Seminar/Talk

Firmament: Fast, Centralized Cluster Scheduling at Scale

Malte Schwarzkopf
MIT, CSAIL

Host: Tom Henzinger

Scheduling tasks on clusters in large-scale datacenters is challenging:
thousands of tasks must be cleverly and quickly placed to achieve good application-level performance and high resource utilization. Centralized datacenter schedulers can make high-quality placement decisions, but today these high-quality placements come at the cost of high latency at scale, which degrades response time for interactive tasks and reduces resource utilization.

In this talk I present Firmament, a centralized scheduler that scales to over ten thousand machines, even though it performs a computationally expensive min-cost max-flow (MCMF) optimization that continuously reschedules all tasks. To achieve this, Firmament automatically chooses between different MCMF algorithms, solves the optimization problem incrementally when possible, and applies problem-specific optimizations.

Experiments with a Google workload trace from a 12,500-machine cluster show that Firmament places tasks in hundreds of milliseconds, and that Firmament improves placement latency by 20x over Quincy, a prior centralized scheduler using the same MCMF optimization. Moreover, even though Firmament is centralized, it matches the placement latency of distributed schedulers for workloads of short tasks. Finally, Firmament exceeds the placement quality of four widely-used centralized and distributed schedulers on a real-world cluster, improving batch task response time by 6x.

Reference:
I. Gog, M. Schwarzkopf, A. Gleave, R. N. M. Watson, S. Hand. "Firmament:
Fast, Centralized Cluster Scheduling at Scale". In Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Savannah, GA, USA, November 2016, p. 99--115.

Wednesday, December 21, 2016 01:45pm - 02:45pm

Mondi Seminar Room 2, Central Building

This invitation is valid as a ticket for the ISTA Shuttle from and to Heiligenstadt Station. Please find a schedule of the ISTA Shuttle on our webpage: https://ista.ac.at/en/campus/how-to-get-here/ The ISTA Shuttle bus is marked ISTA Shuttle (#142) and has the Institute Logo printed on the side.