

Ctrl-A Minute talks

Towards context-awareness for HPC resources management: models,
software architectures, . . .



Raphaël Bleuse

`raphael.bleuse@inria.fr`

`https://research.bleuse.net`

7 mai 2020

How to model parallel systems w.r.t. context ?



platform
(resources)

tasks/jobs



objective

⇒ where ? & when ?


Changing/Augmenting the Model ?

Affinity / $\text{aff} : \text{node} \times \text{task} \mapsto \mathbb{R}^+$

Consider *each* resource : abstract score (e.g., data to transfer)



Performance / GFlop/s
higher is better




Data Transfers / GB
lower is better


Changing/Augmenting the Model ? (again ?)

Geometric constraints


- Bring context-awareness
- Keep model simple (i.e. tractable)



contiguity



connectivity, convexity



locality

Changing *how!*

Closing the loop.

