



NORTHWESTERN
UNIVERSITY

Computer Science Lecture Series

Programming Support and Resource Management for Cluster-based Internet Services

Mr. Hong Tang
University of California at Santa Barbara

12:30-1:30pm
Wednesday, March 5, 2003
Room 381 - Computer Science Dept.

Abstract:

Clusters of workstations are a natural platform for highly available large-scale Internet services due to their cost-effectiveness and incremental scalability. Constructing robust Internet services on clusters is challenging. Software systems need to not only tolerate hardware failures and operation mistakes, but also consider platform and application heterogeneity. Additionally, the highly-concurrent and bursty nature of service traffic calls for a careful design of adaptive request scheduling.

In this talk, I will present two aspects of our research work on cluster-based Internet services: scalable data aggregation and integrated resource management. Aggregating data from a large number of partitions is a fundamental building block of service programming, yet it is poorly supported in previous work. We propose a primitive called DAC (Data Aggregation Call) that hides the complexity of data aggregation operations behind an easy-to-use interface. The runtime system adopts a dynamic reduction tree formation algorithm and an event-driven request scheduling scheme. Our evaluation confirms that the system is effective to deliver low response time and maintain high throughput in the presence of platform heterogeneity and component failures.

I will then describe an integrated resource management framework that provides service-specific resource management and QoS support. Our solution uses yield functions to impose customized management objectives and guide request scheduling. Our decentralized architecture design and adaptive scheduling scheme are effective to achieve service differentiation objectives during demand spikes and node failures while maintaining the scalability.

For a complete calendar, see:
www.cs.northwestern.edu;
click on 'CS Seminars'

Join our e-mail list for notices of upcoming Computer Science lectures:
send an e-mail with the word 'subscribe' in the subject line to:
cs_seminar@cs.northwestern.edu