REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
SEP 2011	**Conference Paper (Post Print)**	

4. TITLE AND SUBTITLE

A COLUMNAR PRIMARY VISUAL CORTEX (V1) MODEL EMULATION USING A PS3 CELL-BE ARRAY

5a. CONTRACT NUMBER
IN-HOUSE

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Michael J. Moore (ITT), Richard Linderman (AFRL), Morgan Bishop (AFRL), Robinson Pino (AFRL)

5d. PROJECT NUMBER
23T1

5e. TASK NUMBER
PR

5f. WORK UNIT NUMBER
OJ

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
AFRL/RITC
525 Brooks Road
Rome, NY 13441-4505

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RITC
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TP-2011-36

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA #: 88ABW-2010-0159
DATE CLEARED: 28 Jan 2010

13. SUPPLEMENTARY NOTES
Publication in The 2010 International Joint Conference on Neural Networks (IJCNN), 18-23 July 2010, pp 1-8, Barcelona, Spain, ISSN: 1098-7576, Print ISBN: 978-1-4244-6916, INSPEC Accession No.: 11593936, Digital Object Identifier: 10.1109 / IJCNN.2010.5596664 Date of Current Version: 14 Oct 2010.

14. ABSTRACT
A model of portions of the cerebral cortex is being developed to explore neuromorphic computing strategies in the context of highly parallel platforms. The interest is driven by the value of applications which can make use of highly parallel architectures we expect to see surpassing one thousand cores per die in the next few years. A central question we seek to answer is what the architecture of hyper-parallel machines should be. We also seek to understand computational methods akin to how a brain deals with sensing, perception, memory, and cognition. The success we encountered with achieving real-time performance is evidence validating the use of Cell-Be architecture in some classes of neuromorphic emulation. In this study we identified a particular gap-fill algorithm for lateral connections within V1 that is suggestive of a learning strategy whereby the lateral network subsumes expectation affect, reducing perception time and improving perception affect.

15. SUBJECT TERMS
Neuromorphic, Cognitive, Computing, Memory, Emerging Technology, Computational Intelligence

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			ROBINSON E. PINO
U	U	U	UU	9	19b. TELEPHONE NUMBER (*Include area code*) N/A

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

A Columnar Primary Visual Cortex (V1) Model Emulation Using a PS3 Cell-BE Array

Michael J. Moore, Richard Linderman, *Fellow, IEEE*

Morgan Bishop, Robinson Pino, *Senior Member, IEEE*

***Abstract*—A model of portions of the cerebral cortex is being developed to explore neuromorphic computing strategies in the context of highly parallel platforms. The interest is driven by the value of applications which can make use of highly parallel architectures we expect to see surpassing one thousand cores per die in the next few years. A central question we seek to answer is what the architecture of hyper-parallel machines should be. We also seek to understand computational methods akin to how a brain deals with sensing, perception, memory, and cognition. The model is being developed incrementally, starting with the primary visual cortex (V1) field. It is based upon structures roughly corresponding to neocortical minicolumn and functional column structures. Gaps in neuroscience, such as inter-cell connectivity, are filled using estimates of functionality that are plausible given current understanding of the micro-anatomy. The success we encountered with achieving real-time performance is evidence validating the use of Cell-Be architecture in some classes of neuromorphic emulation. In this study we identified a particular gap-fill algorithm for lateral connections within V1 that is suggestive of a learning strategy whereby the lateral network subsumes expectation affect, reducing perception time and improving perception affect.**

I. Introduction

THE objective of the project is to investigate architectural issues surrounding neurobiological inspired computational methods based on networks of structures roughly emulating cortical columns. It is the first step in a larger investigation of multiple classes of applications which may be able to take advantage of large scale parallel computing. This multidisciplinary effort focuses on determining how neurological systems perform those aspects of cognition associated with sensing and perception. The work progressed initially on ventral tract (object recognition) aspects of the visual cortex, and is now shifting to include the dorsal tract, theoretically associated with spatial properties.

Manuscript received January 15, 2010. M. J. Moore is with ITT/AES, 775 Daedalian Drive, Rome NY 13441 (305-330-1500; mike.moore@itt.com).
M. Bishop is with USAF AFMC AFRL/RITC, 525 Brooks Road Rome, NY 13441-4505 (315-330-1556, fax 315-330-2953) Morgan.Bishop@rl.af.mil.
R. Pino is with USAF AFMC AFRL/RITC, 525 Brooks Road Rome, NY 13441-4505 (315-330-7109, fax 315-330-2953) Robinson.Pino@rl.af.mil.
R. Linderman is with USAF AFMC AFRL/RI, 525 Brooks Road Rome, NY 13441-4505 (315-330-2208, fax 315-330-2953) Richard.Linderman@rl.af.mil.

A. Anatomy

There are about 1.6 million axonal fibers delivering information from the eyes into the primary visual cortex [12] through the lateral geniculate nucleuses **(**LGNs**).** Each side of the brain receives half of these, organized retinoscopically and stereoscopically. The retinoscopic organization means that the image carried by the fibers is spatially preserved, as if projected through a lens. The stereoscopy characteristic has to do with field of view. Each eye has a left and right field of view. The left side of the brain receives the right field of view from each eye, and the right side receives the left field of view. Thus each hemisphere of V1 receives approximately 800K fibers delivering two partially overlapping fields of view. The neuro-pathway for these, between the LGNs and the visual cortex, is called the optical radiation. There are two; a left and a right. Each of the hemispheres bundles its approximately 800K feed forward axons with approximately 3.2 million feedback axons, terminating at its LGN. The feed forward axons are mostly of two types: Parvocellular (P) axons and Magnocellular (M) axons. The P axons are thought to be associated with shape and color perception; the M axons with motion [18]. P axons account for about 80% of the feed forward; M accounts for about 5%.

V1 itself is part of the neocortex, which in turn is the top layer of a primate brain. The neocortex is thought to be where the essential mechanisms of human cognition reside. It is central to sensation and perception. The neocortex is a sheet of tissue roughly 3 mm thick and 2500 cm^2 in area (2.5 ft^2) [16]. The primary visual cortex is an area roughly 28 cm^2, accounting for both hemispheres [5]. Thus the primary visual cortex is a little more than 1% of the neocortex by area. The total number of neurons in the cerebral cortex is estimated to be 20 billion [11]. The total number of neurons in V1 is estimated as 280 million [11], and thus V1 is about 1.4% of the neocortex by neuron count. The neurons within V1, looking perpendicular to the sheet, are arranged into structures of neurons forming ~30 μm diameter columns extending through the six layers [15]. The columns are called "minicolumns." Estimates for neurons per minicolumn within V1 are in the range of 120 to 200, but using a rule of thumb that the incoming axons from the eyes are roughly evenly distributed, it works out that there is one minicolumn for each afferent (from the eyes) axon, and the neuron count per minicolumn is around 150.

Each parvocellular axon potentially connects to an area whose diameter is approximately 400 microns, which happens to be the scale of a functional column [13]. These "P Channel" fibers provide high contrast, spatially fine grained color information to the brain. Magnocellular fibers overlap a 1,200 micron diameter area, which happens to be on the scale of a hypercolumn [13]. These "M Channel" fibers carry low contrast information on the visual field, are associated with depth and movement perception, and are notably much faster to respond than the "P" channel.

Minicolumns exhibit excitatory and inhibitory interactions with each other. Excitatory communications appear to span a radius of about 3mm [3] while inhibitory is half that [20]. The excitatory span has a reach of about 14 functional columns across the diameter, and the inhibitory about 7 functional columns. Excitatory appear to connect up every other functional column, though there is debate about this. Inhibitory appear to hit every functional column within its reach.

B Levels of Modeling

Neuroscience has provided multiple complexity levels for modeling the cells comprising a brain. There are two general types of cells in the brain (ignoring the circulatory system): neurons and glia cells. The neurons are the cells with axons and dendrites which neuroscientists have historically assumed are the basic functional components of a brain. Glia cells out number neurons 10 to 1. They provide the scaffolding and life support environment for the neurons. They are recently thought to play more of a role in cognition than has been traditionally assumed [9]. Glia cell modeling is accounted for at a molecular level, typically with pharmaceutical interest. They were not included in this study.

The question is how to separate and identify the computationally useful characteristics of neuro-matter from those that are purely life supporting. Neuroscience has developed compartmentalized models (Gerstner and Kistler, 2002) of neurons which capture the intricacies of neuron physical size and shape (morphology), electrochemical dynamics (electrophysiology), molecular interactions between neurons and with glia cells (neurochemistry), and interpretations of information processing thought to be performed by neurons. The more detailed models require significant processing power to emulate. Which characteristics of these cells are harnessed by nature to produce cognition is an open question. It is not clear whether cells are the functional components of cognition. Collections of cells, perhaps cortical columns, may be the key functional building blocks.

Neurons exhibit increasing feature complexity as one looks closer into them. Very detailed compartmental models exhibit up to tens of thousands of individual synapses (connections), each with attributes such as connection strength, type (inhibitory, excitatory), dynamical characteristics, distance from the soma (nucleus), and neurotransmitter type. Simple models of neurons capture only the integrative and non-linearity estimates, ignoring electrophysiological pulse responses and spike timing dependent plasticity; they may have only a few connections. At higher levels of abstraction collections of individual neurons are replaced by "cognitive models" performing the hypothesized functions of the collectives; functions like association, feature perception and memory.

Setting a level of abstraction in a model constrains what the model can do. Accounting for all known cognitive behaviors with a simple model is evidence that the cells are being modeled validly, at least until new behaviors are identified. Levels of feature use may vary across the cortex. For example: detailed dynamical neuron models were not necessary to achieve the efficacy we expected of V1 in this study. We acknowledge they may be needed for other cortical regions or even for V1 itself should Integrate and Fire neurons be an insufficient mechanism.

The "affect" objectives of the V1 model are to account for orientation, color, depth perception (disparity), and motion percepts. The model proposed here has addressed orientation, and partially addressed color. Depth and motion are future plans. Not much is known about how neurons are systematically organized to produce and represent these affects, but there are hints.

Self imposed is the objective to emulate a full scale V1 in real-time. The ability to process in real-time simplifies the use of live video feed and provides a level of practicality reasonable for testing a model over extended durations. Real-time performance adds a "time complexity" challenge to computation, in the "big o" sense[2], restricting the use of algorithms with high time complexity.

C Simulation Facility

At our disposal is a 336 node Play Station 3 CELL-BE cluster organized as 14 subnets each with 24 nodes. Each subnet has a dual 3GHz quad core Xeon processor head node. Network interconnectivity is 10 Gigabit Ethernet amongst the head nodes and 1 gigabit Ethernet to the PS3s. Each PS3 node has six available Synergistic Processing Elements (SPEs) and a dual core 3.2 GHz PPE (Power PC). There are 2116 SPEs in total. Each SPE is capable of slightly more than 25.6 GFLOPS for a total CELL-BE cluster capability exceeding 54 TFLOPS, not accounting for head node and PPE contributions. GNU C++ development tools were used to develop the emulator, and a publish/subscribe message passing system was used for communication within the emulation. The "Pub/Sub" message paradigm loosely couples peer to peer message

passing. A message sender (publisher) does not send to a specific destination. Instead, each message has information in a "header" which describes what it is. This information might take the form of XML, plain text strings, or binary encoded numbers; specifics depend on the individual message system. The point is that the sender is unaware of the destinations. Receivers (subscribers) "sign up" to receive messages based on header content (what the message is) rather than the message source. In a system like a cortical model, inter-process connectivity can then be achieved by subscribing to (for example) axonal fiber names, and publishing on fiber names.

This 2400 core (Xeons + Power PCs + SPEs) facility's processors are somewhat specialized. The head nodes are conventional general purpose platforms with 32 GB of memory (each). The CELL-BE PPEs, also general purpose, each have 228 megabytes of RAM. The SPEs are specialized to be vector processors; they each have about 128Kbytes of useable RAM. Very fast DMA channels within a CELL-BE move data between main memory (PPE) and SPE memory. The Xeons, and PPEs, run Linux; the SPEs are essentially managed by the PPE with only minimal resident executive kernel software, but can interact with each other and the PPEs using DMA channels, interrupts and semaphores.

II. Model Constraints

The model was devised to be close to the anatomical structure of V1. It was also devised to make use of methods our preliminary investigations found compatible with CELL-BE architecture. These included:

- Small collections of neurons, strong localized connectivity, sparse distant connectivity;
- Integrate and fire neurons;
- Spatially tuned receptive fields;
- A localized associative component, possibly a small scale recurrent neural net;
- Feature extraction: max/min calculations, difference calculations, energy estimates, threshold detection;
- Inhibition, excitation interactions.

Methods considered, but avoided initially were:

- Confabulation algorithm [8], on the basis that it required large amounts of memory to support symbol lexicons (this decision was revered after it became apparent Confabulation was useful within the V1 lateral network);
- Spiking neuron models [7]: on the basis that the cognitive mechanisms hypothesized for these, principally dynamical phenomena, are not yet well demonstrated or characterized;
- Bayesian networks [3,5]: on the basis that we are seeking a model more closely aligned to anatomical details;
- Large scale associators, such as Sparse Distributed Memory (SDM), on the basis that we did not feel it was needed for a V1 model.

The challenge of model development was to create a system using just the selected methods that could meet the perception objectives of shape (orientation line), color, motion, and disparity.

III. Model Description

Orientation line perception is the major effort of modeling thus far. It is expected to be the most computationally challenging of all the V1 percepts. Aspects of color perception have been included, and a color percept is produced. It is modeled as the average color and intensity cast onto the field of view of a functional column, and includes an ocular dominance feature which selects the strongest percept in an overlapping (stereoscopic) fields of view. In those cases the functional column with the dominant orientation percept inhibits the other functional column. Motion perception is, like color, part of the objective but not yet emulated. Motion, based on magnocellular information, will produce a percept spatially mapped to the functional columns detecting it; direction and intensity are the intended percepts. The biomorphic model is based on the Reichardt effect[17] using synaptic arrival time differences to excite a neuron. In practice, we are looking at FIR and IIF filters for emulation.

The model is intended for full scale emulation. For this reason parameters are sometimes selected to accommodate the digital environment of the emulation, within the constraint that they represent plausible and reasonable neurological system values. One of these accommodations is powers of two. We have selected the following organizational parameters:

- Number of "ocular axonal fibers" entering V1 hemisphere: 802816;
- Total minicolumns per V1 Hemisphere: 802816;
- Minicolumns per functional column: 64;
- For the sake of emulation, we devised a subunit of a V1 hemisphere which we call a "subfield." A subfield is a collection of 128 functional columns, 64 of which are right FOV and 64 are left FOV. Each (full scale) hemisphere consists of 98 subfields. Note that (98subfields) X (128 FCs per subfield) X (64 minicolumns/FC) works out to 802816 minicolumns per hemisphere.

All minicolumns within a functional column are assumed to have the same parvocellular field of view (aperture). Four functional columns form a macrocolumn; all minicolumns within it are assumed to have the same magnocellular FOV

from both two eyes, and are responsive to all colors and orientations.

The minicolumn model is based on estimates of cell populations in cortical levels II, III, and IV(see Fig. 1). The level IV model component consists of:

- 56 simple cells dedicated to parvocellular inputs
- 10 simple cells dedicated to magnocellular inputs
- 8 complex cells dedicated to (parvocellular) orientation perception from simple cells
- 8 complex cells (not yet modeled) dedicated to (magnocellular) perception.

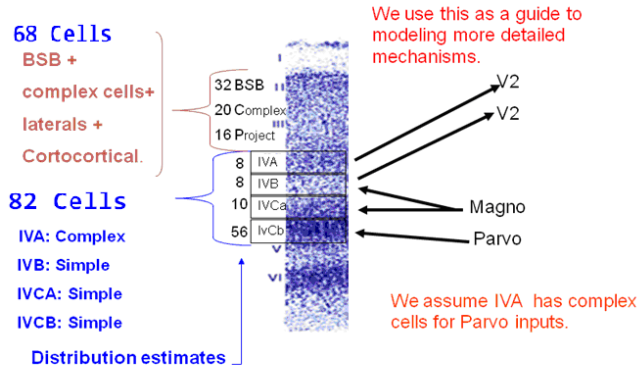

Fig. 1 Plausible cell populations within cortical layers of a V1 minicolumn.

The model currently makes use of parvocellular information; the magnocellular part of the model is not yet completed. Disparity, color and motion are not yet completely modeled, and will likely be modeled by having a subset of minicolumns within a functional column (cytochrome oxidase blob regions[19]) specialized for their perception.

The parvocellular simple cells each make 16 synapses with the afferent fibers. Half are dedicated to dark sensitivity, half to light. The color image is converted to shades of gray before presentation to the simple cells. Each simple cell receptive field has an angle, direction (light to dark, or dark to light), size/shape, and location (see illustration in Fig 2). Variations in size and location provide a degree of invariance.

Each Minicolumn has 56 such parvocellular simple cells, all looking for the same angle, but half looking for light to dark transition and half dark to light. The minicolumns are arranged into 8 columns of 8 (Figure 3), approximating orientation column anatomy[(10)]. Each column is dedicated to a specific angle. The 8X8 structure results in angles that are 22 ½ degrees apart.

The simple cells function by summing their synapse values and "thresholding" the results. The thresholds are presently constant, but variability will be explored in the future as part of a contrast control mechanism.

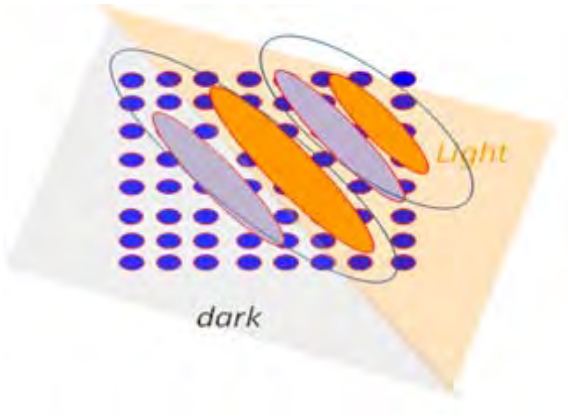

Fig. 2 An illustration of two simple cell receptive fields projected onto the FOV of a functional column. Gray ellipses represent synapses sensitive to dark; yellow, to light. Blue dots represent terminations of afferent fibers.

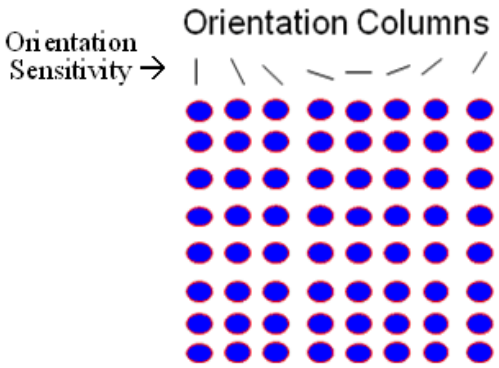

Fig. 3 In this view, dots represent minicolumns. Orientation columns are each a stack of eight minicolumns. Each column is sensitive to one orientation. A functional column is a collection of eight orientations columns.

Complex cells receive simple cell outputs (Figure 4). Four of the eight complex cells form synapses to simple cells that can detect light to dark transitions; the other four to dark to light transitions. Each complex cell makes synapses to 15 simple cells of the 26 available to it. The selection of which simple cells is based on a preference for simple cell receptive fields which center their receptive fields approximately along the same line, at the minicolumn's perception angle. The four regions of perception within the minicolumn's FOV established by this preference, overlap. The complex cells sum their inputs, and normalize the results to be within the range [-1… +1]. For example, a dark to light sensor would issue a -1 for light to dark transition perfectly aligned with it.

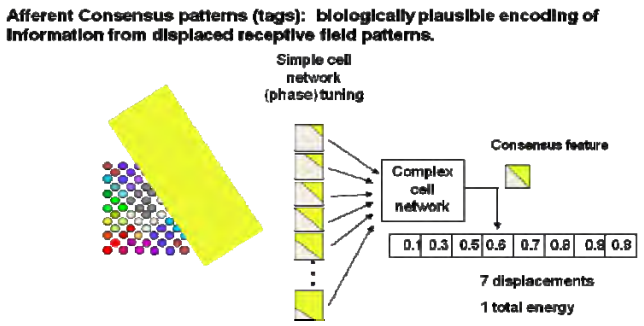

Fig. 4 Light (yellow) passing over a minicolumn's FOV; lower left corner in darkness (shown as colored dots, indicating afferent axon terminations). The simple cells are tuned to all spatial phases.

The outputs of the 8 complex cells are presented to the level II/III part of the model (illustration in Fig. 5).

The level II, III part of the model is called the associative component. It deals with data coming from three sources:

- Afferent detections from level IV,
- Lateral (horizontal) connections to nearby minicolumns,
- Expectation data from other cortical regions such as V2.

The model uses a 32 element recurrent network "Brain State in a Box" (BSB)[1] attractor function to decide whether or not a minicolumn perceives its angular percept. Every minicolumn has its own BSB state vector, but all share the same weight matrix. The common weight matrix is pre-trained to have two basins of attraction; these are set at opposite corners of the BSB hypercube. The basin points correspond to "I see a light to dark transition" and "I see a dark to light transition." Neuromorphically, this may correspond to actual recurrent neural networks, randomly wired but capable of being point attractors. There is no need to involve the BSB in differentiating an angle; the Level IV network does that, and supplies eight elements of "evidence" to the state vector. When the rest of the vector is neutral, afferent inputs alone can drive the BSB to a basin if the angle is fairly well sensed by Level IV. Likewise, Lateral and Extrastriate (expectation) data can singularly drive the BSB to a basin.

The minicolumn concludes its feature perception by computing the (Cartesian) distance of its state vector to each basin of attraction. The shortest distance is selected and is subjected to a threshold criterion. Distances closer than the threshold are converted into a range [0 ... +1] for light to dark, and [0 ... -1] for dark to light by differencing with 1.0 (1.0 – Distance, or -1.0 + Distance, depending on light/dark direction). Subthreshold cases are set to 0.0.

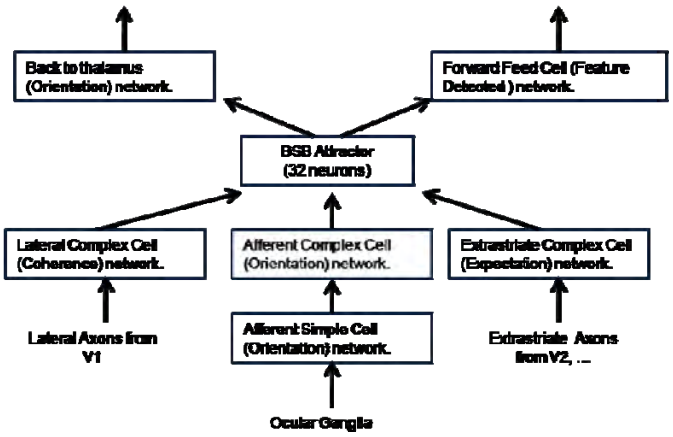

Fig. 5 The associative layer (II/III) has a BSB attractor whose state vector receives inputs from afferent, lateral and extrastriate sources. One of two features is decoded off the state vector and sent as feedback to thalamus and feed forward to extrastriate regions.

Each minicolumn within a functional column contributes to a functional column hypothesis. The strongest perception within each orientation column is selected for the hypothesis. The hypothesis is sent to all neighboring functional columns within a 3 mm reach. The receiving functional column "knows" the distance (hence, a weight) and direction (one of the 8 angles) the hypothesis came from, and uses the information to excite a "token." Tokens in this case are the 8 angles of perception, and their transition direction (light to dark, dark to light). All incoming lateral hypotheses contribute to this excitation. A "winner take all" strategy selects the most excited token and the token is then asserted onto the elements of the state vector dedicated to laterals (same value loaded into an eight elements, having a multiplicative effect on the BSB state vector dynamics).

The whole lateral process is similar to the algorithm reported by Hecht-Nelson which has been demonstrated to generate sentence text based upon noisy data and incomplete sentences[8]. Dubbed "Confabulation Theory," Hecht-Nelson proposes that the brain deals with distinct symbols which are percepts detected by neural networks. These symbols occur in context with other symbols. His example is text: the words would be the percepts, and sentences the contexts. The idea is hierarchical; groups of words (phrases) can be percepts, and paragraphs contexts. Weight matrices ("knowledge links') drive a selection process where a single symbol is selected from a lexicon at each contextual position. Unlike the reported Confabulator, this V1 model uses a large number of lexicons (>500 instead of 20), and each lexicon is small (16 symbols (edge percepts) instead of 10,000 (word symbols)). It gives the model the ability to "see" illusional contours and improve perception in noisy data. Figure 6 illustrates both situations; a diffraction grating is simulated at 135 degrees, with data missing in parts of the field of view passing over three function columns. On the left the upper block is the feed forward

perception, and the lower is perception after lateral data is applied to the minicolumns. A "lateral expectation" based on context tips the minicolumn into perceiving portions of the lines where there are actually blanks. On the right noisy data and limitations of the apertures cause misperceptions of 67.5 degree angles (using feed forward only). Again, the lateral effect corrects the feed forward perceptions (lower). It is plausible that this sort of mechanism can give V1 an ability to "see" combinations of small aperture edge percepts preferred by V2.

A full scale V1, both hemispheres, was emulated using 196 IBM/Sony PS3 Cell-BE processors configured as subclusters of 24 nodes attached to head nodes (*dual quad core Xeon X5450 3GHz)* (Figure 7). At the basis of message communication is IP, but a Publication/subscription service layer was used on top of IP to mitigate the tight binding imposed by socket to socket communication. The Pub/Sub message layer significantly reduced the complexity of regional lateral communications, where functional column hypothesis has to be shared among neighbors within 3 mm. All emulation software was written in C++..

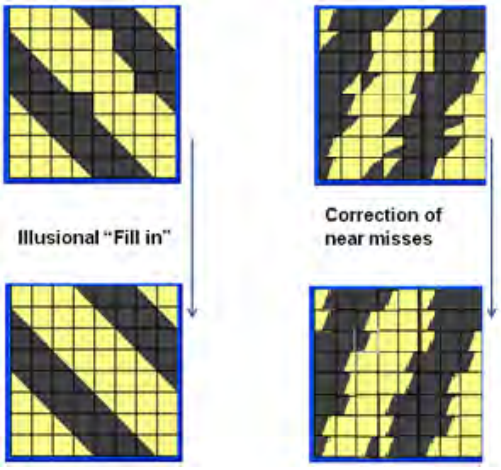

Fig. 6 Two examples of the two dimensional "Confabulation-like" lateral model producing an illusional percept (left column) and correction (right column). Each small grid box represents a functional column (64 minicolumns) in this illustration. The left field of functional columns was exposed to a 135 degree grating pattern. The right side was exposed to a 67.5 degree pattern.

Head-node software consists of stimulation and monitoring which roughly emulate ocular afferent pathways. There is a retina model (one or two may be used) which provides a left and right visual frame (magno and parvo). Output (for the time being) is RGB color pixels. A chiasm model combines frames from retinas and separates them into left and right stereo fields of view. An LGN model is simply a relay which chops up the stereo frames into smaller pieces (essentially subfield FOVs) that get delivered to the PS3 nodes.

Each PS3 node handled 8192 minicolumns and the related functional column model. For development convenience each group of 8192 minicolumns is termed a "subfield," and so each PS3 node handled one subfield. The BSB attractors cycle 5 times for each perception trial. In general, the PPE side of the PS3 nodes handled messaging and orchestration of the SPE processors, and hypothesis generation. The SPEs handled the emulations of Levels II/III and IV. Emulation speed is real-time. Each node is able to complete its processing in about 5.9 milliseconds. The most time demanding aspect is delivery of image fragments to the PS3 units. This takes about 10ms. The entire cycle time for a single frame was measured to be about 18 ms, or 55 Hz.

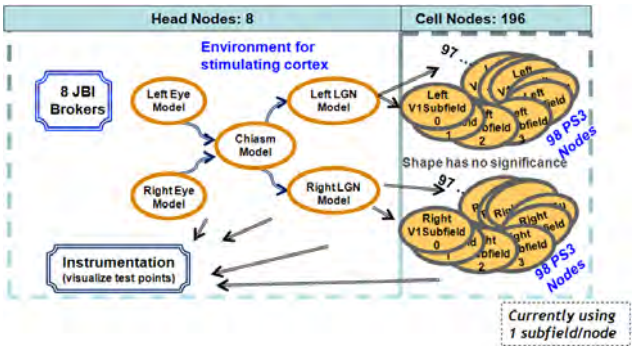

Fig. 7 Schematic of the emulation architecture. "JBI" is the name of the Publication/Subscription message layer used by the emulation.

IV. Sensory Perception Results

To date, only high contrast images are being presented to the system. Natural scene images will be attempted when a contrast control mechanism is in place. The initial test patterns were ideal diffraction grating images spaced to guarantee separation of "bar lines" on functional columns by a distance at least sufficient so no functional column was exposed to two separate lines. No expectation was used during these tests to reinforce perception. The grating patterns were moved across the field of view in steps comparable to the diameter of a minicolumn. There were significant misperceptions of +- 22.5 degrees when image bars were near the spatial limits of the functional column fields of view, but lateral "confabulation effects" corrected these near the end of the perception cycle (see Fig. 6). Sensitivity to contrast was significant, indicating the need for contrast control. However, certain applications, like reading text, are normally high contrast activities which the current model is reasonably suited to pursue.

V. Computational Results

The emulation had two major computational modules: "Layer IV" and "Layer II/III" corresponding to cortical layers. Layer II/III (also called the associative layer) included the 32 element BSB attractor, and a small neuronet which formed functional column perception consensus. The Level IV module emulated the spatially tuned simple cells and the complex cells connecting them to the associative

layer. These all executed on SPEs which, ideally, are able to compute at 25.6 GFLOPS. The associative layer code ran in 2.833 ms, achieving 10.5 GFLOPS. The Level IV code ran in 2.602 ms, achieving 8.6 GFLOPS. The processing of one video frame subfield takes 5.9 ms; approximately 0.47 ms is accounted for in feed-forward message handling. The entire field of 196 subfields required <17 ms. In this case the extra 11.1 ms is due to the serialization of dispatching subfield size pieces of an image. This serialization can be reduced in principle through parallelizing the network feed into V1 from the LGN modules. The 17 ms cycle time for a single frame corresponds to a frame rate of 58 Hz.

VI. Conclusions

A. *Model Efficacy*.

The model perceives lines of orientation well with high contrast images such as text and line drawings. The perception is reduced on natural images but not absent. The diminution is expected because no contrast control is presently incorporated in the model.

The spatially tuned simple cell model only roughly corresponds to actual V1 simple cell spatial tuning. This coarse approximation does indeed provide a useful degree of perception success. It is suggestive that hardware could be devised to likewise be spatially tuned, providing a perception mechanism potentially less computationally complex than Gabor[23] functions.

B. *Efficiency*. We measure efficiency of code executing on a platform in terms of an ideal application executing at 100% efficiency. The metric selected for the Cell-BE platform is floating point operations per second (FLOPS). Ideally, an SPE can achieve 25.6 GFLOPS. The two segments of SPE code, the associative and layer IV, achieved 10.5 FLOPS and 8.6 respectively corresponding to efficiencies of 41% and 33.5. This particular V1 model characteristically has relatively small vectors, typically 32 element or 56 element. Larger vectors are more efficient to manipulate on this platform than smaller ones. For example, large vector space associators, such as Sparse Distributed Memories (SDMs), represent a class of neuromorphic computing algorithms useful for modeling memory structures. SDMs deal in vector spaces on the order of 500 to 1000 elements.

Scaling. A subfield is 128 functional columns represented by a single process running on a PS3 node. The full sale V1 consists of 196 instances of a subfield, all identical, each running on a dedicated PS3 node. The processes are easy to scale because they are embarrassingly parallel[22]. Except for verification of lateral network "Confabulation" all "debugging" was performed using a single node model. Subfields were added in increments of 10 during sessions instrumented to measure message passing latencies, but such incremental scale-ups were not necessary for other types of functional verification.

The system scaled well from individual subfields to full scale V1. Message system latency was the largest source of efficiency degradation as the scaling increased. The message latency can be mitigated by parallelizing the scattering of an image into subfields, and the corresponding distribution of the scatter.

It is clear that the CELL_BE nodes are idle most of the time. It is possible to add significant functionality to the model and still comfortably manage real-time performance. We feel it is likely Contrast management, motion perception, depth perception and an improved color perception capability will be computable within the computational slack time.

C. *Developer's Experience with the platform.*

The CELL-BE architecture fit well, relying on fast DMA communications between node components. Nevertheless, efficiency was sufficiently high so as fewer nodes could have been used for this particular algorithm and still perform in real-time.

The PPE and SPE components each exhibited characteristics which required special attention beyond that normally required on conventional machines to produce highly efficient code. On the PPE special effort was needed to manage the Translation Lookaside Buffers (TLB), which handle the virtual address translation. The problem manifested as slow performance, due to page faults. The solution was to utilize Linux system calls (mlock, mmap, setmntent) to lock a virtual space into physical space. On the SPE, developers experienced a three to four week "ramp up" time to become proficient with the DMA, and vector features. Each code segment development required extra time for optimization review.

D. *Interprocess Communications.*

Feed forward message passing was demonstrated to be a heavy load. The issue is the one to many relationship between the LGN processes and the V1 processes. This can be improved by parallelizing the LGN process over several physical nodes, distributing the message output loading across the parallel components.

Lateral messaging is a smaller scale (than feed forward) "one to many" problem: each of the 196 subfields communicates with each of its 8 neighbors. The point to point characteristics of the Pub/Sub message system applied on this model performed well. One round of lateral communication was accomplished within the subfield process time (5.9 ms). Each message size was on the order of 4Kbytes. Approximately 1500 messages are exchanged between the subfields during this time.

The Pub/Sub features of the messaging system accelerated development time because it removed the tedium of having

to know message destinations, and it provided flexibility for adding and removing message connections.

Finally, we speculate on the address space size that a full brain emulation may need using a model at a similar level of abstraction. The number of messages is mostly determined by surface area that can be emulated by a PS3 node. The V1 model uses about 1,700 messages per frame, but this does not include extrastriate connectivity. Based on anatomical data[21] (surface area of V1 and the extrastriate areas it connects to) the extrastriate message connectivity is likely to be between twice (based on area) and ten times (based on axon count) what is needed for V1 afferent and lateral feed combined. Thus a comfortable overestimate for V1 extrastriate connectivity is 17,000 message types. These numbers suggest that a whole brain emulation message space can be subdivided into regions with modest address space (16 bit) capability.

ACKNOWLEDGMENT

This work was supported in part by the United States Air Force Office of Scientific Research (AFOSR), LRIR# 06IF02COR.

REFERENCES

[1] Anderson, James A. The BSB Model: A Simple Nonlinear Autoassociative Neural Network, Associative Neural Memories, Oxford University Press, Inc., New York, Ny, 1993

[2] Black, Paul E. "big-O notation", in *Dictionary of Algorithms and Data Structures* [online], Paul E. Black, ed., U.S. National Institute of Standards and Technology. 11 March 2005

[3] Buzás P, Eysel UT, Adorján P, Kisvárday ZF (2001) Axonal topography of cortical basket cells in relation to orientation, direction, and ocular dominance maps. J Comp Neurol. 2001 Aug 27;437(3):259-85

[4] Dean, T. Hierarchical Expectation Refinement for Learning Generative Perception Models. Tech. rep., Brown University, Providence, Rhode Island, Aug 2005.

[5] Dougherty, R. F., Koch, V. M., Brewer, A. A., Fischer, B., Modersitzki, J., & Wandell, B. A. (2003). Visual field representations and locations of visual areas V1/2/3 in human visual cortex. Journal of Vision, 3(10):1, 586-598, http://journalofvision.org/3/10/1/, doi:10.1167/3.10.1.

[6] George, Dileep and Hawkins, Jeff. A hierarchical Bayesian model of invariant pattern recognition in the visual cortex. In Proceedings of the International Joint Conference on Neural Networks. IEEE, 2005.

[7] Gerstner and Kistler. Spiking Neuron Models. Single Neurons, Populations, Plasticity. Cambridge University Press, 2002

[8] Hecht-Nielsen R., Mechanism of Cognition. In: Bar-Cohen, Y. [Ed.] Biomimetics: Biologically Inspired Technologies, CRC Press, Boca Raton, FL (2006).

[9] Hertz, Lief. Non-Neuronal Cells of the Nervous System: Function and Dysfunction, Volume 31: Part I: Structure, Organization, Development and Regeneration. 2004, Elsevier. ISBN 0-444-51451-1 (3 volumes).

[10] Hubel, D. H. & Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. Journal of Physiology 148, 574-591

[11] C. Koch Biophysics of Computation. Information Processing in Single Neurons, New York: Oxford Univ. Press, 1999, page 87

[12] G. Leuba, R. Kraftsik "Changes in volume, surface estimate, three-dimensional shape and total number of neurons of the human primary visual cortex from midgestation until old age" Anat Embryo1 (1994) 190:351-366

[13] J. S. Lund, A. Angelucci, and P. C. Bressloff Anatomical Substrates for Functional Columns in Macaque Monkey Primary Visual Cortex Cereb Cortex, Journal of Neuroscience, 2003, Vol 8, 1594-1609.

[14] Marr, D. Vision. W. H. Freeman and Co. 1982

[15] Mountcastle, Vernon B. The columnar organization of the neocortex, Brain, Vol. 20 #4, pp701–722, Oxford University Press, April 1997

[16] Peters A, Jones EG, eds (1984) Cerebral cortex.

[17] Reichert, H., Rowell, C. H. F. and Gris, C. (1985). Course correction circuitry translates feature detection into behavioural action in locusts. Nature, Lond. 315, 142-144.

[18] Tovée, Martin J. (1996), An introduction to the visual system, Cambridge University Press, ISBN 0-521-48339-5 (References, pp.180-198. Index, pp.199-202. 202 pages.)

[19] Wielaard, J. Sajda, P. Neural mechanisms of contrast dependent receptive field size in V1. Advances in Neural Information Processing Systems, 2006, Numb 18, pages 1505-1.

[20] Zolta´ N F. Kisva´ Rday , Ales S. Ferecsko´ , Krisztina Kova´ CS, PE´ Ter Buza´ S, Julian M. L. Budd and lft. Eysel. One axon-multiple functions: Specificity of lateral inhibitory connections by large basket cells. Journal of Neurocytology 31, 255–264 (2002)

[21] Budd Julian ML. Extrastriate feedback to primary visual cortex in primates: A quantitative analysis of connectivity. Proc. R. Soc. Lond. B. June 7, 1998 265:1037-1044; doi:10.1098/rspb.1998.0396.

[22] Designing and Building Parallel Programs, by Ian Foster. Addison-Wesley (ISBN 9780201575941). 1995. Section 1.4.4

[23] Daugman, John: "Complete Discrete 2-D Gabor Transforms by Neural Networks for Image Analysis and Compression", IEEE Trans on Acoustics, Speech, and Signal Processing. Vol. 36. No. 7. July 1988, p.1169-1179

Approved for Public Release, 88AWB-2010-0159