On Scaling Neuronal Network Simulations Using Distributed Computing

Vladimir Omelyusik, Khawar Shehzad, Tyler Banks,
Praveen Rao, Satish S. Nair

University of Missouri, Columbia

Nov 18, 2024

Motivation | Current challenges in neuroscience research

Our final target is construction of end-to-end workflows for interactive collaboration, including allowing users to modify simulations in real time.

Challenges

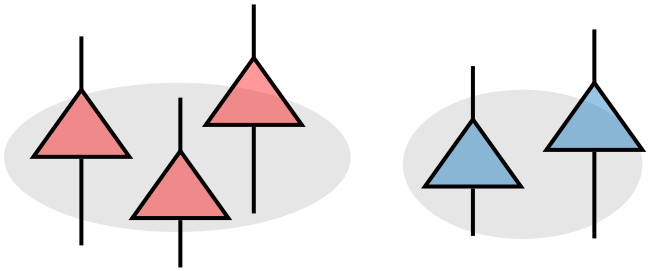

Large-scale simulations of biologically realistic networks

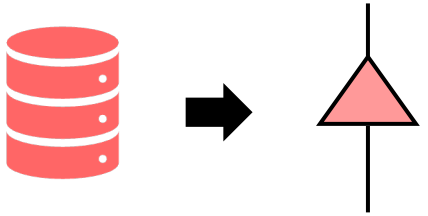

Efficient access to neural databases (e.g., Allen DB)

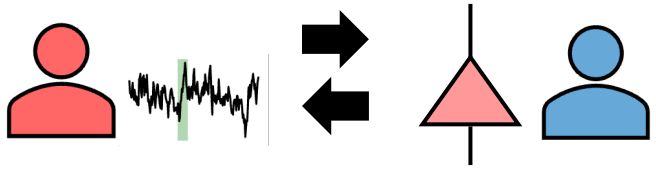

Real-time collaboration and simulation control

Approach

Hardware characterization (single-site)

Compared a single FABRIC site with desktop and server

Hardware characterization (multi-site)

Quantified overheads in a multi-site setup

Software characterization (CoreNEURON)

Tested the effects of connectivity and distribution strategy

A guide for neuroscience users

Instructions are available at *github.com/raopr/neuroscience-on-FABRIC*

Methods | FABRIC

Created in 2019, FABRIC is a nationwide research testbed with high-speed connectivity aiming to push the boundaries of distributed, stateful, 'everywhere' programmable infrastructure.

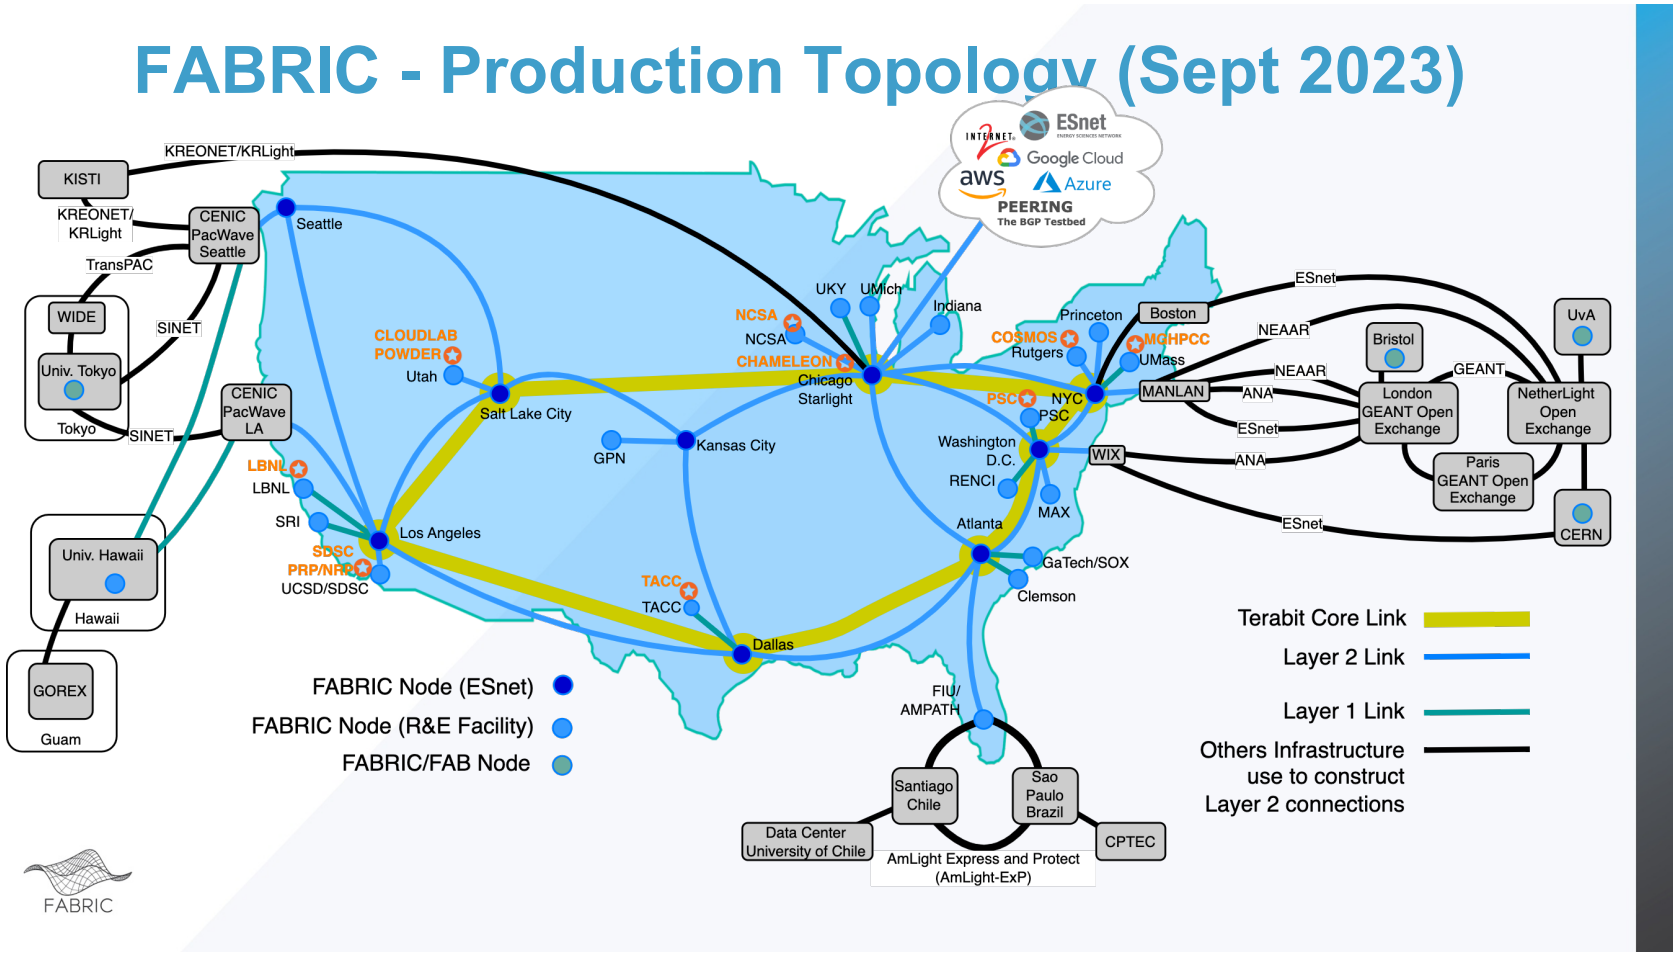

33 sites across the US and Europe

2×32 AMD Rome and Milan with 512 Gb RAM

GPU NVIDIA RTX 6000, T4, A30

1 Tb NVMe drives in servers and 250 Tb rotating storage at each site

(FABIRC, 2023)

Methods | Networks and cell assignment

We utilized different variants of the biologically realistic Pyramidal Interneuron Network – Gamma (PING; Borgers, 2017) consisting of clusters of densely connected cells (assemblies).

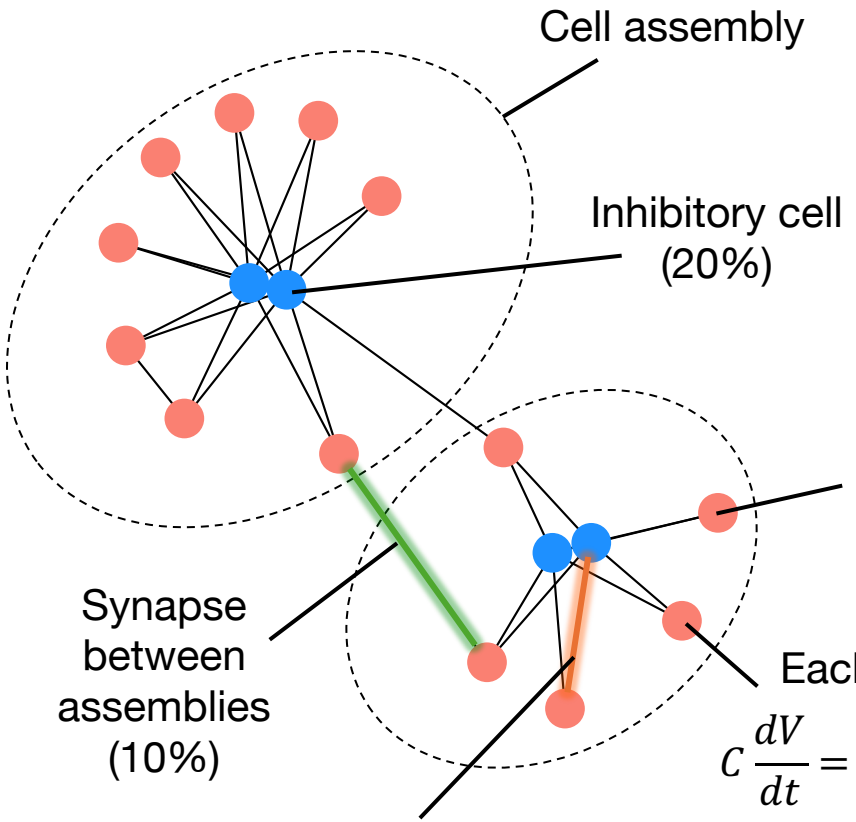

Each cell is modeled as

$$C\frac{dV}{dt} = \sum I_{ion} + \sum I_{syn} - I_{inj}$$

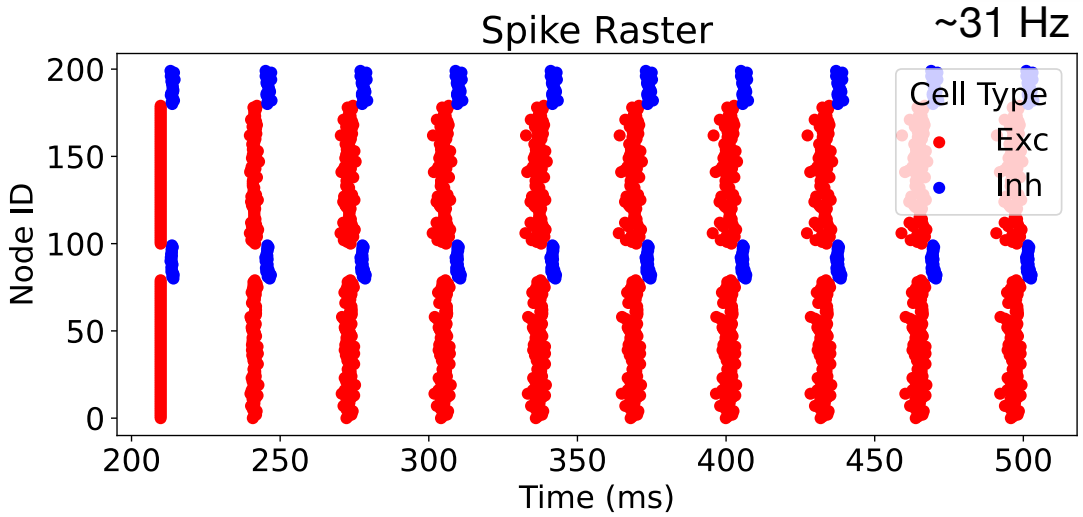

Distribution strategy

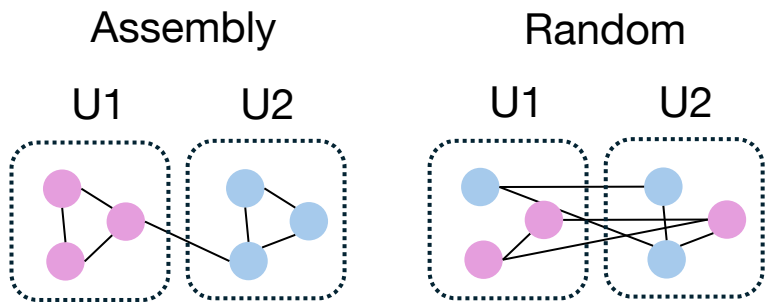

Results | Single-site hardware characterization

FABRIC is a viable alternative to conventional simulation setups, on average being 36% faster in CPU runs and 500% faster in GPU runs.

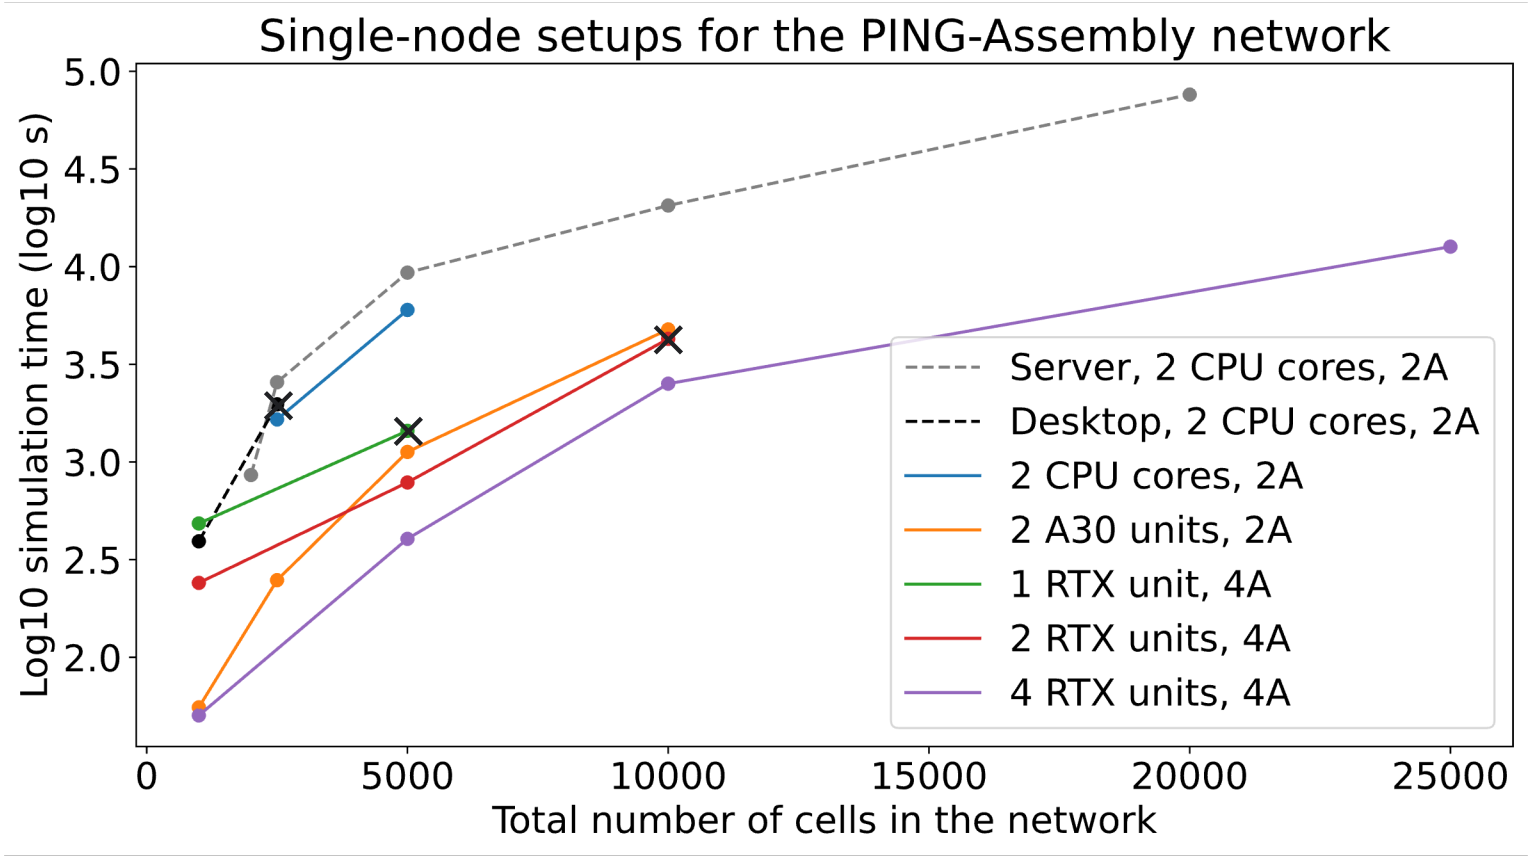

Results | Multi-site hardware characterization

The “natural” delay estimate appeared to be size-dependent and decreasing with the number of cells, with the mean value of around 6.5 times.

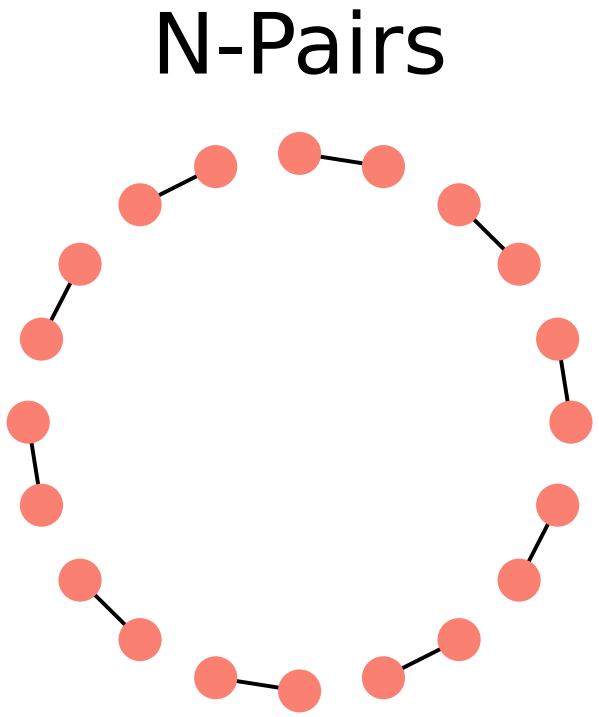

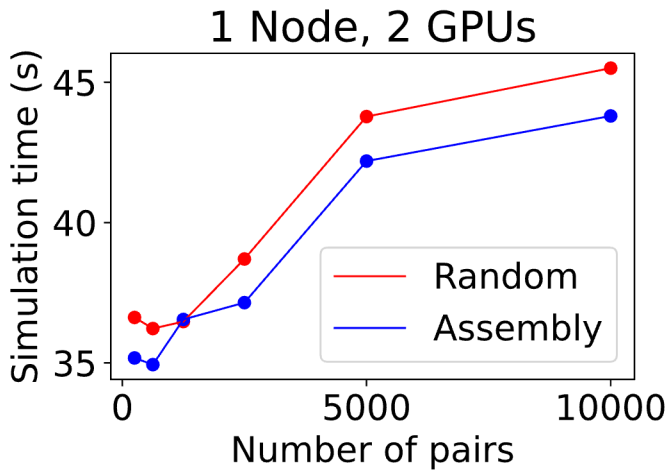

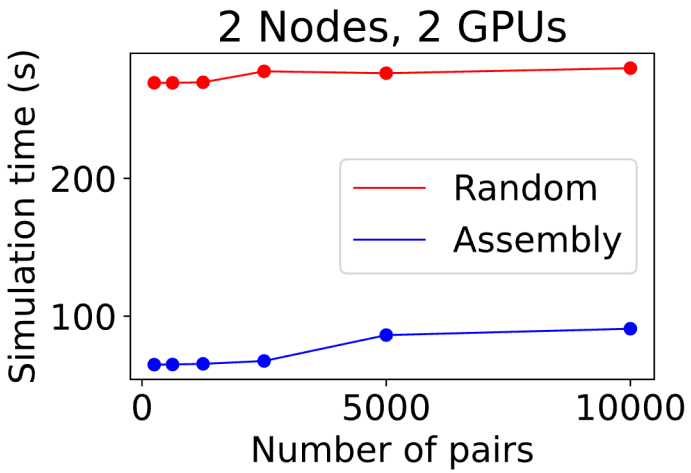

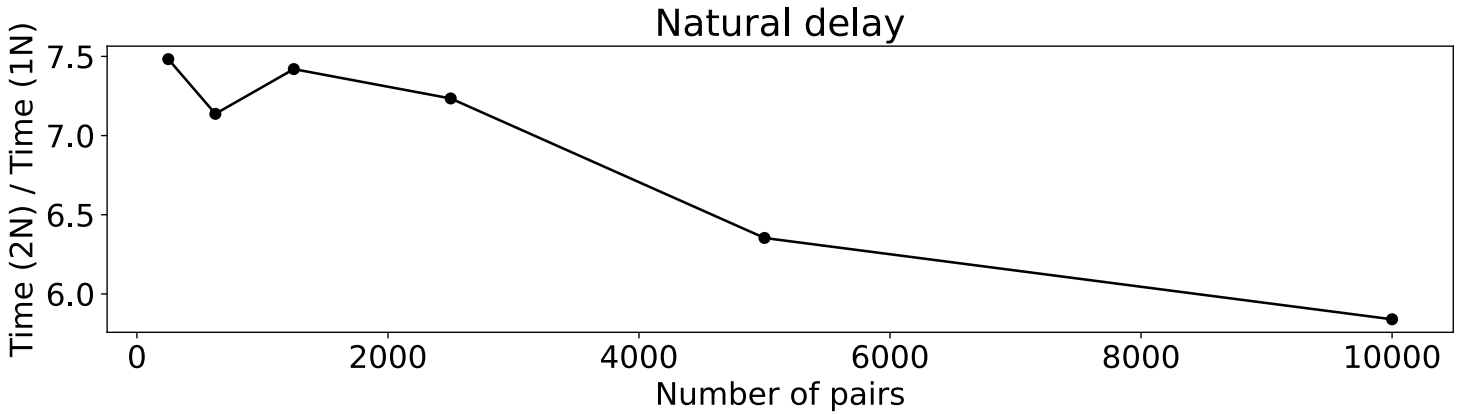

Results | Impact of synaptic connectivity

Under the same connectivity, simulating a high number of smaller assemblies is faster than simulating a small number of bigger assemblies. The number of synaptic connections is one of the crucial factors that affects the simulation time.

N-K-Cliques

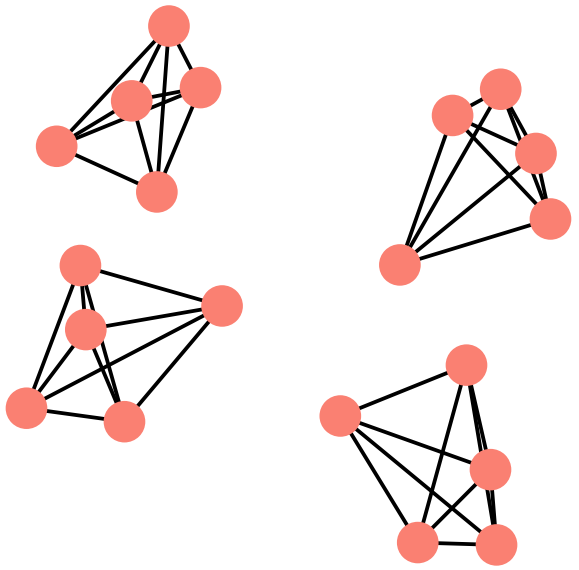

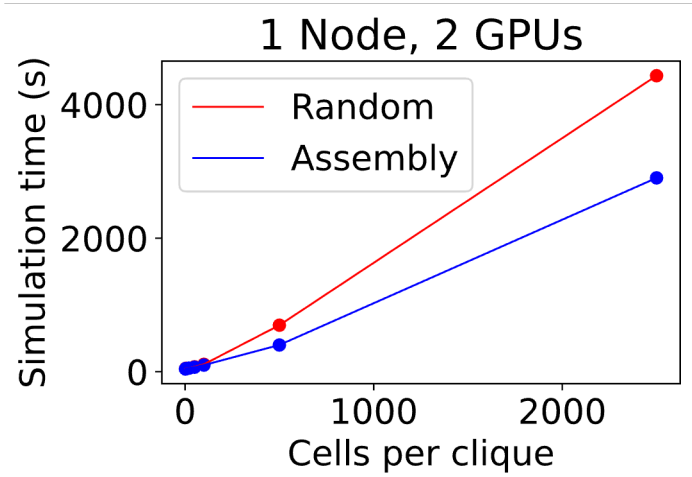

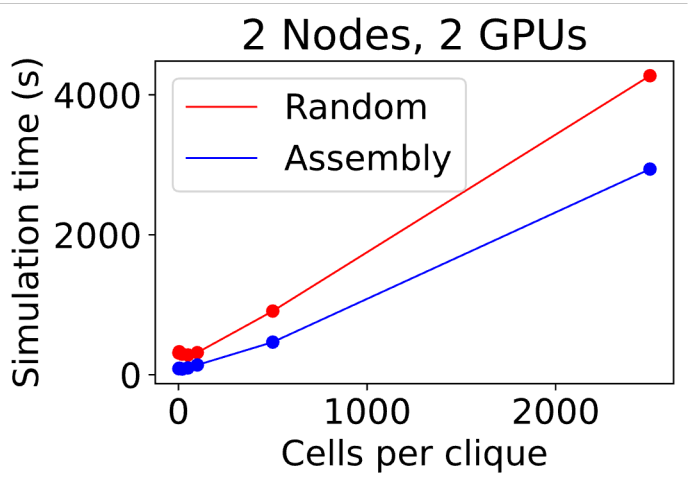

Results | Distribution with biologically realistic connectivity

CoreNEURON's memory optimization algorithm can efficiently handle biologically realistic cases. Others can be addressed with a min k-cut graph partitioning algorithm.

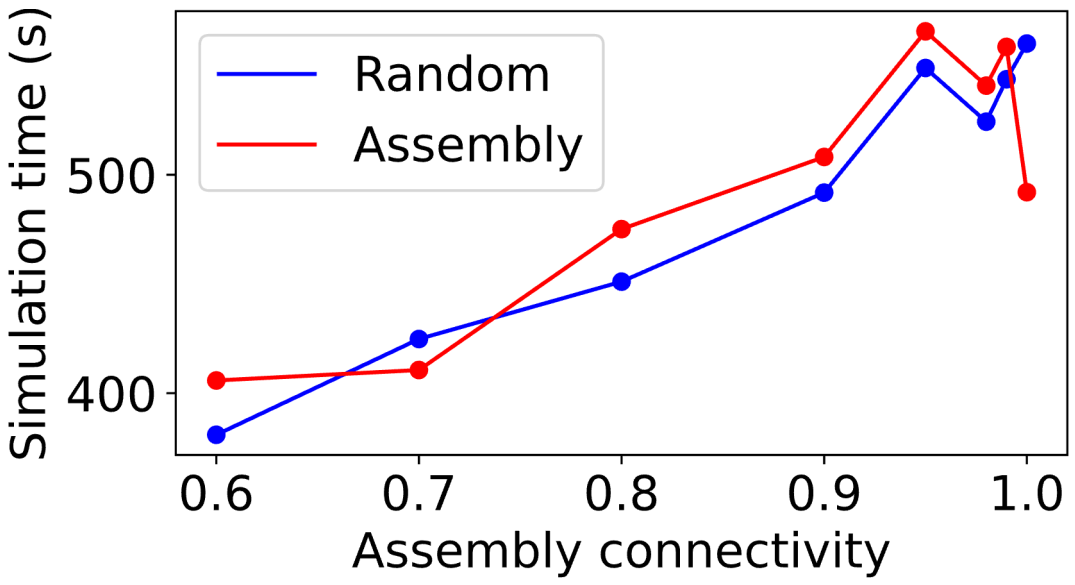

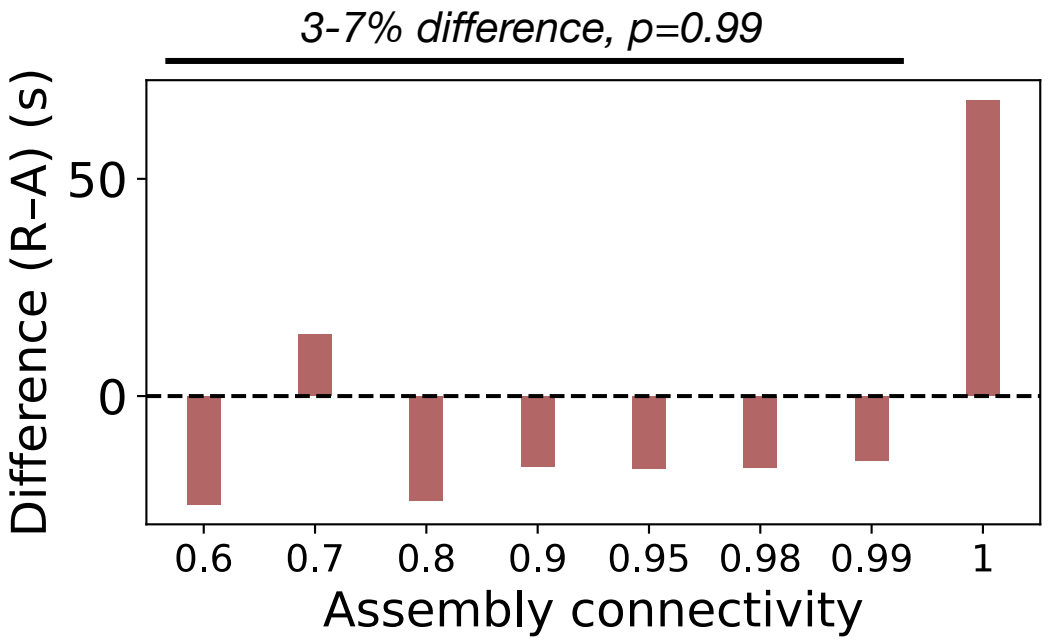

Conclusion

FABRIC facilitates large-scale simulations of biologically realistic networks

- Availability of GPUs allows **reduction** of simulation time for **larger** networks.
- While multi-site configurations are characterized by inherent delays, its effect **decreases** as the network size grows.

CoreNEURON simulator efficiently handles biologically realistic cases regardless of distribution strategy

- Biologically realistic cases can utilize simple **random** assignment.
- By-assembly distribution is practically better only with fully connected assemblies. These cases can be handled with a **min k-cut graph** partitioning algorithm.

Future directions: advancing neuroscience research with FABRIC

- Online simulation control of experimental setups and integration of database searches during ongoing simulations.
- Setup instructions are available at *github.com/raopr/neuroscience-on-FABRIC*

We thank the CoreNEURON and FABRIC teams for their assistance. Praveen Rao would like to acknowledge the partial support of NSF Grant No. 2201583