



An FPGA-Based Massively Parallel Neuromorphic Cortex Simulator

Runchun M. Wang^{1*}, Chetan S. Thakur² and André van Schaik¹

¹ The MARCS Institute, University of Western Sydney, Sydney, NSW, Australia, ² Department of Electronic Systems Engineering, Indian Institute of Science, Bangalore, India

OPEN ACCESS

Edited by:

Gert Cauwenberghs,
University of California, San Diego,
United States

Reviewed by:

Shantanu Chakrabartty,
Washington University in St. Louis,
United States
Jongkil Park,
Electronics and Telecommunications
Research Institute, South Korea
Dan Hammerstrom,
Portland State University,
United States
Theodore Yu,
Texas Instruments, United States

*Correspondence:

Runchun M. Wang
mark.wang@westernsydney.edu.au

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 12 August 2017

Accepted: 16 March 2018

Published: 10 April 2018

Citation:

Wang RM, Thakur CS and van
Schaik A (2018) An FPGA-Based
Massively Parallel Neuromorphic
Cortex Simulator.
Front. Neurosci. 12:213.
doi: 10.3389/fnins.2018.00213

This paper presents a massively parallel and scalable neuromorphic cortex simulator designed for simulating large and structurally connected spiking neural networks, such as complex models of various areas of the cortex. The main novelty of this work is the abstraction of a neuromorphic architecture into clusters represented by minicolumns and hypercolumns, analogously to the fundamental structural units observed in neurobiology. Without this approach, simulating large-scale fully connected networks needs prohibitively large memory to store look-up tables for point-to-point connections. Instead, we use a novel architecture, based on the structural connectivity in the neocortex, such that all the required parameters and connections can be stored in on-chip memory. The cortex simulator can be easily reconfigured for simulating different neural networks without any change in hardware structure by programming the memory. A hierarchical communication scheme allows one neuron to have a fan-out of up to 200 k neurons. As a proof-of-concept, an implementation on one Altera Stratix V FPGA was able to simulate 20 million to 2.6 billion leaky-integrate-and-fire (LIF) neurons in real time. We verified the system by emulating a simplified auditory cortex (with 100 million neurons). This cortex simulator achieved a low power dissipation of 1.62 μ W per neuron. With the advent of commercially available FPGA boards, our system offers an accessible and scalable tool for the design, real-time simulation, and analysis of large-scale spiking neural networks.

Keywords: neuromorphic engineering, neocortex, spiking neural networks, computational neuroscience, stochastic computing

INTRODUCTION

Our inability to simulate neural networks in software on a scale comparable to the human brain (10^{11} neurons, 10^{14} synapses) is impeding our progress toward understanding the signal processing in large networks in the brain and toward building applications based on that understanding. A small-scale linear approximation of a large spiking neural network will not be capable of providing sufficient information about the global behavior of such highly nonlinear networks. Hence, in addition to smaller scale systems with detailed software or hardware neural models, it is necessary to develop a hardware architecture that is capable of simulating neural networks comparable to the human brain in terms of scale, with models with an intermediate level of biological detail, that can simulate these networks quickly, preferably in real time to allow interaction between the simulation and the environment. To this end, we are focusing on a hardware friendly architecture for simulating large-scale and structurally connected spiking neural networks using simple leaky integrate-and-fire (LIF) neurons.

Simulating neural networks on computers has been the most popular method for many decades. Software simulators, such as GENESIS (Bower and Beeman, 1998) and NEURON (Hines and Carnevale, 1997), are biologically accurate and model their components with differential equations and sub-millisecond time steps. This approach introduces tremendous computational costs and hence makes it impractical for simulating large-scale neural networks. Other simulators, such as the NeoCortical simulator (NCS; Hoang et al., 2013), Brian (Goodman, 2008), and Neural Simulation Tool [NEST; (Gewaltig and Diesmann, 2007)] are specifically designed for developing large-scale spiking neural networks. Because these tools perform numerical simulations, they do not scale very well, slowing down considerably for large networks with large numbers of variables. For instance, the Blue Gene rack, a two-million-dollar, 2048-processor supercomputer, takes 1 h and 20 min to simulate 1 s of neural activity in 8 million integrate-and-fire neurons connected by 4 billion static synapses (Wittie and Memelli, 2010). Graphic Processing Units (GPUs) can perform certain types of simulations tens of times faster than a PC (Shi et al., 2015). However, as GPUs are still performing numeric simulations, it can take hours to simulate 1 s of activity in a tiny piece of cortex (Izhikevich and Edelman, 2008).

Along with general hardware solutions, there have been a number of large-scale neuromorphic platforms such as Neurogrid (Benjamin et al., 2014), BrainScaleS (Pfeil et al., 2013), TrueNorth (Merolla et al., 2014), SpiNNaker (Furber et al., 2014), and HiAER-IFAT (Park et al., 2016). In Neurogrid, sub-threshold analog circuits are used to model neuron and synapse dynamics in biological real time, with digital spike communication. A 16-chip board is capable of simulating one million neurons, using a two-compartment model with ion-channels, in real time. In BrainScaleS, a full wafer implementation, each wafer has 48 reticles with eight High-Count Analog Neural Network (HiCANN) dice each. Each HiCANN die has 512 adaptive exponential integrate and fire (AdExp) neurons and over 100,000 synapses. The HiCANN chip runs 10,000 times faster than real time. In SpiNNaker, ARM processors run software neural models. A 48-node SpiNNaker board is capable of simulating 250,000 neurons and 80 million synapses in real time. A recent work has successfully used the SpiNNaker system to implement a spiking neural network model of the thalamic Lateral Geniculate Nucleus (Sen-Bhattacharya et al., 2017). The IBM TrueNorth chip is designed for building large-scale neural networks and a 16-chip board is capable of running 16 million leaky-integrate-and-fire (LIF) neurons in real time. The HiAER-IFAT system has five FPGAs and four custom analog neuromorphic integrated circuits, yielding 262 k neurons and 262 M synapses. The full-size HiAER-IFAT network has four boards, each of which has one IFAT module, serving 1 M neurons and 1 G synapses. However, all these platforms are expensive to build and require proprietary hardware and will not be easily accessible to the computational neuroscience community.

Modern FPGAs provide a large number of logic gates and physical memory, allowing large-scale neural networks to be created at a low cost. Even state-of-the-art FPGAs are affordable for research laboratories. Thus, for the past decade several projects have advanced our knowledge on how to use FPGAs

to simulate neural networks. The BenNuey platform which comprises up to 18 M neurons and 18 M synapses was proposed by Maguire et al. (2007). The EU SCANDLE project created a system with 1 M neurons (Cassidy et al., 2011). Bailey et al. proposed a behavioral simulation and synthesis of biological neuron systems using synthesizable VHDL in Bailey et al. (2011). The Bluehive project (Moore et al., 2012) has achieved a neural network with up-to 256 k Izhikevich neurons (Izhikevich, 2003) and 256 M synapses. A structured AER system, comprising 64 convolution processors, which is equivalent to a neural network with 262 k neurons and almost 32 M synapses, has been proposed by Zamarreno-Ramos et al. (2012). A large-scale spiking neural network accelerator, which is capable of simulating ~400 k AdExp neurons in real-time was proposed by Cheung et al. (2016). A recent system, capable of simulating a fully connected network with up to 1440 Izhikevich neurons, was proposed in Pani et al. (2017). We have previously presented a design framework capable of simulating 1.5 million LIF neurons in real time (Wang et al., 2014c). In that work, the FPGA emulates the biological neurons instead of performing numeric simulations as GPUs perform.

These developments strongly indicate the great potential for using FPGAs to create systems that allow research into large and complex models. Here, we present several advanced approaches based on biological observations to scale our previous system up to 2.6 billion neurons in real time. Though these models may not necessarily be at the level of detail neuroscientists would need yet, our system will present important building blocks and make important contributions in that direction. We will continue work on this system, which will eventually evolve to handle models of the complexity needed, so that our system will become a valuable neuroscience tool.

MATERIALS AND METHODS

Strategy

To emulate the neocortex (hereafter “cortex”) with billions of neurons and synapses in real time, two major problems need to be addressed: the computation problem and the communication problem. The second problem is indeed the major obstacle and far more difficult to address than the first one. Our strategy is to follow fundamental findings observed in the cortex.

Modular Structure

The cortex is a structure composed of a large number of repeated units, neurons and synapses, each with several sub-types. There are six layers in the cortex; layer I is the most superficial layer and mainly contains dendritic arborisation of neurons in the underlying layers and incoming fibers from other brain areas. The distribution of neurons in the different layers is on average: 5% in layer I; 35% in layer II/III; 20% in layer IV; and 40% in layer V/VI (Dombrowski, 2001). Input to the cortex, mainly from the thalamic region, enters in layer IV; however, even in this layer, <10% of the afferent connections come from the thalamic region, the rest being corticocortical afferents (Martin, 2002).

A minicolumn is a vertical volume of cortex with about 100 neurons that stretches through all layers of the cortex (Buxhoeveden and Casanova, 2002). Each minicolumn contains excitatory neurons, mainly pyramidal and stellate cells, inhibitory interneurons, and a large number of internal and external connections. The minicolumn is often considered to be both a functional and anatomical unit of the cortex (Buxhoeveden, 2002), and we use this minicolumn with 100 neurons as the basic building block of the cortex simulator.

As there exist some differences between minicolumns located in different parts of the cortex (such as exact size, structure, and active neurotransmitters), the minicolumn in the cortex simulator is designed to have up to eight different programmable types of neurons. The number of each type of neuron is constrained to be a multiple of four and their parameters are fully configurable while maintaining a small memory footprint. Note, the neuron types do not necessarily correspond to the cortical layers, but can be configured as such.

In the mammalian cortex, minicolumns are grouped into modules called hypercolumns (Hubel and Wiesel, 1977). These are the building blocks for complex models of various areas of the cortex (Johansson and Lansner, 2007). We therefore use the hypercolumn as a functional grouping for our simulator. Biological measurements suggest that a hypercolumn typically consists of about 100 minicolumns (Johansson and Lansner, 2007). The hypercolumn in our cortex simulator is designed to have up to 128 minicolumns. Similar to the minicolumns, the parameters of the hypercolumn are designed to be fully configurable.

Emulating Dynamically

To solve the extensive computational requirement for simulating large networks, we use two approaches. First, it is not necessary to implement all neurons physically on silicon and we can use time multiplexing to leverage the high-speed of the FPGA (Cassidy et al., 2011; Wang et al., 2014a,b, 2017). State-of-the-art FPGAs can easily run at a clock speed of 200 MHz (i.e., a clock period of 5 ns). Therefore, we can time-multiplex a single physical minicolumn (100 physical neurons in parallel) to simulate 200 k time-multiplexed (TM) minicolumns, each one updated every millisecond. Using a pipelined architecture, the result of calculating the update for one time step for a TM neuron only has to be available just before that TM neuron's turn comes around again 1 ms later. Thus, the time to compute a single TM neuron's output is not a limiting factor in the size of the network. For neural network applications that would benefit from running much faster or slower than biological neurons, we can trade off speed and multiplexing ratio (and thus network size). Limited by the hardware resources (mainly the memory), our cortex simulator was designed to be capable of simulating up to 200 k TM minicolumns in real time and 1 M (2^{20}) TM minicolumns at five times slower than real time, i.e., an update every 5 ms.

Second, based on the physiological metabolic cost of neural activity, it has been concluded that fewer than 1% of neurons are active in the brain at any moment (Lennie, 2003). Based on the experimental data in Tsunoda et al.

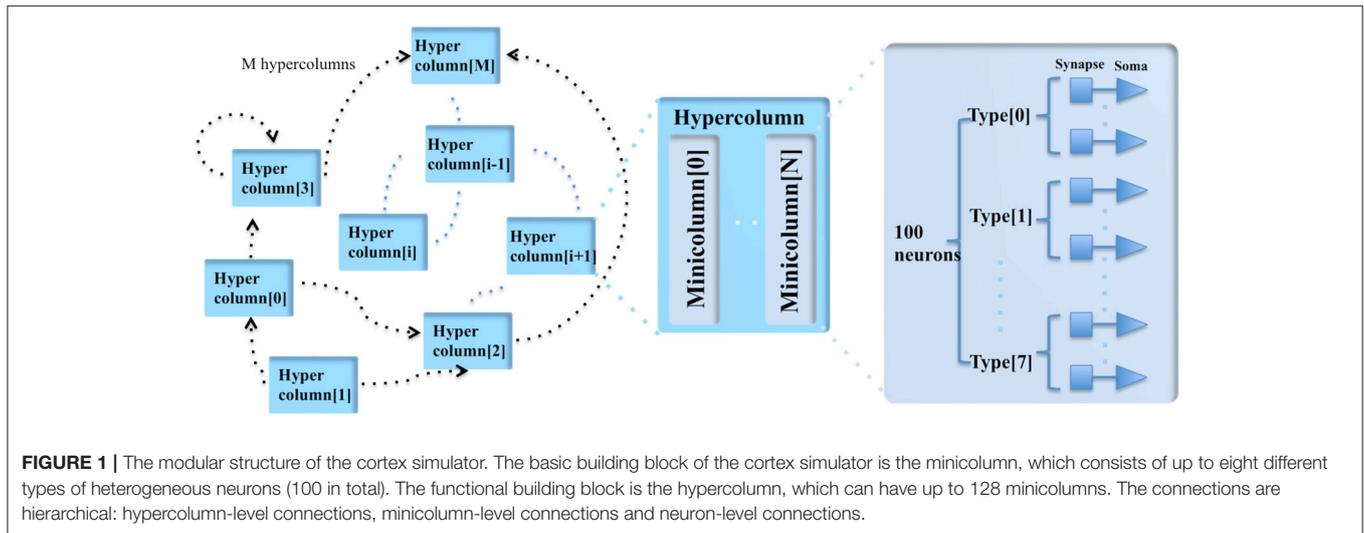
(2001), Johansson and Lansner (2007) concluded that at most a few percent of the hypercolumns and hence only about 0.01% of the minicolumns and neurons are active in a functional sense (integrating and firing) at any moment in the cortex. Hence, in principle, one hardware minicolumn could be dynamically reassigned to simulate 10^4 minicolumns on average. Such a hardware minicolumn will be referred to as a physical minicolumn and the minicolumn to be simulated will be referred to as a dynamically assigned (DA) minicolumn. If a DA minicolumn cannot be simulated in a single time step, the physical minicolumn needs to be assigned to that DA minicolumn for a longer time and the number of DA minicolumns that can be simulated will decrease proportionally.

Theoretically, a TM minicolumn array with 1 M TM minicolumns can be dynamically assigned for $1 \text{ M} \times 10^4 = 10^{10}$ minicolumns, if these minicolumns can be updated every 5 ms and if only 0.01% of the minicolumns are active at any time step. To be able to deal with much higher activity rates, we chose to dynamically assign one TM minicolumn for 128 DA minicolumns. The maximum active rate of the minicolumns that this system can support is therefore $1/128 \approx 0.7\%$, allowing it to support $128 \times 1 \text{ M} = 128 \text{ M}$ (2^{27}) minicolumns (each has a unique 27-bit address). This means the simulator is capable of simulating up to 1 M hypercolumns, each of which has up to 128 minicolumns. We can trade off the active rate and the network size when needed. This dynamic-assignment approach has been validated in our previous work (Wang et al., 2013b, 2014, 2015). Applying this approach with nanotechnology has been fully explored in Zaveri and Hammerstrom (2011).

Hierarchical Communication

Our cortex simulator uses a hierarchical communication scheme (see **Figure 1**) such that the communication cost between the neurons can be reduced by orders of magnitude. It is impractical to use point-to-point connections for cortex-level simulation, as it would require hundreds of terabytes of memory to specify the connections and their weights. Anatomical studies of the cortex presented in Thomson (2003) showed that cortical neurons are not randomly wired together and that the connections are quite structural. Connections of the minicolumns are highly localized so that the connectivity between two pyramidal neurons $<25\text{--}50 \mu\text{m}$ apart is high and the connectivity between two neurons drops sharply with distance (Holmgren et al., 2003). Furthermore, pyramidal neurons in layer II/III and V/VI within a minicolumn are reciprocally connected (Thomson, 2003).

Based on these findings, we chose to store the connection types of the neurons, the minicolumns, and the hypercolumns in a hierarchical fashion instead of individual point-to-point connections. In this scheme, the addresses of the events consist of hypercolumn addresses and minicolumn addresses. Both of them are generated on the fly with connection parameters according to their connection levels, respectively. For example, we can specify that hypercolumns 1 to 10000 are each connected with random weights to the 10 hypercolumns closest to them, while hypercolumns 20000 to 40000 are each connected their nearest 16 hypercolumns with fixed weights. Proximity is defined by



the distance of the hypercolumn index in address space. This method only requires several kilobytes of memory and can be easily implemented with on-chip SRAMs. We will present the details of this hierarchical scheme in section Parameter LUT.

Inspired by observations from neurobiology, where neurons and their connections often form clusters to create local cortical microcircuits (Bosking et al., 1997), the communication between the neurons uses events (spike counts) instead of individual spikes. This arrangement models a cluster of synapses formed by an axon onto the dendritic branches of nearby neurons. The neurons of one type within a minicolumn all receive the same events, which are the numbers of the spikes generated by one type of neuron in the source minicolumns within a time step.

One minicolumn has up to eight types of neurons, and each type can be connected to any type of neuron in the destination minicolumns. Hence, there are up to 64 synaptic connections possible between any two minicolumns. We impose that a source minicolumn has the same number of connections to all of the other minicolumns within the same hypercolumn, but that these can have different synaptic weights. The primary advantage of using this scheme is that it overcomes a key communication bottleneck that limits scalability for large-scale spiking neural network simulations. Otherwise, we would be required to replicate spikes for each group of target neurons.

Our system allows the events generated by one minicolumn to be propagated to up to 16 hypercolumns, each of which has up to 128 minicolumns, i.e., to $16 \times 128 \times 100 = 200$ k neurons. Each of these 16 connections has a configurable fixed axonal delay (from 1 to 16 ms, with a 1 ms step). Since a hypercolumn has up to 128 minicolumns, a hypercolumn can communicate with far more than 16 hypercolumns. Further details will be presented in section Synapse Array.

There two major differences between our scheme and the hierarchical address-event routing scheme used by the HiAER-IFAT system. First, HiAER routes each individual spike while we route the events (the numbers of the spikes generated by one type of neuron (in a minicolumn)). Second, HiAER

uses external memory to store its look up tables (LUTs) such that it is capable of supporting point-to-point connections, each of which has its own programmable parameters, e.g., synaptic weight and axonal delay. This allows HiAER-IFAT to be a general purpose neuromorphic platform for spike-based algorithms. Its scalability is therefore constrained by the size of the external memories. While our scheme takes the advantage of the structural connections of the cortex, we use the on-chip SRAMs to implement LUTs that only store the connection patterns.

Architecture

The cortex simulator was deliberately designed to be scalable and flexible, such that the same architecture could be implemented either on a standalone FPGA board or on multiple parallel FPGA boards. As a proof-of-concept, we have implemented this architecture on a Terasic DE5 kit (with one Altera Stratix V FPGA, two DDR3 memories, and four QDRII memories) as a standalone system. **Figure 2** shows its architecture, consisting of a neural engine, a Master, off-chip memories, and a serial interface.

The neural engine forms the main body of the system. It contains three functional modules: a minicolumn array, a synapse array, and an axon array. The minicolumn array implements TM minicolumns. The axon array will propagate the events generated by the minicolumns with axonal delays to the synapse array. In the synapse array, these events will be modulated with synaptic weights and will be assigned their destination minicolumn address. The synapse array will send these events to the destination minicolumn array in an event-driven fashion. Besides these functional modules, there is a parameter look-up table (LUT), which stores the neuron parameters, connection types, and connection parameters. We will present the details of these modules in the following sections.

Because of the complexity of the system and large number of the events, each module in the neural engine was designed to be a slave module, such that a single Master has full control of the emulation progress. The Master has a memory bus controller that

will control the access of the external memories. As we use time multiplexing to implement the minicolumn array, we will have to store the neural state variables of each TM neuron (such as their membrane potentials). These are too big to be stored in on-chip memory and have to be stored in off-chip memory, such as the DDR memory. Using off-chip memory needs flow control for the memory interface, which makes the architecture of the system significantly more complex, especially if there are multiple off-chip memories. The axon array also needs to access the off-chip memories for storing events. We will present the details of the Memory Bus Controller in section Memory Bus Controller.

The Master also has an external interface module that will perform flow control for external input and output. This module

also takes care of instruction decoding. We will present the details of External Interface in section External Interface. The serial interface is a high-speed interface, such as the PCIe interface, that communicates with the other FPGAs and the host PC. It is board-dependent, and in the work presented here, we use Altera's 10G base PHY IP.

Minicolumn Array

The minicolumn array (see Figure 3) consists of a neuron-type manager, an event generator, and the TM minicolumns, which are implemented by time-multiplexing a single physical minicolumn consisting of 100 physical neurons. These neurons

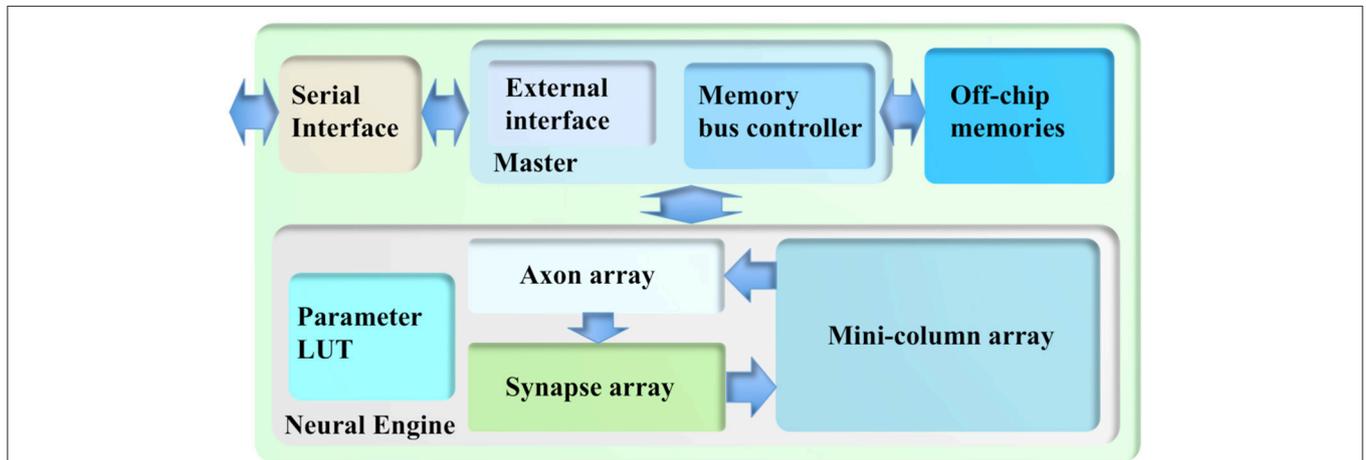


FIGURE 2 | The architecture of the cortex simulator. The system consists of a neural engine, a Master, off-chip memories and a serial interface. The neural engine realizes the function of biological neural systems by emulating their structures. The Master controls the communication between the neural engine and the off-chip memories, which store the neural states and the events. The serial interface is used to interact with the other FPGAs and the host controller, e.g., PCs.

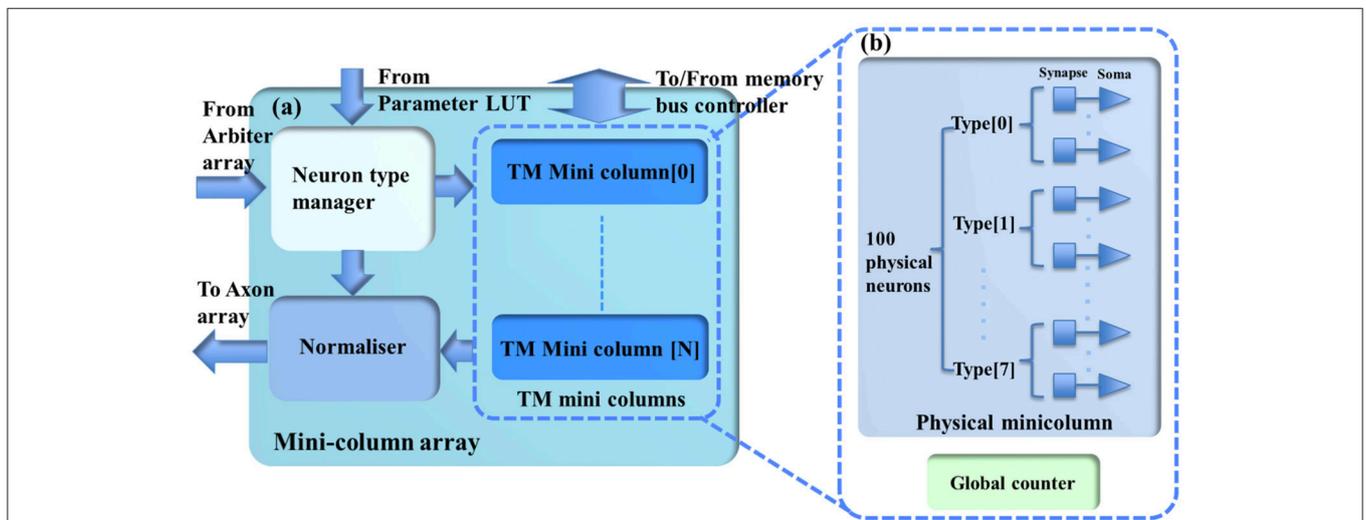


FIGURE 3 | (a) The structure of the minicolumn array. (b) The internal structure of the TM minicolumns. The TM minicolumns are implemented with the time-multiplexing approach and (b) shows its internal structure. It consists of a physical minicolumn and a global counter. The global counter processes the TM minicolumns sequentially. The physical minicolumns consists of 100 physical neurons.

will generate positive (excitatory) and negative (inhibitory) post-synaptic currents (EPSCs and IPSCs) from input events weighted in the synapse array. These PSCs are integrated in the cell body (the soma). The soma performs a leaky integration of the PSCs to calculate the membrane potential and generates an output spike (post-synaptic spike) when the membrane potential passes a threshold, after which the membrane potential is reset and enters a refractory period. Events, i.e., spike counts, are sent to the axon array together with the addresses of the originating minicolumns, the number of connections, and axonal delay values for each connection (between two minicolumns).

Physical Neuron

Rather than using a mathematical computational model with floating-point numbers, the physical neuron has been efficiently implemented using fixed-point numbers to minimize the number of bits that need to be stored for each state variable. The minimization of memory use was effectively achieved using a stochastic method to implement exponential decays for the conductance-based neuron model (Wang et al., 2014c, 2016). We store only the most significant bits (MSBs) and generate the least significant bits (LSBs) on the fly with a random number generator. This method not only reduces the storage needed, but also introduces variability between the (TM) neurons while using the exact same physical neuron model. This makes the network simulations more realistic. In existing implementations, digital neurons are usually perfectly matched, so that two neurons in the network with identical initial conditions receiving the same input will remain synchronized forever.

The stochastic conductance-based neuron is a compact design using only fixed-point numbers. It has a single post-synaptic current generator (see **Figure 4**), which can generate both EPSCs and IPSCs, and a soma to integrate the post-synaptic currents. The PSC generator consists of two multipliers, an adder, a comparator and two multiplexers. Its function is expressed by the following equation:

$$PSC(t + 1) = PSC(t) \times \frac{\tau_{psc}}{\tau_{psc} + 1} + r(t) + g_{syn} \times W(t) \quad (1)$$

where $t + 1$ represents the index of the current time step, $PSC(t)$ represents the value of the PSC (from the memory), and $W(t)$ represents the linear accumulation of the weighted synaptic input from the pre-synaptic events within previous time step. Both $PSC(t)$ and $W(t)$ are signed 4-bit numbers and normalized in the range $[-1,1]$. τ_{PSC} represents the time constant, i.e., 10 ms, controlling the speed with which PSC will decay to 0 exponentially. The PSC generator will use τ_{EPSC} and τ_{IPSC} for EPSC and IPSC, respectively. $r[t]$ is a random number drawn from a uniform distribution in the range $(0,1)$, and it will be different at different time steps, even for the same TM neuron. g_{syn} is the synaptic gain, and is an 8-bit number normalized in the range $[0.06,16]$.

In the hardware implementation, the time constants are efficiently implemented by applying a simple shift operation to the product of $PSC(t)$ and a leakage rate L (an 8-bit number, $L/256 \approx \tau/(\tau+1)$). r is a 5-bit random number and the maximum time constant it can achieve is 30 ms (30 time steps). Its implementation can easily be modified if longer time constants are needed.

The soma is also a stochastic conductance-based model similar to the PSC generator. Its function is expressed similarly by the following equation:

$$V_{mem}(t + 1) = V_{mem}(t) \times \frac{\tau_{soma}}{\tau_{soma} + 1} + r(t) + g_{psc} \times PSC(t + 1) \quad (2)$$

where $V_{mem}(t)$ represents the previous value of the membrane voltage (from the memory). The soma has two states: active state and the refractory state. The soma can only integrate $PSC(t+1)$ in its active state: its membrane voltage is greater than its initial value (a configurable parameter). Otherwise, the soma is in its refractory state and the PSC will be discarded. τ_{soma} represents the time constant controlling the speed with which V_{mem} will decay to the initial value exponentially. The soma will use τ_{mem} and τ_{rfc} for the active and the refractory period, respectively. $r[t]$ is also a random number drawn from a uniform distribution in the range $(0,1)$. g_{psc} is the synaptic gain, and it is an 8-bit number

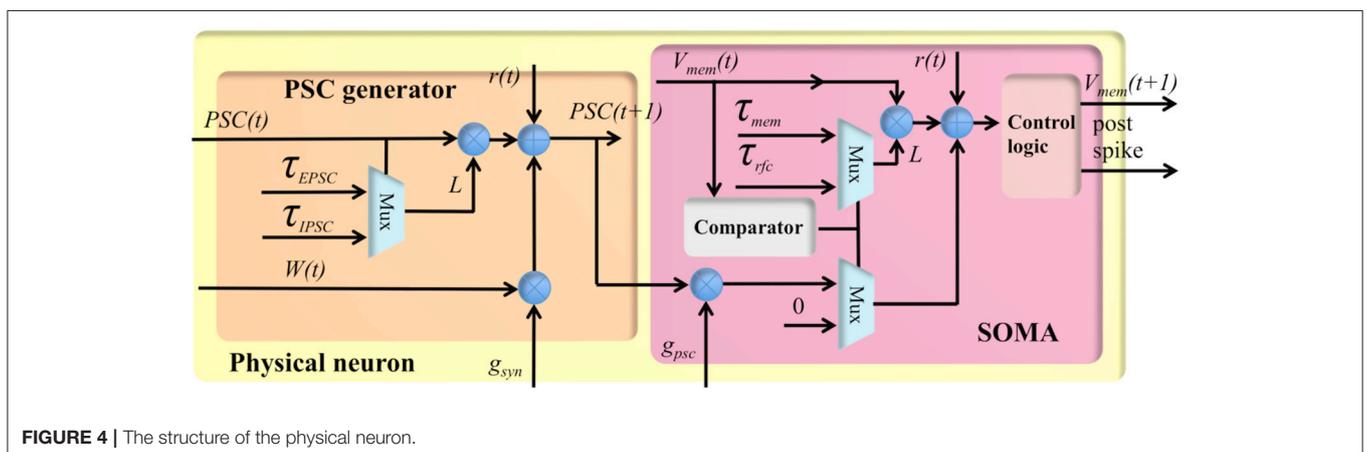


FIGURE 4 | The structure of the physical neuron.

normalized in the range [0.06,16]. The soma will generate a post-synaptic spike when V_{mem} overflows and $PSC(t+1)$ is an EPSC. V_{mem} will be reset, when there is an overflow/underflow caused by EPSC or IPSC, respectively.

The physical neuron has an 11-stage pipeline without halt. Each TM neuron will access each computing module such as the adder and the multipliers for only one clock cycle. In one clock cycle, the adder and the multipliers are all being used, but by different TM neurons. The overhead of the pipeline is negligible, i.e., 22 clock cycles: the first 11 cycles for setting up the pipeline and the last 11 cycles for waiting for the last TM neuron to finish computing.

Time Multiplexing

Time multiplexing enables implementing large-scale spiking neural networks, while requiring only a few physical neurons, by leveraging the high-speed of the FPGA. The bottleneck for time multiplexing is data storage. The on-chip memory is typically highly limited in size (generally only several megabytes). In our system, one physical neuron requires 8 bits and one TM minicolumn thus requires 800 bits. Hence, a network with 100 M TM minicolumns would require 800 Mb, which is too big for on-chip memory. We therefore need to use off-chip memory to store these variables. The use of off-chip memory, however, will be limited by the communication bandwidth and will need flow control for the memory interface, which will make the architecture of the system more complex.

We have developed a time-multiplexing approach that is capable of using off-chip memory efficiently. Here, the whole TM system is segmented into blocks of 1024 (1k) TM neurons. A ping-pong buffer is used in the memory interface such that the system will run one segment at a time, and then pause until the new values from the off-chip memory are ready for the next segment. This allows us to use the external DDR memory in burst mode, which has maximum bandwidth. Thus, the size of a segment was chosen to be equal to the burst size of the DDR memory. Further details about

the memory interface will be presented in section Memory Bus Controller.

Neuron-Type Manager

As each minicolumn can have up to eight different types of neurons, a neuron-type manager is introduced to assign the parameters and weights to each individual neuron. The neuron parameters, which are stored in the parameter LUT, include following parts:

1. The number of types of neurons in the minicolumn
2. The number of each type of neuron
3. The parameters of each type of neuron

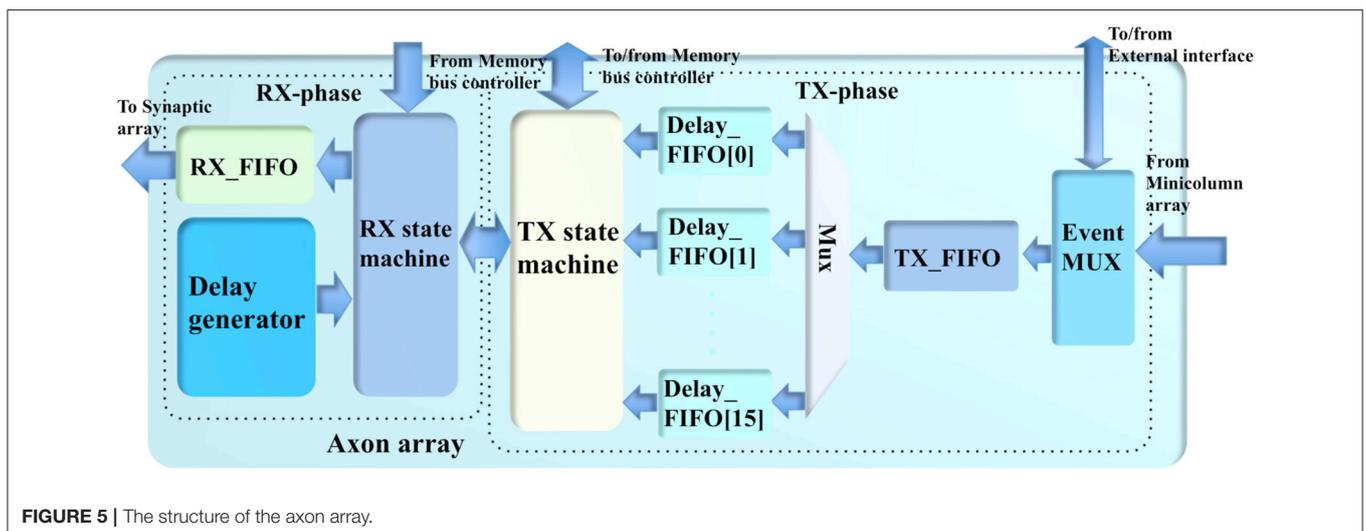
These functions are efficiently implemented with an 8-to-25 multiplexer (25 = 100/4, each minicolumn has 100 neurons and the number of each type of neuron is constrained to a multiple of 4). The neuron-type manager will use linear feedback shift registers (LFSRs) to generate random numbers $[r(t)]$ for each physical neuron such that the neurons within a TM minicolumn, even those with identical parameters, are heterogeneous.

Events Generator

For each type of neuron within a TM minicolumn, the spikes generated by the physical neurons are summed to produce spike counts. Since the source neuron type and the destination neuron types could have different numbers of neurons, these spike counts will be limited to 4-bit numbers. This summation is implemented with eight parallel adders, each of which has 25 inputs, in a 3-stage pipeline. The number of each neuron type is obtained from the neuron-type manager via shift registers, which were designed to work with the pipelines of the physical neuron.

Axon Array

The axon array (see **Figure 5**) propagates the events from the minicolumn array or from the external interface to the synapse array, using programmable axonal delays. To implement this function on hardware, we used a two-phase scheme comprising a TX-phase (TX for transmit) and a RX-phase (RX for receive). In



the TX-phase, the events are written into different regions of the DDR memories according to their programmable axonal delay values. In the RX-phase, for each desired axonal delay value, the events are read out from the corresponding region of the DDR memories stochastically, such that their expected delay values are approximately equal to the desired ones.

The DDR memories are all single-port devices; hence, this two-phase scheme will access them in an alternating fashion such that these two phases do not try to access the same DDR device at the same time. This is guaranteed because the TX-phase has higher priority for accessing the DDR memories: the TX-phase will choose one of the DDR memories to write the events first, and then the RX-phase will choose one of the remaining DDR memories to read. This implementation was chosen because, if the TX-phase cannot move the events generated by the minicolumn array quickly enough, the Master will have to pause the minicolumn array to avoid an overflow of events.

In the system presented here we have two identical 8 GB DDR memories: DDR_A and DDR_B, each of which has eight regions (~ 1 GB each, one for each axonal delay value). DDR_A is for the eight odd-number axonal delay values, e.g., 1 ms, 3 ms, and so on. DDR_B is for the eight even-number axonal delay values. This arrangement is designed to match the alternating access of the DDR memories. Note, the memory bus controller will only grant the axon array access to the DDR memories when the minicolumn array is not accessing them. Further details of this arrangement will be presented in section Memory Bus Controller.

TX-Phase

The TX-phase consists of an event multiplexer, a TX_FIFO (first in first out), a TX state machine and 16 Delay_FIFOs (one per axonal delay value). The event multiplexer transfers the events from the minicolumn array and the external interface to the TX_FIFO. The event multiplexer also sends the events from the minicolumn array to the external interface in case these events need to be sent to other boards (for a multi-FPGA system). The internal events (from the minicolumn array) will be directly written to the TX_FIFO. The event multiplexer simply needs to hold the external events when there are internal events.

The events from the TX_FIFO are moved to the Delay_FIFOs according to their axonal delay values. To avoid losing events, the Master will pause the minicolumn array when the TX_FIFO is almost full. The Delay_FIFO with the maximum usage will be selected by the TX state machine for the next write operation. This selection will decide the target DDR memory automatically. The TX state machine will try to write all the events in the selected Delay_FIFO to the DDR memory as long as the size does not exceed the allowed burst length.

To completely eliminate any potential collisions with the RX-phase, the TX and RX state machines use a cross-lock handshake scheme such that TX-phase will only be allowed to initiate choosing the Delay_FIFO and start writing the selected DDR memory when the RX-phase is inactive (not reading). Otherwise, the TX-phase will wait for the RX-phase to finish reading. The RX-phase will not initiate until the TX-phase has started, unless all the Delay_FIFOs are empty. The TX state machine will notify

the memory bus controller of the completion of one TX-phase. This procedure will continue to repeat as long as the axon array is allowed to access the memory.

RX-Phase

The RX-phase consists of an RX state machine, an RX_FIFO, and a delay generator. The delay generator will enable the RX state machine to stochastically read the events from the DDR memory at different rates to achieve the desired axonal delay values approximately. For each axonal delay value, the events are read out with a probability P_i , which meets the following conditions:

$$\sum_{i=1}^{16} P_i = 1 \quad (3)$$

$$P_i = P_1/i \quad (4)$$

Equation (4) ensures that the ratios between the mean times for reading the same number of events match the ratios between the axonal delay values. For instance, if it takes T time to read out N events with the axonal delay value of 1 ms, it will take $\sim 16T$ time to read the same number of events with an axonal delay value of 16 ms. This assumption requires two conditions to be true: first, the DDR memories have been filled with sufficient events; second, the read-out operations must be evenly distributed across all the 16 axonal delay values. The first assumption is always true as long as the minicolumn array is generating events. To meet the second condition, all the probabilities P_i are normalized to 20-bit integers R_i , which are used to generate a set of subsequent thresholds T_i (17 20-bit integers and $T_0 = 0$) such that $T_{i+1} - T_i = R_i$. In hardware, in each clock cycle, a 20-bit LFSR generates a random value that will fit in one of the threshold ranges (between T_{i+1} and T_i). This implies that in each clock cycle, the delay generator will select one of the 16 axonal delay values (of which the events should be read out). If the selected axonal delay value has no events to be read out, we simply wait for the next clock cycle. With no complicated rules involved, this implementation effectively distributes the read-out operation evenly across all the axonal delay values. Note that using 20-bit integers is an arbitrary choice. In theory, the higher the resolution, the more even the distribution will be; however, this requires more logic gates, and our result shows that a 20-bit resolution is sufficient.

The above approach only provides the ratios of delays; to achieve the actual delay values, we need to scale P_i with a reconfigurable global probability f . This function is implemented with a stochastic approach that does not involve multiplication. The read out is only enabled when the 10-bit normalized f is no greater than the value of a 10-bit LFSR. In hardware, the axon array will access the DDR memories every $1 \text{ k}/200 \text{ MHz} \approx 5 \mu\text{s}$. If f is set to 0, the mean of the actual "1 ms" axonal delay will be $\sim 5 \mu\text{s}$. If f is set to 1023 ($0 \times 3\text{FF}$), the mean of the actual 1 ms axonal delay will be $5 \mu\text{s} \times 1023 \approx 5 \text{ ms}$. The value of f also can be adjusted for the firing rate of the minicolumns for more accurate emulation.

As the TX-phase will select one DDR memory, the RX state machine will initiate a read operation immediately after the start of the TX-phase by selecting the axonal delay values in the other

DDR memory. If there is no event in the region selected by the delay generator, we simply wait for the next selection in the next clock cycle. The events read from the DDR memories are written into the RX_FIFO, which are read out by the synapse array. The RX-phase will not be initiated if the RX_FIFO is full. The data flow in the RX-phase is far less critical than that in the TX-phase since the capacities of the DDR memories are orders of magnitude larger than the on-chip memory.

Synapse Array

The synapse array (see **Figure 6**) emulates the function of biological synaptic connections: it modulates the incoming events from the axon array with synaptic weights and generates destination minicolumn addresses for them. These events are then sent to the TM minicolumns. The synapse array only performs the linear accumulation of synaptic weights of the incoming events, whereas the exponential decay is emulated by the PSC generator in the neuron.

Processing Events

The synapse array consists of a state machine, a synaptic weight modulator, an address mapper, a PRE_FIFO (the FIFO for storing pre-synaptic events), and 16 arbiters. The axon array has already performed the function of converting one post-synaptic event to multiple pre-synaptic events according to its connections. Hence, one incoming pre-synaptic event to this synapse array will be sent to the minicolumns of one and only one hypercolumn, which are referred to as the destination minicolumns and the destination hypercolumn, respectively. For instance, the destination hypercolumn has 100 minicolumns and the size of the connection is 32 (minicolumns); thus, 32 out of 100 minicolumns will be randomly selected as the destination minicolumns.

A serial implementation that performs this function would take up to 128 cycles and would be extremely slow. A parallel implementation that performs this function in one cycle would be too big and unnecessary, since not all connections have 128 minicolumns. As a compromise, this function is performed in a time slot of four clock cycles: the synaptic weight modulator and the address mapper process up to 32 connections in each clock

cycle using a 12-stage pipeline. For instance, if the size of the connections is 80, we will process 32, 32, 16, and 0 connections in the four time slots, respectively. Although the number of connections could be <128 , the time slot has to be fixed to four clock cycles as pipelines can only achieve the best performance with fixed steps, else the overhead of pipeline will be significant.

The state machine reads one pre-synaptic event from the axon array if the RX_FIFO of the axon array is not empty. The minicolumn address of the incoming event is sent to the parameter LUT for its connection parameters. With these parameters, the address mapper first generates the destination hypercolumn address (20 bits) by adding an offset (20 bits) to the hypercolumn address (20 bits) of the incoming event. If it overflows, it will wrap. Recurrent connections can be realized by setting the offset to zero. The addresses of the destination minicolumns are generated randomly based on the size of the connections. The events with a given address are sent to the same destination minicolumns at each time.

Along with the address mapper, the synaptic weight modulator first multiplies the eight spike counts (of the eight neuron types) of this event with the eight corresponding synaptic weights. To update the 64 synaptic connections between two minicolumns, for each type, these eight weighted spike counts are linearly summed with an 8-bit mask (1 bit per connection). The results of this operation are written into the PRE_FIFO, which is read out by the arbiters.

Arbiter

The arbiters read the events from the PRE_FIFO and assign them to the TM minicolumns. To successfully perform this task, it needs to figure out which TM minicolumn to assign (for one incoming event) within a short time, e.g., tens of clock cycles. If not, the FIFO of the arbiter that holds the newly arrived events during processing will soon be full and alert the Master to stop the synapse array from reading events.

In principle, one single arbiter for all the TM minicolumns can achieve the best utilization as any TM minicolumn can be assigned to any incoming event. A straightforward implementation for this approach was to use an address register array (Wang et al., 2013b, 2014), which is effectively a

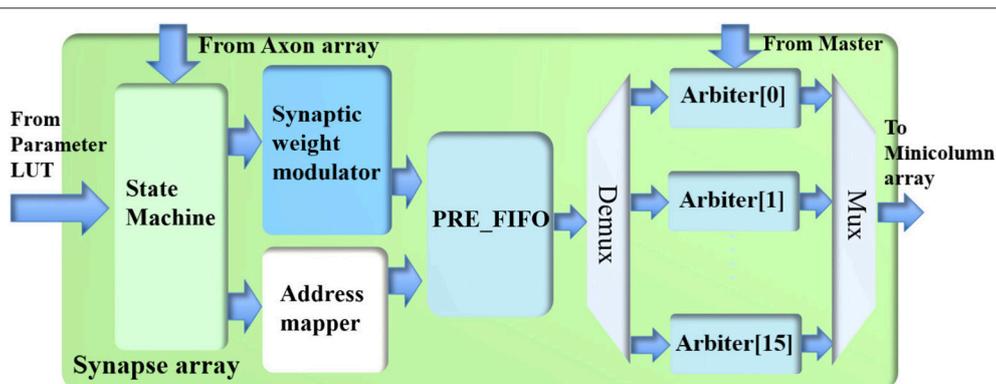


FIGURE 6 | The structure of the synapse array.

content-addressable memory (CAM). It is capable of checking whether there already is a TM minicolumn assigned to the incoming event or not in each clock cycle. However, this approach uses a large number of flip-flops. Since the limited number of flip-flops is one of the bottlenecks for large-scale FPGA designs, this simple method is not possible for 1M TM minicolumns.

To increase parallelism, 16 identical arbiters are used such that each of them can perform the dynamic assignment function for $1\text{ M}/16 = 65536$ (64k) TM minicolumns. The events in the PRE_FIFO are moved to these arbiters according to the most significant four bits of their addresses. The potential loss in utilization of the TM minicolumns will be minimal as long as the (DA) minicolumns are evenly distributed across the whole address range. This requirement is not difficult to meet for large-scale and sparsely connected neural networks with low activity rates.

The arbiter (see **Figure 7**) consists of a state machine, a fast CAM, a synaptic weight buffer, an address buffer, and an IN_FIFO, which is introduced to avoid losing spikes that arrive during the processing time. The fast CAM (assuming its size is $64\text{ k} \times 23$ bit) stores the addresses of minicolumns that have been assigned to the TM minicolumns. As a CAM, it compares input search data against stored data, and returns the address of matching data, which is then used to access the synaptic weight buffer. The latter is a dual-port SRAM with a size of $64\text{ k} \times 32$ bit (one 32-bit value per TM minicolumn).

The state machine reads the IN_FIFO when it is not empty and processes its output before another read. Inside the state machine, there is an index generator implemented by a counter that controls the fast CAM and the synaptic weight buffer. The output value of this index generator is used as the address of the TM minicolumns. The state machine accesses the fast CAM to check whether there is already a TM minicolumn assigned to this event or not. If there is already a TM minicolumn assigned to the incoming event, its values are linearly accumulated to the

values stored in the synaptic weight buffer. Otherwise, the state machine sends this event to an unassigned TM minicolumn, as indicated by the current value of the index generator, by storing the address and weights of this event in the fast CAM and the synaptic weight buffer, respectively. This TM minicolumn is then labeled as assigned by the state machine. After finishing the above actions, the index generator is incremented by one. Its count is reset to zero after it reaches the maximum value.

The most challenging part is the implementation of the fast CAM, which has to complete its search within tens of clock cycles. A straightforward implementation using 64 k 23-bit flip-flops, one per TM minicolumn, would be capable of performing this task. However, the amount of flip-flops required would be too large for the FPGA. An implementation using a single SRAM would be too slow since it can access only one value per clock cycle. In the worst case, the search time would be 64 k clock cycles.

To overcome this problem, we used multiple SRAMs to implement the fast CAM. Because hypercolumns have many minicolumns, we use each value of the fast CAM for eight TM minicolumns simultaneously. In this way, we can reduce the size of the fast CAM from $64\text{ k} \times 23$ bits to $8\text{ k} \times 20$ bits. It can be implemented efficiently using 128 single-port SRAMs, each with a size of 64×20 bits. The maximum search time is thus reduced to 64 clock cycles. The 16-bit address for accessing the synaptic weight buffer is obtained by concatenating the 13-bit address returned by the fast CAM with the least significant 3 bits of the address of the event.

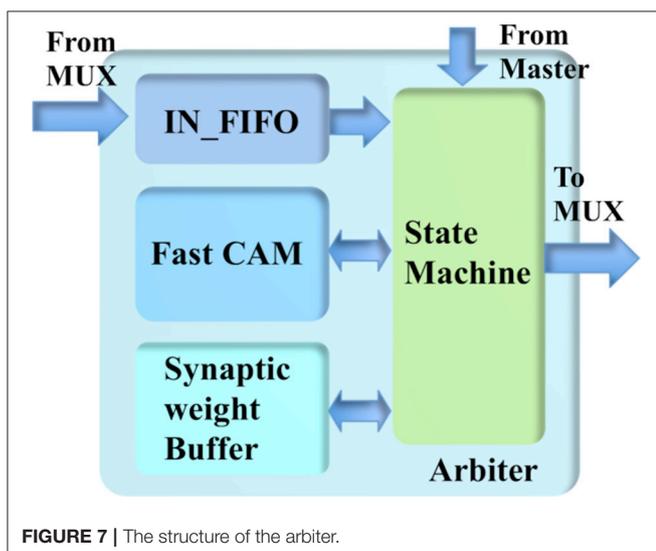
Because the events for the destination minicolumns of the same hypercolumn arrive sequentially, we introduce a bypass mechanism that is capable of significantly reducing the processing time. The newly-read event from the IN_FIFO is first compared with the previous event, before it is sent to the fast CAM. If the most significant 20 bits of their addresses match, the state machine bypasses the fast CAM and performs linear accumulation directly. The values stored in the synaptic weight buffer are accumulated until they are read out and cleared by the minicolumn array.

To obtain the parameters of the neurons and hypercolumn connections, the minicolumn array also needs to read the values stored in the fast CAM, which are sent to the parameter LUT. The read-out of the CAM uses the same scheme as the one for the synaptic weight buffer but without the “clear” operation.

Parameter LUT

The parameter LUT (see **Figure 8**) performs the LUT function for the incoming addresses from the synapse array and minicolumn array within several clock cycles. Its key component is a parallel CAM, which is implemented with 512 27-bit flip-flops, instead of with SRAM, such that it can be accessed by more than one input simultaneously with a fixed latency of only three clock cycles. This fixed latency requirement is critical and must be met, because the system is a highly pipelined design, which will not work with non-fixed latencies. The latency of an SRAM-based CAM is generally not fixed.

The conventional CAMs compare input search data against stored data and return the address of matching data. The parallel CAM differs slightly from the conventional CAMs: it stores a set



of sequential thresholds A_i and its return value tells in which range (between A_{i+1} and A_i) the input address is. This return value is used for accessing the buffers.

The minicolumn array requires two sets of parameters. The first set is for computing the neuron type, time constants and gains of the neurons. Two buffers are used for reducing memory usage. The returned address from the parallel CAM is used to access the parameter-type buffer, the values stored in it are then used to access the minicolumn parameter buffer. The second set of parameters is for routing the post-synaptic events: the axonal delay values and the hypercolumn connections, which are stored in the post-connection buffer.

For the synapse array, the returned addresses from the parallel CAM are used to access the connection address buffer. The values stored in this buffer are then used to access two other buffers. The first buffer is a pre-connection buffer, which stores the size of the connections, offset, synaptic weights, and the size of the destination hypercolumn. The second buffer is a

neuron-connection buffer, which stores the synaptic connections between two minicolumns.

Master

The Master plays a vital role in the cortex simulator: it has complete control over all the modules in the neural engine such that it can manage the progress of the simulation. This mechanism effectively guarantees no event loss or deadlock during the simulation. The Master slows down the simulation by pausing the modules that are running quicker than other modules. The Master has two components (see **Figure 2**): a memory bus controller and an external interface.

Memory Bus Controller

The memory bus controller (see **Figure 9**) has two functions: interfacing the off-chip memory with Altera IPs, and managing the memory bus sharing between the minicolumn array and the axon array. There are multiple off-chip memories in the system: two DDR memories ($128\text{ M} \times 512$ bits with a 512-bit interface) and four QDRII memories, each of which is a high-speed dual-port SRAM but with a limited capacity ($1\text{ M} \times 72$ bits, with a 72-bit interface). Storage of the neural states of the minicolumn array will use both the DDR and QDRII memories. For the DDR memories, the lowest $512\text{ k} \times 512$ bits are used for neural states. The rest of the DDR memories (~ 8 GB) are used to store the post-synaptic events (evenly divided by eight, one per delay), as presented in section Axon Array.

Accessing the QDR memories is quite straightforward as they are all dual-port devices. However, they are running in a different clock domain than the DDR memories, so that we need two asynchronous FIFOs (QDR WR_FIFO and RD_FIFO) for communication between the two clock domains. To efficiently access the DDR memories for the neural states, the memory bus controller accesses them alternately such that while writing 1 k neural states to one DDR memory, 1 k neural states from the other one are read. In the next burst operation, these roles are reversed. This method effectively uses two single-port memories to build a dual-port memory with a bandwidth of 512 bits@200 MHz, whereas four QDR memories can provide 288

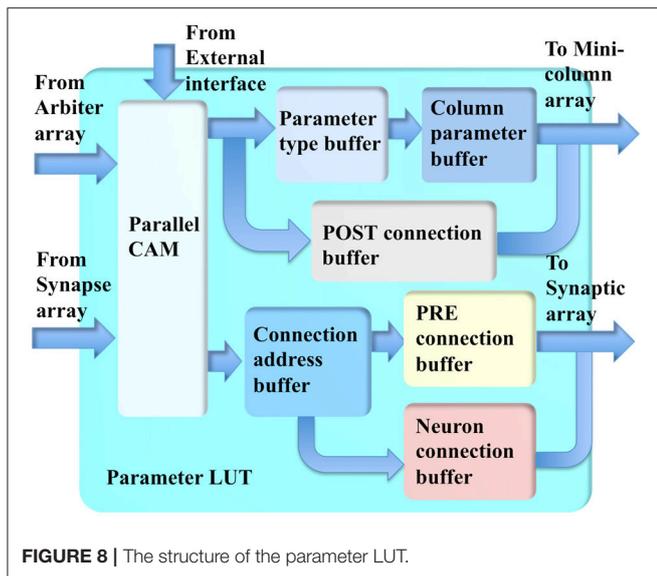


FIGURE 8 | The structure of the parameter LUT.

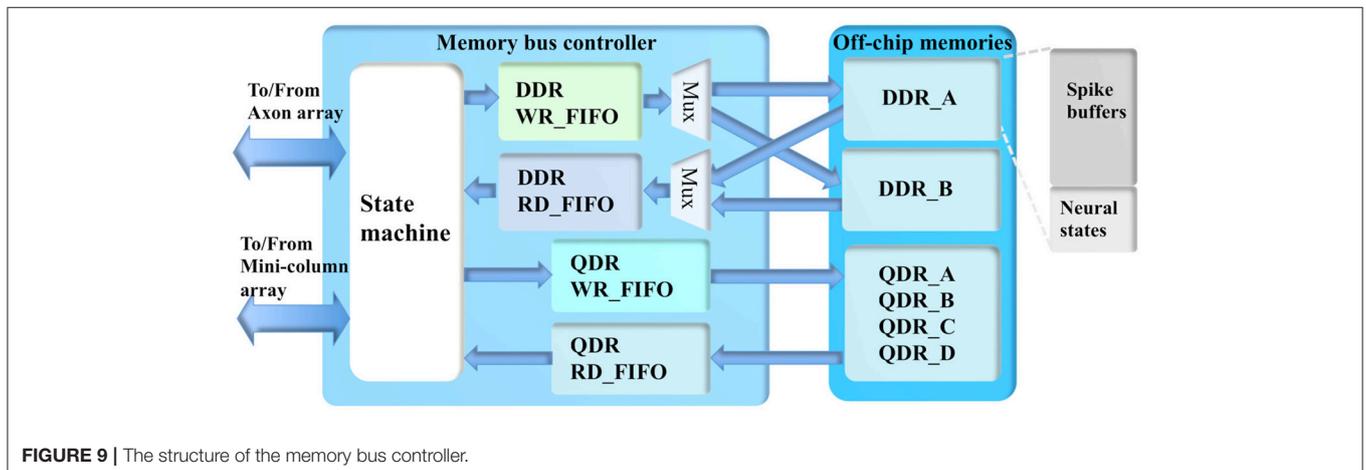


FIGURE 9 | The structure of the memory bus controller.

bits@225 MHz. Hence, we can provide enough bandwidth (800 bits@200 MHz) for the TM minicolumns (100 physical neurons, each neuron needs 8 bits).

There are waiting times when reading the DDR memories. The TM system is segmented using a size of 1024 (1k) TM minicolumns, and there is no waiting time within one segment. To make sure the neural states from the memories are always available during one segment, we use a read-in-advance scheme: the neural states being read out in the current segment are to be used in the next segment. These 1k neural states will be written into a DDR RD_FIFO, which sends its output to the minicolumn for computing. The write operation is relatively less critical because we can choose when to write to the DDR memory. We simply write the neural states from the minicolumn array into the DDR memory via a DDR WR_FIFO. It plays the role of a buffer, since there are waiting times when writing the DDR memory.

The memory bus controller will only grant the axon array access to the DDR memory during the slots between two 1k-burst operations. In the system presented here, this slot is configured to be up to 200 clock cycles. This means if the axon array finishes its operation within 200 clock cycles, it will release this access after that. To avoid the loss of events, when the usage of TX_FIFO (of the axon array) exceeds 95%, the memory bus controller will pause the minicolumn array by not allowing it access to the off-chip memory. Instead, the memory bus controller will allow the axon array to keep using the bus to write events to the DDR memory until the usage of the TX_FIFO is lower than 95%.

External Interface

The external interface module will control the flow of the input and output of events and parameters. The main input to this system are events, which are sent to the minicolumn array via the axon array. This module also performs instruction decoding such that we can configure the parameter LUT and the system registers.

The outputs of the simulator are individual spikes (100 bits, one per neuron) and events generated by the minicolumns. The individual spikes (with their minicolumn addresses) are sent to the host via an SPK_FIFO if their addresses are labeled as being monitored. To support multi-FPGA system, the events that labeled to other boards will also be sent to the host.

RESULTS

Prototype System

The system was developed using the standard ASIC design flow, and can thus be easily implemented with state-of-the-art manufacturing technologies, should an integrated circuit implementation be desired. A bottom-up design flow was adopted, in which we designed and verified each module separately. Once the module-level verification was complete, all the modules were integrated together for chip-level verification. **Table 1** shows the utilization of hardware resources on the FPGA. Note that this utilization table includes the Altera IPs such as the ones for accessing external memories and the 10G transceiver.

TABLE 1 | Device utilization on Altera Stratix 5SGXEA7N2F45C2.

Adaptive logic modules (ALMs)	RAMs	DSPs
157978/234720	41 Mbit/52 Mbit	256/256*

*One DSP slice can be used to implement three 9-bit fixed-point multipliers. The Altera synthesizer tool will only use this function when all the DSPs have been used.

Along with the hardware platform, we also developed a simple application programming interface (API) in Python that is similar to the PyNN programming interface (Davison et al., 2008). This API is very similar to the high-level object-oriented interface that has been defined in the PyNN specification: it allows users to specify the parameters of neurons and connections, as well as the network structure using Python. This will enable the rapid modeling of different topologies and configurations using the cortex simulator. This API allows monitoring of the generated spikes in different hypercolumns. As future work, we plan to provide full support for PyNN scripts and incorporate interactive visualization features on the cortex simulator.

Auditory Cortex Experiment

The results presented here focus on demonstrating the capability of our system for simulating large and structurally connected spiking neural networks in real time by implementing a simplified auditory cortex. The experiment ran the simulation in real time and the support system that generated the stimulus and recorded the results were implemented on a second DE5 board.

The human primary auditory cortex (A1) contains ~100 million neurons (MacGregor, 1993) and receives sensory input via the brainstem and mid-brain from the cochlea. One important aspect of the auditory cortex is its tonotopic map: neurons in auditory cortex are organized according to the frequency of sound to which they respond best from low to high along one dimension. Neurons at one end of a tonotopic map respond best to low frequencies; neurons at the other end respond best to high frequencies. This arrangement reflects the tonotopic organization of the cochlear output, which is maintained throughout the brainstem and midbrain areas.

The simplified auditory cortex model used in this experiment consists of 100 channels, the sensory input to which is synthetic data that emulates the output of a neuromorphic cochlea (Xu et al., 2016). This cochlea implements 100 sections (one per frequency range) in real time with incoming sound sampled at 44.1 kHz. We generated a synthetic signal with 100 channels representing the output of the cochlea in response to a pure tone sweep with exponentially increasing frequency.

Each cochlear channel is connected to a cortical population consisting of 100 hypercolumns with 100 minicolumns each. Our auditory cortex thus contains $100 \times 100 = 10\text{ k}$ hypercolumns, yielding 100 million LIF neurons. We assigned a ratio between inhibitory and excitatory neurons of 1:4. This ratio applies for each minicolumn (80 excitatory neurons and 20 inhibitory neurons). Within a minicolumn there are three groups (i.e., six neuron types) representing the different layers of neurons in cortex: L2/3, L4, and L5/6. The L2/3 group has 40 (32 excitatory

and 8 inhibitory), the L4 group has 20 (16 excitatory and 4 inhibitory) and the L5/6 group has 40 (32 excitatory and 8 inhibitory) neurons, which is reasonably consistent with the number of neurons in these layers in a biological minicolumn. In the experiment, type 1-2, type 3-4, and type 5-6 implemented the excitatory and inhibitory neurons of the L2/3, L4, and L5/6, respectively. The parameters of the excitatory and inhibitory neurons are all set to the same values (see **Table 2**). This was done for simplicity, since we are only using this experiment to verify our hardware implementation, not to reproduce any particular detailed behavior of auditory cortical neurons. For more detailed cortical simulations, different parameters would be chosen for different neuron types.

For real time simulation, we used 176k TM minicolumns and a time slot for the axon array to access the DDR memory (between two 1 k-burst operations), which was configured to be up to 200 clock cycles. Thus, the total time per update cycle is $176 \times (1024+200) \times 5ns \approx 1ms$.

The connectivity of the hypercolumns, minicolumns and groups was also made consistent with known statistics of the anatomy of cortex (see **Figure 10**), which was used to verify the hierarchical communication scheme and the functions of the Parameter LUT. In our simplified auditory cortex, there are no connections between the 100 channels. Within each channel, each minicolumn is randomly connected to 24 other minicolumns: eight of them in the same hypercolumn and eight of them each in two neighboring hypercolumns. The

connections of the groups within the same hypercolumn use the intracortical connection motif, whereas the ones between hypercolumns use intercortical connectivity statistics, as shown in **Figure 10**. For balanced excitation-inhibition, the synaptic weights for excitatory and inhibitory connections are 0.4 and -1.0 , respectively.

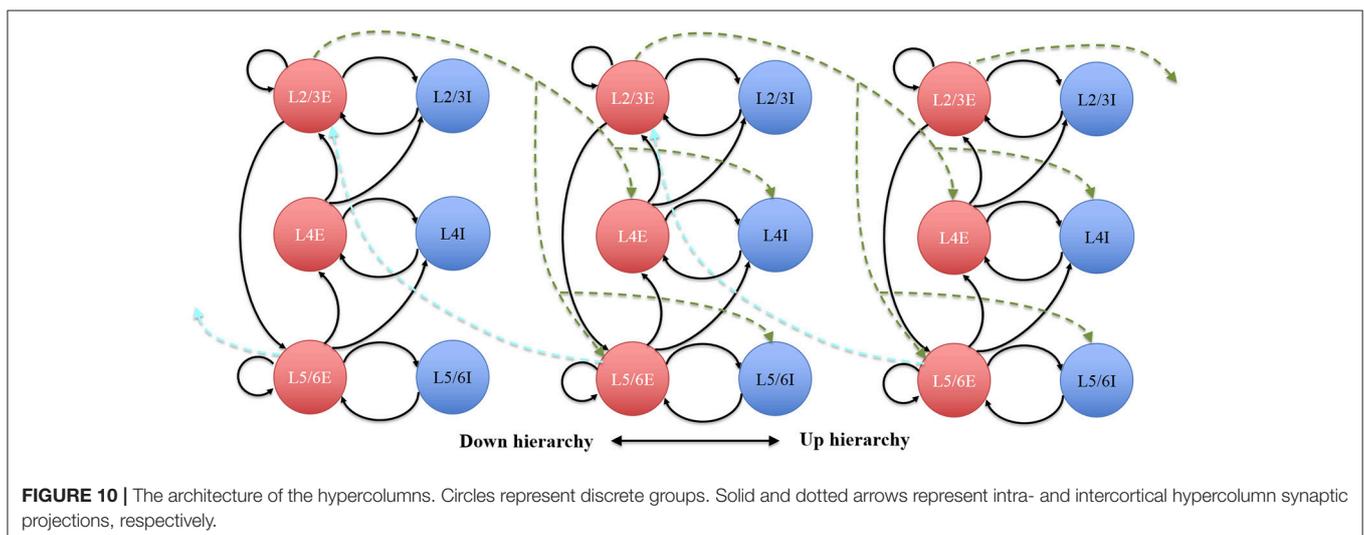
The sensory input is simulated using 1,000 Poisson neurons (10 per channel) driven by the synthetic data that emulates the output of the cochlea. Each Poisson neuron is connected to eight randomly chosen minicolumns in a hypercolumn. The 10 Poisson neurons per channel are thus connected to 10 out of the 100 hypercolumns in that channel. The average firing rate of these Poisson neurons is ~ 10 Hz. The spikes generated by these Poisson neurons are only sent to the excitatory neurons of the L4 group, which is also consistent with known anatomy of cortex. To simulate the cochlear output of a pure tone sweep with exponentially increasing frequency, within every 10 ms, there is one set of 10 Poisson neurons generating spikes. For 100 channels, it will take $100 \times 10ms = 1s$ to go through all the 100 channels. To better monitor the dynamics of the neural networks, the stimulus is repeated 10 times.

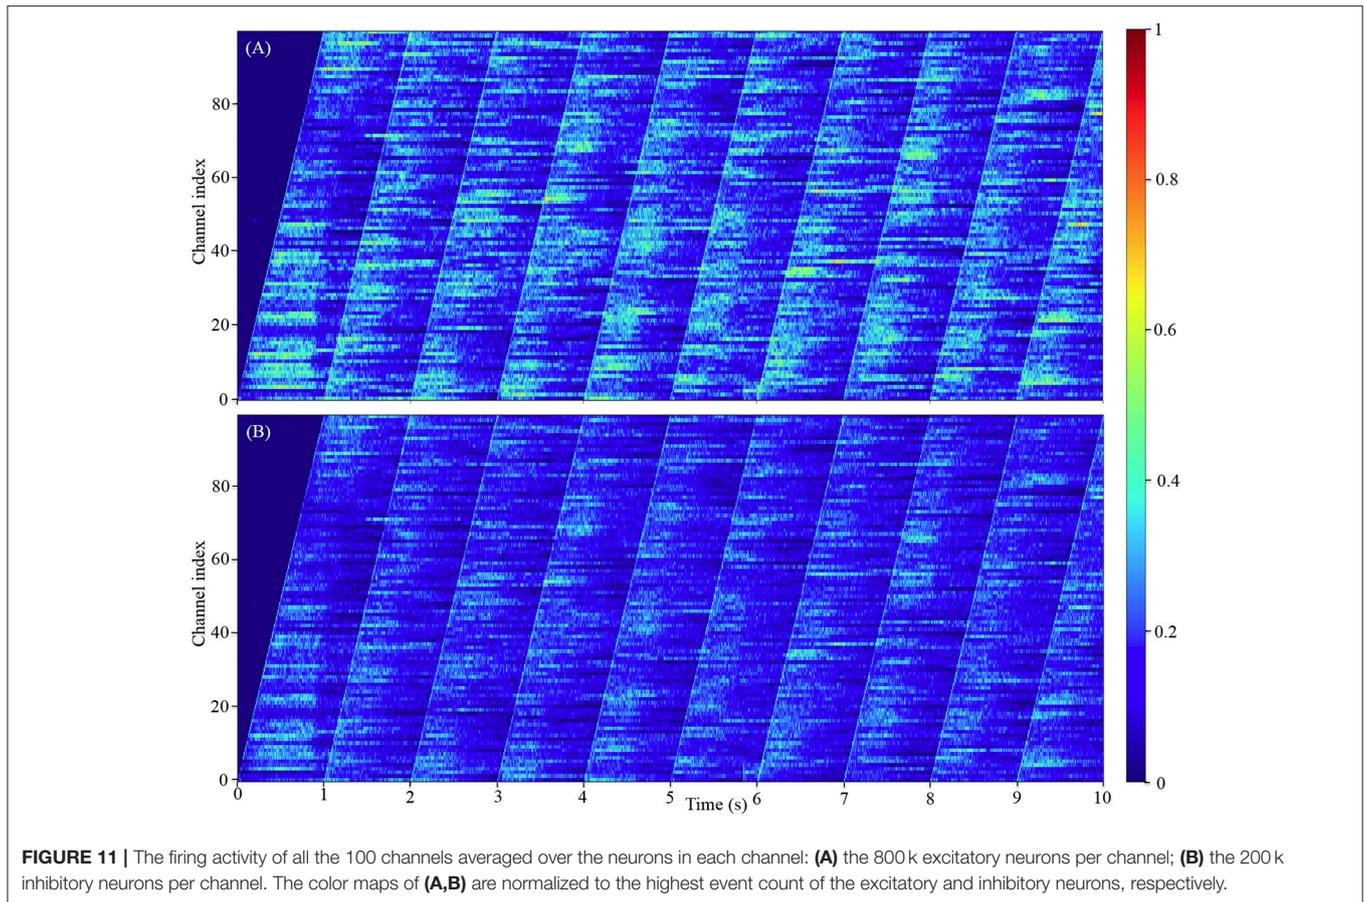
Our model comprising 100 million neurons is to-date the largest spiking neural networks simulated in real-time. Since it is obviously impractical to show the firing activities of all the 100 million neurons, we chose instead to show how the firing rates of all the 100 channels change with the input signals. The results in **Figure 11** show that each channel will remain silent until its input Poisson neurons fire. After that, the channel exhibits the classic asynchronous irregular network states (Brunel, 2000) for hundreds of milliseconds before falling silent or receiving the next input burst from the Poisson neurons. These results also confirm that the channels of this simplified auditory cortex are capable of emulating the heterogenous firing behaviors, due to the randomness introduced by the stochastic approach to decay in our LIF neurons and the randomness of the Poisson neurons.

To obtain the statistics of the dynamic-assignment approach, we ran the experiment for 10 times while monitoring every

TABLE 2 | Parameters used in the BRN.

τ_{EPSC}	5.8ms
τ_{IPSC}	5.8ms
τ_{mem}	5.8ms
τ_{rtc}	3ms
g_{syn}	1
g_{psc}	1





millisecond how many TM minicolumns are being “active,” i.e., at least one of its neurons is integrating or firing and/or the PSC is non-zero. **Figure 12** shows the number of the active (TM) minicolumns, showing that only 80,000 out of the 1 million minicolumns are active at any given time and that this ratio remains fairly stable. It demonstrates the suitability of using the dynamic-assignment approach for sparsely connected large-scale neural networks.

We used PowerPlay, the power analysis tool provided by Altera to estimate the power dissipation of the FPGA, since a direct measurement of the FPGA’s power consumption is not possible on this development kit. PowerPlay estimated the power dissipation (see **Table 3**) from the netlist of the design, using the tool’s default settings. The total power dissipation of the DDR and QDR memories are 6.6 W, which were estimated by the tools of Micron and Cypress. The total power dissipation of the whole system is ~ 32.4 W. Since there are 20 M (TM) neurons, the power dissipation per (TM) neuron is $32.4 \text{ W}/20 \text{ M} \approx 1.62 \mu\text{W}$.

DISCUSSION

Existing neuromorphic platforms were developed for different applications with different specifications, and different trade-offs, and are difficult to compare. Since we are mainly interested in the trade-off between the scale of the implementation, i.e.,

the number of neurons that the system can simulate, and the hardware cost, we will focus here on comparing to solutions that can simulate more than 100 k neurons. As a result of our hardware focus, modeling and simulations of neocortex with biophysically detailed models of each individual neuron as the work studied by Markram et al. (Markram et al., 2015), are out of the scope of this paper and will not be addressed.

There are five large-scale neuromorphic systems: SpiNNaker, BrainScaleS, Neurogrid, HiAER-IFAT, and TrueNorth. Owing to the fact that these systems were implemented with a diverse range of approaches (for diverse objectives), each of them has its own advantages and disadvantages. **Table 4** summarizes some of the main areas for comparison. As our work was aimed at simulating large-scale spiking neural networks, of particular note are:

- **Availability:** Our system uses commercial off-the-shelf FPGAs, while all the other four platforms use specialized hardware that are not easily accessible to other researchers. The availability of our system will enable many more researchers to participate.
- **Scalability:** Our system is scalable in two ways. First, our system can be linearly scaled up using multiple FPGAs without performance loss on the back of its modular architecture and hierarchical communication scheme. Second, our system can leverage the rapid growth in FPGA technology without additional costs.

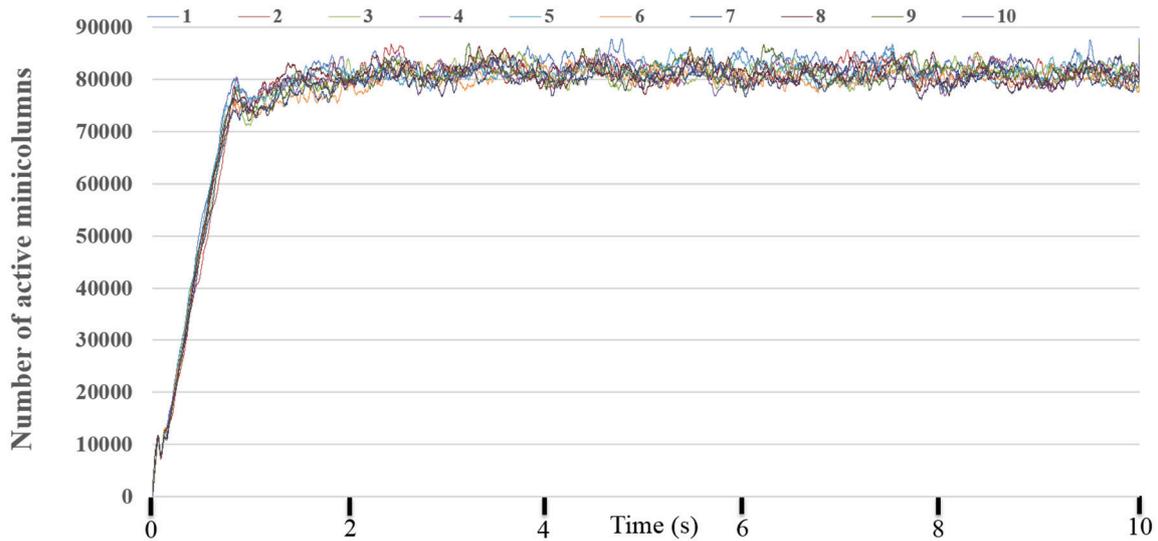


FIGURE 12 | The number of active minicolumns. We run this experiment for 10 times. It clearly demonstrated that the performance of the dynamic-assignment approach is stable. The mean dynamic ratio from 10 runs was about 100:8.

TABLE 3 | FPGA's power dissipation estimated by PowerPlay.

Total power dissipation	25.8 W
I/O power dissipation	8.1 W
Core static power dissipation	3.6 W
Core dynamic power dissipation	14.1 W
DDR interface dynamic power dissipation	502 mW
QDR interface dynamic power dissipation	238 mW
Minicolumn array dynamic power dissipation	369 mW
Synapse array dynamic power dissipation	7.5 W
Axon array dynamic power dissipation	582 mW
Parameter LUT dynamic power dissipation	239 mW
Master dynamic power dissipation	247 mW
Routing dynamic power dissipation*	4.4 W

* This item includes all the dynamic power dissipation caused by the routing of the signals (wires) between modules.

- **Reliability:** Our system is a fully digital design and it is largely insensitive to process variations and device mismatch, which are the major issues for large-scale analog and mixed-signal designs.
- **Flexibility:** Dedicated hardware platforms are generally unchangeable after being fabricated. Hence, they cannot meet the requirement of rapid prototyping. Owing to the programmable nature of FPGAs, our system can be easily reconfigured for different models.

The main contributions in this paper compared to our previous work presented in Wang et al. (2014c) are:

- Using minicolumns (including the sub-type neurons) and hypercolumns, so the parameters can be shared and therefore their number reduced so that only on-chip memory is needed to store them.

- Using structural connections between the neuron types in the minicolumns and hypercolumns, avoiding the need for the terabytes of memory that would be needed to store all the individual connections.
- Communicating spike counts instead of individual spikes; the neurons of a single type within a minicolumn all receive the same input events, so spikes do not need to be replicated individually for each neuron, reducing the required communication bandwidth.
- Improved time-multiplexing approach enabling the use of external memories for neural states.

Despite of the advantages of the FPGAs, most FPGA-based systems have been limited to behavioral tasks and failed to carry out investigations on more complicated biological plausible neural networks. This is due to the fact that an FPGA, as a digital device, does not provide massive numbers of efficient, built-in operators for highly computationally intensive functions, such as exponentials and divisions used in neuronal ion channel models. Although various designs have been presented for this challenge (Graas et al., 2004; Mak et al., 2006; Pourhaj and Teng, 2010), all the large-scale neural networks presented in the literature review were implemented with integrate-and-fire neurons or Izhikevich neuron models. In our future work, we will continue work on this system for investigating the implementation of more complex neuron models.

In addition, future developments of the design also include developing support for synaptic plasticity, e.g., Spike Timing Dependent Plasticity (Gerstner et al., 1996; Magee, 1997; Markram et al., 1997; STDP, Bi and Poo, 1998) and Spike-Timing-Dependent Delay Plasticity (STDDP, Wang et al., 2011a,b, 2012, 2013a). This extension will be straightforwardly carried out by using the scheme introduced in our previous work (Wang et al., 2015), which implements a synaptic plasticity adaptor array that

TABLE 4 | Comparison of with other large-scale neuromorphic systems.

	Neurogrid	BrainScaleS	TrueNorth	HiAER-IFAT	SpiNNaker	This work
Technology	Analog	Analog	Digital	Analog	Digital	Digital
Feature size	180 nm	180 nm	28 nm	130 nm	130 nm	28 nm
Chips	16	352	16	16	48	1
Neurons	1 M	200 k	16 M	1 M	768 k	20 M*
Synapses	4 G	40 M	4 G	1 G	768 M	4 T*
Power	3 W	500 W	3.2 W	40 W	80 W	32 W
Interconnect	Tree-multicast	Hierarchical	2D mesh-unicast	Hierarchical	2D mesh-unicast	Hierarchical
Neuron model	Adaptive quadratic IF	Adaptive exponential IF	Configurable LIF	2-compartment LIF	Programmable	2-compartment LIF
Synapse model	Shared dendrite	4-bit digital	Binary, 4 modulators	Shared conductance-based	Programmable	Shared conductance-based
Axon model	Axonal arbor	Speed-up Programmable delay	Four types of axonal delay	digital programmable delay	4-bit digital programmable delay	4-bit digital programmable delay

*To provide a fair comparison, the dynamic assignment ratio was set to 1:1.

is separate from the neuron array. For each synapse, which remains part of the neuron, a synaptic adaptor will be connected to it when it needs to apply a certain synaptic plasticity rule. The synaptic adaptor will perform the weight/delay adaptation. This strategy enables a hardware neural network to be configured to perform multiple synaptic plasticity rules without needing to change its own structure, simply by connecting the neurons to the appropriate modules in the synaptic plasticity adaptor array. This structure was first proposed by Vogelstein and his colleagues in the IFAT project (Vogelstein et al., 2007).

Our cortex simulator uses 200 k TM minicolumns to simulate 20 million to 2.6 billion LIF neurons in real time. When running at five times slower than real time, it is capable of simulating 100 million to 12.8 billion LIF neurons, which is the maximum network size on the chosen FPGA board, due to memory limitations. Larger networks could be implemented on larger FPGA boards with more external memory. A hierarchical

communication scheme allows one neuron to have a fan-out of up to 200 k neurons. With the advent of inexpensive FPGA boards and compute power, our cortex simulator offers an affordable and scalable tool for design, real-time simulation, and analysis of large-scale spiking neural networks.

AUTHOR CONTRIBUTIONS

RW and AvS: Designed the system and experiment, conducted experiments, and wrote the paper; CT: Contributed writing the paper.

ACKNOWLEDGMENTS

This work was supported by the Australian Research Council Grant DP140103001. The support by the Altera university program is gratefully acknowledged.

REFERENCES

- Bailey, J. A., Wilcock, R., Wilson, P. R., and Chad, J. E. (2011). Behavioral simulation and synthesis of biological neuron systems using synthesizable VHDL. *Neurocomputing* 74, 2392–2406. doi: 10.1016/j.neucom.2011.04.001
- Benjamin, B. V., Gao, P., McQuinn, E., Choudhary, S., Chandrasekaran, A. R., Bussat, J. M., et al. (2014). Neurogrid: a mixed-analog-digital multichip system for large-scale neural simulations. *Proc. IEEE* 102, 699–716. doi: 10.1109/JPROC.2014.2313565
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J. Neurosci.* 17, 2112–2127.
- Bower, J. M., and Beeman, D. (1998). *The Book of Genesis*. New York, NY: Springer.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory neurons. *Comput. Neurosci.* 8, 183–208. doi: 10.1023/A:1008925309027
- Buxhoeveden, D. P. (2002). The minicolumn hypothesis in neuroscience. *Brain* 125, 935–951. doi: 10.1093/brain/awf110
- Buxhoeveden, D. P., and Casanova, M. F. (2002). The minicolumn and evolution of the brain. *Brain. Behav. Evol.* 60, 125–151. doi: 10.1159/000065935
- Cassidy, A., Andreou, A. G., and Georgiou, J. (2011). “Design of a one million neuron single FPGA neuromorphic system for real-time multimodal scene analysis,” in *45th Annual Conference on Information Sciences and Systems* (Baltimore, MD), 1–6.
- Cheung, K., Schultz, S. R., and Luk, W. (2016). NeuroFlow: a general purpose spiking neural network simulation platform using customizable processors. *Front. Neurosci.* 9:516. doi: 10.3389/fnins.2015.00516
- Davison, A. P., Brüderle, D., Eppler, J., Kremkow, J., Müller, E., Pecevski, D., et al. (2008). PyNN: a common interface for neuronal network simulators. *Front. Neuroinform.* 2:11. doi: 10.3389/neuro.11.011.2008
- Dombrowski, S. M. (2001). Quantitative architecture distinguishes prefrontal cortical systems in the Rhesus Monkey. *Cereb. Cortex* 11, 975–988. doi: 10.1093/cercor/11.10.975
- Furber, S. B., Galluppi, F., Temple, S., and Plana, L. A. (2014). The SpiNNaker Project. *Proc. IEEE* 102, 652–665. doi: 10.1109/JPROC.2014.2304638
- Gerstner, W., Kempter, R., Van Hemmen, J. L., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–81. doi: 10.1038/383076a0
- Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430

- Goodman, D. (2008). Brian: a simulator for spiking neural networks in Python. *Front. Neuroinform.* 2:5. doi: 10.3389/neuro.11.005.2008
- Graas, E. L., Brown, E. A., and Lee, R. H. (2004). An FPGA-based approach to high-speed simulation of conductance-based neuron models. *Neuroinformatics* 2, 417–436. doi: 10.1385/NI:2:4:417
- Hines, M. L., and Carnevale, N. T. (1997). The neuron simulation environment. *Neural Comput.* 9, 1179–1209. doi: 10.1162/neco.1997.9.6.1179
- Hoang, R. V., Tanna, D., Jayet Bray, L. C., Dascalu, S. M., and Harris, F. C. (2013). A novel CPU/GPU simulation environment for large-scale biologically realistic neural modeling. *Front. Neuroinform.* 7:19. doi: 10.3389/fninf.2013.00019
- Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J. Physiol.* 551, 139–153. doi: 10.1113/jphysiol.2003.044784
- Hubel, D. H., and Wiesel, T. N. (1977). Ferrier lecture: functional architecture of Macaque Monkey visual cortex. *Proc. R. Soc. B Biol. Sci.* 198, 1–59. doi: 10.1098/rspb.1977.0085
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Netw.* 14, 1569–1572. doi: 10.1109/TNN.2003.820440
- 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
- Johansson, C., and Lansner, A. (2007). Towards cortex sized artificial neural systems. *Neural Netw.* 20, 48–61. doi: 10.1016/j.neunet.2006.05.029
- Lennie, P. (2003). The cost of cortical computation. *Curr. Biol.* 13, 493–497. doi: 10.1016/S0960-9822(03)00135-0
- MacGregor, R. J. (1993). *Theoretical Mechanics of Biological Neural Networks*. London: Academic Press, Inc.
- Magee, J. C. (1997). A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. *Science* 275, 209–213. doi: 10.1126/science.275.5297.209
- Maguire, L. P., McGinnity, T. M., Glackin, B., Ghani, A., Belatreche, A., and Harkin, J. (2007). Challenges for large-scale implementations of spiking neural networks on FPGAs. *Neurocomputing* 71, 13–29. doi: 10.1016/j.neucom.2006.11.029
- Mak, T. S. T., Rachmuth, G., Lam, K.-P., and Poon, C.-S. (2006). A component-based FPGA design framework for neuronal ion channel dynamics simulations. *IEEE Trans. Neural Syst. Rehabil. Eng.* 14, 410–418. doi: 10.1109/TNSRE.2006.886727
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A. et al. (2015). Reconstruction and simulation of neocortical microcircuitry. *Cell* 163, 456–492. doi: 10.1016/j.cell.2015.09.029
- Martin, K. A. C. (2002). Microcircuits in visual cortex. *Curr. Opin. Neurobiol.* 12, 418–425. doi: 10.1016/S0959-4388(02)00343-4
- 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
- Moore, S. W., Fox, P. J., Marsh, S. J. T., Marketos, A. T., and Mujumdar, A. (2012). “Bluehive - a field-programmable custom computing machine for extreme-scale real-time neural network simulation,” in *2012 IEEE 20th International Symposium on Field-Programmable Custom Computing Machines (IEEE)* (Toronto, ON), 133–140.
- Pani, D., Meloni, P., Taveri, G., Palumbo, F., Massobrio, P., and Raffo, L. (2017). An FPGA platform for real-time simulation of spiking neuronal networks. *Front. Neurosci.* 11:90. doi: 10.3389/fnins.2017.00090
- Park, J., Yu, T., Joshi, S., Maier, C., and Cauwenberghs, G. (2016). Hierarchical address event routing for reconfigurable large-scale neuromorphic systems. *IEEE Trans. Neural Netw. Learn. Syst.* 28, 2408–2422. doi: 10.1109/TNNLS.2016.2572164
- Pfeil, T., Grübl, A., Jeltsch, S., Müller, E., Müller, P., and Petrovici, M. A., et al. (2013). Six networks on a universal neuromorphic computing substrate. *Front. Neurosci.* 7:11. doi: 10.3389/fnins.2013.00011
- Pourhaj, P., and Teng, D. H.-Y. (2010). “FPGA based pipelined architecture for action potential simulation in biological neural systems,” in *CCECE* (Calgary, AB), 1–4.
- Sen-Bhattacharya, B., Serrano-Gotarredona, T., Balassa, L., Bhattacharya, A., Stokes, A. B., Rowley, A., et al. (2017). A spiking neural network model of the lateral geniculate nucleus on the SpiNNaker machine. *Front. Neurosci.* 11:454. doi: 10.3389/fnins.2017.00454
- Shi, Y., Veidenbaum, A. V., Nicolau, A., and Xu, X. (2015). Large-scale neural circuit mapping data analysis accelerated with the graphical processing unit (GPU). *J. Neurosci. Methods* 239, 1–10. doi: 10.1016/j.jneumeth.2014.09.022
- Thomson, A. M. (2003). Interlaminar connections in the neocortex. *Cereb. Cortex* 13, 5–14. doi: 10.1093/cercor/13.1.5
- Tsunoda, K., Yamane, Y., Nishizaki, M., and Tanifuji, M. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nat. Neurosci.* 4, 832–838. doi: 10.1038/90547
- Vogelstein, R. J., Mallik, U., Vogelstein, J. T., and Cauwenberghs, G. (2007). Dynamically reconfigurable silicon array of spiking neurons with conductance-based synapses. *IEEE Trans. Neural Netw.* 18, 253–265. doi: 10.1109/TNN.2006.883007
- Wang, R. M., Hamilton, T. J., Tapson, J. C., and van Schaik, A. (2014). A mixed-signal implementation of a polychronous spiking neural network with delay adaptation. *Front. Neurosci.* 8:51. doi: 10.3389/fnins.2014.00051
- Wang, R. M., Hamilton, T. J., Tapson, J. C., and van Schaik, A. (2015). A neuromorphic implementation of multiple spike-timing synaptic plasticity rules for large-scale neural networks. *Front. Neurosci.* 9:180. doi: 10.3389/fnins.2015.00180
- Wang, R., Cohen, G., Hamilton, T. J., Tapson, J., and van Schaik, A. (2013a). “An improved aVLSI axon with programmable delay using spike timing dependent delay plasticity,” in *2013 IEEE International Symposium of Circuits and Systems (ISCAS) (IEEE)* (Beijing), 2–5.
- Wang, R., Cohen, G., Stiefel, K. M., Hamilton, T. J., Tapson, J., and van Schaik, A. (2013b). An FPGA implementation of a polychronous spiking neural network with delay adaptation. *Front. Neurosci.* 7:14. doi: 10.3389/fnins.2013.00014
- Wang, R., Hamilton, T. J., Tapson, J., and van Schaik, A. (2014a). “A compact neural core for digital implementation of the neural engineering framework,” in *BIOCAS* (Lausanne), 2014.
- Wang, R., Hamilton, T. J., Tapson, J., and van Schaik, A. (2014b). “A compact reconfigurable mixed-signal implementation of synaptic plasticity in spiking neurons,” in *2014 IEEE International Symposium on Circuits and Systems (ISCAS) (IEEE)* (Melbourne, VIC), 862–865.
- Wang, R., Hamilton, T. J., Tapson, J., and van Schaik, A. (2014c). “An FPGA design framework for large-scale spiking neural networks,” in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)* (Melbourne, VIC: IEEE), 457–460.
- Wang, R., Jin, C., McEwan, A., and van Schaik, A. (2011a). “A programmable axonal propagation delay circuit for time-delay spiking neural networks,” in *2011 IEEE Int. Symp. Circuits Syst.* (Rio de Janeiro), 869–872.
- Wang, R., Tapson, J., Hamilton, T. J., and van Schaik, A. (2011b). “An analogue VLSI implementation of polychronous spiking neural networks,” in *2011 Seventh International Conference on Intelligent Sensors, Sensor Networks and Information Processing (IEEE)* (Adelaide, SA), 97–102.
- Wang, R., Tapson, J., Hamilton, T. J., and van Schaik, A. (2012). “An aVLSI programmable axonal delay circuit with spike timing dependent delay adaptation,” in *2012 IEEE International Symposium on Circuits and Systems* (Seoul: IEEE), 2413–2416.
- Wang, R., Thakur, C. S., Cohen, G., Hamilton, T. J., Tapson, J., and van Schaik, A. (2017). Neuromorphic hardware architecture using the neural engineering framework for pattern recognition. *IEEE Trans. Biomed. Circuits Syst.* 11, 574–584. doi: 10.1109/TBCAS.2017.2666883
- Wang, R., Thakur, C. S., Hamilton, T. J., Tapson, J., and van Schaik, A. (2016). “A stochastic approach to STDP,” in *Proceedings of the IEEE International Symposium on Circuits and Systems* (Baltimore, MD).
- Wittie, L., and Memelli, H. (2010). “Billion neuron memory models in slender Blue Genes,” in *Program 208.30/MMM21, 2010 Neuroscience Meeting Planner*, (San Diego, CA: Society for Neuroscience), 1.

- Xu, Y., Thakur, C. S., Singh, R. K., Wang, R., Tapson, J., and van Schaik, A. (2016). "Electronic cochlea: CAR-FAC model on FPGA," in *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS) (IEEE)* (Shanghai), 564–567.
- Zamarreno-Ramos, C., Linares-Barranco, A., Serrano-Gotarredona, T., and Linares-Barranco, B. (2012). Multicasting mesh AER: a scalable assembly approach for reconfigurable neuromorphic structured AER systems. Application to ConvNets. *IEEE Trans. Biomed. Circuits Syst.* 7, 82–102. doi: 10.1109/TBCAS.2012.2195725
- Zaveri, M. S., and Hammerstrom, D. (2011). Performance/price estimates for cortex-scale hardware: a design space exploration. *Neural Netw.* 24, 291–304. doi: 10.1016/j.neunet.2010.12.003

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 © 2018 Wang, Thakur and van Schaik. 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 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.