Linaro Forge

Debugging and Optimization Tools for HPC

HPC Development Solutions from Linaro

Best in class commercially supported tools for HPC



Performance Engineering for any architecture, at any scale

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 DDT Debugger Highlights



The Performance Roadmap

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.

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 reveleaed

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.

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.

I/O

 Discover lines of code spending a long time in I/O.

 Trace and debug slow access patterns.

Workloads

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



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



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 MAP Source Code Profiler Highlights



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



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 jsrun -n 2 python3 ./diffusion-fv-2d.py Linaro Forge

Debugging and Performance Engineering for Nvidia and AMD GPUs



Thank you

Go to www.linaroforge.com Beau.Paisley@linaro.org