

RESEARCH ISSUES IN SYMBIOTIC SIMULATION

Heiko Ayd
Stephen John Turner
Wentong Cai
Malcolm Yoke Hean Low

School of Computer Engineering
Nanyang Technological University
Nanyang Avenue, Singapore 639798

ABSTRACT

Symbiotic simulation is a paradigm in which a simulation system and a physical system are closely associated with each other. This close relationship can be mutually beneficial. The simulation system benefits from real-time measurements about the physical system which are provided by corresponding sensors. The physical system, on the other side, may benefit from the effects of decisions made by the simulation system. An important concept in symbiotic simulation is that of the what-if analysis process which is concerned with the evaluation of a number of what-if scenarios by means of simulation. Symbiotic simulation and related paradigms have become popular in recent years because of their ability to dynamically incorporate real-time sensor data. In this paper, we explain different types of symbiotic simulation and give an overview of the state of the art. In addition, we discuss common research issues that have to be addressed when working with symbiotic simulation. While some issues have been adequately addressed, there are still research issues that remain open.

1 INTRODUCTION

Symbiotic simulation is a paradigm which refers to a close relationship between a simulation system and a physical system. It was originally defined at the Dagstuhl seminar on Grand Challenges for Modeling and Simulation in 2002 ([Fujimoto et al. 2002](#)). The simulation system benefits from real-time measurements about the physical system which are provided by corresponding sensors. The physical system, on the other side, may benefit from the effects of decisions made by the simulation system. An important concept in symbiotic simulation is that of what-if analysis. Multiple what-if scenarios, representing alternative decision making scenarios, for instance, are evaluated by means of simulation. Symbiotic simulation is closely related to dynamic data-driven application systems (DDDAS) ([National Science Foundation 2005](#)). DDDAS emphasises the ability of applications to steer the measurement process. While the simulation system in a symbiotic simulation system is also data-driven using real-time sensor data, it does not necessarily steer the measurement process. In addition, while DDDAS is more generally concerned with data-driven applications, symbiotic simulation refers more specifically to data-driven simulations about a physical system. Although the paradigms are not identical, they overlap in some respects. As a result, some symbiotic simulation systems can also be considered as a DDDAS and vice versa.

Another concept which is closely related to symbiotic simulation is that of on-line simulation. The term symbiotic simulation has been coined in 2002. However, the concepts behind symbiotic simulation have been around for longer. In 1998, Davis described an on-line planning application which makes use of on-line simulation in order to evaluate multiple control policies ([Davis 1998](#)). The architecture of this on-line planning application resembles that of a typical symbiotic simulation system. For example, an “alternative control policy generator” is used to create different policies which are evaluated by means of simulation. Each of these policies can be considered as a different what-if scenario. The results of these simulations are evaluated and a corresponding control feedback to the physical system is created. Essentially, what Davis described as an on-line planning/control application is a symbiotic simulation system. Later, multiple scenarios have also been used in an on-line simulation system for military networks ([Perumalla et al. 2002](#)).

The use of the term on-line simulation is problematic because it has never been properly defined. Sometimes on-line simulation refers to a decision support system which is capable of making decisions in near real-time by evaluating multiple scenarios by means of simulation. On other occasions it refers to a simulation that runs in real-time and in parallel with

a physical system (Kamrani 2007). Confusion regarding the term on-line simulation was one of the reasons why the terms symbiotic simulation and what-if analysis have been coined. The term symbiotic simulation emphasises the close, interdependent relationship between a physical system and a simulation system. In addition, what-if analysis highlights the use of multiple what-if scenarios, which is an important concept in symbiotic simulation. Therefore, symbiotic simulation does not refer to a specific type of simulation, such as real-time simulation or discrete event simulation, it is rather an umbrella term which refers to a class of systems in which a number of independent simulations are employed to analyse alternative scenarios regarding a physical system by means of simulation. These simulations make use of information about the physical system provided in real-time by sensors. In addition, the result of a what-if analysis can be used to influence the physical system.

Symbiotic simulation can be applied in a large variety of domains and for different purposes. An important factor is the complexity of the physical system. For some applications, the physical system is very complex and cannot be effectively analysed without simulations. For example, the semiconductor manufacturing process is very complex and involves a large number of processing steps and a variety of different machines. Changes in the settings of the various tools in a semiconductor manufacturing plant (fab) can significantly affect the process flow and thus the performance of the entire fab. For operational decision making it is important to make decisions on time. Manually evaluating alternative decisions is difficult. Symbiotic simulation is capable of automatically evaluating what-if scenarios in a reasonable period of time and thus represents a solution to this kind of problem.

There is an optimisation process at the heart of the what-if analysis for some types of symbiotic simulation systems (e.g., symbiotic simulation for decision support). However, there is more to symbiotic simulation than just optimisation. In contrast to mathematical optimisation, the objective function is represented by a model and evaluated by means of simulation. Executing a simulation is typically significantly more time consuming than evaluating a mathematical objective function. This is particularly true if the simulation model is of a stochastic nature, requiring multiple simulation replications. Once a solution has been found, it may need to be implemented in the physical system, i.e., there is a feedback to the physical system. In addition, real-time data has to be incorporated for high-fidelity simulations. While these issues are typically beyond the scope of optimisation, they play an important role in the context of symbiotic simulation.

Some of the research issues in symbiotic simulation have already been addressed in the various works that can be found in the literature. However, there are still open research issues. In this paper, we give an overview of the state of the art and discuss some of the most important research issues in symbiotic simulation.

2 RELATED WORK

Symbiotic simulation has been applied in the context of various application domains. Low et al. describe a symbiotic simulation system for semiconductor manufacturing backend operations (Low et al. 2005). Upper and lower limits for queue lengths are used to make outsourcing decisions. Optimal upper and lower limits are determined by analysing different settings by means of what-if simulations. Each what-if simulation represents a different set of limits. Another symbiotic simulation system in semiconductor manufacturing has been described by Aydt et al. which is concerned with the control of a particular kind of tool (Aydt et al. 2008). The optimal tool configuration is determined by analysing each possible configuration by means of a what-if simulation.

Symbiotic simulation can also be used for business process optimisation. Low et al. describe a symbiotic simulation system which can be used to monitor real-world operations and optimise a business workflow in the context of high-tech manufacturing and service networks (Low et al. 2007). They consider aerospace spare part logistics in which fillrate performance is considered as the key performance metric. The performance is observed and a what-if analysis is triggered if the fillrate falls below a certain threshold. Various what-if simulations are created and initialised with information from the physical system to reflect its current state. In addition, each what-if simulation uses a different business workflow. The goal is to identify the workflow which results in the best fillrate performance. Once identified, the most appropriate business workflow is implemented in the physical system.

Another domain where symbiotic simulation has been applied is the control of unmanned aerial vehicles (UAVs). Kamrani and Ayani describe how symbiotic simulation can be used for path planning (Kamrani and Ayani 2007). In their application, the path of a UAV is frequently adapted. What-if simulations are used to evaluate alternative paths. Mitchell and Yilmaz have also described a symbiotic simulation system in the context of UAV control (Mitchell and Yilmaz 2008). They explain how symbiotic adaptive multi-simulation (SAMS) can be used for real-time decision making under uncertainty. Each what-if simulation is concerned with a combined model ensemble which combines parameters regarding environmental conditions, not fully known due to uncertainty, and possible system configurations. Unlike other symbiotic simulation applications, SAMS therefore analyses how a particular system configuration performs under different possible environmental conditions.

Symbiotic simulation is often used for decision making or control. For example, the various applications mentioned above use symbiotic simulation for a decision making process. However, symbiotic simulation can also be used for other purposes. For example, a DDDAS which can also be considered as a symbiotic simulation system is described by Kennedy et al. (Kennedy et al. 2007, Kennedy and Theodoropoulos 2006). This application is concerned with the moving behaviour of households depending on affordability, crime level, and other factors. The purpose of the symbiotic simulation system is not to make decisions regarding the physical system but rather to identify inaccuracies in the model. Model revision is suggested if the discrepancy between the simulated and the actual behaviour of the physical system is too large.

Another DDDAS which can be considered as a symbiotic simulation system is WIPER, an emergency response system, described by Madey et al. (Madey et al. 2006, Madey et al. 2007). The communication pattern of millions of mobile phone users are monitored. Anomalies in these patterns are detected and analysed. Symbiotic simulation is used to classify the cause of an anomaly by testing various hypotheses by means of simulations. Once the cause for the anomaly has been determined, the corresponding model can be used to predict how the anomaly will continue to unfold in the near future. This application is another example of a symbiotic simulation system that does not influence the physical system. Instead, it determines a reference model which accurately describes the physical system.

In our earlier work we used symbiotic simulation model validation in the context of a radiation detection application (Aydt et al. 2009). In this application, it is assumed that a moving radiation source (e.g., “dirty bomb” carried by a terrorist) has to be detected. The characteristics of the radiation source and its location are unknown. A highly accurate model of the environment is used to estimate the radioactivity at certain locations in the environment given a particular kind of radiation source and its location. These estimated values are compared with measured values. Similar to WIPER, different hypotheses regarding the radiation source are evaluated by means of simulations. The hypothesis which produces the closest results to the actual measurements, can be considered as the likely location of the radiation source in the real environment.

There is a variety of symbiotic simulation applications that have already been described in the literature. This also includes some applications in the context of DDDAS. A generic framework for this kind of applications would make it significantly easier for researchers and developers to make use of concepts such as symbiotic simulation in their applications. The need for a generic framework for symbiotic simulation systems has been articulated by Huang et al. (Huang et al. 2006). An agent-based implementation for such a generic framework has been developed by us in our previous work (Aydt et al. 2008a). This framework provides generic standard implementations for the various functional components that can usually be found in a symbiotic simulation system. The framework has been designed with requirements regarding applicability, extensibility, and scalability in mind.

3 DIFFERENT TYPES OF SYMBIOTIC SIMULATION SYSTEMS

The original definition of symbiotic simulation is limited to “mutualism” which is a specific form of symbiosis in which all partners benefit from each other. However, there are also other forms of symbiosis, such as “commensalism” and “parasitism” (Douglas 1994). These are not considered by the original definition of symbiotic simulation. Commensalism and parasitism refer to a symbiosis in which some partners benefit while others are not affected or are negatively affected, respectively. An extended definition of symbiotic simulation and the resulting classes of symbiotic simulation systems have been discussed in our earlier work (Aydt et al. 2008b). In this extended definition, we distinguish between closed-loop and open-loop symbiotic simulation systems. In a closed-loop symbiotic simulation system there is a control feedback from the simulation system to the physical system. In an open-loop symbiotic simulation system there is no such feedback. Since the simulation system does not affect the physical system, all open-loop symbiotic simulation systems are therefore of commensalistic nature. A closed-loop symbiotic simulation system, on the other hand, affects the physical system either positively or negatively, which can be referred to as either mutualism or parasitism, respectively.

Closed-loop symbiotic simulation systems can be further characterised by the way they affect the physical system. We distinguish between direct control and indirect control leading to the two classes of closed-loop symbiotic simulations: a symbiotic simulation control system (SSCS) and symbiotic simulation decision support system (SSDSS), respectively. The difference is that an SSCS is directly influencing the physical system by means of actuators while an SSDSS proposes decisions to an external decision maker which can then implement them into the physical system. Apart from this subtle difference, both classes are the same as they make use of what-if simulations to evaluate alternative decisions and their possible effects on the physical system. The purpose of an SSCS/SSDSS is to identify the optimal decision.

Closed-loop symbiotic simulation systems are concerned with decision making. In comparison, the purposes of open-loop symbiotic simulation systems are more diverse. We distinguish between open-loop symbiotic simulation systems for forecasting, model validation, and anomaly detection.

Table 1: Overview of various classes of symbiotic simulation systems

Class	Purpose	Open/Closed Loop	Meaning of What-if Scenarios	Type of Symbiosis
SSCS	Control of a physical system	Closed (Direct)	Control alternatives	Mutualism/Parasitism
SSDSS	Support of an external decision maker	Closed (Indirect)	Decision alternatives	Mutualism/Parasitism
SSFS	Forecasting of a physical system	Open	Different assumptions for environmental conditions	Commensalism
SSMVS	Validation of a simulation system	Open	Alternative models or different parameters	Commensalism
SSADS	Detection of anomalies either in the physical system or in the simulation model	Open	Reference model only	Commensalism

A symbiotic simulation forecasting system (SSFS) is concerned with forecasting the behaviour of a physical system. Several what-if simulations are considered of which each is reflecting different assumptions regarding the environment of the physical system. The purpose of an SSFS is therefore to explore possible predictions by taking uncertainty about the environment into account.

A symbiotic simulation model validation system (SSMVS) is concerned with the identification of a reference model of the physical system. For this purpose, different models and/or model parameters are evaluated by means of simulation. The behaviour of each simulation is compared with the actual behaviour of the physical system. The outcome of each of these hypotheses is compared with sensor data from the physical system. The hypothesis which is closest to the measurements is considered as the reference model and can be used for further forecasting or decision making.

A symbiotic simulation anomaly detection system (SSADS) is concerned with the detection of anomalies either in the simulation model or in the physical system. A reference model of the physical system is continuously compared with measurements from the physical system. Discrepancies between the simulated and actual behaviour of the physical system are considered as anomalies. Anomaly detection can be used for detecting inaccurate models in order to trigger an update process or to identify problems in the physical system. In the latter case, it is assumed that the model is correct and the physical system should behave accordingly. If it does not, then this indicates a problem in the physical system.

An overview of the various classes of symbiotic simulation systems is illustrated in Table 1. This table has been adopted with minor changes from Table 1.2 in (Ayd et al. 2008a).

4 WHAT-IF ANALYSIS RELATED ISSUES

An essential concept of symbiotic simulation is the what-if analysis (WIA) process. It is responsible for evaluating a set of what-if scenarios using simulation. The purpose of the WIA process depends on the class of symbiotic simulation system. For example, the purpose of the WIA process in an SSCS/SSDSS or SSMVS is to identify the optimal decision or reference model among a potentially large number of alternatives, respectively. The purpose of the SSADS is to identify anomalies either in the physical system or in the simulation model. Only a single scenario, using the current reference model, is therefore considered by an SSADS. The purpose of the WIA process in an SSFS is to take account of different possible environmental conditions.

4.1 Triggering

In general, there are various ways of triggering the WIA process. Three triggering mechanisms have been used in various works that can be found in the literature: reactive, preventive, and pro-active triggering.

A reactive WIA is carried out if a triggering condition can be observed in the physical system. Typically, these triggering events are potentially critical conditions that require immediate action. The purpose of the reactive WIA is to find a solution in order to recover from that critical condition. For example, in the context of semiconductor manufacturing reactive WIA has been used to recover from performance problems due to bottleneck situations in (Low et al. 2005) and (Ayd et al. 2008).

Table 2: Overview of the various WIA types

WIA Type	Purpose	Triggering
Reactive WIA	Problem recovery	Observed triggering condition
Preventive WIA	Problem prevention	Forecasted triggering condition
Pro-active WIA	Continuous performance improvement	Periodically

A preventive WIA is similar to a reactive WIA as both are triggered based on the observation of a triggering condition. However, in contrast to reactive WIA, preventive WIA relies on the observation of a triggering condition in the forecast of the physical system. The purpose of the preventive WIA is thus to prevent a critical condition from occurring or at least to reduce its negative effects. Preventive WIA has been introduced in the context of semiconductor manufacturing (Aydı et al. 2008). The applicability of preventive WIA is limited to conditions that can be forecasted. For example, a machine breakdown cannot be forecasted accurately and preventive WIA cannot be applied.

Both reactive and preventive WIA rely on the notion of a triggering condition. This is not the case for pro-active WIA which is triggered periodically. The purpose of pro-active WIA is to continuously improve the performance of a physical system or to maintain certain performance targets. For example, the UAV path planning application, described by Kamrani and Ayani (Kamrani and Ayani 2007), makes use of pro-active WIA. In this applications, the WIA process is responsible for updating the flight path of the UAV. For this purpose, the WIA is triggered in fixed time intervals and used to evaluate alternative paths. An overview of the various WIA types is illustrated in Table 2.

An important characteristic of symbiotic simulation systems is their real-time requirements. This is particularly important in the context of decision support and control, where the simulation system has to respond within a certain period of time. Regardless of the triggering mechanism, the WIA process has to take account of the time that is required to perform the WIA process. In addition, for an SSCS/SSDSS it is important to consider the time required to create feedback and implement a decision into the physical system. While the latter might be negligible for some applications (e.g., SSCS with corresponding actuators to exercise direct control), it might be significant for others. For example, in the case of an SSDSS, the simulation system cannot directly influence the physical system. Instead, it proposes a decision to an external decision maker. This external decision maker may need some time to decide whether the suggested decision should be executed or not. It is therefore important to specify the earliest point of time when a decision can be implemented in the physical system. This point of time has to be taken into account when creating different decision making scenarios.

If the WIA process does not take a long period of time, it might be possible to ignore this issue. For example, in (Aydı et al. 2008) we argued that the entire WIA process takes approximately 1–6 minutes to finish which is sufficiently short in the context of the semiconductor manufacturing application discussed. On the other hand, for the radiation detection SSMVS (Aydı et al. 2009), it is necessary to deal with this issue accordingly. Radiation detection of a moving target is performed in multiple stages. The first stage is concerned with the detection of certain characteristics of the radiation source in addition to the location. Once the characteristics are known, the optimisation process can be simplified. The WIA process in the second stage is only concerned with continuously tracking the location of the target. Since the optimisation problem is simplified for this stage, the WIA process can be executed much faster.

An issue which is specific to preventive WIA is the need for an error handling method. Preventive WIA relies on a forecast which may not be accurate. Therefore, it is possible that the WIA process may either not be triggered at all, because no triggering condition was forecasted, or triggered unnecessarily because a false triggering was forecasted. These two kinds of errors are referred to as a type I error and a type II error, respectively (Aydı et al. 2008). An appropriate error handling method is required to deal with these two errors. For example, one possible error handling method is to use a reactive WIA as backup in case the preventive WIA fails. However, more advanced strategies have to be developed in order to make preventive WIA more reliable. In general, not much work has been done on making the WIA process more robust against errors due to uncertainty. The only work of which we are aware that considers uncertainty in the WIA process is the work by Mitchell and Yilmaz (Mitchell and Yilmaz 2008).

4.2 Efficiency and Effectiveness

Depending on the class and the purpose of the symbiotic simulation system, there can be a potentially large number of what-if scenarios that have to be simulated. While this issue does not apply to an SSADS, which is concerned with a single simulation only, it applies to all other classes of symbiotic simulation systems. The major issue in an SSCS/SSDSS and an SSMVS is not simulation but rather optimisation. The WIA process of an SSCS/SSDSS is essentially an optimisation process which aims to maximise the outcome of decision making by identifying and implementing the optimal decision. Similarly, the WIA process of an SSMVS can be considered as an optimisation process which aims to minimise the error between the simulated and actual behaviour of the physical system by identifying the reference model which describes the physical system most accurately.

In this context, the efficiency and the effectiveness of the WIA process are important. Efficiency refers to the ability of the WIA process to finish in time. For example, in the context of an SSCS/SSDSS, it is important that the WIA process is able to make decisions in time. In practise, it is often the case that solutions which are available too late might not be useful any longer. This can also be true for an SSMVS and is not limited to decision support and control. The time constraints for the WIA process depend on the application context. For example, in a semiconductor manufacturing environment, decisions regarding tool setups have to be made within a few minutes or hours. On the other hand, in a supply chain network, it might be sufficient to make decisions regarding shipping policies within a day.

There are various techniques which might be suitable for improving the efficiency of the WIA process. For example, simulation cloning is a technique that reduces redundancies between multiple simulation scenarios that are very similar (Hybinette and Fujimoto 1997). However, as explained by Hybinette and Fujimoto (Hybinette and Fujimoto 2001, Hybinette and Fujimoto 2002), simulation cloning can be inefficient when facing rapid spreading, i.e., when multiple scenarios diverge rapidly. In this case, there is only little or no redundancy left between these scenarios. In the worst case, simulation cloning is actually slightly worse than traditional simulation replication because of the overhead required by the cloning mechanism. In addition, simulation cloning is often not supported by commercial-of-the-shelf (COTS) simulation packages. Applying this technique is therefore not always possible unless a specialised simulation package is used. However, this is often not desirable for various reasons.

Some simulation models are stochastic and each simulation run may produce a slightly different result because of pseudo random numbers, generated using different seeding values for each replication. This makes it necessary to repeatedly simulate a what-if scenario in order to obtain enough data to derive statistically meaningful conclusions about the performance. The number of required replications depends on the simulation model. The need for simulation replications increases the amount of time required to simulate what-if scenarios significantly. In order to improve the efficiency of the WIA process it might therefore be possible to use techniques that reduce the number of replications. For example, Lee et al. describe a method which reduces the number of simulation replications and may be applicable to symbiotic simulation (Lee et al. 2004).

Another way to improve the efficiency of the WIA process is to make use of computing clusters. What-if scenarios are independent from each other and can therefore be simulated in parallel if enough computing resources are available. Another possible way to improve the efficiency is to make use of parallelisation techniques to partition the simulation model. Whether or not such techniques can be applied depends on various factors, such as the total number of what-if scenarios, the memory characteristics of the simulation, and available hardware resources. For example, if there is a very large number of what-if scenarios that have to be simulated (i.e., significantly more than there are CPUs available), then parallel simulation of scenarios is generally better than parallelisation of the simulation itself. Another issue is the memory characteristics of a single simulation run. If the memory consumption is very high, then it might not be possible to execute many what-if scenarios concurrently. The overall objective is to perform a what-if analysis in the shortest time possible. Depending on the application requirements, available computing resources should be used in the most appropriate way to achieve this objective. An introduction to parallel and distributed simulation can be found in (Fujimoto 2000).

For some applications, there is only a small number of possible what-if scenarios. In this case it is possible to perform an exhaustive search and evaluate all possible what-if scenarios in order to find the best solution. For example, the symbiotic simulation system described in (Aydt et al. 2008) is concerned with finding the best configuration for a semiconductor manufacturing tool. In total, there are only 256 possible scenarios. Similarly, the symbiotic simulation system for business process re-engineering described in (Low et al. 2007), considers only three alternative process flows. In both cases, the search space is very small and exhaustive search can be applied.

Techniques such as simulation cloning and replication reduction are unfortunately only useful for fine tuning the efficiency of the WIA process. If the number of possible what-if scenarios is considerably large, then these techniques are not sufficient to achieve an efficient WIA process. An exhaustive search will take a very long time before it converges. Therefore, more advanced search methods are required. For example, the symbiotic simulation system described in (Mitchell and Yilmaz

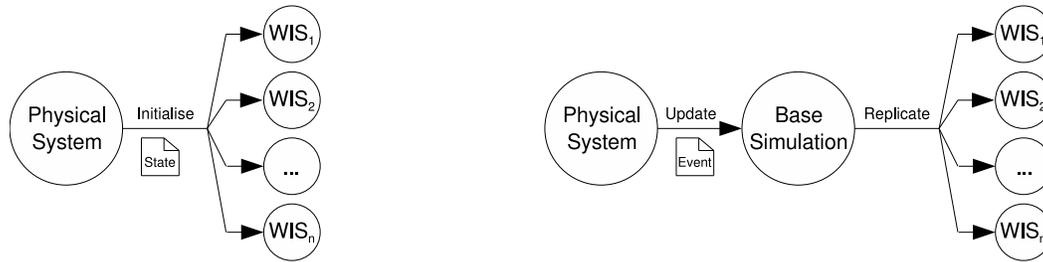


Figure 1: Overview of the state collection method (left) and the base simulation method (right). What-if simulations (WIS) are either initialised with information about the state of the physical system or replicated from the base simulation, respectively.

2008) makes use of a genetic algorithm. Similarly, the radiation detection application, described in (Ayd et al. 2009), uses a hybrid genetic algorithm in combination with a honey bee colony optimisation algorithm.

While an exhaustive search (if applicable) is capable of always finding the best what-if scenario, this is not necessarily the case for meta heuristics such as genetic algorithms. The effectiveness of the WIA process is concerned with the ability of the WIA process to find the optimal scenario. A search heuristic, such as a genetic algorithm, may be able to find the optimum without evaluating all possible what-if scenarios. On the other hand, many meta heuristics do not guarantee to find an optimum. However, this might not be a problem for practical applications as long as the solutions found are good solutions. An important issue in symbiotic simulation decision support and control is to be able to make good decisions in time. A reasonably good decision, found just in time, is often more useful than an optimal decision found too late.

5 MODEL RELATED ISSUES

For symbiotic simulation applications, it is important to use high-fidelity models in order to make meaningful decisions regarding the physical system or forecast the effects of various environmental conditions on the physical system. In general, there are two related issues that have to be considered: initialisation of what-if simulations and model selection.

5.1 Initialisation of What-if Simulations

In symbiotic simulation, each what-if simulation starts by reflecting the current state of the physical system. Two methods have been described in the literature which can be used for this purpose: the state collection method and the base simulation method. Figure 1 illustrates these methods.

The first method initialises each what-if simulation with the current state of the physical system. This requires the simulation system to be able to retrieve all necessary information about the state of the physical system. For example, this method is assumed in the semiconductor manufacturing application in (Ayd et al. 2008). Retrieving the entire state of the physical system may take some time. This delay might represent a serious problem in applications where the symbiotic simulation system does not have much time to react on events. One possible way of implementing the initialisation method is to periodically collect state information of the physical system in a database. If required, the latest complete state of the physical system is used to initialise the what-if simulations. This avoids the need to collect state information when a WIA process is triggered. However, in this case it is necessary to fast forward the what-if simulation from this previous state until it reflects the state of the current time.

The second method uses a base simulation and has been described in (Low et al. 2005). A base simulation emulates the physical system and is realised as a real-time simulation, i.e., it is paced in wallclock time (Fujimoto 2000). This means, that the emulator advances with the same pace as the physical system. When the WIA process is triggered, the base simulation is replicated multiple times and configured to reflect the specific settings that are relevant for the what-if scenario. The advantage of this method is that it is relatively easy to create what-if simulations which reflect the physical system in its current state. However, the base simulation has to be continuously updated. Every event that changes the state of the physical system has to be injected into the base simulation so that it can advance in accordance with the physical system. Events which are not injected may cause the base simulation to diverge from the actual behaviour of the physical system.

The advantage of the state collection method is that there is no need to maintain a base simulation. This reduces the overall complexity of the simulation system. The problem of delays when retrieving a complete state of the physical system can be alleviated to some extent by using the latest available state of the physical system, as described above. In addition,

even if a base simulation is used, it is likely that it has to be re-initialised after some time. This is due to the fact that a model is never perfectly accurate and the base simulation will therefore eventually diverge. Therefore, even when using a base simulation it is necessary to use it in combination with the state collection method. This questions the usefulness of a base simulation.

The use of a base simulation is still justified for some applications in combination with an SSADS. In this case, the base simulation uses the current reference model to emulate the behaviour of the physical system. The simulated behaviour and the actual behaviour of the physical system are compared in real-time. Detection of an anomaly, i.e., a significant discrepancy between the simulation and the actual behaviour, can be interpreted in two ways. It might indicate that the base simulation has to be re-initialised. Depending on the application, occasional re-initialisation is probably necessary and thus normal. However, frequent re-initialisation might indicate problems with the model, i.e., anomalies might indicate that the model is not sufficiently accurate to reflect the behaviour of the physical system. The higher the re-initialisation frequency, the lower the predictive capabilities of the model. In this case, it might be required to modify the model or use another one.

When using the base simulation method in combination with an SSADS, it is therefore possible to detect problems with the model of the physical system. This is not possible when using the state collection method. This represents an important advantage of the base simulation method over the state collection method. Using a base simulation might therefore be preferable in some cases. However, as of the time of writing, such a base simulation has not yet been developed and tested. Although the idea is mentioned in (Low et al. 2005), a base simulation was not actually realised. In general, more research is needed, including the use of a real physical system.

5.2 Model Selection

There might be multiple models that describe the behaviour of the physical system. Depending on the current operating conditions, one model might be more suitable to describe the behaviour than another. In such a case it is necessary to select the most appropriate model. For example, WIPER is an emergency response system which is capable of classifying an emergency event and predicting how it will unfold in the near future (Madey et al. 2006). For this purpose, a number of possible explanations are evaluated by means of simulations (Madey et al. 2007). Once the best explanation has been found, it can be used for forecasting purposes. In terms of symbiotic simulation, WIPER is an SSMVS with the purpose to identify the reference model which explains the behaviour of the physical system. Once found, this reference model is used for forecasting or decision making purposes. For some applications it might be required to select an appropriate model to represent the physical system in its current state. Symbiotic simulation model validation represents one possible approach to this issue. However, depending on the application context other approaches might be required. In general, dynamic model identification in symbiotic simulation is an open research issue.

6 SUMMARY

Symbiotic simulation has been applied in a number of application domains and some of the research issues have already been addressed by previous work. However, there are still open research issues which have not yet been addressed. For example, only little work has been done on robust what-if analysis which takes account of uncertainty. In addition, it is important to understand that the WIA process in an SSCS/SSDSS and an SSMVS is essentially an optimisation process which makes it necessary to apply appropriate methods in order to keep the total number of scenario evaluations small. Efficiency is an important aspect in this context because each scenario is evaluated by means of simulation. This can be very time intensive and performing an exhaustive search is therefore not feasible for many real-life problems. Sophisticated search methods are therefore required. In addition, more advanced simulation techniques are needed. Simulation cloning is an attempt to make the what-if analysis more efficient by removing redundancies. However, as of today simulation cloning has many limitations which have to be resolved first. Another major issue in symbiotic simulation is the initialisation of what-if simulations and related modeling issues. Only very little work has been done so far.

In this paper, we have given an overview of the state of the art and discussed some of the most important research issues in symbiotic simulation. For the future, we expect to see a variety of symbiotic simulation applications. In particular, we expect to see more applications that involve real physical systems.

REFERENCES

Aydt, H., S. J. Turner, W. Cai, and M. Y. H. Low. 2008a. An agent-based generic framework for symbiotic simulation systems. In *Agents, Simulation and Applications*, ed. A. M. Uhrmacher and D. Weyns. Taylor & Francis.

- Aydt, H., S. J. Turner, W. Cai, and M. Y. H. Low. 2008b. Symbiotic simulation systems: An extended definition motivated by symbiosis in Biology. In *Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation*, 109–116.
- Aydt, H., S. J. Turner, W. Cai, M. Y. H. Low, and R. Ayani. 2009. Symbiotic simulation model validation for radiation detection applications. In *Proceedings of the 23rd Workshop on Principles of Advanced and Distributed Simulation*, 11–18.
- Aydt, H., S. J. Turner, W. Cai, M. Y. H. Low, P. Lendermann, and B. P. Gan. 2008. Symbiotic simulation control in semiconductor manufacturing. In *Proceedings of the International Conference on Computational Science*, 26–35.
- Aydt, H., S. J. Turner, W. Cai, M. Y. H. Low, P. Lendermann, B. P. Gan, and R. Ayani. 2008. Preventive what-if analysis in symbiotic simulation. In *Proceedings of the 2008 Winter Simulation Conference*, ed. S. Mason, R. Hill, L. Moench, and O. Rose, 750–758: Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Davis, W. 1998. On-line simulation: Need and evolving research requirements. In *Handbook of Simulation*, ed. J. Banks, 465–516. Wiley-Interscience.
- Douglas, A. 1994. *Symbiotic Interactions*. Oxford University Press.
- Fujimoto, R., D. Lunceford, E. Page, and A. U. (editors). 2002, August. Grand challenges for modeling and simulation: Dagstuhl report. Technical Report 350, Schloss Dagstuhl. Seminar No 02351.
- Fujimoto, R. M. 2000. *Parallel and distributed simulation systems*. Wiley Series on Parallel and Distributed Computing. New York, NY, USA: John Wiley & Sons, Inc.
- Huang, S. Y., W. Cai, S. Turner, W. J. Hsu, S. Zhou, M. Y. H. Low, R. Fujimoto, and R. Ayani. 2006. A generic symbiotic simulation framework. In *Proceedings of the 20th Workshop on Principles of Advanced and Distributed Simulation*, 131–131.
- Hybinette, M., and R. Fujimoto. 1997. Cloning: A novel method for interactive parallel simulation. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K. Healy, D. Withers, and B. Nelson, 444–451: Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Hybinette, M., and R. M. Fujimoto. 2001, October. Cloning parallel simulations. *ACM Transactions on Modeling and Computer Simulation* 11 (4): 378–407.
- Hybinette, M., and R. M. Fujimoto. 2002. Scalability of parallel simulation cloning. In *Proceedings of the 35th Annual Simulation Symposium*, 275–282.
- Kamrani, F. 2007. Using on-line simulation in UAV path planning. Licentiate Thesis in Electronics and Computer Systems, KTH, Stockholm, Sweden.
- Kamrani, F., and R. Ayani. 2007, October. Using on-line simulation for adaptive path planning of UAVs. In *Proceedings of the 11th IEEE International Symposium on Distributed Simulation and Real-time Applications*, 167–174. Chania, Greece.
- Kennedy, C., and G. K. Theodoropoulos. 2006. Intelligent management of data driven simulations to support model building in the social sciences. In *Proceedings of the International Conference on Computational Science*, 562–569.
- Kennedy, C., G. K. Theodoropoulos, V. Sorge, E. Ferrari, P. Lee, and C. Skelcher. 2007. AIMSS: An architecture for data driven simulations in the social sciences. In *Proceedings of the International Conference on Computational Science*, 1098–1105.
- Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2004. Optimal computing budget allocation for multi-objective simulation models. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. Ingalls, M. Rossetti, J. Smith, and B. Peters, 586–594: Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Low, M. Y. H., K. W. Lye, P. Lendermann, S. J. Turner, R. T. W. Chim, and S. H. Leo. 2005. An agent-based approach for managing symbiotic simulation of semiconductor assembly and test operation. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, 85–92. New York, NY, USA: ACM Press.
- Low, M. Y. H., S. J. Turner, D. Ling, H. L. Peng, L. P. Chai, P. Lendermann, and S. Buckley. 2007. Symbiotic simulation for business process re-engineering in high-tech manufacturing and service networks. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. Tew, and R. Barton, 586–576: Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Madey, G. R., A.-L. Barabási, N. V. Chawla, M. Gonzalez, D. Hachen, B. Lantz, A. Pawling, T. Schoenharl, G. Szabó, P. Wang, and P. Yan. 2007. Enhanced situational awareness: Application of DDDAS concepts to emergency and disaster management. In *Proceedings of the International Conference on Computational Science*, 1090–1097.
- Madey, G. R., G. Szabó, and A.-L. Barabási. 2006. WIPER: The integrated wireless phone based emergency response system. In *Proceedings of the International Conference on Computational Science*, 417–424.

- Mitchell, B., and L. Yilmaz. 2008, December. Symbiotic adaptive multisimulation: An autonomic simulation framework for real-time decision support under uncertainty. *ACM Transactions on Modeling and Computer Simulation* 19 (1): 2–31.
- National Science Foundation 2005. DDDAS: Dynamic data driven applications systems. Program Solicitation 05-570, <http://www.nsf.gov/pubs/2005/nsf05570/nsf05570.htm>.
- Perumalla, K., R. Fujimoto, T. McLean, and G. Riley. 2002. Experiences applying parallel and interoperable network simulation techniques in on-line simulations of military networks. In *Proceedings of the 16th Workshop on Parallel and Distributed Simulation*, 88–95.

AUTHOR BIOGRAPHIES

HEIKO AYDT is a Research Associate and PhD student in the School of Computer Engineering at Nanyang Technological University (NTU) in Singapore. He received his MSc from the Royal Institute of Technology (KTH) in Stockholm. His current research interests include: Parallel and Distributed Simulation, Simulation-based Optimisation, and Evolutionary Computing. His e-mail address is [<aydt@ntu.edu.sg>](mailto:aydt@ntu.edu.sg).

STEPHEN JOHN TURNER is Professor of Computer Science and Head of the Computer Science Division in the School of Computer Engineering at Nanyang Technological University (NTU) in Singapore. 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, Grid Computing, High Performance Computing and Multi-Agent Systems. He is steering committee chair of the Principles of Advanced and Distributed Simulation (PADS) conference and an area editor for ACM Transactions on Modeling and Computer Simulation (TOMACS). His e-mail address is [<assjturner@ntu.edu.sg>](mailto:assjturner@ntu.edu.sg).

WENTONG CAI is a Professor in the Division of Computer Science, School of Computer Engineering at Nanyang Technological University (NTU) in Singapore. He is also the Director of the Parallel and Distributed Computing Centre. He received his PhD in Computer Science from University of Exeter (UK) in 1991. His current research interests include: Parallel and Distributed Simulation, Multi-Agent Systems, and Grid and Cluster Computing. He is an associate editor of the ACM Transactions on Modeling and Computer Simulation (TOMACS), and editorial board member of the Multiagents and Grid Systems An International Journal. His e-mail address is [<aswtcai@ntu.edu.sg>](mailto:aswtcai@ntu.edu.sg).

MALCOLM YOKE HEAN LOW is an Assistant Professor in the School of Computer Engineering at the Nanyang Technological University (NTU) in Singapore. Prior to this, he was with the Singapore Institute of Manufacturing Technology, Singapore (SIMTech). He received his Bachelor and Master of Applied Science in Computer Engineering from NTU in 1997 and 1999 respectively. He was awarded a Gintic (now SIMTech) Postgraduate Scholarship in 1999. In 2002, he received his D.Phil. degree in Computer Science from Oxford University. His current research interest is in the application of parallel and distributed computing for the modeling, simulation, analysis and optimisation of complex systems. His email address is [<yhlow@ntu.edu.sg>](mailto:yhlow@ntu.edu.sg).