Event modeling of message interchange in stochastic neural ensembles

Vicenç Gómez[1], Andreas Kaltenbrunner, Vicente López

Abstract— **We propose a modeling framework based on the event-driven paradigm for populations of neurons which interchange messages. Unlike other strategies our approach is focused on the dynamics at the mesoscopic level (spike production and reception) and does not determine the microstates of the neurons. We apply the technique on a discrete model of stochastic ensembles and on extensions of this model to the continuous time domain. Due to the event-driven nature of the method efficient large-scale simulations can be performed without precision errors. The approach uses spike predictions as evidences and a one-step update of the predictions is performed every time an event occurs, resulting in a more efficient solution than the existing strategies.**

I. Introduction

The study of the collective dynamics of networks of interacting spiking neurons is a relevant field in neural computation. It leads to a better understanding of biological information processing and coding and may be applied to build artificial systems inspired in these mechanisms [1], [2].

Some approaches describe the temporal structure of neuronal activity using generic quantitative methods to characterize the statistical properties of the train of action potentials associated with spontaneous activity [3], [4]. These approaches are focused on the mesoscopic level where the activity of a neuron can be regarded as a stochastic point process, because the output signals of most neurons consist of stereotypical electric pulses (spikes). They are useful as a descriptive tool but have a poor physiological foundation to understand the underlying mechanisms at the cell level.

More detailed models consider a microstate variable (the membrane potential) whose temporal evolution can be described by a system of differential equations. For realistic models which account for the inherent stochasticity of the neuronal activity the evolution of this variable can be characterized as a stochastic process. Several theoretical approaches have been developed under this framework, from simple models which consider the microstate as a random walk (Gernstein Mandelbrot Model [3]) to models with much higher complexity such as the Hodking-Huxley Model [5] where four dynamical variables define the microstate, or compartmental models [6] where more than 20 variables are used.

Probably the best compromise between tractability and realism is the arquetypical leaky integrate-and-fire neuron model [7], [1], where the neuron is defined as a leaky integrator that fires if the microstate reaches a threshold. If the contribution of each received impulse to the microstate variable is small and they are many, the microstate evolution can be approximated by a diffusion process (diffusion approximation). The resulting continuous stochastic process is the Ornstein-Uhlenbeck process [8]. Using the Fokker-Planck formalism significant theoretical results have been obtained from a perspective of mean field analysis [9],[10]. This model however is difficult to solve analytically in many situations, since the inter-spike-interval (ISI) distribution for a given input is mathematically equivalent to a first-passage-time problem, which is known to be intractable in the general case.

For the success of those theoretical approaches extensive and efficient numerical simulations are needed to explore the parameter space. Recently, the event-driven paradigm has been introduced as a convenient way to optimize them [11], [12], [13], [14]. Assuming that spikes occur sparsely in time and have pulsed nature, the processing time between two successive events can be avoided under some realistic cortical conditions, namely, small connectivity and low firing rates. In [15] the algorithm was extended to include stochastic dynamics using deterministic linear integrate-and-fire neurons under external background noisy activity using the Fokker-Planck equation to obtain numerically the ISI density function. This strategy overcomes the event saturation problem of a previous approach [12] at the cost of large lookup tables where the ISI density functions are stored after a huge precalculation task.

In the present work, we propose a general event-driven modeling technique based on the spike production and reception which can be applied to models with stochastic dynamics and which decreases the computational demand of the previous stochastic event-driven approach of [15]. Instead of starting with a stochastic differential equation governing the microstate of the neurons we consider each neuron within the population as a stochastic oscillator emitting spikes spontaneously according to some ISI density function which depends on the model under consideration. The isolated activity can be interpreted as an external irregular background activity or also as an intrinsic source of noise. We introduce the general method in section II and apply it to a microscopic model. We first consider discrete-time modeling in section III where the ISI probability distribution of the spontaneous evolution corresponds to a negative binomial distribution and in section IV we introduce two possible extensions of this model to the continuous time domain where the ISI density function corresponds to a gamma density function. In section V we show simulation results which compare the efficiency of the mesoscopic and microscopic simulations and show the

[1]Departament de Tecnologia, Universitat Pompeu Fabra, Passeig de Circumval·lacio 8, 08033 Barcelona, Spain (email: vgomez@iua.upf.es).

accuracy of the extensions to the continuous time domain and the original discrete model. We give some conclusions about the applicability to other models at the end of the paper.

II. Mesoscopic, event-driven modeling

We consider a generic model composed of N coupled spiking units which evolve spontaneously in time. The ensemble can be considered as a directed graph with arbitrary connectivity in which nodes emit and receive messages through the connections. According to the standard event-driven framework [14], to rewrite a particular model of this type four components have to be defined: the state variable a_i, which usually represents the membrane potential of neuron i, two functions, r_i and s_i, which describe how a_i changes after the reception and emission of a message respectively, and the function $\tilde{t}_i$, which gives the predicted time of the next firing, given the present state variable, assuming that no messages will affect i until then. Provided an instance of these four components $\{a_i, r_i, s_i, \tilde{t}_i\}$, an event-driven engine drives the simulation without the necessity to integrate the microstate variable in the interval between two relevant events. The engine basically performs four steps:

1) Select the next event to be processed
2) Process the event
3) Schedule all possible new events generated in the previous step and
4) Return to the first step or end the simulation.

Since messages can be delayed after their emission, two types of events have to be distinguished: the *firing* events and the *reception* events. Step 1) selects the event with minimum time label and Step 2) is essentially a disjunction which evaluates if the selected event is a firing or a reception event, and applies r_i or s_i respectively. Once the value of the microstate a_i is determined, the function $\tilde{t}_i$ is used to predict the next firing time. The previous prediction is then discarded and not used anymore.

Contrary to the classical approach we adopt a different strategy: we define the state of a neuron i as its next predicted firing time, or belief b_i which is reset to an initial value according to a predefined initial ISI probability distribution every time the neuron i fires. If a relevant event occurs, instead of discarding the prediction and determining the value of the microstate a_i, we use b_i as an evidence to obtain its new value b_i' using an inference procedure. Every time a neuron receives an impulse from other unit we update the belief according to Bayes' rule:

$$P(b_i'|b_i) = \frac{P(b_i|b_i')P(b_i')}{P(b_i)} \ . \tag{1}$$

As a consequence, to update the belief of a unit we need to define two probability distributions. An *unconditional* ISI probability distribution $P(b_i)$, used when a unit fires and the next spiking time is independent of previous predictions, and a *conditional* ISI probability distribution $P(b_i'|b_i)$, used when an event affects a unit before its predicted next firing time. A specific model of spiking neurons can be defined using this technique, given that the two following mechanisms are provided: the two probability distributions or *belief updating rules*, which describe how the mesoscopic state b_i changes when i fires or receives a message, and the scheduler of the reception events.

We emphasize the main difference between the approach taken in [15] and our technique. In their approach every time a relevant event occurs the value of the prediction is discarded and two pseudorandom numbers are extracted: the first is required for fixing the microstate a_i, and the second for determining the next possible firing time $\tilde{t}_i$. Our approach does not fix the microstate and uses only one pseudorandom number extraction, resulting in a more efficient approach. The main caveat concerns the implementation of the ISI probability distributions. For the models considered here no lookup tables are required to simulate large networks.

In our strategy the statistical description of the microstate a_i is not retained. The sequence of firing times $\langle T_k^i \rangle$, or equivalently the ISIs are the relevant magnitudes obtained from the simulations. This is not a limitation, since it is widely assumed that the dynamics at the mesoscopic level (spike trains) concern the main aspect of neural information processing [2], whatever the neural coding mechanism adopted (firing rates or exact timing). Moreover, our technique is not limited if dynamics at the level of the synapse is included, since Spike Time Dependent Plasticity (STDP) [16] only considers the spike timing.

To illustrate the previous technique we start with a simple microscopic model and derive the mesoscopic reformulation which reproduces the same dynamics in the temporal domain.

III. An example: Discrete and stochastic ensembles

A. The microscopic model

In this model the microstate of a neuron evolves as a random walk with positive drift towards an absorbing barrier [17], [18]. Neurons are stochastic non-leaky integrate-and-fire and the microstate of neuron i is represented by the activation level a_i. This variable grows at discrete timesteps until a threshold L is reached. When this happens, unit i emits a message to the rest of the units, a_i is reset to the initial value, and the unit enters a refractory period t^{refr} where it remains insensitive to incoming spikes. We set $t^{\text{refr}} = 1$. Formally, the growth of a_i at time $t+1$ is modeled by the following rule, which from now on will be referred to as the spontaneous evolution:

$$\hat{a}_i^{t+1} = \begin{cases} a_i^t + 1 & \text{with probability } p \\ a_i^t & \text{with probability } (1-p) \end{cases} \quad \text{if } a_i^t < L \ ,$$

$$\hat{a}_i^{t+1} = 1 \text{ with probability } 1 \quad \text{if } a_i^t \geq L \ .$$

Within the same timestep, the strength of the possibly received messages from units which fired at the previous timestep t is integrated, $t^{\text{delay}} = 1$, so the activation state at

timestep $t+1$ is:

$$a_i^{t+1} = \hat{a}_i^{t+1} + \epsilon \sum_{j \neq i} H_L(a_i^t)$$

$$\text{where } H_L(x) = \begin{cases} 1 & \text{if } x \geq L \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

ϵ represents a homogeneous synaptic efficacy. Thus, the total evolution of a_i in one timestep is composed of two terms: the stochastic spontaneous evolution and the activity induced by incoming messages of the population.

A characteristic parameter used to describe the degree of interaction between the units is:

$$\eta = \frac{L-1}{(N-1)\epsilon} \tag{2}$$

where ϵ indicates a global coupling term. Parameter η indicates whether the spontaneous dynamics or the message interchange mechanism dominate the behavior of the system. For $\eta \gg 1$, where interactions are small, the system is basically driven by the stochastic, spontaneous activity, and all units behave as independent oscillators. This weak coupling regime is mostly referred as the "Noise-dominated", or spontaneous activity regime, which is recognized as a representative of typical cortical conditions. For values of η near 1 the behavior of the system is dominated by the pulse-exchange mechanism and synchronous behavior starts to emerge. This corresponds to the "driven" activity regime or also called "Drift-dominated" regime.

B. Mesoscopic reformulation

We provide a reformulation of the previous model in terms of the framework presented in section II which reproduces identically the dynamics at the mesoscopic level. We replace the microstate a_i by $m_i = \langle b_i, L_i \rangle$, where b_i is the belief and L_i summarizes all the mesoscopic activity since the last firing time T_i^e. L_i can be viewed as an effective threshold which decreases deterministically every time an incoming impulse is received. It can be calculated from the temporal sequences of spikes $\langle T_k^j \rangle$ but we propagate it as a component of the mesoscopic state m_i for computational efficiency. We now describe how the reception events are handled and which are the rules of belief updating.

The model is defined with global connectivity and homogeneous coupling so each generated message affects equally all the units. Thus N iterations are sufficient to process one event even in the worst case that all units spike in unison. Moreover, since we set delay t^{delay} and refractory period t^{refr} homogeneously and equal to one, all units can be updated consistently at the firing time, avoiding the necessity of a priority queue in the implementation. Rather than being a limitation, this allows us to focus our analysis on the belief updating mechanism instead of the implementation engine. In the continuous time model the priority queue will be required. We now specify the two ISI probability distributions and the model will be fully redefined.

1) Unconditional ISI probability distribution: This probability distribution is used when a unit fires and thus the next spiking time is independent of previous predictions. Consider a unit i which fires at time t, possibly simultaneously with n other units. During the transition from t to $t+1$, i is in its refractory period, insensitive to incoming events. At $t+1$, i receives an initial message of strength $\phi = n\epsilon$ from those units which fired simultaneously with i. A new value of b_i is drawn assuming that evolution is only due to the spontaneous activity. Since the path towards the absorbing barrier L_i which an isolated unit takes can be interpreted as a random walk specified by a Bernoulli process with probability p, the new value of b_i is obtained according to the negative binomial distribution [19]:

$$P_{NB}(b_i') = \binom{b_i - 1}{L_i - 2} p^{L_i - 1} (1-p)^{b_i - L_i + 1} \qquad \text{for } b_i = 1, 2, \ldots \tag{3}$$

where for clarity we have taken temporal values relative to $T_i^e + t^{\text{refr}}$, the last firing time of i plus the refractory period. Note that the quantity $(b_i - L_i + 1)$ is the number of failures, which in our case indicates the number of timesteps that i will delay its firing from $(T_i^e + t^{\text{refr}} + L_i - 1)$. This probability distribution has an infinite support. Thus, the microscopic state has a nonzero probability of crossing the threshold L_i at any timestep greater than $L_i - 1$ in the future. The first two moments of this probability distribution are:

$$\mu_{NB} = \frac{L_i - 1}{p} \qquad \sigma_{NB}^2 = \frac{(L_i - 1)(1-p)}{p^2} \tag{4}$$

2) Conditional ISI probability distribution: In the case of a reception event, the belief b_i of unit i, which was last updated at timestep t, is affected by a subsequent reception event of strength ϕ at timestep t'. Omitting time indices for clarity, we are interested in the rule which describes the update from $m_i = \langle b_i, L_i \rangle$ to $m_i' = \langle b_i', L_i' \rangle$.

The strength of the message ϕ is the number of units which fired at timestep $t' - 1$ multiplied by ϵ. We set $L_i' = L_i - \phi$ and apply Bayes' rule to obtain the new value b_i' using the previous prediction b_i as evidence. For the model under study the three probability distributions involved in the rule are negative binomial distributions. Figure 1 illustrates the procedure. The light grey region covering all possible trajectories of the microscopic state a_i within the interval $[T_i^e + t^{\text{refr}}, b_i]$ can be decomposed in two disjoint sets of possible realizations of a_i. The dark gray area contains the trajectories where unit i fires at time b_i' with effective threshold L_i', and the black one contains the trajectories where a_i starts at L_i' and reaches L_i at time b_i. We are now interested in the probability distribution of the waiting time to reach L_i' state transitions (successes), where the total number of successes L_i and failures $(b_i - L_i + 1)$ are known from the previous prediction. Using the ISI probability distribution of (3) in Bayes' rule we get the negative hypergeometric distribution [20], which is the desired belief update rule (see

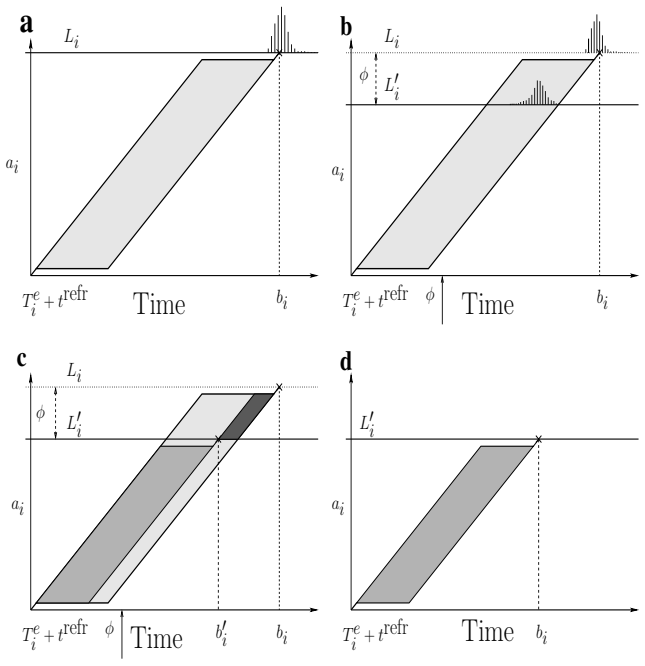

Fig. 1. Illustration of the process involved in the belief updating of a given unit i. Horizontal axis denotes time and vertical axis the microstate variable. (a) After i has fired, the belief is obtained according to the unconditional ISI probability distribution $P(b_i)$. (b) When subsequent events affect i a conditional ISI probability distribution $P(b'_i|b_i)$ is defined. (c) The previous prediction (first packet of trajectories) can be decomposed in two disjoint regions ($P(b'_i)$, $P(b_i|b'_i)$, see text for details). (d) The belief updating is repeated with the new value b'_i.

appendix-A for details) :

$$P_{NH}(b'_i|b_i) = \frac{\binom{b_i-b'_i-1}{L_i-L'_i-1}\binom{b'_i-1}{L'_i-2}}{\binom{b_i-1}{L_i-2}} \tag{5}$$

for $b'_i = L'_i - 1, \ldots, b_i - \phi$. This probability distribution, unlike the one of Eq.(3) has finite support, showing that the number of possible firing times are constrained from previous predictions. The first two moments of this distribution applied to this particular model are [20]:

$$\mu_{NH} = \frac{b_i(L'_i-1)}{L_i-1} \ , \quad \sigma^2_{NH} = \frac{b_i(b_i+L_i-1)(L'_i-1)\phi}{L_i(L_i-1)^2} \ . \tag{6}$$

Note that once the first prediction has been made according to Eq.(3), the rate p, which represents the intrinsic rate of noise, or the external background stationary activity, is no longer relevant in subsequent updates, where only a number counting method is applied until i fires again.

IV. Extension to Continuous time domain models

We extend the previous discrete model by replacing the Bernoulli process governing the evolution of the microstate variable by its continuous counterpart, the Poisson process. The resulting model corresponds to the Poisson excitation model [8], where excitatory inputs occur at random according to a simple Poisson process and cause the microstate to increase. As in the discrete case, this stochastic evolution can also be interpreted as an intrinsic source of noise.

We analyze two possible continuous extensions: the Poisson process of rate p and a version obtained by equating the moments of the gamma density function to those of the negative binomial distribution (method of moments).

To accommodate the dynamics of the discrete case, we choose the same values for the transmission delay and the refractory period, i.e. $t^{\text{delay}} = t^{\text{refr}} = 1$. The fact that events may occur at non-integer times prevents a synchronous update of all the beliefs at firing times as in the discrete case. Therefore, a priority queue where firing and reception events coexist is required as explained in the general framework of section II.

1) Continuous version I: Poisson process approximation: We take the limit in which the size of a timestep Δt tends to zero but the number of state transitions to reach the threshold L remains constant. In this case the probability p of a state transition tends to zero and the amount of timesteps before firing tends to infinity. Taking again values of the belief relative to $T^e_i + t^{\text{refr}}$ we obtain in the limit of $\Delta t \to 0$ the following gamma distribution: [1]

$$p_\Gamma(b_i) = \frac{\beta^\alpha}{\Gamma(\alpha)} b_i^{\alpha-1} e^{-\beta t} \qquad \text{for } b_i \geq 0 \ . \tag{7}$$

where $\alpha = L-1$ and $\beta = p$. This ISI probability distribution can be understood as the sum of $L-1$ exponentially distributed variables with $\lambda = p$. Each of these variables represents the waiting time between two state transitions, which are now Poisson events with rate p. The first two moments of the gamma density applied to the model are:

$$\mu_\Gamma = \frac{L-1}{p} \ , \qquad \sigma^2_\Gamma = \frac{L-1}{p^2} \ . \tag{8}$$

Note that the mean μ_Γ is the same as μ_{NB} of the discrete case, Eq. (4), but the variance is overestimated, and only coincides with the discrete case in the limit of $p \to 0$.

Following a similar procedure as in section II, we apply Bayes' rule to obtain the conditional ISI density function. An equivalent picture as Figure 1 can be used to illustrate the different terms involved in the rule. The resulting density function (see appendix-B) is a rescaled beta distribution, the analogous counterpart of the negative hypergeometric distribution in the continuous time domain:

$$p_\beta(b'_i|b_i) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(1 - \frac{b'_i}{b_i}\right)^{\beta-1} \left(\frac{b'_i}{b_i}\right)^{\alpha-1} \frac{1}{b_i} \ , \quad \text{for } b'_i \in [0, \ldots, b_i] \tag{9}$$

with parameters $\alpha = L'_i - 1$ and $\beta = \phi$. Again, the parameter p is no longer relevant in subsequent updates of the belief. Note also that the previous value of the belief b_i defines the bounded support of the beta distribution, so the new belief is constrained by the previous prediction. The first two moments of this distribution applied to the model are:

$$\mu_\beta = b_i \frac{L'_i - 1}{L_i - 1} \ , \qquad \sigma^2_\beta = b_i^2 \frac{(L'_i-1)\phi}{(L_i-1)^2 L_i} \ . \tag{10}$$

[1] Note that unlike the discrete case, beliefs are not restricted to have values greater than the threshold L. Units in the continuous time domain can advance and fire before an interval of time L.

TABLE I

MOMENTS OF THE RELATED ISI PROBABILITY DISTRIBUTIONS.

Discrete (waiting time)	Continuous (Poisson)
$\mu_{NB} = \frac{L_i-1}{p}$ $\sigma^2_{NB} = \frac{(L_i-1)(1-p)}{p^2}$	$\mu_{\Gamma} = \frac{L-1}{p}$ $\sigma^2_{\Gamma} = \frac{L-1}{p^2}$
$\mu_{NH} = \frac{b_i(L'_i-1)}{L_i-1}$ $\sigma^2_{NH} = \frac{b_i(b_i+L_i-1)(L'_i-1)\phi}{L_i(L_i-1)^2}$	$\mu_{\beta} = \frac{b_i(L'_i-1)}{L_i-1}$ $\sigma^2_{\beta} = \frac{b_i^2(L'_i-1)\phi}{L_i(L_i-1)^2}$
Discrete (failures)	**Continuous (Fitting moments)**
$\mu_{NB} = \frac{(L_i-1)(1-p)}{p}$ $\sigma^2_{NB} = \frac{(L_i-1)(1-p)}{p^2}$	$\mu_{\Gamma'} = \frac{(L_i-1)(1-p)}{p}$ $\sigma^2_{\Gamma'} = \frac{(L_i-1)(1-p)}{p^2}$
$\mu_{NH} = \frac{(L_i-1)(1-p)}{p}$ $\sigma^2_{NH} = \frac{b_i(b_i+L_i-1)(L'_i-1)\phi}{L_i(L_i-1)^2}$	$\mu_{\beta'} = \frac{(L_i-1)(1-p)}{p}$ $\sigma^2_{\beta'} = \frac{b_i^2(L'_i-1)\phi}{((1-p)(L_i-1)+1)(L_i-1)^2}$

We compare them with the moments of the negative hypergeometric distribution in the discrete case, see Eq. (6). The means coincide, but the variance is overestimated by a factor $(b_i + L_i - 1)$ instead of b_i.

2) Continuous version II: method of moments: Other extensions can be obtained depending on how the continuous density functions are chosen, or which limiting process is taken. For example, we get another continuous model if we consider the value of the beliefs relative to $T_i^e + L_i$, instead of T_i^e. This situation is equivalent to considering in the discrete case the belief b_i as the number of *failures*, or delayed timesteps from $T_i^e + L_i$, instead of the waiting time from T_i^e. The discrete dynamics is not modified by this new counting mechanism (the negative binomial is just shifted), but in the continuous case the beliefs are prevented from taking values smaller than the number of required state transitions to reach the threshold. Then, with the following suitable choice of the gamma density parameters: $\alpha = (L_i - 1)(1 - p)$ and $\beta = p$, both continuous and discrete probability distributions will have the same first two moments.

$$\mu_{\Gamma'} = \frac{(L_i - 1)(1 - p)}{p} \qquad \sigma^2_{\Gamma'} = \frac{(L_i - 1)(1 - p)}{p^2} \ . \quad (11)$$

Again, applying Bayes' rule leads to another rescaled beta distribution, but now with parameters $\alpha = (L'_i - 1)(1-p)$ and $\beta = \phi(1-p)$ and rescaled in the time interval $[L_i - \phi, b_i - \phi]$, as the discrete version. The moments of this distribution are:

$$\mu_{\beta'} = b_i \frac{L'_i - 1}{L_i - 1} \quad \sigma^2_{\beta'} = \frac{b_i^2 (L'_i - 1)\phi}{(L_i - 1)^2 \left((1-p)(L_i - 1) + 1\right)} \ . \quad (12)$$

As before there is only one difference with the discrete case in the standard deviation, with the factor $b_i + L_i - 1$ instead b_i, and the factor $((1 - p)(L_i - 1) + 1)$ instead of L_i which coincide when $p \to 0$. To facilitate the comparison between the three different versions of the model we group in Table I all the moments of the used probability distributions and density functions. The negative binomial (NB) and negative hypergeometric (NH) distributions are used in the discrete model (left side of the table). Regardless of whether the waiting time (which is equivalent to take the belief with respect to $T_i^e + t^{\text{refr}}$) or the number of failures (which is equivalent to takeing the belief with respect to $T_i^e + t^{\text{refr}} + L - 1$) is considered as the belief variable, both use the same parameters. In the continuous version (right side of the table), the Poisson approximation uses a gamma density function with parameters $\alpha = (L - 1)$ and $\beta = p$ and a beta density function with parameters $\alpha = L'_i - 1$ and $\beta = \phi$, whereas the approximation which fits the moments uses a gamma density function with parameters $\alpha = (L - 1)(1 - p)$ and $\beta = p$ and a beta density function with parameters $\alpha = (L'_i - 1)(1 - p)$ and $\beta = \phi(1 - p)$.

V. SIMULATION RESULTS

We performed simulations to evaluate the proposed technique. First, we compare the computational complexity of the mesoscopic version of the model against the original microscopic model and then we compare the accuracy of the continuous versions with the discrete one.

Since the model does not assume any particular time units, we quantify the performance of a given model in terms of the number of updates (in cpu-time) within an ISI of length τ. This measure allows us to easily compare both models. Figure 2 shows this quantity for different values of the characteristic parameter η which indicates the degree of interaction between the units of the ensemble. Higher values

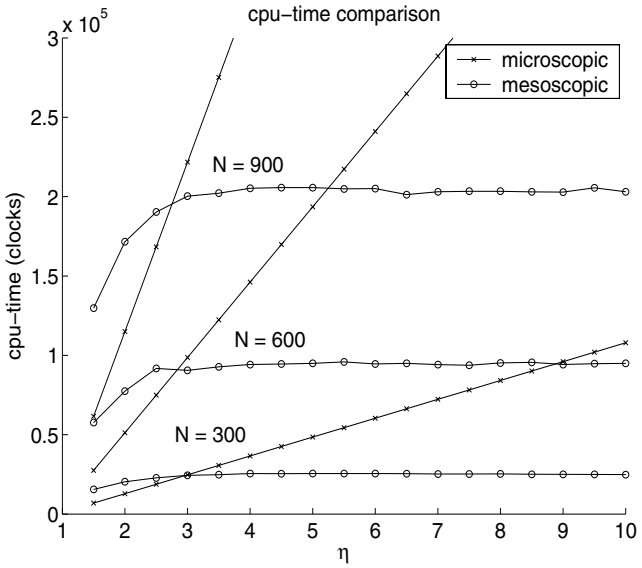

Fig. 2. Comparison of the computational cost between the original microscopic model and the mesoscopic event-based reformulation in function of the degree of interaction for different sizes of the ensemble. We use a full-connected network and the noise rate is $p = 0.9$. No lookup tables were computed a priori. For $\eta > \eta^* \sim 2.5$ the mesoscopic version is more efficient.

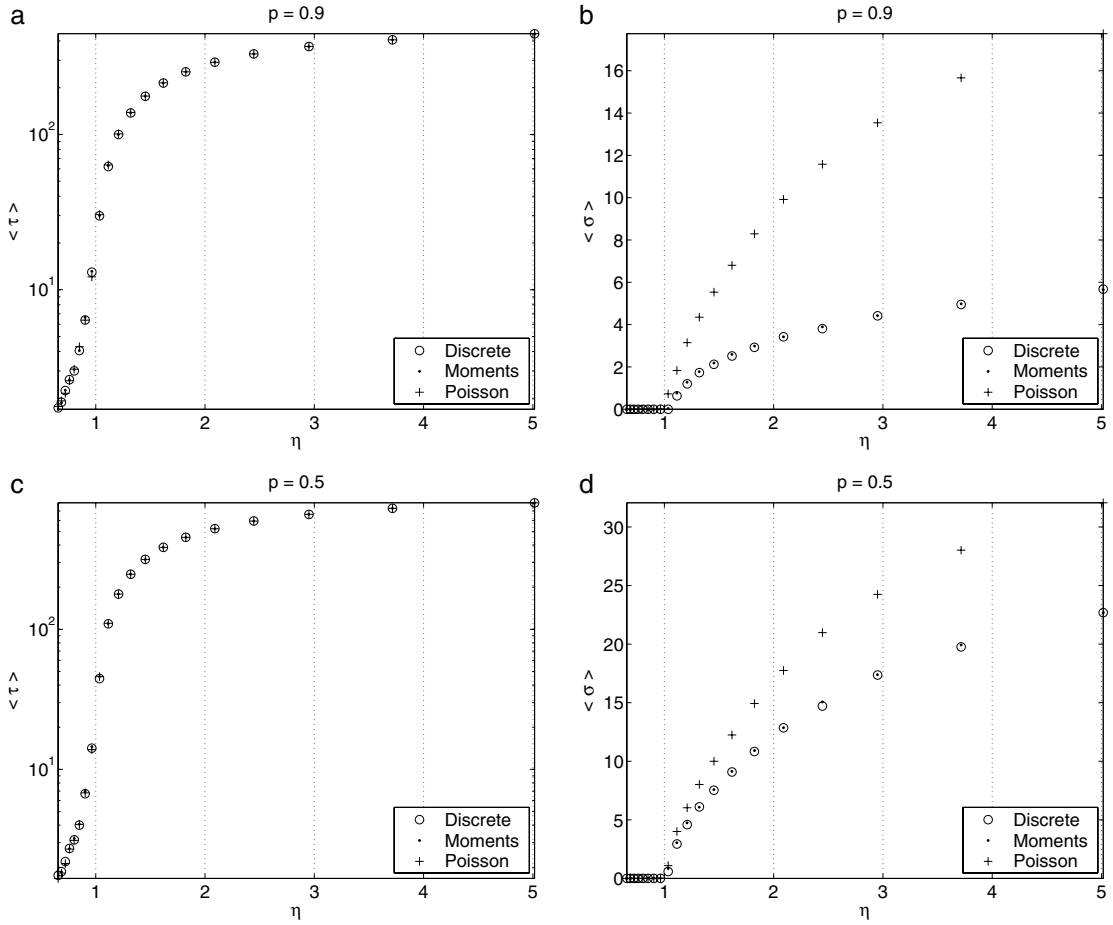

Fig. 3. Simulation results of the event-driven model for the discrete and two possible continuous extensions. The extension which considers the spontaneous evolution as a Poisson process with rate p is indicated by **Poisson** whereas **Moments** indicates the model which uses the method of moments (*see the text*). Upper figures show **(a)** the mean of the ISI in logarithmic scale and **(b)** the respective standard deviation over 500 simulations with $p = 0.9$. Lower figures **(c, d)** show the same statistics but for a noise rate of $p = 0.5$. Simulations were performed over an ensemble of size $N = 200$ and threshold $L = 500$. The coupling strength ϵ was varied around the critical value of $\eta = 1$. Both continuous models fit well the mean ISI. The standard deviation is well fitted by the **Moments** version while the **Poisson** version significantly overestimates it for higher values of the rate p.

of η correspond to small interaction between the units, and vice versa. Clearly, the mesoscopic event-based version of the model is suitable when spikes are rare events in time, since timesteps where neurons do not fire are not simulated. It turns out that for regimes of $\eta < 1$ the event-based version does not perform better than the microscopic version, since spikes occur at every timestep [18] and hence, the number of updates is the same in both approaches.

The computational complexity of the original microscopic model is proportional to the ISI $\mathcal{O}(\tau N)$, since in an ISI of size τ all N states need to be updated at each time step. The event-based strategy, however, requires the update of N predictions every time an event is produced because we are using a full connected network. The homogeneous coupling allows to update a unit in constant time. Thus an average of N^2 updates will be required within an ISI leading to an $\mathcal{O}(N^2)$ computational complexity. Without the requirement of simulating the evolution of the microstate variable, the cost does not depend on the size of the ISI τ unlike the microscopic approach. As can be observed in Fig.(2) the number of updates of the mesoscopic approach is constant and the event-driven approach performs better than the original approach the higher the value of η (the lower the coupling). A decrease is observed for $\eta \sim 2$ and below, which is caused because in this regime simultaneous spikes occur often and can be grouped together into one single impulse. The time independence is a common feature of all event-driven approaches.

Figure 3 shows simulation results of the discrete and both continuous versions of the model, where L and N are fixed and the coupling strength is varied to cover a significant parameter space of η. We simulate the model starting at a regime of small values of ϵ (weak coupling) and progressively increase it until all units emit spikes synchronized in unison. As previously explained, the event-based strategy allows the model to be simulated exactly, without precision errors. The mean ISI $\langle\tau\rangle$ and the dispersion $\langle\sigma\rangle$ for a unit are averaged over all the ensemble and over 500 initial conditions. We report exact agreement between both continuous and discrete versions as expected from the

analytical results (not shown in the figure for clarity) and also the mean of the ISIs is reproduced with accuracy in the continuous versions. As for the standard deviation, the approach based on the method of moments reproduces much better the discrete dynamics than the Poisson approximation, which is more similar if p is low, as theory suggests (note the differences between the upper-right plot and the lower-right plot). Figures 3.b and 3.d show a rapid decrease in the standard deviation which vanishes at $\eta = 1$, where the coupling strength reaches a critical value [21].

VI. Conclusions

The presented framework is useful and convenient for modeling large-scale neural populations when the detailed evolution of the microscopic variables is not the main aspect of analysis and efficient simulations are needed.

In general two critical points arise regarding the computational efficiency of all event-based approaches. On the one hand, the necessity of an ordered data structure for storing the events for models with heterogeneous t^{delay} and t^{refr}. On the other hand, for a general case with heterogeneous efficacies, and high connectivity the cost of processing one event is $\mathcal{O}(N^2)$. These are inherent limitations which can only be alleviated using parallel and distributive computing [22] for which our approach can also be applied.

Despite the simplicity of the model presented, the related ISI probability distributions have been successfully used many times for modeling neural data. See, for example [23], [24] where the gamma density function is used. This is not surprising, since empirical ISI densities under stationary background activity often have three basic forms [8]: for very large excitatory impulses, neurons fire at the reception of one spike and the ISI density is exponential-like; if few impulses occurring in a short enough interval are sufficient to trigger a spike, the ISI density is the gamma density analyzed here. Finally, if the threshold is large compared to the average strength of the impulses, the ISI density has gaussian appearance. We notice the great similarity between the first-passage-time density function of the Ornstein-Uhlenbeck Model and the gamma density function except for the long tail of the first case.

Indeed, the presented technique is not restricted to the particular model analyzed here. Inhibition in the form of negative pulses can easily be incorporated given that no reflecting barrier is imposed at the zero value of the microstate as in the case of Stein's model [25]. For example, the Wiener process (Brownian motion) with positive drift governing the microstate evolution can also be simulated using our approach. In that case, the *inverse gaussian* density function would be the unconditional ISI density function to be used. In general the method is applicable provided the ISI density function has a closed analytical form and when the underlying evolution of the microstate is such that the update is independent of the time at which an event affects the neuron. It is worth emphasizing that any relevant magnitude which can be obtained from the spike history can be captured by the presented approach, given that the required computations can be performed in an event-driven basis. An example in this direction would be to incorporate the dynamics of slow variables whose time constants may span several ISIs. This issue is outside of the scope of this paper and considered as future work.

The approach also has potential applications in the context of pulse-coupled oscillators when a linear (or a piecewise linear) evolution is considered in the phase variable. For example, it can be used for improving the numerical method used in deterministic models [26] as well as in the linear pulse-coupled oscillator where noise is applied only every time an event occurs [27].

Summarizing, under simplifying assumptions on the equation governing the microstate variable the high demand placed on computational resources for simulating stochastic populations of neurons can be alleviated using the proposed technique.

Appendix

A. Bayes' rule in the discrete model

We show how the negative hypergeometric probability distribution (5) is obtained as the posterior probability of the Bayes' Rule: $P_{NH}(b'_i|b_i) = P(b_i|b'_i)P(b'_i)/P(b_i)$ Writing the numerator in terms of the negative binomial distributions:

$$P_{NB}(b_i|b'_i)P_{NB}(b'_i) = \tbinom{b_i-b'_i-1}{L_i-L'_i-1}\tbinom{b'_i-1}{L'_i-2}p^{L_i-1}q^{b_i-L_i+1} \tag{13}$$

The denominator can be obtained summing over all possible values of b'_i:

$$\begin{aligned} P_{NB}(b_i) &= \textstyle\sum_{b'_i=L'_i-1}^{b_i-(L_i-L'_i)} \tbinom{b_i-b'_i-1}{L_i-L'_i-1}\tbinom{b'_i-1}{L'_i-2}p^{L_i-1}q^{b_i-L_i+1} \\ &= \textstyle\sum_{b'_i=0}^{b_i-L_i+1} \tbinom{b_i-b'_i-L'_i}{L_i-L'_i-1}\tbinom{b'_i+L'_i-2}{L'_i-2}p^{L_i-1}q^{b_i-L_i+1} \end{aligned}$$

Decomposing the sum:

$$\begin{aligned} P_{NB}(b_i) = &\Big(\textstyle\sum_{b'_i=0}^{b_i-L'_i} \tbinom{b_i-b'_i-L'_i}{L_i-L'_i-1}\tbinom{b'_i+L'_i-2}{L'_i-2} \\ &- \textstyle\sum_{b'_i=b_i-L'_i+1}^{b_i-L'_i} \tbinom{b_i-b'_i-L'_i}{L_i-L'_i-1}\tbinom{b'_i+L'_i-2}{L'_i-2}\Big) p^{L_i-1}q^{b_i-L_i+1} \end{aligned}$$

The second term of the sum is always zero, since $b_i - (b_i - L'_i + 1) - L'_i < L_i - L'_i - 1$ and $L_i - L'_i = \phi > 0$. Applying the following equality [19]:

$$\sum_{k=0}^{r} \binom{r-k}{m}\binom{s+k}{n} = \binom{r+s+1}{m+n+1}$$

We get:

$$\begin{aligned} P_{NB}(b_i) &= \binom{b_i - L'_i + L'_i - 2 + 1}{L_i - L'_i - 1 + L'_i - 2 + 1} p^{L_i-1}q^{b_i-L_i+1} \\ &= \binom{b_i - 1}{L_i - 2} p^{L_i-1}q^{b_i-L_i+1} \end{aligned} \tag{14}$$

Using (13) and (14) we obtain the negative hypergeometric distribution (5):

$$P_{NH}(b'_i|b_i) = \frac{\tbinom{b_i-b'_i-1}{L_i-L'_i-1}\tbinom{b'_i-1}{L'_i-2}}{\tbinom{b_i-1}{L_i-2}}$$

for b'_i in the interval of values restricted by the previous belief $[L'_i - 1, b_i - \phi]$.

B. Bayes' rule in the continuous model

We derive the beta density function (9) using Bayes' rule in the continuous model:

$$g'(b'_i|b_i) = \frac{g(b_i|b'_i)h'(b'_i)}{h(b_i)} \ . \tag{15}$$

In this case $h'(b'_i)$ is the density function of the waiting time until the threshold lowered by the incoming messages $L'_i = L_i - \phi$ is reached. We use the gamma density function with $\alpha = L'_i - 1$ and $\gamma = p$:

$$h'(b'_i) = \frac{\gamma^\alpha}{\Gamma(\alpha)} {b'_i}^{\alpha-1} e^{-\gamma b'_i} \qquad \text{for } b'_i \geq 0 \tag{16}$$

Similarly $g(b_i|b'_i)$ is the density function of the waiting time until an effective threshold $(L_i - 1)$ is reached at time b_i, assuming that unit i starts with activation state $(L'_i - 1)$ at time b'_i. It represents another gamma density, in this case with $\beta = \phi$ and $\gamma = p$:

$$g(b_i|b'_i) = \frac{\gamma^\beta}{\Gamma(\beta)} (b_i - b'_i)^{\beta-1} e^{-\gamma(b_i - b'_i)} \quad \text{for } (b_i - b'_i) \geq 0 \tag{17}$$

The normalization factor is $h(b_i)$, which is the marginal probability of b_i:

$$h(b_i) = \int_{-\infty}^{\infty} \frac{\gamma^\beta}{\Gamma(\beta)} (b_i - b'_i)^{\beta-1} e^{-\gamma(b_i - b'_i)} \frac{\gamma^\alpha}{\Gamma(\alpha)} {b'_i}^{\alpha-1} e^{-\gamma b'_i} db'_i$$

Using the substitution $t = b'_i/b_i$ and eliminating the regions where the integral is zero we get:

$$h(b_i) = \frac{\gamma^{\alpha+\beta} e^{-\gamma b_i} {b_i}^{\alpha+\beta-1}}{\Gamma(\alpha+\beta)} \tag{18}$$

We use Eqs. (16), (17) and (18) to find the density function:

$$g'(b'_i|b_i) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(1 - \frac{b'_i}{b_i}\right)^{\beta-1} \left(\frac{b'_i}{b_i}\right)^{\alpha-1} \frac{1}{b_i}$$

for $b'_i \in [0, b_i]$, which is the beta distribution with parameters α and β rescaled by the inverse of the previous value of the belief, $1/b_i$.

In the case of the approximation using the method of the moments we can proceed in the same way and get the same probability density function, but with parameters $\alpha = (L'_i - 1)(1 - p)$ and $\beta = \phi(1 - p)$ and rescaled to be within the time interval $b'_i \in [L'_i - 1, b_i - \phi]$, as the discrete version.

ACKNOWLEDGMENT

This work has been supported by Catedra Telefonica UPF de Produccion Multimedia and CiCYT grant number TIN2004-04363-C03-01 from the Spanish Ministry.

REFERENCES

[1] Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, Cambridge, MA, 2002.

[2] F. Rieke, D. Warland, R. de R. van Steveninck, and W. Bialek. *Spikes: Exploring The Neural Code*. The MIT Press, London, England, 1997.

[3] G. L. Gerstein and B. Mandelbrot. Random walk models for the spike activity of a single neuron. *Biophys J.*, 4:41–68, 1964.

[4] D. H. Johnson. Point process models of single-neuron discharges. *Journal Of Computational Neuroscience*, 3:275–299, 1996.

[5] A. L. Hodgkin and A. F. Huxley. Aquantative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 447:235–256, 1952.

[6] R. D. Traub, R. K. S. Wong, R. Miles, and H. Michelson. A model of a ca3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *Journal Of Neurophysiology*, 66:635–650, 1991.

[7] Lapicque L. Recherches quantitatives sur l'excitation electrique des nerfs traitee comme une polarisation. *J. physiol.*, 9:620–635, 1907.

[8] H. C. Tuckwell. *Introduction to Theoretical Neurobiology. Volume 2: Nonlinear and Stochastic Theories*. Cambridge University Press, Cambridge, 1988.

[9] D. J. Amit and N. Brunel. Dynamics of a recurrent network of spiking neurons before and following learning. *Network-Computation In Neural Systems*, 8:373–404, 1997.

[10] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal Of Computational Neuroscience*, 8:183–208, 2000.

[11] Loyd Watts. Event-driven simulation of networks of spiking neurons. *Advances in neural information processing systems*, 6:927–934, 1994. San Mateo, CA. Morgan Kauffman Publishers.

[12] Maurizio Mattia and Paolo Del Giudice. Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. *Neural Comput*, 12(10):2305–2329, 2000.

[13] Takaki Makino. A discrete-event neural network simulator for general neuron models. *Neural Comput & Applic*, 11:210–223, 2003.

[14] Olivier Rochel and Dominique Martinez. An event-driven framework for the simulation of networks of spiking neurons. *ESANN'2003 proceedings*, pages 295–300, 2003.

[15] Jan Reutimann, Michele Giugliano, and Stefano Fusi. Event-driven simulation of spiking neurons with stochastic dynamics. *Neural Comput*, 15:811–830, 2003.

[16] Bi G. Q. and Poo M. M. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal Of Neuroscience*, 18:10464–10472, 1998.

[17] Francisco Borja Rodriguez, Alberto Suarez, and Vicente Lopez. Period focusing induced by network feedback in populations of noisy integrate-and-fire neurons. *Neural Comput*, 13:2495–2516, 2001.

[18] Francisco Borja Rodriguez, Alberto Suarez, and Vicente Lopez. A discrete model of neural ensembles. *Philosophical Transactions: Mathematical, Physical & Engineering Sciences*, 360(1792):559–573, 2002.

[19] William Feller. *An Introduction to probability theory and its applications*. Wiley, 1968.

[20] Norman L Johnson and Samuel Kotz. *Discrete Distributions*. Wiley, 1969.

[21] Andreas Kaltenbrunner, Vicenç Gomez, Alberto Suarez, and Vicente Lopez. Phase transition and hysteresis in an ensemble of stochastic spiking neurons. *(submitted to Neural Computation)*, 2006.

[22] A. Morrison, C. Mehring, T. Geisel, A. Aertsen, and M. Diesmann. Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.*, 17:1776–1801, 2005.

[23] Bishop P. O., Levick W.R., and Williams W.O. Statistical analysis of the dark discharge of lateral geniculate neurones. *Journal Of Physiology*, 170:598–612, 1964.

[24] W. Bair, C. Koch, W. Newsome, and K. Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *Journal Of Neuroscience*, 14:2870–2892, 1994.

[25] R. B. Stein. Some models of neuronal variability. *Biophysical Journal*, 7:37–68, 1967.

[26] J. E. Strang and P. Ostborn. Wave patterns in frequency-entrained oscillator lattices - art. no. 056137. *Physical Review E*, 7205:6137–6137, 2005.

[27] U. Ernst, K. Pawelzik, and T. Geisel. Delay-induced multistable synchronization of biological oscillators. *Physical Review E*, 57:2150–2162, 1998.