Neuromorphic Modeling Abstractions and Simulation of Large-Scale Cortical Networks

Invited Paper

Jeffrey L. Krichmar[1,2], Nikil Dutt[2]
[1]Department of Cognitive Sciences
[2]Department of Computer Science
University of California, Irvine, Irvine, CA, USA
jkrichma@uci.edu, dutt@ics.uci.edu

Jayram M. Nageswaran, Micah Richert
Brain Corporation
San Diego, CA, USA
jayram.nageswaran@braincorporation.com
micah.richert@braincorporation.com

Abstract— Biological neural systems are well known for their robust and power-efficient operation in highly noisy environments. We outline key modeling abstractions for the brain and focus on spiking neural network models. We discuss aspects of neuronal processing and computational issues related to modeling these processes. Although many of these algorithms can be efficiently realized in specialized hardware, we present a case study of simulation of the visual cortex using a GPU based simulation environment that is readily usable by neuroscientists and computer scientists and efficient enough to construct very large networks comparable to brain networks.

***Keywords-** Spiking neural networks, GPU, vision, computational neuroscience, parallel processing, synapse*

I. Introduction

The mammalian nervous system is a network of extreme size and complexity [1], and understanding the principles of brain processing by reverse engineering neural circuits and computational modeling is one of the biggest challenges of the 21st century [2, 3]. Theoretical and computational methodologies will be crucial for healing, understanding, and enriching the mind [4]. Theories and models characterizing the computations underlying perception, memory, decision-making, behavioral intentions, self-regulation, and many other types of mental activity have become essential elements of behavioral and neuroscience research.

Computational neuroscience and neuromorphic engineering are now established subfields of the neurosciences. Computer scientists and engineers are looking to biology and the brain for inspiration when designing systems. Researchers are building large-scale models of the brain utilizing supercomputers [5, 6] and developing genomic brain atlases using high throughput techniques [7]. Current computational models are approaching the size of small mammalian brains [8, 9].

Despite recent increases in computer power, constructing a model the size of a human-brain will require several orders of magnitude increases in computation, communication, and memory capacity. Conventional computer hardware may not be the appropriate architecture for modeling a brain. Unlike a conventional computer, the brain is a massively parallel, analog, fault-tolerant, selective system that does not rely on programmed instructions [10]. Alternative computer architectures and programming paradigms, which are neurobiologically inspired, need investigation [11, 12]. Furthermore, the brain can be modeled and studied at multiple abstraction scales, from microscopic (e.g., biophysical and cellular) to macroscopic (e.g., functional) levels.

There is a need within the computational neuroscience community for abstractions and simulation environments that support modeling at a large-scale (i.e., networks which approach the size of biological nervous systems). In particular, we consider large-scale network models of spiking neurons, which demonstrate important properties of neurobiological information processing, such as temporal dynamics, precise timing, and brain rhythms [13]. Moreover, spiking models, with their digital signaling and sparse coding, are energy efficient and amenable to hardware application development.

Our group has been developing tools to incorporate these brain features into computer models. Specifically, we construct large-scale network models that capture the dynamics of neural signaling at the microcircuit (i.e., within brain areas) and macrocircuit (i.e., between brain areas) levels. Previously, we developed a GPU implementation of Spiking Neural Networks (SNN) that is highly efficient and released it to the modeling community so that researchers would have easy access to large-scale SNN simulations [14]. Our latest simulator extends this prior model to include more biologically plausible descriptions of synaptic connections, and learning rules [15]. In particular, this new simulation environment facilitates the development of very large-scale spiking neural networks that follow the brain's architecture.

In the remainder of the paper, we give an overview of neurobiological systems, discuss some computational aspects of modeling such a system, how these models may be realized in specialized hardware, and present as a case study a model of the visual cortex using our GPU based simulation environment.

II. Brief Neuroscience Primer

The human brain has an estimated 100 billion neural processing elements (neurons) and about 10^{15} connections (synapses) between those neurons. Each neuron is a sophisticated analog processor that not only integrates

This work was supported by the Defense Advanced Research Projects Agency (DARPA) subcontract 801888-BS.

information from other neurons but also exhibits complex internal dynamics. Neurons communicate with other neurons through all-or-none digital signals called action potentials or spikes. These spikes are propagated via long-range fibers called axons. The neurons can be either excitatory or inhibitory, that is they increase or decrease the ability of a downstream neuron to spike in a small window of time. The efficacy of synaptic connections between neurons changes based on the firing pattern of the "sending" (pre-synaptic) neuron and the "receiving" (post-synaptic) neuron.

Anatomically, the vertebrate brain is divided into three regions: the hindbrain, mind-brain and forebrain. The forebrain is further divided into two stages: the thalamic region and cerebrum. All vertebrates, including humans, follow the same basic structural plan [16]. The differences lie in the numbers of neurons and synaptic connections and the relative size of different areas. For example, the frontal cortex has expanded in primates relative to other animals [17]. Furthermore, the brain organization depends on the animal's ecological niche and lifestyle [18]. A monkey, which is a highly visual creature, has a larger proportion of its brain dedicated to vision than a mouse, which is a nocturnal animal having much of its brain dedicated to processing whisker and odor information.

III. NEURONAL MODELING ABSTRACTIONS

The sheer complexity of the brain also requires designers of brain models to choose an appropriate level of modeling abstraction for their study. Brain modeling abstractions can be viewed as a hierarchy, with each level serving a different purpose and incorporating the relevant subset of the brain's features. Lower level models represent biophysical interactions at the cellular level that incorporate molecular and electrical properties of neurons, ionic currents and protein receptors, as well as anatomical details of axonal branching and dendritic arborization [19, 20]. Such detailed models, while biologically accurate, incur tremendous computational costs for simulation and require costly parameter sweeps during modeling. Thus, large-scale simulation of the brain is extremely challenging at this level. At a level up, neural circuit models abstract away many molecular and cellular details and represent the brain as a massive circuit composed of four basic components: neurons for computation, synapses for learning and storage, axons for communication, and neuromodulatory systems to control action selection and learning. These neural circuit models are similar to analog and digital circuits and can take advantage of existing tools and simulation frameworks used in VLSI design. Functional models at higher abstraction levels simplify complex biological processes with more computationally efficient approximations. Although these models sacrifice biological accuracy, they stress functional aspects, have strong explanatory power, and are amenable for use in practical applications.

In the remainder of the paper we focus on brain models at the neural circuit level, using spiking neural networks as an exemplar for these neural circuit models.

IV. MODELING SPIKING NEURAL NETWORKS

In this section, we describe the elements of our spiking neural network models. Our simulator was first published in [14], but has been greatly enhanced to improve functionality and ease of use. More detailed mathematical descriptions can be found in [15].

A. Neuronal Modeling

As described above, we have chosen spiking elements for our model of the neuronal processing. Many modelers use mean firing rate models, which are efficient and can describe many aspects of cognition and neuroscience [21]. However, because these types of models average neural activity over hundreds of milliseconds, they lack the temporal resolution to capture many of the neural dynamics.

To strike the right balance of temporal resolution, biophysical accuracy, and efficient computation, we use the Izhikevich spiking neuron [22]. This neuronal model is a dynamical system in which there is a variable that tracks the neuron's membrane potential and a variable that tracks the recovery from a spiking event. Even though only four open parameters govern the dynamics of the model, it accurately replicates a wide range of neural behaviors.

B. Synaptic Modeling

The synaptic connection between two neurons is quite complex and is critical for neuronal dynamics, learning and memory. Integration of information takes place at the synapse. The state of the neuron, the type of neurotransmitter that is released at the synapse, and the time course of how this neurotransmitter is received are all important to integrating this input signal. This synaptic integration is called conductance and to capture this important aspect of neuronal processing, we model conductances that make the neuron more likely to fire a spike (excitatory) and less likely to fire a spike (inhibitory) over the appropriate timescales.

Learning and memory are mainly due to long lasting changes at the synaptic connections. In general, when a sufficient number of input spikes (i.e., pre-synaptic) arrive "together" (i.e., temporally coincident), the post-synaptic neuron generates a spike of its own. Neurons are continually modifying the strength of their connections based on the timing of their inputs. The synaptic strength of a connection is increased if pre-synaptic spikes come before the post-synaptic spikes; in other words, the temporal order of spikes is consistent with the input causing the output neuron to spike. If however, the order is anti-causal (i.e., input spike occurs after output), the synaptic weight is decreased. This learning rule is called spike timing dependent plasticity (STDP) and can enable a network to learn spatial-temporal patterns and contribute to network stability [23].

Plasticity on a shorter time scale can affect other types of memory. Working memory is the ability to keep information, such as a phone number in mind, over a brief period of time. Alternatively, it may be important to habituate or ignore information that is uninteresting or repetitive. To capture this type of memory we apply Short-Term Plasticity (STP) to the processing of a synaptic event [24]. STP is faster than STDP, on the order of 100ms, and contributes to working memory by pre-synaptic facilitation and habituation by pre-synpatic depression.

C. Creating a Simulation

To create a simulation of the brain, the modeler needs to define the different brain areas they want to investigate, the type of neuronal elements found in these brain areas, and the connections within and between brain areas. In our simulation, the brain areas are defined as groups of Izhikevich neurons. Groups can range from one neuron up to hundreds of thousands of neurons. Once the neuronal groups have been defined, the synaptic connections between them can be defined. Common connection topologies for building networks include: 1) All-to-all connectivity where all neurons in the pre-synaptic group are connected to all neurons in the post-synaptic group. 2) One-to-one connectivity where neuron i in the pre-synaptic group is connected to neuron i in the post-synaptic group assuming the same number of neurons in both groups. 3) Random connectivity where a group of pre-synaptic neurons are randomly connected to a group of post-synaptic neurons with a probability. 4) User-defined connectivity where a topology or projection from one group to another can be specified. For all connection types the user specifies an initial synaptic weight, a maximum synaptic weight, and a range of synaptic delays.

V. HW Architectures for Spike-based Computation

The advent of low-cost, high-performance graphics architectures (e.g., NVIDIA GPUs) opens the door for large-scale SNN simulations on affordable, programmable platforms. Some fundamental benefits and limitations of such graphics architectures for simulating SNNs are:

- Large fine-grained parallelism: Contemporary GPUs with hundreds of scalar processors can execute thousands of threads concurrently. Maximum performance can be achieved as long as a group of threads are executing the same instruction. However, when different threads within the same group require different instructions, *thread divergence* occurs, causing poor parallelism performance.
- Large off-chip memory bandwidth: A typical GPU's off-chip memory bus is based on a 512- (or 256-) bit wide DDR interface, resulting in a 5-fold increase in GPU memory bandwidth over the CPU. Through *memory coalescing*, this GPU memory bandwidth is exploited by clustering memory access patterns from different threads within a 128, 64, or 32-byte memory address space.
- Special Function Units: Each streaming multiprocessor may have multiple special functional units (SFUs) allowing single instruction calculation of exponentials and other mathematical functions.

The mapping of SNNs on to GPUs is non-trivial due to the random memory access structure in SNNs and large memory requirements to store the connectivity and network state (synapses and neurons) information. We proposed various optimization techniques to overcome the above limitations and effectively map SNNs on to GPUs including: (1) Exploiting both neuronal and synaptic parallelism to maximize thread level parallelism, (2) Efficient representation of large-scale SNNs that improves the off-chip memory coalescing, and (3) Minimizing thread divergence by delaying the execution of diverging conditions by buffering them and running them concurrently later. Using these optimization techniques, our SNN simulation with 100,000 neurons and 10 million synapses could be executed close to real-time [14]. Further, the GPU-SNN simulation was about 26 times faster than the CPU for 100K neurons with 50 million synapses.

Large-scale SNN simulations are memory dominated and hence overall speed-up is limited due to saturation of off-chip memory-bandwidth. Off-chip memory requirements can be reduced by providing large, persistent local shared memory to store the neuron state, and also by providing on-chip communication networks for direct spike transfer between multiprocessors. Such improvements have been incorporated into an application specific multiprocessor in the SpiNNaker project for large-scale SNN simulations [25]. But to truly approach the power and area efficiency of brain circuits it is essential to directly model neuronal circuits using hybrid analog-digital architectures [11, 25-27]. Memristor architectures have been proposed to implement neuronal circuits with high-efficiency [28]. These and other approaches are currently being investigated in the EU FACETS (http://facets.kip.uni-heidelberg.de/) and the DARPA SyNAPSE (http://en.wikipedia.org/wiki/SyNAPSE) projects.

Our approach is to design a simulator that is easy to use and yet provide significant computational performance on affordable, programmable platforms. We achieve this by using a PyNN-like interface and abstraction [29]. PyNN is a common programming interface developed by the neuronal simulation community to allow a single script to run on various simulators. In addition, to ensure our simulator can be supported on a wide-range of machines, our simulator runs on both generic x86 CPUs and NVIDIA GPUs under Windows and Unix operating systems.

VI. Challenges for the VLSI-CAD Community

Although significant progress has been made towards specification and simulation of large-scale brain networks on a variety of hardware platforms, many challenges still remain open. Foremost among them is the problem of tuning and stabilization of these large-scale dynamical systems, which are characterized by a large number of state variables and open parameters. Examples of parameters of these large-scale systems include the connectivity profile, synaptic plasticity rules and synaptic strength limits, neuronal properties, etc. These parameters govern the network dynamics, and tuning of these parameters to generate stable, realistic behaviors is extremely challenging. Biology does provide some empirical data that can constrain these systems, but many parameter values must be chosen by the modeler to achieve appropriate neuronal dynamics. Another challenge remains in understanding and interpreting the spike trains generated by these large-scale systems. New approaches based on information theory and data mining need to be used to decipher the temporal aspects of spikes from a neural population over different timescales. Furthermore, new tools and approaches are necessary to visualize the state variables, and connectivity of large cortical networks in real-time. These issues, while challenging, provide research opportunities for the VLSI-CAD community.

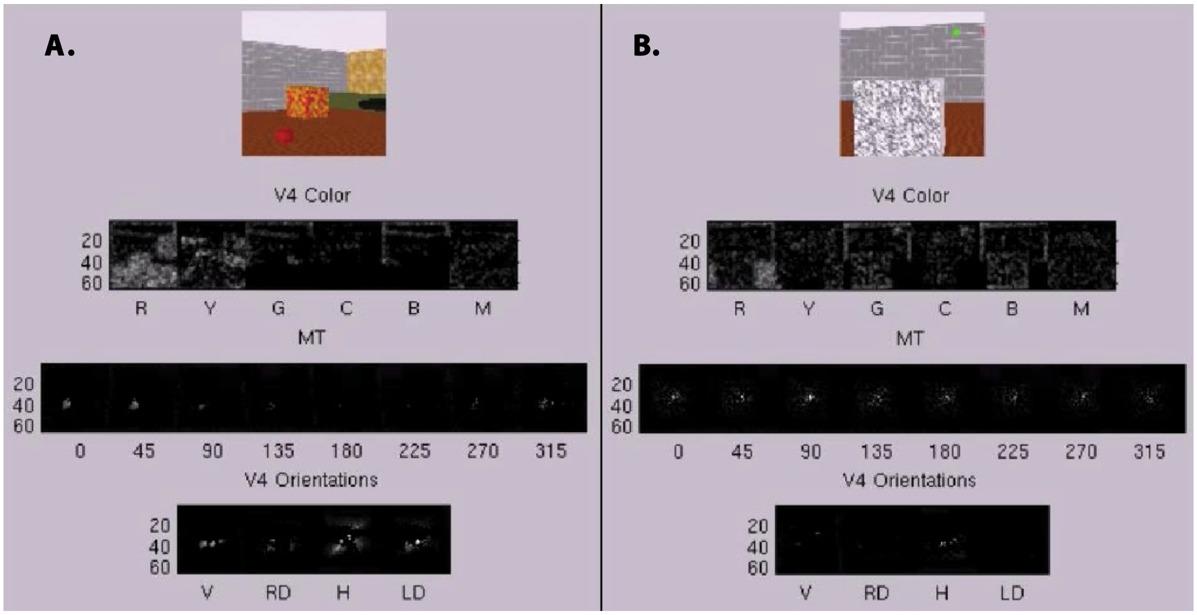

Figure 1. Snapshots from the large-scale simulation of cortical visual processing. The image at the top is a video frame, at 64x64 pixel resolution, provided by the CASTLE virtual environment (http://www.setcorp.com/science-info.html) and fed as input to the V1 layer of the cortical model. Below the image are neuronal responses to color (V4 Color), motion (MT) and edges (V4 Orientations). The complete model contains 552,960 neurons and 118,188,000 synapses.

VII. Case Study: Large-Scale Simulation of Cortical Visual Processing

In order to demonstrate the power and ease of use of our simulator, we have built a large-scale, spiking network to simulate models of the visual cortex for color, orientation selectivity, and motion selectivity [15]. In the model, the primary visual cortex (V1) is a rate-based preprocessor, which calculates color opponency responses [30], as well as motion energy responses [31]. These rate-based responses are converted to Poisson spike trains and fed into higher cortical regions, such as V4 for color and edge detection, and MT for motion detection. All excitatory neurons are Regular Spiking and all inhibitory neurons are Fast Spiking, as defined by Izhikevich [32]. An input image resolution of 32x32 or 64x64 pixels was used; and this resolution was then used for every layer in the network. Figure 1 shows snapshots from the complete model.

A. Cortical model of color selectivity

Following the opponent-color theory, we constructed a rate-based model of area V1 where we had center-surround units that were selective to 1) red center, green surround; 2) green center, red surround; 3) yellow center, blue surround; and 4) blue center, yellow surround. Color opponent signals were then converted to spike-trains using Poisson spike generators and connected to populations that are selective to one of 6 colors: red, green, blue, yellow, magenta (blue+red) and cyan (blue+green). Each color had both an excitatory and inhibitory group, for a total of 12 V4 color groups.

B. Cortical Model of Motion and Orientation Selectivity

To generate motion selective responses, we used the Simoncelli and Heeger motion energy model for V1 [31] and implemented it to run on the GPU. The probability of a connection from V1 to MT was proportional to the projection of the V1 cell's receptive field onto a plane in the spatial frequency-temporal frequency domain. The slope of this plane defined the speed preference of the resulting MT cell and the rotation of the plane around the time axis defined the direction preference. The neurons in our MT model responded preferentially to one of 8 different directions and 3 different speeds at a spatial location. The response of the MT neuron's receptive field was based on connectivity from the V1 motion selective neurons. The motion energy responses were also used to generate orientation selective responses. Since there are units in the 28 space-time filters that are more selective to orientation than motion, their responses can be used to generate a population of V4 orientation selective units.

C. Performance of the Complete Cortical Model

The model demonstrated a wide range of neurophysiological and behavioral responses similar to those found in humans and non-human primates. V4 neuronal units had selectivity to their preferred colors and the appropriate broad tuning to a range of hues similar to those found in the monkey V4 brain region [33]. The model replicated Random Dot Kinematogram (RDK) experiments, performed with monkeys and humans, in which subjects must report the direction of dots moving coherently among a background of randomly moving dots. The model's output was comparable to human psychophysical experiments [34].

The complete model had 138,240 neurons and 29,547,000 synapses at a spatial resolution of 32 by 32. Running these 138K neurons on a single NVIDIA C1060 GPU ran 5% faster than real-time and roughly 65 times faster than on the CPU. At 64 by 64 input resolution network, the model which contained 552,960 neurons and 118,188,000 synapses, did not fit on a

GPU and required being run on the CPU (see Figure 1). An advantage of this model, compared to other large-scale cortical models, is that it can replicate experimental results, as well as be used in application domains. It is constructed such that it can readily receive real image or video data for neurobiologically based machine vision, and so that it can be expanded across multiple GPUs to handle greater resolution in real-time. Moreover, by providing this interface to the real world, our simulation environment can close the loop between sensory input and motor output resulting in interesting behaviors.

While these models are considered large-scale, they are still orders of magnitude smaller than human (10^{11} neurons) or even mouse (10^{7} neurons) brains. Thus, many tools must be developed in order to ultimately reach these scales. Because we believe the complexity and dynamics, which are the hallmark of mammalian brains, are a function of large-scale population dynamics, such large-scale simulations will be necessary for the study of brain function and the construction of artificial brains that are truly intelligent.

VIII. Conclusion

Neuromorphic modeling is emerging due to the hope of brain-like intelligence in artificial systems and the advent of new hardware that is amenable to the parallel and distributed brain architecture. We outlined the major modeling abstractions for the brain and focused on spiking neuronal models that deliver good performance while ensuring biological accuracy. We presented a simulation environment that supports the construction of large-scale models of spiking neural networks on hardware readily available to most researchers. We also presented a case study for simulation of the visual cortex using a large-scale spiking neural network. It is our hope that this environment will benefit both the neuroscience and computer science communities and move us closer to meeting the grand challenge of reverse-engineering the brain.

References

[1] O. Sporns, *Networks of the Brain*. Cambridge, MA: The MIT Press, 2011.

[2] J. M. Nageswaran, M. Richert, N. Dutt and J. L. Krichmar, "Towards reverse engineering the brain: Modeling abstractions and simulation frameworks," in *VLSI System on Chip Conference (VLSI-SoC), 2010 18th IEEE/IFIP*, 2010, pp. 1-6.

[3] NationalAcademyEngineering. (2011, *Grand challenges for engineering: Reverse-engineer the brain, National Academy of Engineering*. Available: http://www.engineeringchallenges.org/cms/8996/9109.aspx

[4] J. S. Albus*, et al.*, "A proposal for a Decade of the Mind initiative," *Science,* vol. 317, p. 1321, Sep 7 2007.

[5] R. Ananthanarayanan, S. K. Esser, H. D. Simon and D. S. Modha, "The cat is out of the bag: cortical simulations with 10^9 neurons, 10^{13} synapses," presented at the Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, Portland, Oregon, 2009.

[6] H. Markram, "The blue brain project," *Nat Rev Neurosci,* vol. 7, pp. 153-60, Feb 2006.

[7] E. S. Lein*, et al.*, "Genome-wide atlas of gene expression in the adult mouse brain," *Nature,* vol. 445, pp. 168-76, Jan 11 2007.

[8] R. Ananthanarayanan, S. K. Esser, H. D. Simon and D. S. Modha, "Cortical Simulations with 10^9 Neurons and 10^13 Synapses.," presented at the Supercomputing 09: Proceedings of the ACM/IEEE SC2009 Conference on High Performance Networking and Computing, Portland, OR, 2009.

[9] E. M. Izhikevich and G. M. Edelman, "Large-scale model of mammalian thalamocortical systems," *Proc Natl Acad Sci U S A,* vol. 105, pp. 3593-8, Mar 4 2008.

[10] G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books, Inc., 1987.

[11] K. Boahen, "Neuromorphic Microchips," *Sci Am,* vol. 292, pp. 56-63, May 2005.

[12] R. Douglas, M. Mahowald and C. Mead, "Neuromorphic analogue VLSI," *Annu Rev Neurosci,* vol. 18, pp. 255-81, 1995.

[13] T. P. Vogels, K. Rajan and L. F. Abbott, "Neural network dynamics," *Annu Rev Neurosci,* vol. 28, pp. 357-76, 2005.

[14] J. M. Nageswaran, N. Dutt, J. L. Krichmar, A. Nicolau and A. V. Veidenbaum, "A configurable simulation environment for the efficient simulation of large-scale spiking neural networks on graphics processors," *Neural Netw,* vol. 22, pp. 791-800, Jul-Aug 2009.

[15] M. Richert, J. M. Nageswaran, N. Dutt and J. L. Krichmar, "An Efficient and Flexible Simulation Environment for Modeling Large-Scale Cortical Processing," *Frontiers in Neuroinformatics,* In press.

[16] L. W. Swanson, "Anatomy of the soul as reflected in the cerebral hemispheres: neural circuits underlying voluntary control of basic motivated behaviors," *J Comp Neurol,* vol. 493, pp. 122-31, Dec 5 2005.

[17] E. C. Bush and J. M. Allman, "The scaling of frontal cortex in primates and carnivores," *Proc Natl Acad Sci U S A,* vol. 101, pp. 3962-6, Mar 16 2004.

[18] L. Krubitzer, "The magnificent compromise: cortical field evolution in mammals," *Neuron,* vol. 56, pp. 201-8, Oct 25 2007.

[19] J. M. Bower and D. Beeman, "Constructing realistic neural simulations with GENESIS," *Methods Mol Biol,* vol. 401, pp. 103-25, 2007.

[20] M. L. Hines and N. T. Carnevale, "NEURON: a tool for neuroscientists," *Neuroscientist,* vol. 7, pp. 123-35, Apr 2001.

[21] J. L. Krichmar and G. M. Edelman, "Brain-Based Devices for the Study of Nervous Systems and the Development of Intelligent Machines," *Artificial Life,* vol. Vol. 11, pp. 63-78, 2005.

[22] E. M. Izhikevich, "Which model to use for cortical spiking neurons?," *IEEE Trans Neural Netw,* vol. 15, pp. 1063-70, Sep 2004.

[23] S. Song, K. D. Miller and L. F. Abbott, "Competitive Hebbian learning through spike-timing-dependent synaptic plasticity," *Nat Neurosci,* vol. 3, pp. 919-26, Sep 2000.

[24] H. Markram, Y. Wang and M. Tsodyks, "Differential signaling via the same axon of neocortical pyramidal neurons," *Proc Natl Acad Sci U S A,* vol. 95, pp. 5323-8, Apr 28 1998.

[25] J. Navaridas, M. Luj, J. Miguel-Alonso, L. A. Plana and S. Furber, "Understanding the interconnection network of SpiNNaker," presented at the Proceedings of the 23rd international conference on Supercomputing, Yorktown Heights, NY, USA, 2009.

[26] K. Boahen, "Neurogrid: emulating a million neurons in the cortex," *Conf Proc IEEE Eng Med Biol Soc,* vol. Suppl, p. 6702, 2006.

[27] V. Rangan, A. Ghosh, V. Aparin and G. Cauwenberghs, "A subthreshold aVLSI implementation of the Izhikevich simple neuron model," *Conf Proc IEEE Eng Med Biol Soc,* vol. 2010, pp. 4164-7, 2010.

[28] S. H. Jo*, et al.*, "Nanoscale memristor device as synapse in neuromorphic systems," *Nano Lett,* vol. 10, pp. 1297-301, Apr 14 2010.

[29] A. P. Davison*, et al.*, "PyNN: A Common Interface for Neuronal Network Simulators," *Front Neuroinformatics,* vol. 2, p. 11, 2008.

[30] M. S. Livingstone and D. H. Hubel, "Anatomy and physiology of a color system in the primate visual cortex," *J Neurosci,* vol. 4, pp. 309-56, Jan 1984.

[31] E. P. Simoncelli and D. J. Heeger, "A model of neuronal responses in visual area MT," *Vision Res,* vol. 38, pp. 743-61, Mar 1998.

[32] E. M. Izhikevich, J. A. Gally and G. M. Edelman, "Spike-timing dynamics of neuronal groups," *Cereb Cortex,* vol. 14, pp. 933-44, Aug 2004.

[33] Y. Kotake, H. Morimoto, Y. Okazaki, I. Fujita and H. Tamura, "Organization of color-selective neurons in macaque visual area V4," *J Neurophysiol,* vol. 102, pp. 15-27, Jul 2009.

[34] A. Resulaj, R. Kiani, D. M. Wolpert and M. N. Shadlen, "Changes of mind in decision-making," *Nature,* vol. 461, pp. 263-6, Sep 10 2009.