

The Colorado State University Computer Science Department presents:

Many-Task Computing: Bridging the Gap between High-Throughput Computing and High-Performance Computing

Dr. Ioan Raicu

Monday, March 2, 2009, 9:00, Computer Science Building 130

CS 692 students are not required to attend

Abstract

Many-task computing aims to bridge the gap between two computing paradigms, high-throughput computing and high-performance computing. Many-task computing is reminiscent to high-throughput computing, but it differs in the emphasis of using many computing resources over short periods of time to accomplish many computational tasks, where the primary metrics are measured in seconds (e.g. tasks per second, I/O per second), as opposed to operations per month (e.g. jobs per month). Many-task computing denotes high-performance computations comprising of multiple distinct activities, coupled via file system operations. Tasks may be small or large, uniprocessor or multiprocessor, compute-intensive or data-intensive. The set of tasks may be static or dynamic, homogeneous or heterogeneous, loosely coupled or tightly coupled. The aggregate number of tasks, quantity of computing, and volumes of data may be extremely large. Many-task computing includes loosely coupled applications that are generally communication-intensive but not naturally expressed using message passing interface commonly found in high-performance computing, drawing attention to the many computations that are heterogeneous but not happily parallel. This talk explores fundamental issues in defining the many-task computing paradigm, as well as theoretical and practical issues in supporting both compute and data intensive many-task computing on large scale systems. We have defined an abstract model for data diffusion an approach to supporting data-intensive many-task computing, have defined data-aware scheduling policies with heuristics to optimize real world performance, and developed a competitive online caching eviction policy. We also designed and implemented the necessary middleware `Falcon` to enable the support of many-task computing on clusters, grids and supercomputers. Micro-benchmarks have shown `Falcon` to achieve over 15K+ tasks/sec throughputs, scale to millions of queued tasks, to execute billions of tasks per day, and achieve hundreds of Gb/s I/O rates. `Falcon` has shown orders of magnitude improvements in performance and scalability across many diverse workloads (e.g heterogeneous tasks from milliseconds to hours long, compute/data intensive, varying arrival rates) and applications (e.g. astronomy, medicine, chemistry, molecular dynamics, economic modeling, and data analytics) at scales of billions of tasks on hundreds of thousands of processors across Grids (e.g. TeraGrid) and supercomputers (e.g. IBM Blue Gene/P and Sun Constellation). For more information, please see <http://people.cs.uchicago.edu/iraicu/>.

Biography

Ioan Raicu holds a Ph.D. in Computer Science (CS) from University of Chicago (March 2009) under the guidance of Dr. Ian Foster, a MS in CS from University of Chicago (2005), a MS in CS from Wayne State University (2002), and a BS in CS from Wayne State University (2000). He is a 3-year award winner of the GSRP Fellowship from NASA Ames Research Center. His research work and interests are in the general area of distributed systems. He has defined the new paradigm Many-Task Computing (MTC), which aims to bridge the gap between two predominant paradigms from distributed systems, High-Throughput Computing (HTC) and High-Performance Computing (HPC). His work for the last five years focused on defining and exploring both the theory and practical aspects of realizing MTC across a wide range of large-scale distributed systems. He is particularly interested in efficient task dispatch and execution systems, resource provisioning, data management, scheduling, and performance evaluations in distributed systems. His future work centers on resource management in large scale distributed systems with a focus on many-task computing, data intensive computing, cloud computing, grid computing, and many-core computing. His research was disseminated in over 45 peer-reviewed publications and proposals (receiving over 550 citations summing to an H-index of 12), and over 60 formal talks at international conferences, universities, government labs, and industry events. Some of the notable publication venues are IEEE/ACM SuperComputing/SC, IEEE Grid, IEEE Internet Computing, and the Journal of Grid Computing. His work has been funded by the NASA Ames Research Center Graduate Student Research Program, as well as the DOE Office of Advanced Scientific Computing Research. Ioan has contributed to the broader community service by being a reviewer, program committee, organizing committee or chair in over 40 workshops, conferences, journals, and book chapters. For more information, please see <http://people.cs.uchicago.edu/iraicu/>.