



OPEN ACCESS

EDITED BY

Sharon Crook,
Arizona State University, United States

REVIEWED BY

Ben-Zheng Li,
University of Colorado Anschutz Medical
Campus, United States

Sen Lu,

The Pennsylvania State University (PSU),
United States

*CORRESPONDENCE

Hans Ekkehard Plesser
✉ hans.ekkehard.plesser@nmbu.no

RECEIVED 10 June 2024

ACCEPTED 30 September 2024

PUBLISHED 23 October 2024

CITATION

Plesser HE (2024) Commentary: Accelerating
spiking neural network simulations with
PymoNNto and PymoNNtorch.
Front. Neuroinform. 18:1446620.
doi: 10.3389/fninf.2024.1446620

COPYRIGHT

© 2024 Plesser. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Commentary: Accelerating spiking neural network simulations with PymoNNto and PymoNNtorch

Hans Ekkehard Plesser^{1,2,3*}

¹Department of Data Science, Faculty of Science and Technology, Norwegian University of Life Sciences, Ås, Norway, ²Institute for Advanced Simulation (IAS-6), Research Centre Jülich, Jülich, Germany, ³Käte Hamburger Kolleg: Cultures of Research (c:o/re), RWTH Aachen University, Aachen, Germany

KEYWORDS

spiking neural network (SNN), comparison, simulator, efficient implementation, exact integration

A Commentary on

[Accelerating spiking neural network simulations with PymoNNto and PymoNNtorch](#)

by Vieth, M., Rahimi, A., Gorgan Mohammadi, A., Triesch, J., and Ganjtabesh, M. (2024). *Front. Neuroinform.* 18:1331220. doi: 10.3389/fninf.2024.1331220

1 Introduction

In a recent paper introducing the PymoNNto and PymoNNtorch simulators, [Vieth et al. \(2024\)](#) compare their own simulators to several publicly available simulation codes using two different neuronal network models. They report that “the native NEST implementation consistently generates slightly fewer spikes in comparison with the other simulators” for their LIF neuron network model (see their [Supplementary Figure S2](#)) but do not investigate this issue beyond confirming that the discrepancy is not caused by the use of double-precision numerics in NEST.

This commentary shows that the observed difference in firing rate is due to an unsuitable numerical integration scheme used in the other simulators while NEST creates the correct model dynamics.

2 Results

[Vieth et al. \(2024\)](#) study a weakly coupled all-to-all network of K excitatory neurons driven by individual random input. The plasticity mechanism in the network is not relevant to the numerical issues investigated here and thus left out. Generalizing the equations of the original paper slightly, the membrane potential of an individual neuron k is given by

$$\dot{v}^{(k)}(t) = -\frac{v^{(k)}(t)}{\tau} + \frac{I_{\text{ext}}^{(k)}(t) + I_{\text{syn}}^{(k)}(t)}{C} \quad (1)$$

with membrane time constant τ and capacitance C . The membrane potential is reset to $v^{(k)} = 0$ when $v^{(k)} = \theta$ and there is no refractory time. The external input current is defined by

$$I_{\text{ext}}^{(k)}(t) = (\mathcal{U}[I_{\text{min}}, I_{\text{max}}] - I) \delta(t - jh) \quad (2)$$

is a piecewise constant current with a new amplitude chosen from a uniform distribution at fixed time intervals h . The recurrent synaptic input is

$$I_{\text{syn}}^{(k)}(t) = \sum_{l,i} w_{kl} \delta(t - t^{(l,i)} - d) \quad (3)$$

with uniformly distributed synaptic weights $w_{kl} \sim \mathcal{U}[0, 1/K]$ mV, the synaptic delay d , and $t^{(l,i)}$ representing the time of the i^{th} spike fired by neuron l .

For the parameters used by Vieth et al. (2024),¹ the mean input current is $\mu = \frac{1}{2}$ pA with variance $\sigma^2 = \frac{1}{12}$ pA², and the average synaptic weight $w = \frac{1}{2K}$ mV = $5 \cdot 10^{-5}$ mV. In this regime, spiking activity is mainly driven by external input with an average firing rate between 100 ms and 295 ms of 11.6 sp/s for the full model and 9.9 sp/s without any connections.

We consider first an individual neuron driven by external noise only. Dropping unnecessary subscripts and discretizing the dynamics on the fixed time grid $t_j = jh$ at which the input current takes on new random values I_j , Equation 1 becomes

$$v_{j+1} = \beta v_j + \alpha I_j. \quad (4)$$

The coefficients α and β depend on the numerical integration method chosen.

Equation 4 will provide the *exact* solution of Equations 1 and 2 if we use the exact integration method (Rotter and Diesmann, 1999) for which

$$\alpha = \frac{\tau}{C} (1 - e^{-h/\tau}) \approx 0.952 \quad (5)$$

$$\beta = e^{-h/\tau} \approx 0.905. \quad (6)$$

In the specific case of piecewise constant input over the integration step as in the given model, this method is equivalent to the exponential integration by MacGregor (1987, Ch. 14.C.5). We will use this solution as reference.

The coefficients for the forward Euler method are

$$\alpha = \frac{h}{C} = 1.0 \quad (7)$$

$$\beta = 1 - \frac{h}{\tau} = 0.9. \quad (8)$$

For first-passage time problems for the process defined by Equation 4 in combination with a spiking threshold, closed-form solutions exist for noise symmetric to the origin (Larralde, 2004), while only bounds are known for more general cases as studied here (Novikov and Kordzakhia, 2008).

¹ $K = 10^4$, $\tau = 10$ ms, $C = 1$ pF, $h = d = 1$ ms, $I_{\text{min}} = 0$ pA, $I_{\text{max}} = 1$ pA, $\theta = 6$ mV.

We thus consider the free membrane potential in the absence of a threshold. Assuming $v_0 = 0$, we obtain by repeated application of the update equation

$$v_{j+1} = \alpha \sum_{k=0}^j \beta^k I_k, \quad (9)$$

where we have renumbered the random currents I_k , exploiting their independence. Averaging over currents and taking the limit $j \rightarrow \infty$, we obtain the mean membrane potential and its standard deviation

$$\langle v \rangle = \frac{\alpha}{1 - \beta} \mu = \frac{\tau}{C} \mu = 5 \text{ mV} \quad (10)$$

$$\begin{aligned} \sqrt{\langle \Delta v^2 \rangle} &= \frac{\alpha}{\sqrt{1 - \beta^2}} \sigma^2 \\ &= \begin{cases} \frac{\tau \sigma}{C} \sqrt{\frac{1 - e^{-h/\tau}}{1 + e^{-h/\tau}}} \approx 0.645 \text{ mV} & \text{(exact integration)} \\ \frac{\tau \sigma}{C} \sqrt{\frac{h}{2\tau - h}} \approx 0.662 \text{ mV} & \text{(forward Euler)}. \end{cases} \end{aligned} \quad (11)$$

The broader distribution of the free membrane potential in the case of forward Euler integration (8.1% more probability mass above threshold θ) suggests that neurons will fire more frequently. This is supported by simulations below.

We confirm this observation quantitatively through numerical Markov analysis. This analysis is in principle continuous in time, but as the noise in the network model switches in intervals h , we consider only time points at which the noise changes. Given v_k , the smallest possible v_{k+1} is obtained for $I = I_{\text{min}}$ during the time step and the largest for $I = I_{\text{max}}$. In between the maximum and minimum possible values, any value of v_{k+1} is attained with equal probability. We can thus write the free transition probability for the membrane potential as

$$\hat{p}(v'|v) = \begin{cases} 0 & \text{if } v' < \alpha I_{\text{min}} + \beta v \Leftrightarrow v > (v' - \alpha I_{\text{min}})/\beta \\ \frac{1}{\tau(I_{\text{max}} - I_{\text{min}})} & \text{else} \\ 0 & \text{if } v' > \alpha I_{\text{max}} + \beta v \Leftrightarrow v < (v' - \alpha I_{\text{max}})/\beta \end{cases} \quad (12)$$

$$= \frac{1}{\tau(I_{\text{max}} - I_{\text{min}})} \left[\Theta \left(v - \frac{v' - \alpha I_{\text{min}}}{\beta} \right) - \Theta \left(v - \frac{v' - \alpha I_{\text{max}}}{\beta} \right) \right] \quad (13)$$

where $\Theta(x)$ is the Heaviside step function. To include the effect of the spiking threshold θ , we define

$$p(v'|v) = \begin{cases} \hat{p}(0|v) + \int_{\theta}^{\infty} \hat{p}(v'|v) dv' & v' = 0 \\ 0 & v' > \theta \\ \hat{p}(v'|v) & \text{else.} \end{cases} \quad (14)$$

Here, the integral in the first clause describes neurons re-inserted at the reset potential after spiking, and the corresponding probability is removed for superthreshold v' by the second clause.

If $q_k(v)$ is the membrane potential distribution at time step k , then the distribution at step $k + 1$ is given by

$$q_{k+1}(v') = \int_{-\infty}^{\infty} p(v'|v) q_k(v) dv. \quad (15)$$

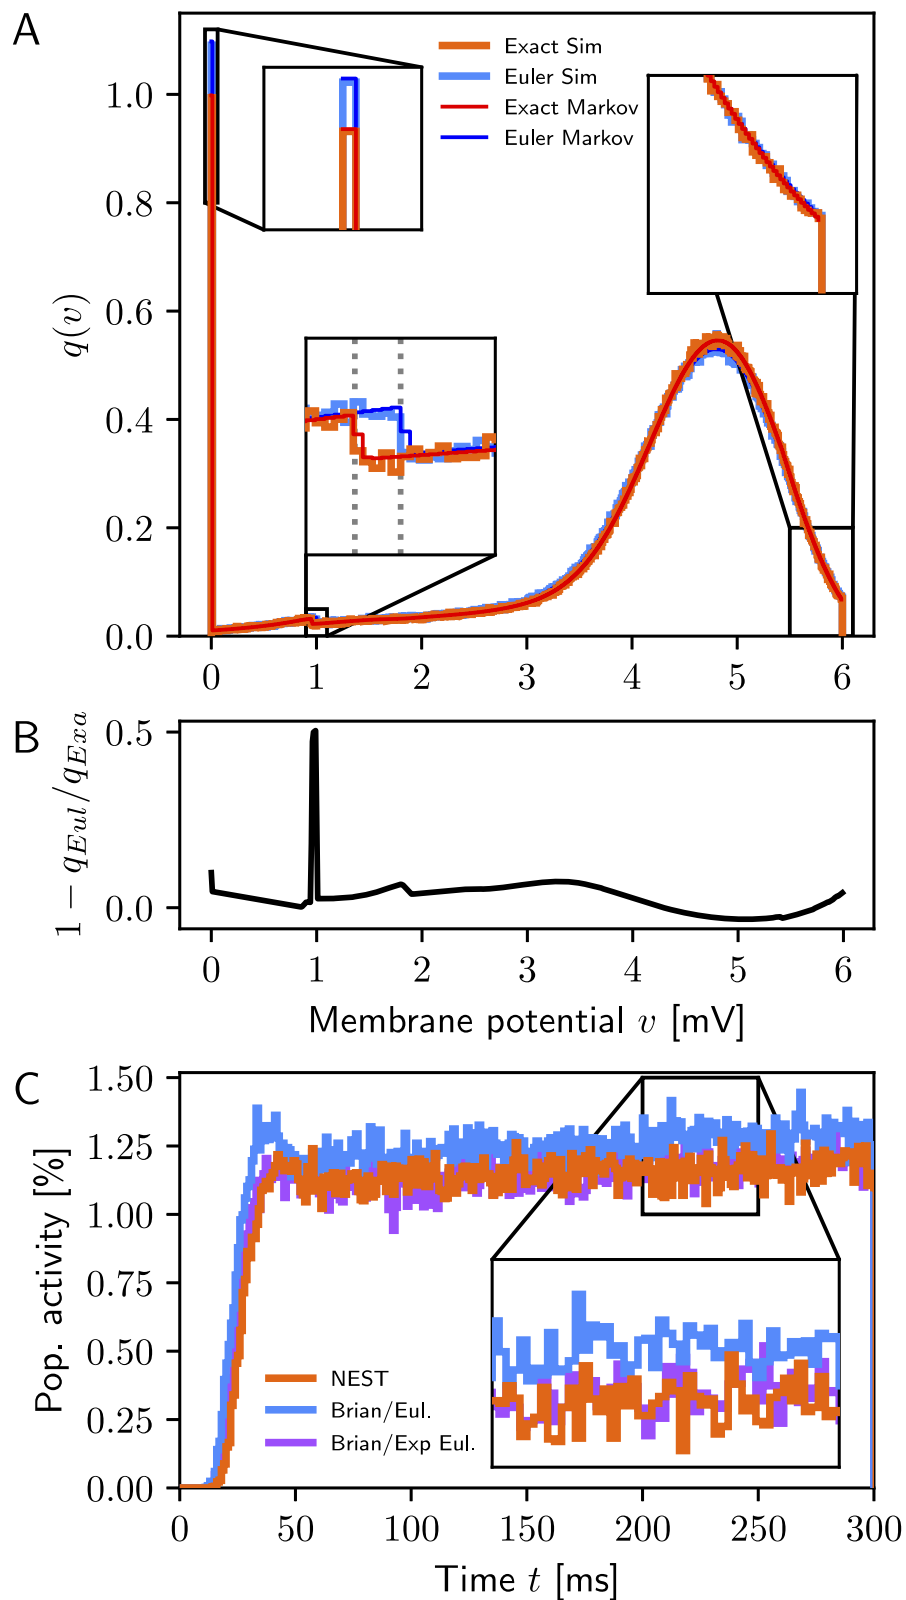


FIGURE 1

(A) Stationary membrane potential distributions $q^{(\infty)}(v)$ for exact integration (simulation: orange, Markov analysis: red) and forward Euler (simulation: light blue, Markov analysis: dark blue). Simulation data are for a single simulation of 1,000 s duration; bin width 0.01 mV. The dotted lines in the bottom-left inset mark $\alpha = 0.952$ for exact integration and $\alpha = 1.0$ for forward Euler, respectively, indicating the difference in the first update step after reset to $v = 0$ mV. (B) Relative difference in membrane potential distributions between exact integration and forward Euler. The peak near 1 mV is due to the difference in the first step after reset. (C) Population activity (1 ms bins) for simulations with NEST (orange) and Brian2 using delta synapses (blue: forward Euler, purple: exponential Euler); data from a single simulation with 10,000 neurons.

It will converge to the stationary membrane potential distribution $q^{(\infty)}(v)$ under reasonable assumptions about $p(v|v)$ (Lasota and Mackey, 1994). The firing probability for an interval h is given by the integral term in Equation 14, yielding the steady-state firing rate

$$r^{(\infty)} = h \int_{-\infty}^{\theta} [p(0|v) - \hat{p}(0|v)] q^{(\infty)}(v) dv. \quad (16)$$

As no closed-form solution is available for these equations, we discretize the transition probability as a Markov matrix M . This matrix has a single eigenvector with eigenvalue 1 which corresponds to $q^{(\infty)}(v)$ up to normalization (Feller, 1970; von Mises, 1964).

In the absence of recurrent connections, Figure 1A shows excellent agreement between simulations and Markov analysis (top) for exact integration and forward Euler, respectively, whereas Figure 1B displays clear differences in the membrane potential distributions between the two models. We obtain from the Markov analysis for exact integration $r^{(\infty)} = 9.93$ sp/s vs. 9.96 ± 0.04 sp/s from simulations and for forward Euler $r^{(\infty)} = 10.92$ sp/s vs. 10.93 ± 0.04 sp/s. This difference of approximately 1 sp/s between exact integration and forward Euler agrees closely with the disparity reported by Vieth et al. (2024, Supplementary Figure S2).

When recurrent connections are included and simulations performed using the original authors' code (PesarAmmehZA and Vieth, 2024, git hash c9c59f0), simulation results differ from results obtained with NEST even if using the exponential Euler method in Brian2. Analysis of their code (Brian_LIF.py) revealed that this neuron model used exponential post-synaptic currents instead of delta pulses. Once this was corrected, results obtained with Brian2 with exponential integration were consistent with NEST results as shown in Figure 1C.

Careful analysis of STDP in a network of 1,000 neurons revealed errors in both NEST and Brian2. NEST missed approximately 10% of weight increases, which could be traced to an error in NESTML code generation which has been fixed.² For Brian2, a small number of synapses experienced up to seven STDP events, while at most three are expected. Details are given in the Supplementary material.

3 Conclusion

Based on the exact solution (Rotter and Diesmann, 1999) to the model defined by Equations 1 and 2, the analysis above confirms that the discrepancy in firing rates between NEST and the other simulation codes observed by Vieth et al. (2024) is due to the use of the unsuitable forward Euler method in the latter codes: using this method to integrate the membrane potential evolution actually changes the model under study. Indeed, inserting α and β for exact integration into Equations 7 and 8 and solving for τ and C shows that the forward Euler method solved the model for $\tau \approx 10.5$ ms and $C \approx 1.05$ pF. Only NEST, the "odd one out," actually simulated the model as defined mathematically.

Interestingly, the investigation of the discrepancy observed by Vieth et al. (2024) helped us to uncover a subtle bug in

NESTML. This shows that careful comparison of results by different simulators provides mutual benefits.

The observed difference in firing rates of approximately 10% most likely will not have more than a 10% effect on the simulation runtimes measured by Vieth et al. (2024) and thus will not affect their overall conclusions. Nonetheless, I find it important to point out that benchmarks are meaningless, unless one defines a desired level of precision (Morrison et al., 2007) and confirms that the simulation codes compared produce statistically equivalent results (see, e.g., Van Albada et al., 2018).

Finally, I would like to question the choice of benchmark model as such. The LIF neuron network model defined by Vieth et al. (2024) is peculiar in that it is a purely excitatory all-to-all network with a plasticity rule that only allows synapses to become stronger. The model can therefore only be benchmarked over a limited simulation time before run-away excitation leads to excessive spike rates. Balanced networks such as the two-population model by Brunel (2000) can, in contrast, be simulated for arbitrary periods of time and usually in different dynamic regimes. Furthermore, comparability with other studies might benefit from using the cortical microcircuit model by Potjans and Diesmann (2014) as a benchmark case, as this model has become widely used as a reference model in recent years (Knight and Nowotny, 2018; Van Albada et al., 2018; Golosio et al., 2021; Kauth et al., 2023).

Author contributions

HEP: Investigation, Methodology, Software, Writing – original draft, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. The Käte Hamburger Kolleg: Cultures of Research (c:ore) was funded by the Federal Ministry of Education and Research under the funding code 01UK2104.

Acknowledgments

I am grateful to Markus Diesmann, Abigail Morrison, and Moritz Helias for discussions and useful comments.

Conflict of interest

HEP is a core developer of NEST and president of the non-profit NEST Initiative.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of

² <https://github.com/nest/nestml/issues/1057>

their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fninf.2024.1446620/full#supplementary-material>

References

- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* 8, 183–208. doi: 10.1023/A:1008925309027
- Feller, W. (1970). *An Introduction to Probability Theory and Its Applications*, volume 1. New York: John Wiley Sons.
- Golosio, B., Tiddia, G., De Luca, C., Pastorelli, E., Simula, F., and Paolucci, P. S. (2021). Fast simulations of highly-connected spiking cortical models using GPUs. *Front. Comput. Neurosci.* 15:627620. doi: 10.3389/fncom.2021.627620
- Kauth, K., Stadtmann, T., Sobhani, V., and Gemmeke, T. (2023). “neuroAIx: FPGA cluster for reproducible and accelerated neuroscience simulations of SNNs,” in *2023 IEEE Nordic Circuits and Systems Conference (NorCAS)* (Aalborg, Denmark: IEEE), 1–7. doi: 10.1109/NorCAS58970.2023.10305473
- Knight, J. C., and Nowotny, T. (2018). GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model. *Front. Neurosci.* 12:941. doi: 10.3389/fnins.2018.00941
- Larralde, H. (2004). A first passage time distribution for a discrete version of the Ornstein-Uhlenbeck process. *J. Phys. A Math. Gen.* 37, 3759–3767. doi: 10.1088/0305-4470/37/12/003
- Lasota, A., and Mackey, M. C. (1994). *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*. New York: Springer 2 edition. doi: 10.1007/978-1-4612-4286-4
- MacGregor, R. J. (1987). *Neural and Brain Modeling*. San Diego: Academic Press.
- Morrison, A., Straube, S., Plesser, H. E., and Diesmann, M. (2007). Exact subthreshold integration with continuous spike times in discrete time neural network simulations. *Neural Comput.* 19, 47–79. doi: 10.1162/neco.2007.19.1.47
- Novikov, A., and Kordzakhia, N. (2008). Martingales and first passage times of AR(1) sequences. *Stochastics* 80, 197–210. doi: 10.1080/17442500701840885
- PesarAmmehZA and Vieth, M. (2024). Fast-SNN-PymoNNto-rch. *GitHub*. Available at: <https://github.com/saeedark/Fast-SNN-PymoNNto-rch>
- Potjans, T. C., and Diesmann, M. (2014). The cell-type specific cortical microcircuit: Relating structure and activity in a full-scale spiking network model. *Cerebr. Cortex* 24, 785–806. doi: 10.1093/cercor/bhs358
- Rotter, S., and Diesmann, M. (1999). Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biol. Cybern.* 81, 381–402. doi: 10.1007/s004220050570
- Van Albada, S. J., Rowley, A. G., Senk, J., Hopkins, M., Schmidt, M., Stokes, A. B., et al. (2018). Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model. *Front. Neurosci.* 12:291. doi: 10.3389/fnins.2018.00291
- Vieth, M., Rahimi, A., Gorgan Mohammadi, A., Triesch, J., and Ganjtabesh, M. (2024). Accelerating spiking neural network simulations with PymoNNto and PymoNNtorch. *Front. Neuroinform.* 18:1331220. doi: 10.3389/fninf.2024.1331220
- von Mises, R. (1964). *Mathematical Theory of Probability and Statistics*. New York: Academic Press.