

PATHSCALE OPTIPATH™ MPI ACCELERATION TOOLS

DRAFT: PRELIMINARY



TUNE MPI APPLICATIONS ON YOUR CLUSTER WITHOUT BECOMING A MPI EXPERT

The PathScale OptiPath™ MPI Acceleration Tools represent the simplest way to identify the root causes preventing an application from scaling on a cluster. Using OptiPath, a novice or a seasoned MPI expert is guided to the most important problems affecting scalability in their applications. The tools present these problems in ranked order, with supporting information. Most importantly, OptiPath makes specific recommendations how to improve the application performance and scalability.

IMPROVE PROGRAMMER PRODUCTIVITY

Seasoned MPI developers have learned to carefully instrument their code with performance counters, catalog performance numbers, plot charts, and tweak their application to get better MPI performance. OptiPath automates this process. With a few key-strokes one can change the data, modify the command line, and vary the cluster size. This is all performed within a managed framework that lets the user plot performance and drill down to the source-code level, showing the "before" and "after" effects.

PRODUCT HIGHLIGHTS

The OptiPath MPI Acceleration Tools include two major components, a Trace Collector which runs on any MPI node, and an Experiment Manager and Analyzer which both run on the user's workstation. Together these automate the setup and running of the MPI application, and the analysis of the results. No recompilation or re-linking is required to use the collector, and it is compatible with batch schedulers such as Torque/PBS, or Sun Grid Engine. The Experiment Manager makes it easy to examine run history, and to re-run complex series of experiments at a single click.

EXPERIMENT-ORIENTED APPROACH

The PathScale OptiPath MPI Acceleration Tools use a guided problem/solution approach. The tools are experiment oriented and capable of analyzing application sensitivity to multiple factors. These factors include different cluster sizes, different data sets, and the before/after effects of application changes. The tools collect extensive performance information in a light-weight fashion and stored it in a database for analysis and comparison to other runs.

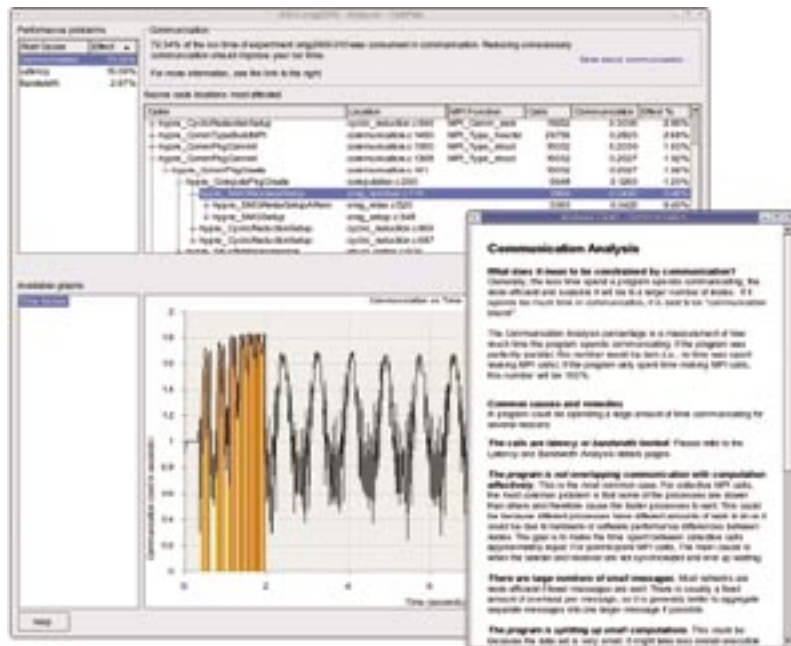
OptiPath allows the user to select intuitive views of an "Experiment" or individual run of an MPI application. The user can also display views of a set of experiments highlighting scalability analysis, which visually reinforce conclusions reached. Additionally, the tools display traces of MPI calls across all nodes, each one annotated with source line information, call context, and impact on overall performance.

OPTIPATH KEY BENEFITS

- Profiles the running MPI application in a light-weight fashion
- Saves the results for post-mortem analysis
- Presents you with a ranked list of performance problems
- Displays graphics indicating when a problem occurs, and how much impact it has on performance
- Pinpoints to the source-code level where the problem occurs
- Explains why such problems frequently occur, and gives suggestions for fixing them

STATE OF THE ART IN MPI ANALYSIS

The PathScale OptiPath MPI Acceleration Tools guide the user to the root causes of performance problems in MPI applications. The tools offer a unique approach to performance analysis, giving source-level details of problems, backed up with graphics, information, and context. A non-expert in MPI can quickly address each problem and return to useful work.



PATHSCALE OPTIPATH™ MPI ACCELERATION TOOLS

DRAFT: PRELIMINARY

THE PATHSCALE OPTIPATH MPI ACCELERATION TOOLS INCLUDE:

DATA COLLECTOR

- Gathers trace data, adding less than 10% overhead during collection, zero at other times
- Command line or GUI driven
- Performs instrumentation, saves results to a database
- Gathers trace data for MPI and interconnect events

EXPERIMENT MANAGER

- Runs the experiment
- Easily examine experiment history
- Re-run complex series of experiments at a single click
- Command line or GUI driven

ANALYZER

- Analyzes data gathered over several collector runs
- Presents ranked hypotheses with potential root cause problems
- Points you to the place in code which is the source of the problem
- Suggests root causes of the problems and how to fix them

ANALYSIS LIFE CYCLE

Using the Experiment Manager:

1. Set up a problem size range and parameters i.e. cluster sizes from 1 to 128, in powers of 2
2. Launch the jobs via the Experiment Manager
3. Wait for the completion of the jobs
4. Start Analyzer and examine the results
5. Make changes as suggested

Start over from step 2; see if changes made a difference

HARDWARE ENVIRONMENT

- AMD AMD64™ and Intel® EM64T processors in 32-bit (x86) and 64-bit (x86-64) modes
- Users can analyze on one architecture, and run on another

INTERCONNECT AWARE

- Support for Fast and Gigabit Ethernet, PathScale InfiniPath™ HTX™ Adapter, and InfiniBand™ host adapters and switches

SOFTWARE ENVIRONMENT

- Red Hat™ Enterprise, Red Hat Fedora Core, SuSE™ Enterprise, and SuSE Professional Linux
- Torque/PBS or Sun Grid Engine batch scheduling systems
- Works with standard MPICH
- Works with C/C++/FORTRAN 77/95
- Also works with other compilers

USER INTERFACE

- Task-based design ensures ease of use
- Designed for the scientist, not just the MPI expert
- GUI and command line interfaces supplied

PERFORMANCE AND SCALABILITY

- No designed-in collector scalability limits
- Less than 10% performance impact during collection

LICENSE MODEL

- Per-cpu license for Collector
- Per-user license for Analyzer and Experiment Manager
- EKOPath Compiler Suite license included

OPTIPATH PRICING AND CONFIGURATION

PathScale provides a variety of configurations to suit any production or development cluster. Please contact sales@pathscale.com for pricing and availability.



PATHSCALE, INC.
2071 STIERLIN COURT, SUITE 200
MOUNTAIN VIEW CA 94043 USA
TEL 650.934.8100
FAX 650.428.1969
PATHSCALE.COM
SALES@PATHSCALE.COM

© 2005 PathScale, Inc. All rights reserved. PathScale, OptiPath and Accelerating Cluster Performance are trademarks of PathScale, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.