

The K computer as an instrument to study brain-scale neuronal networks at microscopic resolution

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

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

Aug 26th 2014 SS-ken HPC Forum, Tokyo



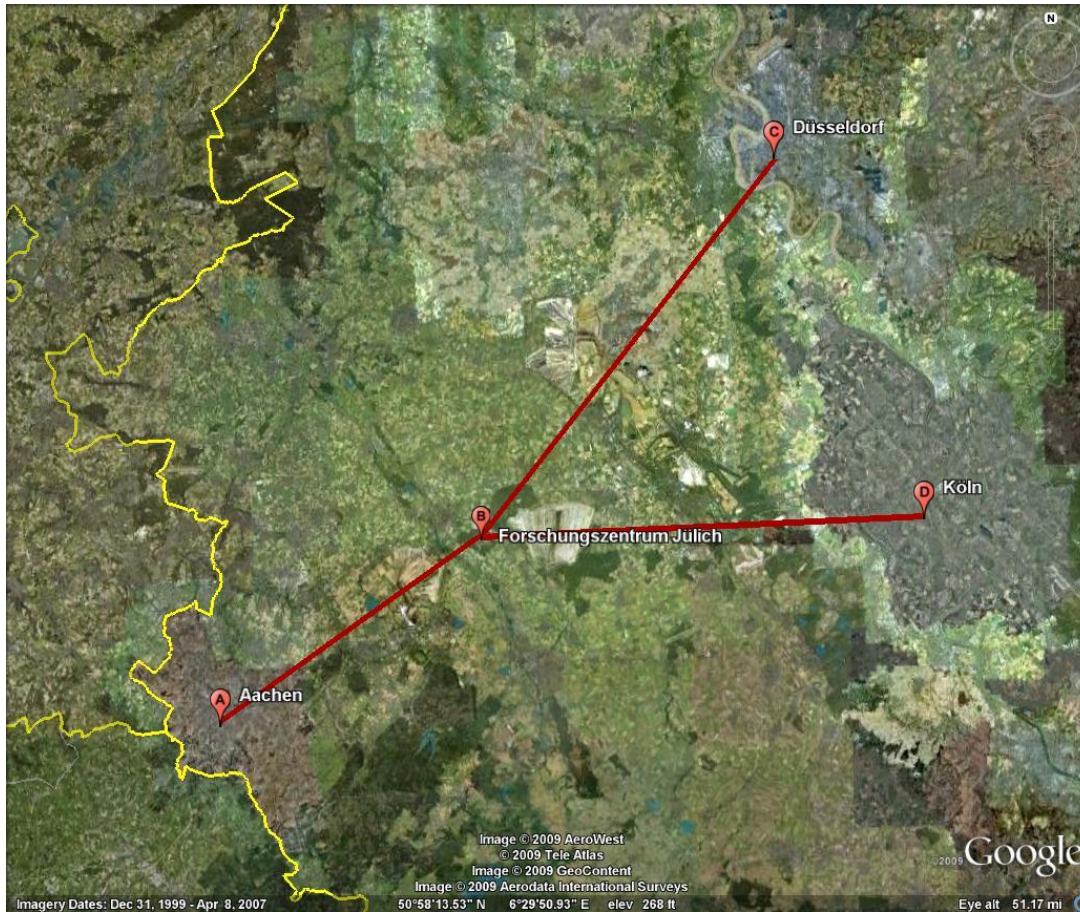
**multi-disciplinary
national research laboratory.**

budget: ~557 Million Euro / year

as per Dec 31st 2012:

- **Personnel: 5,236 in total**
 - **Scientists: 1,658**
 - **PhDs: 469**
 - **Professors: 93**

Region



- 30 km distance to 3 major universities

Research Center Jülich



■ 2.2 km² campus

Helmholtz Mission

We contribute to solving grand challenges which face society, science and industry by performing top-rate research in strategic programmes in the fields of Aeronautics, Space and Transport, Earth and Environment, Energy, Health, Key Technologies as well as the Structure of Matter.

*We research systems of great complexity with our **large-scale facilities and scientific infrastructure**, cooperating closely with national and international partners.*

JUQUEEN supercomputer

- inaugurated, Feb 14th 2013
- 458,752 cores
- 5.9 peta-flop
- 28 compared to 72 racks of JUGENE
- TOP 500, June 2014: 8th (2nd in Europe)





Human Brain Project

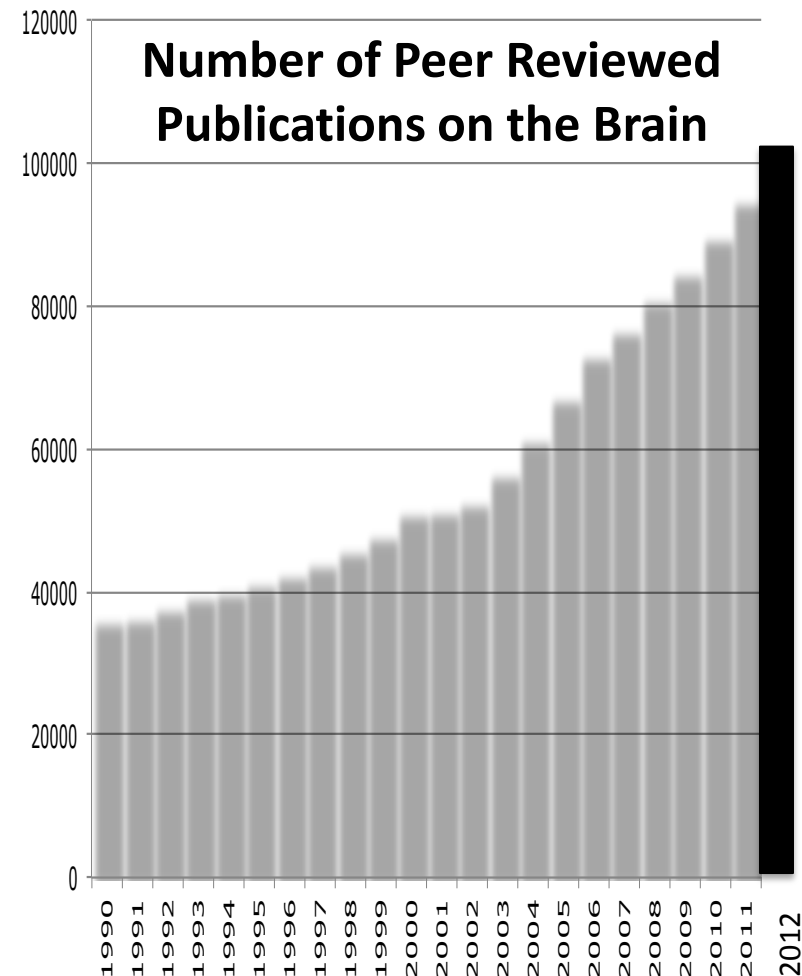
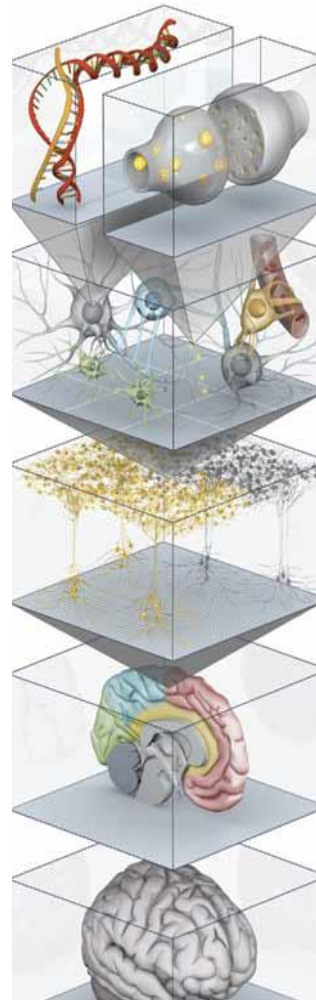
The Human Brain Project (HBP)



- understanding the human brain is still a great challenge
- difficulty: its complexity
- huge number of neurons (10^{11}) and synapses (10^4) per neuron connecting them
- proper description level unclear

Situation in Neuroscience

- data flood
- data fragmentation
- no integration plan
- no curation plan
- no plan to bridge levels
- no plan from animal to human
- lack of reproducibility
- small groups





EU FET (Future and Emerging Technologies) Flagships

- grand scientific and societal challenges which require a common European research effort
- **science - driven, large - scale, multidisciplinary research initiatives oriented towards a unifying goal**
- goals are **visionary and highly ambitious**
- **require cooperation** among a range of disciplines
- boosting European innovation
- sustained support for a period of up to 10 years
- substantial benefits for European competitiveness and society



Human Brain Project

- started in Oct 2013, Kickoff summit in Lausanne
- ramp-up phase (30 months): €54 million from EU Information and Communication Technologies (ICT) 2013 Work Program
- careful monitoring - report: every 6 months
- operational phase: finance will come from the Horizon 2020 Research Program
- Horizon 2020 proposal presently under review



Human Brain Project

Human Brain Project

1.1.1 The objective

The goal of the Human Brain Project is to build a completely new ICT infrastructure for future neuroscience, future medicine and future computing *that will catalyse a global collaborative effort* to understand the human brain and its diseases and ultimately to emulate its computational capabilities.



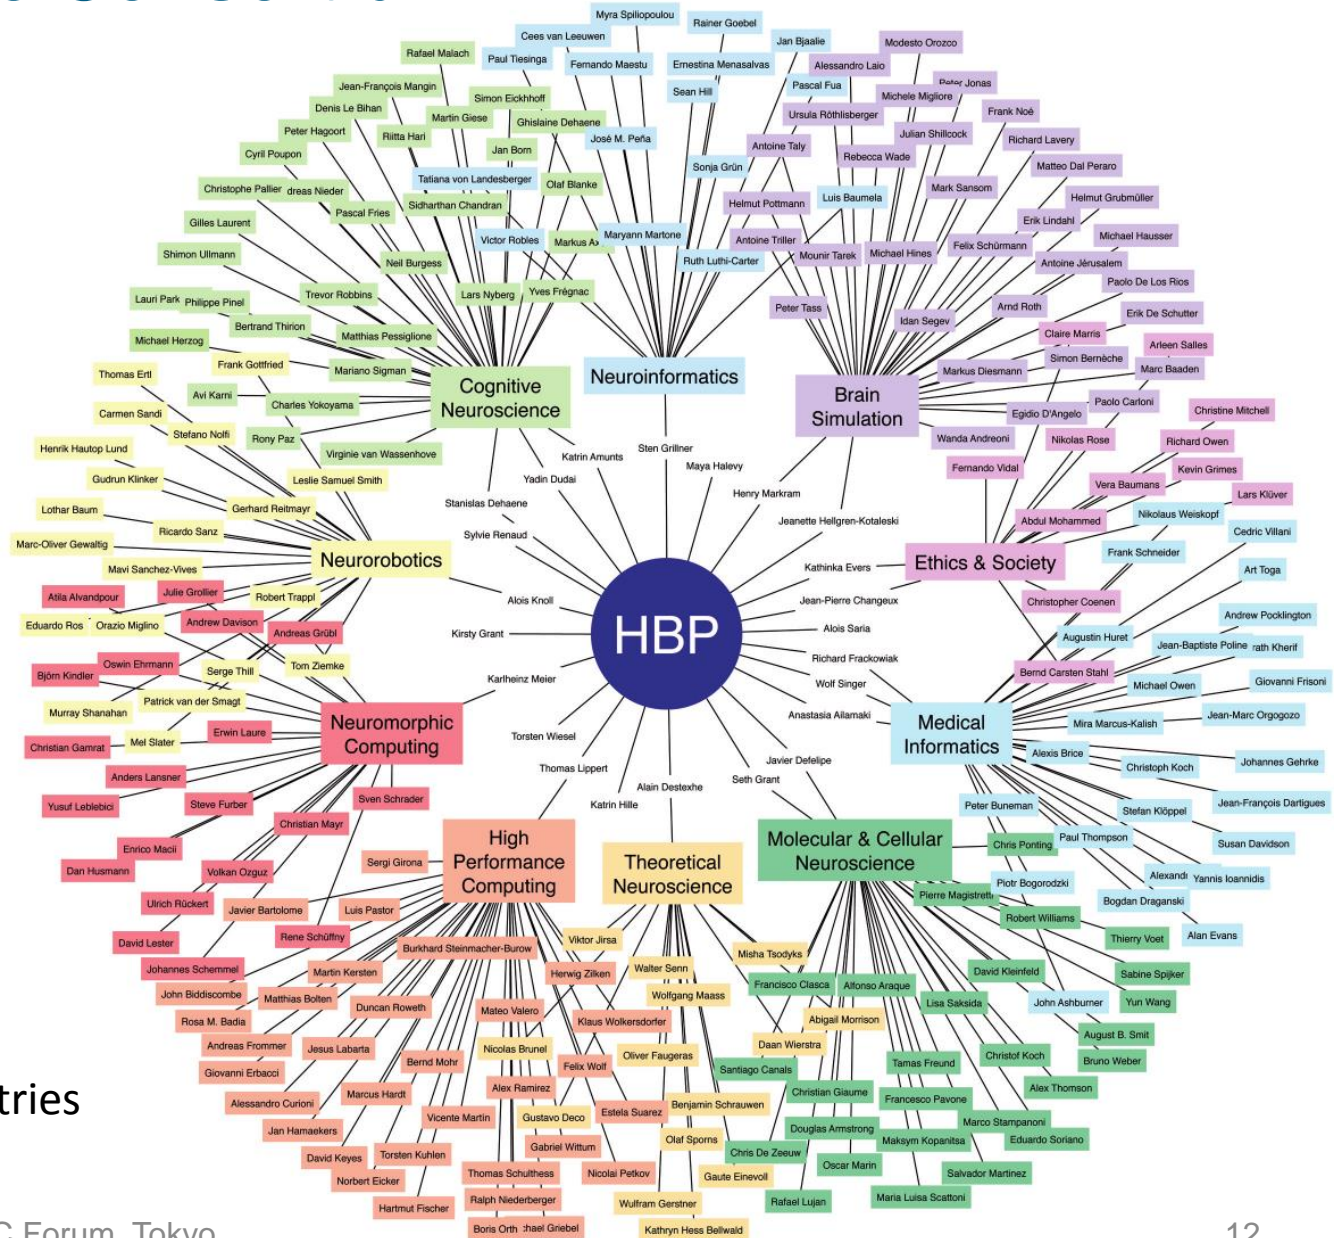
Human Brain Project



EUROPEAN
COMMISSION



The Consortium



- 80 Institutions
- 150 PIs
- 2000 PhDs
- 20 European Countries

HBP as a large-scale scientific project, model: Large Hadron Collider (LHC) at CERN

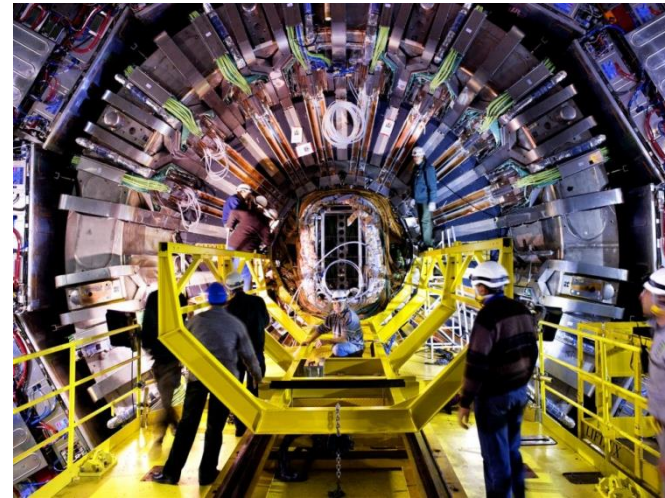
Basic Facts

Construction: 1998-2008

Consists of a 27km ring of superconducting magnets: Thousands of magnets of different varieties and sizes are used to direct the beams around the accelerator. These include 1232 dipole magnets 15m in length which bend the beams, and 392 quadrupole magnets, each 5–7m long, which focus the beams.

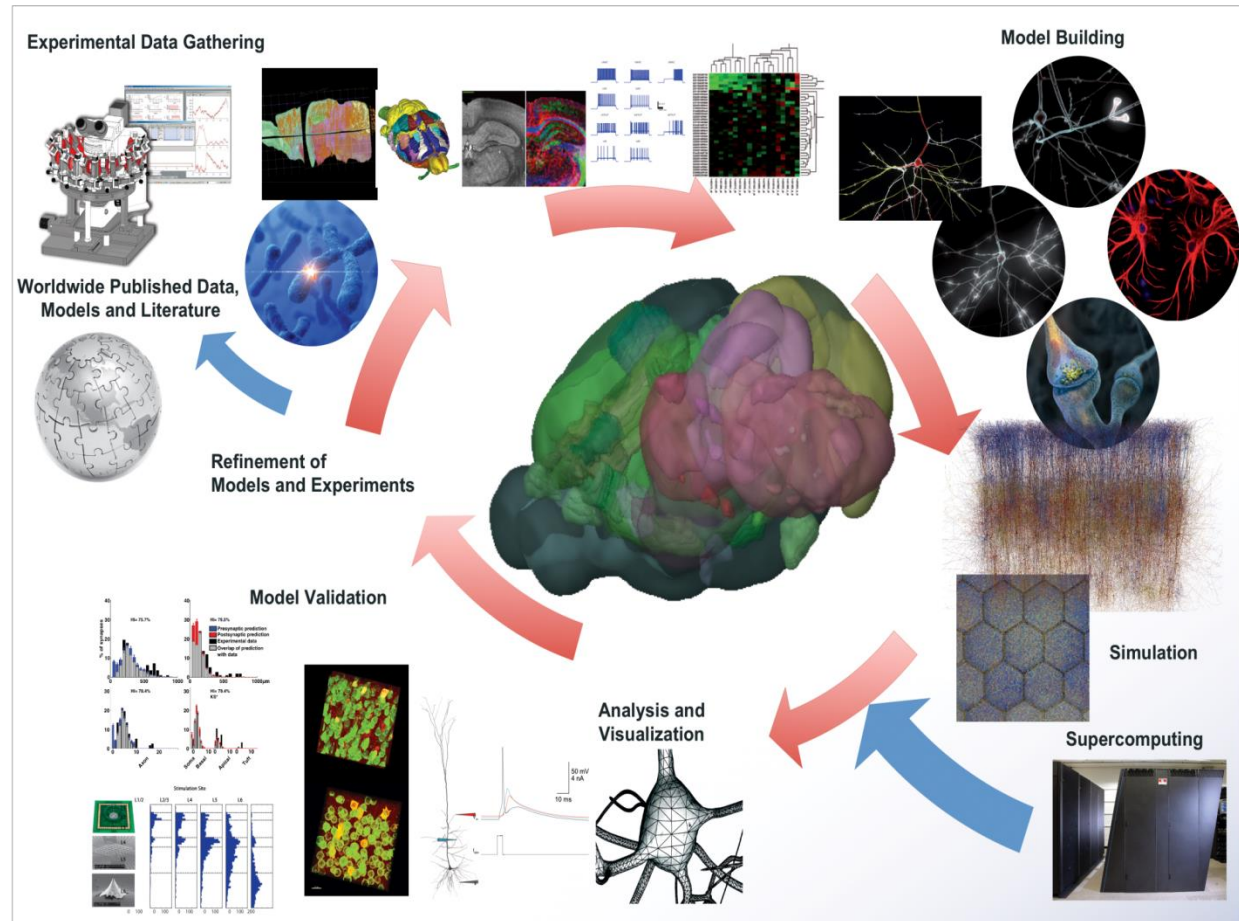
Built in collaboration with over 10,000 scientists and engineers from over 100 countries, as well as hundreds of universities and laboratories.

neuroscience needs to learn to collaborate



Build, Simulate and Validate

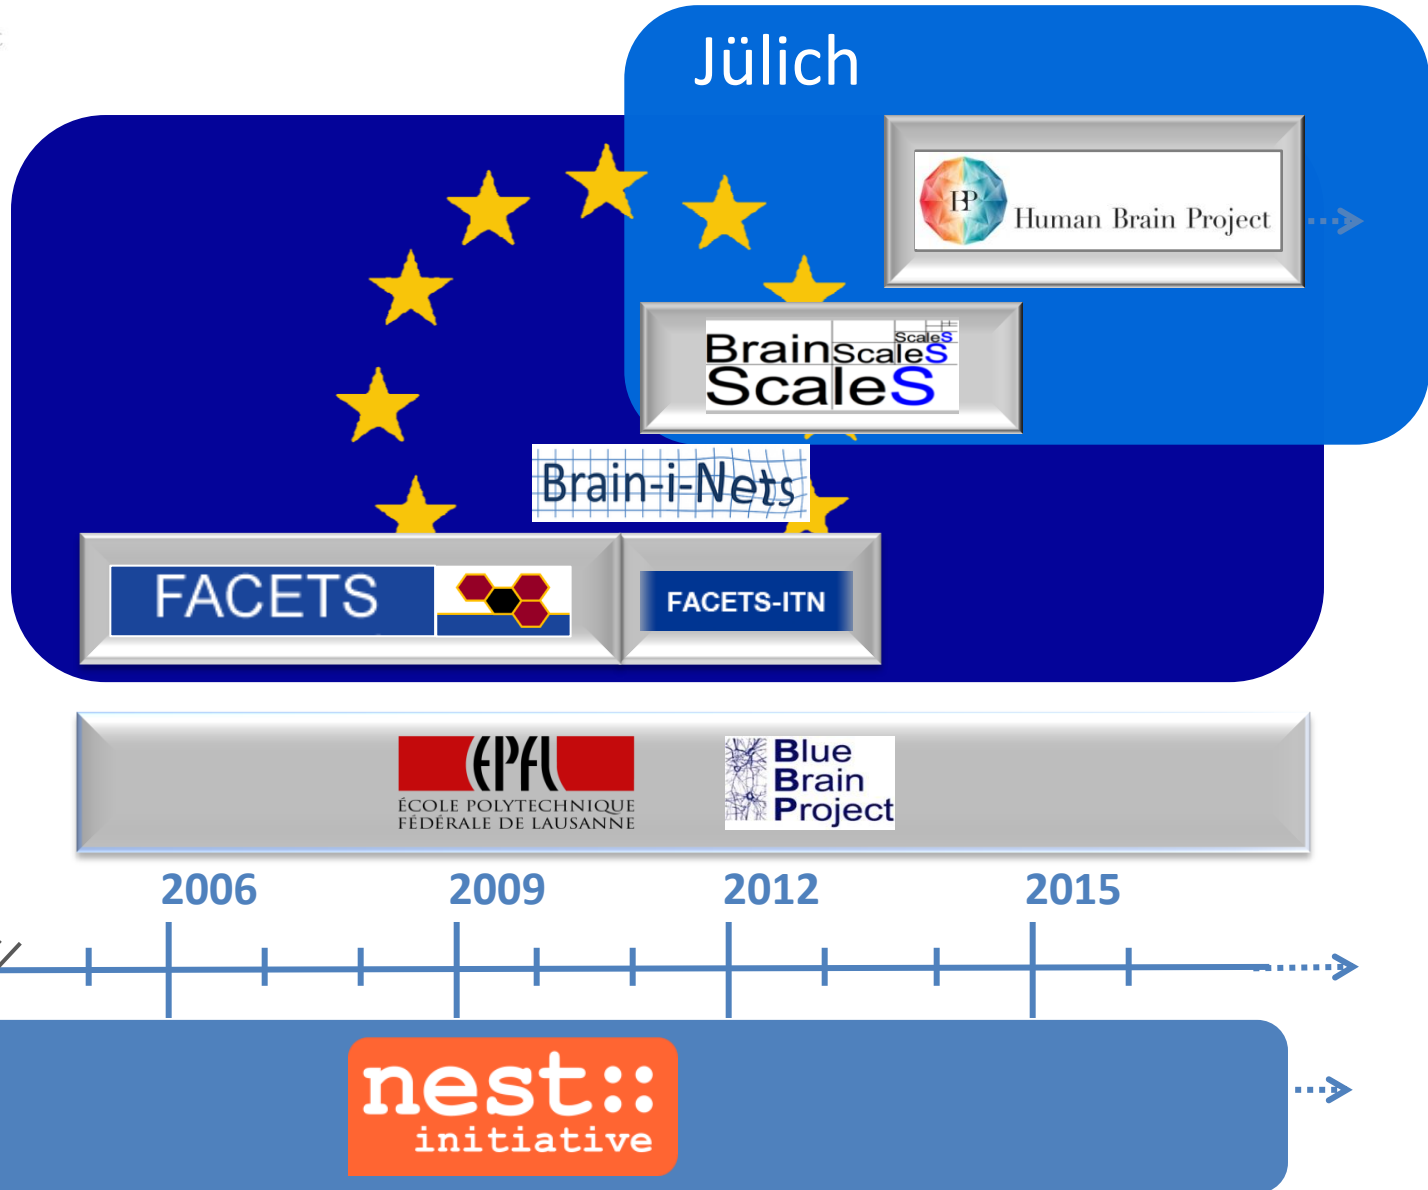
- workbench through web portal
- integrated workflow
- access to platforms
- provenance tracking





Human Brain Project

The EU Computational Neuroscience path towards the Human Brain Project



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]

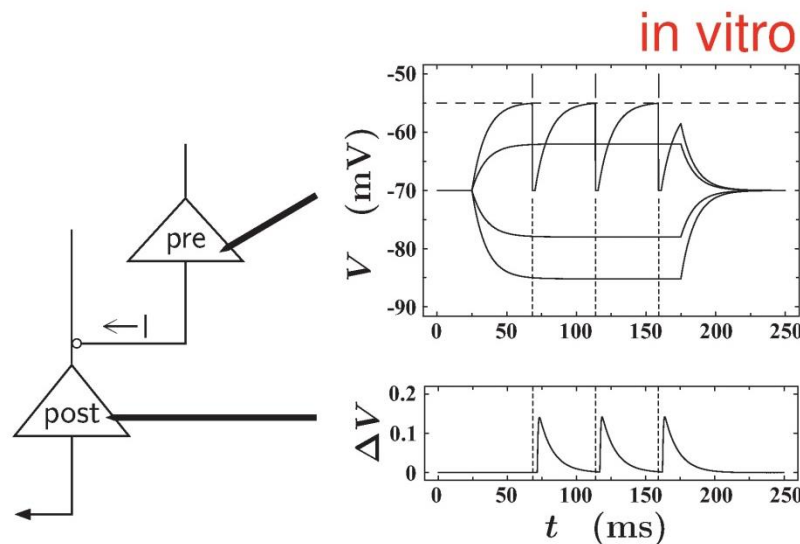


4th generation simulation kernel,
platform in initial phase of HBP

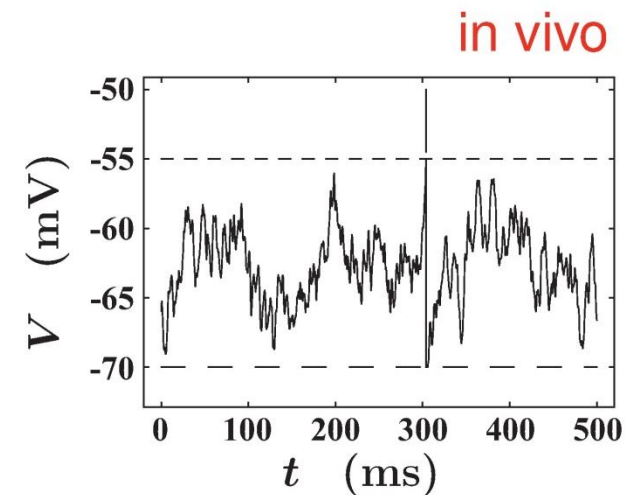
[scale of 100,000 compute nodes]



Interactions between neurons



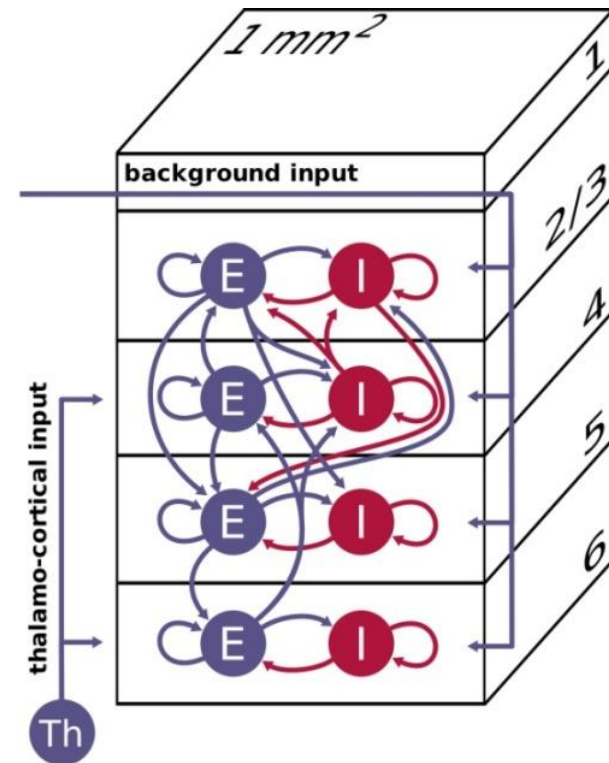
- current injection into pre-synaptic neuron causes excursions of membrane potential
- supra-threshold value causes spike transmitted to post-synaptic neuron
- post-synaptic neuron responds with small excursion of potential after delay
- inhibitory neurons (20%) cause negative excursion



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

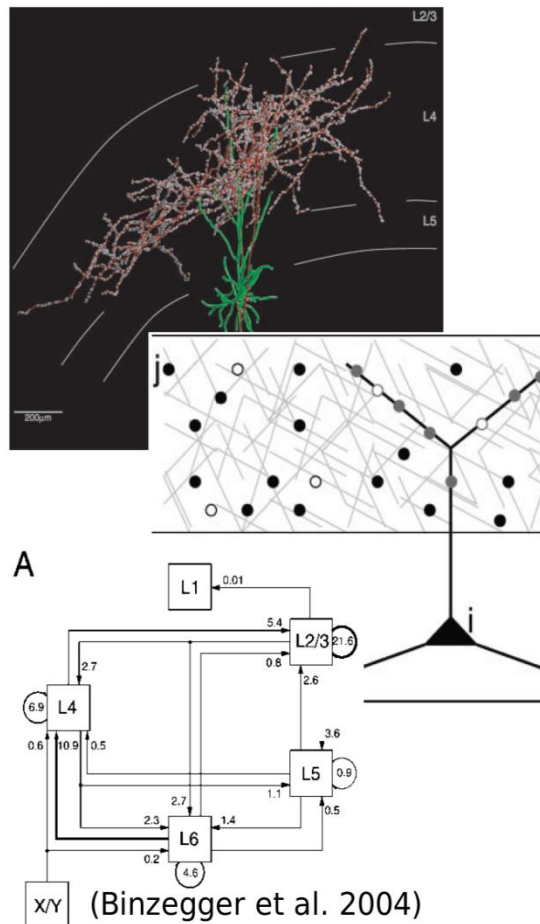
Minimal layered cortical network model

- 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}

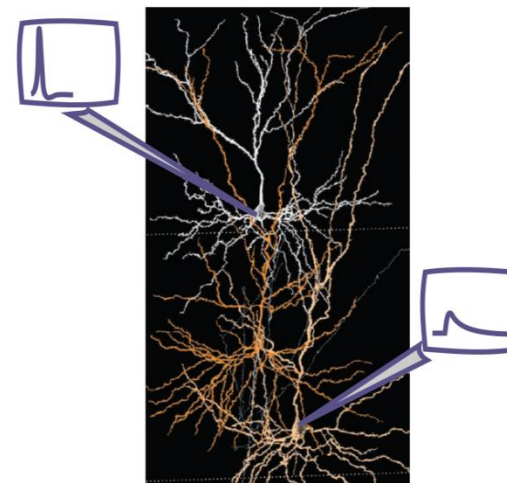


Anatomical data sets

in vivo anatomy



in vitro physiology

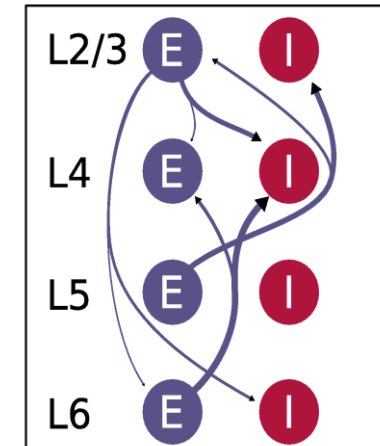
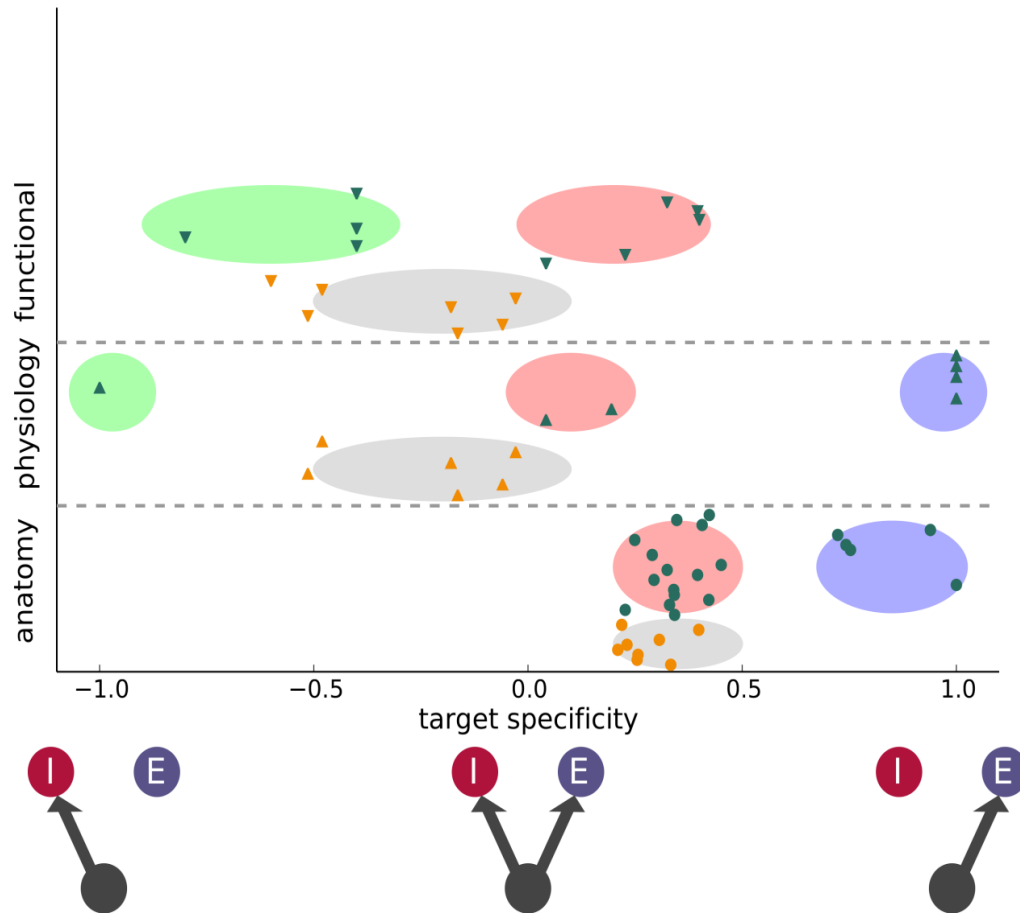


Type of connection^a Connectivity ratio

L5 pyramid to L5 pyramid	1:11 (15:163)
L2/3 pyramid to L2/3 pyramid	1:4 (65:247)
	1:10 (8:81)
L4 excitatory to L4 excitatory	1:5.7 (4:23)
L3 pyramid to L5 pyramid [Postsynaptic apical dendrite]	1:1.8 (16:29)
	1:1 (2:2)
L5 pyramid to L3 pyramid	1:29
L4 excitatory to L3 pyramid	1:3.6 (7:25)
(Presynaptic spiny stellates) ($n = 4$)	1:10 (7:70)
L5 pyramid to L5 interneuron	1:10.4 (7:73)
L5 interneuron to L5 pyramid	1:8 (9:73)

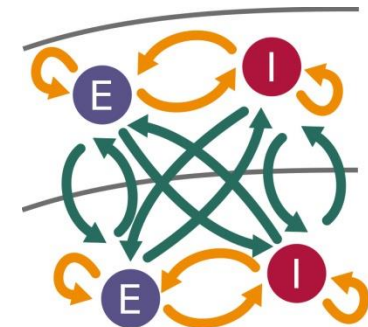
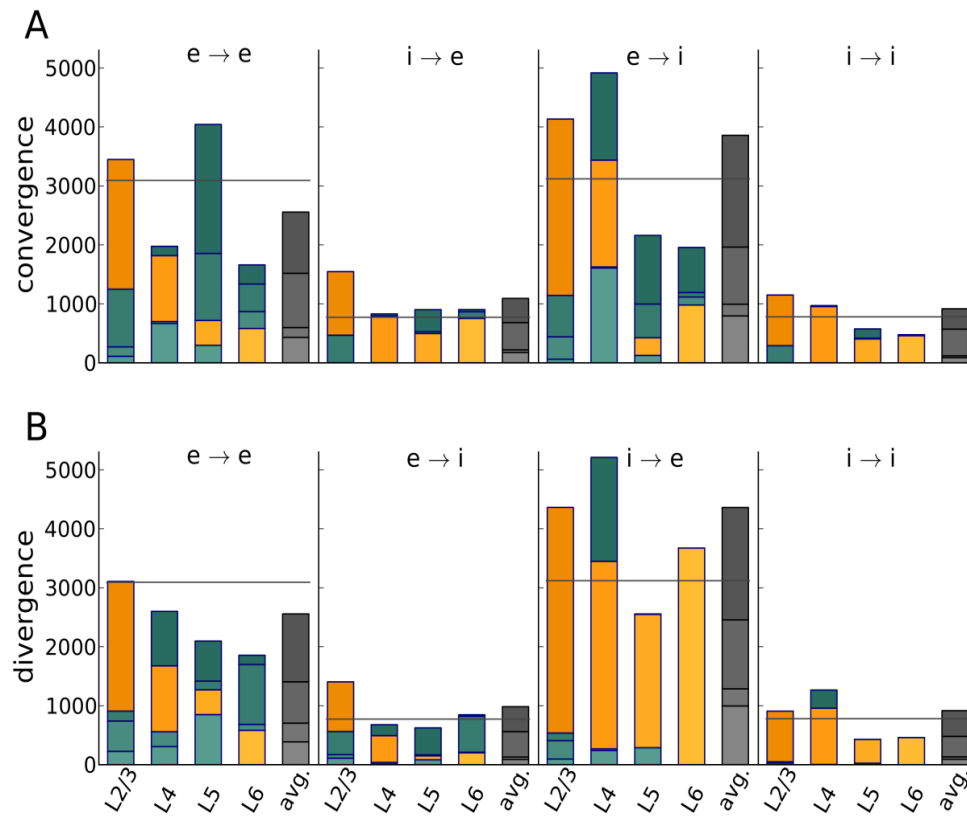
(Thomson et al. 2002)

Target specificity



- correction for bias in anatomical method

Convergence and divergence



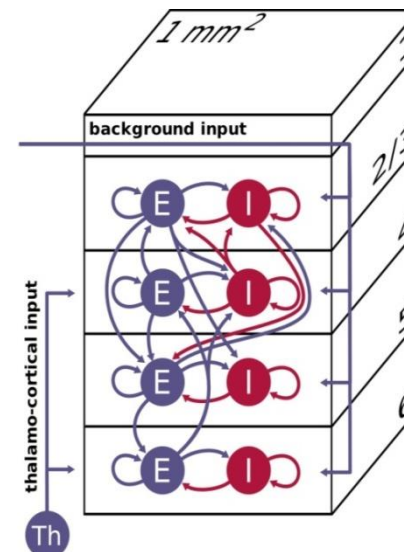
- dominated by within-layer connections
- $e \rightarrow e$ divergence reflects "standard" loop
- $e \rightarrow i$ divergence reflects target-specific feedback

Local cortical microcircuit

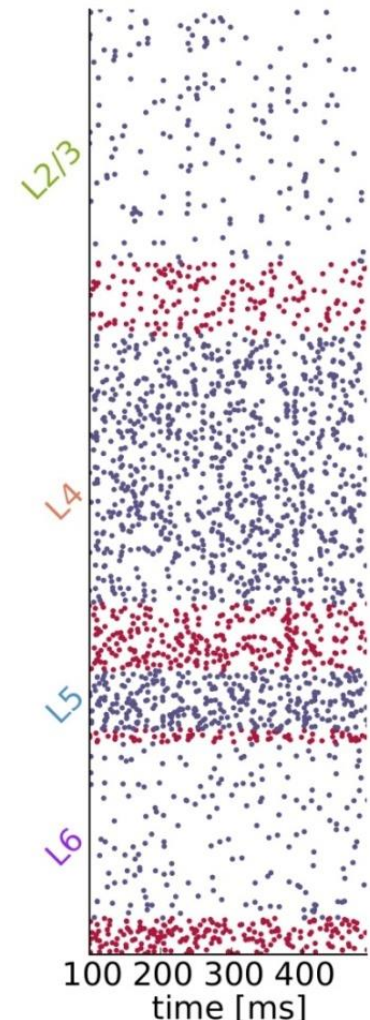
taking into account layer and neuron-type specific connectivity is sufficient to reproduce experimentally observed:

- 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

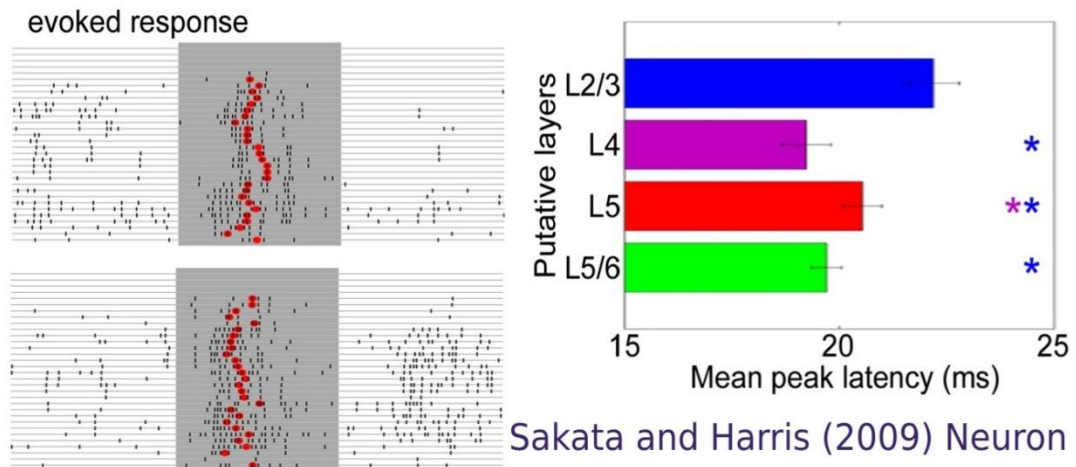
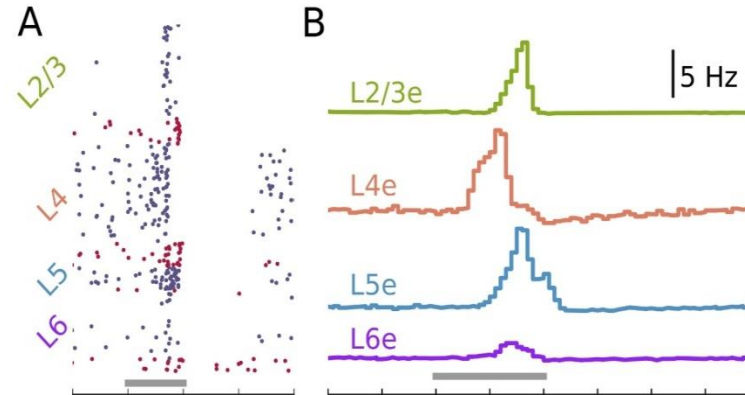


10^5 neurons
 10^9 synapses



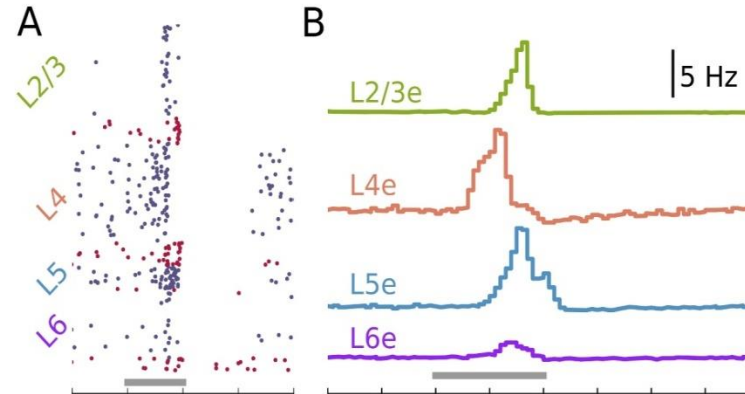
available at: www.opensourcebrain.org

Response to transient inputs

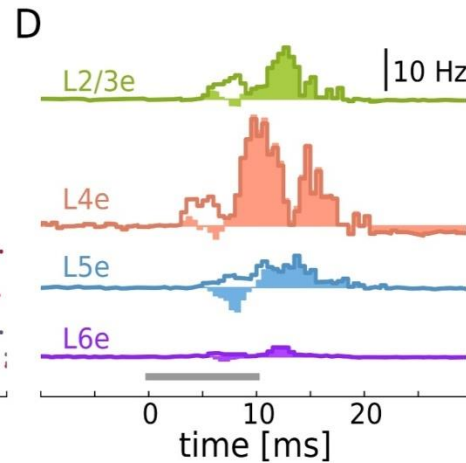
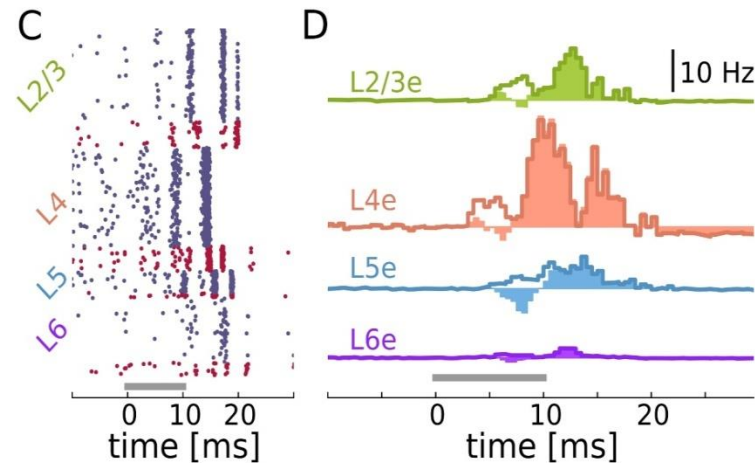


Response to transient inputs

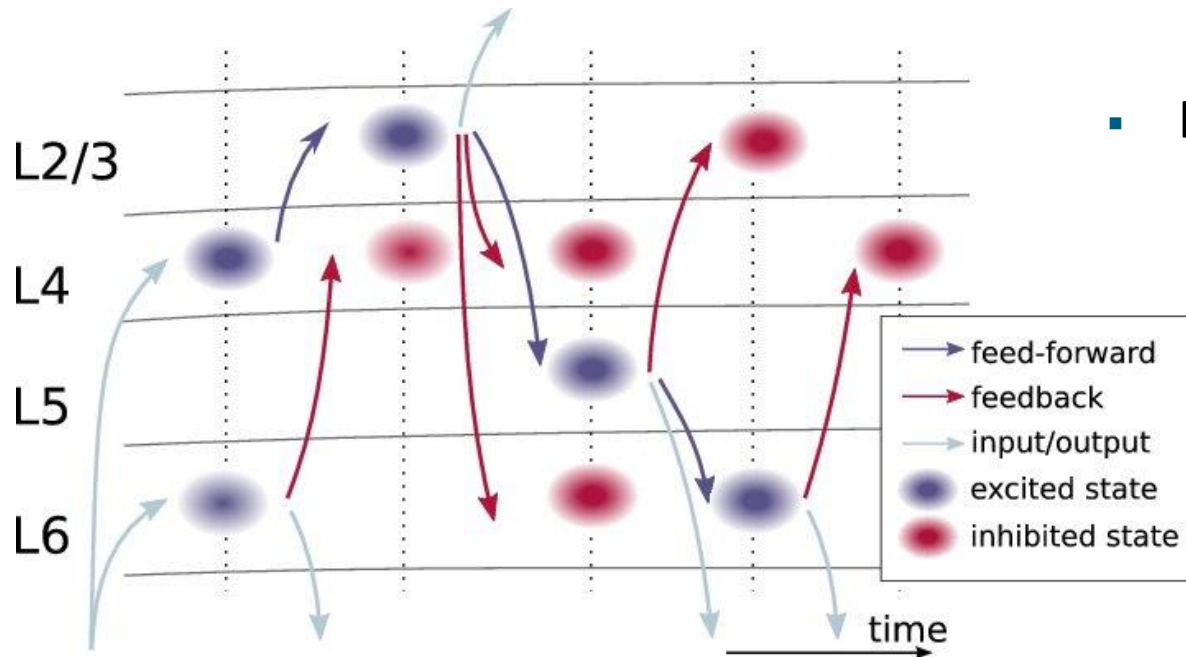
■ $T = -0.4$



■ $T = +0.4$



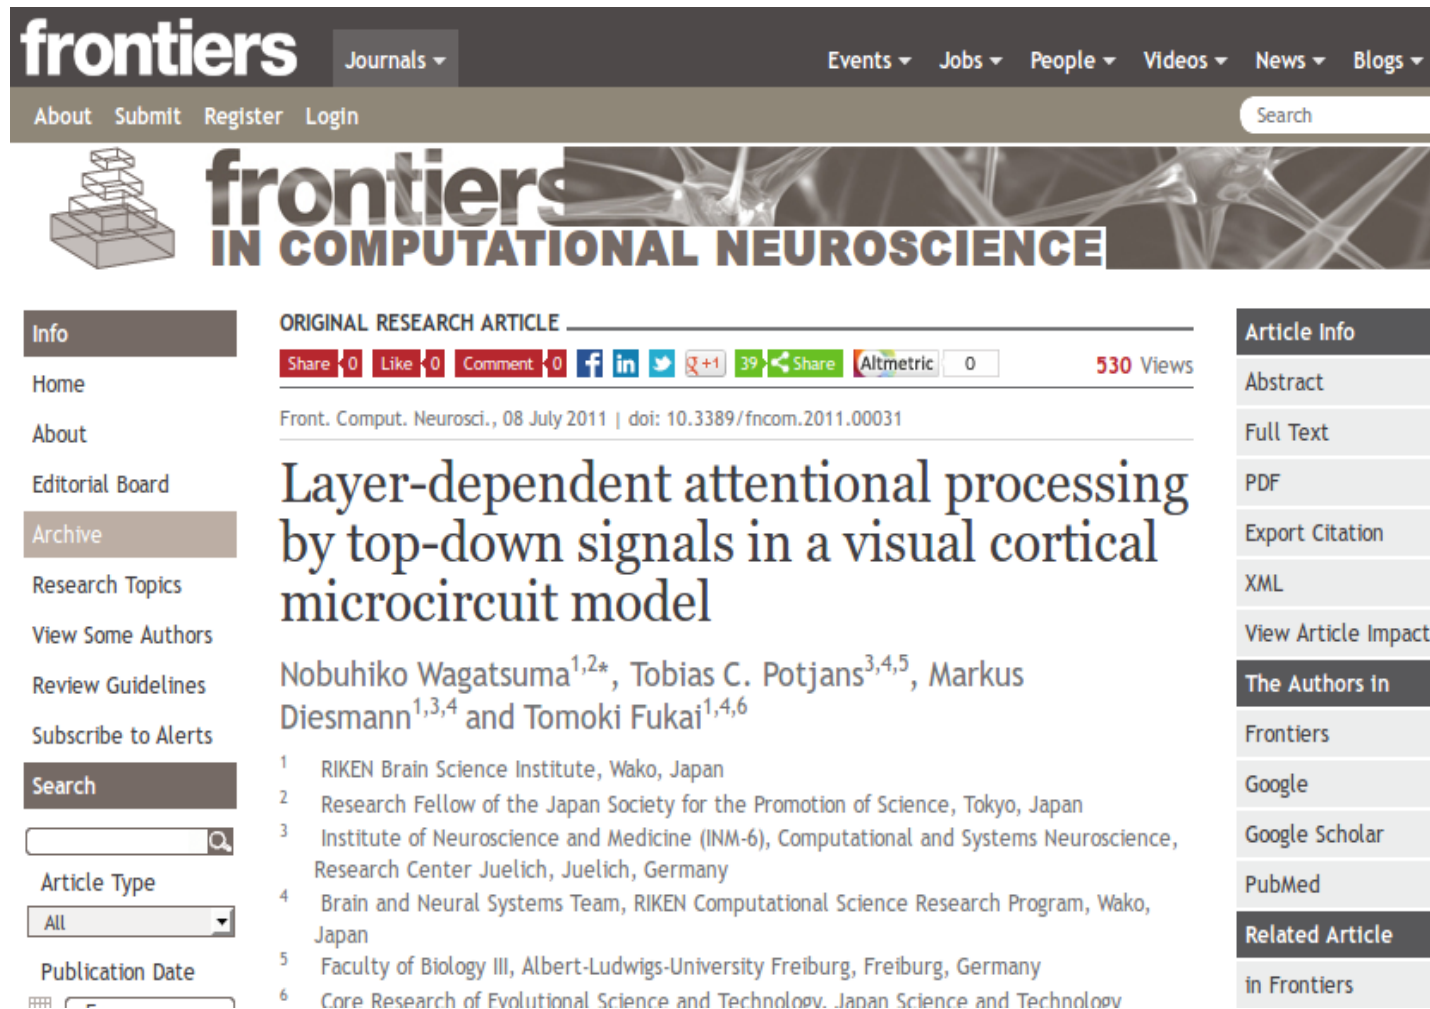
Hypothesis on cortical flow of activity



- handshaking between layers

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

Building block for functional studies



frontiers Journals ▾ Events ▾ Jobs ▾ People ▾ Videos ▾ News ▾ Blogs ▾

About Submit Register Login Search

frontiers
IN COMPUTATIONAL NEUROSCIENCE

Info

- Home
- About
- Editorial Board
- Archive
- Research Topics
- View Some Authors
- Review Guidelines
- Subscribe to Alerts

Search

Article Type: All

Publication Date: [Calendar icon]

ORIGINAL RESEARCH ARTICLE

Share 0 Like 0 Comment 0 +1 39 0 **530 Views**

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

Article Info

- Abstract
- Full Text
- PDF
- Export Citation
- XML
- View Article Impact

The Authors in

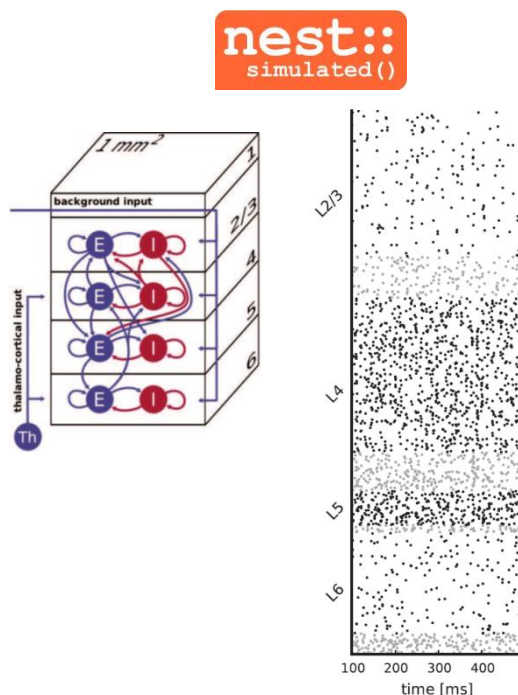
- Frontiers
- Google
- Google Scholar
- PubMed

Related Article

in Frontiers

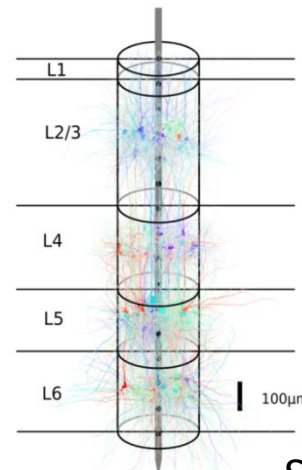
Building block for mesoscopic studies

- collaboration with Gaute Einevoll (UMB, Norway)

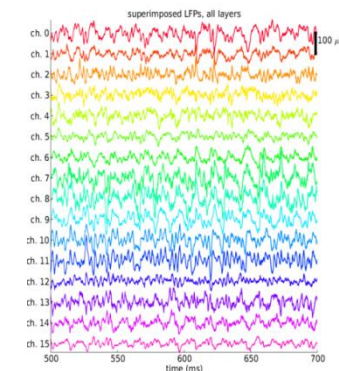


NEURON
for empirically-based simulations of
neurons and networks of neurons

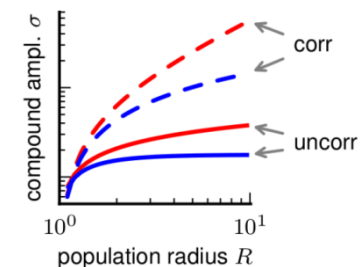
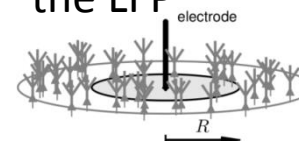
Morphology



Laminar LFP profile



Spatial reach of
the LFP

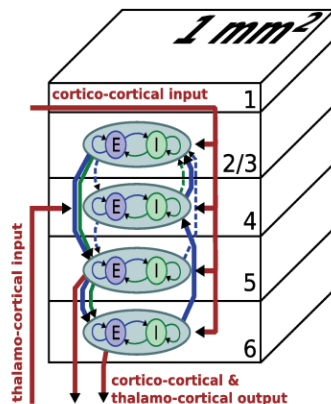


Linden H, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT (2011) Modeling the spatial reach of the LFP. *Neuron* 72(5):859-872

Critique of local network models

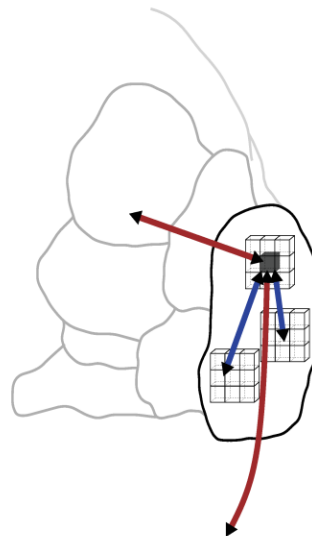
Local Cortical Network Connections:

- ✓ local connections
- ✓ realistic synaptic modeling
- ✗ a major part of synapses missing
- ⇒ input dependent local network dynamics



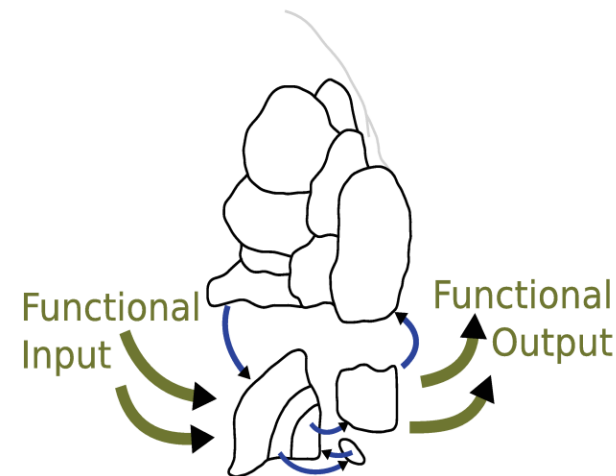
Cortical Area Network Connections:

- ✓ intrinsic connections
- ✗ many synapses missing



Brain-scale Connections:

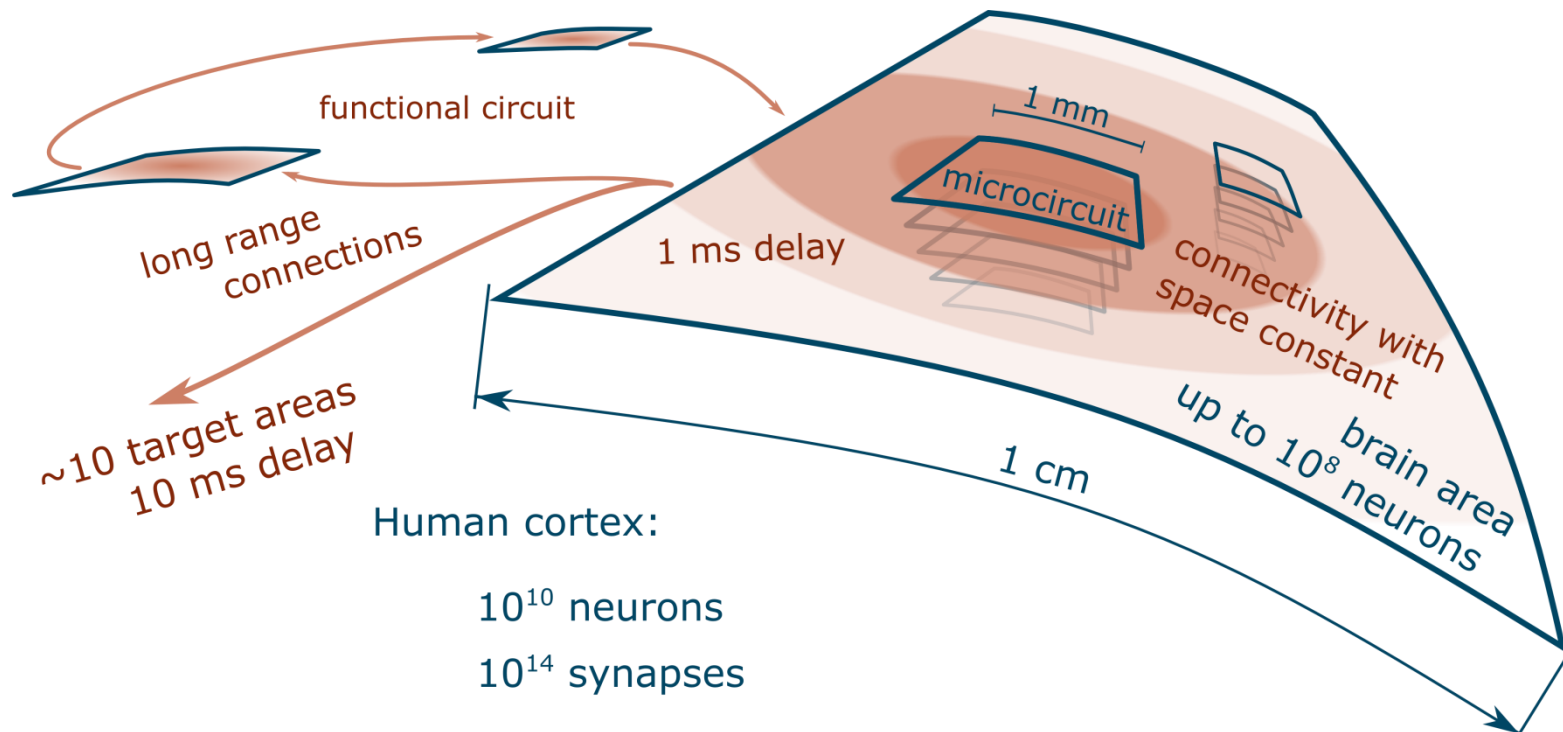
- ✓ all connections
- ⇒ bottom-up and top-down approach meet on this level



Brain and Neural Systems Team, RIKEN Computational Science Research Program
Pilot study: jinb33 (2008) *Jugene Brain-scale simulations* FZ Juelich

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:

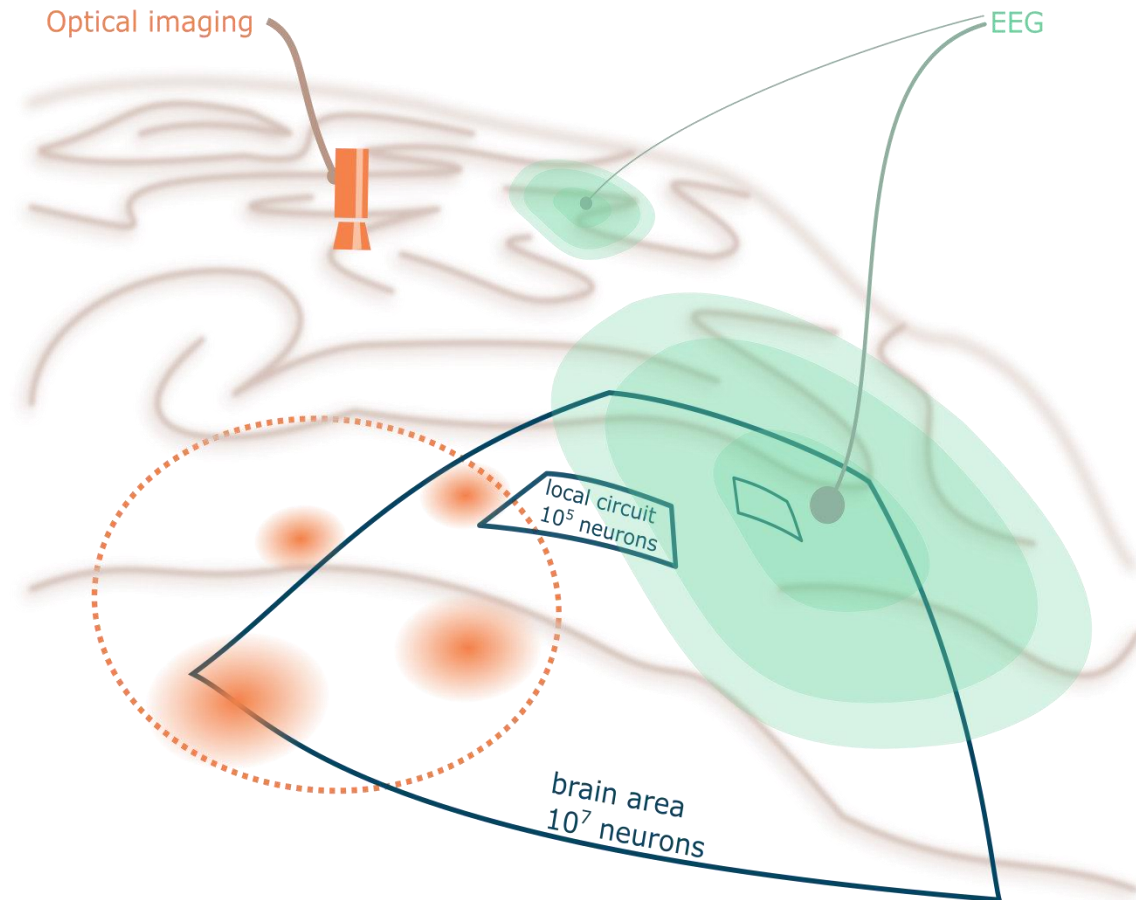
- further measures by forward modeling
- comparison with mean-field models

mesoscopic measures

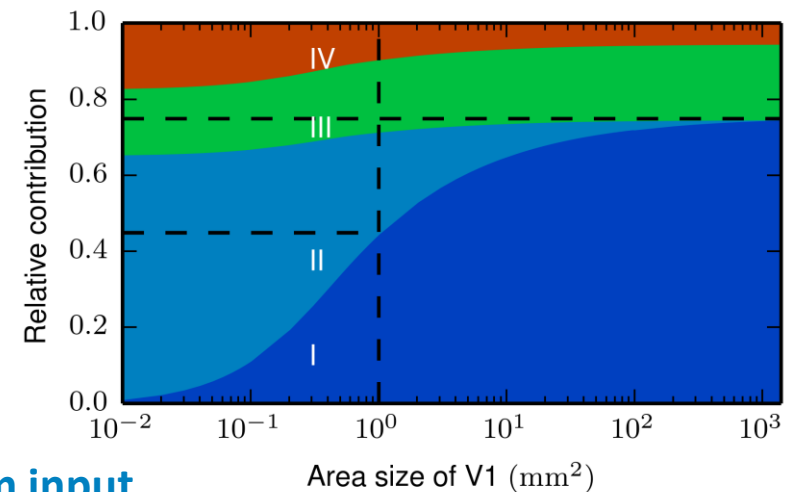
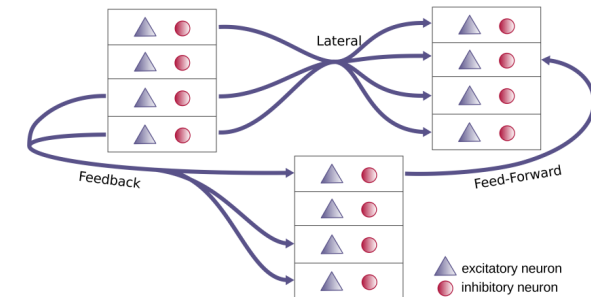
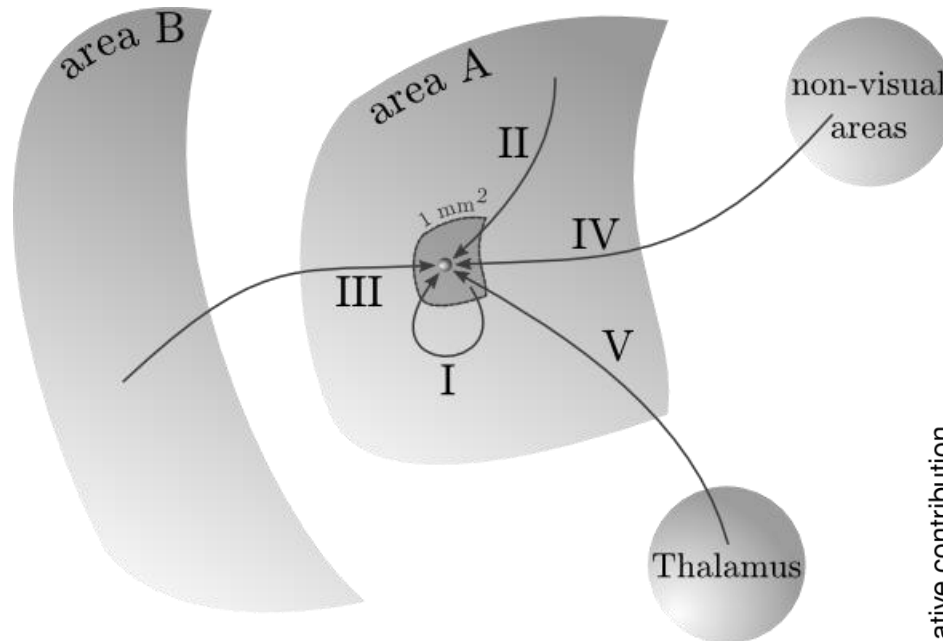
- local field potential (LFP)
- voltage sensitive dyes (VSD)

and macroscopic measures

- EEG, MEG
- fMRI resting state networks



Toward a self-consistent model

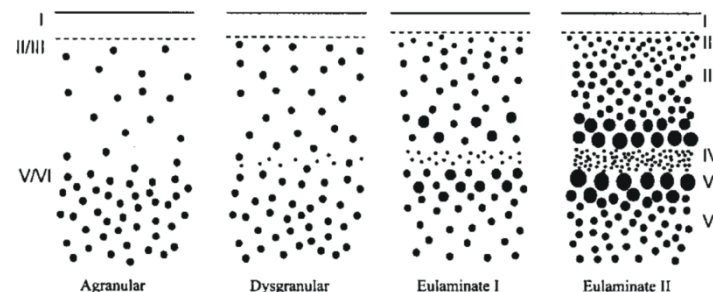
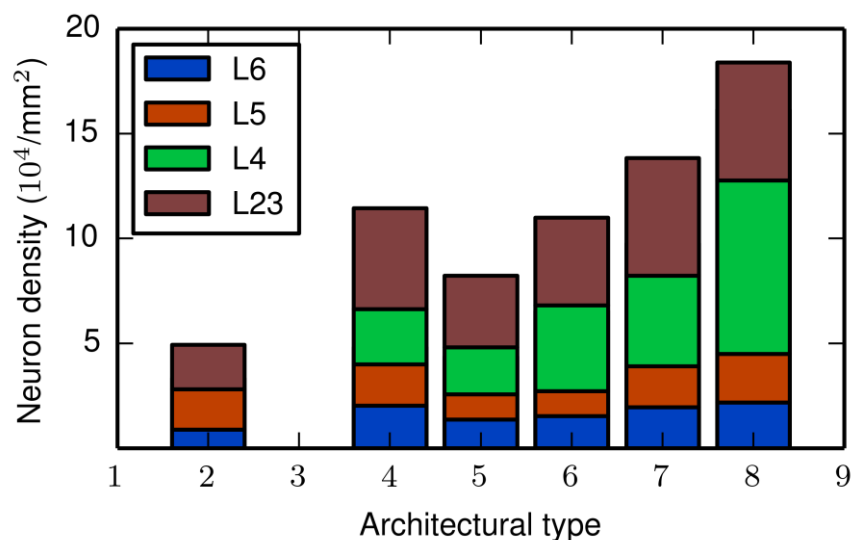
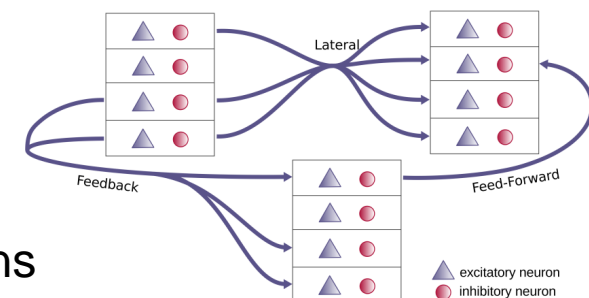


- I. Intra-areal synapses
- II. Intra-areal synapses replaced by random input
- III. Cortico-cortical synapses
- IV. External input represented by random input
- V. Thalamic input

- Sacha van Albada
- Maximilian Schmidt
- Rembrandt Bakker

Multi-area model of macaque visual cortex

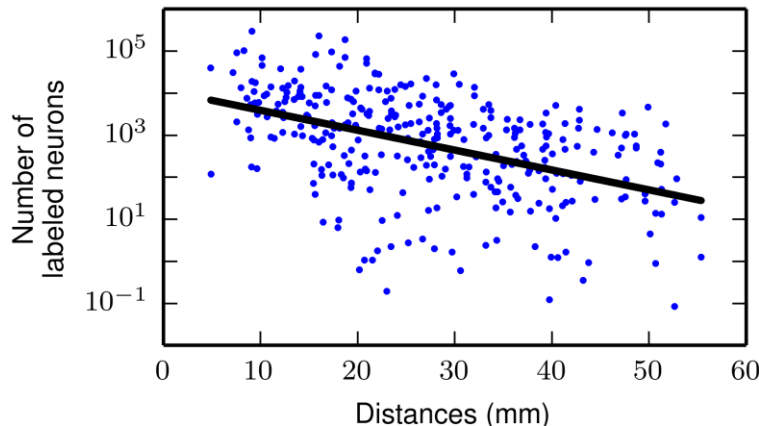
- rich anatomical data sets available (e.g CoCoMac)
- close to human
- 32 areas structured in layers comprising $8 \cdot 10^8$ neurons
- downscaled model with $3.2 \cdot 10^6$ neurons and $3 \cdot 10^{10}$ synapses



From Dombrowski et al. (2001), Cereb Cortex

architectural types provided by C. Hilgetag (private communication)

From Markov et al. (2014),
J. of Comparative Neurology



[cocmac2-gnode.org/cocmac2/services/f99_region_overlap.php](#)
[cocmac2-gnode.org/cocmac2/services/f99_region_overlap.php](#)
[cocmac2-gnode.org/cocmac2/services/f99_region_overlap.php](#)

CoCoMac [services](#) [F99 region overlap](#)

Generate table with region overlap between atlases that are registered to F99

Given a selected pair of atlases that are registered to the F99 cortical surface, this tool counts the vertices that they have in common and returns the overlap as an absolute number or percentage. Note that the surface area represented by a vertex may vary; this is not taken into account.

The following paint files from the Van Essen lab are used:

- [Macaque F99_COMPOSITE.74k_f99.asci.paint](#)
- [M129.R.Markov_CC10_Areas.74k_f99.paint](#)
- [M132.L.areas_F2-injection.74k_f99.paint](#)
- [Macaque F99UA1_RegionalMap.73730.asci.paint](#)

Select atlas 1 :

Select atlas 2 :

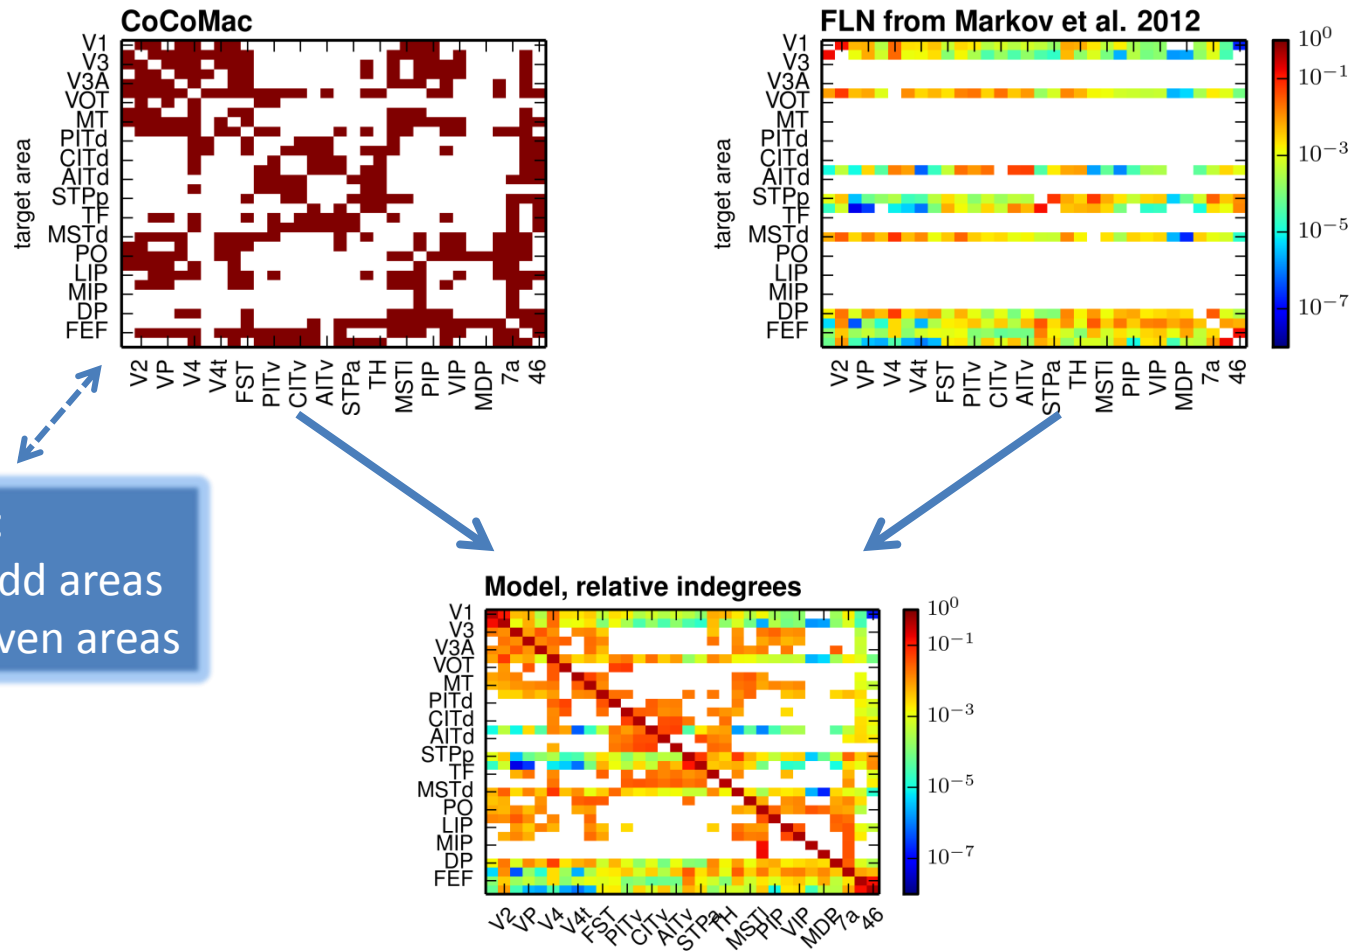
Normalization to compute percentage :

Output type :

[Get overlap table](#)

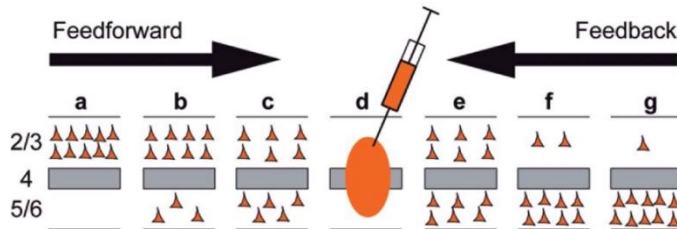
Direct link : http://cocmac2-gnode.org/cocmac2/services/f99_region_overlap.php?atlas1=FV91&atlas2=B05&normalize=avg&output=tsv&run=1

Cortico-cortical connectivity

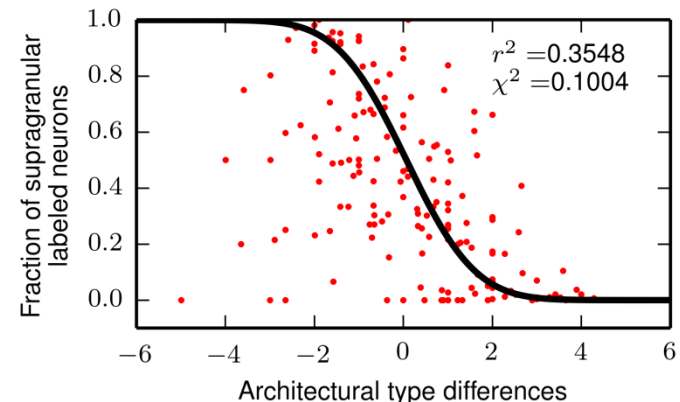


Laminar patterns

Sending side

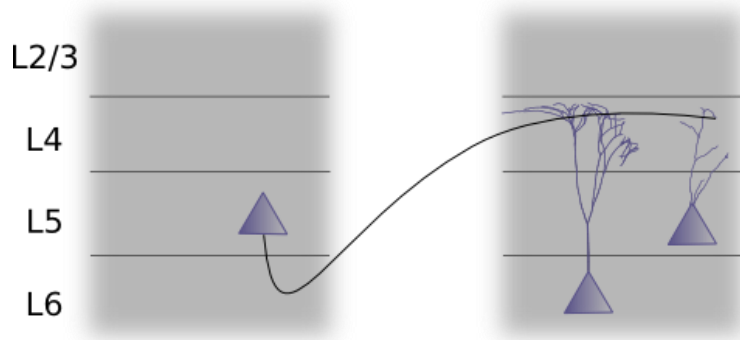


From Markov et al. (2014), J. of Comparative Neurology



Receiving side

- synapse layer: CoCoMac database
- receiving synapse type: Computed from Binzegger et al. (2004)



Fraction of cortico-cortical synapses in each layer

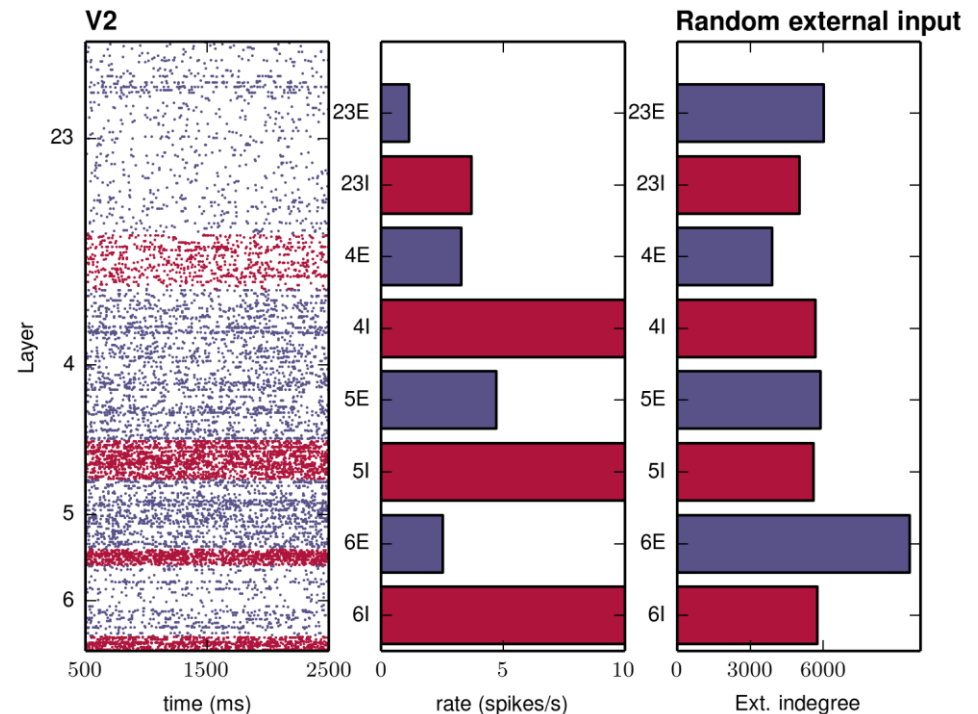
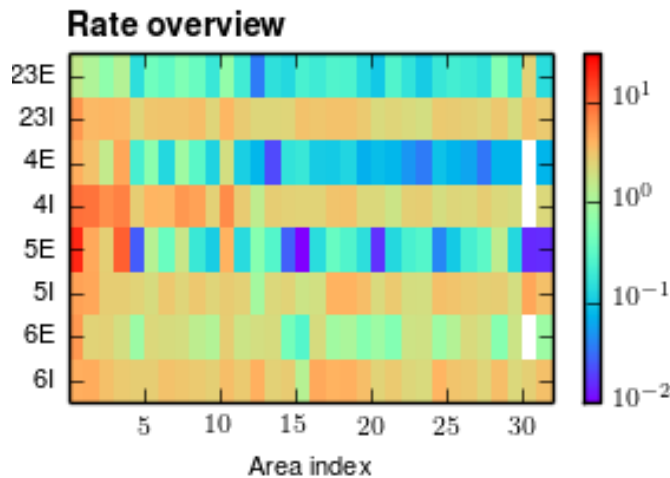
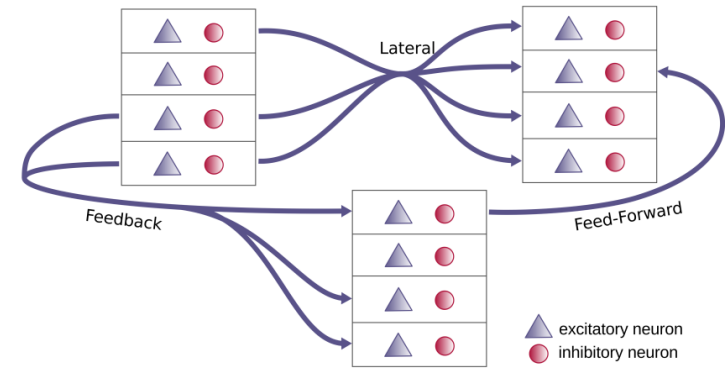
Cell body location

	23E	23I	4E	4I	5E	5I	6E	6I
1	0.57		0.18		0.25		0.003	
23		0.16	0.84					
4			0.73	0.16	0.03		0.09	
5					0.76	0.10	0.14	
6							0.85	0.15

From Binzegger et al. (2004), J Neurosci

Dynamical results

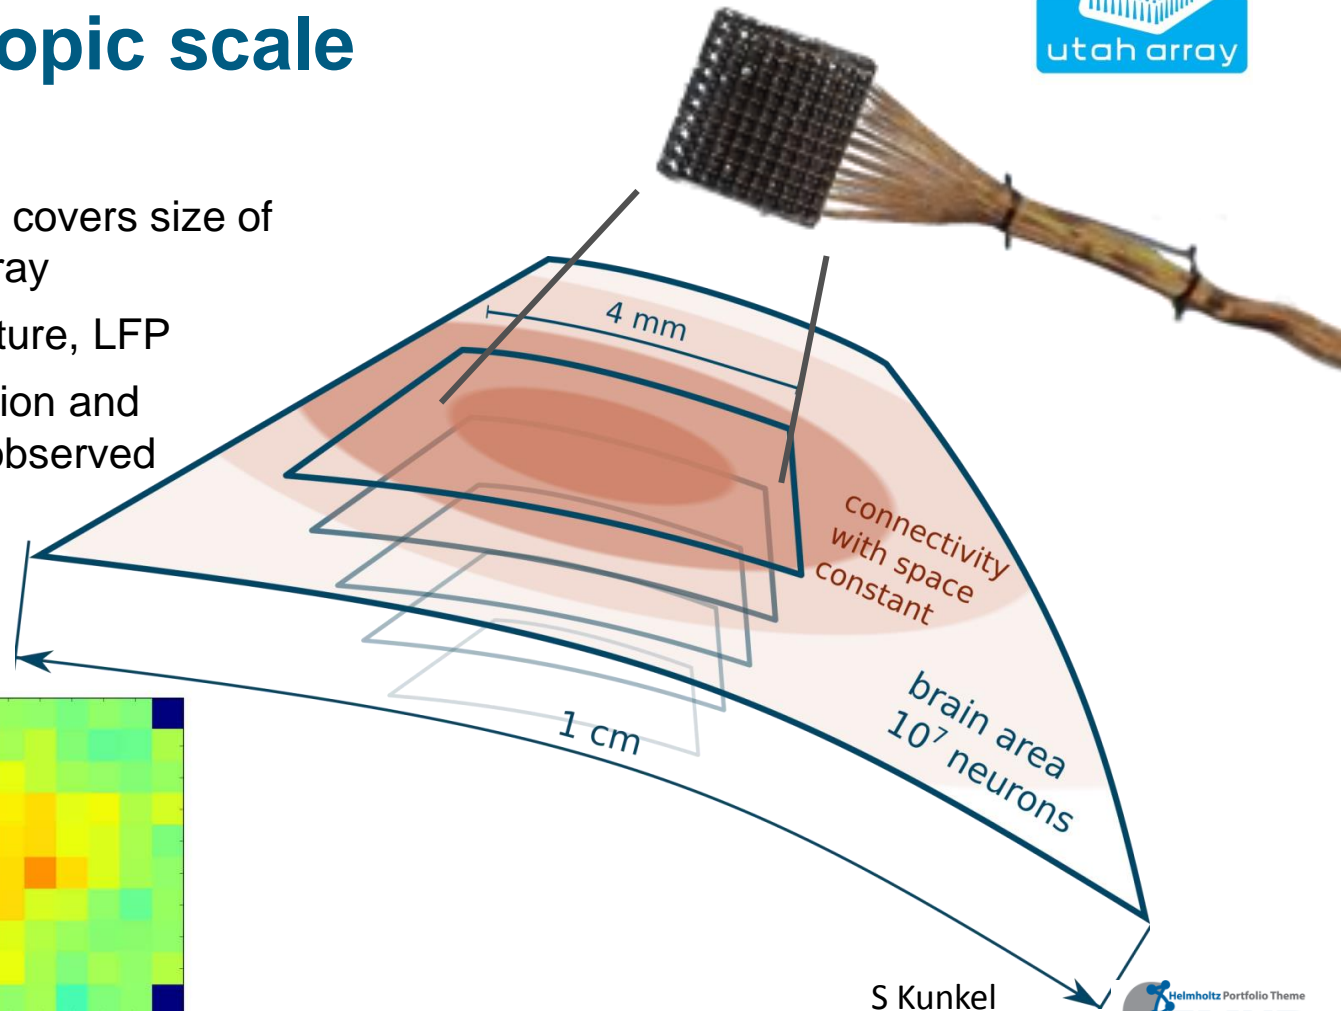
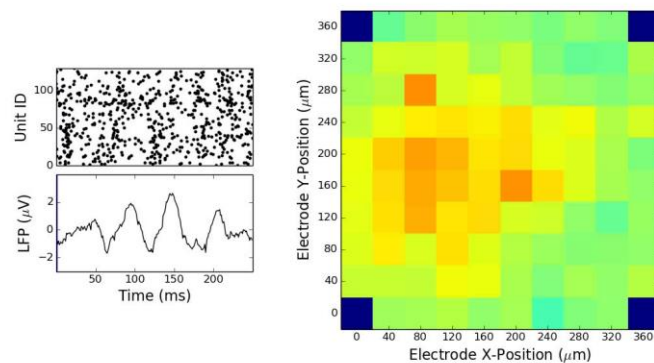
- Heterogeneous laminar firing patterns
- Rates in reasonable range (0.2 – 20 spikes/s)
- Inhibitory rates > excitatory rates
- Broad rate distributions





The mesoscopic scale

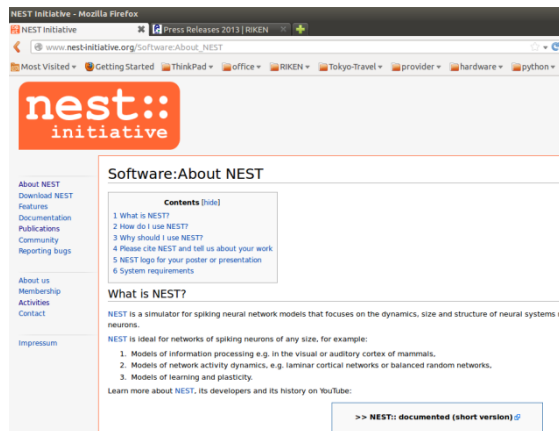
- parallel study
- network of 4mmx4mm covers size of 100 electrode Utah array
- spike correlation structure, LFP
- experimentally, formation and breakdown of waves observed in relation to behavior



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

Simulation Technology: the NEST Initiative

collaborative effort and community building



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* ³⁸

Simulation technology for the brain scale



simulator for large heterogeneous networks of point-neurons and neurons with few compartments

- model of memory usage of NEST
- analysis of memory usage of the 3rd generation simulation kernel (released with NEST 2.2)
- new design of data structures in the 4g kernel
- performance of 4g

www.nest-initiative.org

Model of the memory usage of NEST

- describes the memory usage per MPI process

$$\mathcal{M}(M, T, N, K) = \mathcal{M}_0(M) + \mathcal{M}_n(M, N) \\ + \mathcal{M}_c(M, T, N, K)$$

$$\mathcal{M}_c(M, T, N, K) = TNm_c^0 + TN_c^\emptyset m_c^\emptyset \\ + T(N - N_c^\emptyset)m_c^+ \\ + K_M m_c$$

M total number of MPI processes

T number of threads per MPI process

N total number of neurons

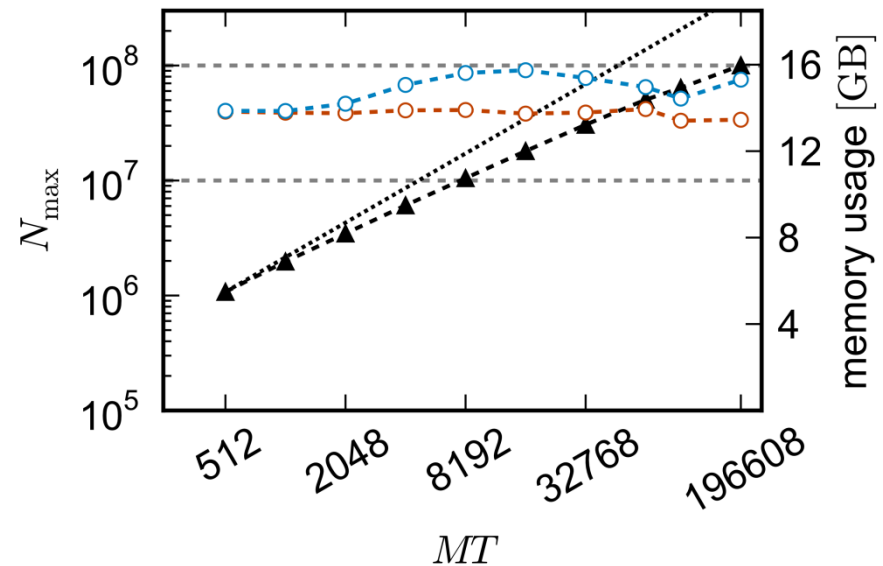
K number of incoming connections per neuron

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

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.

frontiers Journals ▾

Events ▾ Jobs ▾ People ▾ Videos ▾ News ▾ Blogs ▾

About Submit Register Login

Search


frontiers
IN NEUROINFORMATICS

Info

Home

About

Editorial Board

Archive

Research Topics

View Some Authors

Review Guidelines

Subscribe to Alerts

Search

Article Type

All ▾

Publication Date

ORIGINAL RESEARCH ARTICLE

Share 0 Like 0 Comment 0     29   0 **285 Views**

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 Masumoto⁵, Jun Igarashi⁶, Jochen Martin Eppler¹, Shin Ishii⁷, Tomoki Fukai⁶, 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,

Article Info

Abstract

Full Text

PDF

Export Citation

XML

View Article Impact

The Authors in

Frontiers

Google

Google Scholar

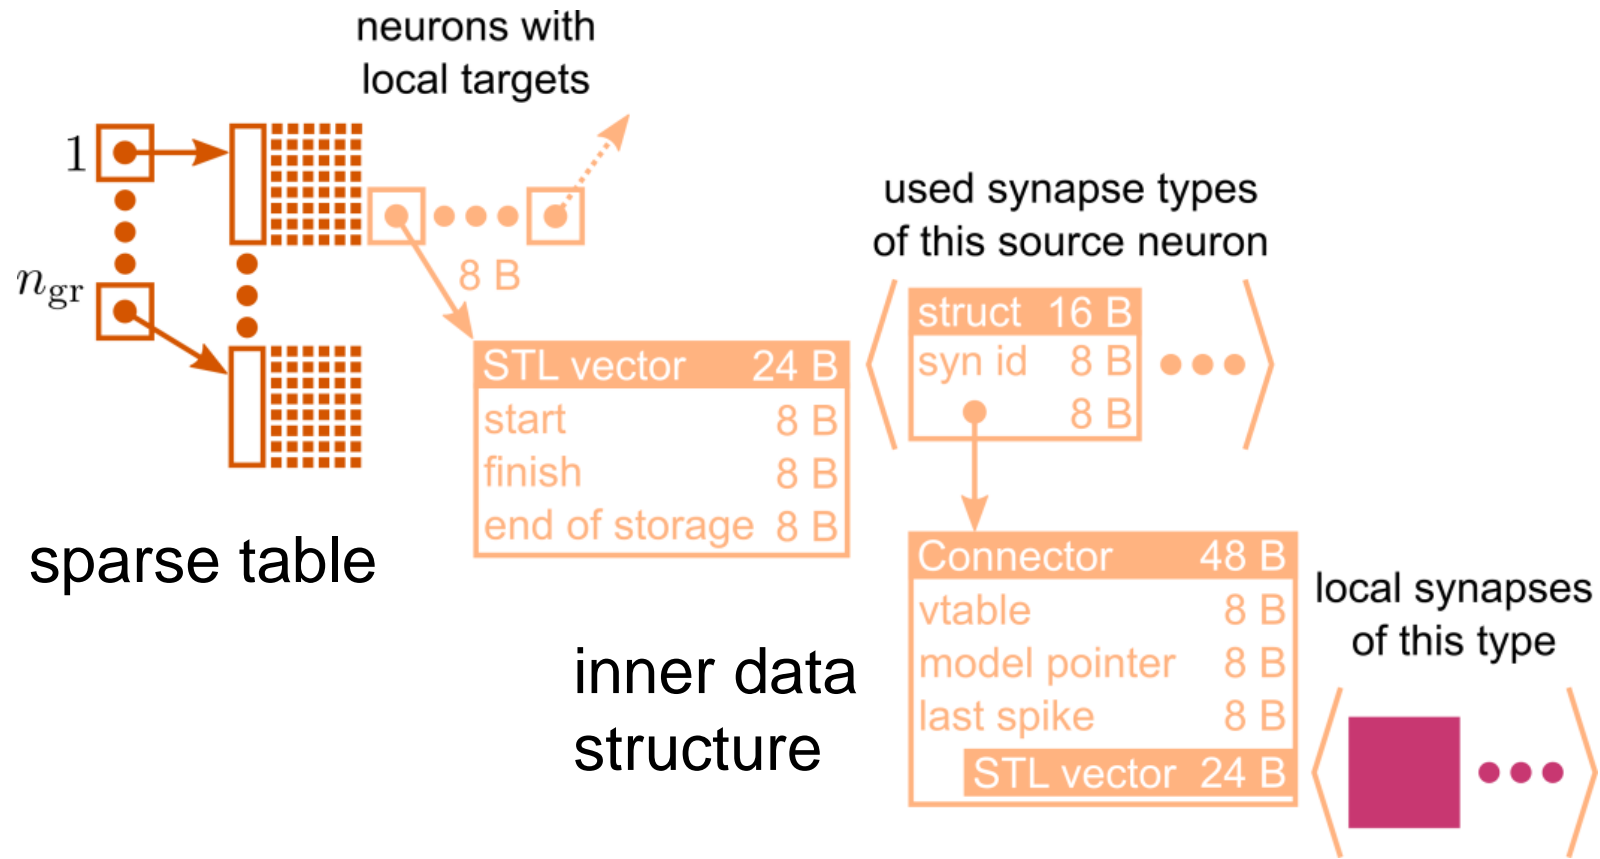
PubMed

Related Article

in Frontiers

Previous connection infrastructure (3g)

- required on each process

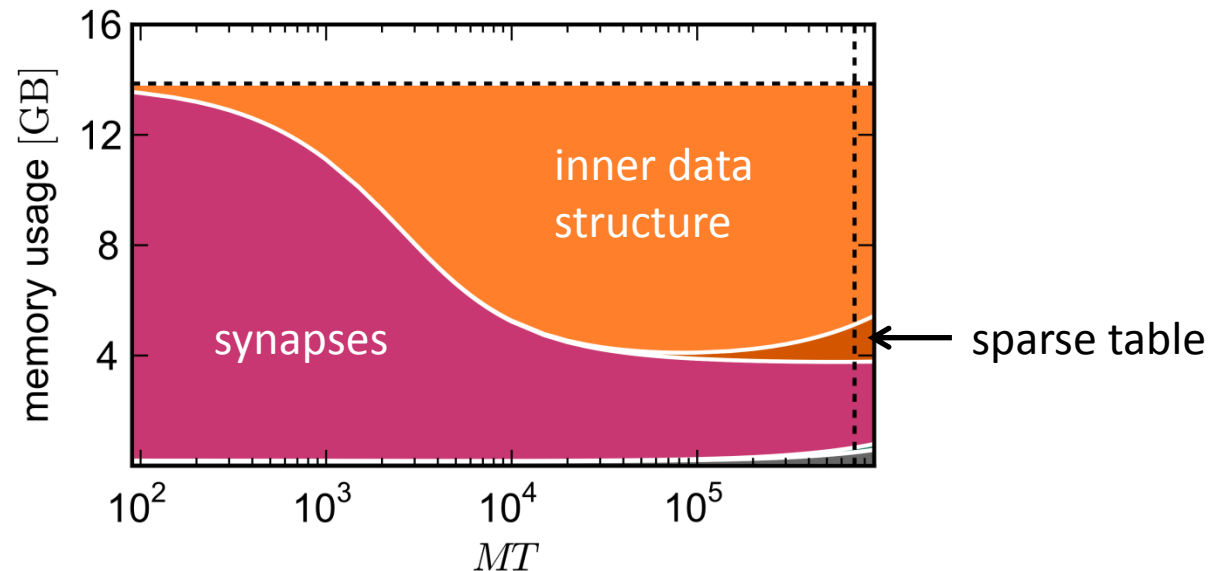


- from 10,000 nodes on collapse along 2 dimensions

3rd generation simulation kernel

analysis of contributions to total memory usage

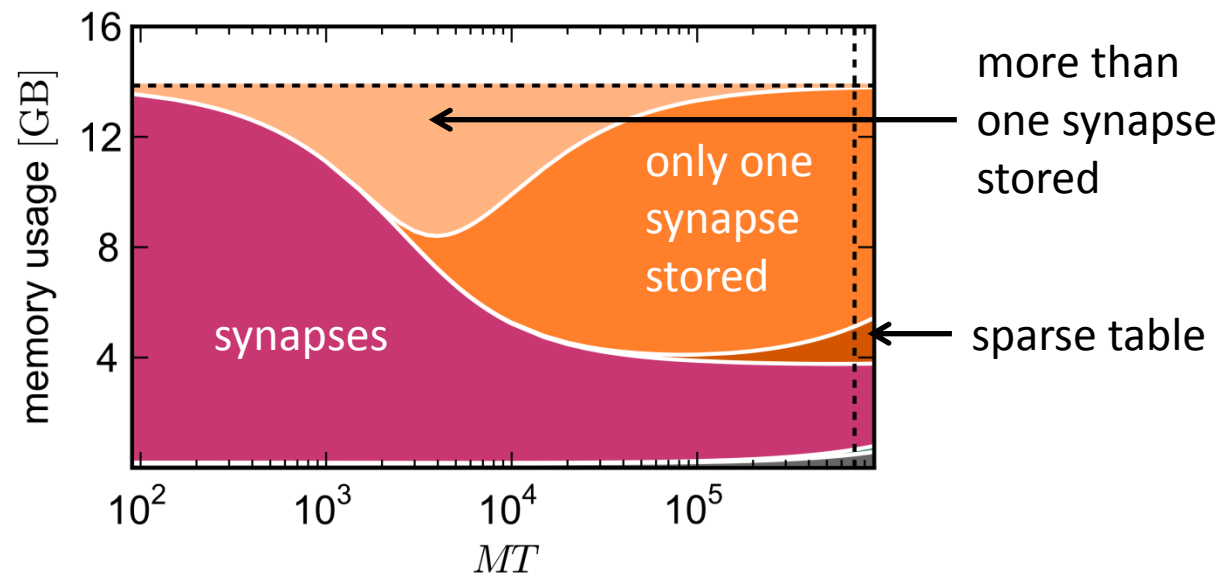
- in the regime of 10k processes and beyond the inner data structure causes severe overhead



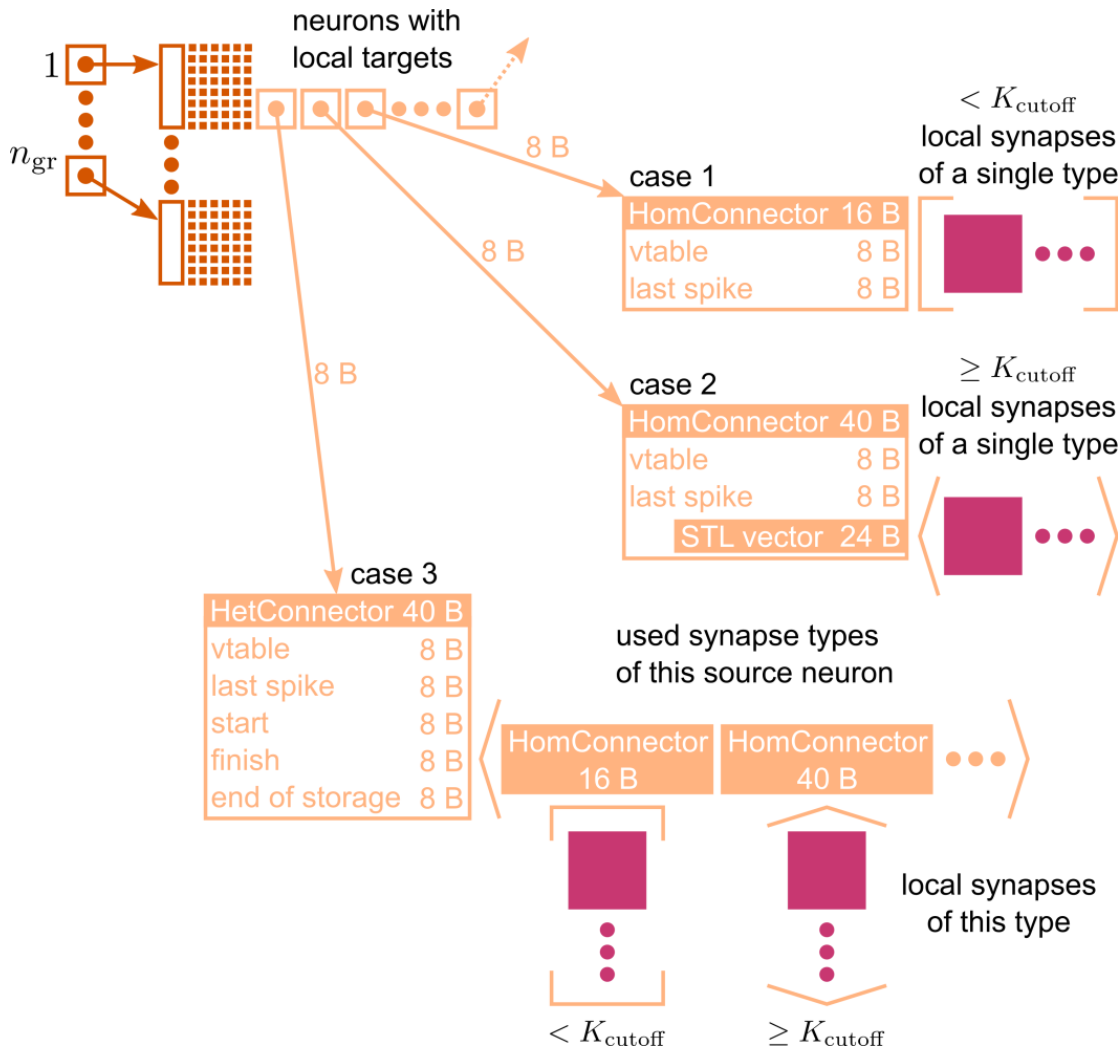
3rd generation simulation kernel

analysis of contributions to total memory usage

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



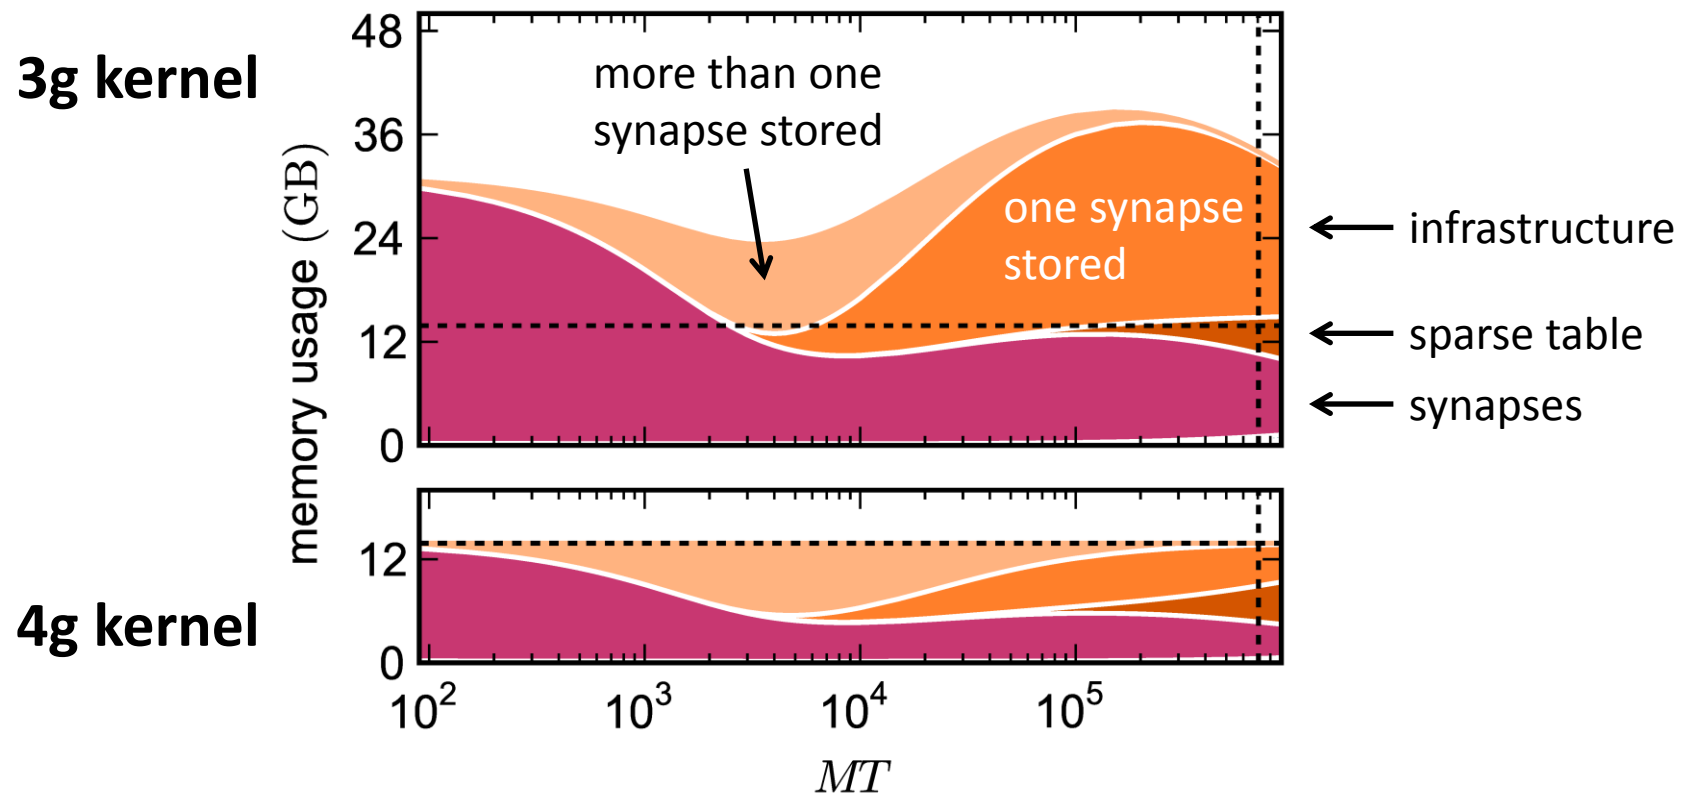
New adaptive connection infrastructure (4g)



low overhead per synapse on supercomputers

full flexibility on laptops and moderately sized clusters

Comparison of 4g to 3g kernel

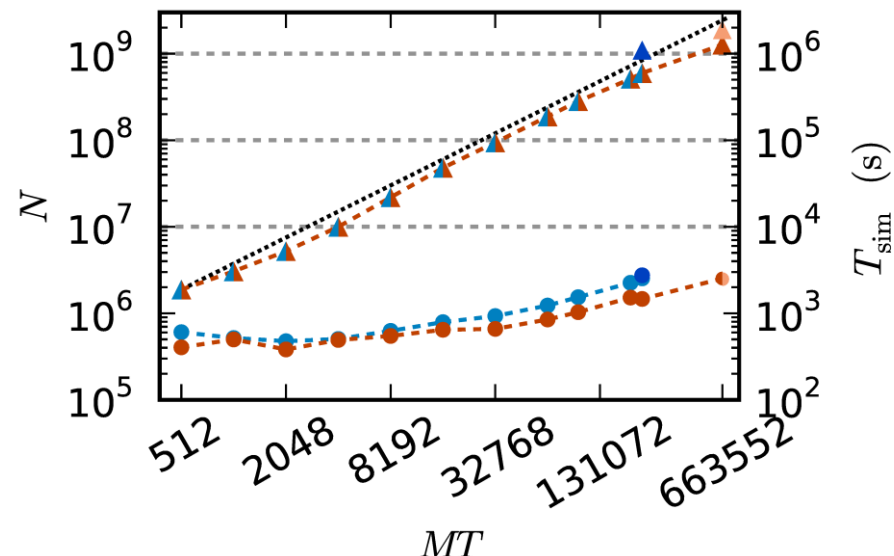


4th generation simulation kernel

- 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 5.73×10^8 neurons on 229,376 cores of JUQUEEN
- up to 1.27×10^9 neurons on 663,552 cores of K
- 11,250 synapses per neuron (exc-exc STDP)

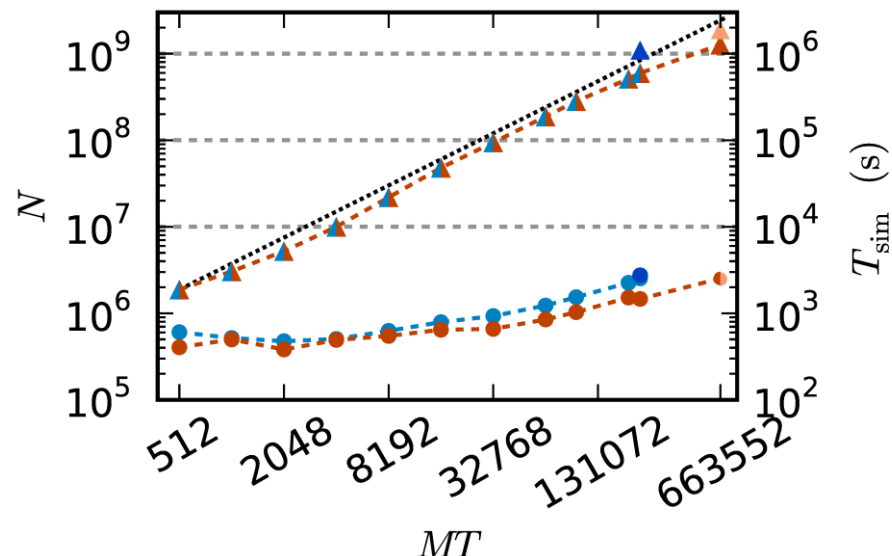


- largest general network simulation performed on K
 1.86×10^9 neurons, 6000 synapses per neuron
(press release July 2013: 1.73×10^9)

- on JUQUEEN 1.08×10^9 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

frontiers in NEUROINFORMATICS

- under review ...

Spiking network simulation code for petascale computers

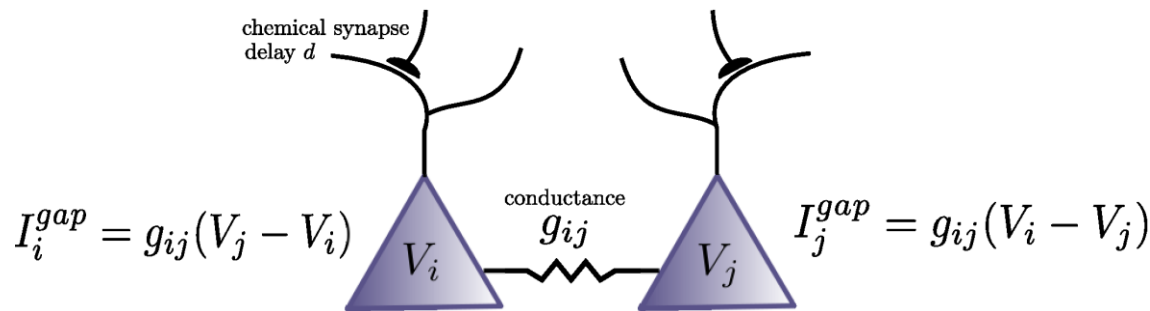
Susanne Kunkel, Maximilian Schmidt, Jochen Martin Eppler, Hans Ekkehard Plesser, Gen Masumoto, Jun Igarashi, Shin Ishii, Tomoki Fukai, Abigail Morrison, Markus Diesmann and Moritz Helias

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

- 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).

Further interactions: gap junctions

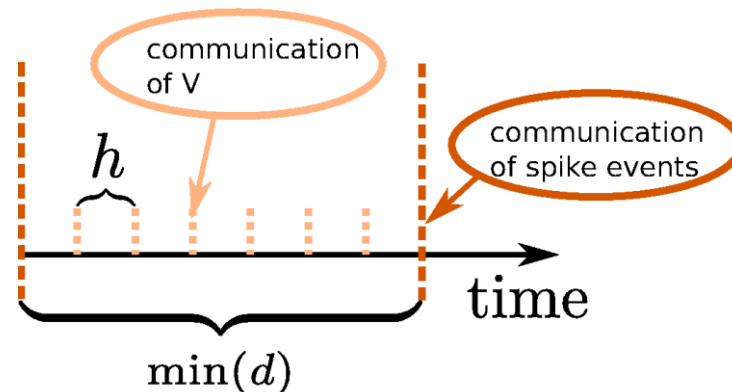


Neuron i (hh_psc_alpha_gap)

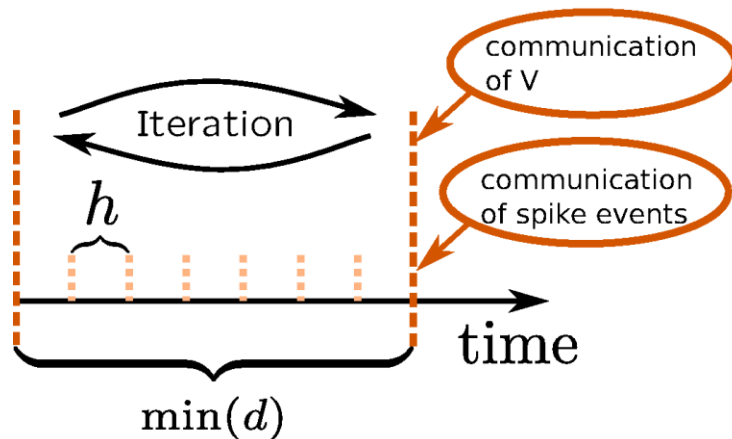
$y'_i(t) = f_i(y_i(t))$, $y_i(t_0)$ given

$$\frac{V'_i}{C_m} = -I_i^{ionic}(V_i, m_i, h_i, n_i, p_i) + I_i^{applied}(I_i^{ex}, I_i^{in}) + I_i^{gap}(V_i, V_j)$$

- at each time point neuron i needs membrane potential of neuron j
- gigantic system of differential equations
- naïve: communication of V in each step



Jacobi waveform relaxation



- communication of approximation only after each $\min(d)$
- iterate until convergence

Neuron i

```
communicate  $V_i^0(t) = V_i(t_0)$ 
```

```
k = 0
```

```
do
```

```
  k = k+1
```

```
  solve the ODE-System
```

```
  using  $V_j^{k-1}(t)$ 
```

```
  compute and communicate
```

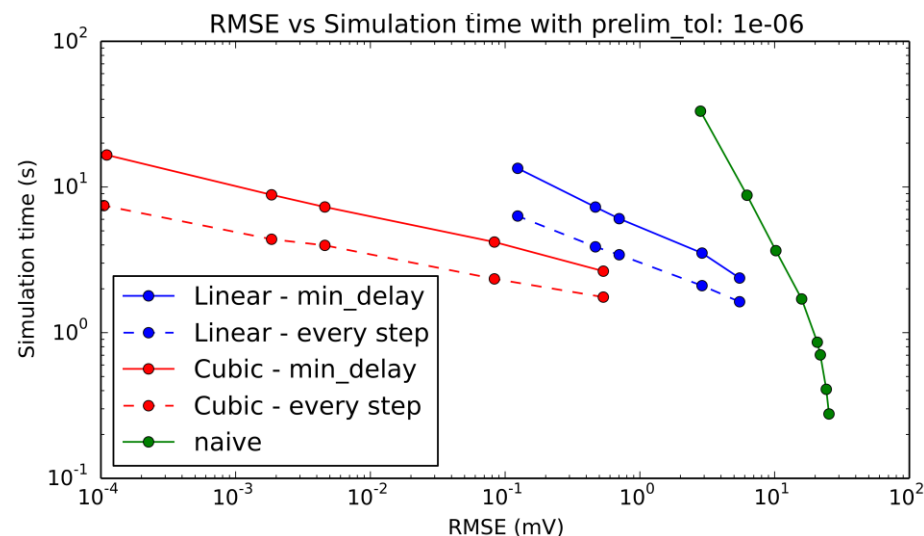
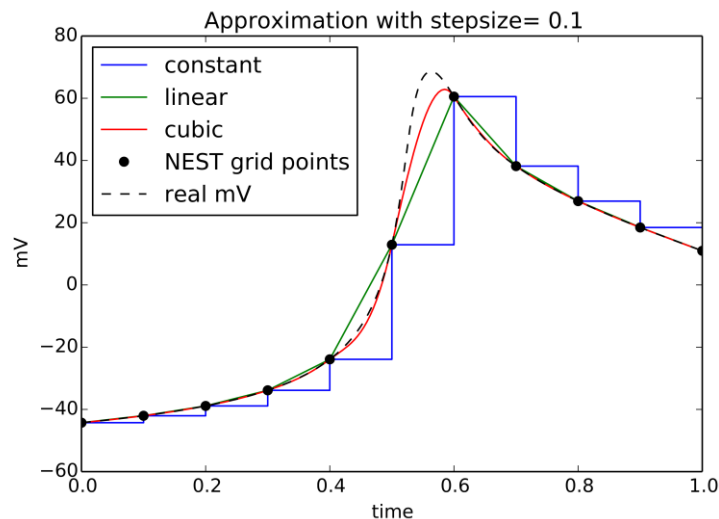
```
  approximation  $V_i^k(t)$ 
```

```
while(  $|V_i^k - V_i^{k-1}| > \text{prelim\_tol}$  )
```

```
continue in time and start over
```

Jan Hahne, Matthias Bolten (University Wuppertal), Jun Igarashi (OIST)

Speed versus accuracy



- pair of neurons with single gap junction
- for practically usable accuracies RMSE < 0.1 mV:
 - waveform relaxation outperforms naïve implementation
 - piecewise-cubic interpolation best
 - very small network: single step faster than min(d) steps
 - reverses for networks larger 1000 neurons

Summary

- full-scale model of local cortical network explains prominent features of neuronal activity
- need for brain-scale models
- simulations require memory only available on supercomputers
- K computer and co-development phase was essential in developing petascale technology in sequence of 2 journal papers
- NEST is simulation engine used by the Human Brain Project (HBP) at the resolution of nerve cells and synapses
- exascale computers present new challenges for data structures
- exascale computers offer new opportunities for simulation speed
- MoU between RIKEN-AICS and Juelich established