

linaroforge

Debugging and Optimizing Parallel Codes with Linaro Forge

Rudy Shand
Field Application Engineer



Agenda

- 10:00 - 12:00 Session 1: Ensuring Program Correctness with Linaro DDT
- 13:00 - 15:00 Session 2: Performance Engineering with Linaro Performance Reports and Linaro MAP

Bug classification

- Crashes
 - One or more processes in application terminates
 - Most common and generally easiest to solve
- Hangs
 - Deadlocks - Stuck waiting for something that never happens
 - Livelocks - Making local progress, but no global progress
- Race conditions
 - One or more threads accessing the same data at the same time in non deterministic way
 - Shows up as incorrect answer or sometimes crashes



Linaro Forge

An interoperable toolkit for debugging and profiling



The de-facto standard for HPC development

- Most widely-used debugging and profiling suite in HPC
- Fully supported by Linaro on Intel, AMD, Arm, Nvidia, AMD GPUs, etc.



State-of-the art debugging and profiling capabilities

- Powerful and in-depth error detection mechanisms (including memory debugging)
- Sampling-based profiler to identify and understand bottlenecks
- Available at any scale (from serial to exascale applications)



Easy to use by everyone

- Unique capabilities to simplify remote interactive sessions
- Innovative approach to present quintessential information to users

Linaro Forge: Where Most of Top Supercomputers turn for Performance Excellence

Build reliable and optimized code on multiple Server and HPC architectures

Linaro Forge combines



Linaro DDT

Market leading, simple to use HPC debugger for C/C++, Fortran and Python applications.



Linaro MAP

Effortless performance analysis for experts and novices alike.




Linaro Performance Reports

At a glance, single-page, application performance summary.

Performance Engineering for any architecture, at any scale

Supported Platforms



DDT UI

- 1 Process controls
- 2 Process groups
- 3 Source Code view
- 4 Variables
- 5 Evaluate window
- 6 Parallel Stack
- 7 Project files
- 8 Find a file or function

DDT UI components and their locations:


- 1 Process controls:** Located in the top toolbar, including buttons for running, pausing, and other process management actions.
- 2 Process groups:** Located in the top panel, showing a list of process groups (All, Group 1, Group 2) with their respective process counts and status.
- 3 Source Code view:** Located in the middle-right panel, showing the source code of the selected file (hello.c) with line numbers.
- 4 Variables:** Located in the middle-right panel, showing a list of variables (argc, argv, beingWatched, bigArray, dest, dynamicArray, environ, i, message, my_rank, p, source, status, t2, tables, tag, test, x, y) and their values.
- 5 Evaluate window:** Located in the bottom-right panel, showing a table of variables and their values (bigArray[3], my_rank, x + y).
- 6 Parallel Stack:** Located in the bottom-left panel, showing a list of processes and functions (Processes, Function).
- 7 Project files:** Located in the middle-left panel, showing a tree view of project files (Application Code, Sources).
- 8 Find a file or function:** Located in the middle-left panel, showing a search bar (Search (Ctrl+K)).

Linaro DDT Debugger Highlights

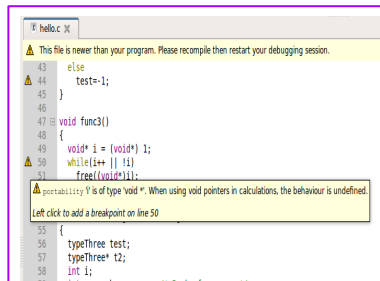
Input/Output	Breakpoints	Watchpoints	Tracepoints	Tracepoint Output	Stacks (All)	
Tracepoint Output						
Tracepoint	Process					Values logged
show 90.05	970, rank0 12, 14, 17, 20, 23, 12	myip	2170-3527	pos	3-40	mod pay
show 90.01	900, rank0 12, 14, 17, 20, 23, 12	no	1	1max	pos	
show 90.05	940, rank0 12, 14, 17, 20, 23, 12	myip	2170-3527	pos	3-40	mod pay
show 90.01	920, rank0 12, 14, 17, 20, 23, 12	no	1	1max	pos	
show 90.05	910, rank0 12, 14, 17, 20, 23, 12	myip	2170-3527	pos	3-40	mod pay
show 90.01	900, rank0 12, 14, 17, 20, 23, 12	no	1	1max	pos	
show 90.05	880, rank0 12, 14, 17, 20, 23, 12	myip	2170-3527	pos	3-40	mod pay
show 90.01	860, rank0 12, 14, 17, 20, 23, 12	no	1	1max	pos	
show 90.05	840, rank0 12, 14, 17, 20, 23, 12	myip	2170-3527	pos	3-40	mod pay
show 90.01	820, rank0 12, 14, 17, 20, 23, 12	no	1	1max	pos	

The scalable print alternative

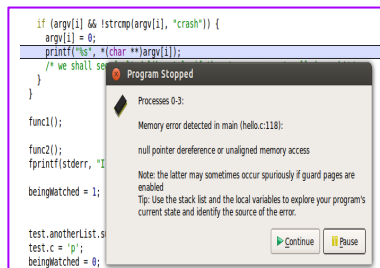
The scalable print alternative




Stop on variable change



Static analysis warnings on code errors



Detect read/write beyond array bounds




Detect stale memory allocations

Core files

You can open and debug one or more core files generated by your application.

Procedure


1. On the Welcome page click **Open Core Files**. The **Open Core Files** window opens.



2. Select an executable and a set of core files, then click **OK** to open the core files and start debugging them.


Note

While Linaro DDT is in this mode, you cannot play, pause, or step, because there is no process active. You are, however, able to evaluate expressions and browse the variables and stack frames saved in the core files.


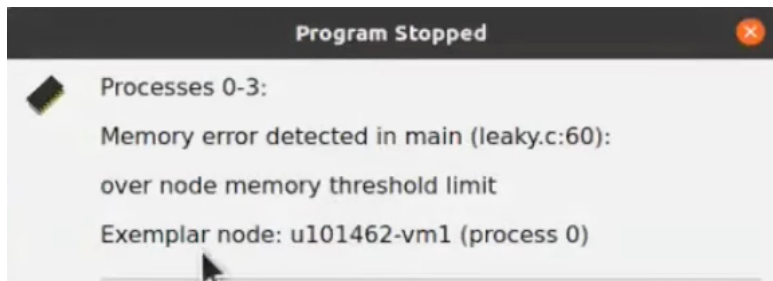


- View core files for CPU's
- View core files for GPU's

Memory debugging menu in Linaro DDT



When manual linking is used,
untick "Preload" box



Multi-dimensional Array Viewer

What does your data look like at runtime?

View arrays


- On a single process
- Or distributed on many ranks

Use metavariables to browse the array

- Example: \$i and \$j
- Metavariables are unrelated to the variables in your program
- The bounds to view can be specified
- Visualise draws a 3D representation of the array

Data can also be filtered

- “Only show if”: \$value>0 for example \$value being a specific element of the array



DDT: Production-scale debugging

Isolate and investigate faults at scale

Who misbehaved?

- Merge stacks from processes and threads
- Sparklines comparing data across processes
- Which MPI rank

Where is the problem?



- Integrated source code editor
- Dynamic data structure visualization

How did it happen?

- Parse diagnostic messages
- Trace variables through execution


Why did it happen?

- Unique “Smart Highlighting”
- Experiment with variable values



Python Debugging

- Debug Features
 - Sparklines for Python variables
 - Tracepoints
 - MDA viewer
 - Mixed language support
- Improved Evaluations:
 - Matrix objects
 - Array objects
 - Pandas DataFrame
 - Series objects
- Python Specific:
 - Stop on uncaught Python exception
 - Show F-string variables
 - Mpi4py, NumPy, SciPy



ddt --connect mpirun -n 8 python3
%allinea_python_debug% ./mmult.py

Debugging Nvidia GPUs

Using Linaro DDT

Debug code simultaneously on the A100 GPU and the CPU


Controlling the GPU execution:

- All active threads in a warp will execute in lockstep. Therefore, DDT will step 32 threads at a time.
- Play/Continue runs all GPU threads
- Pause will pause a running kernel

Key (additional) GPU features:

- Kernel Progress View
- GPU thread in parallel stack view
- GPU Thread Selector
- GPU Device Pane

For NVIDIA's nvcc compiler, kernels must be compiled with the -g and -G flags



Run DDT in offline mode

Run the application under DDT and halt or report when a failure occurs

You can run the debugger in non-interactive mode

- For long-running jobs / debugging at very high scale
- For automated testing, continuous integration...

To do so, use following arguments:

- `$ ddt --offline --output=report.html mpirun ./jacobi_omp_mpi_gnu.exe`
 - **--offline** enable non-interactive debugging
 - **--output** specifies the name and output of the non-interactive debugging session
 - Html
 - Txt
 - Add **--mem-debug** to enable memory debugging **and memory leak detection**

```
ddt --offline -o jacobi_omp_mpi_gnu_debug.txt \  
--trace-at _jacobi.F90:83,residual \  
mpirun ./jacobi_omp_mpi_gnu.exe
```

Report output

12		0:08.188	0-3	Process stopped at breakpoint in update (wave.c:216).																																																																								
13				Additional Information ▼ Stacks <table><thead><tr><th>Processes</th><th>Threads</th><th>Function</th><th>Source</th><th>Variables</th></tr></thead><tbody><tr><td>0-3</td><td>4</td><td>main (wave.c:334)</td><td>► iterations = update(left, right);</td><td>► Rank 0, thread 1</td></tr><tr><td>0-3</td><td>4</td><td>update (wave.c:216)</td><td>► values[j] = newval[j];</td><td>▼ Rank 0, thread 1</td></tr><tr><td></td><td></td><td></td><td></td><td><table><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>i</td><td>0</td></tr><tr><td>iterations</td><td>1</td></tr><tr><td>j</td><td>101</td></tr><tr><td>left</td><td>-2 (from -2 to 2)</td></tr><tr><td>now</td><td><aggregate value></td></tr><tr><td>right</td><td>1 (from -2 to 3)</td></tr><tr><td>stop</td><td>0</td></tr></tbody></table></td></tr><tr><td>0-3</td><td>8</td><td>progress_engine</td><td></td><td></td></tr><tr><td>0-3</td><td>8</td><td>opal_libevent2022_event_base_loop (event.c:1630)</td><td></td><td>► Rank 0, thread 2</td></tr><tr><td>0-3</td><td>4</td><td>poll_dispatch (poll.c:165)</td><td></td><td>► Rank 0, thread 2</td></tr><tr><td>0-3</td><td>4</td><td>poll</td><td></td><td></td></tr><tr><td>0-3</td><td>4</td><td>epoll_dispatch (epoll.c:407)</td><td></td><td>► Rank 0, thread 3</td></tr><tr><td>0-3</td><td>4</td><td>epoll_wait</td><td></td><td></td></tr></tbody></table> Current Stack ▼ Evaluate <table><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>3*j*j</td><td>30603</td></tr><tr><td>j</td><td>101</td></tr></tbody></table>	Processes	Threads	Function	Source	Variables	0-3	4	main (wave.c:334)	► iterations = update(left, right);	► Rank 0, thread 1	0-3	4	update (wave.c:216)	► values[j] = newval[j];	▼ Rank 0, thread 1					<table><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>i</td><td>0</td></tr><tr><td>iterations</td><td>1</td></tr><tr><td>j</td><td>101</td></tr><tr><td>left</td><td>-2 (from -2 to 2)</td></tr><tr><td>now</td><td><aggregate value></td></tr><tr><td>right</td><td>1 (from -2 to 3)</td></tr><tr><td>stop</td><td>0</td></tr></tbody></table>	Name	Value	i	0	iterations	1	j	101	left	-2 (from -2 to 2)	now	<aggregate value>	right	1 (from -2 to 3)	stop	0	0-3	8	progress_engine			0-3	8	opal_libevent2022_event_base_loop (event.c:1630)		► Rank 0, thread 2	0-3	4	poll_dispatch (poll.c:165)		► Rank 0, thread 2	0-3	4	poll			0-3	4	epoll_dispatch (epoll.c:407)		► Rank 0, thread 3	0-3	4	epoll_wait			Name	Value	3*j*j	30603	j	101
Processes	Threads	Function	Source	Variables																																																																								
0-3	4	main (wave.c:334)	► iterations = update(left, right);	► Rank 0, thread 1																																																																								
0-3	4	update (wave.c:216)	► values[j] = newval[j];	▼ Rank 0, thread 1																																																																								
				<table><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>i</td><td>0</td></tr><tr><td>iterations</td><td>1</td></tr><tr><td>j</td><td>101</td></tr><tr><td>left</td><td>-2 (from -2 to 2)</td></tr><tr><td>now</td><td><aggregate value></td></tr><tr><td>right</td><td>1 (from -2 to 3)</td></tr><tr><td>stop</td><td>0</td></tr></tbody></table>	Name	Value	i	0	iterations	1	j	101	left	-2 (from -2 to 2)	now	<aggregate value>	right	1 (from -2 to 3)	stop	0																																																								
Name	Value																																																																											
i	0																																																																											
iterations	1																																																																											
j	101																																																																											
left	-2 (from -2 to 2)																																																																											
now	<aggregate value>																																																																											
right	1 (from -2 to 3)																																																																											
stop	0																																																																											
0-3	8	progress_engine																																																																										
0-3	8	opal_libevent2022_event_base_loop (event.c:1630)		► Rank 0, thread 2																																																																								
0-3	4	poll_dispatch (poll.c:165)		► Rank 0, thread 2																																																																								
0-3	4	poll																																																																										
0-3	4	epoll_dispatch (epoll.c:407)		► Rank 0, thread 3																																																																								
0-3	4	epoll_wait																																																																										
Name	Value																																																																											
3*j*j	30603																																																																											
j	101																																																																											
14		0:11.009	0-3	Play																																																																								

9 Step Guide

Optimizing high performance applications

Improving the efficiency of your parallel software holds the key to solving more complex research problems faster.

This pragmatic, 9 Step best practice guide, will help you identify and focus on application readiness, bottlenecks and optimizations one step at a time.

Bugs

- Correct application

Analyze before you optimize

- Measure all performance aspects. You can't fix what you can't see.
- Prefer real workloads over artificial tests.

Cores

- Discover synchronization overhead and core utilization
- Synchronization-heavy code and implicit barriers are revealed

Vectorization

- Understand numerical intensity and vectorization level.
- Hot loops, unvectorized code and GPU performance revealed

Verification

- Validate corrections and optimal performance

Memory

- Reveal lines of code bottlenecked by memory access times.
- Trace allocation and use of hot data structure

Communication

- Track communication performance.
- Discover which communication calls are slow and why.

Workloads

- Detect issues with balance.
- Slow communication calls and processes. Dive into partitioning code.

I/O

- Discover lines of code spending a long time in I/O.
- Trace and debug slow access patterns.

Linaro Performance Reports

Characterize and understand the performance of HPC application runs



Commercially supported
by Linaro

Gather a rich set of data

- Analyses metric around CPU, memory, IO, hardware counters, etc.
- Possibility for users to add their own metrics



Accurate and
Astute insight

Build a culture of application performance & efficiency awareness

- Analyses data and reports the information that matters to users
- Provides simple guidance to help improve workloads' efficiency



Relevant advice
to avoid pitfalls

Adds value to typical users' workflows

- Define application behaviour and performance expectations
- Integrate outputs to various systems for validation (eg. continuous integration)
- Can be automated completely (no user intervention)


Linaro Performance Reports

A high-level view of application performance with “plain English” insights

Command: `mpirun.hydra -host node-1,node-2 -map-by socket -n 16 -ppn 8 ./Bin/low_freq/../../Src//hydro -i ./Bin/low_freq/../../Src//Input/input_250x125_corner.nml`
Resources: 2 nodes (8 physical, 8 logical cores per node)
Memory: 15 GiB per node
Tasks: 16 processes, OMP_NUM_THREADS was 1
Machine: node-1
Start time: Thu Jul 9 2015 10:32:13
Total time: 165 seconds (about 3 minutes)
Full path: Bin/../../Src

Summary: hydro is **MPI-bound** in this configuration

Compute 20.6% 

MPI 63.2% 

I/O 16.2% 

Time spent running application code. High values are usually good.
This is **very low**; focus on improving MPI or I/O performance first

Time spent in MPI calls. High values are usually bad.
This is **high**; check the MPI breakdown for advice on reducing it

Time spent in filesystem I/O. High values are usually bad.
This is **average**; check the I/O breakdown section for optimization advice

I/O

A breakdown of the 16.2% I/O time:

Time in reads 0.0% |

Time in writes 100.0% 

Effective process read rate 0.00 bytes/s |


Effective process write rate 1.38 MB/s 

Most of the time is spent in **write operations** with a very low effective transfer rate. This may be caused by contention for the filesystem or inefficient access patterns. Use an I/O profiler to investigate which write calls are affected.

Linaro Performance Reports Metrics

Lowers expertise requirements by explaining everything in detail right in the report

Multi-threaded
parallelism



SIMD
parallelism

MPI

Of the 41.3% total time spent in MPI calls:

Time in collective calls	100.0%	■
Time in point-to-point calls	0.0%	■
Estimated collective rate	4.07 bytes/s	■
Estimated point-to-point rate	0 bytes/s	■

All of the time is spent in collective calls with a very low transfer rate. This suggests a significant synchronization overhead in the MPI profiler.

OpenMP


A breakdown of the 99.5% time in OpenMP regions:

Computation	58.9%	■
Synchronization	41.1%	■
Physical core utilization	100.0%	■
System load	99.7%	■

Significant time is spent synchronizing threads in parallel regions. Check the affected regions with a profiler.

This may be a sign of overly fine-grained parallelism (OpenMP regions in tight loops) or workload imbalance.

Load
imbalance



Memory

Per-process memory usage may also affect scaling:

Mean process memory usage	160 Mb	■
Peak process memory		■
Peak node memory		■

The peak node memory is the total number of processes and more.

Lustre

Lustre file operations (per node)

Mean write rate	
Peak write rate	
Mean file operations	
Mean metadata	

Energy


A breakdown of how the 32.3 Wh was used:

CPU	61.9%	■
System	38.1%	■
Mean node power	94.1 W	■
Peak node power	98.0 W	■

Significant time is spent waiting for memory accesses. Reducing the CPU clock frequency could reduce overall energy usage.

OMP
efficiency
System
usage

Performance Improvement



Think,



code,

i, j, k


```
for (int i = 0; i < n; ++i) {
    for (int j = 0; j < n; ++j) {
        for (int k = 0; k < n; ++k) {
            C[i][j] += A[i][k] * B[k][j];
        }
    }
}
```



run, run, run...



...to test and measure many
different implementations




Loop order (outer to inner)	Running time (s)
i, j, k	1155.77
i, k, j	177.68
j, i, k	1080.61
j, k, i	3056.63
k, i, j	179.21
k, j, i	3032.82

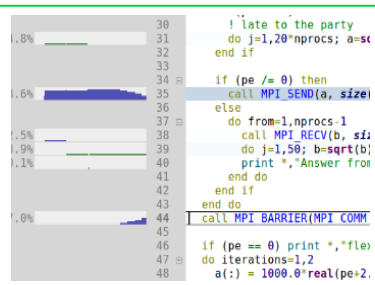
i, k, j

```
for (int i = 0; i < n; ++i) {
    for (int k = 0; k < n; ++k) {
        for (int j = 0; j < n; ++j) {
            C[i][j] += A[i][k] * B[k][j];
        }
    }
}
```


Linaro MAP Source Code Profiler Highlights




Find the peak memory use




Fix an MPI imbalance




Remove I/O bottleneck



Make sure OpenMP regions make sense



Improve memory access



Restructure for vectorization

MAP Capabilities

MAP is a sampling based scalable profiler


- Built on same framework as DDT
- Parallel support for MPI, OpenMP, CUDA
- Designed for C/C++/Fortran

Designed for 'hot-spot' analysis


- Stack traces
- Augmented with performance metrics

Adaptive sampling rate

- Throws data away - 1,000 samples per process
- Low overhead, scalable and small file size



GPU Profiling



Profile

- Supports both AMD and Nvidia GPUs
- Able to bring up metadata of the profile
- Mixed CPU [green] / GPU [purple] application
- CPU time waiting for GPU Kernels [purple]
- GPU Kernels graph indicating Kernel activity

GUI information

- GUI is consistent across platforms
- Zoom into main thread activity
- Ranked by highest contributors to app time

Python Profiling

19.0 adds support for Python

- Call stacks
- Time in interpreter


Works with MPI4PY

- Usual MAP metrics

Source code view


- Mixed language support

Note: Green as operation is on numpy array, so backed by C routine, not Python (which would be pink)



`map --profile mpirun -n 2 python ./diffusion-fv-2d.py`

Toggle percentage-time and core-time in MAP




Use for direct comparisons between runs at the same scale (process/core counts).

- Easily determine if a change has made a portion of code faster, slower, or largely unchanged.
- Performance report automatically includes both percentage-time and core time
- Core-time is an estimation, but should be very close to the application run time

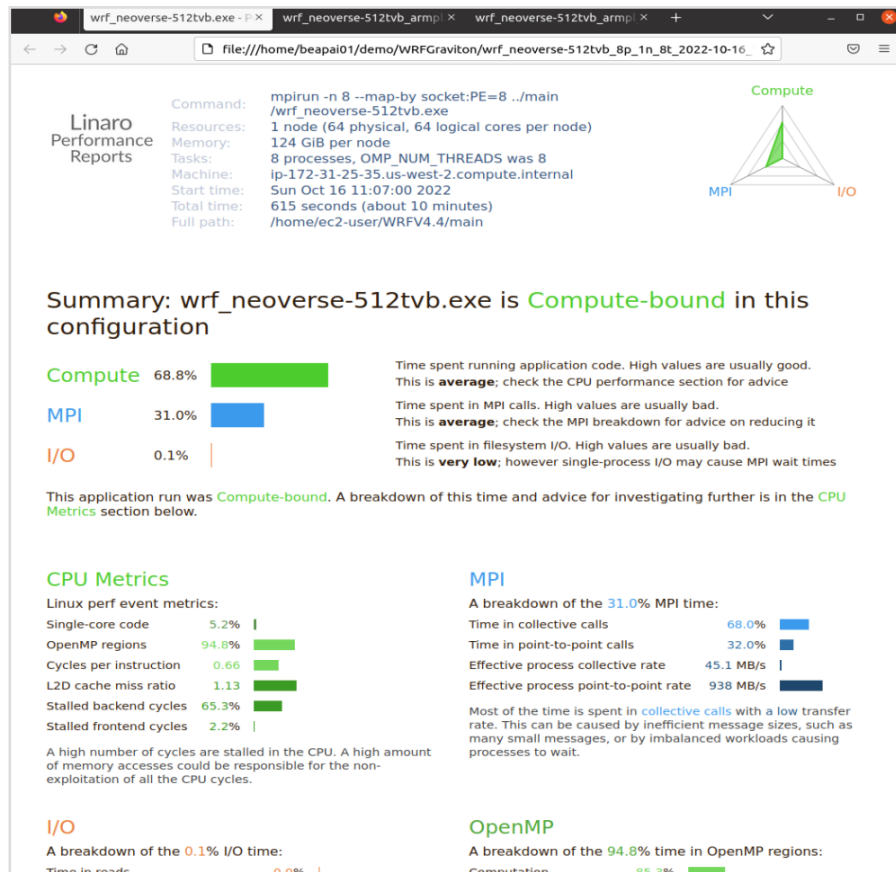
Libraries tab in MAP

- List time spent in shared libraries (left)
- List entry point functions into the selected library (right)

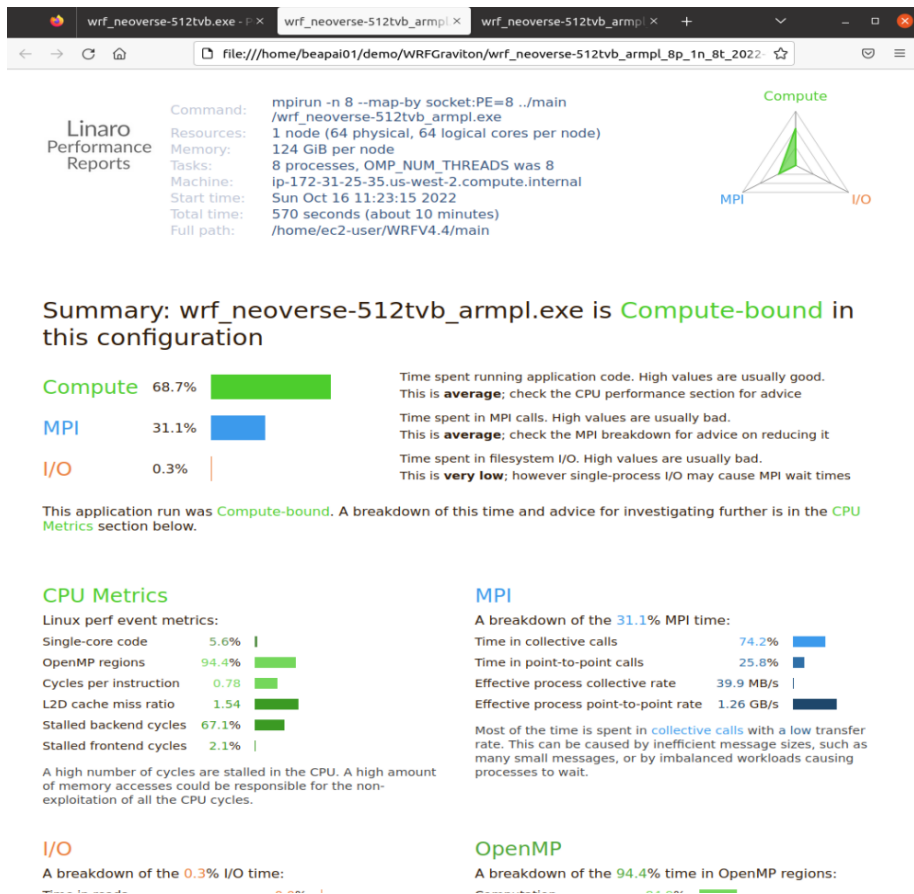


Use to identify the libraries that would benefit the most from optimisation or replacement (e.g. alternative maths library or memory management implementation).

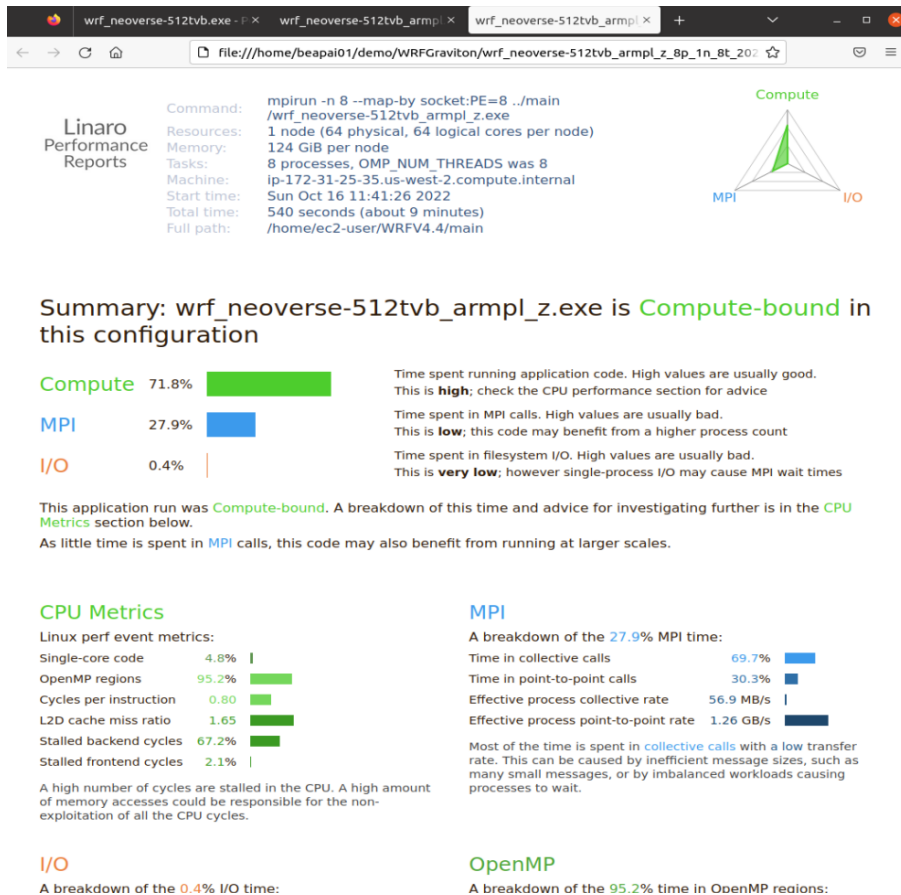
WRF build without enhancements



WRF with Arm Performance Libraries



WRF with Arm Performance Libraries and IO compression libraries



Change at which point of a run the Affinity data is shown (*Library Load, Initialisation, Finalization*).

Process-specific env vars

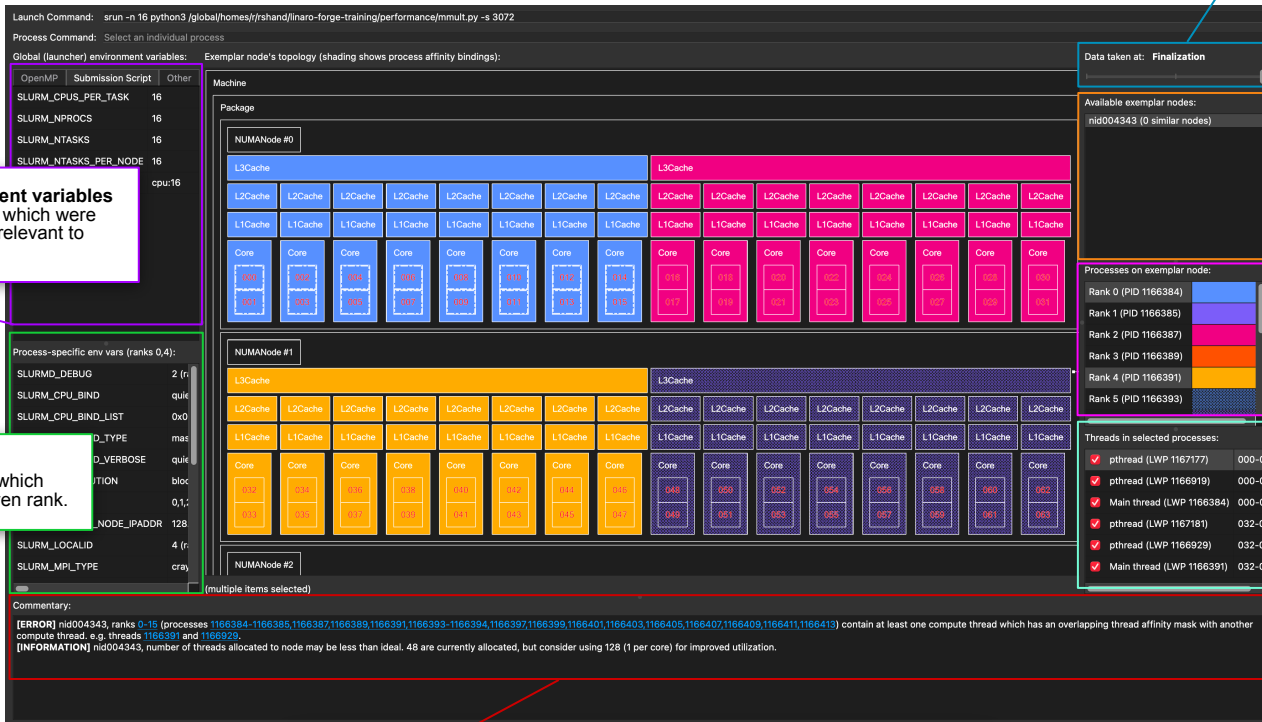
List of Environment Variables which might affect the affinity of a given rank.

A list of commentary, providing information and advice on Memory Imbalance, Core Utilization etc.

Selectable list of exemplars, allowing ability to switch data between nodes of a run. Nodes with similar affinity/structures are merged.

List of processes (by MPI rank) of the selected exemplar. Shows the key for the node topology diagram and selecting one shows all threads for the process.

List of all threads for the selected process. Selecting threads highlights which cores they are bound to in the topology view.



Thank you

rudys.shand@linaro.org

support@forge.linaro.com

Go to www.linaroforge.com