

Currently there is no general solution for managing heterogeneous resources at the operating system-level exploiting a global view of the system state for assigning tasks to accelerator resources (e.g. GPU, FPGA) in a server of the OTF computer center.

In this talk, we will present a user space scheduler that enables scheduling and load balancing on heterogeneous resources in Linux. We determine energy and performance affinities of tasks for each resource using a dedicated measurement framework and perform scheduling decisions based on these metrics and the system state. With a programming pattern that supports the integration of checkpoints into applications, we show the feasibility of preempting tasks and migrating them between accelerators. Our results show, that depending on the selected scheduling policy, we can optimize for different, configurable goals, e.g. maximizing throughput for a task set, or minimizing total energy for executing a task set.