ARM in HPC:
Software and Tools

Geraint North
Architect,
High Performance Ecosystem

geraint.north@arm.com

A brief introduction to ARM

ARM designs scalable, energy efficient-processors and related technologies to deliver the intelligence in applications ranging from sensors to servers, including smartphones, tablets, enterprise infrastructure and the Internet of Things.

- Low power and high efficiency are in our DNA
- For HPC, our novel business model allows for differentiated SoCs:

Our business model

- Technology licensed to our partners (currently over 350), who create their own SoC products.
  - Huge demand for tools to reduce integration time and effort.
- May license an instruction set **architecture** (e.g. "ARMv8") or a specific **implementation** (e.g. "A72").
- **big.LITTLE** architecture allows heterogeneity within a single SoC.
- We also develop and license specifications and implementations of interconnect, GPUs, core radio IP and many more.

ARM

ARM Manchester

- New location, founded July 2014
- 20 engineers focused on runtime, library and compiler development for high-performance computing and server.
- Focus areas:
  - Exploitation of future ARM architectures.
  - Autovectorization improvements for open-source compilers.
  - ARM-tuned numerical libraries.
  - HPC profiling, debugging and programmer productivity tools.

ARM

The ARM Ecosystem Today

Compilers

Open Source

- Significant and ongoing investment in both GCC and LLVM.
- Improving core performance through mostly architectural (not microarchitectural) optimizations.
- Most focus and improvement in floating-point code.
- SPECv6 will likely add focus on OpenMP runtimes.
- New hire into my team later this year will work on LLVM Fortran enablement.

Commercial

- **PathScale, November 2014:** PathScale provides the full EKOPath compiler suite including OpenACC and OpenMP 4.0 C/C++/Fortran support for ARMv8 to support HPC and Enterprise customers exploring the power efficiencies of these devices.

ARM Performance Libraries

In November 2015, we will offer a commercially supported set of 64-bit ARMv8 numerical libraries for scientific computing, built on technology from NAG.

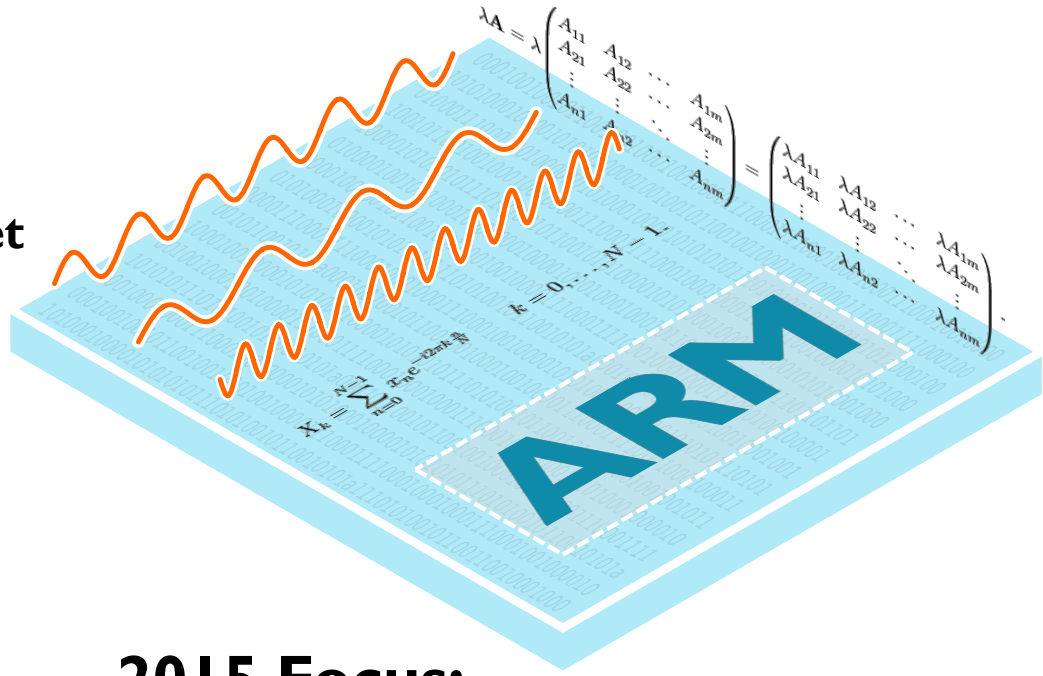

Enable ARM partners' computational kernels tuned for their SoC implementation.

- Unified, validated framework
- A57 and Cavium® ThunderX optimizations available at launch date, others to follow.
- Development model allows for the production of ARM Math Libraries that are highly tuned for specific system configurations.
- All implementations hosted on arm.com

2015 Focus:

- BLAS
- LAPACK
- FFT

By the end of 2015, an HPC-specific ARM microsite will offer downloads, technical reference material, how-to-guides and third-party software recommendations for the scientific computing community.

nag® ARM

Allinea Forge Development Toolsuite for ARMv8

- Development tools for high performance software
  - For multi-threaded and multi-process software
  - C, C++ and F90

Image source

- Allinea DDT
  - debug complex software problems
  - Includes memory debugging, visualization, C++ STL, F90 and more
- Allinea MAP
  - profiler for optimizing code performance
  - Tackle I/O, threads, CPU, synchronization and memory performance issues.

ARM

Research Collaborations for all time horizons

Hartree Centre Energy Efficient Computing	Mont Blanc 3	Fast Forward 2
Established through a £19M capital grant from the Department of Business Innovation and Skills to establish a centre of best practice in the UK that will enable users of computer systems to achieve the same outcomes while minimising the consumption of energy.	Goal is the creation of a new high-end HPC platform (SoC and node) that is able to deliver a new level of performance / energy ratio on real applications.	Evaluation of next-generation architecture in the context of Department of Energy applications.

Image source

Image source

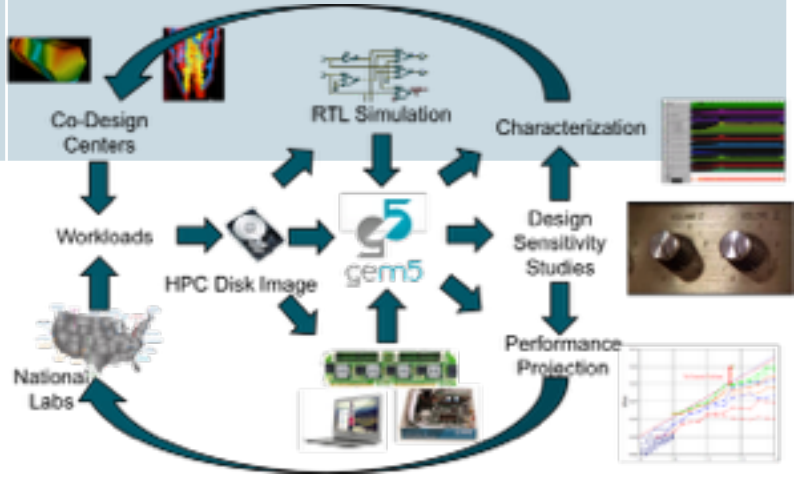

Future Tools Technology

Development cycle for small-scale applications

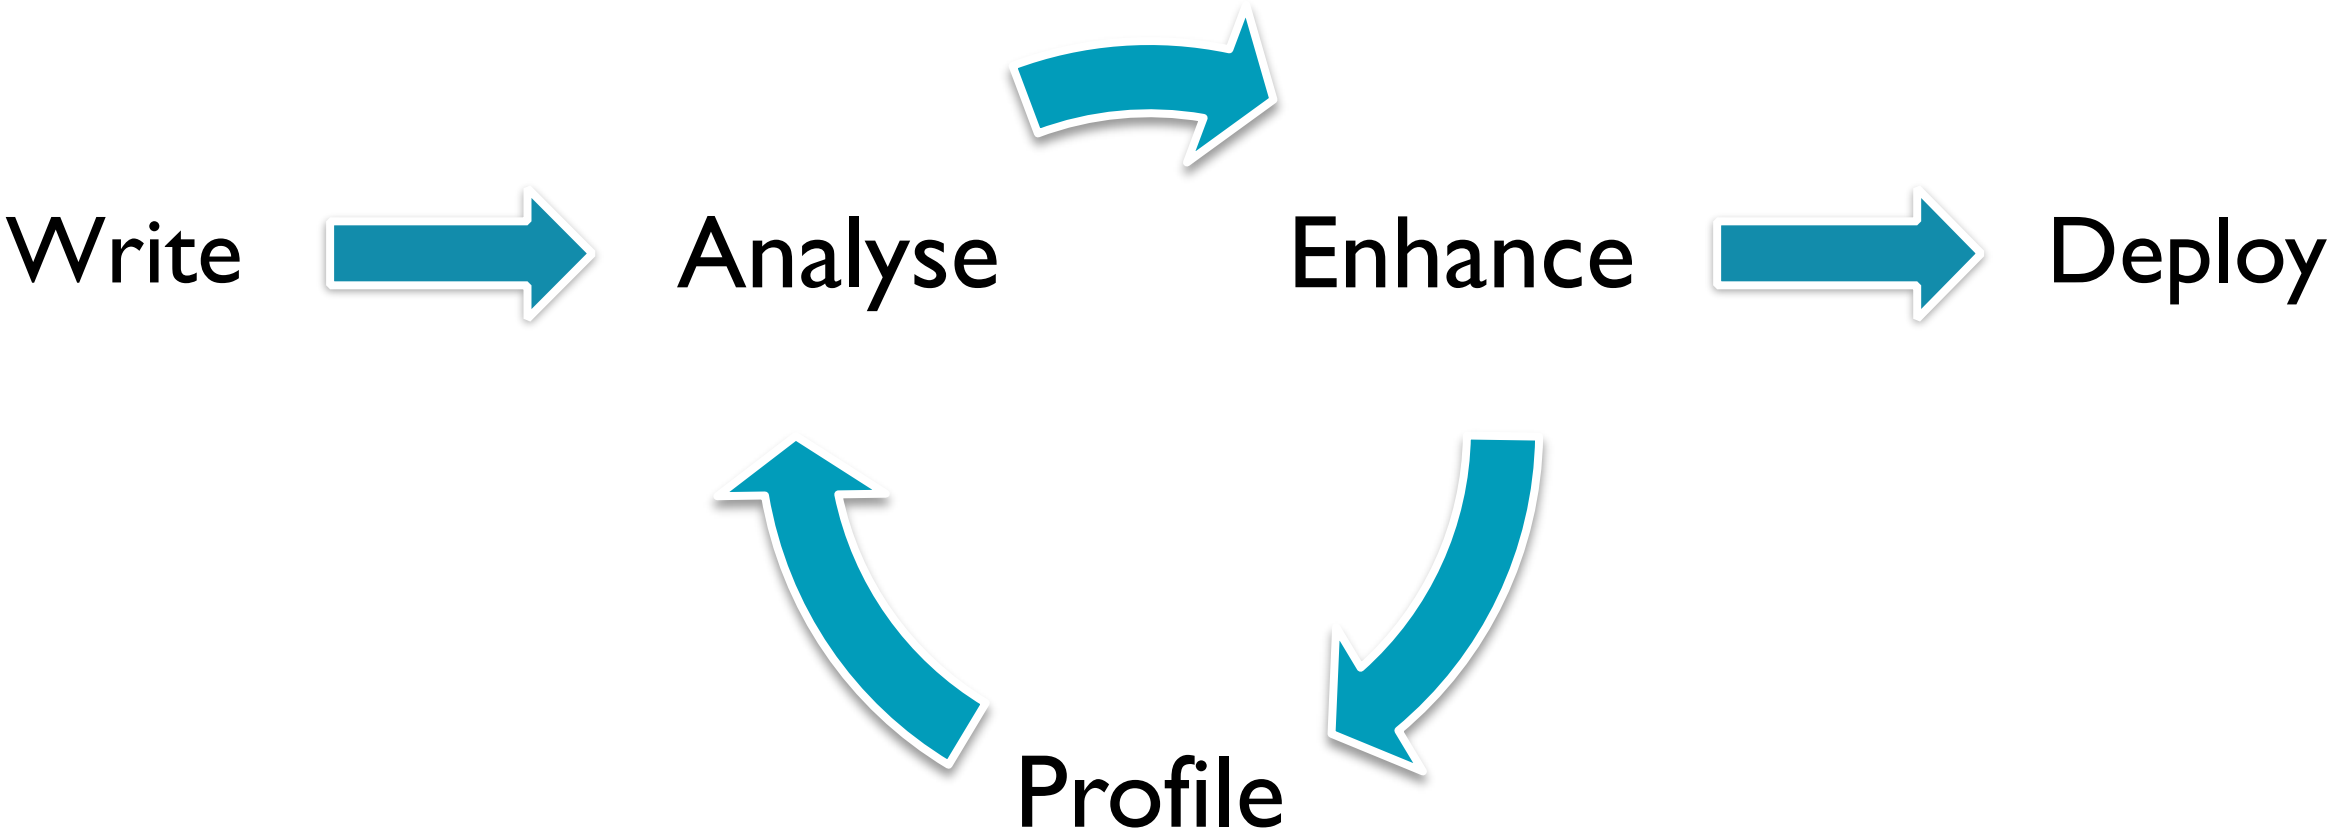

ARM

What's different for at petascale and beyond?

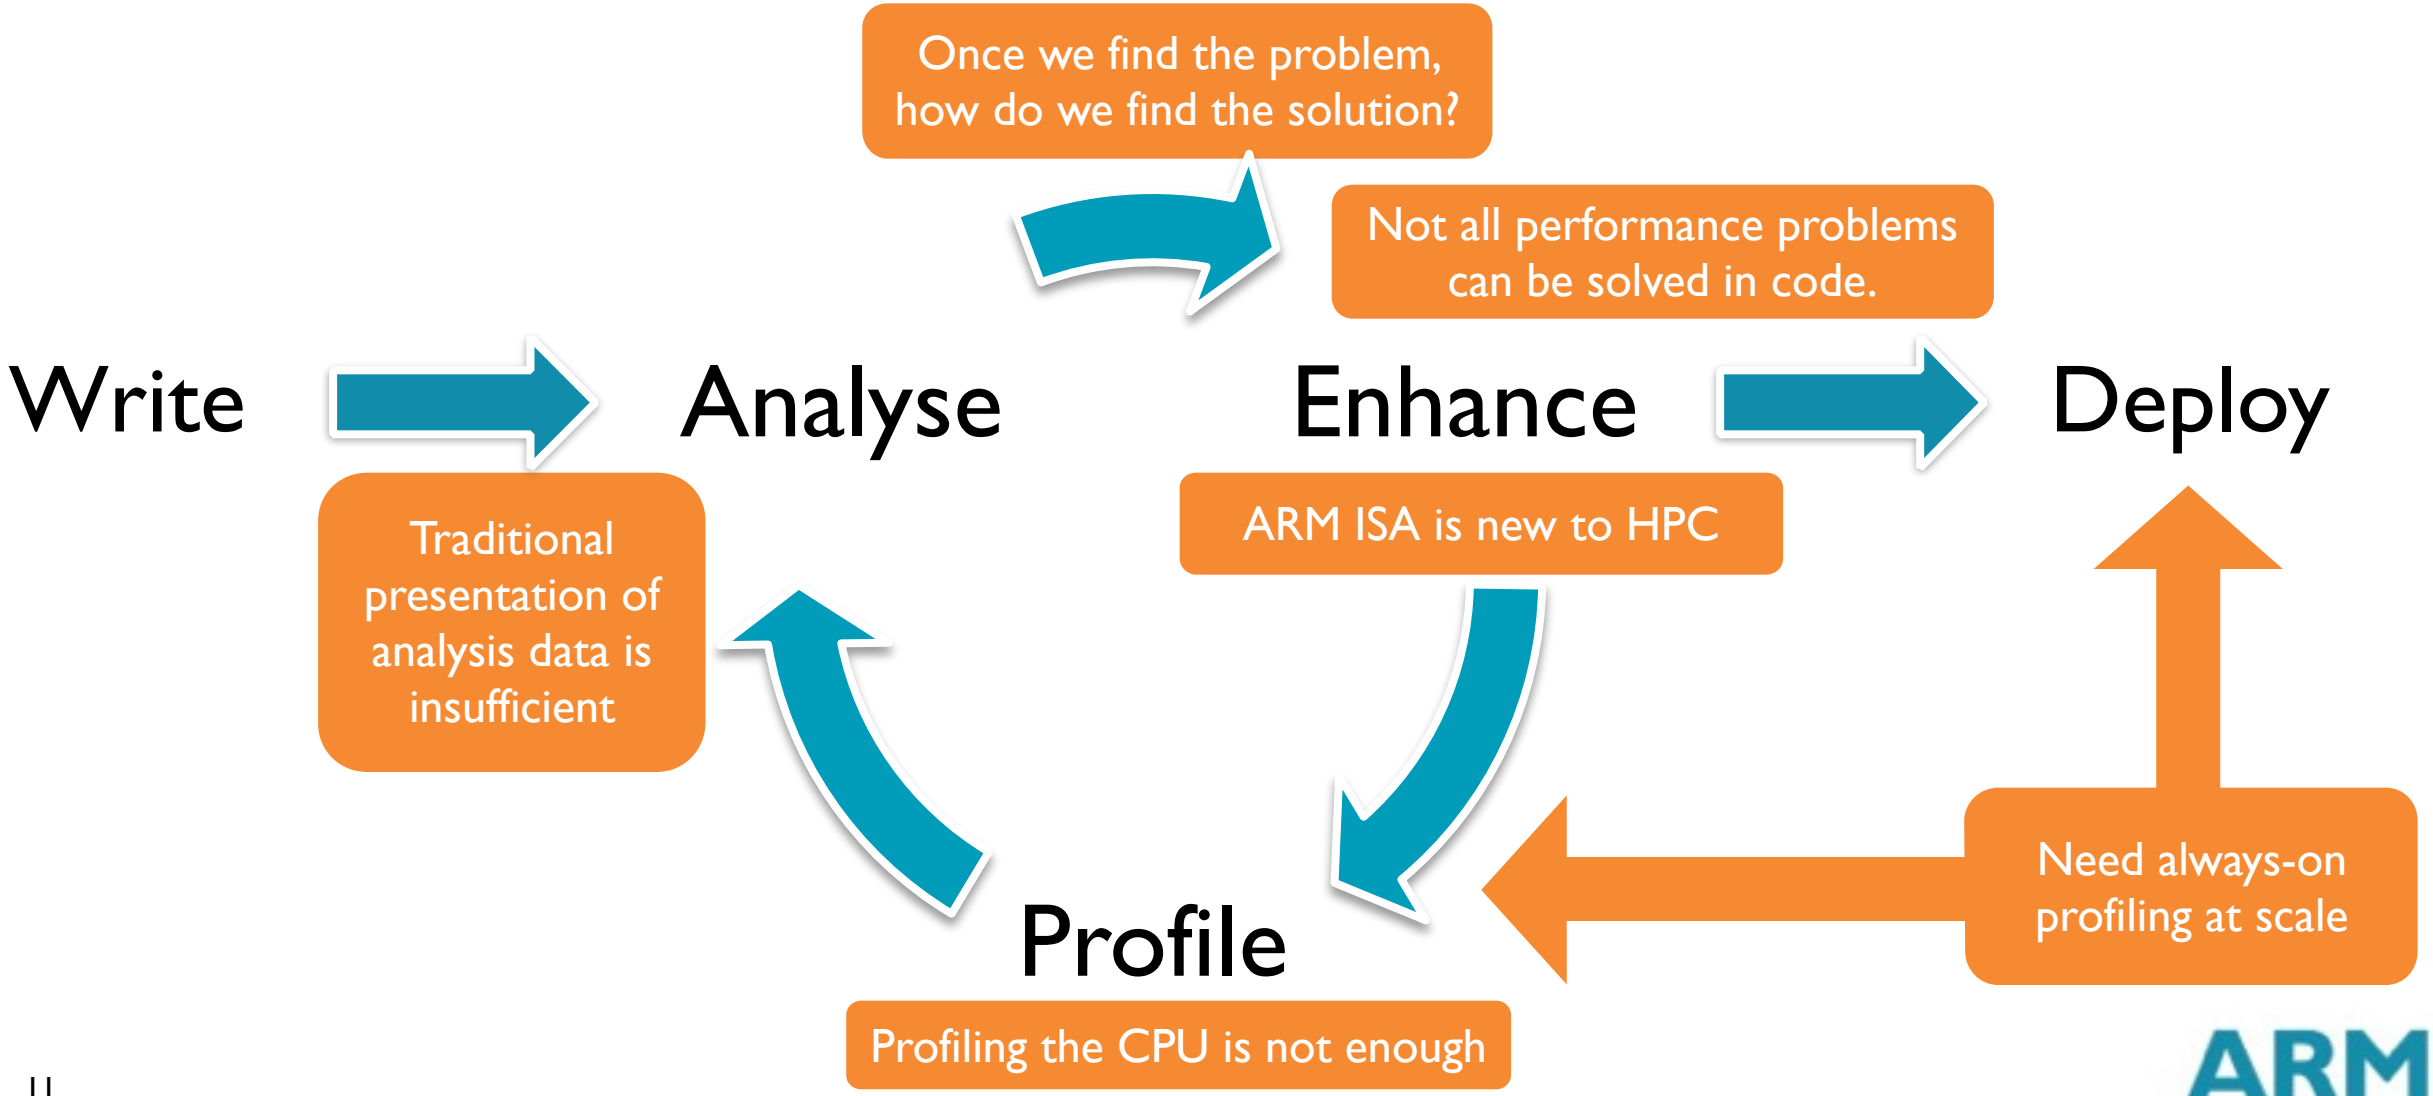

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

- The ARMv8-A instruction set is different to the X86 and POWER ISAs common in HPC.
- The variety that our business model brings is powerful, but…
  - The Architecture Specification tells us *what* an instruction should do.
  - Only the silicon implementor knows *how fast* it might be.

How to close the gap:

- Enable and deliver tuned libraries to reduce the need to write low-level assembly and make code portable across microarchitectures.
- Provide examples and guidance on porting and tuning for ARM.
- Bring together several information sources in one place (e.g. present them in our ARM DS-5 IDE environment)
  - ARM C Language Extensions (e.g. NEON vector intrinsics)
  - ARM Architecture Reference Manual
  - Performance guides where available (e.g. A57)

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

At HPC-scale, we can't expect a small profiling run to be representative of deployed performance.

- Profiling perturbs application performance.
- Scaling up a workload significantly changes the execution profile.

How to close the gap:

- Ensure that preferred profiling APIs are available and high-quality for ARM implementations.
  - Are standard (e.g. PAPI) sufficient? What about aspects like power consumption?
- Design future ARM architectural enhancements to lower the cost of always-on profiling.
- Explore novel heuristics and techniques to obtain useful performance information on today's hardware at low cost.

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

At scale, performance is dictated by:

- Code quality
- Memory bus contention
- Network switch congestion
- Patterns of storage access
- …

How to close the gap:

- Identify the critical system components that require profiling.
- Support and/or develop standards to collect and aggregate performance data.
- This need for standardization is not just an ARM problem – it requires industry-wide involvement.

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

ARM Streamline is our commercial profiling tool, mostly used for embedded and mobile applications.

- Collects profiles via hardware trace or a Linux userspace daemon.
- Displays profile in a timeline format.
- Collects all the usual cycle/ cache miss/branch mispredict etc. counters.

Not ideal for HPC:

- Timeline- and per-core views not appropriate for long-running applications.
- Single node only.

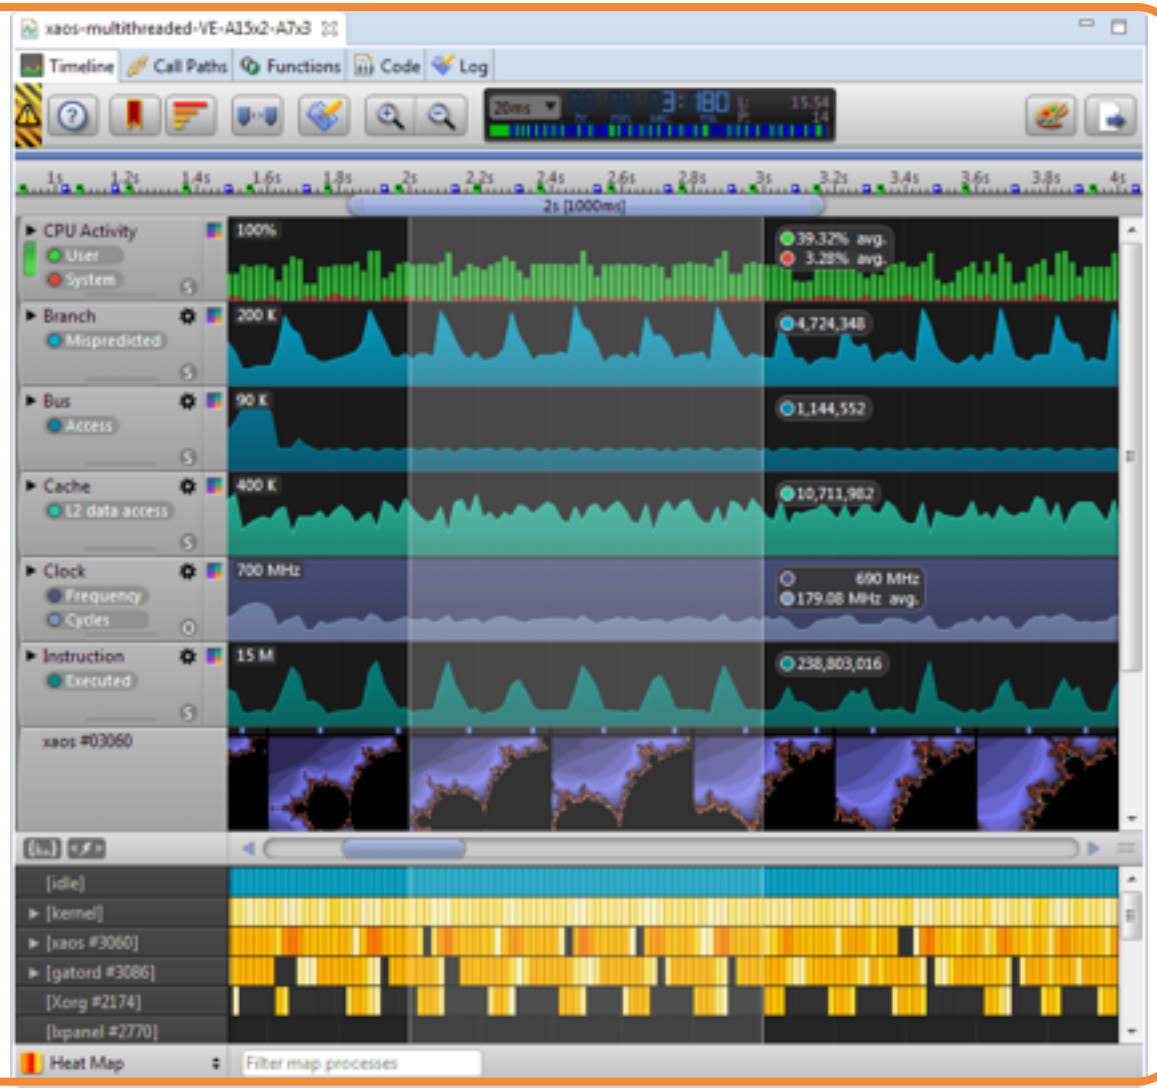

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

An approach key to HPC:
Streamline can correlate performance data from several sources, including libraries and drivers instrumented via a public interface. (OpenCL shown here)

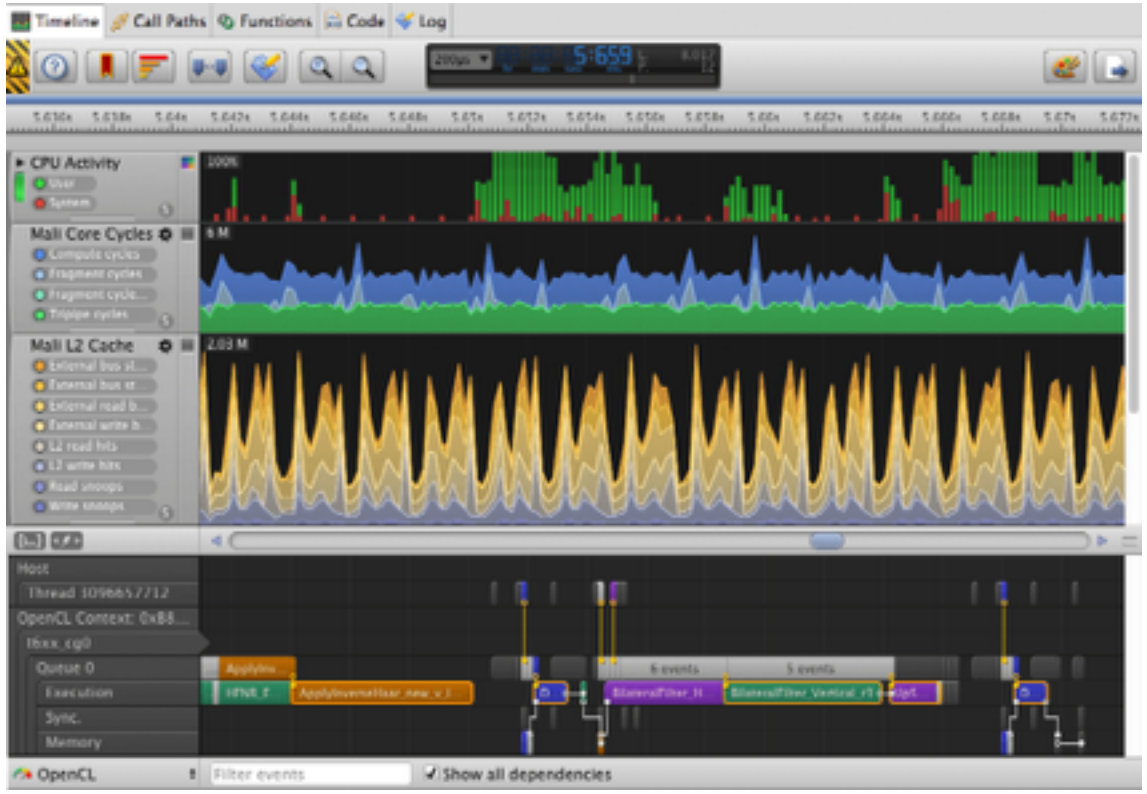

How to close the gap:

- Enhance Streamline and instrument accordingly to make it aware of OpenMP, MPI, ARM Math Libraries.
- Explore novel data presentation methods that put the data in the context of the runtimes and libraries.

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

A problem-focused strategy designed to transfer blame from the threads suffering from the symptoms of inefficiency to the threads causing the inefficiency.

ASCR Tools Challenges for Exascale Computing, 2011

How to close the gap:

- Understanding of MPI/OpenMP runtimes can help to identify the cause of problems such as lock contention.
- ARM Research is investing in Machine Learning, and problem determination is a potential area of investigation.
- For a certain class of problems, system monitoring tools can correlate events that may be the cause of an intermittent performance issue.
  - I/O wait times extended → another storage-intensive workload started on other nodes in the system,

ARM

Closing the gaps

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

We need standardized control-plane infrastructure to tune and reconfigure the whole system.

- Overclock single cores in serial regions.
- Direct work onto the appropriate big.LITTLE core.
- Actively manage cache contents.
- Prioritise network traffic.
- Actively relocate and rebalance storage.

How to close the gap:

- Develop or adopt standard APIs to modify the system configuration. OpenStack-like APIs that can reach deeper into the hardware and reconfigure it?
- If we can't make decisions automatically, expose new programming capabilities (pragmas, hints) to the programmer to exploit

To conclude…

ARM ISA is new to HPC

Need always-on profiling at scale

Profiling the CPU is not enough

Traditional presentation of analysis data is insufficient

Once we find the problem, how do we find the solution?

Not all performance problems can be solved in code.

Enterprise ARM is now very much like any other server system.

The software ecosystem for today's cores and technology is in good shape; we and others are working to address gaps where they exist.

ARM's longer-term HPC problems are no different to HPC's general longer-term problems, but are nuanced by:

- Our goal of power-awareness and energy efficiency in everything that we do.
- The need to manage heterogeneity, both within the SoC and across vendor implementations.
  - Exploit variety
  - Manage complexity

Thank you


Geraint North
Architect,
High Performance Ecosystem
Mobile: +44 7971 989162
Email: geraint.north@arm.com
ARM
Floor 4, York House
York Street,
Manchester, M2 3BB
United Kingdom
Direct: +44 161 234 9421
Fax: +44 1223 400410
www.arm.com
The Architecture for the Digital World®