CS554 Project Ideas

NET:MPSim – Improving Network Throughput through Multipath Network Routing Simulations

Overview

In order to explore potential architectures for future exascale HPC systems, the research community has turned to simulation in order to study the impact of potential architecture decisions. Massively parallel discrete-event (PDE) simulations are a particularly powerful tool for simulations of this size. We will focus on simulation of ultra-large-scale high-end computing systems and their network topologies (e.g. fat-tree, torus, dragonfly) using Rensselaer`s Optimistic Simulator System (ROSS). ROSS is a high-performance, low memory overhead, massively parallel discrete-event simulator. With ROSS, complex network models achieve a peak event rate of 1.3 billion events per second, making their simulation at large-scale tractable. The ROSS simulator has already been validated at small scales up to 2K-cores for a 3D Torus network topology and static single-path routing. The Torus topology was studied up to 1B-nodes for static single-path routing. There have already been studies about the dragonfly network model up to 1M-node scales for static single-path routing.

This work will extend the ROSS simulations by implementing a third network topology, the fat-tree, in addition to the already defined torus and dragonfly. The dynamic multi-path routing protocol will be implemented in the ROSS simulator, and a detailed evaluation study will be done to compare the three different network topologies with different routing strategies (static single-path and dynamic multi-path routing) at various scales (1K-nodes to 1M-nodes), in both single-path and multi-path routing scenarios. We will compare real communication delays, achieved throughput, and load balance across the communication fabric between our simulation models and the actual torus network from the BlueGene/Q supercomputer at Argonne National Laboratory.

Relevant Systems and Reading Material

- Misbah Mubarak, Christopher D Carothers, Robert Ross, and Philip Carns. Modeling a million-node dragonfly network using massively parallel discrete-event simulation. In High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion:, pages 366-376. IEEE, 2012
- T. Anjali, A. Fortin, G. Calinescu, S. Kapoor, N. Kirubanandan, and S. Tongngam. Multipath network flows: Bounded buffers and jitter. In INFOCOM, 2010 Proceedings IEEE, pages 1-7, march 2010
- F. Devetak, J. Shin, T. Anjali, and S. Kapoor. Minimizing path delay in multipath networks. In Proceedings of IEEE ICC 2011, 2011
- Tricha Anjali Junghwan Shin, Fabrizio Devetak and Sanjiv Kapoor. Delay variance optimization in multipath routing networks. GLOBECOM, 2013
- Junghwan Shin Mohammad Sarwat and Sanjiv Kapoor. Utility optimization in multipath routing cdn. ICC, 2013

Preferred/Required Skills

- Programming language choice: C/C++
- Skills/knowledge: Network protocols

Performance Metrics

Throughput, Latency, load balancing; implementation should be done in the ROSS/CODES simulator, and evaluations are expected to tbe done for both single static path and dynamic multi-path protocols at up to 1M node scales on 3D, 4D, and 5D torus network interconnects.

Project Mentor

Ioan Raicu, iraicu@cs.iit.edu