Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

ABOUT THE NEED FOR DISTRIBUTED SIMULATION TECHNOLOGY FOR THE RESOLUTION OF REAL-WORLD MANUFACTURING AND LOGISTICS PROBLEMS

Peter Lendermann

Planning and Operations Management Group
Singapore Institute of Manufacturing Technology
71 Nanyang Drive, Singapore 638075
SINGAPORE

ABSTRACT

Distributed simulation has undergone several cycles of ups and downs in recent years. Although successful in the military domain, it appears that the idea of applying distributed simulation in other fields for modeling and analysis of large-scale, heterogeneous systems such as communication networks or supply chains has still not taken off until today. Is this because of inherent limitations or lack of applicability as such? Or is it because of additional research issues that are yet to be resolved to make distributed simulation applicable? In this paper, the problem is discussed specifically with regard to the application of distributed simulation for design, operation and performance enhancement of manufacturing and logistics systems.

1 INTRODUCTION

Distributed simulation refers to technologies that enable a simulation program to execute on a computing system containing multiple processors that are interconnected by a communication network (Fujimoto 2000). It was originally motivated by needs in the military domain for more effective means to train personnel in distributed virtual environments that mimic actual combat situations (Fujimoto 1998).

Subsequently, the availability of synchronization middleware such as the Runtime Infrastructure of the High Level Architecture (Kuhl et al. 1999) has also inspired research looking at potential application of distributed simulation for modeling and analysis of large-scale, heterogeneous systems such as communication networks or global supply chains.

For example, with industries such as semiconductor and automotive heading towards a paradigm in which the entire manufacturing process will be digitally represented and simulated before any construction work commences, a natural way to achieve a detailed simulation model that covers all causal relationships between the different processes in the factory would be to couple independently designed and developed simulation models (Strassburger et al. 2003).

Alternatively, in the case of the design of a container terminal, several organizations might be involved, each of them using simulation to accomplish their own assignment. Although these organizations focus on simulation models for different parts of the operations, they share the common goal of making the terminal function as a whole which also would require coupling of the respective models (Verbraeck 2004).

Ultimately, when extending the scope beyond the four walls of a factory to an entire supply chain, the notion of distributed simulation might even become indispensable when the participating organizations are not willing to share detailed model information (Gan et al. 2000).

Distributed simulation in the context of supply chain management has also been featured, e.g., by Linn et al. (2002) and Lendermann et al. (2003). These studies, however, have put emphasis on technical feasibility rather than the usage of the technology to solve real-world problems.

According to Boer (2005), *the simulation community in industry* is still looking for an acceptable solution to couple distributed simulation models. This has led to a detailed analysis with the objective to provide an architecture for coupling simulation models and test its appropriateness in industry. However, this study has also not answered whether *industry itself* is also looking for an acceptable solution to couple distributed simulation models.

Similarly, an architecture and interfaces for distributed manufacturing simulation were developed as part of the IMS Mission Project (McLean et al. 2005). For this work, the machine shop operations of a vacuum systems manufacturer were used to help define the requirements for distributed simulation modeling and data interface specifications. However, this paper also does not mention whether distributed simulation was actually required to resolve a specific operational challenge that decision-makers were interested in.

This naturally leads to the question whether and if yes for what kind of "real-world problems" distributed simulation technology is really an indispensable tool? Or is it a "toy" for researchers to address and solve "imaginative" supply network problems? Or is it even just of academic interest to computer scientists without any considerable relevance for any other discipline?

Rather than asking the question "Why has distributed simulation still not found wide application?", should one not rather ask "What kind of challenges are there that can be resolved only with distributed simulation?"

The objective of this paper is to address these questions in more detail: Section 2 re-captures relevant terminology and describes general technical issues related to distributed simulation. This is followed by a review of applicability of simulation as such in view of ongoing paradigm shifts in manufacturing and logistics in Section 3. Some additional conceptual issues relating to simulation modeling of real-world, large-scale distributed systems are discussed in Section 4, followed by a detailed discussion of a few selected applications in Section 5. Based on the limitations arising from Sections 2, 3, and 4, conclusions are drawn with regard to why or why not it makes sense to apply different simulation technology for these specific applications, and − for those applications where it appears to make sense − what are the specific questions that can be answered with the distributed simulation technology as available today, as well as what are additional research issues that will have to be resolved to actually make it happen.

2 DISTRIBUTED SIMULATION TERMINOLOGY AND GENERAL TECHNICAL ISSUES

An integrated distributed simulation model (federation) consists of several model components (federates) that are running on several computers. Material and information flow in between these federates are represented by messages that are exchanged between them through a network, LAN or even the internet.

Even though other architectures such as the FAMAS Simulation Backbone (Boer et al. 2002) have also successfully been applied, the High Level Architecture (HLA, IEEE standard 1516) has emerged as a standard for distributed simulation. The Runtime Infrastructure (HLA-RTI) is used as the middleware for interoperation and synchronization (Kuhl et al. 1999).

Most of the commercially available simulation packages are not yet compliant with HLA standards. However, an international effort led by the Centre for Applied Simulation Modeling at Brunel University (UK), namely the *Commercial-off-the-shelf Simulation Package Interoperability Product Development Group* (CSPI-PDG), has been endorsed by the Simulation Interoperability Standards Organization (Taylor et al. 2006) to drive the development of:

- Standard reference models for distributed simulation,
- A standard for data exchange representation (Object Model Template),
- A standard data exchange mechanism, and
- A standard specification for a distributed simulation co-ordination tool.

3 SIMULATION VERSUS PLANNING AND SCHEDULING

In manufacturing or logistics, simulation has traditionally been used to analyze key performance indicators and refine operations through a steady-state simulation approach using commercial simulation packages in the following manner:

1. Raw material is released at a constant rate and mix into the (simulated) manufacturing or logistics system.
2. After the end of a "warm-up period" (which is required to bring the simulated system from empty state to steady state), collection of KPI-relevant statistics commences. To collect a sufficient number of samples for the required statistical confidence of the result, several replications with different random streams have to be carried out.
3. Depending on the simulation results, system configuration, dispatch rules and other parameters are then refined, and the simulation is repeated to study how the system performance is affected and how it can be enhanced.

In many domains, because of the typically long modeling and validation cycle, analysis is often conducted with rather old system data that does not fully represent any more the operations that are of interest. For this reason, simulation technology has been found to be rather impractical for operational purposes and has therefore been used mainly for tactical and strategic decision support.

At the same time, due to decreasing product lifecycles, increasing number of products and constantly changing demand (quantity and mix), hardly any manufacturing or logistics system still operates in steady state today. In such an environment, a lot of potential benefit that could be gained from simulation analysis can actually not be realized for the real operations. Because of the long modeling-analysis-implementation cycle time, the system would have changed significantly in terms of load, product mix, resource mix by the time measures derived from the simulation analysis can be implemented.

Consequently, simulation of one year of operations of a semiconductor foundry from $t_1 = 0$ onwards with fixed rules and policies would not be a good representation of reality because demand information, resource availability and product mix in the system would already have changed significantly at a time t_2 << 1 year, resulting in a change of production targets that require different dispatch rules and dedication policies. This is illustrated in Figure 1.

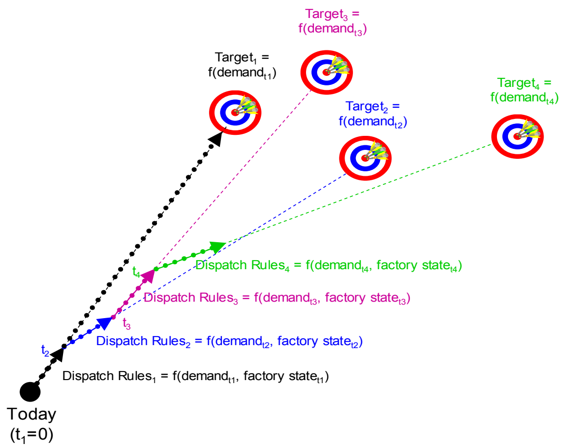

Figure 1: Simulation in a Fast-Changing Environment

A new paradigm for the application of simulation technology for decision-making in manufacturing and logistics is believed to bear more potential in the future: The latest system status will be used as starting point for a high-fidelity simulation, and the performance evolution will be assessed over a relatively short period of time (for example, from t_1 to t_2 in Figure 1). Ideally, the simulated time period would correspond with the frequency of making relevant decisions.

This, however, means that – rather than looking at a steady state – the dynamic evolution of the system from a specific point of time (in most cases probably the latest system state) onwards is now of principal interest. One of the important questions to be addressed is: "What is the best way to react to a contingency (such a contingency could also be the availability of new demand information)?" Since this, however, is one of the major questions that a scheduling system would already address, the legitimate question arises whether discrete event simulation is the appropriate technology to answer such questions.

Deterministic scheduling systems are basically not able to portray random effects; rather they make use of "best" (i.e. most likely) values for each parameter. Inherently they are applicable only for environments that are subject to "little" randomness (or that incorporate buffers to compensate random effects such as variations of processing times) so that such a deterministic schedule becomes "executable". The underlying question of such a scheduling process is: "What is the best-performing schedule (assuming that there is no variability)?" or (after a contingency, see above) "What is the best way to get back on plan?" This also means that once there is a significant deviation from schedule or new information is available, operations would have to be rescheduled.

Generation of factory-wide "schedules" is not deemed possible in an environment where the variability is "faster" (i.e. a predominant effect) compared to the time it takes to generate the schedule because the event density and the variability are so high that by the time a schedule is generated it would already be obsolete. That is the reason why, for example, semiconductor wafer fabrication operations are managed by dispatching (dispatching decisions are made more or less instantly).

In turn, the power of discrete event simulation lies in its ability to portray randomness in such a highly variable environment. Even though it is not possible to generate a schedule, the question that can be answered is: "What are the policies (including dispatching policies) and configurations that are likely to result in the best system performance until the time when the simulation (or simulation optimization) exercise is conducted again (maybe a few hours or days later)?" or (after a contingency) "Is there any, and if yes what policy (including dispatching policy) or configuration change is required to obtain best performance during the time the contingency prevails?" In this context, a "policy" is a rule for operational decisions that needs to be applied in between cycles of simulation (optimization) exercises. The "cycle" time would be the time in between implementations of simulation analysis (optimization) results (once every few hours or days).

4 CHALLENGES FOR SIMULATION MODELING OF REAL-WORLD DISTRIBUTED SYSTEMS

In the setting of the above-described paradigm changes, additional conceptual limitations associated with modeling of systems as complex as a large factory or a supply chain must be taken into consideration when talking about real-world issues that decision-makers would like to be able to address. Since semiconductor manufacturing is a heavy user of simulation already today, examples from the semiconductor domain will be used for the purpose of this discussion.

4.1 Representation of Software-Enabled Decision-Support Processes

A simulation model is not a good representation of reality if whatever drives the underlying system is not represented properly. Traditionally, simulation has successfully been applied for design and performance enhancement of systems that are driven by material release into and material availability within the system. In such a system, for example a wafer fab, once a production lot is available for processing on an appropriate machine that is also available, processing will be executed according to a pre-determined dispatch rule without waiting for any additional demand signal from downstream.

In the case of the semiconductor supply chain, however, from the completion of wafers onwards operations are not driven any more by my material release, rather they are driven by customer demand. To enable meaningful representation of reality, not only the generation of customer demand itself but also the translation of customer demand into material release and movement decisions have to be represented.

Once this is done appropriately, not only the time horizon of a simulation can be extended significantly (beyond the average of t_{i+1}-t_i in Figure 1) but also the scope can be extended from one individual factory to a multiple echelon supply chain.

As long as the translation of customer demand into material release and movement decisions is primarily software-enabled, it is not impossible to realize such a distributed supply chain simulation. For example, Chong et al. (2004) have demonstrated how a virtual experimentation testbed that comprises not only wafer fab and assembly & test models but also a federate that represents a customer order management system can be used to adjust dispatch priorities in the fab to maximize the on-time delivery of finished ICs.

4.2 Heterogeneity of External Drivers

Even though it is possible to generate and run distributed supply chain models that also comprise federates such as a customer order management system, the supply chains that are the basis for the analysis of studies such as Chong et al. (2004) are still highly simplified representations of real-world supply networks that are relevant to decision-makers.

But where in reality do all the wafers produced in a particular wafer fab (W/F) move to one single assembly & test (A/T) facility? And which A/T facility receives all wafers from one single fab? As illustrated in Figure 2, in most if not all real-word semiconductor supply chains, a system of, for example, two W/F and two A/T is not self-contained. To represent reality appropriately, additional demand to the W/F from other A/T facilities and/or additional supply coming in to the A/T from other W/F has to be taken into consideration.

The situation will be further complicated by the fact that each of the "Other suppliers/customers" in Figure 2 could be contract manufactures with additional customers and suppliers, respectively, as well as third party logistics providers executing the material movements between the different nodes that come into place, each with their own policies and constraints.

Because of the heterogeneity of the external drivers and a large number of interfaces, it is hard to imagine that it is actually possible to develop and maintain a high-fidelity simulation model of such a complex system.

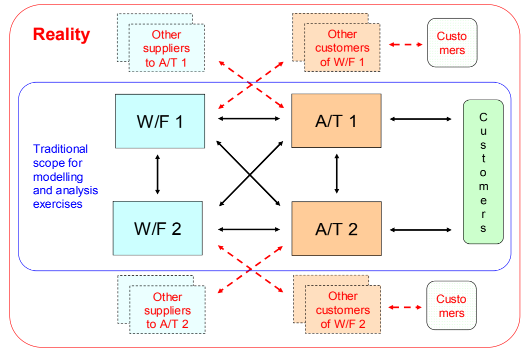

Figure 2: Scope of Supply Network Simulation

4.3 Representation of Human Decision-Making

Unlike as assumed in Section 4.1, many material release and movement decisions in real systems are actually made by humans, especially on the supply chain scale. Humans, however, are not only inherently unstable and unpredictable but also capable of independent actions.

How to represent such factors in a simulation model has been investigated, for example, in Benedettini et al. (2006). Apart from the fact that this research area is still in a rather early stage, it also focuses on workers' performance modeling rather than the representation of complex decisions made by executives.

It seems that the lack of ability to represent such decision-making processes will always be an inherent limitation for generation of high-fidelity models of complex manufacturing and logistics systems.

4.4 Modeling Speed

Decision-makers need tools that are able to give them the solution to their immediate questions rather than at a time when the answer is already obsolete. As discussed in Taylor et al. (2004), many initiatives have been pursued to develop the *information technology* standards that would help speed up the simulation modeling process. However, a really significant reduction of the cycle time for simulation modeling will only be possible if standardization also takes place on the *application* level that would result in archived, re-usable simulation model components that require much less customization effort.

Although the challenges associated with these standard development activities are tremendous, there appears to be no inherent fundamental limitation that would make it impossible to resolve these challenges.

4.5 Execution Speed

In a distributed simulation, federates interact through messages. A message sent from one federate to another represents an event originating in the first federate and affecting the second. Appropriate mechanisms for interoperation and synchronization of different federates have to be in place. "Conservative" synchronization requires that the second federate's local time should not be greater than the time stamp of the arriving message, otherwise the causality principle would be violated.

To make sure that such causality constraint violations are avoided, lookahead is an important parameter to be considered. Its value is determined by the federate's quickest response time to messages it subscribes (Fujimoto 1997). The lookahead value has great implications on the runtime of a distributed simulation. If it is large, federates can potentially achieve a high degree of parallelism in processing events. However, events that have immediate consequences in other federates require near-zero lookahead, resulting in a lot of synchronization overhead that does not allow much parallelism and therefore can slow down the distributed simulation tremendously. Most contemporary distributed systems do comprise such events, especially if information flow between a physical system component represented by one federate and a software-enabled decision support component (as described in Section 4.1) represented by another federate is involved.

Optimistic synchronization mechanisms can help relax synchronization constraints and increase parallelism. But due to the overheads associated with state-savings and rollbacks, significant reductions in simulation runtime can be achieved only for systems where causality errors may occur but in fact seldom occur (Wang et al. 2005a).

Ultimately, events requiring near-zero lookahead will always be a major limiting factor for execution speed of distributed simulation of complex systems.

5 APPLICATION SCENARIOS FOR DISTRIBUTED SIMULATION

Having pointed out some of the major challenges associated with distributed simulation modeling and execution, a selection of application scenarios for distributed simulation that have been proposed in the literature (some of which have already been mentioned in Section 1) will now be discussed. It is important to note that for each scenario the starting point for the discussion is the question "What operational problems need to be resolved?" followed by an analysis with regard to some of the challenges discussed in Section 4, rather than looking at an imaginative technology solution and then asking what kind of problems it could actually solve. The considerations made in this section are summarized in Table 1 at the end of this section.

5.1 Across-Echelon Supply Chain

Other than in the semiconductor industry, where simulation is already a well-established technology to address systems design and operational performance analysis issues (not only in wafer fabs but to some extent also in assembly & test facilities), and to a certain degree in the automotive industry, simulation models of most of the nodes of a typical across-echelon supply chain would not be available, and as a result the simulation modeling effort would be extremely high. Also, as mentioned in Sections 4.3 and 4.2, it is very difficult to represent human decision-making processes with sufficient fidelity and, more importantly, external drivers of such a supply chain are very heterogeneous. Therefore the applicability of distributed simulation technology for across-echelon supply chain management in the real world appears to be very unrealistic.

5.2 Borderless Fab

Many semiconductor manufacturers have several fabs to produce a given device. In a world where effective resource utilization and shorter cycle times become more critical than ever before, exploration of how capacity and cycle time of a system of several fabs can be enhanced on an aggregate level appears to be worthwhile, especially when the fabs are in close proximity to each other, something that it not uncommon in wafer fab parks in Singapore or Taiwan.

The need for re-routing lots from one fab to another might arise either from a temporary resource breakdown or from an inappropriate resource mix in the first fab with regard to the product mix to be processed. Re-routing could be done for a specific production step only, alternatively scenarios in which the re-routed lots remain in the second fab are also possible (see Figure 3).

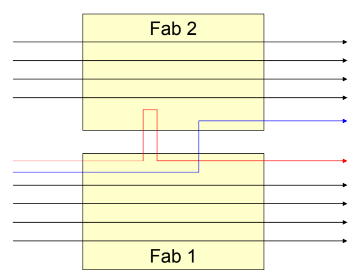

Figure 3: Borderless Fab

In the case of such a "Borderless Fab", the specific operational question to be addressed would be "What re-routing policies should be applied in the next few hours/days to achieve optimal throughput and cycle time?", taking into account the paradigm changes described in Section 3. Connecting simulation models of the individual fabs and executing simulation scenarios in a distributed manner would be the most straightforward way to do this. If the fabs belong to different organizations, application of distributed simulation technology would even be mandatory since it is unlikely that these organizations would be willing to reveal to each other the intellectual property associated with their own fab models when they are integrated into one big model.

The feasibility of using distributed simulation to study Borderless Fab application scenarios has already been demonstrated in Lendermann et al. (2004). Subsequently, doing the same kind of exercise using AutoSched AP, the most commonly used commercial simulation package for wafer fabs, using interoperation mechanisms that are compliant with CSPI-PDG standards has also been enabled as described in Gan et al. (2005a).

Even with regard to the execution time of such distributed scenarios significant progress has been made through a time synchronization mechanism that makes use of the manufacturing process flow information (Gan et al. 2005b). To what extent optimistic synchronization protocols could help to further speed up the simulation, however, remains to be investigated.

The modeling process for a Borderless Fab would not start from scratch since models of the individual fabs would already be available in many cases. Also, external drivers in this single-echelon system are less heterogeneous compared to the across-echelon supply chain. Lastly, in the highly automated environment of a 300mm wafer fabs the need for representation of human decision-making is less critical.

For these reasons, this specific representation of the within-echelon supply chain coordination and optimization problem is probably the most promising application scenario for distributed simulation technology, even though it might not become reality in the near-term future because today many products are still certified by semiconductor manufacturers' customers for production in one specific fab. A paradigm change after which this can be done in two (or more) different fabs (that may even belong to different organizations working together on a project basis, therefore requiring the capability of hiding confidential model data from each other) will certainly take time.

5.3 Modeling and Analysis of 300mm Wafer Fabs with AMHS

Specifically in 300mm wafer fabs, where the representation of automated processes plays an more important role and the representation of human decision-making plays a less important role, manufacturing operations and AMHS (Automated Material Handling System) operations are typically handled in separate simulation models. In many cases, these models would have been developed with simulation packages by Brooks Software (Brooks 2006): AutoSched AP for the manufacturing operations and AutoMod for the AMHS. A Message Communication Module MCM that enables the interoperation between AutoSched AP and AutoMod is also available, however, it makes use of a simple conservative synchronization protocol that does not allow much parallelism. As a result, the execution time is rather slow, and it has therefore has been considered impractical by many 300mm fab operators.

A promising attempt to overcome this problem has been described by Wang et al. (2005b): A "conservative-optimistic" synchronization method involves a compensation in the federate which receives an event in its past. The compensation has the effect of a roll-back in a time warp synchronization, but does not require any anti-messages or other computations because it is assumed that the amount of the rollback is less than the time it takes the receiving federate to react to the received event. This works because in the special case of fab/AMHS system the event density in the AMHS is significantly higher than in the fab. More events can be processed in the AHMS federate if the RTI grants a time accordingly. Because of the lower event density in the fab federate, the compensation would still be smaller than the time it takes the receiving (fab) federate to respond to the received event.

Machine breakdowns, however, have not been taken into consideration in this study. If there are failure events and the virtual rollback sweeps up such an event, the breakdown needs to be compensated, provided that it has an effect on how the receiving federate responds to the received event requiring compensation. Additional overhead, however, would incur because failure events need to be checked for, even though this might be as simple as checking the machine state and scheduling the response to the received event using the virtual rollback if the machine has not failed, or based on the time the machine becomes available if it has failed.

If, on the other hand, the failure event and the received event do not interact, then the virtual rollback should still work. In fact, if the failure does not happen at the tool that receives the event, it might not have any significant effect on the model behavior. In a 300mm wafer fab the likelihood for such a failure happening at the receiving tool is probably quite small. More research will have to be conducted to investigate whether such a model would still be good enough for decision-making. On top of that, questions such as how machine failures are to be modeled or whether or not a tool is reserved before a lot is dispatched need to be addressed.

Typically, AMHS suppliers keep the details of their AHMS models confidential, this makes running the fab model and the AMHS model in a distributed environment mandatory. Distributed simulation is therefore equally relevant for the modeling and the analysis of 300mm fabs with AMHS, also because – similarly to the Borderless Fab – the heterogeneity of external drivers would not be so much of an issue.

5.4 Semiconductor Assembly & Test Schedule Adherence

In semiconductor assembly & test facilities, scheduling systems have been successfully applied in the industry (Quadt and Kuhn 2001, Sivakumar et al. 2001, Chong et al. 2002). However, generation of a new schedule takes a considerable amount of time. At the same time, the system is subject to significant variability which is not portrayed in the scheduling system. Because of this, even if a scheduling system is able to generate a "perfect" schedule with regard to the time when the input data were taken from the manufacturing execution system, the schedule will be already sub-optimal at the time when it is actually generated. As a consequence, production supervisors will have to make independent dispatch decisions on how to react to such variability according to certain rules.

A distributed simulation with one federate representing the scheduling system and another one to represent the manufacturing operations (note that it would be important to continue the simulation of the manufacturing system while the schedule is generated during a finite period of time) appears to be an excellent way to address this and to study questions such as what scheduling policy results in best system performance, what policies should be applied to react to variability of the manufacturing system, and under what circumstances (degree of schedule adherence) should re-scheduling be triggered.

However, semiconductor assembly & test is still subject to significant human decision-making on the shop-floor. Moreover, heterogeneity of external drivers may be an issue if certain production steps are partly outsourced (even though the ability to hide modeling data might be required in this case). Therefore, applicability of distributed simulation appears to be rather unlikely for such systems.

5.5 Aerospace Spare Parts Logistics

Another interesting application for discrete event simulation is emerging in the airlines industry. On-time departure of aircraft is considered as one of the key success factors in this industry. One of the major risks associated with on-time departure is technical delay. To minimise the risk of technical delay, airlines will position a certain number of critical parts (rotables) that have been identified crucial for the dispatch of the aircraft at the destination airports.

In anticipation of new aircraft such as Airbus A380 or Boeing B787 entering into service, a new business paradigm for spare parts management is emerging: Rather than selling spare parts to airlines, an OEM (Original Equipment Manufacturer) or its designated service provider would supply spare parts to airline clients with a guaranteed service level whenever needed. In this setting, new simulation-based decision support tools are required that are able to portray with high fidelity the dynamic implications of advanced business practices for rotables management such as early initiation of delivery logistics actions triggered by advanced failure message transfers and continuous re-balancing of rotables inventory to minimize risk associated with subsequent failures, and enable the determination of:

- The minimum inventory requirements at mainbase and outstations as well as the number of additional warehouses needed,
- The location from where a spare part should be taken, based on availability, expected delivery time and risk, and
- The service level that can be committed with what confidence level to each airline customer.

An important component in the associated service chain are the logistics actions that are required to move critical parts from one location to another to avoid Aircraft-On-Ground situations. Typically, the underlying logistics network would be designed in such a way that parts are moved over relatively short distances, e.g., from Chicago to New York or from Frankfurt to Paris. In turn, logistics movements across continents would be a very rare exception.

Such a network can be represented as a set of subsystems (each representing a "region") that have little interaction with each other. One could expect that a significant degree of concurrent execution can be achieved if these subsystems are distributed over several computers. Because of the few interactions between regions, optimistic protocols as explained in Section 4.5 appear to be particularly applicable. Apart from that, heterogeneity of external drivers, representation of human decision-making and model component availability as such are not a limiting factor. The detailed implications of this particular application, however, have not yet been discussed in the literature.

5.6 Automotive Manufacturing

Unlike in the aerospace domain, the feasibility of using distributed simulation in the automotive domain for reducing the cycle time of layout decisions has already been addressed in the literature (Taylor 2005). In particular, when planning a new engine production line, many complex factors such as machine cost and reliability, partially built engine test, repair and recycle time, and varying operator shift patterns and availability must be taken into account.

Up to this stage, a reduction of simulation execution time has been demonstrated for such systems only with a relatively simple model that only requires asynchronous entity passing (no bounded buffers that would require information exchange with quasi-instant feedback – and therefore near-zero lookahead) between the individual sub-models.

Also, unlike in the aerospace case as described in Section 5.5, it appears to be much more difficult to keep the number of interactions between subsystems small. To what extent optimistic protocols could help to reduce execution time remains to be investigated.

Similar to semiconductor assembly & test, the need for keeping model components separated may arise (although in this case heterogeneity of external drivers could be problematic). Human decision-making as well as model component availability appear to be less of an issue. Overall, further investigation of this particular application scenario appears to be promising.

5.7 Container Terminal

Lastly, distributed simulation has also been used in the context of complex inter-organizational problem solving for the analysis of port extensions (Verbraeck 2004).

To achieve a model for efficient handling of trucks at future container terminals in the port of Rotterdam, several simulation models (container handling with different layout options, truck generation, road traffic) and an agent-based planning and scheduling tool for negotiating truck arrival times at the terminal were integrated into one distributed simulation environment, since it was believed to be impossible to pull together the diversity of model components into one single simulation model.

For this particular application scenario, heterogeneity of external drivers might be an issue though. More severely, availability of model components, especially the road traffic model, as well as representation of human behavior might also be limiting factors. For these reasons, further investigation of this particular application scenario is certainly needed.

Table 1: Overview of Application Scenarios for Distributed Simulation

	Heterogeneity of External Drivers		Human Decision-Making		Model Availability & Modeling Speed		Simulation Execution Speed		Data Hiding potentially required		**Overall potential for application of distributed simulation**
Across-Echelon Supply Chain	Practically impossible in most real-world supply networks	--	Even operational decisions are taken manually	--	Individual models would not exist in most cases	--	Difficult because of federates that require near-zero lookahead	-	Several parties might be involved	+	**Very unlikely**
Borderless Fab	No significant external drivers other than customers, therefore not a limiting factor	0	A limitation only in the 200mm fab environment, especially for inter-fab material flow decisions	0	In many cases models of individual fabs are already available, therefore not a major limitation	0	To be investigated in more detail	0	Applicable in the future when different contract manufacturers might work together on a project basis	++	**Likely**
Fab/AMHS	Not a limiting factor	0	Applicable for 300mm fab environment with high degree of automation	+	Both fab and AMHS model are typically available, therefore not a limiting factor	0	To be investigated in more detail	0	AMHS Suppliers are not necessarily ready to share model details	++	**Likely**
Assembly & Test Schedule Adherence	Potentially a limiting factor if some part of the production is outsourced	-	Variability in the system often requires independent dispatch decisions by production managers, leading to gradually decreasing schedule adherence	-	High-fidelity models are required	-	Not required	0	Not required	0	**Unlikely**
Aerospace Spare Parts Logistics	Not a limiting factor	0	Not a limiting factor	0	Not a limiting factor	0	Optimistic protocols bear significant potential	++	Not required	0	**Likely**
Automotive Manufacturing	Potentially a limiting factor if some part of the production is outsourced	-	Not a limiting factor	0	Unlikely that all models would be available	0	Difficult because of federates that require near-zero lookahead	0	Several parties might be involved	+	**To be investigated in more detail**
Container Terminal	Potentially a limiting factor in application scenarios described in the literature	-	To be investigated in more detail	0	To be investigated in more detail	0	Difficult because of federates that require near-zero lookahead	0	Several parties may be involved	+	**To be investigated in more detail**

6 CONCLUSIONS AND OUTLOOK

As illustrated in this paper, the number of application scenarios for distributed simulation to resolve real-world manufacturing and logistics challenges is actually quite limited. In particular, even though it has been discussed most (with regard to technical feasibility) in the literature, the original across-echelon supply chain management scenario is probably the most unrealistic one due to the inherent limitations associated with it.

In some domains however, not just *the simulation community in industry* but *industry itself* is expected to be looking for acceptable solutions to couple distributed simulation models in future, especially in semiconductor manufacturing where the most promising application scenarios as of today can be found. Not only the Fab/AMHS application appears to make a lot of sense but also the Borderless Fab scenario is likely to become more relevant in the future. But before reasonable simulation execution times become feasible additional research issues are yet to be resolved in both cases.

Other interesting application scenarios have been emerging. For example, a simulation model of the Southampton Process, Testing and Issuing (PTI) centre and four hospital models has recently been realized for the UK National Blood Service (Brailsford 2006). In this study, distributed simulation has been used because only a limited number of hospital simulation models can effectively be executed as a standalone model. To what extent distributed simulation is really required to enable the associated decision-making processes is one of the questions to be addressed in more detail in the future.

ACKNOWLEDGMENTS

The author would like to thank numerous colleagues and collaborators for many fruitful discussions about the topics and challenges covered in this paper: Gan Boon Ping and Nirupam Julka from Singapore Institute of Manufacturing Technology, Associate Professor Stephen J. Turner and Dr Malcolm Low from Nanyang Technological University, Dr Simon Taylor from Brunel University, and especially Professor Leon F. McGinnis from Georgia Institute of Technology.

REFERENCES

Benedettini, O., T. Baines, and J. Ladbrook. 2006. Human Performance Modeling within Manufacturing Systems Design. *Proceedings of the 2006 UK Operational Research Society Simulation Workshop (SW06)*.

Boer, C.A. 2005. *Distributed Simulation in Industry*. Erasmus Research Institute of Management. ERIM Ph.D. Series Research in Management 65.

Boer C. A., A. Verbraeck, and H. P. M. Veeke. 2002. Distributed Simulation of Complex Systems: Application in Container Handling. *SISO European Simulation Interoperability Workshop*. pp.134-142.

Brailsford, S., K. Katsaliaki, N. Mustafee, and S.J.E. Taylor. 2006. Modelling Very Large Complex Systems Using Distributed Simulation: A Pilot Study in a Healthcare Setting. *Proceedings of the 2006 UK ORS Simulation Study Group Workshop*. UK Operational Research Society, Birmingham, UK.

Brooks. 2006. <http://www.brookssoftware.com/> [accessed May 22, 2006].

Chong, C. S., A. I. Sivakumar, and R. Gay. 2002. Design, Development and Application of an Object Oriented Simulation Toolkit for Real-time Semiconductor Manufacturing Scheduling. *Proceedings of the 2002 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

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*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Fujimoto. R. M. 1997. Zero Lookahead and Repeatability in the High Level Architecture. In *Spring Simulation Interoperability Workshop 1997*. <http://www-static.cc.gatech.edu/computing/pads/papers.html#1997> [accessed May 16, 2006].

Fujimoto, R.M. 1998. Parallel and Distributed Simulation. In *Handbook of Simulation: Principles, Methodology, Advance, and Practice*. Ed. J. Banks. John Wiley & Sons Inc. New York NY, USA. pp. 429-464

Fujimoto, R. M. 2000. *Parallel and Distributed Simulation Systems*. John Wiley & Sons Inc. New York, NY, USA.

Gan, B. P., L. Liu, S. Jain, S. J. Turner, W. T. Cai, and W. J. Hsu. 2000. Distributed Supply Chain Simulation Across the Enterprise Boundaries. *Proceeding of the 2000 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Gan, B. P., P. Lendermann, M. Y. H. Low, S. J. Turner, X. G. Wang, and S. J. E. Taylor. 2005a. Interoperating AutoSched AP Using the High Level Architecture. *Proceedings of the 2005 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Gan, B. P., M. Y. H. Low, X. G. Wang, and S. J. Turner. 2005b. Using Manufacturing Process Flow for Time Synchronization in HLA-Based Simulation, *9th IEEE International Symposium on Distributed Simulation and Real-Time Applications*. Montreal, Canada. October 10-12, 2005. pp.148-157.

Kuhl, F., R. Weatherly, and J. Dahmann. 1999. *Creating Computer Simulation Systems: An Introduction to the High Level Architecture.* Prentice Hall PTR.

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), pp.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*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Linn, R. J., C. S. Chen, and J. A. Lozan. 2002. Development of a Distributed Simulation Model for the Transporter Entity in a Supply Chain Process. *Proceedings of the 2002 Winter Simulation Conference.* Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

McLean, C., F. Riddick, and Y.T. Lee. 2005. An Architecture and Interfaces for Distributed Manufacturing Simulation. *SIMULATION: Transactions of the Society for Modeling and Simulation International,* 81(1), pp.126-138.

Quadt, D., and H. Kuhn. 2001. Lot Sizing and Scheduling in Semiconductor Assembly - A Hierarchical Planning Approach. *Proceedings of the International Conference on Modeling and Analysis of Semiconductor Manufacturing (MASM 2002)*. Tempe AZ, USA. April 10-12, 2002. pp. 211-216, 2002.

Sivakumar, A.I., C.S. Chong, and C.C. Wei. 2001. Multi-Objective Scheduling in Semiconductor Test Manufacturing Using a Simulation Based On-Line Near-Real-Time System. *Proceedings of the International Conference on Modeling and Analysis of Semiconductor Manufacturing (MASM 2002)*. pp. 200-205.

Strassburger, S., G. Schmidgall, and S. Haasis. 2003. Distributed Manufacturing Simulation as an Enabling Technology for the Digital Factory. *Journal of Advanced Manufacturing Systems (JAMS)*. 2(1), pp.111-126.

Taylor, S. J. E., P. Lendermann, R. J. Paul, S. W. Reichenthal, S. Strassburger, and S. J. Turner. 2004. Future Challenges in Modeling Methodology. *Proceedings of the 2004 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Taylor, S .J. E., L. Bohli, X. G. Wang, S. J. Turner, and J. Ladbrook. 2005. Investigating Distributed Simulation at the Ford Motor Company. *Proceedings of the Ninth IEEE International Symposium on Distributed Simulation and Real-Time Applications*. IEEE Computer Society. pp. 139-147.

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

Verbraeck, A. 2004. Applying Distributed Simulation for Analyzing Port Extensions. Invited presentation at the *Annual Conference of the Operations Research Society of South-Africa (ORSSA'04)*. Cape Town, South-Africa. September 6-8, 2004.

Wang, X. G., S. J. Turner, M.Y.H. Low, and B. P. Gan. 2005a. Optimistic Synchronization in HLA Based Distributed Simulation. *SIMULATION: Transactions of the Society for Modeling and Simulation International.* 81(4), pp.279-291.

Wang, K., S. Xu, and L. F. McGinnis. 2005b. Time Management in Distributed Factory Simulation, a Case Study Using HLA. *Proceedings of the 2005 Winter Simulation Conference*. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

AUTHOR BIOGRAPHY

PETER LENDERMANN is a Senior Scientist in the Planning and Operations Management Group at Singapore Institute of Manufacturing Technology (SIMTech) where he is leading the research efforts looking at simulation-based decision support technologies for asset-intensive businesses. He is also a Principal Investigator in the Integrated Manufacturing and Service Systems (IMSS) Thematic Strategic Research Programme of the Science and Engineering Research Council (SERC) at A*STAR in Singapore. He holds a concurrent appointment as Adjunct Associate Professor at the Department of Industrial and Systems Engineering at the National University of Singapore. Before joining SIMTech, he was a Managing Consultant with a German consultancy firm where his focus was on business process re-engineering as well as supply chain management. 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 e-mail address is <peterl@SIMTech.a-star.edu.sg>.