A Numerical Study of the Time of Extinction in a Class of Systems of Spiking Neurons

C. Romaro[1] · **F. A. Najman**[2] · **M. André**[2]

Abstract

In this paper we present a numerical study of a mathematical model of spiking neurons introduced by Ferrari et al. (J Stat Phys 172(6):1564–1575, 2018). In this model we have a countable number of neurons linked together in a network, each of them having a membrane potential taking value in the integers, and each of them spiking over time at a rate which depends on the membrane potential through some rate function ϕ. Beside being affected by a spike each neuron can also be affected by leaking. At each of these leak times, which occur for a given neuron at a fixed rate γ, the membrane potential of the neuron concerned is spontaneously reset to 0. A wide variety of versions of this model can be considered by choosing different graph structures for the network and different activation functions. It was rigorously shown that when the graph structure of the network is the one-dimensional lattice with a hard threshold for the activation function, this model presents a phase transition with respect to γ, and that it also presents a metastable behavior. By the latter we mean that in the sub-critical regime the re-normalized time of extinction converges to an exponential random variable of mean 1. It has also been proven that in the super-critical regime the renormalized time of extinction converges in probability to 1. Here, we investigate numerically a richer class of graph structures and activation functions. Namely we investigate the case of the two dimensional and the three dimensional lattices, as well as the case of a linear function and a sigmoid function for the activation function. We present numerical evidence that the result of metastability in the sub-critical regime holds for these graphs and activation functions as well as the convergence in probability to 1 in the super-critical regime.

Keywords Interacting point processes · Stochastic model of spiking neurons · Metastability · Numerical simulation

✉ M. André
    morgan.andre@ic.unicamp.br

1   Departamento de Física, Universidade de São Paulo, São Paulo, Brazil

2   Instituto de Matemática e Estatística, Universidade de São Paulo, São Paulo, Brazil

1 Introduction

This article presents and discusses a simulation study of a system of interacting point processes modeling a biologically inspired neural network; focusing on the time of extinction of a sustained activity exhibited by the network in the absence of external input. The use of point processes, like Hawkes processes [7, 8, 14, 15] or Wold processes [16], as models for biological neural networks has by now a pretty long history. The model considered here is closer to the one introduced by [3], since the process mimicking the neuron membrane potential is reset to a fixed value after each spike emitted by the neuron; it is informally speaking a stochastic version of the integrate and fire model. While network dynamics has been the subject of many studies both from a point process (*i.e.* stochastic) and from a deterministic [5, 21, 26] perspective, the concern for a *rigorous* characterization of the endogenous behavior (that is, in the absence of input) of the first kind of networks is, to the best of our knowledge, recent [1, 2, 10, 11, 13]. The latter studies using Galves-Löcherbach-like models [12].

The exact model we consider here is as in [11], which consists in an infinite number of neurons, each of them being affected by two different point processes representing the spiking times and the leaking times, the latter being an homogeneous Poisson process. This choice of a non-continuous potential loss is a clearly peculiar feature of our model and is introduced for technical convenience. It contrasts with the continuous leakage of the integrate and fire model (in the deterministic family) and of the original Galves-Löcherbach one (in the stochastic family).

Within the framework of our model, a series of rigorous results were established and are reviewed in Sect. 3 of this paper. Considering a one dimensional lattice together with a hard threshold activation function–the neuron's spiking rate switches from zero to one when the membrane potential becomes positive–[11] showed that the network exhibits a phase transition with respect to the leakage rate (Theorem 3.1 in Sect. 3); namely, there is a phase in which every neuron stops spiking in finite time with probability one, while in the other, each neuron continues spiking infinitely many times with positive probability. Considering a *finite* version of the same model, [2] showed that, in the continuously spiking phase of [11]– the sub-critical phase–, the re-normalized time of extinction is asymptotically (with respect to the number of neurons) exponentially distributed (Theorem 3.2 in Sect. 3), a property that we refer to as *metastability*. Informally, the network will present an unpredictably long period of pseudo-stationary activity before being kicked out to the only equilibrium (where all neurons are quiescent) by a sudden fluctuation. [1] also showed that in the other phase, the re-normalized time of extinction goes to one asymptotically (see Theorem 3.3 in Sect. 3). Informally, the network activity will exhibit a sharp membrane potential loss leading to a comparatively fast and predictable extinction time.

The notion of metastability has acquired more and more attention during the last years in the field of neuroscience (see for example [9] and [23]), especially to explain the brain ability to switch periodically between different functional networks. The most common approach consider large scale models to keep track of brainwaves. The approach we adopt here is to the best of our knowledge the first one (along with [2]) to give evidence of metastability in a theoretical model operating at the neuron level. It is inspired by statistical physics and takes its roots in [6], where a formal framework was developed for metastability in interacting particles systems, and where some models such as the Curie-Weiss mean field dynamics and the Harris contact process were proven to be metastable in the sense of this framework. The key point in this framework is to show that under suitable conditions (which here means in the sub-critical phase) the time of extinction of the system is asymptotically memory-less,

or in other words exponentially distributed (Theorem 3.2). The result proven in [2], as well as the results proven in [1] and [11], are nonetheless dependent on the technical assumptions that the structure of the network is a one-dimensional lattice, and that the activation function is a hard threshold, which is convenient mathematically but unrealistic biologically. Here, we develop a numerical study aimed to show that the results related to the time of extinction hold in lattices of higher dimensions $\mathbb{Z}^d$ ($d = 2$ and $d = 3$) and for two other commonly used activation functions.

The choice of such structure for the network was motivated by the fact that lattices of various dimensions has been previously used in the literature that is concerned with mathematical modeling of neural networks (see for example [20, 22] and [25]), and also by the fact that experimental works show that some cortical regions, such as specific regions of the primate visual cortex, present high similarity with multidimensional lattices (see [24] and [28]). We also deliberately chose non-random graphs in order to avoid a double source of stochasticity, which would uselessly complicate the analysis. For each of these multidimensional lattices we investigate the effect of changing the activation function to a linear function and to a sigmoid function. Both activation functions have been used in some form in stochastic models of biological neural networks (see for example [4] and [17]) as well as in artificial neural networks (see for example [19]).

This paper is organized as follows. In Sect. 2 we give a formal definition of the model. In Sect. 3 we give a formal presentation of the rigorous results that were proven in previous papers. In Sect. 4 we give a description of the algorithm used for the simulations. In Sect. 5 we give a precise characterization of the instantiations that we considered. In Sect. 6 we present the results of our simulations. Finally in Sect. 7 we discuss these results.

2 Definition of the Model

Informally the model we consider is as follows. I is a countable set representing the neurons, and to each $i \in I$ is associated a set $\mathbb{V}_i$ of *presynaptic neurons*. If you consider the elements of I as vertices, and draw and edge from j to i whenever $j \in \mathbb{V}_i$, then you obtain the graph structure of the network. The *membrane potential* of neuron i is a stochastic process denoted $(X_i(t))_{t \geq 0}$ taking value in the set $\mathbb{N}$ of non-negative integers. Moreover, we associate to each neuron a Poisson process $(N_i^\dagger(t))_{t \geq 0}$ of some parameter γ, representing the *leak times*. At any of these leak times the membrane potential of the neuron concerned is reset to 0. Another point process $(N_i(t))_{t \geq 0}$ representing the *spiking times* is also associated to each neuron, which rate at time t is given by $\phi(X_i(t))$, where ϕ is the rate function. Whenever a neuron spikes its membrane potential is also reset to 0 and the membrane potential of all of its post-synaptic neurons is increased by one (i.e. the neurons of the set $\{j : i \in \mathbb{V}_j\}$). All the point processes involved are assumed to be are mutually independent.

More formally, beside asking that $(N_i^\dagger(t))_{t \geq 0}$ be a Poisson process of some parameter γ, this is the same as saying that $(N_i(t))_{t \geq 0}$ is the point process characterized by the following equation

$$\mathbb{E}(N_i(t) - N_i(s)|\mathscr{F}_s) = \int_s^t \mathbb{E}(\phi(X_i(u))|\mathscr{F}_s)du$$

where

$$X_i(t) = \sum_{j \in \mathbb{V}_i} \int_{]L_i(t),t[} dN_j(s),$$

$L_i(t)$ being the time of the last event affecting neuron i before time t, that is,

$$L_i(t) = \sup\left\{s \leq t : N_i(\{s\}) = 1 \text{ or } N_i^\dagger(\{s\}) = 1\right\}.$$

$(\mathscr{F}_t)_{t \geq 0}$ is the standard filtration for the point processes involved here. See [11] for more details.

3 Rigorous Results

In [2, 11], and [1] a specific version of the model above was studied. The graph structure chosen there was the one-dimensional lattice, i.e. $I = \mathbb{Z}$ with $\mathbb{V}_i = \{i - 1, i + 1\}$, moreover the activation function was chosen to be a hard threshold of the form $\phi(x) = \mathbb{1}_{x>0}$. In such a paradigm the rate of the point processes representing spiking times can only take two values: 0 or 1, depending on whether the membrane potential is positive or null. In this context whenever the membrane potential of a neuron is different from 0 we say that the neuron is *active*, otherwise we say that it is *quiescent*. More generally we will only consider here function ϕ satisfying $\phi(0) = 0$ and $\phi(x) > 0$ for $x > 0$, so that we will keep this distinction between active and quiescent neurons.

In [11] it was proven that in the case of the one-dimensional lattice with hard threshold the following theorem holds

Theorem 3.1 *Suppose that for any $i \in \mathbb{Z}$ we have $X_i(0) \geq 1$. There exists a critical value γ_c for the parameter γ, with $0 < \gamma_c < \infty$, such that for any $i \in \mathbb{Z}$*

$$\mathbb{P}\left(N_i([0, \infty[) = \infty\right) > 0 \text{ if } \gamma < \gamma_c$$

and

$$\mathbb{P}\left(N_i([0, \infty[) < \infty\right) = 1 \text{ if } \gamma > \gamma_c.$$

In words there is a critical point for the parameter γ, such that below this critical point each neuron stays active forever (with positive probability), and above it each neuron becomes quiescent if you wait long enough. These two mutually exclusive type of behavior for the system are called the sub-critical regime (or phase) and the super-critical regime respectively.

The process as a whole of course never dies completely because of the fact that there is an infinite number of neurons, so that it doesn't makes sense to consider the extinction time. Nonetheless nothing prevents us to consider a finite version of this model. Suppose we're still in the case $\mathbb{V}_i = \{i - 1, i + 1\}$ and $\phi(x) = \mathbb{1}_{x>0}$ and for any $N \geq 0$ consider the system defined on the finite set $I_N = \{-N, \ldots N\}$ instead of the whole lattice. Then you can define the process $(\xi_N(t))_{t \geq 0}$ of the spiking rates of the system, that is, the process taking value in $\{0, 1\}^{I_N}$ defined by $\xi_N(t)_i = \mathbb{1}_{X_i(t)>0}$. This is a Markov process, belonging to the category of interacting particle systems, and by classical results on Markov processes we know that it will always reach the state 0^{I_N}–where all neurons are quiescent–in finite time. We can therefore consider the extinction time of this finite model, which we denote σ_N, and it is natural to ask about its distribution in each of the two phases distinguished by the theorem above. It was proven in [2] that the following holds.

Theorem 3.2 *There exists γ_c' such that if $\gamma < \gamma_c'$, then we have the following convergence*

$$\frac{\sigma_N}{\mathbb{E}(\sigma_N)} \xrightarrow[N \to \infty]{\mathscr{L}} \mathscr{E}(1).$$

In words, the re-normalized time of extinction converges in distribution to an exponential random variable of mean 1.

We know that $\gamma_c' < \gamma_c$, and the fact that the theorem is stated for some γ_c' and not for the critical value γ_c comes from essentially technical reasons related to the way the proof is built. We believe that this metastable result holds for the whole sub-critical region but it is not yet proven. Moreover, it was proven in [1] that we also have the following.

Theorem 3.3 *Suppose that $\gamma > 1$. Then the following convergence holds*

$$\frac{\sigma_N}{\mathbb{E}(\sigma_N)} \xrightarrow[N\to\infty]{\mathbb{P}} 1.$$

We know that $\gamma_c < 1$, so that the result concerns a sub-region of the super-critical region, and as for the previous result, while to the best of our knowledge the result has been proven only for $\gamma > 1$, it is reasonable to expect that it holds in the whole super-critical region.

As stated in the introduction the choice of a one-dimensional lattice for the graph of interaction and of a hard threshold for the activation function were initially essentially motivated by mathematical conveniency, and we're interested in checking that the results hold for a richer class of instantiations of the model. Namely, we are aimed to establish numerically that the results related to the asymptotical distribution of the extinction time stated for the one-dimensional case in Theorems 3.2 and 3.3 remain valid for lattices of dimension 2 and 3. We also investigate the effect of changing the activation function to a linear function and to a sigmoid function on these results.

4 Simulation Algorithm

All simulations were done in Python. The algorithm used to simulate our system of spiking neurons can be informally described as follows:

- **Initial configuration:** The network starts with all neurons actives (by default membrane potential equals 1), and for each neuron an initial spiking value is sampled from an exponential random variable of parameter 1 (corresponding to the first atom of $(N_i(t))_{t\geq0}$) and an initial leaking time value is sampled for each neuron from an another independent exponential random variable of parameter γ (corresponding to the first atom of $(N_i^\dagger(t))_{t\geq0}$).
- **Interaction:** The current time is set to be the smallest of the previously sampled values, the neuron i associated with this value is found and the value of its membrane potential is set to 0. In case the event considered is a spike the membrane potential of the neurons in the set $\{j : i \in \mathbb{V}_j\}$ (the set of post-synaptic neurons) is increased by one. The membrane potential of neuron i being equal to 0 we set the next spiking time for neuron i to be infinite until further notice. If the event was a leaking then we sample an exponential random variable of parameter γ and add it to the current time to get the next leaking time for neuron i. If the event was a spike then we sample an exponential random variable of parameter $\phi(X_j)$ for all post-synaptic neuron j and add this value to the current time to get the next spiking time for these neurons.
- **Stopping condition:** The previous operation is iterated until all neurons are quiescent.

More formally, our simulation algorithm can be described by the following pseudo-algorithm.

Algorithm 1 Simulate the system of spiking neurons and return the extinction time

---

1: I the (finite) set of neurons.
2: ϕ the activation function.
3: γ the rate of the leaking point processes.
4: t the current time.
5: $\mathbb{V}_i$ the set of presynaptic neurons for neuron i.
6: X_i the membrane potential of neuron i at the current time.
7: $\sigma_i^{\dagger}$ the time of the next leaking for neuron i.
8: σ_i the time of the next spike for neuron i.

   INITIALIZATION

9: $t \leftarrow 0$
10: **for** each i in I **do**
11:     $X_i \leftarrow 1$
12: **for** each i in I **do**
13:     $\sigma_i^{\dagger} \leftarrow \mathcal{E}(\gamma)$ $\triangleright \mathcal{E}$ denotes the realization of an exponential random variable
14:     $\sigma_i \leftarrow \mathcal{E}(\phi(X_i))$

   SIMULATION

15: **while** $\sum_{i \in I} X_i \neq 0$ **do**

16:     $MinLeaking \leftarrow \min_{i \in I} \sigma_i^{\dagger}$
17:     $MinSpiking \leftarrow \min_{i \in I} \sigma_i$

18:     **if** $MinLeaking < MinSpiking$ **then**
19:         $t = MinLeaking$
20:         $i \leftarrow \text{argmin}_{j \in I} \sigma_j^{\dagger}$
21:         $X_i \leftarrow 0$
22:         $\sigma_i \leftarrow \infty$
23:         $\sigma_i^{\dagger} \leftarrow t + \mathcal{E}(\gamma)$
24:     **else**
25:         $t = MinSpiking$
26:         $i \leftarrow \text{argmin}_{j \in I} \sigma_j$
27:         $X_i \leftarrow 0$
28:         $\sigma_i \leftarrow \infty$
29:         **for** each j such that $i \in \mathbb{V}_j$ **do**
30:             $X_j \leftarrow X_j + 1$
31:             $\sigma_j \leftarrow t + \mathcal{E}(\phi(X_j))$

32: $\sigma_N \leftarrow t$
33: **return** σ_N

All codes are available at https://github.com/ceciliaromaro/Finite_Spike.

5 Models Investigated

In this section we specify the structure of the graph of interaction and the activation function we're interested in.
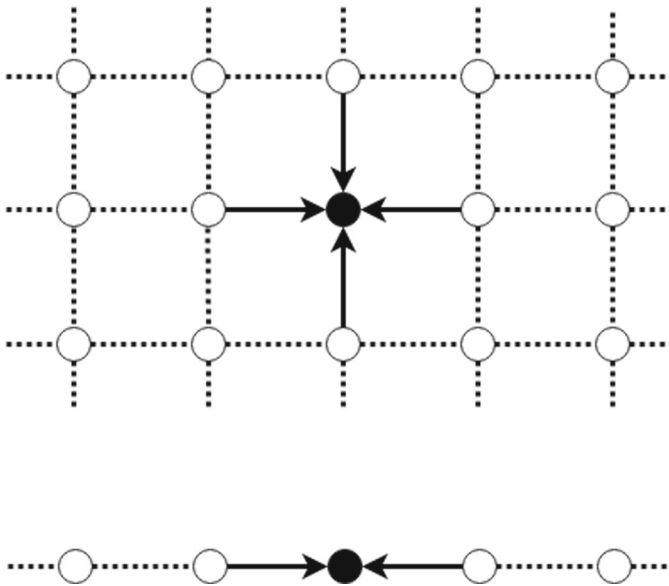
Fig. 1 One-dimensional and two-dimensional lattices. A directed arrow is drawn toward the black neuron from each of its presynaptic neurons

5.1 Multi-dimensional Lattices

For the graph of the network we consider the lattices $\mathbb{Z}^1$, $\mathbb{Z}^2$ and $\mathbb{Z}^3$. For any $d \in \{1, 2, 3\}$, let $\| \cdot \|$ be the norm on $\mathbb{Z}^d$ given for any $j \in \mathbb{Z}^d$ by

$$\|j\| = \sum_{k=1}^{d} |j_k|,$$

where j_k is the k-th coordinate of j. The structure of the network is then given by $I = \mathbb{Z}^d$ and $\mathbb{V}_i = \{j \in I^d : \|i - j\| = 1\}$ for $i \in I$.

Notice that by defining the set of presynaptic neurons as the set of the nearest neighbours we actually have $j \in \mathbb{V}_i$ if and only if $i \in \mathbb{V}_j$. In other words, for a given neuron the set of the presynaptic neurons and the set of the postsynaptic neurons are equal. For this specific choice the graph of interaction is therefore actually undirected.

5.2 Linear and Sigmoid Activation Functions

The activation function considered in [2, 11] and [1] was the hard threshold $\phi(x) = \mathbb{1}_{x>0}$. This choice is convenient mathematically as the system then becomes an additive interacting particle system where any neuron can only have two possible values for the spiking rate at any time: 0 or 1. Nonetheless, from a biological point of view, a hard threshold is a rough choice, and we would like to consider smoother options.

The first option we consider is a linear function of the simplest form: $\phi(x) = x$.

Table 1 Values of the total number of neurons and of the parameter γ used in the simulation for each of the three lattices

Lattice	Number of neurons	Value of γ	Figure
$\mathbb{Z}$	101	0.34	2
		0.85	3
$\mathbb{Z}^2$	121	1.25	2
		5.00	3
$\mathbb{Z}^3$	125	1.80	2
		6.00	3
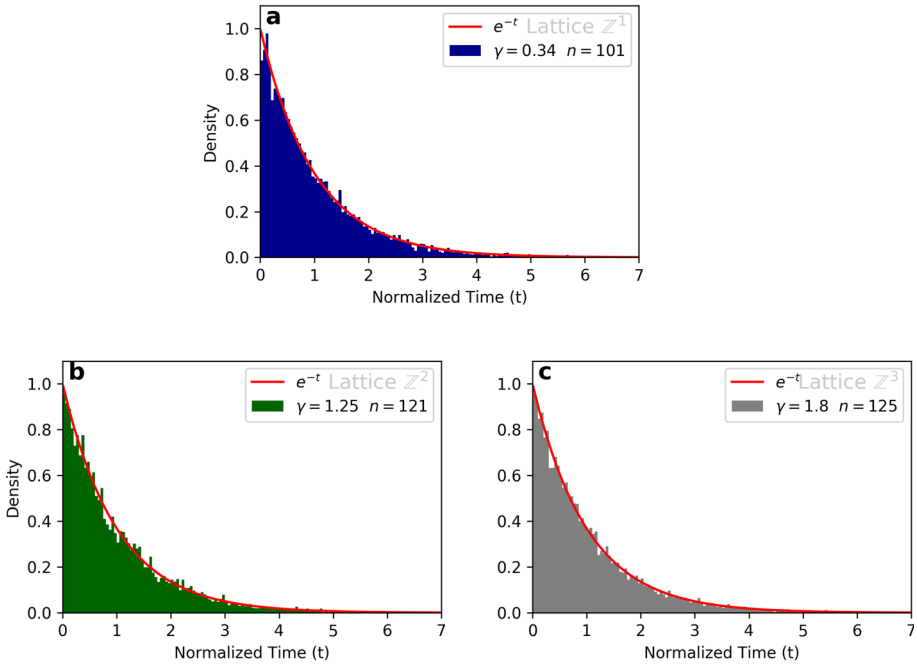
Fig. 2 Histogram of the re-normalized time of extinction σ_N for small values of gamma, and an activation function of the form $\phi(x) = \mathbb{1}_{x>0}$. In **a**, **b** and **c** the blue, green and gray bars are the histograms for the time of extinction in the one-dimensional lattice, two-dimensional lattice and three-dimensional lattice respectively. The red line is the exponential function $t \mapsto e^{-t}$, which corresponds to the density of an exponential law of parameter 1. The parameter n corresponds to the number of neurons (Color figure online)

The second option we consider is a somewhat more sophisticated sigmoid function of the following form

$$\phi(x) = \begin{cases} (1 + e^{-3x+6})^{-1} & \text{if } x > 0, \\ 0 & \text{if } x = 0. \end{cases}$$

Notice that we need to have $\phi(0) = 0$, in order to avoid spontaneous spiking (neuron with null membrane potential that spikes nonetheless). This is the reason why we force this value for the sigmoid function.
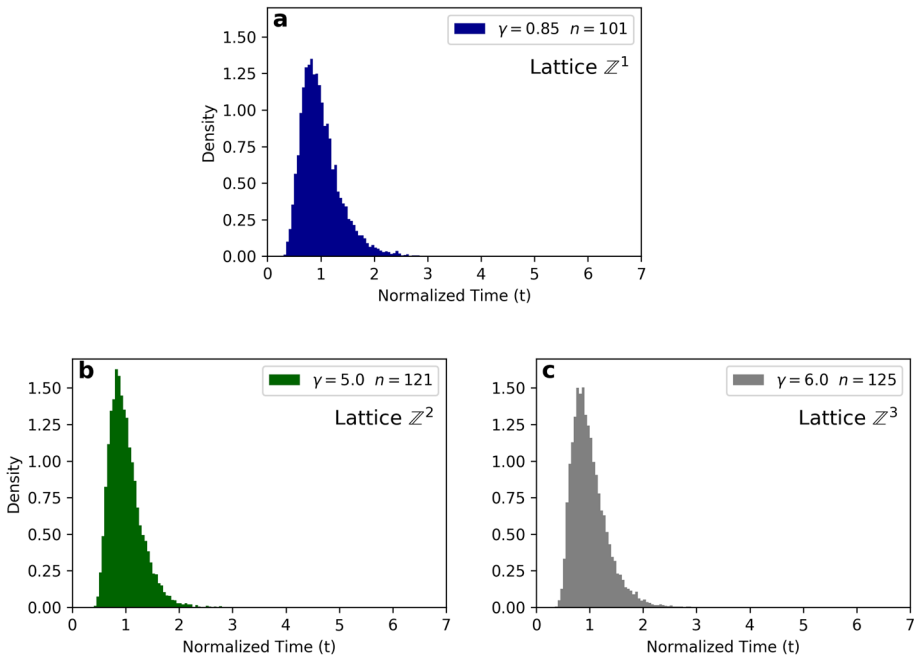
Fig. 3 Histogram of the re-normalized time of extinction σ_N for high values of gamma, and an activation function of the form $\phi(x) = \mathbb{1}_{x>0}$. In **a**, **b** and **c** the blue, green and gray bars are the histograms for the time of extinction in the one-dimensional lattice, two-dimensional lattice and three-dimensional lattice respectively. The parameter n corresponds to the number of neurons (Color figure online)

6 Results

We run simulations for instantiations of the system of spiking neurons consisting of all the possible combinations between the graphs and activation functions described above.

6.1 Simulations with a Fixed Number of Neurons

6.1.1 Multidimensional Lattices and Hard Threshold

For each of the three lattices the Algorithm 1 described in Sect. 4 was run, with an activation function of the form $\phi(x) = \mathbb{1}_{x>0}$. Each of the simulations were run for two different values of γ, 10,000 times for each of these values, using a number of neurons of the order of 100. The mean of the time of extinction σ_N was then computed using these data, and used to build the re-normalized histogram in each of these cases. The exact values for the size of the network and for the parameter γ can be found in Table 1.

The resulting histograms are presented in Figs. 2 and 3.

6.1.2 Multidimensional Lattices, Linear Function and Sigmoid Function

The routine described above was repeated with the two other activation functions. Only the values of γ change, which was necessary as changing the activation function must change
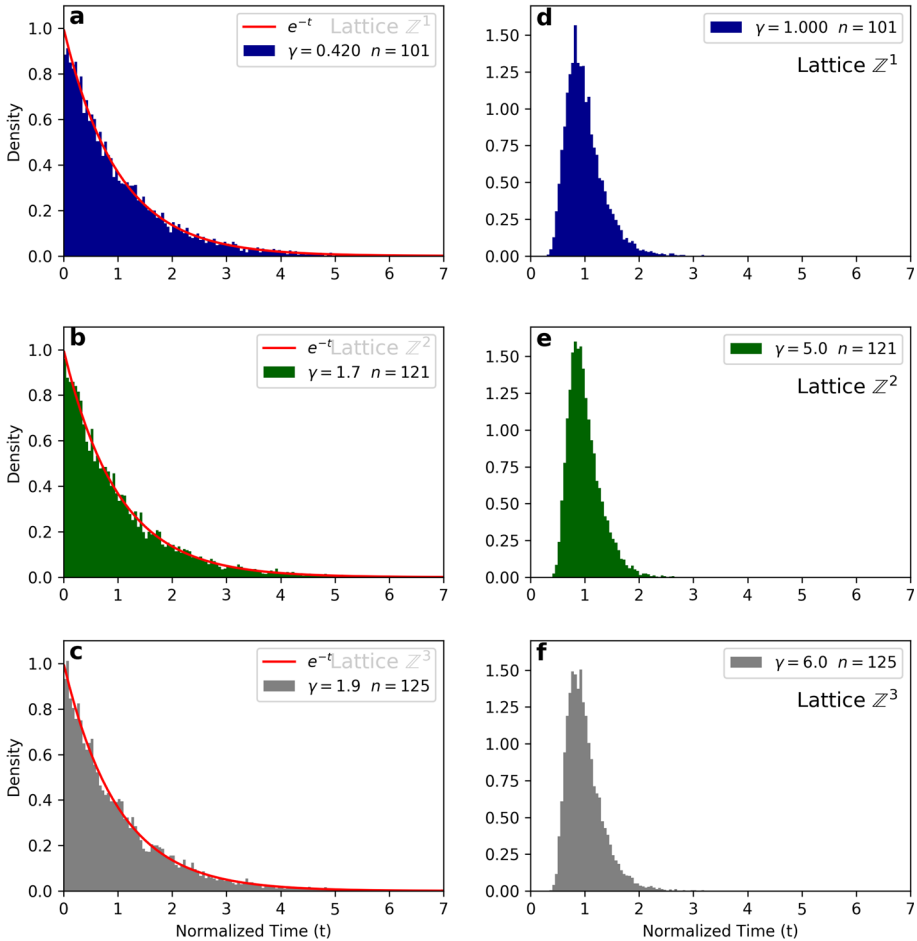
Fig. 4 Histogram of the re-normalized time of extinction σ_N for a linear activation function for each of the three lattices. On the left side are the histograms for small values of γ and on the right side the histograms for high values of γ. The red line on the left side is the exponential function $t \mapsto e^{-t}$, which corresponds to the density of an exponential law of parameter 1. The parameter n corresponds to the number of neurons (Color figure online)

the critical value of the system. These values are given in Table 2. The resulting histograms are presented in Figs. 4 and 5.

6.2 Simulations for a Varying Number of Neurons

To further investigate the behavior of the time of extinction in the super-critical regime, we've run a set of simulation for a (fixed) high value of γ and for a varying number of neurons. Each element of the set consists in 1000 repetitions with $\gamma = 4$. The value of the size of the network varies from 11 to 2000. For each of these values we've estimated the mean and variance of the extinction time σ_N, and the variance of the re-normalized extinction time
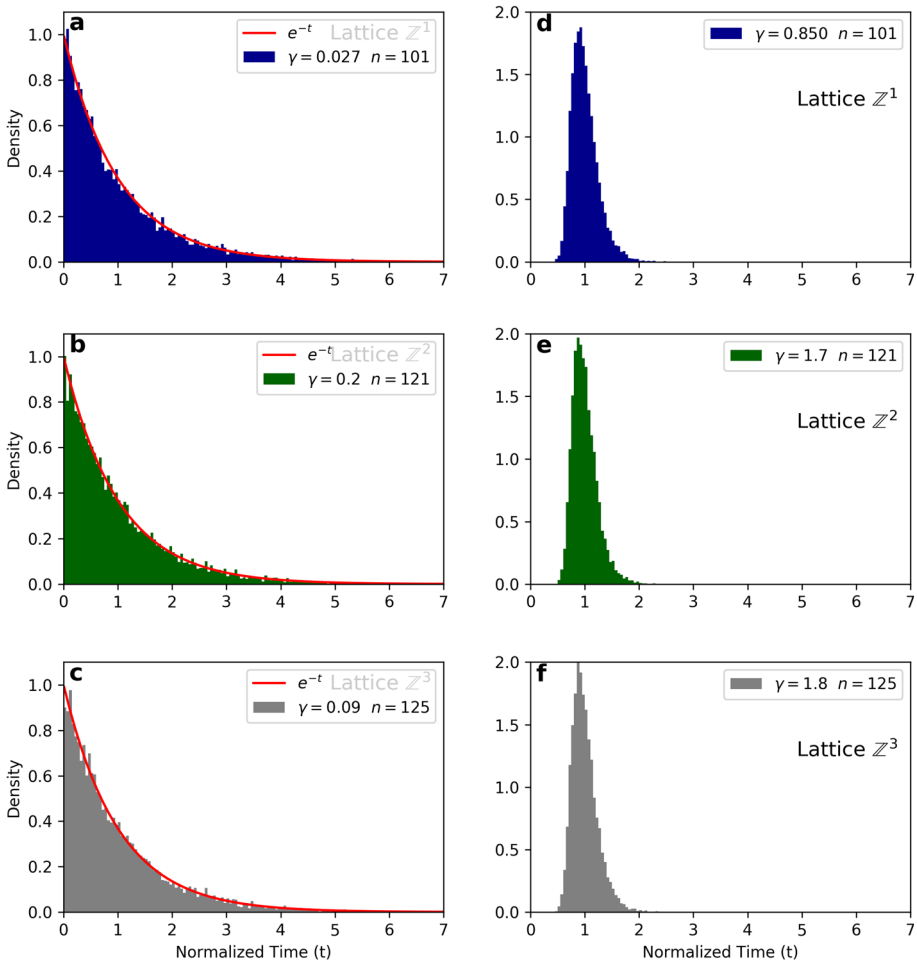
Fig. 5 Histogram of the re-normalized time of extinction σ_N for a sigmoid activation function for each of the three lattices. On the left side are the histograms for small values of γ and on the right side the histograms for high values of γ. The red line on the left side is the exponential function $t \mapsto e^{-t}$, which corresponds to the density of an exponential law of parameter 1. The parameter n corresponds to the number of neurons (Color figure online)

$\sigma_N / \mathbb{E}(\sigma_N)$. These simulations have been done in the one-dimensional lattices for all of the three activation functions The results of these simulations are presented in Figs. 6, 7 and 8.

7 Discussion

7.1 Sub-critical Regime

The histogram built from the simulations for which γ is small (Fig. 2 and left side of Figs. 4 and 5) closely approximates the density of a mean 1 exponential random variable. The fact that the result of the simulations remains identical in all the cases investigated (dimension

Table 2 Values of the total number of neurons and of the parameter γ used in the simulation for each of the three lattices

Lattice	Number of neurons	Activation function	Value of γ	Figure
$\mathbb{Z}$	101	Linear	0.42	4
		Sigmoid	0.028	5
		Linear	1	4
		Sigmoid	0.85	5
$\mathbb{Z}^2$	121	Linear	1.70	4
		Sigmoid	0.2	5
		Linear	5.00	4
		Sigmoid	1.7	5
$\mathbb{Z}^3$	125	Linear	1.90	4
		Sigmoid	0.09	5
		Linear	6.00	4
		Sigmoid	1.8	5
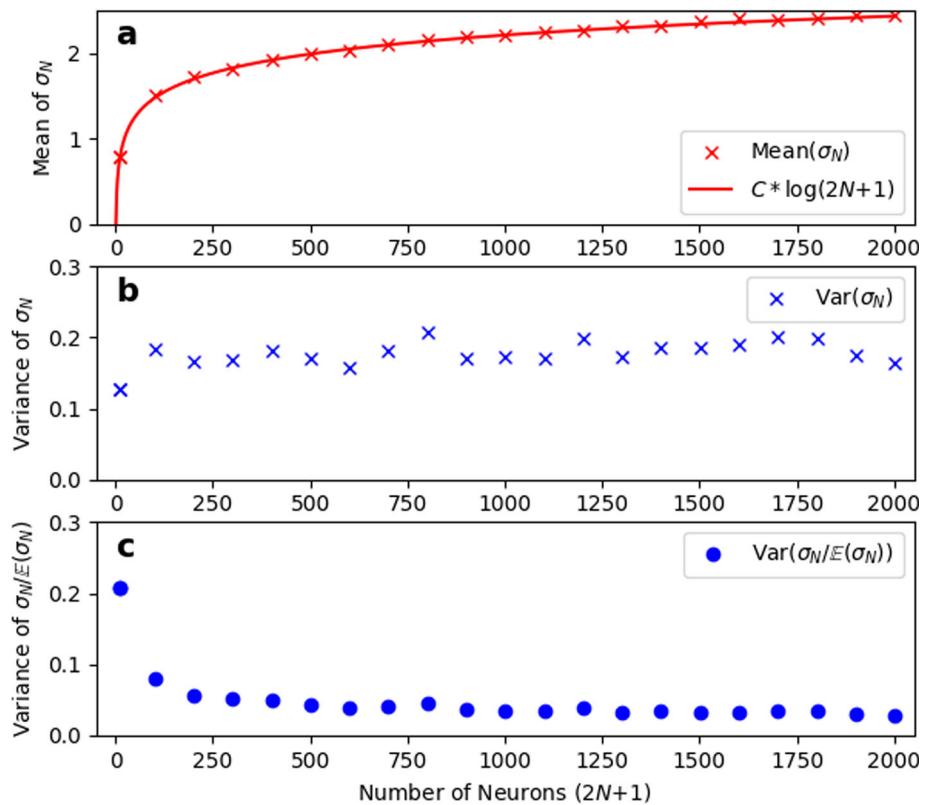
Fig. 6 Mean and variance of σ_N and variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a hard-threshold activation function. In **a** the red dots represents the values of the estimated mean of σ_N for varying numbers of neurons and the red line a logarithmic function fitted over the values of the mean ($C = 0.32$). In **b** the blue crosses represent the variance of σ_N as the number of neurons increases. The blue dots in **c** represent the variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a varying number of neurons. These simulations were run with $\gamma = 4$ (Color figure online)
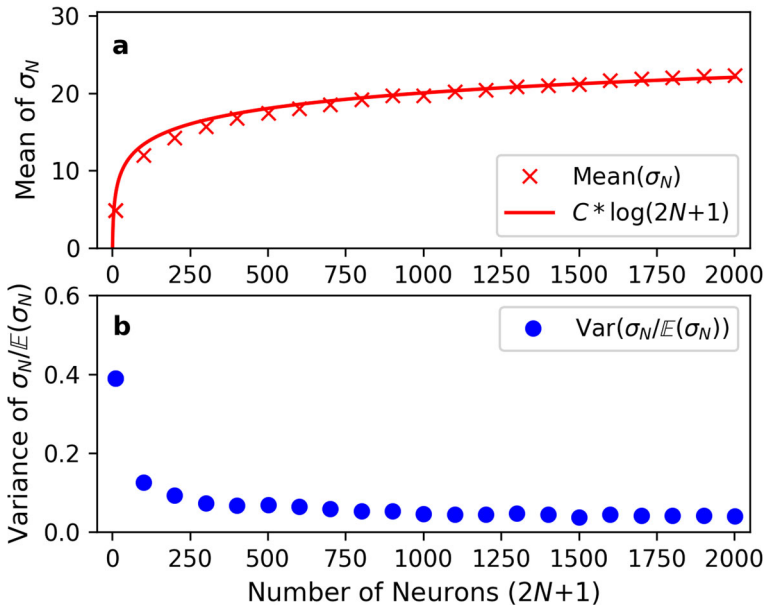
Fig. 7 Mean of σ_N and variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a linear activation function. In **a** the red dots represents the values of the estimated mean of σ_N for varying numbers of neurons and the red line a logarithmic function fitted over the values of the mean ($C = 2.9$). The blue dots in **b** represent the variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a varying number of neurons. These simulations were run with $\gamma = 1$ (Color figure online)

one, two and three, with hard-threshold, linear function and sigmoid function) suggests that Theorem 3.2 doesn't merely hold for the specific instantiation of the model for which is was proven, but for a wide class of systems.

In the one dimensional case with hard threshold the fact that the histogram approximates the density of an exponentially distributed random variable is of course not a surprise as this what Theorem 3.2 predicts asymptotically. Nonetheless it gives us evidences that the approximation by an exponential law holds for relatively small networks (in the simulation concerned the number of neuron in the system is 101). The number of neurons in animals varies from hundreds [27] to billions [18], so that in our model the approximation by an exponential law is observed for all possible biologically realistic number of neurons. This indicates that, in networks where the randomness of the connections between neurons is not of interest, the model presented here might be an interesting choice for the investigation of metastable behaviors in dense cortical regions.

7.2 Super-Critical Regime

The histograms built from the simulations with high values of γ are visibly not approximating any exponential. Instead they show a distribution that is reminiscent of a gamma distribution with mass concentrated around 1 (See Fig. 3 and right side of Figs. 4 and 5).

Moreover the evolution of the variance of the renormalized time of extinction (last graph in Figs. 6, 7 and 8) is seemingly converging toward 0 as the number of neurons grows for all of the three instantiations investigated.
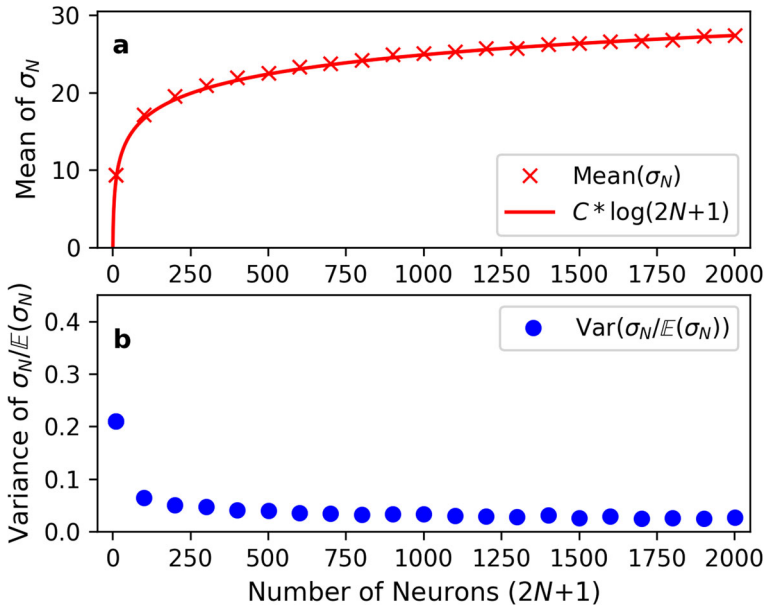
Fig. 8 Mean of σ_N and variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a sigmoid activation function. In **a** the red dots represents the values of the estimated mean of σ_N for varying numbers of neurons and the red line a logarithmic function fitted over the values of the mean ($C = 3.6$). The blue dots in **b** represent the variance of the renormalized extinction time $\sigma_N/\mathbb{E}(\sigma_N)$ for a varying number of neurons. These simulations were run with $\gamma = 0.34$ (Color figure online)

For the simulation of the system with hard-threshold, these facts are not a surprise neither, as this is what Theorem 3.3 predicts. Again the fact that the simulations of the systems with a linear and a sigmoid function for ϕ show similar results suggests that Theorem 3.3 isn't only satisfied for $\phi(x) = \mathbb{1}_{x>0}$, but for a wide class of activation functions.

Moreover the first graph in Figs. 6, 7 and 8 gives us strong evidence that in each of the three instantiations the expectation of the time of extinction grows approximately like a logarithm (up to a multiplicative constant) with respect to the number of neurons. This fact is interesting in itself as the proof of Theorem 3.3 (which can be found in [1]) relies on the fact that for the hard-threshold instantiation of the model we have the following convergence in the super-critical region (at least for $\gamma > 1$)

$$\frac{\mathbb{E}(\tau_N)}{\log(2N+1)} \xrightarrow[N\to\infty]{} C, \tag{1}$$

where C is a strictly positive (and finite) constant. This is an additional hint that the behavior of the time of extinction should be qualitatively identical in the super-critical region as well for any of the choices we proposed here for the activation function.

References

1. André, M., Planche, L.: The effect of graph connectivity on metastability in a stochastic system of spiking neurons. Stoch. Process. Appl. **131**, 292–310 (2021)
2. André, M.: A result of metastability for an infinite system of spiking neurons. J. Stat. Phys. **177**(5), 984–1008 (2019)
3. Brillinger, D.R.: Maximum likelihood analysis of spike trains of interacting nerve cells. Biol. Cybern. **59**(3), 189–200 (1988)
4. Brochini, L., de Andrade Costa, A., Abadi, M., Roque, A.C., Stolfi, J., Kinouchi, O.: Phase transitions and self-organized criticality in networks of stochastic spiking neurons. Sci. Rep. **6**, 35831 (2016)
5. Brunel, N.: Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J. Comput. Neurosci. **8**(3), 183–208 (2000)
6. Cassandro, M., Galves, A., Olivieri, E., Vares, M.E.: Metastable behavior of stochastic dynamics: a pathwise approach. J. Stat. Phys. **35**, 603–634 (1984)
7. Chevallier, J., Caceres, M.J., Doumic, M., Reynaud-Bouret, P.: Microscopic approach of a time elapsed neural model. Math. Models Methods Appl. Sci. **25**(14), 2669–2719 (2015)
8. Chornoboy, E.S., Schramm, L.P., Karr, A.F.: Maximum likelihood identification of neural point process systems. Biol. Cybern. **59**, 265–275 (1988)
9. Deco, G., Kringelbach, M.L., Jirsa, V.K., Ritter, P.: The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. Sci. Rep. (2017). https://doi.org/10.1038/s41598-017-03073-5
10. Duarte, A., Ost, G.: A model for neural activity in the absence of external stimuli. Markov Processes Relat. Fields **22**(1), 35–52 (2016)
11. Ferrari, P., Galves, A., Grigorescu, I., Löcherbach, E.: Phase transition for infinite systems of spiking neurons. J. Stat. Phys. **172**(6), 1564–1575 (2018)
12. Galves, A., Löcherbach, E.: Infinite systems of interacting chains with memory of variable length—a stochastic model for biological neural nets. J. Stat. Phys. **151**(5), 896–921 (2013)
13. Galves, A., Löcherbach, E., Pouzat, C.: A system of interacting neurons with short term synaptic facilitation. J. Stat. Phys. (2019). https://doi.org/10.1007/s10955-019-02467-1
14. Hansen, N.R., Reynaud-Bouret, P., Rivoirard, V.: Lasso and probabilistic inequalities for multivariate point processes. Bernoulli **21**(1), 83–143 (2015)
15. Hawkes, A.G.: Point spectra of some mutually exciting point processes. J. R. Stat. Soc. **33**(3), 438–443 (1971)
16. Kass, R.E., Ventura, V.: A spike-train probability model. Neural Comput. **13**(8), 1713–1720 (2001)
17. Kinouchi, O., Brochini, L., Costa, A.A., Campos, J.G.F., Copelli, M.: Stochastic oscillations and dragon king avalanches in self-organized quasi-critical systems. Sci. Rep. **9**(1), 3874 (2019)
18. Lent, R., Azevedo, F.A., Andrade-Moraes, C.H., Pinto, A.V.: How many neurons do you have? Some dogmas of quantitative neuroscience under revision. Eur. J. Neurosci. **35**(1), 1–9 (2012)
19. Maass, W.: Networks of spiking neurons: the third generation of neural network models. Neural Netw. **10**(9), 1659–1671 (1997)
20. Makarenkov, V.I., Kirillov, A.: Self-organized criticality in neural networks. In: Applications of Artificial Neural Networks II (1991), vol. 1469, International Society for Optics and Photonics, pp. 843–845
21. Millman, D., Mihalas, S., Kirkwood, A., Niebur, E.: Self-organized criticality occurs in non-conservative neuronal networks during up states. Nat. Phys. **6**(10), 801 (2010)
22. Miranda, E., Herrmann, H.: Self-organized criticality with disorder and frustration. Physica A **175**(3), 339–344 (1991)
23. Roberts, J.A., Gollo, L.L., Abeysuriya, R.G., Roberts, G., Mitchell, P.B., Woolrich, M.W., Breakspear, M.: Metastable brain waves. J. Nat. Commun. **10**, 1–16 (2019)
24. Rockland, K.S., Lund, J.S.: Intrinsic laminar lattice connections in primate visual cortex. J. Compar. Neurol. **216**(3), 303–318 (1983)
25. Sporns, O., Kötter, R.: Motifs in brain networks. PLoS Biol. **2**(11), e369 (2004)
26. Vreeswijk, C.V., Sompolinsky, H.: Chaotic balanced state in a model of cortical circuits. Neural Comput. **10**(6), 1321–1371 (1998)
27. White, J.G., Southgate, E., Thomson, J.N., Brenner, S.: The structure of the nervous system of the nematode caenorhabditis elegans. Philos. Trans. R. Soc. Lond. B Biol. Sci. **314**(1165), 1–340 (1986)
28. Yoshioka, T., Levitt, J.B., Lund, J.S.: Intrinsic lattice connections of macaque monkey visual cortical area v4. J. Neurosci. **12**(7), 2785–2802 (1992)