TARCIL: RECONCILING SCHEDULING SPEED AND QUALITY IN LARGE SCALE CLUSTERS

Scriber: Ailing Zhang

Apr 12, 2016

SUMMARY

Motivation

- Current schedulers either aim at scheduling speed or scheduling quality.

Key Contribution

- Use sampling to find available resources & identify resources that best match a job's preference
- Admission Control: queues jobs until appropriate resource is available

PROS:

- Achieves good performance for both short and long jobs.
- Analytical framework provides statistical guarantee for the quality of scheduling decisions
- Evaluation section is thorough and well organized.

CONS:

- Tarcil profiles each job to determine its resource preference.(caused high overhead for short jobs)
- Low priority tasks may suffer high overhead when more than a third of tasks have high priority.

DISCUSSION

- An RU is a fixed unit of resource. Can the resource of RUs be broken down to provide better quality?
- How is network bandwidth allocation achieved?

We currently use Linux containers to partition servers into RUs . Containers enable CPU, memory, and I/O isolation.

Each container is configured to a single core and a fair share of the memory and storage subsystem, and network bandwidth.