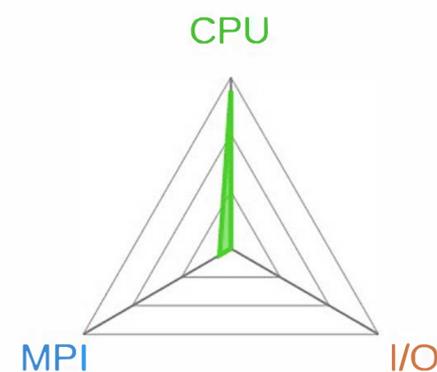
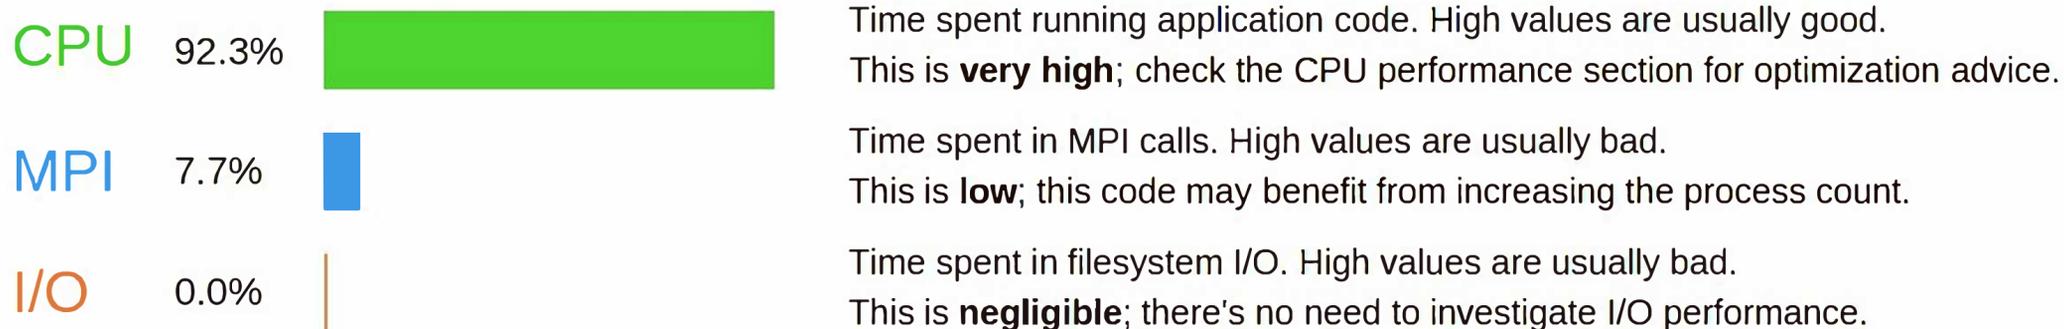


Executable: xhpl
Resources: 256 processes, 16 nodes
Machine: cray-two
Start time: Fri Nov 1 18:01:20 2013
Total time: 169 seconds (3 minutes)
Full path: /users/allinea/hpl/hpl-2.1/bin/
Notes: 256 blocks, P=16, Q=16



Summary: xhpl is CPU-bound in this configuration

The total wallclock time was spent as follows:



This application run was **CPU-bound**. A breakdown of this time and advice for investigating further is in the **CPU** section below. As very little time is spent in **MPI calls**, this code may also benefit from running at larger scales.

CPU

A breakdown of how the 92.3% total CPU time was spent:



The per-core performance is **FPU-bound**. Try to increase the amount of time spent in **vectorized instructions** by analyzing the compiler's vectorization reports.

Significant time is spent on **memory accesses**. Use a profiler to identify time-consuming loops and check their cache performance.

I/O

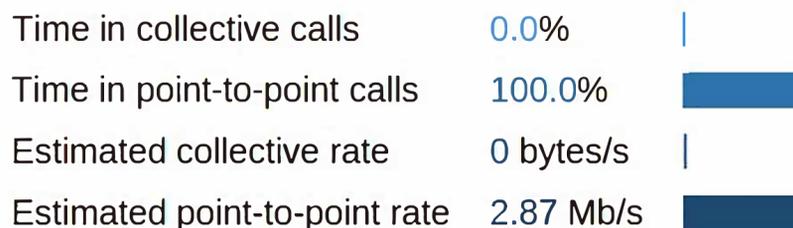
A breakdown of how the 0.0% total I/O time was spent:



No time is spent in **I/O operations**. There's nothing to optimize here!

MPI

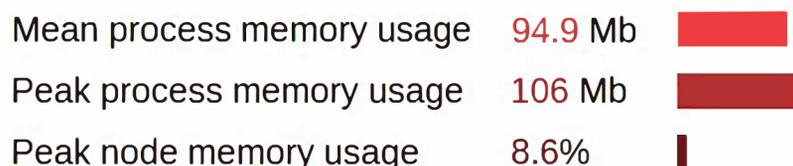
Of the 7.7% total time spent in MPI calls:



The **point-to-point** transfer rate is low. This can be caused by inefficient message sizes, such as many small messages, or by imbalanced workloads causing processes to wait. Use an MPI profiler to identify the problematic calls and ranks.

Memory

Per-process memory usage may also affect scaling:



The **peak node memory usage** is low. You may be able to reduce the total number of CPU hours used by running with fewer MPI processes and more data on each process.