



Universität
Basel

Research Project

Robust Scheduling in High Performance Computing (RS in HPC)

Project funded by own resources

Project title Robust Scheduling in High Performance Computing (RS in HPC)

Principal Investigator(s) [Ciorba, Florina M.](#) ;

Project Members [Mohammed, Ali Omar Abdelazim](#) ;

Organisation / Research unit

Departement Mathematik und Informatik / High Performance Computing (Ciorba)

Project start 01.08.2015

Probable end 31.07.2020

Status Completed

High performance computing systems consume and dissipate a great amount of power. Excessive heat dissipation requires aggressive cooling and extra space that adds to the power consumption and infrastructure cost. Moreover, as the sizes of the system as well as the system temperature rapidly increase, high system failure rates are observed. Thus, a feature of interest for scheduling scientific applications in such environments is support for fault detection and management. This characterizes the quality aspect of the time-to-solution.

A solution to the application-level resilience to faults problem must meet the following requirements: (i) Efficiency, without compromising performance; (ii) The reliability level must be user controlled – greater reliability incurs a higher cost (either in terms of resources, CPU time, energy consumption, or allocation price); and (iii) Minimal code changes in the application. Scheduling algorithms that detect faults and are able to manage them are called fault tolerant (or resilient to faults). The most common fault tolerance strategies include task replication (via double or triple modular redundancy) and application checkpointing. However, it is unclear which of the existing solutions will scale to the size of the exascale computing systems expected by the beginning of the next decade.

Keywords robustness, resiliency, fault tolerance, fault modeling

Financed by

University funds

Other funds

Add publication

Add documents

Specify cooperation partners