PERGAMON

Neural Networks 14 (2001) 755–762

Neural Networks

www.elsevier.com/locate/neunet

2001 Special issue

Spike-based VLSI modeling of the ILD system in the echolocating bat

Timothy Horiuchi[a,*], Kai Hynna[b]

[a]*Electrical and Computer Engineering Department and the Institute for Systems Research, AV Williams Building, University of Maryland, College Park, MD 20742, USA*

[b]*Department of Bioengineering, 120 Hayden Hall, 3320 Smith Walk, University of Pennsylvania, Philadelphia, PA 19104, USA*

Received 1 February 2001; accepted 1 February 2001

Abstract

The azimuthal localization of objects by echolocating bats is based on the difference of echo intensity received at the two ears, known as the interaural level difference (ILD). Mimicking the neural circuitry in the bat associated with the computation of ILD, we have constructed a spike-based VLSI model that can produce responses similar to those seen in the lateral superior olive (LSO) and some parts of the inferior colliculus (IC). We further explore some of the interesting computational consequences of the dynamics of both synapses and cellular mechanisms. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Echolocation; Spike-based; VLSI; Azimuthal localization; Hardware model; Masking

1. Sound localization and bats

While the bat's carefully controlled ultrasonic sonar pulse is not a feature common to most mammals, the neural mechanisms for determining the direction of the reflected sound are thought to be very similar. The azimuthal direction of a sound source can be extracted using two basic concepts: measuring the difference in arrival times of a sound at spatially separated receivers and measuring the difference in intensity level at spatially separated receivers at different frequencies. Practical constraints may limit the usefulness of either or both of these cues, such as small receiver separation or insufficient attenuation between receivers. In humans, the ears are sufficiently separated to produce about 800 μs of time delay between eardrums. Our head and pinnae also attenuate sounds in the upper frequency range of our hearing to provide useful cues about sound direction. Given the critical need of animals to detect and localize predators, prey, mates and conspecifics, nearly all animals have developed some sound localization abilities for survival.

The echolocation systems of bats primarily utilize ultrasonic frequencies to detect objects in three-dimensions. The use of high frequencies is critical for producing echoes from tiny objects (insects), for limiting the range of echoes and for producing a direction-dependent intensity level difference at its two ears. Given the head size of most echolocating bats (~1 cm), the maximum interaural time difference cue (~50 μs) is very small. Additionally, at ultrasonic frequencies, the hair cells of the cochlea in bats are unable to transmit phase information about the waveform, leaving only the envelope modulation to provide interaural time difference cues. The interaural level difference (ILD) cue is, therefore, the dominant acoustic feature for azimuthal echo localization in bats.

Sound waves reaching the head of the animal attenuate as they wrap around towards the back of the head. Wavelengths smaller than the size of the head are significantly attenuated and wavelengths larger than the size of the head are not. This frequency-dependence interacting with various 'shadow producing' features of the head, snout, body, and pinnae produce a complex direction-dependent transfer function that can be used to determine the direction of sound by measuring the attenuation factor at many different frequencies. By subtracting the logarithm of the signal amplitudes (i.e. intensity level) at each ear, we produce a measure of this attenuation factor.

2. The interaural level difference pathway

While it is generally agreed that the inferior colliculus is the final convergence point for most spatially oriented computations done in the brainstem and midbrain, multiple, parallel pathways are known to exist. For example, the sensitivity to ILD found in many cells of the midbrain of

* Corresponding author. Tel.: +1-301-405-7412; fax: +1-301-314-9920.
E-mail address: timmer@isr.umd.edu (T. Horiuchi).

PII: S0893-6080(01)00066-1

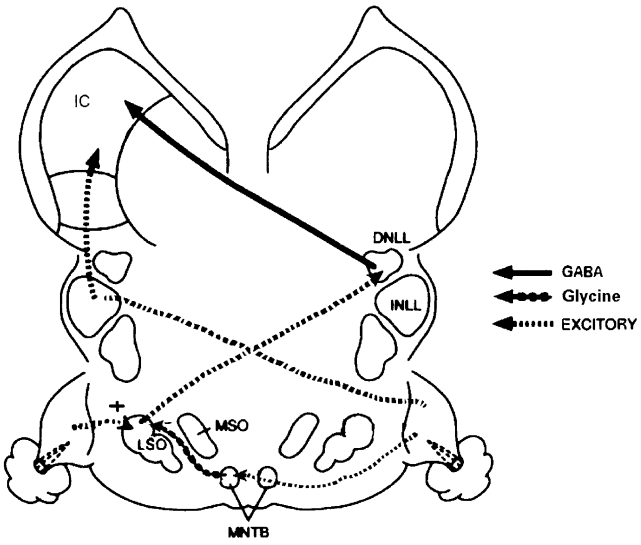

Fig. 1. The auditory midbrain pathway (mustache bat) involved in the computation of interaural level differences. This modeling work focuses on the LSO and IC connections. Figure adapted from Pollak and Park (1995).

the bat are of the excitatory-inhibitory (EI) type. EI-type cells are distinguished by their excitatory response to sounds in one ear and inhibitory response to sounds in the opposite ear. Pollak and Park (1995), using known anatomy and physiology, have identified five different circuits that can produce this behavior in the inferior colliculus. In this paper, we construct and explore one of these circuits. In particular, we mimic experiments by Pollak and Burger (2000) with multiple targets to expose the effects of dynamic synapses on such circuits.

The earliest binaural comparison of sounds from the two ears occurs at the lateral superior olive, or LSO. As shown in Fig. 1, it effectively receives excitation from the ipsilateral ear (via the cochlea and the cochlear nucleus) and inhibition from the contralateral ear, performing a (left–right) operation. This excitatory-inhibitory arrangement is referred to as EI and creates a threshold at a particular interaural intensity level difference. When the sound is ipsilateral to a particular LSO cell, the received excitation will, in general, be stronger than the inhibition coming from the contralateral ear, resulting in a spike. If, however, the sound is contralateral to the LSO cell, then the resulting inhibition can prevent the cell from firing. Recall that echoes will be short in duration (0.5–3 ms) and will generally elicit only one or two spikes when maximally driven.

The outputs of the LSO cells are known to project in many different pathways towards the inferior colliculus (Pollak & Park, 1995). We have focused on one particular pathway that results in the same EI pattern of response. This circuit consists of inhibitory projections originating in the ipsilateral LSO cell (which responds to ipsilateral targets) and excitatory projections originating from the contralateral ear. For the purposes of our simulation, we have simplified the inhibitory projection from the LSO to be directly inhibitory on the IC cell (instead of crossing over twice via the dorsal nucleus of the lateral lemniscus, or DNLL) and we have simplified the excitatory projection to be directly from the cochlear nucleus (instead of through the intermediate nucleus of the lateral lemniscus, or INLL). This type of IC cell thus receives its monaural excitation from the contralateral side and binaural inhibition from the ipsilateral side. An example ILD (or sometimes referred to as IID) curve is shown in Fig. 2.

Using an artificial sonar system, we demonstrate the functionality of this circuit in two stages; first we construct the LSO circuit and then expand the model to include the IC cell. Our working circuit architectures are shown in Fig. 3. Layer 1 represents the LSO connectivity and functionality, receiving excitation from the ipsilateral side and inhibition from the contralateral side. The multiple units shown here represent cells with different weighted connections to produce different inhibitory thresholds. Layer 2 represents one of several possible IC circuits, but one that specifically generates the EI property in the IC using the contralateral monaural excitatory input and ipsilateral inhibition from the LSO. While we will show the basic functionality, we also

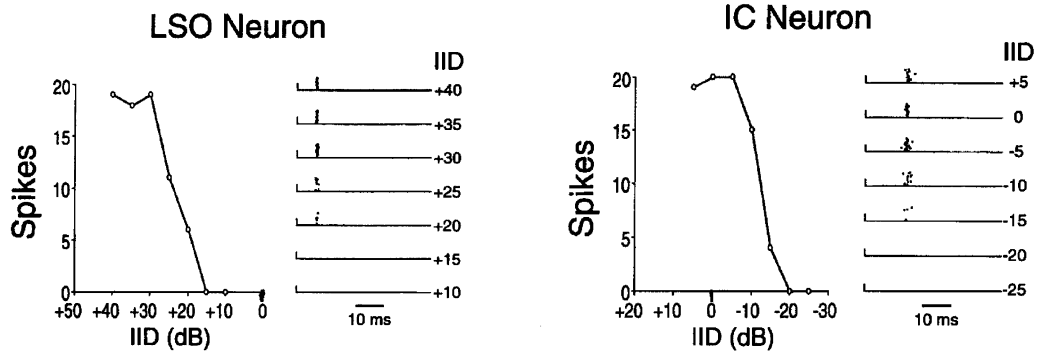

Fig. 2. Spike response (per 20 trials) as a function of interaural intensity difference (same as ILD). Left: LSO cell IID tuning curve to a brief stimulus. Right: IC cell IID tuning to a brief stimulus. Note that the interaural intensity difference corresponds to the direction-dependent attenuation factor. The spike rasters on the right of each trace show the same data but highlight the fact that cells typically fire one or no spikes per trial. (Mexican free-tailed bat) Adapted from Park (1998).

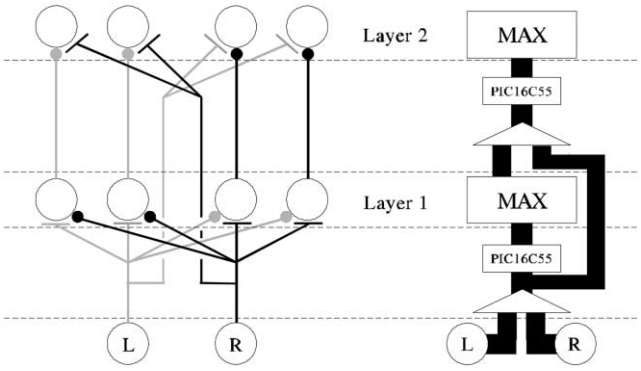

Fig. 3. Bars are excitatory connections and circles are inhibitory connections. Left: Circuit architecture for the LSO and IC cells. Layer 1 receives ipsilateral excitation and contralateral inhibition, while layer 2 effectively receives ipsilateral inhibition and contralateral excitation. Right: circuit architecture for the VLSI-based address-event neural circuitry (see text, Section 4).

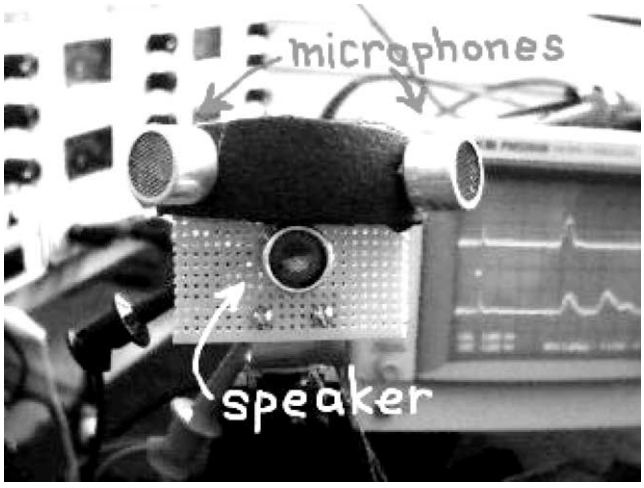

Fig. 4. Photo of the narrowband sonar front-end. The oscilloscope in the background on the right shows an example echo return from the three legs of a tripod. The sonar head is approximately 2.5″ tall and is mounted on a model airplane servo for tracking experiments (not described).

demonstrate the effect of time-dependent inhibition on multiple sonar targets. In particular we will show masking behavior in the LSO and its consequences in the IC.

3. The narrow-band sonar front-end

The experimental front-end of the current system is a narrow-band (39–41 kHz) ultrasonic speaker and microphone pair. The microphones are oriented at nearly 90° from each other to produce sensitivity patterns such that the relative amplitude of signal received in the two microphones is consistent with a particular direction. While this sonar head is larger than a typical bat head and we are able to detect the interaural time difference, this information is not used because it is not available to the bat. This narrow-band sonar front-end will soon be replaced by a tiny (1 cm^3) broadband system using miniature MEMS microphones and speaker (Fig. 4).

A burst of 40 kHz sound (0.5 ms duration) is transmitted from the speaker periodically to generate echoes from targets placed at various points in the room. Each microphone signal is amplified and broadly filtered (the microphones are very narrowly tuned, which provides most of the frequency tuning, but limits our temporal response) before being half-wave rectified. The half-wave rectification drives a low-pass filter to simulate the dynamics of the hair cell in the cochlea. At ultrasonic frequencies, this results in a signal proportional to the echo amplitude with no carrier phase information. This envelope (an example of which is visible in the photo of the sonar head) is used to drive a spike-generation circuit that simulates the spikes seen on the auditory nerve.

The envelope signal is then used to drive a spike generating circuit that fires proportionally to the envelope amplitude. This produces a very high spike rate that is intended to reproduce the *population* response of a pool of cells with different thresholds. Fig. 5 shows an example of this using the echo returns off of the legs of a camera tripod.

Fig. 6 shows the direction-dependent attenuation produced by the sonar head for echoes reflecting off of a cylinder placed approximately 1 m away. These are representative signals that will be used to drive the neurons in the neuron chip described in the next section.

4. Circuit models and network connectivity

Each layer of neurons is represented by a single neuromorphic transceiver chip called the MAX chip. The MAX chip contains a one-dimensional array of 36 silicon neurons. The array is divided into two sections, with each section sharing a common set of circuit biases. All neural connectivity and activity is handled off-chip using *Address Event Representation*. The MAX chip was fabricated using 0.8 μm HP technology, with a total area of 3.2 mm^2.

4.1. Silicon neuron

Our silicon neuron model (Fig. 7) can be divided into three distinct circuits: a spike generator, Ca^{2+} dependant K^+ channels, and two synapses (not shown). The spike generator is a modified version of the axon hillock circuit described by Mead (1989). Input current, I_{in}, from the synapses integrates onto the membrane node capacitance C_m ($= C_1 + C_2$). When the membrane voltage V_m reaches threshold (controlled by V_{pu}), the cascaded inverters transition, causing V_{spk} to switch from Gnd (resting state) to V_{dd}. Due to the coupling capacitor C_2, the transition of V_{spk} increases V_m above the threshold by $V_{dd}\ C_2/(C_1 + C_2)$. With V_{spk} now high, the reset current through the two series transistors begins to remove charge from V_m, at a rate controlled by the gate voltage V_{mpd}. Once V_m decreases

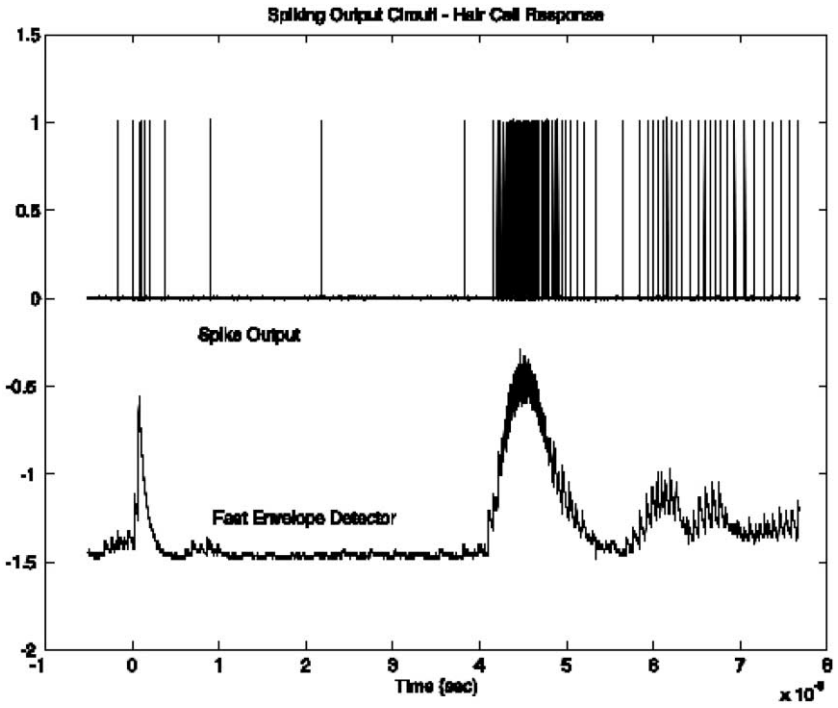

Fig. 5. Bottom: Filtered echo waveform is half-wave rectified and lowpass filtered. The envelope is then used to generate spikes at a rate proportional to the envelope height. The single spiking neuron is intended to represent the firing of a pool of neurons with different firing thresholds. Top: spike output vs. time following a sonar pulse. Note: the first pulse in the envelope and the corresponding spikes are due to the incomplete suppression of the outgoing sonar pulse. These spikes will be suppressed prior to driving the echolocation system.

below threshold, the cascaded inverters transition back to resting state and V_{spk} becomes low, shutting off the reset current and decreasing V_m from threshold by a similar amount V_{dd} $(C_2/C_1 + C_2)$. The neuron is now again ready to integrate the input. The width of the output pulse (t_{spk}) is controlled by the reset current I_{mpd}. For $I_{mpd} \gg I_{in}$, the pulse width is negligible to the interspike interval. In this state, the spike generator behaves like an integrate and fire neuron, with the firing rate dependant on the input current (see Mead, 1989).

The dynamics of Ca^{2+} dependent K^+ channels are modeled using a current-mirror integrator circuit. The voltage, V_{Ca}, at the integrator node is the circuit analog of intracellular calcium concentration. Each time the neuron spikes, a small quanta of charge—controlled by V_{qa}—is dumped onto the capacitor C_3. The diode-connected transistor leaks the accumulated charge away at a rate dependent on the source voltage V_A. If charge enters the integrator node faster than it can decay, then the voltage at that node, V_{Ca}, rises. Equilibrium is reached when amount of charge removed from C_3 during the interspike interval matches the amount added by each spike. A single transistor, whose drain is attached to V_m and whose gate voltage is V_{Ca}, represents the Ca^{2+} dependent K^+ channels. A sudden

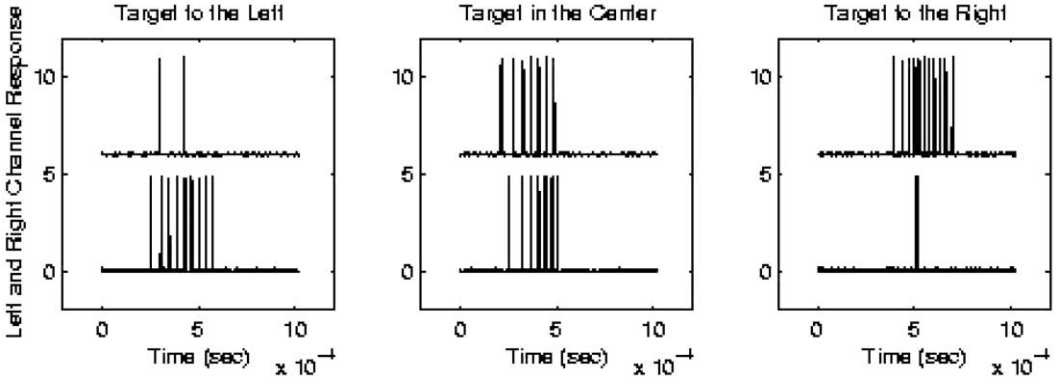

Fig. 6. Left and right ear spike outputs for a distant cylinder presented 45 deg to the left, center, and 45 deg to the right of the sonar's midline. The timescale is in milliseconds (10 e-3).

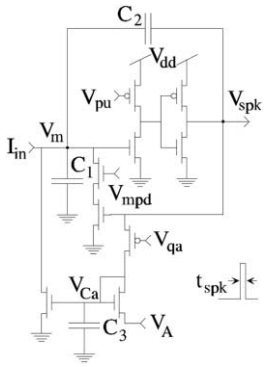

Fig. 7. Schematic of silicon neuron. Not shown here are the synapses (see Fig. 8). See text for details.

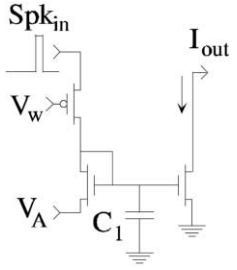

Fig. 8. Inhibitory synapse. An input pulse dumps a quanta of charge (controlled by V_W) onto a current-mirror integrator. The decay of the charge depends on the source voltage V_A. The output is connected to V_m, and thus inhibits the cell. The excitatory synapse is identical except the output passes through a p-type current mirror to reverse the current.

increase in the firing rate of the neuron—due to a step increase in input current for example—floods the integrator node and V_{Ca} rises. As a result, the K^+ channel current increases which then slows down the firing rate of the neuron until a new equilibrium at the new current level is reached. This effect, also found in biology, is called spike-rate adaptation.

Two synapses—one excitatory the other inhibitory—provide input into the silicon neuron. These synapses have the same form of the Ca^{2+} dependent K^+ channels described above, namely a current-mirror integrator (Fig. 8). An incoming presynaptic spike (Spk_{In}) dumps a quanta of charge onto the integrator node through a transistor. The gate of the transistor, V_w, controls the amount of charge added for each incoming pulse and thus acts as the synaptic strength. The charge on the integrator node decays via the diode connected transistor at a rate dependent on the source voltage V_S. The output current I_{out} is controlled by the node voltage and is either inhibitory (directly connected to V_m of the neuron) or excitatory (passing through a P current mirror to reverse the current). It is this finite duration (set to about 1–2 ms) of inhibitory current that utilize in our circuit operation described in Section 5. For more information on the dynamics of the synapses, see Indiveri (2000).

4.2. Address event representation

Input and output activity for the chip is conveyed over a digital bus using *address-event representation* (see Fig. 9). Each silicon neuron is given a unique identifying address so when it fires, an encoder transmits its address onto the output address-event bus. The time of the spike is implicitly encoded in the occurrence of the address on the bus. As long as the activity within the chip is less than the bandwidth of the bus, then the activity is faithfully transmitted down the line. Should two neurons fire 'simultaneously', then an arbiter circuit (similar to that described in the next section) chooses which of the two addresses is sent first. To ensure successful transmission, handshaking occurs between a transmitter and receiver. The transmitter signals the occurrence of an event on the bus, and the receiver acknowledges the receipt of the new address. In the case of two neurons firing nearly simultaneously, the second must wait for the acknowledge from the arbiter before its address can be transmitted. The delay introduced as a result is negligible (see Boahen, 1998).

On the input side, each synapse, rather than each neuron, is given a unique address. A decoder reads the new address-event and sends a pulse to the appropriate synaptic location. Once completed, the receiver circuitry acknowledges the now expired data on the bus and waits for a new event.

4.3. Neural connectivity

Fig. 9 depicts a simple, but not-so-useful, example of

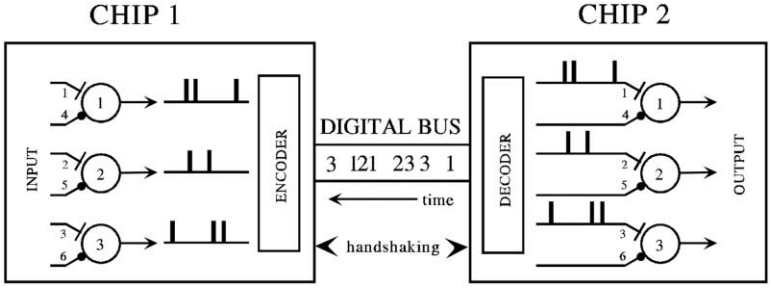

Fig. 9. Communication between to transceiver chips using Address Event Representation. The address of the active neuron is transmitted onto the address event bus by an encoder. A decoder translates the input address event into a pulse at the appropriate synapse. Handshaking occurs between the communicating chips in order to ensure proper transmission of the event. Adapted from Boahen (1998).

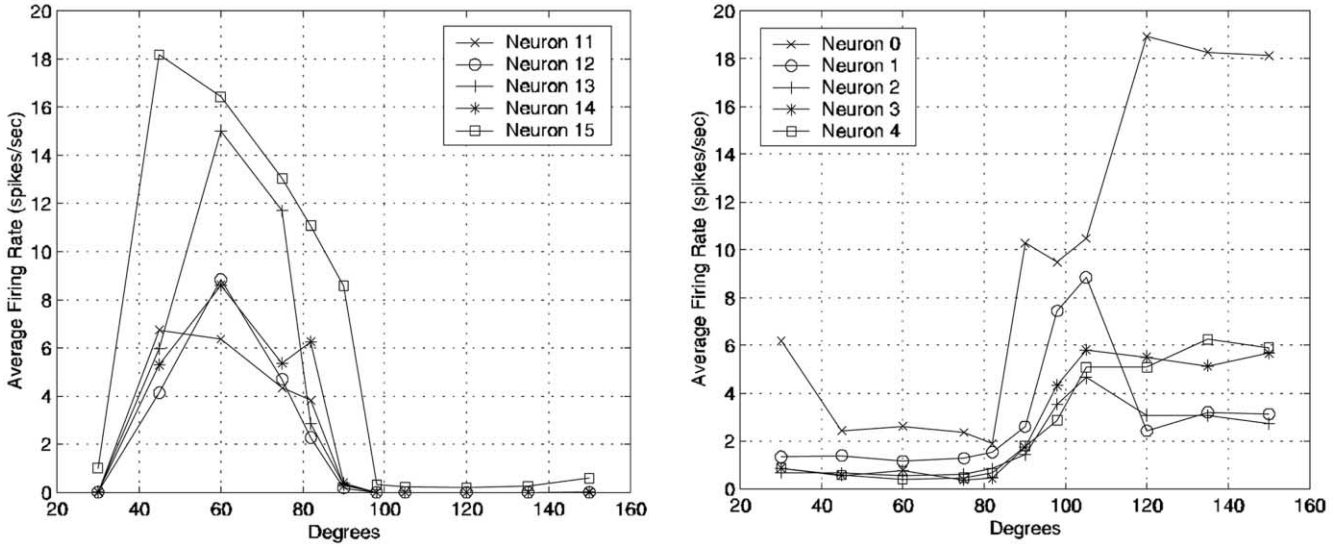

Fig. 10. Responses of the VLSI LSO neurons to echoes from different directions. Left panel: EI neurons with excitation coming from the left and inhibition from the right. This cell fires for echoes coming from the left side. Note that the drop in response for echoes from the far left is a result of the left microphone's drop in sensitivity. Right panel: EI neurons with excitation coming from the right and inhibition from the left. This cells fires for echoes coming from the right side.

connecting two chips together. To illustrate this, imagine neuron 2 from the first chip fires. The encoder transmits its address (2) onto the event bus. The decoder in chip 2 reads the new data and inputs a pulse into synapse two, in this case the excitatory synapse of neuron 2. In this situation, the maximum fan-out of the neuron is one; there would be no point having more synapses than neurons. To increase the fan-out of the neuron (as well as to provide flexibility in the circuits), a lookup table (LUT) is placed onto the digital bus between the two chips. This LUT acts as a translator for new events from chip 1 en route to chip 2, essentially storing the connectivity pattern between the layers. Going back to the previous example, when the LUT receives address 2, it can send out addresses {2,4,6} (a fan-out of 3) to chip 2, exciting the central neuron and inhibiting the neighbors.

The LUT can be any sort of reprogrammable device(s) so long as it can communicate with other chips on the digital bus (i.e. it can be programmed to handshake). The speed of the LUT is also very important, as bandwidth is a major issue in transmitting spikes reliably through the network. Lower speeds results in less complex circuits, as the bandwidth of the bus is dictated by how fast the LUT can translate and communicate with the chips. For example, ten neurons firing at a rate of 100 Hz with a fan-out of 10 requires at least a bandwidth of 10 kHz.

We used a Microchip PIC16C55 microcontroller (PIC) as the lookup table for each neural array. The PIC16C55 has a 200 ns instruction cycle, with 24 bytes of RAM and 768 bytes of program memory. An arbiter circuit provided the input to each PIC. The arbiter is a simple circuit that merges two address-event streams into a single bus, providing all the appropriate handshaking to communicate with the various streams. When there is a new event on one of the buses, the merge circuit allows the data through to the output bus, usually adding an extra bit to the address indicating the source. Should an event arrive at the other input, it waits until the first has been transmitted before it can pass through. The arbiter is a simplified version of the arbiter circuitry found within the MAX chip (see Boahen, 1998).

In the first layer, the activity from each ear was translated into excitatory inputs for ipsilateral neurons and inhibitory inputs for contralateral neurons (as in Fig. 3). Since the MAX chips provides the same synaptic weight for all neurons, an attempt to create variable weights was performed in the PIC using pseudo-probabilistic synapses. Given the complexity of the PIC code (for these synapses), the fan-out of each connection and the firing rate of the ear neurons, we were limited to using only 10 neurons (five for each side) in order stay within bandwidth limits of the bus. Only the excitatory synapses (five for each ear) were made probabilistic, each allowing only a certain percentage of the activity through (100, 90, 80, 67 and 50%). The second layer connections did not attempt any sort of variable weighting. Ear inputs excited all five contralateral neurons while neurons from layer I inhibited their corresponding partner in layer II.

5. Experimental results

In Fig. 6, we show the responses of various units to echoes returning to the sonar system from different directions. These spiking inputs were used to drive the LSO circuitry defined by the address-event infrastructure. The EI circuit arrangement for the LSO units produces the expected threshold in ILD (Fig. 10), however, the expected

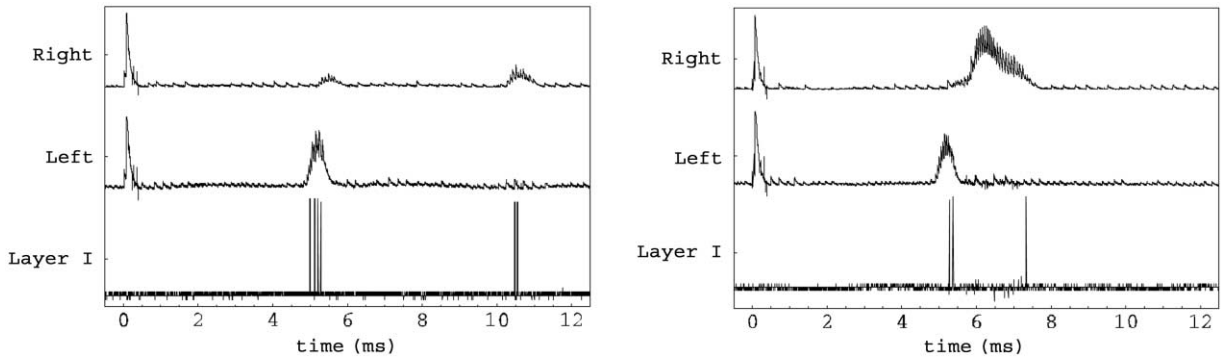

Fig. 11. Masking behavior of inhibition. Two targets on opposite sides of the midline are moved close to each other. The top trace is the envelope output of the right channel, the middle trace is the envelope of the left channel and the bottom trace is a digital output that fires whenever any neuron on the chip fires. As the target on the right approaches the first target (even though the echo signal is much larger) the neuronal response is significantly lower.

variation in threshold position due to the differential weightings is extremely weak.

An interesting aspect of the circuit is revealed when we observe the response of the LSO cells to multiple targets. Fig. 11 shows two experiments where a distant target on the opposite side of the sonar's midline is moved closer until the spiking response that it normally elicits (when presented alone) is diminished. This is a consequence of the inhibition produced on the LSO cells by the closer target. This masking effect produced a shadow behind the first target in spite of a clear, strong echo visible in the envelope trace. It should be noted that not all LSO neurons would be disabled by this inhibition, rather a limited population would provide a differential response for this type of target configuration.

For the second layer, excitatory inputs came from the contralateral ear directly (via the AER bypass pathway) and inhibition came from one of the neurons in the first layer. No probabilistic neurons were used in this case. While this configuration simply reproduces the EI result seen in layer 1, an additional feature emerged that we describe in below.

Fig. 12 shows an experiment (a control and a probe case) where a second target is introduced in front of an existing target, but on the other side of the midline. In the control case (left column), a cell that responds to targets on the left is stimulated by a target on the left. The corresponding cell on the right is suppressed (by the LSO cell that responds to targets on the left). When a target on the right is placed in front, the LSO that normally provides suppression of that cell is momentarily disabled by the duration of inhibition.

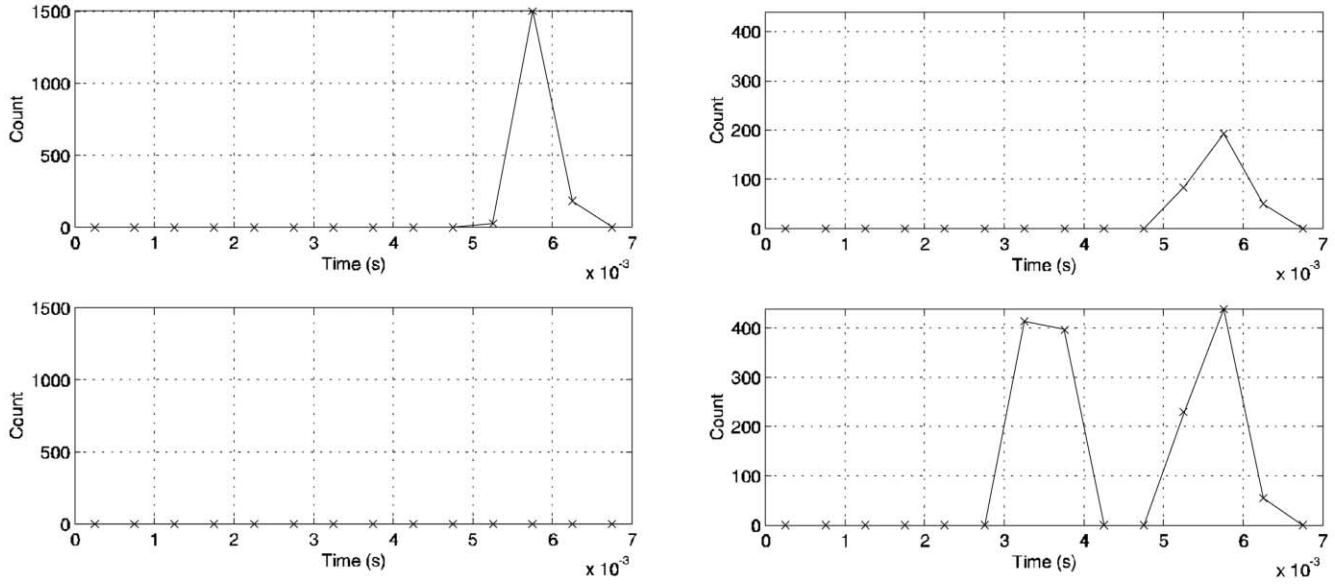

Fig. 12. Layer 2 (IC) post-stimulus histograms (PSTH) for two stimulus conditions. The top row represents the output spikes from the left-side IC cell, and the bottom row represents the output spikes from the right-side IC cell. Left column: a single target is placed to the right of the midline. The left IC cell responds while the right side is inhibited by the right LSO cell. Right column: a second target is added in front of the first target and to the left of the midline, eliciting a response from the right IC cell. The first target (left side) also inhibits the right LSO strongly enough that the second target cannot inhibit the right IC cell (through the right LSO) as in the case before.

The functional knockout of inhibition by a preceding stimulus has serious consequences on the spatial receptive field of the cell. Instead of limiting the response to a particular direction, it will now respond to a larger region of space. Since its excitation comes from a monaural source, the cell has been rendered momentarily monaural. This transient enlargement could possibly be used to define complex receptive fields. When combined with the masking effects seen in the basic LSO unit, a richer picture of the bat's feature detection space emerges.

6. Discussion and conclusions

In this paper we describe our ongoing efforts to construct a large-scale model of the neural processing in the echolocation system of the bat. To date, we have focused on the processing of a single frequency band, however, our expansion into broadband processing has already begun. Using a spike-based representation for echo amplitude, we have demonstrated the EI processing in an integrate-and-fire type of neuron to produce units that only respond to targets from a limited range of directions, consistent with the types of responses seen in LSO cells of the bat. We have gone further to demonstrate a second layer of processing that combines information from the ipsilateral LSO layer and the contralateral signal inputs to produce one type of IC cell's EI response. We investigate the dynamic effects of inhibition (1–2 ms) on the response of the various cells and suggest a possible utility for these dynamics.

One of the many questions that arise when sifting through the myriad possible circuits that could be constructed from the known neuroanatomy of the pathways leading to the inferior colliculus is, 'what is the difference between an EI response computed in LSO and passed up to the IC and an EI response computed de novo in the IC?'. Recently, Pollak and Burger (2000) demonstrated the construction of EI responses in IC using excitation from one ear and a GABAergic projection from the dorsal nucleus of the lateral lemniscus (DNLL). When two sounds were played in quick succession, the second sound which normally produces inhibition on the IC, was shown to be functionally inhibited by the first sound. As they suggest, the cell has a more complex spatial receptive field due to the dynamics of the circuitry used to drive it.

While we have not exactly duplicated all of the known elements of the bat midbrain pathway in our silicon model, we have been able to demonstrate similar phenomena that are only present in models that preserve the dynamics of the original neural circuit. This type of modeling hints that the EI cells in the bat IC have much more diversity in their spatial response fields that is not well described by experiments using only single targets and static models of computation.

Acknowledgements

The authors would like to thank Professor Kwabena Boahen for continuing support of Mr Hynna, Francis Chew and Gideon Romm who designed the early version of the AER neuron chip, David Goldberg (Johns Hopkins University) and Socrates Deligeorges (Boston University) for their involvement in the early stages of this project at the 2000 Telluride Neuromorphic Engineering Workshop. We thank the Telluride Neuromorphic Engineering Workshop for making this collaboration possible. Mr Hynna is supported by a NSF LIS/KDI Grant (KDI-ECS98-74463) and Dr Horiuchi is supported by a NSF Learning and Intelligent Systems Grant (BES-972053) and by the Minta Martin Aeronautical Engineering Foundation at the University of Maryland.

References

Boahen, K. A. (1998). Communicating neuronal ensembles between neuromorphic chips Chap. 11. In T. S. Lande, *Neuromorphic systems engineering* (pp. 229–261). Boston, MA: Kluwer Academic Publishers.

Indiveri, G. (2000). Modeling Selective Attention Using a Neuromorphic Analog VLSI Device. *Neural Computation*, *12* (12), 2857–2880.

Mead, C. (1989). Axons. *Analog VLSI and neural systems* (pp. 193–204). Menlo Park: Addison-Wesley.

Park, T. (1998). IID Sensitivity differs between two principal centers in the interaural intensity difference pathway: the LSO and the IC. *J. Neurophysiol.*, *79* (5), 2416–2431.

Pollak, G., & Burger, R. M. (2000). Reversible inactivation of the dorsal nucleus of the lateral lemniscus reveals its role for processing multiple sound sources in the inferior colliculus, *Poster Presentation at the Symposium on Synaptic Function in Hearing and Balance*, Johns Hopkins University.

Pollak, G., & Park, T. (199). Inferior colliculus. In A. Popper & R. R. Fay, *Hearing by bats* (pp. 296–367). New York: Springer-Verlag.