

Derya Cavdar

Thesis Supervisor: Prof. Fatih Alagöz

Priority Scheduling For Heterogeneous Workloads in Computing Clusters

Abstract

Nowadays, large computing clusters constantly strive for an optimal tradeoff between resource efficiency and performance. In this thesis, we are concerned with the efficient use of system resources and we also aim to improve response time of the tasks. We tackle with the challenges of task scheduling on large heterogeneous clusters executing highly heterogeneous bursty workloads with different priorities, resource demands, and performance objectives. Firstly, we propose a scheduling algorithm for tasks with communication needs which improves the response time and resource utilization by controlling communication and computation resources at the same time. Secondly, we propose a novel scheduling framework for exploring various aspects of priority scheduling with heterogeneous workloads while investigating the tradeoff between evictions and response time. To better understand the impact of evictions, we first analyze simple eviction policies and wasted resources associated with evictions by using trace-driven simulations. Furthermore, by exploiting the heterogeneity of the workload, we propose a workload-aware slot configuration and task assignment methodology incorporated with slot-based priority scheduling to improve class-based response time and resource efficiency. Finally, we introduce a task scheduling policy, which aims to provide scheduling and execution guarantees for low priorities while preserving the performance benefits of high priority tasks. The proposed scheduling policy effectively handles both prioritization and performance issues of low priorities by utilizing a combination of preemptive and non-preemptive scheduling.

PUBLICATIONS

Journals

1. **Derya Cavdar**, Lydia Chen, Fatih Alagöz, “Priority Scheduling for Heterogeneous Workloads: Tradeoff between Evictions and Response Time”, IEEE Systems Journal, vol.PP, no.99, pp.1-12, 2015.
2. **Derya Cavdar**, Robert Birke, Lydia Chen, Fatih Alagöz, “A Simulation Framework for Priority Scheduling on Heterogeneous Clusters”, Future Generation Computer Systems (Elsevier), Vol. 52, pp. 37-48, 2015.
3. **Derya Cavdar**, Andrea Rosa, Lydia Chen, Walter Binder, and Fatih Alagöz, “Quantifying the Brown Side of Priority Schedulers: Lessons from Big Clusters”, ACM Sigmetrics Performance Evaluation Review, Vol. 42 Issue 3, pp. 76-81, 2014.

Conferences

1. **Derya Cavdar**, Lydia Chen, Fatih Alagöz, “Green MapReduce for Heterogeneous Data Centers”, IEEE Global Communications Conference (GLOBECOM), Austin, USA, 2014.
2. **Derya Cavdar**, Fatih Alagöz, “Network Aware Job Scheduling in Green Data Centers”, International Conference on Computing, Networking and Communications (ICNC), Honolulu, USA, 2014.

3. **Derya Cavdar**, Fatih Alagöz, “A Survey of Research on Greening Data Centers”, IEEE Global Communications Conference (GLOBECOM), Anaheim, USA, 2012.

Defense Jury Members

- | | |
|-------------------------------------|------------------------|
| 1. Prof. Fatih Alagöz | Boğaziçi University |
| 2. Prof. Özgür Barış Akan | Koç University |
| 3. Prof. Mehmet Ufuk Çağlayan | Boğaziçi University |
| 4. Assist. Prof. Özlem Durmaz Incel | Galatasaray University |
| 5. Prof. Tuna Tuğcu | Boğaziçi University |

Defense Date: 08.06.2015