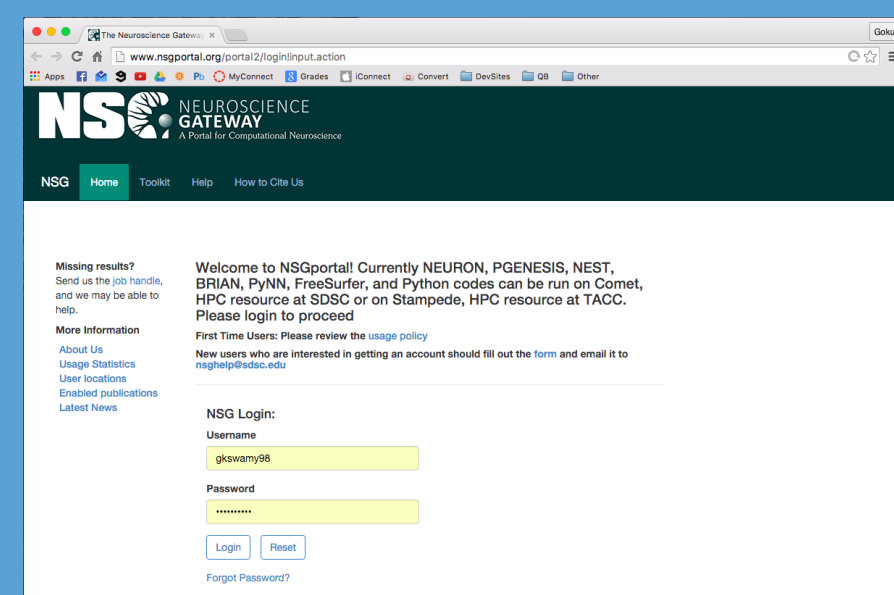


Abstract

Our work this summer at the San Diego Supercomputer Center can be divided into three sections: running models, improving the NSG Portal, and a scaling study. First, we explored how supercomputers can be used to simulate neurons firing by running models directly on the Comet supercomputer. After we achieved a greater understanding of computationally modeling the brain, we proceeded to work on improving the NSG Portal. The portal, which gives researchers free access to XSEDE resources like SDSC's Comet, allows the uploading of models and the review of their outputs. I worked with my mentors to improve the usability and design of the portal. We made it easier to access current NSG news and generally modernized and simplified the appearance of the portal. After improving the portal, we proceeded to run a balanced and unbalanced model through the portal of the brain's mu rhythms on an increasing number of cores in a scaling study. We discovered that both model's runtimes decreased polynomially with the number of cores, with the balanced model running quicker.

Introduction

The main chunk of our research focused on improving the Neuroscience Gateway, an online portal for neuroscientists to easily run their models on supercomputers. Our work made it easier for scientists to conduct independent computer-assisted research. We also performed a scaling study to see how quickly models run on different numbers of cores which showed how the large number of cores supercomputers possess make complex computations feasible.



Methods

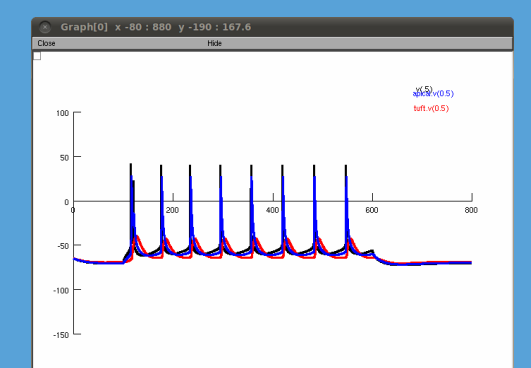
For running models, we used the SLURM (Simple Linux Utility for Resource Management) to run bash scripts on Comet. Slurm took our requests and submitted them to a queue where they were eventually run.



For improving the Portal, we used the web development languages of HTML, CSS, and Javascript to make the UX (User Experience) less complex.



For our scaling study, we used a language called Neuron, developed by researchers at Yale, to model the brain performing different types of activity.



Story

The goal of our project was to explore how computers can be used to model the brain, make it easier for professionals to do so, and then see how models reacted to different conditions.



```
#!/bin/bash
#SBATCH --job-name=jones
#SBATCH --partition=compute
#SBATCH --time=00:30:00
#SBATCH --nodes=1
#SBATCH --tasks-per-node=8

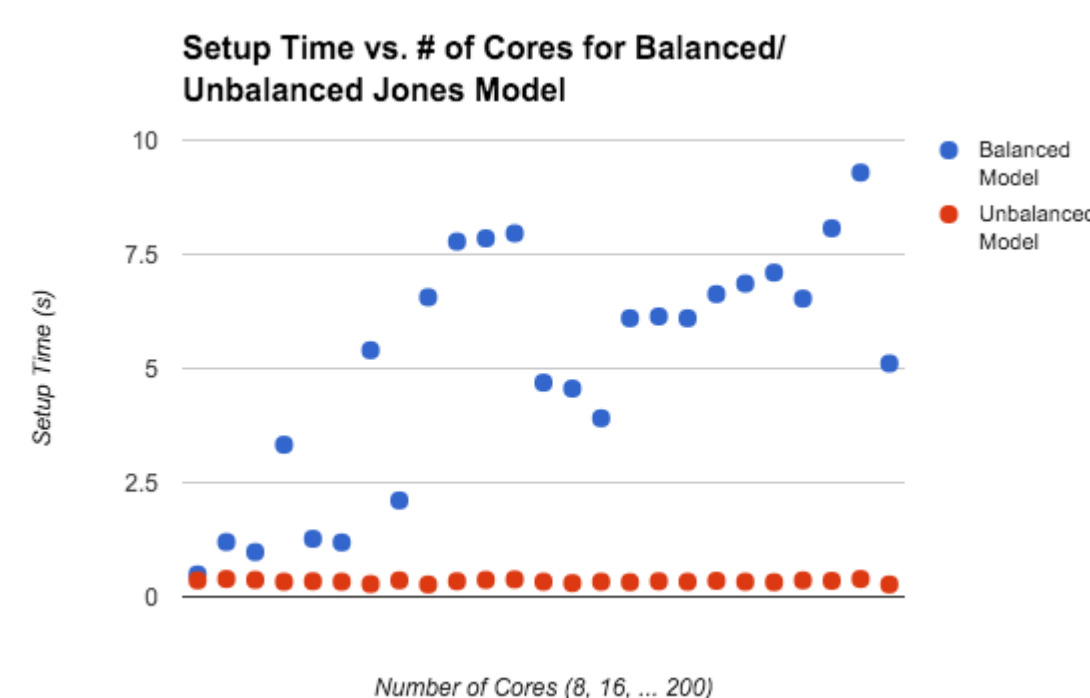
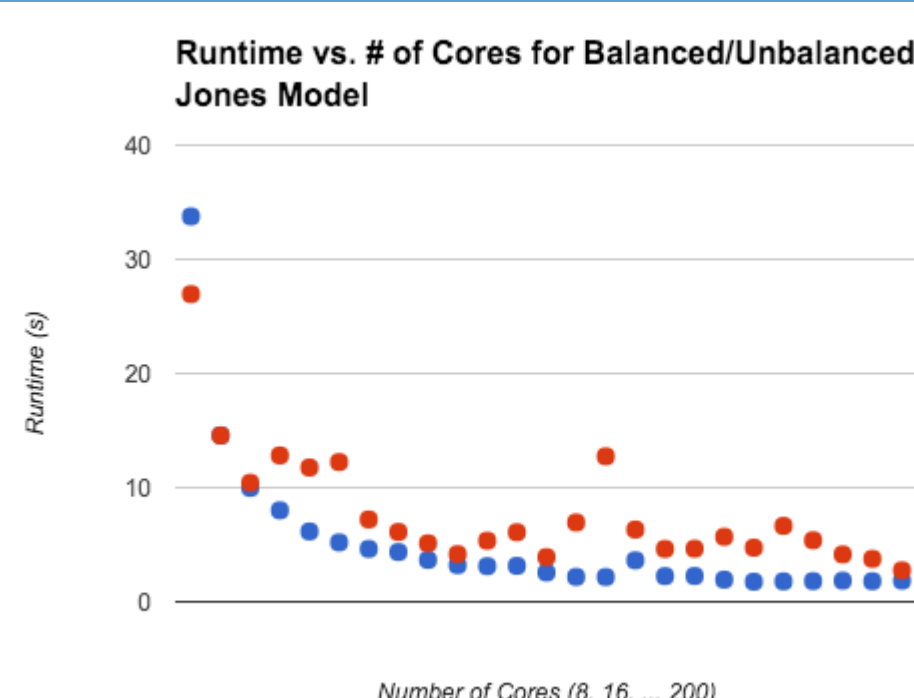
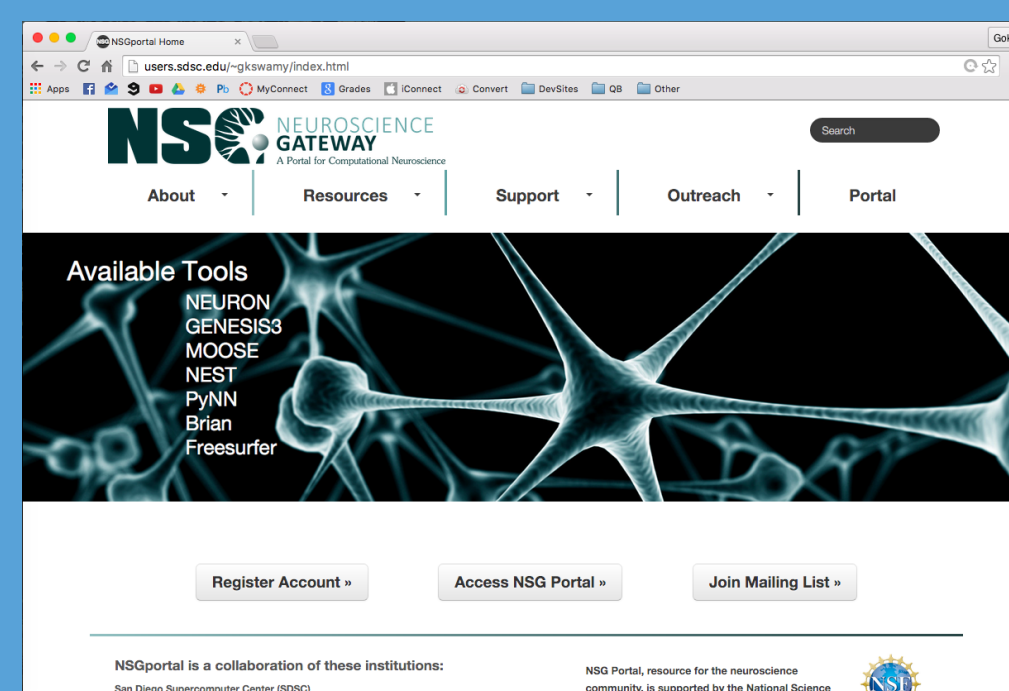

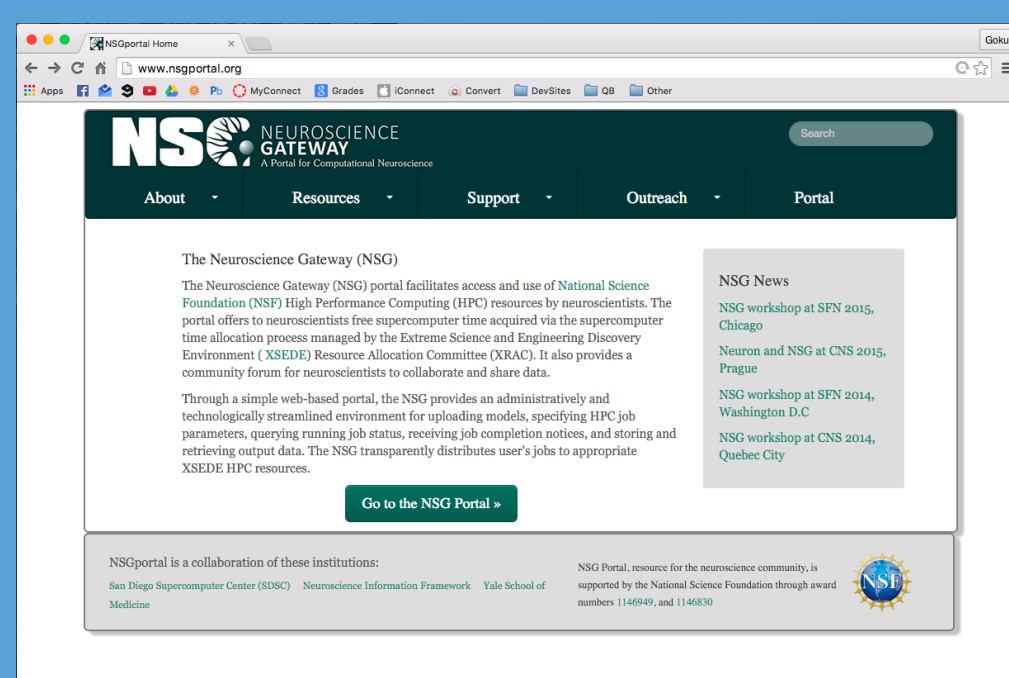
module purge
module load gnu
module load openmpi_lib
module load python
module load gsl
module load intel
cd /home/gkswamy/JonesETAL2009
ibrun -npe=nodes1 /projects/ps-nsf/home/nsuser/applications/neuron73/nrn/INSTALL/
x86_64/bin/nrniv -mpi Batch.hoc
```

Step 1: We ran a model directly on Comet to see how the brain can be simulated using a computer. We submitted a batch job to a resource manager called SLURM and then collected the output

My mentors are interested to see the relationship between runtime and the number of cores in which the model is run on. In addition, my mentors wanted a more clean and easily updatable platform for neuroscientists looking for information on the NSG website.

Step 3: We tested over 200 models for compatibility with the NSG. However, only models that created a file output were compatible as we could not access the GUI's (graphical user interfaces) that were pulled up. Balanced and unbalanced versions of the Jones model were run on Comet to test how model run time was related to the number of cores, a relationship we predicted was linear. Instead run-time was found to be cubically decreasing for both while the setup time (time for the model to read files and initialize structures) increased for the balanced model while staying roughly constant for the unbalanced version. The balanced model ran quicker.

Step 2: Using our knowledge of running models, we redesigned the NSG Portal website to be more user and admin friendly. We opted to go for a more clean and modernized look to capture the cutting-edge nature of the neuroscience performed through the portal. We used HTML, CSS, and Javascript to modify the files hosted on our personal SDSC directories.



Conclusion

Our research project answered much needed calls for a revamped NSG website making it easier for neuroscientists to access information and for my mentors to update information. We learnt about how to model the brain in Yale's Neuron language, learnt what kind of models can be run through the Portal, and discovered how models scale on supercomputers. We were able to improve a resource and then use it to answer our own questions.