactions of the Society for Modeling and Simulation International.* 79(3): 126-138.

Lendermann, P., B.P. Gan, Y.L. Loh, H.K. Tan, S.K. Lieu, L.F. McGinnis, and J.W. Fowler. 2004. Analysis of a borderless fab scenario in a distributed simulation testbed. *Proceedings of the 2004 Winter Simulation Conference*, eds. R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 1896-1901.

Low, M.Y.H., K.W. Lye, P. Lendermann, S.J. Turner, S.H. Leo, and R.T.W. Chim. 2005. An agent-based approach for managing symbiotic simulation of semiconductor assembly and test operations. Accepted by *International Conference on Autonomous Agent and Multiagent Systems (AAMAS).* Utrecht University, The Netherlands, July 25-29, 2005.

McLean, Charles, and F. Riddick. 2000. The IMS mission architecture for distributed manufacturing simulation. *Proceedings of the 2000 Winter Simulation Conference*, eds. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fiswick. 1539-1548.

NGPP. 2005. <http://www.ngpp.ngp.org.sg> [accessed March 31, 2005].

ProModel. 2005. <http://www.promodel.com/products/promodel/> [accessed March 31, 2005].

Strassburger, S., T. Schulze and U. Klein. 1999. Migration of HLA into civil domains. *SIMULATION*. 73(5): 296-303.

Swain, J.J. 2003. Simulation reloaded: Sixth biennial survey of discrete event software tools. *OR/MS Today*, 30(4): 46-57.

Taylor, S.J.E., Turner, S.J. and Low, M.Y.H. 2005a. The COTS simulation interoperability product development group. *Proceedings of the 2005 European Interoperability Workshop*. Simulation Interoperability Standards Organization, Institute for Simulation and Training, Florida, 05E-SIW-056.

Taylor, S.J.E., Turner, S.J., Mustafee, N., Ahlander, H. and Ayani, R. 2005b COTS distributed simulation: A comparison of CMB and HLA interoperability approaches to type I interoperability reference model problems. *SIMULATION: Transactions of the Society for Modeling and Simulation International.* 81(1): 33-43.

Taylor, S.J.E., X.G. Wang, S.J. Turner, and M.Y.H. Low. 2005c. Integrating heterogeneous distributed COTS discrete-event simulation packages: An emerging standards-based approach. Accepted by *IEEE Transactions on Systems, Man and Cybernetics*.

Terzi, S. and S. Cavalieri. 2004. Simulation in the supply chain context: A survey. *Computers in Industry*. 53: 3-16.

Turner, S.J., W.T. Cai, and B.P. Gan. 2001. Distributed supply chain simulation using high level architecture. *Transactions of The Society for Modeling and Simulation International*. 18(2): 98-109.

Wang, X.G., S.J. Turner, M.Y.H. Low, and B.P. Gan. 2005a. Optimistic Synchronization in HLA based distributed simulation. Accepted by *SIMULATION: Transactions of the Society for Modeling and Simulation International*.

Wang, X.G., S.J. Turner, M.Y.H. Low, and B.P. Gan. 2005b. A COTS simulation package emulator (CSPE) for investigating COTS simulation package interoperability. *Proceedings of the 2005 Winter Simulation Conference*, eds. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines.

AUTHOR BIOGRAPHIES

PETER LENDERMANN is a Senior Scientist and Head of the D-SIMLAB Programme at Singapore Institute of Manufacturing Technology (SIMTech). Previously he was a Managing Consultant with agiplan in Germany where his focus was on the areas of supply chain management and production planning. He also worked as a Research Associate at the European Laboratory for Particle Physics CERN in Geneva (Switzerland) and Nagoya University (Japan). He obtained a Diploma in Physics from the University of Munich (Germany), a Doctorate in Applied Physics from Humboldt-University in Berlin (Germany) and a Master in International Economics and Management from Bocconi-University in Milan (Italy).

His research interests include modelling and analysis of manufacturing and logistics systems as well as distributed simulation. His email address is `<peterl@SIMTech.a-star.edu.sg>`.

MALCOLOM YOKE HEAN LOW is a Research Engineer with the D-SIMLAB Programme at Singapore Institute of Manufacturing Technology. He received his Ph.D. from Oxford University in 2002. His research interests are in the areas of adaptive tuning and load-balancing for parallel and distributed simulation systems, and the application of multi-agent technology in supply chain logistics coordination. His e-mail address is `<yhlow@SIMTech.a-star.edu.sg>`.

BOON PING GAN is a Research Engineer with the D-SIMLAB Programme at the Singapore Institute of Manufacturing Technology. He received a Bachelor of Applied Science in Computer Engineering and a 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 application of genetic algorithms. His email address is `<bpgan@SIMTech.a-star.edu.sg>`.

NIRUPAM JULKA is a Research Engineer with the D-SIMLAB Programme at Singapore Institute of Manufacturing Technology. He received a Bachelor of Technology (Honors) from Indian Institute of Technology (IIT) Kharagpur in 1999 and a Master of Engineering from the National University of Singapore in 2002. His Masters dissertation work included investigation of chemical supply chains and development of proper frameworks for decision support systems to handle them. He is presently pursuing a Singapore-based Engineering Doctorate with Cranfield University in the field of Manufacturing Systems. His research interests are supply chain management and optimization, interoperability of decision support systems with simulation components and capacity expansion processes in global manufacturing networks. His email address is `<nirupamj@SIMTech.a-star.edu.sg>`.

LAI PENG CHAN is a Senior Research Engineer in the D-SIMLAB Programme at Singapore Institute of Manufacturing Technology. She received her B.Sc and M.Tech (Computer Science) degrees from the National University of Singapore. Her research interests include distributed computing, simulation modeling and optimisation. Her email address is `<lpchan@simtech.a-star.edu.sg>`.

STEPHEN J. TURNER joined Nanyang Technological University (Singapore) in 1999 and is Director of the Parallel and Distributed Computing Centre in the School of Computer Engineering. Previously, he was a Senior Lecturer in Computer Science at Exeter University (UK). He received his MA in Mathematics and Computer Science from Cambridge University (UK) and his MSc and PhD in Computer Science from Manchester University (UK). His current research interests include parallel and distributed simulation, distributed virtual environments, grid computing and multi-agent systems. He is steering committee chair and program co-chair of the Workshop on Principles of Advanced and Distributed Simulation and advisory committee member and general chair of the Distributed Simulation and Real Time Applications Symposium. His e-mail address is `<ASSJTurner@ntu.edu.sg@ntu.edu.sgedu.sg>`.

WENTONG CAI is an Associate Professor with the School of Computer Engineering at Nanyang Technological University (NTU), Singapore, and Head of the Computer Science Division. He received his B.Sc. in Computer Science from Nankai University (P. R. China) and Ph.D., also in Computer Science, from University of Exeter (U.K.). He was a Post-doctoral Research Fellow at Queen's University (Canada) before joining NTU in February 1993. His current research interests include: Grid and Cluster Computing, Parallel and Distributed Simulation and Parallel and Distributed Programming Environments. His email address is `<aswtcai@ntu.edu.sg@ntu.edu.sg>`.

XIAOGUANG WANG is a Ph.D student at the School of Computer Engineering at Nanyang Technological University, Singapore. She received her B.Sc in Computer Science form Nanjing University of Aeronautics and Astronautics, China in 1997. Her research interests include Distributed Simulation and the High Level Architecture. Her e-mail address is `<xgwang@pmail.ntu.edu.sg>`.

LOO HAY LEE is an Assistant Professor in the Department of Industrial and Systems Engineering, National University of Singapore. He received his B.S. (Electrical Engineering) degree from the National Taiwan University in 1992 and his S.M. and Ph.D. degrees in 1994 and 1997 from Harvard University. He is currently a senior member of IEEE, vice president of ORSS, and a member of INFORMS. His research interests include simulation-based optimisation, production scheduling, logistics and supply chain planning, and vehicle routing. His email address is `<iseleelh@nus.edu.sg>`.

TERENCE HUNG is the Manager of the Software & Computing Programme at the Institute of High Performance Computing in Singapore where he is responsible for R&D in the areas of grid computing and visualisation. He was formerly VP Technology of Commerce Exchange Pte

Ltd and has consultancy experience in HPC, e-commerce and e-financial. His current research focus is on grid middleware to address needs in computational science and engineering problems. In addition to leading the national Physical Science Virtual Grid Community efforts, Terence is also a member of the national Manufacturing Virtual Grid Community and Middleware workgroup. He obtained his Ph.D. in Electrical Engineering from the University of Illinois at Urbana-Champaign in 1993. His email address is <terence@ihpc.a-star.edu.sg>.

SIMON J.E. TAYLOR is a Senior Lecturer in the Department of Information Systems and Computing and is a member of the Centre for Applied Simulation Modeling, both at Brunel University, UK. He is also a Visiting Associate Professor at the Parallel and Distributed Computing Centre at Nanyang Technological University in Singapore. He is also Information Director of ACM SIGSIM, ACM SIGSIM PADS Liaison Officer and Chair of the Simulation Study Group of the UK Operational Research Society. He is a steering committee member of PADS and general co-chair of the UK Simulation Workshop Series. His main research interests are distributed simulation and applications of ICT to simulation modeling. His email address is <simon.taylor@brunel.ac.uk>.

LEON F. McGINNIS is Eugene C. Gwaltney Professor of Manufacturing Systems at Georgia Institute of Technology, where he also serves as Associate Director of the Manufacturing Research Center and founding Director of the Keck Virtual Factory Lab. His research focuses on the application of operations research and computer science to solve decision problems arising in the design and operation of industrial logistics systems. His email address is <leon.mcginnis@isye.gatech.edu>.

STEPHEN BUCKLEY has been a Research Staff Member at the IBM Thomas J. Watson Research Center in Yorktown Heights, NY since 1987, and a manager at that facility since 1995. His most recent interest is in Sense & Respond systems; his team has implemented Sense & Respond for several IBM lines of business. He was one of a group of IBM researchers who received the prestigious Franz Edelman award from INFORMS in 1999 for the successful deployment of supply chain simulation and optimization technology in IBM. He received the Ph.D. degree in Computer Science from MIT in 1987. His e-mail address is <sbuckley@us.ibm.com>.