Performance Analysis for Scalable MPI/OpenMP Applications

The Department of Energy's Accelerated Strategic Computing Initiative (ASCI) requires performance analysis capabilities and scalability in software tool development to keep pace with the rapid improvements in processing power of ASCI computer platforms. To accelerate the development and commercialization of these software technologies, ASCI has joined forces with Intel Corporation's KAI Software Lab and Pallas GmbH in the ASCI PathForward Parallel System Performance Project (PSP). The resulting software tools combine proven features of Vampir and Vampirtrace for MPI analysis, and the KAP/Pro Toolset for OpenMP analysis into the new PSP toolset called VGV (Vampir/GuideView). The assimilation of MPI/OpenMP software-performance-analysis tools has introduced the following new capabilities:

- Scalability to thousands of processors
- Integration of MPI and OpenMP analysis
- Specific support for clustered SMP systems
- Integration of hardware performance data
- Support for top–down performance-analysis experiments

Versatility and Portability

The VGV tools are supported on the IBM, HP-Compaq, IA32 Linux, and Cplant ASCI systems with the same user interface and core capabilities. Applications written in Fortran 90, C, and C++ are all supported. For object-oriented applications, specific tracing of individual objects or a set of objects is possible like in the TAU system. The objects can share user-defined characteristics.

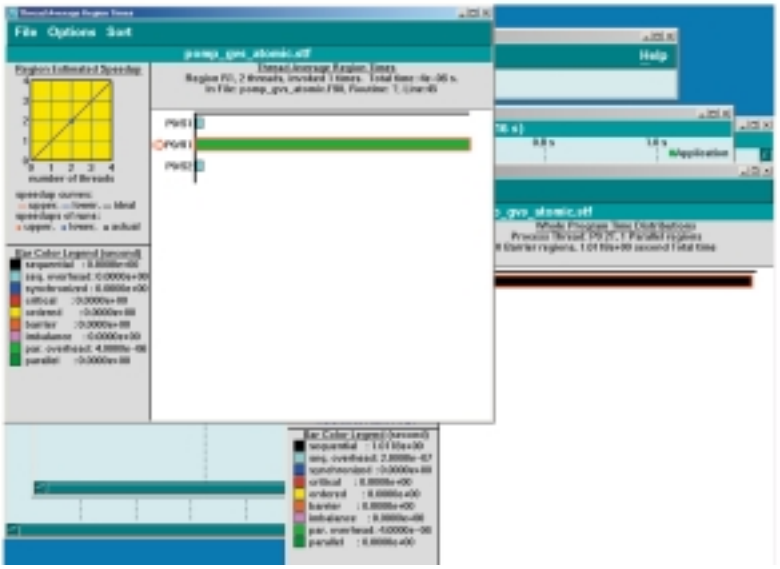

Figure 1. Vampir window display for an OpenMP application showing the Thread Average Region Times: the associated OpenMP application was instrumented with POMP.

New VGV Features for 2002

The Guide OpenMP compilation system has been extended to support the OpenMP performance introspection interface known as POMP. Using this interface, performance tools can interact with OpenMP compilation and runtime systems in a standardized and portable way without having to tackle vendor compatibility issues (Figure 1). The POMP interface is designed to be flexible, supporting a spectrum of performance data at several levels of granularity.

Vampir now contains a statistics viewer that displays profiling statistics derived from the precomputed statistics in the structured trace file (STF) database (Figure 2). The new color–coded time line display (Figure 3) improves scalability in the process dimension by mapping time–dependent performance data to a color gamut. It

Figure 2. OpenMP Region statistics for sweep3d: the execution time of OpenMP regions is severely unbalanced across processes.

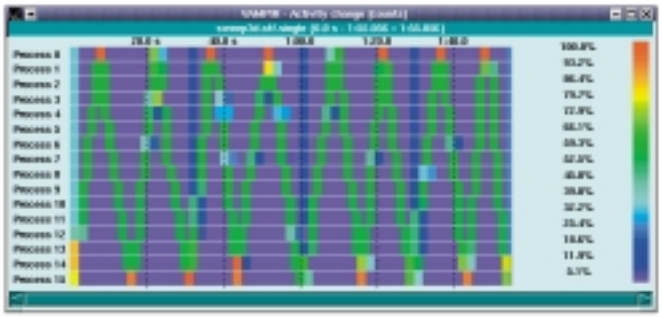

Figure 3. Color–coded timeline: the colors indicate the frequency of routine calls over time – red highlights areas with many short routines, blue shows places where routines take a lot of time, e.g. by being blocked waiting for MPI communication.

also provides filter functionality with the coloring scheme highlighting where performance metrics fall below user–defined thresholds or exceed them.

Vampirtrace introduces the trace agent to assist a user in tailoring the data acquisition options to limit the trace information to the amount for identifying a performance problem:

1. With an initial configuration, the agent lets the user generate execution statistics in a very compact trace database.
2. Using the collected statistics, the agent tries to identify performance problems and then recommends options for additional experiments, collecting more detailed event data. In turn, upon analyzing this data, the agent directs the user towards detecting a bottleneck, or towards verifying that performance is okay.

The GuideView displays within Vampir provide an intelligent filter that presents the relevant and important OpenMP performance information. The critical performance issues for OpenMP are first sorted, and then displayed in a hierarchical manner. Two analysis capabilities have been implemented through an expert system to provide the following hierarchical functionality:

1. Identify the OpenMP performance problems according to a rule set.
2. Sort the performance information to emphasize important problems and tuning tips.

What's Next?

We are launching Phase Four of the project and it will bring several improvements to the VGV tools. These include: a tool architecture that leverages parallelism for the analysis process, new capabilities for the trace agent and integration with the statistics viewer, support for analyzing the VGV performance data from scripting languages, definition of execution scopes for analysis of code components or libraries, and enhancements in the handling of hardware performance data. In addition, the VGV project will contribute to the definition of MPI and OpenMP performance introspection interfaces, and prototype support for them.

Further Information

For more information about the ASCI PathForward PSP Project contact: Jeffrey Vetter, vetter3@llnl.gov; Bob Kuhn, Bob.Kuhn@intel.com; or Hans-Christian Hoppe, Hans-Christian.Hoppe@pallas.com.

Useful Links

PAPI performance counter interface
(see http://icl.cs.utk.edu/projects/papi)
TAU instrumentation and performance analysis system
(http://www.cs.uoregon.edu/research/paracomp/tau/)
TotalView Debugger *(http://www.etnus.com)*
Vampir and Vampirtrace tools
(http://www.pallas.com/e/products/vampir/index.htm)
POMP OpenMP Performance Interface
(http://www.fz-juelich.de/zam/kojak/opari/docs/pomp-talk-ewomp02.ppt)