Blue
Brain
Project

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE
CENTER FOR
BRAIN SIMULATION

Simulating Morphologically Detailed Neuronal Networks at Extreme Scale

Aleksandr Ovcharenko
(aleksandr.ovcharenko@epfl.ch)

Pramod Kumbhar, Michael Hines, Francesco Cremonesi, Timothee Ewart, Stuart Yates, Felix Schuermann and Fabien Delalondre

Blue Brain Project Approach

Constructing virtual brain model by reverse engineering biological components

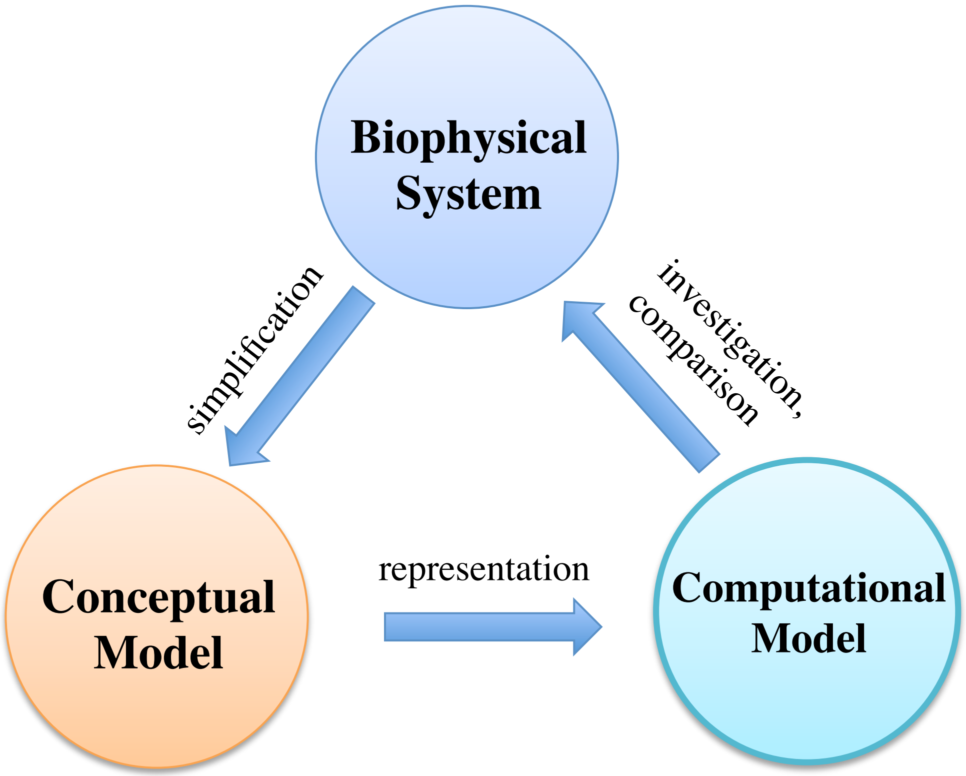

Different Scale … Different Simulators …

Blue Brain Project

Nest

nest:: initiative

A simulator for spiking neural network models that focuses on dynamics, size, structure

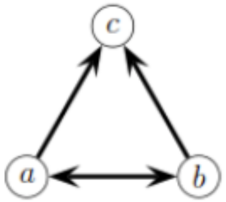

Blue Brain Project

Different Scale … Different Simulators …

Nest

nest:: initiative

A simulator for spiking neural network models that focuses on dynamics, size, structure

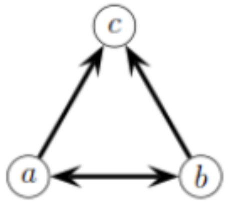

Neuron

A simulator for cells with complex anatomical and biophysical properties.

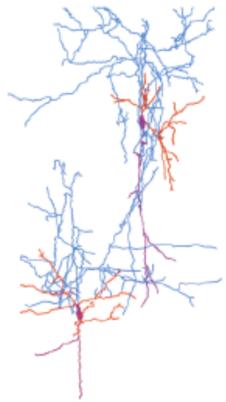

Blue Brain Project

Different Scale … Different Simulators …

Nest

nest:: initiative

A simulator for spiking neural network models that focuses on dynamics, size, structure

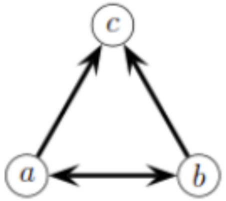

Neuron

A simulator for cells with complex anatomical and biophysical properties.

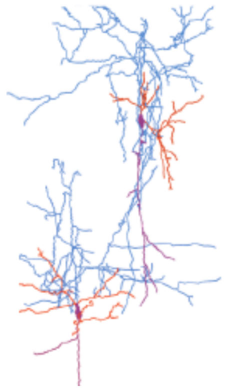

STEPS

A simulator for detailed models of neuronal signaling pathways at molecular level

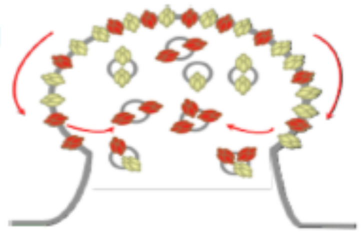

Blue Brain Project

Different Scale of Neuronal Simulators

Nest

nest:: initiative

A simulator for spiking neural network models that focuses on dynamics, size, structure

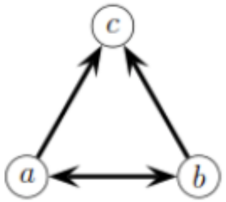

Neuron

A simulator for cells with complex anatomical and biophysical properties.

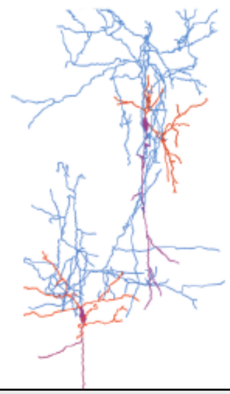

STEPS

A simulator for detailed models of neuronal signaling pathways at molecular level

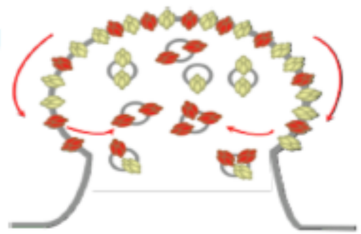

Model Reconstruction

Blue Brain Project

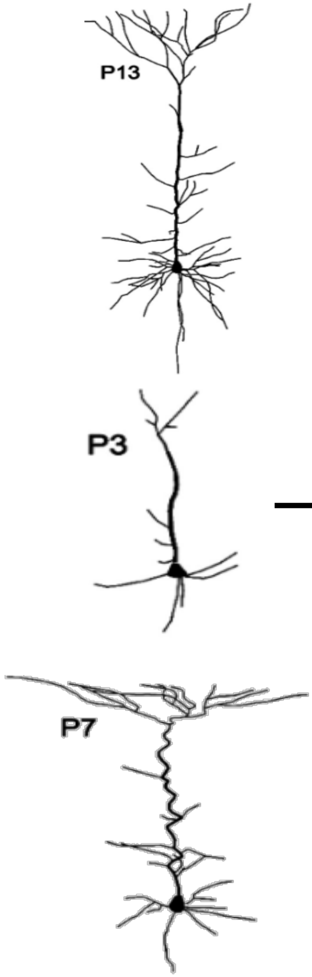

Morphologically detailed neurons

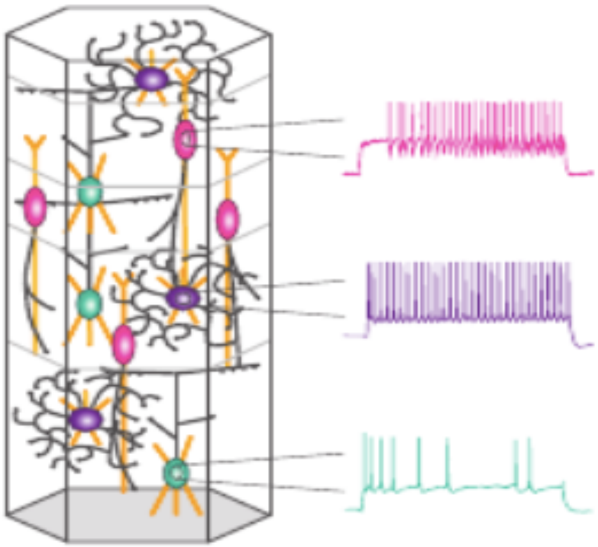

Neuronal circuit and membrane voltage activity

NEURON Modeling

Blue Brain Project

400 Compartments / Neuron

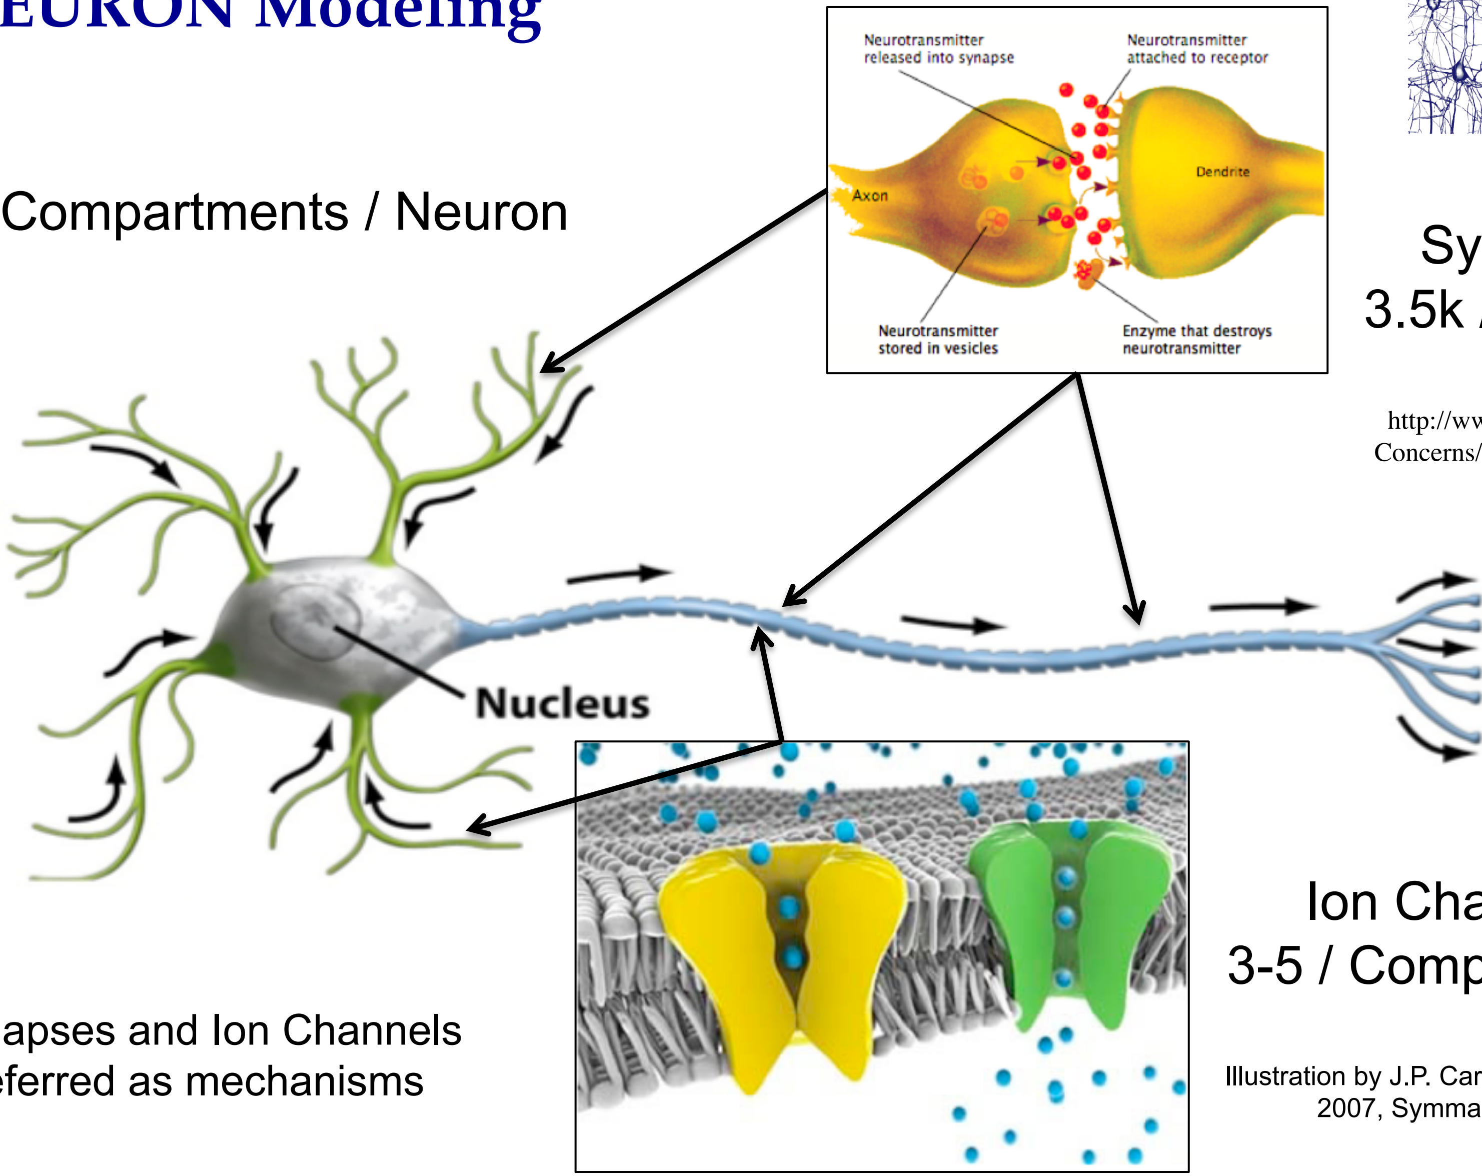

Synapse
3.5k / Neuron

http://www.sailhome.org/Concerns/Excitotoxins.html

Ion Channel
3-5 / Compartment

Illustration by J.P. Cartailler. Copyright 2007, Symmation LLC

Synapses and Ion Channels referred as mechanisms

Need to solve O(5k) non-linear mechanisms to assemble O(400) dof 3-diagonal sparse matrix

NEURON Data Structures

Blue Brain Project

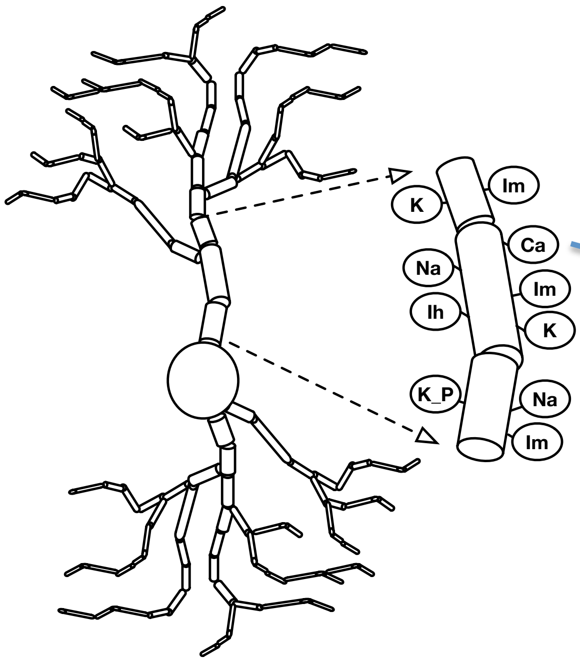

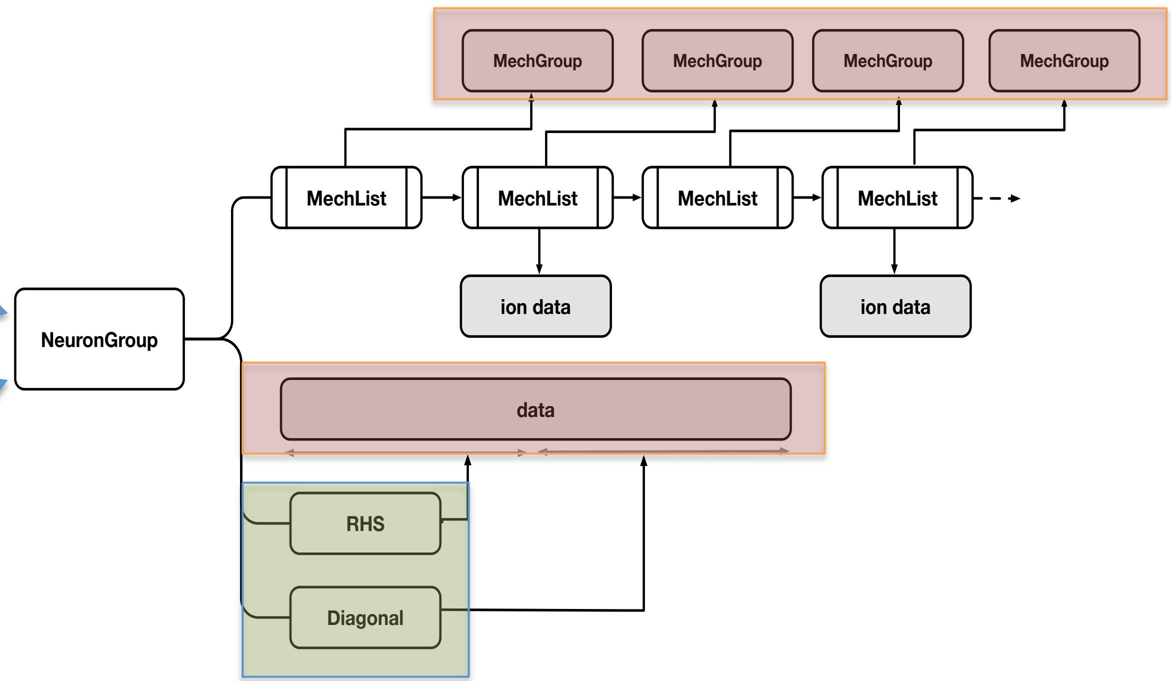

Biologist view:
compartment model

Memory View: In memory
representation of neurons

NModl DSL to .C to Support Scientists

Blue Brain Project

```
DERIVATIVE states {
     LOCAL mAlpha, mBeta, mInf, mTau, lv, qt

     qt = 2.952882641412121
     lv = v

     if(lv == -32){
          lv = lv+0.0001
     }

     mAlpha = mAlphaf(lv)
     mBeta = mBetaf(lv)
     mInf = mAlpha/(mAlpha+mBeta)
     mTau = (1/(mAlpha+mBeta))/qt
     m' = (mInf-m)/mTau

     v = lv
}
```

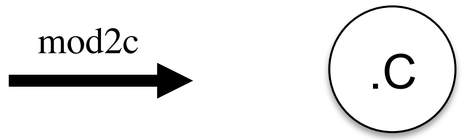

Scientists use domain specific language NModl

Each .mod mechanism file is converted to .c

NEURON Workflow

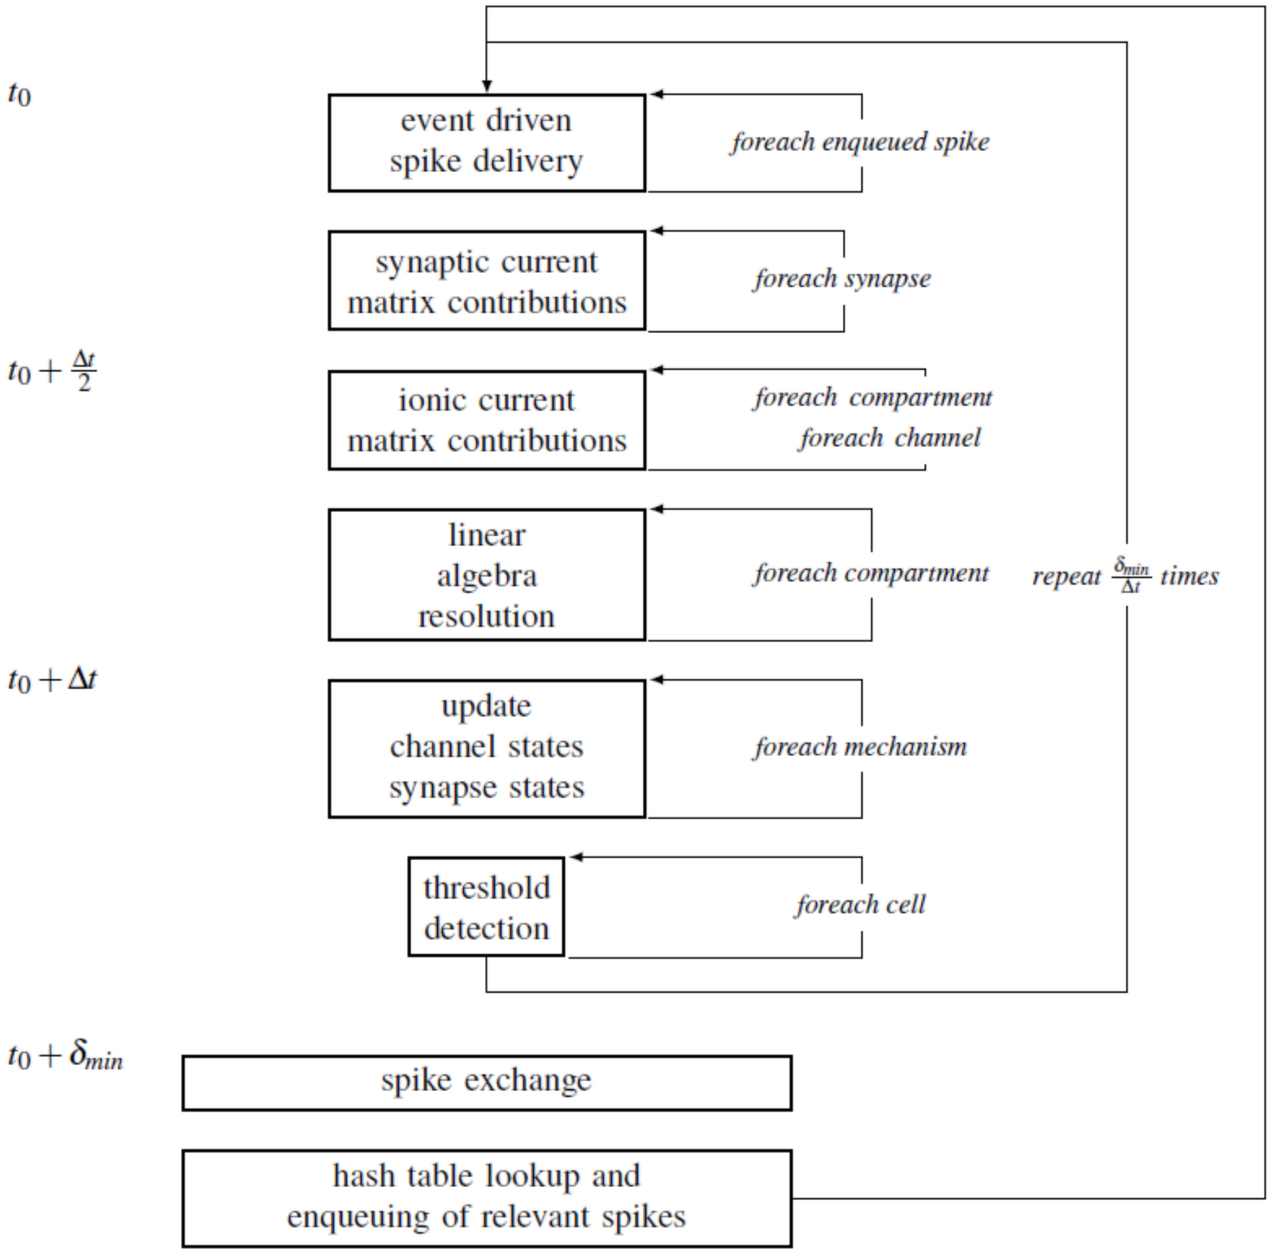

Static Load Balancing Workflow

Blue Brain Project

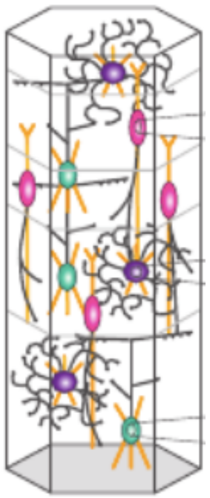

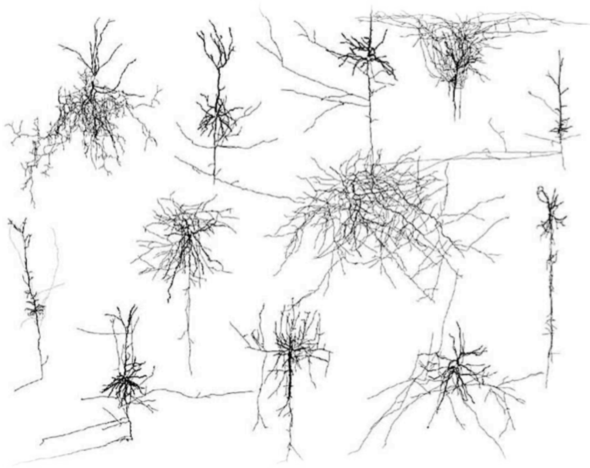

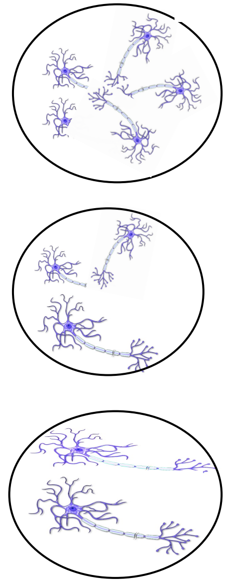

Circuit Building

Neuronal problem domain created by neuroscientists

Cell Computational Complexity

LPT algorithm calculates number of compartments & channels and their computational complexity

Neuron Groups

Construct neuron groups based on complexity factors

Less than 2% load imbalance on IBM Blue Gene/Q

Going Further in Scale

Blue Brain Project

Problem size and memory requirements for future simulations

parameters / brain type	rat	monkey	human
number of neurons	1×10^{8}	1×10^{9}	8×10^{10}
number of synapses	5×10^{11}	1×10^{13}	1×10^{15}
number of state variables	3.3×10^{12}	6.3×10^{13}	6×10^{15}
estimated size in memory	100 TiB	1 PiB	80 PiB

NEURON/Neurodamus: 12 MiB/Neuron

Blue Brain Project

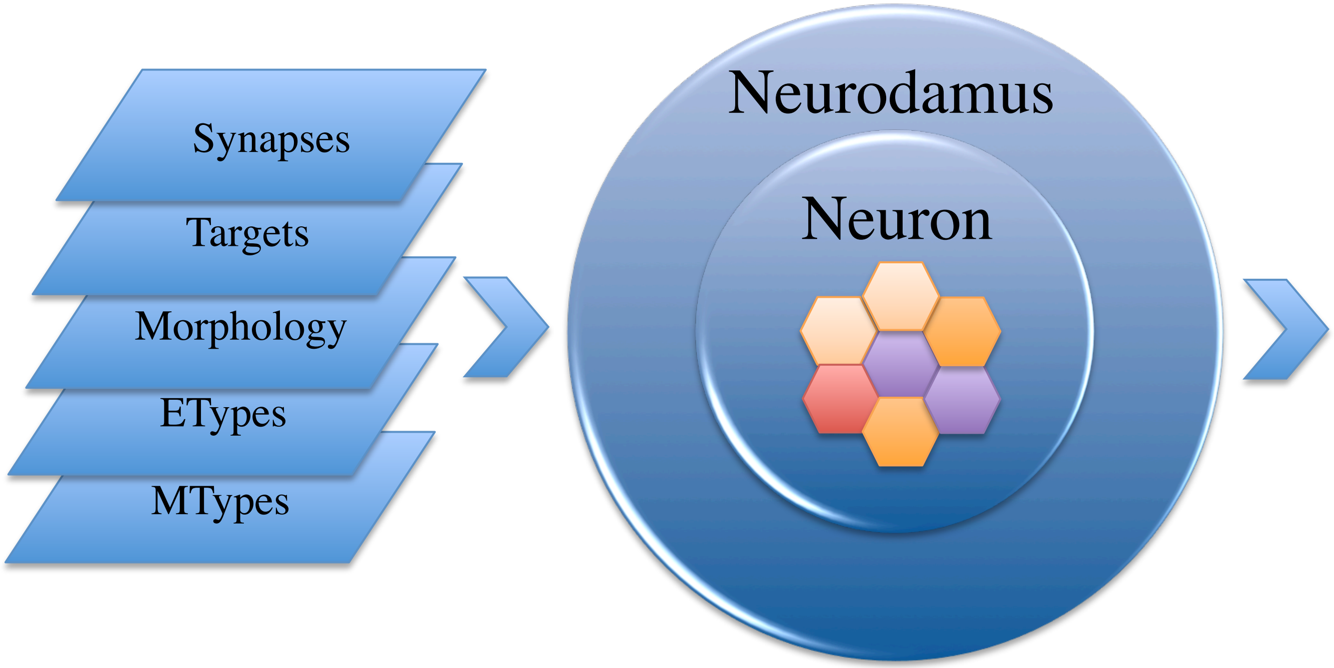

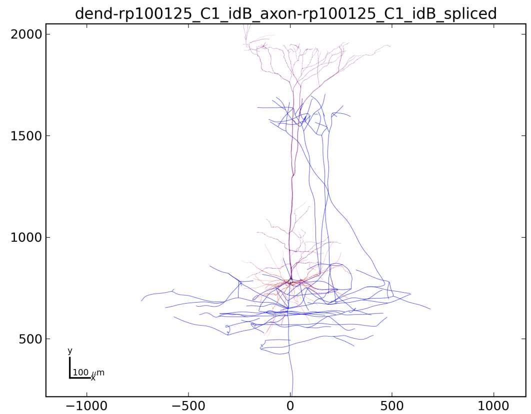

Fails at 756k neurons
Every process stores 900 MiB of global information

2008	2010	2012
IBM Blue Gene/L	IBM Blue Gene/P	IBM Blue Gene/Q
8,192 cores	65,536 cores	65,536 cores
10k neurons	217k neurons	**756k neurons**

CoreNeuron Development Decision

Blue
Brain
Project

Why?

- Simulate bigger models
- Decrease time to solution
- Support scalability at extreme scale
- Portability & extensibility on any HPC platform

CoreNeuron Development Decision

Blue Brain Project

How?

- Have reduced & optimized data structures
- Vectorization using autovectorization
- Hybrid MPI / OpenMP with single MPI process per node
- Code reduction
  - Model configuration data structured are removed
  - Support of interpreter languages is not included
  - 15k lines vs 300k lines in NEURON

Highlights of CoreNeuron

Blue Brain Project

- Simulation functionality of NEURON
- Reduced memory footprint (2MB/neuron vs 12MB/neuron)
- Three levels of parallelism
  - Nodes: collection of cell groups
  - Threads: each cell group has its (OpenMP) thread
  - Vectorization: computed mechanisms per cell group
- Spike delivery is done via MPI_Allgather(v)

Where is CoreNeuron in Scale

Blue Brain Project

Problem size & memory requirements for future simulations

parameters / brain type	rat	monkey	human
number of neurons	1×10^{8}	1×10^{9}	8×10^{10}
number of synapses	5×10^{11}	1×10^{13}	1×10^{15}
number of state variables	3.3×10^{12}	6.3×10^{13}	6×10^{15}
estimated size in memory	100 TiB	1 PiB	80 PiB

Today: the order of the rat brain size on a full JUQUEEN

NEURON & Neurodamus vs CoreNeuron

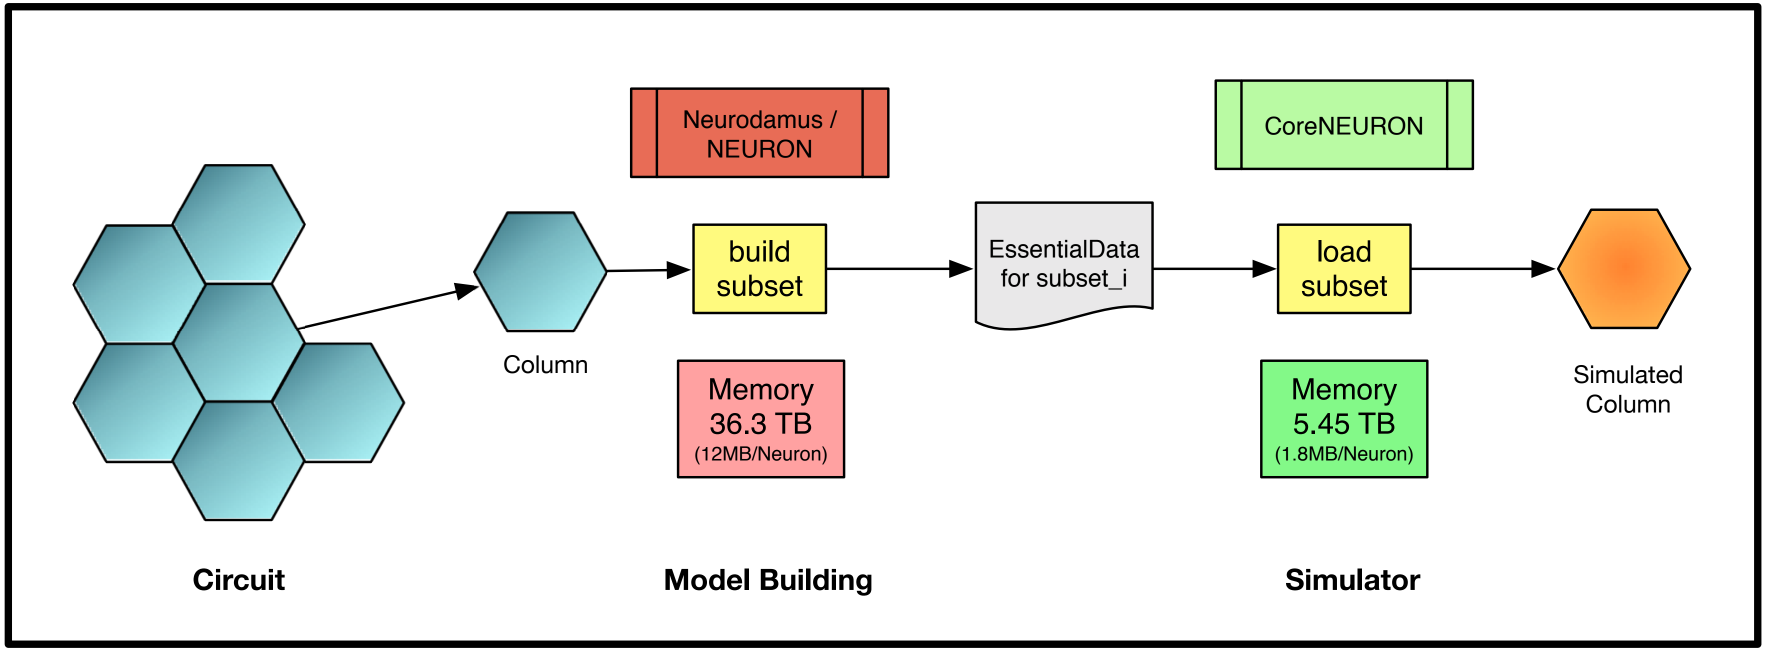

NEURON & Neurodamus vs CoreNeuron

Blue Brain Project

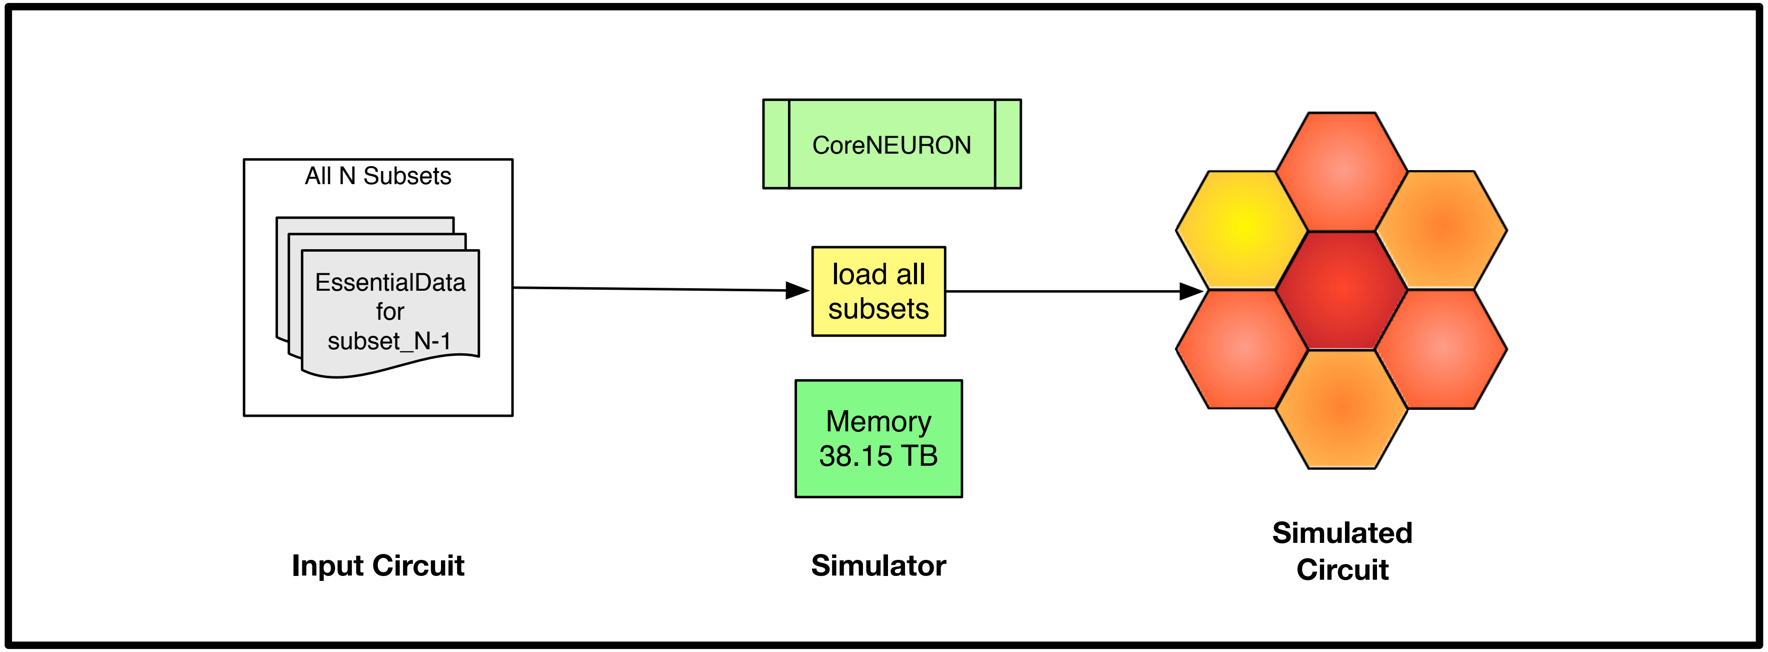

Use Case Description

Blue Brain Project

Single column

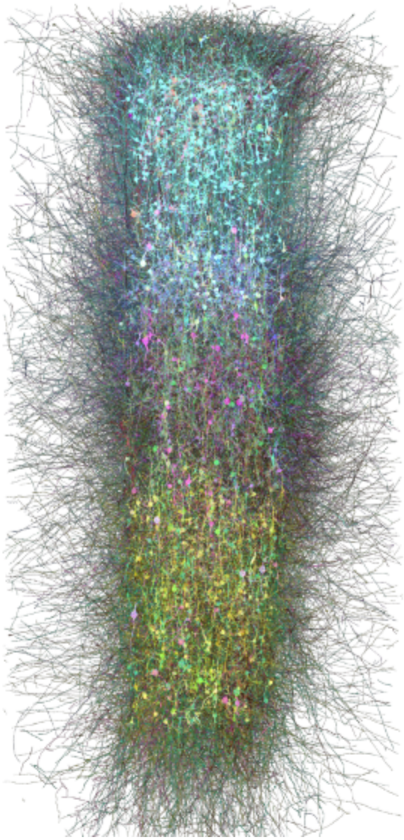

Top-side view of the circuit, 100 columns

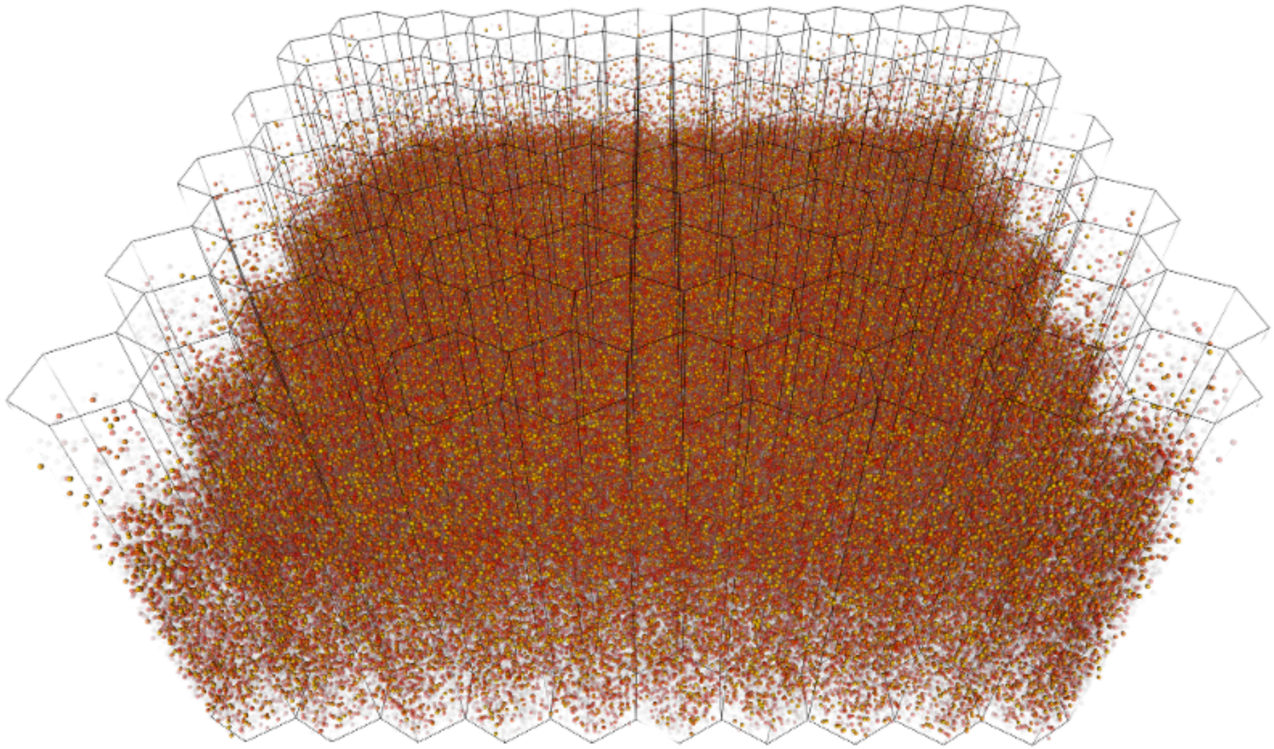

- Initial circuit: 3 million neurons and 9 billion synapses
- Ready for in-memory duplication. Default size: 24 million neurons
- I/O size to read: 5 TB

IBM Blue Gene/Q Node Performance Analysis

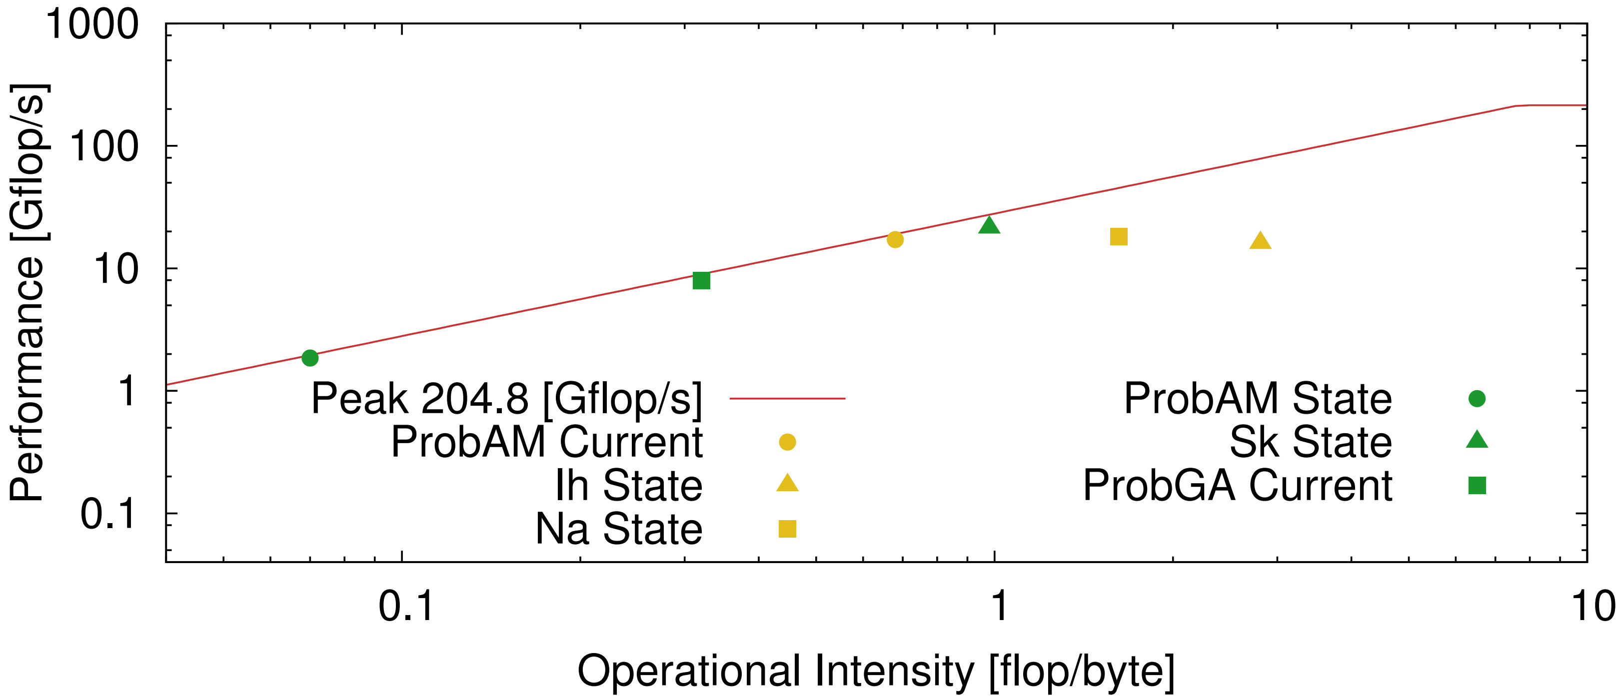

- Most kernels (ProbAM Current, etc.) memory bandwidth limited
- Some (*Ih State*, *Na State*) can be vectorized to get better performance

Strong Scaling Studies

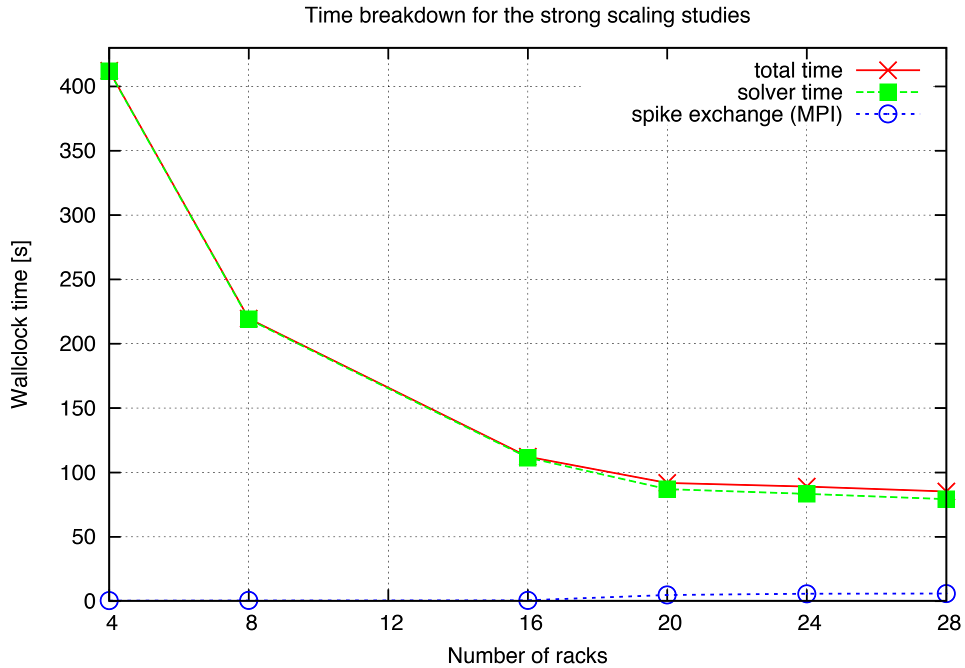

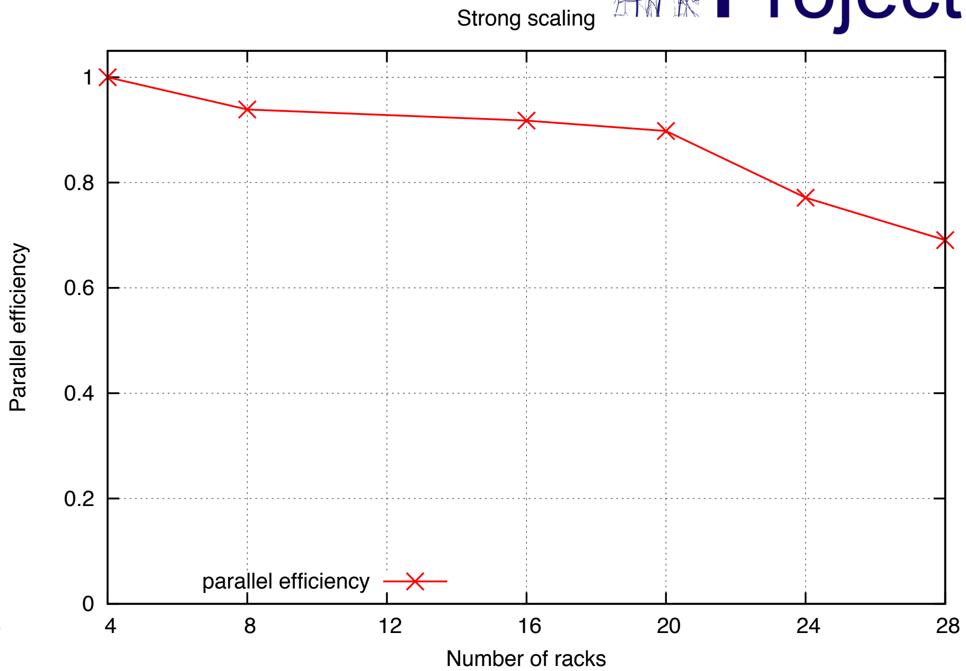

- 1 MPI task per node, 64 OpenMP threads per node
- 10 milliseconds of biological time
- 24 million neurons for each run (from ~90% of DRAM to ~10% of DRAM)
- No uniform distribution of data across 20, 24, 28 racks <=> artificial load imbalance
- Up to 16 racks: 10% of strong scaling efficiency loss

Weak Scaling Studies

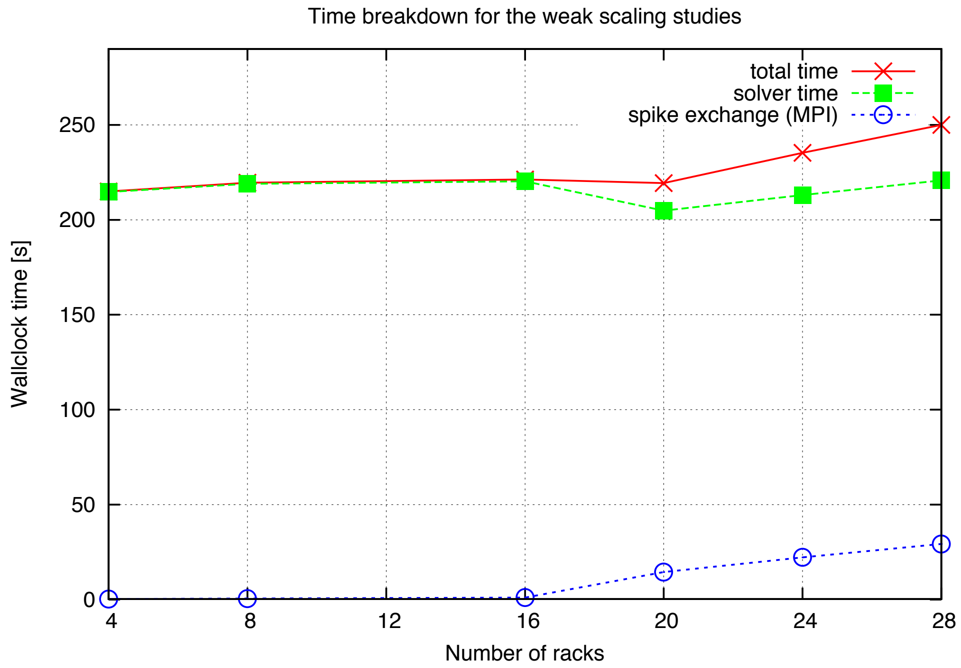

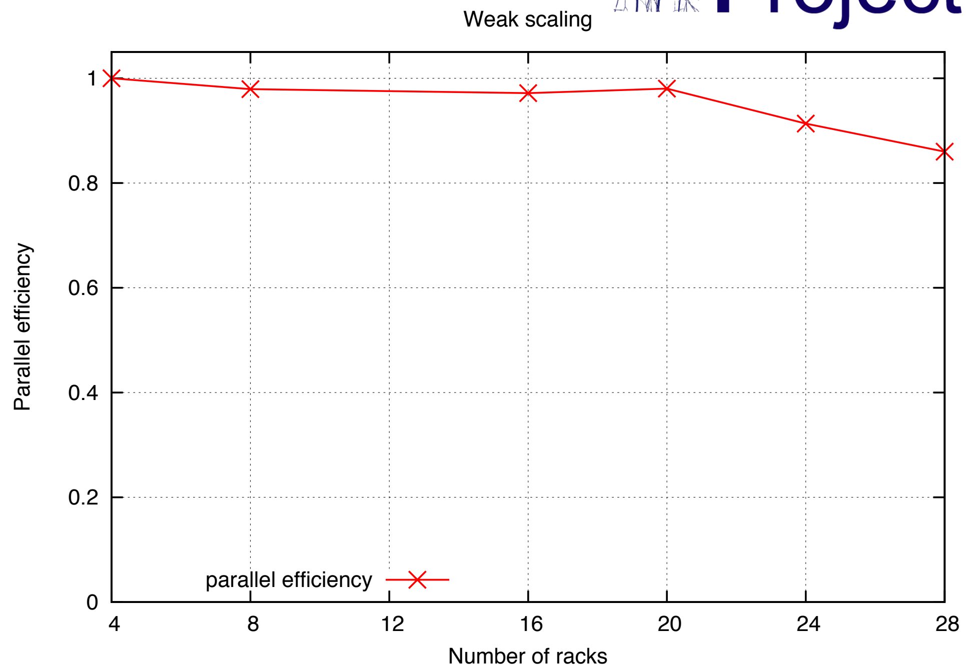

- 1 MPI task per node, 64 OpenMP threads per node
- 10 milliseconds of biological time
- 2906 neurons per node: from 12 million to 82 million neurons (~50% of DRAM)
- Parallel efficiency is nearly optimal up to 20 racks

Closing Remarks

- Full JUQUEEN machine simulation
  - 28 racks, utilizing all 1,835,008 threads
  - 15.9 GB of node DRAM, 155 million neurons (duplicated circuit)
- Memory reduction 6-8 times comparing to NEURON
- Improved on-node performance, ready for larger scale

Further Steps

- Ongoing reduction of memory footprint
- Disk-to-memory data management: utilize HDF5
- Implementation simplification
- Introduction of clear C++ API
- Light-weight python interface for high-level API
- Exposing more parallelism in mechanisms & spike exchange