CS554 Project Ideas

ZHT:Mon: Investigation of the Communication Overheads for Distributed Monitoring with Aggregation Tree

Overview

A distributed monitoring framework serves as an important building block for large-scale data aggregation and continuous event monitoring applications, such as cluster and grid resource monitoring. The underlying aim is to provide a global view of information in the system at a reasonable cost and within a specified precision bound.

The main activities involved in maintaining a global view include one to all multicast communication and all to one reduction communication. The communication is usually based on a hierarchical design where an aggregation tree of communicating components is built, and the edges represent the data flow. The multicast communication is from the root spreading the information all the way down to all the leaf nodes, while the reduction communication is aggregating the information in all the leaf nodes to the root. The basic problem is to find the optimal fan out and height of the tree to achieve the best (smallest) latency for large-scale systems; the more extensive problem is about resiliency, which means how to rebuild the tree efficiently when node failures happen.

Relevant Systems and Reading Material

Paper: MRNet: A Software-Based Multicast/Reduction Network for Scalable Tools. Access at: ftp://ftp.cs.wisc.edu/paradyn/papers/Roth03MRNet.pdf
ZHT: http://datasys.cs.iit.edu/publications/2013_IPDPS13_ZHT.pdf

Methodology

Option 1: Simulation the aggregation tree up to Exascale with millions of nodes.

Option 2: Implementation of a real system at moderate scales (e. g. thousands of nodes).

Preferred/Required Skills

Required: Java for Simulation; C/C++, sockets and multi-threading for real implementation

Preferred: TCP/IP, algorithms to operate tree (add/delete nodes)

Parameters

fan out, tree height, and failure rate

Performance Metrics

latency, speedup (vs flat or centralized pattern)

Project Mentor

Ke Wang <kwang22@hawk.iit.edu>
Tonglin Li <tli13@hawk.iit.edu>