MRNet: Multicast/Reduction Network

Efficient Communication for Large Scale Systems

MRNet Overview

Overlay network for efficient data multicasts and reductions.

Tree topology for distribution of data analysis and management.

Customizable components for use in diverse tool applications.

MRNet Features

Flexible process-tree topologies.

Dynamic loading of user-defined data aggregations.

Built-in filters for elementary aggregations.

High-throughput data transfers.

Efficient multicasts for scalable control.

Concurrent data channels with different aggregation schemes.

Open-source licensing.

MRNet Improves Tool Performance by Distributing Data Processing and Analyses

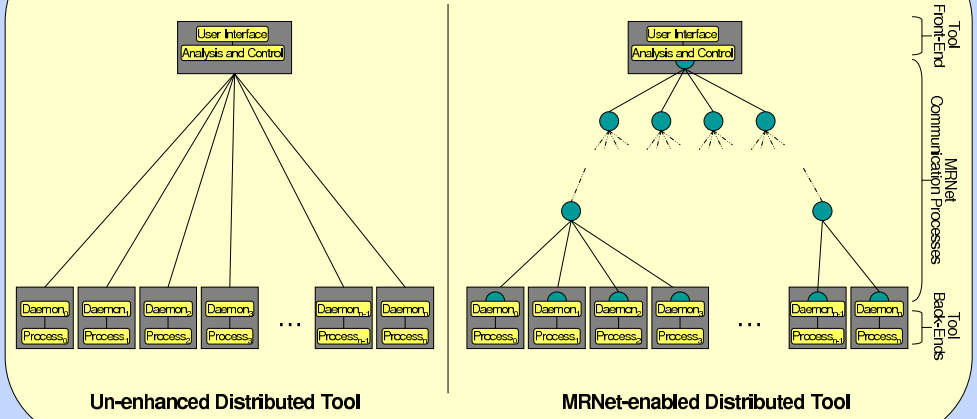

Sample Performance

Improving the Performance of Paradyn's Data Analysis

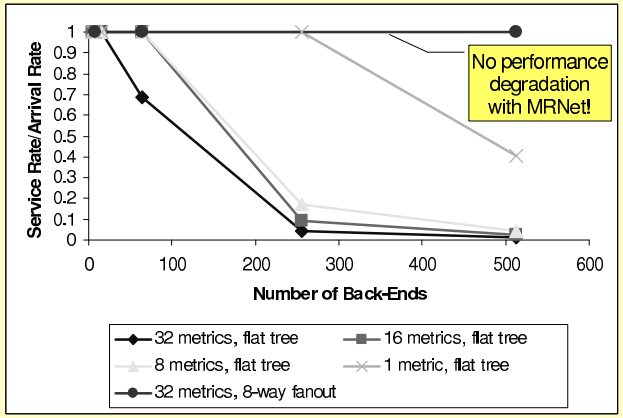

Microbenchmark Results

- ASCI Blue Pacific:
  - 256x4 PowerPC
- 512 back-ends
  - ~64 byte packet sizes
- 8-way vs. Flat Tree
  - Startup Latency:
    - 13.9 Speedup
  - Round-Trip Latency:
    - 13.4 Speedup
- Reduction Throughput
  - 8-way: 55.23 ops/sec.
  - Flat tree: 4.03 ops/sec (256 back-ends).

MRNet API

C++ based API for integration into tool front-end and daemons.

End-points, retrieved from instantiated networks, are grouped into communicators.

Streams are virtual channels that bind the end-points in a communicator for group communication.

Streams are created with filters that perform specified data aggregations on the data flow.

Tool-specific filters can be loaded into MRNet for custom aggregations.

MRNet Filters

Native MRNet Filters

MIN, MAX, SUM, AVG

Concatenation

Custom Paradyn Filters

Clock synchronization.

Equivalence classes for application data.

Time-based data synchronization.

Status

Release v1.1:

November 2004

Tested Platforms:

Linux, Aix, Solaris

Additional Information:

www.paradyn.org/mrnet/