



# Scaling of a Large-Scale Simulation of Synchronous Slow-Wave and Asynchronous Awake-Like Activity of a Cortical Model With Long-Range Interconnections

Elena Pastorelli<sup>1,2\*</sup>, Cristiano Capone<sup>1,3</sup>, Francesco Simula<sup>1</sup>, Maria V. Sanchez-Vives<sup>4,5</sup>, Paolo Del Giudice<sup>3</sup>, Maurizio Mattia<sup>3</sup> and Pier Stanislaw Paolucci<sup>1</sup>

<sup>1</sup> INFN, Sezione di Roma, Rome, Italy, <sup>2</sup> PhD Program in Behavioural Neuroscience, "Sapienza" University, Rome, Italy, <sup>3</sup> National Center for Radiation Protection and Computational Physics, Istituto Superiore di Sanità, Rome, Italy, <sup>4</sup> Systems Neuroscience, IDIBAPS, Barcelona, Spain, <sup>5</sup> Department of Life and Medical Sciences, ICREA, Barcelona, Spain

## OPEN ACCESS

### Edited by:

Preston E. Garraghty,  
Indiana University Bloomington,  
United States

### Reviewed by:

Sacha Jennifer van Albada,  
Julich Research Centre, Germany  
Sergio E. Lew,  
University of Buenos Aires, Argentina

### \*Correspondence:

Elena Pastorelli  
elena.pastorelli@roma1.infn.it

**Received:** 20 March 2019

**Accepted:** 08 July 2019

**Published:** 23 July 2019

### Citation:

Pastorelli E, Capone C, Simula F, Sanchez-Vives MV, Del Giudice P, Mattia M and Paolucci PS (2019) Scaling of a Large-Scale Simulation of Synchronous Slow-Wave and Asynchronous Awake-Like Activity of a Cortical Model With Long-Range Interconnections. *Front. Syst. Neurosci.* 13:33. doi: 10.3389/fnsys.2019.00033

Cortical synapse organization supports a range of dynamic states on multiple spatial and temporal scales, from synchronous slow wave activity (SWA), characteristic of deep sleep or anesthesia, to fluctuating, asynchronous activity during wakefulness (AW). Such dynamic diversity poses a challenge for producing efficient large-scale simulations that embody realistic metaphors of short- and long-range synaptic connectivity. In fact, during SWA and AW different spatial extents of the cortical tissue are active in a given timespan and at different firing rates, which implies a wide variety of loads of local computation and communication. A balanced evaluation of simulation performance and robustness should therefore include tests of a variety of cortical dynamic states. Here, we demonstrate performance scaling of our proprietary Distributed and Plastic Spiking Neural Networks (DPSNN) simulation engine in both SWA and AW for bidimensional grids of neural populations, which reflects the modular organization of the cortex. We explored networks up to  $192 \times 192$  modules, each composed of 1,250 integrate-and-fire neurons with spike-frequency adaptation, and exponentially decaying inter-modular synaptic connectivity with varying spatial decay constant. For the largest networks the total number of synapses was over 70 billion. The execution platform included up to 64 dual-socket nodes, each socket mounting 8 Intel Xeon Haswell processor cores @ 2.40 GHz clock rate. Network initialization time, memory usage, and execution time showed good scaling performances from 1 to 1,024 processes, implemented using the standard Message Passing Interface (MPI) protocol. We achieved simulation speeds of between  $2.3 \times 10^9$  and  $4.1 \times 10^9$  synaptic events per second for both cortical states in the explored range of inter-modular interconnections.

**Keywords:** spiking neural network, slow wave activity, asynchronous activity, long range interconnections, distributed simulation, strong scaling, weak scaling, large-scale simulation

## 1. INTRODUCTION

At the large scale, the neural dynamics of the cerebral cortex result from an interplay between local excitability and the pattern of synaptic connectivity. This interplay results in the propagation of neural activity. A case in point is the spontaneous onset and slow propagation of low-frequency activity waves during the deep stages of natural sleep or deep anesthesia (Hobson and Pace-Schott, 2002; Destexhe and Contreras, 2011; Sanchez-Vives and Mattia, 2014; Reyes-Puerta et al., 2016).

The brain in deep sleep expresses slow oscillations of activity at the single-neuron and local network levels which, at a macroscopic scale, appear to be synchronized in space and time as traveling waves (slow-wave activity, SWA). The “dynamic simplicity” of SWA is increasingly being recognized as an ideal test bed for refining and calibrating network models composed of spiking neurons. Understanding the dynamical and architectural determinants of SWA serves as an experimentally grounded starting point to tackle models of behaviorally relevant, awake states (Han et al., 2008; Curto et al., 2009; Luczak et al., 2009). A critical juncture in such a logical sequence is the description of the dynamic transition between SWA and asynchronous, irregular activity (AW, asynchronous wake state) as observed during fade-out of anesthesia, for instance, the mechanism of which is still a partially open problem (Curto et al., 2009; Steyn-Ross et al., 2013; Solovey et al., 2015). To help determine the mechanism of this transition, it may be of interest to identify the factors enabling the same nervous tissue to express global activity regimes as diverse as SWA and AW. Understanding this repertoire of global dynamics requires high-resolution numerical simulations of large-scale networks of neurons which, while keeping a manageable level of simplification, should be realistic with respect to both non-linear excitable local dynamics and to the spatial dependence of the synaptic connectivity (as well as the layered structure of the cortex) (Bazhenov et al., 2002; Hill and Tononi, 2005; Potjans and Diesmann, 2014; Krishnan et al., 2016).

Notably, efficient brain simulation is not only a scientific tool, but also a source of requirements and architectural inspiration for future parallel/distributed computing architectures, as well as a coding challenge on existing platforms. Neural network simulation engine projects have focused on: flexibility and user friendliness, biological plausibility, speed and scalability [e.g., NEST (Gewaltig and Diesmann, 2007; Jordan et al., 2018), NEURON (Hines and Carnevale, 1997; Carnevale and Hines, 2006), GENESIS (Wilson et al., 1989), BRIAN (Goodman and Brette, 2009; Stimberg et al., 2014)]. Their target execution platforms can be either homogeneous or heterogeneous (e.g., GPGPU-accelerated) high-performance computing (HPC) systems, (Izhikevich and Edelman, 2008; Nageswaran et al., 2009; Modha et al., 2011), or neuromorphic platforms, for either research or application purposes [e.g., SpiNNaker (Furber et al., 2013), BrainScaleS (Schmitt et al., 2017), TrueNorth (Merolla et al., 2014)].

From a computational point of view, SWA and AW pose different challenges to simulation engines, and comparing the simulator performance in both situations is an important element

in assessing the general value of the choices made in the code design. During SWA, different and limited portions of the network are sequentially active, with a locally high rate of exchanged spikes, while the rest of the system is almost silent. On the other hand, during AW the whole network is homogeneously involved in lower rate asynchronous activity. In a distributed and parallel simulation framework, this raises the question of whether the computational load on each core and the inter-process communication traffic are limiting factors in either cases. We also need to consider that activity propagates for long distances across the modeled cortical patch, therefore the impact of spike delivery on the execution time depends on the chosen connectivity. Achieving a fast and flexible simulator, in the face of the above issues, is the purpose of our Distributed and Plastic Spiking Neural Networks (DPSNN) engine. Early versions of the simulator (Paolucci et al., 2013) originated from the need for a representative benchmark developed to support the hardware/software co-design of distributed and parallel neural simulators. DPSNN was then extended to incorporate the event-driven approach of Mattia and Del Giudice (2000), implementing a mixed time-driven and event-driven strategy similar to the one introduced in Morrison et al. (2005). Here, we report the performances of DPSNN in both slow-wave (SW) and AW states, for different sizes of the network and for different connectivity ranges. Specifically, we discuss network initialization time, memory usage and execution times, and their strong and weak scaling behavior.

## 2. MATERIALS AND METHODS

### 2.1. Description of the Distributed Simulator

DPSNN has been designed to be natively distributed and parallel, with full parallelism also exploited during the creation and initialization of the network. The full neural system is represented in DPSNN by a network of C++ processes equipped with an agnostic communication infrastructure, designed to be easily interfaced with both Message Passing Interface (MPI) and other (custom) software/hardware communication systems. Each C++ process simulates the activity of one or more clusters of neighboring neurons and their set of incoming synapses. Neural activity generates spikes with target synapses in other processes; the set of “axonal spikes” is the payload of the associated exchanged messages. Each axonal spike carries the identity of its producing neuron and its original time of emission [AER, Address Event Representation (Lazzaro et al., 1993)]. Axonal spikes are only sent to target processes where at least one target synapse exists.

The memory cost of point-like neuron simulation is dominated by the representation of recurrent synapses which, in the intended biologically plausible simulations, are numbered in thousands per neuron. When plasticity support is switched off, the local description of each synapse includes only the identity of the target and source neurons, the synaptic weight, the transmission delay from the pre- to the post-synaptic neuron,

and an additional optional identifier for possible different types of synapse (see **Table 1**).

When the synaptic plasticity support is switched on, each synapse takes note of its own previous activation time. For every post-synaptic neuron spike and synaptic event a Spike-timing-dependent plasticity (STDP) contribution is computed using double floating point precision. Individual STDP events contribute to synaptic Long Term Potentiation (LTP) and Depression (LTD) through a first order low pass filter operating at a longer timescale than that of neural dynamics. Intermediate computations are performed in double precision, while the status of the low-pass filter is stored in the synaptic representation using single floating point precision. In the present work, synaptic plasticity is kept off.

During the initialization phase only, each synapse is represented at both the processes storing the source and target neuron because, in the current implementation, each process sends to each member in the subset of processes to be connected a single message requesting the creation of all needed synaptic connections. This memory overhead could be reduced by splitting the generation of connection requests in group, at the price of a proportional increase in the number of communication messages needed by the initialization phase. Indeed this is the moment of peak memory usage. Afterwards, the initialization synapses are stored exclusively in the process-hosting target neurons. For each process, the list of all incoming synapses is maintained. The synaptic list is double-ordered: synapses with the same delay are grouped together and are further ordered by presynaptic neuron index as in Mattia and Del Giudice (2000). Incoming spikes are ordered according to the identity of the presynaptic neurons. The ordering of both the synaptic list and incoming spikes helps a faster execution of the demultiplexing stage that translates incoming spikes into synaptic events because it increases the probability of contiguous memory accesses while exploring the synaptic list.

Because of the high number of synapses per neuron the relative cost of storing the neuron data structure is relatively small. It mainly contains the parameters and status variables used to describe the dynamics of the single-compartment point-like neuron. Moreover, it also contains the queue of input spikes (both recurrent and external) that arrived during the previous simulation time step, with their associated post-synaptic current value and time of arrival. External spikes are generated

as a Poissonian train of synaptic inputs. In the neuron data structure, double-precision floating-point storage is adopted for the variables used in the calculation of the membrane potential dynamics, while all the other variables and constant parameters related to the neuron are stored using single precision. Also, we opted for double precision floating point computations, because of the dominance of the cost related to the transport and memory accesses while distributing neural spikes to synapses.

In DPSNN there is no memory structure associated with external synapses: for each neuron, for each incoming external spike the associated synaptic current is generated on the fly from a Gaussian distribution with assigned mean and variance (see section 2.2); therefore external spikes have a computational cost but a negligible memory cost. The described queuing system ensures that the full set of synaptic inputs, recurrent plus external, are processed using an event-driven approach.

### 2.1.1. Execution Flow: Overview of the Mixed Time- and Event-Driven Simulation Scheme

There are two phases in a DPSNN simulation: (1) the creation of the structure of the neural network and its initialization; (2) the simulation of the dynamics of neurons (and synapses, if plasticity is switched on). For the simulation phase we adopted a combined event-driven and time-driven approach partially inspired by Morrison et al. (2005). Synaptic events drive the neural dynamics, while the message passing of spikes among processes happens at regular time steps, which must be set shorter than the minimum synaptic delay to guarantee the correct causal relationship in the distributed simulation. The minimum axo-synaptic delay sets the communication time step of the simulation, in our case 1ms.

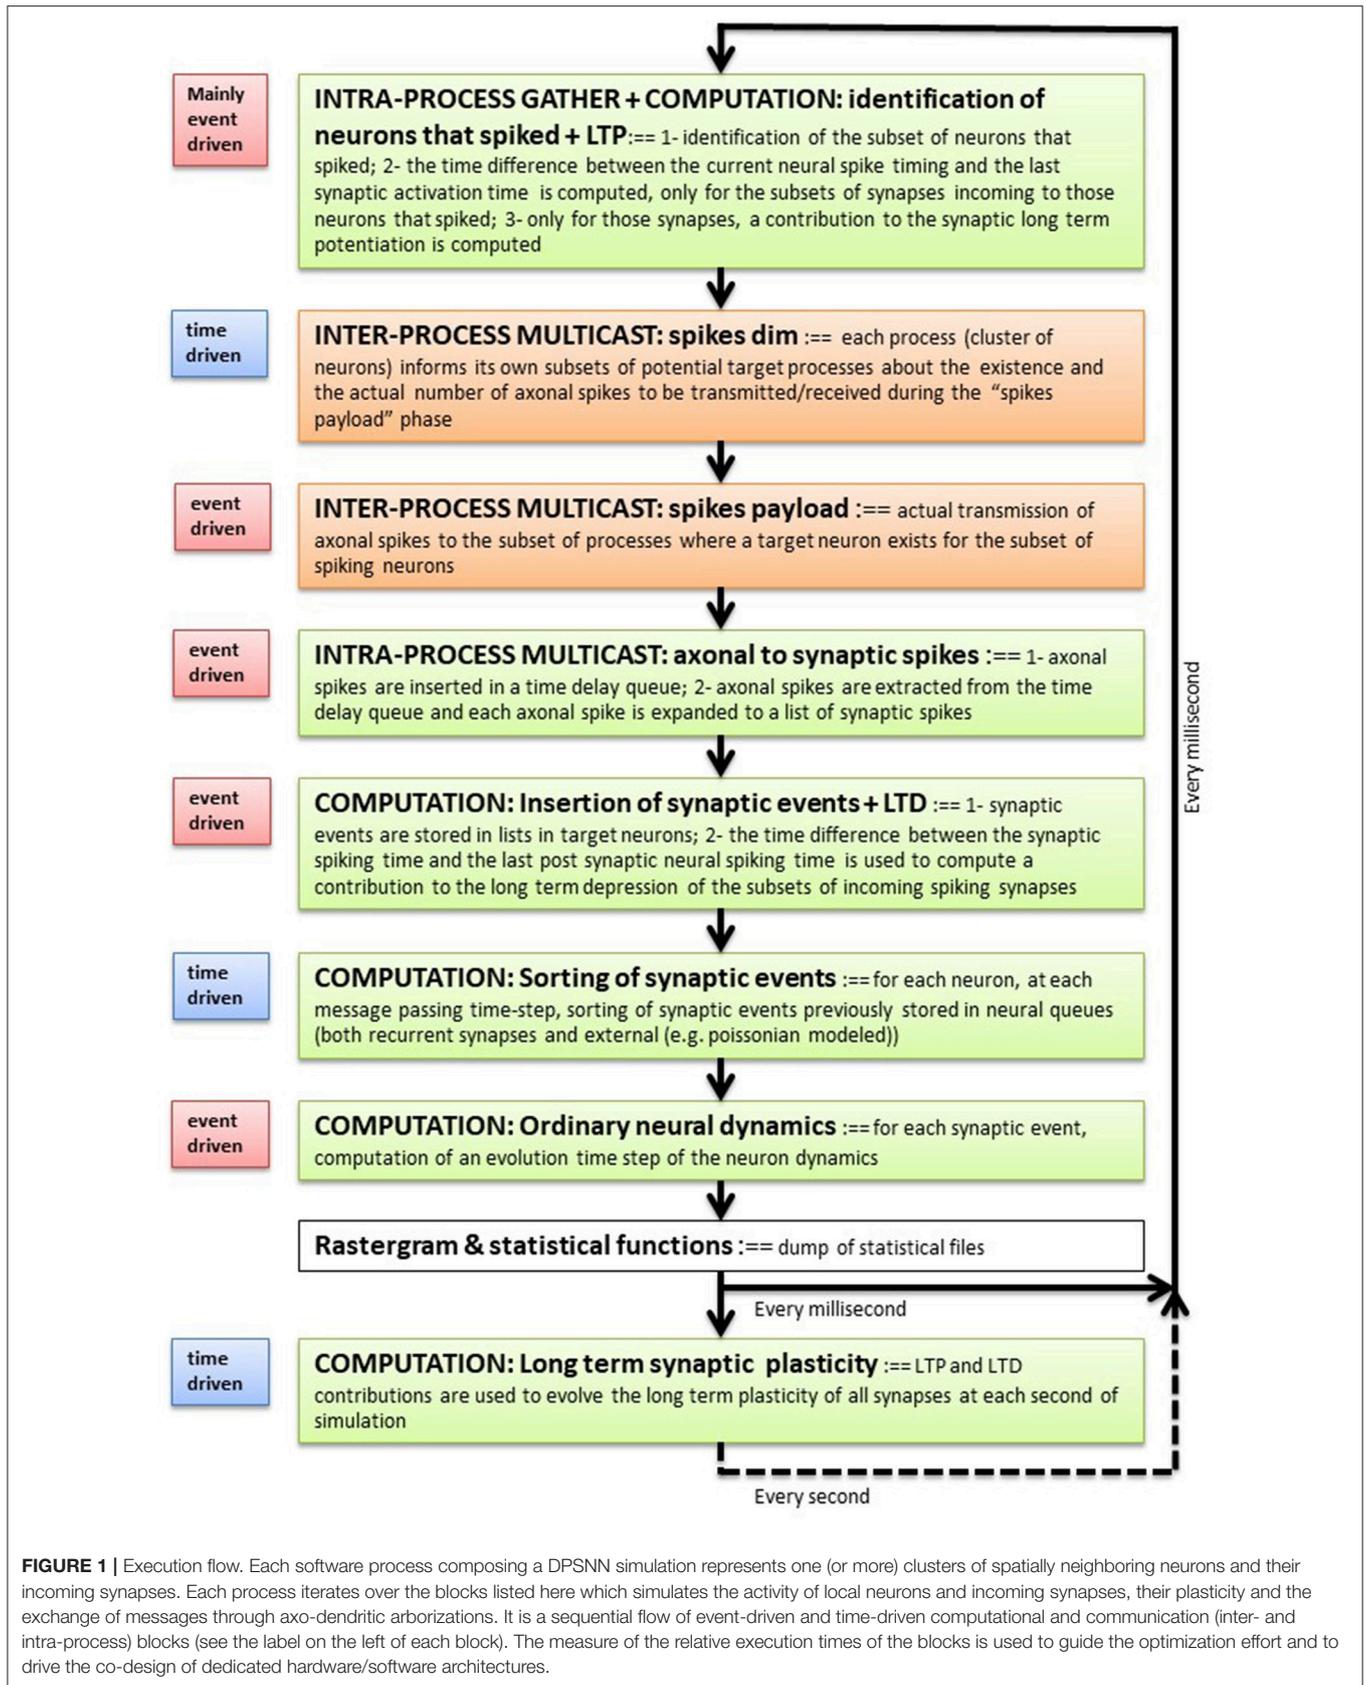
**Figure 1** describes the main blocks composing the execution flow and the event- or time-driven nature of each block. The simulation phase can be broken down into the following phases: (1) identification of the subset of neurons that spiked during the previous time step, and (when plasticity is switched on) computation of an event-driven STDP contribution; (2) spikes are sent to the cluster of neurons where target synapses exist (inter-process communication blocks in the figure); (3) the list of incoming spikes to each process is placed into the double-ordered synaptic list, waiting for a number of time steps that match the synaptic delays, at which point the corresponding target synapses are activated; (4) synapses inject their event into queues that are local to their post-synaptic neuron and compute the STDP plasticity contribution; (5) each neuron sorts the lists of input events produced by recurrent and external synapses; (6) for each event in the queue, the neuron integrates its own dynamic equations using an event-driven solver. Periodically, at a slower time step (1 s in the current implementation), all synapses modify their efficacies using the integrated plasticity signal described above. Later sections describe the individual stages and data representation in further detail.

### 2.1.2. Spike Messages: Representation and Communication

Spike messages are defined according to an AER, where each spike is represented by the identity of the spiking neuron, and the spiking time. Spikes sharing a target process (i.e., targeting

**TABLE 1** | Representation of recurrent synapses. Static synapses cost 12 bytes per synapse. Plasticity support adds a cost of 8 bytes per synapse.

Synaptic representation							
Static (12 bytes/synapse)						Plasticity (8 bytes/synapse)	
Field	Source neuron ID	Target neuron ID	Weight	Delay	Kind	Last spiking time	Derivative
Size (byte)	4	4	2	1	1	4	4



**FIGURE 1 |** Execution flow. Each software process composing a DPSNN simulation represents one (or more) clusters of spatially neighboring neurons and their incoming synapses. Each process iterates over the blocks listed here which simulates the activity of local neurons and incoming synapses, their plasticity and the exchange of messages through axo-dendritic arborizations. It is a sequential flow of event-driven and time-driven computational and communication (inter- and intra-process) blocks (see the label on the left of each block). The measure of the relative execution times of the blocks is used to guide the optimization effort and to drive the co-design of dedicated hardware/software architectures.

the cluster of neurons managed by a process) are packed into “axonal spikes” messages that are delivered from the source process to its target processes during the communication phase. During execution, the “axon” arborization is managed by the target process in order to reduce the network communication load. Its construction during the initialization phase involves the following steps: in a “fan-out” step, each process notifies the processes which will host the target synapses of its own neurons. In the next “fan-in” step, each target process stores its own list of source neurons, and the associated double-ordered list of incoming synapses. For details about the initial construction of the connectivity infrastructure and the delivery of spiking messages see **Supplementary Material**.

### 2.1.3. Bidimensional Grids of Cortical Columns and Their Mapping Onto Processes

Neural networks are organized in this study as bidimensional grids of modules (mimicking cortical columns) as in Capone et al. (2019b). Each module is composed of 1,250 point-like neurons, and further organized into subpopulations. The set of cortical columns can be distributed over a set of processes (and processors): each process can either host a fraction of a column (e.g.,  $\frac{1}{2}$ ,  $\frac{1}{4}$ ), a whole single column, or several columns (up to 576 columns per process for the experiments reported in this paper). In this respect, our data distribution strategy aims to store in contiguous memories neighboring neurons and their incoming synapses. As usual in parallel processing, for a given problem size the performance tends to saturate and then to worsen, as it is mapped onto an excessively large number of processors, so that for each grid size an optimal partitioning onto processes is heuristically established, in terms of the performance measure defined in section 2.4.

Each process runs on a single processor; if hyperthreading is active, more than one process can be simultaneously executed on the same processor. For the measures in this paper, hyperthreading was deactivated, so that one hardware core always runs a single process. In DPSNN, distributing the simulation over a set of processes means that both the initialization and the run phases are distributed, allowing the scaling of both phases with the number of processes, as demonstrated by the measures reported in section 3.

## 2.2. Model Architecture and Network States

For each population in the modular network, specific parameters for the neural dynamics can be defined, as well as specific intra- and inter-columnar connectivity and synaptic efficacies. Connectivity among different populations can be modeled with specific laws based on distance-dependent probability, specific to each pair of source and target subpopulation. By suitably setting the available interconnections between different populations, cortical laminar structures can also be potentially modeled in the simulation engine.

The grids (see **Table 2**) are squares in a range of sizes ( $24 \times 24$ ,  $48 \times 48$ ,  $96 \times 96$ ,  $192 \times 192$ ). Each local module is always composed of  $K = 1,250$  neurons, further subdivided into subpopulations. We implemented a ratio of 4:1 between

**TABLE 2** | Configurations used for the scaling measures of DPSNN.

Grid	Number of columns	Number of neurons	Number of synapses		MPI proc	
			Recurrent	Total	Min	Max
$24 \times 24$	576	0.72M	0.8G	1.1G	1	64
$48 \times 48$	2,304	2.88M	3.2G	4.4G	4	256
$96 \times 96$	9,216	11.52M	12.9G	17.6G	64	1,024
$192 \times 192$	36,864	46.08M	51.8G	70.3G	1024	1,024

excitatory and inhibitory neurons ( $K_I = 250$  neuron/module); in each module, excitatory neurons ( $E$ ) were divided into two subpopulations:  $K_F = 250$  (25%) strongly coupled “foreground” neurons ( $F$ ), having a leading role in the dynamics, and  $K_B = 750$  (75%) “background” neurons ( $B$ ) continuously firing at a relatively low rate. Populations on the grid are connected to each other through a spatial connectivity kernel. The probability of connection from excitatory neurons decreases exponentially with the inter-module distance  $d$ :

$$C_{ts\lambda}(d) = C_{ts\lambda}^0 \times \exp\left(\frac{-d}{\lambda}\right). \quad (1)$$

More specifically,  $d$  is the distance between the source ( $s = \{F, B\}$ ), and target ( $t = \{F, B, I\}$ ) module, and  $\lambda$  a characteristic spatial scale of connectivity decay.  $d$  and  $\lambda$  are expressed using inter-modular distance units (imd). For the simulations here described, the translation to physical units sets imd in the range of a few hundreds micrometers. Simulations are performed considering different  $\lambda$  values (0.4, 0.5, 0.6, 0.7) imd, but  $C_{ts\lambda}^0$  is set so as to generate the same mean number of projected synapses per neuron ( $M_t = 0.9 * K_t, t = \{F, B, I\}$ ) for all  $\lambda$  values. Connections originating from inhibitory neurons are local (within the same local module) and also in this case  $M_t = 0.9 * K_t, t = \{F, B, I\}$ . All neurons of the same type (excitatory/inhibitory) in a population share the same mean number of incoming synapses. The connectivity has open boundary conditions on the edges of the two-dimensional surface.

Synaptic efficacies are randomly chosen from a Gaussian distribution with mean  $J_{ts}$  and SD  $\Delta J_{ts}$ , chosen in different experiments so as to set the system in different working regimes and simulate different states. The procedure for the selection of the efficacies is based on a mean-field method described below. Each neuron also receives spikes coming from neurons belonging to virtual external populations, collectively modeled as a Poisson process with average spike frequency  $\nu_{ext}$  and synaptic efficacy  $J_{ext}$ . Excitatory neurons are point-like leaky integrate-and-fire (LIF) neurons with spike frequency adaptation (SFA) (Gigante et al., 2007; Capone et al., 2019b). SFA is modeled as an activity-dependent self-inhibition, described by the fatigue variable  $c(t)$ . The time evolution of the membrane potential  $V(t)$ , and  $c(t)$ , of

excitatory neurons between spikes is governed by:

$$\begin{cases} \dot{V} = -\frac{V-E}{\tau_m} - g_c \frac{c}{C_m} + \sum J_i \delta(t - t_i - \delta_i) + \sum J_{ext}^i \delta(t - t_i^{poiss}) \\ \dot{c} = -\frac{c}{\tau_c} + \alpha_c \sum_k \delta(t - t_{sp}^{(k)}) \end{cases} \quad (2)$$

$\tau_m$  is the membrane characteristic time,  $C_m$  the membrane capacitance, and  $E$  the resting potential. SFA is not considered for inhibitory neurons; that is, in (2), the second equation and the  $g_c \frac{c}{C_m}$  term in the first are absent. Incoming spikes, generated at times  $t_i$ , reached the target neuron with delay  $\delta_i$  and provoked instantaneous membrane potential changes of amplitude  $J_i$ . Alike, external stimuli produce a  $J_{ext}^i$  increment in the membrane potential, with  $t_i^{poiss}$  representing the spike times generated by a Poisson distribution of average  $\nu_{ext}$ . An output spike at time  $t_{sp}^{(k)}$  was triggered if the membrane potential exceeded a threshold  $V_\theta$ . On firing, the membrane potential was reset to  $V_r$  for a refractory period  $\tau_{arp}$ , whereas  $c$  was increased by the amount  $\alpha_c$ . Once the network connectivity and the neural dynamics have been defined, synaptic efficacies and external stimuli can be set to determine the dynamical states accessible to the system, by means of mean-field theory.

We elaborated a dynamical mean-field description for our simulations. We assume that the input received by the different neurons in a module are independent but are conditioned by the same (and possibly time-dependent) mean and variance of the input synaptic current (Brunel and Hakim, 1999; Mattia and Del Giudice, 2002).

The gain function  $\phi_i$  for this kind of neurons was firstly found by Ricciardi (1977) as the first-passage time (FPT) of the membrane potential to the firing threshold, for an integrate-and-fire neuron with stationary white noise input current, in the diffusion approximation.

Spike frequency adaptation is an important ingredient for the occurrence of slow oscillations in the mean field dynamics (Gigante et al., 2007; Capone et al., 2019b).

The mean-field dynamics for the average activity  $\nu_i, i = \{F, B, I\}$  is determined by the gain functions  $\phi_i$  as follows:

$$(F, B): \begin{cases} \dot{\nu}_i = \frac{\phi_i(\vec{\nu}, \vec{c}) - \nu_i}{\tau_E} \\ \dot{c}_i = -\frac{c_i}{\tau_c} + \alpha_c \nu_i \end{cases} ; \quad (I): \nu_i = \frac{\phi_i(\vec{\nu}) - \nu_i}{\tau_I} \quad (3)$$

where  $\tau_E$  and  $\tau_I$  are phenomenological time constants. The interplay between the recurrent excitation embodied in the gain function, and the activity-dependent self-inhibition, is the main driver of the alternation between a high-activity (Up) state and a low-activity (Down) state (Mattia and Sanchez-Vives, 2012; Capone et al., 2019b).

In the mean-field description of the neural module, synaptic connectivity is chosen so as to match the total average synaptic input the neurons would receive inside the multi-modular network. The fixed points of the dynamics expressed by Equation (3) can be analyzed using standard techniques (Strogatz, 2018). The nullclines of the system (where  $\dot{\nu} = 0$  or  $\dot{c} = 0$ ) cross at fixed points that can be predicted to be either stable or unstable. In the

simulations described here, the strengths of recurrent synapses  $J_{t,s}$  connecting source population  $s$  and target population  $t$ , and of external synapses  $J_{t,ext}$ , is randomly chosen from a Gaussian distribution with mean  $J_{t,s}$  and variance  $\Delta J_{t,s} = 0.25 \times J_{t,s}$ .

We relied on mean-field analysis to identify neural and network parameters setting the network's modules in SW or AW dynamic regimes. **Figure 2A** shows an example of nullclines for the mean-field equations (3) of a system displaying SW. The black S-shaped line is the nullcline for the rate  $\nu$  while the red straight line is the one for the fatigue variable  $c$  (for details see Mattia and Sanchez-Vives, 2012; Capone et al., 2019b). The stable fixed point, at the intersection of the nullclines, has a low level of activity and is characterized by a small basin of attraction: the system can easily escape from it thanks to the noise, and it gets driven toward the upper branch of the  $\nu$  nullcline from which, due to fatigue, it is attracted back to the fixed point, thus generating an oscillation (see **Figure 2B**).

Network parameters can also be set in order to have an asynchronous state, mainly by setting a lower Foreground-to-Foreground (FF) synaptic efficacy, which generates a more linear  $\nu$  nullcline close to the fixed point (see **Figure 2C**). In this case the basin of attraction of the fixed point is larger, and oscillations do not occur, resulting in a stationary asynchronous state (see **Figure 2D**), in which neural activity fluctuates around the mean-field fixed point. **Table S1** reports the complete list of synaptic parameters for both SW and AW states.

### 2.3. Neural Activity Analysis

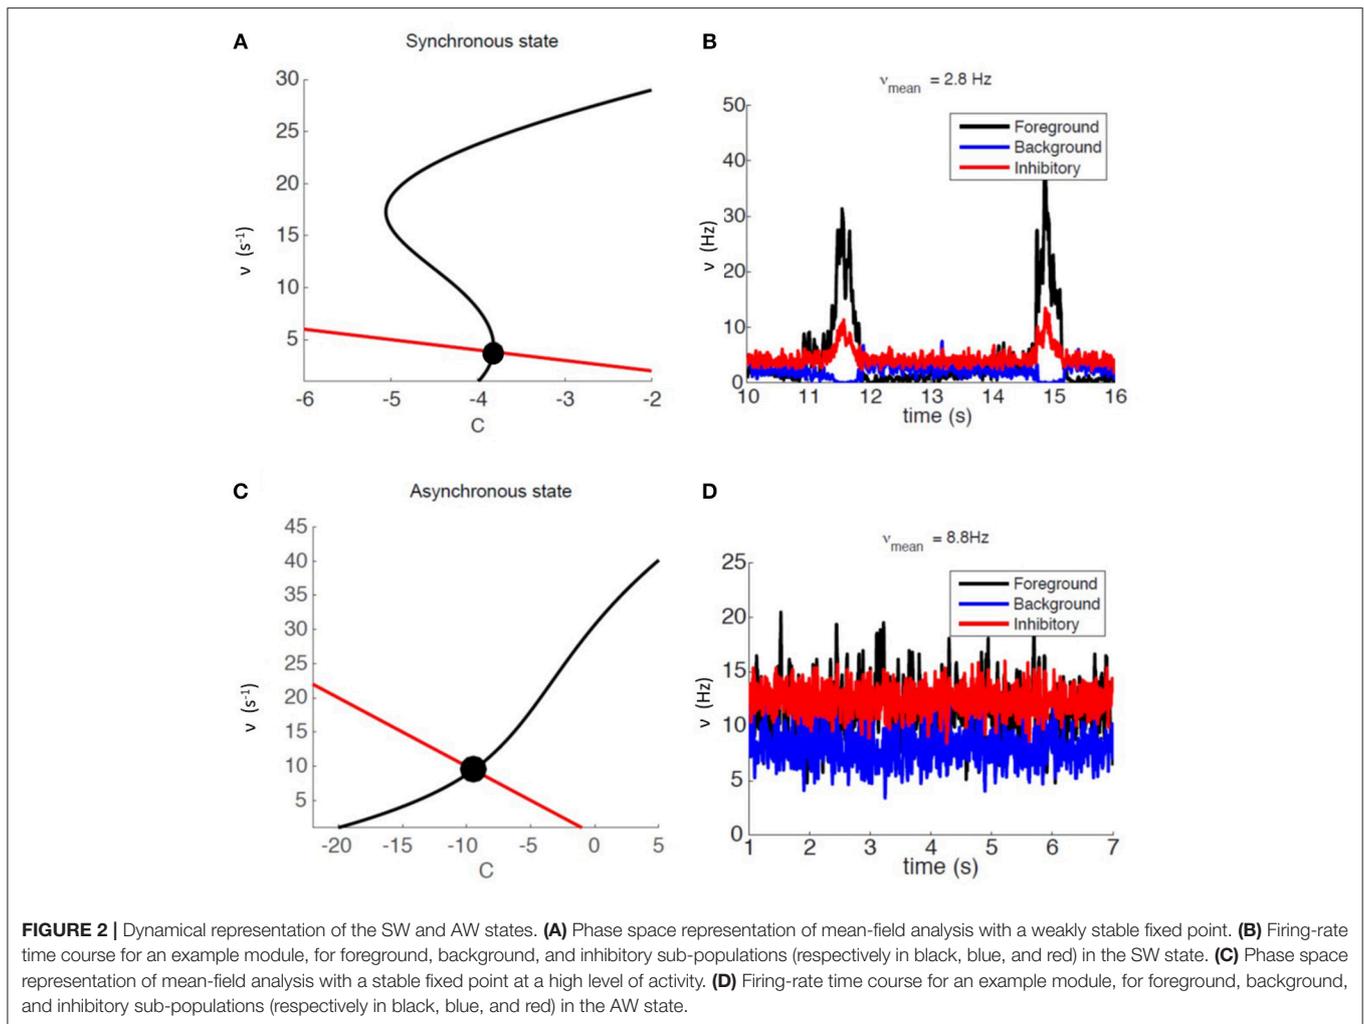
The simulation generates the spikes produced by each neuron in the network. From these, the time course of the average firing rate for each subpopulation in the network can be computed using an arbitrary sampling step, which here we set to be 5 ms. Power spectra are computed using the Welch method (see **Figures 6B,E** for examples of SW and AW power spectra including delta band).

A proxy to experimentally acquired multi-unit activity (MUA(t)) is computed as the average firing rate of a simulated module. In similarity with experimental data analysis (Capone et al., 2019b) we considered the logarithm of such signal ( $\log \frac{MUA(t)}{MUA_{down}}$ ) where the  $MUA_{down}$  is the average MUA(t) in the Down state. A zero mean white noise is added to emulate unspecific background fluctuations from neurons surrounding the module. The variance of such noise (0.5) was chosen to match the width of experimental  $\log \frac{MUA(t)}{MUA_{down}}$  distributions for bistable neural populations in their Down state.

For SW, the MUA(t) signal alternated between high and low activity states (Up and Down states).  $T(x, y)$ , the Down-to-Up transition time, is determined by a suitable threshold and is detected for each point in the spatial grid, defining the propagation wavefront.  $V(x, y)$ , the absolute value of wavefront speed, can be computed as

$$V(x, y) = \frac{1}{\sqrt{\left(\frac{\partial T(x,y)}{\partial x}\right)^2 + \left(\frac{\partial T(x,y)}{\partial y}\right)^2}} \quad (4)$$

The average speed during the collective Down-to-Up transitions is obtained by averaging  $V(x, y)$  over all the simulated positions  $(x, y)$ .



## 2.4. Performance Measures

Performance is quantified in terms of “strong scaling” and “weak scaling”: the former refers to keeping the problem size fixed and varying the number of hardware cores, while the latter refers to the increase in equal proportions of the problem size and the number of hardware cores.

As a performance measure we computed the equivalent synaptic events per second as the product of the total number of synapses (recurrent and external) and the number of spikes occurred across the whole simulation, divided by the elapsed execution time. This way, a comparison of the simulation cost among different problem sizes and hardware/software resources (core/processes) can be captured in a single graph. Similarly, we defined a convenient metric to evaluate the memory efficiency of a simulation: by dividing the total memory required by the simulation by the number of recurrent synapses to be represented. Indeed, as stated in section 2.1 we expect the memory usage to be dominated by the representation of synapses which are thousands per neuron.

Scalability measurements were taken on different neural network sizes, varying the size of the grid of columns and, for

each size, distributing it over a different span of MPI processes. We selected four grid sizes:  $24 \times 24$  columns, including 0.7M neurons and 1.1G synapses;  $48 \times 48$  columns including 2.8M neurons and 4.4G synapses;  $96 \times 96$  columns including 12M neurons and 17.6G synapses;  $192 \times 192$  columns including 46M neurons and 70G synapses. The number of processes over which each network size was distributed varied from a minimum, bounded by memory, and a maximum, bounded by communication or HPC platform constraints (see **Table 2**). For the  $192 \times 192$  configuration, only one measure was taken because of memory requirements, corresponding to a run distributed over 1,024 MPI processes/hardware cores.

Execution times for SW were measured across the time elapsed between the rising edges of two subsequent waves, and for AW across a time span of 3 s. In both cases, initial transients were omitted.

## 2.5. Validation of Results and Comparison of Performances

We used NEST version 2.12, the high-performance general purpose simulator developed by the NEST Initiative, as a

reference for the validation of results produced by the specialized DPSNN engine and for the comparison of its performances (speed, initialization time, memory footprint). The comparison was performed for both SW and AW dynamical states, for a network of 4.4G synapses ( $48 \times 48$  grid of columns). We assessed the correctness of results using power spectral density (PSD) analysis of the temporal evolution of the firing rates of each subpopulation. **Figure 3** reports an example of this comparison, showing the PSD match of NEST and DPSNN for each single subpopulation in a randomly chosen position inside the grid of columns. In AW state the average firing rate is about 8.8Hz with about 1125 synapses per neuron. Not surprisingly, NEST started to converge to the production of stable Power Spectral Densities only when the integration time step was reduced down to 0.1ms. We used this time step for all simulations used for comparison with DPSNN.

## 2.6. Execution Platforms

The benchmark measures assessing the DPSNN scalability in terms of run time, initialization time, and memory usage herein described, were performed on the Galileo server platform provided by the CINECA consortium. Galileo is a cluster of 516 IBM nodes, each of which includes 16-cores, distributed on two Intel Xeon Haswell E5-2630 v3 octa-core processors clocked at 2.40 GHz. The nodes are interconnected through an InfiniBand network. Due to the specific configuration of the server platform, the maximum-allowed partition of the server includes 64 nodes. Hyperthreading is disabled on all cores, therefore the number of MPI processes launched during each run exactly matches the hardware cores, with a maximum of 1,024 hardware cores (or equivalently MPI processes) available for any single run. In the present version of the DPSNN simulator, parameters cannot be changed at runtime, and dynamic equations are explicitly coded (i.e., no meta-description is available as an interface to the user).

Simulations used for validation and comparison between DPSNN and NEST were run on a smaller cluster of eight nodes interconnected through a Mellanox InfiniBand network; each node is a dual-socket server with a six-core Intel(R) Xeon(R) E5-2630 v2 CPU (clocked at 2.60 GHz) per socket with HyperThreading enabled, so that each core is able to run two processes.

## 3. RESULTS

### 3.1. Initialization Times

The initialization time, in seconds, is the time required to complete the setup phase of the simulator, which is necessary to build the whole neural network. Measures reported in **Figure 4A** show the scaling of the DPSNN initialization time for four network grid sizes, distributed over a growing number of hardware cores, which also correspond to MPI processes. Dashed colored lines represent strong scaling. Weak scaling is represented by the four points connected by the dotted black line (each point refers to a 4-fold increase in the network size and number of used cores with respect to the previous one). For the explored network sizes and hardware resources, the

initialization time speedup is almost ideal for fewer cores, while for the highest numbers of cores it is sub-ideal by a factor between 10 and 20%.

**Figure 4B** reports the scaling of the initialization time for two different values of the characteristic spatial scale of the decay of the connection probability ( $\lambda$ ), for a single grid size ( $96 \times 96$  columns). As already explained, for all simulations we kept the total number of synapses per neuron constant. Going from  $\lambda = 0.4$  to  $\lambda = 0.6$  imd, each excitatory neuron in a column has synaptic connections with neurons in a number of other columns growing from 44 to 78 (neglecting variations to proximity to columns' boundaries). Almost the same scaling is observed for both values of  $\lambda$ , with expected higher values for the initialization of a network with longer-range connectivity. Indeed, a dependence of the initialization time on synaptic connectivity range is expected, because it affects the proportion of target synapses residing in different columns, for which MPI messaging is needed for the connectivity setup, as explained in Materials and Methods.

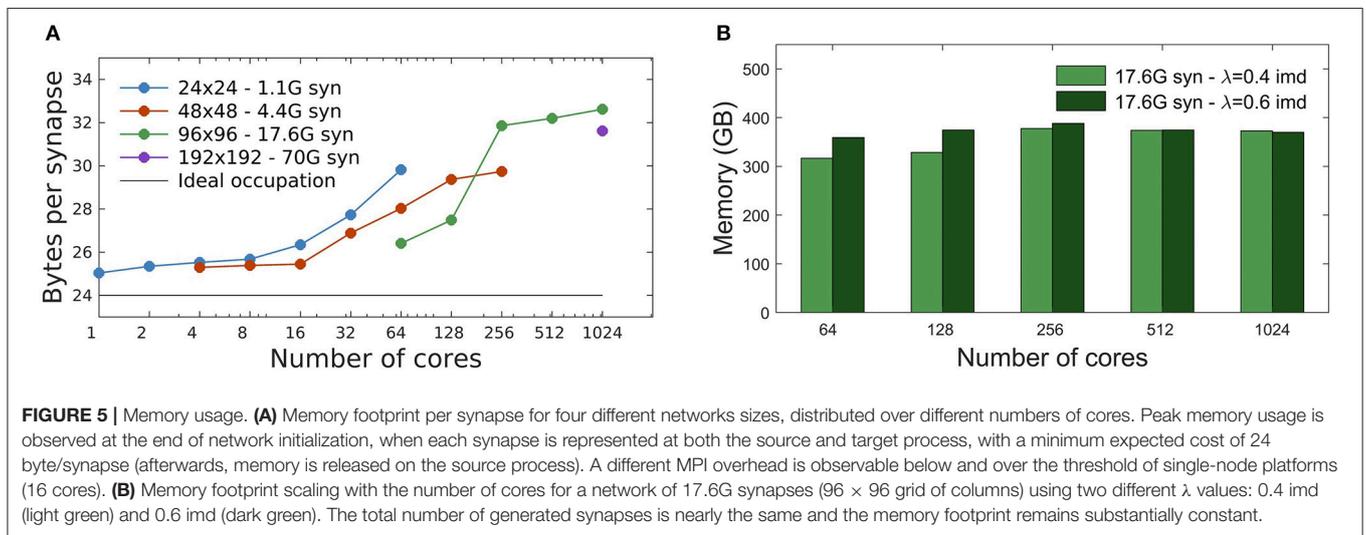
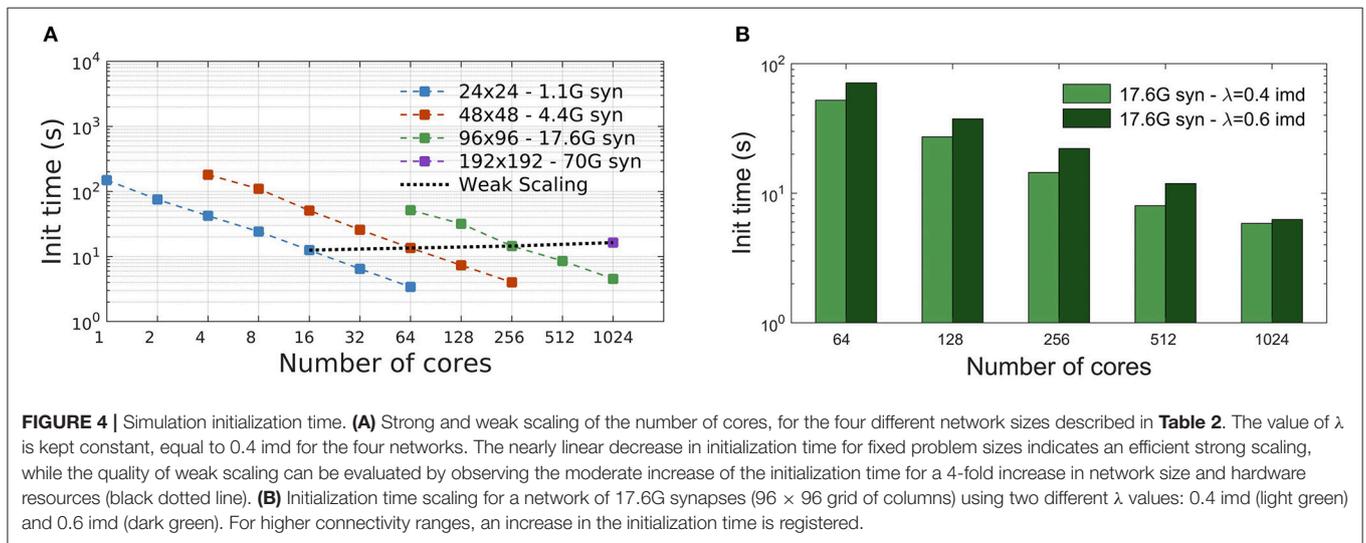
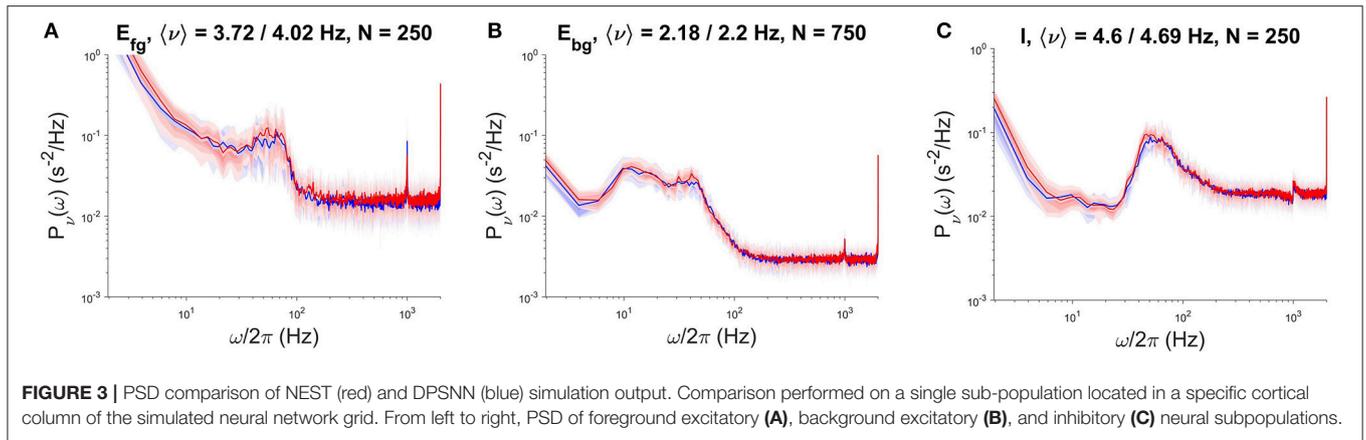
### 3.2. Memory Occupation

In DPSNN, we expect the memory usage to be dominated by the representation of synapses, which are thousands per neuron, with only a minor impact of the representation of neurons and external synapses. Therefore, we defined a memory consumption metric as the total required memory divided by the number of recurrent synapses (see section 2.4).

Each static synapse needs 12 bytes (see **Table 1**). As described in section 2.1, during network initialization synapses are generated by the process storing the source neuron. Subsequently, their values are communicated to the process containing the target neuron. The result is that, only during the initialization phase, each synapse is represented at both source and target processes. Therefore, we expect a minimum of 24 bytes/(recurrent synapse) to be allocated during initialization, which is the moment of peak memory usage.

The measured total memory consumption is between 25 and 32 bytes per recurrent synapse, for all simulations performed, from 1 to 1,024 MPI processes (**Figure 5**). From **Figure 5A** we observe a different trend of the memory cost vs. number of cores, below and above the threshold of single-node platforms (16 cores, see section 2.6). Beyond a single node, additional memory is mainly required by the buffers allocated by MPI interconnect libraries that adopt different strategies for communications over shared memory (inside a node) or among multiple nodes (**Figure 5A**). The memory occupation per synapse has been measured for different network sizes. **Figure 5A** shows that, for a given number of cores, the memory overhead typically decreases for increasing size of the network, as expected.

**Figure 5B** shows the impact of the spatial scale of connectivity decay on the memory footprint. Here the network has size  $96 \times 96$  columns (for a total of 17.6G synapses), and  $\lambda = 0.4, 0.6$  imd. As expected the memory footprint, mainly dependent on the total number of synapses, essentially remains constant.



Note that the relative differences for  $\lambda = 0.4, 0.6$  imd decrease as the number of cores increases; this is consistent with the fact that, given the network size, for larger core

numbers more columns get distributed on processes residing in different cores. This dilutes the effect of  $\lambda$ , which goes in the same direction.

### 3.3. Simulation Speed and Its Scaling

#### 3.3.1. Impact of Simulated State (SW or AW)

DPSNN execution times for SW and AW state simulations are reported in **Figure 6**. **Figures 6A,B** show example snapshots of the simulated network activity in the SW and AW states, respectively; **Figures 6C,D** show corresponding power spectra of the network activity, confirming the main features predicted by theory for such states (such as the low-frequency power increase for SW, the high-frequency asymptote proportional to the average firing rate, the spectral resonances related to SFA, and delays of the recurrent synaptic interaction). **Figures 6E,F** represent strong scaling for SW and AW: the wall-clock time required to simulate 1 s of activity, vs. number of processes, for four network sizes (see **Table 2**). In the same plots, weak scaling behavior is measured by joining points referring to a 4-fold increase in both the network size and the number of processes.

Simulation speeds are measured using the equivalent synaptic-events-per-second metric, defined in section 2.4. Departures from the ideal scaling behavior (linearly decreasing strong scaling plots, horizontal weak scaling plots) are globally captured in **Figure 7A**, for both SW and AW states and for all the problem sizes. For the simulations reported in this study, the simplified synaptic-events-per-second metric is a good approximation of a more complex metric separating recurrent and external events (Poisson noise), as demonstrated by **Figure 7B**: looking in more detail, the simulation of recurrent synaptic events is slower than that of external events that are locally generated by the routine that computes the dynamics of individual neurons. In our configurations there is about one order magnitude more recurrent synaptic event than external.

As numbers representative of the good scalability of the simulator, we remark that in the worst cases (1,024 processes;  $96 \times 96$  for SW,  $192 \times 192$  for AW), the speedup with respect to the reference case ( $24 \times 24$  on a single process) is about 590 for AW and 460 for SW, compared with the ideal speedup of 1,024.

We also notice that performance scaling is slightly better for AW than SW, which is understood in terms of workload balance (due to the more homogeneous use of resources in AW); it is in any case remarkable that, although SW and AW dynamic states imply very different distributions of activity in space and time, the simulator provides comparable performance scaling figures.

Starting from the simulation speed expressed in terms of equivalent synaptic events per wall-clock second, DPSNN simulations presents a slow-down factor with respect to the real time, variable with the simulated network state and the available hw resources used for simulation. For example, considering a commodity cluster as the one used in this paper, able to allocate up to 64 processes, a network of size  $24 \times 24$  is about 12 times slower than real-time when simulating a SW state and about 23 times slower for the AW state. For larger simulations the slow-down factor increases, due to the not perfect scaling of the simulator performances. In this case, a  $96 \times 96$  network distributed on 1,024 processes presents a slow-down factor of about 23 for a SW simulation and of about 31 for the AW state.

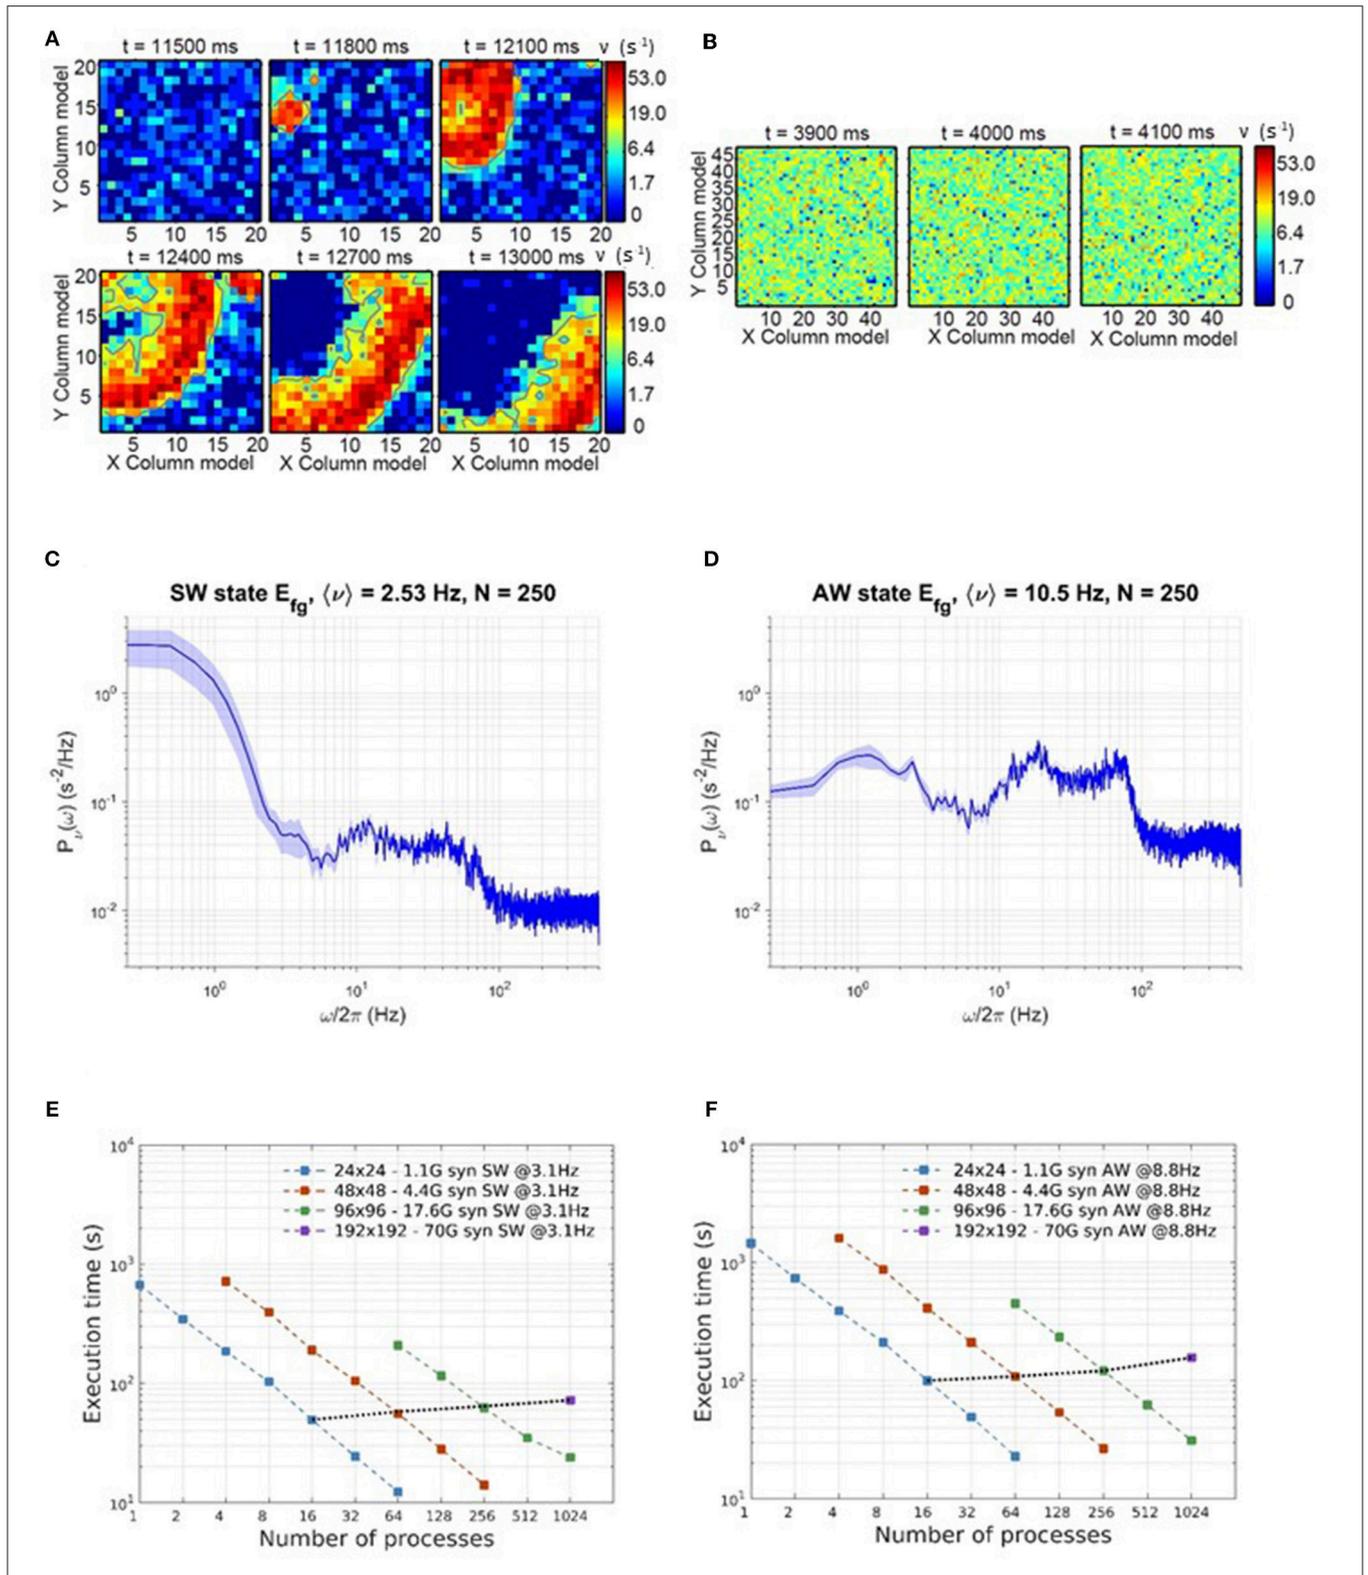
#### 3.3.2. Impact of Communication

Simula et al. (2019) studied the relative impact of computation and communication on the performance of DPSNN applied to simulations of AW states. For neural network sizes and number of processes in the range of those reported in this paper, the time spent in computation is still dominant, while communication grows to be the dominant factor when the number of neurons and synapses per process is reduced. In this paper, we evaluated the impact of communications on DPSNN performances both on SW and AW state simulations, using two different approaches. In the case of SW simulations, the analysis has been carried out with varying  $\lambda$  (therefore varying the ratio of local vs. remote excitatory synaptic connections, at a fixed total number of synapses per neuron):  $\lambda = 0.4$  imd (60% local connectivity),  $\lambda = 0.6$  imd (35% local connectivity); clearly, higher  $\lambda$  results in higher payload in communication between processes. Also, **Figure 8A** shows the known linear dependence of the slow-wave speed on  $\lambda$  (Coombes, 2005; Capone and Mattia, 2017). As the wave speed increases, the duration of the Up states stays approximately constant (not shown), so that an increasing portion of the network is simultaneously activated, which in turn may impact the simulation performance. In physical units, for an inter-modular distance (imd) of 0.4 mm,  $\lambda = 0.6$  imd implies a spatial decay scale of 0.24 mm, and the corresponding speed is  $\sim 15$  mm/s, which is in the range of biologically plausible values (Sanchez-Vives and McCormick, 2000; Ruiz-Mejias et al., 2011; Wester and Contreras, 2012; Stroh et al., 2013; Capone et al., 2019b).

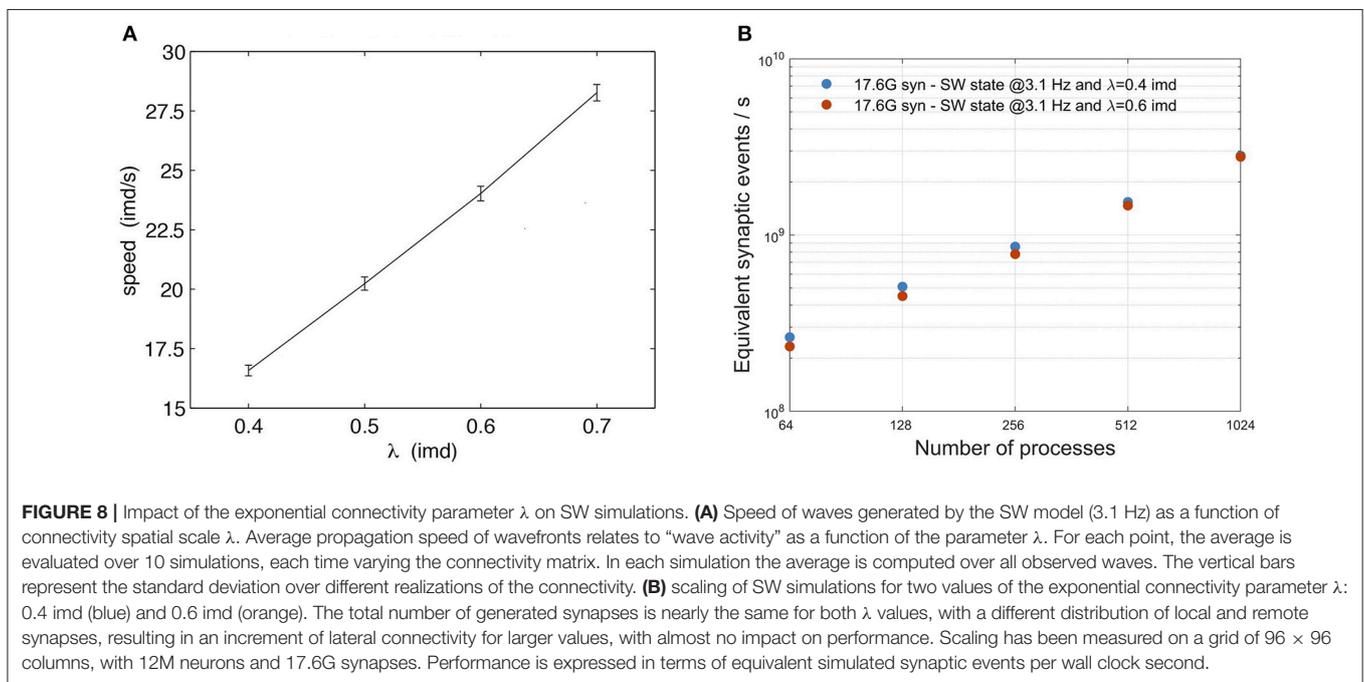
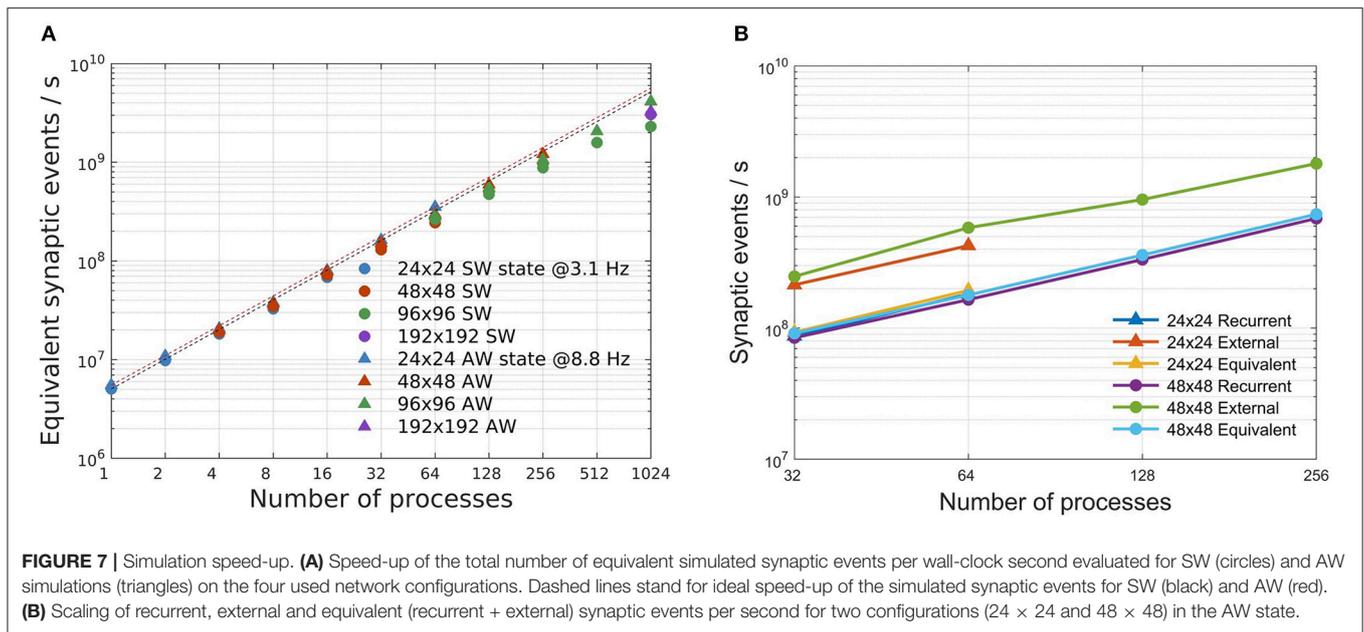
**Figure 8B** shows that the impact of  $\lambda$  on SW simulations is almost negligible. **Figure 9** shows, instead, the impact of different mean firing rates on AW simulations. Higher firing rates imply higher payloads. DPSNN also demonstrates good scaling behavior in this case, with a slight performance increase for systems with a higher communication payload; this latter feature is due to communication costs being typically dominated by latencies and not bandwidth in spiking network simulations.

### 3.4. DPSNN and NEST: Comparison of Performances

**Table 3** reports a performance comparison between DPSNN and NEST using the configuration, described in section 2.5. DPSNN is about 130 times faster than NEST for SW simulations and about 80 times faster for AW cases, and its initialization is about 19 times faster. The memory footprint of DPSNN is about 2.5 times lower due to the decisions of representing the identities of presynaptic and post-synaptic neurons with only 4 bytes per neuron and the storage of weights using only 2 bytes per synapse (see **Table 1**). For the comparison of execution speed, we selected the best execution time for each simulator, for a fixed amount of used hardware resources; that is, the number of nodes of the cluster server. On the NEST simulator, we explored the space of all possible combination of MPI processes and number of threads during the AW simulations, in order to find the configuration performing better on a fixed number of hardware resources. In the same configuration, we also compared the initialization phase, the memory usage, and the SW execution times.



**FIGURE 6 |** Simulation time scaling and phenomenological behavior. **(A)** Time consecutive snapshots of the activity distribution in space, showing the wavefront propagation during a simulation expressing SW states. **(B)** Consecutive snapshots of the whole network activity in an asynchronous state. **(C,D)** Power spectra of network activity, respectively in SW **(C)** and AW **(D)** states. **(E,F)** Scaling of wall-clock execution time for 1 s of SW **(E)** and AW **(F)** simulated activity. In both SW and AW states, the scaling has been measured on four different network sizes, as in **Table 2**.



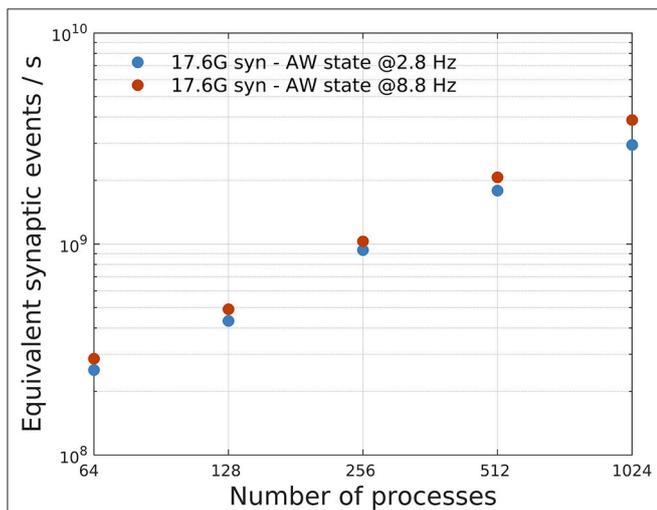
### 4. DISCUSSION

We presented a parallel distributed neural simulator, with emphasis on the robustness of its performance and scaling with respect to quite different collective dynamical regimes. This mixed time- and event-driven simulation engine (Figure 1) has been used to simulate large-scale networks including up to 46 million point-like spiking neurons interconnected by 70 billion instantaneous current synapses.

The development of DPSNN originated from the need for a simple, yet representative benchmark (i.e., a mini-application) developed to support the hardware/software co-design of distributed and parallel neural simulators. Early versions of DPSNN (Paolucci et al., 2013) have been used to drive the development of the EURETILE system (Paolucci et al., 2016) in which a custom parallelization environment and the APEnet hardware interconnect were tested. DPSNN was then extended to incorporate the event-driven approach

of Mattia and Del Giudice (2000), implementing a mixed time-driven and event-driven strategy inspired by Morrison et al. (2005). This simulator version is also currently included among the mini-applications driving the development of the interconnect system of the EXANEST ARM-based HPC architecture (Katevenis et al., 2016). In the framework of the Human Brain Project (<https://www.humanbrainproject.eu>), DPSNN is used to develop high-speed simulation of SW and AW states in multi-modular neural architectures. The modularity results from the organization of the network into densely connected modules mimicking the known modular structure of cortex. In this modeling framework, the inter-modular synaptic connectivity decays exponentially with the distance.

In this section, we first discuss the main strengths of the simulation engine and then put our work into perspective by comparing the utility and performances of such a specialized engine with those of a widely used general-purpose neural simulator (NEST).



**FIGURE 9 |** Impact of different mean firing rates on AW state simulations. A network with 17.6G synapses (96 × 96 grid of columns) has been simulated with two different values of mean firing rates, 2.8 Hz (blue) and 8.8 Hz (orange), corresponding to a different payload in communications. The plot shows a slight performance increase, in terms of equivalent simulated synaptic events per second, for systems with higher firing rate (and higher communication payload).

### 4.1. Speed, Scaling, and Memory Footprint at Realistic Neural and Synaptic Densities

DPSNN is a high-speed simulator. The speed ranged from  $3 \times 10^9$  to  $4.1 \times 10^9$  equivalent synaptic events per wall-clock second, depending on the network state and the connectivity range, on commodity clusters including up to 1,024 hardware cores. As an order of magnitude, the simulation of a square cortical centimeter ( $\sim 27 \times 10^9$  synapses) at realistic neural and synaptic densities is about 30 times slower than real time on 1024 cores (Figure 7).

DPSNN is memory parsimonious: the memory required for the above square cortical centimeter is 837 GB (31 byte/synapses, including all library overheads), which is in the range of commodity clusters with few nodes. The total memory consumption ranges between 25 and 32 byte per recurrent synapse for the whole set of simulated neural networks and all configurations of the execution platform (from 1 to 1,024 MPI processes, Figure 5). The choice of representing with 4 bytes the identities of presynaptic and post-synaptic neurons limits the total number of neurons in the network to 2 billions. However, this is not yet a serious limitation for neural networks including thousands of synapses per neuron to be simulated on execution platforms including few hundreds of multi-core nodes. The size of neural ID representation will have to be enlarged for simulation of systems at the scale of human brains. Concerning synaptic weights, as already discussed, static synapses are stored with only two bytes of precision, but injection of current and neural dynamics is performed with double precision arithmetic. When plasticity is turned on, single precision floating-point storage of LTP and LTD contributions is adopted.

The engine has very low initialization times. DPSNN takes about 4 s to set up a network with  $\sim 17 \times 10^9$  synapses (Figure 4). We note that the initialization time is relevant, especially when many relatively short simulations are needed to explore a large parameter space.

Its performance is robust: good weak and strong scaling have been measured in the observed range of hardware resources and for all sizes of simulated cortical grids. The simulation speed was nearly independent from the mean firing rate (Figure 9), the range of connectivity (Figure 8), and from the cortical dynamic state (AW/SW) (Figure 7).

### 4.2. Key Design Guidelines of the Simulation Engine

A few design guidelines contribute to the speed and scaling of the simulation engine. We kept as driving criteria the

**TABLE 3 |** Memory footprint, initialization, and execution times for 10 s of activity in SW and AW states required by DPSNN and NEST to simulate a neural network made of 2.8M neurons and 4.4G synapses (48 × 48 grid of cortical columns).

Execution platform		Init time (s)		SW Exec time (s)		AW Exec time (s)		Memory (GB)	
Nodes	Cores/Procs.	DPSNN	NEST	DPSNN	NEST	DPSNN	NEST	DPSNN	NEST
2	48	29.2	603.4	1,143.3	164,555.3	2,371.3	164,008.8	78.8	208.0
4	96	16.3	299.3	620.0	83,305.0	1,257.7	85,025.2	78.4	200.8
8	192	8.3	149.1	360.0	42,601.7	684.9	42,425.2	96.3	197.6

The second column reports the number of processes over which simulations are distributed. In case of DPSNN it corresponds to the number of pure MPI processes, while for NEST it is the number of Virtual Processes (VP). Each VP is calculated as MPI processes × number of threads, where six threads are used for all NEST simulations.

goals of increasing the locality in memory accesses, the reduction of interprocess communication, an ordered traversal of lists and a reduction of backward searches and random accesses, the minimization of the number of layers in the calling stack, and a complete distribution of computation and storage among the cooperating processes with no need for centralized structures.

We stored in the memory of a process a set of spatially neighboring neurons, incoming synapses, and outgoing axons. For the kind of spatially organized neural networks described in this paper, this reduces the size of the payload and the required number of interprocess communications. Indeed, many spiking events will need to reach target neurons (and synapses) stored in the same process of the spiking neuron itself.

An explicit ordering in memory is adopted for outgoing and ingoing communication channels (representing the first branches in an axonal like arborizations). Explicit ordering in memory is also adopted for other data structures, like: the lists of incoming spiking events and recurrent synapses, the set of incoming event queues associated to different synaptic delays. Lists are implemented as arrays, without internal pointers. Also explicit ordering is used for the set of target neurons. Neurons are progressively numbered with lower bits in their identifiers associated to their local identity and higher bits conveying the id of the hosting process. The explicit ordering of neurons and synapses reduces the time spent during the sequence of memory accesses that will be followed during the simulation steps. In particular, lists are traveled only once per communication step, e.g., first searching for target synapses that are targets of spiking events and then through an ordered list of target neurons.

A third design criteria has been to keep the stack of nested function calls short, following the scheme described in **Figure 1**, and attempting to group at each level multiple calls to computational methods accessing the same memory structure. As an example, both the generation of external synaptic events (e.g., Poissonian stimulus) and the temporal reordering of recurrent and external synaptic events targeting a specific neuron, is deferred to the computation of the individual neural dynamics. The execution of this routine is supported by local queues storing all the synaptic currents targeting the individual neuron. This queue of events is accessed during a single call of the routine computing the dynamics of the neuron. In this case the data structure supports both locality in memory and in computation.

Similar design guidelines would be problematic to follow for general purpose simulation engines that are supposed to support maximum flexibility in the description and simulation of the data structures and of the dynamics of neurons, axonal arborizations and synapses. Moreover, higher abstraction requires separating the functionality of the simulation engine into independent modules. This would dictate a higher number of layers in the calling stack, more complex interfaces and data structures that hide details like their internal memory ordering.

### 4.3. Comparison With a State-of-the-Art User-Friendly Simulator and Motivations for Specialized Engines

There is a widely felt need for versatile, general-purpose neural simulators that offer a user-friendly interface for designing complex numerical experiments and provide the user with a wide set of models of proven scientific value. This boosted a number of initiatives (notably the NEST initiative, now central to the European Human Brain Project). However, such flexibility comes at a price. Performance-oriented engines, missing all the layers required for offering user generality and flexibility, contain the bare minimum code. In the case of DPSNN, this resulted in higher simulation speed, reduced memory footprint, and diminished initialization times (see section 3.4 and **Table 3**). In addition, optimization techniques developed for such engines on use-cases of proven scientific value can offer a template for future releases of general-purpose simulators. Indeed, this is what is happening in the current framework of cooperation with the NEST development team. Finally, engines stripped down to essential kernels constitute more easily manageable mini-application benchmarks for the hardware/software co-design of specialized simulation systems, because of easier profiling and simplified customizations on system software environments and hardware targets under development.

### 4.4. Future Works

The present implementation of DPSNN demonstrated to be efficient for homogeneous bidimensional grids of neural columns and for their mapping of up to 1,024 processes, and this facilitates a set of interesting scientific applications. However, further optimization could improve DPSNN performance, either in the perspective of moving simulations toward million-core exascale platforms or targeting real-time simulations at smaller scale (Simula et al., 2019), in particular addressing sleep-induced optimizations in cognitive tasks like classification (Capone et al., 2019a). For instance, we expect that the delivery of spiking messages will be a key element to be further optimized (e.g., using a hierarchical communication strategy). This will also be beneficial for an efficient support of white-matter long-range connectivity (brain connectomes) between multiple cortical areas.

A full exploitation of the model requires parameters tuned exploiting information provided by experimental data. Addressing this goal our team is improving tools for the analysis of both electrophysiological (De Bonis et al., 2019) and optical (Celotto et al., 2018) recordings (micro-ECoG and wide-field Calcium Imaging). The main aim is the spatio-temporal characterization of SWA. We also plan to apply inference methods to obtain refined maps of connectivity and excitability for insertion in the simulated model.

## DATA AVAILABILITY

The source code of the DPSNN engine and the data that support the findings of this study are openly available in GitHub at <https://github.com/APE-group/201903LargeScaleSimScaling>.

The DPSNN code also corresponds to the internal svn release 961 of the APE group repository.

## AUTHOR CONTRIBUTIONS

EP, PP, and FS improved the simulation engine, measured and analyzed the scaling. CC, MM, and PD provided the cortical model. MS-V provided the experimental data for calibration of simulations. All authors contributed to the final version of the manuscript and approved it for publication.

## FUNDING

This research has received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreements No. 785907 (HBP SGA2), No. 720270 (HBP SGA1), and No. 671553 (EXANEST).

## REFERENCES

- Bazhenov, M., Timofeev, I., Steriade, M., and Sejnowski, T. J. (2002). Model of thalamocortical slow-wave sleep oscillations and transitions to activated states. *J. Neurosci.* 22, 8691–8704. doi: 10.1523/JNEUROSCI.22-19-08691.2002
- Brunel, N., and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.* 11, 1621–1671.
- Capone, C., and Mattia, M. (2017). Speed hysteresis and noise shaping of traveling fronts in neural fields: role of local circuitry and nonlocal connectivity. *Sci. Rep.* 7:39611. doi: 10.1038/srep39611
- Capone, C., Pastorelli, E., Golosio, B., and Paolucci, P. S. (2019a). Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model. *Sci. Rep.* 9. doi: 10.1038/s41598-019-45525-0
- Capone, C., Rebollo, B., Muñoz, A., Illa, X., Del Giudice, P., Sanchez-Vives, M. V., et al. (2019b). Slow waves in cortical slices: how spontaneous activity is shaped by laminar structure. *Cereb. Cortex* 29, 319–335. doi: 10.1093/cercor/bhx326
- Carnevale, N. T., and Hines, M. L. (2006). *The NEURON Book*. Cambridge, UK: Cambridge University Press.
- Celotto, M., De Luca, C., Muratore, P., Resta, F., Allegra Mascaro, A. L., Pavone, F. S., et al. (2018). Paolucci, analysis and model of cortical slow waves acquired with optical techniques. *arXiv:1811.11687*.
- Coombes, S. (2005). Waves, bumps, and patterns in neural field theories. *Biol. Cybern.* 93, 91–108. doi: 10.1007/s00422-005-0574-y
- Curto, C., Sakata, S., Marguet, S., Itskov, V., and Harris, K. D. (2009). A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J. Neurosci.* 29, 10600–10612. doi: 10.1523/JNEUROSCI.2053-09.2009
- De Bonis, G., Dasilva, M., Pazienti, A., Sanchez-Vives, M. V., Mattia, M., and Paolucci, P. S. (2019). Slow waves analysis pipeline for extracting features of slow oscillations from the cerebral cortex of anesthetized mice. *arXiv:1902.08599*.
- Destexhe, A., and Contreras, D. (2011). “The fine structure of slow-wave sleep oscillations: from single neurons to large networks,” in *Sleep and Anesthesia*, Chapter 4, ed A. Hutt (New York, NY: Springer New York), 69–105.
- Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., et al. (2013). Overview of the SpiNNaker system architecture. *IEEE Trans. Comput.* 62, 2454–2467. doi: 10.1109/TC.2012.142
- Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430
- Gigante, G., Mattia, M., and Giudice, P. D. (2007). Diverse population-bursting modes of adapting spiking neurons. *Phys. Rev. Lett.* 98:148101. doi: 10.1103/PhysRevLett.98.148101

## ACKNOWLEDGMENTS

Large-scale simulations have been performed on the Galileo platform, provided by the CINECA in the frameworks of HBP SGA 1 and 2 and of the INFN-CINECA collaboration on the Computational Theoretical Physics Initiative. We acknowledge Hans Ekkehard Plesser and Dimitri Plotnikov for their support in setting up NEST simulations. We are grateful to the members of the INFN APE Parallel/Distributed Computing laboratory for their strenuous support. This manuscript has been released as a pre-print at Pastorelli et al. (2019).

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnsys.2019.00033/full#supplementary-material>

- Goodman, D., and Brette, R. (2009). The brian simulator. *Front. Neurosci.* 3:26. doi: 10.3389/neuro.01.026.2009
- Han, F., Caporale, N., and Dan, Y. (2008). Reverberation of recent visual experience in spontaneous cortical waves. *Neuron* 60, 321–327. doi: 10.1016/j.neuron.2008.08.026
- Hill, S. L., and Tononi, G. (2005). Modeling sleep and wakefulness in the thalamocortical system. *J. Neurophysiol.* 93, 1671–1698. doi: 10.1152/jn.00915.2004
- Hines, M. L., and Carnevale, N. T. (1997). The neuron simulation environment. *Neural Comput.* 9, 1179–1209.
- Hobson, J. A., and Pace-Schott, E. F. (2002). The cognitive neuroscience of sleep: neuronal systems, consciousness and learning. *Nat. Rev. Neurosci.* 3, 679–693. doi: 10.1038/nrn915
- Izhikevich, E. M., and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. U.S.A.* 105, 3593–3598. doi: 10.1073/pnas.0712231105
- Jordan, J., Ippen, T., Helias, M., Kitayama, I., Sato, M., Igarashi, J., et al. (2018). Extremely scalable spiking neuronal network simulation code: from laptops to exascale computers. *Front. Neuroinform.* 12:2. doi: 10.3389/fninf.2018.00002
- Katevenis, M., Chrysos, N., Marazakis, M., Mavroidis, I., Chaix, F., Kallimanis, N., et al. (2016). “The exanest project: interconnects, storage, and packaging for exascale systems,” in *2016 Euromicro Conference on Digital System Design (DSD)*, 60–67.
- Krishnan, G. P., Chauvette, S., Shamie, I., Soltani, S., Timofeev, I., Cash, S. S., et al. (2016). Cellular and neurochemical basis of sleep stages in the thalamocortical network. *eLife* 5, 1–29. doi: 10.7554/eLife.18607
- Lazzaro, J., Wawrzyniec, J., Mahowald, M., Sivillotti, M., and Gillespie, D. (1993). Silicon auditory processors as computer peripherals. *IEEE Trans. Neural Netw.* 4, 523–528.
- Luczak, A., Barthó, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425. doi: 10.1016/j.neuron.2009.03.014
- Mattia, M., and Del Giudice, P. (2000). Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. *Neural Comput.* 12, 2305–2329. doi: 10.1162/089976600300014953
- Mattia, M., and Del Giudice, P. (2002). Population dynamics of interacting spiking neurons. *Phys. Rev. E* 66:051917. doi: 10.1103/PhysRevE.66.051917
- Mattia, M., and Sanchez-Vives, M. V. (2012). Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity. *Cognit. Neurodyn.* 6, 239–250. doi: 10.1007/s11571-011-9179-4
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with

- a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Modha, D. S., Ananthanarayanan, R., Esser, S. K., Ndirango, A., Sherbondy, A. J., and Singh, R. (2011). Cognitive computing. *Commun. ACM* 54, 62–71. doi: 10.1145/1978542.1978559
- Morrison, A., Mehring, C., Geisel, T., Aertsen, A., and Diesmann, M. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.* 17, 1776–1801. doi: 10.1162/0899766054026648
- Nageswaran, J. M., Dutt, N., Krichmar, J. L., Nicolau, A., and Veidenbaum, A. V. (2009). A configurable simulation environment for the efficient simulation of large-scale spiking neural networks on graphics processors. *Neural Netw.* 22, 791–800. doi: 10.1016/j.neunet.2009.06.028
- Paolucci, P. S., Ammendola, R., Biagioni, A., Frezza, O., Lo Cicero, F., Lonardo, A., et al. (2013). Distributed simulation of polychronous and plastic spiking neural networks: strong and weak scaling of a representative mini-application benchmark executed on a small-scale commodity cluster. *arXiv:1310.8478*.
- Paolucci, P. S., Biagioni, A., Murillo, L. G., Rousseau, F., Schor, L., Tosoratto, L., et al. (2016). Dynamic many-process applications on many-tile embedded systems and HPC clusters: the EURETILE programming environment and execution platforms. *J. Syst. Archit.* 69, 29–53. doi: 10.1016/j.sysarc.2015.11.008
- Pastorelli, E., Capone, C., Simula, F., Sanchez-Vives, M. V., Del Giudice, P., Mattia, M., et al. (2019). Scaling of a large-scale simulation of synchronous slow-wave and asynchronous awake-like activity of a cortical model with long-range interconnections. *arXiv:1902.08410*.
- Potjans, T. C., and Diesmann, M. (2014). The cell-type specific cortical microcircuit: Relating structure and activity in a full-scale spiking network model. *Cereb. Cortex* 24, 785–806. doi: 10.1093/cercor/bhs358
- Reyes-Puerta, V., Yang, J.-W., Siwek, M. E., Kilb, W., Sun, J.-J., and Luhmann, H. J. (2016). Propagation of spontaneous slow-wave activity across columns and layers of the adult rat barrel cortex *in vivo*. *Brain Struct. Funct.* 221, 4429–4449. doi: 10.1007/s00429-015-1173-x
- Ricciardi, L. M. (1977). *Diffusion Processes and Related Topics in Biology*. Berlin; Heidelberg; New York, NY: Springer-Verlag Berlin Heidelberg.
- Ruiz-Mejias, M., Ciria-Suarez, L., Mattia, M., and Sanchez-Vives, M. V. (2011). Slow and fast rhythms generated in the cerebral cortex of the anesthetized mouse. *J. Neurophysiol.* 106, 2910–2921. doi: 10.1152/jn.00440.2011
- Sanchez-Vives, M., and Mattia, M. (2014). Slow wave activity as the default mode of the cerebral cortex. *Arch. Ital. Biol.* 152, 147–155. doi: 10.12871/000298292014239
- Sanchez-Vives, M. V., and McCormick, D. A. (2000). Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat. Neurosci.* 3:1027. doi: 10.1038/79848
- Schmitt, S., Klähn, J., Bellec, G., Grübl, A., Güttler, M., Harte, A., et al. (2017). “Neuromorphic hardware in the loop: training a deep spiking network on the BrainScaleS wafer-scale system,” in *2017 International Joint Conference on Neural Networks (IJCNN)* (Anchorage, AK: IEEE), 2227–2234. doi: 10.1109/IJCNN.2017.7966125
- Simula, F., Pastorelli, E., Paolucci, P. S., Martinelli, M., Lonardo, A., Biagioni, A., et al. (2019). “Real-time cortical simulations: energy and interconnect scaling on distributed systems,” in *2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)* (Pavia), 283–290.
- Solovey, G., Alonso, L. M., Yanagawa, T., Fujii, N., Magnasco, M. O., Cecchi, G. A., et al. (2015). Loss of consciousness is associated with stabilization of cortical activity. *J. Neurosci.* 35, 10866–10877. doi: 10.1523/JNEUROSCI.4895-14.2015
- Steyn-Ross, M. L., Steyn-Ross, D. A., and Sleight, J. W. (2013). Interacting Turing-Hopf instabilities drive symmetry-breaking transitions in a mean-field model of the cortex: a mechanism for the slow oscillation. *Phys. Rev. X* 3:21005. doi: 10.1103/PhysRevX.3.021005
- Stimberg, M., Goodman, D., Benichoux, V., and Brette, R. (2014). Equation-oriented specification of neural models for simulations. *Front. Neuroinform.* 8:6. doi: 10.3389/fninf.2014.00006
- Strogatz, S. H. (2018). *Nonlinear Dynamics and Chaos With Student Solutions Manual: With Applications to Physics, Biology, Chemistry, and Engineering*. Boca Raton, FL: CRC Press.
- Stroh, A., Adelsberger, H., Groh, A., Rühlmann, C., Fischer, S., Schierloh, A., et al. (2013). Making waves: initiation and propagation of corticothalamic Ca<sup>2+</sup> waves *in vivo*. *Neuron* 77, 1136–1150. doi: 10.1016/j.neuron.2013.01.031
- Wester, J. C., and Contreras, D. (2012). Columnar interactions determine horizontal propagation of recurrent network activity in neocortex. *J. Neurosci.* 32, 5454–5471. doi: 10.1523/JNEUROSCI.5006-11.2012
- Wilson, M. A., Bhalla, U. S., Uhley, J. D., and Bower, J. M. (1989). “Genesis: a system for simulating neural networks,” in *Advances in Neural Information Processing Systems 1*, ed D. S. Touretzky (San Francisco, CA: Morgan-Kaufmann), 485–492.

**Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Pastorelli, Capone, Simula, Sanchez-Vives, Del Giudice, Mattia and Paolucci. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.