

DYNAMIC I/O-AWARE LOAD BALANCING AND RESOURCE MANAGEMENT FOR CLUSTERS

Xiao Qin, Ph.D.

Department of Computer Science and Engineering
University of Nebraska, 2004

Advisor: Hong Jiang

In the last decade, clusters have become increasingly popular as powerful and cost-effective platforms for executing parallel and distributed applications. Dynamic load balancing techniques have been investigated to achieve efficient utilizations of resources on clusters. Many load balancing policies achieve high system performance by increasing the utilization of CPU or memory resources. However, these load-balancing policies suffer significant performance drop on clusters when workload comprises a large number of data-intensive applications, and when disk- or network-I/O resources exhibit imbalanced load. This is because in modern cluster systems, the performance gap between CPU and I/O subsystems is rapidly growing. Therefore, any dynamic load balancing scheme has to be “I/O-aware” in order to sustain high performance in this new application environment.

In this thesis, we propose a series of disk-I/O-aware and communication-aware load-balancing schemes, which dynamically detect disk- and network-I/O load imbalance on nodes of a cluster. One approach is to assign data-intensive jobs to nodes with light I/O load, whereas the second approach preemptively migrates some currently running jobs

with large I/O demands from overloaded nodes to other less- or under-loaded nodes. In addition, these schemes take into account both CPU and memory load sharing in the system. Since heterogeneity in disks inevitably imposes a performance constraint when coupled with an imbalanced disk-I/O load, we propose a load balancing scheme to support heterogeneous clusters by hiding the heterogeneity of resources. Extensive trace-driven simulations, which are based on a set of real applications and synthetic jobs with various I/O characteristics, show that the I/O-aware load-balancing schemes consistently improve the performance over existing non-I/O-aware load-balancing schemes for a wide spectrum of workload conditions on both homogeneous and heterogeneous clusters. Furthermore, we have presented a feedback control mechanism to adaptively manipulate disk-I/O buffer sizes. The empirical results have demonstrated that the feedback control mechanism can not only be leveraged to enhance the performance of load-balancing schemes, but also be applied to clusters where workload conditions exhibit dynamic behaviors.