

## Improved Job Scheduling using SLURM and Tasklets

The goal of this thesis is to explore the potential of introducing Tasklets to SLURM, which is a job scheduler often used in HPC (high performance computing) clusters to manage the workload. While SLURM is very scalable and extensive, often resources go unused because larger jobs with higher priority are waiting in a queue already. Currently, it is not possible to notify SLURM that a user has a low priority job, that may be terminated in favor of more important ones. The Tasklet system is a middleware for distributed applications that allows developers to offload computation to remote resources via self-contained units of computation the so-called Tasklets. These Tasklets are lightweight and best-effort by design and can be displaced by higher-priority tasks at any time. This allows for the efficient use of excess capacities in computing clusters where no guarantees can be given regarding resource availability. By pairing SLURM with Tasklets, the resource utilization should be optimized.

- ▶ Tasklets
  - ▶ <https://ieeexplore.ieee.org/abstract/document/7568580>
  - ▶ <https://dl.acm.org/doi/pdf/10.1145/3007592.3007597>
- ▶ SLURM
  - ▶ <https://github.com/SchedMD/slurm>
  - ▶ <https://slurm.schedmd.com/overview.html>

Contact: › Michael Kuhn (<https://parcio.ovgu.de/People/Michael+Kuhn.html>) and › Janick Edinger (UHH) (<https://www.inf.uni-hamburg.de/inst/ab/dos/people/edinger.html>)