

SIMULATION

<http://sim.sagepub.com/>

Quantized state simulation of spiking neural networks

Guillermo L Grinblat, Hernán Ahumada and Ernesto Kofman

SIMULATION published online 17 March 2011

DOI: 10.1177/0037549711399935

The online version of this article can be found at:

<http://sim.sagepub.com/content/early/2011/03/09/0037549711399935>

Published by:



<http://www.sagepublications.com>

On behalf of:



Society for Modeling and Simulation International (SCS)

Additional services and information for *SIMULATION* can be found at:

Email Alerts: <http://sim.sagepub.com/cgi/alerts>

Subscriptions: <http://sim.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>



Quantized state simulation of spiking neural networks

Guillermo L Grinblat, Hernán Ahumada and Ernesto Kofman

Abstract

In this work, we explore the usage of quantized state system (QSS) methods in the simulation of networks of spiking neurons. We compare the simulation results obtained by these discrete-event algorithms with the results of the discrete time methods in use by the neuroscience community. We found that the computational costs of the QSS methods grow almost linearly with the size of the network, while they grows at least quadratically in the discrete time algorithms. We show that this advantage is mainly due to the fact that QSS methods only perform calculations in the components of the system that experience activity.

Keywords

discrete event system, quantized state systems, spiking neural networks

1. Introduction

Continuous system simulation is a topic that has rapidly evolved in recent decades. The development of faster and more powerful computers has allowed the representation of larger and more complex models that require the usage of more efficient and sophisticated simulation algorithms.

Hundreds of numerical integration methods for ordinary differential equations (ODEs) can be found in the literature.^{1–4} These algorithms, according to their features, are classified as one-step or multistep, implicit or explicit, fixed or variable step, fixed or variable order, etc. In spite of their differences, all the methods share a property: they are all based on time discretization, that is, given the solution at a time instant, they compute the approximate solution for the next discrete time point.

In recent years, a new family of algorithms was developed that replace the time discretization by the quantization of the state variables. Based on an original idea of Bernard Zeigler, who showed that continuous time systems with their inputs and outputs being quantized can be represented by discrete event system (DEVS) models,^{5–7} a general method for the numerical integration of ODEs called the quantized state system (QSS) was proposed by Kofman and Junco.⁸

The formulation of the QSS method was followed by second- and third-order accurate algorithms (QSS2⁹ and QSS3¹⁰, respectively). The QSS family has nice stability, convergence, and error bound theoretical properties,^{4,8,9,11} and, from a practical point of view, offers important advantages to detect and handle discontinuities.¹²

QSS methods are also very efficient for simulating large sparse systems,^{4,9} since they only invoke calculations on the states that experience sensible changes in their values or derivatives. In other words, QSS algorithms intrinsically exploit the *activity* of the system.^{13–15}

In this work, we explore the usage of QSS methods to simulate deterministic ODE models of spiking neural networks (SNNs). These models are emerging as a plausible paradigm for characterizing neural dynamics

CIFASIS – CONICET, FCEIA, Universidad Nacional de Rosario, Argentina.

Corresponding author:

Ernesto Kofman, Laboratorio de Sistemas Dinámicos, CIFASIS-CONICET, FCEIA-UNR, Riobamba 245 bis, 2000 Rosario, Argentina
Email: kofman@fceia.unr.edu.ar

in the cerebral cortex. The SNN models have high biological fidelity, and can model many characteristics of brain architecture.¹⁶

SNN models are usually large (they are composed of several neurons), sparse (each neuron is generally connected to a small subset of neurons), and each spike represents a discontinuity in the ODE. The presence of discontinuities and their large dimensions poses several difficulties to the different conventional numerical methods used by the SNN community. Consequently, the simulation of large networks of spiking neurons becomes slow and demands more suitable numerical methods.

Taking into account that QSS methods are good at discontinuity handling and sparsity exploitation, they are – in principle – good candidates to improve the performance of SNN simulation. This hypothesis is reinforced by the fact that previous works have shown some advantages of using the DEVS methodology in the simulation of non-ODE models of SNNs.^{17,18} In addition, in Wainer,¹⁹ a discrete version of the Hodgkin–Huxley model (which could be used to model SNNs) is efficiently implemented with Cell–DEVS and a similar approach, which approximates the continuous time behavior, is reported by Wainer and Giambiasi.²⁰

The goal of this work is then to implement simulations of ODE models of SNNs based on QSS methods and to compare their performance with those of conventional numerical methods currently in use for those models.

As a main contribution, we shall show that the principal advantage of using QSS methods to simulate SNNs is that the computational costs grow linearly with the number of neurons in the network, while classic integration techniques lead to quadratic growth (in the best case). In large networks, this property is translated into a sensible reduction of the central processing unit (CPU) time required to complete a simulation.

The article is organized as follows. Section 2 provides the background for the work, presenting different ODE models for SNNs, analyzing the different numerical methods used in the literature, and introducing the QSS methods, as well as the software tools for their implementation. Section 3 then describes the work done to model and to simulate SNNs in a software tool that implements the complete family of QSS algorithms. After that, Section 4 presents the simulation results and compares those obtained with different QSS methods with those obtained by classic methods for different SNN configurations. Finally, Section 5 finishes the work with conclusions and some ideas to continue with this line of research.

2. Modeling and simulation of spiking neural networks

This section provides the background on which the rest of the work is based. We first introduce different ODE models of SNNs and then we present the classic numerical methods used to simulate them. After that, we introduce the family of QSS methods and the software tools that implement them.

2.1. Models of spiking neural networks

A single spiking neuron can be described by a system of ODEs with discontinuities at the firing times. Several models have been proposed, with varying complexity. Among the most used we can mention the following ones.

- Hodgkin–Huxley model:²¹ this model was developed in the 1950s based on the experiments performed on the squid giant axon. Due to this fact, the model parameters have a clear observed meaning. The main inconvenience is that its simulation is very expensive due to the model complexity, as each neuron is represented by four equations governed by 10 parameters. Thus, its usage is limited to networks formed by few neurons.²²
- Integrate and fire:^{23,24} Contrary to Hodgkin–Huxley, this is an extremely simple model. A neuron is modeled by the equation

$$\dot{v}(t) = I(t) + a - bv(t), \quad \text{if } v(t) \geq v_u \text{ then } v \leftarrow c,$$

where $v(t)$ is the membrane potential, $I(t)$ is the input current, and a , b , c , and v_u are user-defined parameters to obtain different behaviors.

Due to its simplicity, the equation can be analytically integrated between discontinuities and large neural networks can be simulated. However, the model is not rich enough to represent many features that are usually observed in real neurons.

Several modifications have been proposed to improve this model, such as the inclusion of a quadratic term on $v(t)$ (quadratic integrate and fire²⁵), or the addition of a second state variable in order to represent more complex behaviors (integrate and fire or burst²⁶).

- The Izhikevich model:^{27,28} Recently, Izhikevich proposed a rather simple but versatile model that can represent different behaviors according to their parameter configuration. Its equations are

$$\begin{aligned} C\dot{v} &= kv(v - v_t) - u + I \\ \dot{u} &= a(bv - u) \end{aligned} \quad (1)$$

where $v(t)$ represents, as before, the membrane potential and $u(t)$ models the membrane restitution phenomenon. Variable $I(t)$ is the input current and the rest are user-defined constant parameters that allow one to obtain different types of behavior.

When at time t variable $v(t)$ reaches a given threshold value, the firing is produced and the state variables are reset as follows:

$$\begin{aligned} v(t^+) &= c \\ u(t^+) &= u(t) + d \end{aligned} \quad (2)$$

where c and d are user-defined constant parameters.

The parameters of Equation (1) are usually selected so that the equation becomes

$$\begin{aligned} \dot{v}(t) &= 0.04 \cdot v^2 + 5 \cdot v + 140 - u + I(t) \\ \dot{u}(t) &= a \cdot (b \cdot v - u) \end{aligned} \quad (3)$$

As can be observed in the three models described, there is always a trade-off between simulation costs and the possibility of reproducing the different behaviors observed in real neurons. There are several other models for SNNs and a comparison among them can be found in Izhikevich.²²

We shall use Izhikevich's model in this work, as it combines rich dynamics with fair computational effort requirements.

The interconnection between neurons (i.e. synapses) can be also modeled in different ways. In this work, we selected the synaptic current approach,^{16,29} also taking into account excitatory and inhibitory currents.³⁰

Thus, the original Izhikevich's model of Equation (1) was transformed into Equation (4):

$$\begin{aligned} C\dot{v} &= kv(v - v_t) - u - \eta(v - E_\eta) - \gamma(v - E_\gamma) + I \\ \dot{u} &= a(bv - u) \end{aligned} \quad (4)$$

Here, η and γ are the excitatory and inhibitory conductances, respectively, while E_η and E_γ are the *reversal potentials*. When a neuron receives the firing of an excitatory neuron, the excitatory conductance η increases its value in 6 nS, while when the firing comes from an inhibitory neuron, then γ is increased in 67 nS. The rest of the time, η and γ decay exponentially following a first-order dynamic:

$$\begin{aligned} \dot{\eta} &= -\lambda_\eta \cdot \eta \\ \dot{\gamma} &= -\lambda_\gamma \cdot \gamma \end{aligned} \quad (5)$$

where λ_η and λ_γ are the decay rate parameters.

2.2. Numerical integration of SNN models

When a classic numerical method is used to simulate a SNN a problem appears. The models of SNNs exhibit a discontinuous behavior each time a neuron fires.

Conventional numerical methods cannot integrate across a discontinuity. The numerical integration methods are always based on polynomial approximations of some functions that are no longer valid when those functions are discontinuous. Thus, when they perform an integration step that jumps through a discontinuity, the error committed is unacceptable. To solve this problem, the methods must detect the discontinuity first, advance the simulation time until that instant, and then restart the simulation from the new situation.⁴

SNN discontinuities belong to the *state event* type, that is, their occurrence depends on a condition on the state variables (typically, a zero-crossing of some signal). The detection of this type of discontinuity requires iterations, and it is computationally expensive. The simulation restart is also time consuming, as the simulation step size must be restarted, typically from a small value.

From a practical point of view, the interesting SNN models are those composed of several neurons. As firings at different neurons occur at different times, the rate of firings in a network grows linearly with the number of neurons.

This is, if each neuron provokes on average 100 firings per second, a network of 1000 neurons will provoke about 100,000 firings per second. Then, any classic numerical method simulating a single neuron will have an upper bound for its step size of 1/100 in order to properly handle discontinuities (accuracy and stability considerations may impose a lower upper bound). However, when simulating a network of 1000 neurons that upper bound will be as low as 1/100,000. In addition, each step will be 1000 times more expensive as it involves the evaluations on the derivatives of 1000 times more states.

Due to these facts, the computational cost grows at least quadratically with the number of neurons in the network.

Yet, classic numerical methods are widely used for simulating SNNs. Among them, Runge-Kutta and Bulirsch-Stoer³ appear frequently reported in the literature.

A recent work proposes a discrete event-like solution based on the linearization of the ODE and its analytical solution between firings,³¹ yet for neuron models involving more than one state variable this approach is only first-order accurate. Thus, for the case of Izhikevich's model – which is a second-order system – this approach cannot offer decent results. The fact that the algorithm results are only first-order accurate implies that it will not be able to achieve an acceptable accuracy without increasing enormously the computational costs. However, it is worth mentioning as this solution, called *voltage stepping*, has some connection with the QSS methods that we describe below.

2.3. Quantized state system simulation

While all classic methods are of the *discrete time* type, a new approach was recently developed that replaces time discretization by state quantization.

The first of these algorithms was the QSS method of first order (QSS1), which is defined below.

Consider a time invariant ODE:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \quad (6)$$

where is $\mathbf{x}(t) \in R^n$ the state vector and $\mathbf{u}(t) \in R^m$ is an input vector, which is a known piecewise constant function.

The QSS1 method^{4,8} analytically solves an approximate ODE, which is called QSS:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{q}(t), \mathbf{u}(t)) \quad (7)$$

where $\mathbf{q}(t)$ is a vector of quantized variables that are quantized versions of the state $\mathbf{x}(t)$. Each component $q_j(t)$ of $\mathbf{q}(t)$ follows a piecewise constant trajectory, related with the corresponding component $x_j(t)$ of $\mathbf{x}(t)$ by a hysteretic quantization function so that (we denote $q_j(t^-) = \lim_{\tau \rightarrow t^-} q_j(\tau)$, that is, the limit from the left-hand side of $q_j(t)$)

$$q_j(t) = \begin{cases} x_j(t) & \text{if } |q_j(t^-) - x_j(t)| = \Delta Q_j \\ q_j(t^-) & \text{otherwise} \end{cases} \quad (8)$$

and $q_j(t_0) = x_j(t_0)$. That is, $q_j(t)$ only changes when it differs from $x_j(t)$ by $\pm \Delta Q_j$. The magnitude ΔQ_j is called quantum. Figure 1 shows a typical QSS1 quantization function.

The piecewise constant evolution of the quantized variables $q_j(t)$ implies that the state derivatives $\dot{x}_j(t)$ follow piecewise constant trajectories, and then the states $x_j(t)$ evolve in a piecewise linear way. These facts permit the analytical integration of the system of Equation (7) in a straightforward manner.

The QSS1 method has some nice stability and global error bound properties.^{4,9,11} Yet, it performs only a first-order approximation and a good accuracy cannot be obtained without significantly increasing the number of steps.

A second-order accurate method called QSS2 was proposed by Kofman.⁹ QSS2 has the same definition as QSS1, except that the components of $q_j(t)$ are now calculated to follow piecewise linear trajectories. Figure 2 shows a typical evolution of state and quantized variables.

In QSS2 the state derivatives $\dot{x}_j(t)$ are computed as piecewise linear trajectories and then the states $x_j(t)$ follow piecewise parabolic trajectories. Like in QSS1, the analytical solution of the quantized system of

Equation (7) can be easily obtained following a simple algorithm.

QSS1 and QSS2 ideas were also extended to obtain a third-order accurate method called QSS3.¹⁰ In QSS3, quantized variables follow piecewise parabolic trajectories and states are piecewise cubic.

Both methods, QSS2 and QSS3, share the stability and error bound properties of QSS1. The family of QSSs is completed with a Backward QSS algorithm (BQSS) and the Centered QSS (CQSS) conceived to integrate stiff and marginally stable systems, respectively.³²

Since Izhikevich's model is not stiff or marginally stable, we shall use QSS3 as it offers the maximum accuracy order among all the QSS methods. Thus, it will obtain the best relationship between accuracy and computational costs.

The quantization functions shown in Figures 1 and 2 use a uniform quantum ΔQ_j . Alternatively, QSS




Figure 1. Hysteretic quantization.




Figure 2. QSS2 quantization.

methods can use logarithmic quantization, where the quantum is proportional to the corresponding state magnitude.³³

When uniform quantization is used, the absolute global simulation error is bounded by a linear function of the quantum.⁴ In the case of logarithmic quantization, the relative global simulation error is intrinsically controlled.³³

Steps in QSS methods are only produced when some quantized variable $q_j(t)$ changes, that is, when the corresponding state $x_j(t)$ differs from $q_j(t^-)$ in a quantum. That change implies also that some state derivatives (those that depend on x_j) are also changed. Then, each step involves a change in only one quantized variable and in some state derivatives.

Thus, when a large sparse system experiences activity only in a few states while the rest of the system remains unchanged, the QSS methods intrinsically exploit this fact, performing computations only when and where the changes occur.

Another important advantage of the QSS methods is that they handle discontinuities in a straightforward and very efficient manner.¹² According to the order of the method, the state variables follow piecewise linear, parabolic, or cubic trajectories. Then, detecting zero crossings is straightforward, as it involves solving a cubic equation in the worst case. Once a discontinuity is detected, the algorithm handles it as an ordinary step, since each step is in fact a discontinuity in a quantized variable. Hence, the occurrence of a discontinuity implies only some local calculations to recompute the state derivatives that are directly affected by that event.

2.4. QSS methods and DEVS

Although QSS algorithms can be easily coded in any programming language, they are usually implemented as *discrete event systems* within the DEVS formalism framework.⁷

Notice that each component of Equation (7) can be considered as the coupling of two elementary subsystems. A static one:

$$d_j(t) = f_j(q_1, \dots, q_n, u_1, \dots, u_m) \quad (9)$$

and a dynamical one:

$$q_j(t) = Q_j(x_j(\cdot)) = Q_j\left(\int d_j(\tau)d\tau\right) \quad (10)$$

In the case of the first-order QSS1 method, 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 $u_j(t)$ and $q_j(t)$ are piecewise constant, the output of subsystem (9), that is, $d_j(t)$, will

be piecewise constant. In this way, both subsystems have piecewise constant input and output trajectories.

If we represent every change of a piecewise constant trajectory as an *event*, then the trajectories can be considered sequences of events.

Thus, the subsystems of Equations (9) and (10) can be seen as DEVs that process event sequences.

The DEVS formalism⁷ allows one to describe any model that processes event sequences and the representation of the models of Equations (9) and (10) is in fact very simple. The DEVS equivalent of Equation (9) is called the *static function* and the DEVS equivalent of Equation (10) is called the *hysteretic quantized integrator*. Their definitions can be found in Cellier and Kofman.⁴

Figure 3 shows the block diagram representation of Equation (7). It is composed by n static subsystems, like Equation (9), and n dynamic subsystems, like Equation (10).

By connecting the DEVS models corresponding to static functions and quantized integrators following the block diagram of Figure 3, a coupled DEVS model is obtained that exactly represents the dynamics of Equation (7).

The same idea can be applied to represent QSS2 approximations as DEVS models, but now the events carry two variables with the initial value and slope of each segment of a piecewise linear trajectory. The DEVS models corresponding to static functions and quantized integrators are more complex, since they take into account also the slopes of the corresponding trajectories.

The QSS3 method can be also be implemented following these ideas.

The family of QSS methods, including QSS1, QSS2, QSS3, BQSS, and CQSS algorithms, were implemented in PowerDEVS,³⁴ a DEVS-based simulation software.

Figure 4(a) shows the QSS continuous library of PowerDEVS. The blocks contained in this library implement the quantized integrators and static functions for the mentioned QSS methods.

Hybrid systems can be represented and simulated in PowerDEVS coupling blocks of the continuous library with the blocks of the QSS hybrid library shown in Figure 4(b). The hybrid blocks are DEVS models that handle different types of discontinuities, making use of the piecewise polynomial features of quantized variable trajectories in QSS methods.

For instance, the hybrid system corresponding to Izhikevich's model of Equations (2) and (3) can be implemented by the PowerDEVS model of Figure 5.

The models of continuous and hybrid systems in PowerDEVS can be built using the classic block diagram approach. That is, we include an integrator for each state variable and then we build the expression of

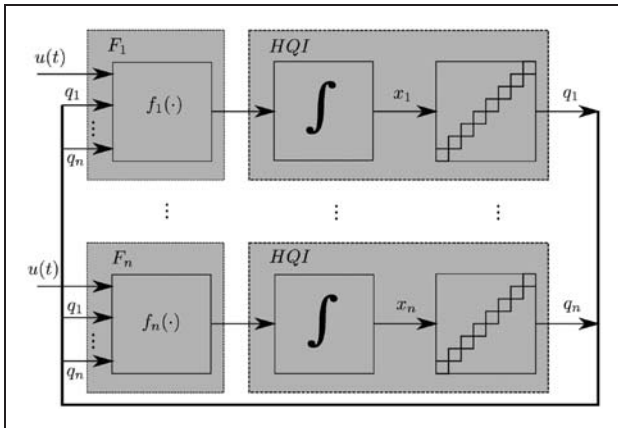


Figure 3. Block diagram representation of a QSS1 approximation.

the state derivatives using continuous and/or hybrid blocks. In addition, for input signals, we use source blocks.

The first-order accurate QSS1 algorithm was also implemented in other simulation tools: CD++,³⁵ VLE,³⁶ and Dymola/Modelica.³⁷ However, those implementations do not include higher order methods, such as QSS3 (with the exception of ModelicaDEVS, but this implementation is not efficient enough³⁷), so we shall focus only on PowerDEVS.

3. PowerDEVS modeling and simulation of spiking neural networks

This section describes the work done to simulate SNNs in PowerDEVS using QSS methods. We first describe

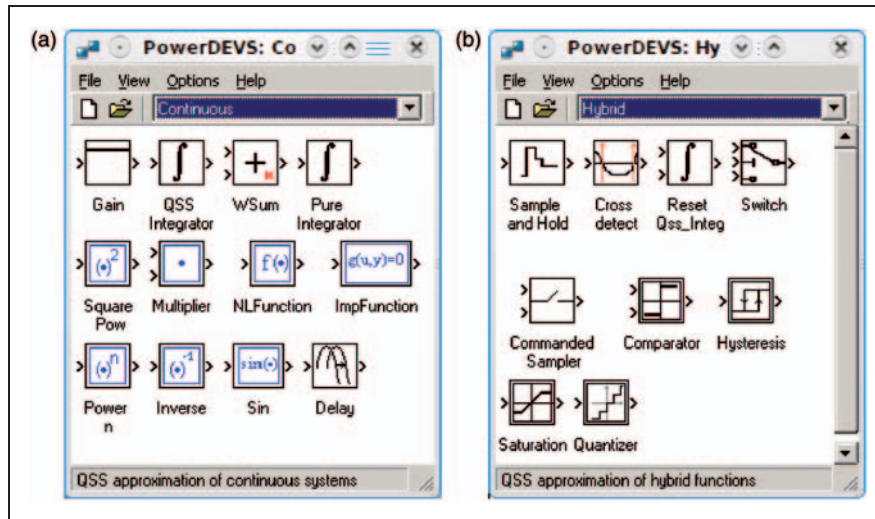


Figure 4. PowerDEVS QSS continuous and hybrid libraries.

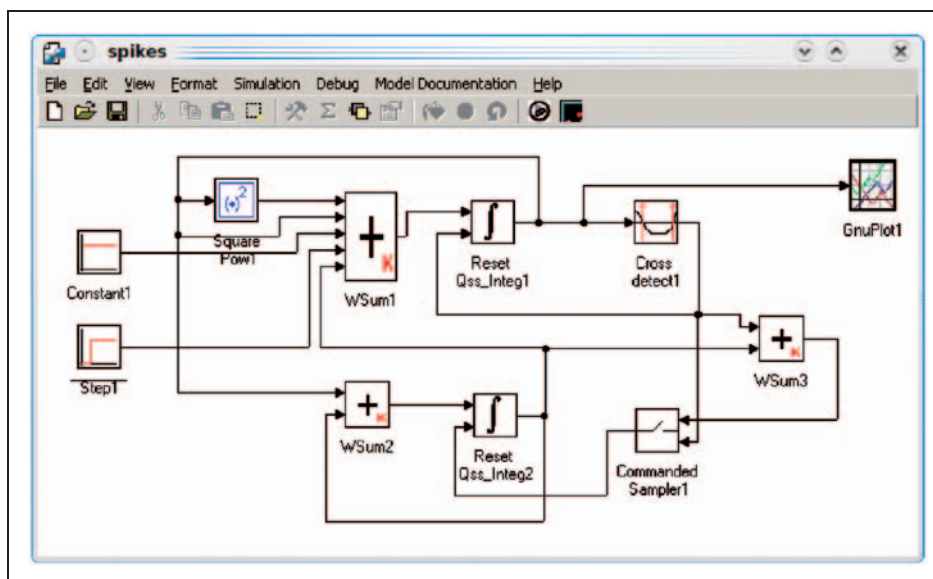


Figure 5. Izhikevich's model in PowerDEVS.

the PowerDEVS model of a single isolated neuron, then the model of a neuron with synapsis interface, and finally the model of a network of neurons and the modifications introduced in PowerDEVS in order to efficiently simulate large coupled systems.

3.1. PowerDEVS model of a single isolated neuron

A single isolated neuron, following Izhikevich's model of Equations (1) and (2), was modeled in PowerDEVS as Figure 5 shows. This block diagram was built following the classic procedure mentioned above, and it is the direct representation in block diagrams of the corresponding differential equations.

There, the blocks 'Reset QSS_Integ1' and 'Reset QSS_Integ2' compute the states v and u , respectively. Similarly, the static blocks 'Wsum1', 'Square Pow1', and 'Wsum2' calculate the state derivatives, making use also of the source blocks 'Constant1' and 'Step1'. The block 'Cross detect1' produces an event each time its input signal crosses a given level, in this case when the condition $v = 60$ is reached. This event is used to reset the first integrator to the value $v = -30$

and to compute the signal $u - 30$ and reset the second integrator.

3.2. PowerDEVS model of a neuron with synapsis interfaces

The addition of the synapsis ports transforms Equations (1) and (2) into Equations (4) and (5). Thus, we modified the PowerDEVS model of Figure 5 by adding new blocks, as Figure 6 shows.

In this new PowerDEVS block diagram, the new blocks 'Reset QSS_Integ3' and 'Reset QSS_Integ4' calculate the new states η and γ , and the blocks 'Gain1' and 'Gain2' compute the corresponding state derivatives. The blocks 'Commanded Sampler2' and 'Commanded Sampler3' are in charge of resetting the integrators of η and γ when they receive the events that indicate the firing of other neurons.

The blocks 'WSum4', 'WSum5', 'Wsum6', 'Multiplier1', and 'Multiplier2' calculate the additional terms of the derivative $\dot{v}(t)$ at the right-hand side of Equation (4), and finally, the blocks 'WSum7' and 'WSum8' compute the reset values for η and γ .

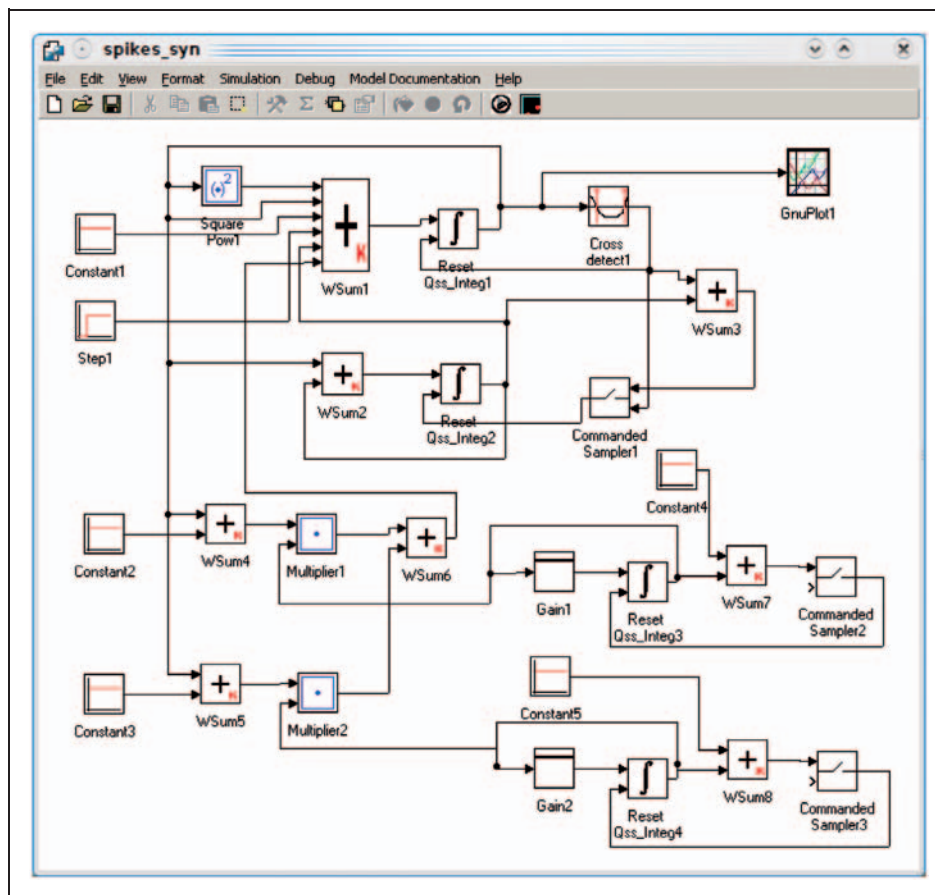


Figure 6. Izhikevich's model with synapsis in PowerDEVS.

3.3. Modeling and simulation of a large neural network

Modeling and simulating a large network of neurons, such as those of Figure 6 using PowerDEVS requires one to solve some problems first.

3.3.1. Model construction. The first of the problems is related to the construction of a very large model with a graphical user interface (GUI), such as that of PowerDEVS. Although it allows one to encapsulate complex models so that the entire block diagram of Figure 6 appears as a single block, copying and coupling 1000 of these blocks (in order to simulate a network of 1000 neurons) is impossible.

We also wanted to generate the connection structure in a random way in order to perform multiple simulation runs of different networks and the GUI of PowerDEVS does not have that capability.

So, we used the PowerDEVS GUI only to build the model of a single neuron (with and without synapsis ports) and then we wrote a C++ program that replicates that model as many times as we want and that generates random interconnections (if necessary) between the different neurons.

3.3.2. Model simulation. The second problem we faced was related to the size and the structure of the model and the way PowerDEVS searches the next event time and the next transitioning model.

Notice that the model of a single neuron with synapsis interfaces (Figure 6) contains 27 blocks (without taking into account the Gnuplot block in charge of plotting the results). Thus, a model of 1000 connected neurons has 27,000 atomic blocks.

After producing and propagating an event, the simulation engine must find which is the block that provokes the next event and when that event is being produced. As PowerDEVS was not originally designed to simulate large networks, it simply looks at the next event time of each block and then it takes the one having the smallest (in case the that two or more blocks share the minimum time, PowerDEVS uses a priority list to decide among them).

This solution, although being appropriate for small models, is completely inefficient for a large network, as it implies searching along the whole structure at each step. Thus, the cost of each search for the next event time grows linearly with the size of the system.

Taking into account that the number of neuron firings grows linearly with the size of the network, so will grow the number of simulation events. If the search for the next event time also grows linearly, then the total cost will grow at least quadratically with the network dimension. Consequently, we shall obtain something

similar to conventional numerical methods. Thus, if we want to improve this, we need to avoid performing a linear search for the next event time.

The problem of optimizing the search for the next event time in large models has been previously discussed in the literature.

The DEVS abstract simulator of Zeigler et al.⁷ proposes a tree structure for the DEVS model. That way, the search for the next event time is no longer linear. However, in a SNN model, we can have connections from any neuron to any other neuron and there is no natural way of splitting the model in a tree-like form at the modeling stage. The best we can do to reduce the size of the higher hierarchy level is to build a single coupled model for each neuron. However, the number of neurons is still very large and the cost of performing a linear search among them is still linear.

In order to improve the performance of the time scheduling procedure some authors have proposed parallelizing part of the algorithms and also reducing the search space among a subset of atomic models having a minimum time advance.^{38,39} As we want to show an algorithm working in a single processor, we cannot take the first solution. The second one, although it would improve the performance, will not avoid the linear growth.

Another solution consists in leaving aside from the search those atomic models in passive states, that is, with time advance equal to infinite.^{40,41} As before, this solution in a SNN network will not prevent the search for the next event time from growing linearly.

Finally, to accelerate the search among the models with finite time advance, a heap structure is usually implemented.^{40,41}

As this solution reduces the cost from linear to logarithmic, we made our implementation based on that. We could have also combined this idea with the other solutions analyzed above. That way, we might have improved further the results. However, we only wanted to reduce the time search from linear to logarithmic in order to prevent the total simulation cost from growing quadratically with the size of the network.

We organized the atomic models with a binary tree structure. Every leaf of the tree contains a reference to an atomic model and to its next event time. Every node of the tree contains the minimum time of its children and the reference to the corresponding atomic model. Thus, the main node of the tree contains the minimum time of the network and the reference to the corresponding atomic model.

After an event, for each atomic model that changed its next event time, we only need to compare it with the time of its brother. If the smaller time does not change, nothing else has to be done. Otherwise, we need to

propagate the new minimum time to its father node that will in turn compare it back with its brother node. In the worst case, the propagation will reach the main node. In that worst case, only about $\log_2(N)$ comparisons are performed (where N is the number of blocks) for each atomic model that changed its time advance.

As the number of connections per neuron is bounded by a constant in these models, each event is propagated to a maximum number of blocks that does not depend on the total number of blocks. Thus, the number of comparisons needed after each transition is bounded in a logarithmic way.

According to some preliminary results, the usage of the binary tree allowed us to reduce the total simulation time by two orders of magnitude in the simulation of a model with 1000 neurons.

4. Simulation results

This section shows and discusses the simulation results. Here, we compare the results obtained with the QSS3 method (in PowerDEVS) using constant and logarithmic quantization, with those obtained using Runge–Kutta–Fehlberg (RK45) and Bulirsch–Stoer (BS) (for both methods we used the code provided in Press et al.³ with the addition of discontinuity handling routines).

We worked on three examples: a single isolated neuron, a large network of isolated neurons, and finally a large network of interconnected neurons.

4.1. Simulation of a single neuron

We first simulated a single neuron with a final simulation time of 1 second. We used the model of Equation (3). For the parameters, we followed Izhikevich,²⁷ choosing $a=0.02$, $b=0.2$, $c=-65+15\cdot r^2$, and $d=8-6\cdot r^2$, where r is a random variable with uniform distribution in the interval $[0,1]$. Similarly, $I(t)$ was taken as a constant function with random value.

We ran each simulation 100 times in order to obtain meaningful statistical data. We measured, for different tolerance and quantum settings, the mean absolute error and the total simulation time. We used as reference solution that obtained with QSS3 using a quantum $\Delta Q = 10^{-11}$.

We used this solution as a reference, since the analytical solution cannot be obtained. In order to check that the different methods converge to the reference solution, we included simulations performed with those methods using small tolerances.

Figure 7 shows the state trajectories for a particular simulation run. They coincide with the typical solutions of Izhikevich's model.

The difference between the solutions given by the three numerical integration methods cannot be distinguished by the naked eye.

Tables 1–4 summarize the results. Figures 8 and 9 illustrate the evolution of the CPU time and error as a function of the tolerance or quantum parameter for the different methods.

For small tolerances, we see that the error of BS and RK45 (i.e. the difference with the reference solution of QSS3) is around $1E-8$. This means that the reference solution is useful up to this tolerance.

The results exhibit a clear advantage of the RK45 and BS methods over QSS3. The dynamics do not exhibit many discontinuities and the system is fully interconnected. Thus, QSS3 does not have much to offer here. Moreover, the PowerDEVS implementation (Figure 5) is quite inefficient, as every event of the integrators is propagated to several static atomic blocks, which are in turn propagated back along the coupled structure. So, the simulation engine spends more time handling the coupling structure than doing useful calculations. In addition, the PowerDEVS engine performs some

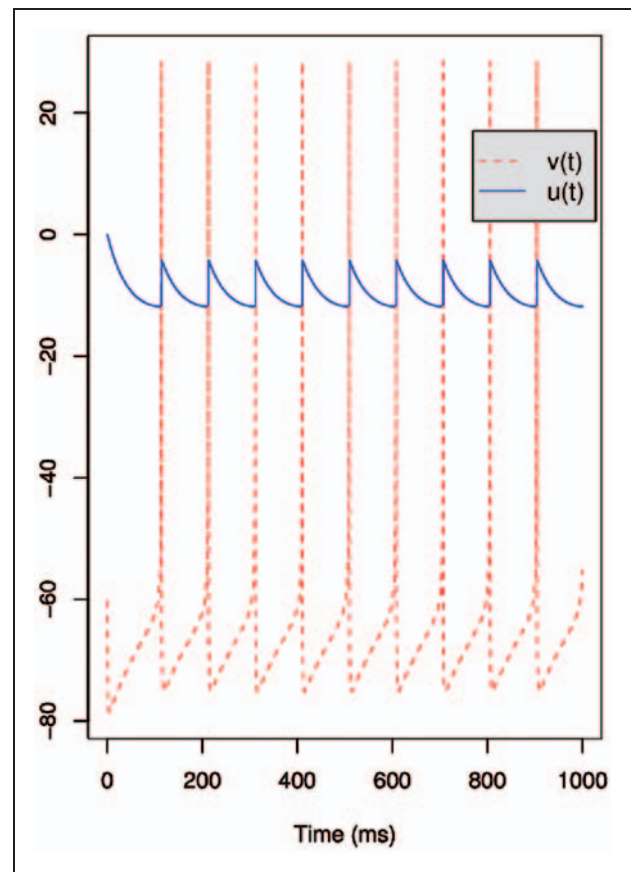


Figure 7. State trajectories of a single neuron simulation using QSS3.

Table 1. Simulation with RK45. Parameter = relative tolerance

Parameter	Mean CPU time in sec. (variance)	Mean error (variance)
10^{-03}	7.21E-04 (1.69E-04)	6.51E-01 (4.17E-01)
10^{-04}	7.93E-04 (1.76E-04)	9.95E-02 (6.91E-02)
10^{-05}	8.85E-04 (1.93E-04)	6.62E-03 (7.19E-03)
10^{-06}	1.05E-03 (2.09E-04)	3.78E-04 (5.43E-04)
10^{-07}	1.32E-03 (2.41E-04)	4.71E-05 (5.71E-05)
10^{-08}	1.78E-03 (3.00E-04)	4.96E-06 (6.07E-06)
10^{-09}	2.26E-03 (3.42E-04)	6.31E-07 (5.35E-07)
10^{-10}	3.22E-03 (4.63E-04)	8.13E-08 (4.43E-08)
10^{-11}	4.80E-03 (6.50E-04)	7.16E-09 (7.76E-09)

CPU: central processing unit.

Table 2. Simulation with BS. Parameter = relative tolerance

Parameter	Mean CPU time in sec. (variance)	Mean error (variance)
10^{-03}	1.12E-03 (2.41E-04)	1.41E + 00 (6.84E-01)
10^{-04}	1.44E-03 (3.03E-04)	3.67E-01 (2.31E-01)
10^{-05}	1.53E-03 (4.59E-04)	1.40E-01 (9.88E-02)
10^{-06}	2.17E-03 (4.40E-04)	1.05E-02 (1.69E-02)
10^{-07}	2.23E-03 (4.93E-04)	1.70E-04 (2.30E-04)
10^{-08}	2.72E-03 (6.53E-04)	1.53E-05 (1.11E-05)
10^{-09}	2.88E-03 (5.71E-04)	1.26E-06 (1.04E-06)
10^{-10}	3.32E-03 (7.19E-04)	1.57E-05 (8.49E-05)
10^{-11}	3.66E-03 (7.38E-04)	1.24E-08 (1.31E-08)

CPU: central processing unit.

Table 3. Simulation with QSS3. Parameter = ΔQ

Parameter	Mean CPU time in sec. (variance)	Mean error (variance)
10^{-03}	2.35E-02 (1.53E-03)	1.16E-01 (7.50E-02)
10^{-04}	3.17E-02 (3.29E-03)	1.74E-02 (9.74E-03)
10^{-05}	5.05E-02 (6.64E-03)	1.52E-03 (2.11E-03)
10^{-06}	8.88E-02 (1.42E-02)	2.64E-04 (1.04E-03)
10^{-07}	1.74E-01 (3.11E-02)	7.74E-06 (2.94E-06)
10^{-08}	3.58E-01 (6.71E-02)	7.63E-07 (3.12E-07)
10^{-09}	7.55E-01 (1.53E-01)	8.26E-08 (5.21E-08)
10^{-10}	1.62E + 00 (3.16E-01)	6.71E-09 (6.37E-09)
10^{-11}	3.47E + 00 (6.32E-01)	0.00E + 00 (0.00E + 00)

CPU: central processing unit.

initialization procedures before each simulation that take some fixed time (it takes about 0.01 seconds), so there is a lower bound for the total simulation time.

When it comes to errors and tolerance, we found that the performance of QSS3 compared with RK45

and BS worsens when the tolerance becomes more stringent (Figure 8). This is due to the fact that QSS3 is only a third-order accurate algorithm and thus the CPU time grows with the cubic root of the required accuracy.

Table 4. Simulation with logarithmic QSS3. Parameter = ΔQ_{rel}

Parameter	Mean CPU time in sec. (variance)	Mean error (variance)
10^{-03}	1.92E-02 (8.78E-04)	1.15E+00 (3.08E-01)
10^{-04}	2.26E-02 (3.74E-03)	1.66E-01 (1.10E-01)
10^{-05}	2.82E-02 (2.97E-03)	6.73E-03 (8.38E-03)
10^{-06}	4.12E-02 (5.60E-03)	1.41E-03 (3.55E-03)
10^{-07}	6.96E-02 (1.08E-02)	4.67E-04 (1.39E-03)
10^{-08}	1.31E-01 (2.31E-02)	1.89E-04 (8.97E-04)
10^{-09}	2.66E-01 (5.33E-02)	3.02E-06 (5.27E-06)
10^{-10}	5.51E-01 (1.09E-01)	3.02E-07 (5.46E-07)
10^{-11}	1.18E+00 (2.50E-01)	1.89E-08 (2.52E-08)

CPU: central processing unit.




Figure 8. CPU time versus tolerance (quantum) for different methods.

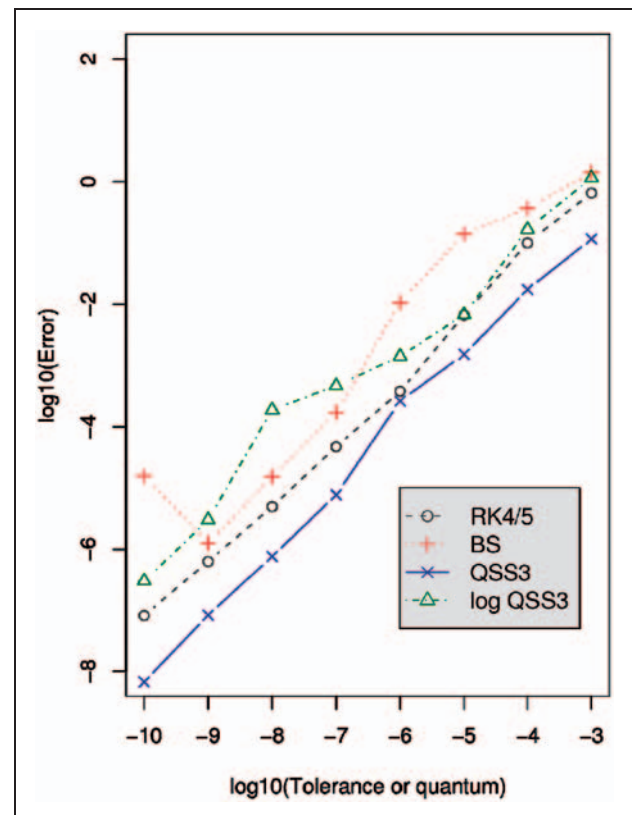


Figure 9. Error versus tolerance (quantum) for different methods.

4.2. Simulation of networks of isolated neurons

In this second experiment, we simulated different systems of disconnected neurons. We ran different simulations of systems with 10, 100, 1000, and 2000 neurons using the same methods of the previous example. In order to have similar errors with the different

algorithms, we selected a tolerance of 10^{-4} for RK45 and 10^{-5} for BS. Similarly, we selected a quantum $\Delta Q = 10^{-3}$ for QSS3 and $\Delta Q_{rel} = 10^{-4}$ for QSS3 with logarithmic quantization. With these values, according to the simulation results of the previous example, all the methods have a mean absolute error around 10^{-1} .

Table 5. Mean CPU time in seconds for simulating 10, 100, 1000, and 2000 disconnected neurons. The variance is reported between parentheses

Method	10 neurons	100 neurons
RK45	1.41E-02 (1.78E-02)	1.31E-01 (3.00E-02)
BS	1.30E-02 (5.27E-03)	3.70E-01 (8.33E-02)
QSS3	2.08E-01 (2.02E-01)	7.70E-01 (2.19E-01)
QSS3 log	1.24E-01 (8.39E-02)	5.93E-01 (9.40E-02)
Method	1000 neurons	2000 neurons
RK45	1.04E + 01 (8.36E-01)	3.88E + 01 (3.30E + 00)
BS	3.71E + 01 (4.33E + 00)	1.35E + 02 (1.53E + 01)
QSS3	7.89E + 00 (2.09E + 00)	1.72E + 01 (3.49E + 00)
QSS3 log	7.07E + 00 (1.74E + 00)	1.61E + 01 (3.51E + 00)

Table 5 exhibits the mean time (after 30 simulation runs) required by each method to complete the simulation. The final simulation time was set to 1 second.

Figure 10 plots the CPU time required by each method for the different number of neurons.

The results become now more interesting. For small systems (10–100 neurons), QSS3 methods show a poor performance compared with BS and RK45. However, when the number of neurons increases, QSS3 rapidly outperforms both discrete time algorithms.

This fact can be easily explained taking into account the way in which QSS methods exploit activity. As we mentioned in Section 2.2, discrete time algorithms must evaluate all the state derivatives at each step, and they must restart the whole simulation after each discontinuity (i.e. after each firing). As the number of firings, as well as the size of the right-hand side function of the ODE, grow linearly with the size of the network, the discrete time algorithms experience a quadratic growth of the CPU time.

Although the QSS3 method is inefficient to simulate a single neuron, every step and every firing only provokes calculations at one neuron. Thus, the algorithm only performs computations where the changes occur, that is, where the system shows some activity. In consequence, the CPU time grows almost linearly with the number of neurons. This fact can be observed in Table 5.

4.3. Simulation of networks of interconnected neurons

In this last test we simulated different networks of interconnected neurons. To this end, we followed the scheme proposed in the Benchmark I of Brette et al.¹⁶

We considered networks composed of 400, 1000, 2000, and 4000 neurons. In each case, 80% of the neurons are of the excitatory type while the remaining 20% are of the inhibitory type. We also established that each


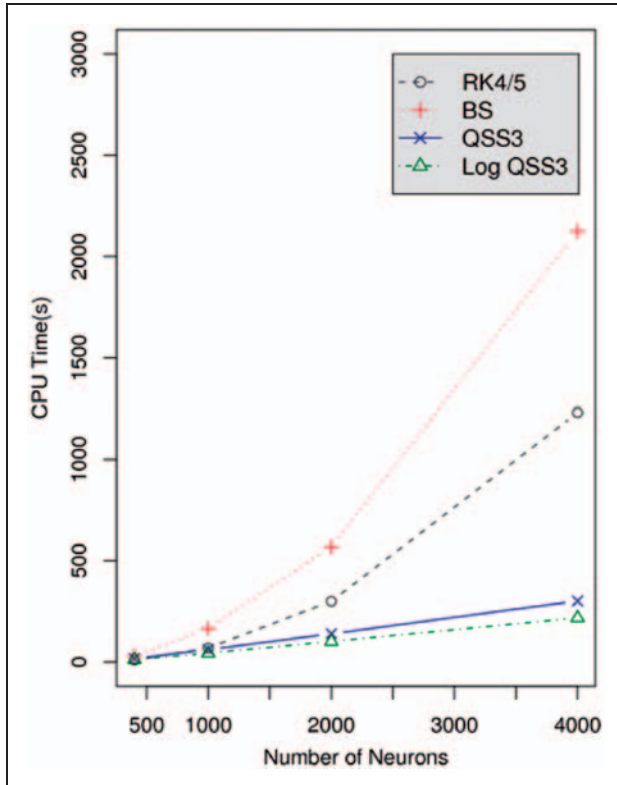
**Figure 10.** CPU time (seconds) for RK45, BS, QSS3, and QSS3 logarithmic in the simulation of a system of 10, 100, 1000, and 2000 disconnected neurons.

Table 6. Mean CPU Time (seconds) for 400, 1000, 2000, and 4000 interconnected neurons. The variance is reported between parentheses

Method	400	1000	2000	4000
RK45	9.79 (1.62)	68.89 (7.94)	300.29 (25.84)	1231.24 (76.98)
BS	30.08 (3.69)	165.58 (16.40)	568.40 (43.50)	2124.14 (293.23)
QSS3	17.74 (1.59)	61.33 (2.71)	139.76 (4.79)	301.76 (11.00)
QSS3 log	12.56 (0.85)	44.15 (2.19)	101.80 (3.55)	218.19 (7.61)

**Figure 11.** Mean CPU time versus number of neurons.

neuron is connected (in a random way) with 80 neurons.

We compare the performance of the four algorithms of the previous example, with the same tolerance and quantum settings, but now we selected a final simulation time of 250 milliseconds. As before, we ran each simulation 30 times and computed the mean CPU time required for each algorithm. Table 6 summarizes the results.

Figure 11 plots the CPU time as a function of the number of neurons for each algorithm.

These new results agree with the previous case of disconnected neurons. When the number of neurons becomes larger, QSS3 methods become more efficient than both discrete time algorithms.

As we observed in the previous example, Figure 11 shows that the CPU time grows almost linearly in QSS3 methods, while it grows at least quadratically in the other algorithms.

5. Conclusions

In this article we studied the use of QSS algorithms in the simulation of SNNs. We found that their efficient discontinuity handling and their activity-driven features offer very important advantages over the most widely used discrete time methods.

The greatest advantage exhibited by the QSS methods is that their computational costs grow almost linearly with the number of neurons, while they grow at least quadratically for discrete time methods. Consequently, the performance of the QSS algorithms in the simulation of large networks of neurons is clearly superior.

As a side result, we implemented an algorithm to manage the time advance of the simulation engine based on a binary tree of the submodels. While the cost for computing the next event time is linear with the number of submodels in most DEVS implementations, in our algorithm its results are only logarithmic.

Although the results are promising and the QSS methods noticeably improve the performance of discrete time algorithms in the systems analyzed, some work has to be done before claiming that this discrete event approach constitutes a valid and general alternative for SNN simulation.

Firstly, we need to perform experiments with different models (so far, we have used only Izhikevich's model with a particular type of synapse). We conjecture that we shall find the same advantages, but we cannot affirm that without running simulations.

In the current work, as we explained in Section 3, the models were built using an ad hoc C++ program that generated the PowerDEVS model structure. If we want to convince people from the neurosciences community to use our algorithms, we definitely need to develop a better end-user interface.

Moreover, it is possible that PowerDEVS (or any other DEVS-based modeling tool) is not the best

choice to implement these simulations. Any general-purpose DEVS simulation engine performs several tasks that are not useful in the context of simulating a QSS approximation of an ODE. We are currently working on the development of standalone QSS solvers. Some preliminary results show a simulation speedup of one order of magnitude with respect to the same QSS approximation executed by PowerDEVS. Thus, if we can develop a specific end-user interface for modeling large SNNs and then we integrate them with these standalone solvers, we can expect a significant improvement of the results shown in this article.

Finally, for larger networks of neurons, we will need to tackle the problem of parallelization of the algorithms.

Funding

This work was supported by CONICET (grant number PIP-2009/2011-00183).

Conflict of interest statement

None declared.

References

- Hairer E, Norsett S and Wanner G. *Solving ordinary differential equations I*. Berlin: Springer, 1991.
- Hairer E and Wanner G. *Solving ordinary differential equations II. Stiff and differential-algebraic problems*. Berlin: Springer, 1991.
- Press WH, Teukolsky SA, Vetterling WT, and Flannery BP. *Numerical Recipes in C: The Art of Scientific Computing*. Second Edition. Cambridge University Press, New York, 1992.
- Cellier FE and Kofman E. *Continuous system simulation*. New York: Springer, 2006.
- Zeigler B. DEVS representation of dynamical systems: Event-based intelligent control. *Proc IEEE* 1989; 77: 72–80.
- Zeigler B and Lee JS. Theory of quantized systems: formal basis for DEVS/HLA distributed simulation environment. In: *Proceedings of SPIE*, 1998, pp.49–58.
- Zeigler B, Kim TG and Praehofer H. *Theory of modeling and simulation*, 2nd ed. New York: Academic Press, 2000.
- Kofman E and Junco S. Quantized state systems. A DEVS approach for continuous system simulation. *Trans SCS* 2001; 18: 123–132.
- Kofman E. A second order approximation for DEVS simulation of continuous systems. *Simulation* 2002; 78: 76–89.
- Kofman E. A third order discrete event simulation method for continuous system simulation. *Lat Am Appl Res* 2006; 36: 101–108.
- Nutaro J and Zeigler B. On the stability and performance of discrete event methods for simulating continuous systems. *J Comput Phys* 2007; 227: 797–819.
- Kofman E. Discrete event simulation of hybrid systems. *SIAM J Sci Comput* 2004; 25: 1771–1797.
- Jammalamadaka R. Activity characterization of spatial models: application to discrete event solution of partial differential equations. Master's thesis, The University of Arizona, 2003.
- Muzy A, Aiello A, Santoni PA, Zeigler BP, Nutaro JJ, and Jammalamadaka R. Discrete event simulation of large-scale spatial continuous systems. In: *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 2005, Vol. 4, pp.2991–2998.
- Muzy A, Nutaro J, Zeigler BP, and Coquillard P. Modeling and simulation of fire spreading through the activity tracking paradigm. *Ecological Modelling*, 2008; 219(1–2): 212–225.
- Brette R, Rudolph M, Carnevale T, Hines M, Beeman D, Bower JM, et al. Simulation of networks of spiking neurons: a review of tools and strategies. *Journal of Computational Neuroscience*, 2007; 23(3): 349–398.
- Mayrhofer R, Affenzeller M, Prähofer H, Hfer G, and Fried A. DEVS simulation of spiking neural networks. In: *Proceedings of Cybernetics and Systems (EMCSR)*, 2002, Vol. 2, pp.573–578.
- Boiko Y and Wainer G. Modeling of neural decoder based on binary spiking neurons in DEVS. In: *Proceedings of the 2009 Spring Simulation Multiconference*, San Diego, CA, 2009.
- Wainer G. Performance analysis of continuous cell-DEVS models. In: *Proceedings of High Performance Computing & Simulation Conference (HPC&S), 18th European Simulation Multiconference*, Magdeburg, Germany, 2004.
- Wainer G and Giambiasi N. Cell-DEVS/GDEVs for complex continuous systems. *Simulation* 2005; 81: 137–151.
- Hodgkin AL and Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 1952; 117: 500–544.
- Izhikevich EM. Which model to use for cortical spiking neurons? *IEEE Trans Neural Networks* 2004; 15: 1063–1070.
- Lapicque L. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarization. *J Physiol Pathol Gen* 1907; 9: 620–635.
- Stein RB. Some models of neuronal variability. *Biophys J* 1967; 7: 37–68.
- Ermentrout B. Type I membranes, phase resetting curves, and synchrony. *Neural Comput* 1996; 8: 979–1001.
- E. M. Izhikevich. Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5): 1063–1070, September 2004.
- Izhikevich EM. Simple model of spiking neurons. *IEEE Trans Neural Networks* 2003; 14: 1569–1572.
- Izhikevich EM. *Dynamical systems in neuroscience: The geometry of excitability and bursting*. The MIT Press, 2007.
- Abbott TP and Vogels LF. Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci* 2005; 25: 10786–10795.

30. Stewart RD and Bair W. Spiking neural network simulation: Numerical integration with the parker-sochacki method. *J Comput Neurosci* 2009; 27: 115–133.
31. Zheng G, Tonnelier A and Martinez D. Voltage-stepping schemes for the simulation of spiking neural networks. *J Comput Neurosci* 2009; 26: 409–423.
32. Cellier F, Kofman E, Migoni G, and Bortolotto M. Quantized state system simulation. In: *Proceedings of Summer Sim 08 (2008 Summer Simulation Multiconference)*, Edinburgh, Scotland, 2008.
33. Kofman E. Relative error control in quantization based integration. *Lat Am Appl Res* 2009; 39: 231–238.
34. Bergero F and Kofman E. PowerDEVS. A tool for hybrid system modeling and real time simulation. *Simulation* 2011; 87: 113–132.
35. D'Abreu M and Wainer G. M/cd++: modeling continuous systems using modelica and DEVS. In: *Proceedings of MASCOTS 2005*, Atlanta, GA, 2005.
36. Quesnel G, Duboz R, Ramat E and Traoré M. Vle: a multimodeling and simulation environment. In: *Proceedings of the 2007 Summer Computer Simulation Conference*, San Diego, CA, 2007, pp.367–374.
37. Beltrame T. Design and development of a dymola/mod- elica library for discrete event-oriented systems using DEVS methodology. Master's thesis, ETH Zurich, Zurich, Switzerland, 2006.
38. Liu Q and Wainer G. Accelerating large-scale DEVS- based simulation on the cell processor. In: *Proceedings of the 2010 Spring Simulation Multiconference: Symposium on Theory of Modeling and Simulation. DEVS Integrative M&S Symposium (DEVS 2010)*, Orlando, FL, 2010.
39. Liu Q, Wainer L, Lu G and Perrone M. Novel perfor- mance optimization of large-scale discrete-event simula- tion on the cell broadband engine. In: *Proceedings of the 2010 International Conference on High Performance Computing & Simulation (HPCS 2010)*, Caen, France, 2010.
40. Hu X and Zeigler B. A high performance simulation engine for large-scale cellular DEVS models. In: *Proceedings of the 2004 High Performance Computing Symposium (HPC'04). Advanced Simulation Technologies Conference*, Arlington, VA, 2004.
41. Muzy A and Nutaro J. Algorithms for efficient imple- mentations of the DEVS & DSDEVS abstract simulators. In: *Proceedings of the 1st Open International Conference on Modeling & Simulation*, ISIMA/Blaise Pascal University, France, 2005, pp.401–407.

Guillermo L Grinblat currently holds a Fellowship at the French Argentine International Center for Information and Systems Sciences (CIFASIS–CONICET), Rosario, Argentina. He is also a teaching assistant at the National University of Rosario, Faculty of Engineering and Exact Sciences, Department of Computer Science, Rosario, Argentina.

Hernán Ahumada currently holds a Fellowship at the French Argentine International Center for Information and Systems Sciences (CIFASIS–CONICET), Rosario, Argentina.

Ernesto Kofman is an Adjunct Professor at the Department of Control of the National University of Rosario (FCEIA -UNR). He also holds a research position at the National Research Council of Argentina (CIFASIS –CONICET). His research interests include Automatic Control Theory and Numerical Simulation of Hybrid Systems.