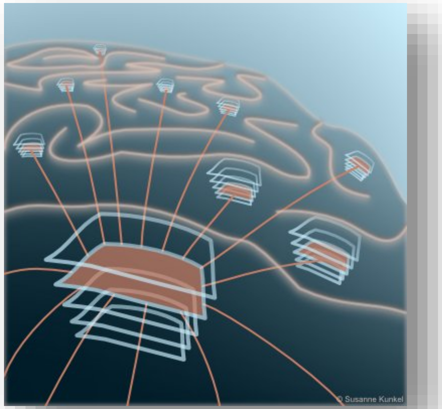
© Susanne Kunkel

Jülich **A**achen **R**esearch **A**lliance

Spiking network simulation code for the peta scale

www.csn.fz-juelich.de
www.nest-initiative.org

Oct 2nd 2014, BrainScaleS, Heidelberg

Markus Diesmann
Institute of Neuroscience and Medicine (INM-6)
Institute for Advanced Simulation (IAS-6)
Jülich Research Centre

Overview

- models of large-scale networks

  cortical microcircuit

  ⇩

  need for brain-scale models

- designing simulation software for the brain scale

  3rd generation simulation kernel

  [scale of 10,000 compute nodes]

  ⇩

  4th generation simulation kernel, platform in initial phase of HBP

  [scale of 100,000 compute nodes]

nest::

Interactions between neurons

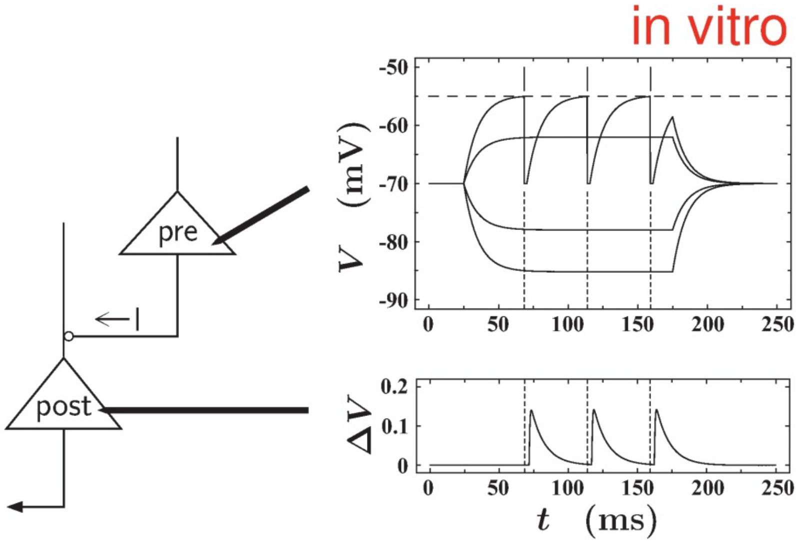

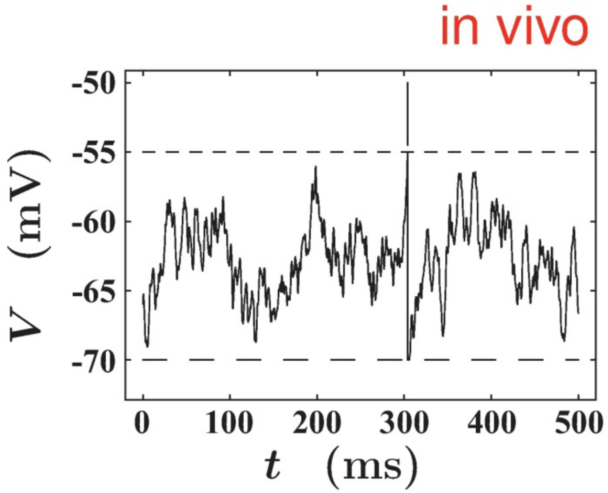

- current injection into pre-synaptic neuron causes excursions of membrane potential
- supra-threshold value causes spike transmitted to post-synaptic neuron
- post-synaptic neuron responds with small excursion of potential after delay
- inhibitory neurons (20%) cause negative excursion

- each neuron receives input from 10,000 other neurons
- causing large fluctuations of membrane potential
- emission rate of 1 to 10 spikes per second

Precision

- NEST supports spike interaction without discretization of time
- feature accessible via PyNN
- unrelated to whether algorithm is globally event-driven

Morrison A, Straube S, Plesser H E, Diesmann M (2007) Exact subthreshold integration with continuous spike times in discrete time neural network simulations Neural Computation 19: 47-79

Hanuschkin A, Kunkel S, Helias M, Morrison A and Diesmann M (2010) A general and efficient method for incorporating precise spike times in globally time-driven simulations Front. Neuroinform. 4:113

- feature regularly used for high-precision simulations

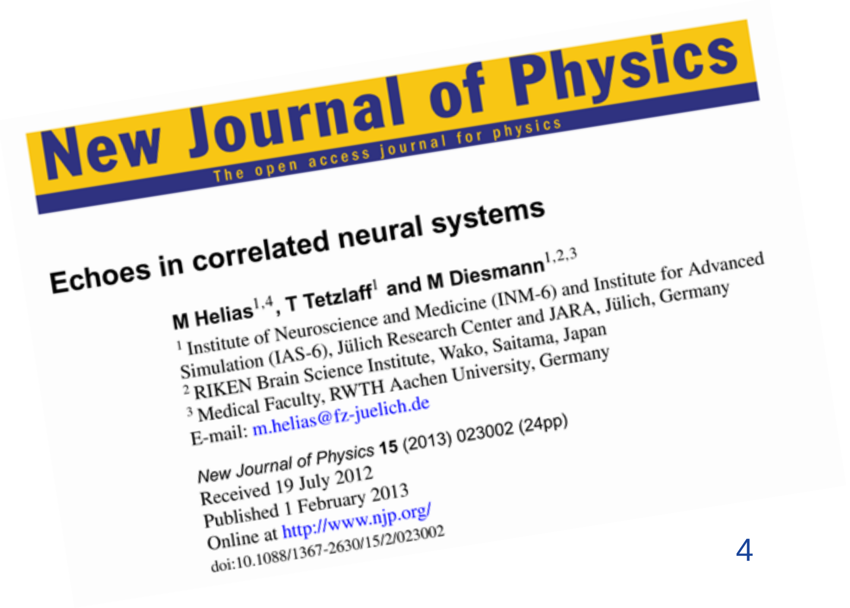

New Journal of Physics
The open access journal for physics

Echoes in correlated neural systems

M Helias[1,4], **T Tetzlaff**[1] **and M Diesmann**[1,2,3]

[1] Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6), Jülich Research Center and JARA, Jülich, Germany
[2] RIKEN Brain Science Institute, Wako, Saitama, Japan
[3] Medical Faculty, RWTH Aachen University, Germany
E-mail: m.helias@fz-juelich.de

New Journal of Physics **15** (2013) 023002 (24pp)
Received 19 July 2012
Published 1 February 2013
Online at http://www.njp.org/
doi:10.1088/1367-2630/15/2/023002

Minimal layered cortical network model

- 1 mm^3
- 1 billion synapses, 100,000 neurons
- 2 populations of neurons per layer:
    - E: Excitatory
    - I: Inhibitory
- E and I identical neuronal dynamics
- laterally homogeneous connectivity
- layer- and type-specific C_{ij}^{xy}

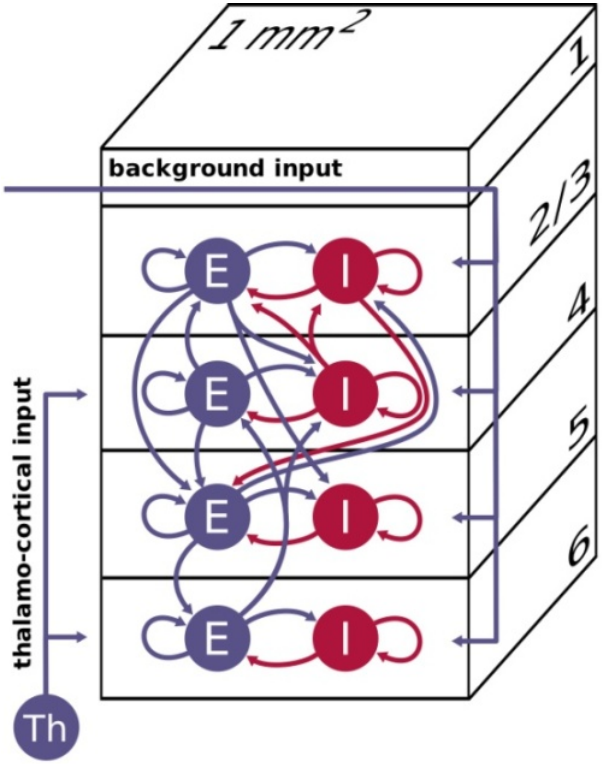

Convergence and divergence

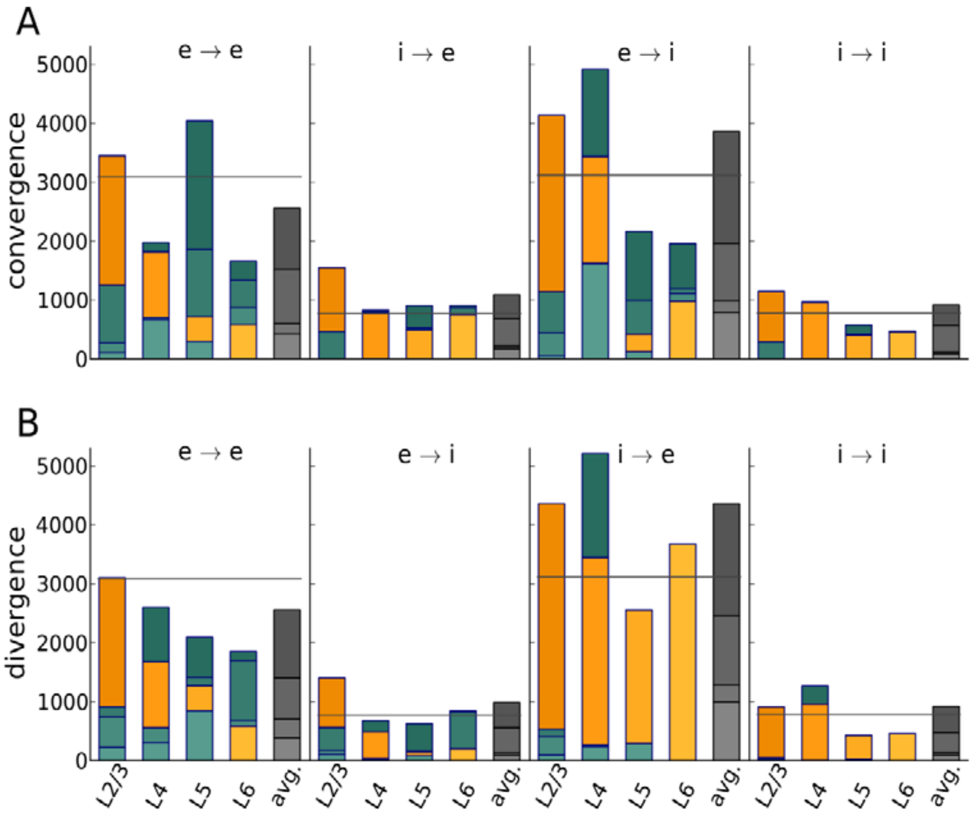

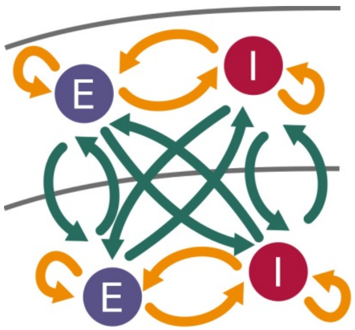

- dominated by within-layer connections
- e ⟶ e divergence reflects ”standard” loop
- e ⟶ i divergence reflects target-specific feedback

Local cortical microcircuit

taking into account layer and neuron-type specific connectivity is sufficient to reproduce experimentally observed:

- asynchronous-irregular spiking of neurons
- higher spike rate of inhibitory neurons
- correct distribution of spike rates across layers
- integrates knowledge of more than 50 experimental papers

Potjans TC & Diesmann M (2014) The cell-type specific connectivity of the local cortical network explains prominent features of neuronal activity. *Cerebral Cortex* 24 (3): 785-806

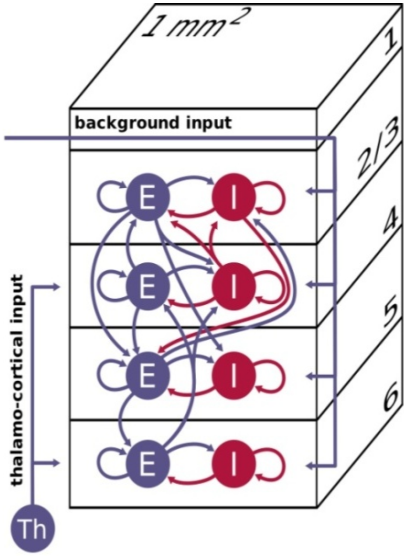

10^5 neurons
10^9 synapses

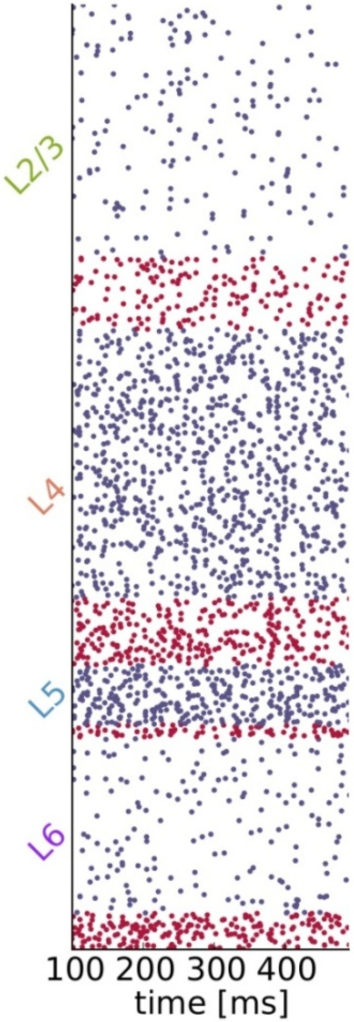

available at: www.opensourcebrain.org

Building block for functional studies

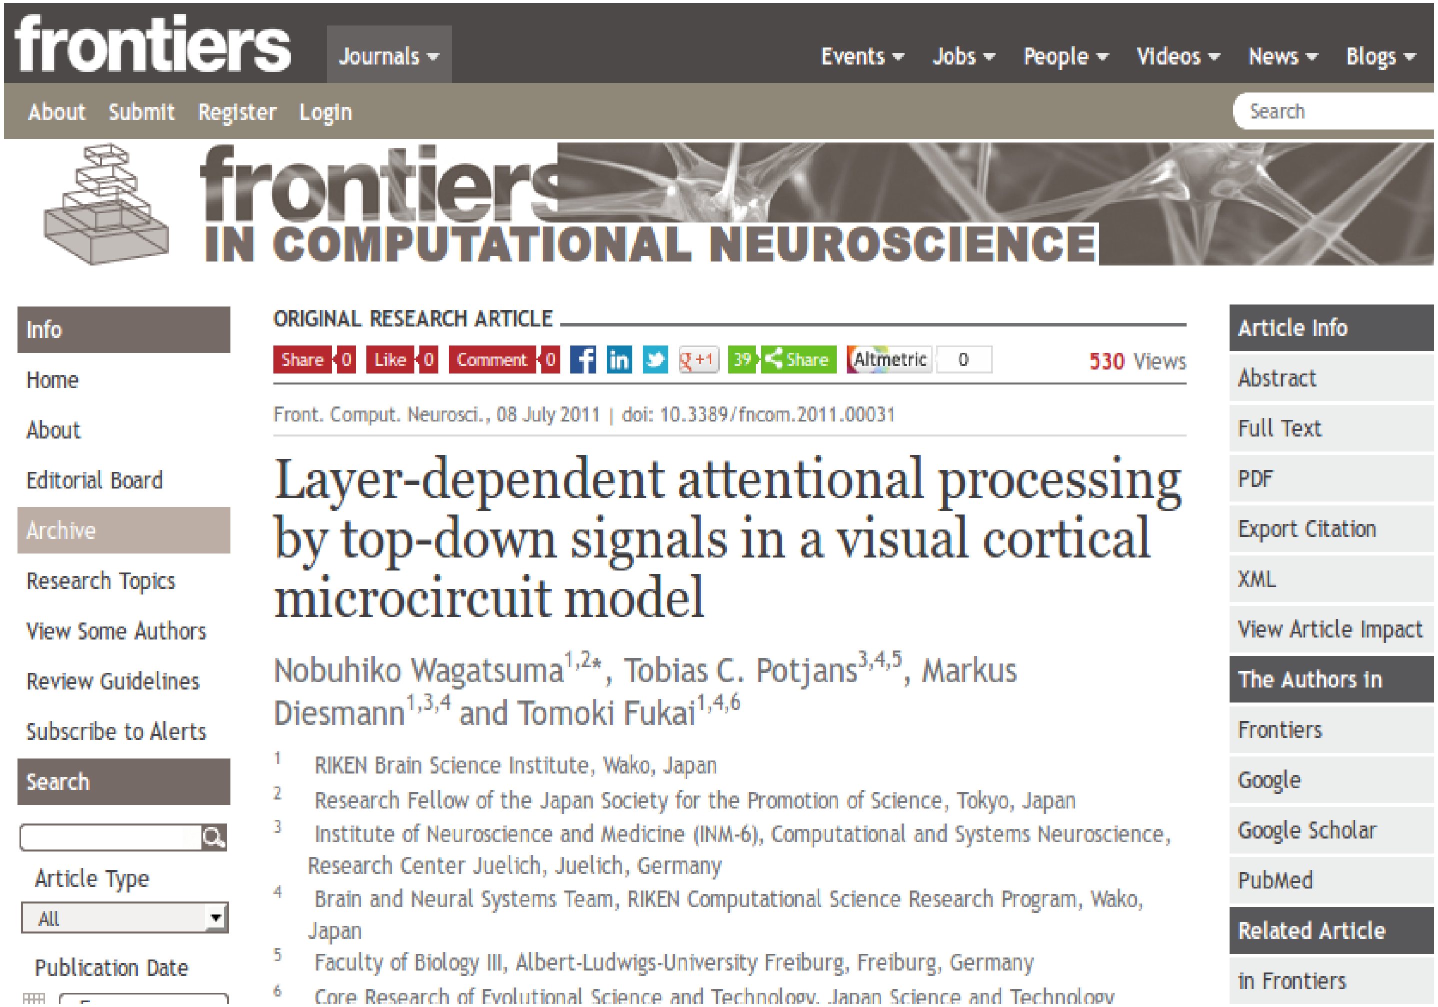

ORIGINAL RESEARCH ARTICLE

Front. Comput. Neurosci., 08 July 2011 | doi: 10.3389/fncom.2011.00031

Layer-dependent attentional processing by top-down signals in a visual cortical microcircuit model

Nobuhiko Wagatsuma[1,2*], Tobias C. Potjans[3,4,5], Markus Diesmann[1,3,4] and Tomoki Fukai[1,4,6]

[1] RIKEN Brain Science Institute, Wako, Japan
[2] Research Fellow of the Japan Society for the Promotion of Science, Tokyo, Japan
[3] Institute of Neuroscience and Medicine (INM-6), Computational and Systems Neuroscience, Research Center Juelich, Juelich, Germany
[4] Brain and Neural Systems Team, RIKEN Computational Science Research Program, Wako, Japan
[5] Faculty of Biology III, Albert-Ludwigs-University Freiburg, Freiburg, Germany
[6] Core Research of Evolutional Science and Technology, Japan Science and Technology

Building block for mesoscopic studies

- collaboration with Gaute Einevoll (UMB, Norway)

nest::simulated()

LFPy
Local Field Potentials in Python

NEURON
for empirically-based simulations of neurons and networks of neurons

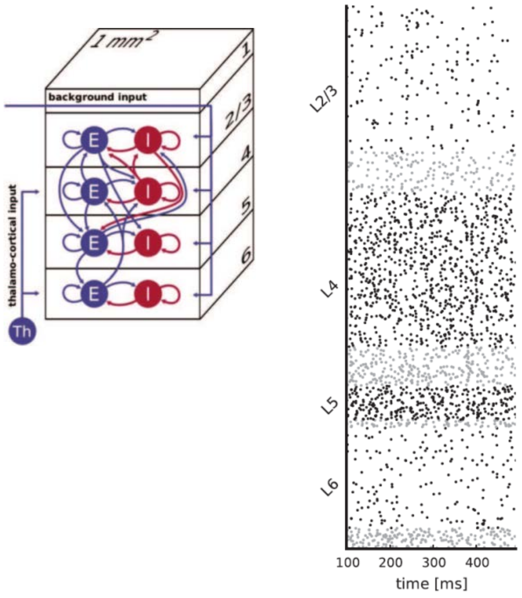

+

´Morphology

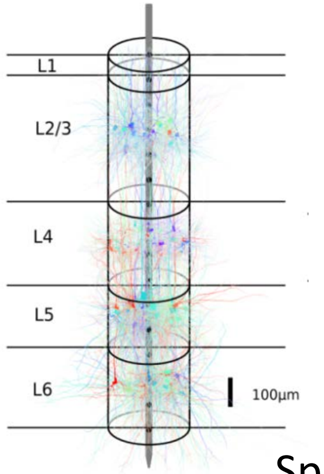

=

Laminar LFP profile

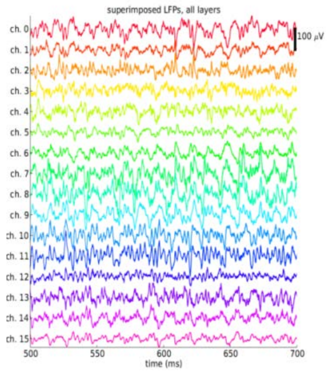

Spatial reach of the LFP

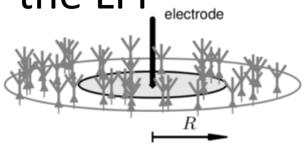

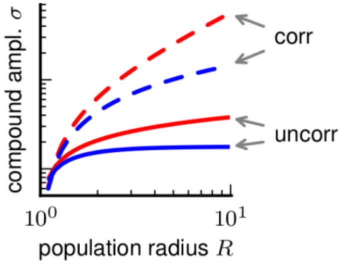

Linden H, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT (2011) Modeling the spatial reach of the LFP. *Neuron* 72(5):859-872

Critique of local network model

a network of networks with at least three levels of organization:

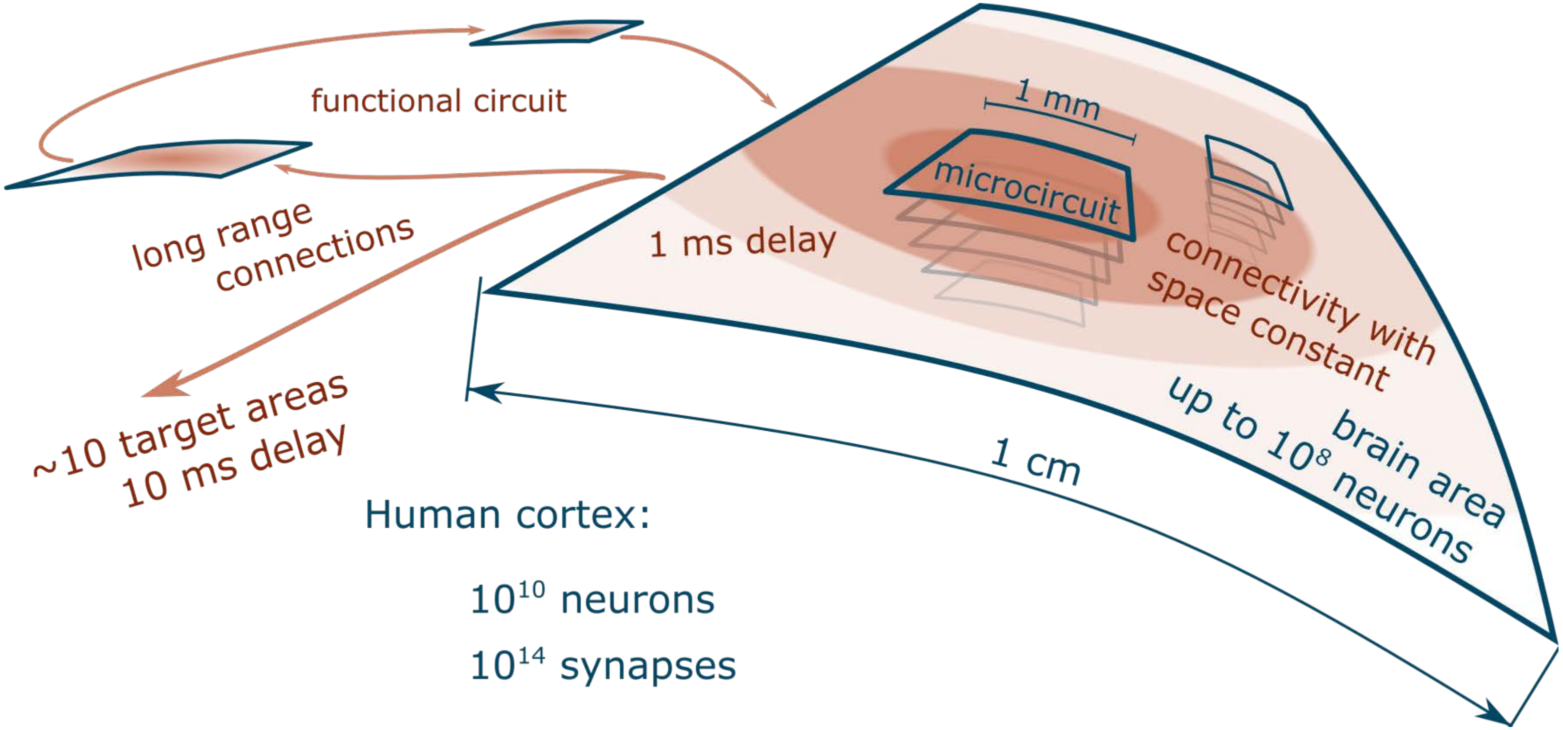

- neurons in local microcircuit models are missing 50% of synapses
- e.g., power spectrum shows discrepancies, slow oscillations missing
- solution by taking brain-scale anatomy into account

Meso- and macro-scale measures

brain-scale networks basis for:
- further measures by forward modeling
- comparison with mean-field models

mesoscopic measures
- local field potential (LFP)
- voltage sensitive dyes (VSD)

and macroscopic measures
- EEG, MEG
- fMRI resting state networks

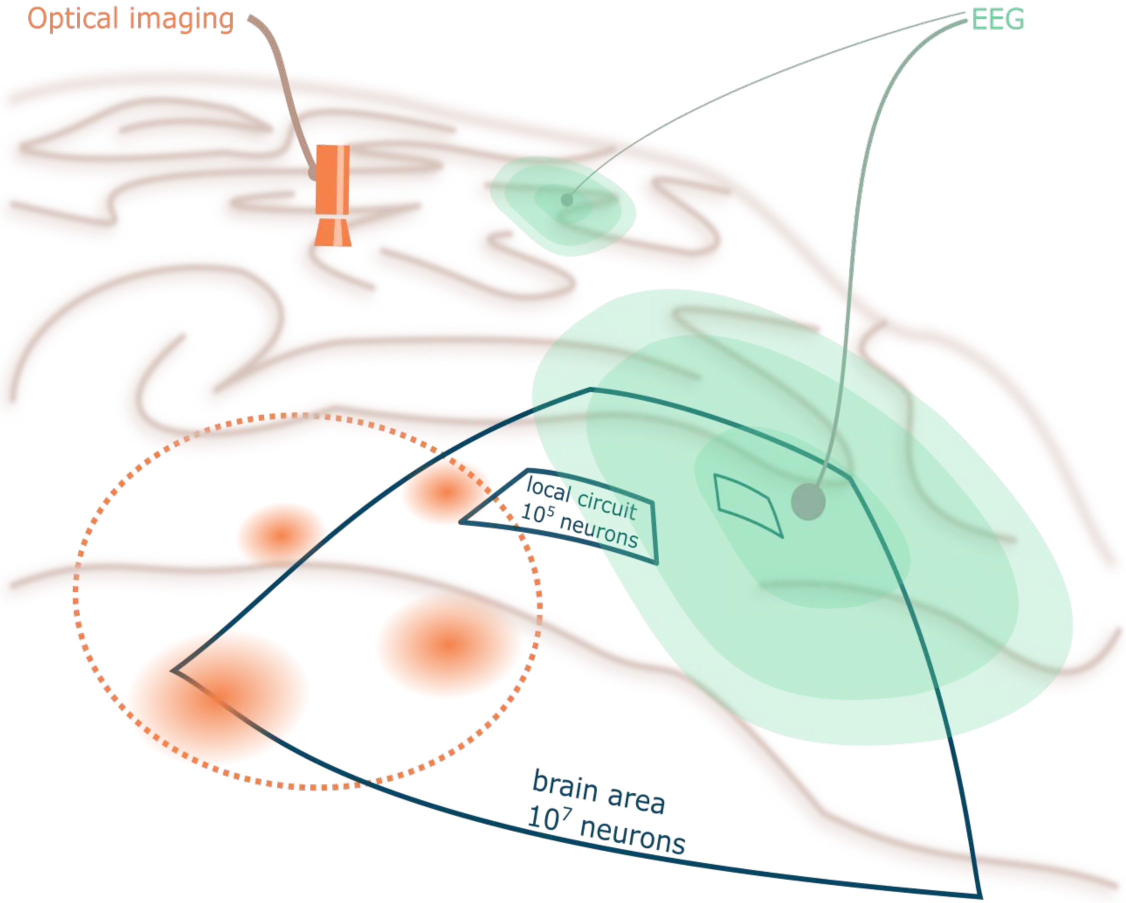

Toward a self-consistent model

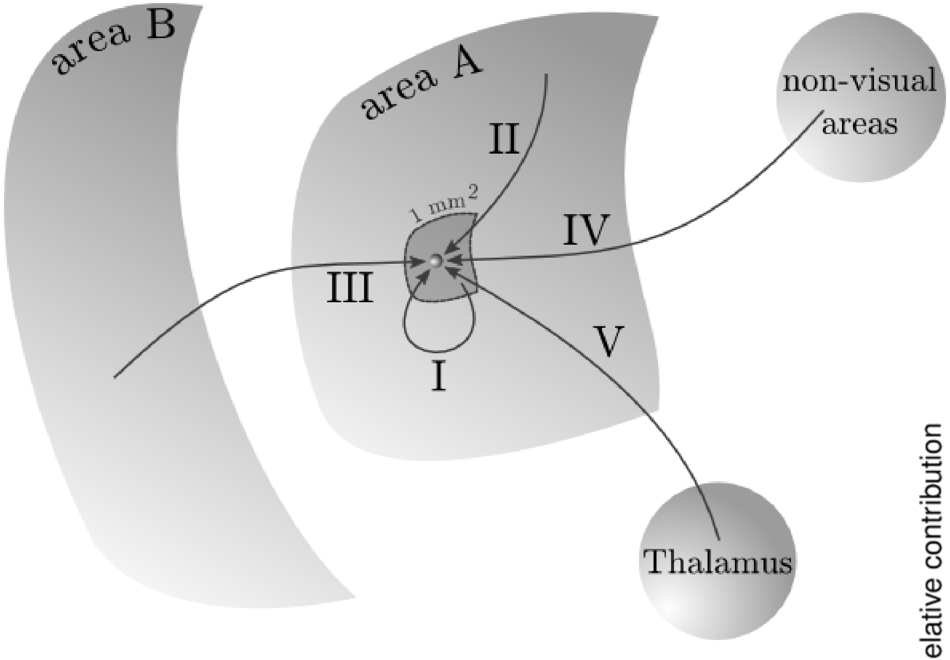

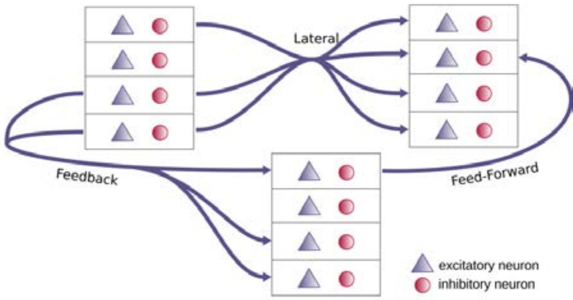

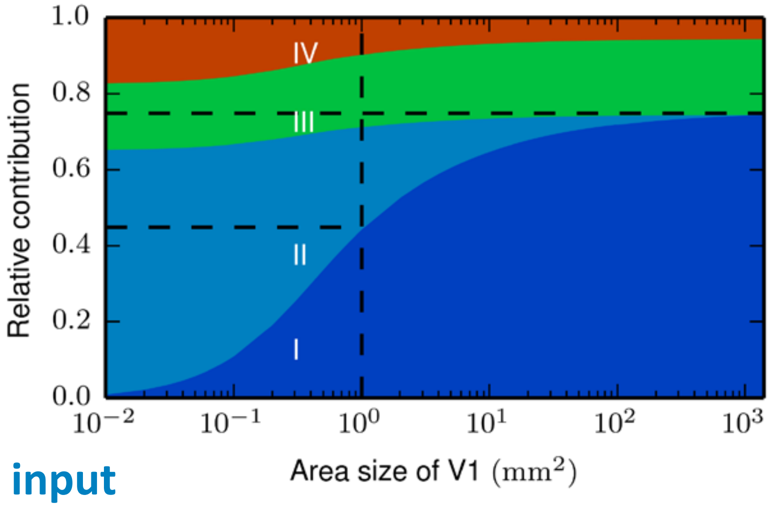

I. **Intra-areal synapses**
II. **Intra-areal synapses replaced by random input**
III. **Cortico-cortical synapses**
IV. **External input represented by random input**
V. **Thalamic input**

- Sacha van Albada
- Maximilian Schmidt
- Rembrandt Bakker

Multi-area model of macaque visual cortex

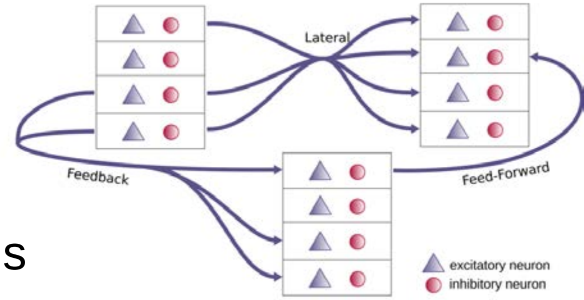

- rich anatomical data sets available (e.g CoCoMac)
- close to human
- 32 areas structured in layers comprising $8 \cdot 10^8$ neurons
- downscaled model with $3.2 \cdot 10^6$ neurons and $3 \cdot 10^{10}$ synapses

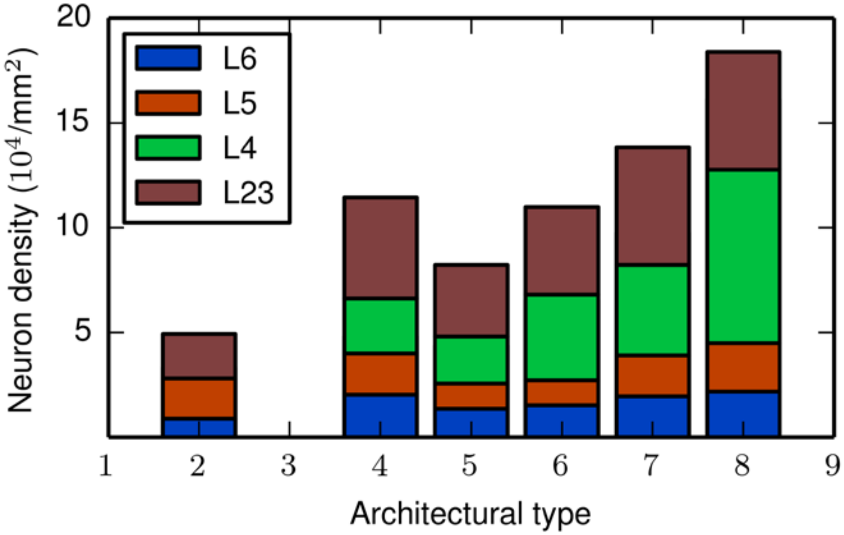

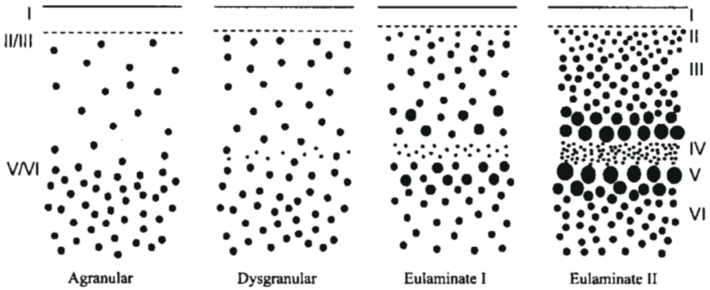

From Dombrowski et al. (2001), Cereb Cortex

architectural types provided by C. Hilgetag (private communication)

Cortico-cortical connectivity

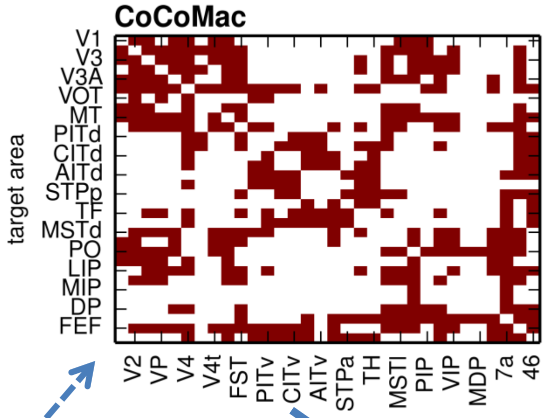

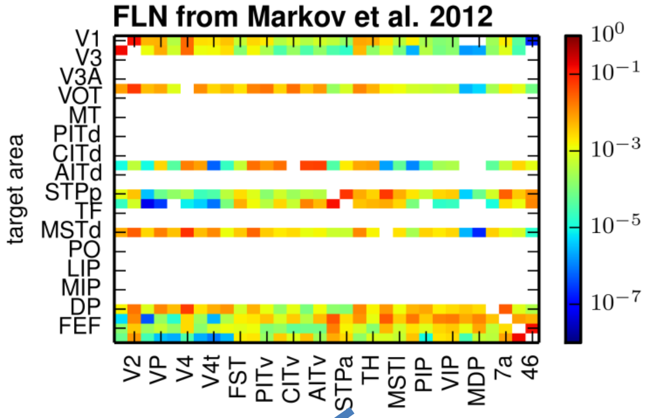

labeling:
y-axis: odd areas
x-axis: even areas

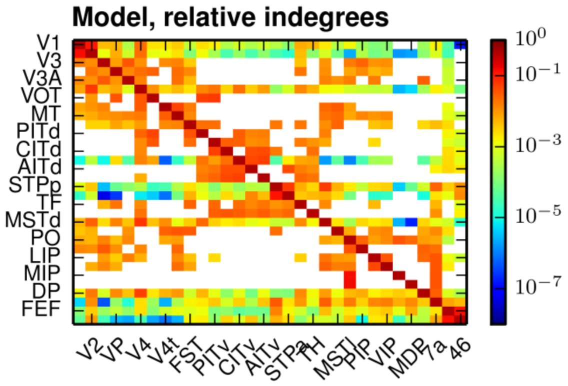

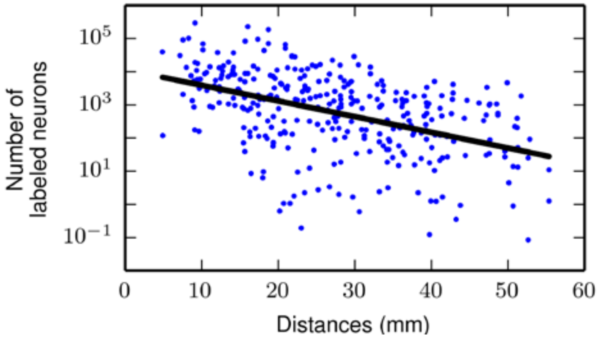

Markov et al. (2012)
Cereb Cortex

Dynamical results

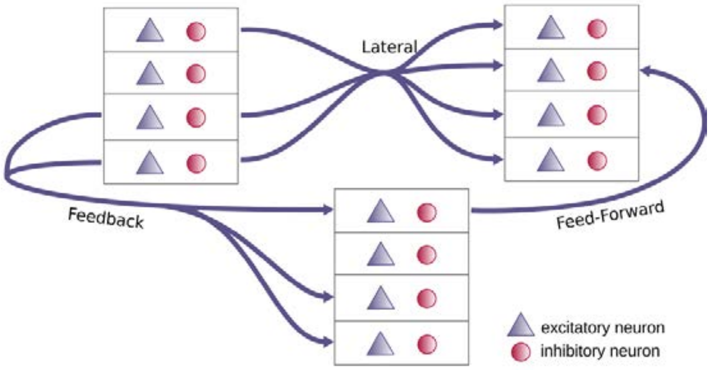

- Heterogeneous laminar firing patterns
- Rates in reasonable range (0.2 – 20 spikes/s)
- Inhibitory rates > excitatory rates
- Broad rate distributions

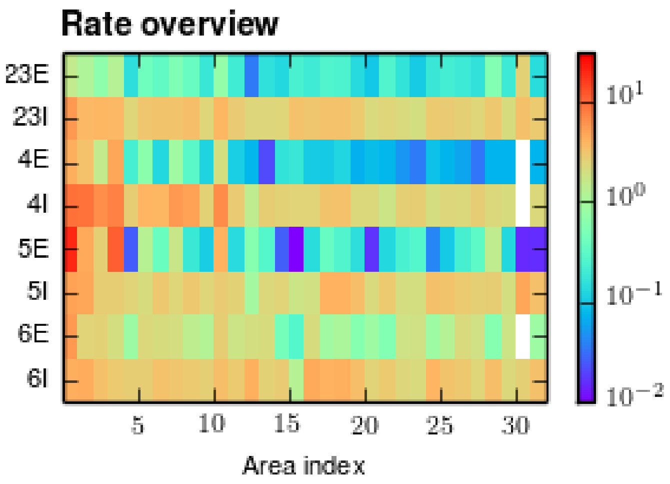

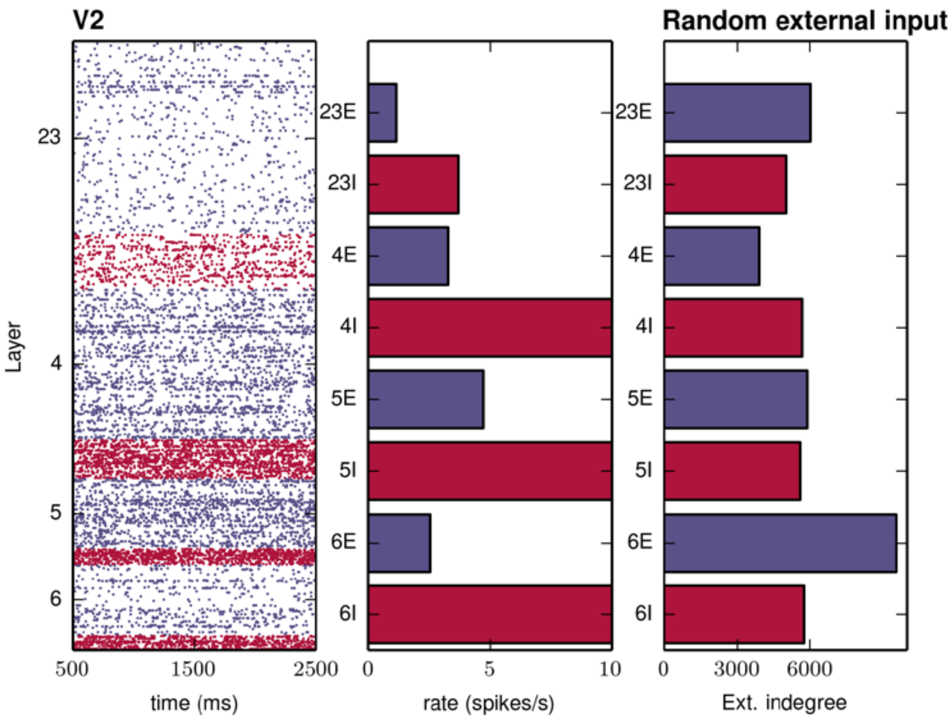

The mesoscopic scale

- parallel study
- network of 4mmx4mm covers size of 100 electrode Utah array
- spike correlation structure, LFP
- experimentally, formation and breakdown of waves observed in relation to behavior

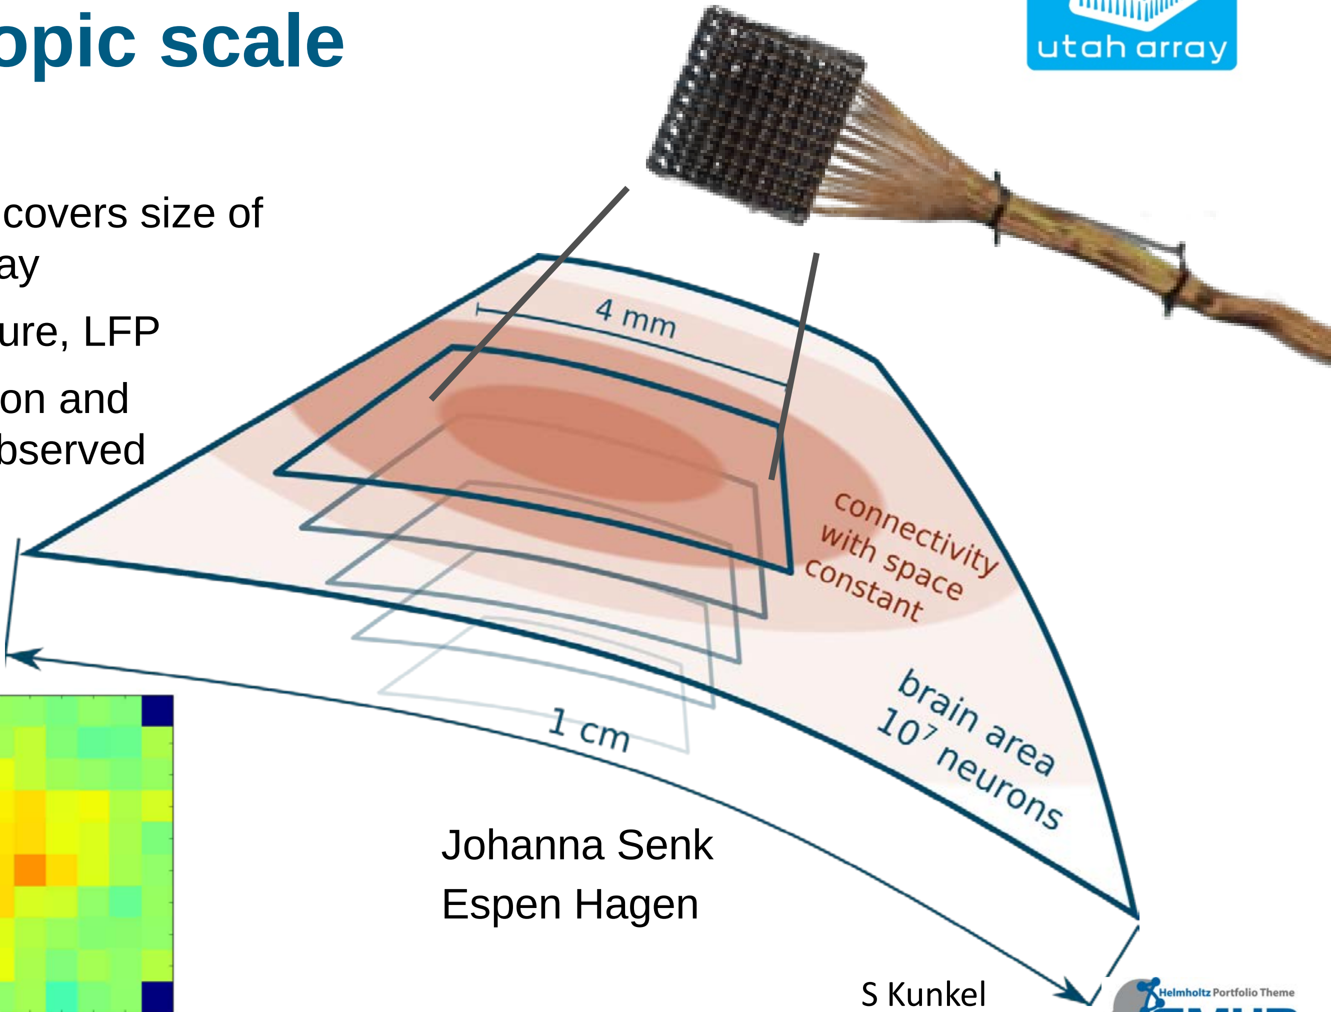

Johanna Senk
Espen Hagen

S Kunkel

data: INT/Juelich

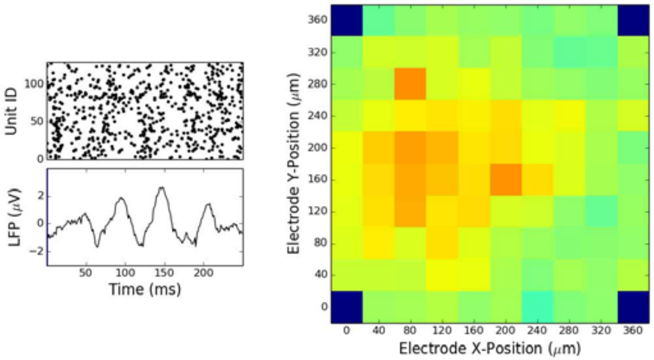

- about 1 million neurons with 10 billion synapses
- requires supercomputing but not incredible resources

Simulation Technology: the NEST Initiative

collaborative effort and community building

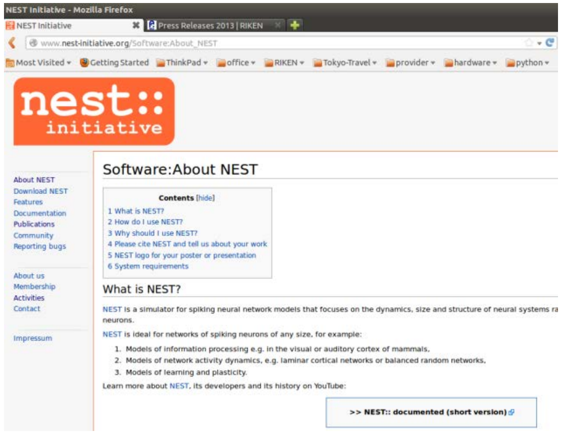

Software:About NEST

Contents [hide]
1 What is NEST?
2 How do I use NEST?
3 Why should I use NEST?
4 Please cite NEST and tell us about your work
5 NEST logo for your poster or presentation
6 System requirements

What is NEST?

NEST is a simulator for spiking neural network models that focuses on the dynamics, size and structure of neural systems ra
neurons.

NEST is ideal for networks of spiking neurons of any size, for example:

1. Models of information processing e.g. in the visual or auditory cortex of mammals,
2. Models of network activity dynamics, e.g. laminar cortical networks or balanced random networks,
3. Models of learning and plasticity.

Learn more about NEST, its developers and its history on YouTube:

>> NEST:: documented (short version)

Major goals:

systematically publish
new simulation technology

produce public releases
under GPL

- origins in 1994, collaboration of several labs (since 2001)
- registered society (since 2012)
- teaching in international advanced courses:
    - Okinawa Computational Neuroscience Course OCNC, Japan
    - Advanced Course in Computational Neuroscience ACCN, Europe
    - Latin American School on Computational Neuroscience LASCON, South America

www.nest-initiative.org

- core technology of Human Brain Project

e.g.: Morrison et al. (2005) *Neural Computation*
Zaytsev, Morrison (2013) *Frontiers in Neuroinformatics*

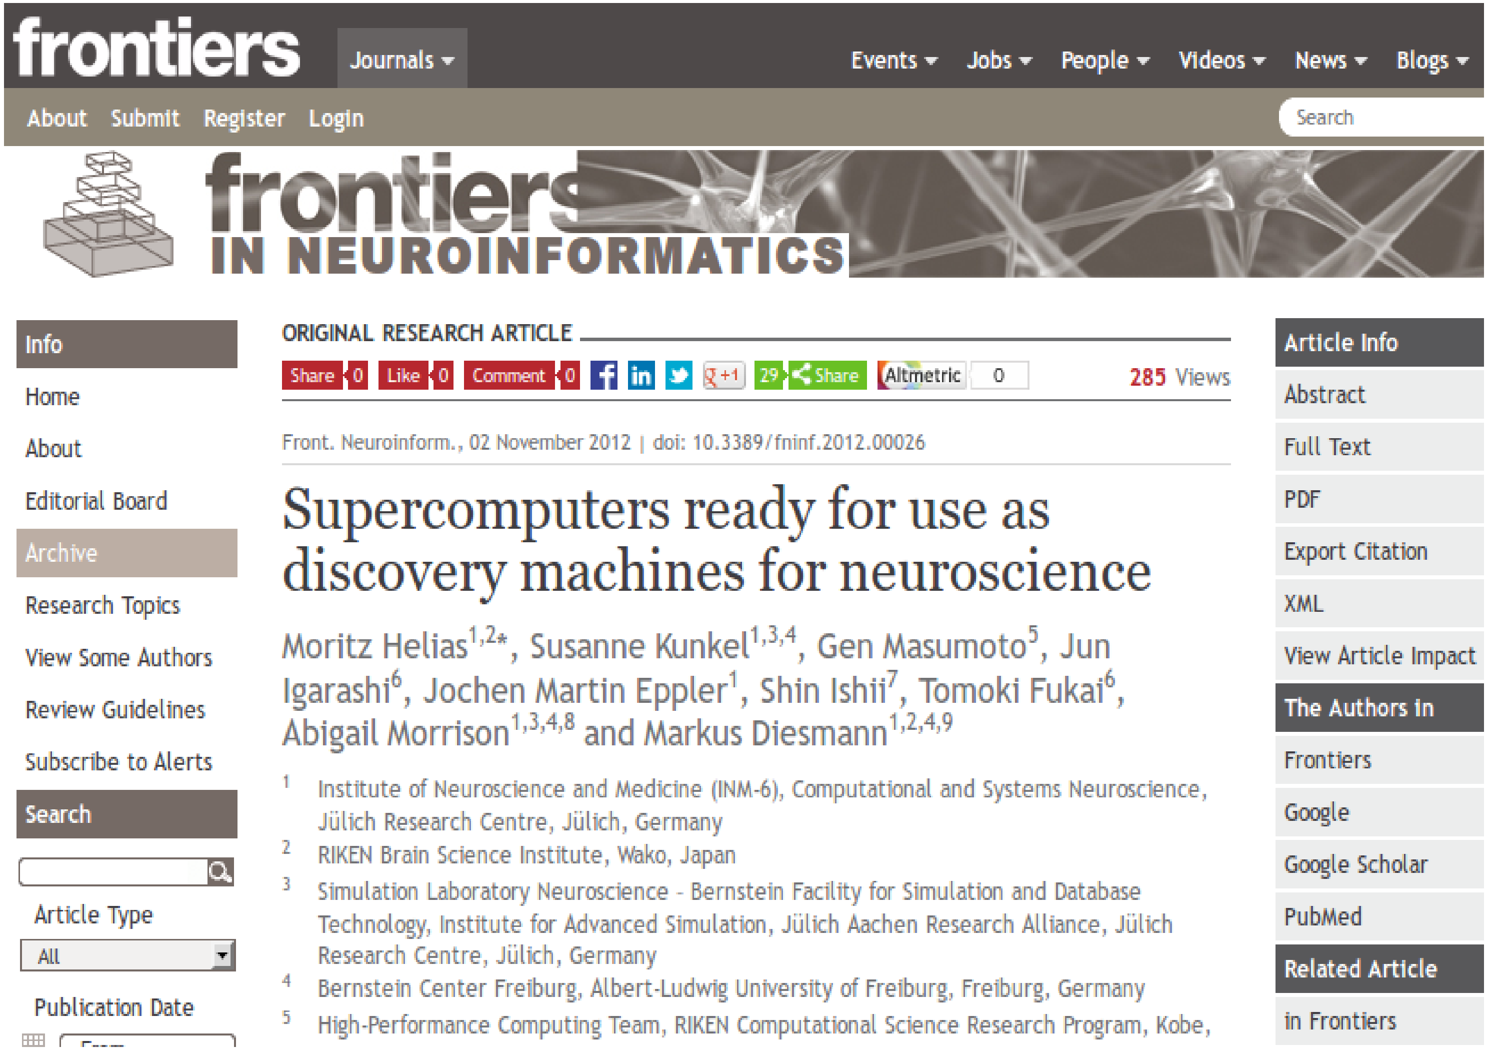

ORIGINAL RESEARCH ARTICLE

Front. Neuroinform., 02 November 2012 | doi: 10.3389/fninf.2012.00026

Supercomputers ready for use as discovery machines for neuroscience

Moritz Helias[1,2]*, Susanne Kunkel[1,3,4], Gen Masumoto[5], Jun Igarashi[6], Jochen Martin Eppler[1], Shin Ishii[7], Tomoki Fukai[6], Abigail Morrison[1,3,4,8] and Markus Diesmann[1,2,4,9]

[1] Institute of Neuroscience and Medicine (INM-6), Computational and Systems Neuroscience, Jülich Research Centre, Jülich, Germany

[2] RIKEN Brain Science Institute, Wako, Japan

[3] Simulation Laboratory Neuroscience – Bernstein Facility for Simulation and Database Technology, Institute for Advanced Simulation, Jülich Aachen Research Alliance, Jülich Research Centre, Jülich, Germany

[4] Bernstein Center Freiburg, Albert-Ludwig University of Freiburg, Freiburg, Germany

[5] High-Performance Computing Team, RIKEN Computational Science Research Program, Kobe,

3rd generation simulation kernel

(released with NEST 2.2 in December 2012)

- up to 10^8 neurons on K (and JUQUEEN)
- 11,250 synapses per neuron (exc-exc STDP)
- using up to MT=196,608 threads and T=8 threads per node
- 16 GB of memory per node

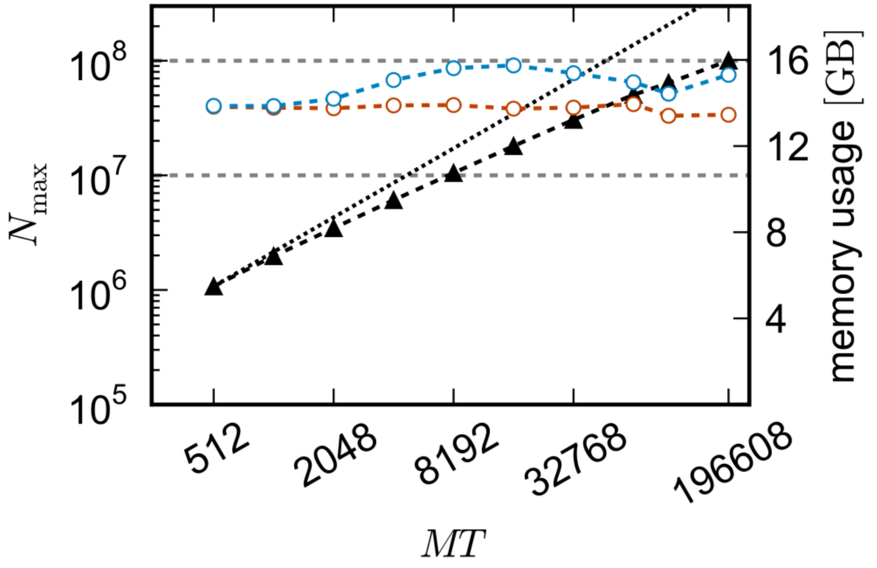

Helias M et al. (2012) *Front. Neuroinform.* **6**:26.

Model of memory usage of NEST

- describes the memory usage per MPI process

$$\begin{aligned}\mathcal{M}(M,T,N,K) &= \mathcal{M}_0(M) + \mathcal{M}_{\mathrm{n}}(M,N)\\ &\quad + \mathcal{M}_{\mathrm{c}}(M,T,N,K)\\ \mathcal{M}_{\mathrm{c}}(M,T,N,K) &= TNm_{\mathrm{c}}^{0} + TN_{\mathrm{c}}^{\emptyset}m_{\mathrm{c}}^{\emptyset}\\ &\quad + T\left(N - N_{\mathrm{c}}^{\emptyset}\right)m_{\mathrm{c}}^{+}\\ &\quad + K_M m_{\mathrm{c}}\end{aligned}$$

M total number of MPI processes
T number of threads per MPI process
N total number of neurons
K number of incoming connections per neuron

Kunkel S et al. (2012)
Front. Neuroinform. **5**:35.

Previous connection infrastructure (3g)

- required on each process

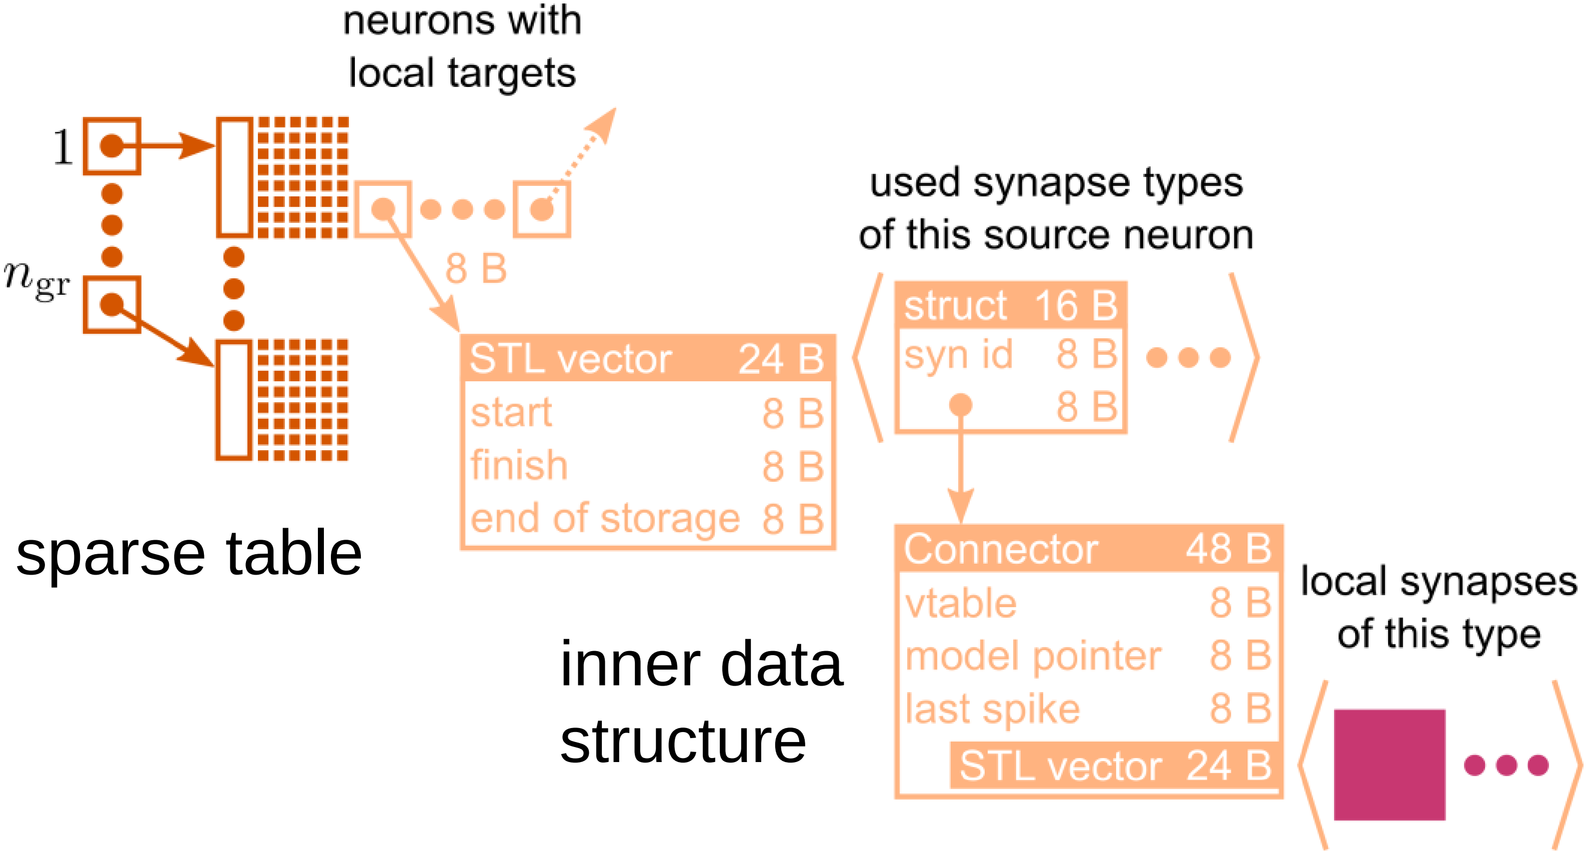

neurons with local targets

used synapse types of this source neuron

STL vector	24 B
start	8 B
finish	8 B
end of storage	8 B

struct	16 B
syn id	8 B
	8 B

Connector	48 B
vtable	8 B
model pointer	8 B
last spike	8 B
STL vector	24 B

local synapses of this type

sparse table

inner data structure

- from 10,000 nodes on collapse along 2 dimensions

3^{rd} generation simulation kernel

analysis of contributions to total memory usage

- in the regime of 10k processes and beyond the inner data structure causes severe overhead

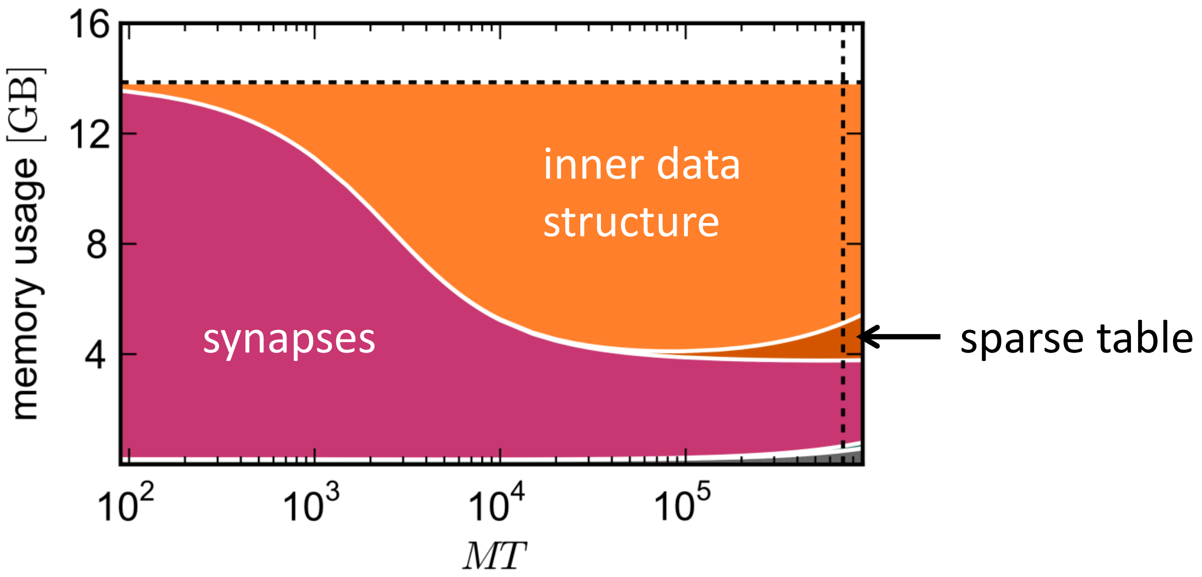

3rd generation simulation kernel

analysis of contributions to total memory usage

- adapt the model to account for short target lists
- potential solution: low-overhead data structure on supercomputers

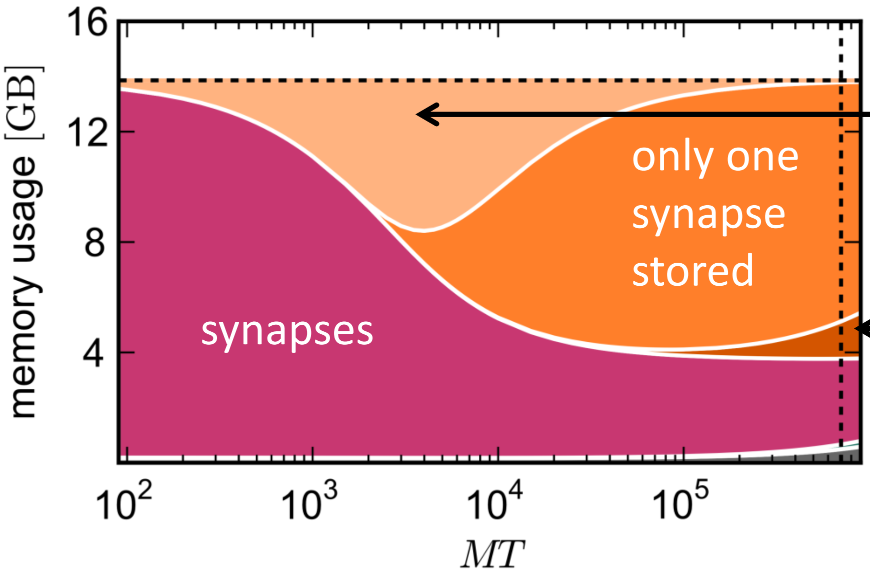

more than
one synapse
stored

sparse table

New adaptive connection infrastructure (4g)

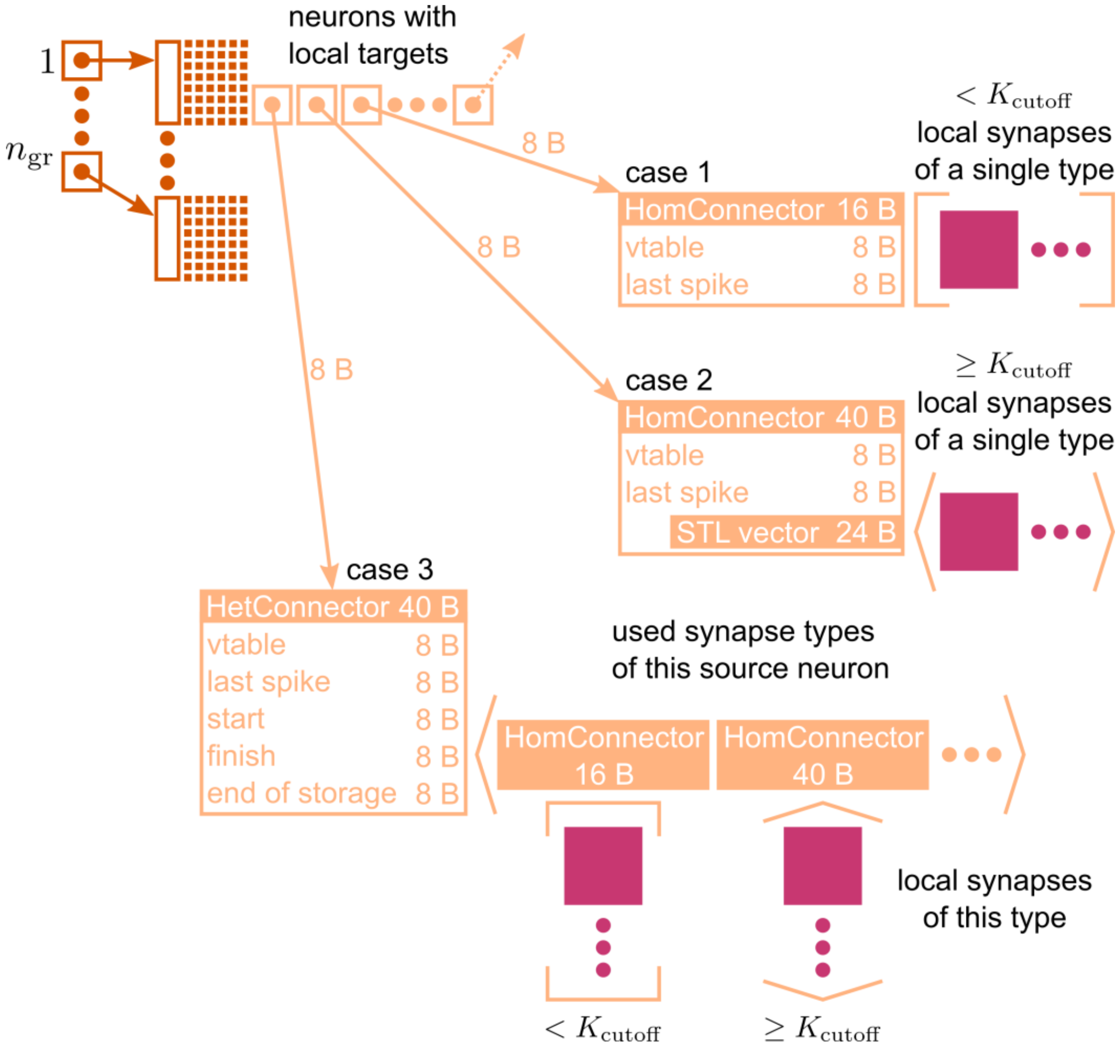

low overhead per synapse on supercomputers

full flexibility on laptops and moderately sized clusters

Comparison of 4g to 3g kernel

3g kernel

4g kernel

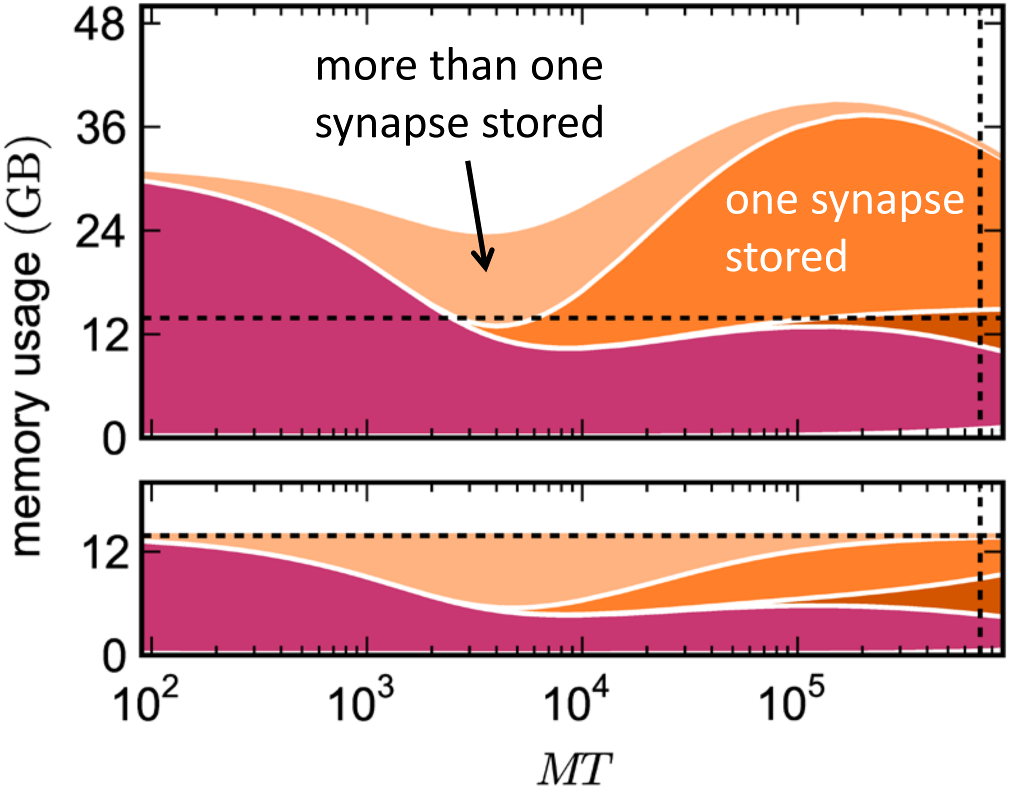

← infrastructure

← sparse table

← synapses

4th generation simulation kernel

- adaptive connection infrastructure with low overhead in case of short target lists
    - achieved by metaprogramming
    - no compromise on generality
- reduced memory usage of synapses
    - e.g. removed vtable pointers
    - no compromise on precision of synaptic state variables
- reduced setup time
- reduced simulation time

Maximum network size

- up to 5.73x10^{8} neurons on 229,376 cores of JUQUEEN
- up to 1.27x10^{9} neurons on 663,552 cores of K
- 11,250 synapses per neuron (exc-exc STDP)

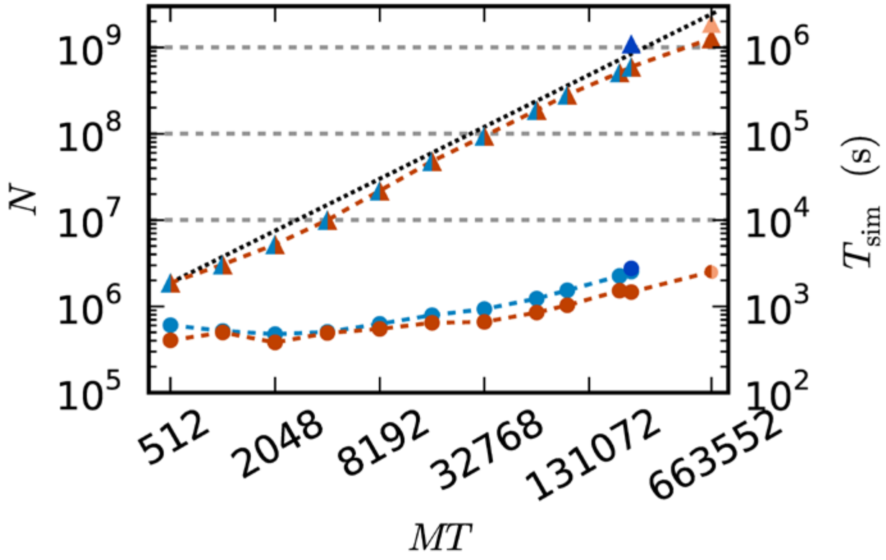

- largest general network simulation performed on K
  1.86x10^{9} neurons, 6000 synapses per neuron
  (press release July 2013: 1.73x10^{9})
- on JUQUEEN 1.08x10^{9} neurons, 6000 synapses per neuron

Runtime for a simulation of 1s

- between 8 and 41 min on JUQUEEN
- between 6 and 42 min on the K computer
- setting up the network takes between 3 and 15 min

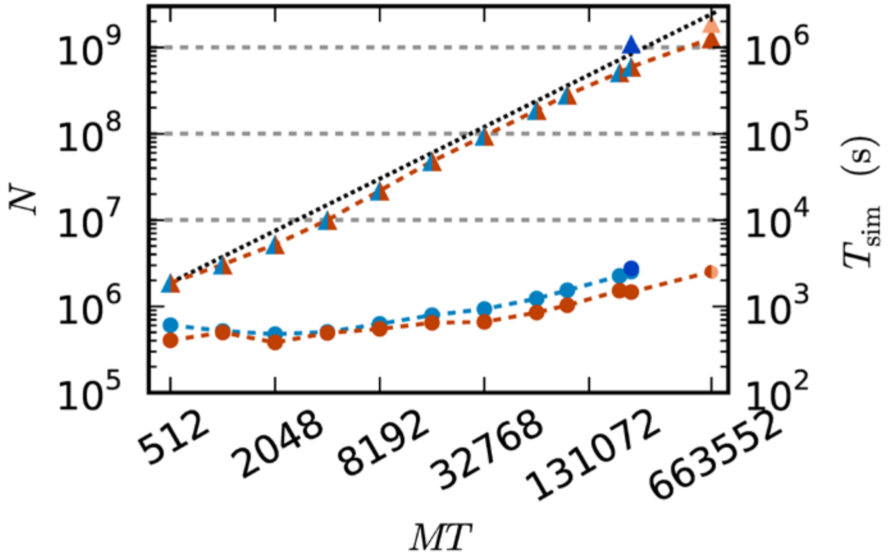

- still not fast enough for studies of plasticity
- need to increase multi-threading

- in press

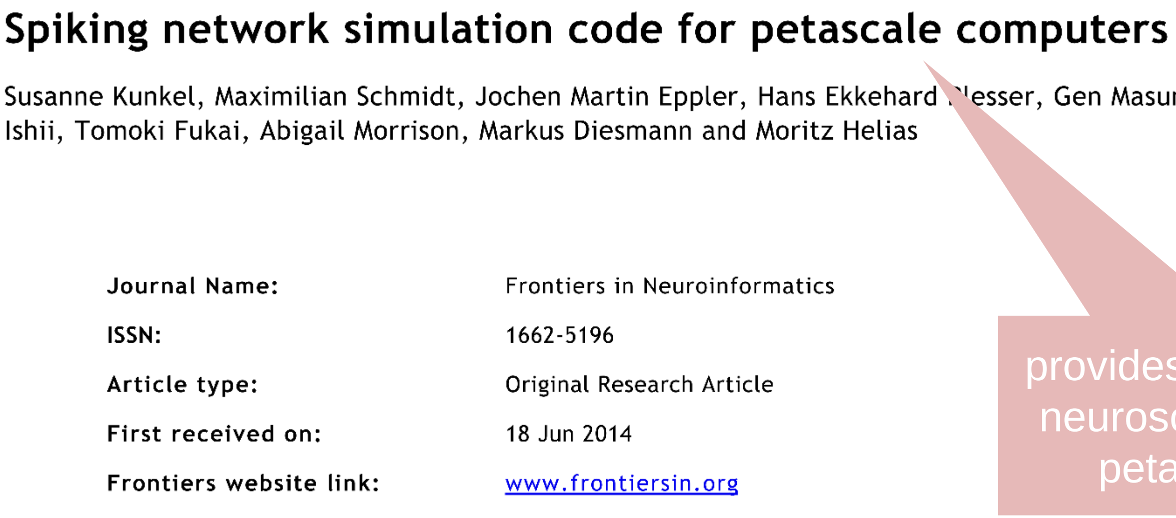

Spiking network simulation code for petascale computers

Susanne Kunkel, Maximilian Schmidt, Jochen Martin Eppler, Hans Ekkehard lesser, Gen Masumoto, Jun Igarashi, Shin Ishii, Tomoki Fukai, Abigail Morrison, Markus Diesmann and Moritz Helias

Journal Name:	Frontiers in Neuroinformatics
ISSN:	1662-5196
Article type:	Original Research Article
First received on:	18 Jun 2014
Frontiers website link:	www.frontiersin.org

provides the evidence that neuroscience can exploit petascale systems

- last paragraph of introduction:

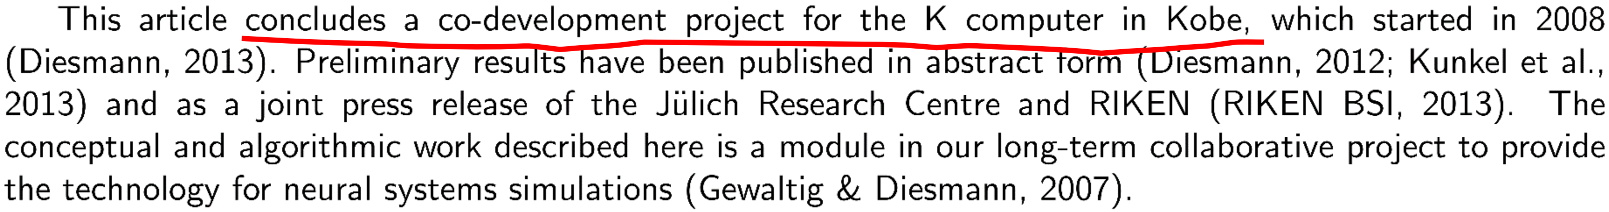

This article concludes a co-development project for the K computer in Kobe, which started in 2008 (Diesmann, 2013). Preliminary results have been published in abstract form (Diesmann, 2012; Kunkel et al., 2013) and as a joint press release of the Jülich Research Centre and RIKEN (RIKEN BSI, 2013). The conceptual and algorithmic work described here is a module in our long-term collaborative project to provide the technology for neural systems simulations (Gewaltig & Diesmann, 2007).

Summary

- full-scale model explains prominent features of network activity
- is building block of further studies (www.opensourcebrain.org)
- need for brain-scale models
- require memory only available on supercomputers
- machines ready for use by neuroscience (www.nest-initiative.org)
- co-development phase was essential for petascale technology
- NEST is simulation engine of HBP at cellular and synaptic resolution
- exascale computers present new challenges for data structures
- exascale computers offer new opportunities for simulation speed