Hummingbird Computational Cluster

Overview and Introduction

Hummingbird Support Team hummingbird@ucsc.edu

Hummingbird Etiquette

The DOs and DON'Ts of using the Hummingbird Cluster

Software - the `module` system

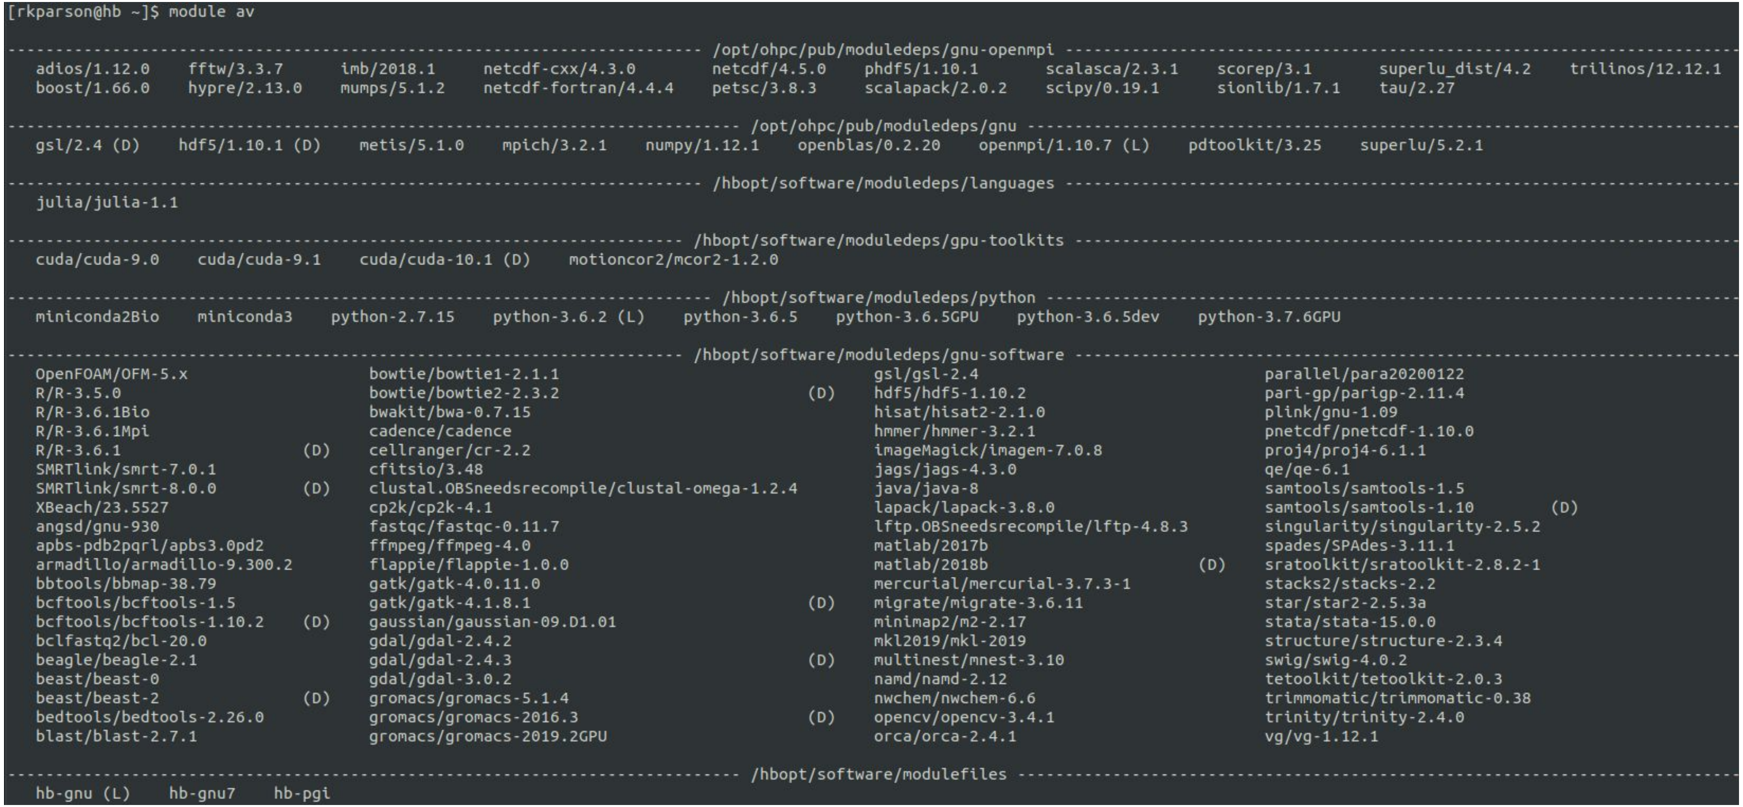

```
[rkparson@hb ~]$ module av

------------------------------------------------------------------------- /opt/ohpc/pub/moduledeps/gnu-openmpi -----------------------------------------------------------------------
   adios/1.12.0    fftw/3.3.7      imb/2018.1     netcdf-cxx/4.3.0        netcdf/4.5.0    phdf5/1.10.1       scalasca/2.3.1    scorep/3.1       superlu_dist/4.2    trilinos/12.12.1
   boost/1.66.0    hypre/2.13.0    mumps/5.1.2    netcdf-fortran/4.4.4    petsc/3.8.3     scalapack/2.0.2    scipy/0.19.1      sionlib/1.7.1    tau/2.27

----------------------------------------------------------------------------- /opt/ohpc/pub/moduledeps/gnu ---------------------------------------------------------------------------
   gsl/2.4 (D)    hdf5/1.10.1 (D)    metis/5.1.0    mpich/3.2.1    numpy/1.12.1    openblas/0.2.20    openmpi/1.10.7 (L)    pdtoolkit/3.25    superlu/5.2.1

------------------------------------------------------------------------- /hbopt/software/moduledeps/languages -----------------------------------------------------------------------
   julia/julia-1.1

----------------------------------------------------------------------- /hbopt/software/moduledeps/gpu-toolkits ----------------------------------------------------------------------
   cuda/cuda-9.0    cuda/cuda-9.1    cuda/cuda-10.1 (D)    motioncor2/mcor2-1.2.0

-------------------------------------------------------------------------- /hbopt/software/moduledeps/python -------------------------------------------------------------------------
   miniconda2Bio    miniconda3    python-2.7.15    python-3.6.2 (L)    python-3.6.5    python-3.6.5GPU    python-3.6.5dev    python-3.7.6GPU

----------------------------------------------------------------------- /hbopt/software/moduledeps/gnu-software ----------------------------------------------------------------------
   OpenFOAM/OFM-5.x                   bowtie/bowtie1-2.1.1                                 gsl/gsl-2.4                              parallel/para20200122
   R/R-3.5.0                          bowtie/bowtie2-2.3.2                          (D)    hdf5/hdf5-1.10.2                         pari-gp/parigp-2.11.4
   R/R-3.6.1Bio                       bwakit/bwa-0.7.15                                    hisat/hisat2-2.1.0                       plink/gnu-1.09
   R/R-3.6.1Mpi                       cadence/cadence                                      hmmer/hmmer-3.2.1                        pnetcdf/pnetcdf-1.10.0
   R/R-3.6.1                   (D)    cellranger/cr-2.2                                    imageMagick/imagem-7.0.8                 proj4/proj4-6.1.1
   SMRTlink/smrt-7.0.1                cfitsio/3.48                                         jags/jags-4.3.0                          qe/qe-6.1
   SMRTlink/smrt-8.0.0         (D)    clustal.OBSneedsrecompile/clustal-omega-1.2.4        java/java-8                              samtools/samtools-1.5
   XBeach/23.5527                     cp2k/cp2k-4.1                                        lapack/lapack-3.8.0                      samtools/samtools-1.10        (D)
   angsd/gnu-930                      fastqc/fastqc-0.11.7                                 lftp.OBSneedsrecompile/lftp-4.8.3        singularity/singularity-2.5.2
   apbs-pdb2pqrl/apbs3.0pd2           ffmpeg/ffmpeg-4.0                                    matlab/2017b                             spades/SPAdes-3.11.1
   armadillo/armadillo-9.300.2        flappie/flappie-1.0.0                                matlab/2018b                      (D)    sratoolkit/sratoolkit-2.8.2-1
   bbtools/bbmap-38.79                gatk/gatk-4.0.11.0                                   mercurial/mercurial-3.7.3-1              stacks2/stacks-2.2
   bcftools/bcftools-1.5              gatk/gatk-4.1.8.1                             (D)    migrate/migrate-3.6.11                   star/star2-2.5.3a
   bcftools/bcftools-1.10.2    (D)    gaussian/gaussian-09.D1.01                           minimap2/m2-2.17                         stata/stata-15.0.0
   bclfastq2/bcl-20.0                 gdal/gdal-2.4.2                                      mkl2019/mkl-2019                         structure/structure-2.3.4
   beagle/beagle-2.1                  gdal/gdal-2.4.3                               (D)    multinest/mnest-3.10                     swig/swig-4.0.2
   beast/beast-0                      gdal/gdal-3.0.2                                      namd/namd-2.12                           tetoolkit/tetoolkit-2.0.3
   beast/beast-2               (D)    gromacs/gromacs-5.1.4                                nwchem/nwchem-6.6                        trimmomatic/trimmomatic-0.38
   bedtools/bedtools-2.26.0           gromacs/gromacs-2016.3                        (D)    opencv/opencv-3.4.1                      trinity/trinity-2.4.0
   blast/blast-2.7.1                  gromacs/gromacs-2019.2GPU                            orca/orca-2.4.1                          vg/vg-1.12.1

----------------------------------------------------------------------------- /hbopt/software/modulefiles ----------------------------------------------------------------------------
   hb-gnu (L)    hb-gnu7    hb-pgi
```

Using the cluster - Scheduling

Template SLURM scripts

Premade templates for your convenience. Just make a copy to your working directory and edit to meet your needs

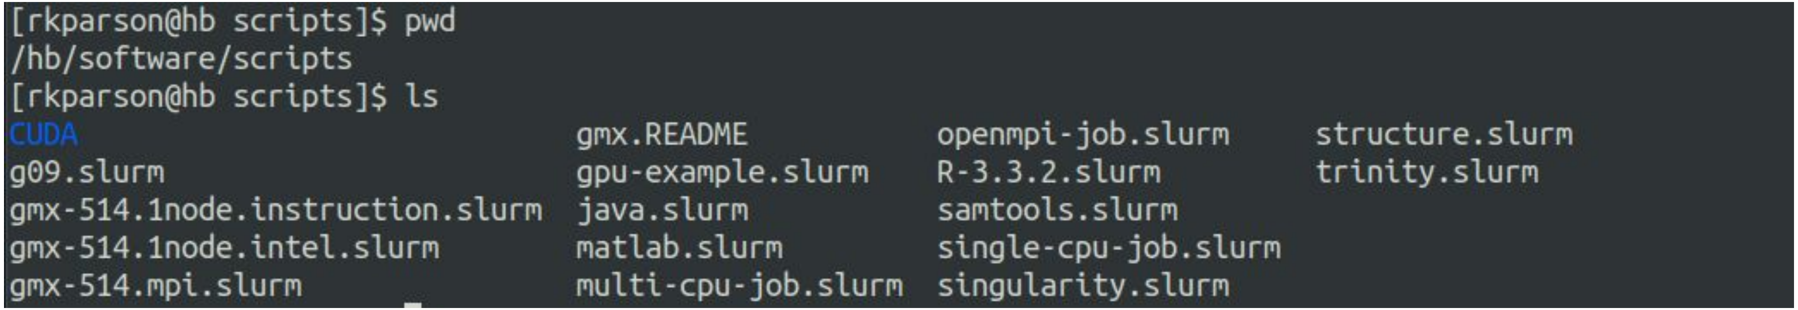

```
[rkparson@hb scripts]$ pwd
/hb/software/scripts
[rkparson@hb scripts]$ ls
CUDA                             gmx.README           openmpi-job.slurm     structure.slurm
g09.slurm                        gpu-example.slurm    R-3.3.2.slurm         trinity.slurm
gmx-514.1node.instruction.slurm  java.slurm           samtools.slurm
gmx-514.1node.intel.slurm        matlab.slurm         single-cpu-job.slurm
gmx-514.mpi.slurm                multi-cpu-job.slurm  singularity.slurm
```

If you have a template that is useful to more than just yourself, let us know and we can include it here.

Anatomy of a basic SLURM Script

Name of Partition to use

Name to give your run

Name for output/error logs
(use for troubleshooting)

Number of nodes to request

Number of tasks your
program will require

Amount of RAM requested
for your program

Time limit for the run (!!!)

Get status emails about your job

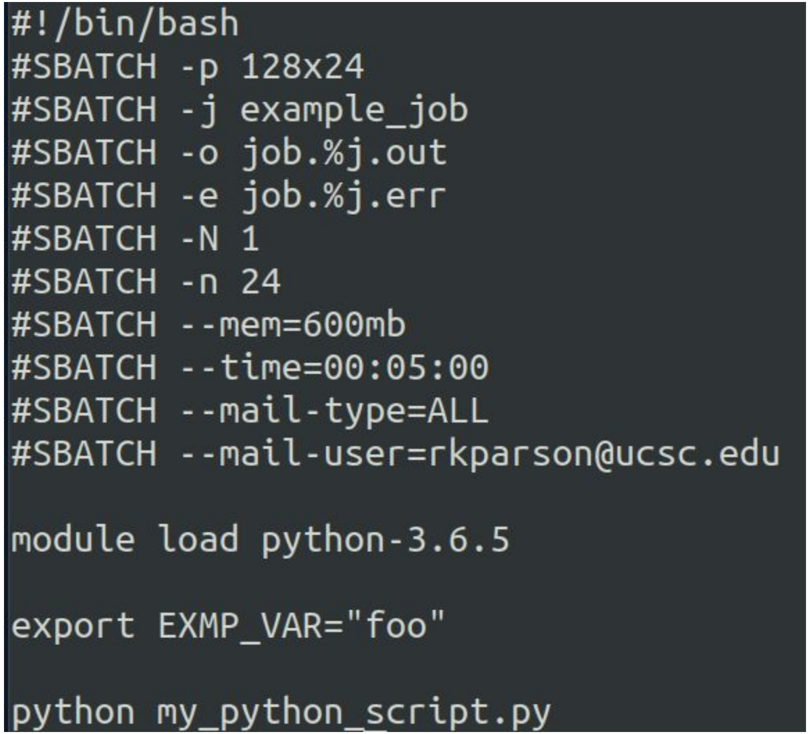

```
#!/bin/bash
#SBATCH -p 128x24
#SBATCH -j example_job
#SBATCH -o job.%j.out
#SBATCH -e job.%j.err
#SBATCH -N 1
#SBATCH -n 24
#SBATCH --mem=600mb
#SBATCH --time=00:05:00
#SBATCH --mail-type=ALL
#SBATCH --mail-user=rkparson@ucsc.edu

module load python-3.6.5

export EXMP_VAR="foo"

python my_python_script.py
```

Anatomy of a basic SLURM Script

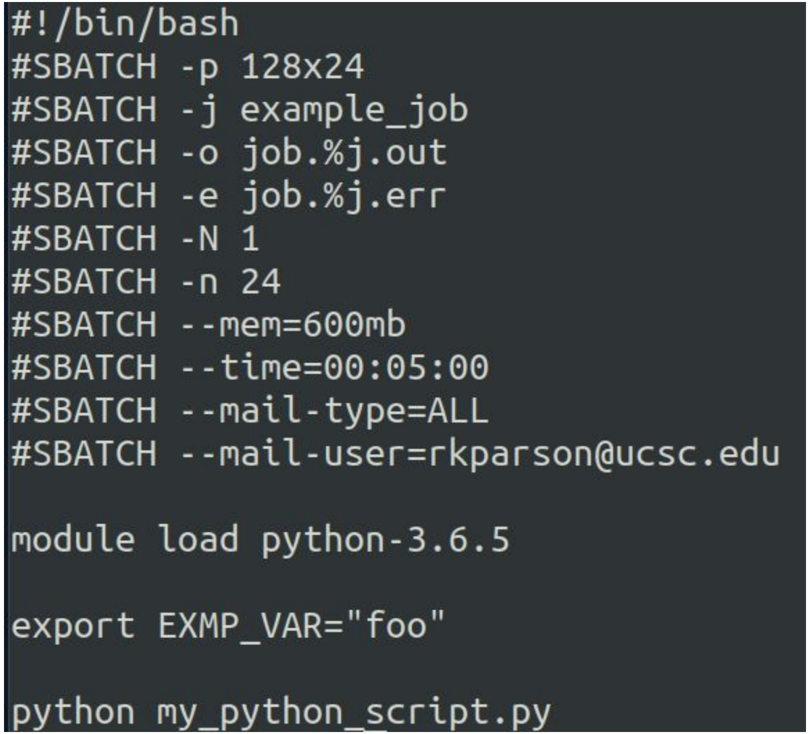

```
#!/bin/bash
#SBATCH -p 128x24
#SBATCH -j example_job
#SBATCH -o job.%j.out
#SBATCH -e job.%j.err
#SBATCH -N 1
#SBATCH -n 24
#SBATCH --mem=600mb
#SBATCH --time=00:05:00
#SBATCH --mail-type=ALL
#SBATCH --mail-user=rkparson@ucsc.edu

module load python-3.6.5

export EXMP_VAR="foo"

python my_python_script.py
```

Any modules your program requires in order to run

Any system variables your program might require

The program to run

Submitting your job

```
sbatch your_job_script.slurm
```

Verifying your job's settings

```
scontrol show job <jobid>
```

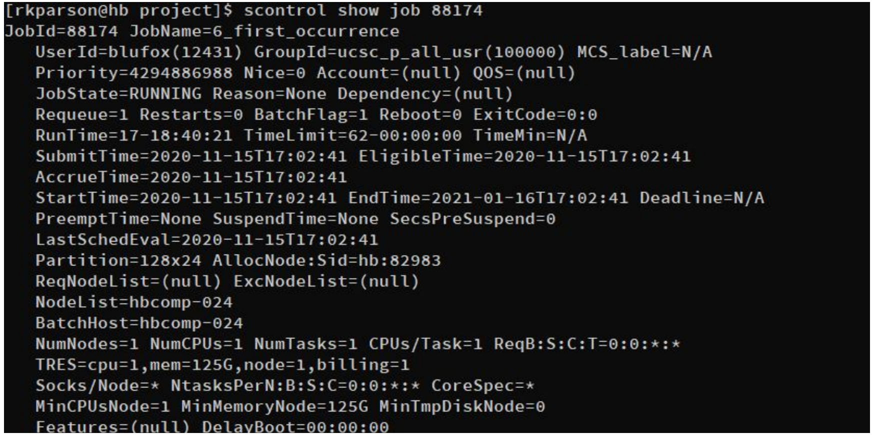

```
[rkparson@hb project]$ scontrol show job 88174
JobId=88174 JobName=6_first_occurrence
   UserId=blufox(12431) GroupId=ucsc_p_all_usr(100000) MCS_label=N/A
   Priority=4294886988 Nice=0 Account=(null) QOS=(null)
   JobState=RUNNING Reason=None Dependency=(null)
   Requeue=1 Restarts=0 BatchFlag=1 Reboot=0 ExitCode=0:0
   RunTime=17-18:40:21 TimeLimit=62-00:00:00 TimeMin=N/A
   SubmitTime=2020-11-15T17:02:41 EligibleTime=2020-11-15T17:02:41
   AccrueTime=2020-11-15T17:02:41
   StartTime=2020-11-15T17:02:41 EndTime=2021-01-16T17:02:41 Deadline=N/A
   PreemptTime=None SuspendTime=None SecsPreSuspend=0
   LastSchedEval=2020-11-15T17:02:41
   Partition=128x24 AllocNode:Sid=hb:82983
   ReqNodeList=(null) ExcNodeList=(null)
   NodeList=hbcomp-024
   BatchHost=hbcomp-024
   NumNodes=1 NumCPUs=1 NumTasks=1 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
   TRES=cpu=1,mem=125G,node=1,billing=1
   Socks/Node=* NtasksPerN:B:S:C=0:0:*:* CoreSpec=*
   MinCPUsNode=1 MinMemoryNode=125G MinTmpDiskNode=0
   Features=(null) DelayBoot=00:00:00
```

That's a lot of details!

A view into what's running - `squeue`

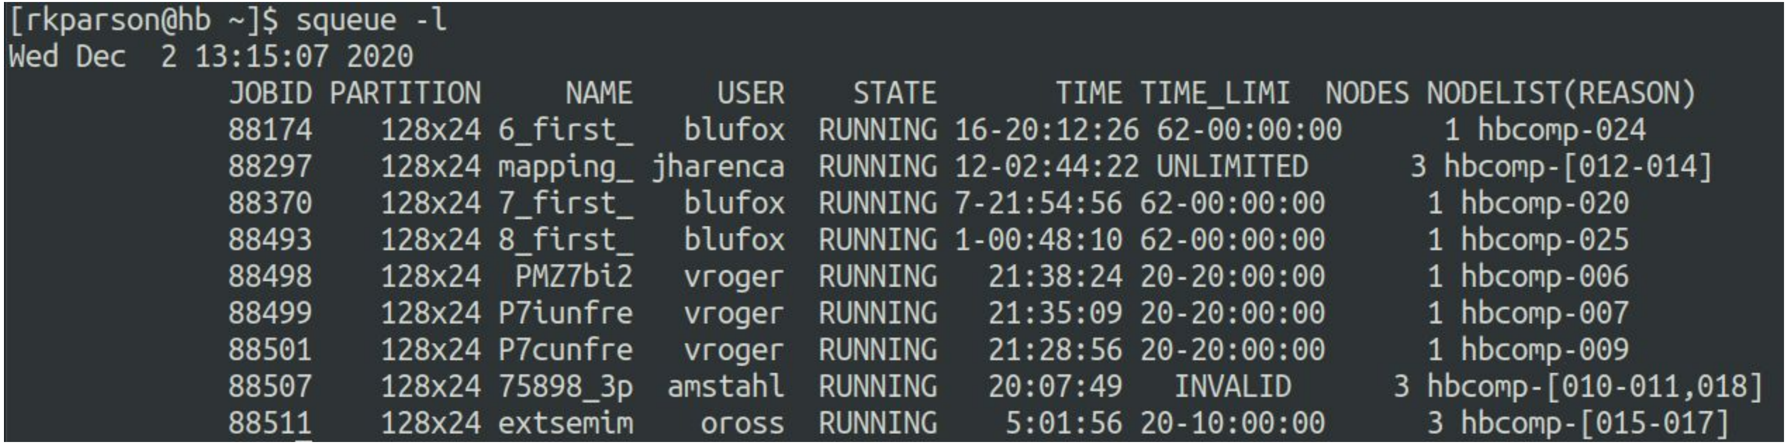

```
[rkparson@hb ~]$ squeue -l
Wed Dec  2 13:15:07 2020
             JOBID PARTITION     NAME     USER    STATE       TIME TIME_LIMI  NODES NODELIST(REASON)
             88174    128x24 6_first_   blufox  RUNNING 16-20:12:26 62-00:00:00      1 hbcomp-024
             88297    128x24 mapping_ jharenca  RUNNING 12-02:44:22 UNLIMITED      3 hbcomp-[012-014]
             88370    128x24 7_first_   blufox  RUNNING 7-21:54:56 62-00:00:00      1 hbcomp-020
             88493    128x24 8_first_   blufox  RUNNING 1-00:48:10 62-00:00:00      1 hbcomp-025
             88498    128x24  PMZ7bi2   vroger  RUNNING   21:38:24 20-20:00:00      1 hbcomp-006
             88499    128x24 P7iunfre   vroger  RUNNING   21:35:09 20-20:00:00      1 hbcomp-007
             88501    128x24 P7cunfre   vroger  RUNNING   21:28:56 20-20:00:00      1 hbcomp-009
             88507    128x24 75898_3p  amstahl  RUNNING   20:07:49   INVALID      3 hbcomp-[010-011,018]
             88511    128x24 extsemim    oross  RUNNING    5:01:56 20-10:00:00      3 hbcomp-[015-017]
```

Much easier to parse!

Hummingbird is NOT a SUPERCOMPUTER!

- Hummingbird does NOT use a high speed backplane, thus it is not suitable for highly parallelized jobs requiring many cores/nodes acting synchronously (FEM, FDTD, etc...)
- Humming IS good for serialized style processing (Monte Carlo, batch processing, etc...)
- Users are limited to using three nodes for a single job (a maximum of 72 cores)
- Most nodes only have 128 GB or RAM
- Hummingbird works on First-In, First-Out with Backfilling
- We don't *enforce* a MaxTime, so you *should*

A quick way to check the cluster - `sinfo`

Some queues are restricted and are not for general use

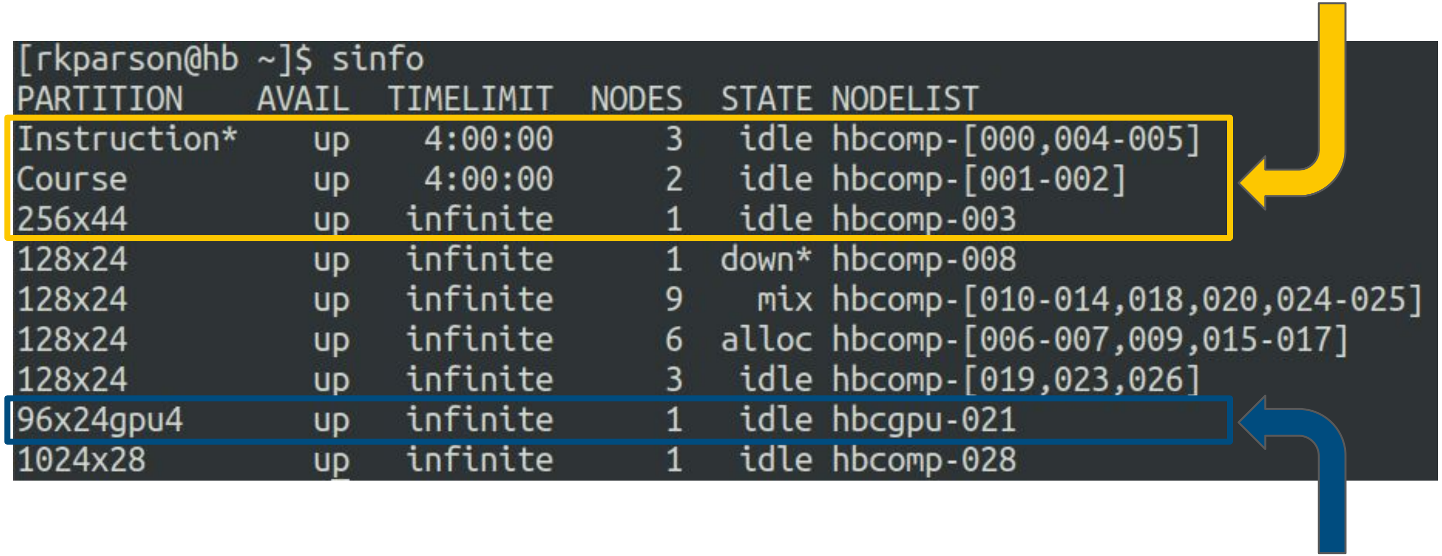

```
[rkparson@hb ~]$ sinfo
PARTITION     AVAIL  TIMELIMIT  NODES  STATE NODELIST
Instruction*     up    4:00:00      3   idle hbcomp-[000,004-005]
Course           up    4:00:00      2   idle hbcomp-[001-002]
256x44           up   infinite      1   idle hbcomp-003
128x24           up   infinite      1  down* hbcomp-008
128x24           up   infinite      9    mix hbcomp-[010-014,018,020,024-025]
128x24           up   infinite      6  alloc hbcomp-[006-007,009,015-017]
128x24           up   infinite      3   idle hbcomp-[019,023,026]
96x24gpu4        up   infinite      1   idle hbcgpu-021
1024x28          up   infinite      1   idle hbcomp-028
```

We have four NVIDIA GPUs, and are setup for CUDA and TensorFlow!

Knowing your Hummingbird variants

hb.ucsc.edu

- Cluster login node
- Used for:
  - Compiling your code
  - Checking the status on your submitted jobs

hbfeeder.ucsc.edu

- Cluster storage node
- Used for:
  - Accessing your data when you don't need the cluster
  - Transferring your data to or from the cluster

"Violet-tailed Sylph resting on a branch in northwestern Ecuador"
https://en.wikipedia.org/wiki/File:Violet-tailed_Sylph_2_JCB.jpg

Male Anna's Hummingbird
https://en.wikipedia.org/wiki/File:AnnasHummingbirdPaloAltoNorvig.jpg

How much space are you using?

- Everyone is limited to 1TB of storage in their home folders.
- If you exceed this, your jobs won't run correctly.
- How to check:
  - `du -sh ~`

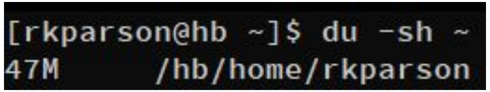

```
[rkparson@hb ~]$ du -sh ~
47M     /hb/home/rkparson
```

total used space summary

  - `du -d1 -h ~`

```
[rkparson@hb project]$ du -d1 -h ~
1.0K    /hb/home/rkparson/.config
1.5K    /hb/home/rkparson/.cache
63K     /hb/home/rkparson/.ssh
512     /hb/home/rkparson/.local
10K     /hb/home/rkparson/project
216K    /hb/home/rkparson/.lmod.d
47M     /hb/home/rkparson/.conda
47M     /hb/home/rkparson
```

used space per-folder (one-layer deep)

Key points to take away

- ✓ DO Check the Message of the Day for useful updates
- ✓ DO Check in the modules before you ask for software
- ✓ DO Use the scheduler when submitting jobs
- ✗ DON'T Use the cluster for heavily parallel style workflows
- ✓ DO Set time limits on your jobs
- ✗ DON'T Overprovision when submitting a job
- ✓ DO regularly check your disk space usage
- ✗ DON'T Leave data in Scratch long-term
- ✓ DO Reach out for help!

Reaching out for help!

Open a ticket!

https://ucsc.service-now.com/ess/

Come to our weekly open-hours sessions!

Find the link in our Message of the Day on the Hummingbird login node:

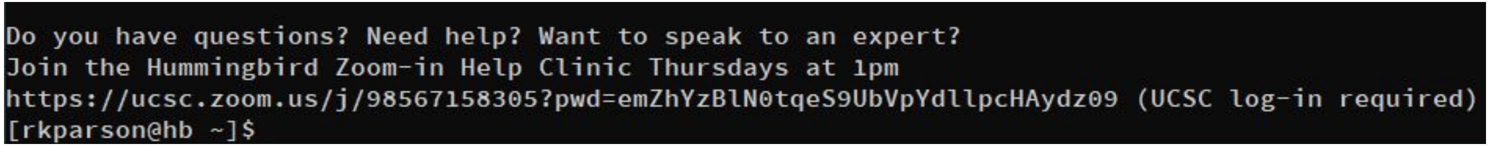

```
Do you have questions? Need help? Want to speak to an expert?
Join the Hummingbird Zoom-in Help Clinic Thursdays at 1pm
https://ucsc.zoom.us/j/98567158305?pwd=emZhYzBlN0tqeS9UbVpYdllpcHAydz09 (UCSC log-in required)
[rkparson@hb ~]$
```

Check out the website for more information!

https://www.hb.ucsc.edu/

Please take a moment and fill out our post-session survey

https://docs.google.com/forms/d/e/1FAIpQLSeaLMqv3yPcgxy3-sfwZsnCLu5PHvyXg4i60rYnSRxJVVEZog/viewform?usp=sf_link