

Neuromorphic robot vision with mixed analog-digital architecture

Robotic vision is the most fascinating and feasible application for neuromorphic engineering, since processing images in real time with low power consumption is the field's most critical requirement. Conventional machine vision systems have been implemented using CMOS (complimentary metal-oxide semiconductor) imagers or CCD (charge-coupled device) cameras that are interfaced to digital processing systems operating with serial algorithms. These systems often consume too much power, are too large, and compute with too high a cost.¹

Though neuromorphic technology has advantages in these areas, there are some disadvantages too: current implementations have less-programmable architectures, for example, than digital processing technologies. In addition, digital image processing has a long history and highly-developed hardware and software for pattern recognition are readily available. We therefore think it is practical—at least at the current stage of progress in neuromorphic engineering—to combine neuromorphic sensors with con-

ventional digital technology to implement, for robot vision, the computational essence of what the brain does. On this basis, we designed a neuromorphic vision system consisting of analog VLSI (very-large silicon integration) neuromorphic chips and field-programmable gate array (FPGA) circuits.

Figure 1 shows a block diagram of the system, which consists of silicon retinas, 'simple-cell' chips (named after the simple cell in the V1 area of the brain) and FPGA circuits. The silicon retina is implemented with active pixel sensors (conventionally-sampled photo sensors)² and has a concentric center-surround Laplacian-Gaussian-like receptive field.² Its output image is transferred to the simple-cell chips serially. These chips then aggregate analog pixel outputs from the silicon retina to generate an orientation-selective response similar to the simple-cell response in the primary visual cortex.³ The architecture mimics the feed-forward model proposed by Hubel and Wiesel,⁴ and efficiently computes the two dimensional Gabor-like receptive field using the concentric center-surround receptive field.

The signal transfer from the silicon retina to the simple cell chip is performed using analog voltage, aided by analog memories embedded in each pixel of the simple cell chip. The output image of the simple-cell chip is then converted into a digital signal and fed into the FPGA circuits, where the image is further processed with program-

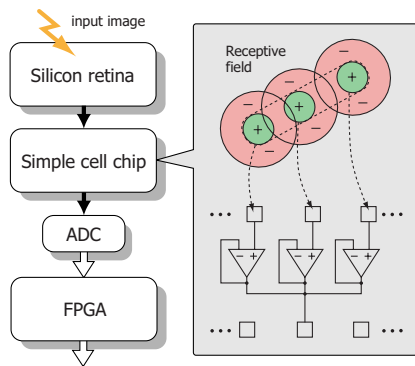


Figure 1. Block diagram of the mixed analog-digital system and analog multi-chip circuit for an orientation-selective response.

Figure 2. Photograph of the binocular vision system.



In this issue:

- Telluride Workshop Highlights
- 'Grand Challenge' Robot Race
- Audiosapiana: a biography
- High speed/precision winner-take-all circuit
- Bio-inspired robot vision for a grasping task
- Wide-dynamic-range imaging
- Cortically-inspired active binocular vision system

Laboratory Notes

- The USB Revolution




Figure 3. Real time disparity computation with the binocular vision system. (A) Sketch of the experimental setup. (B) Neural images of the complex cells tuned to the disparity of near, fixation and far.

mable logic in parallel (not in serial). An example of this real-time computation is shown in Figure 3.

The system emulates the responses of complex cells that are tuned to particular binocular disparities based on the disparity energy model.⁵ Here, the emulation is carried out with a binocular platform as shown in Figure 2. In the experiment, a hand was

LABORATORY NOTES

The USB revolution

We have all heard the quote, “Revolutions do not happen overnight”. Indeed, the USB revolution was years in the making: but soon after the USB (universal serial bus) board fired its first spike, nearly every member of our lab had abandoned their chip-testing setups to join the insurgency. This custom tool had ousted logic analyzers, pattern generators, digital I/O cards, and other bourgeois equipment. My comrades, we wish to spread the revolution.

We have two aims in this article: first, we want to demonstrate that, with the right components and driver, USB 2.0 offers a high-performance (7 million spikes/sec), portable (compatible with any PC), inexpensive (under \$500 in parts and printed circuit board), and flexible (coded in C++) computer interface. Second, we want to describe how this interface has revolutionized the way in which we interact with our neuromorphic systems. Our obvious enthusiasm stems from the sentiment that our chips are no longer limited by the surrounding technology, but instead by our own creativity.

Hardware

To achieve the interactivity that we envision, the USB link must seamlessly integrate with our neuromorphic chips. A typical neuromorphic chip with 10,000 neurons¹ can demand bandwidths exceeding 500,000 address-events²/s (considering both the transmitter and receiver), and proposed multi-chip systems³ can increase this demand by ten-fold. In theory, USB2.0 can support 480Mb/s: to blast address-events at these rates, we must enlist the appropriate chip set.

We chose the USB Cypress FX2 transceiver chip because its configurable datapath allows us to meet our performance requirement of streaming millions of address events per second. Specifically, we can bypass the FX2's on-chip microcontroller (through a firmware assembly configuration file), which would otherwise introduce a devastating bottleneck for real-time data transfers. We use a pair of 4×256×16 bit (quadruple buffered) FIFOs (first in, first outs) on the FX2 for streaming address events to and from our chip (one for each direction).

The transfers are asynchronous and we match the address-event representation protocol to the FX2's protocol using a CPLD (complex programmable logic device from *Lattice Semiconductor*). And indeed, we real-




Figure 1. USB 2.0 communicates between a PC and a neuromorphic system at 7 million 16bit events per second (duplex).

ize a high-performance bidirectional link, achieving 7 million 16bit events per second, or 112Mb/s (see Figure 1).

Software

With our high-performance USB2.0 link, we now have the ability to stream spikes (and other data) in real-time between a neuromorphic system and a software application (GUI, graphical user interface): the challenge is ensuring that the software can keep up. Since coding drivers is anathema to us hardware guys and gals, we decided to use the commercial USB driver from Thesycon. Naturally, Thesycon is not the only game in town, but by following their Microsoft Visual C++ examples we were able to integrate real-time events into our software apps in no time. Essentially what Thesycon provides is the ability to probe

the status of the FX2's FIFOs (both input and output). For instance, when an input FIFO buffer is full, the driver copies it to memory. At low event rates, this takes a while, so the CPLD injects timing events every 50 μ s. For the majority of the time, however, our GUI is free to process, plot, and save spike data.

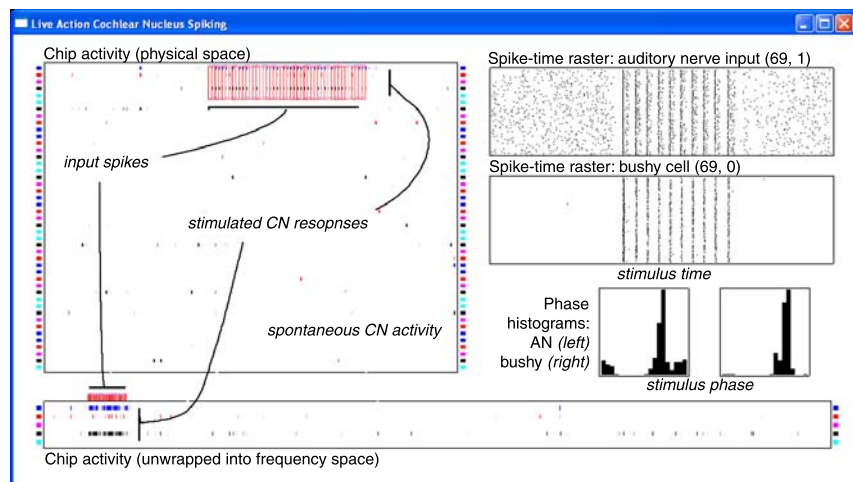
The GUI enables us to visualize chip activity, processing and plotting incoming spikes in real-time. Visualization, however, has the potential to hold up the link if FIFOs fill up faster than the screen can update. For this reason, we use OpenGL (the open graphics language), which plots extremely quickly.⁴ Of course, some of the more advanced plotting features (e.g., in three dimensions) are still too slow: it is not a general panacea. In practice, OpenGL suffices for the two dimensional plots we use.

¡Viva la revolución!

The USB2.0 interface has become an essential part of testing, visualizing, and interacting with the neuromorphic systems in our lab. We briefly highlight some of the systems we have visualized using this link in Table 1. Note that, for each project, the GUI was specifically tailored to monitor the essential computation in that system. For instance, if we are interested in the phase-locking

Merolla, continued on p. 11

Figure 2. The cochlear nucleus chip responds to sound encoded by auditory nerve spikes (inputs spikes). Chip activity displays neurons that fired recently according to chip layout (physical space) and cochlear mapping (frequency space). Spike rasters display 100 trials of a neuron's spike times during a 50ms stimulus, while phase histograms quantify spike-time precision. The selected bushy cell (address 69,0) phase-locks better than its inputs: its phase-histogram is tighter.



Shefer, continued from p. 8

that, for such variations, the LPF output remains constant. Denoting this constant by \bar{E}_o we have

$$G_{AC} \equiv \frac{dE_o}{dE_i} \Big|_{AC} = K(1 - \bar{E}_o) = K \left(1 - \frac{E_i}{1/K + E_i}\right) = \frac{1/K}{1/K + E_i}$$

from which we get

$$G_{AC} / G_{DC} = 1 + KE_i$$

Taking the quotient (G_{AC}/G_{DC}) as a measure of the detail enhancement of the retinal DRC, we see that the amount of this detail enhancement, or the ‘effective visual acuity’ increases linearly with the average luminance of the viewed scene: a well known property of human and animal vision. The lower half of Figure 2 is *Room* after retinal




Figure 3. The spatial feedback automatic-gain-control (fb-AGC) model.

DRC to 5bpc.

Conclusions

On the capture side of wide-dynamic-range imaging, we have demonstrated that dealing with saturation alone cannot significantly increase the DR of image sensors. We therefore conclude that this can only be achieved via significant noise reduction. On the display side of the problem, Figure 2 above demonstrates how retinal DRC can increase the potential number of perceived distinguished colors—by a factor of 512 in the current example—and shows how this affects our watching experience.




Figure 4. The Weber's law curve.

Thanks to Dr. Jacques Benkoski for his helpful remarks.

Moti Shefer
Trusight Ltd.,
Haifa, Israel
E-mail: moti@trusight.com

References

1. A. Belenky, E. Artyomov, A. Fish and O. Yadid-Pecht, *Wide dynamic Range (WDR) Imaging, The Neuromorphic Engineer* 1 (1), p. 4, Spring 2004.
2. M. Shefer, *Increasing the dynamic range of image sensors: Critiques of three commonly practiced methods*, www.trusight.com/motis_place.htm, Sept. 2005.
3. M. Shefer, *What's DR?*, www.trusight.com/motis_place.htm, Sept. 2005.
4. M. Shefer, *Neuromorphic-retinal dynamic range compression as applied to image display*, www.trusight.com/motis_place.htm, Sept. 2005.
5. M. Shefer, *AGC models for Retinal signal processing*, M.Sc. Thesis, EE Dept., Technion, I.I.T., November 1979.
6. S. Polyak, *The Retina*, University of Chicago Press, 1941.
7. F. Ratliff, *Mach Bands: Quantitative studies on neural networks in the Retina*, Holden-Day Inc., 1965.
8. G. Westheimer, *Visual acuity and spatial modulation thresholds*, *Handbook of sensory physiology* 7 (4), pp. 170-187, Springer Verlag, 1972.
9. J. E. Dowling, B. Ahinger and W. L. Hadden, *The interplexiform cell: A new type of Retinal neuron*, *Investigations in Ophthalmology* 15, pp. 916-926, 1976.
10. S. Chen and R. Ginosar, *Adaptive sensitivity CCD image sensor*, *Proc. SPIE* 2415, pp. 303-309, 1995.

Merolla, continued from p. 10

property of neurons, we can directly plot a metric that quantifies phase-locking: the task of finding neuron parameters that enhance phase-locking can now occur interactively

(see Figure 2). The ostensible real-time advantage of neuromorphic systems over other modeling techniques is now apparent.

The high-performance USB2.0 interface we have described here has revolutionized the way in which we interact with our neuromorphic systems. Our strategy for designing this interface was to optimize the hardware so that the much of the computational burden was transparent to the software. This strategy has paid off, as evident from the rapid development (three weeks) of the

five customized real-time applications.

Paul Merolla, John Arthur, and John Wittig Jr
University of Pennsylvania
Bioengineering Department
Philadelphia, PA
E-mail: {pmerolla, jarthur, jwittig}@seas.upenn.edu

References

1. P. Merolla and K. Boahen, *A Recurrent Model of Orientation Maps with Simple and Complex Cells*, *Advances in Neural Information Processing Systems* 16, S. Thrun and L. Saul, Eds., MIT Press, pp. 995-1002, 2004.
2. K. A. Boahen, *Point-to-Point Connectivity Between Neuromorphic Chips using Address-Events*, *IEEE Trans. on Circuits & Systems II* 47 (5), pp. 416-434, 2000.
3. T. Choi, P. Merolla, J. Arthur, K. Boahen and B. Shi, *Neuromorphic Implementation of Orientation Hypercolumns*, *IEEE Trans. on Circuits & Systems I* 52 (6), pp. 1049-1060, 2005.
4. http://nehe.gamedev.net
5. http://www.neuroengineering.upenn.edu/boahen/index.htm

Table 1: *Doing it on the fly*⁵

Revolutionary	Chip	Visualization (all real time)
Bo Wen	Inner ear	Cochleagram (spectrum vs. time) with real sound input
Brian Taba	Axon Guidance	Activity and histograms of ON-OFF retinal inputs for balanced topographic refinement
John Arthur	Hippocampus	Rasters and chip activity phase-coded in color
John Wittig Jr	Cochlear Nucleus	Rasters and histograms showing phase locking to sound waveforms.
Paul Merolla	Visual Cortex	Orientation maps generated with drifting gratings