

A Neural Model for Sonar-Based Navigation in Obstacle Fields

Timothy K. Horiuchi

Electrical and Computer Engineering Department,
Institute for Systems Research & the Neuroscience and Cognitive Science Program
University of Maryland, College Park, MD 20742, timmer@isr.umd.edu

Abstract — The rapid control of sonar-guided vehicles through obstacle fields has been a goal of robotics for decades. How sensory data is represented strongly affects how obstacles and goal information can be combined to select a direction of travel. While typical approaches combine attractive and repulsive effects to directly determine steering, we are investigating an algorithm that evaluates multiple directions simultaneously followed by a winner-take-all (WTA) function which then guides steering. In this paper we describe a neuromorphic VLSI implementation of this algorithm using the inherent echo delay to create a range-dependent gain in a ‘race-to-first-spike’ neural WTA circuit. The chip was fabricated in a commercially-available 0.5 μ m CMOS process and in this paper we present preliminary test results.

Index Terms— step inhibition, spike-timing, winner-take-all, bat echolocation, multiple obstacles, robot navigation.

I. INTRODUCTION

Echolocating bats are capable of flying through dense forests in complete darkness during their hunt for flying insects and other prey. How biological systems can turn the storm of sensory information into short-term motion plans amidst 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 short-term obstacle avoidance problem have utilized the summation of repulsive or attractive forces and torques to steer a particular creature (e.g., [1]). While this approach has proven to be quite successful in many cases, there are a few drawbacks to consider. An extended object, if detected as two closely-spaced objects, could produce twice the repulsive force. This force-field approach can also produce trajectories between closely-spaced objects, such as in room corners. Part of the problem is that an obstacle that produces a repulsive force (or torque) attempts to turn the creature in a particular direction, whereas obstacles should only inform the creature of where it should *not* go.

A. The ‘Open Space’ Algorithm

The approach we are taking is a risk-minimization view of

navigation, using the sonar system to evaluate the desirability of different directions of travel simultaneously and selecting the direction with the highest evaluation.

The evaluation process begins with a field of evaluation units that receive an initial evaluation that represents the prior assumptions about the desirability of a given direction. This can incorporate information about actuation limits, 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). A winner-take-all (WTA) process then selects the direction with the maximum evaluation. In this approach, echo strength could further modulate suppression such that weaker echoes produced weaker suppression, eliminating the problem of detection thresholds. From the selected direction, we assume that a motor control subsystem will steer our creature onto the desired heading. This process of simultaneous direction evaluation is similar to other mean-field-theory approaches to robot navigation [2].

With this approach, closely-spaced objects produce only deep suppressions, but not necessarily very wide suppressions in the evaluation function. In addition, 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 neurons recorded in the bat echolocation system, the neurons in our model produce the greatest activity when there is only open space.

The evaluation function for each direction θ can be described by the equation:

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

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. The coefficient g is the amplitude of an additive Gaussian term 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, r_i . Notice that the σ parameter is also

a function of range.

B. A Spiking Neuron Implementation

While there are many implementation possibilities, we are interested in dedicated VLSI approaches that could operate in real-time on a model aircraft. Because our laboratory’s principal interest is in the neural implementation of navigation in bats, we have considered a neural model.

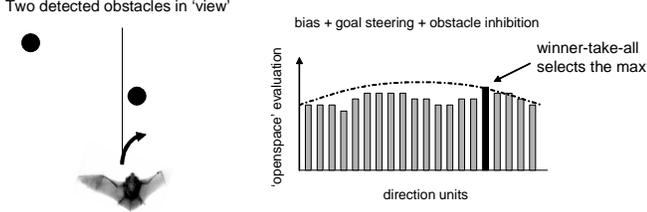


Figure 1. Left: An echolocating bat that is attempting to fly directly forward detects two obstacles (filled circles). Right: The evaluation pattern consists of a constant plus a wide, low-amplitude Gaussian with two dips created by the suppression from the two obstacles. A WTA function selects the direction with the highest evaluation. 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>

An obvious neural implementation of the “open space” algorithm is to use a field of neurons that fire tonically to a uniform input bias. 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) 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 neuron spiking. By incorporating a global inhibitory feedback connection, the well-known winner-take-all (WTA) function can be implemented on this field of neurons [3].

Although the mean firing rate could be used to represent the evaluation, the inter-pulse-interval also carries the information, but on a shorter timescale. If we had a time-zero reference and simultaneously reset (i.e., strongly inhibit) all neurons, the input currents would be inversely expressed in the spike latency across the field of neurons (see Fig. 2a). The neurons which integrate to threshold first are considered to be the winners. Temporal WTA circuits like this have recently been fabricated [3].

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 (see Fig. 2b), the latency will increase as inhibitory pulses start earlier. The use of such step currents in neural computation is described by Maass [4] and has recently been used in a VLSI circuit for visual processing [5].

For a neuron with a membrane capacitance C_{mem} , a spike threshold V_{thresh} , that receives a constant excitatory bias current E_o and a step-inhibition current at time t_i , the latency of the spike, T , is given by Equation 1.2.

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

We assume that $E_o > I$ and that $t_i < \frac{C_{mem} \cdot V_{thresh}}{E_o}$. Thus we

obtain increased latency for closer obstacles *without* explicitly computing the range or increasing the synaptic strength.

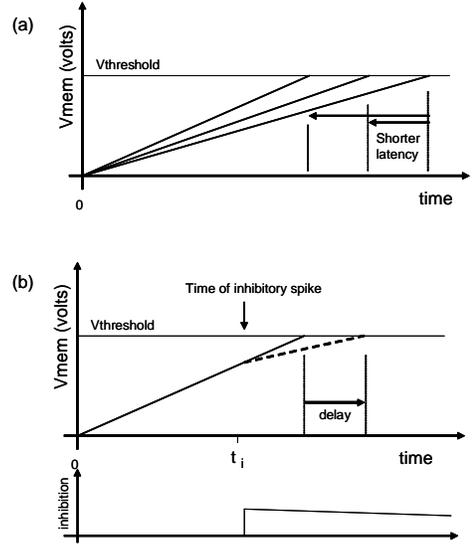


Figure 2. (a) Increasing the strength of the excitatory inputs to a neuron shortens the latency of the spike following a reset pulse. By determining the neuron that fires first, we find the neuron with the largest average input. (b) A long-lasting inhibitory current delays the spike or prevents firing altogether. Inhibitory currents that start earlier will produce a longer added delay in firing.

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

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 [6]. In this system, an asynchronous digital bus provides the address of a target synapse and produces a handshaking pulse, delivering a brief ($\sim 1\mu s$) voltage spike to the target address. This same system is used to transmit neuron spikes out of the chip. The inputs to the

system (Fig. 3) are the *goal direction* and the echo-triggered address-event spikes corresponding to different obstacle directions. The outputs are the spikes from the evaluation neuron array. The test chip consists of 25 evaluation neurons and one global inhibitory neuron.

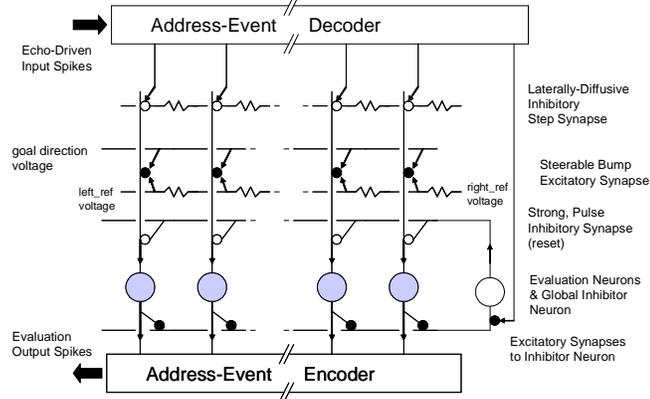


Figure 3. System Block Diagram. The *goal direction* input biases the field of neurons to fire, while the address-event inhibitory spikes (echo-triggered) increase the spike latency following a reset pulse. The global inhibitor implements a WTA function once the first set of spikes occur.

B. A Steerable Excitatory “Bump” Bias

The bias term E_o and the steerable Gaussian excitation term in equation 1.1 are provided as a current to the neuron by the “bump” [7] circuit in Fig. 4. The parameter *fixedbias* controls the excitatory DC current and *goalv* corresponds to the *goal direction*. The drain current of M2 as a function of (*goalv* – *res_R*) is approximately a Gaussian.

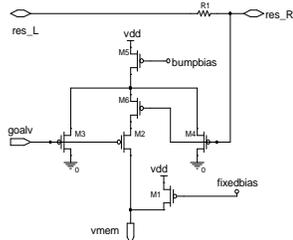


Figure 4. A resistive ladder (created by connecting *res_R* to the neighboring circuit’s *res_L* and different endpoint voltages) gives each neuron a unique reference voltage. A global signal, *goalv*, is compared to this signal by the “bump” circuit and produces a maximum current when *goalv* equals the reference voltage.

C. The Spiking AER Neuron

The integrate-and-fire neuron circuit in Fig. 5 is based on a neuron design [8] that utilizes the inverter (M1-M3) threshold and decouples the main integration capacitor during the spike allowing fast, but low-power operation. This neuron also has a refractory period controlled by the parameter *refr*.

D. Excitatory and Inhibitory Synapses

The synapse circuits used to interconnect the evaluation neurons and the inhibitory neuron are shown in Fig. 6. For these synapses, the input signal is a logical high voltage pulse.

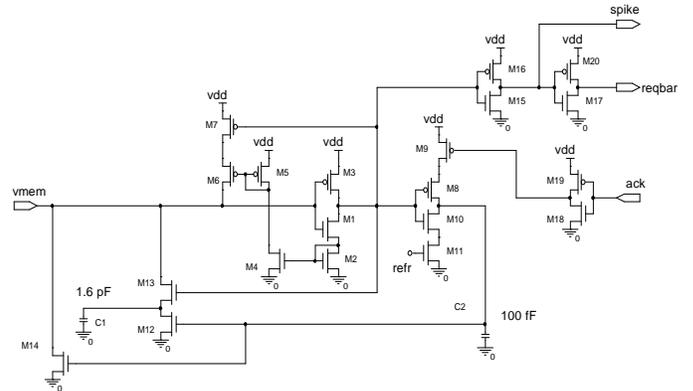


Figure 5. The neuron circuit receives current from the synapses through the node labeled *vmem* and produces a digital voltage spike at *slope*. The spike is passed to the address-event transmitter system with the handshake signals *reqbar* and *ack*. The signal *slope* is used to drive on-chip synapses. C1 was made large to allow long time constants (~100 ms).

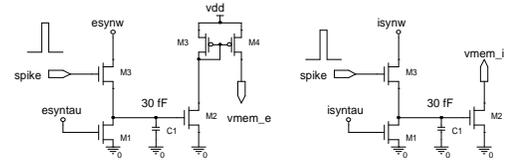


Figure 6. Excitatory synapse circuit (left) and inhibitory synapse circuit (right). The input spike drives the gate of M2 to the voltage *esynw* (or *isynw*) and the linear leak by M1 controlled by *esyntau* (or *isyntau*) creates a quick-onset current with an exponential decay that is injected into (or drawn from) the membrane capacitance of the postsynaptic neuron.

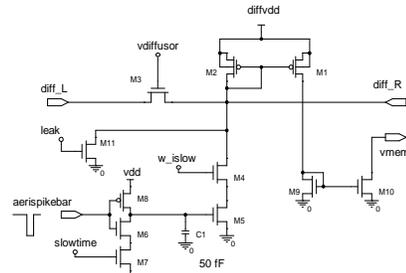


Figure 7. Echo-triggered inhibition (via address-event) produces laterally-spread inhibitory currents through the diffusor transistor M3.

E. A Diffusive AER Inhibitory Synapse

The obstacle-dependent inhibition begins with an address-event spike input, *aerspikebar*, that triggers a long-lasting inhibitory current step (Fig. 7). The voltage at C1 is abruptly charged to V_{dd} and then slowly discharged by M7 (controlled by *slowtime*) down to 0V. M5 acts as a switch for the current defined by M4 (controlled by *w_islow*). The inhibitory current is then subtracted from the neuron. The pMOS mirror is connected via nMOS transistors to adjacent inhibitory circuits, creating a “diffusor” network [9] that shares inhibitory current with neighboring neurons. The point-spread function for the diffusor is an exponentially-decaying function in each direction, instead of the desired Gaussian.

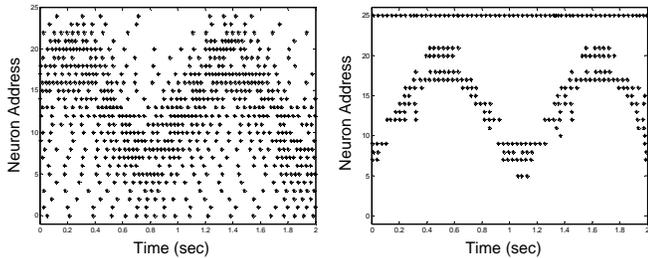


Figure 8. Neuron spikes vs. time. The steerable “bump” of excitatory bias current is driven by a dynamic input voltage, *goalv*. Neuron spike rasters show the firing of different neurons over time. (Left) inhibition turned off, (Right) weak feedback inhibition activated.

III. TESTING RESULTS

The chip was fabricated in a commercially-available 0.5 μm , 2-poly, 3-metal CMOS process. For testing, we used a software-generated pattern of address-event spikes as input to the chip and captured the address-event output spikes with timing resolution of 1 μs . To first demonstrate the steerable bias currents, a sinusoidal voltage was used at the *goalv* input, producing rapid firing at changing locations (Fig. 8, left). In Fig. 8 (right), feedback inhibition was activated by using weak excitatory synapses and weak feedback inhibitory synapses. With weak inhibition, the global inhibitor does not reset the field of neurons and this WTA behavior operates in a mean firing rate regime.

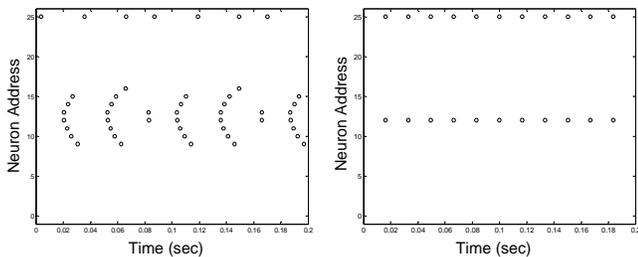


Figure 9. (Left) Strong excitation of the inhibitory neuron resets the neuron field repeatedly, allowing only the most activated neurons to fire. (Right) Further strengthening of the excitation of the inhibitor reduces the size of the winning group to one.

Under conditions of strong feedback inhibition, interesting firing patterns emerge at short time scales (Fig. 9), revealing that the spike latencies following a reset pulse are related to their input strength. By changing the strength of the excitatory synapse onto the global inhibitor the size of the winning group can be modulated. Fig. 10 shows that, following the sonar-triggered reset of the neurons, the timing of echo-triggered inhibition can indeed control the race-to-first-spike competition. The diffusive inhibitory synapse is also shown to create a neighborhood of suppression.

IV. DISCUSSION

The use of input spike timing to modulate the efficacy of a

connection in an on-going computation can be an effective mechanism that does not rely on the modulation of synaptic strength or on increased spike rates. We have found a natural match of this computational mechanism with the echolocation-based “open space” algorithm proposed above.

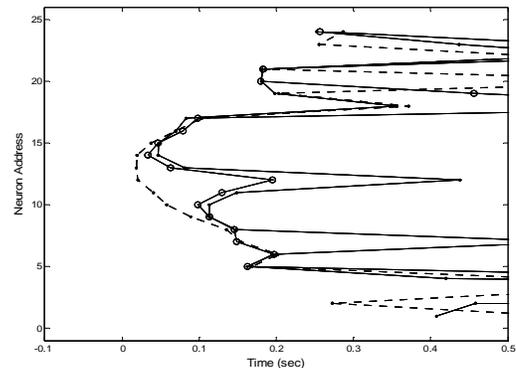


Figure 10. With feedback inhibition off, all neurons integrate to their first spike (other spikes removed). With no echo-triggered inhibition (dots with dotted lines) neurons 8 through 17 form an arc of first spikes consistent with the “bump” input currents. When an inhibitory spike is delivered 40 ms after the reset (circles with solid lines) to neuron 12, the neuron and its neighbors are delayed. If the inhibitory spike at neuron 12 occurs earlier (i.e., a closer obstacle) at 25 ms (dots with solid lines), the effect of the inhibition is greater and the neurons fire even later.

ACKNOWLEDGMENT

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

REFERENCES

- [1] B. R. Fajen and W. H. Warren, "Behavioral dynamics of steering, obstacle avoidance, and route selection," *J Exp Psychol Hum Percept Perform*, vol. 29, pp. 343-62, 2003.
- [2] W. Erlhagen and G. Schoner, "Dynamic field theory of movement preparation," *Psychological Review*, vol. 109, pp. 545-572, 2002.
- [3] J. P. Abrahamsen, P. Hafliger, and T. S. Lande, "A time domain winner-take-all network of integrate-and-fire neurons," *Proc. Intl Symp Circ and Sys*, V-361-V-364, 2004.
- [4] W. Maass, *Pulsed Neural Networks*. Cambridge: MIT Press, 1998.
- [5] V. Ravinuthula and J. G. Harris, "Time-based arithmetic using step functions," *Proc Intl Symp Circ and Sys*, I-305-308, 2004.
- [6] K. A. Boahen, "Retinomorphic Vision Systems II: Communication Channel Design," *Intl Symp Circ and Sys*, Supplement, 9-14, 1996.
- [7] T. Delbrück, "'Bump' circuits for computing similarity and dissimilarity of analog voltages," *Intl Joint Conf Neur Networks*, I-475-479, 1991.
- [8] E. Culurciello, R. Etienne-Cummings, and K. A. Boahen, "A biomorphic digital image sensor," *IEEE Journal of Solid-State Circuits*, vol. 38, pp. 281-294, 2003.
- [9] A. G. Andreou and K. A. Boahen, "A 48,000 pixel, 590,000 transistor silicon retina in current-mode subthreshold CMOS," *Proc 37th Midwest Symp Circ and Sys*, I-97-102, 1994.