Using the Jean Zay AI cluster

Mauricio Diaz

Café de la Neuroimagerie - ICM - Feb. 2020

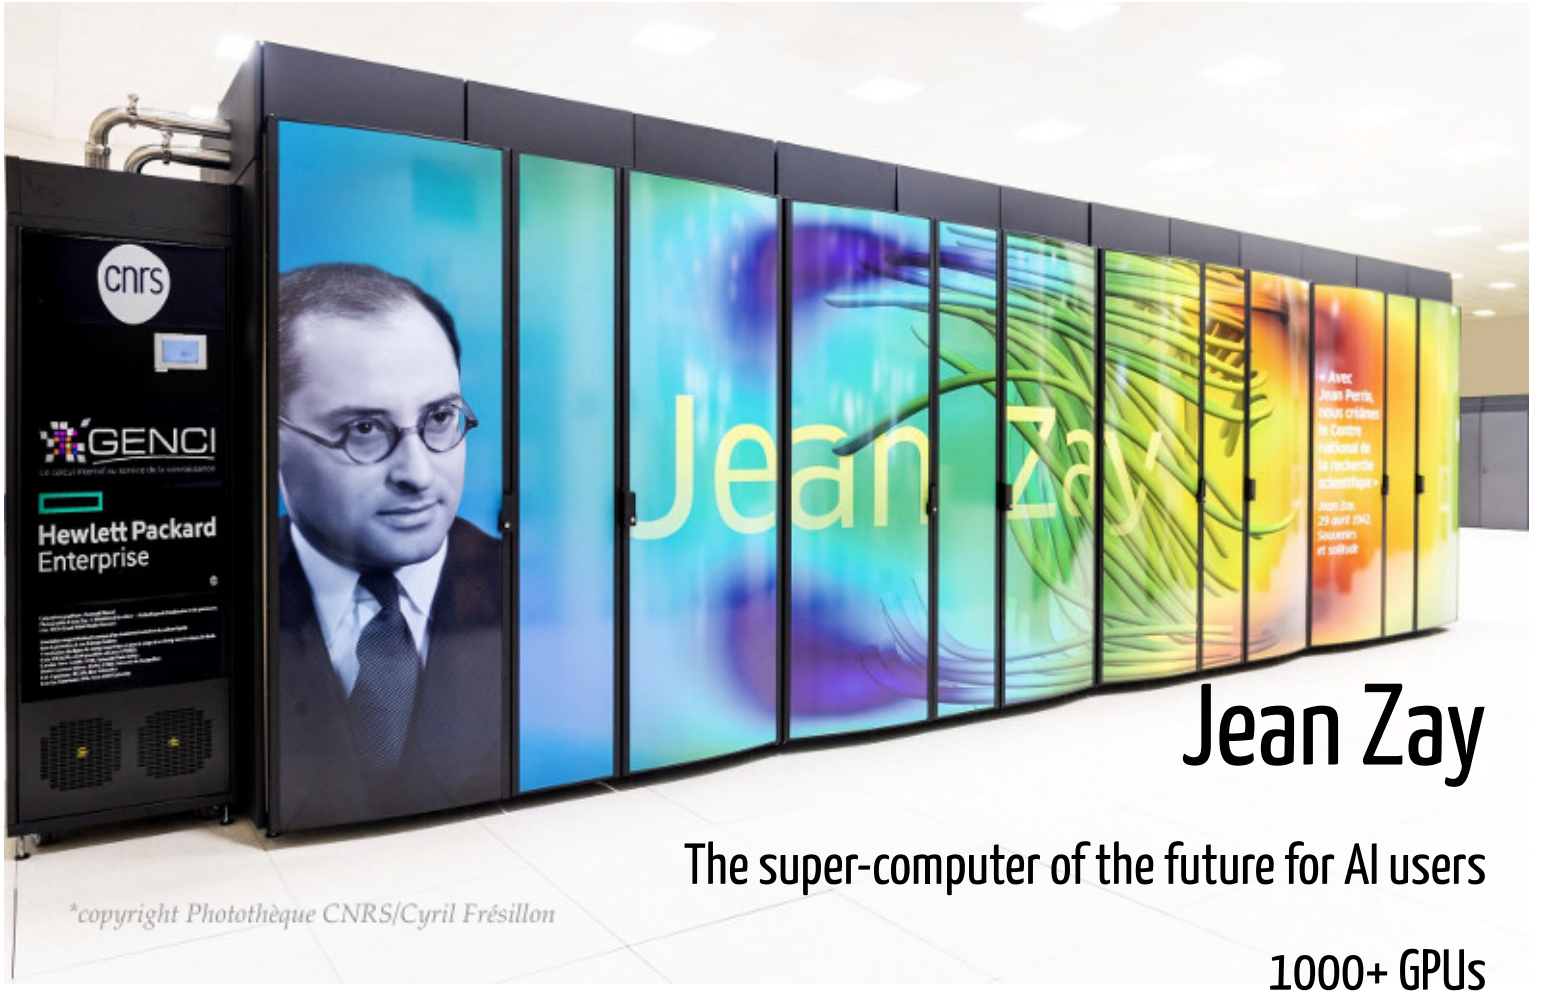

Jean Zay

The super-computer of the future for AI users

1000+ GPUs

Half of the super-computer is for more traditional HPC usages

Specs Hardware

- HPE SGI 8600 computer whith two partitions:
  - 1528 scalar nodes (CPU only - traditional HPC) 4.9 Pflops
  - 261 converged nodes (CPU+GPU - AI+traditional HPC) 9.02 Pflops
- Each node has:
  - 192 GB of memory
  - Omni-PAth interconnection network 100 Gb/s (1 link per scalar node and 4 links per converged node)
  - 4 GPU nVIDIA V100 32 GB (converged nodes)
- In addition: 5 frontal nodes, 5 visualisation nodes, 4 large memory nodes (1 GPU Nvidia V100 and 3 TB of memory per node)
- Spectrum Scale parallel file system (ex-GPFS)
- Parallel storage device with SSD disks with a capacity of 1 PB

How to get access

Possible access modes

- Dynamic access: individual access, simplified process (<=10000 GPU hours).
- Project access: group access, more administrative forms.
- Grand Challenges access (legacy): beta testers last summer within the Aramis team.

Important details:

- Everything is in French: filling the form in English is fine.
- Ask for <= 10000 GPU hours and they will be granted easily.
- 5-10 lines is enough for the project description.
- "Directeur structure de recherche": for ICM: Alexis Brice (ask to his assistant).
- "Responsable sécurité informatique": depending on the host institution (for ICM: see with the DSI).

Docs

Jean Zay Hardware:

- http://www.idris.fr/eng/jean-zay/cpu/jean-zay-cpu-hw-eng.html

Jean Zay Documentation:

- http://www.idris.fr/eng/jean-zay

Users collaborative doc:

- https://github.com/jean-zay-users/jean-zay-doc

Environment description

OS:

- RedHat version 7.6

Scheduler:

- Slurm

Modes:

- monoGPU, multiGPU and multiGPU MPI.

Interactive mode is available.

Environment description - storage

`$HOME` (config files)

Permanent. Usage: e.g. config files, 3Gb.

`$WORK` (quota for the project)

Permanent. Usage: code, databases, 3M inodes, 10Tb

`$SCRATCH` (large quota)

Temporary. Usage: output data files, experiments results.

`$STORE` (quota for project)

Permanent. Usage: large size but occasional consultation, 100k inodes, 51Tb

`$DSDIR` (popular databases)

Environment description - SW

Libraries ready and optimized, via module:

- Intel compilers, with Intel(R) Math Kernel Library
- PGI comipler
- Python 2/3,
- Tensorflow (1.4, 1.8, 1.13, 1.14, 2.0.0-beta1),
- Pytorch (1.1),
- Caffe (1.0)
- Cuda (10.1.1)
- nccl (2.4.2)
- cudnn (10.1-v7.5.1.10)
- OpenACC
- MPI Cuda
- GPUDirect

You can also install your own libraries via Conda.

Environment description - Scheduler

Partitions:

- CPU: two queues - `cpu_p1` et `prepost`.
- GPU: two queues - `gpu_p1` et `prepost`. In the `gpu_p1` queue some nodes are reserved for dev jobs (fast access but limited to 2h).

To be aware:

- Maximal duration of a job (walltime): for batch jobs 20h and for dev jobs 2h (!!).
- Since last december a dedicated queue allows to submit long jobs (less than 120 hours).
- Dealing with checkpoints.

Early feedback - Access

- Jean Zay admins try to propose a "lightweight" procedure (compared with HPC traditional procedure).
- FYI:
  - ARAMIS process -> it took more than 1 month to get access (mode "Grand challenge/beta testers").
  - If everything goes well, it could take around 3 weeks to get access to Jean Zay from filling out the first form to be able to ssh to Jean Zay (mode "Dynamic access").
  - It should be possible to get access faster (in particular if signatures can be obtained faster).

Early feedback - Aramis project

- Goal: Reproduce experiments to guarantee reproducibility of the results of a publication:
- What we did: Train several CNN networks on 3D MRI data.
- Before using JZ: It took (on a local cluster) more than 3 months (about 1 month of full cluster usage) to get the results.
- After using JZ: All the models were re-trained within a week (about 24h of full cluster usage).
- Next: search of hyper-parameters to improve our models.
- Next: access to a new database (*UKBIOBANK*). More data to re-train our models.

Early feedback - General

- Easy and fast access (at least during the GC period). After opening to the public, some peaks >10 min for the long queue.
- Huge computing ressources, impossible to access locally.
- Perfect setup: local cluster (with slurm) for dev -> then JZ
- Idea: increase some developement efforts to optimize use of ressources:
  - handle correctly checkpoints,
  - develop code multiGPU,
  - use tensorcores,
  - handle mixed-precision,
  - etc...
- Idea: At ICM an output proxy is already configured (you can connect to JZ from all the computers connected to the ICM internal network).
- Needs to improve communication about maintenance.

```
upd53tc@jean-zay.idris.fr password:
*************************************************************************
*                                                                       *
*                        Maintenance sur Jean Zay                       *
*                                                                       *
*               Le mardi 29 octobre 2019 de 8h30 a 12h00                *
```

Thanks!