

Folding@Home jobs in ALICE

costin.grigoras@cern.ch

F@H execution

<https://download.foldingathome.org/releases/public/release/fahclient/centos-6.7-64bit/v7.5/>
→ /cvmfs/alice.cern.ch/fah/

3 stage pilot job

- fetch a previously saved piece of work (or start from scratch)
- run it for 12h (using the `timeout` command)
- upload the new snapshot (`cores/` and `work/`) back to the same place

`FAHClient` configured to start immediately and run on a single CPU core
migrating work between slots requires the same number of cores

F@H config.xml

```
<config>  
  <!-- Client Control -->  
  <idle-seconds v='0'/>  
  
  <!-- Folding Slot Configuration -->  
  <cpus v='1'/>  
  <gpu v='false'/>  
  
  <!-- User Information -->  
  <team v='38188'/>  
  <user v='ALICE-${ALIEN_SITE}'/>  
  
  <!-- Folding Slots -->  
  <slot id='0' type='CPU'/>  
</config>
```

External management

Two conditions that trigger actions:

Kill and restart the process if it cannot fetch a payload (`:ERROR:` in the last few log lines)

Kill, `rm -rf work/ cores/`, restart if

```
FahCore returned: INTERRUPTED (102 = 0x66)
```

Activity and accounting

Running up to 6k concurrent jobs Grid-wide

- ~25% of the team contribution in WorkUnits
- ~10% in Credits

Not targeting particular resources, use any available CPU

Excluding a few sites that dedicate other resources to F@H

Stats collected from this JSON (updated @ ~15min):

<https://stats.foldingathome.org/api/team/38188>

ALICE Online cluster

Largest contribution of our sites (%Credit)

Commissioning of the the new online nodes

- 149 machines for First Level Processing (FLP)

 - 137x Xeon(R) Silver 4210 CPU @ 2.20GHz (2x 10 cores)

 - 12x Xeon(R) Gold 6230 CPU @ 2.10GHz (2x 20 cores)

 - No GPUs

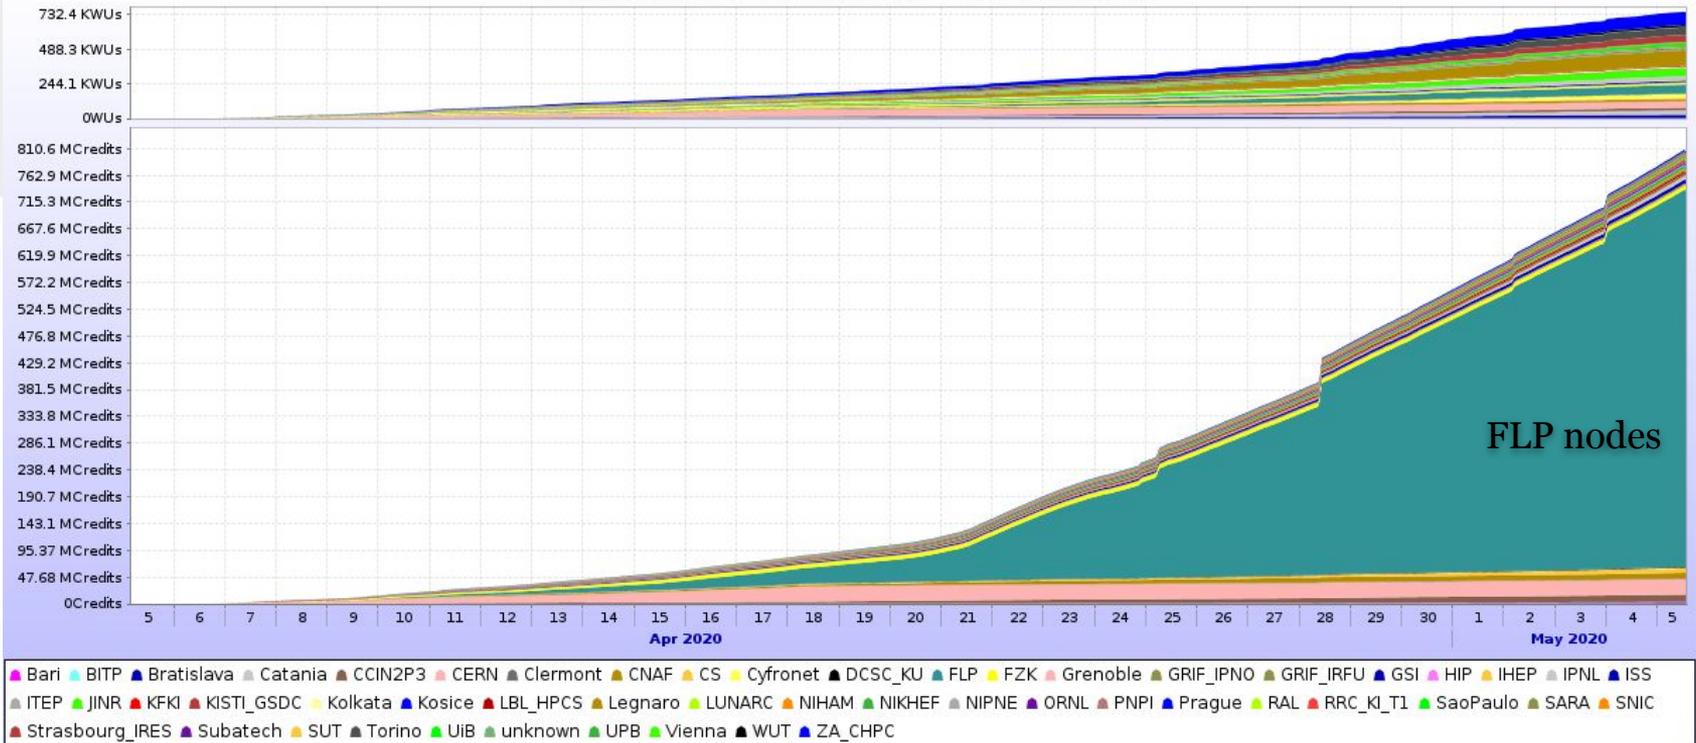
- Fully running F@H as part of the hardware commissioning

 - Machine stability under load, cooling performance

 - Planning to stop running F@H after the commissioning

Breakdown of contributions

ALICE sites contribution to Folding@Home

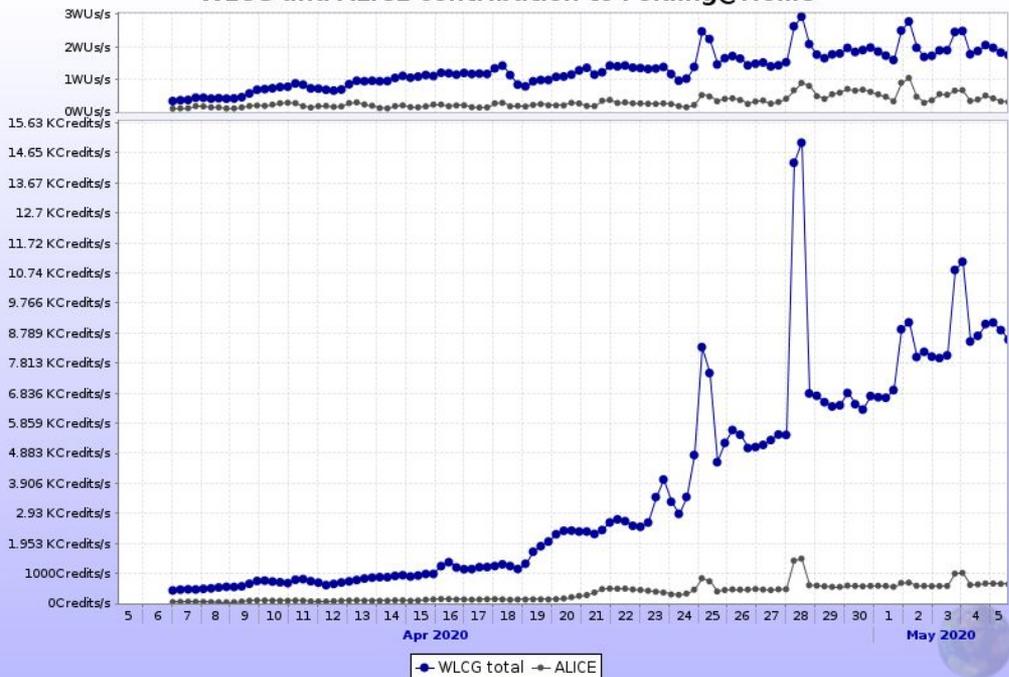


Some plots

Active Folding@Home jobs



WLCG and ALICE contribution to Folding@Home



WLCG team ranking in F@H



Final remarks

We plan to continue running these jobs on our Grid resources at a level of ~6k jobs (~5% of our # of CPU cores)

Plus any fixed services that can spare CPU cycles

Minus the FLP nodes