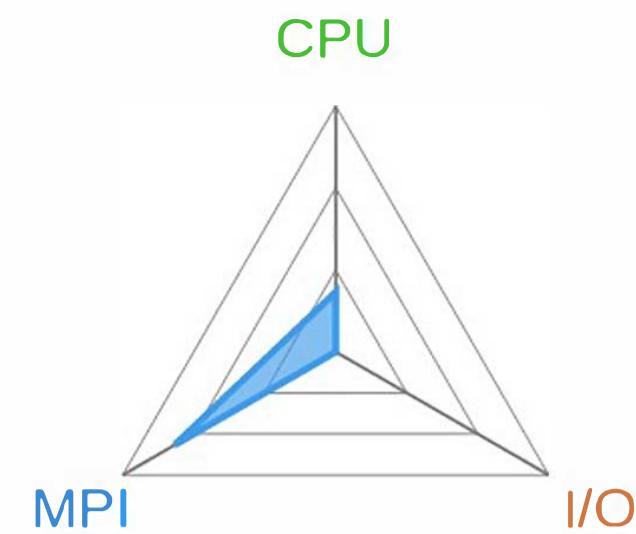
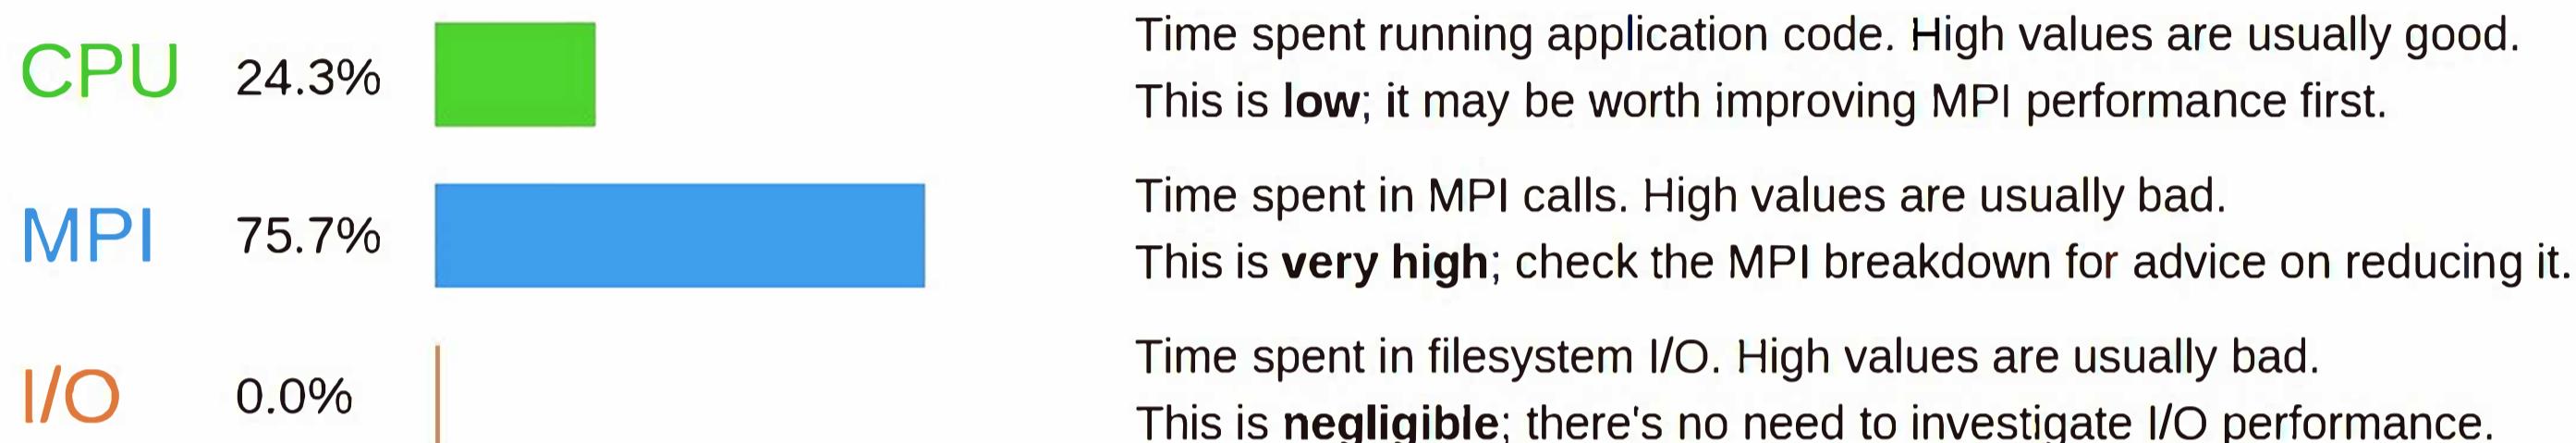


Executable: xhpl
 Resources: 256 processes, 16 nodes
 Machine: cray-two
 Start time: Thu Oct 31 09:57:52 2013
 Total time: 468 seconds (8 minutes)
 Full path: /users/allinea/hpl/hpl-2.1/bin/
 Notes: 64 blocks, P=8, Q=8



Summary: xhpl is MPI-bound in this configuration

The total wallclock time was spent as follows:



This application run was **MPI-bound**. A breakdown of this time and advice for investigating further is in the [MPI](#) section below.

CPU

A breakdown of how the **24.3%** total CPU time was spent:

Scalar numeric ops	33.4%	<div style="width: 33.4%; background-color: #90EE90;"></div>
Vector numeric ops	32.2%	<div style="width: 32.2%; background-color: #90EE90;"></div>
Memory accesses	34.4%	<div style="width: 34.4%; background-color: #4682B4;"></div>
Other	0.0	<div style="width: 0.0%; background-color: #E69138;"></div>

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.

MPI

Of the **75.7%** total time spent in MPI calls:

Time in collective calls	99.1%	<div style="width: 99.1%; background-color: #4682B4;"></div>
Time in point-to-point calls	0.9%	<div style="width: 0.9%; background-color: #E69138;"></div>
Estimated collective rate	0 bytes/s	<div style="width: 0%; background-color: #E69138;"></div>
Estimated point-to-point rate	282 Mb/s	<div style="width: 282%; background-color: #4682B4;"></div>

Most of the time is spent in **collective calls** with a very low **transfer rate**. This suggests a significant load imbalance is causing synchronization overhead. You can investigate this further with an **MPI profiler**.

I/O

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

Time in reads	0.0%	<div style="width: 0.0%; background-color: #E69138;"></div>
Time in writes	0.0%	<div style="width: 0.0%; background-color: #E69138;"></div>
Estimated read rate	0 bytes/s	<div style="width: 0%; background-color: #E69138;"></div>
Estimated write rate	0 bytes/s	<div style="width: 0%; background-color: #E69138;"></div>

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

Memory

Per-process memory usage may also affect scaling:

Mean process memory usage	69.5 Mb	<div style="width: 69.5%; background-color: #FF8C00;"></div>
Peak process memory usage	266 Mb	<div style="width: 266%; background-color: #A52A2A;"></div>
Peak node memory usage	16.2%	<div style="width: 16.2%; background-color: #A52A2A;"></div>

There is significant variation in **memory usage** between processes. This may be a sign of workload imbalance or a memory leak.

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.