JÜLICH
FORSCHUNGSZENTRUM

The NEST 4g kernel:
highly scalable simulation code from laptops to supercomputers

Susanne Kunkel

28/01/2014 Code Jam

SimLab Neuroscience, JSC

Overview

Model of the memory usage of NEST

3rd generation simulation kernel

Analysis of memory usage of 3g kernel

New design of data structures in 4g

Performance of 4g

In collaboration with:

Moritz Helias

Maximilian Schmidt

Jochen M. Eppler

Markus Diesmann

(INM6/IAS6)

Abigail Morrison

(INM6/IAS6 & SimLab)

Hans Ekkehard Plesser

(UMB, Norway)

Jun Igarashi

(OIST, Japan)

Gen Masumoto

(RIKEN, Japan)

Model of the memory usage of NEST

- describes the memory usage per MPI process

$$\mathcal{M}(M,T,N,K) = \mathcal{M}_0(M) + \mathcal{M}_\mathrm{n}(M,N) + \mathcal{M}_\mathrm{c}(M,T,N,K)$$

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

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 et al. (2012) Meeting the memory challenges of brain-scale network simulation. *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 M=196,608 compute nodes and T=8 threads per node
- 8 GB of memory per node

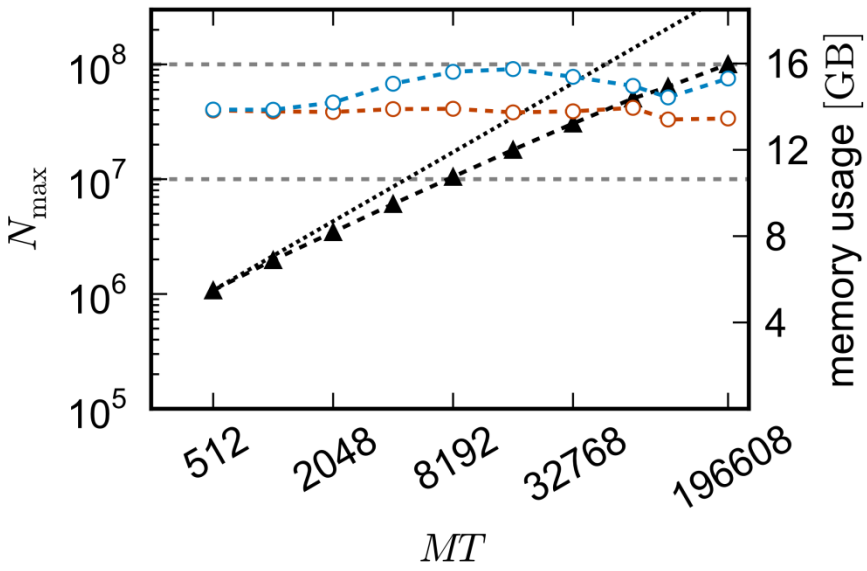

Helias et al. (2012) Supercomputers ready for use as discovery machines for neuroscience. *Front. Neuroinform.* **6**:26.

Previous connection infrastructure (3g)

- required on each process

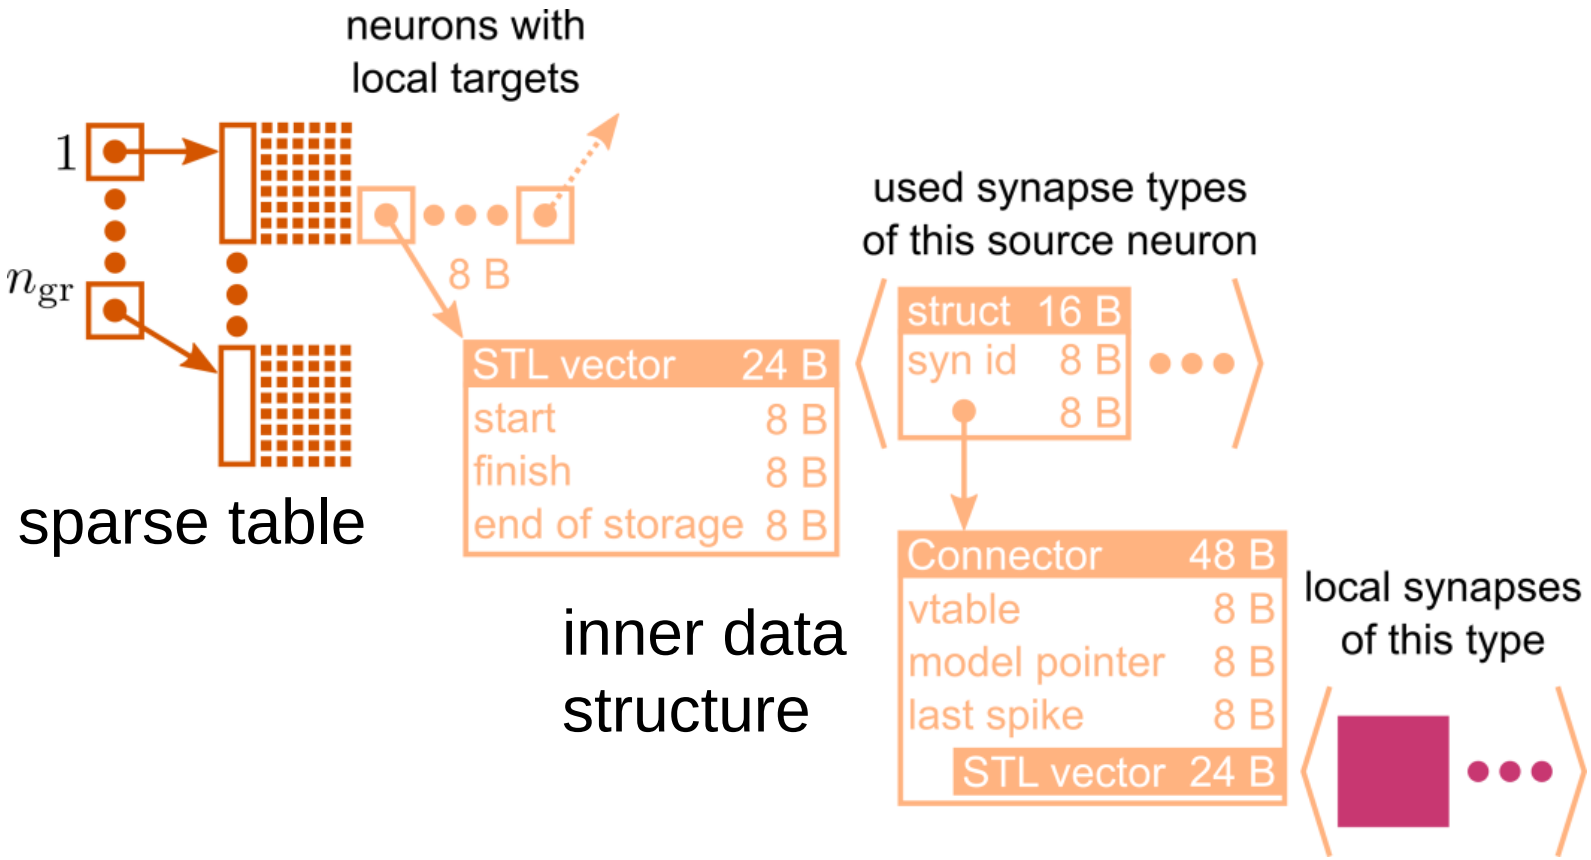

3rd generation simulation kernel

Analysis of contributions to total memory usage

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

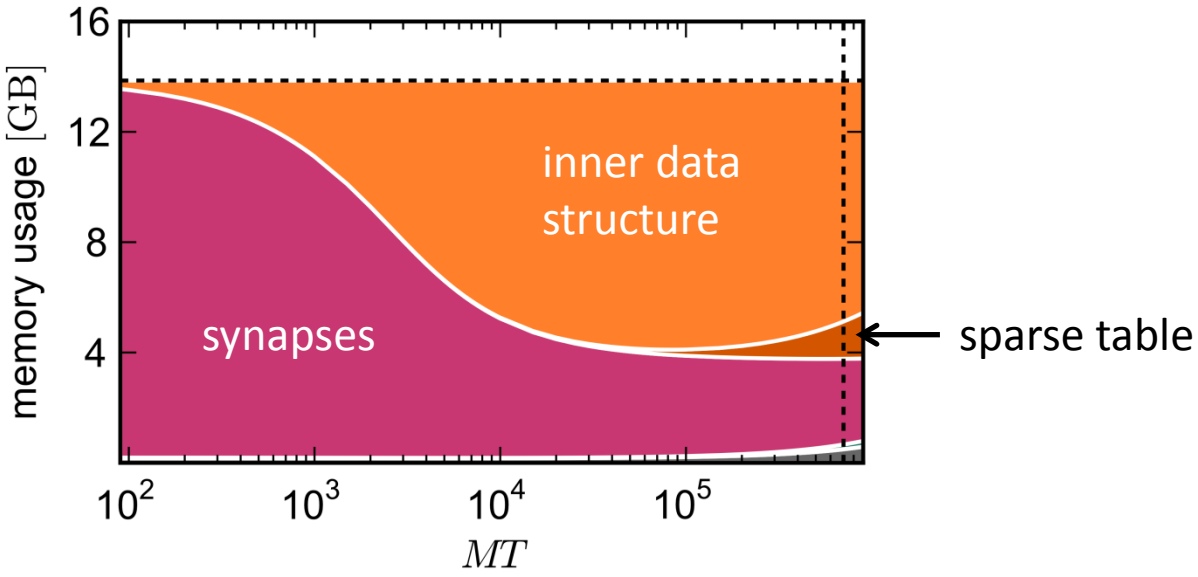

Previous connection infrastructure (3g)

- on supercomputers either no local targets or only few

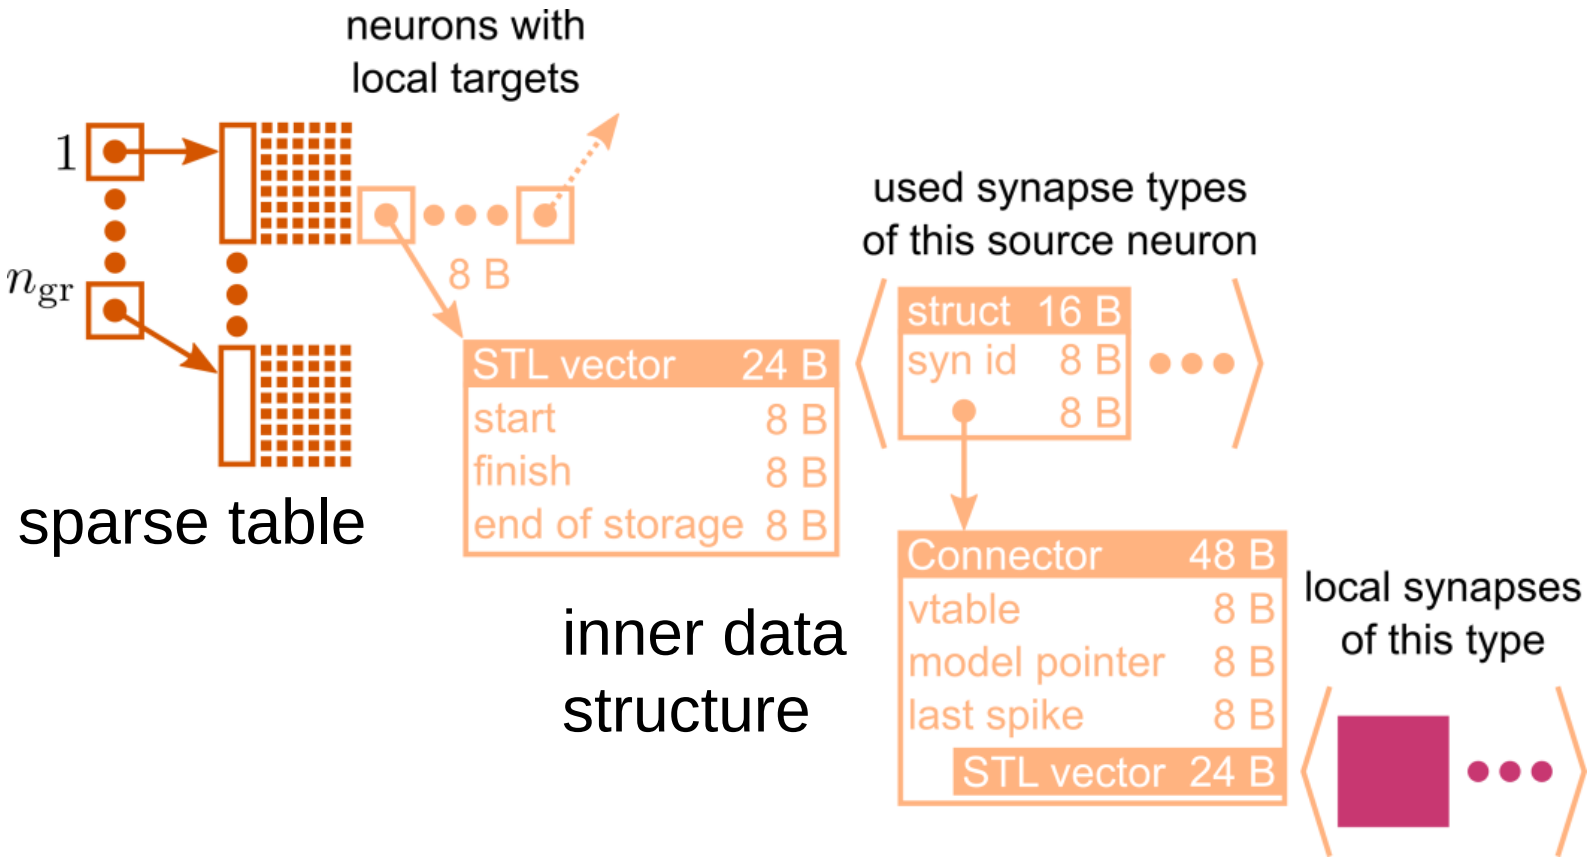

3rd generation simulation kernel

Analysis of contributions to total memory usage

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

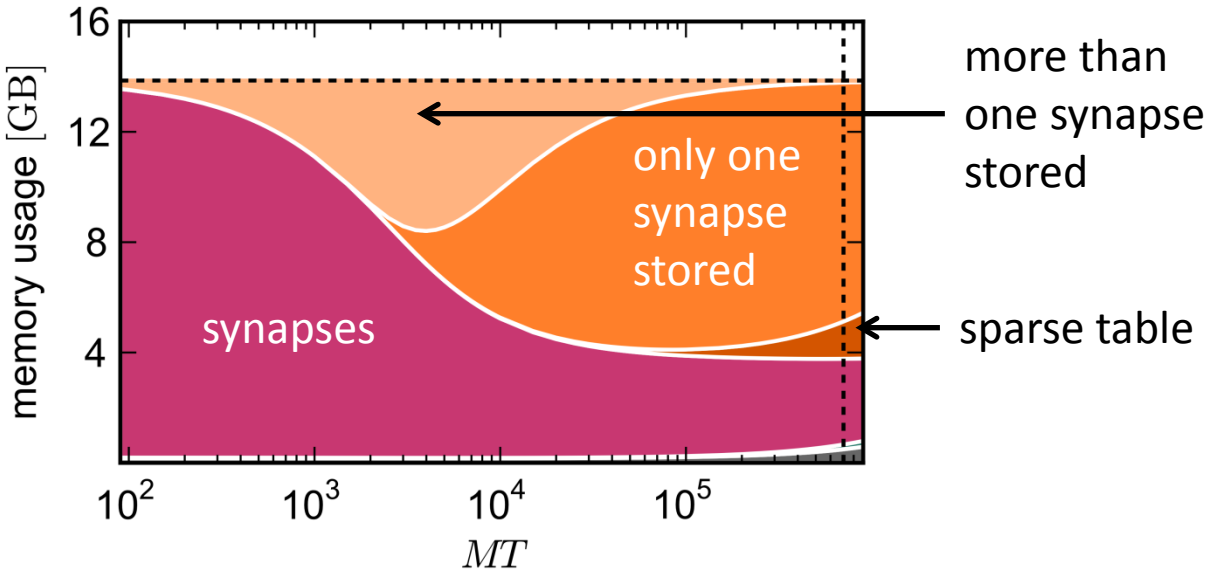

New adaptive connection infrastructure (4g)

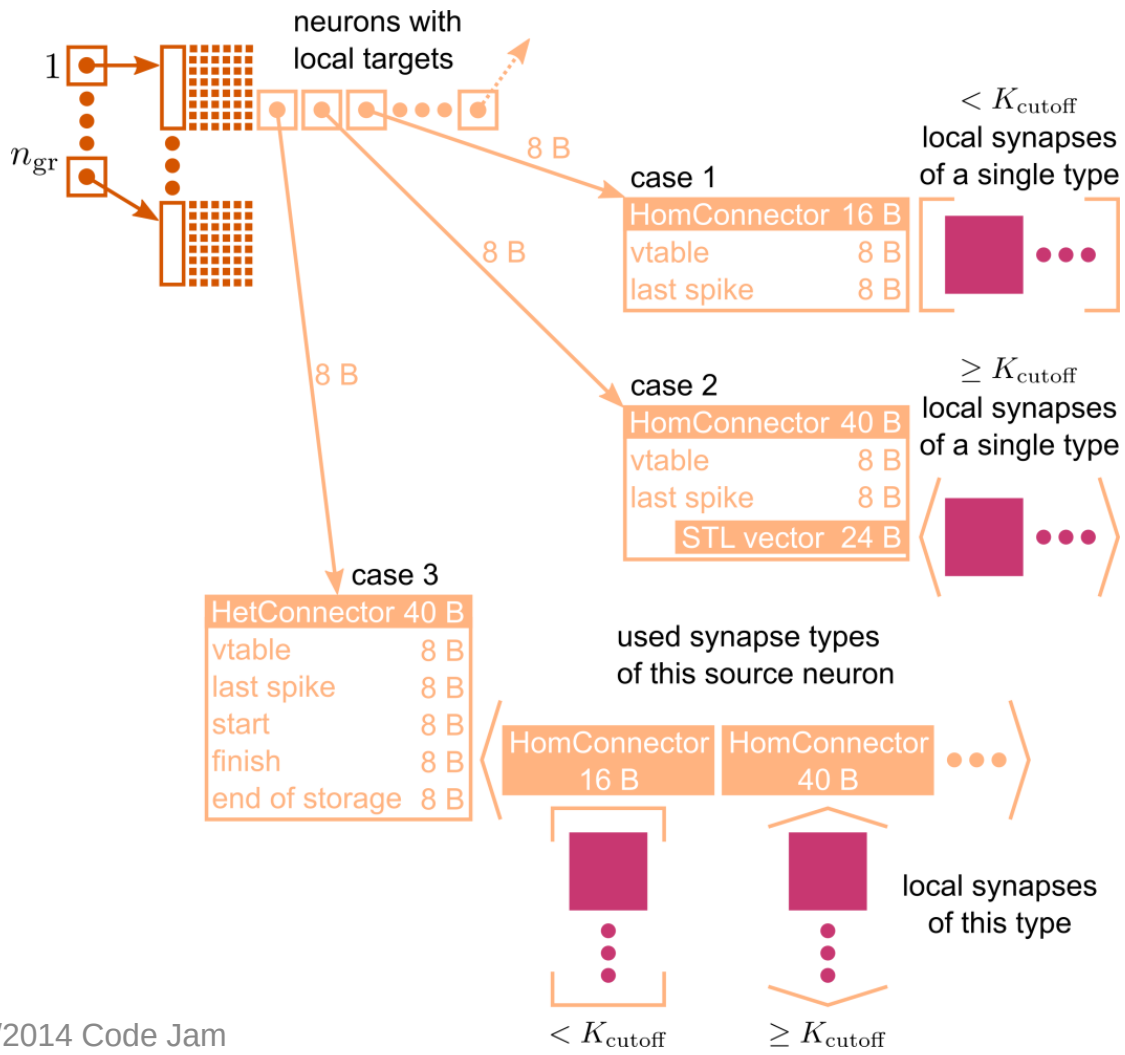

low overhead per synapse on supercomputers

full flexibility on laptops and moderately sized clusters

New adaptive connection infrastructure (4g)

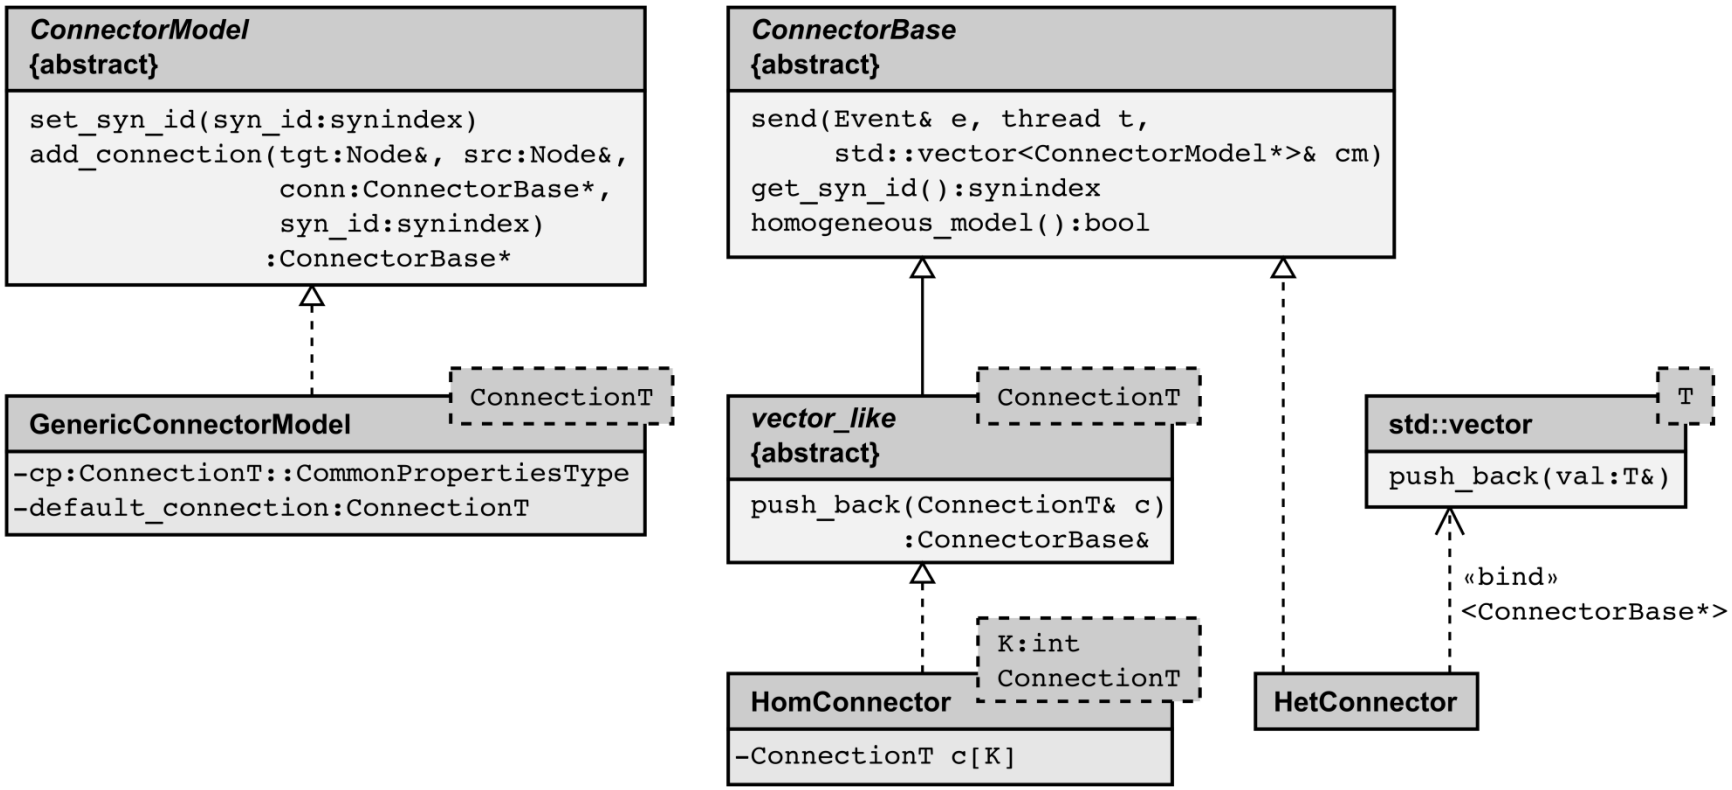

4th generation simulation kernel

- new adaptive connection infrastructure which causes only low overhead in the case of short target lists
- reduced memory usage of synapses
  - e.g. removed vtable pointer
  - no compromise on precision of synaptic state variables
- SparseNodeArray replaces sparse table in neuronal infrastructure
  - memory overhead only per local neuron
  - exploits round-robin distribution of neurons
  - enables fast access to neuron model info during setup

Kunkel et al. Spiking network simulation code for petascale computers (in prep)

Comparison of 4g to 3g kernel

- simulation of same network using 3g and 4g kernel

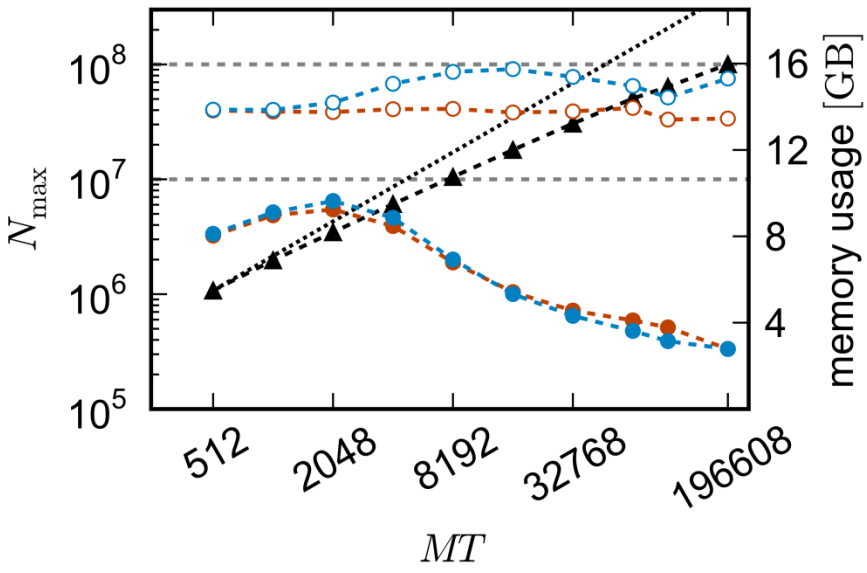

Reduced memory usage

- in all regimes of number of processes
- especially in the regime of 10k processes and beyond

(less than 1/3 at 100,000 cores)

Comparison of 4g to 3g kernel

Reduced setup time

- optimization of wiring routines
- faster memory allocation using dedicated pool allocator

Reduced simulation time

- smaller objects in connection infrastructure enable more efficient use of cache

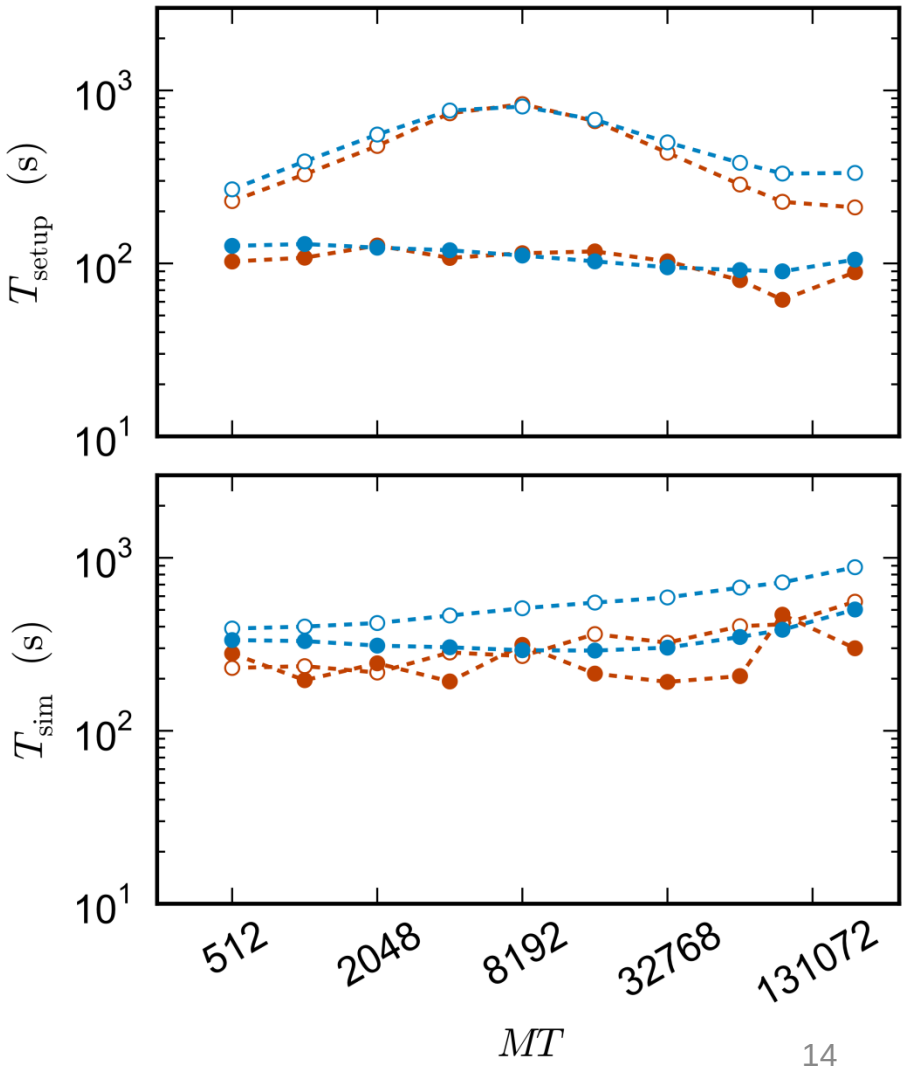

Maximum network size

- 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)

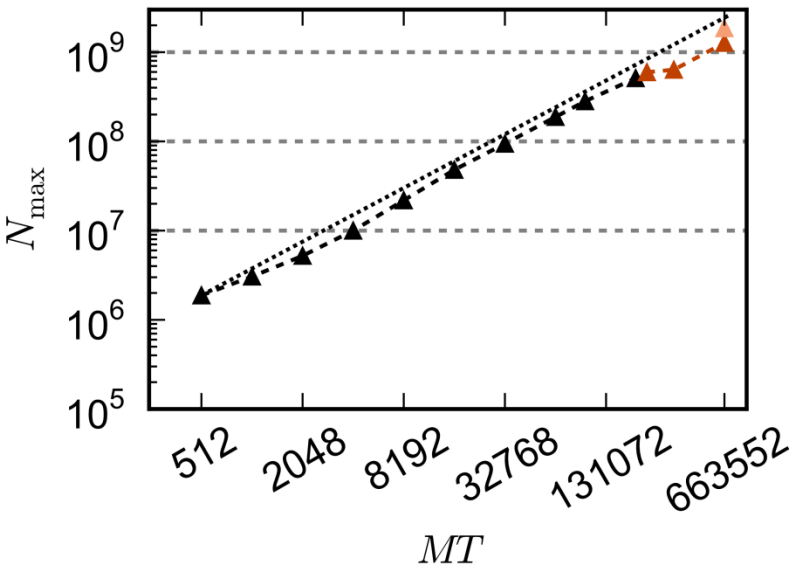

- largest general network simulation performed on K in July 2013

  (1.73x10^9 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 took between 3 and 15 min

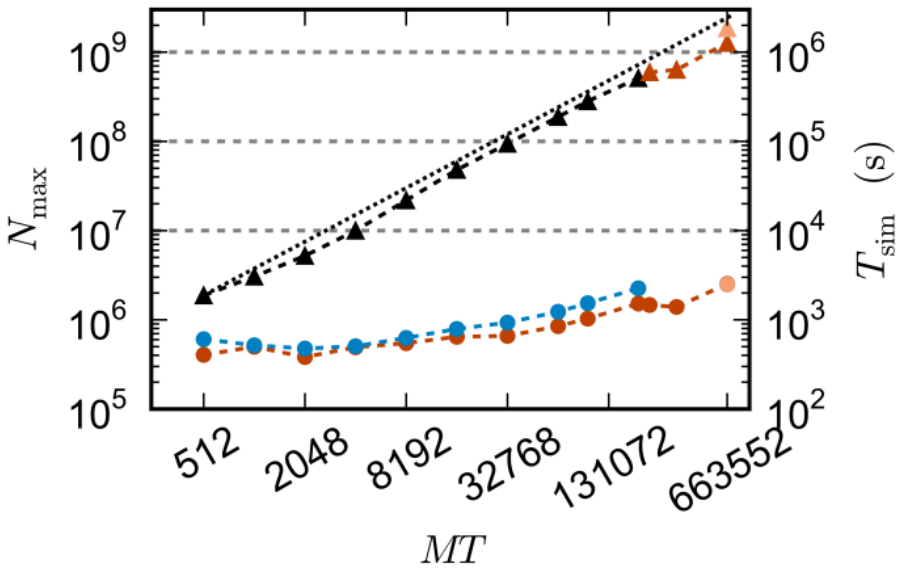

- largest general network simulation performed on K in July 2013

  (1.73x10^9 neurons, 6000 synapses per neuron) → 56 min in total

Power consumption

- JUQUEEN requires per rack max. 100 kW
- full K computer requires max. 9.89 MW

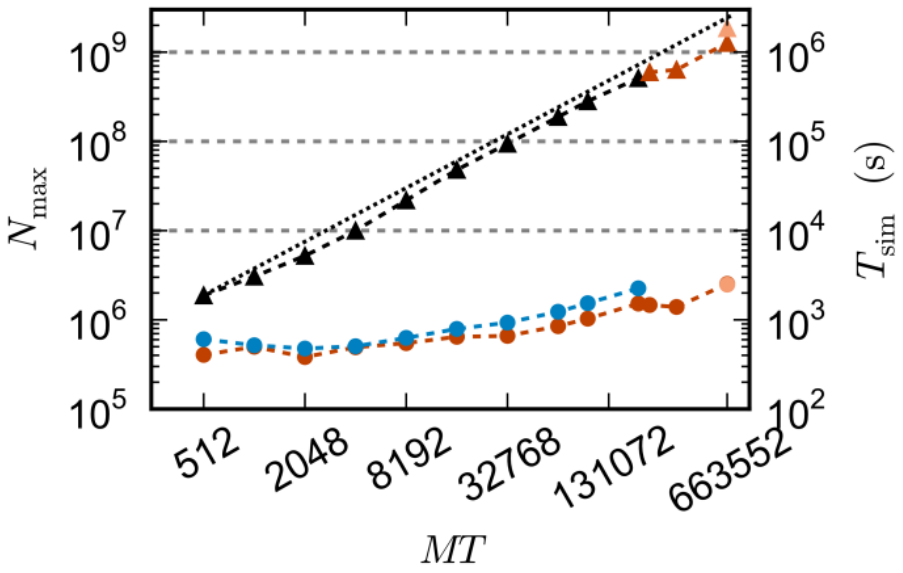

Power consumption

- simulations on JUQUEEN consumed 2714 kWh
- simulations on K consumed 25,408 kWh

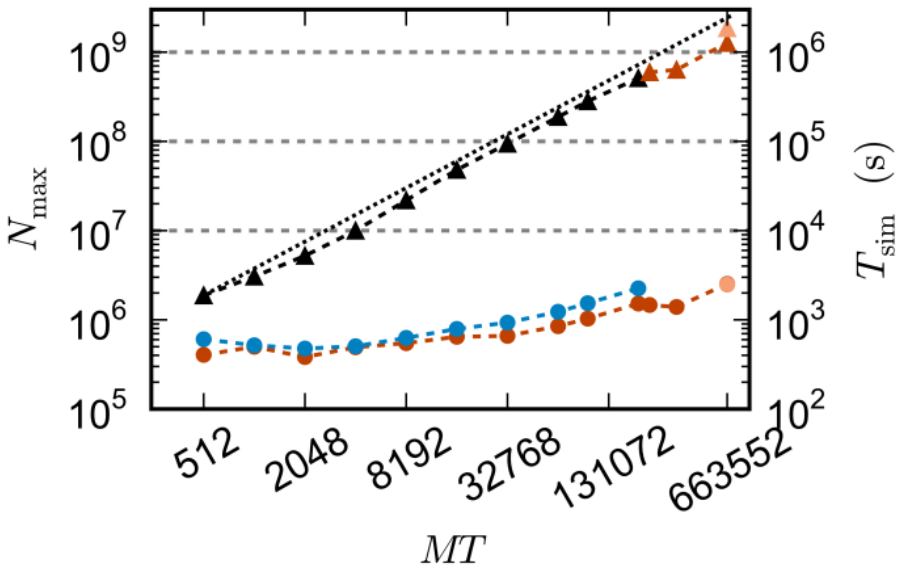

- largest general network simulation performed on K in July 2013

  (1.73x10^9 neurons, 6000 synapses per neuron) → 9253 kWh

Power consumption

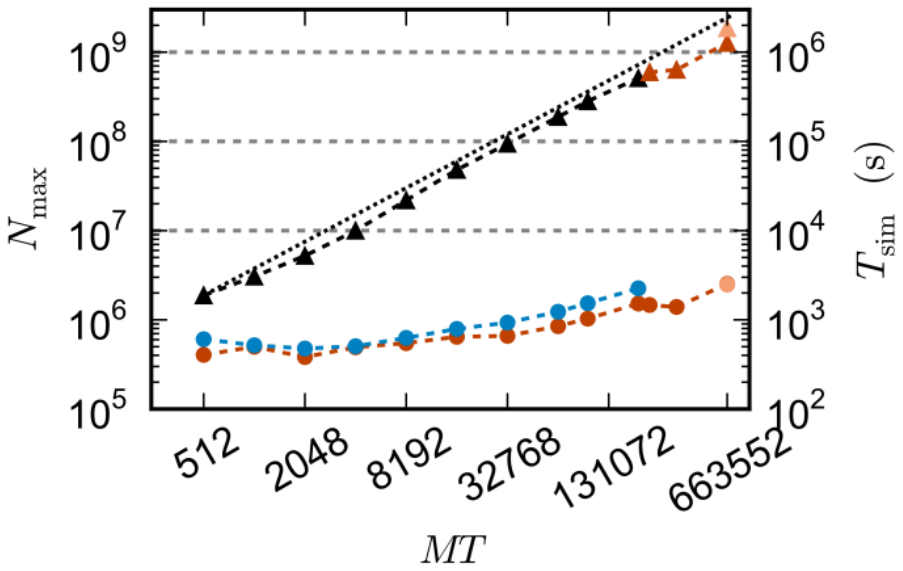

Power consumption

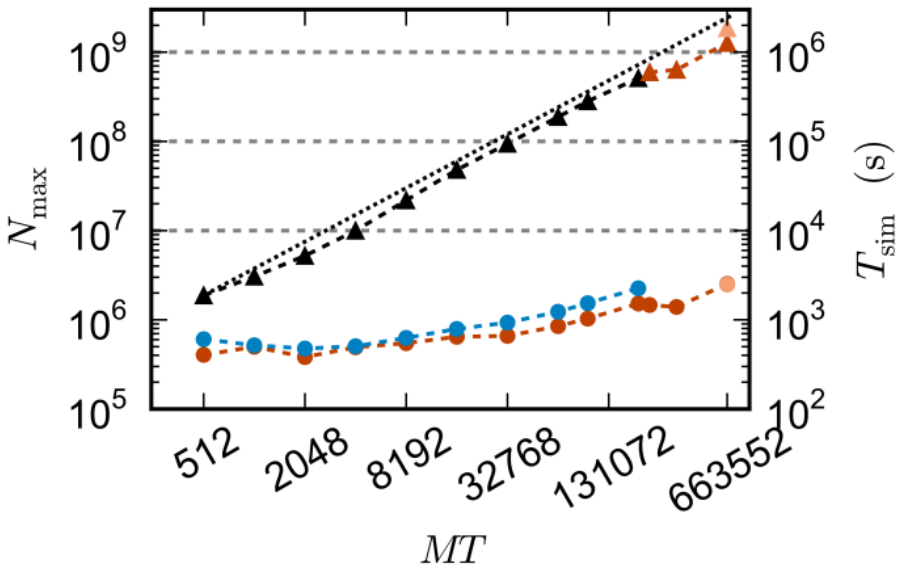

- largest general network simulation performed on K in July 2013 (1.73x10^9 neurons, 6000 synapses per neuron) → 14 hours

Comparison of 4g to 3g kernel

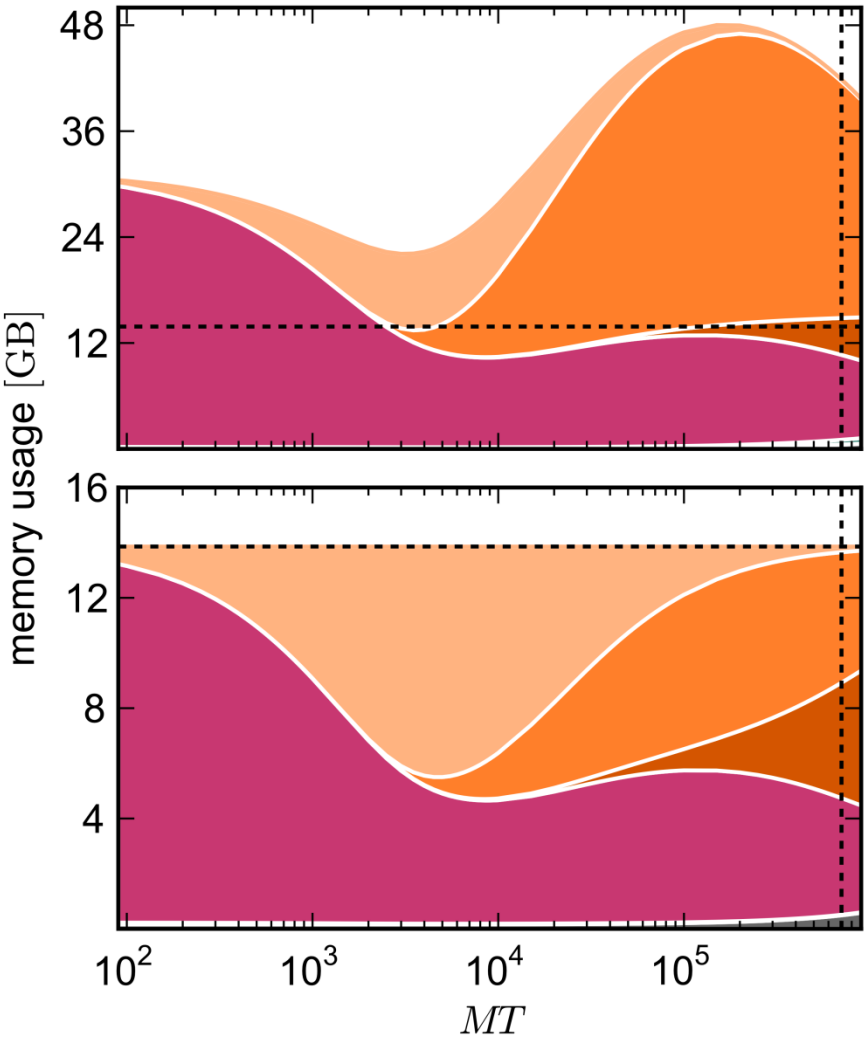

3rd generation simulation kernel

Setup times

- between 4.5 and 13.5 min on JUQUEEN
- between 3.5 and 13.9 min on the K computer

Simulation times (for 1s)

- between 6.5 and 14.7 min on JUQUEEN
- between 3.6 and 9.3 min on the K computer