

High-Performance Computing with RELION

David Hoover, PhD

HPC @ NIH

Biowulf Seminar Series



Sagar Chittori LCB, NCI

**Cryo-EM studies of glutamate receptors
and nucleosomes**

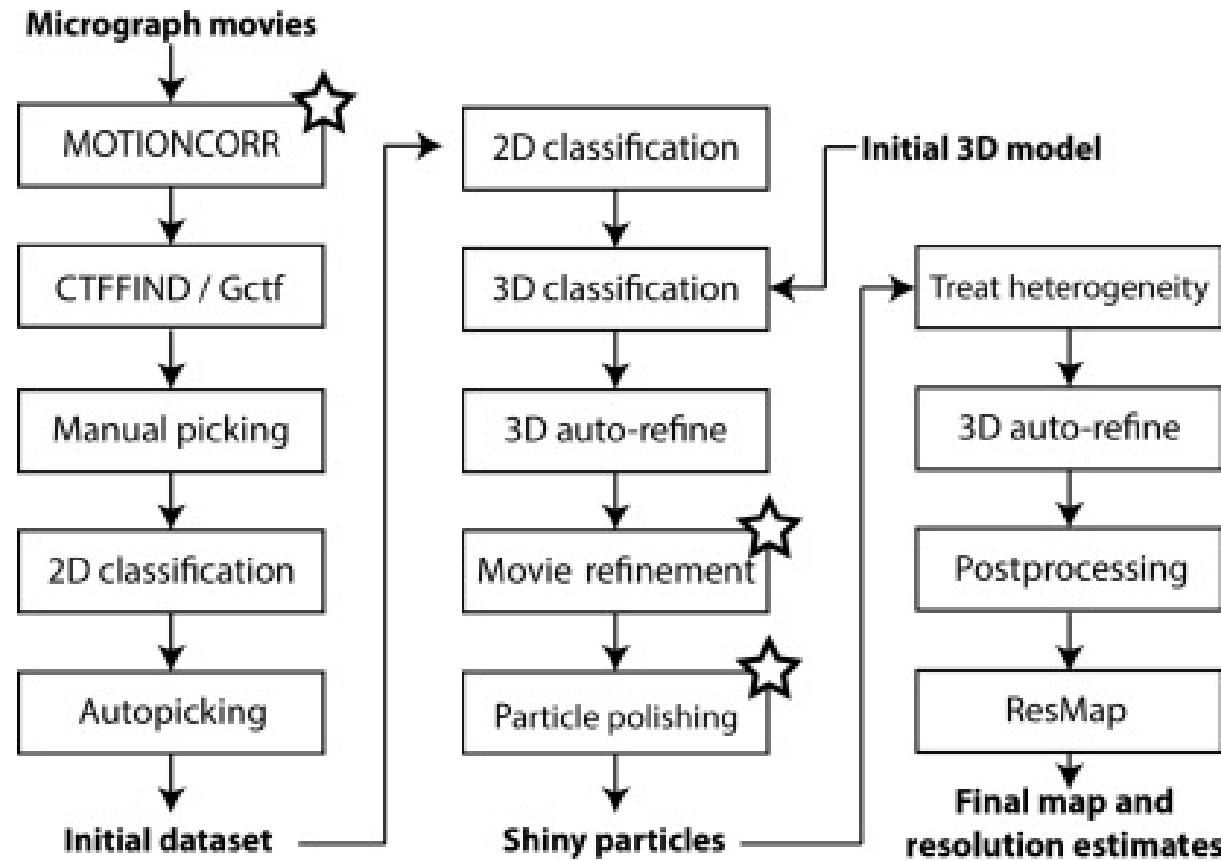
Wed May 09, 2018, 11 am – noon
Bldg 50, Rm 1227 (next to the coffeshop)

TOMORROW ONLY!

Brief Overview of RELION

- *Purify macromolecules and freeze onto platform*
- *Collect TEM images (2D projections)*
- Clean up images (defocus, tilt, motion, blur)
- Isolate individual particle projection images
- Classify particles
- Predict original orientation, maximum likelihood
- Refine prediction with “smooth” maps
- Polish particles to improve signal-to-noise

Brief Overview of RELION



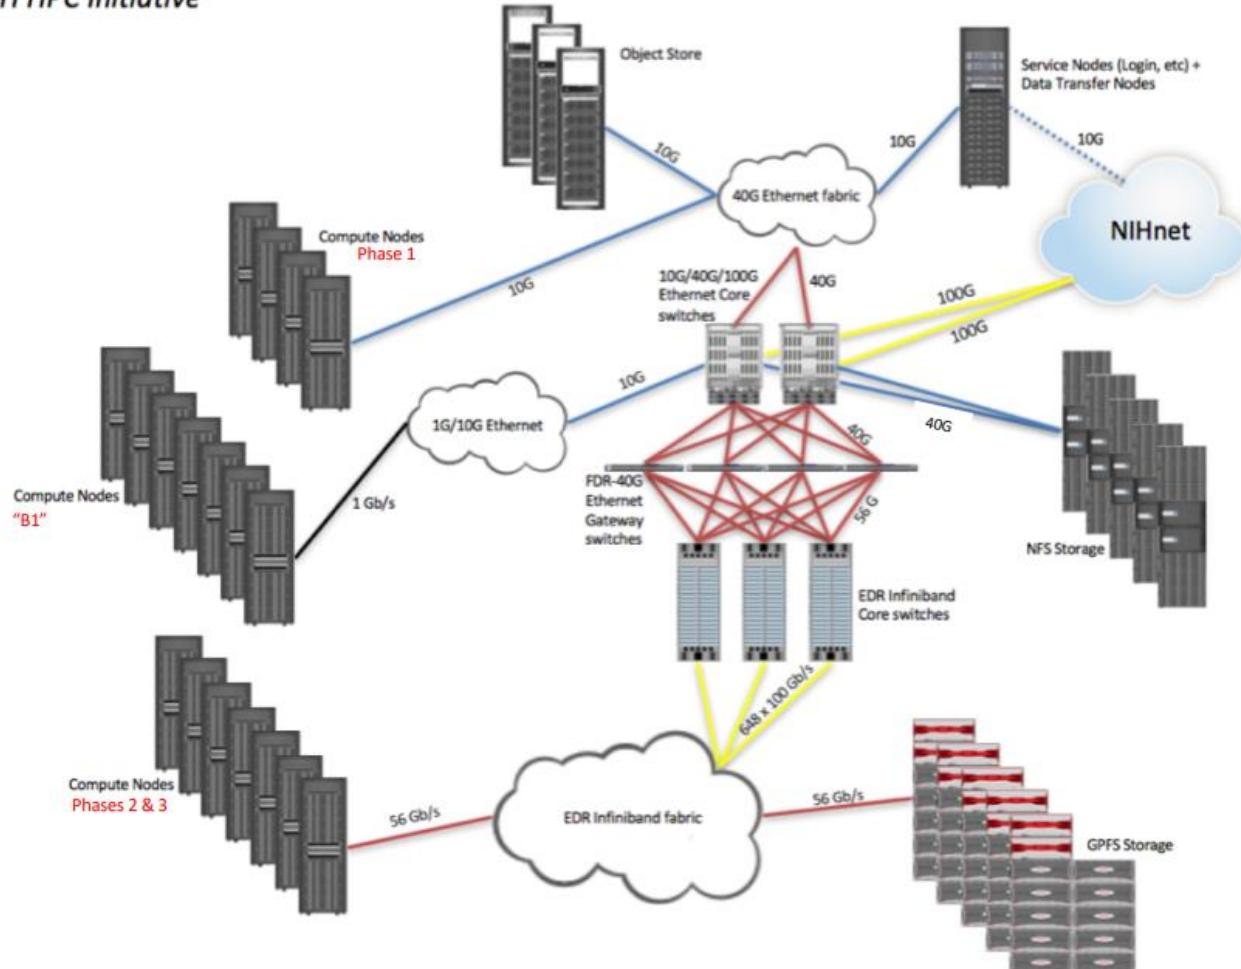
Brief Overview of RELION

- Many steps can be parallelized
- A few steps can be accelerated using GPUs

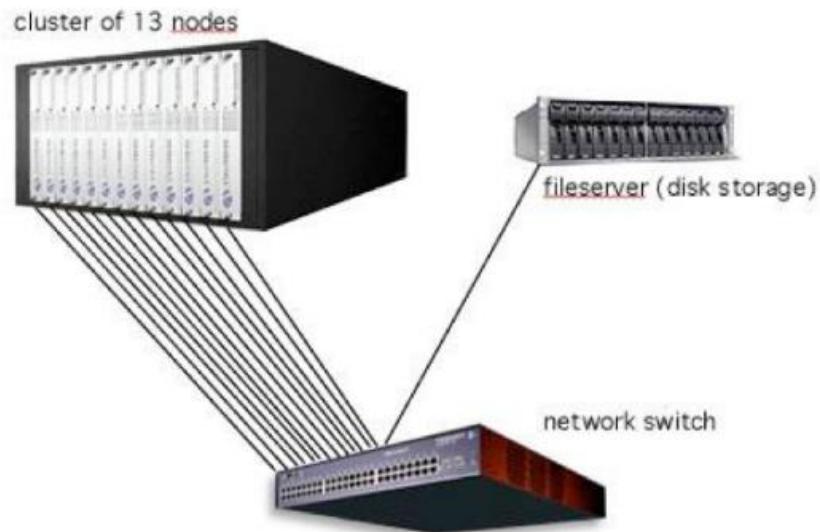
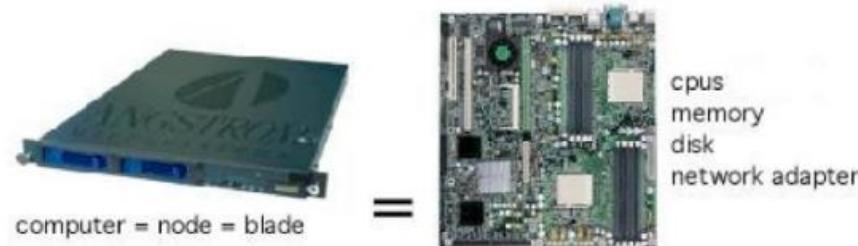
Cluster and Hardware Overview

Biowulf Cluster Overview

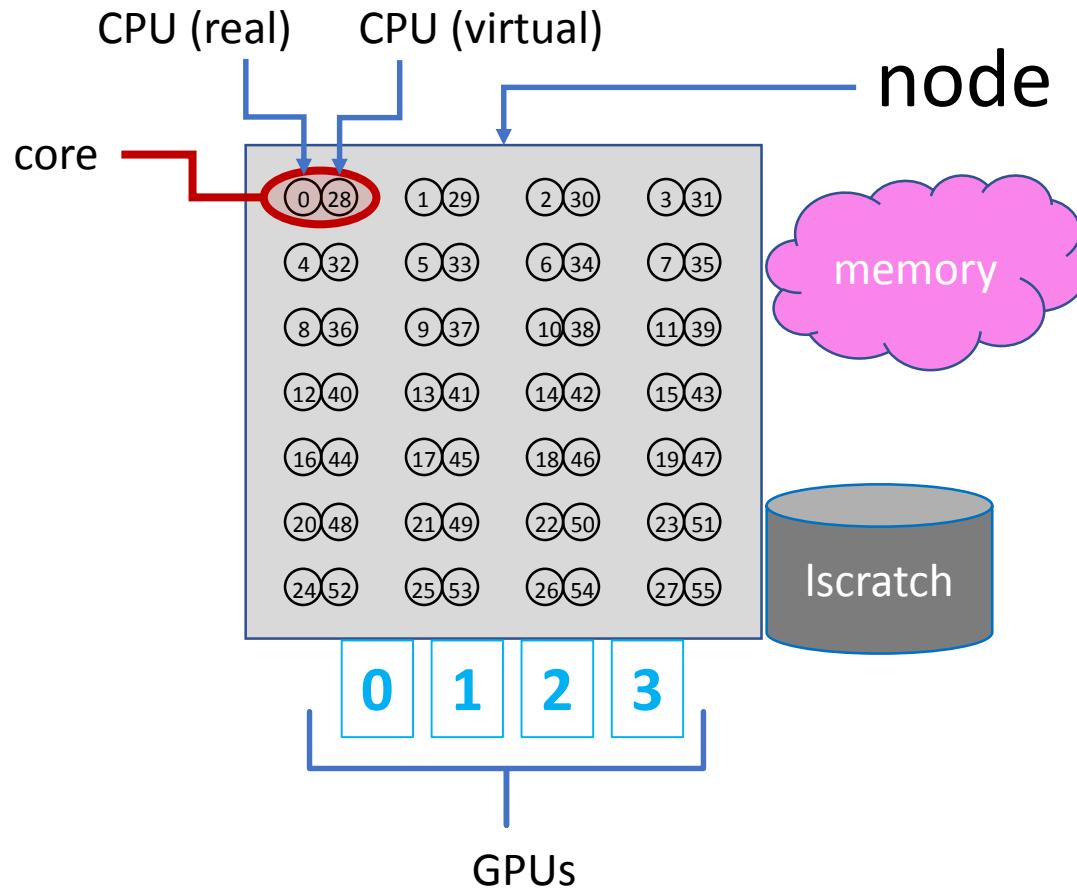
NIH HPC Initiative



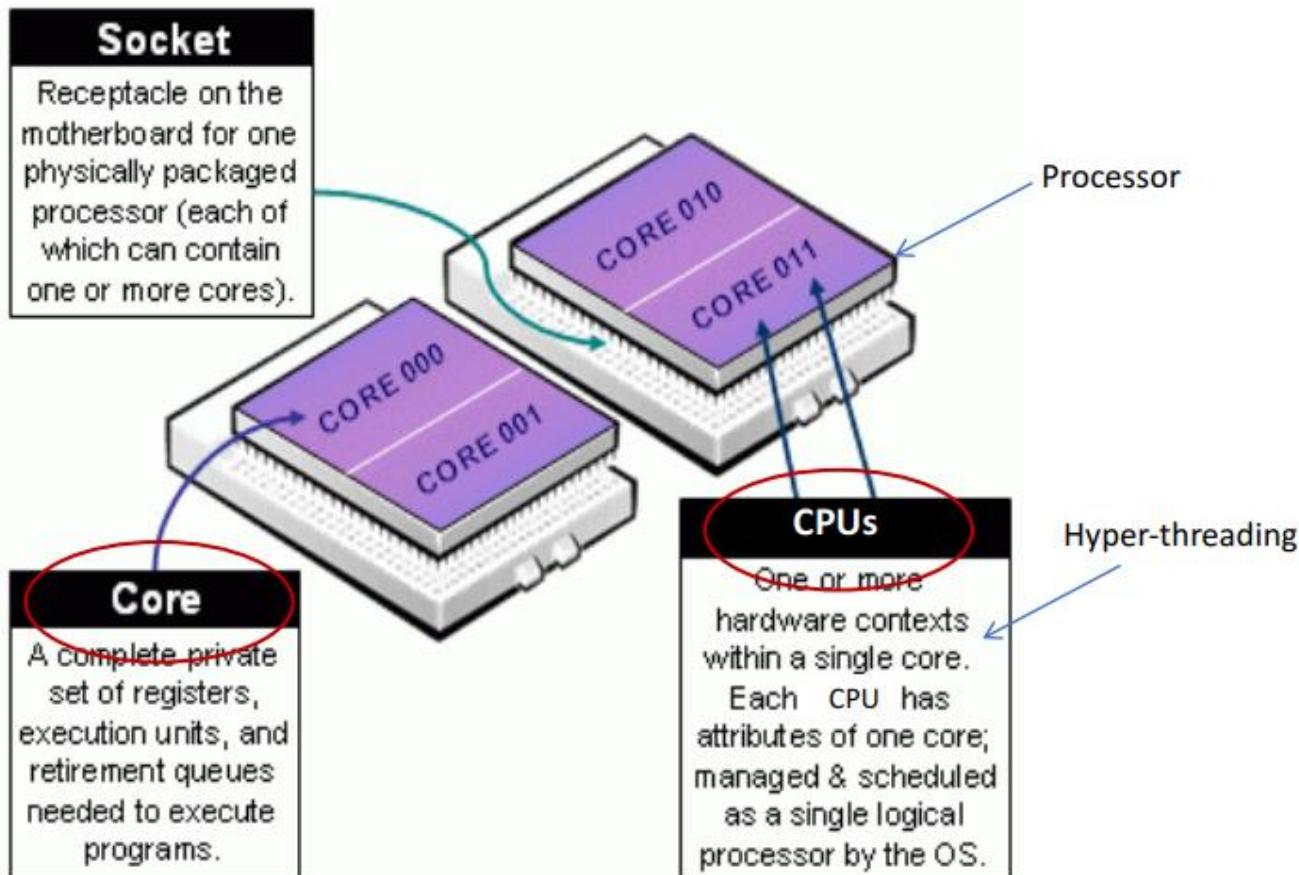
Biowulf Cluster Overview



Node Anatomy



Processor Hardware



Parallelization on the Biowulf Cluster

Parallelization on Biowulf using Slurm

- Allocation of resources
- Distribution of resources
- Distribution of tasks on the allocated and distributed resources

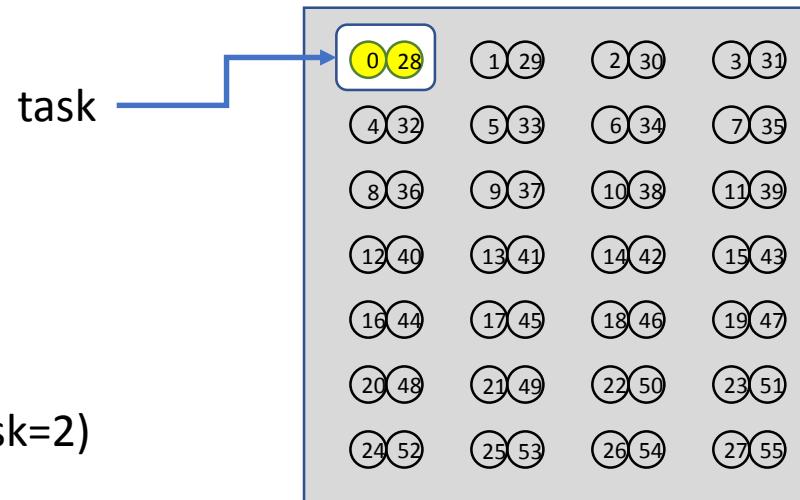


Multi-user Time-sharing Batch System



Single-Threaded Job

Slurm allocates in units of **cores**, not CPUs



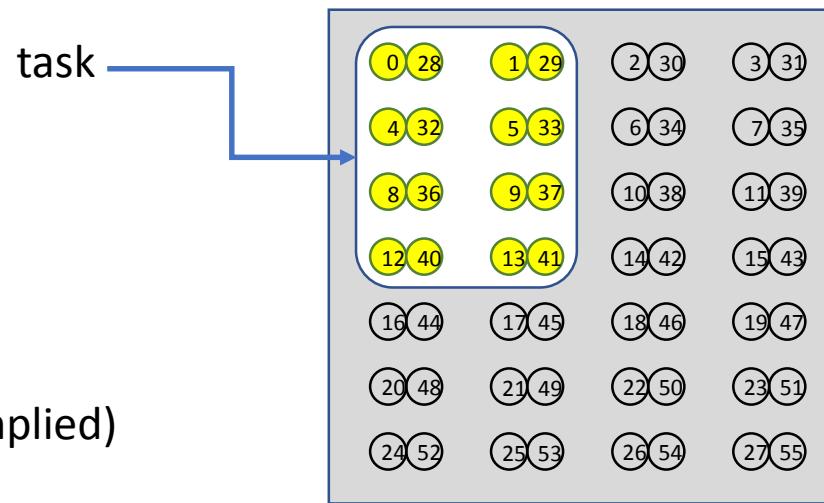
`sbatch`

executable arg arg arg ...

`sinteractive`

executable arg arg arg ...

Multi-threaded Job



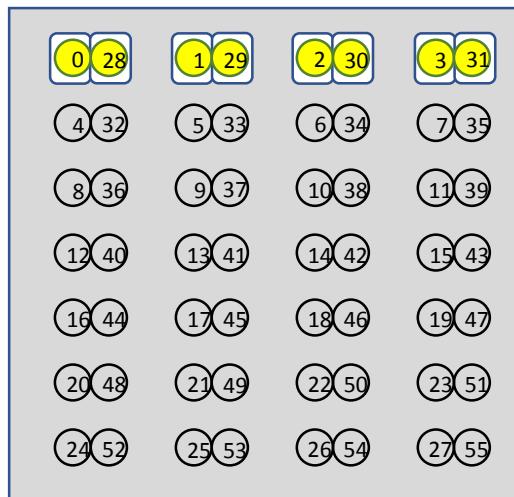
```
sbatch --cpus-per-task=16
```

```
executable -p ${SLURM_CPUS_PER_TASK} arg arg arg ...
```

```
sinteractive --cpus-per-task=16
```

```
executable -p ${SLURM_CPUS_PER_TASK} arg arg arg ...
```

MPI Job: single node



```
sbatch --ntasks=8 --nodes=1
```

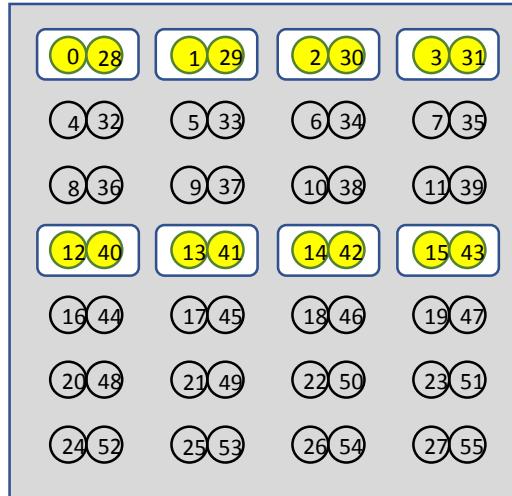
```
srub --mpi=pmi2 executable arg arg arg ...
```

```
sinteractive --ntasks=8 --nodes=1
```

```
mpirun executable arg arg arg ...
```

launcher

Multi-threaded MPI Job: single node



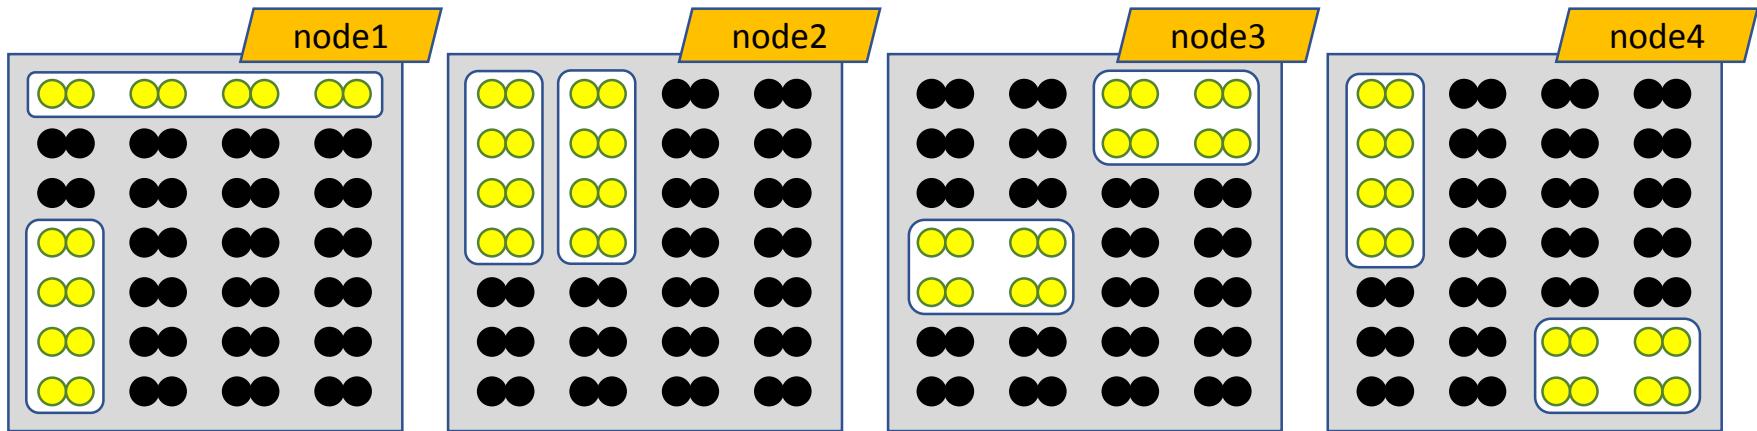
```
sbatch --ntasks=8 --nodes=1 --cpus-per-task=2
```

```
srun --mpi=pmi2 executable -p $SLURM_CPUS_PER_TASK arg arg arg ...
```

```
sinteractive --ntasks=8 --nodes=1 --cpus-per-task=2
```

```
mpirun executable -p $SLURM_CPUS_PER_TASK arg arg arg ...
```

Multi-threaded MPI Job: multiple nodes



```
sbatch --cpus-per-task=8 --ntasks=8 --partition=multinode
```

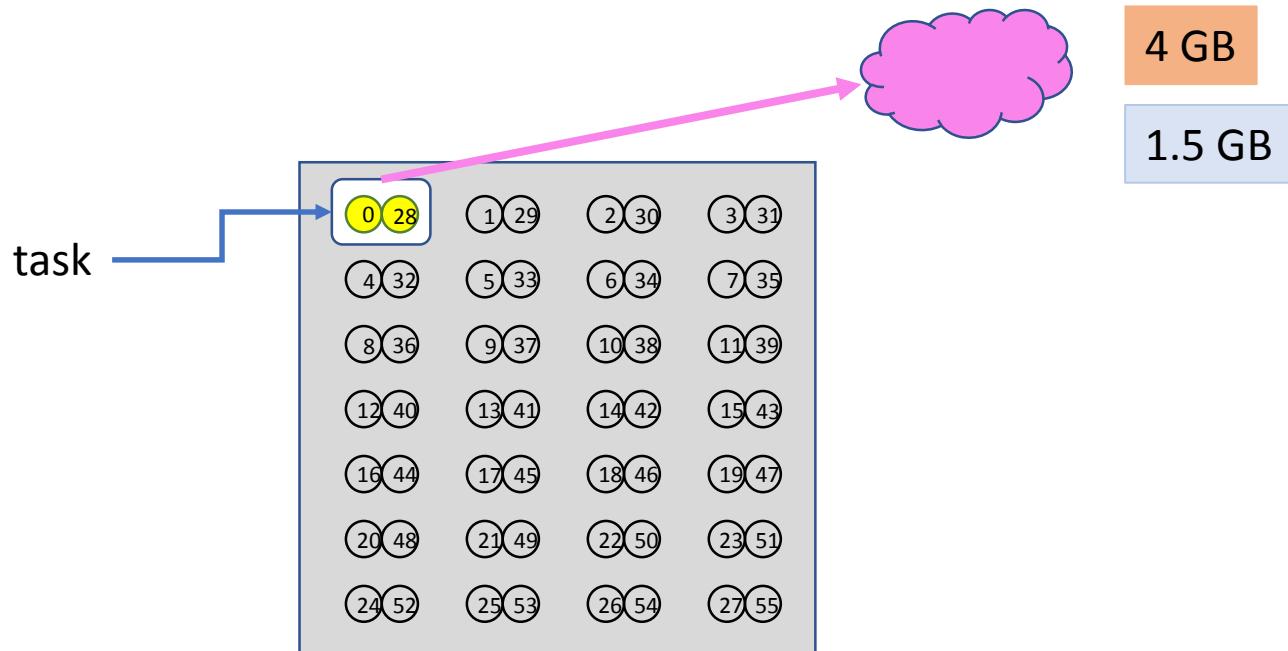
```
srun --mpi=pmi2 executable -p $SLURM_CPUS_PER_TASK arg arg arg ...
```

```
sinteractive --cpus-per-task=8 --ntasks=8 --partition=multinode
```

```
mpirun executable -p $SLURM_CPUS_PER_TASK arg arg arg ...
```

Slurm Resources

Memory: Single-threaded Job



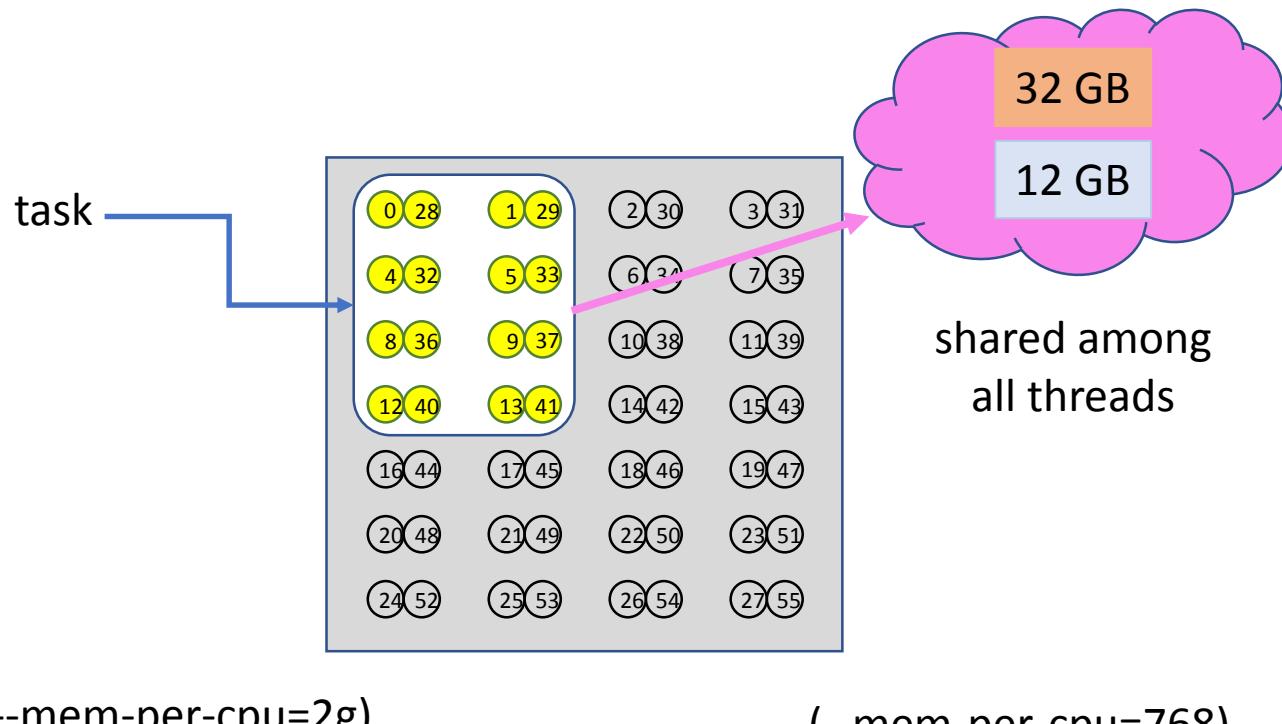
`sbatch`

(`--mem-per-cpu=2g`)

`sinteractive`

(`--mem-per-cpu=768`)

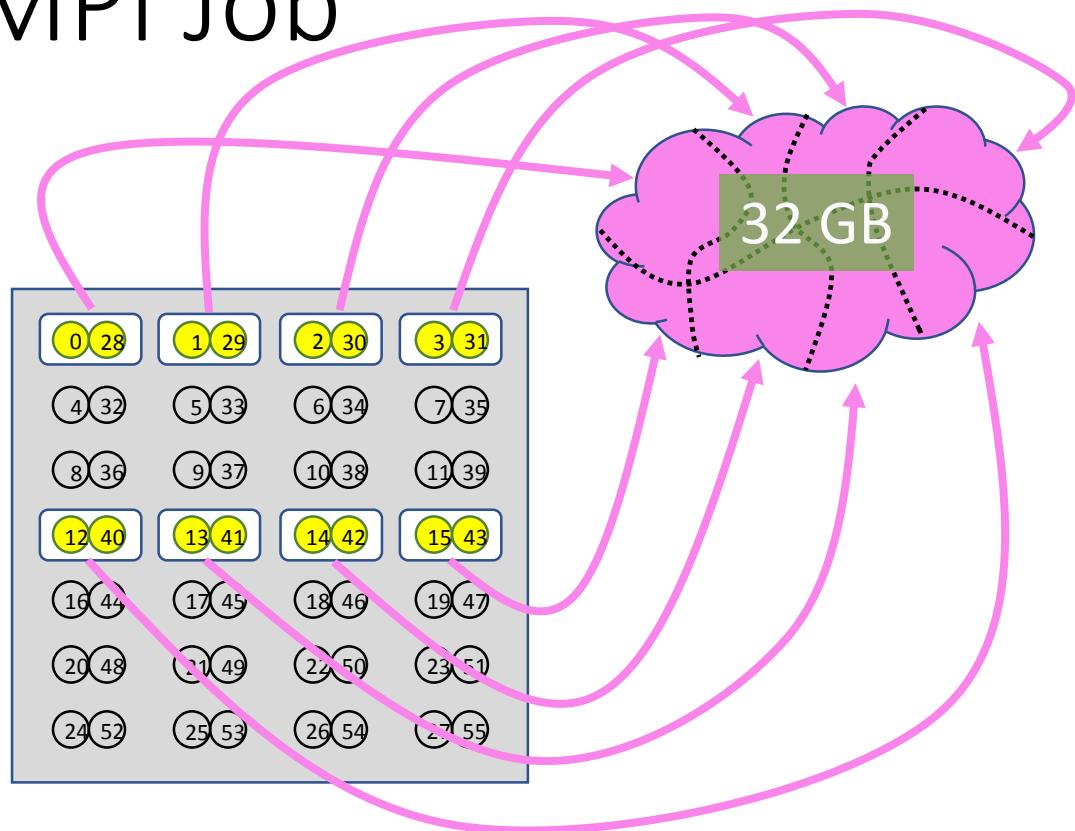
Memory: Multi-threaded Job



```
sbatch --cpus-per-task=16
```

```
sinteractive --cpus-per-task=16
```

Memory: MPI Job



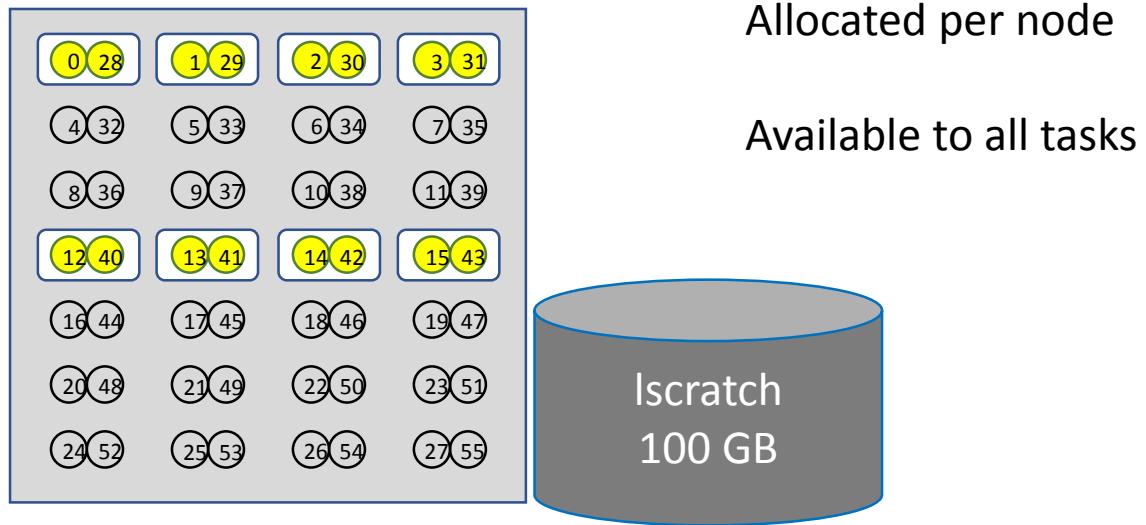
```
sbatch --ntasks=8 --nodes=1 --cpus-per-task=2 --mem-per-cpu=2g
```

\$SLURM_MEM_PER_CPU=2048

```
sbatch --ntasks=8 --nodes=1 --cpus-per-task=2 --mem=32g
```

\$SLURM_MEM_PER_NODE=32768

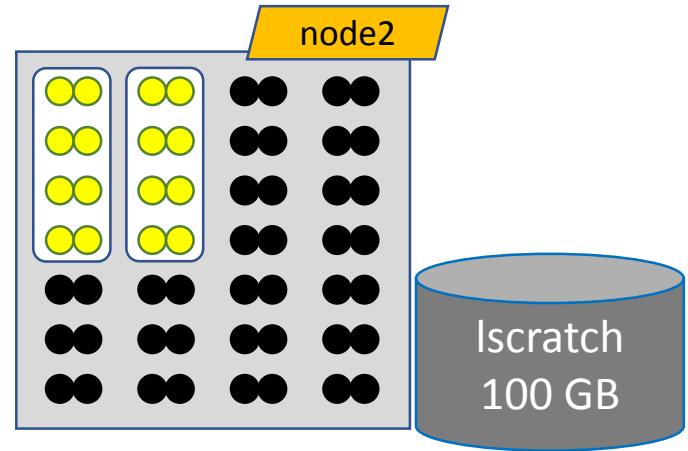
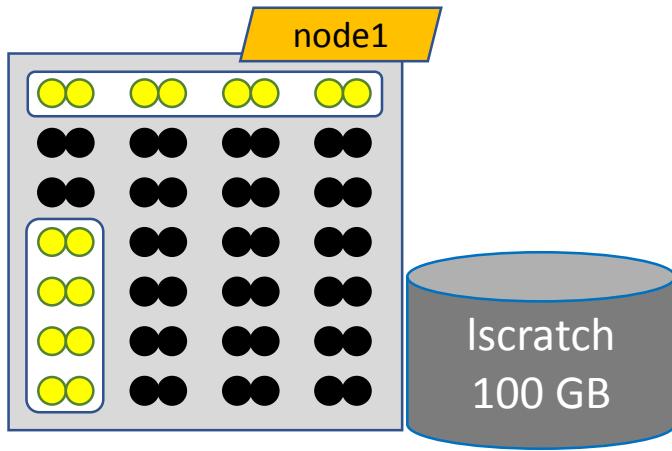
Local Scratch: Single Node



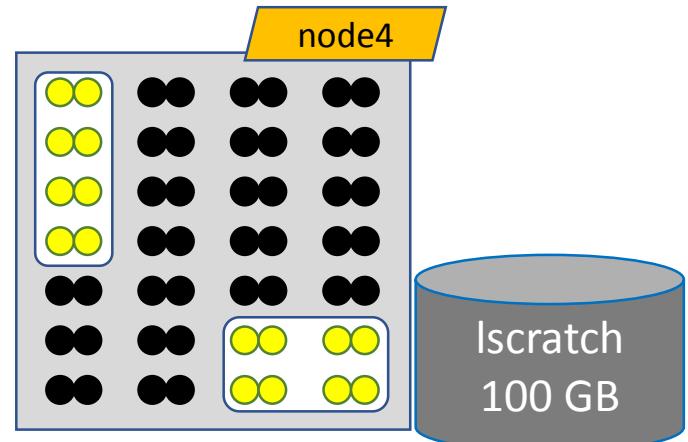
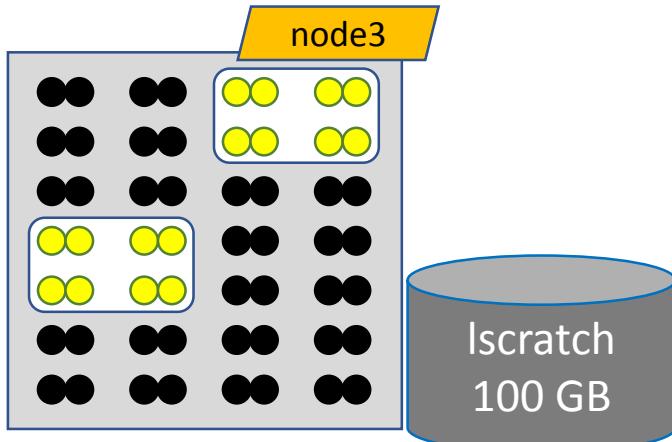
```
sbatch --cpus-per-task=2 --ntasks=8 --gres=lscratch:100
```

`/lscratch/$SLURM_JOB_ID`

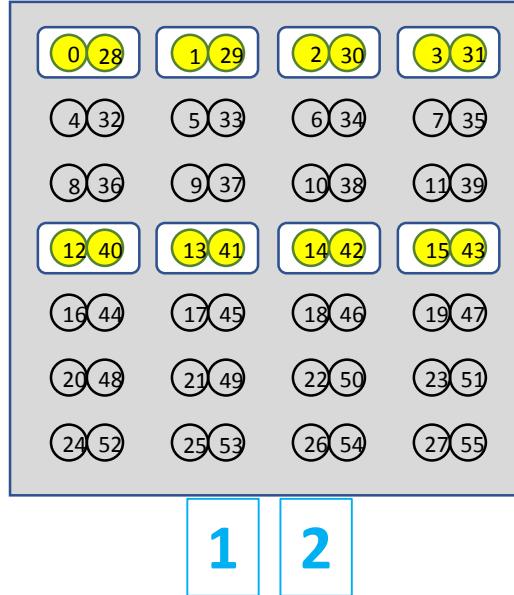
Local Scratch: Multiple Nodes



```
sbatch --ntasks=8 --nodes=4 --cpus-per-task=8 --gres=lscratch:100 --ntasks-per-node=4
```



GPUs



Allocated per node

Accessible to all tasks

```
sbatch --ntasks=8 --nodes=1 --cpus-per-task=2 --gres=gpu:p100:2
```

CUDA_VISIBLE_DEVICES=1,2

GPU Resource Types

- k20x: NVIDIA Kepler 20 (10 GB VRAM, 2x PN)
- k80: Kepler 80 (12 GB VRAM, 4x PN)
- p100: Pascal 100 (16 GB VRAM, 4x PN)

```
--gres=gpu:GPUtype:NGPU
```

Time

- Time is a resource – don't waste it!
- Default depends on partition

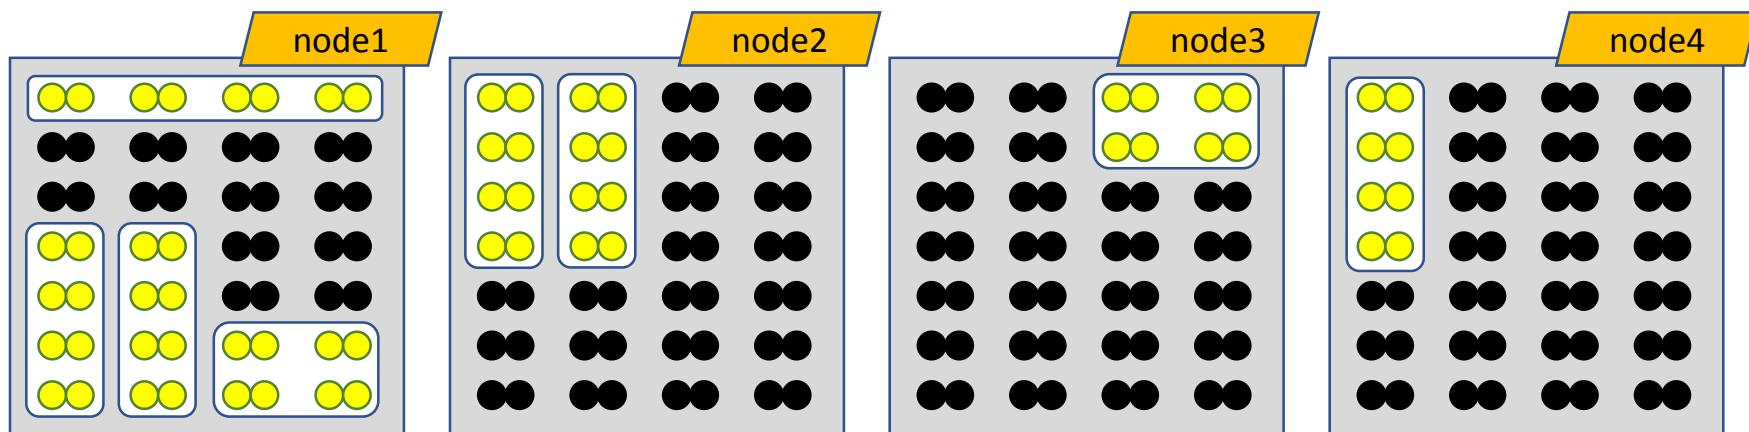
```
sbatch --time=60
```

```
sbatch --time=4:00:00
```

```
sbatch --time=2-00:00:00
```

MPI Job: Distribution

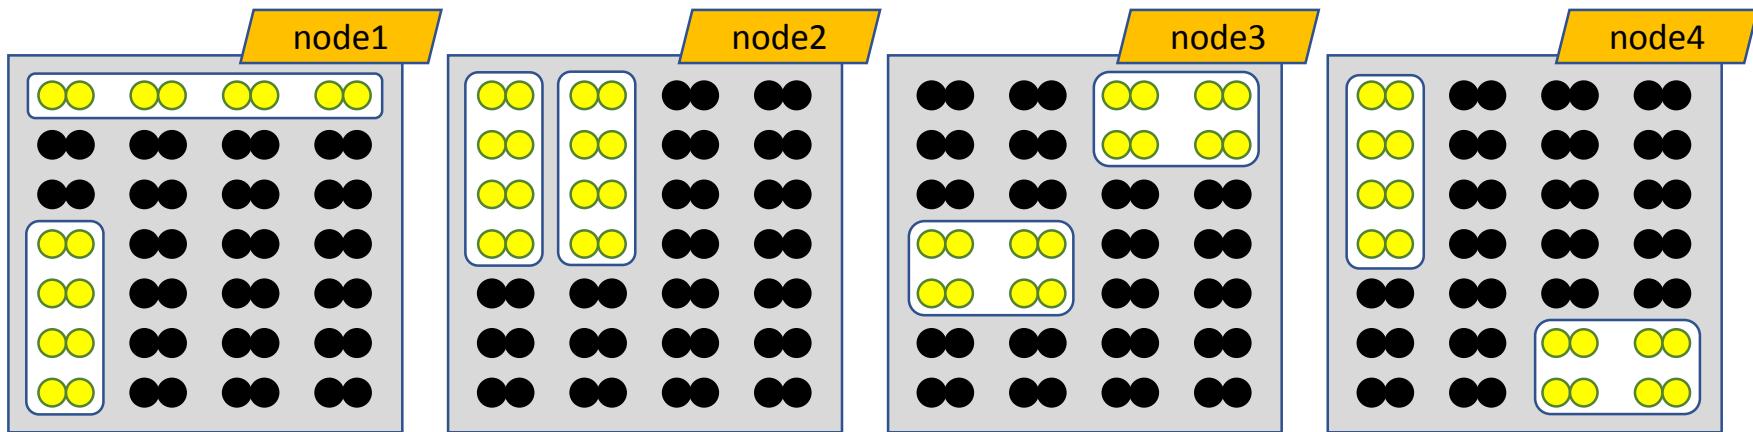
No guarantee of task placement



```
sbatch --ntasks=8 --cpus-per-task=8
```

MPI Job: Distribution

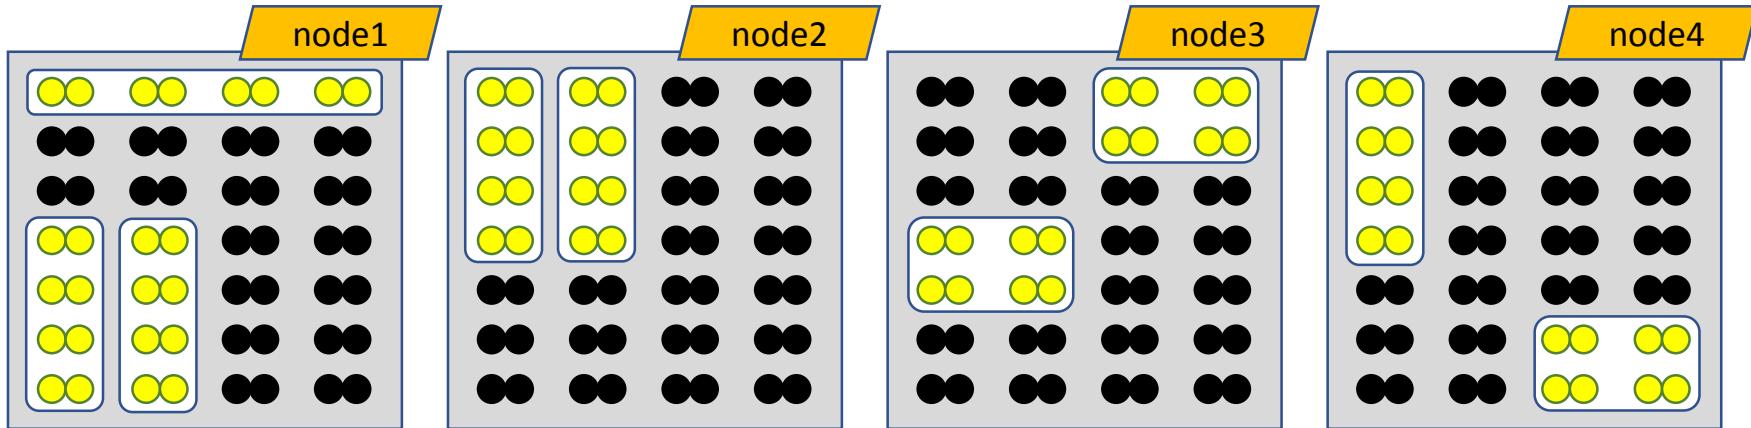
“Egalitarian”, “even” or “flat” distribution



```
sbatch --ntasks=8 --cpus-per-task=8 --ntasks-per-node=2
```

MPI Job: Distribution

arbitrary distribution



```
sbatch --ntasks=8 --cpus-per-task=8 --ntasks-per-node=2  
--distribution=arbitrary
```

file:

node1
node1
node1
node2
node2
node3
node3
node4
node4

```
export SLURM_HOSTFILE=file  
unset SLURM_NTASKS_PER_NODE
```

Other Resources

- Network type (ethernet vs. infiniband)
- Processor type (x2680, x2695, x2650)

RELION Run Methods

There are four options

- --no_parallel_disc_io
- --preread_images
- --scratch_dir=*directory path*
- --dont_combine_weights_via_disc

--no_parallel_disc_io

- Giving this option causes master process to do all the reading and writing
- Don't set this

--preread_images

- Giving this option causes all MPI processes to read particle and image data into memory
- Requires quite a bit of memory

*number of particles * box_size * box_size * 4 / (1024 * 1024 * 1024)*

- About the same speed as local scratch

`--scratch_dir=directory path`

- Giving this option causes a single process per node to copy micrographs to local scratch directory
- `--scratch_dir=/lscratch/$SLURM_JOB_ID`
- Usually takes less time than `--preread_images` to get started
- In general total runtime is about the same

--dont_combine_weights_via_disc

- If this option is NOT given, then after each iteration temporary files will be written to the work directory
- Total waste of time and I/O, unless you are debugging the program
- Always give this option

RELION Job Types

There are 5 Job Types

- Non-MPI, single-threaded, CPU-only
- Flat MPI, single-threaded, CPU-only
- Flat MPI, single-threaded, CPU + GPU (MotionCor2)
- Master-slave MPI, multi-threaded, CPU-only
- Master-slave MPI, multi-threaded, CPU + GPU

Single-threaded, no worries

```
#!/bin/bash
#SBATCH --partition=norm
#SBATCH --mem-per-cpu=8g
#SBATCH --error=Sort/job012/run.err
#SBATCH --output=Sort/job012/run.out
#SBATCH --time=60

module load RELION
relion_particle_sort --i ./Extract/job011/particles.star --ref ./Select/job007/class_averages.star \
--o Sort/job012/particles_sort.star --ctf
```

RELION Parallelization

- Always done with MPI
- A few types use flat, single-threaded MPI (CTFFIND, particle sorting, MotionCor2)
- The rest use both MPI and pthreads (hybrid-parallel)
- Master-Slave method for hybrid-parallel
- Master = MPI rank 0, all others are slaves

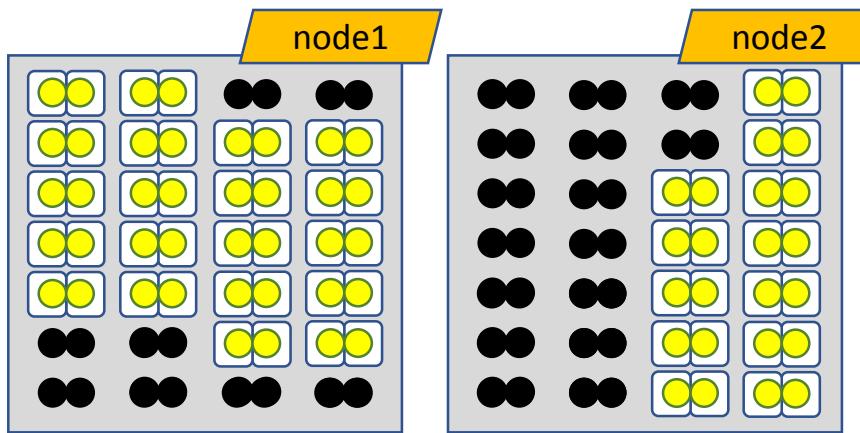
Flat MPI, CPU-only: CTFFIND4

```
#!/bin/bash
#SBATCH --ntasks=64
#SBATCH --partition=multinode
#SBATCH --cpus-per-task=1
#SBATCH --error=CtfFind/job003/run.err
#SBATCH --output=CtfFind/job003/run.out
#SBATCH --time=240
#SBATCH --mem-per-cpu=8g

module load RELION
srun --mem-per-cpu=8g --mpi=pmi2 relion_run_ctffind_mpi \
--i MotionCorr/job002/corrected_micrographs.star --o CtfFind/job003/ \
--CS 2 --HT 300 --AmpCnst 0.1 --XMAG 10000 --DStep 3.54 --Box 512 --ResMin 30 --ResMax 7.1 \
--dFMin 5000 --dFMax 50000 --FStep 500 --dAst 100 --use_noDW \
--ctffind_exe $RELION_CTFFIND_EXECUTABLE --ctfwin -1 --is_ctffind4 --only_do_unfinished
```

Flat, unfixed distribution of 64 tasks with 1 CPU per task, unknown number of nodes, memory allocated per CPU

Flat MPI, CPU-only



Flat, unfixed distribution of 64 tasks with 1 CPU per task, unknown number of nodes, memory allocated per CPU

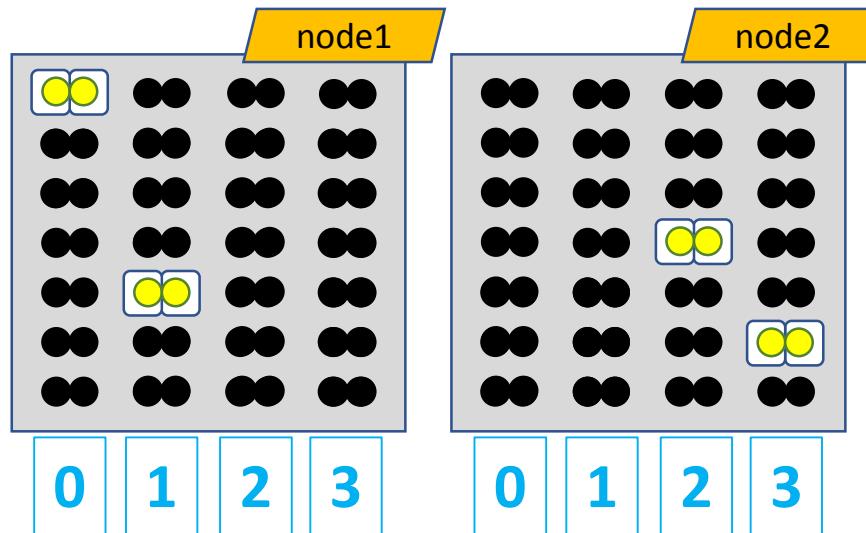
Flat MPI, CPU+GPU: MotionCorr2

```
#!/bin/bash
#SBATCH --ntasks=8
#SBATCH --partition=gpu
#SBATCH --cpus-per-task=1
#SBATCH --error=MotionCorr/job002/run.err
#SBATCH --output=MotionCorr/job002/run.out
#SBATCH --time=240
#SBATCH --mem=10g
#SBATCH --gres=gpu:p100:4
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=4

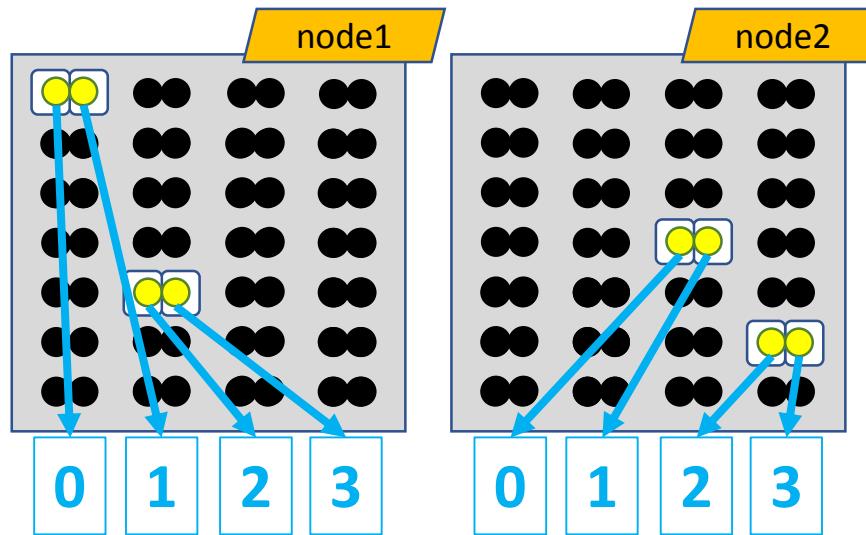
module load RELION
srun --mem-per-cpu=8g --mpi=pmi2 relion_run_motioncorr_mpi --i ./Import/job001/movies.star \
--o MotionCorr/job002/ --save_movies --first_frame_sum 1 --last_frame_sum 16 \
--use_motioncor2 --bin_factor 1 --motioncor2_exe $RELION_MOTIONCOR2_EXECUTABLE \
--bfactor 150 --angpix 3.54 --patch_x 5 --patch_y 5 --gpu "" --dose_weighting \
--voltage 300 --dose_per_frame 1 --preexposure 0 --only_do_unfinished
```

Flat, fixed distribution of 8 tasks on 2 GPU nodes, 4 CPUs and 4 GPUs each node, memory allocated per node, GPUs unassigned

Flat MPI, CPU+GPU: MotionCorr2



Flat MPI, CPU+GPU: MotionCorr2



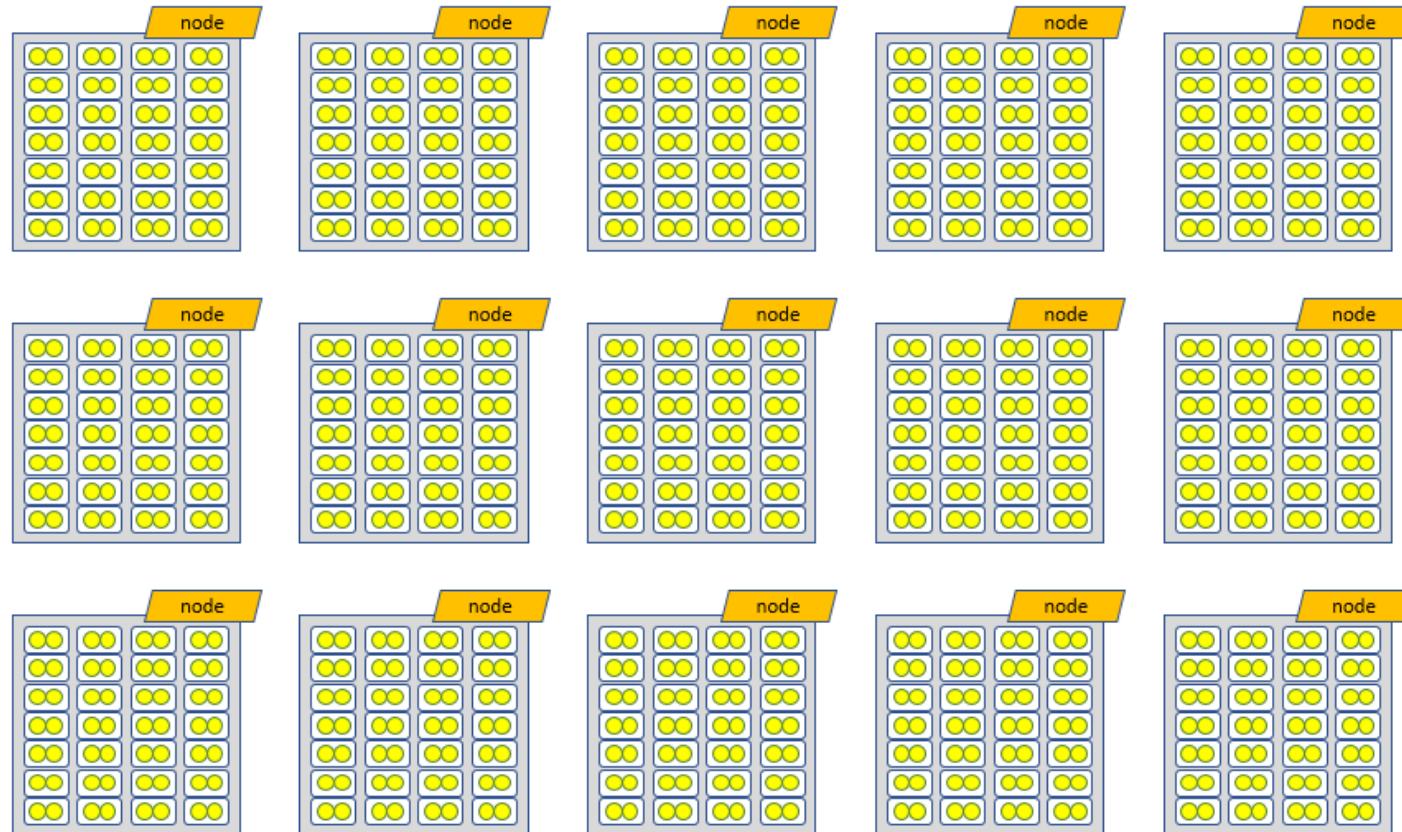
Master-Slave CPU-only: 2D Classification

```
#!/bin/bash
#SBATCH --gres=lscratch:200
#SBATCH --partition=multinode
#SBATCH --ntasks=513
#SBATCH --cpus-per-task=2
#SBATCH --mem-per-cpu=8g
#SBATCH --error=Class2D/job004/run.err
#SBATCH --output=Class2D/job004/run.out
#SBATCH --time=2-00:00:00

source restrict_tcp_for_mpi.sh

srun --mpi=pmi2 relion_refine_mpi --o Class2D/job004/run --i Particles/shiny_2sets.star
--dont_combine_weights_via_disc --scratch_dir=/lscratch/$SLURM_JOB_ID --pool 100 --ctf --iter 25
--tau2_fudge 2 --particle_diameter 360 --k 200 --flatten_solvent --zero_mask --oversampling 1
--psi_step 6 --offset_range 5 --offset_step 2 --norm --scale --j 2
```

Master-Slave CPU-only: 2D Classification



This is a BIG job... likely about 50 nodes

Master-Slave CPU-only: 2D Classification – all in memory

```
#!/bin/bash
#SBATCH --partition=multinode
#SBATCH --ntasks=64
#SBATCH --cpus-per-task=16
#SBATCH --mem-per-cpu=15g
#SBATCH --ntasks-per-node=1
#SBATCH --error=Class2D/job011/run.err
#SBATCH --output=Class2D/job011/run.out
#SBATCH --time=2-00:00:00

source restrict_tcp_for_mpi.sh

srun --mpi=pmi2 relion_refine_mpi --o Class2D/job011/run --i Particles/shiny_2sets.star
--dont_combine_weights_via_disc --preread_images --pool 100 --ctf --iter 25 --tau2_fudge 2
--particle_diameter 360 --K 200 --flatten_solvent --zero_mask --oversampling 1 --psi_step 6
--offset_range 5 --offset_step 2 --norm --scale --j 16
```

Memory is shared among threads

Time required to read into memory can be substantial (hours?)

Not always true that --preread_images is faster than --scratch_dir

Tidbits

- RELION does not scale much beyond 2000 CPUs
- It takes a LONG time to read images into memory of every single MPI rank
- 50GB images took ~60 minutes to read into memory for ~500 ranks

Understanding GPU Jobs

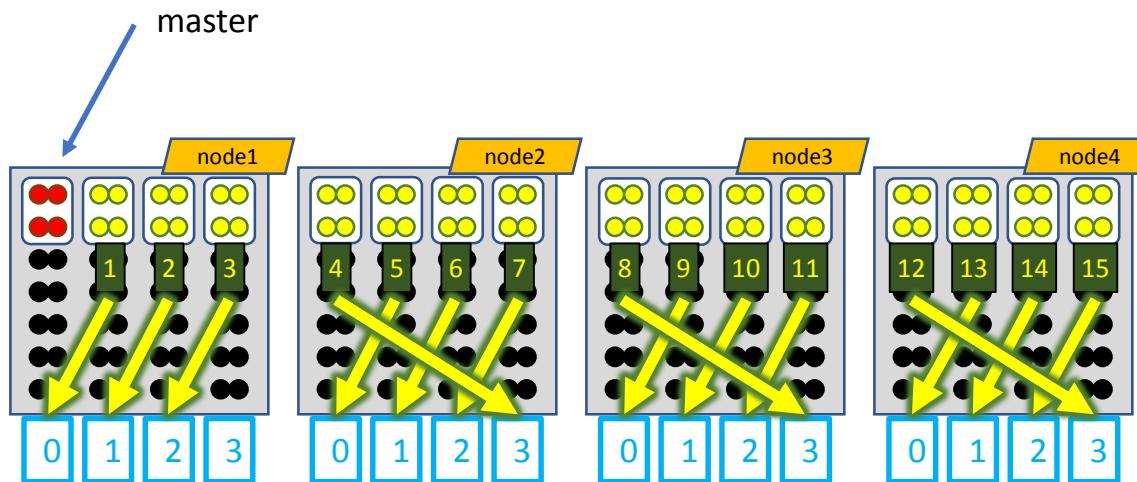
Understanding GPU Jobs

- GPUs accelerate particular execution steps 80x-100x
- Master – Slave model of parallelization conflicts with Slurm resource allocation and MPI rank distribution for GPU nodes
- GPU devices mapped to slaves on the basis of MPI rank id and device id

MPI rank - GPU Mapping

- Default --gpu ""
- GPUs are mapped to MPI ranks in order across ALL MPI ranks, not just those on a node
- You could set --gpu manually, but good luck, stockpile aspirin
- --gpu 0,0,0,0:1,1,1,1:2,2,2,2:3,3,3,3

Class2D, Class3D, Refine – Problems



```
sbatch  
--cpus-per-task=4  
--ntasks=16  
--ntasks-per-node=4  
--mem=240g  
--partition=gpu  
--gres=gpu:k80:4,1scratch:100  
--constraint=gpuK80
```

Why waste GPU when it is worth 80-100 CPUs?

Class2D, Class3D, Refine – Problems

```
sbatch  
--cpus-per-task=4  
--ntasks=17  
--ntasks-per-node=4
```

not possible due to GPU limit (16 per job)

```
sbatch  
--cpus-per-task=4  
--ntasks=17  
--nodes=4
```

random distribution of tasks on nodes

```
sbatch  
--cpus-per-task=4  
--ntasks=17  
--nodes=4  
--distribution=plane=1
```

complete overutilization of GPUs

```
sbatch  
--cpus-per-task=4  
--ntasks=17  
--nodes=4  
--mincpus=16
```

unpredictable which node receives 5 tasks

Arbitrary Distribution

- Over allocate on purpose

```
sbatch  
  --cpus-per-task=4  
  --ntasks=20  
  --nodes=4  
  --ntasks-per-node=5  
  --distribution=arbitrary
```

`$SLURM_HOSTFILE`

```
cn4000  
cn4000  
cn4000  
cn4000  
cn4000  
cn4001  
cn4001  
cn4001  
cn4001  
cn4002  
cn4002  
cn4002  
cn4002  
cn4003  
cn4003  
cn4003  
cn4003
```

- Create `$SLURM_HOSTFILE` and unset `$SLURM_NTASKS_PER_NODE`

```
array=( $( scontrol show hostname $SLURM_JOB_NODELIST ) )  
file=$(mktemp)  
echo ${array[0]} > $file  
for ((i=0;i<${SLURM_JOB_NUM_NODES};i++)); do  
    for ((j=0;j<${(SLURM_NTASKS_PER_NODE-1)};j++)); do  
        echo ${array[$i]} >> $file  
    done  
done  
export SLURM_HOSTFILE=$file  
unset SLURM_NTASKS_PER_NODE
```

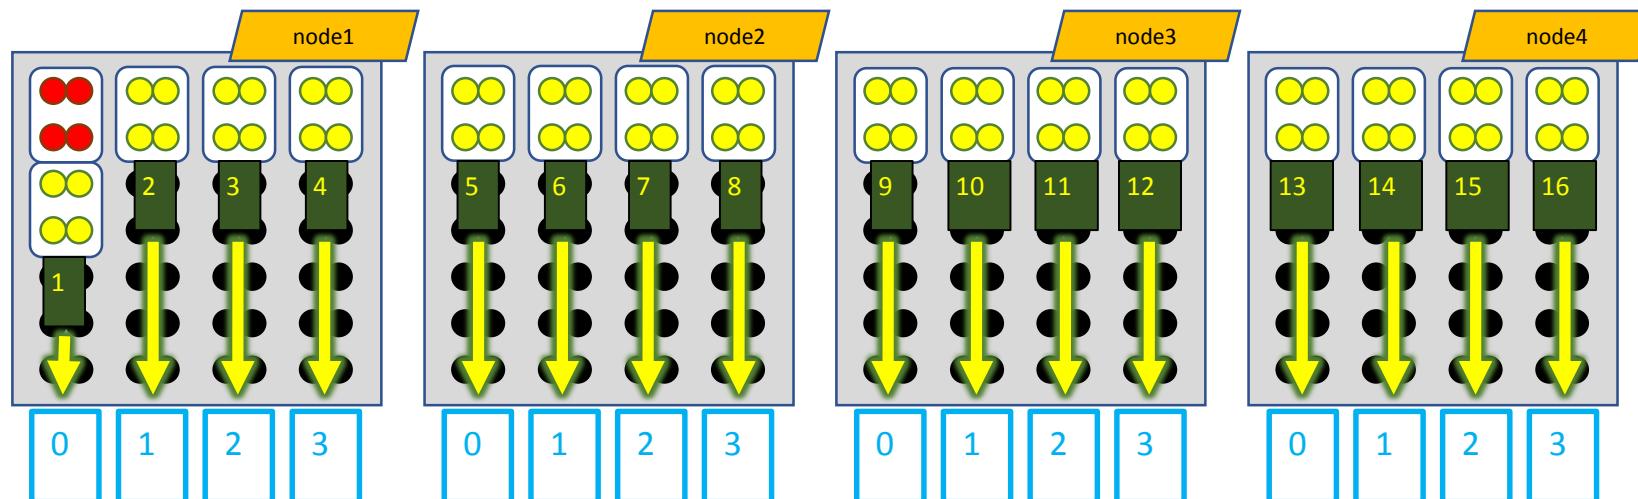
RELION Utility Script

```
#!/bin/bash
#SBATCH --partition=gpu
#SBATCH --ntasks=20
#SBATCH --cpus-per-task=4
#SBATCH --mem-per-cpu=30g
#SBATCH --ntasks-per-node=5
#SBATCH --nodes=4
#SBATCH --gres=gpu:p100:4,lscratch:200
#SBATCH --distribution=arbitrary
#SBATCH --error=Class2D/job011/run.err
#SBATCH --output=Class2D/job011/run.out
#SBATCH --time=2-00:00:00

source restrict_tcp_for_mpi.sh
source add_extra_MPI_task.sh

srun --mpi=pmi2 relion_refine_mpi --o Class2D/job011/run --i Particles/shiny_2sets.star
--dont_combine_weights_via_disc -scratch_dir=/lscratch/$SLURM_JOB_ID --pool 100
--ctf --iter 25 --tau2_fudge 2 --particle_diameter 360 --K 200 --flatten_solvent --zero_mask
--oversampling 1 --psi_step 6 --offset_range 5 --offset_step 2 --norm --scale --j 4 --gpu ""
```

Class2D, Class3D, Refine



Tidbits

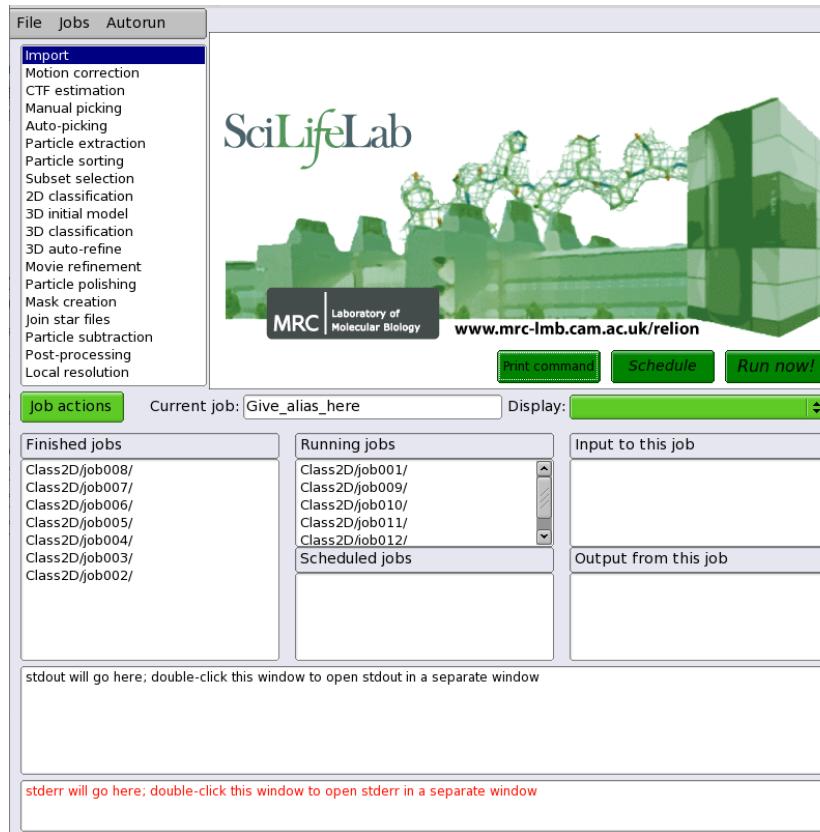
- Oversubscription of GPUs generally NOT a good idea
- Maintain 1 MPI rank per GPU, but 4 or 8 threads per GPU not always bad
- The more GPUs, the faster (to a point...)
- Run like Goldilocks

GPUs are not constantly used, and
all steps not parallelizable

1. Noise estimation, parallelized over MPI-ranks, not threads
2. Expectation setup, parallelized over MPI-ranks, not threads
3. Expectation, parallelized over MPI-ranks and threads,
GPU-accelerated
4. Maximization, parallelized over both MPI-ranks and
threads
5. Estimate orientational accuracy, not parallelized
6. Next iteration, repeat from step 2

Maximize MPI, not threads!

Using the GUI



Don't run GUI on the login node

- Some job types run quickly
- You will forget you are on the login node

Be aware of options passed to sinteractive

- Default --mem-per-cpu=768
SLURM_MEM_PER_CPU=768
- Subsequent batch MPI jobs may get confused

```
#!/bin/bash
#SBATCH --constraint=gpuK80
#SBATCH --ntasks=16
#SBATCH --nodes=4
#SBATCH --cpus-per-task=4
#SBATCH --mem=240g
#SBATCH --partition=XXXqueueXXX
#SBATCH --gres=gpu:k80:4,1scratch:xxxextra1xxx
#SBATCH --error=XXXerrfilexxx
#SBATCH --output=XXXoutfilexxx
#SBATCH --time=XXXextra2xxx

unset SLURM_MEM_PER_CPU
srun --mpi=pmi2 relion_refine_mpi --i Particles
```

command line > environment variable > within batch script

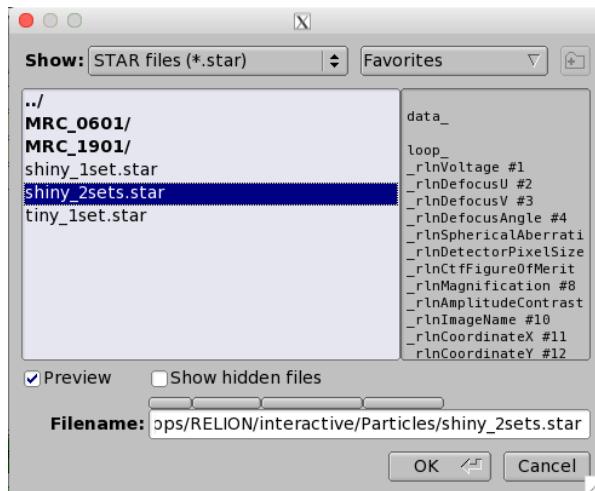
```
SLURM_MEM_PER_CPU=768
SLURM_MEM_PER_NODE=245760
```

```
srun: Force Terminated job step 58421330.0
srun: error: cn2692: task 0: Killed
srun: error: cn2692: task 1: Killed
srun: Force Terminated job step 58421330.1
```

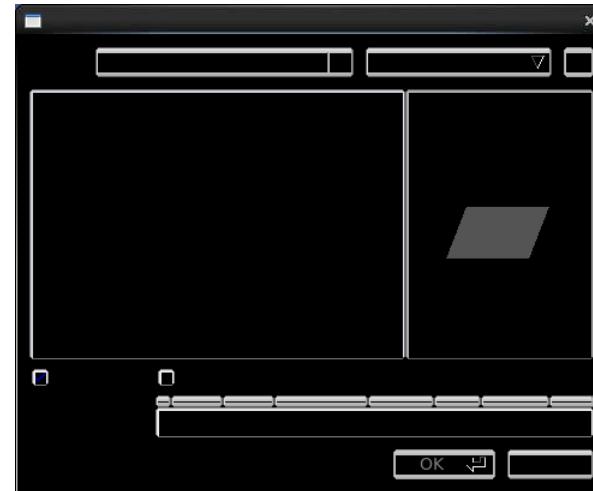
Always use --mem-per-cpu, unless you know what you are doing!

NX messes up file browsing

Xquartz:

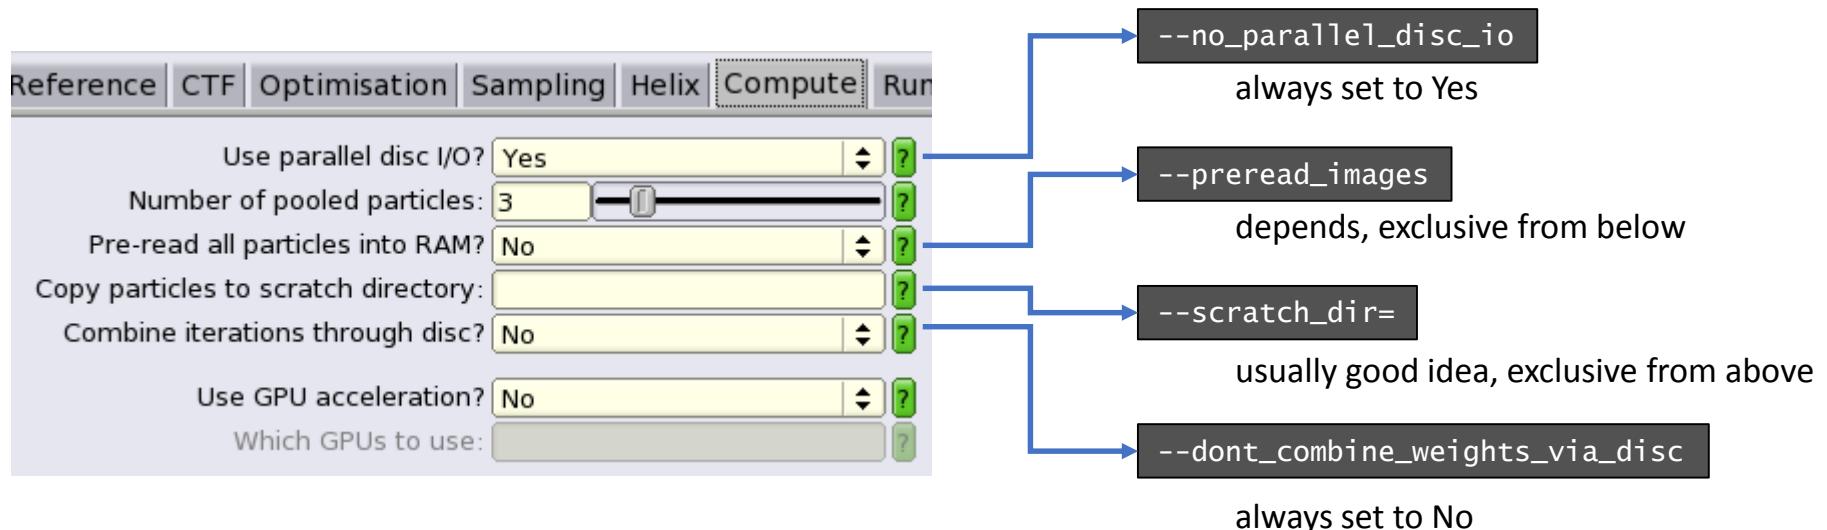


NX:



RELION Compute Options

- Four major choices to be made:



```
relion_command ... --dont_combine_weights_via_disc --scratch_dir=/lscratch/$SLURM_JOB_ID ...
```

or

```
relion_command ... --dont_combine_weights_via_disc --preread_images ...
```

Choice

- Fastest execution (under ideal conditions) is to NOT use --scratch_dir or --preread_images
- OK for short jobs, but long (>2h) and large jobs will have I/O hassles
- --preread_images is perhaps a little faster
- --scratch_dir most reliable

RELION/2.1.0-new

- GUI code hacked to be more Slurm friendly
- Allows a single batch script template
- Will be the default under CentOS7

Running/Batch System Options

Optimisation | Sampling | Helix | Compute | Running

Number of MPI procs:	20	?
Number of threads:	4	?
Submit to queue?	Yes	?
Queue name:	gpu	?
Queue submit command:	sbatch	?
Walltime	1-00:00:00	?
Memory Per Thread	8g	?
Gres	lscratch:200,gpu:p100:4	?
Addl. (ex4) SBATCH Directives	--N 4 --ntasks-per-node=5	?
Addl. (ex5) SBATCH Directives	--distribution=arbitrary	?
Addl. (ex6) SBATCH Directives	--job-name=Class2D	?
Standard submission script:	template_scripts/common.sh	?
Additional arguments:		
	Print command	Schedule
	Run now!	

1 = non-MPI
>1 = MPI

must match job requirements
(norm, multinode, gpu)

--time must be set
--mem-per-cpu must be set
--gres is optional

Room for other sbatch options

/usr/local/apps/RELION/batch_template_scripts/common.sh

Running/Batch System Options

F | Optimisation | Sampling | Helix | Compute | Running

Number of MPI procs:	1025	[?]	XXXmpinodesXXX, --ntasks
Number of threads:	2	[?]	XXXthreadsXXX, --cpus-per-task
Submit to queue?	Yes	[?]	XXXqueueXXX, --partition
Queue name:	multinode	[?]	XXXextra1XXX, --time=
Queue submit command:	sbatch	[?]	
Walltime	8-00:00:00	[?]	XXXextra2XXX, --mem-per-cpu=
Memory Per Thread	4g	[?]	
Gres	lscratch:200	[?]	XXXextra3XXX, --gres=
Addl. (ex4) SBATCH Directives	--constraint=x2695	[?]	XXXextra4XXX (optional)
Addl. (ex5) SBATCH Directives	--mail-type=BEGIN,END	[?]	
Addl. (ex6) SBATCH Directives	--job-name=Refinement	[?]	XXXextra5XXX (optional)
Standard submission script:	template_scripts/common.sh	[?]	
Additional arguments:		[?]	XXXextra6XXX (optional)

Print command **Schedule** **Run now!**

/usr/local/apps/RELION/batch_template_scripts/common.sh

common.sh batch script template

```
#!/bin/bash
#SBATCH --ntasks=XXXmpinodesXXX
#SBATCH --partition=XXXqueueXXX
#SBATCH --cpus-per-task=xxxthreadsXXX
#SBATCH --error=XXXerrfileXXX
#SBATCH --output=XXXoutfileXXX
#SBATCH --time=XXXextra1XXX
#SBATCH --mem-per-cpu=XXXextra2XXX
#SBATCH --gres=XXXextra3XXX
#SBATCH XXXextra4XXX
#SBATCH XXXextra5XXX
#SBATCH XXXextra6XXX

source restrict_tcp_for_mpi.sh
source add_extra_MPI_task.sh

srun --mem-per-cpu=XXXextra2XXX --mpi=pmi2 XXXcommandXXX
```

Memory

- Minimum amount of memory (per MPI rank) for --preread_images =

$$\text{number of particles} * \text{box_size} * \text{box_size} * 4 / (1024 * 1024 * 1024)$$

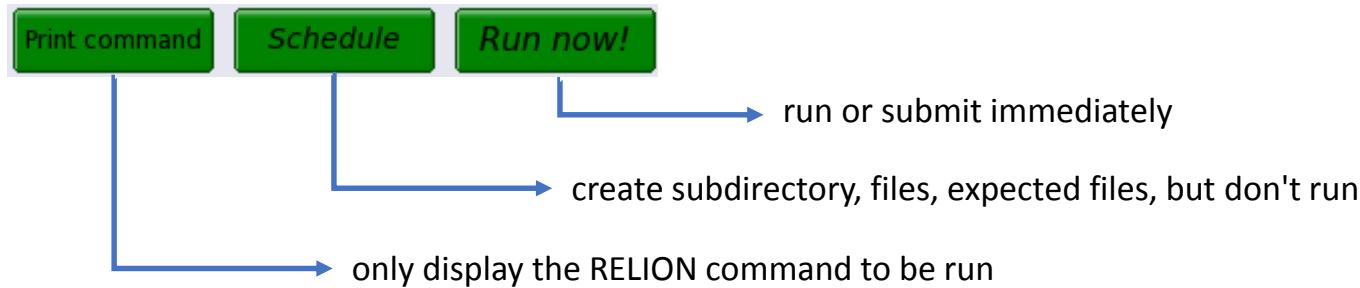
- E.g. 100K particles, 360 pixel box will need at least 48GB per task
- Memory use will grow 10-20% over time
- Convert to --mem-per-cpu = [*calc RAM required*]/--cpus-per-task
- Same example above with -j 4 (--cpus-per-task=4) needs 14g per cpu

"Dumb down" the network interface

```
export OMPI_MCA_bt1="self,sm,tcp"
export OMPI_MCA_bt1_tcp_if_include="10.2.0.0/18"
```

- No difference in execution time between Infiniband and plain 'ol TCP
- MPI RELION can get confused on management network interfaces (10.90)
- Will give warnings when a network interface doesn't exist
- Automated with **restrict_tcp_for_mpi.sh** script

Check twice, run once



RELION Job Monitoring

- run.out and run.err files
- jobhist, jobload, jobdata
- dashboard
- nvidia-smi

Dashboard

The screenshot shows the HPC @ NIH User Dashboard. At the top, there's a navigation bar with links for Status, Applications, Reference Data, Storage, User Guides, **User Dashboard** (which is circled in red), How To, and About. Below the navigation is a search bar and social media sharing icons. A large "TOP 500 66" logo is on the left. The main content area contains a paragraph about the NIH HPC group and a "Quick Links" sidebar with links to System Status, How To, and Application/DB updates. There's also a section for Recent Papers that used Biowulf & HPC Resources.

Jobs remain listed on the dashboard for 10 days after ending

Search by jobid, jobname, timestamp, state

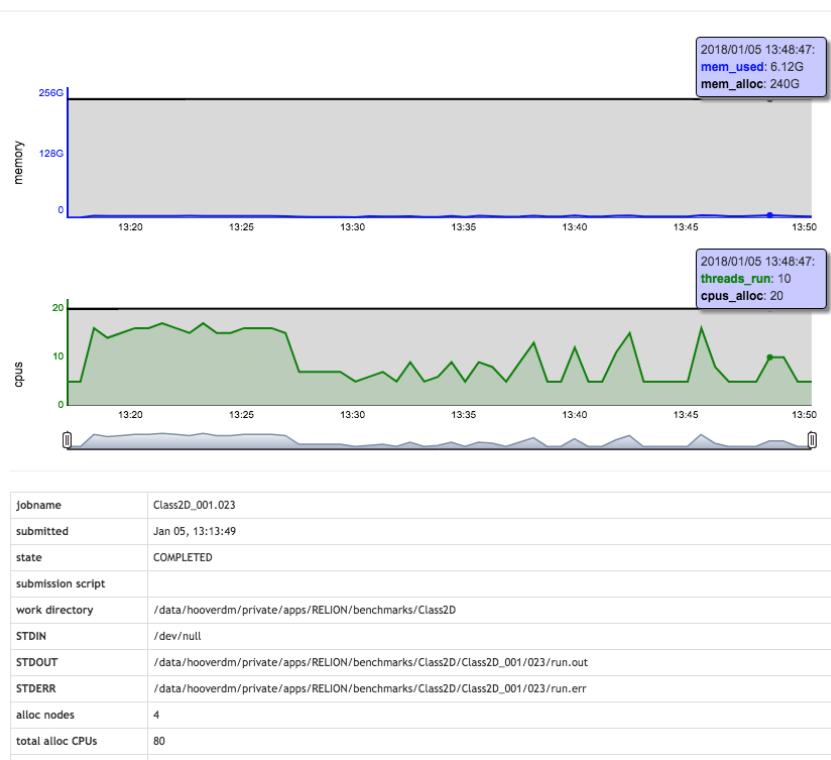
User Dashboard

last page refresh: 2018-01-11 12:30:10 PM

User Dashboard						
Accounts		Disk Usage		Job Info		
Show 10 entries		Search: <input type="text"/>				
jobid	jobname	submit	start	end	state	
57881627	Class2D_001.030	Jan 05, 13:44:12	Jan 05, 13:45:43	Jan 05, 14:10:41	COMPLE...	
57881621	Class2D_001.029	Jan 05, 13:44:02	Jan 05, 13:44:03	Jan 05, 14:08:28	COMPLE...	
57881620	Class2D_001.028	Jan 05, 13:43:50	Jan 05, 13:43:51	Jan 05, 14:09:09	COMPLE...	
57881612	Class2D_001.027	Jan 05, 13:43:33	Jan 05, 13:43:35	Jan 05, 14:12:13	COMPLE...	
57881610	Class2D_001.026	Jan 05, 13:43:19	Jan 05, 13:43:27	Jan 05, 14:12:23	COMPLE...	
57881608	Class2D_001.025	Jan 05, 13:43:09	Jan 05, 13:43:27	Jan 05, 14:12:35	COMPLE...	
57879127	Class2D_001.024	Jan 05, 13:13:54	Jan 05, 13:21:32	Jan 05, 13:46:30	COMPLE...	
57879126	Class2D_001.023	Jan 05, 13:13:49	Jan 05, 13:16:32	Jan 05, 13:50:04	COMPLE...	
57879125	Class2D_001.022	Jan 05, 13:13:43	Jan 05, 13:14:40	Jan 05, 13:39:03	COMPLE...	
57879119	Class2D_001.021	Jan 05, 13:13:37	Jan 05, 13:14:32	Jan 05, 13:38:43	COMPLE...	

Dashboard

BioWulf Job 57879126



nodelist	cn4186,cn4187,cn4188,cn4189
ended	Jan 05, 13:50:04
CPU hours	44.71113
alloc mem per cpu	12.0 GB
max mem per cpu	296 MB
exit code	0
script	<pre>#!/bin/bash mkdir /lscratch/\${SLURM_JOB_ID}/TMPDIR export TMPDIR=/lscratch/\${SLURM_JOB_ID}/TMPDIR export OMPI_MCA_btl="self,sm,tcp" export OMPI_MCA_btl_tcp_if_include="10.2.0.0/18" echo SLURM_JOB_ID=\${SLURM_JOB_ID} echo RELION=2.1.0 echo NT=20 echo CPT=4 echo TPN=5 echo SLURM_JOB_NUM_NODES=\${SLURM_JOB_NUM_NODES} echo GPN=4 echo date=\$(date) array=(\$(control show hostname \${SLURM_JOB_NODELIST})) file=\$(mktemp) echo \${array[0]} > \$file for ((i=0;i<\${SLURM_JOB_NUM_NODES};i++)); do for ((j=0;j<\${SLURM_NTASKS_PER_NODE}-1;j++)); do echo \${array[\$i]} >> \$file done done export SLURM_HOSTFILE=\$file unset SLURM_NTASKS_PER_NODE srun --mpi=pmi2 relion_refine_mpi --o Class2D_001/023/run --i Particles/tiny_iset.star --pool 100 --ctf --iter 10 --tau2_fudge 2 --particle_diameter 360 --K 200 --f latten_solvent --zero_mask --oversampling 1 --psi_step 6 --offset_range 5 --offset_step 2 --norm --scale --j 4 --no_parallel_disc_io --scratch_dir /lscratch/\${SLURM_JOB_ID} --random_seed 0 --gpu for i in \$(ls Class2D_001/023/*.{star,mrcs}) ; do hollow_file.sh \$i ; done echo date=\$(date)</pre>

nvidia-smi

- While job is running, display snapshot of GPU utilization

```
$ jobload -j 11111111
      JOBID      TIME      NODES   CPUS   THREADS   LOAD      MEMORY
                     Elapsed / Wall          Alloc   Active
 11111111    00:52:53 / 1-00:00:00 cn4235     20      11    55%    4.0 / 240.0 GB
               00:52:53 / 1-00:00:00 cn4236     16      15    94%    3.7 / 240.0 GB
               00:52:53 / 1-00:00:00 cn4237     16      16   100%    3.6 / 240.0 GB
               00:52:53 / 1-00:00:00 cn4238     16      14    88%    3.6 / 240.0 GB
Nodes: 4 CPUs: 68 Load Avg: 84%
$ ssh cn4236 nvidia-smi
Thu Jan 11 12:39:16 2018
+-----+
| NVIDIA-SMI 375.26           Driver Version: 375.26 |
+-----+
| GPU  Name      Persistence-MI Bus-Id      Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap| Memory-Usage | GPU-Util  Compute M. |
+-----+
|  0  Tesla K80      On  0000:83:00.0    Off |                  Off |
| N/A  78C   P0  137W / 149W | 11481MiB / 12205MiB | 100%     Default |
+-----+
|  1  Tesla K80      On  0000:84:00.0    Off |                  Off |
| N/A  55C   P0  148W / 149W | 11481MiB / 12205MiB | 86%     Default |
+-----+
|  2  Tesla K80      On  0000:8A:00.0    Off |                  Off |
| N/A  79C   P0  120W / 149W | 11481MiB / 12205MiB | 89%     Default |
+-----+
|  3  Tesla K80      On  0000:8B:00.0    Off |                  Off |
| N/A  65C   P0  148W / 149W | 11481MiB / 12205MiB | 90%     Default |
+-----+
+-----+
| Processes:                               GPU Memory |
| GPU  PID  Type  Process name             Usage    |
+-----+
|  0   37242  C  .../local/relion_2.1.0/bin/relion_refine_mpi 11477MiB |
|  1   37243  C  .../local/relion_2.1.0/bin/relion_refine_mpi 11477MiB |
|  2   37240  C  .../local/relion_2.1.0/bin/relion_refine_mpi 11477MiB |
|  3   37241  C  .../local/relion_2.1.0/bin/relion_refine_mpi 11477MiB |
+-----+
$
```

nvidia-smi dmon -d 10

- Continuously monitor GPU utilization for defined periods of time
- dmon is highly configurable

```
$ ssh cn4236 nvidia-smi dmon -d 10
# gpu  pwr   temp    sm     mem    enc    dec    mclk   pcclk
# Idx   W     C      %      %      %      %      MHz    MHz
  0    146   77    100     4      0      0      2505   810
  1    143   54     15     1      0      0      2505   797
  2    70    79      6     0      0      0      2505   810
  3    147   65     12     1      0      0      2505   875
  0    143   77    100     1      0      0      2505   810
  1    148   55    100     1      0      0      2505   810
  2    144   80    100     2      0      0      2505   823
  3    125   66     44     2      0      0      2505   797
  0    146   78     64     4      0      0      2505   810
  1    148   55    100     1      0      0      2505   810
  2    142   80    100     2      0      0      2505   810
  3    87    66     81     4      0      0      2505   875
  0    87    78    100     1      0      0      2505   810
  1    87    55    100     4      0      0      2505   810
  2    141   80    100     2      0      0      2505   810
  3    145   67    100     4      0      0      2505   797
  0    145   78    100     4      0      0      2505   810
  1    148   56    100     2      0      0      2505   810
  2    141   80    100     2      0      0      2505   810
  3    147   67    100     4      0      0      2505   797
^C$ ^C
```

Troubleshooting

- RELION touches every weak and sensitive point in the cluster
- Many different ways things can go south

Insufficient Memory

- Insufficient memory causes individual MPI ranks to be killed, leading to "zombie" RELION jobs
- run.err shows something like this:

```
srun: error: cn3068: task 3: killed
```

- Job continues to run, but run.out does not progress
- Other causes for rank failure include network hiccups, quota overrun, generic hardware failure

Insufficient GPU Memory

- k80 GPUs have access to 12 GB VRAM
- p100 GPUs have 16 GB VRAM

WARNING

with the current settings and hardware, you will be able to use an estimated image-size of 533 pixels during the last iteration...

...but your input box-size (image_size) is however 659. This means that you will likely run out of memory on the GPU(s), and will have to then re-start from the last completed iteration (i.e. continue from it) *without* the use of GPUs.

You are also splitting at least one GPU between two or more mpi-slaves, which might be the limiting factor, since each mpi-slave that shares a GPU increases the use of memory. In this case we recommend running a single mpi-slave per GPU, which will still yield good performance and possibly a more stable execution.

Mixing RELION Versions

```
PipeLine::read: cannot find name or type in pipeline_nodes table
```

- The GUI creates hidden files for maintaining state between sessions under .Nodes
- Running a different version of RELION on top of .Nodes may cause corruption

```
$ tree .Nodes
.Nodes
  0
    Import
      job001
        movies.star
    MotionCorr
      job002
        corrected_micrograph_movies.star
  1
    CtfFind
      job003
        micrographs_ctf.star
    ManualPick
      job004
        micrographs_selected.star
    MotionCorr
      job002
        corrected_micrographs.star
  10
    Refine3D
      job012
        run_half1_class001_unfil.mrc
  11
    LocalRes
      first3dref
        relion_locres_filtered.mrc
    PostProcess
      first3def
        postprocess_masked.mrc
        postprocess.mrc
  12
    LocalRes
      first3dref
        relion_locres.mrc
  13
    MotionCorr
      job002
```

Inappropriate RELION Options

- Almost always diagnosable from .out and .err files
- Contact RELION developers or CCP4EM mailing list if warning is unclear
- Mixing --preread_images and --scratch_dir
- Too many classes (--K)
- --write_fom_maps from within AutoPick

I/O

- Chronic problem nowadays
- Certain steps (copying to /lscratch, reading into memory) take inordinately long
- Don't launch 1000's of nodes simultaneously – stagger
- Use GPUs whenever possible
- ALWAYS SET --dont_combine_weights_via_disc

Bad Input

- Filter and prune your data!
- Choose only the best images, particles, and classes
- More data does NOT mean better results

Questions? Comments?

staff@hpc.nih.gov

