A Bird's-Eye View of PetaVision, the World's First Petaflop/s Neural Simulation*

Portland State UNIVERSITY

Dan Coates

Portland State University,
Maseeh College of Engineering and Computer Science,
Portland OR

Los Alamos NATIONAL LABORATORY EST. 1943

Garrett Kenyon,
Craig Rasmussen

Los Alamos National Laboratory,
Los Alamos, NM

Parallel Implementations of Learning Algorithms:
"What Have You Done For Me Lately?"

NIPS08
Whistler, BC
December 13, 2008

Neural Information Processing Systems Conference

* The authors acknowledge the support of the National Science Foundation, under a grant administered by the New Mexico Consortium

PetaVision Project at LANL

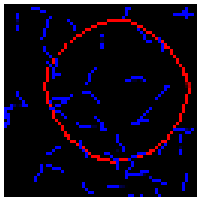

Goal: Achieve human-level performance in a “synthetic visual cognition” system

On: IBM/DOE Roadrunner petascale supercomputer (or a multicore PC)

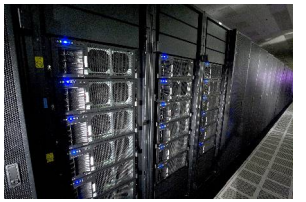

Running: A spiking LIF neural network inspired by visual cortex.

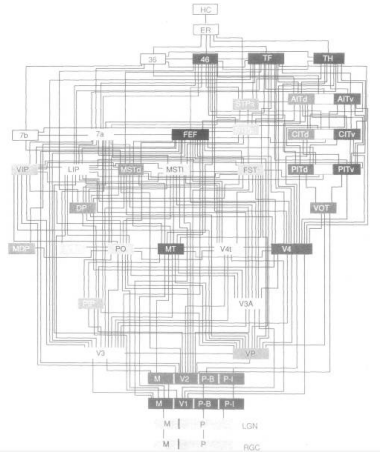

What level of abstraction?

Emulate the cortical circuits for mid/low-level visual processing.

We model the gross architecture of visual cortex, trying not to violate proven neural science.

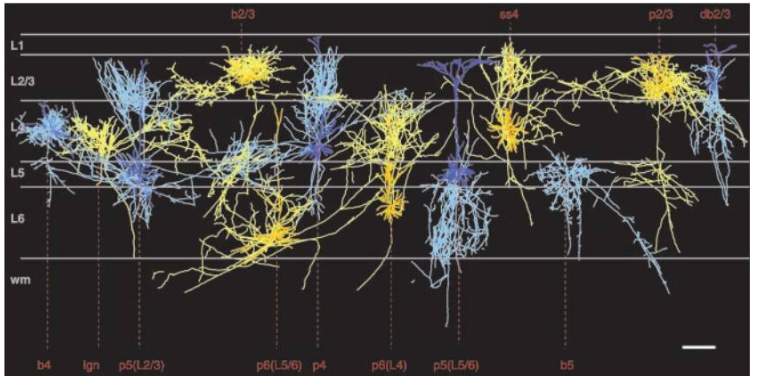

Binzegger, et. al.

What are the crucial features of V1?

Retinotopic mapping.

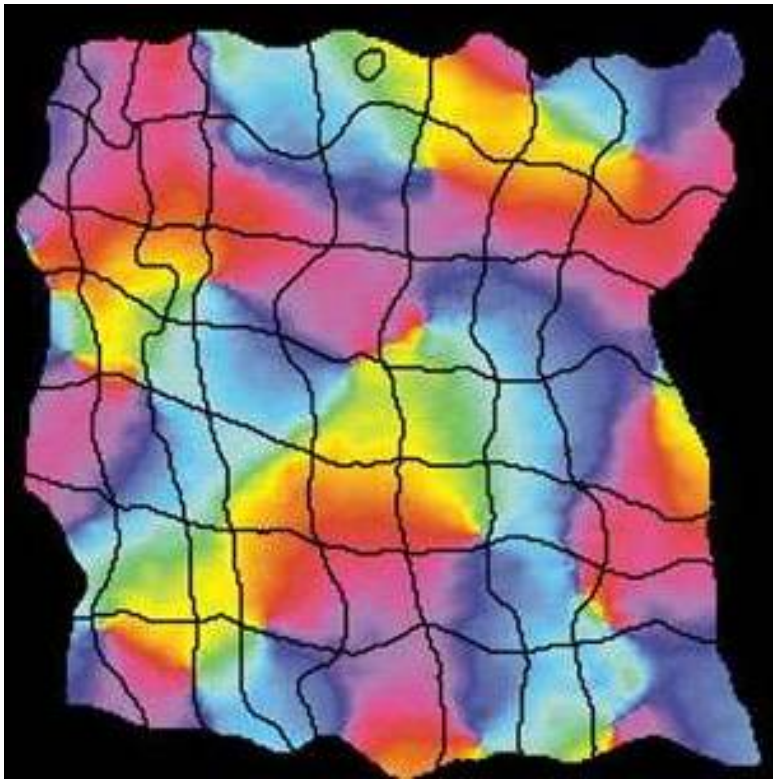

Edge detectors
of Hubel & Wiesel

Distinct laminar neural
populations.

What are the crucial features of V1?

Retinotopic mapping.

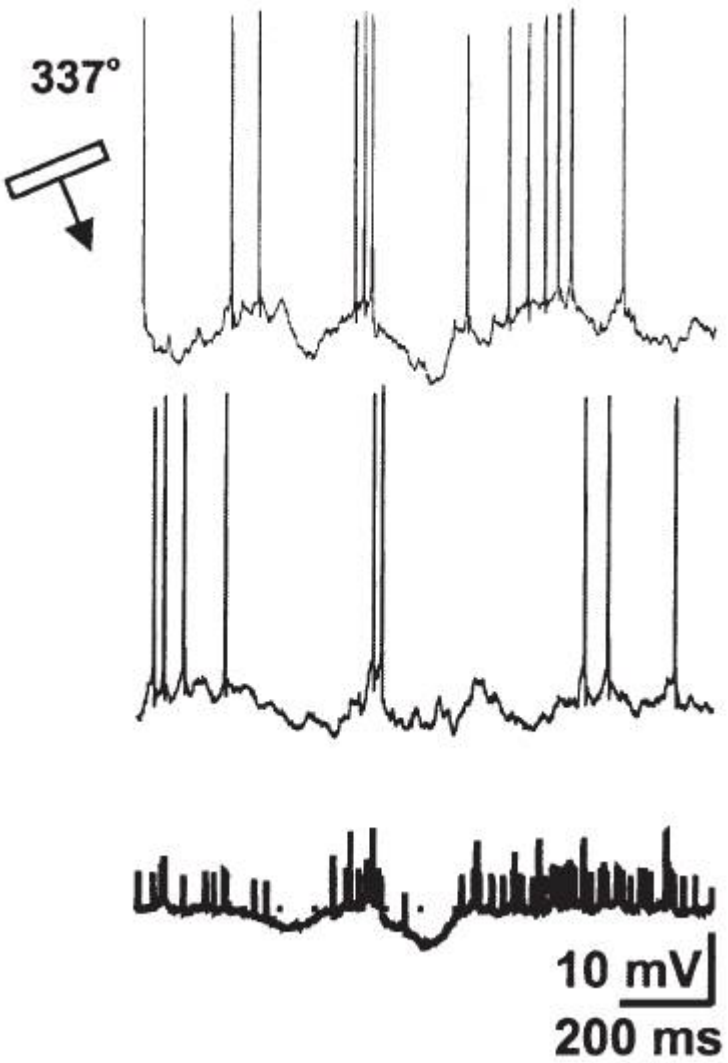

Edge detectors of Hubel & Wiesel

Distinct laminar neural populations.

What are the crucial features of V1?

Retinotopic mapping.

Edge detectors
of Hubel & Wiesel

Distinct laminar neural populations.

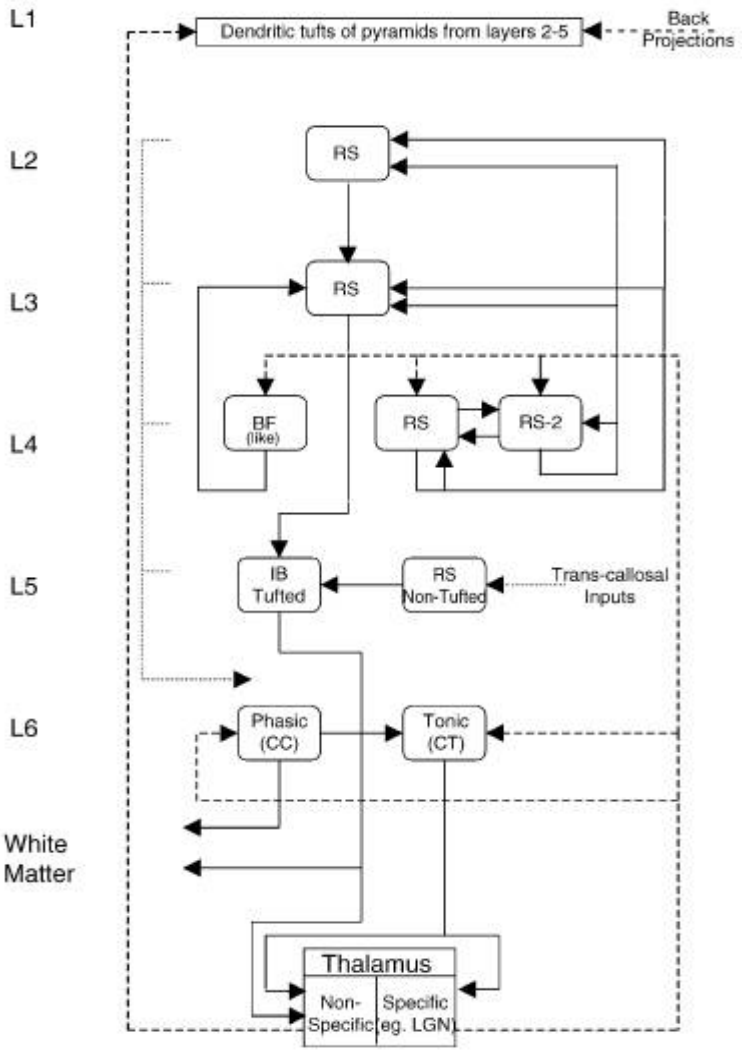

Bannister. Laminar circuit.

What are the elements, and how does that help us?

Spiking neurons and *specific* connectivity

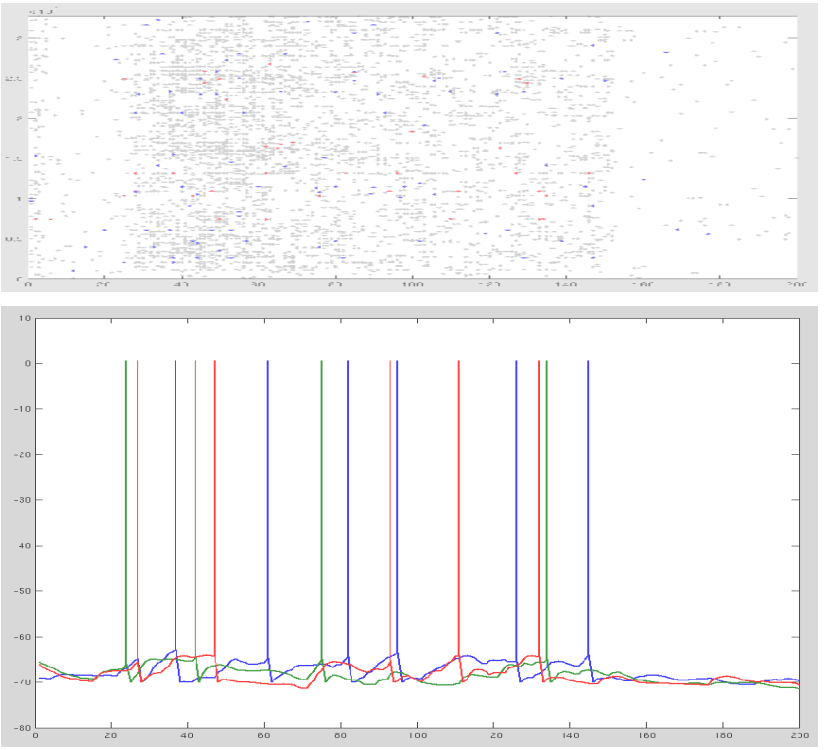

- efficient, possibly asynchronous operation
- sparse inter-node communication

What are the elements, and how does that help us?

Spiking neurons and ***specific*** **connectivity**

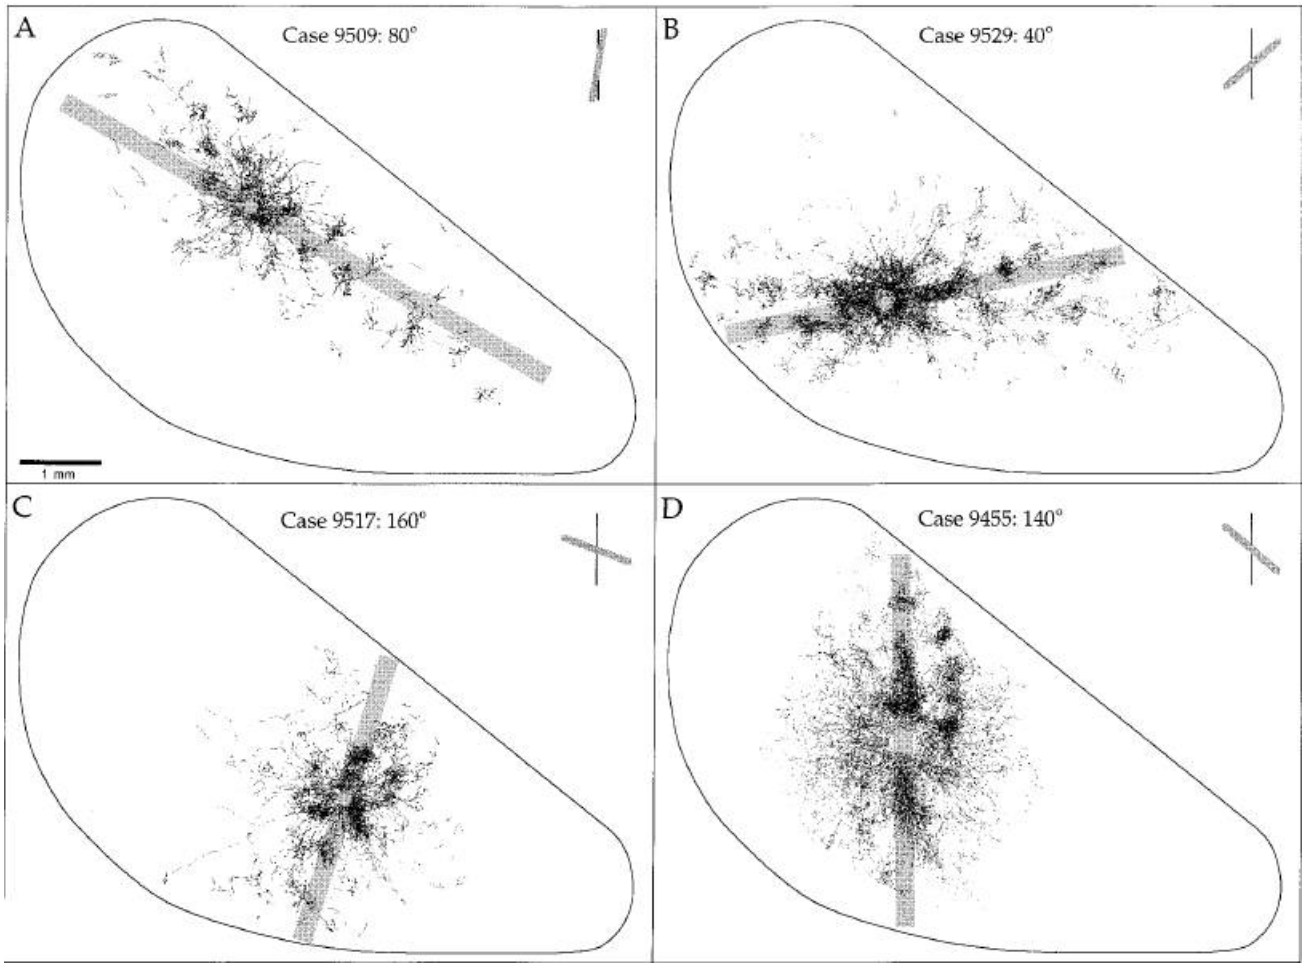

- connections are primarily local
- function inherent in wiring

Bosking, et al. “Patchy” connectivity expresses orientation preference of horizontal connections.

Example: edge detection

V1 simple cells have been shown to respond like a Gabor functions. We use 8 orientations.

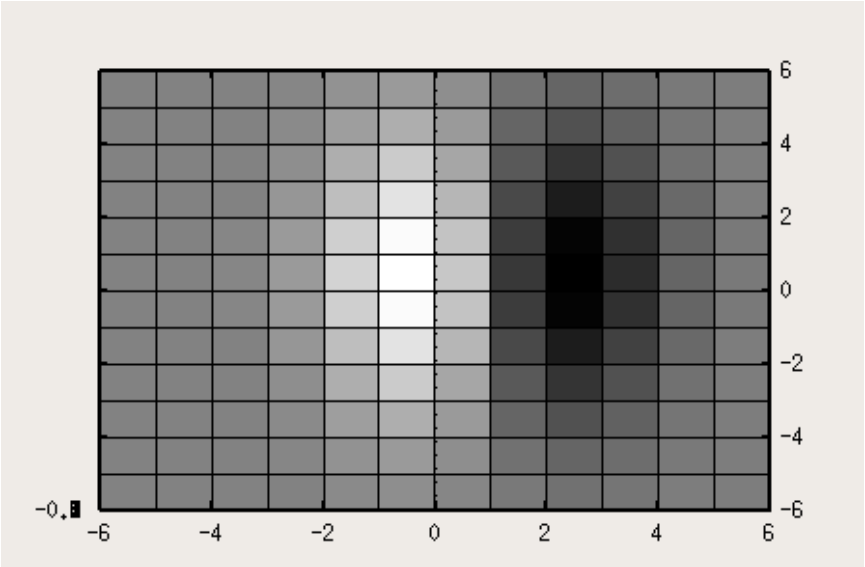

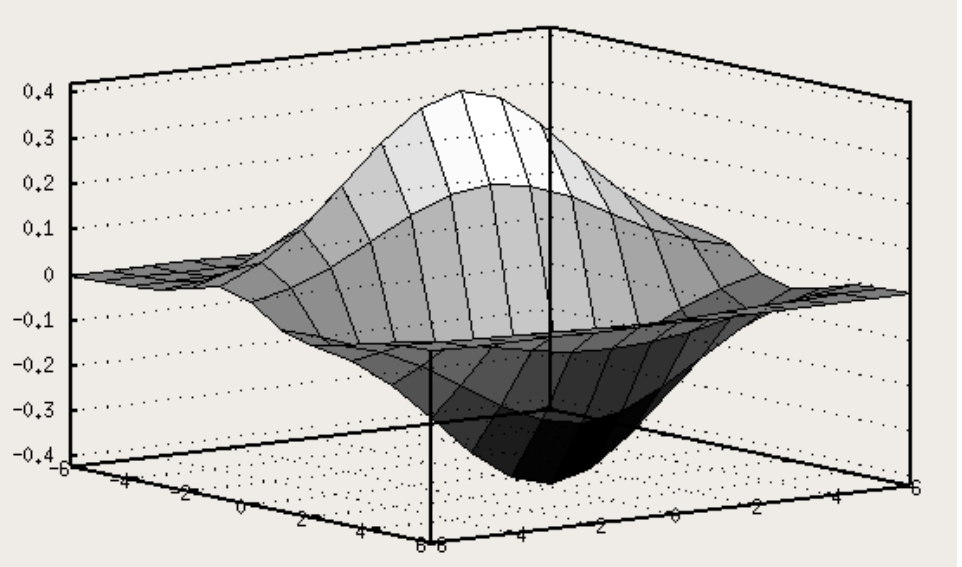

Beyond edges: long-range association field

Ben-Shahar and Zucker have proposed additional connectivity patterns formalized using differential geometry. [Neural Computation, 2004]

Besides curve integration, such a scheme could also be used for shape-from-shading and natural texture identification.

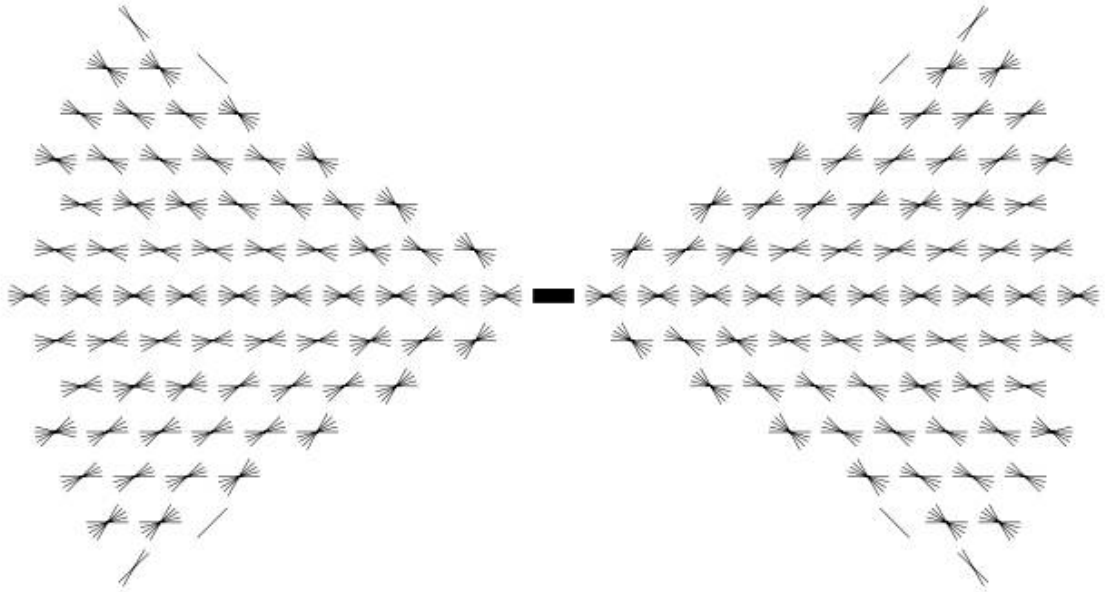

Beyond edges: long-range association field

Ben-Shahar and Zucker have proposed additional connectivity patterns formalized using differential geometry. [Neural Computation, 2004]

Besides curve integration, such a scheme could also be used for shape-from-shading and natural texture identification.

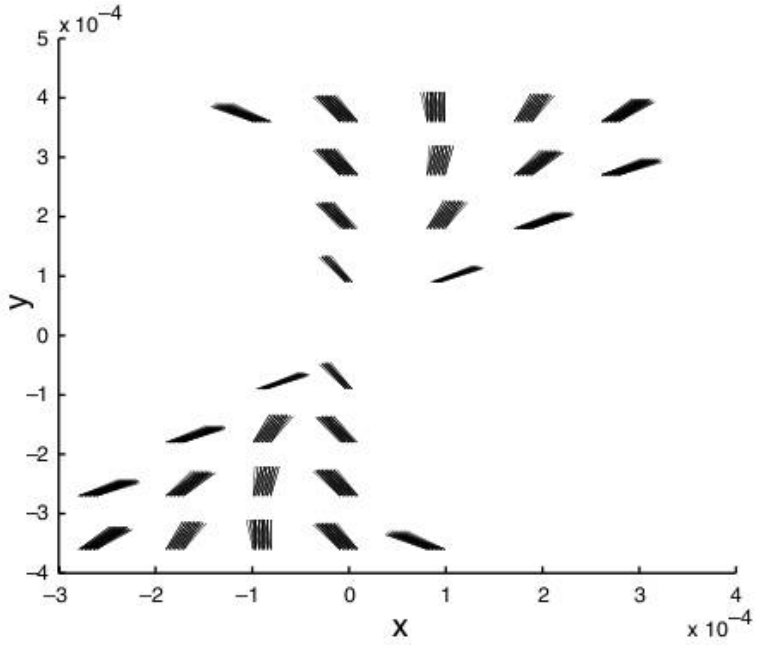

Beyond edges: long-range association field

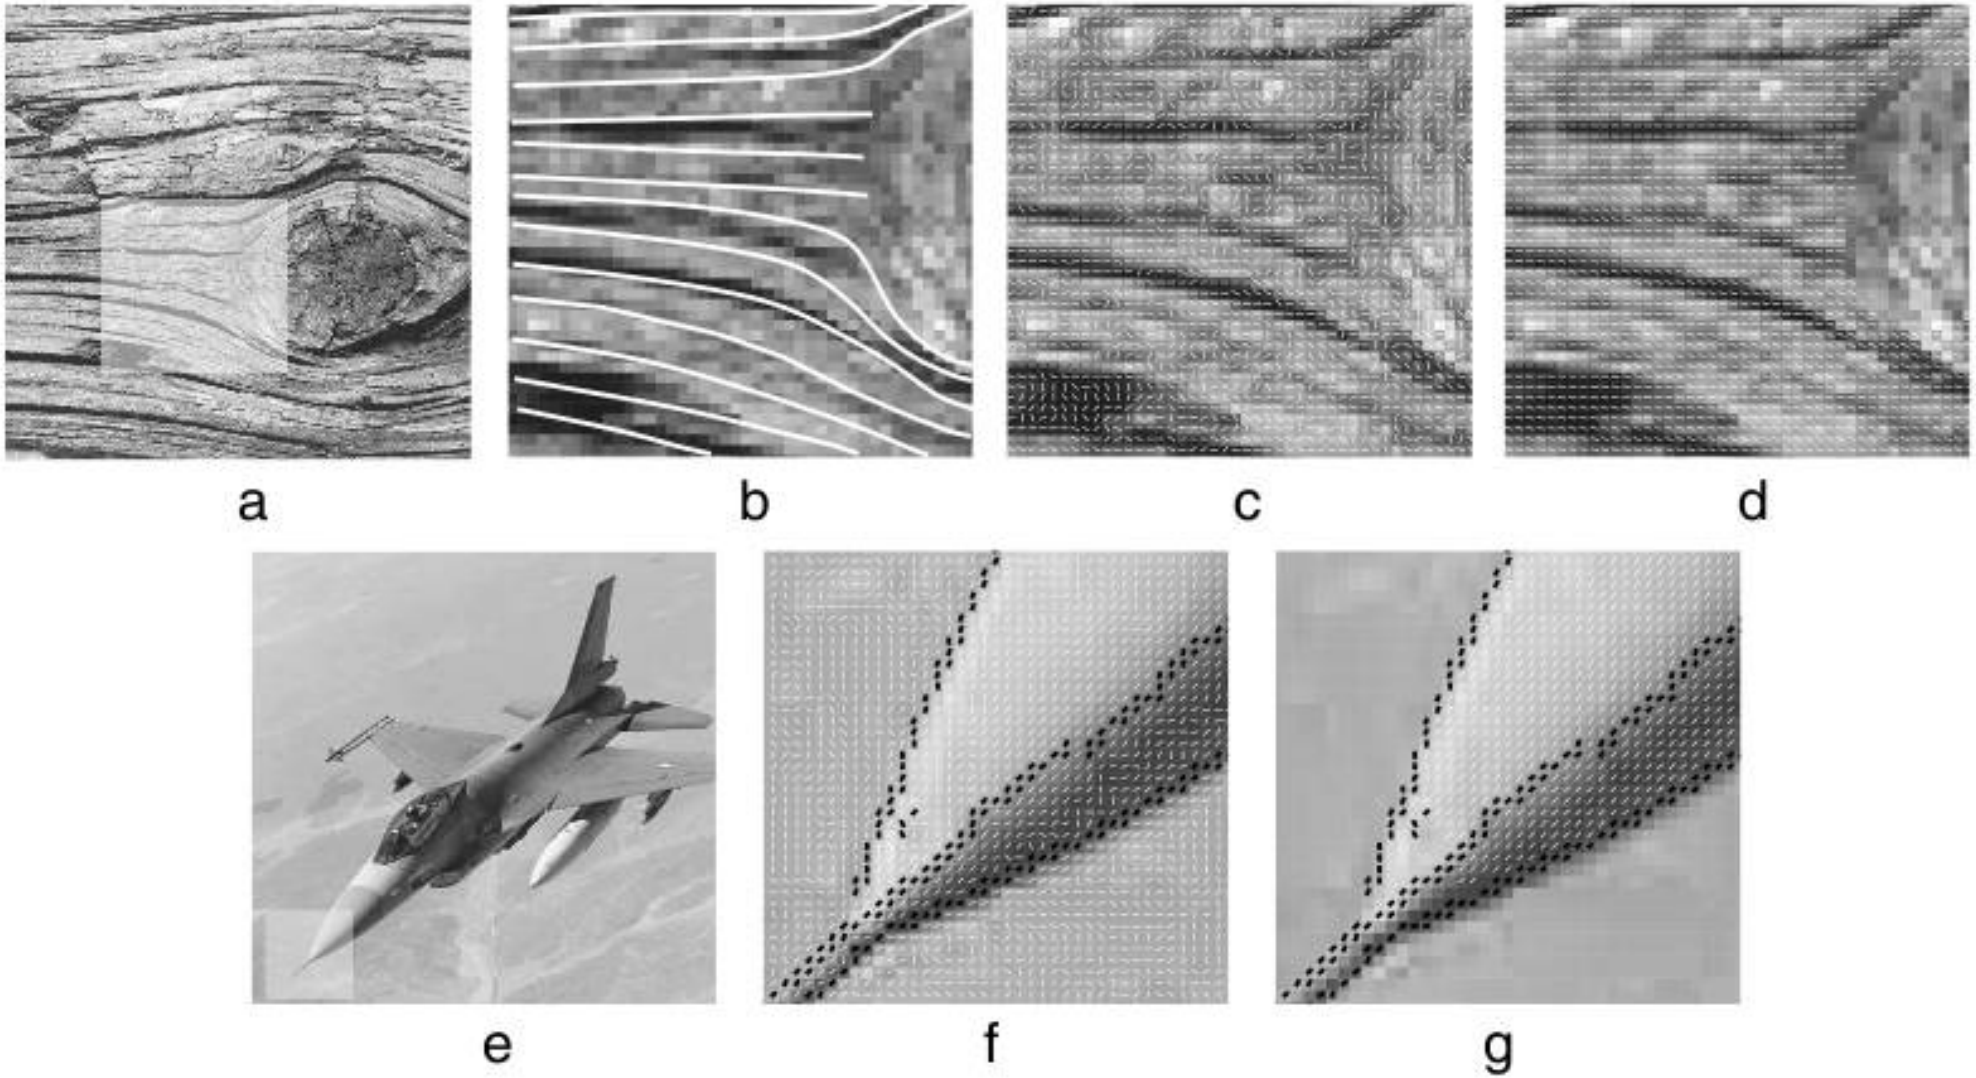

Summary of Biological Inspiration

Network structure for a computer vision system can be modeled after architecture of mammalian visual cortex.

There are analytic correlates of these techniques, although closed-form derivations are difficult.

Note: these connections have been shown to be **learnable**, although we hard-code as mathematical functions.

Implementation: software abstractions

PVLayer: Population of neurons. Retina, LIF.

PVConnection: Connectivity pattern, represented by a mathematical weight function. Anything-to-anything routing possible.

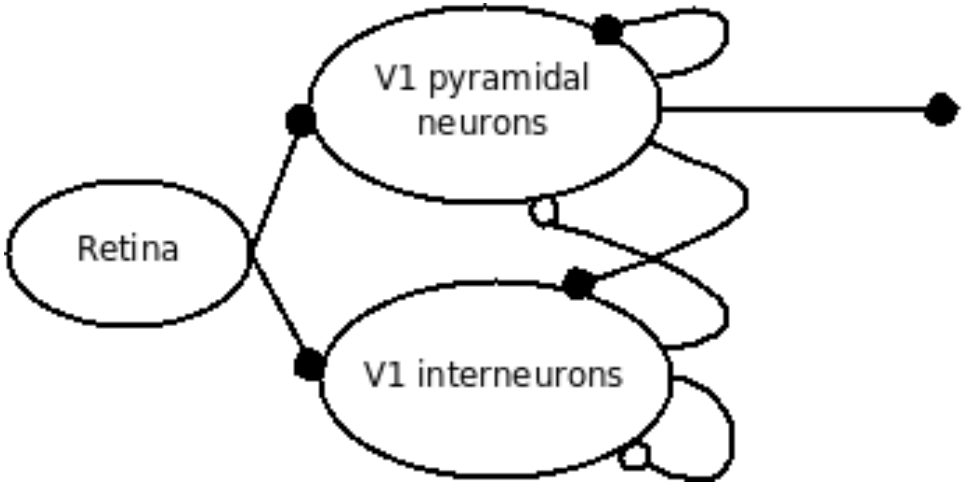

Implementation: LIF

$$\frac{dV_m(t)}{dt} = \sum g_{ex}(t - t_i)(V_m(t) - E_{ex}) + \sum g_{inh}(t - t_i)(V_m(t) - E_{inh}) + g_{leak}(V_m(t) - E_{leak})$$

Implementation: PVConnection

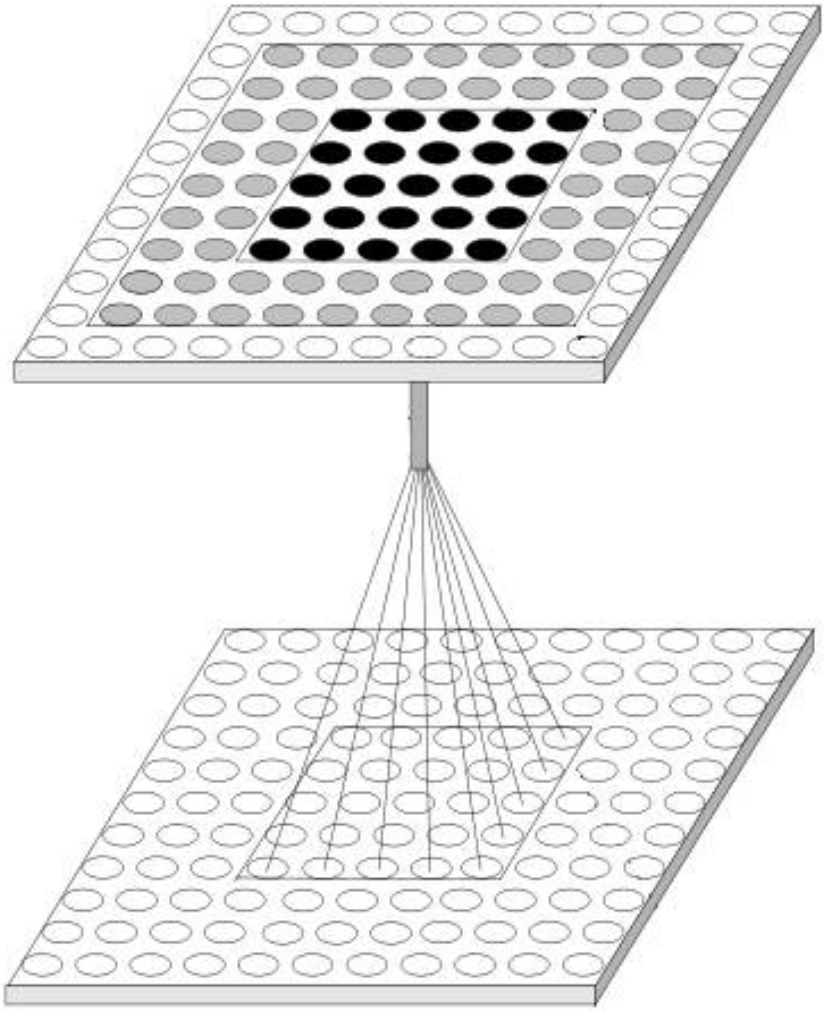

Connection kernels are translation-invariant.

Implementation: Parallel Algorithm

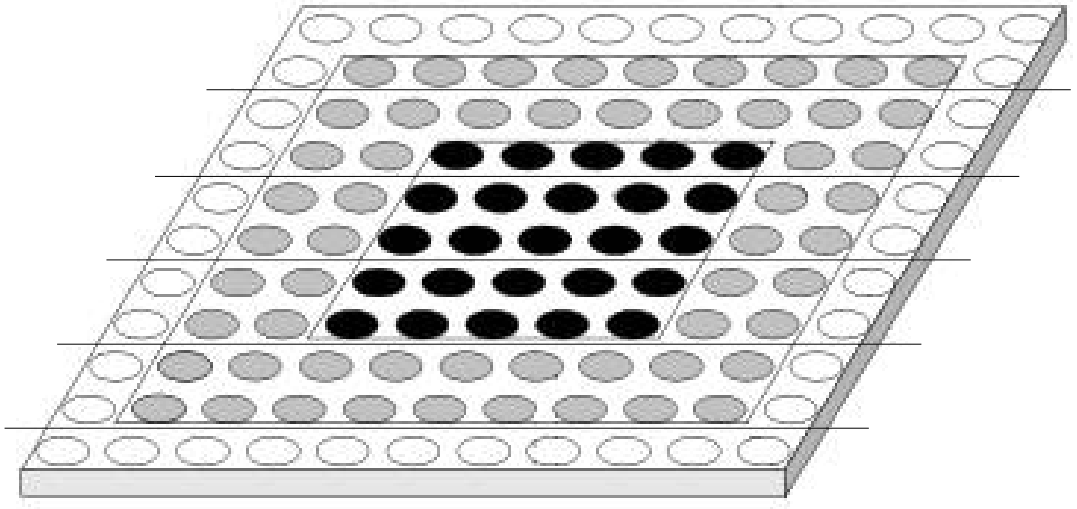

```
Process each PVConnection:
    Process each presynaptic event
        Update effected postsynaptic neurons

Update each layer

Perform I/O
```

How to interpret results?

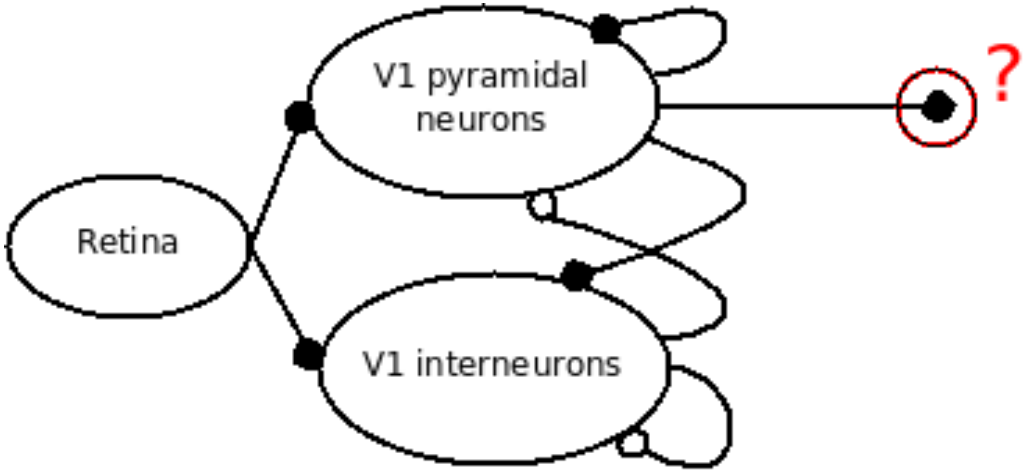

Readout: Spike trains are post-processed for firing rate. Temporal correlations such as synchrony and oscillatory power are also measured.

Roadrunner

Roadrunner “core”

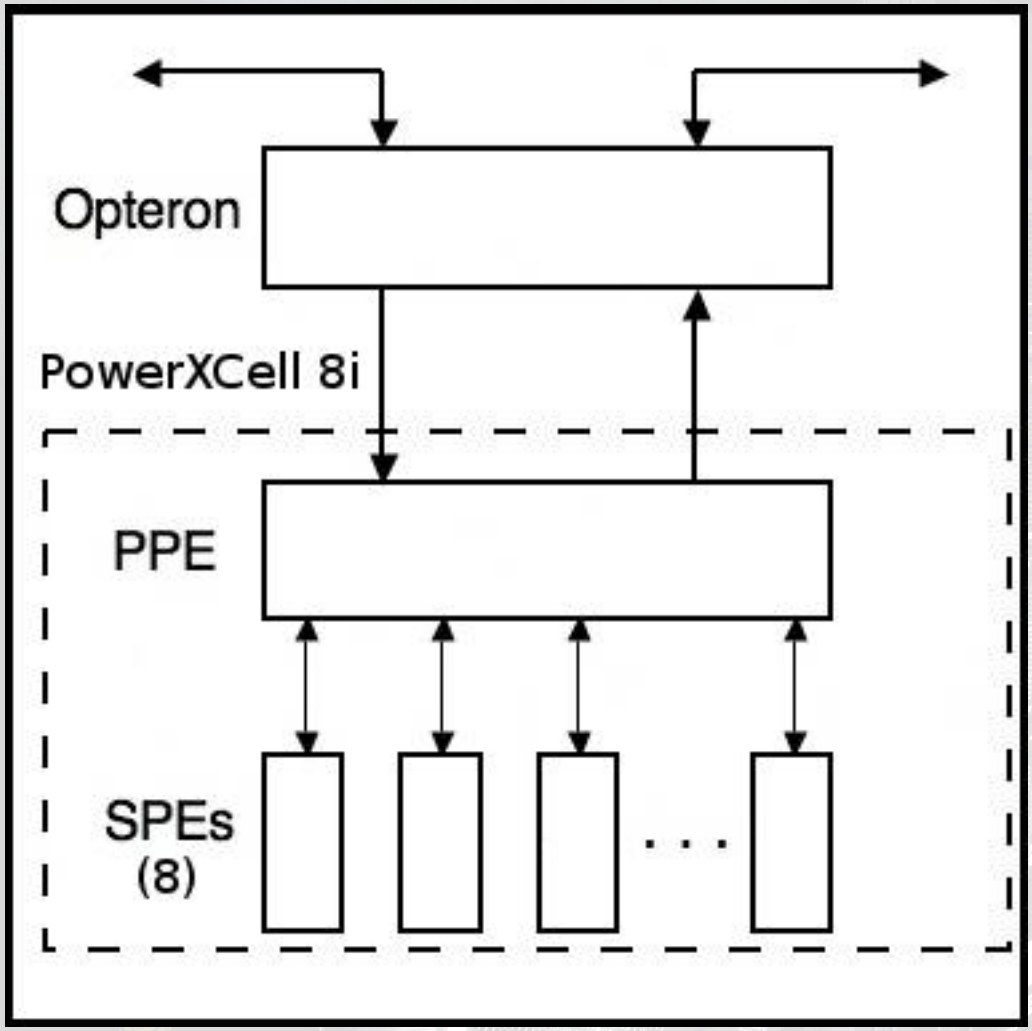

1 Opteron Core
- 1.8 Ghz
- 4 GB DDR2

IBM PowerXCell 8i
- 3 Ghz clock speed
- 200 Gflops w/single precision & pipelining
- 4 GB DDR2
- SPEs = 256k

Roadrunner node: triblade

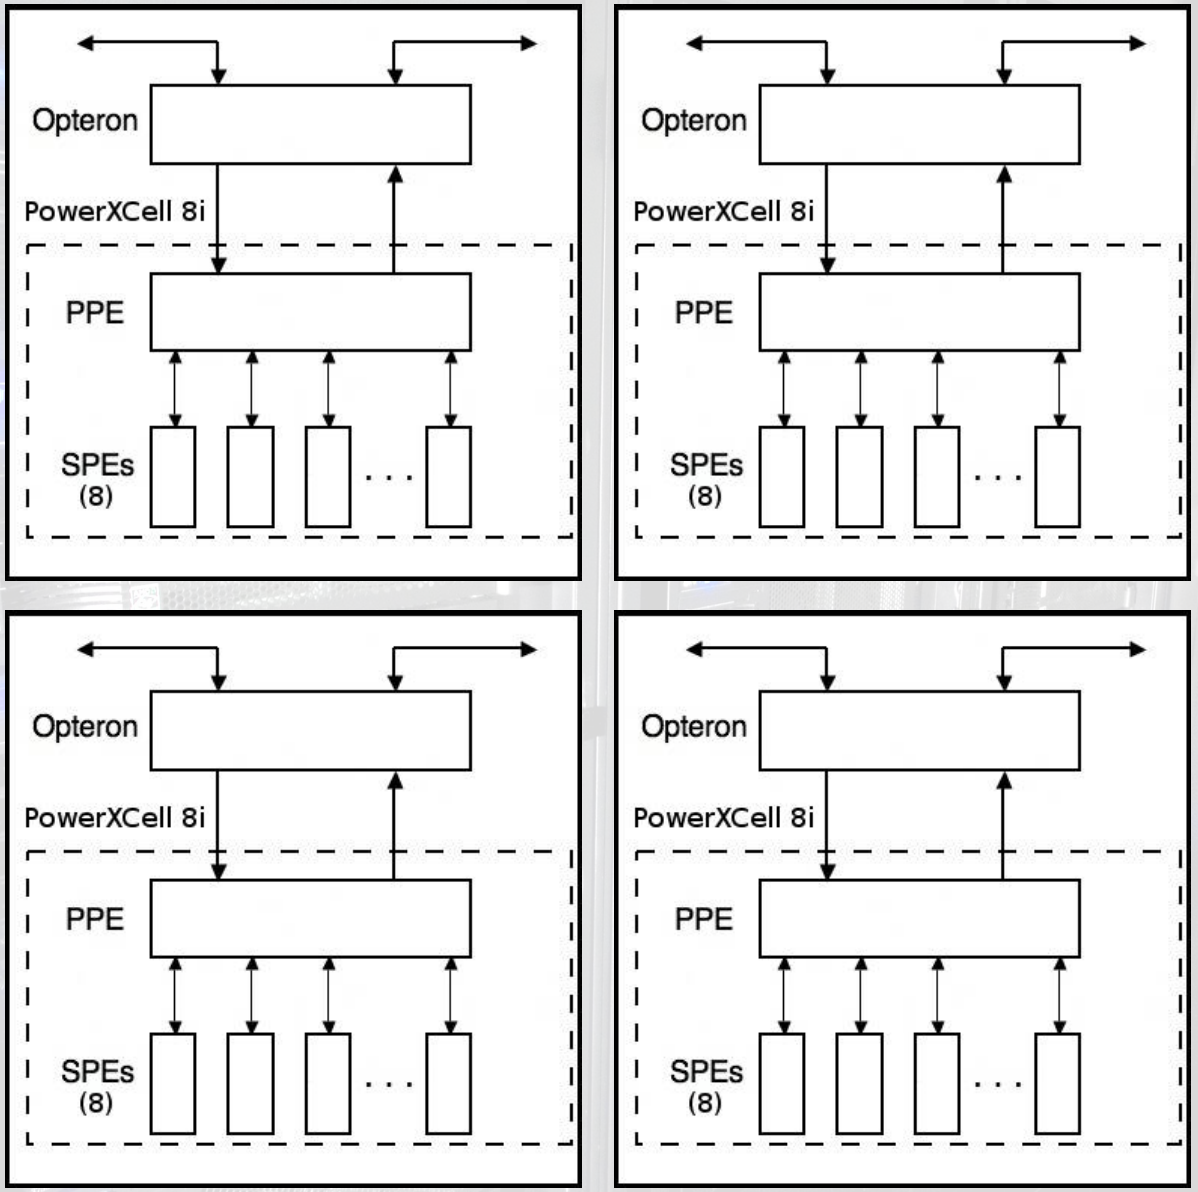

Roadrunner

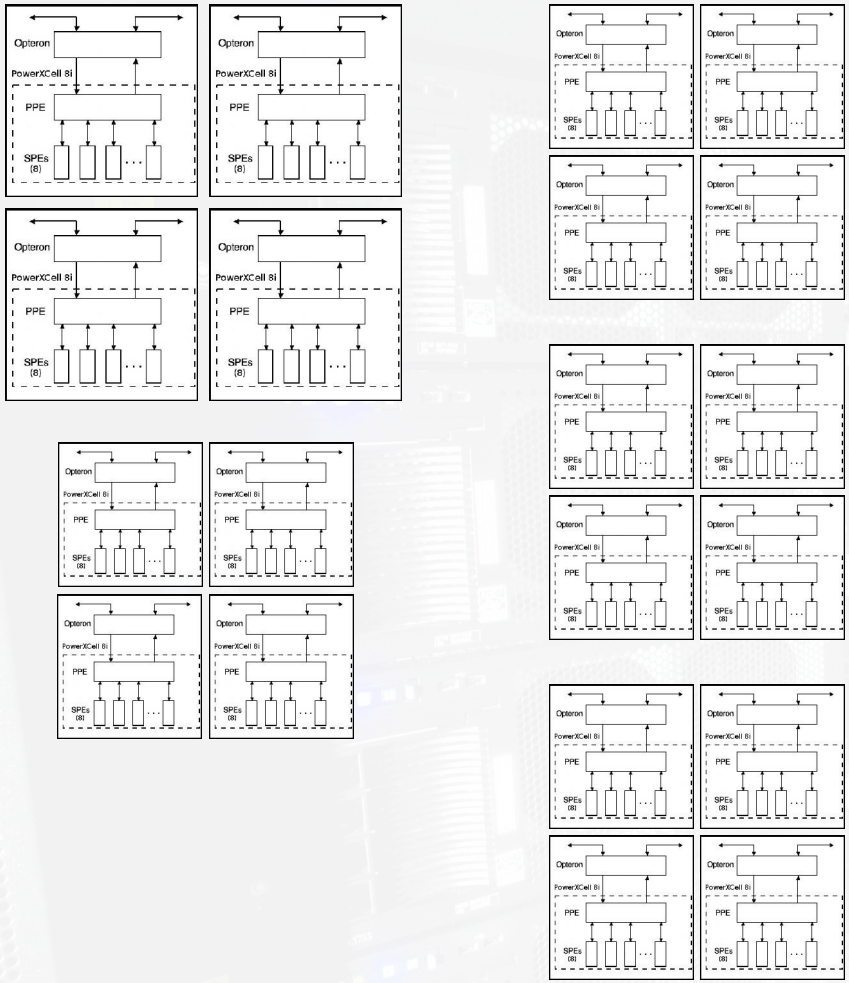

3,240 nodes:
- 2 Opteron dual-cores 1.8Ghz,16 GB memory
- 4 PowerXCell 8i

Infiniband connections

Peak system performance: ~1.7 Petaflop/s.

PetaVision SPMD on Roadrunner

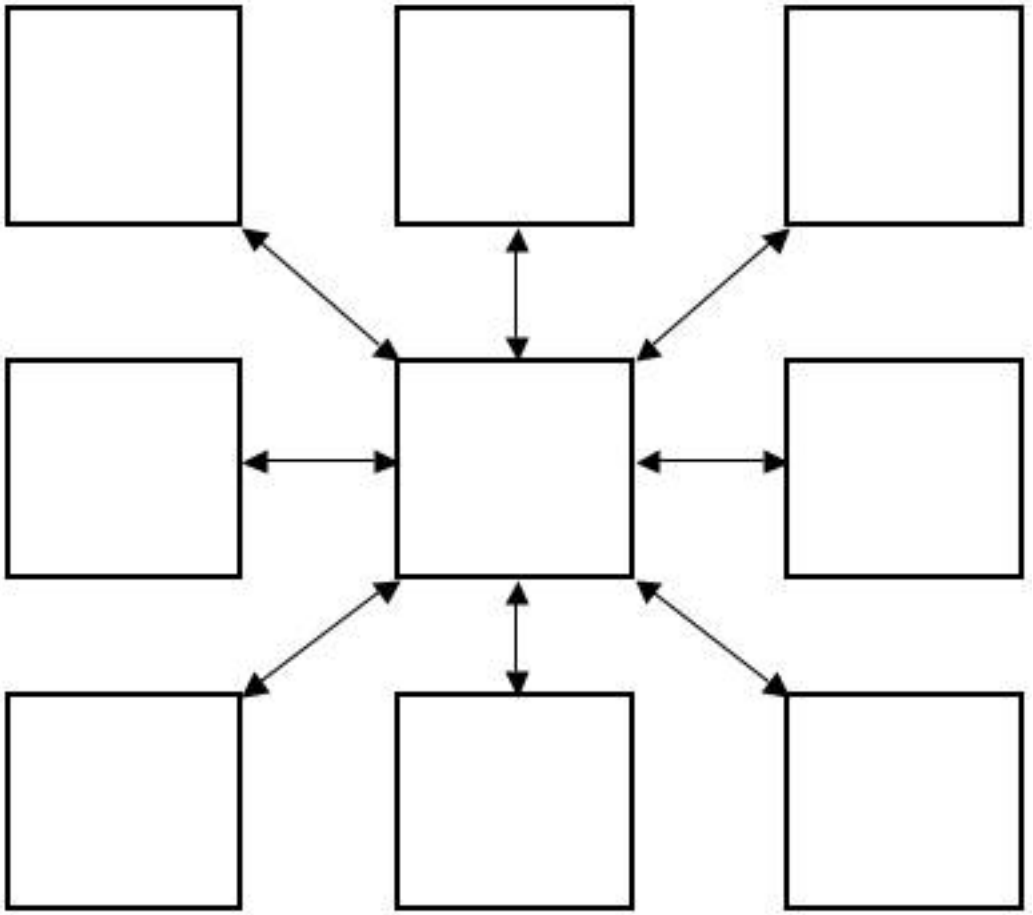

Each node handles an image patch.

PetaVision SPMD on Roadrunner

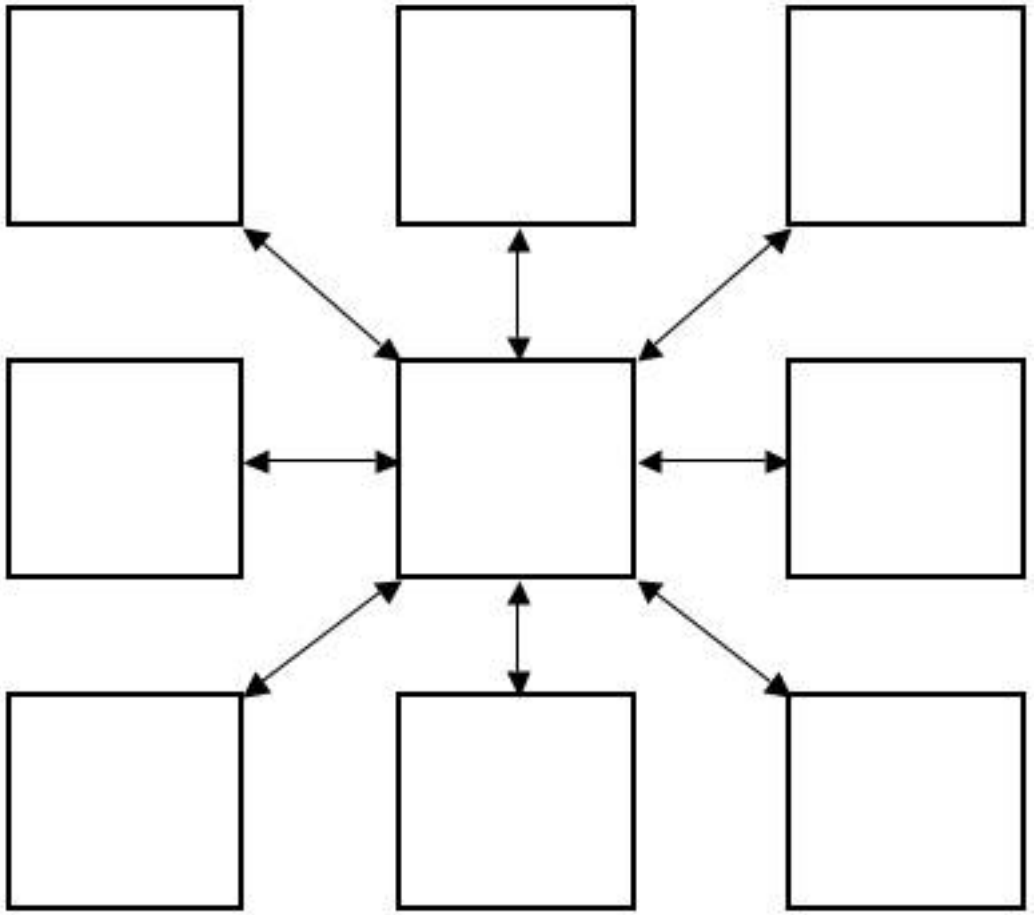

Process local activity

Process remote activity

Update layer

Send output spikes

Roadrunner SPMD Components

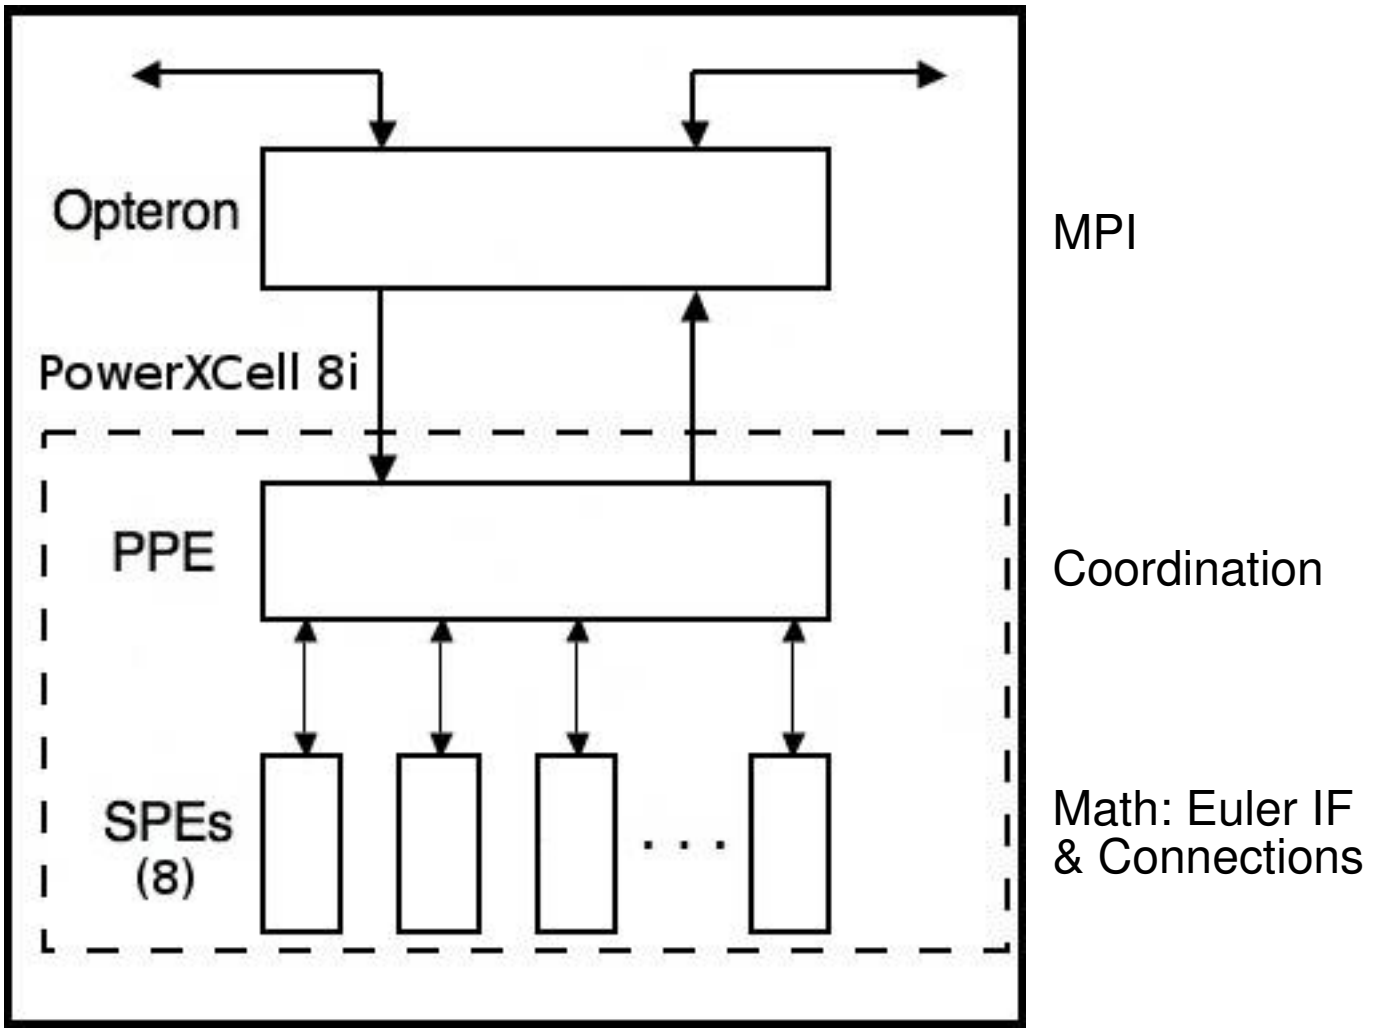

Visual task

“Closed contour present?”

- No need for higher-level knowledge

- Nontrivial

- Humans can solve effortlessly. (psychophysics)

Visual task - results

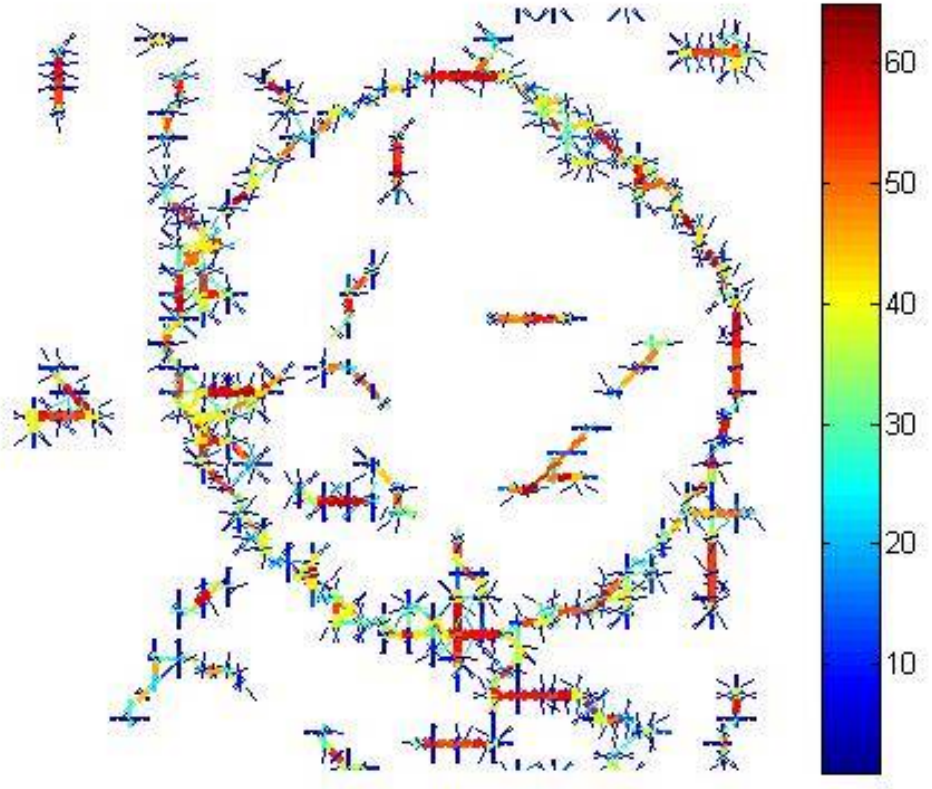

Prototypical network response.
Color represents average firing rate.

Thank you!

Questions?