A Spike-Latency Model for Sonar-Based Navigation in Obstacle Fields

Timothy K. Horiuchi, *Member, IEEE*

***Abstract*—The rapid control of sonar-guided vehicles through obstacle fields has been a goal of robotics for decades. How sensory data are represented strongly affects how obstacles and goal information can be combined to select a direction of travel. Many approaches combine attractive and repulsive effects to steer; we have implemented an algorithm that first evaluates the desirability of different directions followed by a winner-take-all (WTA) mechanism to guide steering. We describe a neuromorphic VLSI implementation of this algorithm using the inherent echo delay of obstacles to produce a range-dependent gain in a "race-to-first-spike" neural WTA circuit.**

***Index Terms*—Collision avoidance, echolocation, neuromorphic VLSI, obstacles, spike latency, spike timing, step inhibition, winner-take-all (WTA).**

I. Introduction

THE EASE with which echolocating bats appear to use ultrasonic echoes to perceive their 3-D world has long been the fascination of scientists and engineers. Capable of flying through dense forests in complete darkness during their hunt for flying insects and other prey, echolocating bats must do more than simply home in on the closest object. How biological systems can transform the storm of sensory information into short-term motion plans amid multiple obstacles and goals is an ongoing quest for many roboticists and neuroscientists. Noisy, ambiguous sensory data, limited time to make decisions, and the tricky question of what an obstacle is, all make this a difficult task.

Recent approaches to this local obstacle-avoidance problem have utilized the summation of repulsive or attractive torques generated by obstacles and goals to steer a particular creature (e.g., [1] and [2]). While these approaches are successful and interesting because they also include vehicle dynamics into the steering choice, there are a few considerations that suggest a different philosophical approach. There are situations where obstacles on either side of the creature and a goal straight ahead can force the creature into a collision with a narrow gap. In these approaches, an obstacle produces a repulsive torque in a specific direction, whereas it seems that obstacles should not turn the creature in any particular direction but only indicate to the creature where it should *not* go.

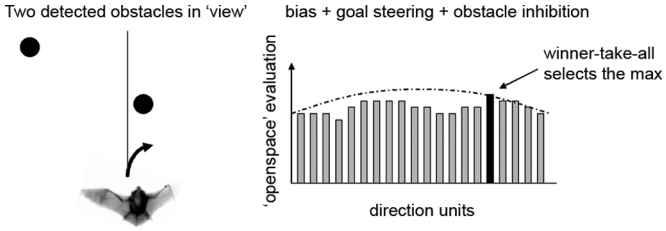

Fig. 1. (Left) An echolocating bat that is attempting to fly directly forward detects two obstacles (filled circles). (Right) The evaluation pattern consists of a constant plus a wide low-amplitude Gaussian with two dips created by the suppression from the two obstacles. Multiple choices are available; however, a WTA function selects the direction with the highest evaluation. In a simpler reflex-based system, this right turn toward an object would generally not be an option. The dotted line indicates the default evaluation with no obstacles present. Simulation and robot movies can be viewed at http://www.isr.umd.edu/~timmer.

A. "Openspace" Algorithm

The approach presented here is similar to the work by Borenstein and Koren [3] which takes a risk-minimization view of navigation, using the sonar system to first evaluate the *desirability* of different directions of travel simultaneously and then to select the direction with the highest evaluation.

The evaluation process begins with a field of evaluation units that receive an initial input pattern that represents prior assumptions about the desirability of various directions. This pattern could incorporate information about actuation limits, vehicle dynamics, energy conservation, single or multiple goal directions, history of previous choices, etc. Obstacles then produce a pattern of suppression (i.e., inhibition) on the evaluation pattern such that close objects produce deep, wide suppressions and faraway objects produce only narrow, shallow suppressions (see Fig. 1). In this context, the inhibition depth represents the confidence with which a direction should *not* be selected. A winner-take-all (WTA) process then selects the direction with the maximum evaluation. In this approach, echo amplitude (which can represent the confidence about the existence of an object) could further modulate suppression such that weaker echoes produce weaker suppression, eliminating the problem of determining detection thresholds to decide when to suppress. From the selected direction, we assume that a motor control subsystem (not described) will steer our creature onto the desired heading. This process of simultaneous direction evaluation is similar to other

Manuscript received August 22, 2008; revised November 03, 2008. First published February 18, 2009; current version published November 04, 2009. This work was supported in part by the Air Force Office of Scientific Research under Grant FA95500710446 and in part by the National Institute of Biomedical Imaging and Bioengineering, National Institutes of Health, under Grant R01 EB004750-01. This paper was recommended by Associate Editor A. van Schaik.

The author is with the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland, College Park, MD 20742 USA (e-mail: timmer@umd.edu; timmer@isr.umd.edu).

Digital Object Identifier 10.1109/TCSI.2009.2015597

mean-field-theory approaches to robot navigation and behavior selection [4].

By utilizing a narrow suppression profile, a closely spaced cluster of objects produces deep suppressions with only a modest growth in the width of suppression. Another advantage to the WTA approach is that an open direction on the far side of an obstacle from the direction of travel can be selected (e.g., Fig. 1). Unlike the activity of most neurons recorded in the bat echolocation system which fire in response to echoes, the neurons in our model produce the greatest activity for directions where there are no obstacles (i.e., lack of stimulus). While little is known about how or where obstacle avoidance and goal pursuit is performed in the brain of the bat, studies in other mammals suggest a role for the superior colliculus (SC). We discuss this in further detail in Section IV.

The evaluation function for each direction θ can be described by

$$E(\theta) = E_o + g \cdot \mathrm{e}^{\frac{-(\theta-\theta_g)^2}{\sigma_g^2}} - \sum_{i=1}^{N} \frac{1}{r_i} \cdot \mathrm{e}^{\frac{-(\theta-\theta_i)^2}{\sigma(r_i)^2}}. \qquad (1.1)$$

The first term E_o is a constant bias term to allow the evaluation to remain positive following subtraction by other terms. In general, this term does not need to be constant, but could incorporate information about the desirability of certain directions due to actuation limits. In the spiking neural network in later sections, E_o represents the baseline firing rate. The coefficient g is the amplitude of an additive Gaussian term (with its center at θ_g and width controlled by σ_g) which represents an increase in the desirability due to a known target location. The center of the Gaussian should be steerable with changing goal directions. The index i refers to the N obstacles that suppress the evaluation with a subtractive Gaussian term that is scaled inversely with the range (i.e., distance) r_i of each obstacle. The lateral spread of inhibition was range dependent, i.e., $\sigma(r_i) = \Delta\sigma((r_{\max} - r_i)/r_{\max}) + \sigma_0$ for $r_i < r_{\max}$, where $r_{\max}$ was the maximum sensory range; this range dependence was not a critical feature and was made constant for the circuit described later.

Simulations of this simple algorithm (see Fig. 2 for an example trajectory) have shown that, for light to moderate tree densities, successful parameters for collision avoidance are easy to find. Software implementations of this algorithm have been tested on wheeled mobile robots at speeds of up to 0.5 m/s. To utilize this algorithm on a flying mobile robot, however, we are interested in developing a very low power high-speed VLSI implementation.

B. Spiking Neuron Implementation

Our laboratory's principal interest is the neuromorphic VLSI implementation of the neural circuits of bat echolocation and thus we have developed a neurally plausible model. An obvious neural implementation of the openspace algorithm is to start with a field of neurons that fire tonically to a uniform input bias current. These neurons also receive a steerable Gaussian-shaped excitatory input pattern with the peak centered on the desired goal direction. The obstacle detection system (i.e., sonar)

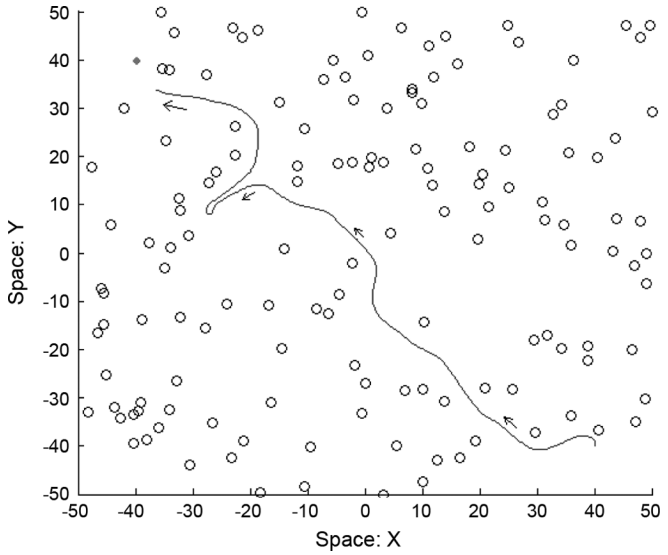

Fig. 2. In this MATLAB simulation of the openspace algorithm, a "bat" moves through an imaginary "forest of trees" toward a target at $(-40, 40)$ and evades collisions. The sensory field of view used in this simulation is 180 ° in front of the bat and out to ten units of range. In this simulation, the target bearing is always known.

projects inhibition onto this field with a strength and width inversely proportional to the range. Thus, the evaluation for each possible direction is represented as the input to each neuron. If each neuron fires monotonically with the strength of its input, the evaluation pattern is observable in the pattern of the neuron firing rate. By incorporating a global inhibitory feedback connection based on the evaluation neuron activity, the well-known WTA or "soft-WTA" function can be implemented on this field of neurons [5].

Although the mean firing rate could be used to represent the evaluation, the interpulse interval also carries the same information but on a shorter timescale. If we had a time-zero reference and simultaneously reset (i.e., strongly inhibit and then release) all neurons, the input currents would be inversely expressed in the spike *latency* across the field of neurons [see Fig. 3(a)]. The neurons which integrate to threshold first are considered to be the winners. Temporal WTA circuits like this have previously been fabricated and used in neuromorphic VLSI contexts (e.g., [6]).

In echolocation, the returning echoes from obstacles arrive at different times according to their range. If the field of neurons is reset at the time of the sonar pulse and echoes trigger long-lasting but weak inhibitory currents, the latency will increase as inhibitory pulses start earlier [see Fig. 3(b)]. The use of such step currents in neural computation is described by Maass [7] and has also been used in a VLSI circuit for visual processing [8].

For a neuron with a membrane capacitance C_{mem}, a spike threshold V_{thresh}, a constant excitatory bias current E_o, and a step-inhibition current I at time t_i, the latency of the spike T is given by

$$T = \frac{C_{\mathrm{mem}} V_{\mathrm{thresh}} - I \cdot t_i}{E_o - I}. \qquad (1.2)$$

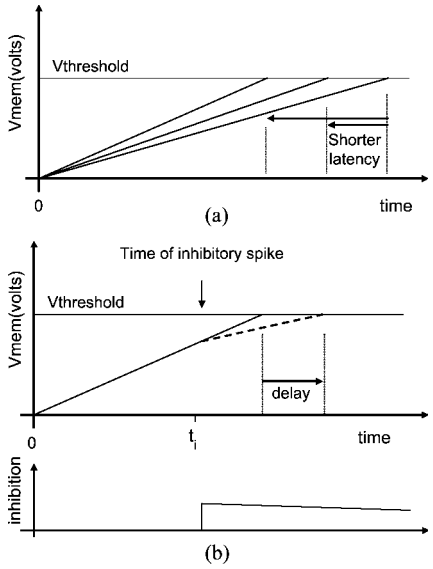

Fig. 3. (a) Increasing the strength of the excitatory inputs to a neuron shortens the latency of the spike following a reset pulse. By determining the neuron that fires first, we find the neuron with the largest average input. (b) Long-lasting inhibitory current delays the spike or prevents firing altogether. Inhibitory currents that start earlier will produce a longer added delay in firing.

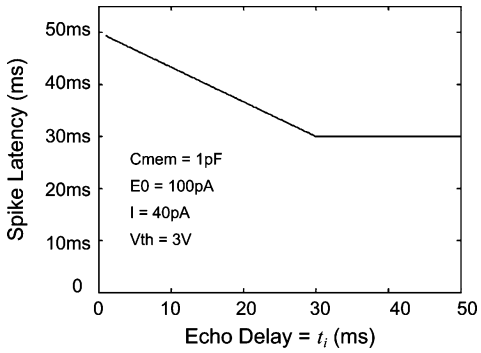

Fig. 4. Example plot of calculated spike latency versus echo delay (target range) for an example parameter setting. Without inhibition, the neuron will fire at 30 ms.

We assume that $E_o > I$ and that $t_i < (C_{\mathrm{mem}} \cdot V_{\mathrm{thresh}})/E_o$. Thus, we obtain an increased latency for closer obstacles *without* explicitly computing the range or increasing the synaptic strength. A plot of this dependence is shown in Fig. 4.

To obtain WTA functionality, evaluation neurons excite a global inhibitory cell that, in turn, fires the global reset pulse. In the extreme case, if the connections between the evaluation neurons and the inhibitory cell are sufficiently strong, the first neuron in the evaluation field to fire will trigger a spike from the inhibitory cell, preventing any other cell from firing and allowing only a single neuron (or a small number of neurons) to fire.

In the sections to follow, we describe the design and testing of a neuromorphic VLSI implementation of a slightly modified version of the openspace algorithm as described previously. Portions of this work have been previously reported in conference proceedings [9].

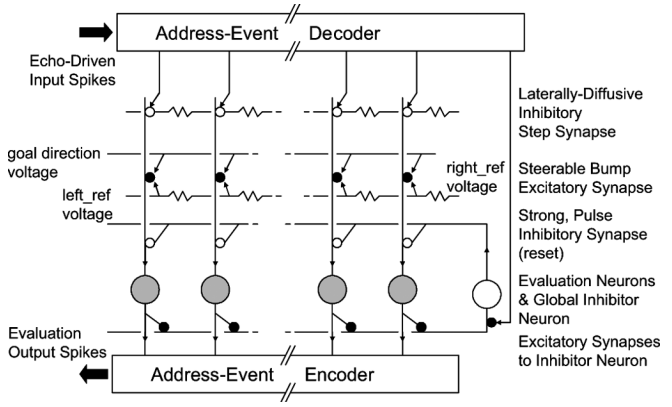

Fig. 5. System block diagram. The *goal-direction voltage* input biases the field of neurons to fire, while the address–event inhibitory spikes (echo triggered) increase the spike latency following a reset pulse. The global inhibitor implements a WTA function once the first set of spikes occurs.

II. Circuits

A. System Design

To facilitate the communication of spikes in and out of the chip, we use a communication protocol known as the "address–event" representation [10]. In this communication system, an asynchronous digital bus provides the address of a target synapse and produces a handshaking pulse, delivering a brief ($\sim$ 1 μs or shorter) digital voltage pulse or spike to the target address. This same signal representation is used to transmit neuron spikes out of the chip. The dynamic inputs to the system (Fig. 5) are the *goal direction* and the echo-triggered address–event spikes corresponding to different obstacle directions. The outputs from the chip are the spikes from the evaluation neuron array. The test chip consists of 25 evaluation neurons and one global inhibitory neuron.

B. Steerable Excitatory "Bump" Bias

The bias term E_o and the steerable Gaussian excitation term in (1.1) are provided as a current to the neuron by the "bump" [11] circuit in Fig. 6. The parameter *fixedbias* controls the excitatory DC current, and *goalv* (provided by the user from off-chip) corresponds to the *goal direction* (from Fig. 5). The drain current of M2 as a function of (*goalv* − *res_R*) is approximately a Gaussian.

C. Spiking AER (Evaluation) Neuron

The integrate-and-fire neuron circuit in Fig. 7 is based on several different neuron designs [12], [13] that utilize the inverter (M1–M3) threshold and decouple the large integration capacitor from *vmem* during the outgoing spike. This allows fast but low-power operation by avoiding the need to charge C1 up to *vdd* and then back down to 0 V in a very short time. This neuron also has a refractory period controlled by the parameter labeled *refr*.

In this circuit, input currents from the synapses and bias sources enter the circuit onto the node labeled *vmem*. If we first consider the neuron in its nonfiring state with $vmem = 0$ V, we can see that the output of the inverter defined by M1 and M3 is at *vdd*. This turns *off* M7 (that provides the positive feed-

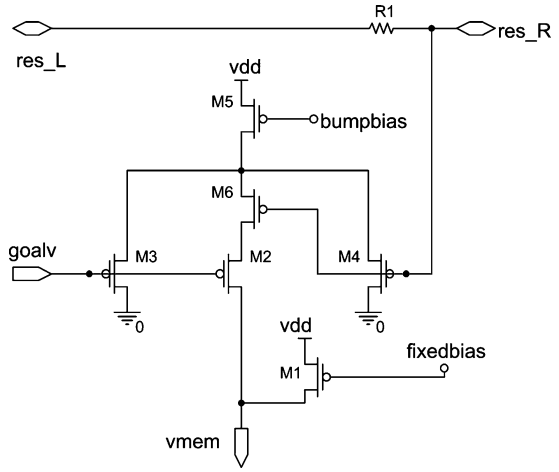

Fig. 6. Resistive ladder (created by connecting *res_R* to the neighboring circuit's *res_L* and different endpoint voltages) gives each neuron a unique reference voltage. A global signal *goalv* is compared to this signal by the bump circuit and produces a maximum current when *goalv* equals the reference voltage. All transistors have a W over L ratio equal to $3.6/3.6$, given in micrometers.

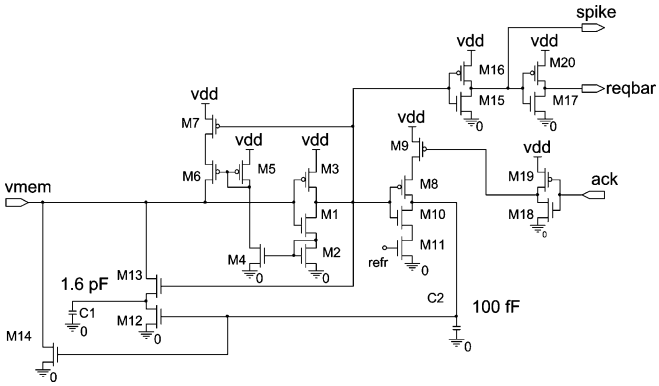

Fig. 7. Neuron circuit receives current from the synapses into the node labeled *vmem* and produces a digital voltage spike at *spike*. The spike is passed to the address–event transmitter system with the handshake signals *reqbar* (active low) and *ack* (active high). The signal *spike* is used to drive the on-chip synapses onto the global inhibitory neuron. C1 was made large to allow long integration times (~ 100 ms). W/L ratios were 1.8/1.8 for M15–M20 (inverters), 3.6/1.8 for M12 and M14, and 3.6/3.6 for other transistors. Measurements are given in micrometers.

back pathway for spike generation) and turns *on* M13 (which connects C1 to *vmem*). The input current charges C1 until the M1/M3 inverter flips state and drives *spike* high. When *spike* goes high, a request (active low) is sent to the address–event representation (AER) arbiter (not shown) through the inverters M15/M16 and M17/M20, the positive feedback current is switched on (M7), and the capacitor C1 is disconnected from the input node. With the capacitor disconnected, the positive feedback can very rapidly charge *vmem*. Once the AER arbiter provides the active-high acknowledge signal *ack*, the capacitor C1 is discharged (M12), and *vmem* is pulled to ground (M14), withdrawing the AER request and ending the spike.

D. Excitatory and Inhibitory Synapses

The synapse circuits used to interconnect the evaluation neurons and the inhibitory neuron are shown in Fig. 8. The evaluation neurons produce short digital voltage spikes with durations

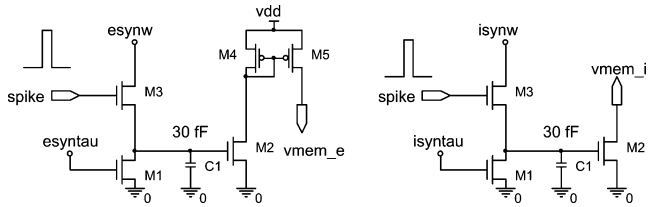

Fig. 8. (Left panel) Excitatory synapse circuit (that connects evaluation neurons to the inhibitory neuron). (Right panel) Inhibitory synapse circuit (that connects the global inhibitory neuron to the evaluation neurons). The input spike drives the gate of M2 to the voltage *esynw* (or *isynw*), and the linear leak by M1 controlled by *esyntau* (or *isyntau*) creates a quick-onset current with an exponential decay that is injected into (or drawn from) the membrane capacitance of the postsynaptic neuron. Transistors M2 and M3 have a $W/L = 3.6/1.2$, and M1 has a $W/L = 1.8/1.8$. Measurements are given in micrometers.

largely defined by the speed of the address–event arbiter circuit (a circuit that determines which neuron is allowed to transmit off the chip first when there are timing conflicts) for transmitting spikes off of the chip.

These voltage spikes charge the upper plate of the capacitor C1 (via transistor M3) to the voltage *esynw*. Transistor M2 operates in saturation and sets the output synaptic current. The parameter *esyntau* defines a constant current in transistor M1 to linearly reduce the gate voltage of M2. If M2 operates in the subthreshold region of operation, the output synaptic current decreases exponentially with time. These synapses can be biased to operate as either as long-duration weak synapses or short-duration, strong synapses (see Section III.B).

E. Diffusive AER Inhibitory Synapse

The obstacle-dependent inhibition begins with an address–event input spike, namely, *aerispikebar* (active low), that triggers a long-lasting inhibitory current pulse (Fig. 3, bottom). In the circuit that generates this pulse (Fig. 9), the voltage on the upper node of C1 is abruptly charged to *vdd* and then slowly discharged by M7 (controlled by *slowtime*) down to 0 V. M5 primarily acts as a switch for the current defined by M4 (controlled by *w_islow*). The inhibitory current is then subtracted from the neuron via the two mirrors M1/M2 and M9/M10. The pFET mirror is connected via nFET transistors to adjacent inhibitory circuits, creating a "diffusor" network [14] that shares inhibitory current with neighboring neurons. The point-spread function for the diffusor is an exponentially decaying function in both directions (left and right) instead of the desired Gaussian in the original algorithm. Although set to zero in these experiments, the *leak* input control allows a constant DC leak current to be applied equally to all neurons in the array.

Our evaluation equation now becomes

$$E(\theta) = E_o + g \cdot \mathrm{e}^{\frac{-(\theta-\theta_g)^2}{\sigma_g^2}} - \sum_{i=1}^{N} \alpha \cdot (r_{\max} - r_i) \cdot \mathrm{e}^{\frac{-|\theta-\theta_i|}{w_o}} \quad (2.1)$$

where the range dependence of the inhibition is linear and that of the lateral spread of inhibition is a decaying exponential instead of a Gaussian.

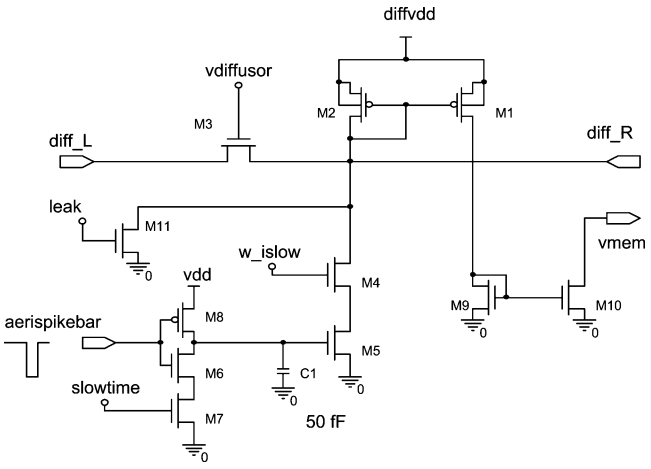

Fig. 9. Echo-triggered inhibition (via address–event) produces laterally spread inhibitory currents through the diffusor transistor M3. Transistors M6 and M8 have a $W/L = 3.6/1.2$, and all others have a $W/L = 3.6/3.6$. Measurements are given in micrometers.

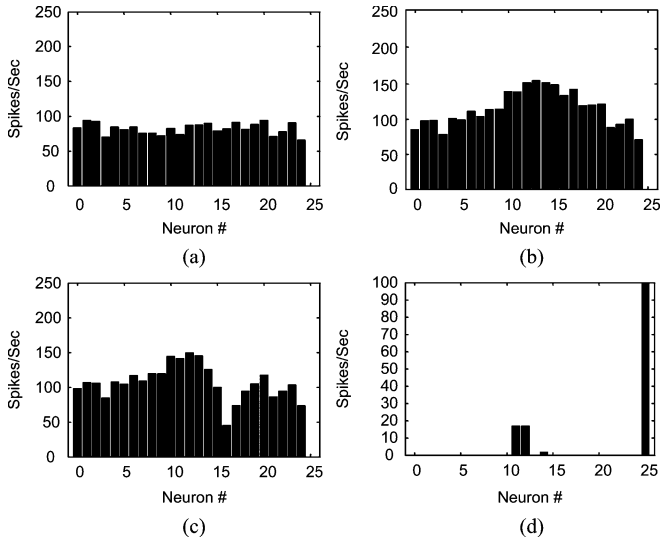

Fig. 10. Operation in mean-rate firing mode. Panels (a)–(c) are generated with the inhibitory neuron disabled. (a) Constant DC current is applied equally to all neurons ($fixedbias = 4.33$ V). (b) Bump circuit provides increased current to the central set of neurons ($bumpbias = 4.27$ V). (c) Inhibitory current profile is applied to neuron #16 ($w\_islow = 0.52$ V and $vdiff = 4.27$ V). (d) Inhibitory neuron is activated ($esyntau = 0.18$ V, $esynw = 0.66$ V, $isynw = 0.55$ V, and $isyntau = 0.17$ V), and a small group of neurons with the strongest activation continues to fire. Neuron #25 is the global inhibitory neuron.

III. Test Results

The chip was fabricated in a commercially available 0.5-μm 2-poly 3-metal CMOS process. This process had an nFET threshold voltage estimated at 0.74 V and a pFET threshold voltage estimated at -0.92 V. This chip consisted of an array of 25 evaluation neurons and one inhibitory neuron. For a sonar system, viewing a full 180 ° in front of the animal would have an angular resolution of 7.2 °, sufficient for the task of collision avoidance given repeated measurements. Testing was performed using both computer-generated signals and, later, with an echolocation system (ultrasonic and air coupled). Input and output spikes were communicated using the address–event protocol; output spikes were captured with a timing resolution of 1 μs.

A. Mean-Firing-Rate Mode

As a basic test of its operation, the mean spiking-rate behavior was evaluated. Spike trains were collected for approximately 10 s, and the mean firing rate was reported. First, a constant identical DC current ($fixedbias = 4.32$ V, see Fig. 6) was provided to each neuron to estimate the functional mismatch in the array. All inhibitory currents were disabled and the neurons had very short refractory periods ($refr = 1.99$ V). As shown in Fig. 10(a), for relatively low firing rates, mismatch is visible. This firing-rate mismatch is likely to be the result of mismatch in the DC current-source transistors and neuron threshold mismatch. In this experiment, the mean firing rate was 81.9 spikes/s with a standard deviation of 8.0 spikes/s. This level of fixed-pattern mismatch can produce a bias toward some neurons to be consistently preferred over others; this will not prevent the successful operation of the system but will affect the sensitivity of the system to respond to subtle changes in the sensory input. Next, the bump circuit, diffusive inhibitory synapses, and the global inhibitory neuron (i.e., WTA behavior) are demonstrated in mean-rate operation [Fig. 10(b)–(d)].

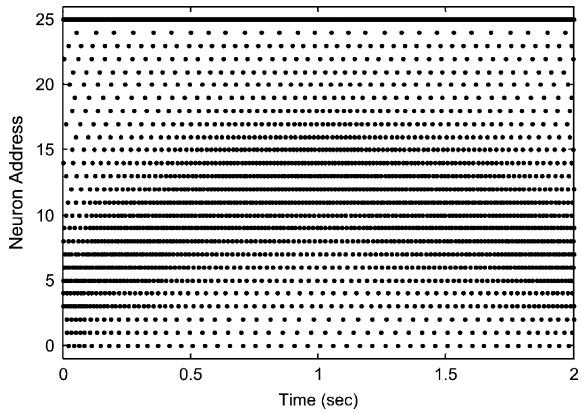

Fig. 11. Neuron spikes versus time. The steerable bump of excitatory bias current is driven by a dynamic input voltage *goalv*. In this case, although the inhibitory neuron is firing vigorously (neuron #25), all of the global feedback inhibitions have been turned off.

To demonstrate the steerable bias currents (bump circuit, Fig. 6), a sinusoid-like voltage was used at the *goalv* input, producing rapid neural firing at changing locations (Fig. 11). All feedback inhibitory currents were turned off in this experiment. For this experiment, $fixedbias = 4.34$ V, $refr = 0.92$ V, $bumpbias = 4.20$ V, $left\_ref\_voltage = 2.31$ V, and $right\_ref\_voltage = 2.70$ V.

In Fig. 12, the global feedback inhibition was activated by the field of neuron using weak excitatory synapses and weak feedback inhibitory synapses. For this experiment, $isynw = 0.60$ V, $isyntau = 0.30$ V, $esynw = 0.60$ V, and $esyntau = 0.96$ V. With this weak inhibition, the global inhibitory neuron does not reset the field of neurons with a single spike, and this WTA behavior operates in a mean-firing-rate regime.

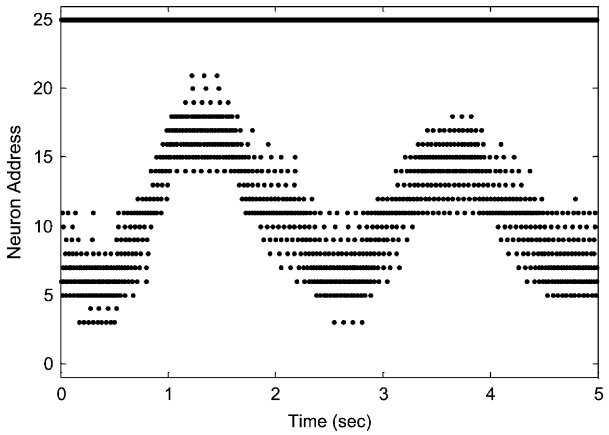

Fig. 12. Neuron spikes versus time. In this case, weak feedback inhibition was activated, shutting off weakly driven neurons. In this example of soft-WTA behavior, the inhibitory neuron fires rapidly, but provides inhibition using long-duration weak inhibitory synaptic currents. Notice that neurons near the center location of the bump generally fire at a higher frequency than neurons at the edges of the bump. Neuron #25 is the global inhibitory neuron.

B. Spike-Latency Mode

As discussed earlier, to properly implement the openspace algorithm in the mean-rate mode of operation, obstacle range information would need to be incorporated into the strength of inhibition applied. In particular, range information would therefore need to be extracted. In the spike-latency mode, however, echo delay can functionally implement our range-dependent inhibition.

Under conditions of strong feedback inhibition (i.e., quick reset of all neurons), interesting firing patterns emerge at short timescales (Fig. 13), revealing that the spike latencies following a reset pulse are related to their input strength, whereas the firing frequency is identical for all neurons that are firing. In the left panel, weak excitation from the neurons in the array to the global inhibitory means that many neurons will fire before the inhibitor fires and resets the whole field of neurons. By changing the strength of the excitatory synapse onto the global inhibitory neuron, the size of the winning group can be modulated by changing how many input spikes it takes to produce a spike. Fig. 13 (right) shows the extreme case where a single spike from a neuron stops the competition by driving the global inhibitory cell with a very strong excitatory synaptic current. An example of a moving target is shown in Fig. 14, where mismatch produces different group sizes.

To show the lateral spread of inhibition and its effectiveness in delaying spike times, we provided a DC bias current ($fixedbias = 4.35$ V) along with a bump profile of current centered on the array ($bumpbias = 4.27$ V). An inhibitory input spike targeted at neuron 12 for three different simulated echo delays results in changes of the first-spike latency (Fig. 15) for neuron 12 and its neighbors. Example first-spike latencies are shown for no inhibition, 25-ms delay, and 35-ms delay.

By taking the difference in first-spike times between the case where no inhibition was applied and the other two delayed inhibition cases, we obtain a measure of the increase in latency produced by our laterally spread inhibitory step current (Fig. 16) targeted at neuron #12.

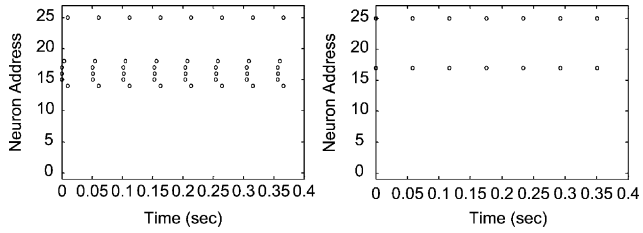

Fig. 13. (Left) Strong excitation of the inhibitory neuron resets the neuron array repeatedly, allowing only the most strongly activated neurons to fire. For this experiment, $isynw = 1.15$ V, $isyntau = 1.10$ V, $esynw = 1.10$ V, and $esyntau = 0.97$ V. Note that the synapse transistors are initially (at the moment of the input spike) operating above threshold. (Right) Further strengthening of the excitation of the inhibitory neuron reduces the size of the winning group to one. For this experiment, *esyntau* was lowered to 0.67 V, increasing the duration (and, thus, efficacy) of the excitatory synaptic current pulse.

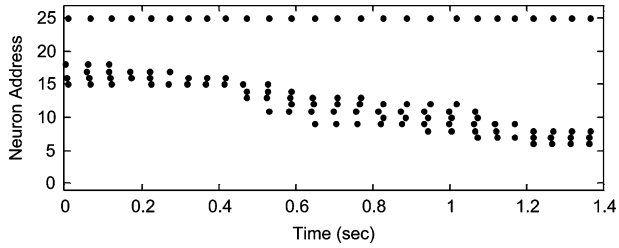

Fig. 14. Example response of the neuron population (spike-latency mode) to a moving goal (*goalv* is swept). Neuron #25 is the global inhibitory neuron.

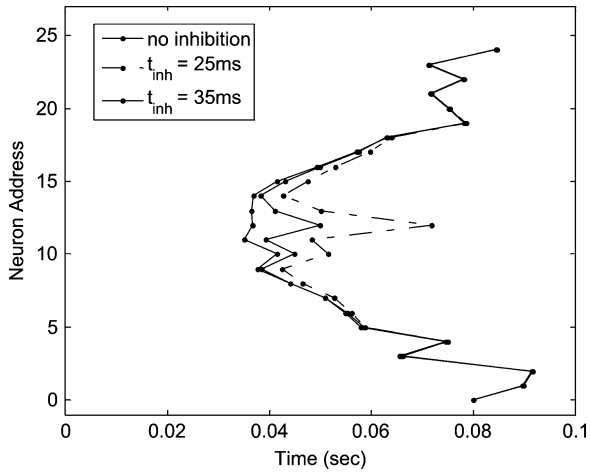

Fig. 15. A single inhibitory spike produces differential effects with its latency. With feedback inhibition off, all neurons integrate to their first spike (other spikes removed). (Short-dashed line with dots) With no echo-triggered inhibition, neurons 4–18 form a prominent arc of first spikes consistent with the bump input currents. (Solid lines with dots) When an inhibitory spike is delivered 35 ms after the reset to neuron 12, the neuron and its neighbors are delayed. (Dashed-dotted line with dots) If the inhibitory spike at neuron 12 occurs earlier (simulating a closer obstacle) at 25 ms, the effect of the inhibition is greater, and the neurons fire even later.

C. Sonar-Triggered Behavior

To demonstrate the use of the chip with live sonar-triggered signals, we connected a simple custom binaural sonar system [15] that produces a spike signal to the openspace chip for resetting all of the evaluation neurons at the time of the outgoing sonar emission and generates a spike at the time of an echo with an address that is proportional to the direction from which it arrived. An important feature of this sonar (and any active sensory

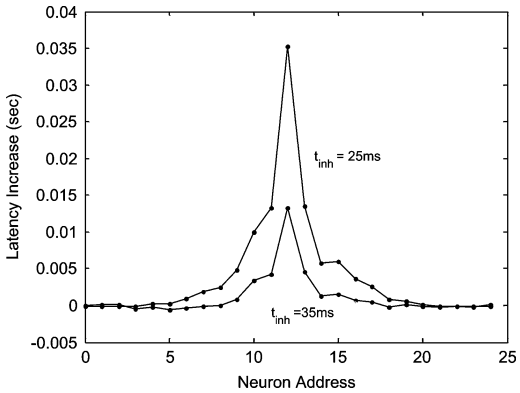

Fig. 16. Neighborhood size and amount of latency increase are larger for inhibitory spikes that arrive earlier.

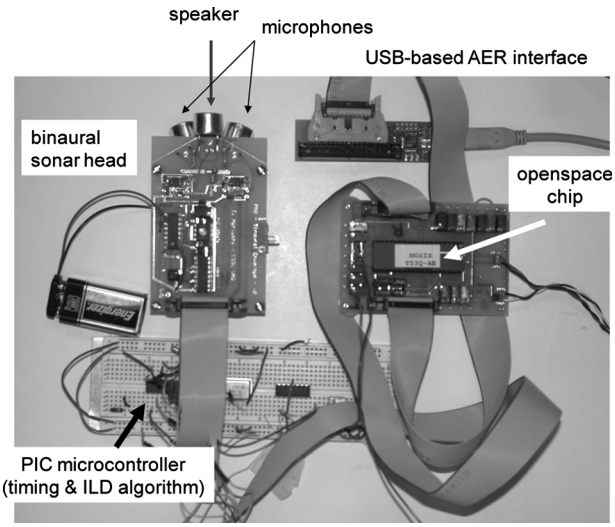

Fig. 17. Photograph of the test system with binaural sonar, the angle calculation on a PIC microcontroller, and the openspace chip.

system) is the suppression of sensory-driven spikes in response to the outgoing sound.

Fig. 17 shows a photo of the demonstration system with the binaural sonar head triggered by an onboard microcontroller (PIC12CE674, PIC Microchip Inc.) that generates the transmission and initiates all timing in the system. The two analog ultrasonic microphone signals are amplified and sent to another PIC microcontroller (PIC18F2620) on which the direction is estimated. Address–event spikes representing the direction of the sound are generated as echoes arrive and are sent to the openspace chip to produce the desired inhibition pattern.

It should be noted that, in testing this chip, the time constant of the global inhibitory synapses (onto each neuron) was actively modulated to simulate two different inhibitory synapses: one that is strong and short in duration to reset all of the neurons and synchronize the race-to-first-spike and one that is strong and longer in duration that stops the race-to-first-spike and holds all of the neurons quiet.

Using the sonar to trigger inhibition due to an obstacle, Fig. 18 shows the latency-delaying effect of the step inhibition on the rise of the *vmem* node voltage.

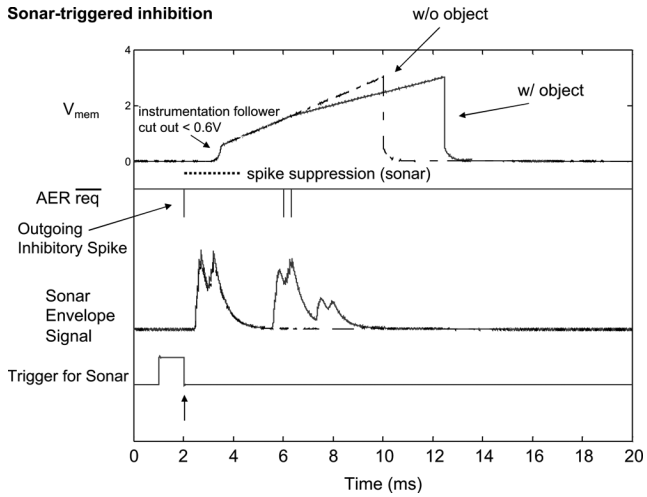

Fig. 18. Demonstration of the increase in latency due to a sonar-triggered inhibition. The bottom trace indicates the digital signal used to trigger the sonar emission. The sonar envelope signal compares examples from two different emissions: (Dashed line) "Without object' and (solid line) 'with object.' A suppression signal (see dotted line, top trace, *vmem*) produced by the sonar transmitter suppresses any neural response to the outgoing sound. When an object is detected (see sonar envelope, second bump), two sonar-derived spikes are sent through the AER interface system to inhibit the neuron. The resulting inhibition (starting at the time of the object echo) delays the time when *vmem* reaches the threshold voltage and resets to 0 V.

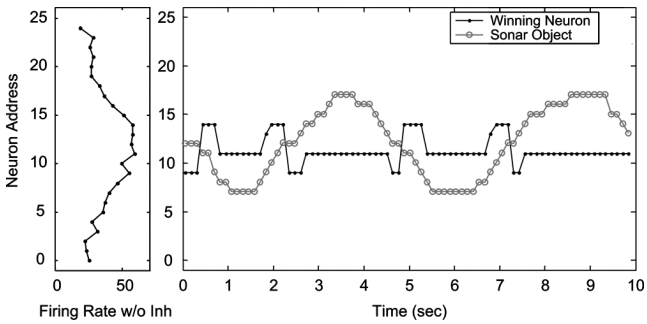

Fig. 19. WTA behavior in the network in response to an object oscillating in front of the sonar. In this demonstration, the bump circuit is biased to select a central location. The left panel shows the firing rates of the neurons without the WTA inhibitory feedback activated. The peak firing rate occurs at neuron #11. The sonar system detects the echo, determines the direction of the object (fixed distance), and sends an AER spike to the diffusive inhibitory synapse of the appropriate neuron approximately 25 ms after the outgoing pulse. The object is moved back and forth to interfere with the selection of the winning neuron. As the (red circles) inhibitory activity generated by the object slides across the array, the (black dots) winning location is, at first, deflected away from the inhibition, but as it passes, the winning location snaps to the opposite side of the inhibition.

To demonstrate the sonar and chip working together to select a winning openspace direction, an object was moved side to side in front of the sonar at approximately 75-cm distance. It was moved specifically to occlude the goal direction specified by the bump circuit. The binaural sonar calculates the echo direction based on intensity comparisons and provides direction-specific inhibition to the neuron array (using the AER interface) to delay spikes in that direction. Fig. 19 shows the effect of a moving sonar obstacle on the winning location as it oscillates in front of the sonar. In this case, the goal direction was fixed, and global inhibition was activated for the full WTA behavior.

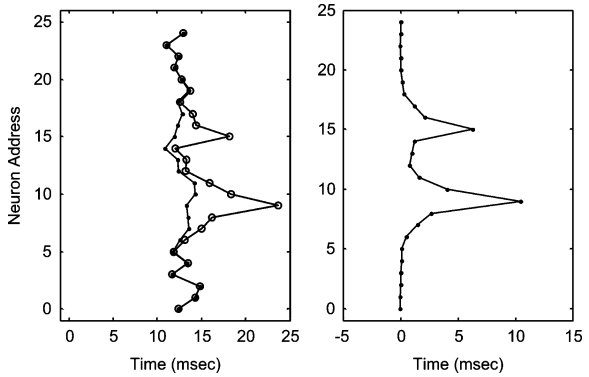

Fig. 20. In this example, two *sonar* targets are presented at 40 cm (neuron #9) and 55 cm (neuron #15), and the resulting latencies in the first spike are compared. The same dc bias current was applied to all neurons to allow a better comparison. (Left) Dark solid line is the response *with* targets, and the thin dotted line is the response latency *without* targets. (Right) Latency difference. The closer target (at neuron #9) provides more inhibition.

In our final demonstration (Fig. 20), the sonar system was used to provide multiple inhibitory events corresponding to multiple targets at different ranges. We demonstrate that closer targets produce deeper and wider inhibition and that multiple obstacles can be integrated.

Although the source of goal-direction information was not discussed, it is assumed that another system (presumably based on information coming from the echolocation system) would determine goal directions. In bats, goals would typically include detected insects (temporally modulated echoes), other bats, navigational waypoints or flyways, and its home or resting roosts.

IV. Discussion

The use of input spike timing to modulate the efficacy of a synaptic connection in an ongoing computation can be an effective mechanism that does not rely on the modulation of synaptic strength or on increased spike rates. This type of candidate mechanism for spike-based neural computation has potential in systems where a race-to-first-spike is a fundamental aspect of the problem due to a globally synchronizing signal. We have found a natural match of this computational mechanism with the echolocation-based openspace algorithm proposed earlier in this paper.

An important issue that is commonly faced when using WTA systems is the effect of device mismatch. Often overlooked in theoretical or software-based models of these networks, (gain and/or offset) mismatch can make it difficult to compare inputs "fairly" to determine the input with the largest input. While the use of synaptic modification in biological systems seems like a likely choice for compensation, the use of *populations* of neurons that display symmetrical zero-mean mismatch distributions may also be a biologically plausible solution. In this case, neurons within a population would have lateral excitatory connections within the group. This type of spiking neuron network with lateral excitation and global inhibitory feedback is well known and has been constructed in mean-rate coded neuromorphic VLSI in several different contexts (e.g., [16] and [17]). In the race-to-first-spike WTA, the outcome depends more heavily on individual neuron characteristics because the averaging effects from the neighborhood are not available. Toward this end, a floating-gate calibration approach seems most appropriate for the analog VLSI system described in this paper.

While our specific application is focused on bat echolocation where the echo delay can directly generate desired inhibitory scaling effects, the basic openspace algorithm in the mean-rate mode of operation can also be used to merge information from multiple sources of spatial information (vision, somatosensory, whiskers, etc.). In the brains of mammals, an obvious candidate area for merging multiple sensory modalities for the purpose of orienting the animal is the superior colliculus (*SC*). Studied in many different animals, the SC has a well-documented role in motor control of the orienting behavior of the eye, pinnae, head, and body, as well as a lesser studied role in avoidance behavior [18]. While little is known about the role of the SC in collision avoidance or goal pursuit during locomotion, eye-movement studies in primates suggest that the SC *is* involved in target selection when multiple choices are presented and that inhibitory circuits are involved in this selection [19]. The SC likely operates in concert with the frontal cortex to integrate fast reflexive behavior (SC) with slower volitional control of navigation (frontal cortex).

Acknowledgment

The author would like to thank T. Delbrück for the USB-based AER capture system critical for obtaining the data for this paper, S. Liu and M. Oster for the helpful discussions, and P. S. Krishnaprasad for the encouragement and advice throughout this project.

References

[1] W. H. Huang, B. R. Fajen, J. R. Fink, and W. H. Warren, "Visual navigation and obstacle avoidance using a steering potential function," *Robot. Auton. Syst.*, vol. 54, no. 4, pp. 288–299, Apr. 2006.
[2] B. Hamner, S. Singh, and S. Scherer, "Learning obstacle avoidance parameters from operator behavior," *J. Field Robot.*, vol. 23, no. 11/12, pp. 1037–1058, Nov./Dec. 2006.
[3] J. Borenstein and Y. Koren, "Histogramic in-motion mapping for mobile robot obstacle avoidance," *IEEE Trans. Robot. Autom.*, vol. 7, no. 4, pp. 535–539, Aug. 1991.
[4] W. Erlhagen and G. Schoener, "Dynamic field theory of movement preparation," *Psychol. Rev.*, vol. 109, no. 3, pp. 545–572, 2002.
[5] D. Hansel and H. Sompolinsky, "Modeling feature selectivity in local cortical circuits," in *Methods in Neuronal Modeling: From Synapse to Networks*, 2nd ed. Cambridge, MA: MIT Press, 1998, pp. 499–567.
[6] J. P. Abrahamsen, P. Hafliger, and T. S. Lande, "A time domain winner-take-all network of integrate-and-fire neurons," in *Proc. Int. Symp. Circuits Syst.*, 2004, pp. V-361–V-364.
[7] W. Maass, *Pulsed Neural Networks*. Cambridge, MA: MIT Press, 1998.
[8] V. Ravinuthula and J. G. Harris, "Time-based arithmetic using step functions," in *Proc. Int. Symp. Circuits Syst.*, 2004, pp. I-305–I-308.
[9] T. Horiuchi, "A neural model for sonar-based navigation in obstacle fields," in *Proc. ISCAS*, Kos, Greece, May 2006, pp. 4543–4546.
[10] K. A. Boahen, "Retinomorphic vision systems II: Communication channel design," in *Proc. Int. Symp. Circuits Syst.*, 1996, pp. 14–17.
[11] T. Delbruck, "'Bump' circuits for computing similarity and dissimilarity of analog voltages," in *Proc. Int. Joint Conf. Neural Netw.*, 1991, pp. I-475–I-479.
[12] E. Culurciello, R. Etienne-Cummings, and K. A. Boahen, "A biomorphic digital image sensor," *IEEE J. Solid-State Circuits*, vol. 38, no. 2, pp. 281–294, Feb. 2003.
[13] G. Indiveri, E. Chicca, and R. J. Douglas, "A VLSI reconfigurable network of integrate-and-fire neurons with spike-based learning synapses," in *Proc. ESANN*, 2004, pp. 405–410.
[14] A. G. Andreou and K. A. Boahen, "A 48,000 pixel, 590,000 transistor silicon retina in current-mode subthreshold CMOS," in *Proc. 37th Midwest Symp. Circuits Syst.*, 1994, pp. I-97–I-102.

[15] R. Z. Shi and T. K. Horiuchi, "A neuromorphic VLSI model of bat interaural level difference processing for azimuthal echolocation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 54, no. 1, pp. 74–88, Jan. 2007.
[16] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature*, vol. 405, no. 6789, pp. 947–951, Jun. 2000.
[17] E. Chicca, A. M. Whatley, P. Lichtsteiner, V. Dante, T. Delbruck, P. Del Giudice, R. J. Douglas, and G. Indiveri, "A multichip pulse-based neuromorphic infrastructure and its application to a model of orientation selectivity," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 54, no. 5, pp. 981–993, May 2007.
[18] B. E. Stein and M. A. Meredith, *The Merging of the Senses*. Cambridge, MA: MIT Press, 1992.
[19] R. M. McPeek and E. L. Keller, "Deficits in saccade target selection after inactivation of superior colliculus," *Nat. Neurosci.*, vol. 7, no. 7, pp. 757–763, Jul. 2004.

Timothy K. Horiuchi (M'89) received the B.S. degree in electrical engineering and the Ph.D. degree in computation and neural systems from the California Institute of Technology, Pasadena, in 1989 and 1997, respectively.

He was a Postdoctoral Scholar with the Johns Hopkins University, Baltimore, MD, until he moved to the University of Maryland, College Park, in 1999, where he is currently an Associate Professor with the Department of Electrical and Computer Engineering and the Institute for Systems Research. His main research interests are in the design and fabrication of neuromorphic VLSI circuits and the implementation of neural computation in silicon. His primary focus has been on the modeling of spike-based models of the auditory system and higher-order functions of behavior in the echolocating bat.