

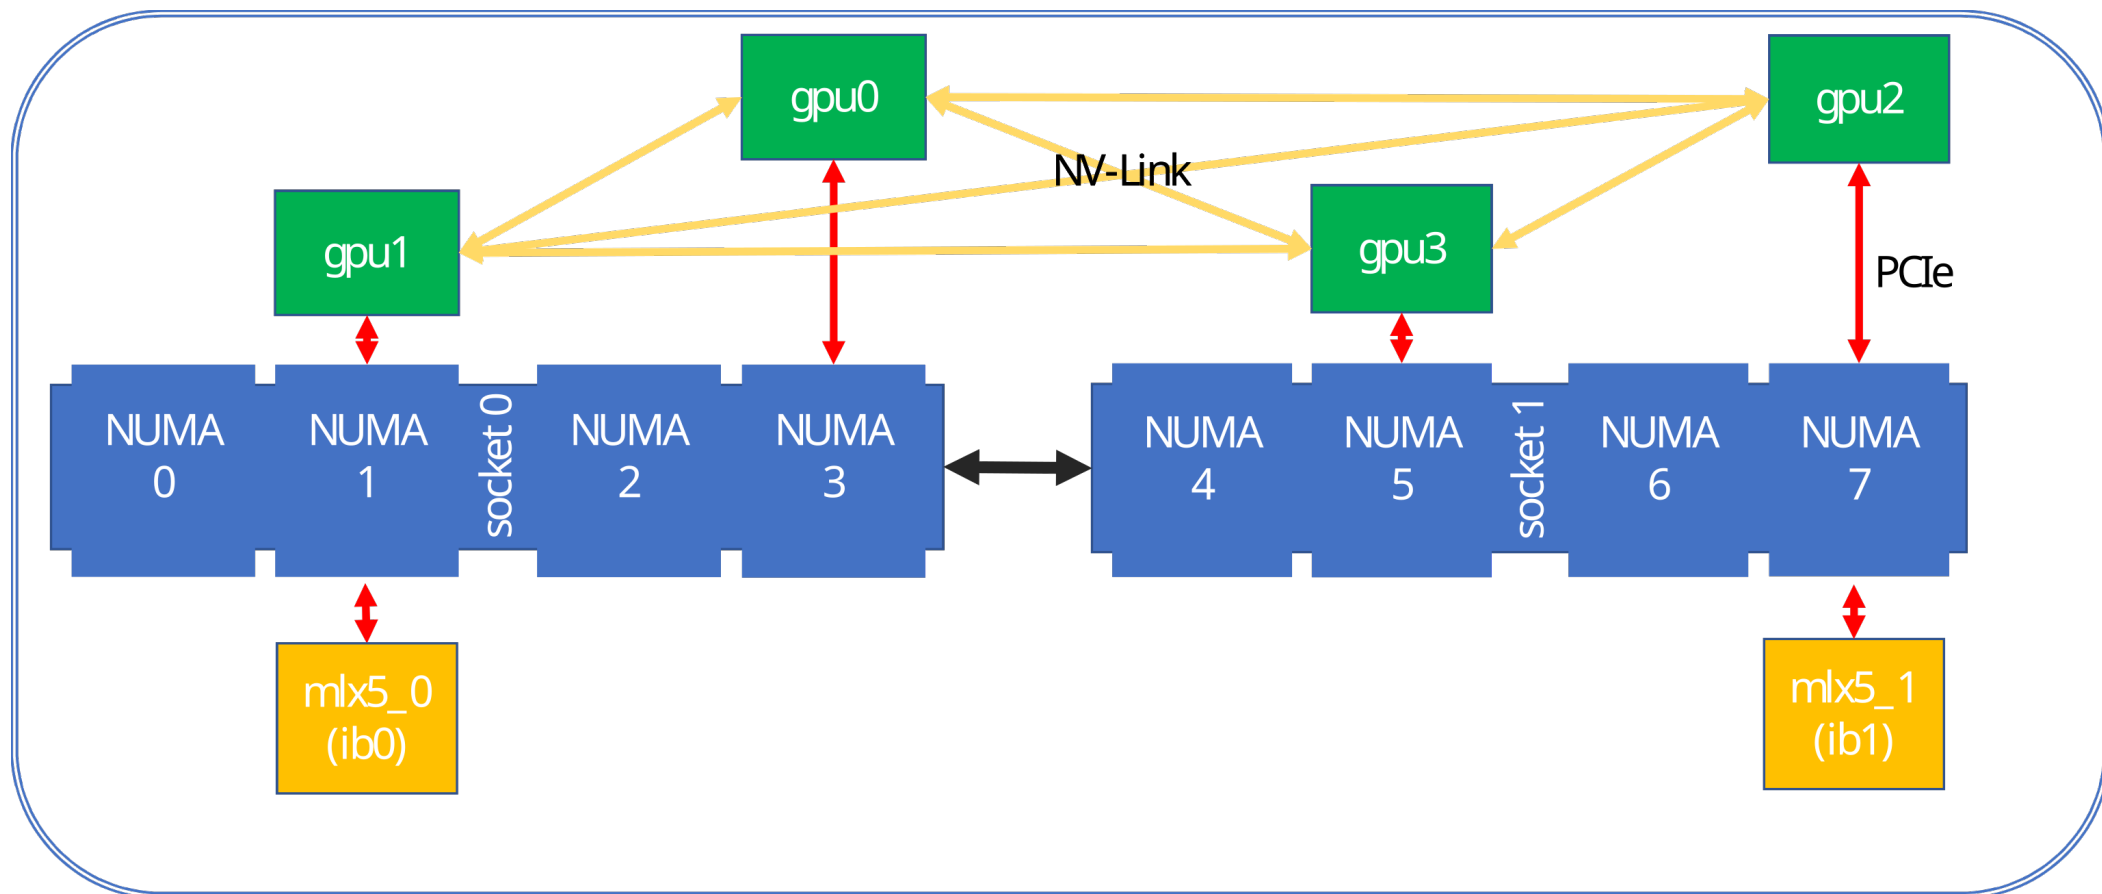
Levante GPUs

Jan Frederik Engels, Harald Braun

Goal of this session

- What is specific about Levante GPU nodes?
- What does one need to take care of to use them efficiently?

Configuration of a Levante GPU node



Types of GPU nodes in Levante

GPUs	CPUs	#Nodes	feature
4x NVIDIA A100 40GB	2x AMD 7763 CPU 128 cores	4	a100_40
4x NVIDIA A100 80GB	2x AMD 7763 CPU 128 cores	56	a100_80

Most nodes have 512GB CPU RAM, two have 1024GB.

Allocating resources on a GPU node

One GPU

```
1 #SBATCH --partition=gpu
2 #SBATCH --gpus=1
3 #SBATCH --account=<prj-account>
4 #SBATCH --constraint=<feature> #Optional
5 #SBATCH --reservation=<reservation> #Only for this course.
```

One or multiple GPU-nodes

```
1 #SBATCH --partition=gpu
2 #SBATCH --gpus-per-node=4
3 #SBATCH --nodes=<number of nodes>
4 #SBATCH --account=<prj-account>
5 #SBATCH --constraint=<feature> #Optional
6 #SBATCH --reservation=<reservation> #Only for this course.
7 #SBATCH --exclusive
8 #SBATCH --mem=0
```

SLURM script for GPU jobs

One GPU

```
1 #!/bin/bash  
2 #SBATCH ...  
3  
4 srun <executable>
```

SLURM script for GPU jobs

One or more nodes

```
1  #!/bin/bash
2  #SBATCH ...
3
4  rm -f gpu_wrapper.sh
5  cat > gpu_wrapper.sh << EOF
6  #!/bin/bash
7
8  mygpu=\$SLURM_LOCALID
9  export CUDA_VISIBLE_DEVICES=\$mygpu
10 cpus=\`nvidia-smi topo -i \$mygpu -C | awk -F ':' '*' '{print \$2}'\`
11 numa=\`nvidia-smi topo -i \$mygpu -M | awk -F ':' '*' '{print \$2}'\`
12
13
14 numactl --cpunodebind=\$cpus --membind=\$numa <executable>
15 EOF
16 chmod 700 ./gpu_wrapper.sh
17
18 srun -l --ntasks-per-node=4 ./gpu_wrapper.sh
```