# MAELSTROM ON AMD

## MI250x vs V100 vs A100 vs H100

November 7, 2023 | Stepan Nassyr | JSC



JÜLICH
Forschungszentrum

V100 nodes (JUWELS Cluster):

| | |
|---|---|
| CPU | 2 × Intel Xeon Gold 6148, (2×20 cores@2.4 GHz) |
| GPU | 4 × NVIDIA V100 GPU, 16 GB HBM |
| Memory | 196GiB DDR4-2666 |
| NIC | 2 × Mellanox EDR InfiniBand ConnectX 4 |

A100 nodes (JUWELS Booster):

| | |
|---|---|
| CPU | 2 × AMD EPYC 7402 (2×24 cores@2.8Ghz) |
| GPU | 4 × NVIDIA A100 GPU, 40 GB HBM |
| Memory | 512GiB DDR4-3200 |
| NIC | 4 × Mellanox HDR InfiniBand ConnectX 6 |

H100 node (JURECA DC Eval.):

| | |
|---|---|
| CPU | 2 × Intel Xeon Platinum Sapphire Rapid 8452Y (2×36 cores@2.0Ghz) |
| GPU | 4 × NVIDIA H100 PCIe GPU, 80 GB HBM |
| Memory | 512GiB DDR5-4800 |
| NIC | 1 × BlueField-2 ConnectX-6 DPU |

AMD GPU nodes @ JSC:

| | |
|---|---|
| CPU | 2 × AMD EPYC 7443 (2×24 cores@2.85Ghz) |
| GPU | 4 × AMD MI250 GPUs, 8 GCDs with 64GB HBM each |
| Memory | 512GiB DDR4-3200 |
| NIC | 1 × Mellanox HDR InfiniBand ConnectX 6 |

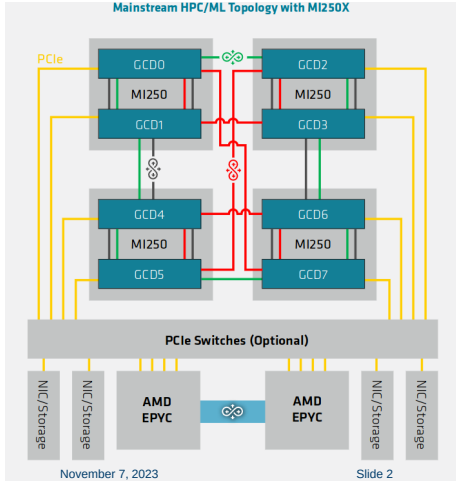

# AMD GPU nodes @ JSC:

- 2 GCDs per card
- assymmetric chip-to-chip bandwidth

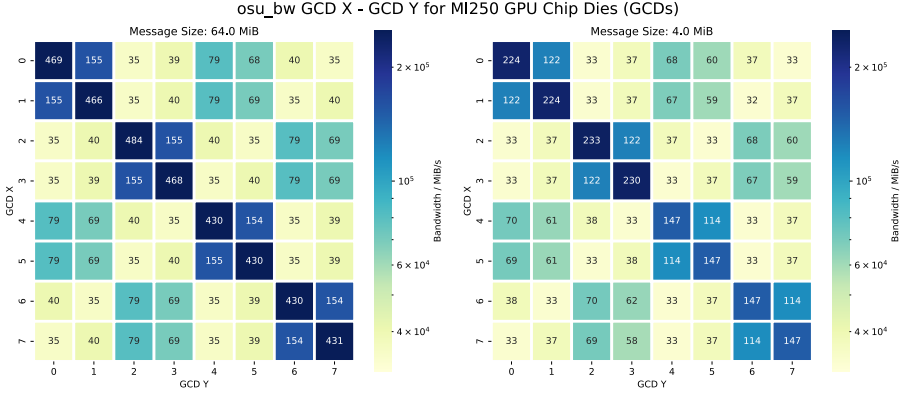


Maps show multiple runs of bandwidth tests between various sets of two GPU GCDs (via HIP_VISIBLE_DEVICES=A,B) on single AMD MI250 node of JURECA DC Evaluation Platform.
Values in cells in GiB/s, Software versions: OSU MB 5.9, ROCm 5.2.0, ROCm driver 5.16.9.22.20, UCX 1.12.1 (UCX_TLS=rc_x,self,sm,rocm_copy,rocm_ipc), OpenMPI 4.1.2,
OSU MB compiled as per AMD recommendation (→ https://go.fzj.de/jsc-amdgpu-osu); osu_bw launched with '-d rocm -m SIZE:SIZE D D' with $2^{26}$ or $2^{22}$ for SIZE; Note the relative color scales.
Detailed description: https://go.fzj.de/jsc-amdgpu-eval.

Theoreticals Mi250 vs A100:

| Metric | MI250 | A100 |
|---|---|---|
| FP64 | 45.3 TFLOP/s | 19.5 TFLOP/s |
| FP32 | 90.5 TFLOP/s | 19.5 TFLOP/s |
| TF32 | - | 156\|312 (dense\|sparse) TFLOP/s |
| FP16 | 362 TFLOP/s | 312\|624 (dense\|sparse) TFLOP/s |

Software:

- MI250x: container - tensorflow_rocm5.7-tf2.13-dev.sif
- V100: JWC EasyBuild modules
- A100: JWB EasyBuild modules
- H100: container - tensorflow_23.10-tf2-py3.sif

Containers + VENV tricky:

- MI250x
  - run container
  - create venv without pip
  - install into container with prefix
- H100
  - run container
  - venv doesn’t work at all
  - install into directory with prefix
  - set PYTHONPATH manually
- Guide/sourceable env scripts for D3.7 WIP

Measuring Energy:

- AMD: energy counter

```
~]$ rocm-smi --showenergycounter --csv
```

- NVIDIA: Measure power and integrate

```
~]$ nvidia-smi --query-gpu=index,timestamp,name,power.draw --format=csv --loop-ms=5
```

```
with GetPower() as measured_scope:
    print('Measuring Energy during main() call')
    try:
        main(args)
    except Exception as exc:
        print(f"Errors during training: {exc}")
print("Energy data:")
f = open(f"EnergyFile-NVDA-{args.id}", 'a')
print(measured_scope.df.groupby('index').get_group(0))
f.write(str(measured_scope.df.groupby('index').get_group(0)))
print("Energy-per-GPU-list:")
print(measured_scope.energy())
f.write(str(measured_scope.energy()))
```

AP5 Results:

- 1 x MI100 (scaled from D3.6) : 218s/epoch
- 0.5 x MI250 (1 GCD): 198-212s/epoch
- 1x V100: 255-260s/epoch
- 1x A100: 125-130s/epoch
- 1x H100: 128-145s/epoch (unoptimized?)

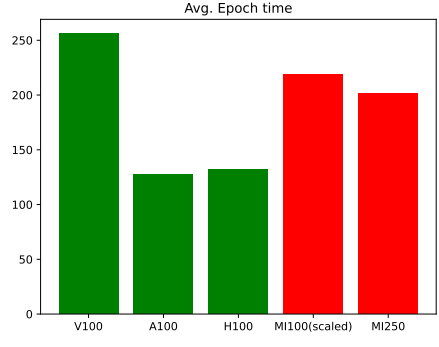

AP5 Results (total training time):

- 1 x MI100 (from D3.6) : 4962s
- 0.5 x MI250 (1 GCD) (scaled): 4691s
- 1x V100 (from D3.6): 5294s
- 1x A100 (from D3.6): 2882s
- 1x H100 (scaled): 3011s

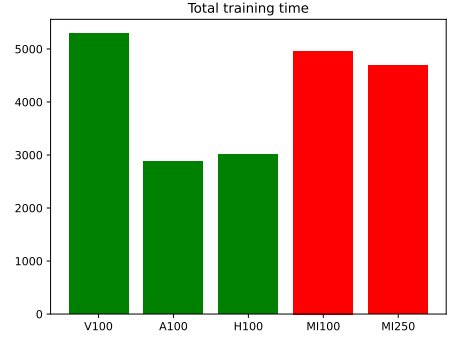

AP5 Results (Energy):

- MI100 : no GPU measurement
- MI250 : 3723Wh ??
- V100 (from D3.6): 131Wh
- A100 (from D3.6): 83Wh
- H100: crash

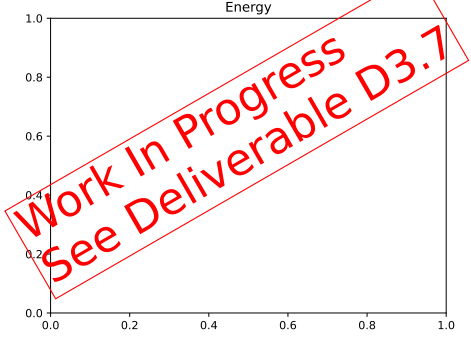

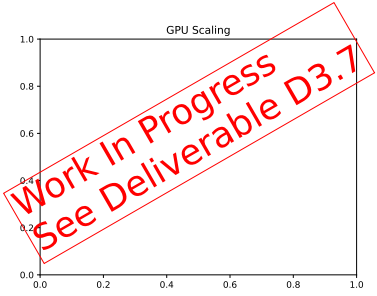
GPU Scaling
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
Work In Progress
See Deliverable D3.7

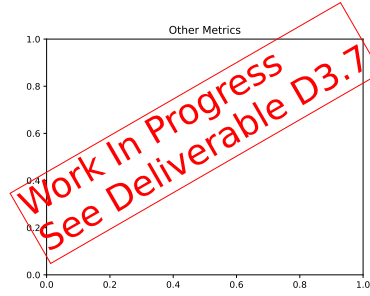
Other Metrics
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
Work In Progress
See Deliverable D3.7

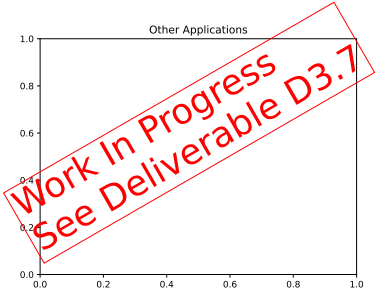
Other Applications
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
Work In Progress
See Deliverable D3.7

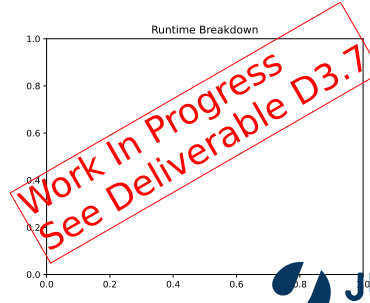
Runtime Breakdown
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
Work In Progress
See Deliverable D3.7

WIP - ironing out kinks

- container issues (both AMD+NVDA)
- containers mentioned before work well
- 0.5 $\times$ Mi250 results promising
- H100 appears underutilized - worth investigating
- Expect more interesting results in D3.7!

# Thank you for your attention!