

Author Query Form

Journal Title : Simulation: Transactions of the Society for Modeling and Simulation International [SIM]
Article Number : 454931

Dear Author/Editor,

Greetings, and thank you for publishing with SAGE. Your article has been copyedited, and we have a few queries for you. Please respond to these queries when you submit your changes to the Production Editor.
Thank you for your time and effort.

Please ensure that you have obtained and enclosed all necessary permissions for the reproduction of artistic works (e.g. illustrations, vcharts, maps, other visual material, etc.) not owned by yourself, and ensure that the Contribution contains no unlawful statements and does not infringe any rights of others, and agree to indemnify the Publisher, SAGE Publications Ltd, against any claims in respect of the above warranties and that you agree that the Conditions of Publication form part of the Publishing Agreement.

Please assist us by clarifying the following queries:

Sl. No.	Query
1	AQ: Please insert name of the publishing institution
2	AQ: Please insert name of the institution



A novel parallelization technique for DEVS simulation of continuous and hybrid systems

Federico Bergero¹, Ernesto Kofman¹ and François Cellier²

Abstract

In this paper, we introduce a novel parallelization technique for Discrete Event System Specification (DEVS) simulation of continuous and hybrid systems. Here, like in most parallel discrete event simulation methodologies, the models are first split into several sub-models which are then concurrently simulated on different processors. In order to avoid the cost of the global synchronization of all processes, the simulation time of each sub-model is locally synchronized in a real-time fashion with a scaled version of physical time, which implicitly synchronizes all sub-models. The new methodology, coined Scaled Real-Time Synchronization (SRTS), does not ensure a perfect synchronization in its implementation. However, under certain conditions, the synchronization error introduced only provokes bounded numerical errors in the simulation results. SRTS uses the same physical time-scaling parameter throughout the entire simulation. We also developed an adaptive version of the methodology (Adaptive-SRTS) where this parameter automatically evolves during the simulation according to the workload. We implemented the SRTS and Adaptive-SRTS techniques in *PowerDEVS*, a DEVS simulation tool, under a real-time operating system called the Real-Time Application Interface (RTAI). We tested their performance by simulating three large-scale models, obtaining in all cases a considerable speedup.

Keywords

parallel simulation, real-time systems, discrete event systems, QSS, hybrid systems

1 Introduction

Computer simulation has become a crucial tool for studying, developing, and evaluating mathematical models of systems of different domains (physics, economics, biology, etc.).

Continuous systems are usually represented by sets of differential equations. Since most differential equations lack analytical solutions, they must be discretized to simulate the continuous systems they represent. To this end, most numerical integration methods approximate the differential equations by discretizing the time variable transforming them into sets of difference equations.¹

Some recently developed numerical methods, however, are based on state discretization. These methods, called Quantized State Systems (QSS)¹ transform the originally continuous systems into discrete event systems that can be represented by the Discrete Event System Specification (DEVS) formalism.² The usage of QSS algorithms then allows the DEVS formalism to represent and to simulate continuous and hybrid systems. DEVS models can be easily simulated, but when the models are large, the computational requirements become critical and parallelization techniques must be employed.

Parallel Discrete Event Simulation (PDES) is a technique that can be used for the simulation of large DEVS models in multi-core or cluster environments. In PDES, the model is first split into several sub-systems called physical processes. Then, the processes are simulated concurrently on different logical processors (LPs).

Compared with simulating a large model in a sequential manner, PDES reduces the computational cost but it introduces a new problem related to the need of synchronization between processes. Synchronization between different LPs is required since each process needs to know the results of other processes in order to correctly simulate its own sub-system. If the LPs are not correctly synchronized, the simulation may receive events out of order (with time-

¹Laboratorio de Sistemas Dinámicos, FCEIA - UNR, CIFASIS-CONICET, Rosario, Argentina

²Modeling and Simulation Research Group, ETH, Zürich, Switzerland

Corresponding author:

Federico Bergero, Laboratorio de Sistemas Dinámicos, FCEIA - UNR, CIFASIS-CONICET, 27 de Febrero 210 Bis, S2000EYP Rosario, Argentina.

Email: bergero@cifasis-conicet.gov.ar

stamps lower than the actual simulation time), which can lead to an incorrect result. This is called a causality constraint.

There are many approaches to overcome this problem. One of the earliest was the CMB algorithm proposed by Chandy, Misra, and Bryant.³⁻⁵ CMB implements a conservative synchronization where none of the LPs advances its simulation time until it is safe. A different approach is that proposed by Jefferson,⁶ where an optimistic synchronization mechanism is introduced relaxing the causality constraint. The LPs are allowed to advance their simulation time as fast as they can, and proper actions are taken when inconsistencies are detected. Finally, a technique called NOTIME⁷ explores the effects of not synchronizing the processes at all.

In this article, we present a novel PDES technique, specialized for DEVS approximations of continuous and hybrid system, called Scaled Real-Time Synchronization (SRTS). The basic idea is to synchronize each LP's simulation time with a scaled version of the physical time or wall-clock time. As all LPs are synchronized with the physical time, they are indirectly synchronized against each other.

As our approach is intended for continuous and hybrid systems, events represent continuous trajectories. Thus, synchronization errors provoke numerical errors. Provided that the errors are bounded, they do not invalidate the final results, hence we can relax the causality constraints.

While in SRTS the user must provide the real-time scaling factor as a simulation parameter, we also developed an Adaptive-SRTS (ASRTS) algorithm where the scaling factor is adjusted dynamically depending on the system workload.

We have implemented SRTS and ASRTS in *PowerDEVS*,⁸ a general-purpose open-source DEVS simulator. *PowerDEVS*⁸ implements the complete family of QSS methods on which this work is based.^{1,19-13} As we need a precise synchronization with physical time in order to achieve a precise synchronization between the processes, the implementation runs under a real-time operating system called the Real-Time Application Interface (RTAI).¹⁴

This article is organized as follows. Section 2 introduces the main concepts used in the rest of the paper, and Section 3 describes some related work in the field. Then Sections 4 and 5 introduce and describe the SRTS and ASRTS algorithms, the main contributions of this article. Finally, Section 6 shows some application examples, and Section 7 concludes the article and discusses some lines of future work.

2 Background

In this section, we introduce a few basic concepts and ideas that will be used in the article. We first describe the

DEVS formalism and the QSS methods, a family of numerical integration methods for continuous and hybrid system simulation. Then, we describe the RTAI real-time operating system and the DEVS simulation tool called *PowerDEVS*. Finally, we provide a brief introduction to parallel architectures.

2.1 DEVS

DEVS stands for Discrete EVent System specification, a formalism introduced originally by Bernard P Zeigler.¹⁵

A DEVS model processes an input event trajectory and, according to that trajectory and its own initial conditions, provokes an output event trajectory.

An atomic DEVS model is defined by the following structure:

$$M = (X, Y, S, \delta_{\text{int}}, \delta_{\text{ext}}, \lambda, ta)$$

where:

- X is the set of input event values, i.e. the set of all possible values that an input event can adopt;
- Y is the set of output event values;
- S is the set of state values;
- $\delta_{\text{int}}, \delta_{\text{ext}}, \lambda$ and ta are functions that define the system dynamics.

Each possible state s ($s \in S$) has an associated time advance computed by the time advance function $ta(s)$ ($ta(s) : S \rightarrow \mathbb{R}_0^+$). The time advance is a non-negative real number specifying how long the system remains in a given state in the absence of input events.

Thus, if the state adopts the value s_1 at time t_1 , $ta(s_1)$ units of time later (i.e. at time $t_1 + ta(s_1)$) the system undergoes an internal transition, changing its state from s_1 to s_2 . The new state is calculated as $s_2 = \delta_{\text{int}}(s_1)$. Function δ_{int} ($\delta_{\text{int}} : S \rightarrow S$) is called the internal transition function.

When the state transitions from s_1 to s_2 , an output event is produced with value $y_1 = \lambda(s_1)$. Function λ ($\lambda : S \rightarrow Y$) is called the output function. In this way, the functions ta , δ_{int} , and λ define the autonomous behavior of a DEVS model.

When an input event arrives, the state changes instantaneously. The new state value depends not only on the input event value but also on the previous state value and the elapsed time since the last transition. If the system entered state s_2 at time t_2 and then an input event arrives at time $t_2 + e$ with value x_1 , the new state is calculated as $s_3 = \delta_{\text{ext}}(s_2, e, x_1)$ (note that $ta(s_2) > e$). In this case, we say that the system performs an external transition. Function δ_{ext} ($\delta_{\text{ext}} : S \times \mathbb{R}_0^+ \times X \rightarrow S$) is called the external transition function. No output event is produced during an external transition.

The formalism presented is also called classic DEVS to distinguish it from parallel DEVS,² which consists of an extension of the classic DEVS formalism designed to improve the handling of simultaneous events.

Atomic DEVS models can be coupled. DEVS theory guarantees that the coupling of atomic DEVS models defines new DEVS models (i.e. DEVS is closed under coupling) and, thus, complex systems can be represented by DEVS in a hierarchical way.²

Coupling in DEVS is usually represented through the use of input and output ports. With these ports, the coupling of DEVS models becomes a simple block-diagram construction. Figure 1 shows a coupled DEVS model N that results from coupling models M_a and M_b .

According to the closure property, the model N can itself be interpreted as an atomic DEVS and can be coupled with other atomic or coupled models.

2.2 Quantized state systems

A continuous time system can be written as a set of ordinary differential equations (ODEs):

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \quad (1)$$

where $\mathbf{x} \in \mathfrak{R}^n$ is the state vector and $\mathbf{u} \in \mathfrak{R}^m$ is a vector of known input functions.

The mathematical model (1) can be simulated using a numerical integration method. While conventional integration algorithms are based on time discretization, a new family of numerical methods was developed based on state quantization.¹ The new algorithms, called QSS methods, approximate ODEs like that of Equation (1) by DEVS models.

Formally, the first-order accurate QSS method (QSS1) approximates Equation (1) by

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{q}(t), \mathbf{v}(t)) \quad (2)$$

where each pair of variables q_j and x_j are related by a hysteretic quantization function.

The presence of a hysteretic quantization function relating $q_j(t)$ and $x_j(t)$ implies that $q_j(t)$ follows a piecewise constant trajectory that only changes when the difference

with $x_j(t)$ becomes equal to a parameter Δq_j , called quantum. The variables q_j are called quantized variables. They can be viewed as piecewise constant approximations of the corresponding state variables x_j . Similarly, the components of $\mathbf{v}(t)$ are piecewise constant approximations of the corresponding components of $\mathbf{u}(t)$.

Since the components $q_j(t)$ and $v_j(t)$ follow piecewise constant trajectories, it results that the state derivatives $\dot{x}_j(t)$ also follow piecewise constant trajectories. Then, the state variables $x_j(t)$ have piecewise linear evolutions.

Each component of Equation (2) can be thought of as the coupling of two elementary sub-systems, a static one,

$$\dot{x}_j(t) = f_j(q_1, \dots, q_n, v_1, \dots, v_m) \quad (3)$$

and a dynamical one

$$q_j(t) = Q_j(x_j(\cdot)) = Q_j\left(\int \dot{x}_j(\tau) d\tau\right) \quad (4)$$

where Q_j is the hysteretic quantization function (it is not a function of the instantaneous value $x_j(t)$, but a functional of the trajectory $x_j(\cdot)$).

Since the components $v_j(t)$, $q_j(t)$, and $\dot{x}_j(t)$ are piecewise constant, both sub-systems have piecewise constant input and output trajectories that can be represented by sequences of events. Then, subsystems (3) and (4) define a relation between their input and output sequences of events. Consequently, equivalent DEVS models can be found for these systems, called static functions and quantized integrators, respectively.¹ The piecewise constant input trajectories $v_j(t)$ can be also represented by sequences of events, and source DEVS models that generate them can be easily obtained.

Then, the QSS approximation (2) can be exactly simulated by a DEVS model consisting of the coupling of n quantized integrators, n static functions, and m signal sources. The resulting coupled DEVS model looks identical to the block diagram representation of the original system of Equation (1), as Figure 2 shows.

Based on the idea of QSS1, a second-order accurate method was developed replacing the piecewise constant approximations by piecewise linear ones. The method, called QSS2, can be implemented using DEVS in the same way as QSS1. However, the trajectories are now piecewise linear instead of piecewise constant. Thus, the events carry two numbers that indicate the initial value and the slope of each segment. Also, the static functions and quantized integrators are modified with respect to those of QSS so that they can take into account the slopes.

Following the idea of QSS2, the third-order accurate QSS3 method¹⁶ uses piecewise parabolic trajectories. The family of QSS methods has been further augmented with methods for stiff and marginally stable systems (Backward QSS [BQSS] and Centered QSS [CQSS] of first order,¹³




Figure 1. Coupled DEVS model.

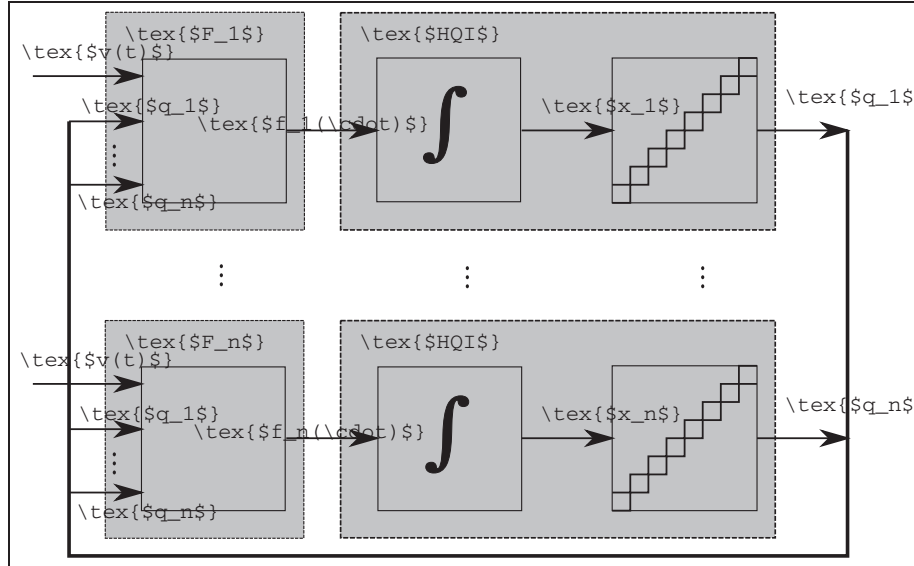


Figure 2. Coupled DEVS model of a QSS approximation.

and Linearly Implicit stiff QSS methods of orders 1 and 2 [LIQSS1 and LIQSS2]¹¹).

2.3 RTAI

RTAI is a real-time operating system (RTOS) that supports several architectures (i386, PowerPC, ARM). RTAI is an extension of the Linux kernel to support real-time tasks to run concurrently with Linux processes. In order to enable real-time simulation, RTAI employs a method that was first used in RT-Linux.¹⁷

The Linux kernel is not by itself a RTOS. It does not provide real-time services, and in some parts of the kernel, interrupts are disabled as a method of synchronization (to update internal structures in an atomic way). The periods of time when interrupts are disabled lead to a scenario where the response time of the system is unknown and time deadlines may therefore be missed.

To avoid this problem, RTAI inserts an abstraction layer underneath the Linux kernel. In this way, Linux never disables the *real* hardware interrupts. The Linux kernel is run on top of another micro-kernel (RTAI + Adeos; Adeos^{18,19} is the abstraction layer used in RTAI, Adeos implements the concept of an interrupt domain, see <http://home.gna.org/adeos/>) as user processes. All hardware interrupts are captured by this micro-kernel and are forwarded to Linux (if Linux has interrupts enabled).

Another problem running real-time tasks under Linux is that the Linux scheduler can take control of the processor from any running process without restrictions. This is unacceptable in a real-time system. Thus, in RTAI, all real-time tasks are run in the micro-kernel without any supervision by Linux. Moreover, these processes are not seen by the Linux scheduler.

Real-time tasks are managed by the RTAI scheduler. There are two different kinds of processes, Linux processes and RTAI processes. RTAI processes cannot make use of Linux services (such as the File system) and vice versa. To avoid this problem, RTAI offers various IPC mechanisms (Pipes, Mailboxes, Shared Memory, etc.).

RTAI provides the user with basic services to implement real-time systems:

Deterministic interrupt response time: An interrupt is an external event to the system. The response time to this event varies from one computer to another, but RTAI guarantees an upper limit (on each computer), which is necessary to communicate with external hardware, for example a data acquisition board.

Inter-process communication (IPC): RTAI supports a number of different methods of IPC, such as semaphores, pipes, mailboxes, and shared memory. These IPC mechanisms can be used for the communication between processes running in real time and normal Linux processes.

High-accuracy timers: When developing real-time systems, the time handling accuracy is very important. RTAI offers clocks and timers with nanosecond (1.0×10^{-9} s) precision. On i386 architectures, these timers use the processor time stamp; therefore, they are very precise.

Interrupt handling: RTAI allows the user to capture hardware interrupts and treat them with custom user handlers. Normal operating systems (such as Linux, Windows) hide hardware interrupts from the user, making the development of communication with external hardware more cumbersome (one needs to write a kernel driver).

Multiprocessor support: RTAI supports also multiprocessor architectures enabling the user to run simultaneously different real-time tasks on separate processors.

2.4 PowerDEVS and real time

*PowerDEVS*⁸ is a general-purpose tool for DEVS simulation. It consists of two main modules: the graphic interface that allows the user to describe DEVS models in a block-diagram fashion, and the simulation engine that simulates the models.

In order to simulate a model, *PowerDEVS* first translates it into a C++ description and compiles it together with the simulation engine. In spite of being usable as a general-purpose DEVS simulator, *PowerDEVS* was primarily designed to simulate continuous and hybrid systems through the use of QSS methods. The standard distribution comes with a comprehensive library of atomic models for hybrid system simulation based on these methods. Users can add new atomic models by defining the dynamics of the DEVS model in C++.

PowerDEVS can be run under RTAI synchronizing the simulation with the physical (wall-clock) time obtaining an accuracy of the order of 1 μ s.⁸

2.5 Parallelization platforms

Parallel computing is a form of computation, in which many calculations are simultaneously carried out on separate processors. Large problems can usually be divided into smaller ones that can then be solved concurrently.

There are different levels of parallelism: instruction level, data level, task level, etc. They differ in many aspects, such as shared or distributed memory, single or multiple processors, the type of communication mechanism they employ, such as message passing, or shared memory, etc.

We shall focus on task-level parallelism, where each LP runs a different piece of code with different data, i.e. according to Flynn's taxonomy,²⁰ we operate on multiple instructions and multiple data (MIMD) streams.

In particular, we shall work with a multi-core architecture²¹ as this offers many desirable features for our implementation, such as shared and unified memory, local communication (on chip), and a generic instruction set.

There exist also other architectures, such as the Message Passing Interface (MPI)²² and general-purpose computing on graphics processing units (GPGPU),²³ but they both have their drawbacks. MPI is a message-passing protocol. This protocol is particularly well suited for clusters. GPGPU does not fit our purposes well, since the architecture lacks shared memory, the code executed on each processor must be identical, synchronization is very expensive, and the instruction set is not as generic as for a PC.

3 Related work

In this section, we review previous work and similar developments.

3.1 Parallel discrete event simulation

Since the beginning of the 1980s, much research has been published on algorithms for synchronizing parallel simulation of discrete event systems.

One of the first approaches to the problem was proposed by Chandy, Misra, and Bryant.³⁻⁵ They suggested a conservative distributed algorithm, whereby processes communicate only by message passing without shared memory and with a centralized coordinator.

The synchronization is achieved by enforcing the correct order of all messages, making all LPs wait until it is safe to produce the next event. By safe we mean that the LP producing the next event is certain that it will never receive a message with an earlier time-stamp than its own logical simulation time.

The authors prove formally the correctness of the algorithm and show that no deadlock can occur by the use of null messages. The memory consumption of the algorithm is fixed as it does not require state saving at all. Chandy and Misra²⁴ also proposed an alternative scheme that avoids the use of null messages and allows deadlocks to happen, which are then broken.

The major drawback of this approach is that the amount of actual parallel computing done is greatly diminished by the mechanism used to enforce the causality constraint. By including lookahead and structure analysis, the efficiency of the algorithm can be improved, but the achievable speedup is highly dependent on the model.

Another solution was proposed by Jefferson⁶ who introduced an optimistic synchronization method for parallel simulation, called Time Warp. Each LP carries its own logical time and advances as if there were no need for synchronization at all. If an event with a smaller time-stamp than the simulation time is received (a straggler event), the LP rolls back to a previously saved state where it is safe to receive that event. It also un-sends (by using anti-messages) all events that were sent out in the rollback period. This involves saving state values in each LP for possible rollbacks and, thus, the memory requirements are larger than for other algorithms.

Both the state saving and the anti-message scheme (that may produce a cascade of rollbacks) are computationally expensive and, consequently, many extensions and additions have been proposed to the original algorithm to avoid rollbacks and reduce the number of state savings required. A complete review of these extensions can be found in an article by Liu.²⁵

More recently, an algorithm called NOTIME⁷ was proposed where the simulation is performed without any synchronization. Using this approach, the different LPs simulate their physical processes as fast as they can. Straggler events are allowed and accepted as correct. This introduces some error in the simulation since the logical time in the different LPs will vary from one processor to another.

While this approach exploits the parallelism optimally well, the technique can produce incorrect results. Yet it may be an acceptable technique in situations where a user is willing to trade simulation accuracy for simulation speed.

The correctness of the results depends indirectly on how the work load is balanced across different LPs. If the work load is unevenly distributed, some LPs will run faster, while those with a heavier workload will run more slowly.

Tang et al. studied an optimistic PDES using reverse computations applied to partial differential equation problems obtaining linear speedup for up to four processors.²⁶ Similar studies have been reported that applied PDES algorithms to the simulation of physical systems, such as an article by Bauer and Page²⁷ (based on a paper by Nutaro²⁸) concluding that for certain kinds of problems optimistic PDES is not suitable if not accompanied by dynamic load balancing.

Results have also been reported concerning the distributed real-time simulation of systems other than differential equations systems. Adelantado et al.²⁹ present a real-time capable HLA middleware and a formal model for its validation. As we shall see, Adelantado et al.'s technique is related to our work as it also uses a local wall clock to synchronize different federates. On the other hand, he assumes in his formal model that each federate executes periodically with a known worst-case execution time. This assumption does not hold for our types of applications since the execution time of each federate is often not known in advance and may vary during the simulation.

3.2 Parallel and real-time DEVS simulation

Many articles address the problem of simulating DEVS models in parallel using any of the aforementioned approaches. Jafer and Wainer³⁰ present CCD + +, a conservative DEVS and Cell-DEVS simulator, and they compare this approach with an optimistic one. Kim et al.³¹ present a Time Warp-based methodology for distributed DEVS simulation. A mixed approach combining conservative and optimistic strategies that is able to optimally exploit the lookahead is investigated by Praehofer and Reisinger.³² Liu³³ explores a light version of the Time Warp mechanism for a DEVS (and Cell-DEVS) simulator on the Cell processor. He introduces the Lightweight Time Warp (LTW) and Multicore Accelerated DEVS System (MADS) as means to overcome Time Warp bottlenecks and to hide the complexity of multicore programming from the general user.

Finally Hong et al.³⁴ proposed an extension to the DEVS formalism, called Real-Time DEVS (RTDEVS), for modeling real-time systems. The main idea of RTDEVS is to fill time advances with executable activities and to specify time bounds required for each activity. They also

implemented a DEVS executive (or DEVS simulator) that synchronizes the wall clock with the simulation time. As shall be discussed (Section 5.3), their technique is also related to our work.

3.3 Parallel simulation of QSS approximations of continuous systems

Some previous results have also been reported concerning the parallelization of simulations of DEVS models resulting from the application of QSS algorithms to continuous systems.

Nutaro³⁵ investigated an implementation of the QSS1 algorithm on a parallel DEVS simulator. The author studied the use of the Time Warp approach and concluded that it is unsuited for simulating large continuous systems. By using a (first-order accurate) QSS, he simulated a sod shock tube problem achieving a speedup that increased to a maximum factor of three when using up to six processors, but quickly decreased when using more processors.

Another implementation of QSS on a multi-processor architecture has been investigated by Maggio et al.³⁶ This implementation was realized using CUDA (a GPGPU language). Unfortunately as was already mentioned earlier in this article, graphical processing units (GPUs) offer a limited instruction set, and synchronization tasks are very expensive. Moreover, GPUs follow a single instruction multiple data (SIMD) stream under Flynn's taxonomy,²⁰ where all processors execute the same code in parallel, which significantly reduces the performance when the code has diverging branches.

These issues limit the classes of applications where GPUs can be efficiently used. In the example analyzed in the cited reference, a speedup of up to a factor of 7.2 has been reported for a system with 64 states using 64 processors.

4 Scaled Real-Time Synchronization

In this section, we present the SRTS technique for PDES synchronization.

4.1 Basic idea

As mentioned already, several different PDES techniques have been proposed in the literature that are based on either conservative, optimistic, or unsynchronized algorithms.

DEVS models resulting from the application of QSS methods to large-scale continuous systems have particular features:

1. Each event originates at a quantized integrator and is instantaneously transmitted to other quantized

integrators through the static functions depending on the corresponding state variable.

2. The events received by quantized integrators change their future evolution, but they do not provoke instantaneous output events.
3. Different quantized integrators provoke output events at different times.
4. Each event represents a section of a piecewise polynomial function.

On the one hand, the second feature implies that synchronization is usually necessary and that techniques such as NOTIME can rarely be used.

On the other hand, if the synchronization is strict so that it respects the ordering of all events, there will be very few calculations performed in parallel. The reason for this assertion is simple. Each time an integrator provokes an event, the event is transmitted to a few static functions, and through them to a few quantized integrators.

A conservative algorithm will not allow any other LP to advance the simulation time until those computations have been completed. Consequently, almost nothing is gained by parallelization.

An optimistic algorithm will allow the simulation time of other LPs to advance, but as soon as the effect of an event is propagated to them, a rollback mechanism will need to take place. Rollback in a large continuous system is impractical, as this will consume a large amount of memory and time and, furthermore, it is difficult to assess in advance, how much time is going to be used up by the rollback.

To overcome these difficulties, we propose a technique, in which we perform a non-strict synchronization. Recalling that events represent sections of polynomials, errors in the synchronization imply a bounded numerical

error in the trajectories transmitted from one process to another.

As with other PDES techniques, SRTS divides the overall model into sub-models representing physical processes, and each of these processes is simulated on a different LP.

In order to avoid the inter-process synchronization overhead, instead of synchronizing the simulation time between all LPs, we synchronize in SRTS the simulation time of each LP with the physical (wall-clock) time.

The only communication between different LPs takes place when events are transmitted from one LP to another. These events are transmitted using RTAI's mailbox IPC mechanism.¹⁴

Each LP synchronizes with the wall clock in the following way. If the next event in the LP is scheduled to occur after τ units of simulation time, the LP waits until the wall clock advances by τ/r units of physical time. We call r the real-time scaling factor. It is a parameter that must be chosen according to the speed, at which the system can be simulated.

Once the waiting period is completed, the LP computes the output event and recomputes the state. When the output event needs to be propagated to sub-models belonging to different LPs, a time-stamped message containing the event is sent to the corresponding mailboxes.

During the waiting period, each LP checks for messages arriving at its mailboxes. Whenever a message arrives, the LP processes the event making use of the time-stamp, recomputes the state, and, if the recomputed next event is scheduled after the current time, it waits again.

In Figure 3, a time diagram is shown where two LPs are synchronized using SRTS. Initially both LPs are waiting. At logical simulation time t_1 (scaled wall time t_1/r), LP A computes its next event. This computation takes Δ_1




Figure 3. Time diagram of the SRTS technique

units of wall time. Then LP A waits until t_2 is reached (at t_2/r of scaled wall time) when it emits its output event, i.e. makes an internal transition, and then goes back to wait. This event is propagated to LP B that computes its external transition processing the received event (consuming Δ_2 units of wall time), schedules its next event to occur at t_3 , and goes back to wait. This cycle is repeated until the final (logical) time is reached.

Implementing this scheme in an abstract simulator, there can be no straggler events; the only problem that could arise concerns simultaneous events in different LPs, but they have no consequences when simulating a QSS approximation. In a real implementation, however, straggler events can appear, since neither the synchronization nor the communication methods are perfect, and also the computation of each event consumes some time.

Yet, as was mentioned before, bounded synchronization errors in the context of QSS simulation provoke in most cases only bounded numerical errors in the trajectories transmitted between LPs.

4.2 SRTS technique

As with other PDES techniques, SRTS requires that the model first be split into as many physical processes as there are different LPs to be used. In this article, we shall not discuss any methodology for efficiently partitioning a model, but many publications can be consulted that address this issue.^{37–40}

We consider that each LP holds a coupled DEVS sub-model, and those sub-models are coupled through input

and output ports modeling the interactions between the sub-systems. Figure 4 shows the coupling structure of a model M divided into four sub-models simulated on four different LPs.

4.2.1 Simulation startup. A sequential DEVS simulation invokes first an initialization routine that sets the initial state of the model, the parameters, etc. Then, the simulation starts.

In SRTS, after each sub-model has been initialized, all simulation threads must start at the same wall-clock time in order to ensure synchronization between different LPs.

To accomplish a simultaneous startup, a barrier is used. All simulation threads are launched and the i th thread starts in the following way:

1. Set the shared memory variable $flag[i] = true$.
2. Wait until $flag[j] = true$ for all $j \neq i$.
3. Measure the initial wall-clock time ($t_0[i]$) of the LP.
4. Start the simulation routine.

The initial wall-clock time is then used to compute the physical time from the beginning of the simulation, by subtracting it from the measured wall-clock time.

4.2.2 Simulation algorithm. The SRTS simulation algorithm can be described as follows:

1. Compute the time of next event t_n of the sub-model.




Figure 4. Coupling structure of SRTS.

2. If $t_n > t_f$ (final simulation time), then set $t_n := t_f$.
3. Wait until the scaled physical time reaches the simulation time (i.e. $t_p/r = t_n$) or until a message arrives.
4. If a message arrives, go to step 10.
5. If $t_n = t_f$ end of simulation.
6. Advance the simulation time to $t := t_n$.
7. Compute and propagate the output event inside the sub-model (external transition).
8. If the event propagates to other sub-models, send messages to the corresponding LPs.
9. Recompute the state of the model causing the event (internal transition) and return to step 1.
10. Advance the simulation time t to the value contained in the message time-stamp.
11. Propagate the event inside the sub-model (external transition) and return to step 1.

This algorithm works well provided that every LP sends and receives the messages at the correct time. To send a message at the correct time, a LP must finish the corresponding calculations before that time, i.e. it should not be in overrun. We shall discuss this issue next.

4.2.3 Overrun policy. Overrun situations can be minimized by setting a sufficiently small value for the real-time factor r . However, we want r to be as big as possible in order to simulate fast. Also, even choosing a very small value for r , real operating systems have latency, and overrun situations occur often.

When a LP is in overrun, it continues simulating without checking for input events. The reason is that input events arriving are in the future with respect to the logical time of the LP in overrun. Thus, it is better to process arriving messages only after the overrun situation has finished.

4.2.4 Message handling. As mentioned earlier, events sent to sub-models belonging to different LPs are transmitted through time-stamped messages. It can happen that a sub-model is in overrun or busy computing state transitions while a message is sent to it. If it does not finish its job before a new message arrives, we face two options. We can queue the messages until the LP is able to process them, or we can discard some of the messages.

While in the context of general DEVS models queueing the messages seems the most reasonable choice, we recall that SRTS is intended to simulate continuous systems where events represent changes in trajectories.

If at time t the LP finishes its job and realizes that there are two input messages in the same mailbox waiting to be processed, one corresponding to time t_1 and the other to time t_2 with $t_1 < t_2 < t$, it serves no purpose to process the older message. It is already late, and the change in the

trajectory it represented has been ignored. However, it makes sense to process the change in t_2 that contains more recent information about the trajectory.

For this reason, each sub-model in SRTS uses one mailbox for each input port (i.e. one mailbox per input trajectory), and each mailbox has capacity 1. No matter if the LP was able to read it or not, each message at the mailbox is overridden by the arrival of a new message.

4.3 SRTS and numerical errors in QSS methods

The use of SRTS implies that messages sent between LPs may be received at incorrect times. Since LPs do not check for input messages when they are in overrun, they will not receive messages with time-stamps greater than the logical time. In other words, they never receive messages from the future. Synchronization errors always imply that messages arrive late.

Let us suppose that we use SRTS with two processors to simulate a QSS approximation of a high-order system:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{q}(t), \mathbf{v}(t))$$

To this end, we split the system as follows:

$$\begin{aligned} \dot{\mathbf{x}}_a(t) &= \mathbf{f}_a(\mathbf{q}_a(t), \mathbf{q}_b(t), \mathbf{v}(t)) \\ \dot{\mathbf{x}}_b(t) &= \mathbf{f}_b(\mathbf{q}_a(t), \mathbf{q}_b(t), \mathbf{v}(t)) \end{aligned}$$

where

$$\mathbf{x} = [\mathbf{x}_a \ \mathbf{x}_b]^T; \quad \mathbf{q} = [\mathbf{q}_a \ \mathbf{q}_b]^T; \quad \mathbf{f} = [\mathbf{f}_a \ \mathbf{f}_b]^T.$$

The delay introduced by the SRTS mechanism implies that \mathbf{q}_a is received with a delay by the LP that computes \mathbf{x}_b , and vice versa. Thus, the QSS approximation takes the form

$$\begin{aligned} \dot{\mathbf{x}}_a(t) &= \mathbf{f}_a(\mathbf{q}_a(t), \mathbf{q}_b(t - \tau_b(t)), \mathbf{v}(t)) \\ \dot{\mathbf{x}}_b(t) &= \mathbf{f}_b(\mathbf{q}_a(t - \tau_a(t)), \mathbf{q}_b(t), \mathbf{v}(t)) \end{aligned}$$

where $\tau_a(t)$ and $\tau_b(t)$ are the communication delays between the LPs (including latency and overrun effects, and delays affected by the real-time scaling factor).

We shall discuss later some implementation issues that ensure that the delays are bounded. Depending on some features of the original ODE (input-to-state stability properties¹²), the presence of bounded delays will only provoke a numerical error that will be superposed to that introduced by the QSS approximation itself. Although this analysis can be easily extended to systems with N LPs, we do not intend to perform a formal study of stability and numerical error bounds related to the usage of SRTS in this article. It is, however, worth mentioning that the numerical error introduced by SRTS is minimized when latency and overrun are minimized.

4.4 Implementation of SRTS on PowerDEVS -RTAI

We implemented the SRTS algorithm in *PowerDEVS* on a multicore architecture. We chose *PowerDEVS* because it implements the whole family of QSS methods and can run under a RTOS (RTAI). The usage of a RTOS is crucial for SRTS since, as explained before, the delay introduced by the latency must be bounded and minimized.

The PDES version of *PowerDEVS* generates a simulation executable that performs the following steps:

- It initializes the *PowerDEVS* model (see Section 2.4).
- It starts the RTAI real-time system, initializing the real-time clocks and a main real-time task.
- For each LP, it creates and sets up as many mailboxes as there are input ports.
- It starts one real-time thread per physical process. It then forks these threads with the simulation loop. An index parameter allows each thread to know, which physical processes it corresponds to.
- It waits for all real-time threads to finish and then exits while un-allocating all mailboxes and threads.

We have developed a live-CD from where one can try out (or install) *PowerDEVS* -RTAI on any i386 processor.

4.4.1 Synchronization and CPU allocation. We implemented all synchronization tasks by a busy-waiting strategy, wasting CPU cycles without releasing the processor until the waiting time ends or a new message arrives. Although this strategy wastes computing power, it offers a more accurate synchronization. A busy-waiting strategy under RTAI does not allow any other task to use the processor that is running the thread. In this way, the latency is minimized. However, with this strategy, each processor can only run one thread, which limits the number of processes that can run in parallel.

RTAI also gives the user the ability to select the processor to run a given thread, i.e. we can allocate one thread per processor. As a consequence, we can divide the model into as many sub-models representing physical processes as there are processors available. On our hardware, we can allocate up to 12 LPs (as hyper-threading technology divides each core processor into two virtual processors).

4.4.2 Latency and overrun. In the absence of latency and neglecting the processing time for messages, transitions, and other calculations, the algorithm described earlier in this section correctly simulates the overall model. However, in a real implementation, we cannot avoid some latency, and sometimes calculations take longer than the physical time allowed, and the simulation then enters an overrun situation.

For example, if the physical processor needs 1 wall-clock second to simulate 1 second of logical time, we cannot use a real-time scaling factor greater than 1. If we do that, the thread will never be able to catch up with the scaled wall-clock time. While this type of overrun situation can be avoided by using a sufficiently small real-time scaling factor, one event generated by a quantized integrator in a QSS method usually triggers many simultaneous events through the static functions. Whereas the first event will be correctly synchronized, the remaining events will be invariably in overrun.

Latency also introduces some error in the synchronization, which affects the communication mechanism. If an event scheduled for t_1 is sent out at time $t_1 + \Delta t$, the other threads will receive it with a time delay of Δt . If during that delay period the other threads computed an event, an error appears that affects the global result.

The hardware platform we use under RTAI exhibits an average latency that ranges from 150 to 500 ns, with a maximum value of around 10 μ s. These values limit the maximum real-time factor that we may use. For instance, if we use a real-time factor of 1000, a latency of 10 μ s is translated to 10 ms of maximum simulation time error in events transmitted between different LPs. If that simulation time error provokes an unacceptable simulation error, then we would be enforced to use a smaller real-time factor.

5 Adaptive real-time scaling

The main drawback of SRTS is that the real-time scaling factor r must be chosen by the user. Unless the user performs some previous trial and error experiments, he would be hard-pressed to know a priori which value to assign to the real-time scaling factor. Another problem is that r remains constant during the entire simulation. In many cases, the simulation workload changes with time, and it makes sense to adapt r according to the actual workload. In this section, we introduce an adaptive version of SRTS, called ASRTS, where the real-time scaling factor is automatically adapted to optimize the efficiency of the simulation.

5.1 Basic idea

ASRTS attempts to automatically change the real-time scaling factor to minimize the time that processors spend waiting but without provoking overrun situations. In this way, ASRTS improves the overall efficiency by lowering the simulation time.

The idea behind ASRTS is rather simple. It is identical to SRTS, except that it periodically changes the real-time scaling factor. ASRTS cuts the simulation time into sampling periods of equal duration. During each such period, each LP measures the total time that it spends waiting. At the end of the period, one of the threads collects the

accumulated waiting times of all LPs and determines their minimum (i.e. the waiting time of the process that had the heaviest workload). Using that information, it computes the ratio between the smallest accumulated waiting time and the length of the sampling period. If the minimum waiting ratio w is greater than a prescribed desired waiting ratio w_0 (we use values around 10%), the algorithm is allowed to increase the real-time scaling factor and simulate faster. Otherwise, the algorithm decreases r . In this way, ASRTS tries to keep the minimum waiting ratio around the desired waiting ratio w_0 .

Changes in the real-time scaling factor are done synchronously, that is, all LPs change their value of r at the same time (logically and physically) and to the same value, since the scaling factor affects directly the implicit synchronization among the entire group of LPs. Once r has been changed, the LPs continue with the simulation using the SRTS algorithm until the next sampling period has ended.

5.2 ASRTS algorithm

As mentioned above, ASRTS is identical to SRTS but includes periodic checkpoints (with a sampling period of ΔT), when the scaling factor r may be changed. During each sampling, ASRTS gathers statistics about the workload and the time spent waiting by each LP. When they reach the next checkpoint, the LPs stop simulating and wait until the new real-time scaling factor has been computed. This computation is done by a thread coordinator (thread 0), while the remaining threads wait in a synchronizing barrier that holds them until the new scaling factor has been computed. At the end of each sampling period, after each LP computed its accumulated waiting time T_{w_i} , the coordinator computes the new real-time scaling factor as follows:

1. It finds the minimum waiting ratio

$$w = \frac{\min_i \{T_{w_i}, \sigma\}}{\Delta T},$$

where σ is a value lower but close to 1 to avoid division by zero. In a real implementation not all of the T_{w_i} can be 1 because of overheads.

2. It computes the optimal real-time scaling factor \hat{r} for the desired waiting ratio w_0 as

$$\hat{r} = r \frac{1 - w_0}{1 - w}. \quad (5)$$

3. If $\hat{r} < r$, it computes the new real-time scaling factor as $r := \hat{r}$, slowing down the speed of the simulation.

4. Otherwise if $\hat{r} \geq r$, it computes the new real-time scaling factor as

$$r := \lambda r + (1 - \lambda)\hat{r}, \quad (6)$$

which smoothly increases the simulation speed according to the discrete eigenvalue λ . (Note that Equation (6) is a difference equation with λ as eigenvalue.)

The ASRTS algorithm has three tuning parameters: the sampling period ΔT , the desired waiting ratio w_0 , and the discrete eigenvalue λ . We discuss the most suitable values of these parameters later in this section. ASRTS makes use of Equation (5) that computes the real-time scaling factor \hat{r} , for which a waiting ratio w_0 is obtained. This expression is derived below.

5.2.1 Optimal real-time scaling factor. As mentioned above, ASRTS attempts to drive the minimum waiting ratio w to the desired waiting ratio w_0 by adjusting the real-time scaling factor r . Suppose that during a sampling period ΔT the LP with the maximum workload had a waiting ratio w while the real-time scaling factor was r . Thus, during that period, it spent computing a total of

$$T_c = \Delta T - w \cdot \Delta T = (1 - w) \cdot \Delta T$$

units of time, and the simulation time advanced by $r \cdot \Delta T$ units of time.

Thus, in order to advance $r \cdot \Delta T$ units of simulation time, the LP needed T_c units of physical time. We can define the workload as follows

$$W \triangleq \frac{\text{physical time}}{\text{simulation time}} = \frac{T_c}{r \cdot \Delta T} = \frac{1 - w}{r}.$$

If we assume that the workload changes slowly, we can expect that, in the next sampling period, it will remain almost unchanged. If we use the appropriate scaling real-time factor \hat{r} in the new sampling period, we shall obtain the desired waiting ratio w_0 and, then,

$$W = \frac{1 - w}{r} = \frac{1 - w_0}{\hat{r}}$$

from which we obtain Equation (5).

5.2.2 Parameter selection. In order to properly set the values of the method parameters of ASRTS, we may consider the following points.

- The sampling period ΔT should be large compared with the physical time required by the re-synchronization routine and also compared with the time necessary to compute several events. In this way, the time spent in the synchronization routine

is negligible, and every sampling occurs after several events were computed so that workload statistics make sense. However, ΔT must be small enough so that the real-time scaling factor is changed several times during the simulation to be able to react to workload changes in a timely fashion.

In the *PowerDEVS* implementation, we used $\Delta T = 100$ ms. In our platform, the synchronization routine takes from 10 to 100 μ s, and we are interested in simulations that take at least several seconds (otherwise parallelization might not be necessary). Thus, a period of $\Delta T = 100$ ms is a reasonable choice for most cases.

- The desired waiting ratio w_0 should be greater than 0 and less than 1. We want simulations to run fast, minimizing the waiting time, so a small value of w_0 should be chosen. However, if w_0 is very close to 0, several types of overrun situations can occur that may cause large numerical errors.

In the *PowerDEVS* implementation, we normally use $w_0 = 0.15$, which means that the real-time scaling factor is adjusted such that the LP with the heaviest workload wastes approximately 15% of the time in a waiting routine.

- The discrete eigenvalue λ (with $0 < \lambda < 1$) determines how fast the real-time scaling factor r increases its value to reach the optimal value \hat{r} according to Equation (6). When λ is close to 0 the adaptation is fast. When it is close to 1, the adaptation is slow and smooth. In the presence of fast changes in the workload, a slow adaptation prevents the real-time scaling factor from assuming very large values that might cause frequent overruns.

For this reason, we normally choose $\lambda = 0.9$. With this value, the real-time scaling factor reaches 80% of its final value after 10 periods (i.e. 1 second in our implementation).

5.3 Relation with other approaches

The present work can be seen as a combination of two previously formulated ideas of the PDES and DEVS worlds. First, it is related to the RTDEVS executive³⁴ since simulations run synchronized with the wall clock, although in our case, a scaled version of the wall clock. However, the RTDEVS formalism is not used for model description, and regular DEVS is used instead. In our case, the time-consuming activities are, in fact, the computation of the output events and the busy waiting associated with synchronizing them.

SRTS and ASRTS are also related to the NOTIME technique⁷ since the causality constraint may be violated. Violations of the causality constraint are, however, minimized by the introduction of busy waiting associated with the synchronization of the LPs with the wall clock. These

waiting periods act as an artificial load that help balance the workload among LPs, thus they implicitly solve the unbalancing problem of NOTIME.

6 Applications and examples

In this section, we present some simulation results. We simulated three separate large-scale systems using QSS methods, first in a sequential way and then using SRTS and ASRTS techniques.

All examples were run using an Intel i7 970, six-core machine with hyper-threading.⁴¹ RTAI's test suite reports 416 ns average and 18,000 ns maximum latencies on this setup. In all examples, each simulation was repeated 10 times, and the results did not exhibit any noticeable change.

6.1 Power control of an air conditioner population

This example, taken from Perfumo et al.,⁴² is a model to study the dynamics and control of the power consumption of a population of air conditioners (ACs).

We consider here a large population of ACs used to control the temperature of different rooms. The temperature of the i th room $\theta_i(t)$ follows the equation

$$\frac{d\theta_i(t)}{dt} = -\frac{1}{C_i \cdot R_i} [\theta_i(t) - \theta_a + R_i \cdot P_i \cdot m_i(t) + w_i(t)], \quad (7)$$

where R_i and C_i are parameters representing the thermal resistance and capacity of the i th room, respectively. Here P_i is the power of the i th air conditioner when it is in its on state, θ_a is the ambient temperature, and $w_i(t)$ is a noise term representing thermal disturbances.

Variable $m_i(t)$ is the state of the i th air conditioner, which takes a value of 1 when the unit is in its on state, and 0 otherwise. It follows the hysteretic on-off control law:

$$m_i(t^+) = \begin{cases} 0 & \text{if } \theta_i(t) \leq \theta_r(t) - 0.5 \text{ and } m_i(t) = 1 \\ 1 & \text{if } \theta_i(t) \leq \theta_r(t) + 0.5 \text{ and } m_i(t) = 0 \\ m_i(t) & \text{otherwise} \end{cases} \quad (8)$$

where $\theta_r(t)$ is the reference temperature that is calculated by a global control system.

The power consumption of the entire AC population is computed as

$$P(t) = \sum_{i=1}^N m_i(t) \cdot P_i$$

and a global control system regulates it, so it follows a desired power profile $P_r(t)$. In order to accomplish this, a




Figure 5. Power consumption in the AC system (sequential simulation result).

proportional integral (PI) control law is used to compute the reference temperature:

$$\theta_r(t) = K_P \cdot [P_r(t) - P(t)] + K_I \cdot \int_{\tau=0}^t [P_r(\tau) - P(\tau)] d\tau$$

where K_P and K_I are the parameters of the PI controller.

In this example, we chose $N = 2400$ air conditioner units, and the set of parameters given in the article.⁴² Thus, we are dealing with a fairly large hybrid system consisting of a set of 2401 differential equations, with 2400 zero-crossing conditions.

We first simulated the system using the QSS3 method with a quantization of $\Delta Q_{\text{rel}} = 10^{-3}$ and $\Delta Q_{\text{min}} = 10^{-6}$ in a sequential fashion. Figure 5 plots the average power consumption $P(t)/P_{\text{max}}$.

The final simulation time was $t_f = 3000$ s, and the sequential simulation consumed 58 s of physical time.

We then divided the system into 12 sub-systems with 200 ACs in each sub-model. The first sub-model also included the PI controller. We then used the SRTS technique with different real-time scaling factors. We started with a scaling factor of $r = 200$, which consumes $t_f/r = 3000/200 = 15$ s of physical time. We obtained decent results using scaling factors of up to $r = 950$, which takes $t_f/r = 3000/950 = 3.16$ s. For each scaling factor, we compared the simulation results with those of the sequential simulation and measured the relative root mean square (RMS) error.

Here and in the following examples, the RMS error is computed as follows:

$$\text{RMS} = \frac{\text{mean}(x - \hat{x})^2}{\text{mean}(x)}$$

where x is the sequential result and \hat{x} is the parallel result.

Figure 6 plots the measured error in function of the real-time scaling factor.

In this example, using $r = 950$, SRTS was more than 18 times faster than the sequential simulation without introducing unacceptably large numerical errors.

However for $r > 500$, we observed that the LPs were most of the time in overrun. Yet, in this case, the overrun situation does not affect the numerical error in major ways.

Then we simulated the system using ASRTS with parameters $\Delta T = 100$ ms, $w_0 = 0.15$, and $\lambda = 0.9$. The ASRTS simulation took 6.6 s, reaching a maximum real-time scaling factor of $r = 480$. In consequence, ASRTS was almost 9 times faster than the sequential simulation using 12 processors. The real-time scaling factor reached was close to the limit, at which large overruns are observed.

We repeated the simulation with ASRTS for different values of λ . Figure 7 shows the way in which ASRTS adapts the scaling real-time factor during each simulation run.

We notice that r converges very slowly to the optimal scaling factor for large values of λ (close to 0.99). Consequently the simulation is slow. On the other hand, r converges quickly for small values of λ (around 0.8) but needs to reduce its value frequently. A reduction in the real-time scaling factor indicates that some LP waited less than the desired waiting ratio w_0 , which is unsafe as this might indicate overrun. A value of $\lambda \approx 0.9$ provides a

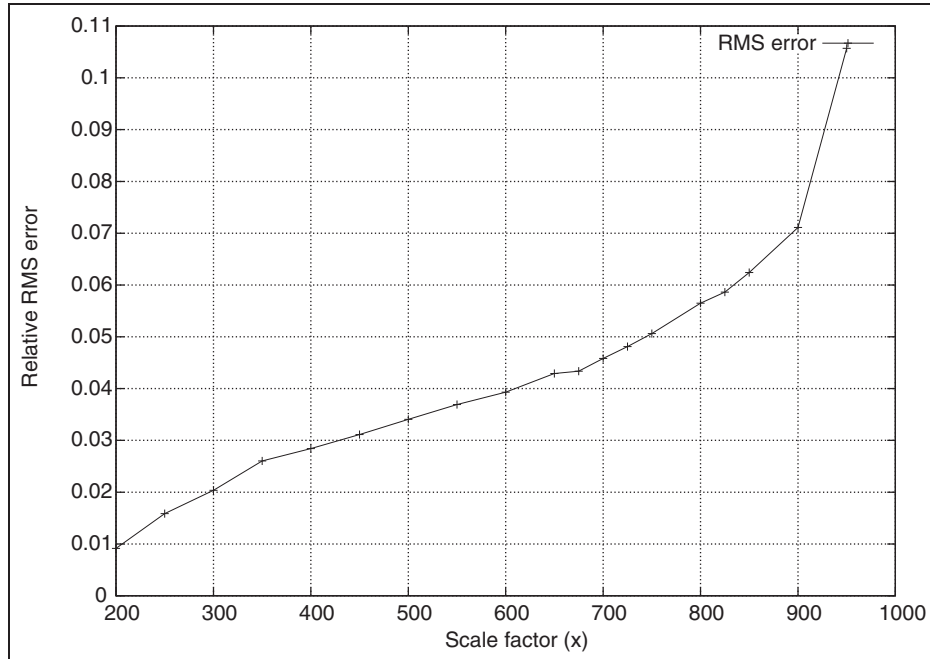


Figure 6. SRTS: error versus real-time scaling factor, AC example




Figure 7. ASRTS: real-time scaling factor adaptation for different values of λ, AC example

good compromise between fast adaptation and overrun avoidance.

We then repeated the ASRTS simulation while varying the number of processors, splitting the model into two,

four, six, and eight sub-models. Figure 8 shows the simulation speedup obtained for these cases.

We can observe a close to linear increase in the simulation speedup with the number of processors.

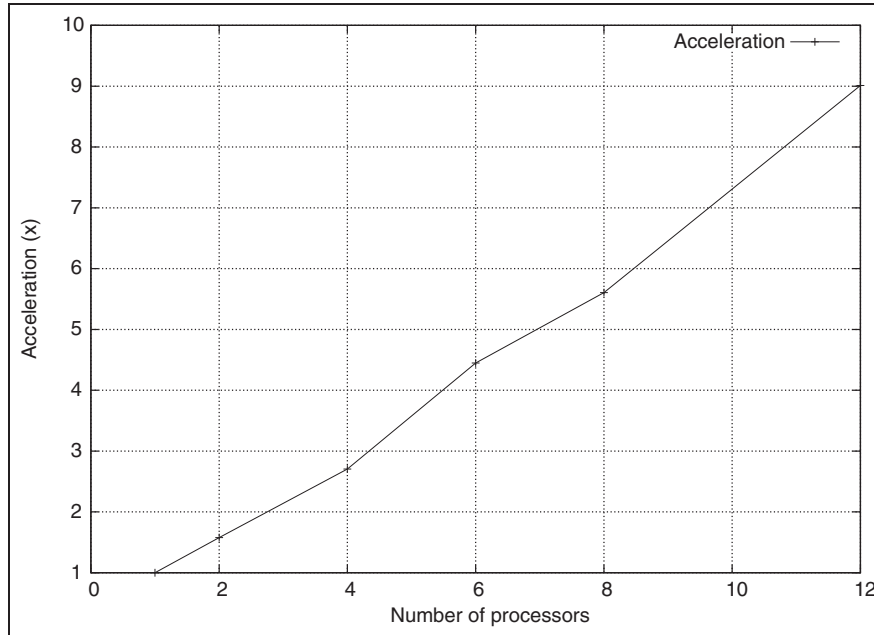


Figure 8. ASRTS: speedup versus number of processors, AC example.

6.2 Logical inverter chain

A logical inverter performs a logical operation on a signal. When the signal assumes a high level, it outputs a low level, and vice versa.

Logical inverters are implemented by electrical circuits. They exhibit a non-ideal response since the rise and fall time of the signal output is limited by physical characteristics, and the correct output level is not immediately obtained, i.e. it is delayed.

An inverter chain is a concatenation of several inverters, where the output of each inverter acts as the input to the next one. Making use of the aforementioned physical limitations, inverter chains can be used to obtain delayed signals. We consider here a chain of m inverters according to the model given in Savcenca and Mattheij,⁴³ which is characterized by the following equations:

$$\begin{cases} \dot{\omega}_1(t) = U_{\text{op}} - \omega_1(t) - \Upsilon g(u_{\text{in}}(t), \omega_1(t)) \\ \dot{\omega}_j(t) = U_{\text{op}} - \omega_j(t) - \Upsilon g(\omega_{j-1}(t), \omega_j(t)), \\ j = 2, 3, \dots, m \end{cases} \quad (9)$$

where

$$g(u, v) = (\max(u - U_{\text{thres}}, 0))^2 - (\max(u - v - U_{\text{thres}}, 0))^2. \quad (10)$$

We used the set of parameters given in the article: $\Upsilon = 100$ (which results in a very stiff system), $U_{\text{thres}} = 1$, and $U_{\text{op}} = 5$. The initial conditions are, as in the given reference, $\omega_j = 6.247 \times 10^{-3}$ for odd values of j and $\omega_j = 5$ for even values of j . The input is a periodic trapezoid signal,

with parameters $V_{\text{up}} = 5$ V, $V_{\text{low}} = 0$ V, $T_{\text{down}} = 10$, $T_{\text{up}} = 5$, $T_{\text{rise}} = 5$, and $T_{\text{fall}} = 2$.

In this case, we considered a system of $m = 504$ inverters, so we have a set of 504 differential equations with 1008 discontinuity conditions due to the ‘max’ functions in Equation (10).

As in the previous example, we first simulated the system in a sequential fashion using the stiff LIQSS3 solver, which has been added to *PowerDEVS* after publication of Migoni and Kofman.¹¹ The resulting voltage at the last inverter is shown in Figure 9.

We then divided the system into 12 sub-models with 42 inverters each and simulated with SRTS and different real-time scaling factors. This time around, we obtained decent results for up to $r \approx 35$.

Figure 10 shows the error in function of the real-time scaling factor used. The RMS error looks large because it is mainly due to a small time shift in the output voltage wave. As the wave shows abrupt changes, a small time shift can cause a large RMS error.

We then applied the ASRTS algorithm and, using 12 processors, the simulation was almost 4.5 times faster than the sequential one.

In this example, the workload is not evenly distributed over the system. As the wave travels, it provokes a high workload at the processor that computes the states that are switching. In contrast, processors that compute states that do not change show a very low workload. Consequently in this case, neither SRTS nor ASRTS can exploit parallelization as efficiently as in the previous example. However, the results are still good.




Figure 9. Voltage on the last inverter (sequential simulation result).

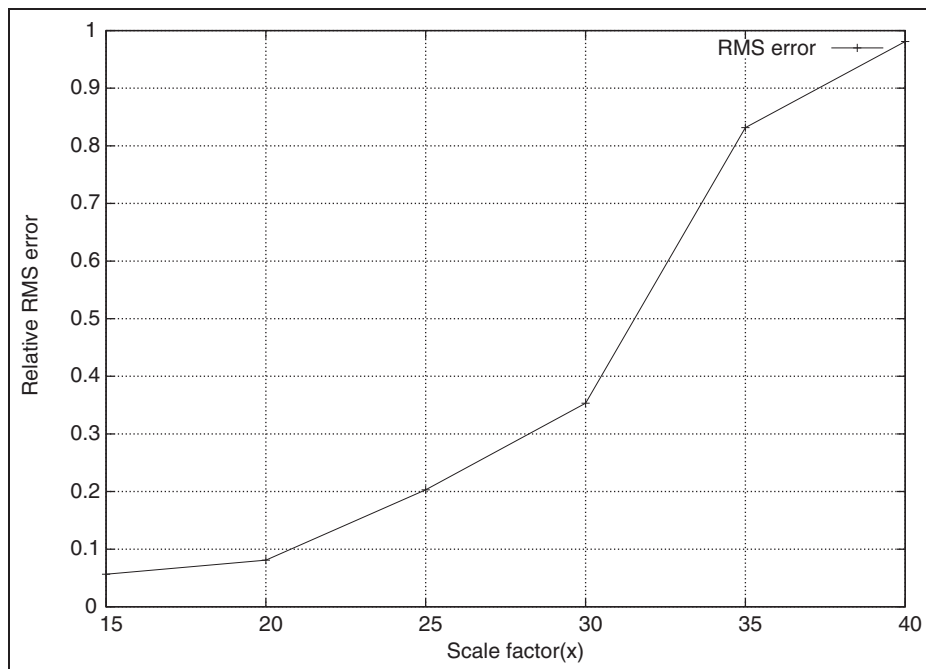


Figure 10. SRTS: error versus real-time scaling factor, inverter chain example.

We then repeated the simulations using ASRTS while varying the number of processors. Figure 11 shows the speedup obtained in each case.

Once again, the simulation speed increases almost linearly with the number of processors in use.

6.3 LC transmission line

The following system of equations represents a lumped model of a lossless transmission line, formed by 600 sections of LC circuits with parameters $L = C = 1$:




Figure 11. ASRTS: speedup versus number of processors, inverter chain example.

$$\begin{aligned}
 \dot{\phi}_1(t) &= u_0(t) - u_1(t) \\
 \dot{u}_1(t) &= \phi_1(t) - \phi_2(t) \\
 &\vdots \\
 \dot{\phi}_j(t) &= u_{j-1}(t) - u_j(t) \\
 \dot{u}_j(t) &= \phi_j(t) - \phi_{j+1}(t) \\
 &\vdots \\
 \dot{\phi}_n(t) &= u_{n-1}(t) - u_n(t) \\
 \dot{u}_n(t) &= \phi_n(t).
 \end{aligned} \tag{11}$$

We consider a sinusoidal input signal:

$$u_0(t) = \sin(\omega t) \tag{12}$$

with $w_0 = 0.13$ Hz. We also set zero initial conditions $u_i = \phi_i = 0$, $i = 1, \dots, n$. This is a marginally stable system, of order 1200 as there are two state variables per segment.

As in the previous examples, we first simulated the system in a sequential fashion using the QSS3 solver. Figure 12 shows the resulting voltage at the end of the transmission line $u_{600}(t)$.

Then we divided the system into 12 sub-models with 50 segments each and simulated it using SRTS with different values of the real-time scaling factor. Figure 13 shows the RMS error as a function of the real-time scaling factor.

This time around, we obtained good results for $r < 400$. For larger values of the real-time scaling factor, the results soon become numerically unstable. This is not surprising as the addition of delays in marginally stable systems tends to cause instability.

We then simulated the system using ASRTS. With 12 processors, ASRTS managed to accelerate the simulation so it ran more than 5.5 times faster than the sequential simulation.

As in the previous example, the workload is not evenly distributed while the wave travels from the input to the output, which limits the advantages of parallelization.

We then used ASRTS with a varying number of processors, obtaining again an almost linear relationship between the number of processors in use and a speedup as shown in Figure 14.

7 Conclusions and future work

We have introduced SRTS and ASRTS, two new techniques for parallel DEVS of continuous and hybrid systems.

The techniques are based on the idea of synchronizing each LP simulation time with a scaled version of the physical time. As all LPs are synchronized with the physical time, they are indirectly synchronized also against each other. This indirect synchronization permits running each process independently of the remaining processes, avoiding the overhead imposed by inter-process synchronization and coordination.

We implemented these new techniques in *PowerDEVS* using the RTOS RTAI. The RTOS was needed to ensure a precise synchronization, to properly perform the CPU allocation, and for avoiding that the OS (Linux) takes control of the processor while the thread is running.




Figure 12. Voltage on last LC segment (sequential simulation result).




Figure 13. SRTS: RMS error versus real-time scaling factor, LC line example.

We tested these new techniques with three application examples and studied their performance and limitations. Using 12 processors, we found a major improvement on the simulation speed. In the examples analyzed, ASRTS simulated from 4.5 to 9 times faster than the sequential

simulation. The non-adaptive SRTS is able to simulate even slightly faster but requires prior experimentation to determine the most suitable real-time scaling factor to use.

We have also analyzed the additional numerical error introduced by the techniques. We found that, provided that




Figure 14. ASRTS: speedup versus number of processors, LC line example.

the real-time scaling factor r remains small enough to avoid frequent overruns, the error remains bounded. Moreover, ASRTS adjusts r in order to avoid overruns so the error introduced is kept small.

This work opens up several lines of research to be considered in the future. A few of them are listed below:

- In the present work, we manually split the model trying to obtain a well-balanced division. It is interesting to investigate an automatic way to split the model and also to study techniques to dynamically re-balance the workload among the available processors.

We are currently working on an extension to the DEVS formalism, called vectorial DEVS, that will simplify this task.

- A fall-back mechanism for overrun situations is also a desirable feature. By implementing a roll-back mechanism when a large overrun is found, we could bound the errors introduced.

To this end, we could use the ASRTS checkpoints as roll-back stages. In this way, we would only need to save the state of the system at the last checkpoint consuming a bounded and constant amount of memory. Also, we can know exactly how much time we waste at each roll-back: ΔT .

- We have developed the implementation on a multi-core system running a RTOS. A very important issue is to study how these algorithms scale with a greater number of processors or in a cluster

environment. To this end, several problems have to be solved first, such as the clock synchronization among different nodes or CPUs, the communication between nodes (in the current multicore approach it is implemented using shared memory), and the CPU allocation.

- We will work on a formal analysis of how the delays in the communication affect the numerical results, following the approach outlined in Section 4.3. The idea is to represent the effects of the delays as bounded perturbations and to analyze the perturbed systems to ensure numerical stability and error boundedness.
- We can also investigate new heuristics for the ASRTS. So far, we have only included the time spent waiting as a measurement of how the simulation is doing but including overrun statistics or delay data might result in an improved performance.
- Although we developed these new techniques for discrete event simulation of continuous systems, it may be interesting to study their usage also with other numerical integration methods for continuous and hybrid systems (particularly with multi-rate algorithms). Also, it is worth considering the study of SRTS and ASRTS in more general applications of discrete event simulation.

The complete code of the SRTS and ASRTS implementation in *PowerDEVS* can be found and downloaded from the real-time branch of the SVN.⁴⁴

References

1. Cellier F and Kofman E. *Continuous System Simulation*. New York: Springer, 2006.
2. Zeigler B, Praehofer H and Kim TG. *Theory of Modeling and Simulation*, 2nd edn. New York: Academic Press, 2000.
3. Bryant RE. *Simulation of Packet Communication Architecture Computer Systems*. Technical Report, Cambridge, MA, USA, 1977.
4. Chandy KM and Misra J. Distributed simulation: a case study in design and verification of distributed programs. *IEEE Trans Softw Engng* 1979; 5: 440–452.
5. Misra J. Distributed discrete-event simulation. *ACM Comput Surv* 1986; 18: 39–65.
6. Jefferson DR. Virtual time. *ACM Trans Program Lang Syst* 1985; 7: 404–425.
7. Rao DM, Thondugulam NV, Radhakrishnan R and Wilsey PA. Unsynchronized parallel discrete event simulation. In *Simulation Conference Proceedings, 1998, Winter*, vol. 2, 1998, pp. 1563–1570.
8. Bergero F and Kofman E. PowerDEVS: a tool for hybrid system modeling and real-time simulation. *Simulation* 2011; 87: 113–132.
9. Kofman E. Quantization-based simulation of differential algebraic equation systems. *Simulation* 2003; 79: 363–376.
10. Kofman E. Discrete event simulation of hybrid systems. *SIAM J Sci Comput* 2004; 25: 1771–1797.
11. Migoni G and Kofman E. Linearly implicit discrete event methods for stiff ODEs. *Latin Am Appl Res* 2009; 39: 245–254.
12. Castro R, Kofman E and Cellier F. Quantization based integration methods for delay differential equations. *Sim Modell Practice Theory* 2011; 19: 314–336.
13. Migoni G, Kofman E and Cellier F. Quantization-based new integration methods for stiff ordinary differential equations. *SIMULATION* 2012; 88: 387–407.
14. DIAPM. Programming Guide 1.0, 2000. <http://www.aero.polimi.it/~rtai/documentation/>.
15. Zeigler B. *Theory of Modeling and Simulation*. New York: John Wiley & Sons, 1976.
16. Kofman E. A third order discrete event simulation method for continuous system simulation. *Latin Am Appl Res* 2006; 36: 101–108.
17. Barabanov M. *A Linux-based RealTime Operating System*. Master's thesis, New Mexico Institute of Mining and Technology, New Mexico, 1997.
18. Yaghmour K. Adaptive Domain Environment for Operating Systems, 2002. <http://www.opersys.com/adeos/>.
19. Yaghmour K. Building a Real-Time Operating Systems on top of the Adaptive Domain Environment for Operating Systems, 2003. <http://www.opersys.com/adeos/>.
20. Flynn MJ. Very high-speed computing systems. *Proceedings of the IEEE* 1966; 54: 1901–1909.
21. Blake G, Dreslinski RG and Mudge T. A survey of multicore processors. *IEEE Signal Processing Magazine* 2009; 26(6): 26–37.
22. Snir M, Otto S, Huss-Lederman S, Walker D and Dongarra J. *MPI-The Complete Reference, Volume 1: The MPI Core*, 2nd (revised) edn. Cambridge, MA: MIT Press, 1998.
23. Randima F. *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*. Pearson Higher Education, 2004.
24. Chandy KM and Misra J. Asynchronous distributed simulation via a sequence of parallel computations. *Commun ACM* 1981 24: 198–206.
25. Liu J. *Parallel Discrete-Event Simulation (Wiley Encyclopedia of Operations Research and Management Science)*. New York: John Wiley & Sons, 2010.
26. Tang Y, Perumalla KS, Fujimoto RM, Karimabadi H, Driscoll J and Omelchenko Y. Optimistic parallel discrete event simulations of physical systems using reverse computation. In *Workshop on Principles of Advanced and Distributed Simulation, 2005 (PADS 2005)*, 2005, pp. 26–35.
27. Bauer DW and Page EH. Optimistic parallel discrete event simulation of the event-based transmission line matrix method. In *Winter Simulation Conference, 2007*, pp. 676–684.
28. Nutaro J. A discrete event method for wave simulation. *ACM Trans Model Comput Simul* 2006; 16: 174–195.
29. Adelantado M, Siron P and Chaudron J. Towards an HLA run-time infrastructure with hard real-time capabilities. In *International Simulation Multi-Conference*, Ottawa, Canada, 2010.
30. Jafer S and Wainer G. Conservative DEVS: a novel protocol for parallel conservative simulation of DEVS and cell-DEVS models. In *Proceedings of the 2010 Spring Simulation Multiconference (SpringSim '10)*, San Diego, CA, 2010. Society for Computer Simulation International, 2010, pp. 140:1–140:8.
31. Kim KH, Seong YR, Kim TG and Park KH. Distributed simulation of hierarchical DEVS models: hierarchical scheduling locally and time warp globally. *Trans Soc Comput Sim Int* 1996; 13: 135–154.
32. Praehofer H and Reisinger G. Distributed simulation of DEVS-based multiformalism models. In *Proceedings of the Fifth Annual Conference on AI, Simulation, and Planning in High Autonomy Systems, 1994. Distributed Interactive Simulation Environments*, 1994, pp. 150–156.
33. Liu Q. *Algorithms for Parallel Simulation of Large-Scale DEVS and Cell-DEVS Models*. PhD Thesis, Systems and Computer Engineering Department, Carleton University, 2010.
34. Hong JS, Song H-S, Kim TG and Park KH. A real-time discrete event system specification formalism for seamless real-time software development. *Discrete Event Dynamic Syst* 1997; 7: 355–375. DOI: 10.1023/A:1008262409521.
35. Nutaro J. *Parallel Discrete Event Simulation with Application to Continuous Systems*. PhD thesis, 2003. AAI3119971.
36. Maggio M, Stavåker K, Donida F, Casella F and Fritzon P. Parallel simulation of equation-based object-oriented models with quantized state systems on a GPU. In *Proceedings of 7th Modelica Conference*, Como, Italy, 2009.
37. Deelman E and Szymanski BK. Dynamic load balancing in parallel discrete event simulation for spatially explicit problems. In *Proceedings Twelfth Workshop on Parallel and Distributed Simulation, 1998 (PADS 98)*, 1998, pp. 46–53.

AQ1

AQ2

38. Som TJ and Sargent RG. Model structure and load balancing in optimistic parallel discrete event simulation. In *Proceedings of the 14th Workshop on Parallel and Distributed Simulation (PADS '00)*. Washington, DC: IEEE Computer Society, 2000, pp. 147–154.
39. Glazer DW and Tropper C. On process migration and load balancing in Time Warp. *IEEE Trans Parallel Distributed Syst* 1993; 4: 318–327.
40. Boukerche A and Tropper C. A static partitioning and mapping algorithm for conservative parallel simulations. *SIGSIM Sim Dig* 1994; 24: 164–172.
41. Marr DT, Binns F, Hill DL, Hinton G, Koufaty DA, Miller JA, et al. *Hyper-Threading Technology Architecture and Microarchitecture*. Technical report, Intel Technology Journal, 2002.
42. Perfumo C, Kofman E, Braslavsky J and Ward JK. Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. *Energy Conversion Management* 2012; 55: 36–48.
43. Savcenco V and Mattheij RMM. A multirate time stepping strategy for stiff ordinary differential equations. *BIT Numer Math* 2007; 47: 137–155.
44. PowerDEVS site at SourceForge, <http://sourceforge.net/projects/powerdevs/>.