



Contribution ID: 140

Type: **Sectional reports**

## HPC workload balancing algorithm for co-scheduling environments.

*Friday, 9 July 2021 12:00 (15 minutes)*

Commonly used job schedulers in high-performance computing environments do not allow resource oversubscription. They usually work by assigning the entire node to a single job even if the job utilises only a small portion of nodes' available resources. This may lead to cluster resources under-utilization and to an increase of job wait time in the queue. Every job may have different requirements for shared resources (e.g. network, memory bus, IO bandwidth or cpu cores) and they may not overlap with requirements of other jobs. Because of that, running non-interfering jobs simultaneously on shared resources may increase resource utilization.

Without accounting for jobs resource requirements and their performance degradation due to shared resources, co-scheduling may only decrease job performance and overall scheduler throughput. In this work, we propose a method for measuring job run-time performance and an algorithm for selecting and running combinations of jobs simultaneously on shared resources.

Performance metrics were validated on experimental data and an algorithm was derived from a mathematical model, tested on numerical simulations and implemented in the scheduler.

### Summary

**Primary authors:** KUCHUMOV, Ruslan (Saint Petersburg State University); KORKHOV, Vladimir (St. Petersburg State University)

**Presenter:** KUCHUMOV, Ruslan (Saint Petersburg State University)

**Session Classification:** HPC

**Track Classification:** 5. High Performance Computing