FALKON:

A FAst and Light-weight tasK executiON Framework
http://datasys.cs.iit.edu/projects/Falkon/

Illinois Institute of Technology
Computer Science Department
Data-Intensive Distributed Systems Laboratory

Falkon aims to enable the rapid and efficient execution of many tasks on large compute clusters, and to improve application performance and scalability using novel data management techniques. Falkon combines three techniques to achieve these goals: (1) multi-level scheduling techniques to enable separate treatments of resource provisioning and the dispatch of user tasks to those resources; (2) a streamlined task dispatcher able to achieve order-of-magnitude higher task dispatch rates than conventional schedulers; and (3) performs data caching and uses a data-aware scheduler to leverage the co-located computational and storage resources to minimize the use of shared storage infrastructure. Falkon's integration of multi-level scheduling, streamlined dispatchers, and data management delivers performance not provided by any other system. Falkon has been deployed and tested in a wide range of environments, from 100 node clusters, to Grids (TeraGrid), to specialized machines (SiCortex with 5832 CPUs), to supercomputers (IBM BlueGene/P with 160K CPUs). Micro-benchmarks have shown Falkon to achieve over 15K+ tasks/sec throughputs, scale to millions of queued tasks, and to execute billions of tasks per day. Large-scale applications from many domains have been successfully executed using the Falkon framework. Data diffusion has also shown to improve applications scalability and performance, with its ability to achieve hundreds of Gb/s I/O rates on modest sized clusters, with Tb/s I/O rates on the horizon. Falkon is actively being developed at Illinois Institute of Technology (Professor Ioan Raicu) and the University of Chicago / Argonne National Laboratory (Professor Ian T. Foster and Mike Wilde) with funding from NSF, DOE, and NASA, and has been instrumental in several other proposals to DOE and NSF for additional funding.

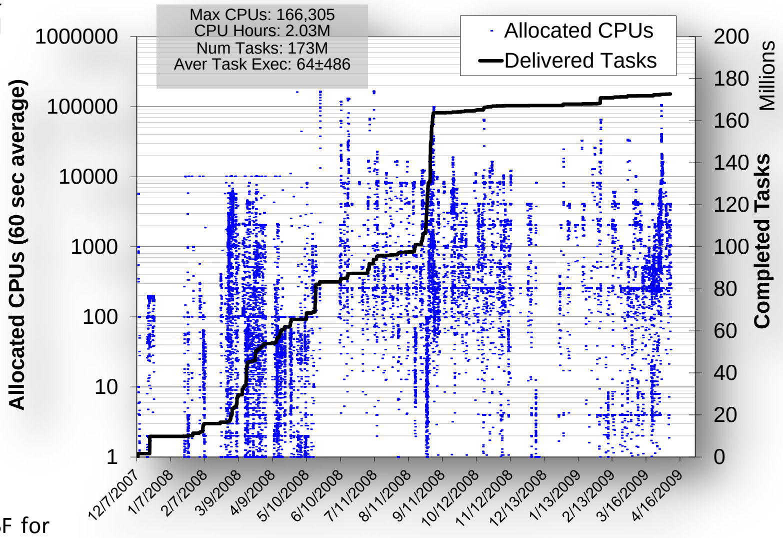

Applications

Field	Description	Characteristics	Status
Astronomy	Creation of montages from many digital images	Many 1-core tasks, much communication, complex dependencies	Experimental
Astronomy	Stacking of cutouts from digital sky surveys	Many 1-core tasks, much communication	Experimental
Biochemistry*	Analysis of mass-spectrometer data for post-translational protein modifications	10,000-100 million jobs for proteomic searches using custom serial codes	In development
Biochemistry*	Protein structure prediction using iterative fixing algorithm; exploring other biomolecular interactions	Hundreds to thousands of 1- to 1,000-core simulations and data analysis	Operational
Biochemistry*	Identification of drug targets via computational docking/screening	Up to 1 million 1-core docking operations	Operational
Bioinformatics*	Metagenome modeling	Thousands of 1-core integer programming problems	In development
Business economics	Mining of large text corpora to study media bias	Analysis and comparison of over 70 million text files of news articles	In development
Climate science	Ensemble climate model runs and analysis of output data	Tens to hundreds of 100- to 1,000-core simulations	Experimental
Economics*	Generation of response surfaces for various economic models	1,000 to 1 million 1-core runs (10,000 typical), then data analysis	Operational
Neuroscience*	Analysis of functional MRI datasets	Comparison of images; connectivity analysis with structural equation modeling, 100,000+ tasks	Operational
Radiology	Training of computer-aided diagnosis algorithms	Comparison of images; many tasks, much communication	In development
Radiology	Image processing and brain mapping for neuro-surgical planning research	Execution of MPI application in parallel	In development

Note: Asterisks indicate applications being run on Argonne National Laboratory's Blue Gene/P (Intrepid) and/or the TeraGrid Sun Constellation at the University of Texas at Austin (Ranger).