Neuromorphic Systems: past, present and future

Leslie S. Smith

Abstract Neuromorphic systems are implementations in silicon of elements of neural systems. The idea of electronic implementation is not new, but modern microelectronics has provided opportunities for producing systems for both sensing and neural modelling that can be mass produced straightforwardly. We review the the history of neuromorphic systems, and discuss the range of neuromorphic systems that have ben developed. We discuss recent ideas for overcoming some of the problems, particularly providing effective adaptive synapses in large numbers.

1 Introduction

Neuromorphic systems are electronic implementations of neural systems. Such implementations may take place at a number of different levels. For example, one may model sensory or sensorimotor systems, or one may model specific neural systems at many different levels, ranging through (at least) whole brain, brain region, cortical column, mid-brain or brainstem nucleus, neural microcircuits, single neurons, structural parts of neurons (dendrites, axons, soma), patches of membrane, down to ion channels encased in the neural bilipid membrane. Some go further, suggesting modelling quantum effects at synapses and in the dendrite [Ham07] [Hir91], but there do not appear to be electronic (as opposed to software) implementations of these models.

Different technologies have been used for these implementations at different times, reflecting the prevailing electronic technologies, and current systems are (of course) implemented in either analogue or digital very large scale implementation (aVLSI or dVLSI). According to Wikipedia (http://en.wikipedia.org/wiki/Neuromorphic) the term was coined by Carver Mead (see below), in the late 1980's, but the ideas

Leslie S. Smith
Department of Computing Science and Mathematics, University of Stirling, e-mail: l.s.smith@cs.stir.ac.uk

have roots that go back well before this. The idea of an equivalent circuit for a neuron goes back at least to 1907 ([Lap07]), where a neuron is modelled by a resistor and a capacitor. Since then, various technologies have been used to model both neurons and sensory surfaces, and these are discussed in sections 2 to 4. There are particular issues that arise in these different types of models, such as matching the speed of the implementation to the environmental changes of interest, or finding ways to store large numbers of possibly adaptive synaptic weights. More recently, a number of new methods of using VLSI technology have been proposed for implementation, as well as some novel techniques for storing adaptable weights, and these are reviewed in section 5.

A different form of electronic implementation of neural systems is to directly implement a formal model of a neuron, or of a neural system, for example the McCulloch-Pitts neuron [MP43], or the weightless systems used by Aleksander and Stonham [AS79]. These have been the basis for silicon implementations: in the McCulloch-Pitts case, one example (of many) is the Intel chip [Cor90], and in the weightless case, the work of Aleksander et al simply using memory technology [ATB84], and Austin using more specialised silicon implementations [AK98]. We do not discuss these or related work further in this chapter, since they are not strictly "neuromorphic": nonetheless, these remain interesting approaches.

2 Early forms of neuromorphic systems: systems built from discrete components

The 20th century saw huge leaps in the capabilities of electronic components. Although these were used initially primarily for communication purposes, and later for digital calculation and computation, a small number of researchers also saw electronic technology as a mechanism for modelling neural systems. Indeed, Hodgkin and Huxley's work on the dynamics of active cells [HH52] is often presented both as a set of equations and as an electrical equivalent circuit. In the 1960's for example, FitzHugh used analog computers to model single neurons, gaining detailed insights into their operation from these implemented models [Fit66] (see http://www.scholarpedia.org/article/FitzHugh-Nagumo_model). Others followed with simple electronic models of neural processes [Lew68][JH69],[Roy72] [Bro79]. On the larger scale, Runge [RUV68] built a fairly detailed model of the pigeon retina using 145 Cadmium Sulphide sensors and 381 neural analogue circuits taking up 50 circuit boards. Early electrical models of the cochlea are discussed in [KS59], including a transmission line of 178 sections, with each section comprising two inductors and four capacitors. Although this type of work did permit the building of electronic models, and their testing (and comparison with real neural and sensory systems), they were not intended for direct incorporation into equipment, but were purely research models of particular systems.

The approach of using discrete components is limited by the complexity and expense of building such systems. It is not practical to build large numbers of such

systems, and as a result, they were useful for research purposes, attempting to improve understanding of the original system by re-creating it in hardware, but not for more general applications.

3 Modern Neuromorphic Systems

In 1989, Carver Mead published his second seminal book (his first was [MC80]), "Analog VLSI and Neural Systems" [Mea89]. This book brought together ideas on the ways in which analogue field-effect transistor based circuits, particularly those operating in the subthreshold domain (where currents change exponentially with gate voltage change), were similar operationally to neural membranes. The book describes at length how circuits which emulate elements of neural systems may be constructed. It also contains some highly influential system examples, drawn from the auditory and visual domain. Even now, most current works in systems reference some of the content of this book. It laid down a path for those who wanted to implement models of neural systems directly in VLSI hardware. This has the major advantage of being easy to replicate, so that a successful design can be built relatively cheaply in bulk and possibly become a low-cost system component. However, unlike systems built from discrete components, it is not possible to alter the circuitry once it has been built, nor even to monitor internal voltages unless that voltage is made available off-chip. Further, there is generally a long delay between submission of a design to a chip foundry and receiving testable chips back. These problems make developing such systems much more difficult.

Nonetheless, researchers in a number of groups have worked to develop circuits for a number of different neural types of applications, based partly on the ideas and circuits in Mead's book. Quite a number of novel circuits developed by the analogue designers in the neuromorphic systems community are discussed in [LKI+02]. Some of these are essentially related to sensory systems, whilst others are more related to the actual underlying neurons themselves. In addition there is a body of work on emulating aspects of the cortex itself, whether modelling the components of the cortex or attempting to build complete models of small parts of cortex [MG07], as in the Blue Brain project (http://bluebrain.epfl.ch/: these are primarily software simulations, and are outside the scope of this review. Below, we present a brief review of the current range of neuromorphic systems: a full-scale review is beyond the scope of this article.

3.1 Neuromorphic systems for sensory processing

Implementations of neural systems are generally only considered to be neuromorphic if they work in real time. Real time operation is an absolute requirement for sensory systems, and many the neuromorphic systems have been targeted specifi-

cally at sensory systems. Below, we review particularly visual, auditory, olfactory and tactile neuromorphic systems. There has recently also been interest in proprioceptary sensors as well [WZPF06]. So far as the author is aware, there have not been any neuromorphic implementation of systems for taste! The other characteristic of neuromorphic implementations is that they operate in parallel, and this has particular advantages for sensory systems as discussed in section 3.1.4.

3.1.1 Neuromorphic systems for vision

The earliest neuromorphic visual system appears to be that of Runge, discussed earlier. In terms of VLSI based designs, one of the earliest is described in Mead's book, chapter 15, which was a version of [MM88]. This paper discusses a "silicon retina" which implements both the transduction (basic photoreceptor circuit which has a near-logarithmic response), and a horizontal resistive layer which models the outer plexiform layer of the retina. The original retina had 48 by 48 pixels. The system produces an output which is the difference between the centre intensity and a weighted average of the the surrounding intensities, which is quite unlike what happens on a digital camera. The overall effect is that the response to a static edge is a spatial derivative, rather like what happens in a real retina. In addition, the response to a featureless surface is essentially zero independent of the brightness.

Many other papers have developed the ideas within this paper. Better photoreceptors are proposed and analysed in detail in [DM96]. These are used to develop a more effective contrast-sensitive retina in [AMSB95]. By performing additional processing at the silicon retina, one can make the system able to model specific visual capabilities. For example, a time to collision detector[Ind98], and a model of the fly elementary motion detector has been built [HK98], and more recently a depth from motion system [YMW$^+$06]. One particularly interesting recent advance has been a system which detects intensity changes in an overall brightness independent way, and transmits these serially using AER, thus enabling sensing of rapidly altering visual scenes without the huge data rates implied by the need to process whole frames [LPD08].

3.1.2 Neuromorphic systems for audition

There has been considerable interest in cochlear models since the work of Helmholtz and von Bekesy on the nature of the transduction of the pressure wave in the cochlea. Unlike the case in vision, pressure wave to electrical signal transduction is external to the device, and uses a (traditional) microphone. Early electrical models of the cochlea [KS59] discussed earlier were large and unwieldy. Building neuromorphic auditory subsystems that might actually be used to provide auditory capabilities for whole systems really had to wait for VLSI capabilities of the type discussed by Mead. Some of the earliest integrated neuromorphic systems are discussed in Mead's book. Lyon and Mead [LM88] describe the implementation of a sequence

of filters used to build what they called a silicon cochlea: implementation techniques have been codified in [FA95] . This work has been extended to be more biologically realistic in [LM89a], and applied to auditory localization [LM89c], pitch perception [LM89b], voiced speech [LAJ93b] and speech analysis [LW97], [LAJ93a]. The real cochlea is active, and attempts have been made to implement this in [FvSV97], and much more recently in[HJvST08]. This is particularly important because of the very wide dynamic range of the biological cochlea, and the way in which the selectivity alters with changing sound pressure level.

It is well known that early auditory processing is not simply a matter of transduction, and researchers have considered the processing that occurs in the auditory brainstem as well. Considerable work has been done on silicon modelling of the cochlear nucleus, the first port of call of the axons of the auditory nerve [SFV96]. Some of the neurons in the brainstem nuclei respond to what are clearly features, for example amplitude modulation [SV97] or onsets [GHS02].

3.1.3 Other sensory neuromorphic systems

Neuromorphic systems have been designed for other sensory modalities as well. There has been considerable interest in olfaction: the electronic nose is a device that could have considerable application in various industries (such as brewing and perfumery). In this sensory domain, the sensors detect electrical changes due to the odorant molecules [SG92]. It is possible to integrate the sensors on to the actual CMOS chip [Pea97]. Although the idea is relatively straightforward (altering the electrical properties of an insulating (polymer) layer as a result of the arrival of air-borne molecules), there are difficulties both in delivery, and, because of the chemical nature of the sensing, due to issues of drift and poisoning. A spike-based implementation is described in [KHT$^+$05].

Tactile sensing systems transduce pressure or motion into electrical signals. Neuromorphic motion based systems based on models of rodent vibrissae have been developed: earlier versions do not generally have the vibrissae directly incorporated on to the CMOS VLSI system, but use the outputs from these sensors directly, and are intended for robot based applications [PNG$^+$06]. More recently, both the sensor and the electronics have been integrated [AHW$^+$07]. Skin-like sensor arrays have been developed as well: these are of interest as sensors for grippers, and more generally, as general tactile sensors. A capitative technique is used in [Roy06], to produce an array of 59 sensing units which he calls *taxels*, and a polycrystalline silicon technique is used in [VAV$^+$06].

3.1.4 Parallelism in sensory neuromorphic systems: greedy processing

Hardware implementations, particularly fully-implemented ones, in the sense used by Hecht-Nielsen [HN90] (page 267), have the advantage of permitting true parallelism, unlike software implementations, or even systems which are only partially

implemented (where, for example, a digital floating point multiplier is shared between a number of synapses). As as result, they can perform numerous different transformations on incoming (parallel) sensory signals simultaneously. This processing can proceed all the time, whether it is needed or not. This "greedy" processing means that values are available immediately should they be required. Such processing appears to be the case in animal systems, where (for example) in the brainstem auditory nuclei, all the auditory nerve signals appear to be continuously processed to produce a representation which arrives at the mid-brain inferior colliculus. The visual domain is a little different, since the foveal section of the retina can only examine a small visual angle at a time: there, the parallelism seems to occur in the early stages of cortical visual processing (particularly V1).

In neuromorphic systems, this greediness takes the form of, for example, continuously processing signals from all the pixels, to produce event based signals which may then be sent down an AER bus (as in [LPD08]). In the auditory domain, the numerous band-passed signals are simultaneously processed to search for onsets or for amplitude modulation, mirroring processing processing the the auditory brainstem [SV97]. Greedy processing is useful for computing features in ays that are invariant under expected alterations in the sensory processing environment (such as overall illumination changes in vision, or level variation in audition). In addition, using greedy processing should mean that a higher-level system which required particular information about some part of an image, or some part of the spectrum would be able to retrieve this immediately without having to wait of it to be computed.

3.2 Neuromorphic models of neural circuitry, neurons and membranes

In [Mea89] Mead draws an analogy between ionic currents passing across cell membranes through ion channels and electron currents through field effect transistors operating in subthreshold mode. He also provides a description of an implementation of neural axons (chapter 12 of [Mea89]). Since that time there has been considerable interest in neuromorphic models of neurons, implemented at a range of different levels, from simple single compartment models of neurons through to patches of cell membrane which could be assembled into multi-compartment neuron models. We delay discussion of neuromorphic models of synapses to section 5. At a rather higher level, there has been work on modelling neural circuits at a range of different levels: mostly this work has been in software, but there is growing interest in hardware based implementations.

3.2.1 Single compartment neuromorphic models

Single compartment neural models treat the neuron's state variable (sometimes called activation, and sometimes described as the depolarisation of the neuron, de-

pending on the level of neural realism intended) as a single value: they thus ignore the complexity of the dendrites, and consider all the conductances as lumped together. These models generally generate a spike (which may be a realistic neural spike, or simple an event characterised purely by time of occurrence) when this activation crosses some positive threshold from below. Such models are useful (and indeed, commonplace in simulations, since both integrate-and-fire neurons and spike response models (see sections 4.1 and 4.2 of [GK02]) are of this form. Even such straightforward model neurons can have different degrees of faithfulness to reality: for example, they may (or may not) implement a refractory period (time after firing when the neuron cannot fire), relative refractory period (time after firing when it is possible, but more difficult to make the neuron fire), and may have all the conductances from the membrane gathered into a single conductance (or "leak"), or may consider a number of conductances independently.

One of the earliest neuromorphic single compartment models was [WMT96] in which a quite detailed leaky integrate and fire (LIF) neuron was implemented: it also exhibited other characteristics of real neurons, such as facilitation, accommodation and post-inhibitory rebound. A more recent model built in silicon exhibits the spiking behaviour of a number of classes of cortical neurons [WD08]. Others have been more interested in using LIF neurons as system components, aiming to use their computational properties rather than model real neurons: these are used by [SV97] as part of his amplitude modulation detecting system, and by [GHS02] as part of a sound analysis system. In these cases, the neuron implemented was a much simpler form of the leaky integrate and fire neuron. The primary advantage of these simpler models is that they still display useful computational capabilities (such as synchronization and co-incidence detection), but require less circuitry to implement. As a result it becomes possible to implement larger numbers of them on a single chip.

3.2.2 Neuromorphic models of elements of neurons

Single compartment models entirely ignore the dendrites of neurons, yet these are frequently large and complex structures. There has been interest in neuromorphic implementations of dendrites as the timings of the different inputs to these, and the way in which they are combined can provide powerful computational capabilities even before signals reach the soma of the the neuron. In the models developed in [Eli93][NE96] straightforward linear summation can occur, but interaction between nearby synapses can permit discrimination between different pulse intervals and patterns, as well as detecting correlations between spike trains. A different aspect of dendrites is their ability to transmit signals integrated from synapses forward (towards the soma) at the same time as transmitting action potentials backwards, and this has been demonstrated in a neuromorphic circuit in [RD01]. This capability is important for determining when synaptic characteristics should be altered. At the other end of the neuron, [MHDM95] describe a neuromorphic implementation of

an axon, permitting low power transmission of a pulse at slow speeds: this could be useful for matching the speed of silicon with that of events in the real world.

At the level of the interaction of the ion channels on the membrane itself, the first major model demonstrating both the sodium and potassium conductances was implemented in [MD91]: this model demonstrated that spike generation in a neuron-like way could be emulated electronically using subthreshold aVLSI circuitry. This work has been extended and improved in [RDM98][RD00]. A different approach which uses novel semiconductor fabrication techniques is adopted by [NLBA04]: this approach uses properties of he semiconductors more directly, rather than using circuitry. In [SNT+08] they extend this approach and use a tunnel diode to regenerate electronic signals. The effect is like that of neural axonic conduction, but at a rather higher speed. Logically, both of these approaches would allow a hardware implementation of a multi-compartment neuron to be built up from these patches of membrane.

3.2.3 Modelling neural circuitry

As well as modelling neurons or parts of neurons, there is also interest in neuromorphic implementations of neural circuitry. There are two different motives for this: firstly engineers would like to be able to use circuits of neurons to achieve particular processing functions, and secondly, modellers would like to be able to model particular arrangements of neurons that they find from neuroanatomy. An example of the first of these is a model of a *winner takes all* (WTA) network. These essentially choose one (or possibly more than one) of a set of interconnected neurons with different inputs, and select the one with the greatest activation (the winner). This is a useful capability, and models of these have been around for some time [LRMM89][PAS97]. Improving these by making them smaller, or by adjusting the timescale of the inputs is still an area of research [MG07].

Developing models of small volumes of cortex in silicon is an avowed aim of a number of groups (for example the *Brains in silicon* group at Stanford University, as well as Rodney Douglas at the Zurich Institute for Neuroinformatics (personal communication), and the subject of at least one PhD thesis [Mer06]. As matters stand, however, this are is still dominated by huge software simulations [DLJ+08].

4 Implementation technologies for neuromorphic systems

Emulation of neural systems can be implemented in many different ways, ranging from software on standard digital computers, through to application specific integrated circuits (ASICs) implemented either in digital or analogue technology, possibly with additional technologies piggybacked on to implement elements that are difficult in CMOS. In general, to be called neuromorphic, there has to be some element of direct hardware implementation: otherwise one simply has a software

model. (There is some discussion about whether an implementation based on efficient software on a multiple core processor might yet be called neuromorphic: we will not enter this discussion!) The other factor that is required in neuromorphic implementations is real-time operation, as discussed in section 3.1.

In Mead's original work, the implementation technology was sub-threshold analogue VLSI: however, many systems which are taken to be neuromorphic use digital technologies. There remain many choices still about the form of the implementation. One possibility is to use field programmable devices, whether field-programmable digital arrays, or field-programmable analogue arrays. These have the major advantages of being much easier and faster to configure, and avoiding the long delays inherent in the fabrication of ASICs. In addition, they may be reprogrammed unlike ASICs. However, the circuit density achievable is much less than that of ASICs, and the power consumption is much higher: this can be a particular problem for small or mobile devices. Further, where one is integrating sensors as well, field-programmable devices are not an option: one needs to choose an ASIC based implementation. It is important to note that the designers of neuromorphic ASICs generally have to use the technologies and foundries developed for digital chips, since this market is far bigger, and it is not currently possible to develop manufacturing technologies for this small niche market.

Even ASICs do not solve all the problems. In particular issues relating to interconnectivity and adaptiveness present problems. The planarity of CMOS devices places severe limits on interconnection possibilities. Further, most CMOS technologies are intended to produce components which are stable and always behave the same way, rather than components whose characteristics can evolve. There are also issues of ensuring that the speeds within the system (for example integration times) actually fit with events in the world whose timescales are often of the order of hundreds of milliseconds, rather than the nano- to micro-second range more usually encountered in the ASIC domain. Recently, a number of groups have developed novel technologies to address the issues of interconnectivity and adaptivity, and these are discussed in section 5.

4.1 Signal coding

An important issue for implementation technologies is the nature of the signal representation both on-chip, and for transfer between chips. Essentially there are three overall possibilities: analogue, digital, and pulse-based signals. Analogue representations essentially imply analogue implementation, which has the advantages of implicit real-time operation and perhaps lower power. Further, signals in the world are normally analogue, simplifying interfacing. However, transistor variation across chip can make reliable circuit design very difficult, and the standard digital fabrication processes may not be as reliable for analogue as for digital designs. Digital designs, on the other hand, offer the usual digital advantages: noise immunity, ease of manufacture, high density, and effective use of the standard manufacturing

processes. However real-time operation needs to be ensured in the design, and the digital signals require to be interfaced to the analogue world: in addition, some operations which can be small and simple in analogue implementation, such as multiplication, can require relatively large amounts of chip real estate to implement digitally.

One compromise is to use pulse or spike based representations. It is well known that spikes are used for communication between most neurons in animal brains, although the precise nature of the representations used in these is still under debate. Spikes provide a mechanism for asynchronous communication between electronic circuits: the actual spike is binary (it is an all or nothing event), but the precise time of the spike is essentially analogue. Between the elements on a single chip, spikes may be sent directly, but the planar nature of chips means that being able to send spikes between any pair of circuits on chip may not always be possible. One solution to this is to use internal routing as in [MHH+08]. Equally clearly, the small number of connections coming off a chip means that simple spike coding cannot in general, be used for communication between arbitrary circuits on different chips. One technique for overcoming this limitation is to use a bus. For such a bus to work effectively, it needs to be standardised. The address-event bus[Boa00] (see http://www.pcmp.caltech.edu/aer/) is the current standard in this area. Of course, such a bus implies a maximal rate at which spiking events may be transferred, and a maximal precision to the timing of transmitted events: however, digital transmission busses do have a very high (and known) bandwidth so that one can calculate whether this is likely to be a problem at the design stage.

5 Adaptivity and interconnectivity for neuromorphic systems

Two characteristics of neural and neuromorphic systems that set them aside from traditional computer systems are adaptivity and a high degree of interconnection between the elements. The adaptive elements are normally at the interconnections between the neural circuit elements, since they are generally modelled on synapses whose plasticity is an important element in learning in real neural systems. If there are n neurons there are normally $O(n^2)$ interconnecting synapses, so that it is important that the circuitry modelling the synapse is small and low-power. (There are other possibilities: if synapse implementation circuitry is sufficiently fast, it may be shared in time between a number of emulated synapses (partially implemented in the terminology of [HN90]). This does, however, make the circuitry more complex, but may be appropriate in digital implementations where synapses can include large multipliers.)

Synapses need to provide some specific capabilities: they need to be able to transmit a signal from the pre-synapse neuron (however this signal is coded), whilst modulating this signal, providing some particular alteration (in voltage or current, again depending on the nature of the implementation) at the post-synaptic neuron. The size of this alteration will need to be adjustable, if the synapse is adaptive. The

earliest adaptive elements for the first generation of hardware neural networks included novel devices such as the memistor [Wid60], an electrochemical adaptive resistive element (named for being a memory resistor). By modern standards, these are large and slow. However, the production of reliable adaptive devices which are small enough to be able to be deployed in large numbers remains difficult. One possibility is to use a memory word, implemented digitally. The value coded in the word defines the alteration at the post-synaptic neuron. Though practical for truly digital implementations, this requires both digital to analog conversion and a technique for altering the value in analogue implementations. Another possibility is to use floating gate devices: this is the same technology used to create flash memory devices, and has been shown to be highly reliable for creating digital memories. It has been used for synapses [DPMM96][HFD02], using electron injection and electron injection to increase and decrease the voltage on the floating gate. However, these have not been used in large numbers. In addition to requiring a technique for long-term storage of synaptic effectiveness, synapses also change their effects on a shorter time scale (called synaptic facilitation and depression). A simple short-term adaptive synapse permitting both facilitation and depression is described in [CBD$^+$03], and a more sophisticated implementation including NMDA voltage-gated channels is described in [BI07].

Recently, there have been some new technologies which have shown promise as new implementation techniques for synapses. These are based on nanowire crossbars whose junctions are built from metal oxides with hysterectic resistances: tin oxides and zirconium oxides are often used[SPS07]. These devices are actually memristors [Bus08][SSSW08], "a class of passive two-terminal circuit elements that maintain a functional relationship between the time integrals of current and voltage" (Wikipedia), which makes them able to be resistors with memory. These are built from a technology known as CMOL (CMOS and molecular electronics), and are two terminal devices which can be made very small. If these can be implemented appropriately, they offer for the first time a technology which could provide the appropriate number of synapses which cold be adaptive. However, this is still very much a matter of research.

6 Looking forward: where are neuromorphic systems headed?

There remains considerable interest in auditory and visual neuromorphic systems as technologies for eventually producing synthetic sensing systems with the same types of capabilities as biological auditory and visual systems. The rapid response of the neuromorphic camera in[LPD08] without the use of large-scale frame technology represent a real step forward. However, neuromorphic auditory systems have yet to prove themselves capable: this may be because they have yet to properly incorporate the brainstem processing in more than very simple way on to the filtering technology.

Another area of progress is likely to be in the integration of different types of sensors on to CMOS systems. Light sensors have been around for a long time, and polymer based sensors are in use in olfactory neuromorphic systems. In addition, microelectromechanical systems (MEMS) microphones have now been developed (e.g. the Infineon SMM310, or the Akustika AKU1126), and these seem to be good candidates for integration directly on to CMOS substrates. Different types of sensors for olfactory and other senses may be based on ion sensitive FETs (ISFETs)[HAC04] and chemical sensitive FETs (chemFETs) [HN02]: these new technologies are based directly on FETs and so are clearly integratabtle on to CMOS systems, although dealing the the nearness of liquids presents some novel problems for such electrical equipment. There is also work ongoing in the development of proprioceptary sensors [WZPF06]. Being able to directly incorporate the transduction on to the CMOS system both reduces complexity and component count, and permits processing to be applied directly to the signal allowing the chip to produce usable outputs directly.

6.1 When will neuromorphic systems come out of the lab?

It is now almost 20 years since Mead's book [MC80] was published. At this point neuromorphic systems have had some applications, in robotics, and in some sensors, as well as in a rather interesting system for training neurophysiologists in how cells in the visual system respond (see http://www.ini.ethz.ch/ tobi/friend/chip/index.php). They have also been applied to toys: some of the Wowwee toys use neuromprphic hardware (http://www.wowwee.com/. They have yet to really catch on even in the autonomous robotics area. Why is this? So far they have been very much a low budget interest area for researchers, built using technologies developed for other purposes. However, there are signs that this is changing. Recently, DARPA announced they SyNAPSE initiative (http://www.darpa.mil/baa/BAA08-28.html), and this may lead to rather larger sums being available for development.

Part of the problem has been the simple capabilities of standard processor technologies: it has become quite practical to place full-scale computer systems on even quite small autonomous robots. However, it remains very difficult to perform the sorts of sophisticated processing that truly autonomous robots require to be useful in the relatively unpredictable real world (as opposed to on top of a table in an experimenter's laboratory). For real applications the capabilities of truly parallel neuromorphic systems (for example in dealing with varying light levels and real acoustic situations) may become more important.

The other likely application area is in interfacing computer-based systems to both users and the environment. Currently, most computers still use only the keyboard and mouse for input, and a screen and loudspeaker for output. There are improvements, for example touch screen and multi-touch screens which are being introduced. Yet the hardware of the user interface still conforms to the "make the user adjust to the machine" paradigm that disappeared from the software for the user interface many years ago. Further, if one is building systems that interact directly

with their environment (without human mediation), then the system must sense its environment, and make sense of the sensory data directly. This implies both richness and complexity of the sensory interface (as well as real-time operation) and sophisticated processing that can cope with variation in this sensory data, and extract the important (invariant) information that is required for behaviour in the environment. Hardware solutions, as well as integrated sensors appear appropriate for his area, and it may be that this is where neuromorphic systems will finally make their mark.

Acknowledgements This work started off as a talk at BICS 2008, and was revised through being given as a seminar at Stirling University and the University of Surrey, and in the light of the referees' comments.

References

[AHW+07] P. Argyrakis, A. Hamilton, B. Webb, Y. Zhang, T. Gonos, and R. Cheung. Fabrication and characterization of a wind sensor for integration with neuron circuit. *Microelectronic Engineering*, 84(1749-1753), 2007.

[AK98] J. Austin and J. Kennedy. PRESENCE, a hardware implementation of bunary neural networks. In M. Boden L. Niklasson and T. Ziemke, editors, *ICANN98: Proceedings of the 8'th international conference on artificial neural networks*, volume 1, pages 469–474, 1998.

[AMSB95] A.G. Andreou, R.C. Meitzler, K. Strohben, and K.A. Boahen. Analog VLSI neuromorphic image acquisition and pre-processing systems. *Neural Networks*, 8(7-8):1323–1347, 1995.

[AR88] J.A. Anderson and E. Rosenfeld, editors. *Neurocomputing: Foundations of Research*. MIT Press, Cambridge, 1988.

[AS79] I Aleksander and T.J. Stonham. Guide to pattern recognition using random access memories. *IEE Proceedings: Computers and Digital Techniques*, 40:2–29, 1979.

[ATB84] I. Aleksander, W.V. Thomas, and P.A. Bowden. WISARD a radical step forward in image recognition. *Sensor Review*, 4(3):120–124, 1984.

[BI07] C. Bartolozzi and G. Indiveri. Synaptic dynamics in analog VLSI. *Neural Computation*, 19:2581–2603, 2007.

[Boa00] K.A. Boahen. Point-to-point connectivity between neuromorphic chips using address-events. *IEEE Transactions on Curcuits and Systems II*, 47(5):416–434, 2000.

[Bro79] W.H. Brockman. A simple electronic neuron model incorporating both active and passive responses. *IEEE Transactions on Biomedical Engineering*, BME-26:635–639, 1979.

[Bus08] S. Bush. HP nano device implements memristor. *Electronics Weekly*, May 2008.

[CBD+03] E. Chicca, D. Badoni, V. Dante, M. D'Andreagiovanni, G. Salina, L. Carota, S. Fusi, and P. Del Giudice. A vlsi recurrent network of integrate-and-fire neurons connected by plastic synapses with long term memory. *IEEE Transactions on Neural Networks*, 14(5):1409–1416, 2003.

[Cor90] Intel Corporation. 80170NN electrically trainable analog neural network. *Datasheet*, 1990.

[DLJ+08] M. Djurfeldt, M. Lundqvist, C. Johansson, M. Rehn, O .Ekeberg, and A. Lansner. Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer. *IBM Journal of Research and Development*, 52(1/2), 2008.

[DM96] T. Delbruck and C.A. Mead. Analog VLSI transduction. Technical Report CNS Memo 30, California Institute of Technology Computation and neural Systems Program, Pasadena California, USA, April 1996.

[DPMM96] C. Diorio, P.Hasler, B.A. Minch, and C.A. Mead. A single-transistor silicon synapse. *IEEE Transactions on Electron Devices*, 43(11):1982–1980, 1996.

[Eli93] J.G. Elias. Artificial dendritic trees. *Neural Computation*, 5(4):648–664, 1993.

[FA95] P. Furth and A.G. Andreou. A design framework for low power analog filter banks. *IEEE Transactions on Circuits and Systems*, 42(11):966–971, November 1995.

[Fit66] R. Fitzhugh. An electronic model of the nerve membrane for demonstration purposes. *Journal of applied physiology*, 21:305–308, 1966.

[FvSV97] E. Fragniére, A. van Schaik, and E.A. Vittoz. Design of an analogue VLSI model of an active cochlea. *Analog Integrated Circuits and Signal Processing*, 12:19–35, 1997.

[GHS02] M. Glover, A. Hamilton, and L.S. Smith. Analogue VLSI leaky integrated-and-fire neurons and their use in a sound analysis system. *Analog Integrated Circuits and Signal Processing*, 30(2):91–100, 2002.

[GK02] W. Gerstner and W. Kistler. *Spiking Neural Models*. Cambridge, 2002.

[HAC04] P.A. Hammond, D. Ali, and D.R.S. Cumming. Design of a single-chip pH sensor using a conventional 0.6-μm CMOS process. *IEEE Sensors Journal*, 4(6):706–712, 2004.

[Ham07] S.R. Hameroff. The brain is both a neurocomputer and a quantum computer. *Cognitive Science*, 31:1033–1045, 2007.

[HFD02] D. Hsu, M. Figueroa, and C. Diorio. Competitive learning with floating-gate circuits. *IEEE Transactions on Neural Networks*, 13:732–744, 2002.

[HH52] A.L. Hodgkin and A.F. Huxley. Currents carried by sodium and potassium ions through the membrane of the giant squid axon of loligo. *Journal of Physiology*, 116:449–472, 1952.

[Hir91] N. Hirokawa. Molecular architecture and dynamics of the neuronal cytoskeleton. In R.D. Burgoyne, editor, *The neuronal cytoskeleton*. Wiley-Liss, 1991.

[HJvST08] T.J. Hamilton, C. Jin, A. van Schaik, and J. Tapson. An active 2-d silicon cochlea. *IEEE Transactions on biomedical circuits and systems*, 2(1):30–43, 2008.

[HK98] R.R. Harrison and C. Koch. An analog VLSI model of the fly elementary motion detector. In M.I.Jordan, J. Kearns, and S.A. Solla, editors, *Adcanves in Neural Information Processing Systems 10*, pages 880–886. MIT Press, 1998.

[HN90] R. Hecht-Nielsen. *Neurocomputing*. Addison-Wesley, 1990.

[HN02] A.M. Hodge and R.W. Newcomb. Evaluation of the VLSI adaptation of the chemfet, a biosensor forfluid analysis. In *IEEE International Symposium on Circuits and Systems, ISCAS 2002.*, volume 2, pages II580–II583, 2002.

[Ind98] G. Indiveri. Analogue VLSI model of locust DCMD neuron response for computation of object aproach. In L.S. Smith and A. Hamilton, editors, *Neuromorphic Systems: engineering silicon from neurobiology*, pages 47–60, 1998.

[JH69] R.H. Johnson and G.R. Hanna. Membrane model:a single transistoranalog of excitable membrane. *Journal of Theoretical Biology*, 22:401–411, 1969.

[KHT+05] T.J. Koickal, A. Hamilton, S.L. Tan, J.A. Covington, J.W. Gardner, and T.C. Pearce. Analog VLSI circuit implementation. of an adaptive neuromorphic olfaction chip. In *IEEE International Symposium on Circuits and Systems (ISCAS) IEEE International Symposium on Circuits and Systems (ISCAS),Kobe, JAPAN*, volume 54, pages 60–73, 2005.

[KS59] W.J. Karplus and W.W. Soroka. *Analog Methods: Computation and Simulation*. McGraw-Hill, 1959.

[LAJ93a] W. Liu, A.G. Andreou, and M.H. Goldstein Jr. Analog cochlear model for multiresolution speech analysis. In *Advances in Neural Information Processing Sytems 5*, pages 666–673, 1993.

[LAJ93b] W. Liu, A.G. Andreou, and M.H. Goldstein Jr. Voiced speech representation by an analog silicon model of the auditory periphery. *IEEE Trans. Neural Networks*, 3(3):477–487, 1993.

[Lap07] L. Lapique. Sur l'excitation electrique des nerfs. *Journal of Physiology, Paris*, pages 620–635, 1907.

[Lew68] E. R. Lewis. An electronic model of the neuroelectric point process. *Kybernetik*, 5:30–46, 1968.

[LKI+02] S-C. Liu, J. Kramer, G. Indiveri, T Delbruck, and R. Douglas. *Analog VLSI: Circuits and Principles*. MIT Press, 2002.

[LM88] R.F. Lyon and C. Mead. An analog electronic cochlea. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 36(7):1119–1134, 1988.

[LM89a] J. Lazzaro and C. Mead. Circuit models of sensory transduction in the cochlea. In *Analog VLSI implementations of neural networks*, pages 85–101. Kluwer, 1989.

[LM89b] J. Lazzaro and C. Mead. Silicon modeling of pitch perception. *Proceedings of the National Academy of Sciences of the United States*, 86(23):9597–9601, 1989.

[LM89c] J Lazzaro and C.A. Mead. A silicon model of auditory localization. *Neural Computation*, 1(1):47–57, 1989.

[LPD08] P. Lichtsteiner, C. Posch, and T. Delbruck. A 128 x 128 120db 15 μs latency asynchronous temporal contrast vision detector. *IEEE Journal of Solid-state Circuits*, 43(2):566–576, 2008.

[LRMM89] J. Lazzaro, S. Ryckebush, M.A. Mahowald, and C.A. Mead. Winner-take-all networks of O(n) complexity. In D. Touretzky, editor, *Advances in Neural Inforpation Processing Systems*, pages 703–711, 1989.

[LW97] J. Lazzaro and J. Wawrzynek. Speech recognition experiments with silicon auditory models. *Analog Integrated Circuits and Signal Processing*, 13:37–51, 1997.

[MC80] C. Mead and L.A. Conway. *Introduction to VLSI Systems*. Addison-Wesley, 1980.

[MD91] M. Mahowald and R. Douglas. A silicon neuron. *Nature*, 354(6354):515–518, 1991.

[Mea89] C. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.

[Mer06] P. Merolla. *A Silicon Model of the Primary Visual Cortex: Representing Features Through Stochastic Variations*. PhD thesis, Department of Bioengineering, University of Philadelphia, PA, USA, 2006.

[MG07] C.M Markan and P. Gupta. Neuromorphic building blocks for adaptable cortical feature maps. In *IFIP International conference on VLSI*, pages 7–12, 2007.

[MHDM95] B.A. Minch, P. Hasler, C. Diorio, and C. Mead. A silicon axon. In G. Tesauro, D.S. Touretzky, and T.K. Leen, editors, *Advances in Neural Information Processing Systems 7*. MIT Press, 1995.

[MHH+08] L. McDaid, J. Harkin, S. Hall, T. Dowrick, Y. Chen, and J. Marsland. EMBRACE: emulating biologically-inspired architectures on hardware. In *NN'08: Proceedings of the 9th WSEAS International Conference on Neural Networks*, pages 167–172, Stevens Point, Wisconsin, USA, 2008. World Scientific and Engineering Academy and Society (WSEAS).

[MM88] C. Mead and M.A. Mahowald. A silicon model of early visual processing. *Neural networks*, 1(1):91–97, 1988.

[MP43] W.S. McCulloch and W. Pitts. A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 1943. Reprinted in [AR88].

[NE96] D.P.M. Northmore and J.G. Elias. Spike train processing by a silicon neuromorph: The role of sublinear summation in dendrites. *Neural Computation*, 8(6):1245–1265, 1996.

[NLBA04] A. Nogaret, N.J. Lambert, S.J. Bending, and J. Austin. Artificial ion channels and spike computation in modulation-doped semiconductors. *Europhysics Letters*, 68(6):874–880, 2004.

[PAS97] P.O. Pouliquen, A.G. Andreou, and K. Strohben. Winner-takes-all associative memory. *Analog Integrated Circuits and Signal Processing*, 13(1-2):211–222, 1997.

[Pea97] T.C. Pearce. Computational parallels between the biological olfactory pathway and its analogue 'the electronic nose'. 2. sensor-based machine olfaction. *Biosystems*, 41(2):69–90, 1997.

[PNG+06] M. Pearson, M. Nibouche, I. Gilhespy, K. Gurney, C. Melhuish, B. Mitchison, and A.G. Pipe. A hardware based implementation of a tactile sensory system for neuromorphic signal processing applications. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, May 2006.

[RD00] C. Rasche and R.J. Douglas. An improved silicon neuron. *Analog Integrated Circuits and Signal Processing*, 23(3):227–236, 2000.

[RD01] C. Rasche and R.J. Douglas. Forward- and backpropagation in a silicon dendrite. *IEEE Transactions on Neural Networks*, 12(2), 2001.

[RDM98] C. Rasche, R.J. Douglas, and M. Mahowald. Characterization of a silicon pyramidal neuron. In L.S. Smith and A. Hamilton, editors, *Neuromorphic Systems: Engineering Silicon from Neurobiology*. World Scientific, 1998.

[Roy72] G. Roy. A simple electronic analog of the squid axonmembrane: the neuroFET. *IEEE Transactions on Biomedical Engineering*, BME-18:60–63, 1972.

[Roy06] D. Roy. Design and developmental metrics of a 'skin-like' multi-input quasi-compliant robotic gripper sensor using tactile matrix. *Journal of Intelligent and Robotic Systems*, 46(4):305–337, 2006.

[RUV68] R.G. Runge, M. Uemura, and S.S. Viglione. Electronic synthesis of the avian retina. *IEEE Transactions on Biomedical Engineering*, BME-15:138–151, 1968.

[SFV96] Andr Van Schaik, Eric Fragniere, and Eric Vittoz. An analogue electronic model of ventral cochlear nucleus neurons. In *5th International Conference on Microelectronics for Neural Networks and Fuzzy Systems (MicroNeuro '96)*, 1996.

[SG92] H.V. Shurmer and J.W Gardner. Odor discrimination with an electronic nose. *Sensors and Actuators B-Chemical*, 8(11):1–11, 1992.

[SNT+08] A. Samardak, A. Nogaret, S. Taylor, J. Austin, I. Farrer, and D.A. Ritchie. An analogue sum and threshold neuron based on the quantum tunnelling amplification of neural pulses. *New Jpurnal of Physics*, 10, 2008.

[SPS07] G.S. Snider, P.J.Kuekes, and D.R. Stewart. Nanoscale lacth-array processing engines. US Patent number 7227993, June 2007.

[SSSW08] D.B. Strukov, G.S. Snider, D.R. Stewart, and S.R. Williams. The missing memristor found. *Nature*, 453:80–83, 2008.

[SV97] Andre Van Schaik and Eric Vittoz. A silicon model of amplitude modulation detection in the auditory brainstem. In *Advances in NIPS 9*, pages 741–747. MIT Press, 1997.

[VAV+06] G. Vasarhelyi, M. Adam, E. Vazsonyi, A. Kis, I. Barsony, and C. Ducso. Characterization of an integrable single-crystalline 3-d tactile sensor. *IEEE Sensors journal*, 6(4):928–934, 2006.

[WD08] J.H.B. Wijekoon and P. Dudek. Compact silicon neuron circuit with spiking and bursting behaviour. *Neural Networks*, 21:524–534, 2008.

[Wid60] B. Widrow. An adaptive "adaline" neuron using chemical "memistors". Technical Report 1553-2, Stanford University, Solid-state electronics laboratory, Stanford electronics laboratories, Stansted University, Stansted, California, October 1960.

[WMT96] S. Wolpert and E. Micheli-Tzanakou. A neuromime in VLSI. *IEEE Transactions on Neural Networks*, 7(2):300–306, 1996.

[WZPF06] D.H.B. Wicaksono, L-J. Zhang, G. Pandraud, and P.J. French. Fly's proprioception-inspired micromachined strain-sensing structure: idea, design, modeling and simulation, and comparison with experimental results. In *J. Physics: Conference Series: International MEMS Conference 2006*, volume 34, pages 336–341, 2006.

[YMW+06] Z. Yang, A.F. Murray, F. Woergoetter, K.L. Cameron, and V. Boonsobhak. A neuromorphic depth-from-motion vision model with stdp adaptation. *IEEE Transactions on Neural Networks*, 17(2):482–495, 2006.