

A Quasi-Stationary Approach to Metastability in a System of Spiking Neurons with Synaptic Plasticity

Morgan André* Christophe Pouzat†

Abstract

After reviewing the behavioral studies of working memory and of the cellular substrate of the latter, we argue that metastable states constitute candidates for the type of transient information storage required by working memory. We then present a simple neural network model made of stochastic units whose synapses exhibit short-term facilitation. The Markov process dynamics of this model was specifically designed to be analytically tractable, simple to simulate numerically and to exhibit a quasi-stationary distribution (QSD). Since the state space is finite this QSD is also a Yaglom limit, which allows us to bridge the gap between quasi-stationarity and metastability by considering the relative orders of magnitude of the relaxation and absorption times. We present first analytical results: characterization of the absorbing region of the Markov process, irreducibility outside this absorbing region and consequently existence and uniqueness of a QSD. We then apply Perron-Frobenius spectral analysis to obtain any specific QSD, and design an approximate method for the first moments of this QSD when the exact method is intractable. Finally we use these methods to study the relaxation time toward the QSD and establish numerically the memorylessness of the time of extinction.

Keywords: Working memory ; Markov processes, Stochastic simulation; Mathematical Neuroscience.

MSC2020 subject classifications: Primary 60K35, Secondary 37H30 ; 60J28 ; 92C20 ; 92C42.
Submitted to MNA on July 12, 2021, final version accepted on November 15, 2024.
Supersedes [HAL:hal-03281732](https://hal.archives-ouvertes.fr/hal-03281732).

1 Introduction

1.1 The neurobiological horizon

As described by Fuster in the introduction of his 1973 article [18]: *A delayed-response trial typically consists of the presentation of one of two possible visual cues, an ensuing period of enforced delay and, at the end of it, a choice of motor response in accord with the cue. The temporal separation between cue and response is the principal element making the delayed response procedure a test of an operationally defined short-term memory function.* In that article Fuster described, in the monkey prefrontal

*Instituto de Matemática e Estatística - USP, University of São Paulo, São Paulo Brasil. E-mail: morgan.andre@usp.br

†IRMA, Strasbourg University and CNRS UMR 7501, Strasbourg, France. E-mail: christophe.pouzat@math.unistra.fr

cortex, neurons that switch between no activity and a *sustained activity at constant rate during the delay period* when the animal had to perform a delayed-response task. He showed moreover—using distracting stimuli that interrupted the sustained activity of these neurons—that the monkey errors at the end of the delay period were positively correlated with the interruption of sustained activity. Since then, many experimental investigations reviewed in [19, Chap. VII] and [11] have confirmed this basic finding and showed that some of the “sustained activity neurons” are insensitive to the type of cue (color, shape, location, sound) and seem to encode the “abstract” notion of remembering “any” cue until the expiration of a delay, while others, especially outside the prefrontal cortex, are sensitive to the type of cue. These “sustained activity neurons” are relatively easy to record from, *implying that they are fairly abundant* [28]. This sustained activity has been intriguing modelers for a long time, leading them to explore first network models with subgroups of strongly reciprocally coupled excitatory neurons [44, 1]. The sustained activity has then been interpreted as a local attractor of some dynamical system (reviewed in [41]). Stability issues when the transiently memorized item is a continuous quantity—like an angle—, led modelers to include some “slow” and “use dependent” coupling, initially in the form of NMDA receptors [10, 41]—for a review of basic neurophysiology, see [30]. But [42] described a subclass of pyramidal (and therefore excitatory) cells in the prefrontal cortex that are strongly interconnected and whose synapses are unusual, since they exhibit a marked short-term facilitation—synapses between neurons of this type usually exhibit short-term depression. This has led to several studies giving a more or less central role to short-term facilitation in sustained activity generation or stabilization, *e.g.* [4, 24, 21]—or even proposing a working memory mechanism without sustained activity [34]—, reviewed in [5]. But the secondary status of the “noise” in these studies, where variability comes into play mostly at the neurons input level, is at odd with basic empirical observations. It is indeed well known [13, 43, 30] that neurons depend on ion channels that are randomly going back and forth between closed and opened states both for the action potential generation [40] and the synaptic transmission [25]; that (chemical) synaptic transmission involves the release of a variable number of transmitter packets / quanta [16, 14], giving rise to the rather noisy membrane voltage trajectories that are actually observed.

These considerations strongly suggest an alternative model construction strategy: working with stochastic units / neurons instead of deterministic ones. Continuing and simplifying [20], we therefore developed a minimal model of the sub-network of reciprocally coupled pyramidal cells with facilitating synapses [42]; a model that is both amenable to analytical solutions and that can be easily simulated. This model is made of stochastic neurons that accumulate their inputs until a threshold is reached. The synapses exhibit short-term facilitation, represented by a binary random variable, enabling the sub-network to exhibit a transient *memoryless* sustained activity—that is, to reach a quasi-stationary distribution (QSD) leading to genuine metastability—, reminiscent of what is observed in working memory experiments.

As pointed out by one of the anonymous referees, models that explain this kind of persistent activity by introducing a stochastic binary facilitation state have already been considered in the literature. One of the main references is [35], in which the persistent activity is quantified via mean-field theory. Nonetheless, both our model and the analytic approach we develop are quite different from what was proposed in the latter. Indeed, while the facilitation state is the only stochastic element in [35]—the spikes being produced in a deterministic fashion—in the model we propose the spike generation itself is inherently stochastic. More importantly, our analysis doesn’t require mean-field approximation: QSD theory allows an exact characterization of the metastable phase at the mesoscopic scale, which is our main contribution. We nonetheless also

propose a mean-field analysis of our system in a second time, since the QSD approach is currently intractable at the macroscopic scale.

It shall be noticed that the model we propose does not include any leak currents, as it is common in biological neural networks models. This is done on purpose, since the metastable properties of leaky stochastic models has already been investigated in several articles (see [2, 3, 37, 31, 27]), and one of our intentions is to show that leakage is not the only possible biophysical mechanism by which metastability might be implemented in the brain.

1.2 Metastability and Quasi-stationarity

The expressions “metastability” and “quasi-stationary distribution” used in the previous section have proper mathematical meanings, which deserve explanations. The notion of metastability on one hand has a long history, and appears in a wide variety of fields to describe various apparently unrelated phenomena, from nuclear physics to the study of avalanches and super-cooling water. More recently it has gained popularity in the field of neuroscience, as in many respects the brain seems to exhibit metastable-like behaviors. In the present article the notion of metastability is to be understood as in the rigorous characterization introduced in [8] for interacting particle systems, which requires that, for a Markov process evolving in a space which includes an absorbing state: (i) the time to reach the absorbing state from any other state is memoryless (i.e. follows an exponential distribution), asymptotically with respect to the number of particles in the system; (ii) before reaching this absorbing state (the only real equilibrium), the system exhibits *thermalization*, i.e. an apparent stabilization, temporary but long, in a region of the state space away from the actual equilibrium. The notion of quasi-stationary distribution on the other hand refers to the stationary distribution of a modified Markov process which has been conditioned to stay away from its absorbing state (see [23, 33] for an introduction). Since here we consider a system which is essentially irreducible and evolves in a finite state space, the QSD is unique and corresponds to the Yaglom limit, that is, the unique limit distribution to which the system conditioned on non-absorption relaxes, starting from any given state of the irreducibility region. Moreover it is well-known that, in this case, if the initial state is distributed with respect to the QSD then the time of extinction is indeed memoryless. In the following we argue that the thermalization referred to above (that is, the metastable phase) can be understood through the theory of quasi-stationarity—a possibility which, quite surprisingly, seems to have been mostly ignored in the literature, with the notable exceptions of [6] and [22], in specific settings which are not applicable in our case. This approach gives us analytical tools to quantitatively characterize the metastable phase, as well as a quite straightforward (but not yet rigorous) way of establishing the memorylessness evoked in point (i).

1.3 Overview

The present paper is organized as follows. In Section 2 we give the basic definition of our model: first from the modeling viewpoint; then in a somewhat more mathematical form, as a continuous-time Markov chain. In Section 3 we partition the state space into relevant sub-regions; in particular we give an explicit characterization of the absorbing region, and of the support of the (yet to be proven) metastable phase; irreducibility is also established. In Section 4, we introduce briefly the relevant elements of the theory of quasi-stationarity distributions; the key role of the Perron-Frobenius spectral analysis is emphasized. In Section 5, these ideas are implemented for a simple example with 5 neurons, which allows us to compare numerically the relative orders of magnitude

of the relaxation toward the QSD and then toward the absorbing region, establishing the memorylessness of the time of extinction. Since this exact method becomes rapidly intractable, we then study numerically the extinction time in Section 6. An heuristic calculation that provides an approximate alternative to the spectral approach is presented in Section 7.

2 Model definition

We start by defining informally a stochastic system of interacting spiking neurons—a formal definition will be given in the next section. The system consists in a finite set of N neurons. A *membrane potential process*, denoted $(U_t(i))_{t \geq 0}$, taking values in the set of non-negative integers is associated to every neuron $i \in \{1, \dots, N\}$. The spiking activity of the neurons depends on a threshold value $\theta \in \mathbb{Z}^+$. When $U_t(i) < \theta$ neuron i cannot spike, and we say that it is *quiescent*, while when $U_t(i) \geq \theta$ we say that neuron i is *active*: it spikes at rate β —i.e. it waits a time Δ_t distributed as an exponential random variable of parameter β and spikes at time $t + \Delta_t$. Every neuron is connected to every other neuron of the network with a uniform *synaptic strength* and the effect of a spike depends on the state of the synapse between the spiking neuron and the other neurons of the network at that time. At any time the synapse can be either *facilitated* or *not facilitated*, meaning that if a spike occurs, it will or will not be "transmitted" to the other neurons. The facilitation state of the synapse of neuron i is denoted $(F_t(i))_{t \geq 0}$; it is a stochastic process taking value in $\{0, 1\}$. Whenever $F_t(i) = 1$ we say that the synapse of neuron i is *facilitated* at time t . The facilitated synapse loses its facilitation (i.e. goes back to the unfacilitated state $F_{t+\Delta_t}(i) = 0$) at a given rate λ . Now the fact that a given neuron j spikes at time t means the following: (i) its membrane potential is *reset to 0*, $U_{t+}(j) = 0$; (ii) its synapse becomes facilitated if it was not already, $F_{t+}(j) = 1$; (iii) if its synapse was facilitated at the time of the spike, the spike is said to be *efficient* and the membrane potential process of *all the other neurons* in the system increases by one unit, while nothing happens to these neurons if the synapse was not facilitated at the moment of the spike, $U_{t+}(i) = U_{t-}(i) + F_{t-}(j)$, for all $i \neq j$; in the latter case the spike is said to be *inefficient*. All exponential random variables involved are assumed to be independent. Then a trajectory of the system is entirely characterized by the family of interacting stochastic processes $(U_t(i), F_t(i))_{t \geq 0}$ (for $i \in \{1, \dots, N\}$).

2.1 Formal definition as a continuous-time Markov chain

It is clear that although the membrane potential of any neuron i can get arbitrarily large in the above model formulation, from the dynamics viewpoint all that matters is to know whether $U_t(i) < \theta$ or not. We can therefore consider that there are only $\theta + 1$ relevant membrane potential states: $U_t(i) = 0, U_t(i) = 1, \dots, U_t(i) = \theta - 1$ and $U_t(i) \geq \theta$. The effective state space of a neuron is therefore finite with $2(\theta + 1)$ elements:

$$\Phi \equiv \{0, \dots, \theta\} \times \{0, 1\},$$

where the set $\{0, 1\}$ corresponds to the synaptic facilitation. Figure 1 illustrates the effective states accessible to an arbitrary neuron $i \in \{1, \dots, N\}$ of the network as well as the possible transitions among those states. The state of any neuron of the network can therefore be fully specified by placing a token on one of the spaces (circles/squares, by analogy to the board of the game of the goose) of Fig. 1 and the state of the whole network can be represented by placing N tokens labelled 1 to N on the $2(\theta + 1)$ spaces. The state of the network is then specified by selecting a single element of the set: $\chi = \Phi^N$. We adopt the following notation: for any $x \in \chi$ we write $x = (x_1, \dots, x_N)$, with $x_i = (x_i^U, x_i^F)$ for any $i \in \{1, \dots, N\}$, where $x_i^U \in \{0, \dots, \theta\}$ corresponds to the value of the

membrane potential for neuron i while x_i^F corresponds to the facilitation state. Notice that the state space size increases very quickly with N and θ :

- for $N = 5$ and $\theta = 1$, the size is, $4^5 = 1024$,
- for $N = 50$ and $\theta = 10$, it becomes roughly $2.2^{43} \mathcal{N}_A$, where \mathcal{N}_A is Avogadro's constant.




Figure 1: An illustration where $\theta = 5$. Each node represents a possible membrane potential value, the first element of the pair (being understood that "5" should be interpreted as " ≥ 5 ") and a synaptic facilitation value in the second element of the pair. If the neuron has its synapse facilitated it sits in the outer circle (light blue circle), if its synapse is not facilitated it sits in the inner circle (light orange squares). The transition rates are encoded by the color: brown is the rate of "effective spikes" generated by the network (spike occurring while the synapse of the spiking neuron is still facilitated), dark green is λ and red is β .

Formally we're considering a continuous time Markov process $(X_t)_{t \geq 0}$ taking value in the finite state space \mathcal{X} . Its dynamic can be characterized by an infinitesimal generator $Q = (q_{x,y})_{x,y \in \mathcal{X}}$, which is defined by the following requisites. Let $x, y \in \mathcal{X}$ such that $x \neq y$ and let $i \in \{1, \dots, N\}$, then¹:

- **Loss of facilitation** : Suppose y is such that $y_j = x_j$ for any $j \neq i$ and $y_i = (x_i^U, 0)$ then $q_{x,y} = \lambda \delta_1(x_i^F)$.
- **Inefficient spike** : Suppose y is such that $y_j = x_j$ for any $j \neq i$ and $y_i = (0, 1)$, then $q_{x,y} = \beta \delta_{(\theta,0)}(x_i)$.
- **Efficient spike** : Suppose y is such that $y_j = (\min(\theta, x_j^U + 1), x_j^F)$ for any $j \neq i$ and $y_i = (0, 1)$, then $q_{x,y} = \beta \delta_{(\theta,1)}(x_i)$.

Moreover for any $y \neq x$ not considered above $q_{x,y} = 0$ and of course $q_{x,x} = -\sum_{y \neq x} q_{x,y}$. The key features of this model dynamics are illustrated next with simulations.

¹Here and in the rest of this article δ denotes the Kronecker delta function.

2.2 Basic dynamics features

2.2.1 Process trajectories

Simulation details are provided in Se. A.1. Figure 2 shows the trajectories of the “membrane potentials”, $u_t(i)$ a realization of $U_t(i)$, and the synaptic facilitation, $f_t(i)$ a realization of $F_t(i)$, of all the (50) neurons of a simulation ($N = 50$, $\theta = 5$, $\beta = 10$, $\lambda = 6.7$) during one time unit.




Figure 2: Trajectories between time units 1 and 2 of the membrane potentials of the fifty neurons of a simulated network. The traces are blue when the synapse of the neuron is facilitated and orange otherwise.

This figure displays the complete state of the network. Notice that at any given time, most of the neurons are in the susceptible state (their membrane potential is $\geq \theta$). Notice also that the membrane potentials of the neurons that have not yet reached θ evolve in parallel. Spikes are emitted when the membrane potential of one neuron goes from θ to zero, this is the only way the membrane potential can decrease. A finer time display is proposed on Fig. 3.

The features of the model are clearly visible:

- When a neuron with an un-facilitated synapse spikes (the trace is orange when the membrane potential is at the threshold level just before dropping to 0):
 - its synapse gets facilitated (the trace turns blue) immediately after the spike,
 - the membrane potential of all the other neurons remains the same.
- When a neuron with a facilitated synapse spikes (the trace is blue when the membrane potential is at the threshold level just before dropping to 0):
 - its synapse remains facilitated (the trace stays blue) immediately after the spike,
 - the membrane potential of *all* the other neurons that are below threshold increases by 1.




Figure 3: Enlarge display between times 1.20 and 1.25 of the data shown of Fig. 2.

2.2.2 It is the same but it is not the same

We now turn to the key property our model was designed to exhibit. The next two figures (4 and 5) show spike trains displayed as raster plots (every spike is represented by a dot) of the same network of 50 neurons started from the same initial state but using two different sequences of (pseudo) random numbers. Fig. 4 shows an abrupt disappearance of the activity (after 13 time units).



Figure 4: Raster plots of a 50 neurons network, with $\lambda = 6.7$, $\beta = 10$ and $\theta = 5$. Left, from time 0 to 14; right from time 12 to 14. Dots are blue when the synapse is facilitated and orange otherwise

The dots color is blue when the synapse is facilitated and orange otherwise. We see on the right side of Fig. 4 that the last spikes occurring before the “network death” are all with an un-facilitated synapse. Fig. 5 shows the same network as Fig. 4, starting from the same state and remaining active for whole simulation (50 time units).

Judging from the dots pattern, the activity looks regular with a constant ratio of blue dots over orange dots. But a better way to graphically assess the network activity (network spiking frequency) is provided by the observed counting process (a step function that



Figure 5: Same as Fig. 4 but different random numbers sequence. The scale bar is drawn between time 10 and time 15.

increases by one every time an event occurs) as shown on Fig. 6 for the two simulations of Fig. 4 and 5. Extracting the slope by eye, we see that the network generates roughly




Figure 6: Observed counting processes for the simulations of Fig. 4 (black) and Fig. 5 (red).

375 events per time unit (before it reaches the quiescent state in the case of the black trace). *We have, qualitatively at least, the behavior we are interested in: the activity seems “stationary” until it abruptly vanishes.*

2.3 The aggregated process

Notice now that the dynamics of our system is invariant with respect to the permutation of the neurons: they are all equivalent members of the network. We are moreover interested in the network state, as opposed to the individual neuron states. For these reasons we might focus on the number of neurons in each possible state at each time. In

order to do this, for any $x \in \mathcal{X}$ and any $(i, j) \in \{0, \dots, \theta\} \times \{0, 1\}$ we define

$$z^{(i,j)}(x) = \sum_{k=1}^N \delta_i(x_k^U) \delta_j(x_k^F).$$

Then $z^{(i,j)}$ is simply a function from \mathcal{X} to $\{0, \dots, N\}$ counting the number of neurons having membrane potential i and facilitation state j . For any $(i, j) \in \{0, \dots, \theta\} \times \{0, 1\}$ and any time $t \geq 0$ we might now define the stochastic version of these counting functions:

$$Z_t(i, j) = z^{(i,j)}(X_t).$$

Furthermore, for any $t \geq 0$ we write $Z_t = (Z_t(i, j))_{(i,j) \in \{0, \dots, \theta\} \times \{0, 1\}}$, and the resulting process $(Z_t)_{t \geq 0}$ is a continuous-time Markov chain taking value in the set $S = \{0, \dots, N\}^{\{0, \dots, \theta\} \times \{0, 1\}}$. For any $z \in S$ we adopt the following notation: $z_{i,j}$ denotes the value in $\{0, \dots, N\}$ corresponding to the coordinates $i \in \{0, \dots, \theta\}$ and $j \in \{0, 1\}$, while z_i denotes the ordered pair $(z_{i,0}, z_{i,1})$. Moreover for any $i \in \{0, \dots, \theta\}$ we write

$$z_{i,\bullet} = z_{i,0} + z_{i,1}.$$

In the following, if a state z is such that $z_{i,\bullet} = k$ we adopt the terminology of saying that there are k neurons at level i . Furthermore notice that because the number of neurons in the system is fixed to N we may narrow down the state space a bit; the range of the process (the actual state space) is $R \subset S$, defined as:

$$R = \left\{ z \in S : \sum_{i=0}^{\theta} z_{i,\bullet} = N \right\}.$$

The size of this state space is much smaller than the one we started from, since instead of having $(2(\theta + 1))^N$ it has “only” [17, Sec. II.5, p. 38]: $\binom{N+2\theta+1}{2\theta+1}$ elements. Compared to the previous considered cases we get:

- for $N = 5$ and $\theta = 1$, the size is 56,
- for $N = 50$ and $\theta = 10$, it becomes roughly 5.5×10^{17} .

It is easy to obtain the dynamics of the process $(Z_t)_{t \geq 0}$ from the definition of $(X_t)_{t \geq 0}$. We write this dynamic explicitly in term of maps from R to R corresponding to the three possible events susceptible to affect the system. The map corresponding to efficient and inefficient spikes will be denoted respectively π and π^* , and the map corresponding to a loss of facilitation on a neuron at level i will be denoted π_i^\dagger . Suppose the current state is $z \in R$.

- **Loss of facilitation:** When it happens at level $i \in \{0, \dots, \theta\}$ the number of un-facilitated neurons at level i increases by 1, and the number of facilitated neurons decreases by 1. The map π_i^\dagger is therefore defined by:

$$\left(\pi_i^\dagger(z) \right)_j = \begin{cases} (z_{i,0} + 1, z_{i,1} - 1) & \text{if } j = i, \\ (z_{j,0}, z_{j,1}) & \text{otherwise.} \end{cases}$$

- **Inefficient spike:** This event leads to a decrease of the non-facilitated neurons at level θ by one and an increase of the number of facilitated neurons at level 0 by one (red arrow from square (5, 0) to circle (0, 1) in Fig. 1). The map π^* is therefore defined by:

$$\left(\pi^*(z) \right)_j = \begin{cases} (z_{\theta,0} - 1, z_{\theta,1}) & \text{if } j = \theta, \\ z_j & \text{if } j \in \{1, \dots, \theta - 1\}, \\ (z_{\theta,0}, z_{\theta,1} + 1) & \text{if } j = 0. \end{cases}$$

- Efficient spike:** This event leads to a decrease of the number of facilitated neurons at level θ by one and an increase by $z_{\theta-1,1}$ (when the spike comes from a neuron with a facilitated synapse, all neurons get their membrane potential increased by one except the one that spiked, whose membrane potential drops to 0), while the number of non-facilitated neurons at level θ increases by $z_{\theta-1,0}$. The number of facilitated neurons at level 0 is set to 1, and the number of non-facilitated neurons at level 0 is set to 0. Moreover for $i \in \{1, \dots, \theta - 1\}$, the number of facilitated (resp. non-facilitated) neurons at level i is set to $z_{i-1,1}$ (resp. $z_{i-1,0}$). On Fig. 1, the contents of all the circles and squares rotate by one step counter clockwise, except for element (0, 0) that becomes 0, element (0, 1) that becomes 1 and elements (5, 0) that adds the content of element (4, 0) to its own and element (5, 1) that also adds the content of (4, 1) to its own and decrease by 1. The map π is defined by:

$$\left(\pi(z)\right)_j = \begin{cases} (z_{\theta,0} + z_{\theta-1,0}, z_{\theta,1} - 1 + z_{\theta-1,1}) & \text{if } j = \theta, \\ (z_{j-1,0}, z_{j-1,1}) & \text{if } j \in \{1, \dots, \theta - 1\}, \\ (0, 1) & \text{if } j = 0. \end{cases}$$

Then the infinitesimal dynamic of the process $(Z_t)_{t \geq 0}$ is given by the following: starting from some state $z \in R$, inefficient spikes occur at rate $\beta z_{\theta,0}$, efficient spikes occurs at rate $\beta z_{\theta,1}$, and losses of facilitation at level $i \in \{0, \dots, \theta\}$ occur at rate $\lambda z_{i,1}$. A slightly less formal but perhaps more intuitive way of describing the above dynamics follows, writing z the network state at time t in a matrix form (we show here the transpose in order to save space):

$$z^T \equiv \begin{bmatrix} z_{0,0} & z_{1,0} & \dots & z_{i,0} & \dots & z_{\theta,0} \\ z_{0,1} & z_{1,1} & \dots & z_{i,1} & \dots & z_{\theta,1} \end{bmatrix}$$

A *loss of facilitation* can occur leading to (changes appear in red and the occurrence rate appears above the arrow):

$$z^T \xrightarrow{\lambda z_{i,1}} \begin{bmatrix} z_{0,0} & z_{1,0} & \dots & z_{i,0} + 1 & \dots & z_{\theta,0} \\ z_{0,1} & z_{1,1} & \dots & z_{i,1} - 1 & \dots & z_{\theta,1} \end{bmatrix}$$

Alternatively, an *inefficient spike* will lead to:

$$z^T \xrightarrow{\beta z_{\theta,0}} \begin{bmatrix} z_{0,0} & z_{1,0} & \dots & z_{i,0} & \dots & z_{\theta,0} - 1 \\ z_{0,1} + 1 & z_{1,1} & \dots & z_{i,1} & \dots & z_{\theta,1} \end{bmatrix}$$

While an *efficient spike* will give:

$$z^T \xrightarrow{\beta z_{\theta,1}} \begin{bmatrix} 0 & z_{0,0} & \dots & z_{\theta,0} + z_{\theta-1,0} \\ 1 & z_{0,1} & \dots & z_{\theta,1} + z_{\theta-1,1} - 1 \end{bmatrix}$$

In the following this will be expressed alternatively as an infinitesimal generator Q in matrix form, or an infinitesimal generator \mathcal{L} in functional form (in Section 7), depending on what is more suitable for the current purpose.

3 Cutting the state space into pieces

Later we will be interested in studying the quasi-stationary distribution of our system, that is the stationary distribution associated with the Markov process obtained when conditioning $(Z_t)_{t \geq 0}$ on non-absorption. In order to ensure that this actually makes sense we have two concerns: (i) we would like to properly define the absorbing region; (ii) prove that the process restricted to the complement of this absorbing region is irreducible. Once this is established, the existence and uniqueness of the quasi-stationary distributions follow from classical results (see [12]).

3.1 Absorbing region

First notice that there is an obvious absorbing set of states for $(Z_t)_{t \geq 0}$, which are the states $D \subset R$ defined by $D = \{z \in R : z_{\theta, \bullet} = 0 \text{ and } \sum_{i=0}^{\theta-1} z_{i,1} = 0\}$, that is the states with no neuron at level θ and no facilitated neurons. Then neither a spike, nor a facilitation loss can happen and the process stays there for eternity.

Nonetheless we would be short-sighted if we stopped there, and took this set D to be the absorbing region. It is clear that one can find other states which, while not being properly absorbing, can only lead to D in a bounded number of steps with probability one. We write $A_\theta = \{z \in R : z_{\theta,1} = 0\}$. Notice that $D \subset A_\theta$, moreover if $(Z_t)_{t \geq 0}$ reaches some $a' \in A_\theta$ at some point, then after some inefficient spikes and after the facilitated neurons lose their facilitation one after the other (which is the only thing that can happen) the process hits D , and again stays there for eternity. This happens in a finite number of transitions—at most $2N$ of them actually, which corresponds to the case in which $a'_{\theta,0} = N$. More generally, for any $i \in \{1, \dots, \theta\}$ define:

$$A_i = \left\{ z \in R : \sum_{j=i}^{\theta} z_{j,1} \leq \theta - i \right\}.$$

Notice that, for $i = \theta$, this agrees with the previous definition. Furthermore it is easy to see that if at some point the process $(Z_t)_{t \geq 0}$ reaches A_i for some $i \in \{1, \dots, \theta\}$, then with probability one it then reaches A_θ after a maximum of $\theta - i$ efficient spikes, and again reaches D after some inefficient spikes and facilitation losses.

For the case $i = 0$ we need to be a little bit more careful, since when non-facilitated neurons at level θ spike the number of facilitated neurons at level 0 increases, which in turn could prevent the process from being absorbed. Therefore, we define

$$A_0 = \left\{ z \in R : z_{\theta,0} + \sum_{j=0}^{\theta} z_{j,1} \leq \theta \right\}.$$

Again it is easy to see that, if $(Z_t)_{t \geq 0}$ reaches A_0 at some point it then hits A_θ in a maximum of θ spikes and then gets absorbed. Finally we define the absorbing region to be:

$$A = \bigcup_{i=0}^{\theta} A_i.$$

We claim that this is the correct definition of the absorbing region, meaning that:

1. once A has been reached the system cannot get out of it, and eventually gets definitively absorbed in D , which happens with probability one in a bounded number of steps²,
2. and if the system is in A^c then it has a positive probability of staying out of A for an arbitrarily big number of steps.

The first point is simply a reformulation of the discussion above. The second point is a consequence of the important fact that, whenever the process starts outside A , there is always a positive probability of having an arbitrarily big number of (efficient) spikes in any time interval of arbitrary length, with no facilitation loss in the meantime. We will see in the next section that we can actually go further than that, and show that once the state space has been restricted a little bit, then when starting in A^c any other state from A^c can be reached.

²For example one can easily see from the considerations above that the number of transition to reach D from any point in A has to be less than $\theta + 2N$, even if it is certainly not the sharpest possible bound.

3.2 Irreducibility

One can easily see that, even when A has been taken out of the state space R , the process $(Z_t)_{t \geq 0}$ is not properly irreducible. Consider for example the class $R' \subset R$ of states defined by

$$R' = \{z \in R : z_{i,\bullet} = 0 \text{ for some } i \in \{0, \dots, \theta - 1\}\}.$$

Let $z \in R' \setminus A$ and suppose that $(Z_t)_{t \geq 0}$ starts from z . First suppose that $z_{0,\bullet} = 0$. Notice that if no spike occurs in the future, then all neurons lose their facilitation and soon the system reaches A . Otherwise, after the first spike the system reaches some state z' which is such that $z'_{0,\bullet} \geq 1$, and then this inequality remains true for eternity, as the only way for the number of neurons at level 0 to decrease is if there is an efficient spike, moving neurons from level 0 to level 1, but in this case the number of neurons at level 0 is immediately reset to 1 (because the membrane potential of the spiking neuron is reset to 0). Now if we suppose that $z_{0,\bullet} \geq 1$ but $z_{1,\bullet} = 0$, then at the instant of the first efficient spike the system reaches a state z' such that $z'_{1,\bullet} \geq 1$, and again, the only way the number of neurons at level 1 can then be affected is when there is an efficient spike, which push the value at level 0 to level 1, so that it can never be less than 1 (because we already know that at level 0, there cannot be 0 neuron anymore). If both $z_{0,\bullet}$ and $z_{1,\bullet}$ are equal to 0, then it takes at maximum two spikes (the second one being efficient) to obtain the same result. More generally it is easy to see how the same argument applies recursively to prove the following statement.

Lemma 3.1. *let $z \in R' \setminus A$ and let $(Z_t^z)_{t \geq 0}$ be the process starting from z . Then either $(Z_t^z)_{t \geq 0}$ is absorbed in A before reaching $R' \setminus R'$, or it reaches $R' \setminus R'$ and then never hits R' again. In the second case $R' \setminus R'$ is reached after a maximum of $\theta - 1$ efficient spikes.*

While a recent article [9] has treated the existence of a quasi-stationary measure for discrete-time Markov chains in such a case, that is when the state space is not properly irreducible but consists in two successive classes, to the best of our knowledge results are still lacking in the continuous-time framework. Nonetheless, as R' is always left after a maximum of only $\theta - 1$ efficient spikes, it might simply be discarded from the state space. Indeed what the lemma above is essentially saying is that states in R' are atypical states, artifacts that are only possible if we force the system to start from there, which are soon left, with no return possibility. Thus we will simply get rid of R' and consider the restricted state space $\widehat{R} = R \setminus R'$, and then define the subset $R^* = \widehat{R} \setminus A$, which is the support of $(Z_t)_{t \geq 0}$ previous to extinction. We are now set to prove the main result of this section.

Proposition 3.2. $(Z_t)_{t \geq 0}$ is irreducible on R^* .

Proof. Fix $x, y \in R^*$ and some $t > 0$. We shall find a sequence of states $y^0, y^1, \dots, y^n \in R^*$ such that for $s > 0$ we have

$$\begin{aligned} \mathbb{P}(Z_s^x = y^0) &> 0, \\ \mathbb{P}(Z_s^{y^i} = y^{i+1}) &> 0 \text{ for any } i \in \{0, \dots, n-1\}, \\ \text{and } \mathbb{P}(Z_s^{y^n} = y) &> 0. \end{aligned}$$

Of course the exact value of s is unimportant, as in our continuous time setting if this is true for some s then it is true for any s . Once this has been established, the result evidently follows from Markov property:

$$\begin{aligned} \mathbb{P}(Z_t^x = y) &\geq \mathbb{P}\left(Z_{\frac{t}{n+2}}^x = y^0, Z_{\frac{2t}{n+2}}^x = y^1, \dots, Z_{\frac{(n+1)t}{n+2}}^x = y^n, Z_t^x = y\right) \\ &= \mathbb{P}\left(Z_{\frac{t}{n+2}}^x = y^0\right) \mathbb{P}\left(Z_{\frac{t}{n+2}}^{y^0} = y^1\right) \dots \mathbb{P}\left(Z_{\frac{t}{n+2}}^{y^n} = y\right) > 0. \end{aligned}$$

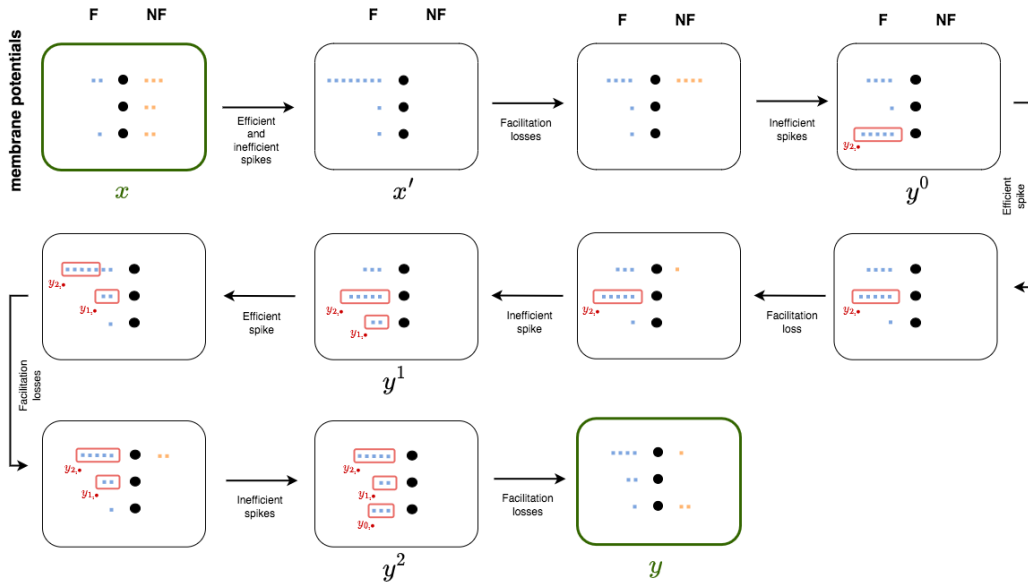


Figure 7: Illustration of the idea behind the proof of Proposition 3.2 in a minimal example with $N = 10$ and $\theta = 2$.

We now define our sequence. It is advisable to take a glance at Figure 7 before reading the following text, which is a somewhat laborious (but necessary) translation of a nonetheless simple idea. For the sake of notation clarity the first state in the sequence will actually be written x' , the next state being y_0 and so on. x' is the state in R^* which is such that $x'_{\theta,1} = N - \theta$, and $x'_{i,1} = 1$ for all $i \in \{0, \dots, \theta - 1\}$ (and of course all other coordinates to 0). One can easily see that there is always a positive probability to reach x' from x . Indeed suppose that, starting from x , the system undergoes exactly θ efficient spikes (and that nothing else happens in-between). This sequence of events has positive probability starting from A^c , moreover this has the effect of pushing all non-facilitated neurons to level θ . Then suppose that all these non-facilitated neurons at level θ spike—and thus become facilitated—and that then the systems undergoes again exactly θ effective spikes (and that no facilitation loss happens in the mean time). This sequence of events is of positive probability and whatever x is, the process ends up on x' .

Now we define the other elements of the sequence. For any $k \in \{0, \dots, \theta - 1\}$ the state y^k is defined by:

$$\begin{aligned} y_{i,1}^k &= y_{\theta-k+i,\bullet} \text{ for } i \in \{0, \dots, k\}, \\ y_{\theta,1}^k &= N - \theta + k + 1 - \sum_{i=\theta-k}^{\theta} y_{i,\bullet}, \end{aligned}$$

and if $k < \theta - 1$ then $y_{i,1}^k = 1$ for $i \in \{k + 1, \dots, \theta - 1\}$.

Moreover the penultimate state in our sequence (before y) is y^θ , defined by $y_{i,1}^\theta = y_{i,\bullet}$ for $i \in \{0, \dots, \theta\}$. First we show that for any $k \in \{0, \dots, \theta\}$ the state y^k is indeed outside A . In order to do this it is sufficient to show that the number of facilitated neurons at any level is greater or equal to 1. The fact that $y \notin R'$ implies (by definition) that $y_{i,\bullet} \geq 1$ for any $i \in \{0, \dots, \theta\}$, and thus we have $y^\theta \notin A$ as well. Now pick some $k \in \{0, \dots, \theta - 1\}$. For $i \in \{0, \dots, k\}$ we have $y_{i,1}^k \geq 1$ for the same reason, and for $i \in \{k + 1, \dots, \theta - 1\}$ we have $y_{i,1}^k = 1$. Finally to check $y_{\theta,1}^k \geq 1$ use again the fact that $y \notin R'$ and notice that therefore

$$\sum_{i=\theta-k}^{\theta} y_{i,\bullet} = n - \sum_{i=0}^{\theta-k-1} y_{i,\bullet} \leq n - \theta + k.$$

This implies that

$$y_{\theta,1}^k = n - \theta + k + 1 - \sum_{i=\theta-k}^{\theta} y_{i,\bullet} \geq 1.$$

All levels are checked, and therefore $y^k \notin A$.

Let see how the system might reaches y^0 from x' . y^0 is the state such that $y_{\theta,1}^0 = n - \theta - (y_{\theta,\bullet} - 1)$, $y_{1,1}^0 = y_{2,1}^0 = \dots = y_{\theta-1,1}^0 = 1$ and $y_{0,1}^0 = y_{\theta,\bullet}$. Then it is clear that y^0 is reached from x' if exactly $y_{\theta,\bullet} - 1$ neurons at level θ lose their facilitation and then emit a non-efficient spike one after the other. This happens with positive probability. The fact that there is enough facilitated neurons at level θ in x' for this to happen follows from the fact that $y_{\theta,\bullet} = n - \sum_{i=0}^{\theta-1} y_{i,\bullet} \leq n - \theta$ (remember that $y \notin R'$, so that $y_{i,\bullet} \geq 1$ for any $i \in \{0, \dots, \theta - 1\}$).

Now we shall find a way of going from y^k to y^{k+1} for any $k \in \{0, \dots, \theta - 1\}$. We know that $y^k \notin A$ so that with positive probability there is an efficient spike occurring in the system before anything else. Let z^k denote the state reached after this spike. The number of active neurons at level θ in the current state of the process cannot decrease, since one neuron was taken out (the spiking neuron) and at least one was added (exactly one for $k < \theta - 1$ and possibly more if $k = \theta - 1$), that is $z_{\theta,1}^k \geq y_{\theta,1}^k$. For the levels from 1 to $\theta - 1$ the values are simply pushed one-step upward, that is $z_{i,1}^k = y_{i-1,1}^k$ for any $i \in \{1, \dots, \theta - 1\}$. Finally evidently $z_{0,1}^k = 1$. In other words we have $z_{i,1}^k = y_{i,1}^{k+1}$ for any $i \in \{1, \dots, \theta - 1\}$ and to go from z^k to y^{k+1} it only remains to find a way to make it agree at level 0 and θ as well. But this is easy using the same trick as in the previous paragraph: it suffices that the system undergoes exactly $y_{\theta-(k+1),\bullet} - 1$ successive facilitation losses at level θ and then the same number of inefficient spikes. This obviously fixes the number of neurons at level 0, and therefore the number at level θ as well, since the different levels have to sum up to n . Moreover there is always enough facilitated neurons at level θ in z^k for this to happen with positive probability, as shown by the following computation:

$$\begin{aligned} z_{\theta,\bullet}^k &= y_{\theta,\bullet}^k = n - \theta + k + 1 - \sum_{i=\theta-k}^{\theta} y_{i,\bullet} \\ &= \sum_{i=0}^{\theta-k-1} y_{i,\bullet} + \sum_{i=\theta-k}^{\theta} y_{i,\bullet} - \theta + k + 1 - \sum_{i=\theta-k}^{\theta} y_{i,\bullet} \\ &= \sum_{i=0}^{\theta-(k+1)} y_{i,\bullet} - \theta + k + 1 \\ &= y_{\theta-(k+1),\bullet} + \sum_{i=0}^{\theta-(k+1)-1} y_{i,\bullet} - \theta + k + 1 \\ &\geq y_{\theta-(k+1),\bullet}. \end{aligned}$$

To obtain the last inequality we have used the fact that, as $y \notin R'$, the elements in the sum need to be greater or equal to 1.

It only remains to go from y^θ to y , but this is very easy: this will happen if there is exactly $y_{0,0}$ facilitation loss(es) at level 0 in the process starting from y^θ , then $y_{1,0}$ facilitation loss(es) at level 1 and so on, with no spike in the mean time. This happens with positive probability and the proof is over. □

4 Quasi-stationary distribution

As a consequence of the irreducibility proven in the previous section, classical results³ guarantee that there exists a unique probability measure μ supported by R^* , called the *quasi-stationary distribution* of $(Z_t)_{t \geq 0}$, which is such that for any $z \in R^*$ and any $t \geq 0$:

$$\mathbb{P}(Z_t^\mu = z \mid Z_t^\mu \notin A) = \mu(z).$$

Furthermore, not only the theory guarantees the existence and uniqueness of such probability measure, but it also provides a direct method to actually compute it. In order to explain how this is done we consider the infinitesimal generator Q of $(Z_t)_{t \geq 0}$. Suppose we re-indexed the states by mapping the elements in A into $\{1, \dots, |A|\}$ and the ones of R^* into $\{|A| + 1, \dots, |A| + |R^*|\}$. Then the matrix Q looks like

$$Q = \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{R} & \mathbf{T} \end{pmatrix},$$

where the matrix \mathbf{A} is the sub-generator corresponding to the transitions occurring inside A , \mathbf{T} is the matrix corresponding to the transitions occurring inside R^* , \mathbf{R} is the matrix corresponding to the transitions occurring from R^* to A , and $\mathbf{0}$ is a matrix full of 0 with $|A|$ rows and $|R^*|$ columns. As explained earlier we might as well identify A to a single absorbing state and hence write

$$Q = \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{r} & \mathbf{T} \end{pmatrix},$$

where $\mathbf{0}$ designates the null row vector of length $|R^*|$ and \mathbf{r} is a non-null column vector, of length $|R^*|$ as well. Then, following [26] (pages 214 and 215), the appropriate version of Perron-Frobenius theorem ensures the existence of a left-eigenvector $\mu \in [0, 1]^{|R^*|}$ associated to the Perron-Frobenius eigenvalue α_{PF} :

$$\mu \mathbf{T} = \alpha_{PF} \mu. \tag{4.1}$$

We know moreover that $\alpha_{PF} \in \mathbb{R}$ and that it is the largest eigenvalue in the real part, that is $\text{Re}(\alpha_{PF}) > \text{Re}(\alpha)$ for any eigenvalue $\alpha \neq \alpha_{PF}$. On the other hand we also know that μ is unique up to a multiplicative constant, so that there exists only one vector satisfying both (4.1) and

$$\mu \mathbf{1} = 1, \tag{4.2}$$

where $\mathbf{1}$ denotes the column vector of length $|R^*|$ with only 1's. These facts clearly suggest a simple procedure to obtain a particular QSD:

1. first compute the eigenvalues of the matrix \mathbf{T} and keep only the one that is maximal in the real part,
2. then obtain the left eigenvector associated to this Perron-Frobenius eigenvalue and normalize it so that it satisfies (4.2).

³good references on the subject are [12], [33] and section 4.6 in [26]

However, as n and θ grow, the computational cost of using such method for our model rapidly explodes, so that we ought to design a more efficient approximate method. This is actually the subject of Section 7. Now, writing $\mathbf{T}(t)$ for the (lossy) transition matrix associated to the transitions occurring in R^* up to time t we have

$$\mathbf{T}(t) = \sum_{k=0}^{\infty} \frac{\mathbf{T}^k t^k}{k!},$$

so that applying (4.1) and denoting $\gamma = -\alpha_{PF}$ gives

$$\mu \mathbf{T}(t) = \sum_{k=0}^{\infty} \frac{(\alpha_{PF} t)^k}{k!} \cdot \mu = e^{-\gamma t} \mu. \tag{4.3}$$

Then taking the sum of all the elements of both vectors on the left-hand side and on the right-hand side respectively in the equation above gives that for any $t \geq 0$

$$\mathbb{P}(Z_t^\mu \notin A) = e^{-\gamma t}. \tag{4.4}$$

In words, the time of extinction of the system when the initial state is chosen according to the QSD has an exponential distribution of some parameter γ . As γ is simply $-\alpha_{PF}$, a by-product of the two-steps procedure defined earlier is that we immediately obtain the rate of the extinction time when the system start from the QSD. Moreover, a classical result states that not only μ is quasi-stationary, but it is also the unique Yaglom limit, meaning that for any states x and z in R^* the following holds⁴

$$\lim_{t \rightarrow \infty} \mathbb{P}(Z_t^x = z \mid Z_t^x \notin A) = \mu(z). \tag{4.5}$$

In other words, whatever the distribution of the initial state, there is always *relaxation* toward the QSD for the unabsorbed process. This two last points are of crucial importance, since one of our goal is to establish the metastable nature of our system, which requires that we show that the time of extinction, starting from any given state in R^* , is exponentially distributed. But thanks to (4.4) and (4.5) this will be established at the condition that we show that the relaxation time toward the QSD is negligible relative to the extinction time. This last point will be studied by numerical means.

5 Solution for N small

A concrete example of an explicit QSD construction (Sec. 4) for the model of Sec. 2.3 is presented next. Since the size of model state space grows very rapidly with N and θ , we consider here rather small values for these parameters, namely $N = 5$ and $\theta = 1$.

5.1 The Q matrix as a graph

The state space has only 56 elements and all the states can be displayed as shown on Fig. 8. On this figure the states are explicitly represented as z^T (columns correspond to membrane potential values: 0 for the first column and 1 for the second; rows correspond to the synaptic facilitation value: 0 for the first row and 1 for the second). Each element of z^T contains the “headcounts”: how many of the five neurons of the network have the corresponding membrane potential and synaptic facilitation state. The thickness of the arrows represents the type of transition: synaptic facilitation loss for thin arrows and spike for thick ones. The states belonging to R^* are shown with a light blue background on Fig. 8. Each such state has at least one arrow leading to it from another “blue” state and at least one arrow leaving it for another “blue” state.

⁴See Theorem 4.27 in [26].

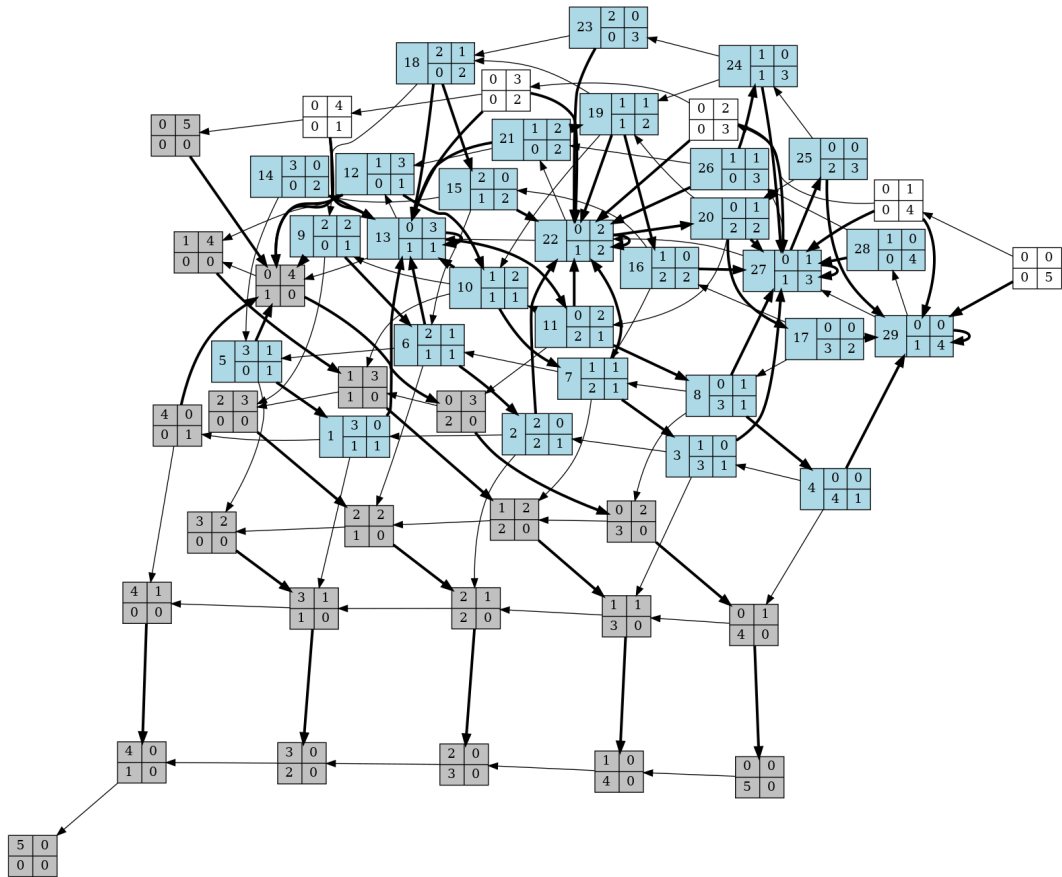



Figure 8: A network with 5 neurons and $\theta = 1$. Configurations that belong to set A appear in gray filled boxes, the ones that belong to set R^* appear in light blue filled boxes with the index of the configuration on the left (in order to read the next matrix), while the ones belonging to set $R' \setminus A$ appear in white filled boxes. Transitions due to a spike appear as thick arrows, while transitions due to a synaptic facilitation loss appear as thin arrows.

5.2 The T matrix

Matrix **T** of Sec. 4 can be here explicitly represented as follows (the row and column indices correspond to the numbers of the “blue” states on Fig. 8):



The zeros are not explicitly written but replace by a "•". The factor with which λ should be multiplied appear in orange, while the factor with which β should be multiplied appear in blue. We see for instance that state 29 (last row) can make a transition to state 28 (last row and penultimate column) by a *single* synaptic facilitation loss or to state 27 (last row and antepenultimate column) by one among 4 possible synaptic facilitation losses. These two transitions correspond to what is displayed on Fig. 8. The actual values of the diagonal are given next:

index	rate	index	rate	index	rate
1	$-2\lambda - \beta$	11	$-3\lambda - 3\beta$	21	$-2\lambda - 4\beta$
2	$-3\lambda - \beta$	12	$-\lambda - 4\beta$	22	$-3\lambda - 2\beta$
3	$-4\lambda - \beta$	13	$-2\lambda - 3\beta$	23	$-3\lambda - 3\beta$
4	$-5\lambda - \beta$	14	$-2\lambda - 2\beta$	24	$-4\lambda - 3\beta$
5	$-\lambda - 2\beta$	15	$-3\lambda - 2\beta$	25	$-5\lambda - 3\beta$
6	$-2\lambda - 2\beta$	16	$-4\lambda - 2\beta$	26	$-3\lambda - 4\beta$
7	$-3\lambda - 2\beta$	17	$-5\lambda - 2\beta$	27	$-4\lambda - \beta$
8	$-4\lambda - 2\beta$	18	$-2\lambda - 3\beta$	28	$-4\lambda - 4\beta$
9	$-\lambda - 3\beta$	19	$-3\lambda - 3\beta$	29	-5λ
10	$-2\lambda - 3\beta$	20	$-4\lambda - 3\beta$		

We see that matrix **T** is very sparse, having at most 3 positive non-diagonal elements per row.

5.3 Case $N = 5, \theta = 1, \beta = 10, \lambda = 4$

We compare now the theoretical/direct results provided by solving Eq. (4.1) and (4.2) with a straightforward simulation of the Markov process (see Sec. A for implementation details). 10^5 independent replicates were simulated, all starting with each of the 5 neurons in the susceptible state ($U=1$), with a facilitated synapse ($F=1$). The simulation of a given replicate was stopped when one of the A states was reached or, if it survived till then, when time 4 was reached. Fig. 9 shows this comparison. The left part of the figure displays, on a log ordinate scale, the empirical survival function (a replicate is “dead” when it reaches one of the A states), as well as the theoretical one provided by Eq. (4.4). The right part of the figure displays the empirical mean headcounts for each

state that a single neuron can have—namely: $(U = 0, F = 0)$ (orange), $(U = 0, F = 1)$ (blue), $(U = 1, F = 0)$ (red) and $(U = 1, F = 1)$ (black)—, together with the expected values computed from the QSD (horizontal lines):

$$\mu_{i,j} = \int_{R^*} z_{i,j} d\mu = \sum_{z \in R^*} z_{i,j} \mu(z), \quad i, j \in \{0, 1\},$$

where $\mu(\bullet)$ is the QSD (Eq. (4.1) and (4.2)). Notice the *fast relaxation* of the empirical means towards the corresponding expected values; remember that each replicate starts with the 5 neurons being susceptible ($U = 1$) with a facilitated synapse, a state that is not in R^* (see Fig. 8, rightmost state). Sec. B presents additional simulations showing that the case we just illustrated is at least reasonably representative. Our stochastic

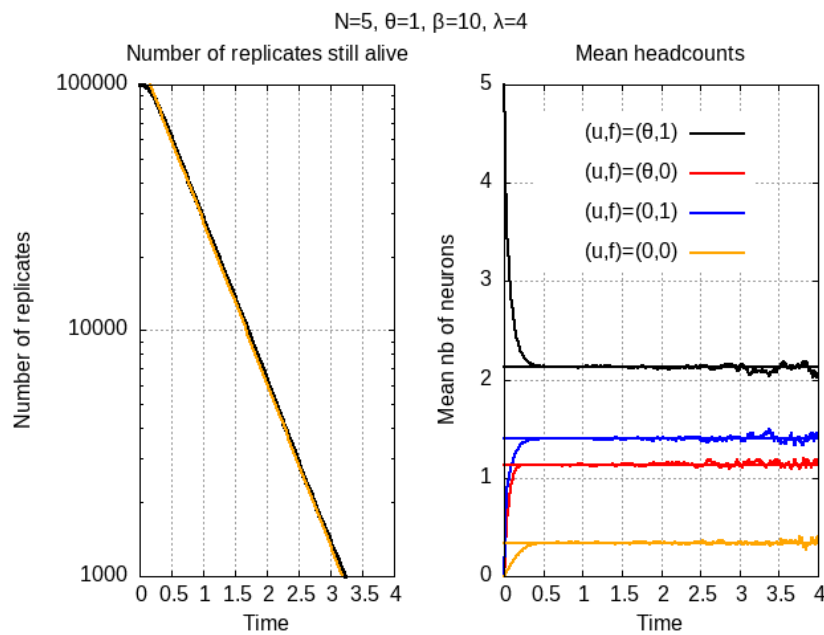


Figure 9: Stochastic simulations and direct solution comparison for a network with $N = 5$, $\theta = 1$, $\beta = 10$, $\lambda = 4$. 10^5 replicates were used for the simulations. Left, the empirical survival function (black), together with the theoretical straight line (orange) whose slope is given by γ in Eq. (4.1), (4.3) and (4.4); right, the empirical mean headcounts of the four states (see legend), together with the expected values computed from the QSD (Eq. (4.2)) shown as horizontal lines.

simulations can be used for computing the empirical distributions, among the R^* states, of the alive replicates at different times. This is what is shown on the upper part of Fig. 10. There a “line plus glyph” display is used instead of the vertical bars that would have been more appropriate for a *single* distribution. When looking at these graphs, the reader should therefore keep in mind that the only meaningful values are the ones indicated by the glyphs. The latter are located at integer values of the abscissa and correspond to the R^* states indexed on Fig. 8. The lines in-between the glyphs are here to help seeing the *profile* of the distributions: the distributions at the four different times are (essentially) a scaled copy of a *fundamental* one, the QSD. The decaying values of the empirical distributions with time is just the consequence of replicates leaving the R^* set for the A set (Fig. 9, left). The profile identity of the empirical distributions is what is illustrated on the bottom part of Fig. 10: here the empirical distributions have

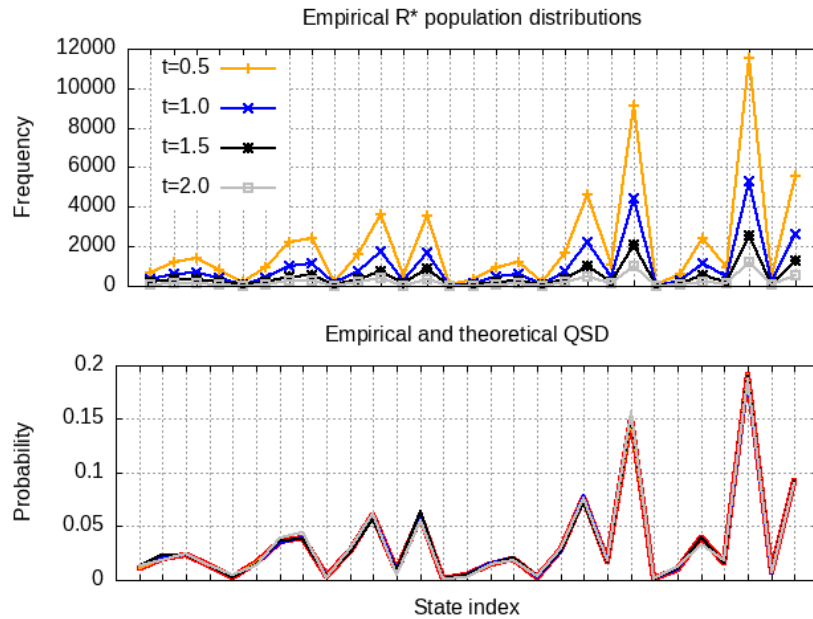


Figure 10: Stochastic simulations of a network with $N = 5$, $\theta = 1$, $\beta = 10$, $\lambda = 4$. 10^5 replicates were used for the simulations. Top, empirical distributions of the still alive replicates among the 29 states of R^* at four different times: 0.5, 1.0, 1.5 and 2.0. Bottom, scaled versions of the above distributions (they sum to 1) and theoretical QSD (thick red line).

been normalized to sum to 1, and the QSD as been added (thick red line). The key feature justifying the name of the *quasi-stationary distribution* (QSD) is illustrated here: considering the *alive replicates only*, the fraction of them in each state of R^* *does not change with time*.

6 Solution for N big

When N and θ grow, the number of states in R^* , $\binom{N+2\theta+1}{2\theta+1}$, becomes very quickly, very large. Finding the Perron-Frobenius eigenvalue α_{PF} (Eq. (4.1)) and the QSD (Eq. (4.1) and (4.2)), as we did in the previous section, is not doable anymore and we have to stick to stochastic simulations. The direct comparison of the theoretical and simulation based results presented in the last section does nevertheless give us confidence in our stochastic simulation code—we can add to that the simplicity of this code of course. We consider here networks with a fixed ratio $N/\theta = 5$, $N = 50$ and $N = 500$. Like in the previous section, 10^5 replicates are used and all replicates start with all the neurons in the susceptible state with a facilitated synapse. Fig. 11 shows the increase of the survival time with the network size (N at constant N/θ ratio). Looking at the graph we see that the level of $10^5/e$ (the level at one time constant) is reached at time ~ 0.5 with $N = 5$, ~ 1.5 with $N = 50$ and ~ 3.8 with $N = 500$.

On the other hand, as shown by Fig. 12, the relaxation from the initial state, that is not in R^* , towards the QSD value is “fast” (< 1). It does not depend strongly on the network size (the empirical mean displayed on the figure for $N = 50$ has been scaled to facilitate the comparison). This figure only shows the evolution of the mean number of susceptible neurons with a facilitated synapse in order to reduce clutter/improve




Figure 11: Empirical survival function from stochastic simulations of networks with $N = 5, \theta = 1$ (black), $N = 50, \theta = 10$ (blue), $N = 500, \theta = 100$ (orange) and $\beta = 10, \lambda = 5$. 10^5 replicates were used for each parameter set. Each replicate with each parameter set started with all the neurons susceptible and all synapses facilitated. Dotted horizontal line, level at one time constant.

readability⁵. Seeing this figure, the reader should abstain from jumping to the conclusion that the mean QSD value, $\mu_{\theta,1}$, is proportional to the network size. As the approximate solution developed in the next section will make clear, this is only true when $\beta \mu_{\theta,1} \gg \lambda$ and this order of magnitude difference becomes more accurate as N grows. Looking at the middle right panel of Fig. 13 in Sec. B the reader can see that for $N = 5$, the QSD mean value is closer to 2 (and would therefore be at 200 using the proper scaling on Fig. 12).

7 Approximation of the first moment of the Quasi-stationary distribution: mean-field analysis

7.1 What is approximated?

The method presented in Section 4 to obtain the QSD of the system using Perron-Frobenius Theorem becomes rapidly intractable as the number of neurons in the system grows. To detect the emergence of the QSD along a particular trajectory it might therefore be useful to design some less costly approximate method.

In order to do so we would like to compute the expected values of the number of elements in each possible combinations of membrane potential and facilitation states under the QSD. While in the finite setting the only actual invariant measure is trivial⁶, and therefore unrelated to the QSD, solving the steady-state equations outside of the absorbing region shall give a good approximation of the QSD at the macroscopic scale (when N is big). Moreover, in accordance with classical mean-field theory, we might

⁵Showing all empirical means would require 22 curves for $\theta = 10$ and 102 for $\theta = 100$.

⁶In other words it concentrates all the mass on the absorbing states.




Figure 12: Scaled mean number of susceptible neurons with a facilitated synapse as a function of time. Parameter used: $N = 50, \theta = 10$ (blue), $N = 500, \theta = 100$ (orange) and $\beta = 10, \lambda = 5$. 10^5 replicates were used for each parameter set. Each replicate with each parameter set started with all the neurons susceptible and all synapses facilitated. The dotted lines show the approximate QSD mean values. They are the numerical solutions of the implicit Eq. (7.23) obtained in the next section.

neglect pairwise correlations between neurons in order to solve these equations, and therefore proceed to obtain an explicit macroscopic approximation of the QSD.

More precisely, for any $(i, j) \in \{0, \dots, \theta\} \times \{0, 1\}$, let $\mu_{i,j}$ denote the QSD mean value:

$$\mu_{i,j} = \int_{R^*} z_{i,j} d\mu.$$

We adopt the same notation as for elements of R and for any $i \in \{0, \dots, \theta\}$ we write

$$\mu_{i,\bullet} = \mu_{i,0} + \mu_{i,1}. \tag{7.1}$$

In the following we write the infinitesimal generator of $(Z_t)_{t \geq 0}$ as a functional operator \mathcal{L} , defined by:

$$\mathcal{L}f(z) = \beta z_{\theta,1} (f(\pi(z)) - f(z)) + \beta z_{\theta,0} (f(\pi^*(z)) - f(z)) + \sum_{i=0}^{\theta} \lambda z_{i,1} (f(\pi_i^\dagger(z)) - f(z)) \tag{7.2}$$

where f is any suitable test function and $z \in R$. In order to get the infinitesimal rate of change of the number of neuron at coordinate $(i, j) \in \{0, \dots, \theta\} \times \{0, 1\}$ in $(Z_t)_{t \geq 0}$, starting from some state $z \in R$, we consider the functions from R to $\{0, \dots, N\}$ defined by $f_{i,j} : z \mapsto f_{i,j}(z) = z_{i,j}$. We perform the computation explicitly for $f_{0,0}$ and let the reader check the other cases, which follow from similar calculations. The second term in the right hand side of (7.2) equals 0, as $f_{0,0}(\pi^*(z)) - f_{0,0}(z) = 0$ —an inefficient spike doesn't changes the number of neurons at the coordinate $(0, 0)$. In the third term, all the elements of the sum equals 0 except for $i = 0$, where $f_{0,0}(\pi_0^\dagger(z)) - f_{0,0}(z) = 1$ —as a facilitation loss affects the number of non-facilitated neurons at level 0 only if it occurs on a neuron of null membrane potential, and when it does so it increases this number by one unit. For the first term, we have $f_{0,0}(\pi(z)) = 0$, and therefore $f_{0,0}(\pi(z)) - f_{0,0}(z) = -z_{0,0}$. Finally the complete expression is

$$\mathcal{L}f_{0,0}(z) = -\beta z_{\theta,1} z_{0,0} + \lambda z_{0,1}. \tag{7.3}$$

Now, if μ were a stationary distribution for $(Z_t)_{t \geq 0}$, by standard results on Markov processes⁷ we should have $\int \mathcal{L}f_{i,j} d\mu = 0$. Of course, strictly speaking, μ is not a stationary distribution for $(Z_t)_{t \geq 0}$, as it is only a quasi-stationary distribution, or in other words a stationary distribution for the modification of $(Z_t)_{t \geq 0}$ in which we have prohibited any transition to the absorbing region A . We will assume nonetheless here that μ is close enough to be a stationary distribution for $(Z_t)_{t \geq 0}$ and try to solve $\int \mathcal{L}f_{i,j} d\mu = 0$ for $\mu_{i,j}$. Using (7.3), this gives

$$0 = -\beta \int z_{\theta,1} z_{0,0} d\mu + \lambda \mu_{0,1}.$$

We proceed to further approximations by neglecting pairwise correlations, and we finally obtain:

$$0 = -\beta \mu_{\theta,1} \mu_{0,0} + \lambda \mu_{0,1} \tag{7.4}$$

⁷See for example theorem 3.37 in [29].

7.2 Getting the approximate expected values of the QSD

One can proceed similarly for all $f_{i,j}$, writing the steady-state equation $\int \mathcal{L} f_{i,j} d\mu = 0$ as a quadratic multilinear polynomial equation by neglecting pairwise correlations. This leads to the following system of equations.

$$0 = -\beta \mu_{\theta,1} \mu_{0,1} - \lambda \mu_{0,1} + \beta \mu_{\theta,\bullet}. \quad (7.5)$$

For $i \in \{1, \dots, \theta - 1\}$:

$$0 = -\beta \mu_{\theta,1} (\mu_{i,0} - \mu_{i-1,0}) + \lambda \mu_{i,1}, \quad (7.6)$$

$$0 = -\beta \mu_{\theta,1} (\mu_{i,1} - \mu_{i-1,1}) - \lambda \mu_{i,1}. \quad (7.7)$$

And for $\mu_{\theta,0}$ and $\mu_{\theta,1}$:

$$0 = -\beta \mu_{\theta,0} + \beta \mu_{\theta,1} \mu_{\theta-1,0} + \lambda \mu_{\theta,1}, \quad (7.8)$$

$$0 = -\beta \mu_{\theta,1} + \beta \mu_{\theta,1} \mu_{\theta-1,1} - \lambda \mu_{\theta,1}. \quad (7.9)$$

Now we introduce:

$$\rho_i = \frac{\mu_{i,1}}{\mu_{i,\bullet}}, \quad i = 0, \dots, \theta. \quad (7.10)$$

In what follows we assume $N > \theta$, since otherwise the absorbing region covers the whole state space. The quantity $\mu_{i,\bullet}$ is the total number of neurons whose membrane potential is i in the network, regardless of their synaptic facilitation status. ρ_i is the fraction of neurons with a facilitated synapse among the $\mu_{i,\bullet}$ neurons whose membrane potential is i . We have obviously: $\mu_{i,1} = \mu_{i,\bullet} \rho_i$ and $\mu_{i,0} = \mu_{i,\bullet} (1 - \rho_i)$. If we add Eq. 7.4 and 7.5 we get with Eq. 7.1 and 7.10:

$$0 = -\beta \mu_{\theta,\bullet} (\rho_\theta \mu_{0,\bullet} - 1).$$

The QSD must have a non null value of $\mu_{\theta,1}$ and therefore of $\mu_{\theta,\bullet}$ otherwise it would be identical to the absorbing state. In a more useful way we can then write the last equation as:

$$\rho_\theta \mu_{0,\bullet} = 1. \quad (7.11)$$

If we add Eq. 7.6 and 7.7 we get with Eq. 7.1:

$$\mu_{i,\bullet} = \mu_{i-1,\bullet} \quad i = 1, \dots, \theta - 1. \quad (7.12)$$

If we add Eq. 7.8 and 7.9 we get with Eq. 7.1:

$$\rho_\theta \mu_{\theta-1,\bullet} = 1. \quad (7.13)$$

Eq. 7.12 implies that we can define κ such that:

$$\kappa = \mu_{0,\bullet} = \mu_{1,\bullet} = \dots = \mu_{\theta-1,\bullet}. \quad (7.14)$$

In words, all the states below θ are equally populated. Since $N = \sum_{i=0}^{\theta} \mu_{i,\bullet}$, the QSD must satisfy:

$$N = \kappa \theta + \mu_{\theta,\bullet}. \quad (7.15)$$

Combining Eq. 7.11 or 7.13 with Eq. 7.14, we get:

$$\rho_\theta = \mu_{\theta,1} / \mu_{\theta,\bullet} = 1 / \kappa. \quad (7.16)$$

Using Eq. 7.1, 7.10, 7.14 and 7.16, we can rewrite Eq. 7.5 as:

$$0 = -\beta \mu_{\theta,1} \rho_0 \kappa - \lambda \rho_0 \kappa + \beta \mu_{\theta,\bullet},$$

or

$$\rho_0 = \frac{\beta\mu_{\theta,\bullet}}{(\lambda + \beta\mu_{\theta,1})\kappa},$$

but $1/\kappa = \mu_{\theta,1}/\mu_{\theta,\bullet}$ (Eq. 7.16), therefore:

$$\rho_0 = \frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}}. \quad (7.17)$$

Using Eq. 7.1, 7.10, 7.14 and 7.16, we can rewrite Eq. 7.7 as:

$$0 = -\beta\mu_{\theta,1}(\rho_i - \rho_{i-1})\kappa - \lambda\kappa\rho_i, \quad i = 1, \dots, \theta - 1,$$

leading to:

$$\rho_i = \frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}}\rho_{i-1}, \quad i = 1, \dots, \theta - 1. \quad (7.18)$$

Combining Eq. 7.17 and Eq. 7.18, we get:

$$\rho_i = \left(\frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}} \right)^{i+1}, \quad i = 0, \dots, \theta - 1. \quad (7.19)$$

Now using Eq. 7.1, 7.10, 7.14 and 7.16, we can rewrite Eq. 7.9 as:

$$\rho_{\theta-1} = \frac{\lambda + \beta}{\kappa\beta}. \quad (7.20)$$

Eq. 7.19 and 7.20 together lead to:

$$\left(\frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}} \right)^\theta = \frac{\lambda + \beta}{\kappa\beta}. \quad (7.21)$$

But Eq. 7.15 and 7.16 imply that:

$$\frac{1}{\kappa} = \frac{\theta + \mu_{\theta,1}}{N}, \quad (7.22)$$

leading to:

$$\mu_{\theta,1} = \frac{N\beta}{\lambda + \beta} \left(\frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}} \right)^\theta - \theta. \quad (7.23)$$

$\mu_{\theta,1}$ must be solution of Eq. 7.23. If we find such a $\mu_{\theta,1}$, we immediately get κ (Eq. 7.22) and $\mu_{\theta,\bullet}$, the number of neurons at or above the spiking threshold θ (Eq. 7.16). With Eq. 7.19, we get the successive $\mu_{i,0}$ and $\mu_{i,1}$ for $i = 0, \dots, \theta - 1$, that is, the whole description of the QSD in terms of its expected values. A comparison between this approximation for $\mu_{\theta,1}$ and empirical values computed from stochastic simulations is shown on Fig. 12 in the previous section.

7.3 Remarks

The $\mu_{\theta,1}$ solution(s) of Eq. (7.23) depend on β and λ only through their ratio: $\eta = \lambda/\beta$. So multiplying β and λ by the same factor does not change the solution(s) of Eq. (7.23). We can consider limiting cases to check if the argument we just developed makes sense. A network without facilitation loss ($\lambda = 0$) is an instance of a limiting case. For such a network, as soon as a neuron has spiked, its synapse gets facilitated and remains facilitated forever. Then Eq. (7.23) leads to, $\mu_{\theta,1} = N - \theta$; Eq. (7.22) gives, $\kappa = 1$ and Eq. (7.17) implies that, $\mu_{\theta,0} = 0$. The use of Eq. (7.14) together with Eq. (7.19), yields: $\mu_{i,0} = 0$ and $\mu_{i,1} = 1$ for $i = 0, \dots, \theta - 1$. We have therefore one neuron (with a facilitated synapse) at each possible membrane potential value below threshold and all the other

neurons (with a facilitated synapse) at threshold. This is the state x' defined in Sec. 3.2 and in Fig. 7 and this is precisely what we expect from the argument developed in that section.

If there is no threshold ($\theta = 0$), every neuron of the network has two accessible states: facilitated synapse and un-facilitated synapse; it goes from the first to the second with rate β and from the second to the first with rate λ . Elementary Markov process results tell us that the probability of finding a neuron in the facilitated state is $\frac{\beta}{\beta+\lambda}$ leading to $\mu_{0,1} = \frac{N\beta}{\beta+\lambda}$ as given by Eq. (7.23) when $\theta = 0$.

7.4 Back-of-the-envelope calculation leading to ρ_θ

We can get an approximate value of $\rho_\theta = \frac{\theta+\mu_{\theta,1}}{N}$ (Eq. (7.11) and 7.22) with an even bolder approach. Quantity $\beta\mu_{\theta,1}$ is the rate of efficient spikes; spike making all the other neurons climb one step on the membrane potential ladder. A neuron that just spiked has a null membrane potential and a facilitated synapse. That neuron needs θ efficient spikes in order to be able to spike again and that will require on average $\theta/(\beta\mu_{\theta,1})$ time units. Once the neuron membrane potential reaches θ , that neuron can spike again and will do so with a rate β . The question is: what is the “probability” for the synapse of that neuron to be still facilitated when its next spike comes? In other words, what is the value of ρ_θ ? Facilitation is lost with rate λ , that probability is therefore:

$$\int_0^\infty e^{-\lambda(\theta/(\beta\mu_{\theta,1})+t)}\beta e^{-\beta t} dt = \frac{\beta}{\lambda + \beta} \exp\left(-\frac{\lambda\theta}{\beta\mu_{\theta,1}}\right).$$

Now assuming that $\beta\mu_{\theta,1} \gg \lambda$, writing

$$\exp\left(-\frac{\lambda\theta}{\beta\mu_{\theta,1}}\right) = \left(\frac{1}{\exp\frac{\lambda}{\beta\mu_{\theta,1}}}\right)^\theta$$

and using a first order Taylor expansion for the exponential function leads to:

$$\rho_\theta = \frac{\beta}{\lambda + \beta} \left(\frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}}\right)^\theta$$

that is what we already obtained (we just have to combine Eq. (7.23), 7.22 and 7.11).

7.5 How good is this approximation

We can first compare the first moments of the QSD derived in Sec. 7.2 with the exact values provided by the complete QSD of Sec. 4 for the case illustrated in Sec. 5.3: $N = 5$, $\theta = 1$, $\beta = 10$, $\lambda = 4$. The exact results are (also shown as horizontal lines on the right panel of Fig. 10):

$$\mu_{0,0} = 0.342, \mu_{0,1} = 1.398, \mu_{1,0} = 1.135, \mu_{1,1} = 2.125$$

The approximate values are:

$$\mu_{0,0} = 0.285, \mu_{0,1} = 1.400, \mu_{1,0} = 1.347, \mu_{1,1} = 1.968$$

Next, we only look at $\mu_{\theta,1}$ for networks with $\beta = 10$, $\lambda = 5$, a ratio $N/\theta = 5$ (first table) or $N/\theta = 10$ (second table) and a sequence of (N, θ) pairs. We compare the output of stochastic simulation with the approximate calculation. Our numerical solution of the implicit Eq. (7.23) gives us 4 significant digits after the decimal point (Sec. A.3). The simulation results are shown with their associated standard error. 10^5 replicates were used for the first 2 entries of each table, 10^4 for the third and 5×10^3 for the fourth.

Three time units were simulated and the mean values computed at time 2 are reported in the tables. We used fewer replicates with the larger N values in order to save time. With $N/\theta = 5$ we observe:

(N, θ)	Approx.	Simulation
(50,10)	12.563	10.76 ± 0.05
(100,20)	24.526	20.20 ± 0.06
(500,100)	119.738	101.4 ± 0.5
(1000,200)	238.661	212.3 ± 0.9

We see that in this setting the relative approximation error is about 20%. If we change slightly the setting, $N/\theta = 10$, in order to get higher spiking rates, we observe:

(N, θ)	Approx.	Simulation
(50,5)	25.216	24.91 ± 0.02
(100,10)	50.400	50.14 ± 0.02
(500,50)	251.866	251.8 ± 0.2
(1000,100)	503.700	503.6 ± 0.3

The standard errors are smaller because: i) fewer replicates reach the “dead set” A per time units—the statistics are therefore computed with more observations—; ii) when the spiking rate is larger, the QSD is less spread around its first moments (its standard deviation is smaller) and our key approximation (Sec. 7.1) is more accurate. We see that when the network spiking rate is large enough (because N/θ is large as shown here, but that also holds if λ is “small”), the approximate solution becomes very close to the truth.

8 Conclusions

We proposed a simple stochastic model of strongly interconnected neurons exhibiting synaptic facilitation, and suggested that the metastable properties of such system might give an interesting explanation of the mechanism behind the sustained activity observed experimentally in networks of neurons involved in short-term memory.

The metastability of this system was established using arguments based on the notion of quasi-stationarity. These arguments are partially based on simulations and are therefore not entirely rigorous from the mathematical point of view. Nonetheless it allows a very simple and intuitive comprehension of the phenomenon of metastability: the system started in any given state rapidly converges to the Yaglom limit, which happens to be the QSD as well, and, once there, it evolves in an almost stationary manner for an exponentially distributed random time, due to classical results about QSD. The fact that the relaxation toward the QSD is much faster than the extinction of the system implies that the extinction time is exponentially distributed as well—asymptotically with respect to the number of neurons—starting from any state out of the absorbing region, i.e. genuine metastability. Establishing that fact is actually the only missing piece to make the argument rigorous, and it should be closely related to a large spectral gap between the eigenvalues with largest and second-largest real parts, because, while extinction is dominated by the former, relaxation ought to be dominated by the latter. This shall be further investigated in future work. It shall also be noticed that such arguments would then apply to a large class of Markov processes (at least any irreducible Markovian system having an absorbing state)—including the famous Contact process for example, which has been rigorously proven to exhibit metastability by various means [39, 8, 15, 36], but with technical complications making the understanding of the fundamental reasons behind the phenomenon difficult to interpret.

A Some numerical details

This section provides a brief outline of the numerical methods implemented to simulate our model, to “solve the QSD” (Eq. (4.1) and (4.2)) and the implicit equation (7.23). The actual codes and their comprehensive documentation can be found on our GitLab repository⁸.

A.1 Simulations

The simulation of the Markov processes defined in Sec. 2.1 and 2.3 is straightforward once the intensity matrix Q is known. Our Fortran code implements algorithm 2.7.2 of [38, p. 80]. This algorithm is just the constructive definition of a Markov process of [23, Sec. 8.3.2, p. 243]. Our simulation code requires a uniform (pseudo)random number generator; we used the implementation of the xoshiro128plus generator [7] provided by Jannis Teunissen’s `rng_fortran`⁹.

A.2 QSD computation

The eigenvalues and left eigenvectors of T (Eq. (4.1) and (4.2)) were obtained with subroutine `dgeev` of LAPACK-3.11.0¹⁰. Since the LAPACK code is written in FORTRAN 77, we used the modern Fortran interface provided on GitHub by the *Numerical Algorithms Group* (NAG)¹¹.

A.3 Solving the implicit Eq. (7.23)

We used a modification of Arjen Markus robust Newton method¹² [32]. This code explicitly brackets the roots of the equation:

$$\mu_{\theta,1} + \theta - \frac{N\beta}{\lambda + \beta} \left(\frac{\beta\mu_{\theta,1}}{\lambda + \beta\mu_{\theta,1}} \right)^\theta$$

and provides therefore an upper bound of its error.

B Additional comparisons between theoretical QSD and simulations

We consider again the setting of Sec. 5, with $N = 5$, $\theta = 1$, $\beta = 10$. As a complement to the case $\lambda = 4$ of Sec. 5.3, we illustrate here cases where λ take values 3, 5 and 8. The equivalent of Fig. 9 is then obtained and shown on Fig. 13. The basic features are the same as what is seen on Fig. 9. Not surprisingly, the mean survival depends strongly on λ ; the smaller the latter, the larger the former. The mean number of neurons in the susceptible state with a facilitated synapse (black curves on the right panels) exhibits a similar trend. The “noisy” aspect of the empirical means displayed on Fig. 13 bottom, right ($\lambda = 8$) for a time larger than 1.25 are due to the small sample size: few replicates are still alive past that time.

⁸https://gitlab.com/c_pouzat/metastability-in-a-system-of-spiking-neurons-with-synaptic-plasticity

⁹https://github.com/jannisteunissen/rng_fortran.

¹⁰<https://www.netlib.org/lapack/>.

¹¹https://github.com/numericalalgorithmsgroup/LAPACK_Examples.

¹²`robust_newton.f90`: https://flibs.sourceforge.net/robust_newton.f90.




Figure 13: Stochastic simulations and direct solution comparison for networks with $N = 5$, $\theta = 1$, $\beta = 10$ and $\lambda = 3$ (top), $\lambda = 5$ (middle), $\lambda = 8$ (bottom). 10^5 replicates were used for the simulations. Left, the empirical survival function (black), together with the theoretical straight line (orange) whose slope is given by γ in Eq. (4.1), (4.3) and (4.4); right, the empirical mean headcounts of the four states (see legend), together with the expected values computed from the QSD (Eq. (4.2)) shown as horizontal lines.

References

- [1] Daniel Amit and Nicolas Brunel, *Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex*, *Cereb. Cortex* **7** (1997), no. 3, 237–252, [doi](#).
- [2] Morgan André, *A result of metastability for an infinite system of spiking neurons*, *J. Stat. Phys.* **177** (2019), no. 5, 984–1008, [doi](#).
- [3] Morgan André and Léo Planche, *The effect of graph connectivity on metastability in a stochastic system of spiking neurons*, *Stochastic Process. Appl.* **131** (2021), 292–310, [doi](#).
- [4] Omri Barak and Misha Tsodyks, *Persistent activity in neural networks with dynamic synapses*, *PLoS Comput. Biol.* **3** (2007), no. 2, 323–332, [doi](#).
- [5] Omri Barak and Misha Tsodyks, *Working models of working memory*, *Current Opinion in Neurobiology* **25** (2014), 20–24, [doi](#).
- [6] Alessandra Bianchi and Alexandre Gaudillière, *Metastable states, quasi-stationary distributions and soft measures.*, *Stochastic Process. Appl.* **126** (2016), no. 6 1622–1680.
- [7] David Blackman and Sebastiano Vigna, *Scrambled linear pseudorandom number generators*, *ACM Trans. Math. Software* **47** (2021), no. 4, 1–32, [doi](#).
- [8] Marzio Cassandro, Antonio Galves, Enzo Olivieri, and Maria Eulália Vares, *Metastable behavior of stochastic dynamics: A pathwise approach*, *J. Statist. Phys.* **35** (1984), no. 5-6, 603–634, [doi](#).
- [9] Nicolas Champagnat and Denis Villemonais, *Quasi-stationary distributions in reducible state spaces*, arXiv:2201.1015.
- [10] Albert Compte, Nicolas Brunel, Patricia S. Goldman-Rakic, and Xiao-Jing Wang, *Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model*, *Cereb. Cortex* **10** (2000), no. 9, 910–923, [doi](#).
- [11] Christos Constantinidis, Shintaro Funahashi, Daeyeol Lee, John D. Murray, Xue-Lian Qi, Min Wang and Amy F.T. Arnsten, *Persistent spiking activity underlies working memory*, *The J. Neurosci.* **38** (2018), no. 32, 7020–7028, [doi](#).
- [12] John Darroch and Eugène Seneta, *On quasi-stationary distributions in absorbing continuous-time finite Markov chains*, *J. Appl. Probability* **4** (1967), 192–196.
- [13] Louis J. DeFelice, *Introduction to Membrane Noise*, Plenum Press, 1981.
- [14] José del Castillo and Bernard Katz, *Quantal components of the end-plate potential*, *J. Physiol.* **124** (1954), no. 3, 560–573, [doi](#).
- [15] Richard Durrett and Roberto Schonmann, *The Contact Process on a Finite Set. ii*, *Ann. Probab.* **16** (1988), no 4, 1570–1583.
- [16] Paul Fatt and Bernard Katz, *Spontaneous subthreshold activity at motor nerve endings*, *J. Physiol.* **117** (1952), no. 1, 109–128, [doi](#).
- [17] William Feller, *An introduction to probability theory and its applications.*, third ed., vol. 1, John Wiley & Sons, 1968.
- [18] Joaquín M. Fuster, *Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory.*, *J Neurophysiol.* **36** (1973), no. 1, 61–78, [doi](#).
- [19] ———, *The prefrontal cortex*, Elsevier Science, 2015.
- [20] Antonio Galves, Eva Löcherbach, Christophe Pouzat, and Errico Presutti, *A System of Interacting Neurons with Short Term Synaptic Facilitation*, *J. Stat. Phys.* **178** (2020), no. 4, 869–892, [doi](#).
- [21] David Hansel and G. Mato, *Short-term plasticity explains irregular persistent activity in working memory tasks*, *J. Neurosci.* **33** (2013), no. 1, 133–149, [doi](#).
- [22] Wilhelm Huisinga, Sean Meyn, and Christof Schütte, *Phase transitions and metastability in markovian and molecular systems*, *Ann. Appl. Probab.* **14** (2004), no 1, 419–458. [doi](#)
- [23] Marius Iosifescu, *Finite markov processes and their applications*, Dover books on mathematics, Dover Publications, 2007.
- [24] Vladimir Itskov, David Hansel, and MishaTsodyks, *Short-term facilitation may stabilize parametric working memory trace*, *Front. Comput. Neurosci.* **5** (2011), 40, [doi](#).

- [25] Bernard Katz and Ricardo Miledi, *Membrane noise produced by acetylcholine*, *Nature* **226** (1970), no. 5249, 962–963, [doi](#).
- [26] Masaaki Kijima, *Markov processes for stochastic modeling*, Springer New York, 1997.
- [27] Kadmo S. Laxa, *Metastability in a stochastic system of spiking neurons with leakage*, *J. Stat. Phys.* **190** (2023), no. 29, [doi](#).
- [28] Matthew L. Leavitt, Diego Mendoza-Halliday, and Julio C. Martinez-Trujillo, *Sustained Activity Encoding Working Memories: Not Fully Distributed*, *Trends in Neurosciences* **40** (2017), no. 6, 328 – 346, [doi](#).
- [29] Thomas Liggett, *Continuous time Markov processes: An introduction*, *Grad. Stud. Math.*, 113, American Mathematical Soc, 2010.
- [30] Liqun Luo, *Principles of neurobiology*, Garland Science, Jul 2015, [doi](#).
- [31] Eva Löcherbach and Pierre Monmarché, *Metastability for systems of interacting neurons*, *Ann. Inst. Henri Poincaré Probab. Stat.* **58** (2022), no 1, 343–378. [doi](#).
- [32] Arjen Markus, *Modern Fortran in Practice*, Cambridge University Press, 2012.
- [33] Sylvie Méléard and Denis Villemonais, *Quasi-stationary distributions and population processes*, *Probab. Surv.* **9** (2012), 340–410. [doi](#)
- [34] Gianluigi Mongillo, Omri Barak, and Misha Tsodyks, *Synaptic theory of working memory*, *Science* **319** (2008), no. 5869, 1543–1546, [doi](#).
- [35] Gianluigi Mongillo, David Hansel, and Carl van Vreeswijk, *Bistability and Spatiotemporal Irregularity in Neuronal Networks with Nonlinear Synaptic Transmission*, *Phys. Rev. Lett.* **108** (2012), no. 15, 1–5, [doi](#).
- [36] Thomas Simon Mountford, *A metastable result for the finite multidimensional contact process*, *Canad. Math. Bull.* **36** (1993), no 2, 216–226.
- [37] Cecilia Romaro, Fernando Najman, and Morgan André, *A Numerical Study of the Time of Extinction in a Class of Systems of Spiking Neurons*, *J. Stat. Phys.* **190** (2023), no. 41, [doi](#).
- [38] Reuven Y. Rubinstein and Dirk P. Kroese, *Simulation and the monte carlo method*, *Wiley Series in Probability and Statistics*, Wiley, 2016, [doi](#).
- [39] Roberto H. Schonmann, *Metastability for the contact process*, *J. Stat. Phys.* **41** (1985), no. 3-4, 445–464.
- [40] A.A. Verveen and H.E. Derksen, *Fluctuation phenomena in nerve membrane*, *Proceedings of the IEEE* **56** (1968), no. 6, 906–916, [doi](#).
- [41] Xiao-Jing Wang, *Synaptic reverberation underlying mnemonic persistent activity.*, *Trends Neurosci.* **24** (2001), no. 8, 455–63, [doi](#).
- [42] Yun Wang, Henry Markram, Philip H. Goodman, Thomas K Berger, Junying Ma, and Patricia S. Goldman-Rakic, *Heterogeneity in the pyramidal network of the medial prefrontal cortex*, *Nat. Neurosci.* **9** (2006), no. 4, 534–542, [doi](#).
- [43] Yosef Yarom and Jorn Hounsgaard, *Voltage fluctuations in neurons: Signal or noise?*, *Physiol. Rev.* **91** (2011), no. 3, 917–929, [doi](#).
- [44] David Zipser, Brandt Kehoe, Gwen Littlewort, and Joaquín M. Fuster, *A spiking network model of short-term active memory*, *J. Neurosci.* **13** (1993), no. 8, 3406–3420, [doi](#).

Acknowledgments. This article was produced as part of the activities of FAPESP Research, Innovation and Dissemination Center for Neuromathematics (grant number 2013/07699-0, S.Paulo Research Foundation, Brazil), and Morgan André was supported by FAPESP, Brazil scholarships (grant number 2020/12708-1). Christophe Pouzat was supported by the ANR project Simbadnesticost (ANR-22-CE45-0027-01). Eva Löcherbach saw that we were dealing with a quasi-stationary distribution and put us on the right track. We thank our two anonymous referees for gently pushing us to refine our analytical and numerical arguments. Finally and very sadly, this article is dedicated to the memory of our friend and colleague Antonio Galves. Antonio was in many ways—through his contributions on metastability and his constant drive for simple (but not too simple) models amenable to analytical treatment—at the origin of the present work.