Mor Harchol-Balter

Associate Professor
Computer Science Department
SCS
August 29th, 2007 2:30pm

Research Topic
Queuing Theory

Research Problem
What is the most efficient scheduling algorithm for a network of supercomputers which handle jobs of enormous size?

Problem Statement
Given a network of supercomputers and a description of the processing and size requirements of jobs which are considered to be "very large," construct a scheduling algorithm which will process any random input sequence of "very large" jobs of varying size in the most efficient way.

Network: a collection of some types of elements, usually computers, which communicate with each other, usually in a distributed manner.

Supercomputers: computers which can perform calculations at a rate far above the normal rate for regular commercial computers.

Processing & size requirements of jobs: the amount of physical and virtual storage required to process a job, and the length of time that would be required to fully process the job.

Scheduling algorithm: A mathematical function which, given a set of resources and a set of tasks, assigns the tasks to the resources in a predefined and (hopefully) intelligent manner.

Problem Description
Prof. Harcol-Balter primarily studies and teaches queuing theory through the use of probability. She is currently working on a number of problems. The primary project she is currently working on is:

TeraGrid
For a network of supercomputers that handle jobs with enormous processing requirements, how do you schedule jobs to be completed? These are computers which handle very large simulations, such as projects to determine if there was life on Mars. Previously, job scheduling was done manually. Prof. Harcol-Balter is attempting to use fairness and auction-based concepts, along with predictability, to find a better solution.

Computer Science Perspective
This problem represents traditional (if complex) computer science concepts such as queuing theory. Determining the most efficient way to perform some set of distributed tasks, such as scheduling jobs, allocating power, or prioritizing requests, is a well-established problem in the computer science field due to our numerous constraints on performance and resources, among other quality attributes.

Disciplines actively involved
Mathematics; Probability; Algorithms; Queuing Theory.

Actively Involved Discipline: a discipline in which further research can result as a result of the successful completion of this research.

Description of Disciplines Involved
Mathematics, and specifically probability theory, is involved in this research. The ultimate focus is to expand on Queuing Theory, and Algorithms will likely be used to perform experiments and prove or disprove the research.

References
Presenter web page:
http://www.cs.cmu.edu/~harchol/
TeraGrid effort
http://www.cs.cmu.edu/~harchol/teragrid.pdf

By mnovakou
Updated mnovakou