

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE Final		3. DATES COVERED (From - To) 03/01/04 - 02/28/07	
4. TITLE AND SUBTITLE Neuromorphic VLSI-based Bat Echolocation for Micro-aerial Vehicle Guidance				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA95500410130	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Timothy K. Horiuchi and P. S. Krishnaprasad				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Systems Research, University of Maryland, College Park, MD 20742 A. V. Williams Bldg, Paint Branch Drive				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research (AFOSR) Dr. Willard Larkin AFOSR/NL AFOSR/NL Suite 325 875 N. Randolph St Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT N/A Approved for public release. Distribution is unlimited AFRL-SR-AR-TR-07-0483					
13. SUPPLEMENTARY NOTES N/A					
14. ABSTRACT We summarize the state of the various projects our laboratories have pursued during the course of this support. This includes multiple efforts related to a VLSI-based echolocation system being developed in one of our laboratories from algorithm development, bat flight data analysis, to VLSI circuit design and testing of these algorithms. We have pursued investigations into the spike-based implementation of the interaural-intensity processing regions of the bat brainstem and midbrain with our VLSI modeling of the lateral superior olive (LSO), the dorsal nucleus of the lateral lemniscus (DNLL), and the inferior colliculus (IC). We summarize an exciting new perspective that neurobiologically-realistic, conductance-based synaptic integration is better suited for the particular computations we are seeking. In collaborative work with Cynthia Moss, we have been analyzing previously-captured bat flight data to understand the strategy the bat appears to use in the capture of moving targets. The analysis and computations underlying this problem is closely aligned to the techniques used in describing a control law for explaining a hunting behavior ('motion camouflage') observed in the visually-guided dragonfly. We continue to develop new circuits for an ultrasonic cochlea and have uncovered interesting new issues in our choice for representing the intensity of signals. We have just finished testing the first chip version of an echo-timing-based algorithm ('openspace') for sonar-guided navigation amidst multiple obstacles.					
15. SUBJECT TERMS Neuromorphic VLSI, bat echolocation, biologically-inspired algorithms & electronic circuits, geometric control theory, pursuit algorithm, 3D formations, motion camouflage					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON Timothy K. Horiuchi
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) (301) 405-7412

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

20071108151

Final Performance Report for AFOSR (Grant # FA95500410130)
Spike Neuromorphic VLSI-based Bat Echolocation
for Micro-Aerial Vehicle Guidance

PI: Timothy K. Horiuchi, University of Maryland, College Park
Co-PI: P. S. Krishnaprasad, University of Maryland, College Park
(March 31, 2007)

Summary:


We summarize the state of the various projects our laboratories have pursued during the course of this support. This includes multiple efforts related to a VLSI-based echolocation system being developed in one of our laboratories from algorithm development, bat flight data analysis, to VLSI circuit design and testing of these algorithms. All efforts are focused on the understanding of the bat echolocation system and circuits that can be used to implement such a system. We have pursued investigations into the spike-based implementation of the interaural-intensity processing regions of the bat brainstem and midbrain with our VLSI modeling of the lateral superior olive (LSO), the dorsal nucleus of the lateral lemniscus (DNLL), and the inferior colliculus (IC). We summarize an exciting new perspective that neurobiologically-realistic, conductance-based synaptic integration is better suited for the particular computations we are seeking in comparison to the current-injection synapses we had been using (section 1). In collaborative work with Cynthia Moss and our co-advised graduate student (now postdoctoral scholar) Kaushik Ghose, we have been analyzing previously-captured bat flight data to understand the strategy the bat appears to use in the capture of moving targets (section 2). The analysis and computations underlying this problem is closely aligned to the techniques used in describing a control law for explaining a hunting behavior ('motion camouflage') observed in the visually-guided dragonfly (section 3). We have further developed our understanding of interactions between the brainstem and midbrain auditory areas and how inhibition can produce unexplained response changes (section 4). We continue to develop new circuits for an ultrasonic cochlea and have uncovered interesting new issues in our choice for representing the intensity of signals (section 5). We have new results in a project that investigates a constraint-minimization circuit which we are hoping to soon apply to the echolocation domain ('visual horizon chip') (section 6). We have just finished testing the first chip version of an echo-timing-based algorithm ('openspace') for sonar-guided navigation amidst multiple obstacles (section 7). Finally, in an effort to provide ready-to-use sensors for the Telluride Neuromorphic Engineering Workshop we have also developed a relatively low-cost sonar experimentation board for a project group focused on echolocation.

In what follows, we have also included information on work immediately prior to the reporting period for the sake of completeness and coherence of the narrative.

1. A Synaptic Conductance-based Model of Interaural Level Differences

In parallel with an integrate-and-fire model of interaural level difference computation which computes the difference of two log-encoded intensities, we have begun investigating the effectiveness of conductance-based synaptic computation.

The anatomy of the brainstem afferents to the lateral superior olive show the convergence of excitation and inhibition driven by the left and right ears, respectively. Although we have interpreted this in the past as addition and subtraction in an integrate-and-fire



model, we are now considering them as depolarizing and hyperpolarizing conductances driven linearly by intensity.

Figure 1 shows the circuit schematic for an abstract neuron with the intensity-dependent synaptic conductances g_L and g_R . We begin by representing the sound intensity as α , the transfer characteristic of the head to each ear as $h_L(\theta)$ and $h_R(\theta)$ and assume that each pathway has a gain that may not be exactly matched, denoted by k_L and k_R .

$g_L = \alpha \cdot k_L \cdot h_L(\theta)$ and $g_R = \alpha \cdot k_R \cdot h_R(\theta)$ and therefore in steady-state,

$$V_{mem} = \frac{g_R}{g_L + g_R} E_{inh} + \frac{g_L}{g_L + g_R} E_{exc}$$

$$V_{mem} = \frac{k_R \cdot h_R(\theta)}{k_L \cdot h_L(\theta) + k_R \cdot h_R(\theta)} E_{inh} + \frac{k_L \cdot h_L(\theta)}{k_L \cdot h_L(\theta) + k_R \cdot h_R(\theta)} E_{exc}$$

so V_{mem} (in steady-state) will be intensity-independent and will depend entirely on the relationship between $h_R(\theta)$ and $h_L(\theta)$.

Dividing both terms by $h_L(\theta)$ we see that we can express V_{mem} in terms of the ratio: $\frac{h_R(\theta)}{h_L(\theta)}$ or in

terms of the ILD expressed in dB, where $ILD = 20 \cdot \log\left(\frac{h_R(\theta)}{h_L(\theta)}\right)$

$$V_{mem} = \frac{k_R \cdot \frac{h_R(\theta)}{h_L(\theta)}}{k_L + \frac{k_R \cdot h_R(\theta)}{h_L(\theta)}} E_{inh} + \frac{k_L}{k_L + \frac{k_R \cdot h_R(\theta)}{h_L(\theta)}} E_{exc} = \frac{k_R \cdot 10^{\frac{ILD}{20}}}{k_L + k_R \cdot 10^{\frac{ILD}{20}}} E_{inh} + \frac{k_L}{k_L + k_R \cdot 10^{\frac{ILD}{20}}} E_{exc}$$


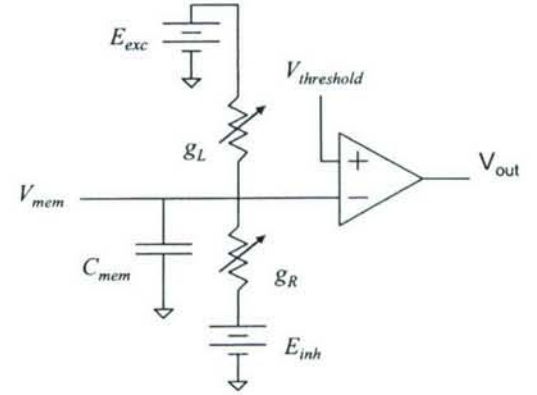
From the equation we can see that V_{mem} (in steady-state), depends only on ILD and other fixed parameters.

This provides us the insight that the neuron charges to the same steady-state voltage independent of the intensity of the sound, a critical feature of the computation. Sound intensity, however, does affect how quickly it charges to this steady-state voltage and thus affects the rate at which it will fire action potentials. This approach properly models the intensity-rate characteristic of real neurons, utilizes more realistic synaptic behavior, and solves a number of “initial condition” problems present in the purely integrate-and-fire neuron model.

We have just submitted a new version of the LSO model that contains an elegant current-mode circuit model that captures this conductance synapse computation.

$$I_m = I_o \cdot e^{\frac{\kappa(V_{dd} - V_b)}{V_T}} \left(\frac{I_e}{I_e + I_i} \right)$$

In this circuit, conductance is represented by two currents I_e and I_i and the final output current (which represents the membrane voltage) is a ratio of the two conductances that is



independent of sound intensity. I_o is a transistor constant and V_b is an externally-defined voltage. Not shown, is the current-based threshold that triggers an action potential in the neuron.

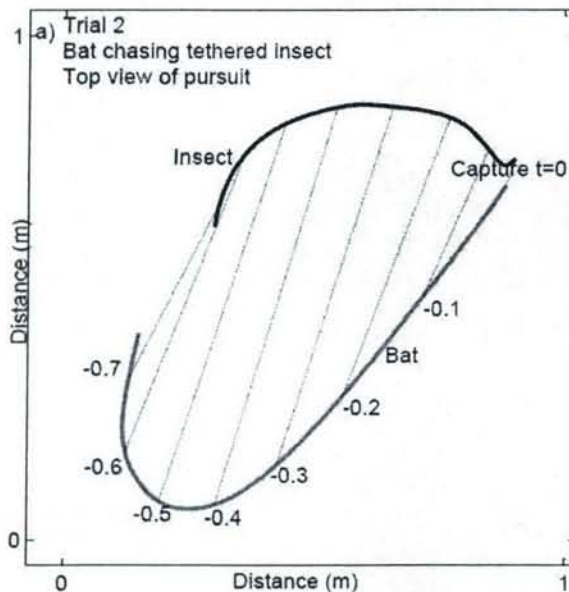
2. A Nearly Time-Optimal Strategy for Capturing Unpredictable Prey

To consider the behavioral strategies that bats use to navigate obstacles and capture prey, a collaborative study of 3D bat flight data (also supported by a grant from the NIH/NIBIB – Collaborative Research in Computational Neuroscience) involving Kaushik Ghose, Timothy Horiuchi, P. S. Krishnaprasad and Cynthia Moss has been analyzing the relationship of the steering choices of an echolocating bat to the bearing of a moving target. This work has produced some new insights into how effectively animals can utilize sensory information for the capture of prey.

It has been shown that the big brown bat's capture strategy is different from the "constant bearing" strategy often observed in nature. While many studies have shown that humans and many animals are capable of determining and controlling their movements to detect or generate collision paths, these studies have mostly focused on the case where prey move along straight-line trajectories. In such a scenario, by moving such that the target remains at a constant bearing relative to the direction of movement, a time-optimal path is generated. When prey move in unpredictable directions, this strategy breaks down.

Flight experiments in the laboratory of Cynthia Moss, investigating capture of free flying mantis and tethered insects by the echolocating bat *Eptesicus fuscus*, provide strong evidence that the strategy used by bats is a "constant absolute target direction" strategy; this strategy does not keep the target at a fixed bearing relative to the direction of movement, rather the orientation of the line connecting the pursuer and target in absolute coordinates remains constant. This is the same as the "parallel navigation" strategy used by missile guidance algorithms in the face of unpredictable target trajectories.

In the figure to the left, the flight of the bat (grey line, moving from $t = -0.7$ to $t = 0$) and the flight of the insect (black line) are shown together with lines connecting their positions at different times. Notice that the bearing of the insect with respect to the flight direction of the bat changes throughout the flight, yet the connecting lines remain remarkably parallel.



Notice that the bearing of the insect with respect to the flight direction of the bat changes throughout the flight, yet the connecting lines remain remarkably parallel.

This work has strong connections to the work on motion camouflage (described below) and further investigations of sensorimotor feedback laws are under way. It is anticipated that such investigations will prove useful in bat-inspired navigation algorithms for UAVs, implementable in analog VLSI technology.

Papers

Ghose, K., Horiuchi, T. K., Krishnaprasad, P. S., and Moss, C. F., "Echolocating Bats Use a Nearly Time-Optimal Strategy to Intercept Prey", **Public Library of Science Biology**, Vol. 4., Issue 5. (e108) pp 865-873, May 2006

3. Cooperative Control, Formation Control, and Motion Camouflage

P. S. Krishnaprasad has pursued questions of cooperative control using methods from differential geometry of curves and Lie group theory. This work is based on the idea of designing interaction laws for particles in two and three dimensions. Interacting particles are viewed as trajectory level models of unmanned aerial vehicles (UAVs). This body of work is in collaboration with colleague Eric Justh and former Ph.D. student Fumin Zhang (now a post-doctoral fellow at Princeton University). Initial work on planar formations (also supported by a previous AFOSR grant) has been successfully extended to three dimensions using the method of natural or relatively parallel adapted frames for curves (particle trajectories). In contrast to the well-known Frenet-Serret frames, natural frames have the advantage of being well-defined even on segments of a curve where the curvature vanishes. The natural curvatures are viewed as control variables, defining gyroscopic control forces on the particles via feedback laws that respect Euclidean group symmetry and additional discrete symmetries. Novel Lyapunov functions are used in convergence analysis of the resulting interacting particle systems to show collision avoidance, and convergence to relative equilibrium states. In a test-bed under development in the Intelligent Servosystems Laboratory, algorithms based on these feedback laws have been implemented on mobile robots to simulate UAV interactions. In the summer of 2005, Kevin Galloway, a Ph.D. student in Electrical and Computer Engineering, worked in the Naval Research Laboratory developing hardware-in-the-loop tests of the control laws on models of the Dragoneye UAV. He used Piccolo autopilots and a proprietary user interface in this successful test of three vehicle formation control. In the summer of 2006 a group of three undergraduates developed demonstrations of pursuit laws with mobile robots in the Intelligent Servosystems Laboratory at the Institute for Systems Research. They used a type of indoor GPS known as Cricket that employs RF and ultrasonic signals for ranging.

The methodology for cooperative control developed under this project has also proved useful in understanding certain situations of conflict in nature. It has been argued that various insects (e.g. dragonflies and hoverflies) employ a strategy of parallel navigation similar to the one known to missile guidance engineers, for the purpose of *motion camouflage*. Pioneering work on this intriguing biological phenomenon has been carried out by Mandyam Srinivasan and his collaborators. Under partial support from AFOSR we have shown the correctness of certain high gain feedback laws to realize motion camouflage. This work promises to be of use in developing cooperative attack strategies for UAVs and preliminary studies in this direction are under way.

During summer 2006, a graduate student Reddy did a detailed analysis of 3D bat flight data to see if they support the motion-camouflage inspired control laws that provably lead to the same geometric strategy as in bats pursuing insects. The results have been strongly supportive of the proposed control laws, but with a servo delay of 100 ms, consistent with other analysis of the data. This work is being written up.

This work has also been supported by Army Research Office under a MURI grant and a new single investigator award on *Gyroscopic Many Body Problems* (Program officer: R. Zachery).

Papers

F. Zhang, A. O’connor, D. Luebke and P. S. Krishnaprasad (2004), “Experimental study of curvature-based control laws for obstacle avoidance”, **Proceedings of the IEEE International Conference on Robotics and Automation**, pp. 3849-3854, IEEE, New York.

F. Zhang, E. W. Justh and P. S. Krishnaprasad (2004), “Boundary following using gyroscopic control”, **Proceedings of the 43rd IEEE Conference on Decision and Control**, pp. 5204-5209, IEEE, New York.

E. W. Justh and P. S. Krishnaprasad (2005), “Natural frames and interacting particles in three dimensions,” arXiv:math.OA/0503390, <http://www.arxiv.org/abs/math.OA/0503390> (8 pages), also in **Proceedings of the 44th IEEE Conference on Decision and Control**, pp. 2842-2846.

E. W. Justh and P. S. Krishnaprasad (2005), “Steering laws for motion camouflage,” arXiv:math.OA/0508023, <http://www.arxiv.org/abs/math.OA/0508023> (8 pages), also in **Proceedings of the Royal Society of London A**, FirstCite Early Online Publishing, pp 1-15 (2006), <http://www.journals.royalsoc.ac.uk/openurl.asp?genre=article&id=doi:10.1098/rspa.2006.1742>

Reddy, P.V., E.W. Justh, and P.S. Krishnaprasad (2006). “Motion camouflage in three dimensions,” arXiv:math.OA/0603176, <http://www.arxiv.org/abs/math.OA/0508023> (8 pages), also to appear in **Proceedings of 45th IEEE Conference on Decision and Control**, December.

Invited Lectures by P. S. Krishnaprasad

Spatial patterns in cooperation and conflict – Workshop on Swarming in Natural and Engineered Systems, Napa Valley, California, August 4, 2005.

Control, Observation and Feedback - Quantum Control Summer School, Caltech, Pasadena, August 8, 2005. (hour lecture)

Natural Frames, Interacting Particles and Stealth – Applied Mathematics Seminar, University of Illinois, Urbana-Champaign, September 19, 2005.

Natural Frames, Interacting Particles and Stealth – IPAM-UCLA workshop on swarming in natural and engineered systems, March 2, 2006.


Natural Frames, Interacting Particles and Stealth – Electrical Engineering and Computer Science Seminar, Harvard University, March 10, 2006.

Natural Frames, Interacting Particles and Stealth – Department of Mathematics Colloquium, North Carolina State University, April 7, 2006.

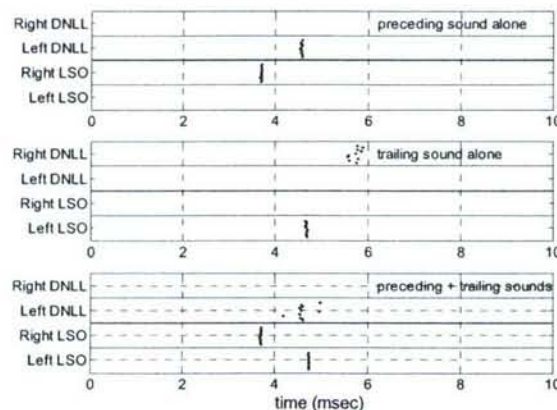
4. Modeling of the Interaural Intensity Comparison Pathway

In our previous progress report we showed data pertaining to the effects of long-lasting inhibition and refractory period on the cells of the lateral superior olive. These included the masking effects that prevent a cell from responding to new stimuli. Recently, we have been characterizing how these inhibitory effects manifest themselves at the next layer in the ascending information pathway towards the inferior colliculus. The next layer is the dorsal nucleus of the lateral lemniscus (DNLL).

In this work we explore the complicated interaction between the left and right halves of the LSO to DNLL pathway. In previous work we have shown that when the LSO neurons respond to an echo, the group on the contralateral side is inhibited while those on the ipsilateral side are excited. After the response to the first sound, the inhibited LSO neurons remain in this state for a short time, during which a second sound may arrive. These inhibited cells may then fail to pass along excitation to the contralateral DNLL and inhibition to the ipsilateral DNLL. This failure to respond can mean that inferior colliculus cells may fail to be inhibited and can respond in a non-selective or monaural fashion as described by Burger and Pollack (2001).



One half of the ILD pathway to DNLL. The ILD pathway starts from cochlear nucleus (CN). The first ILD processing center is the LSO, which receives excitation from the ipsilateral CN but inhibition from the contralateral CN via MNTB. The second ILD processing center is DNLL, which receives inhibition from the ipsilateral LSO (as well as from cochlear nucleus (CN)), but the excitation from the contralateral LSO (CN). Solid line for excitation, dotted line for inhibition. Revised from Yang and Pollack 1994.



Long-lasting inhibition in the mutually-inhibitory DNLL-DNLL synapses produces DNLL masking effects. Raster plot of four cells (Right DNLL, Left DNLL, Right LSO, and Left LSO) over 10 trials of an experiment. Top: An ultrasonic echo is presented to the right hand side (35 deg right of center). The right LSO responds (while left LSO is inhibited), which in turn excites the left DNLL neuron. Middle: A trailing sound is presented alone (30 deg left of center). The trailing sound excites the left LSO cell which in turn drives the right DNLL cell to fire. Bottom: When both sounds are present, the right DNLL cell is inhibited from the opposite DNLL cell (left DNLL).

Papers

(submitted) – Shi, R., and Horiuchi, T. K., “A Neuromorphic VLSI Model of Bat Interaural Level Difference Processing for Azimuthal Echolocation”, **Transactions in Circuits and Systems I**, (15 pages).

Shi, R., and Horiuchi, T., “A VLSI Model of the Bat Dorsal Nucleus of the Lateral Lemniscus for Azimuthal Localization”, **Proceedings of the International Symposium on Circuits and Systems (ISCAS 2005)** May 23-26, 2005, in Kobe, Japan pp. 4217-4220.

Shi, R., and Horiuchi, T., “A VLSI model of the bat lateral superior olive for azimuthal echolocation”, **Proceedings of the 2004 International Symposium on Circuits and Systems (ISCAS'04)**, May 23- 26, 2004.

5. The Ultrasonic Cochlea Chip

In our previous progress report we presented data for our first ultrasonic cochlea chip that was based on a parallel bank of resonators, different from most other implementations of silicon cochleae along with a novel current-mode neuron and address-event interface. These differences were all attempts to address the need for tightly-spaced, high-Q filters that did not suffer from electrical interference from the spiking neurons. We have since redesigned parts of this cochlea circuit to produce more control to obtain uniform Q-values (tuning sharpness) and a faster-spiking neuron for representing the sound intensity with higher dynamic range. At the same time, we have begun to investigate the use of latency of the neuron response to estimate the sound intensity. At the present time, we are in the midst of characterizing the new cochlea chip and its improved properties.

We are also testing a third cochlea chip that uses improvements in neuron circuit design for lower spike-induced power supply noise (see openspace chip below), a more precise filter biasing circuit, and incorporates a new approach to intensity coding with a heterogeneous population of cells at every tap in the cochlea. This most recent chip has been demonstrated to operate properly and we are now investigating different neural coding strategies for signaling intensity. The temporal demands of bat echolocation (short vocalizations) puts special demands on intensity encoding strategies because any one auditory nerve fiber can only fire a few spikes for a given echo. We anticipate the submission of a conference paper (ISCAS 2007) soon.

Papers

Abdalla, H., and Horiuchi, T., “An Ultrasonic Filterbank with Spiking Neurons”, **Proceedings of the International Symposium on Circuits and Systems (ISCAS 2005)**, May 23-26, 2005, in Kobe, Japan, pp. 4201-4204.

6. The Openspace Algorithm and Spike-based VLSI Implementation

How bats (and other animals) combine information about multiple obstacles and target directions to determine their movement is a general problem about navigation that has worked on heavily for many decades. While older, artificial intelligence approaches assumed an accurate map of objects and distances in an environment, recent approaches to this short-term obstacle avoidance problem have gone to the other extreme, utilizing the summation of imaginary


repulsive or attractive forces and torques (generated by objects) to steer a particular creature, requiring less-stringent accuracy on sensory data and relying on constant feedback with the environment. While this approach is clearly an improvement in most realistic cases, there are still a few drawbacks to consider. A large object, if detected as two closely-spaced objects, could produce twice the repulsive force. This force-field approach tends to produce trajectories that are easily trapped in local minima, such as in room corners. Part of the problem is that an obstacle that produces a repulsive force effectively “tells” the creature to turn away, whereas a more reasonable view to take is that obstacles only “tell” the creature where it should *not* go.

5.1 The Openspace Algorithm

The approach we are taking is a hypothesis-testing view of navigation, using the sonar system to evaluate the desirability of all possible 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 figure below). 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.

With this approach, closely-spaced objects produce only deep suppressions, but not necessarily very wide depressions in the evaluation function. The depth of suppression represents the confidence in avoiding that particular direction. In addition, an open direction on the far side of an obstacle from the direction of travel can be



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>

selected. Most intriguing from a neural modeling perspective, our model proposes that our creature actively represents empty space.

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

$$E(\theta) = E_o + g \cdot e^{-\frac{(\theta - \theta_g)^2}{\sigma_g^2}} - \sum_{i=1}^N \frac{1}{r_i} \cdot e^{-\frac{(\theta - \theta_i)^2}{\sigma(r_i)^2}} \quad (0.1)$$

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 corresponds 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.

5.2 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.

An obvious neural implementation of the openspace 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 is implemented on this field of neurons.

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.

Ya). 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. Yb), the latency will increase as inhibitory pulses start earlier. The use of such step currents in neural computation is described by Wolfgang Maass and has recently been used in a VLSI circuit for visual processing by Ravinuthula and Harris.

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

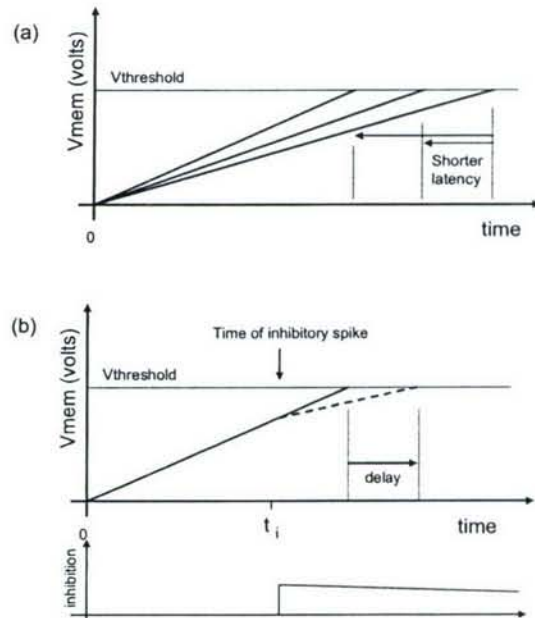


Figure Y. (a) Increasing the strength of the excitatory inputs to a neuron shortens the latency of the spike following a reset pulse. By determining the neuron that fires first, we find the neuron with the largest average input. (b) A long-lasting inhibitory current delays the spike or prevents firing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

by Equation 1.2.

$$T = \frac{C_{\text{mem}} V_{\text{thresh}} - I \cdot t_i}{E_o - I} \quad (0.2)$$

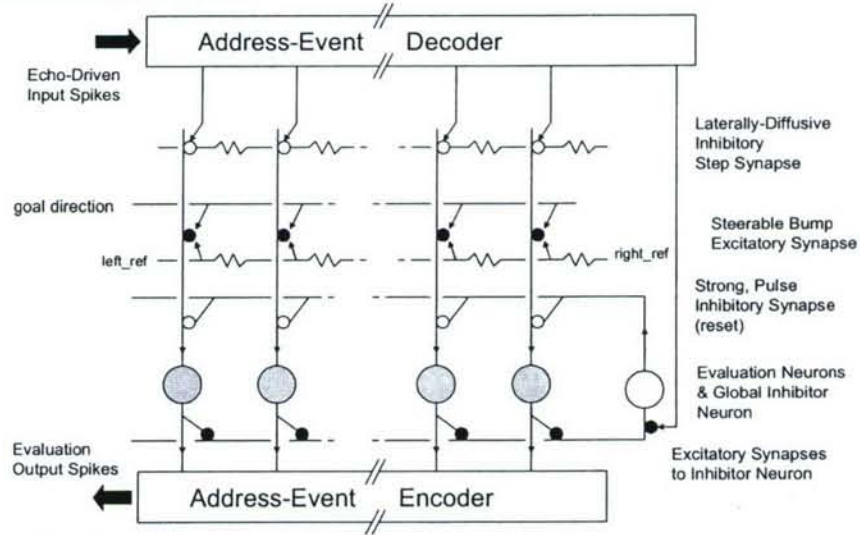
We assume that $E_o > I$ and that $t_i < \frac{C_{\text{mem}} \cdot V_{\text{thresh}}}{E_o}$. Thus, we obtain increased latency for closer obstacles *without* explicitly computing the range nor increasing the synaptic strength.

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.

5.3 Circuits

To facilitate the communication of spikes in and out of the chip, we use a communication protocol known as the “address-event” representation. In this system, an asynchronous digital bus provides the address of a target synapse and produces a handshaking pulse, delivering a brief ($\sim 1\mu\text{s}$) voltage spike to the target address. This same system is used to transmit neuron spikes out of the chip. The inputs to the system (see figure below) are the *goal direction* and 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.

There are two novel circuits used in this chip: a new low-power neuron circuit and a diffusive inhibitory synapse that takes one input spike and generates a spread of injected current in a limited neighborhood of neurons.



5.4 Testing Results

The chip was fabricated in a commercially-available $0.5 \mu\text{m}$, 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 $1 \mu\text{s}$.

We have successfully demonstrated that the array of neurons can operate in the race-to-first spike WTA mode, implementing the openspace algorithm. We have demonstrated the time-dependent inhibitory properties of the race to first spike for step-like inhibitory currents. The use

of input spike timing to modulate the efficacy of a connection can be an effective mechanism that does not rely on the modulation of synaptic strength nor on increased spike rates. We have found a natural match of this computational mechanism with the echolocation-based openspace algorithm proposed above.

Papers

T. Horiuchi, "A Neural Model for Sonar-based Navigation in Obstacle Fields" **Proceedings of the 2006 International Symposium on Circuits and Systems**, Kos, Greece, pp. 4543–4546, May 2006.

Horiuchi, T. and Cheely, M., "A Systems View of a Neuromorphic VLSI Echolocation System" **Proceedings of the 2007 International Symposium on Circuits and Systems (ISCAS 2007)** May 27-30, 2007 in New Orleans, LA. pp. 605-608.

7. The Horizon Estimation Chip – Constraint Minimization

Unmanned micro-aerial vehicles are rapidly being developed for use as a low-cost, portable, aerial surveillance platform for semi-autonomous operation. While they are successfully achieving flight, the sensors needed for autonomous flight (in contrast to long-range navigation) are lacking. Obstacle avoidance and flight stability remain a problem for such small vehicles with tiny weight and power budgets. Their small size makes them susceptible to tiny wind gusts, making the speed of processing critical for stability.


While many visual motion approaches to stabilizing aerial vehicles with low-power chips are being pursued, detection of the horizon is also desirable for high-altitude pitch and roll stabilization. Several low sensor-count horizon sensing systems have been developed for assisting aircraft pilots that utilize contrast in the infrared spectrum or visible light. In recent years, a team from the University of Florida (UF) has demonstrated an automatic visual-horizon finding algorithm operating on a high-speed computer on the ground that receives a transmitted color video-feed from the airplane. We have devised a similar algorithm that uses an optimization approach embedded in an analog VLSI vision chip to find the best visual horizon as well as provide a measure of confidence. The VLSI implementation has the potential for real-time, low-power performance and operation over a wide dynamic range of image intensities.

Most importantly, this project is intended to pave the way towards a new outlook on computation in neurally-inspired chips by exploring ways that VLSI circuits can solve minimization problems that could be ill-posed inverse problems, such as the problem of determining echo direction from multiple overlapping echoes.

I. HORIZON DETECTION ALGORITHM

A. The Horizon Vector, h

Consider an image where each pixel is assigned a horizontal and vertical coordinate (x, y) with the origin in the center of the image (see figure to the



Left: The sign of the dot-product of a horizon vector (h) and the pixel vector (p) plus a bias determines the horizon line. The dotted line represents the horizon line with a positive bias. Right: Block diagram of processing in a single pixel. v_x , v_y , and v_b represent the horizon vector and each pixel's coordinate (x, y) is determined by i_x and i_y .

right). We can think of this coordinate as the pixel vector, p^μ , where μ is the index. We introduce the horizon vector, h , and define the horizon as the boundary separating the two polarities (or “classes”) resulting from the dot product between pixel vectors and the horizon vector (plus a bias parameter, b , and fixed threshold, θ).

$$\text{class}(p^\mu, h) = \text{sign}((p^\mu \cdot h) + b - \theta)$$

The two classes represent ‘sky’ and ‘ground’. In the horizon detection algorithm to follow, exactly which class represents sky or ground will not be specified, but can be determined after the horizon is found by measuring the average intensity of the sky and ground classes. The horizon *line* is thus perpendicular to the horizon *vector* and is offset from the origin by a distance defined by the ‘bias’ parameter, b . The horizon vector is available at each pixel location and the class assignment is computed in parallel at each pixel.

The goal of finding a visual horizon requires a working definition of the differences between ‘sky’ and ‘ground’. We use an approach similar to the UF team, noting that a histogram of pixel intensities (in their case, the RGB vector) will show a bimodal distribution with the sky pixels bright and the ground pixels dark. This is obviously not always true, but describes most situations adequately. The goal is to find the line that best separates the two intensity distributions.

We begin by computing the average intensity of each class for the current state of the horizon line. At each pixel, the absolute *differences* between the pixel’s intensity and the class averages are computed. We decide that a pixel is ‘misclassified’ if the pixel intensity is closer in value to the opposite class average. The goal of the horizon detection algorithm is to find the horizon vector that *minimizes* the total number of misclassified pixels. For any realistic image, the horizon will not be perfectly straight due to trees, buildings, canyons, mountains, or lens distortions; by monitoring the total number of misclassified pixels, however, we will have an ongoing estimate of the success or failure to fit a straight line.

The horizon line calculation operates as a linear discriminant function over a two-dimensional input space. By utilizing neural network learning algorithms, we can achieve adaptation of the horizon vector to minimize the total number of misclassified pixels.




Figure X. Time-evolution of the horizon estimate. Top-Left: Static image, Top-right: Final sky/ground class image, Bottom-Left: Horizon vector end-points during iteration (line-circle is the final iteration), Bottom-Right: Total mismatch current during the evolution of the horizon estimate.




Figure XX. Three example snapshots of internal state following horizon detection. Top Row: Photodiode image. Bottom Row: sky/ground class image following settling.

II. TESTING RESULTS

The chip was fabricated in a commercially-available 0.5 μm , 2-poly, 3-metal CMOS process using the top metal layer as a light shield with holes over the photodiodes. Horizon-like images were projected onto the chip through a lens mounted on the chip package. The photocurrent and selected class for each pixel were scanned off using a current-sense amplifier. An example is shown in the top two panels of Figure 5. Transistor mismatch and light-leakage currents in the class determination circuit produce a ragged horizon boundary.

To observe the adaptation process, the v_x , v_y , and v_b voltages were held externally by computer-controlled digital-to-analog converters (DAC) while the total vector adaptation currents on these lines were measured. After reading the photodiode image and class image, the horizon vector voltages were iteratively changed in proportion to the measured adaptation current, simulating the time evolution of the v_x , v_y , and v_b voltages if they had been left floating.

Fig. X shows an example image where the horizon vector was initially pointing towards the bottom-left corner, producing an incorrect horizon line (Fig. X, bottom-left panel). Initial iterations show a rapid rotation of the horizon vector followed by a slow lengthening. The total mismatch current (Fig. X, bottom-right panel) which reflects the number of mismatched pixels rapidly decreases resulting in a stable solution. This final mismatch current is typical for good horizon solutions with the particular parameter settings used.

We then allowed the horizon vector and bias value to freely adapt to produce rapid horizon solutions. Three example images with the resulting class separations are shown in Figure XX. External 0.022 μF capacitors are attached to the v_x , v_y , and v_b lines for stability, creating a time constant of about 20 ms.


Roll angle estimation accuracy is shown in the figure below where horizon images were presented to the chip and the resulting horizon vector was transformed into an angle using: $\theta = \arctan((v_y - v_{ref}) / (v_x - v_{ref}))$.

While the power consumption varies dynamically with the image, long-term observations show the power supply current to be less than 500 μA (or 2.5mW with a 5V power supply).

III. DISCUSSION

While this chip satisfies the basic goals of the project, there are many places for improvement. In particular, the sensitivity of the sensor can be improved by using phototransistors and transistor mismatch can be reduced by using larger transistors. The contrast between ultraviolet (UV) and green light is known to be a more reliable measure and could be used here by adding optical filters and UV-sensitive photosensors. We are excited about the success of the operation of the chip and are attempting to apply some of the concepts to the problem of echolocation.

Papers



Measured horizon line roll angle vs. actual horizon roll angle. Error bars represent one standard deviation from 25 measurements.

Horiuchi, T., "A Low-Power Visual Horizon Estimation Chip", **Proceedings of the International Symposium on Circuits and Systems (ISCAS 2005)** May 23-26, 2005, in Kobe, Japan, pp. 4755-4758.

8. Binaural Sonar Experiment Board for the Telluride Neuromorphic Engineering Workshop

During the 2005 Telluride Neuromorphic Engineering Workshop, Giacomo Indiveri and Timothy Horiuchi organized a small robot race that required the use of sensors, robots, and various software controllers. One criticism we received was the inordinate amount of time that the participants spent to get even basic hardware functional and that not enough time was spent on the intellectual challenges. As we hope to create a new competition for the coming years, It was decided that it would be valuable to design and construct easy-to-use echolocation sensor boards that could be used by workshop participants next year.

Immediately after the workshop this summer, an undergraduate student (Jeffrey Lee) and Horiuchi designed and tested a narrow-band (40 kHz \pm 1 kHz) binaural sonar board that could be easily interfaced to MATLAB using an inexpensive (\sim \$150) USB-connected data acquisition board (Measurement Computing). In addition to being useful for our own research projects, this board makes algorithmic experimentation easy. We have built four of these boards: one that is in Krishnaprasad's laboratory and three of which are in use in Horiuchi's laboratory. In 2006, one of these sensors was used in Telluride (although the author did not attend) and during the 2007 Neuromorphic Engineering Workshop, 3 were used for different projects. We anticipate more sophisticated versions of these modules and hope to make sonar a more consistent sensory domain used in projects of the future.

