Frontiers in neuromorphic engineering

Giacomo Indiveri[1]* and Timothy K. Horiuchi[2]

[1] *Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland*
[2] *Department of Electrical and Computer Engineering, Institute for Systems Research, University of Maryland, College Park, MD, USA*
*Correspondence: giacomo@ini.phys.ethz.ch

Neurobiological processing systems are remarkable computational devices. They use slow, stochastic, and inhomogeneous computing elements and yet they outperform today's most powerful computers at tasks such as vision, audition, and motor control, tasks that we perform nearly every moment that we are awake without much conscious thought or concern. Despite the vast amount of resources dedicated to the research and development of computing, information, and communication technologies, today's fastest and largest computers are still not able to match biological systems at robustly accomplishing real-world tasks. While the specific algorithms and representations that biological brains use are still largely unknown, it is clear that instead of Boolean logic, precise digital representations, and synchronous operations, nervous systems use hybrid analog/digital components, distributed representations, massively parallel mechanisms, combine communications with memory and computation, and make extensive use of adaptation, self-organization, and learning. On the other hand, as with many successful man-made systems, it is clear that biological brains have been co-designed with the body to operate under a specific range of conditions and assumptions about the world.

Understanding the computational principles used by the brain and how they are physically embodied is crucial for developing novel computing paradigms and guiding a new generation of technologies that can combine the strengths of industrial-scale electronics with the computational performance of brains.

NEUROMORPHIC ENGINEERING

While the history of implementing electronic models of neural circuits extends back to the construction of perceptrons (Rosenblatt, 1958) and retinas (Fukushima et al., 1970), the modern wave of research utilizing VLSI technology and emphasizing the non-linear current characteristics of the transistor began in the mid 1980s with the collaboration that sprung up between prominent scientists Max Delbrück, John Hopfield, Carver Mead, and Richard Feynman (Hey, 1999). Inspired by graded synaptic transmission in the retina, Mead sought to use the graded (analog) properties of transistors, rather than simply operating them as on–off (digital) switches. He showed that analog neuromorphic circuits share many common physical properties with protein channels in neurons (Mead, 1989). As a consequence, these types of circuits require far fewer transistors than digital approaches to emulating neural systems.

Through the *Physics of Computation* course at Caltech (led by Carver Mead, John Hopfield, and Richard Feynman), Mead (1989)'s textbook *Analog VLSI and Neural Systems*, and the creation of the *Telluride Neuromorphic Engineering Workshop*, the field of Neuromorphic Engineering was established. Prominent in the early expansion of the field were scientists and engineers such as Christof Koch, Terry Sejnowski, Rodney Douglas, Andreas Andreou, Paul Mueller, Jan van der Spiegel, and Eric Vittoz, training a generation of cross-disciplinary students.

It has been argued that neuromorphic circuits are ideal for developing a new generation of computing technologies that use the same organizing principles of the biological nervous system (Douglas et al., 1995; Boahen, 2005; Sarpeshkar, 2006). In addition to the computations of a single neuron, many neuromorphic circuits also utilize spiking representations for communication, learning and memory, and computation. The use of asynchronous spike- or digital event-based representations in electronic systems can be energy-efficient and fault-tolerant, making them ideal for building modular systems and creating complex hierarchies of computation.

The most successful neuromorphic systems to date have been single chip devices that emulate peripheral sensory transduction such as silicon retinas, visual motion sensors, and silicon cochleas for a wide variety of applications. In recent years, many larger multi-chip neuromorphic systems have begun to emerge that have raised new issues and challenges. These systems typically comprise one or more neuromorphic sensors, interfaced to general-purpose neural network chips using spiking silicon neurons and dynamic synapses.

The method used to transmit spikes across chip boundaries in these systems is based on the address-event representation (AER; Mahowald, 1994). It is an asynchronous digital communication protocol that sends the *address* of the neuron that emitted the event in real-time or close to real-time. The information being transmitted may be analog or digital, but must be communicated via spikes, thus raising the critical and exciting issue of signal encoding that is currently a very active topic in neuroscience. Signals can be encoded in the mean frequency (rate) of spikes, in their precise timing with respect to a time reference, or in the population response. In fact, multiple signals can be simultaneously encoded in a single spike train. Once on a digital bus, the address-events can be remapped to multiple destinations using commercially available synchronous or custom asynchronous processing. Digital AER infrastructures allow us to construct large multi-chip networks with nearly arbitrary connectivity and to dynamically reconfigure the network topology for experimentation. By using analog circuits for local computations on-chip and digital circuits for long-distance communication (off-chip), neuromorphic systems can exploit the best of both worlds.

Another distinguishing feature of neuromorphic engineering has been the integration of fine-grained synaptic modification mechanisms that both enable these networks to change their behavior with experience (as is ubiquitous in biological nervous systems) and to implicitly overcome the inherent device parameter variability

found in all manufacturing technologies whether silicon or mechanical. The most prominent storage mechanisms have been: on-chip capacitance, on-chip floating-gate charge storage, and off-chip AER remapping of the network to either dynamically change the connectivity or to implement stochastic spike delivery. To implement biologically plausible learning rules (e.g., spike-timing dependent plasticity), many of these implementations also incorporate learning circuits directly at the synapse.

FRONTIERS IN NEUROMORPHIC ENGINEERING
At its heart, neuromorphic engineering is about the real-time interaction of the algorithm with its physical implementation and the environment in solving tasks. This synergy is easy to appreciate at the sensory and motor interfaces with the world, but more subtle and interesting when considering cognitive-level tasks.

With increasing knowledge of what single neurons and their synapses can do computationally, the desire for more sophisticated implementation technologies has grown. At present, new technologies such as nano-scale transistors, quantum devices, organic electronics, memristors, phase-change materials, 3D integrated circuits, and electro-active polymers for actuation are all promising directions for research.

Neuromorphic engineering now aims to use these technologies for developing larger-scale neural processing systems and move from the predominantly feed-forward, reactive neuromorphic systems of the past to adaptive behaving ones that can be considered cognitive. For example, a key mechanism in cognition, selective attention, has long been part of the neuromorphic engineering toolkit, but has largely operated as a bottom-up process, operating on short-term information and memory. Expanding its role in top-down behavior (e.g., guiding the learning of more abstract concepts) will be important for understanding and implementing context-dependent behavior.

While the majority of cognitive architectures and their software implementations have avoided detailed neural implementations, due to limited computational power and the assumption that the details of single spiking neurons are not important at this level, a growing number of research groups worldwide have begun to consider the consequences of biologically plausible implementations at both the level of neural fields and single spiking neurons. By providing real-time spiking implementations of core neural circuits, neuromorphic engineering will play an important role in the development and fielding of biologically relevant working models of cognition interacting with the real-world.

One of the Grand Challenges of Neuromorphic Engineering is to demonstrate cognitive systems using hardware neural processing architectures integrated with physical bodies (e.g., humanoid robots) that can solve everyday tasks in real-time. To be successful in this ambitious endeavor, an integrated multi-disciplinary approach is critical that brings together research in:

- *VLSI circuits and systems* for implementing hardware models of neural processing systems, mixed analog/digital asynchronous AER communication infrastructures, spike-based sensory–motor systems, and event-driven processing methods;
- *emerging technologies* including 3D VLSI, nanotechnologies, phase-change materials, and memristive devices, applied to the construction of low-power neuromorphic systems;
- *robotic platforms and control* with particular focus on new actuators and materials, compliant systems, contraction theory and controllability of complex systems, and on the computational role of the physical body in locomotion and active sensing;

- *neural computation,* involving studies of spiking winner-take-all networks, attractor networks, mean-field theory, spike-based learning mechanisms, probabilistic graphical models, cortical development, and self-constructing principles;
- *biologically plausible cognitive architectures* for studying attention, working memory, state-dependent computation, action selection mechanisms, planning, and multi-agent interaction.

Through this journal, we intend to encourage the presentation of these diverse perspectives, technical approaches, and goals, to facilitate the development of neuromorphic cognitive systems, and reach new frontiers in neuromorphic engineering.

REFERENCES
Boahen, K. A. (2005). Neuromorphic microchips. *Sci. Am.* 292, 56–63.
Douglas, R. J., Mahowald, M. A., and Mead, C. (1995). Neuromorphic analogue VLSI. *Annu. Rev. Neurosci.* 18, 255–281.
Fukushima, K., Yamaguchi, Y., Yasuda, M., and Nagata, S. (1970). An electronic model of the retina. *Proc. IEEE* 58, 1950–1951.
Hey, T. (1999). Richard Feynman and computation. *Contemp. Phy.* 40, 257–265.
Mahowald, M. (1994). *An Analog VLSI System for Stereoscopic Vision.* Boston, MA: Kluwer.
Mead, C. A. (1989). *Analog VLSI and Neural Systems.* Reading, MA: Addison-Wesley.
Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* 65, 386–408.
Sarpeshkar, R. (2006). Brain power – borrowing from biology makes for low power computing – bionic ear. *IEEE Spectr.* 43, 24–29.