

Scalasca

Introduction

Scalasca is a software tool that supports the performance optimization of parallel programs by measuring and analyzing their runtime behavior. The analysis identifies potential performance bottlenecks – in particular those concerning communication and synchronization – and offers guidance in exploring their causes.

Scalasca supports profiling of MPI, OpenMP and hybrid MPI+OpenMP applications.

Installed versions

There are currently two versions of Scalasca 2.0 modules installed on Anselm:

- class="s1"> scalasca2/2.0-gcc-openmpi, for usage with GNU Compiler and OpenMPI,
- class="s1"> scalasca2/2.0-icc-impi, for usage with Intel Compiler and Intel MPI.

Usage

Profiling a parallel application with Scalasca consists of three steps:

1. Instrumentation, compiling the application such way, that the profiling data can be generated.
2. Runtime measurement, running the application with the Scalasca profiler to collect performance data.
3. Analysis of reports

Instrumentation

Instrumentation via "scalasca -instrument" is discouraged. Use Score-P instrumentation.

Runtime measurement

After the application is instrumented, runtime measurement can be performed with the "scalasca -analyze" command. The syntax is :

`scalasca -analyze [scalasca options][launcher] [launcher options] [program] [program options]`

An example :

```
$ scalasca -analyze mpirun -np 4 ./mympprogram
```

Some notable Scalasca options are:

- t Enable trace data collection. By default, only summary data are collected.
- e <directory> Specify a directory to save the collected data to. By default, Scalasca saves the data to a directory with prefix >scorep_, followed by name of the executable and launch configuration.

Scalasca can generate a huge amount of data, especially if tracing is enabled. Please consider saving the data to a scratch directory.

Analysis of reports

For the analysis, you must have Score-P and CUBE modules loaded. The analysis is done in two steps, first, the data is preprocessed and then CUBE GUI tool is launched.

To launch the analysis, run :

```
scalasca -examine [options] <experiment_directory>
```

If you do not wish to launch the GUI tool, use the “-s” option :

```
scalasca -examine -s <experiment_directory>
```

Alternatively you can open CUBE and load the data directly from here. Keep in mind that in that case the preprocessing is not done and not all metrics will be shown in the viewer.

Refer to CUBE documentation on usage of the GUI viewer.

References

1. <http://www.scalasca.org/>