Semester: Fall 2009
MS Student: Hira K. Shrestha
Chair: Dr. John Antonio
Thesis title: SCHEDULING WORKFLOWS ON A CLUSTER OF MEMORY MANAGED MULTICORE MACHINES

Abstract:
Workflows are modeled with directed acyclic graphs in which vertices represent computational tasks, referred to as requests, and edges represent precedent constraints among requests. Associated with each workflow is a deadline that defines the time by which all computations of a workflow should be complete. Workflows are submitted by numerous clients to a scheduler that assigns workflow requests to a cluster of memory managed multi-core machines for execution. The objective of the scheduler is to minimize missed workflow deadlines. The characteristics of workflows are assumed to vary along several dimensions, including: arrival rate, periodicity, degree of parallelism, and number of requests. For the purpose of this thesis, an overall scheduling policy is defined by two underlying components: (1) a request selection policy and (2) a machine selection policy. A total of forty-two scheduling policies are evaluated, which are defined according to the cross-product of seven underlying request selection policies and six machine selection policies. The performance of each policy is determined through extensive simulation studies. All of the machine selection policies rely on a specified loading threshold. The simulation studies conducted reveal the existence of an optimal threshold value for each machine selection policy that results in relatively comparable performance across all machine selection policies, for the scenarios considered. Of the seven request selection policies evaluated, three are newly introduced, and the remaining are derived from previously known policies from the literature. From the studies conducted, it is determined that one of the newly introduced request selection policies outperforms the others evaluated in terms of minimizing a measure of normalized tardiness.