



Contribution ID: 281

Type: oral presentation

## Scheduling multicore workload on shared multipurpose clusters

*Tuesday 14 April 2015 14:30 (15 minutes)*

With the advent of workloads containing explicit requests for multiple cores in a single grid job, grid sites faced a new set of challenges in workload scheduling. The most common batch schedulers deployed at HEP computing sites do a poor job at multicore scheduling when using only the native capabilities of those schedulers. This talk describes how efficient multicore scheduling was achieved at the three sites represented in the author list, by implementing dynamically-sized multicore partitions via a minimalistic addition to the Torque/Maui batch system already in use at those sites.

The first part of the talk covers the theory related to this particular problem, which is also applicable to e.g. the scheduling of large-memory jobs or data-aware jobs similarly comprising part of a highly heterogeneous workload. The system design is also presented, linking it to previous work at Nikhef on grid-cluster scheduling. The second part of the talk presents an evaluation of several months of production operation at the three sites.

**Author:** TEMPLON, Jeff (NIKHEF (NL))

**Co-authors:** Dr FORTI, Alessandra (University of Manchester); Dr PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES); ACOSTA SILVA, Carlos (Universitat Autònoma de Barcelona (ES)); FLIX MOLINA, Jose (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES); STARINK, Ronald (Nikhef)

**Presenter:** TEMPLON, Jeff (NIKHEF (NL))

**Session Classification:** Track 6 Session

**Track Classification:** Track6: Facilities, Infrastructure, Network