Hopper Cluster

Hopper is our group's high-performance computing cluster, which we share with the Reed, Steinschneider, and Anderson groups.

Zach Brodeur wrote a user guide for Hopper and a more detailed guide which you should read.

Access

To get access to Hopper, message Vivek (this should be part of your onboarding process).

You can access Hopper for file transfers through Globus (see the Cornell CAC wiki for information on how to use Globus).

Hardware

- 22 compute nodes (c0001-c0022) with dual 20-core Intel Xeon Gold 5218R CPUs @ 2.1 GHz, 192 GB of RAM
- 225 TB of total usable storage. This should be a healthy amount per person. We'll figure out a more formal way to partition storage resources over time.
- Hyperthreading is available on all nodes and enabled by default (so you can treat each physical core as two logical cores, if the software allows; in other words, think of each node as consisting of 80 CPUs). If you don't want to use hyperthreading, including the following in your submission script: `#SBATCH --hint=nomultithread`

Software

Filesystem

Use your home directory (`~`) for data you want to keep. These directories are not backed up by default.

Scheduler

Hopper uses the SLURM scheduler. There is only one queue (`normal`) with no resource or time limits.