Simulating large-scale models of brain neuronal circuits using Google Cloud Platform

Subhashini Sivagnanam^{1,2}
sivagnanam@sdsc.edu

Wyatt Gorman³
wyattgorman@google.com

Donald Doherty¹
donald.doherty@actionpotential.com

Samuel A Neymotin⁵
Samuel.Neymotin@nki.rfmh.org

Stephen Fang³
sfang@google.com

Hermine Hovhannisyan³
hermine@google.com

William W Lytton^{1,4}
bill.lytton@downstate.edu

Salvador Dura-Bernal^{1,5}
Salvador.Bernal@downstate.edu

1. State University of New York Downstate, Brooklyn NY 2. San Diego Supercomputer Center / University California San Diego, La Jolla CA, 3. Google Corporation, Mountain View CA 4. King's County Hospital, Brooklyn NY 5. Nathan Kline Institute for Psychiatric Research, Orangeburg NY USA

ABSTRACT

Biophysically detailed modeling provides an unmatched method to integrate data from many disparate experimental studies, and manipulate and explore with high precision the resulting brain circuit simulation. We developed a detailed model of the brain motor cortex circuits, simulating over 10,000 biophysically detailed neurons and 30 million synaptic connections. Optimization and evaluation of the cortical model parameters and responses was achieved via parameter exploration using grid search parameter sweeps and evolutionary algorithms. This involves running tens of thousands of simulations, with each simulated second of the full circuit model requiring approximately 50 cores hours. This paper describes our experience in setting up and using Google Compute Platform (GCP) with Slurm to run these large-scale simulations. We describe the best practices and solutions to the issues that arose during the process, and present preliminary results from running simulations on GCP.

CCS CONCEPTS

• Computing methodologies ~ Modeling and simulation ~ Simulation types and techniques ~ Massively parallel and high-performance simulations • Applied computing ~ Life and medical sciences ~ Computational biology ~ Biological networks

KEYWORDS

Brain modeling, computational neuroscience, large-scale simulations, Google Cloud Platform

ACM Reference format:

Subhashini Sivagnanam, Wyatt Gorman, Donald Doherty, Hermine Hovhannisyan, Stephen Fang, William W Lytton, Salvador Dura-Bernal. 2020. Simulating large-scale models of brain neuronal circuits using Google Cloud Platform. In *Proceedings of Practice and Experience in Advanced Research Computing 2020 (PEARC'20)*. Portland OR, USA, 6 pages. <https://doi.org/10.1145/1234567890>

1 Introduction

The cerebral cortex is the outermost layer of the brain and is responsible for most high-level functions like vision, language and reasoning, each implemented by circuits in different “bespoke” regions. Advances in brain imaging and recording technology have yielded vast amounts of data on the structure and dynamics of cortical circuits [1][2][3][4]. These data span multiple scales, each providing important pieces of the puzzle: from molecular processes in ionic channels and synapses [5] to electrical properties along the dendrites of neurons to dynamics and oscillations of large networks. Biophysically detailed modeling provides an unmatched method to integrate data from many disparate experimental studies, and manipulate and explore with high precision the resulting brain circuit simulation [6][7][8]. Some examples of brain regions simulated using detailed computational multiscale models include motor cortex (M1), visual (V1) [9], somatosensory (S1), and auditory (A1) cortices, as well as other cortical areas such as hippocampus [8] and cerebellar cortex [10]. We have developed a highly detailed model of M1, simulating a volume of cortical tissue with over 10,000 biophysically detailed neurons and 30 million synaptic connections [11], including 15 neural populations with cell density, distribution and synaptic connectivity patterns derived from experimental data. The model has been used to investigate circuit information pathways, oscillatory coding mechanisms and the role of noradrenaline in modulating output to spinal cord and movement initiation.

Running such simulations and data analysis requires significant computational resources, such as the High Performance Computing (HPC) resources that are available through the National Science Foundation’s Extreme Science and Engineering Discovery Environment (XSEDE) program [12], or cloud computing resources provided by commercial cloud providers such as Google Cloud Platform (GCP) [13]. This is because calculating the electrical currents within each neuron -- including synaptic inputs received from thousands of other neurons --

requires solving thousands of differential equations every fraction of a millisecond: simulating one second of the full M1 cortical circuit model requires approximately 50 HPC core hours. Distributing cells across compute nodes parallelizes each network simulation, and many of these simulations are typically executed in parallel to explore different parameter combinations. Finding model parameters that reproduce biology and investigating neural coding via in silico experiments involve exploring large parameter spaces [14][15]. Optimization and evaluation of the cortical model parameters and responses is achieved via parameter exploration methods including grid search parameter sweeps and evolutionary algorithms [14–18]. This involves running tens of thousands of simulations, with each simulated second of the full circuit model, requiring approximately 1M core hours for 20,000 instantiations.

We have previously used evolutionary algorithm optimization to train a cortical model to control a virtual arm using XSEDE supercomputers[14]. We were recently awarded the NSF/Internet2 Exploring Clouds for Acceleration of Science (E-CAS) grant [19] to continue our research by simulating large-scale models of brain circuits using GCP. This paper describes our experience in setting up and using GCP with Slurm to run these large-scale simulations, best practices and solutions to the issues that arose during the process, and preliminary results from running simulations on GCP.

2 Simulation Software

The simulations require NEURON [20] and NetPyNE [21] to be installed in the computational resource. Below is a brief description of the software and their role in the simulation.

NEURON: NEURON was used as the backend simulator to run both the single cell and network simulations. NEURON is the leading simulator in the neural multiscale-modeling domain. For our simulations NEURON was compiled to work with MPI, using either MVAPICH2 or OpenMPI (depending on the resource). We also employed NMODL, the language used to implement custom synaptic and membrane mechanisms in NEURON; this language is compiled to C thus increasing simulation performance.

NetPyNE: The NetPyNE python package facilitates the design, parallel simulation, optimization analysis and visualization of data-driven neuronal networks in NEURON. The complexity of these models requires that we use this high-level declarative language to organize parameters across scales: molecular, cellular, network. The NetPyNE tool also allows users without programming expertise to build sophisticated brain models using either the declarative format or the graphical user interface. NetPyNE supports two parameter optimization methods: grid search parameter sweep and evolutionary algorithms (using the Inspyred package [22]). In both cases, NetPyNE generates the job scheduler (Slurm or PBS Torque) scripts corresponding to a simulation with a specific set of parameters and submits the required jobs.

3 Setting up the Cloud Environment

3.1 Requirements

Slurm [23] is an open source, fault-tolerant, and highly scalable cluster management and job scheduling system for HPC resources. Integration of Slurm with GCP compute resources is integral for our project computing needs for the following reasons:

Model characteristics: One mode of running the model requires launching multiple jobs for high throughput computing (HTC), e.g., 1000 jobs in parallel, where each job is an MPI-based HPC job requiring multiple cores. As described in section 1, our study involves performing parameter optimizations, which requires significant HPC resources. For example, grid search parameter optimization or exploration involves searching exhaustively through a specified subset of parameters, evaluating all combinations of parameter values. This requires submitting thousands of jobs simultaneously, one for each parameter combination. Similarly, evolutionary algorithm optimization requires running hundreds or thousands of generations, each evaluating many individuals, where each individual corresponds to a simulation with a parameter combination. Slurm is well suited to handle this task by enabling the user to submit multiple jobs at once to the queue. Slurm then assigns computing resources as they become available.

Elasticity/Walltime: While the XSEDE resources were free of cost to use, there were limitations, whose specifics varied by system. In general, we were restricted in walltime (hours that a job is allowed to be resident on the systems), number of running jobs at a given time per user, and total number of cores available. GCP supports elastic computing and can auto-scale according to job requirements, and has no specific walltime, number of jobs, or core number restrictions. Using Slurm with GCP enables us to submit multiple such jobs in a batch environment.

Ease of Use: We had used Slurm for managing our jobs on XSEDE resources, and were familiar with the scripts for usage and management. Additionally, NetPyNE already had built-in support for automatically generating Slurm scripts and submitting Slurm jobs. Setting up GCP with Slurm makes it easier to be used as a central lab resource to run simulations and share results.

Figure 1: GCP-Slurm setup (reproduced from [24])

3.2 GCP-Slurm Setup

Once the GCP project account was created and credits transferred, user accounts were created for the lab members to access GCP setup. Similar to XSEDE HPC setup, users have to log in to a login node where they can submit the jobs to compute nodes using Slurm scheduler. The Slurm controller node runs the Slurm controller and the database. Figure 1 shows the basic architectural diagram of a Slurm Cluster in Google Cloud Platform and Figure 2 shows this setup instantiated in the VM instances panel in the GCP console.

A custom image from the centos-7 image was created to install and setup our custom software setup: NEURON, NetPyNE, and Python3. To deploy the cluster, we first downloaded the git repository that contained the Slurm for GCP deployment-manager files and followed the instructions outlined in the google cloud lab user guide [24]. The YAML file, from the git repository, contains the details of the configuration of the deployment, the slurm version to deploy, and the machine instance types to deploy. The YAML file was modified to include the newly created disk image and the type of instances required for launching the compute nodes. The controller instances were initially created with 8 cores and later changed to 16 or 64 cores with high memory to prevent the node from crashing frequently due to processing and memory limitations. Table 1 shows the GCP instance setup for running 10 minutes of the M1 model vs performing evolutionary optimization. Once the YAML file was updated we ran the shell scripts that execute the gcloud commands to deploy the cluster, including login, controller and compute nodes.

Figure 2: GCP console showing the deployed ‘slurm’ cluster in the Deployments panel (top) and the controller, login and compute image nodes in the VM instances panel (bottom).

Table 1 shows the GCP instance setup for running 10 minutes of the M1 model vs performing evolutionary optimization. Once the YAML file was updated we ran the shell scripts that execute the gcloud commands to deploy the cluster, including login, controller and compute nodes.

Table 1: YAML configuration parameters for running M1 model for 10 minutes and performing evolutionary optimization

	M1 model	Evolutionary Optimization
Compute machine type	n1-standard-32	n1-standard-96
Compute disk size (gb)	10	10
Login machine type	n1-standard-2	n1-standard-8
Login disk size (gb)	10	10
Controller machine type	n1-standard-32	n1-highmem-64
Controller disk type (gb)	2000	2000

3.3 Experiences in setting up GCP-Slurm

Some of the initial challenges that we faced while setting up GCP-Slurm integration ranged from disk mounting to latency issues. We were able to resolve most of the issues by communicating with customer support and initiating discussions on the google discussion forum [25]. Being able to run on a large number of cores also helped us to identify issues that were present in our NEURON/NetPyNE software, which led to modification of the code to fix bugs. Some of the issues we encountered during our learning experience with the GCP-slurm setup are captured in Table 2.

Table 2: Our learning experience with cloud setup and Slurm integration

Issue Description	Mitigation/Learning
Error automatically mounting /home folder via NFS	With helpful discussions from GCP-Slurm forum, the problem was identified as a result of incompatibility between python versions in the disk image and the GCP startup scripts (disk image having Python 3 but the GCP startup scripts required Python 2).

Running simulations on large clusters (> 16 nodes) resulted in job scheduler crashing / stop responding.	The issue was resolved when we increased the number of cores/memory for the controller node (changed from 8 cores -> 16 cores; standard -> highmem).
~25% simulations crash with a timeout error from the NEURON simulator that happens on random simulations at random times.	After working closely with NEURON developers to debug this issue, and communicating with Google support, we concluded that the timeout error happens due to the internode latency. Currently there is no existing solution that will enable workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication which is required for our simulations.. However, running each job on a single 96-core avoided the timeout errors.
Different results obtained across different numbers of cores while running the M1 model.	The reproducibility issue was fixed by reducing the integration step (dt) from 0.05 ms to 0.025 ms in NEURON/NetPyNE. The availability and flexibility provided by GCP, facilitated identifying and fixing this issue in our software.
Shutdown of node instances took over 3 hours	During the shutdown process, Slurm waits for each node to return that it's shutdown and deleted from the list. An increase in the number of nodes caused an increase in the overall shutdown time. This was fixed by decreasing the Suspend/Resume timeout settings in the configuration and removing the option to wait for confirmation from the nodes (wait_for_operation).

We explored the use of preemptible VMs that are significantly cheaper (approximately by a factor of 4) than the regular instances, but requires the code to be fault tolerant in case of preemption. GCP-Slurm was set up to automatically resubmit jobs that get preempted. To use preemptible cores we enabled the option to perform preemptible bursting in the YAML file.

4 Simulation Results

We show here preliminary results from running the M1 model on GCP using Slurm.

4.1 Single Simulation Output

We present here the results from a single simulation of the M1 cortical model, with 10k cells and over 30M synapses. Figure 3 shows a 3D representation of the cortical network, including the detailed morphologies of several pyramidal neurons. Figure 4 shows the spike time raster plot for all neurons, grouped and colored by population, during 2 seconds of model time. Simulating the network required solving the differential equations that describe the electrical properties of each neuron, including their ionic and synaptic currents and membrane voltage (order 10ms), and how they generate action potentials (spikes, order 1 ms), a stiff problem. Membrane voltage traces of two of 2 typical network neurons are shown in Figure 4 (bottom).

4.2 Long duration simulation

We then investigated self-organizing criticality and avalanches in the M1 model. Neuronal avalanches are a cascade of bursts of activity in neuronal networks whose size distribution follows a power law. This is a robust and reproducible phenomenon observed experimentally across cortex, but whose physiological mechanisms are not currently known. Published data shows that avalanches exhibited a scale free pattern with a constant power law relation across different time scales. Therefore, investigating this phenomenon required running simulations for long durations, beyond the simulation durations ranging 1-5s we typically employed. The longest simulation we had previously ran was for 60s. Using GCP we were now able to run a simulation of 10 minutes, an order of magnitude higher. The simulation ran on 512 cores (32 nodes x 16 cores/node), took approximately 90 hours, and produced 22,610,549 spikes and 1.2 TB of data. This was possible due to the unconstrained number of hours available on this platform.

Figure 3: 3D representation of M1 cortical network, illustrating the detailed morphologies of multiple pyramidal neurons.

Figure 4: Output of single simulation of the motor cortex (M1) model. Top. Raster plot showing spikes (points) over time (x-axis) for 10k neurons. Bottom. Membrane voltage of two example cells in the network, showing action potentials (brief spikes of increased voltage over baseline).

4.3 Grid search parameter optimization

We ran simulations on approximately 30K cores simultaneously to perform grid-search parameter optimization in the M1 model. The aim was to find the values for 7 model parameters that produced average population firing rates that approximately matched experimental data. Model parameters included the overall connection strengths between excitatory and inhibitory populations, and the strength of long-range inputs driving the model. A total of 1152 jobs were run to evaluate different parameter combinations, where each job simulated 2 sec of the M1 model. Each job ran on 48 cores (8 nodes * 8 cores/node) for approximately 2.5 hours, making a total of approximately 150K core hours. Since preemptible nodes were employed, approximately 20% of the jobs were automatically resubmitted after being preempted. The grid search simulations generated a total of 21GB of output data. Exploration of the 1152 output simulations allowed us to find the parameter combination that best reproduced cortical experimental data. The results were presented at Google Next 2018 Cloud talks in San Francisco and London.

4.4 Evolutionary algorithm optimization

Evolutionary algorithms optimization was performed on a reduced version of the M1 microcircuits model, with over 2.5k multicompartiment neurons with conductance-based ionic channels, 7k point neurons for long-range inputs, and over 5M synapses. Ten parameters at different scales were optimized within biologically constrained ranges: firing rates of populations providing long-range inputs, weight scaling factors of local

microcircuit connections, and maximum conductance of HCN channels in corticospinal cells.

The fitness function quantified the similarity between experimental and model firing rates for the 15 network populations as shown in Figure 5. Each simulation (fitness evaluation) was parallelized on 128 cores and all 60 evaluations were executed in parallel, resulting in approx. 20 minutes per generation.

Figure 5: Evolutionary algorithm optimization fitness error (best, median, average and std) across generations. The best fitness error gradually decreased with generations

After 120 generations the best fitness had increased by a factor of four, which lead to a model with all population rates within a biologically reasonable range – in contrast to initial solutions where several populations were silent or overexcited. The best model solution and the associated parameters are shown in table 3. This optimization required approx. 40 hours on 7.7k cores (=308k core hours) of GCP.

Table 3: Best model solution

Population	Rate (Hz)	Fitness error
CT6	4.0	3
PV2	7.1	2
PV6	0.4	4
SOM5A	22.4	2
IT5A	26.8	29
IT5B	5.5	2
PT5B	36.4	194
PV5A	22.0	2
PV5B	24.3	3
IT6	16.5	4
IT4	7.9	2
SOM2	26.6	3
IT2	8.6	1
SOM5B	8.0	2
SOM6	19.0	1

5 Conclusion and Future Directions

This paper describes our preliminary efforts at using the GCP-Slurm integration to run simulations and optimize the parameters of our large-scale biophysically detailed cortical models. We have provided details of our setup as well as solutions to multiple issues, in the hope that this will be useful to future researchers wishing to use GCP-Slurm for similar projects.

More rigorous work is still needed to make sure that the setup is fully functional for multiple lab members to use it. Given the access to a large number of simultaneous cores and unrestricted run times through GCP, we are planning to run large scale models of other brain regions, including auditory thalamocortical circuits. Google is working on providing us access to a Placement group, where all nodes are in the same rack, which would potentially address the high latency issue. We will work with the Google team to be early testers of this feature and provide feedback. We are also planning to do further scaling studies to understand the performance characteristics of NetPyNE/NEURON. This can be done by carrying out detailed profiling and tracing analysis to understand processor and memory level performance characteristics, communication patterns and characteristics, and I/O behavior.

ACKNOWLEDGMENTS

We would like to acknowledge funding support from NSF E-CAS grant 1904444, NIH grants U01EB017695, U24EB028998 and R01DC012947, NYS SCIRB grant DOH01-C32250GG-3450000, ARO grant W911NF-19-1-0402, and cloud resources from Google.

REFERENCES

- Gao R, Asano SM, Upadhyayula S, Pisarev I, Milkie DE, Liu T-L, et al. Cortical column and whole-brain imaging with molecular contrast and nanoscale resolution. *Science*. 2019;363. doi:10.1126/science.aau8302
- Gal E, London M, Globerson A, Ramaswamy S, Reimann MW, Muller E, et al. Rich cell-type-specific network topology in neocortical microcircuitry. *Nature Neuroscience*. 2017. pp. 1004–1013. doi:10.1038/nn.4576
- Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, et al. Saturated Reconstruction of a Volume of Neocortex. *Cell*. 2015. pp. 648–661. doi:10.1016/j.cell.2015.06.054
- Grosenick L, Marshel JH, Deisseroth K. Closed-Loop and Activity-Guided Optogenetic Control. *Neuron*. 2015. pp. 106–139. doi:10.1016/j.neuron.2015.03.034
- McDougal RA, Hines ML, Lytton WW. Reaction-Diffusion in the NEURON Simulator. *Front Neuroinform*. 2013;7:28. doi:10.3389/fninf.2013.00028
- Markram H, Muller E, Ramaswamy S, Reimann MW, Abdellah M, Sanchez CA, et al. Reconstruction and Simulation of Neocortical Microcircuitry. *Cell*. 2015. pp. 456–492. doi:10.1016/j.cell.2015.09.029
- Hawrylycz M, Anastassiou C, Arkhipov A, Berg J, Buice M, Cain N, et al. Inferring cortical function in the mouse visual system through large-scale systems neuroscience. *Proceedings of the National Academy of Sciences*. 2016. pp. 7337–7344. doi:10.1073/pnas.1512901113
- Bezaire MJ, Raikov I, Burk K, Vyas D, Soltesz I. Interneuronal mechanisms of hippocampal theta oscillations in a full-scale model of the rodent CA1 circuit. *eLife*. 2016. doi:10.7554/eLife.18566
- Arkhipov A, Gouwens NW, Billeh YN, Gratiy S, Iyer R, Wei Z, et al. Visual physiology of the Layer 4 cortical circuit in silico. *bioRxiv*. 2018. doi:10.1101/292839
- Ito M. Bases and implications of learning in the cerebellum — adaptive control and internal model mechanism. *Progress in Brain Research*. 2005. pp. 95–109. doi:10.1016/S0079-6123(04)48009-1
- Dura-Bernal S, Neymotin SA, Suter BA, Shepherd GMG, Lytton WW. Multiscale dynamics and information flow in a data-driven model of the primary motor cortex microcircuit. doi:10.1101/201707
- Home - XSEDE. [cited 8 Feb 2020]. Available: <http://www.xsede.org>
- Cloud Computing Services | Google Cloud. In: Google Cloud [Internet]. [cited 8 Feb 2020]. Available: <https://cloud.google.com/>
- Dura-Bernal S, Neymotin SA, Kerr CC, Sivagnanam S, Majumdar A, Francis JT, et al. Evolutionary algorithm optimization of biological learning parameters in a biomimetic neuroprosthesis. *IBM J Res Dev*. 2017;61: 6.1–6.14. doi:10.1147/JRD.2017.2656758
- Neymotin SA, Suter BA, Dura-Bernal S, Shepherd GMG, Migliore M, Lytton WW. Optimizing computer models of corticospinal neurons to replicate in vitro dynamics. *J Neurophysiol*. 2017;117: 148–162. doi:10.1152/jn.00570.2016
- Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nat Neurosci*. 2004;7: 1345–1352. Available: <https://www.ncbi.nlm.nih.gov/pubmed/15558066>
- Achard P, De Schutter E. Complex Parameter Landscape for a Complex Neuron Model. *PLoS Computational Biology*. 2006. p. e94. doi:10.1371/journal.pcbi.0020094
- Rumbell TH, Draguljić D, Yadav A, Hof PR, Luebke JI, Weaver CM. Automated evolutionary optimization of ion channel conductances and kinetics in models of young and aged rhesus monkey pyramidal neurons. *J Comput Neurosci*. 2016;41: 65–90.
- Exploring Clouds for Acceleration of Science | Internet2. [cited 8 Feb 2020]. Available: <https://www.internet2.edu/vision-initiatives/initiatives/exploring-clouds-acceleration-science/>
- NEURON | empirically-based simulations of neurons and networks of neurons. [cited 8 Feb 2020]. Available: www.neuron.yale.edu
- Welcome to NetPyNE's documentation! — NetPyNE documentation. [cited 8 Feb 2020]. Available: www.netpyne.org
- Tonda A. Inspyred: Bio-inspired algorithms in Python. *Genetic Programming and Evolvable Machines*. 2019. doi:10.1007/s10710-019-09367-z
- Slurm Workload Manager - Documentation. [cited 8 Feb 2020]. Available: <https://slurm.schedmd.com>
- Deploy an Auto-Scaling HPC Cluster with Slurm. [cited 8 Feb 2020]. Available: <https://codelabs.developers.google.com/codelabs/hpc-slurm-on-gcp/#0>
- Google Groups. [cited 8 Feb 2020]. Available: <https://groups.google.com/forum/#!topic/google-cloud-slurm-discuss/>