

# Job scheduling with jobs' energy profiles

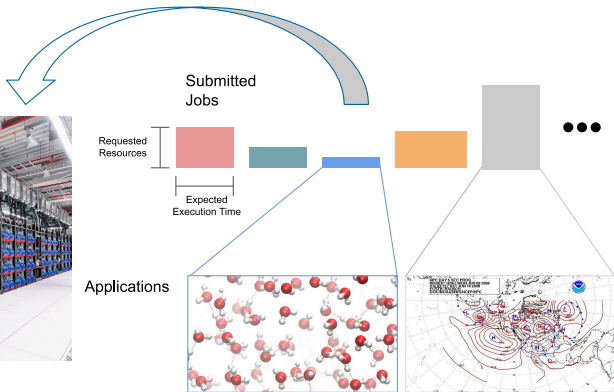
**Danilo Carastan-Santos**  
Inria project team DataMove

<sup>1</sup>Université Grenoble Alpes, Grenoble INP, Inria, LIG, France  
email:[danilo.carastan-dos-santos@inria.fr](mailto:danilo.carastan-dos-santos@inria.fr)

December 19, 2023

# Overview of the problem

High-Performance Computing (HPC) resource management



# Our use-case

- Gricad<sup>a</sup> large-scale computing platform
- **Dahu** cluster (Grenoble site)
  - Each node: dual-socket Intel Xeon Gold 6130 (16 physical cores, 32 virtual)
  - Nodes' energy data collected with **Colmet**<sup>b</sup> (Oar-team in Grenoble)

---

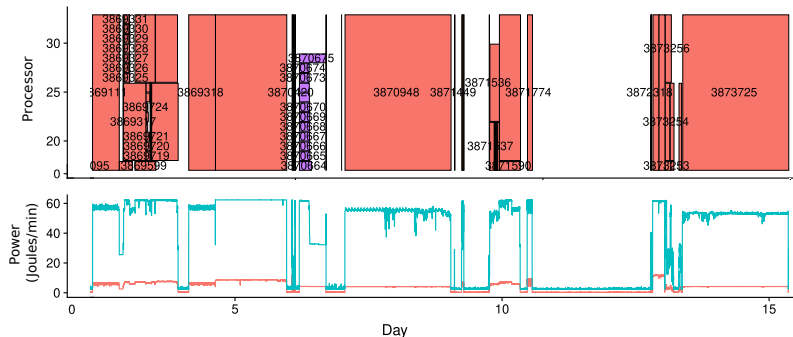
<sup>a</sup><https://gricad.univ-grenoble-alpes.fr/>

<sup>b</sup><https://github.com/oar-team/colmet>



# Monitoring the energy consumption of the Dahu Cluster<sup>2</sup>

- Two sources of data:
  - **Jobs** (OAR, upper graph): processing time and number of processors
  - **Energy consumption** (Colmet, lower graph)






**Problems:** Jobs that share a node, jobs that run in multiple nodes, incomplete energy traces, container jobs<sup>1</sup>

<sup>1</sup>Jobs that host other jobs inside. This is a standard OAR feature

<sup>2</sup>Example illustrating a single socket of a Dahu node, with hyperthreading enabled.

# Some job energy profiles

- We can see what's inside the “job box” regarding the energy consumption
- More data to do scheduling decisions
- **Challenges:**
  - Predicting the jobs' energy profile
  - Using energy predictions when scheduling jobs



# Job scheduling with jobs' energy profiles

Danilo Carastan-Santos

## Contact

- Email:  
danilo.carastan-dos-santos@inria.fr
- Website: <https://danilo-carastan-santos.github.io/> (**QR code on the right**)

