Analogue VLSI leaky integrate-and-fire neurons and their use in a sound analysis system.

Mark Glover (mglover@dera.gov.uk) and Alister Hamilton (alister.hamilton@ee.ed.ac.uk)
Department of Electronics and Electrical Engineering, University of Edinburgh, Edinburgh EH9 3JL, Scotland, UK

Leslie S. Smith (l.s.smith@cs.stir.ac.uk)
Department of Computing Science and Maths, University of Stirling, Stirling FK9 4LA, Scotland, UK

Abstract. Integrate-and-fire neurons are simple model neurons which can handle continuously time-varying signals. We have applied them to problems in real-time analysis of sounds. Two different chips have been built: the first had a fixed network architecture with all synaptic weights identical, and the second is reconfigurable with individually programmable weights. We present results characterising the latter chip, and results from processing real data from the earlier chip. We note that the second chip provides a more general integrate-and-fire neuron implementation.

1. Aims

We describe work aimed at developing a hardware leaky integrate-and-fire neuron suitable for use in the sound analysis domain with new results from a recently fabricated analogue device. Leaky integrate-and-fire neurons (reviewed below) provide a neural model of complexity and sophistication between that of McCulloch-Pitts neurons and Perceptrons, and biologically realistic multiple compartmental models. The sound analysis domain provides interesting (and difficult) problems which require the capability of processing continuously varying signals over time and demand real-time results. It is thus an appropriate domain for experimenting with neurons implemented in hardware.

2. Background

We present a brief introduction firstly to the leaky integrate and fire neuron, and secondly to the problems to which we have applied them.

2.1. Integrate-and-fire neurons

Integrate -and-fire neurons have been proposed as model neurons for a long time starting with (Lapique, 1905). They, and a related model, the spike response model, have been growing in popularity in recent years (Gerstner, 1995; Ho et al., 1998). There are a number of reasons for this: they can be used for processing time-varying signals (Smith, 1996; Smith, 1998; van Schaik et al., 1997; Zahn et al., 1998), they are the simplest neural models that can support synchronisation as part of a representation (Glunder et al., 1993; Ernst et al., 1994; Nishwitz et al., 1995), they have been shown to be able to be used as powerful computing systems (Maass, 1996; Maass, 1997). They are also quite straightforward to implement either in software or hardware (Jahnke et al., 1998; Wolpert et al., 1996).

Between spikes, the activity (voltage) of a leaky integrate-and-fire neuron is governed by

$$\frac{dV(t)}{dt} = -\frac{V(t)}{RC} + \frac{I(t)}{C} \tag{1}$$

where the leakiness is measured by the dissipation ($= \frac{1}{RC}$), and $I(t)$ is the external input current. The neuron fires when the activity reaches some predefined threshold. After firing, the neuron may have a refractory period, during which it cannot fire. In some models, this is followed by a relative refractory period, during which the threshold is heightened, but gradually reduces to the original threshold.

2.2. Applying leaky integrate-and-fire neurons

Although leaky integrate-and-fire neurons can be applied anywhere that simpler models are applied, they are particularly useful in processing signals which vary continuously over time. The analysis of sound presents a good source of such signals. However, because these neurons operate continuously on these time-varying signals, their time constants must match those of the signal itself. It is impractical to make large capacitances on an analogue device and difficult to make large precision resistances. Making the RC value programmable adds a further degree of complexity. In addition, the time constant associated with the refractory period (and relative refractory period) is of the same order of magnitude, and should also be programmable.

The time constants required in sound analysis depend on the precise application. They vary from the short (10-100 microseconds) required for binaural direction finding, to the medium (100 microseconds- 2 milliseconds) for monaural source grouping, to the relatively long (10-100 milliseconds) required for sound segmentation, to the very long

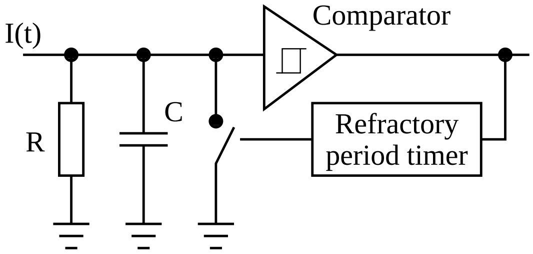

Figure 1. Integrate and Fire Neuron Circuit Model. An external current is integrated by the leaky integrator. Once the integrated voltage reaches the threshold of the comparator, the output of the comparator goes high and starts the refractory period timer which holds the integrated voltage at a reset potential until the end of the refractory period. During this reset period as the integrated voltage approaches the lower switching threshold of the comparator, the comparator output voltage goes low. Thus a fixed length voltage pulse or spike is produced at the output of the neuron.

(upwards of 100 ms) required for rhythm discrimination. The first three of these are called fine time structure, periodicity and envelope levels in (Rosen, 1992), and called fine, medium, and coarse short-term time structure in (Smith, 1998). The fourth is described in (Todd et al., 1999). In this paper we are considering only relatively long time constants for sound segmentation.

3. The leaky integrator neural hardware

The initial motivation for the work reported here was to provide a hardware implementation for a (software) monaural sound segmentation algorithm (Smith, 1996). Sound segments are in the 20-200ms range, and this defines the timescale at which the integrate-and-fire unit must operate. Such timescales are orders of magnitude greater than the natural time constants associated with analogue VLSI designs.

The basic circuit model of the integrate and fire neuron used in the implementation is illustrated in figure 1.

The leaky integrator formed by the RC network integrates the input current, $I(t)$, until the upper threshold of the comparator is reached. The comparator output goes high starting the refractory period timer and discharging C via the switch. Once the lower threshold of the comparator output is reached, the comparator output goes low completing the fixed length output spike from the neuron. The switch remains closed during the refractory period preventing further integration during this time. The input, $I(t)$, is formed by an external input to the network (in our application this signal is formed by the preprocessing stages of the sound segmentation system and represents compressed

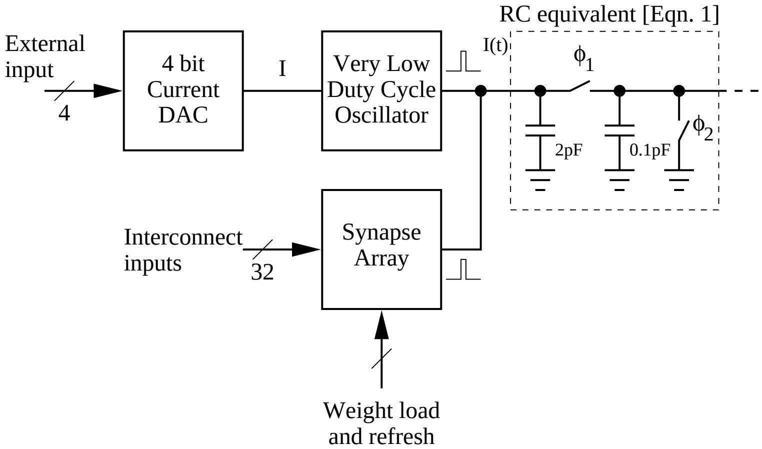

Figure 2. Integrate and Fire Neuron Input Stages as implemented on-chip. The input current to each neuron is supplied externally as a 4 bit value. It is converted to a current by a DAC. The very low duty cycle oscillator converts this current into a pulse frequency coded current (see text). Spikes from other neurons are converted into current pulses of fixed length but varying sign and magnitude by the synapse array. The switched capacitor arrangement forms the RC network of Figure 1 and Equation 1 where ϕ_1 and ϕ_2 are non-overlapping clocks derived from an oscillator external to the chip. Note that the comparator and refractory period timer of figure 1 are not shown.

sound onset/offset signals) plus a contribution from synapses connected to other integrate and fire neurons. In earlier work (Glover et al., 1998; Smith et al., 1998; Glover et al., 1998) these synapses had equal excitatory strength (globally variable by an external current source). In the work reported here for the first time, the 32 synapses connected to each neuron are independently programmable with excitatory or inhibitory connection strengths.

The main design problem is to implement large time constants. The RC time constant was designed to be programmable in the range 1-20ms, and the refractory period timer in the range 10ms to greater than 100ms. Both programmable time constants have been implemented with a switched capacitor structure (Glover et al., 1998) which is programmable in the required ranges. The external input is applied as a 4-bit digital signal and a current digital-to-analogue converter (DAC) is used to control a very low duty cycle oscillator which pulses current into the integrating node. This approach has been taken to keep the area of the integrating capacitor down, while allowing the (pulsed) current to be large enough to enable good matching. Pulsing $I(t)$ with a very low duty cycle (0% - 0.21% in steps of 0.014%) is practical due to the relatively slow nature of the auditory input signals. The digital interface

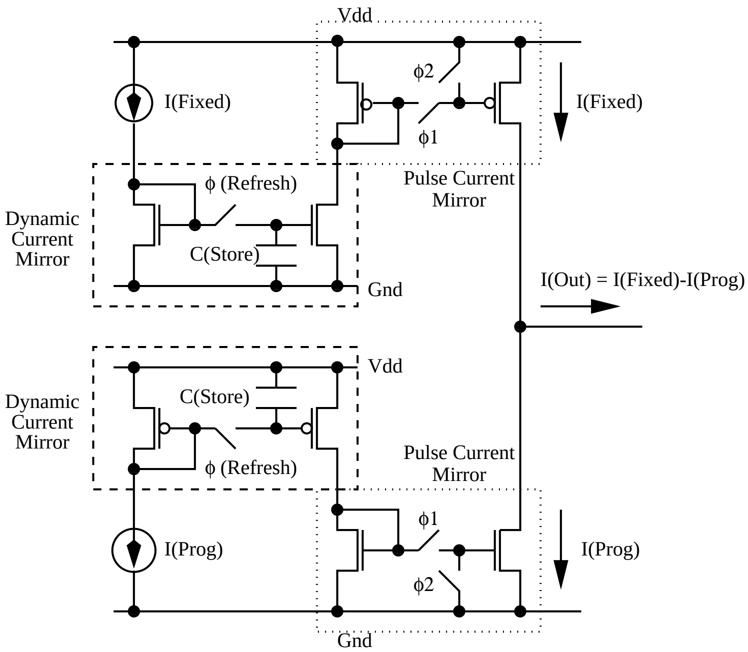

Figure 3. Circuit diagram of the fully programmable synapse used in the synapse array of Figure 2. The circuit uses two dynamic current mirrors at each synapse to store the synaptic weight. The upper dynamic current mirror stores a fixed current of $I_{max}/2$. The lower current mirror is programmable and refreshed via a single on-chip current DAC. This current is programmable in the range $0 - I_{max}$. When a pulse or spike arrives from the pre-synaptic neuron, the switches controlled by ϕ_1 are closed and those controlled by ϕ_2 are opened and the pulsed current mirrors output a current spike to the post-synaptic neuron. This current is in the range from $-I_{max}/2$ to $I_{max}/2$. Note that this pulsed current mirror technique results in minimal switching noise injection into the integrating node.

allows data from the software model of the preprocessing stages of the sound segmentation system to be fed directly into the chip.

The programmable synapse design (Hamilton, 1994) is illustrated in figure 3 and implements the synapse array of figure 2. One on-chip 5 bit weight DAC is used to set I(Prog) and I(Fixed). Weight values are stored at each synapse (as the voltages on C(store)) and are periodically refreshed using the weight DAC. Each synapse contains two dynamic current mirrors which are periodically refreshed. One current mirror is programmed with a fixed current ($I_{max}/2$) and the other with a programmable current ($0 - I_{max}$) from a single current DAC. Each dynamic current mirror drives a pulsed current mirror output stage which is pulsed when a spike arrives from the pre-synaptic neuron.

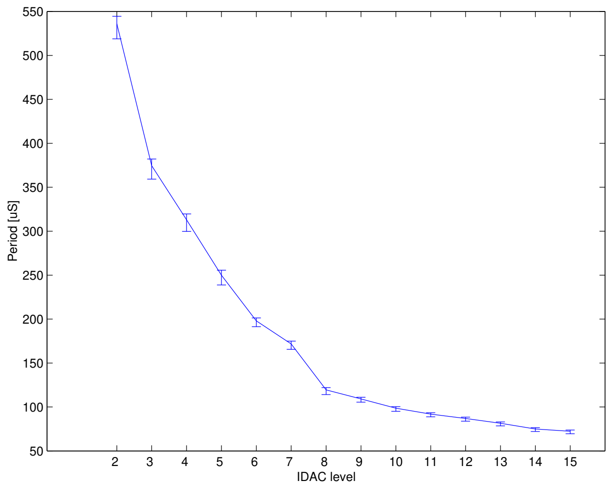

Figure 4. Very Low Duty Cycle Oscillator Performance. Inter-pulse interval versus current DAC setting for one neuron. Error bars are standard deviation of inter-pulse interval measured on a single neuron.

4. Results

We present results (i) from static measurements of the new chip, showing the range of achievable input strengths, dissipations, refractory periods, and synaptic strengths, and (ii) from testing the original chip with real data derived from sound inputs.

4.1. Static measurements

The performance of the low duty cycle oscillator is illustrated in figure 4. The input is a four bit digital word converted into a current to control the output duty cycle of the oscillator. This shows the expected trend. Testing on other neurons shows good matching except at low current values. The inter-pulse interval shows a jitter of approximately 5% throughout.

The RC time constant has been measured by observing the charging time of the integrating node for different input levels (expressed in terms of the current DAC input) and different switched capacitor clock frequencies (ϕ_1 and ϕ_2 in figure 2). The results of this test are shown in figure 5 where the time taken to reach threshold by neuron 0 on chips A and B is measured, with switched capacitor clock frequencies of 100Hz, 1kHz, 10kHz and 20kHz. This graph shows that there is little difference between the performance of chips A and B due to process variation.

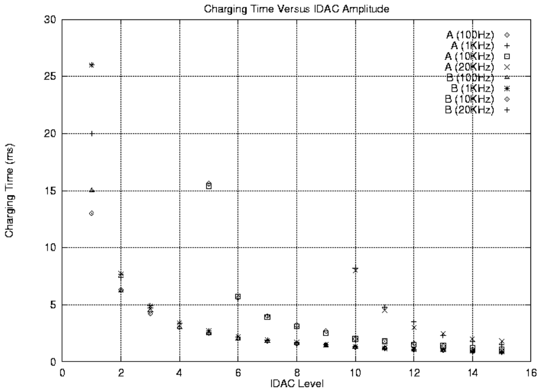

Figure 5. Switched capacitor RC network charging times for varying Very Low Duty Cycle Oscillator inputs and time constants

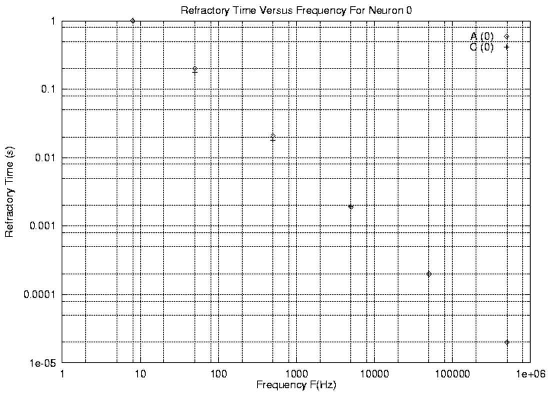

Figure 6. Refractory period for different switched capacitor clock frequencies

The refractory period which is generated by a switched capacitor circuit similar to that used for the RC time constant ranges over 5 orders of magnitude from 10 microseconds to 1 second (figure 6).

The performance of the synapse array was tested by applying a constant external input to the neuron (current DAC level 8), switched capacitor frequencies of 5kHz and 12kHz for the RC time constant and refractory period timer respectively, and a simulated pre-synaptic pulse frequency of 4.5kHz on synapse 1. Figure 7 illustrates the effect of the varying synapse interconnect strength on the post-synaptic neuron firing frequency. This figure illustrates the correct operation of

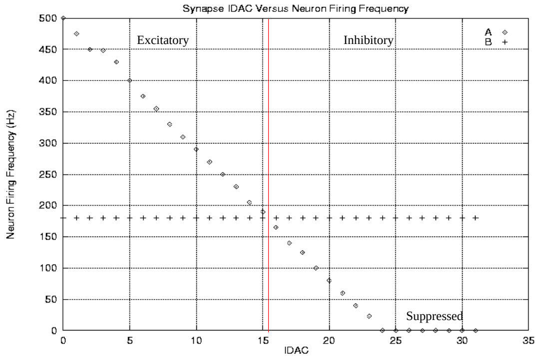

Figure 7. Effect of synapse interconnect strength on post-synaptic neuron firing rate. Trace A shows the firing rate of neuron 0 with increasing levels of synapse current DAC input. Trace B shows the neuron output with no synaptic input.

the synapse and the correct interaction with the neuron. The synapse behaviour at low levels of input is strongly excitatory decreasing as it approaches the mid-point (effectively zero synaptic weight - current DAC value 15), as it crosses the mid-point the synapse behaviour becomes inhibitory. The neuron firing rate is gradually reduced until it is finally suppressed.

4.2. Processing real sound input

A network of integrate and fire neurons was constructed with a linear array of 32 neurons, each neuron output being connected to 10 adjacent neurons, 5 on each side, with a fixed excitatory interconnect strength. In practice, an earlier version of the devices presented here was used which had fixed interconnect strength rather than the fully programmable synapses as described earlier. Otherwise the devices are similar.

The original chip has been tested with data from processed real sounds: see figure 8. The processing consists of digitisation followed by cochlear bandpass filtering into 31 bands using the Gamma-tone filter-bank (Patterson et al., 1995), followed by rectification, smoothing, onset enhancement and logarithmic compression as in (Smith, 1996). These emulate some of the transformations applied by the organ of Corti, spiral ganglion neurons and auditory nerve. This preprocessing was implemented in software. Signals were sent to the 32 neurons on 4 chips every 1ms.

We report the results of testing the systems with two different sound signals, namely a (locally generated) series of hand-claps (sampled at 22050 samples/second) and an utterance from the TIMIT database

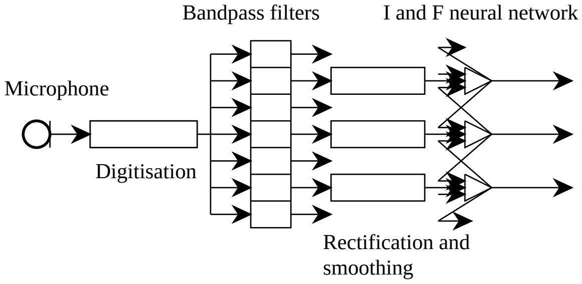

Figure 8. Block diagram of the sound processing system. Sound input was digitised and then bandpass filtered (see text). The output of each filter was rectified and compressed, then used as input to the hardware integrate-and-fire neural network. Each neuron receives input from exactly one channel, and outputs to the 5 adjacent neurons on each side (not all neurons are shown). The synapses have equal fixed excitatory strengths.

(dr1/fsjk1/sa1) (sampled at 16000 samples/second). Both the simulated and the hardware integrate-and-fire neural networks give similar results, as can be seen in figures 9. The primary difference in two onset volleys generated from the claps data occur at the first clap: the hardware appears to have been affected by some noise. In the onset volleys generated from the utterance, the low frequency channels are less in evidence in the hardware implementation. However, increasing the strength of the input results in too many pulses appearing.

The way in which the integrate-and-fire neural network clusters the data can be seen if we examine a 50ms section of the claps data. Looking at the 50ms section starting just before the second clap, we can examine the effect of turning off the excitatory weight in the simulation, and at the effect of altering the dissipation of the integrate-and-fire neuron whilst adjusting the input strength to ensure that the neurons still fire.

With the earlier chip, the excitatory weights are fixed; however, we can alter the strength of the input to the chip, adjusting the dissipation to ensure roughly the same number of spikes are generated. The results of this are shown in figure 11.

Applying the same techniques to the TIMIT utterance, we see some changes occurring as the strength of the input increases and the dissipation decreases: this can be seen in figure 9f, g, and h. In figure 12 we show an enlargement of a 50ms section of the TIMIT spikes. Comparing figures 10 and 11, one sees spatio-temporal clustering, and an decrease in latency of the onset of the first spike as the strength of the input is increased. However, the spikes produced are closer together in the hardware model (figure 11): this is believed to be due

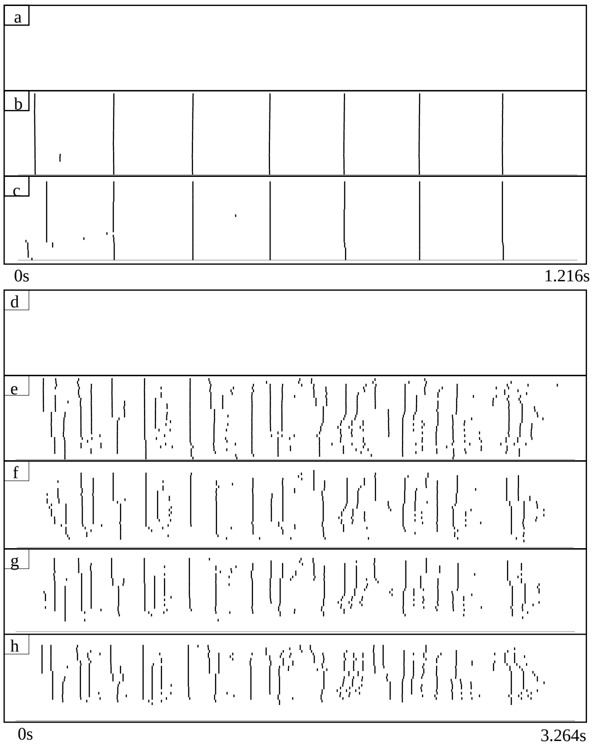

Figure 9. (a) envelope of the claps sound. (b) onset volleys found by the simulated system. Low frequency channels are at the bottom, and high frequency channels at the top. (dissipation $(=1/RC) = 50$) (c) onset volleys found by the hardware system. (dissipation = 40). (d) envelope of the TIMIT utterance. (e) onset volleys found by the simulated system (dissipation=50). (f-h) onset volleys found by the hardware system (f) input at quarter of full strength, dissipation=15 (g) input at half full strength, dissipation = 50 (h) input at full strength, dissipation = 120.

to the differences in neuron to neuron delays between the simulation and the hardware: in the simulation, there is a delay of 1 timestep (0.1ms), whereas spikes arrive at their destination almost instantly in the hardware implementation.

5. Discussion and further work

We have presented results characterising both the older and the newer versions of the chip. From these, it is clear that we can achieve a wide range of input strengths, a wide range of dissipations, and, in the newer chip, a range of five orders of magnitude of refractory period, and a wide range of both excitatory and inhibitory synapse strengths. For the older

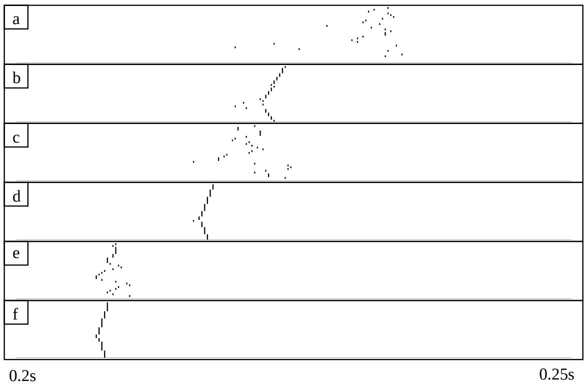

Figure 10. Results on a 50 ms section of the claps sound using the simulator. (a) Integrate-and-fire output with dissipation = 0, high input attenuation (weight = 0.001), and no excitatory weights between neurons. (b) as (a), except that excitatory weights are on (c) integrate-and-fire output with dissipation = 50, lower input attenuation (weight = 0.0025), and no excitatory weights between neurons. (d) as (c), except that excitatory weights are on. (e) integrate-and-fire output with dissipation = 500, low input attenuation (weight = 0.0175), and no excitatory weights. (f) as (e) except that excitatory weights are on.

Figure 11. Results on a 50 ms section of the claps sound using the silicon neurons. (a) input attenuated to quarter original strength, dissipation = 10 (b) input attenuated to half original strength, dissipation = 20 (c) full strength input, dissipation = 40.

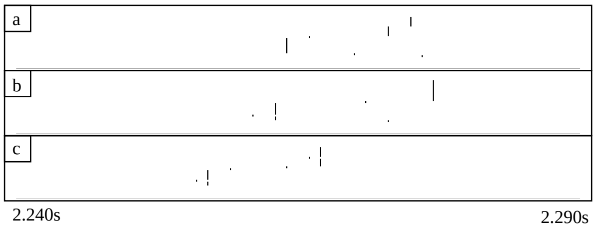

Figure 12. Results on a 50 ms section of the TIMIT sound using the silicon neurons. (a) input attenuated to quarter original strength, dissipation = 15 (b) input attenuated to half original strength, dissipation = 50 (c) full strength input, dissipation = 120.

chip, we have shown that the results produced from processing real data are similar to those produced by a software simulator.

We have not tested the newer chip on real auditory data. We had expected that by setting variable values to the corresponding values, we would see the same results as for the older chip. The new chip is more versatile because one can set its inter-neuron weights to a wide range of values. However, testing this chip has shown up problems with noise. These manifest themselves as variation in inter-spike interval for each neuron, even with constant input. This has restricted the range of possible application of the new chip. The results from the new chip are described and discussed in (Smith et al., 1999) implementing oscillatory synchronizing networks.

Acknowledgements

At the time this research was undertaken, Mark Glover was supported by the UK EPSRC.

References

U. Ernst, K. Pawelzik, and T. Geisel. Multiple phase clustering of globally pulse coupled neurons with delay. In M. Marinaro and P.G. Morasso, editors, *ICANN94: Proceedings of the international conference on artificial neural networks, Sorrento, Italy, 26-29 May 1994*, pages 1063–1066. Springer-verlag, 1994.

W. Gerstner. Time structure of the activity in neural networks. *Physical Review E*, 51(1):738–758, 1995.

L.S. Smith M.A. Glover and A. Hamilton. A comparison of a hardware and a software integrate and fire neural network for clustering onsets in cochlear filtered sound. In *Neural Networks for Signal Processing VIII: Proceedings of the 1998 Workshop*, pages 516–523, 1998.

M. Glover, A. Hamilton, and L.S. Smith. , an analog VLSI integrate-and-fire neural network for sound segmentation. In *Proceedings of the International ICSC/IFAC Symposium on Neural Computation (NC'98)*, 1998.

M.A. Glover, A. Hamilton, and L.S. Smith. Analogue VLSI integrate and fire neural network for clustering onset and offset signals in a sound segmentation system. In L.S. Smith and A. Hamilton, editors, *Neuromorphic Systems: engineering silicon from neurobiology*, pages 238–250, 1998.

H. Glunder and A. Nischwitz. On spike synchronization. In A. Aertsen, editor, *Brain Theory*, pages 251–258. Elsevier, 1993.

A. Hamilton. Synaptic weight representation in pulsed neural network VLSI: Voltage or current? In *Proc. Int. Conf. on Neural Information Processing, Seoul, South Korea*, pages 353–358, 1994.

T.V. Ho and J. Rouat. A novelty detector using integrate and fire neurons. In W. Gerstner, A. Germond, M. Hasler, and J-D Nicoud, editors, *Artificial Neural Networks - ICANN-97*, pages 103–108. Springer-Verlag, 1998. Lecture Notes in Computer Science 1327.

A. Jahnke, T. DSchonauer, U. Roth, K. Mohraz, and H. Klar. Simulation of spiking neural networks on different hardware platforms. In W. Gerstner, A. Germond, M. Hasler, and J-D Nicoud, editors, *Artificial Neural Networks - ICANN-97*, pages 1187–1192. Springer-Verlag, 1998. Lecture Notes in Computer Science 1327.

L. Lapique. Sur l'excitation electrique des nerfs. *J. Physiology (Paris)*, 9:620–635, 1905.

W. Maass. Lower bounds for the computational power of networks of spiking neurons. *Neural Computation*, 8(1):1–40, 1996.

W. Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.

A. Nishwitz and H. Glunder. Local lateral inhibition - a key to spike synchronization. *Biological Cybernetics*, 73(5):389–400, 1995.

R.D. Patterson, M.H. Allerhand, and C. Giguere. Time-domain modelling of peripheral auditory processing: A modular architecture and a software platform. *Journal of the Acoustical Society of America*, 98:1890–1894, 1995.

S. Rosen. Temporal information in speech: acoustic, auditory and linguistic aspects. *Phil. Trans. R. Soc. London B*, 336:367–373, 1992.

L.S. Smith. Onset-based sound segmentation. In D.S. Touretzky, M.C. Mozer, and M.E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 729–735. MIT Press, 1996.

L.S. Smith. Extracting features from the short-term structure of cochlear filtered sound. In J.A. Bullinaria, D.W. Glasspool, and H. Houghton, editors, *4th Neural Computation and Psychology Workshop, London, 9-11 April 1997*, pages 113–125. Springer Verlag, 1998.

L.S. Smith, B.E. Eriksson, A. Hamilton, and M.Glover. SPIKEII: an integrate and fire AVLSI chip. *International Journal of Neural Systems*, 9(5):479–484, 1999.

N.P. Todd, C.S. Lee, and D.J. O'Boyle. A sensory-motor theory of rhythm, time perception and beat induction. *Journal of New Music Research*, 28(1):5–29, 1999.

T.P.Zahn, R. Izak, K.Trott, and P. Pashke. A paced analog model of auditory attention. In L.S. Smith and A. Hamilton, editors, *Neuromorphic Systems: engineering silicon from neurobiology*, pages 238–250, 1998.

A. van Schaik, E. Fragnière, and E. Vittoz. A silicon model of amplitude modulation detection in the auditory brainstem. In M.C. Mozer et al., editor, *Advances in Neural Information Processing Systems 9*, pages 741–747. MIT Press, 1997.

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

Address for Offprints: Dr Alister Hamilton, Department of Electronics and Electrical Engineering, University of Edinburgh, King's Buildings, Mayfield Road, Edinburgh EH9 3JL, Scotland, UK