Cross-platform Performance Engineering

Enabling cross-platform developers

ISC Workshop “X86, ARM, GPUs, Oh My!”

John C. Linford,
Chris Goodyer, Will Lovett, Ashok Bhat, Patrick Wohlschlegel,
David Lecomber, Sandra Boynton, et al.

OpenFOAM and ParaView across the Arm ecosystem

Cross-platform ecosystem and standards make this possible

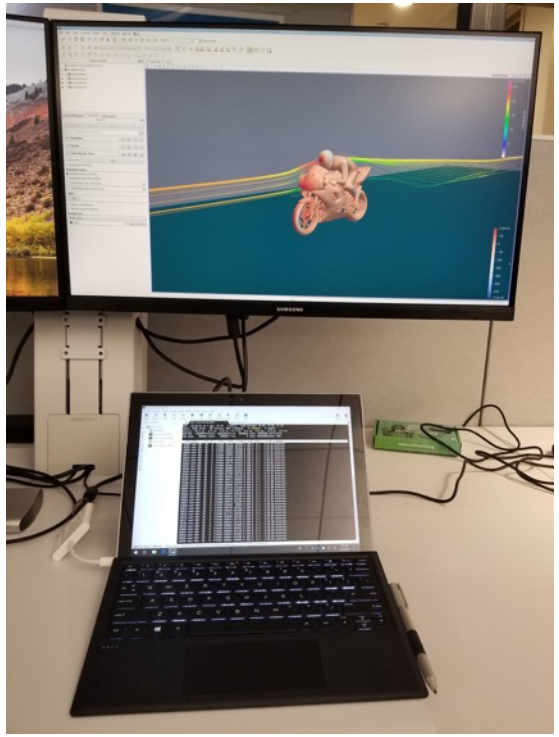

Hewlett Packard Enterprise

OpenFOAM

ParaView

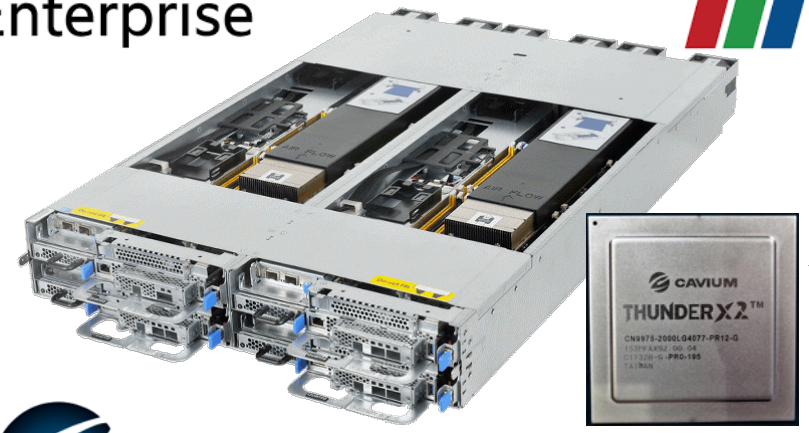

CAVIUM

ubuntu

Windows 10

arm

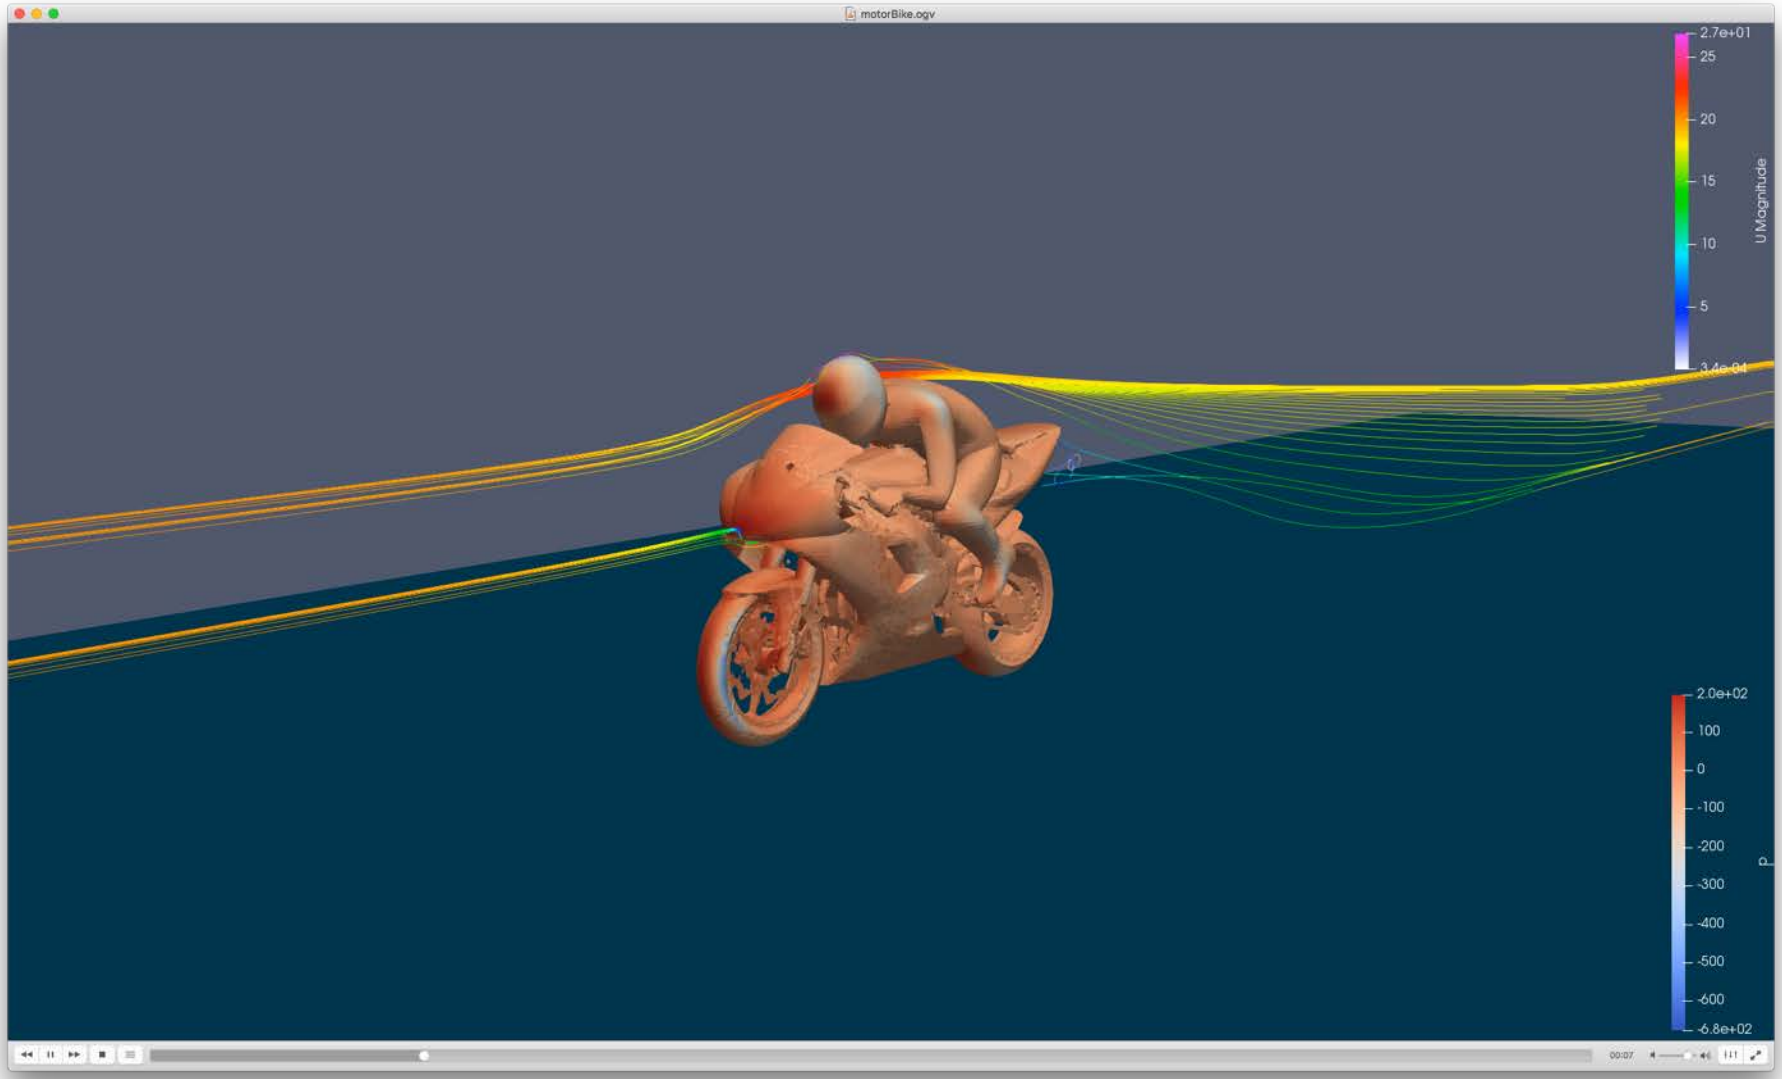

arm

Can you tell me how to port my code?

arm

Arm Allinea Studio and Friends

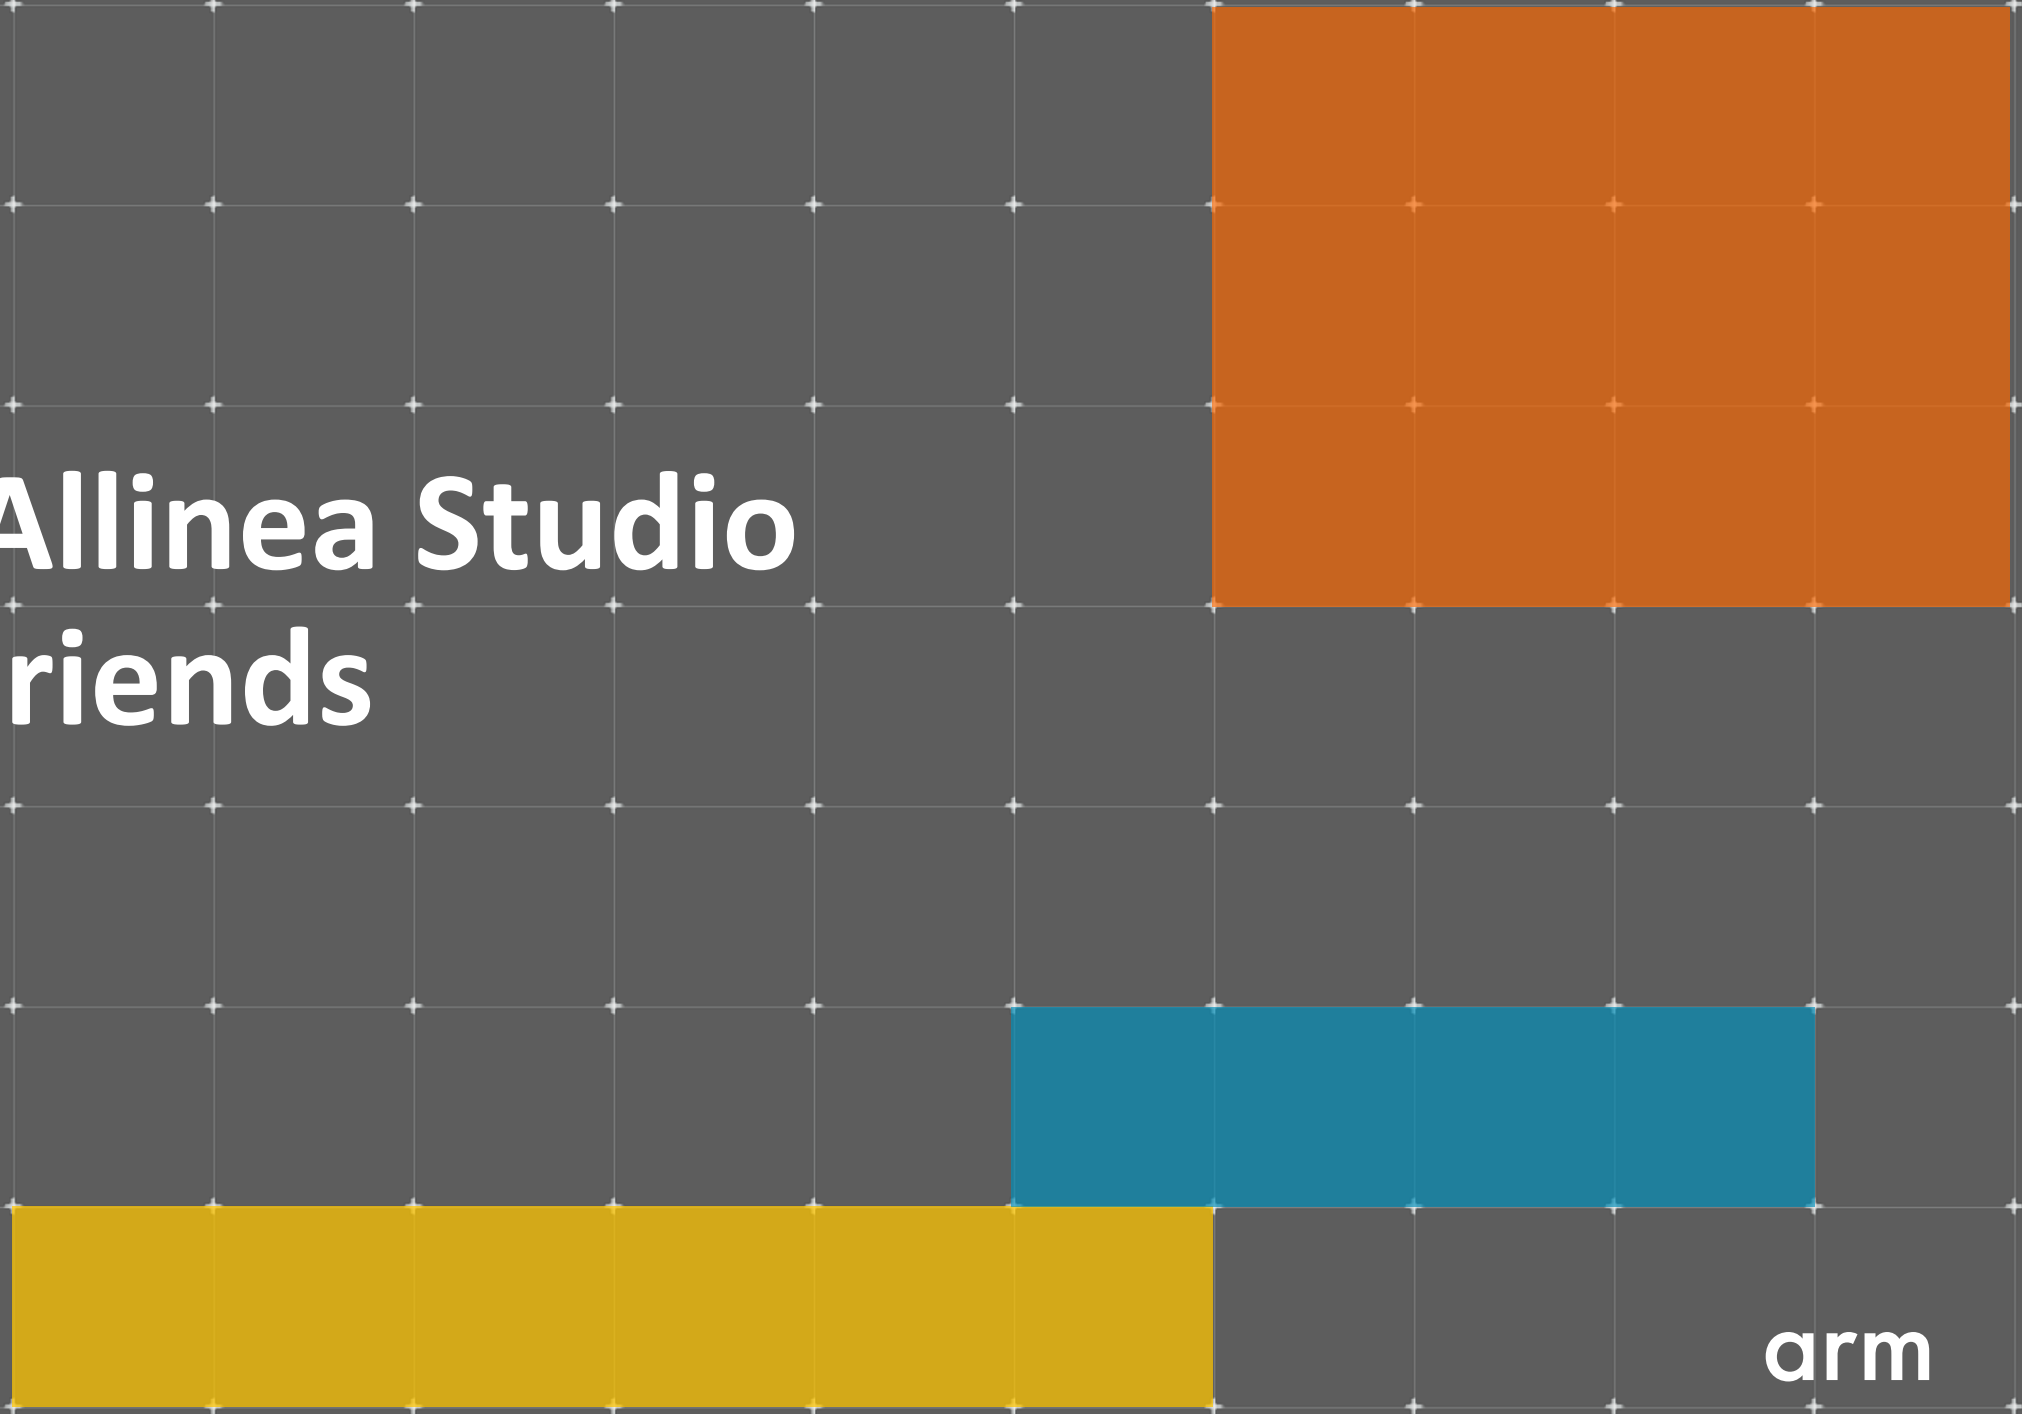

arm

Our solution for *any* architecture, at *any* scale

Commercial tools for aarch64, x86_64, ppc64 and accelerators

Arm Cross-Platforms Tools

Debug, optimize and analyze any platform

Arm DDT Professional	Arm MAP Professional	Arm Forge Professional	Arm Performance Reports
Slash your time to debug on any hardware, at any scale.	Speed-up applications with a lightweight scalable profiler	Arm DDT and MAP in One Single Package	Find the most efficient settings for your workloads.

Arm Allinea Studio

All-inclusive development toolkit for Arm hardware

Arm Compiler for HPC	Arm Performance Libraries	Arm Forge Professional	Arm Performance Reports
Linux user space compiler for HPC applications	BLAS, LAPACK and FFT	Multi-node interoperable profiler and debugger	Interoperable application performance insight

arm

arm COMPILER

Commercial C/C++/Fortran compiler with best-in-class performance

Compilers tuned for Scientific Computing and HPC

Latest features and performance optimizations

Commercially supported by Arm

Tuned for Scientific Computing, HPC and Enterprise workloads

- Processor-specific optimizations for various server-class Arm-based platforms
- Optimal shared-memory parallelism using latest Arm-optimized OpenMP runtime

Linux user-space compiler with latest features

- C++ 14 and Fortran 2003 language support with OpenMP 4.5*
- Support for Armv8-A and SVE architecture extension
- Based on LLVM and Flang, leading open-source compiler projects

Commercially supported by Arm

- Available for a wide range of Arm-based platforms running leading Linux distributions – RedHat, SUSE and Ubuntu

Arm Compiler – Building on LLVM, Clang and Flang projects

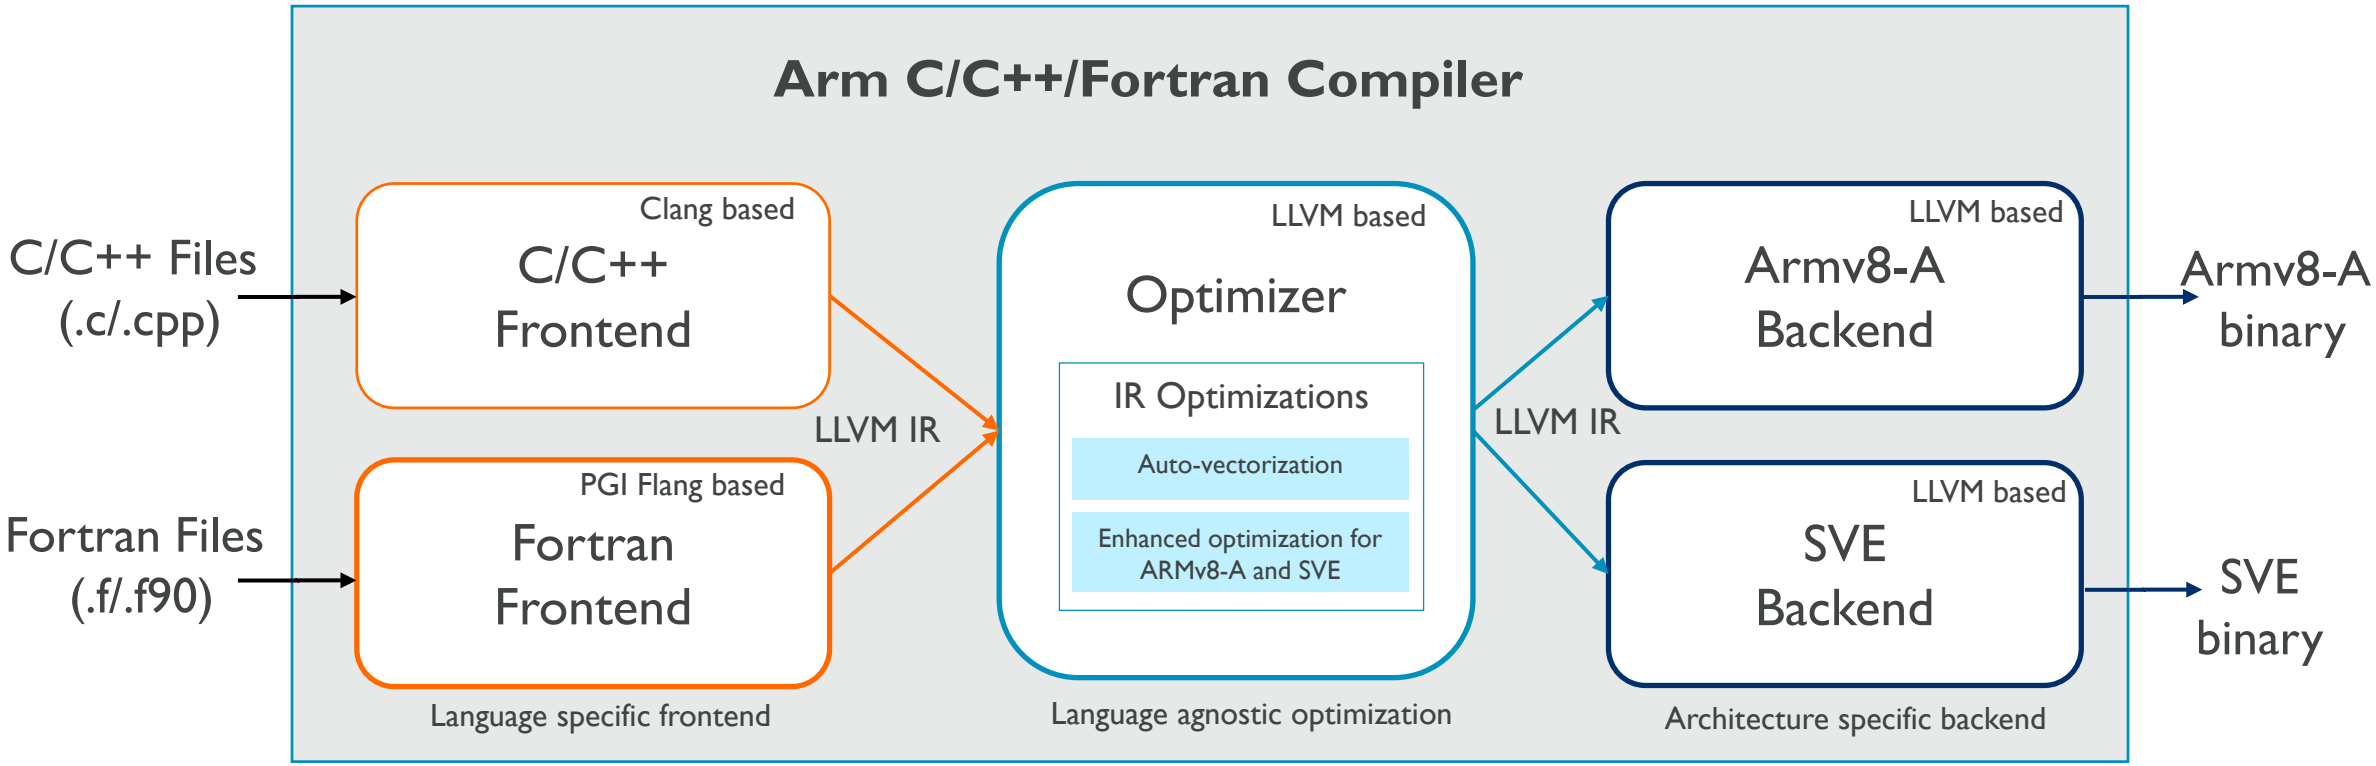

arm

arm PERFORMANCE LIBRARIES

Optimized BLAS, LAPACK and FFT

Commercially supported by Arm

Best in class performance

Validated with NAG test suite

Commercial 64-bit Armv8-A math libraries

- Commonly used low-level math routines - BLAS, LAPACK and FFT
- Provides FFTW compatible interface for FFT routines
- Batched BLAS support

Best-in-class serial and parallel performance

- Generic Armv8-A optimizations by Arm
- Tuning for specific platforms like Cavium ThunderX2 in collaboration with silicon vendors

Validated and supported by Arm

- Available for a wide range of server-class Arm-based platforms
- Validated with NAG's test suite, a de-facto standard

Arm Forge

An interoperable toolkit for debugging and profiling

Commercially supported by Arm

Fully Scalable

Very user-friendly

The de-facto standard for HPC development

- Available on the vast majority of the Top500 machines in the world
- Fully supported by Arm on x86, IBM Power, Nvidia 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 petaflopic applications)

Easy to use by everyone

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

Arm Performance Reports

Characterize and understand the performance of HPC application runs

Commercially supported by Arm

Accurate and astute insight

Relevant advice to avoid pitfalls

Gathers a rich set of data

- Analyses metrics 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

Adds value to typical users' workflows

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

Arm Instruction Emulator 18.0

Develop your user-space applications for future hardware today

Develop software for tomorrow's hardware today

Runs at close to native speed

Commercially Supported by ARM

Start porting and tuning for future architectures early

- Reduce time to market, Save development and debug time with Arm support

Run 64-bit user-space Linux code that uses new hardware features on current Arm hardware

- SVE support available now. Support for 8.x planned.
- Tested with Arm Architecture Verification Suite (AVS)

Near native speed with commercial support

- Integrates with DynamoRIO allowing arbitrary instrumentation extension
- Emulates only unsupported instructions
- Integrated with other commercial Arm tools including compiler and profiler
- Maintained and supported by Arm for a wide range of Arm-based SoCs

Arm Instruction Emulator

Develop your user-space applications for future hardware today

Run Linux user-space code that uses new hardware features (SVE) on current Arm hardware

Simple “black box” command line tool

```
$ armclang hello.c --march=armv8+sve
$ ./a.out
Illegal instruction
$ armie -a=armv8+sve ./a.out
Hello
```

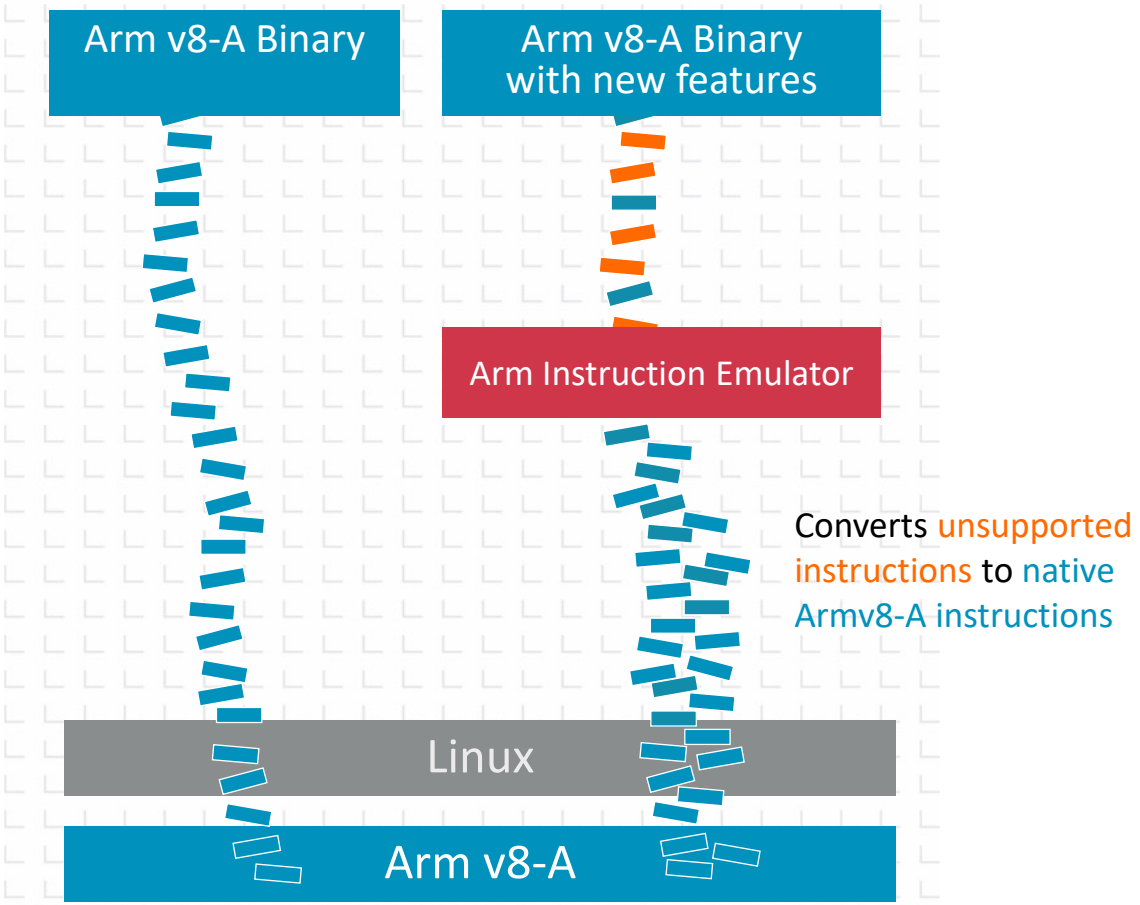

arm

VI-HPS and the tools ecosystem

See the http://www.vi-hps.org/tools/ for an excellent view of the tools ecosystem.

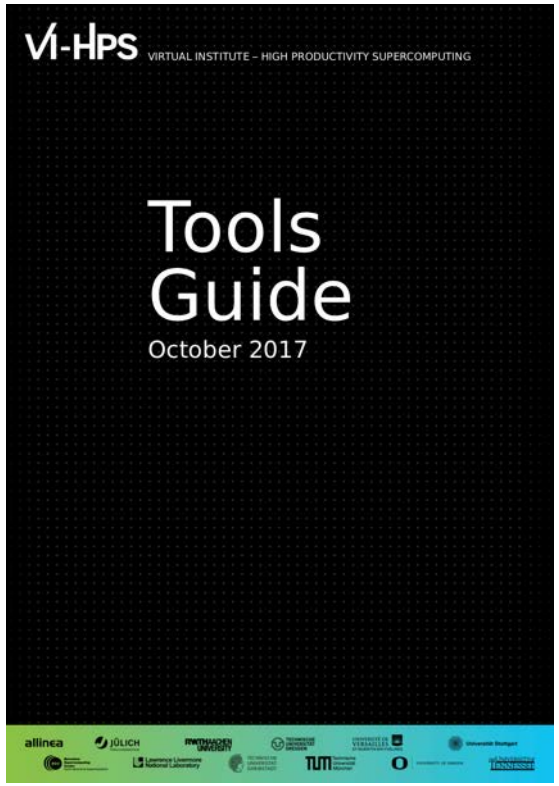

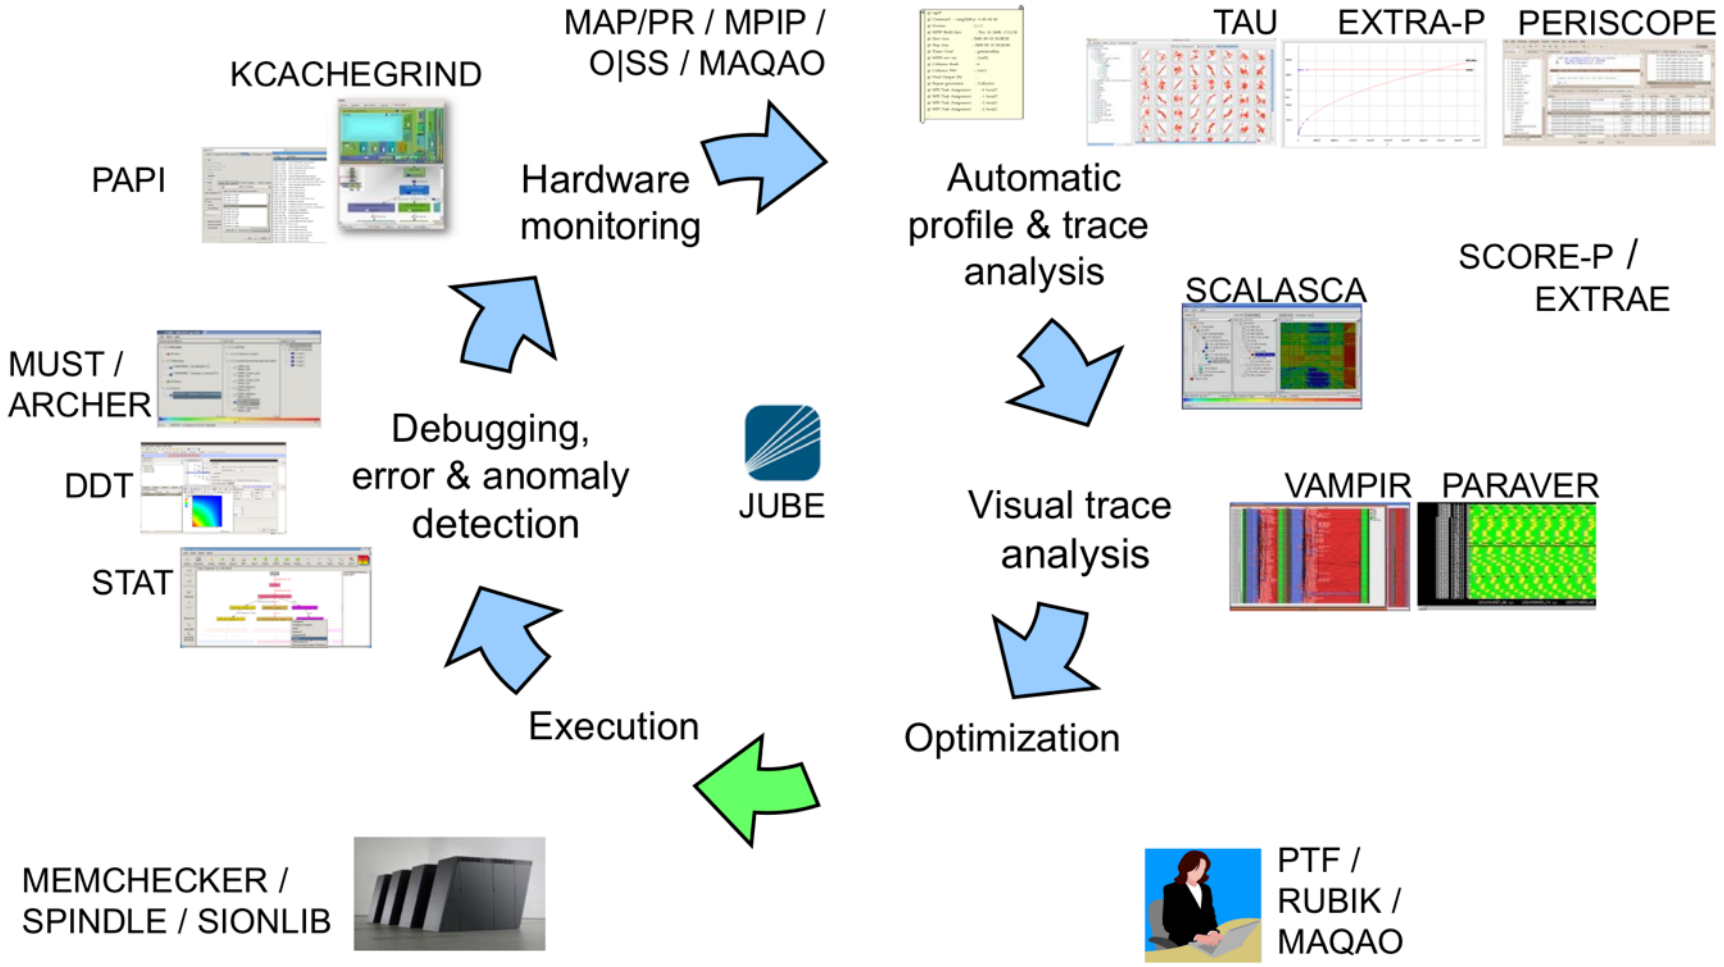

arm

Step 0: Get it to run

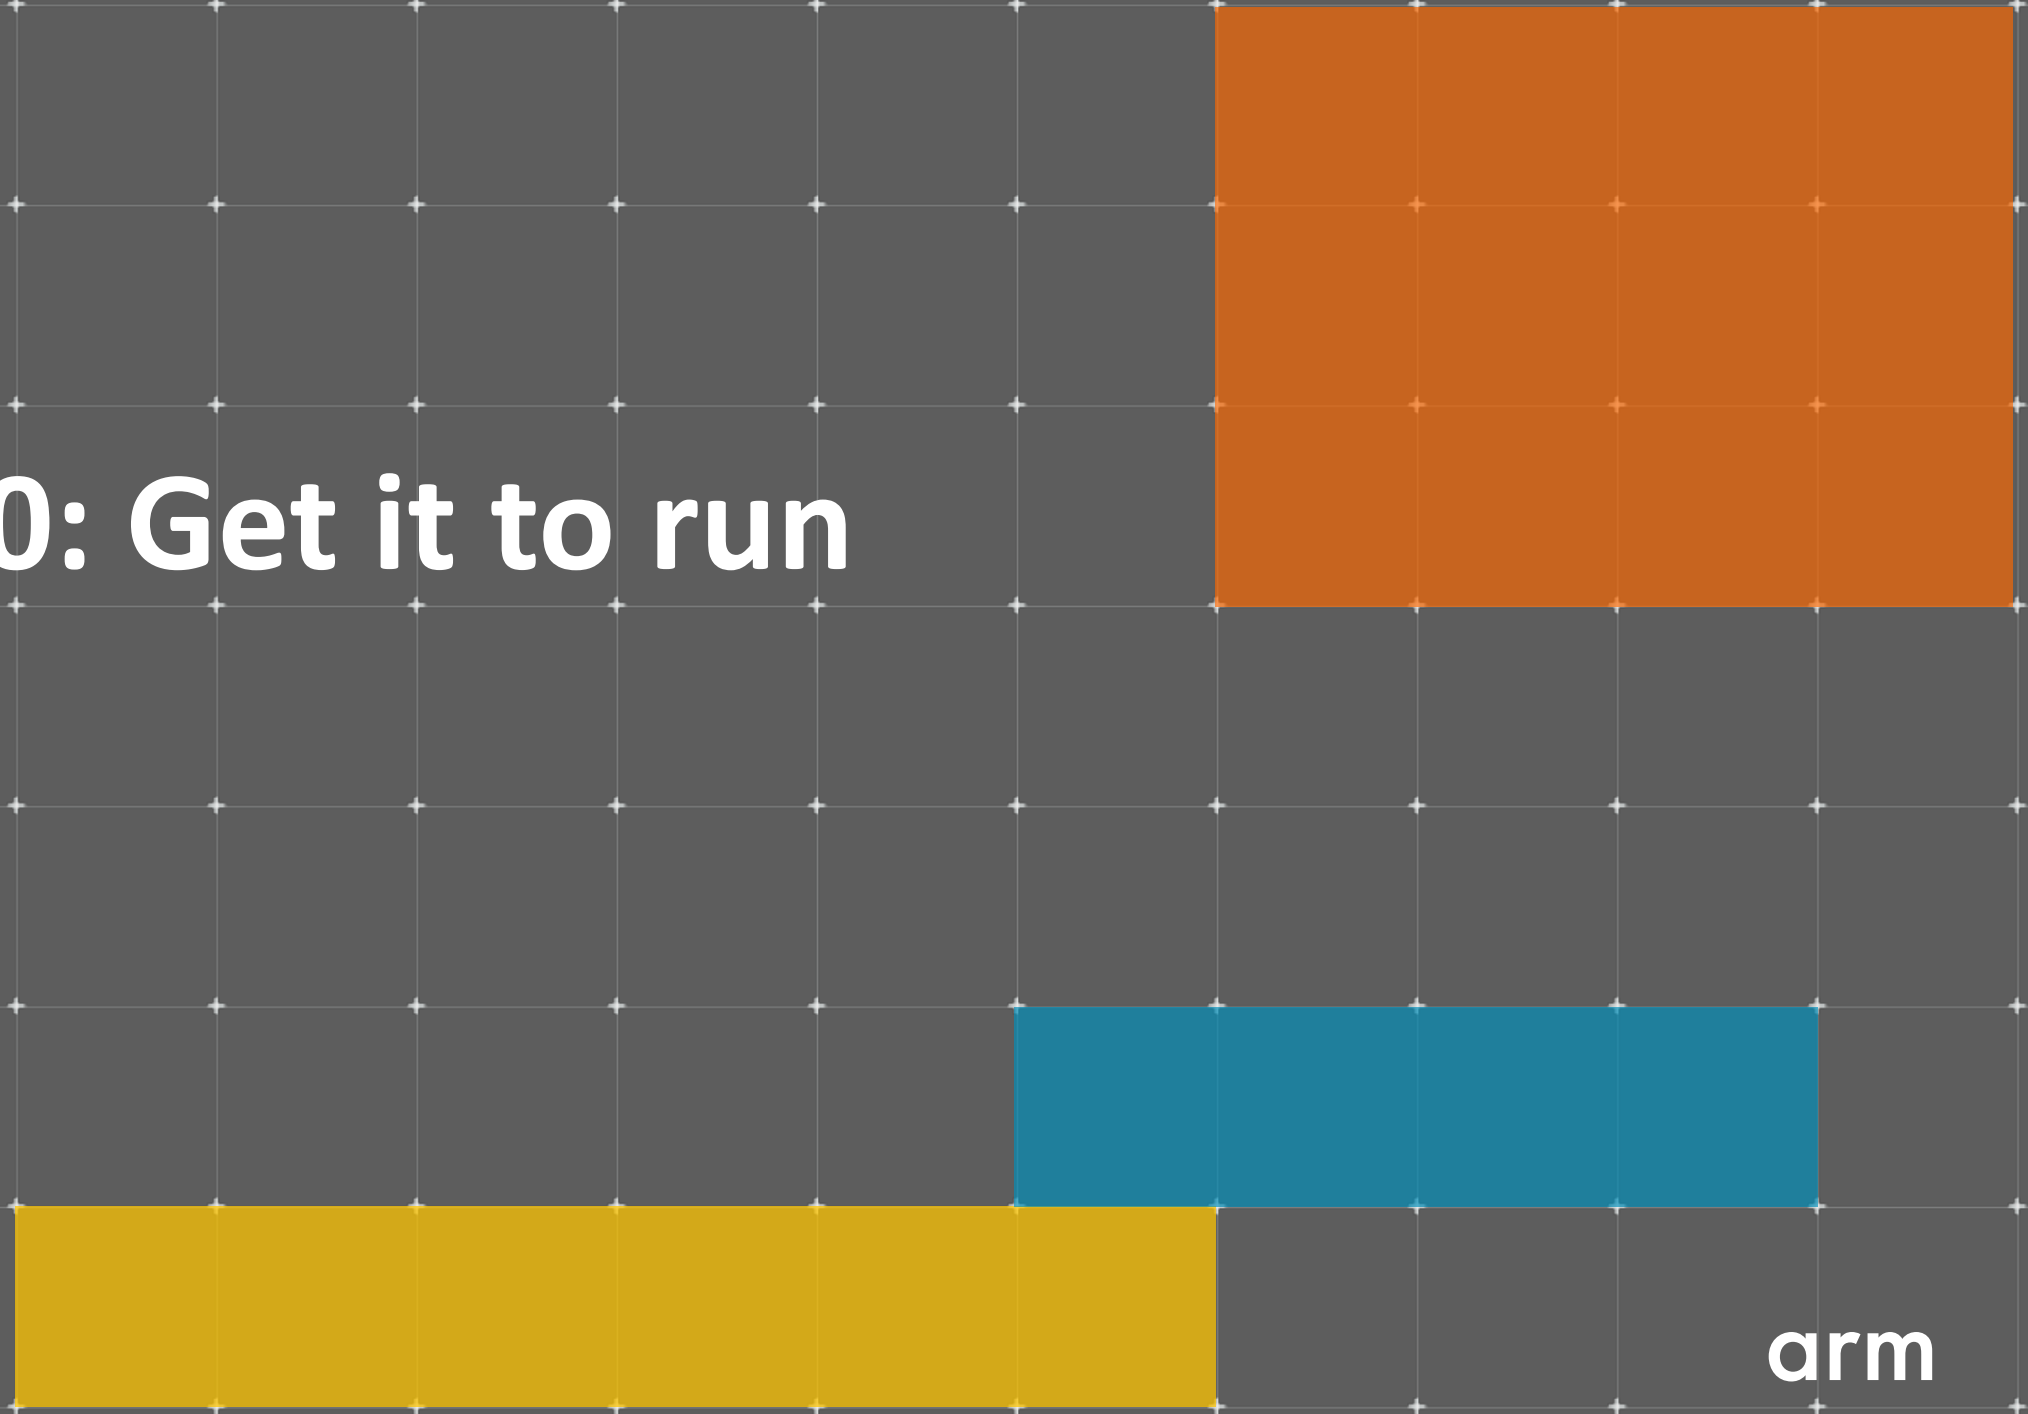

arm

Arm Porting Cheat Sheet

Ensure all dependencies have been ported.

- Arm HPC Packages Wiki: https://gitlab.com/arm-hpc/packages/wikis/categories/allPackages

Update or patch autotools and libtool as needed

```
•wget 'http://git.savannah.gnu.org/gitweb/?p=config.git;a=blob_plain;f=config.guess;hb=HEAD' -O config.guess
•wget 'http://git.savannah.gnu.org/gitweb/?p=config.git;a=blob_plain;f=config.sub;hb=HEAD' -O config.sub
•sed -i -e 's#wl=""#wl="-Wl,"#g' libtool
•sed -i -e 's#pic_flag=""#pic_flag=" -fPIC -DPIC"#g' libtool
```

Update build system to use the right compiler and architecture

- Check #ifdef in Makefiles. Use other architectures as a template.

Use the right compiler flags

- Start with `-mcpu=native -Ofast`.
- See slides further on for details.

Avoid non-standard compiler extensions and language features

- Arm compiler team is actively adding new “unique” features, but it’s best to stick to the standard.

Update hard-wired intrinsics for other architectures

- https://developer.arm.com/technologies/neon/intrinsics
- Worst case: default to a slow code.

Update, and possibly fix, your test suite

- Regression tests are a porter’s best friend.
- Beware of tests that expect exactly the same answer on all architectures!

Know architectural features and what they mean for your code

- Arm’s weak memory model.
- Division by zero is silently zero on Arm.

arm

“A maze of twisty little passages, all alike” -- ADVENT, 1976

...I’ve got dependencies, and *those* dependencies have dependencies!

Scientific software may be quite monolithic - but it is rarely self-contained.

Use of external libraries is increasingly common, and a conscious design choice for many projects.

- IO libraries are very common: HDF5, NetCDF (C, parallel and Fortran flavours)
- as are linear solvers: (PETSc, HYPRE, Trilinos...)*
- and FFTs: (FFTW...)*
- Some applications utilise a separate communications layer or parallel execution environment: Charm++, GA...
- Some go even further to try and deliver performance portability and memory abstraction: Kokkos, RAJA...
  - The physics kernels can end up being abstracted some way from the hardware.

Ultimately, the more applications that are ported to arm, the more of these packages get ported, and the less likely you are to encounter an unknown one!

The **Arm UG GitLab wiki** is a great place to look for recipes for building libraries and dependencies, and the **OpenHPC** spec files (which support Arm hardware and LLVM) may also help.

**Arm Performance libraries provide optimized BLAS, LAPACK and FFT routines (`-L${ARMPL_DIR}/lib -larmpl_lp64 -lflang -lflangrti`)*

arm

Arm HPC Packages Wiki

https://gitlab.com/arm-hpc/packages/wikis/home

Dynamic list of common HPC applications

Provides focus for porting progress

Community driven.

Maintained by Arm, but anyone can join and contribute.

Allows developers to share recipes, and learn from progress on other applications

Provides a mechanism for tracking status of applications and package sets (e.g. OpenHPC packages, Mantevo, etc.)

Up-to-date summary of package status

Package	External URL	Last Wiki Update	BuildMaturity	CompilesARMCompiler	CompilesGCC	NEONOptimized
EPOCH	http://www.ccpp.ac.uk	19/10/17 22:10:20	NeedsPatch	Yes	Yes	
SDF	https://github.com/keithbennett/SDF	20/10/17 00:13:45	NeedsPatch	Yes	Yes	
VPIC	https://github.com/lanl/vpic	19/10/17 22:10:20		Yes	Yes	
adios	http://www.olcf.ornl.gov/center-projects/adios/	17/07/17 23:33:11		Yes	Yes	
arpack	http://www.caam.rice.edu/software/ARPACK/	17/07/17 23:33:11		Yes		
autoconf	http://www.gnu.org/software/autoconf/autoconf.html	01/08/17 21:48:30		Yes	Yes	
automake	http://www.gnu.org/software/automake	18/07/17 13:41:43		Yes	Yes	
bookleaf	https://uk-mac.github.io/BookLeaf/	17/07/17 23:33:11		Yes	Yes	
boost	http://www.boost.org	18/07/17 11:29:00		Yes	Yes	
ccs-qcd	https://github.com/fiber-miniapp/ccs-qcd	01/08/17 21:43:40		Yes		
cloverleaf	http://uk-mac.github.io/CloverLeaf/	19/10/17 22:10:20	Upstream	Yes	Yes	
cloverleaf3d	http://uk-mac.github.io/CloverLeaf3D/	24/07/17 21:41:31	Upstream	Yes	Yes	
comd	http://exmatex.github.io/CoMD	24/07/17 21:36:31	Upstream	Yes	Yes	
dgemm	http://www.nersc.gov/research-and-development/apex/apex-benc	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
ffb	http://www.ciss.iis.u-tokyo.ac.jp/riss/english/project/fluid/	24/07/17 21:36:31		Yes		
fftw	https://github.com/FFTW/fftw3	17/07/17 23:33:11		Yes	Yes	Yes
gnu-scientific-library	http://www.gnu.org/software/gsl/	18/07/17 07:08:24		Yes	Yes	
gromacs	http://www.gromacs.org/	24/07/17 21:36:31		Yes	Yes	
hdf5	http://www.hdfgroup.org	17/07/17 23:33:11		Yes	Yes	
hpc-challenge	http://icl.cs.utk.edu/hpcc/index.html	24/07/17 21:36:31	Upstream	Yes	Yes	
hpccg	http://mantevo.org/downloads/HPCCG-1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
hpcg	http://www.nersc.gov/research-and-development/apex/apex-benc	24/07/17 21:48:22	Upstream	Yes	Yes	
hypre	https://computation.llnl.gov/project/linear_solvers/software.php	17/07/17 23:33:11		Yes	Yes	
imb		17/07/17 23:33:11		Yes	Yes	
lammps	http://lammps.sandia.gov/	19/10/17 22:10:20		Yes	Yes	
libtool		18/07/17 13:41:43		Yes	Yes	
lulesh	https://codesign.llnl.gov/lulesh.php	17/07/17 23:33:11		Yes		
metis		17/07/17 23:33:11		Yes	Yes	
miniaero	http://mantevo.org/downloads/miniAero_1.0.html	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
miniamr	http://mantevo.org/downloads/miniAMR_1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
minife	http://www.nersc.gov/users/computational-systems/cori/nersc-8-a	24/07/17 21:36:31	Upstream	Yes	Yes	
minighost	http://www.nersc.gov/users/computational-systems/cori/nersc-8-a	24/07/17 21:36:31	Upstream	Yes	Yes	
minimd	http://mantevo.org/downloads/miniMD_1.2.html	24/07/17 21:36:31	NeedsPatch	Yes	Yes	
minixyce	http://mantevo.org/downloads/miniXyce_1.0.html	24/07/17 21:36:31	Upstream	Yes	Yes	
mpich		19/10/17 22:10:20	NeedsPatch	Yes	Yes	
mumps		17/07/17 23:33:11		Yes	Yes	
mvapich-2	http://mvapich.cse.ohio-state.edu	21/08/17 13:26:19	Upstream	Yes	Yes	
namd	http://www.ks.uiuc.edu/Research/namd/	24/07/17 21:36:31	NeedsPatch	Yes	Yes	

arm

Older autotools need an update

...I'm relying on a config.guess that's *way* out-of-date!

Often, the config.guess supplied with an application and used by configure will not correctly identify the platform.

This can be true for a config.guess already installed on the system and used by some configure scripts.

Obtaining up-to-date versions will fix this problem:

```
wget 'http://git.savannah.gnu.org/gitweb/?p=config.git;a=blob_plain;f=config.guess;hb=HEAD' -O config.guess

wget 'http://git.savannah.gnu.org/gitweb/?p=config.git;a=blob_plain;f=config.sub;hb=HEAD' -O config.sub
```

What is this armclang of which you speak?

...I'm relying on libtool, but it knows nothing of this "Arm compiler"

configure may not correctly identify the Arm compiler. It may not set the correct flags for libtool to use for position independent code and passing arguments through to the linker. When building libraries, this can cause problems down-the-road

Following **`configure`**, patch libtool as follows:

```
sed -i -e 's#wl=""#wl="-Wl,"#g' libtool

sed -i -e 's#pic_flag=""#pic_flag=" -fPIC -DPIC"#g' libtool
```

Use the right compiler

...I've got the compiler binary defined in several different ways!

$CC, $CXX, $FC set?

How about $F77 and $F90?

Maybe $MPICC, $MPIF90 and $MPIFORT?

...is it a parallel application that also wants to know where the serial compilers are?

...I've got gcc/icc/... hard-coded into a Makefile somewhere!

Very difficult to spot... But not to worry, I'll silently soldier on and select GCC

And since your architecture didn't match I'll just accumulate some random flags that didn't get overridden - *and then I'll continue to compile*. . .

Arm HPC Compiler OpenMP scaling

Better scaling at higher thread count

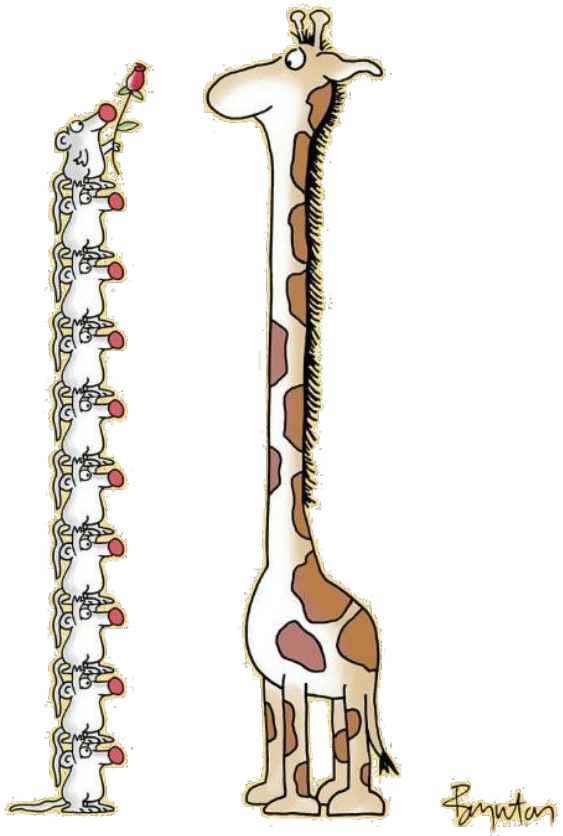

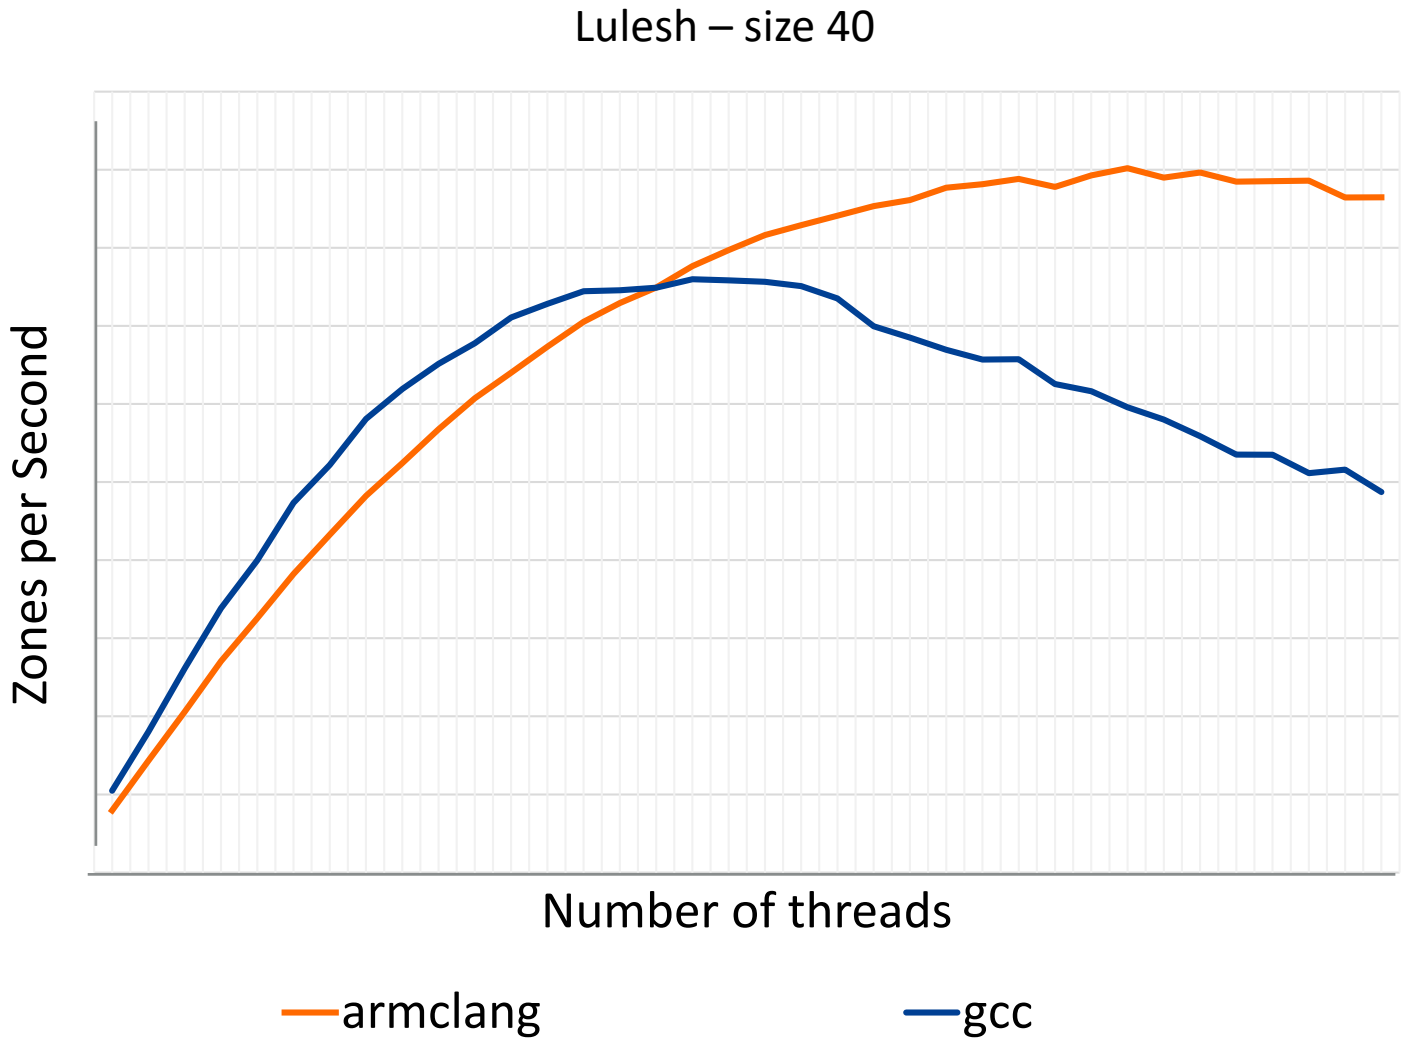

arm

Use the right compiler flags

General guidance for all Arm architectures when building with Arm HPC compilers.

1. Start with `-Ofast -mcpu=native`.
2. If Fortran application runs into issues with `-Ofast`, try `-Ofast -fno-stack-arrays` to force automatic arrays on the heap.
3. If `-Ofast` is not acceptable and produces wrong results due to reordering of math operations, use `-O3 -ffp-contract=fast`.
4. If `-ffp-contract=fast` does not produce correct results, then use `-O3`.

Power users: `armflang -###` shows the expanded compile line.

Stick to the standard

...but I'm relying on non-standard extensions…

For example ISNAN, COSD, or very very long lines...

Or compiler-specific intrinsics, mm_prefetch, SSE calls etc.

There may be an alternate code path that can be used already. Of possibly the code isn't critical and can be deactivated for now, or an equivalent call can be used, or you could write one?

...I'm relying on a very forgiving, and non-pedantic compiler!

some compilers let you get a away with an awful lot.

A developer can get used to that.

Pragmas to control vectorization

#pragma clang loop vectorize(assume_safety)

- Allows the compiler to assume that there are no aliasing issues in a loop

#pragma clang loop unroll_count(_value_)

- Forces a scalar loop to unroll by a given factor

#pragma clang loop interleave_count(_value_)

- Forces a vectorized loop to be interleaved by a given factor

Do you support language feature X?

...I'm relying on some language features you don't support!

For example, ArmFlang has excellent support for Fortran 2003:

https://developer.arm.com/products/software-development-tools/hpc/arm-fortran-compiler/fortran-2003-status

But the 2008 standard isn't fully supported yet:

https://developer.arm.com/products/software-development-tools/hpc/arm-fortran-compiler/fortran-2008-status

...while the 2018 standard is on the horizon. And Fortran 202X will add yet more capability.

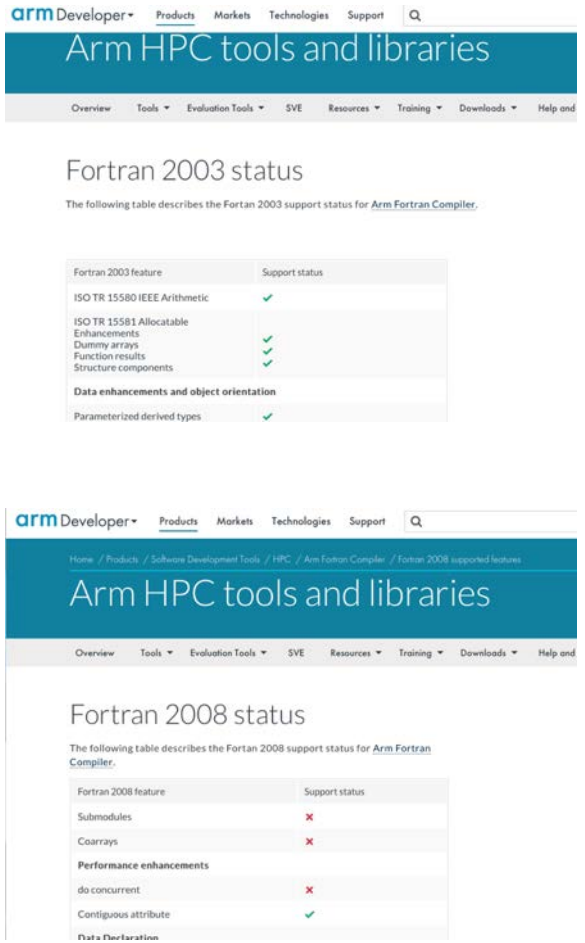

arm

./configure && make && sudo make install … almost

If root has a minimal environment, using sudo can break compiler license verification

If your application uses libtool during installation, you may see something like this:

/home/user.0004/johlin02/openmpi-3.1.0/build/libtool: line 10554: **armclang: command not found**

Or maybe this:

clang-5.0: error: **Failed to check out a license.** See below for more details.

Don’t panic! Break it into steps. Instead of `` `sudo make install` `` just do:

```
$ sudo -i
$ module load <compiler module>
$ make install
```

Test your tests; talk to the expert

...my test suite never passes, *everyone* knows that!

Often the test suites are a work in progress.

For example, out-of-the-box test appears to have the wrong reference solution. Earlier commits give the conditions used for reference solutions (intel, IEEE etc.), repeating gives a new reference solution, for which Arm and GCC agree!

...I've got *some* Arm support, but no one is looking after it!

Someone had a go, a while back, possibly just-for-fun.

Committed it to the repo and moved on.

Hasn't been maintained, and doesn't actually work.

...but looks like it might, briefly.

arm

Test your tests; talk to the expert

...my test suite never passes, *everyone* knows that!

Often the test suites are a work in progress.

For example, out-of-the-box test appears to have the wrong reference solution. Earlier commits give the conditions used for reference solutions (intel, IEEE etc.), repeating gives a new reference solution, for which Arm and GCC agree!

...I've got *some* Arm support, but no one is looking after it!

Someone had a go, a while back, possibly just-for-fun.

Committed it to the repo and moved on.

Hasn't been maintained, and doesn't actually work.

...but looks like it might, briefly.

arm

Arm uses a weak memory model

...*I'm* not doing anything wrong but seemingly get a weird race condition!

Some HPC codes we came across had their own parallelization implementations

- Usually based directly on top of pthreads
- Written to have more control over the threads of execution and how they synchronize
- Some had no problems working with AArch64's **weakly ordered memory system**
- Others exhibited issues in multi-threaded modes that were particularly hard to diagnose without a detailed investigation into how the multi-threaded mode was implemented
  - Problems are almost always down to a lock-free thread interaction implementation
  - Key symptom: correct operation on a strongly ordered architecture, failure on weakly ordered

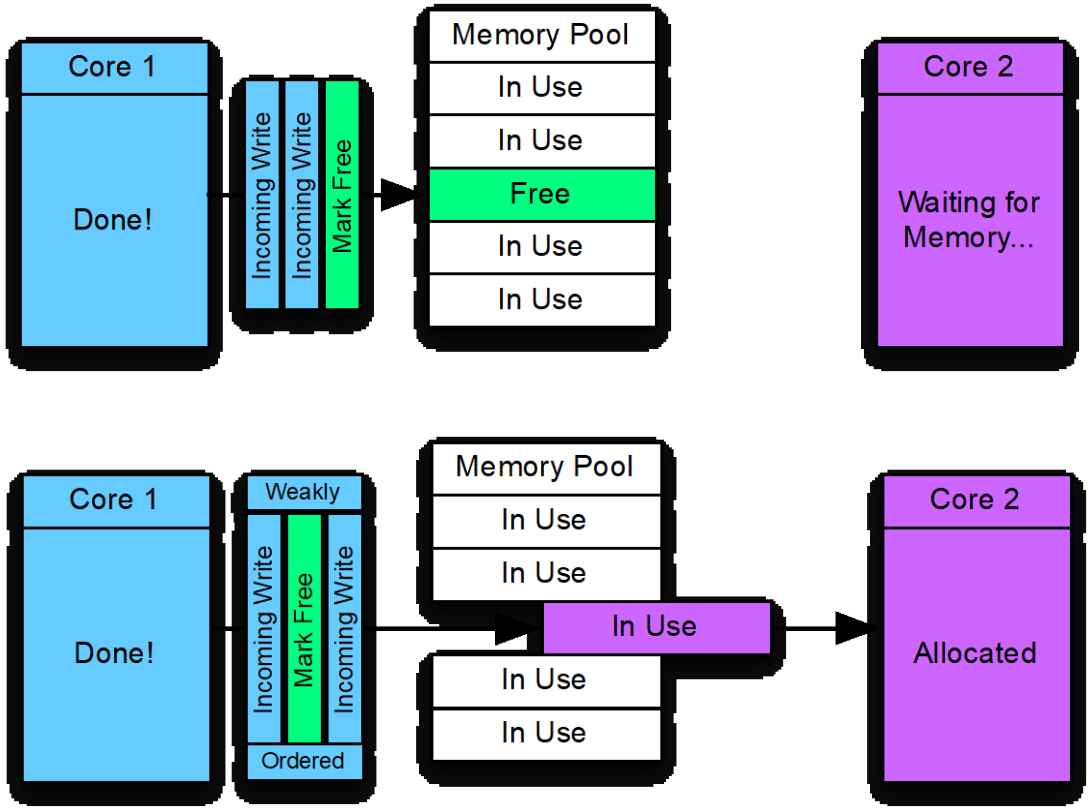

arm

Psst! 1/0 == 0 on ARM

… my results are all zero, but tests for division by zero never fail?

For example…

```
#include <stdio.h>



int main(int argc, char ** argv)
{
  int x = argc - 1;
  printf("%d\n", 1 / x);
  return 0;
}
```

Skylake

$ gcc x.c && ./a.out

Floating point exception: 8

ThunderX2

$ gcc x.c && ./a.out

0

arm

Step 1: Optimization by Linker

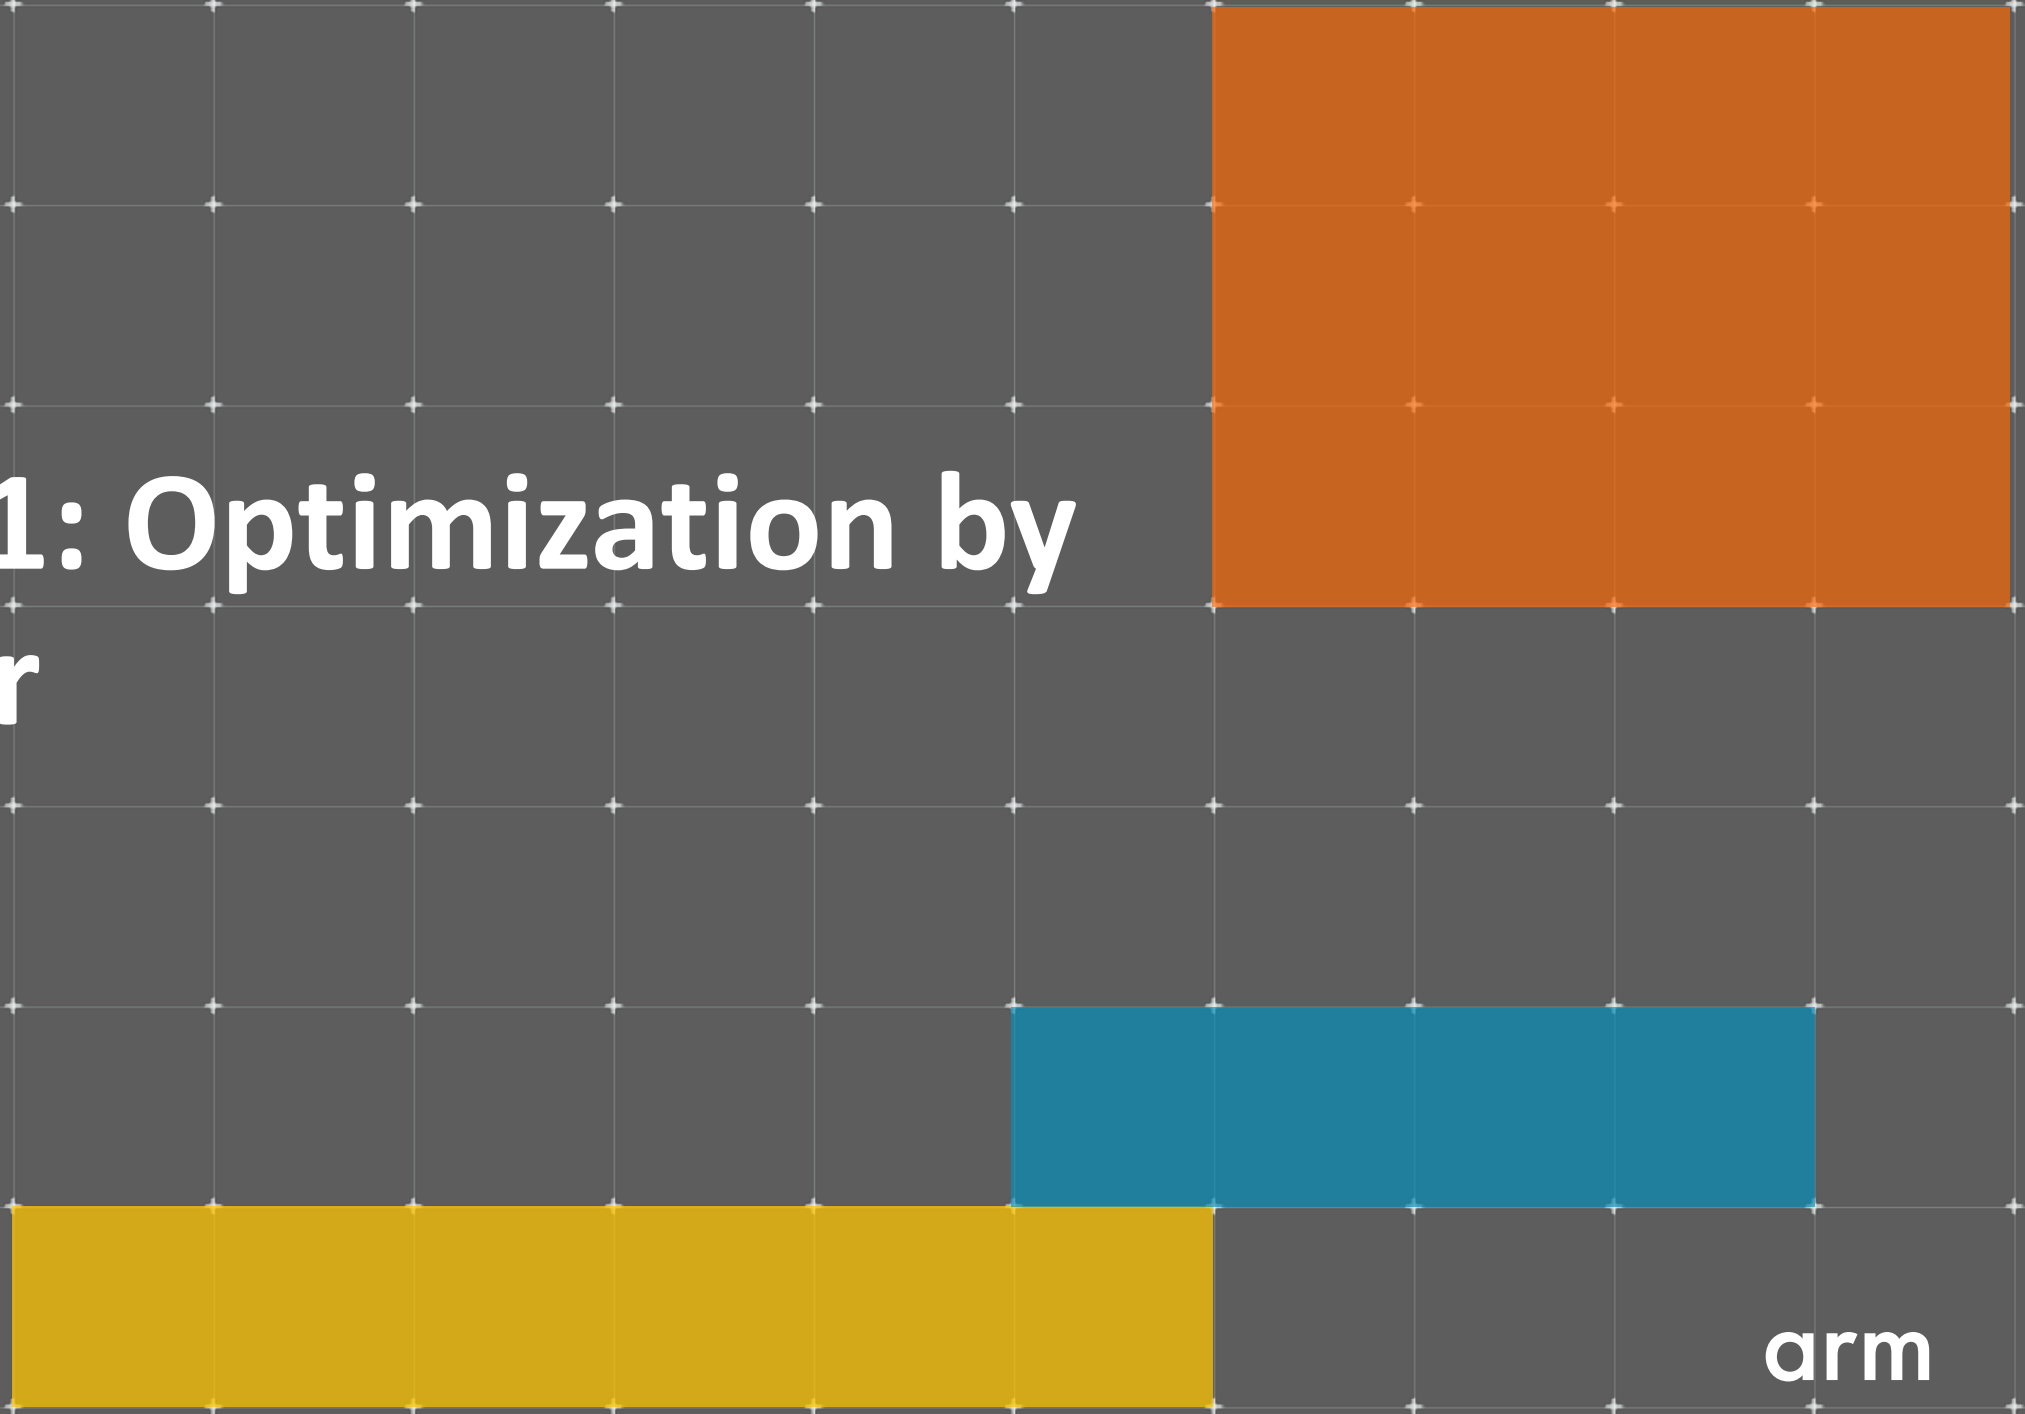

arm

arm PERFORMANCE LIBRARIES

Optimized BLAS, LAPACK and FFT

Commercially supported by Arm

Best in class performance

Validated with NAG test suite

Commercial 64-bit Armv8-A math libraries

- Commonly used low-level math routines - BLAS, LAPACK and FFT
- Provides FFTW compatible interface for FFT routines
- Batched BLAS support

Best-in-class serial and parallel performance

- Generic Armv8-A optimizations by Arm
- Tuning for specific platforms like Cavium ThunderX2 in collaboration with silicon vendors

Validated and supported by Arm

- Available for a wide range of server-class Arm-based platforms
- Validated with NAG's test suite, a de-facto standard

DGEMM performance on Cavium ThunderX2

Excellent serial and parallel performance

Achieving very high performance at the node level leveraging high core counts and large memory bandwidth

Single core performance at 95% of peak for DGEMM

Parallel performance significantly higher than OpenBLAS

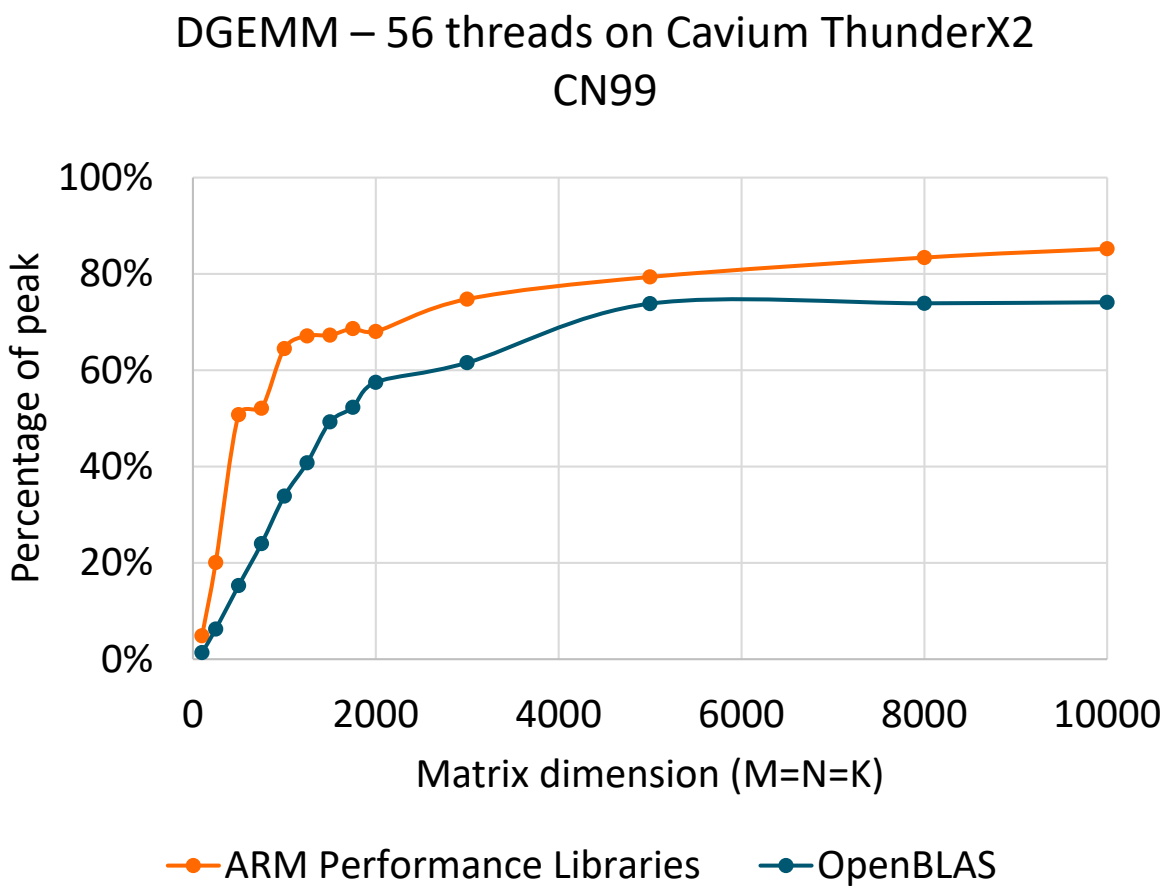

arm

DGEMM performance

Arm Performance Libraries using Arm Compiler+Arm PL vs GCC+OpenBLAS

DGEMM shows good single socket performance from tiny cases upwards.

Peak performance of serial cases at over 97% is over 10% better than OpenBLAS.

Parallel scaling is also high, +90% of peak

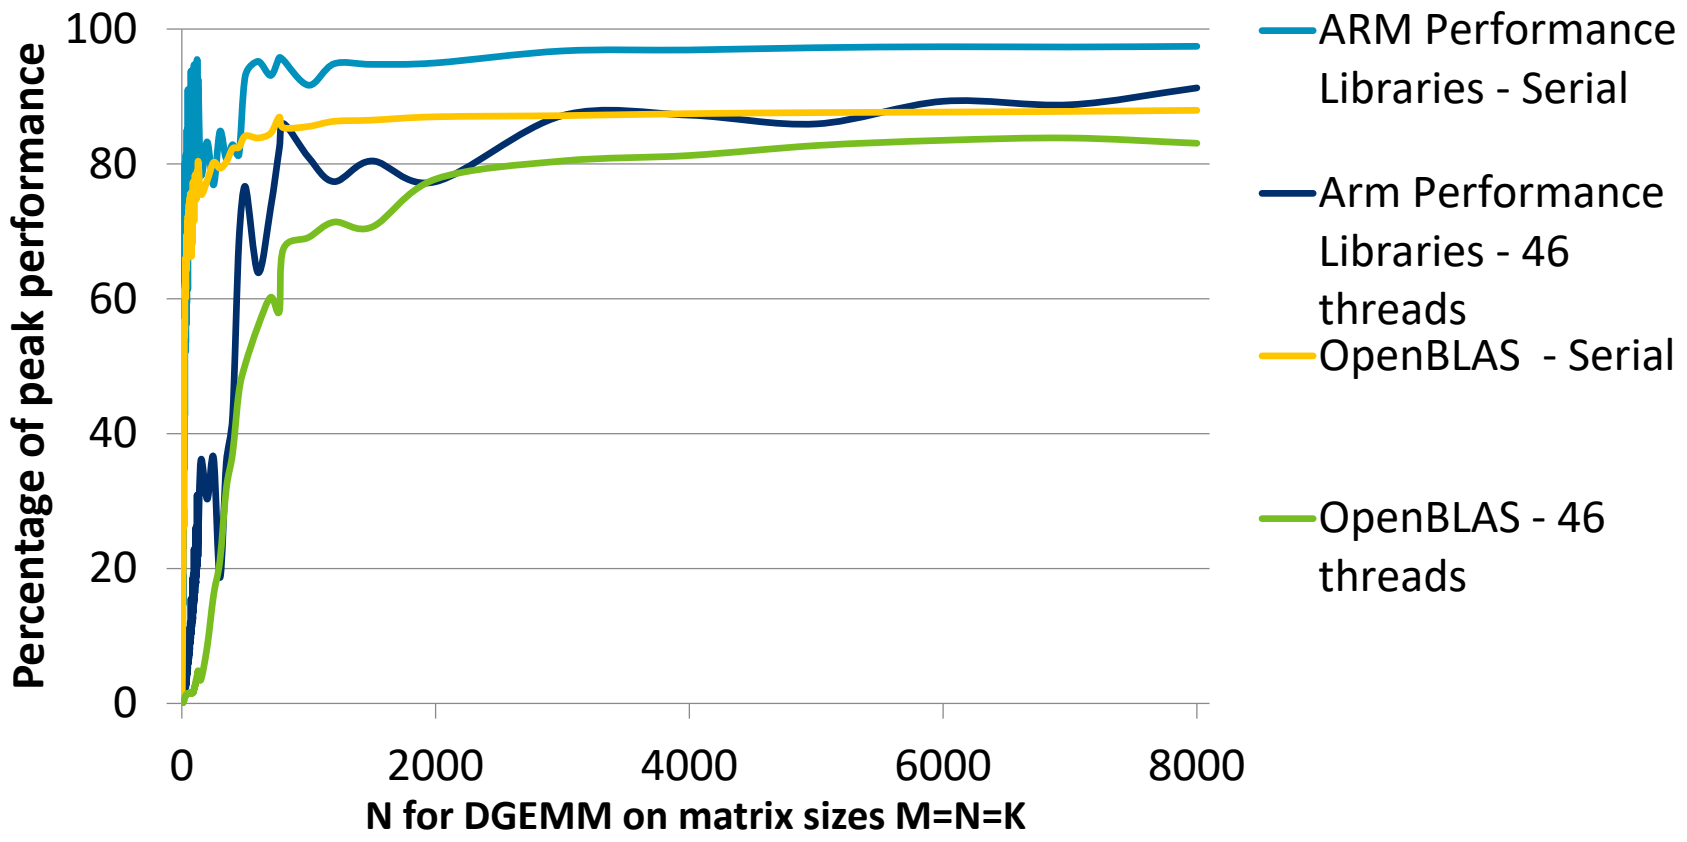

arm

Arm Performance Libraries

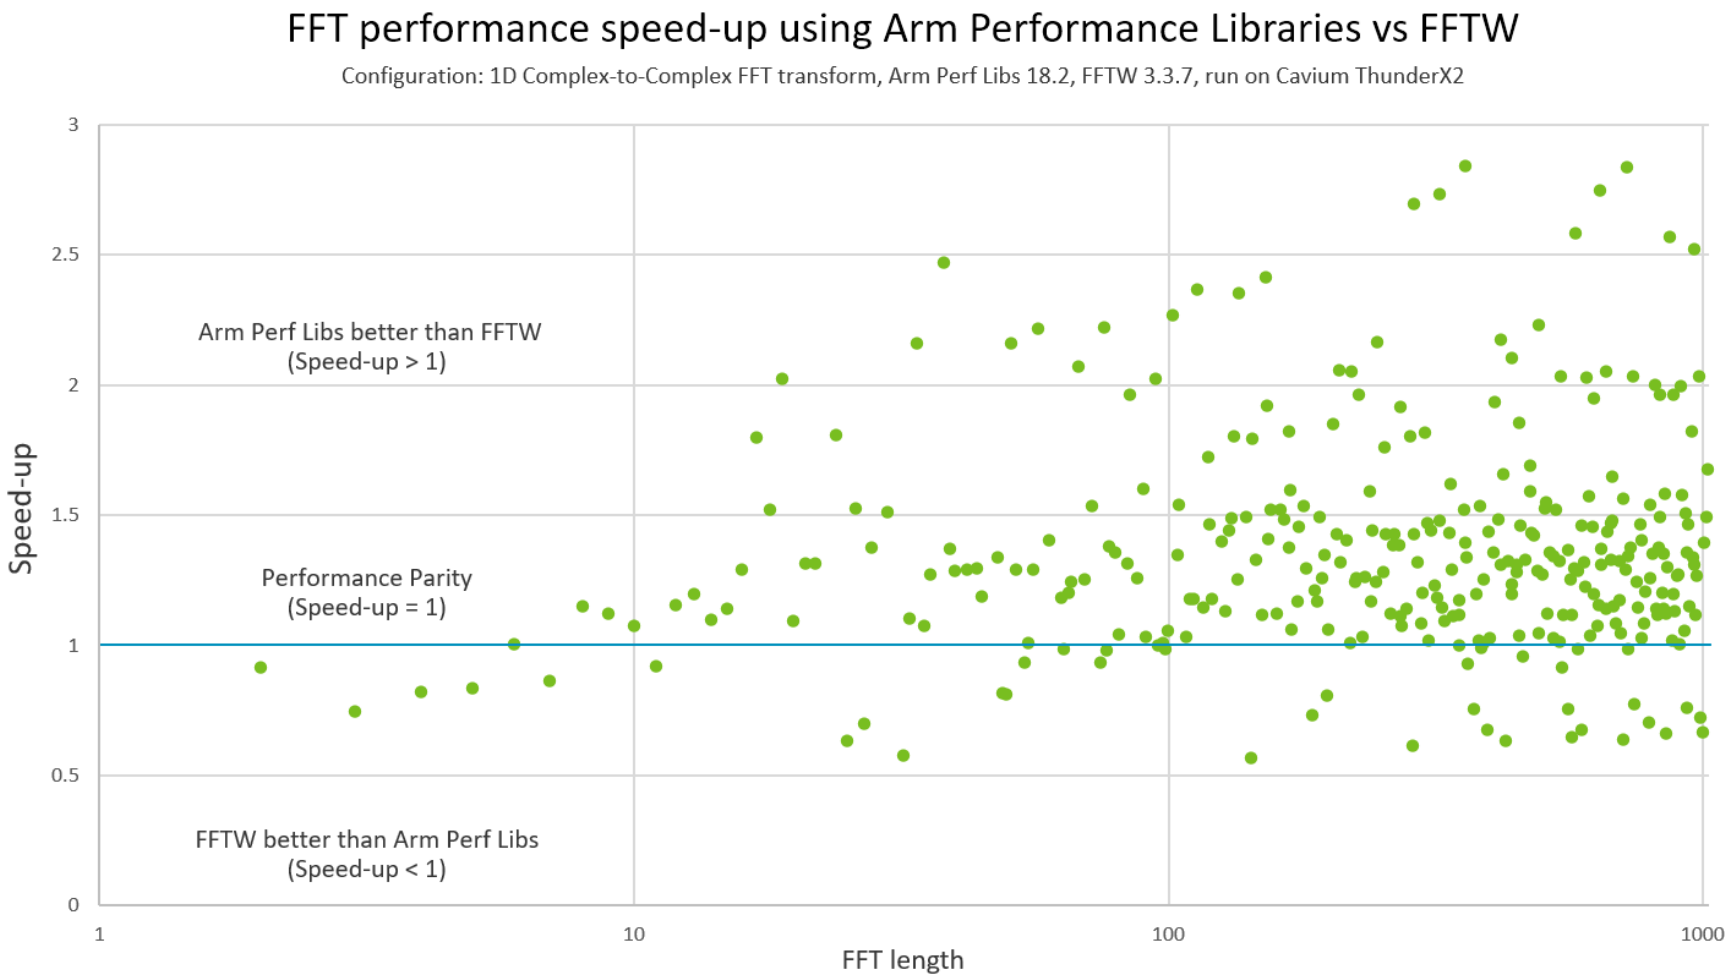

arm

Micro-architectural tuning

In order to achieve the best performance possible on all partner systems we need to do different micro-architectural tuning

All BLAS kernels are handwritten in *assembly code* in order to maximise overall performance

Different micro-architectures sometimes need fundamental differences in the instruction ordering – or even the instructions used

At run-time this work should all be transparent to the user

However multiple packages are typically available for users to choose from, and they need to load the appropriate module to set up their paths

Currently available are versions for:

- A57
- A72
- Cavium ThunderX
- Cavium ThunderX2
- Generic AArch64

Step 2: Optimization by Iteration

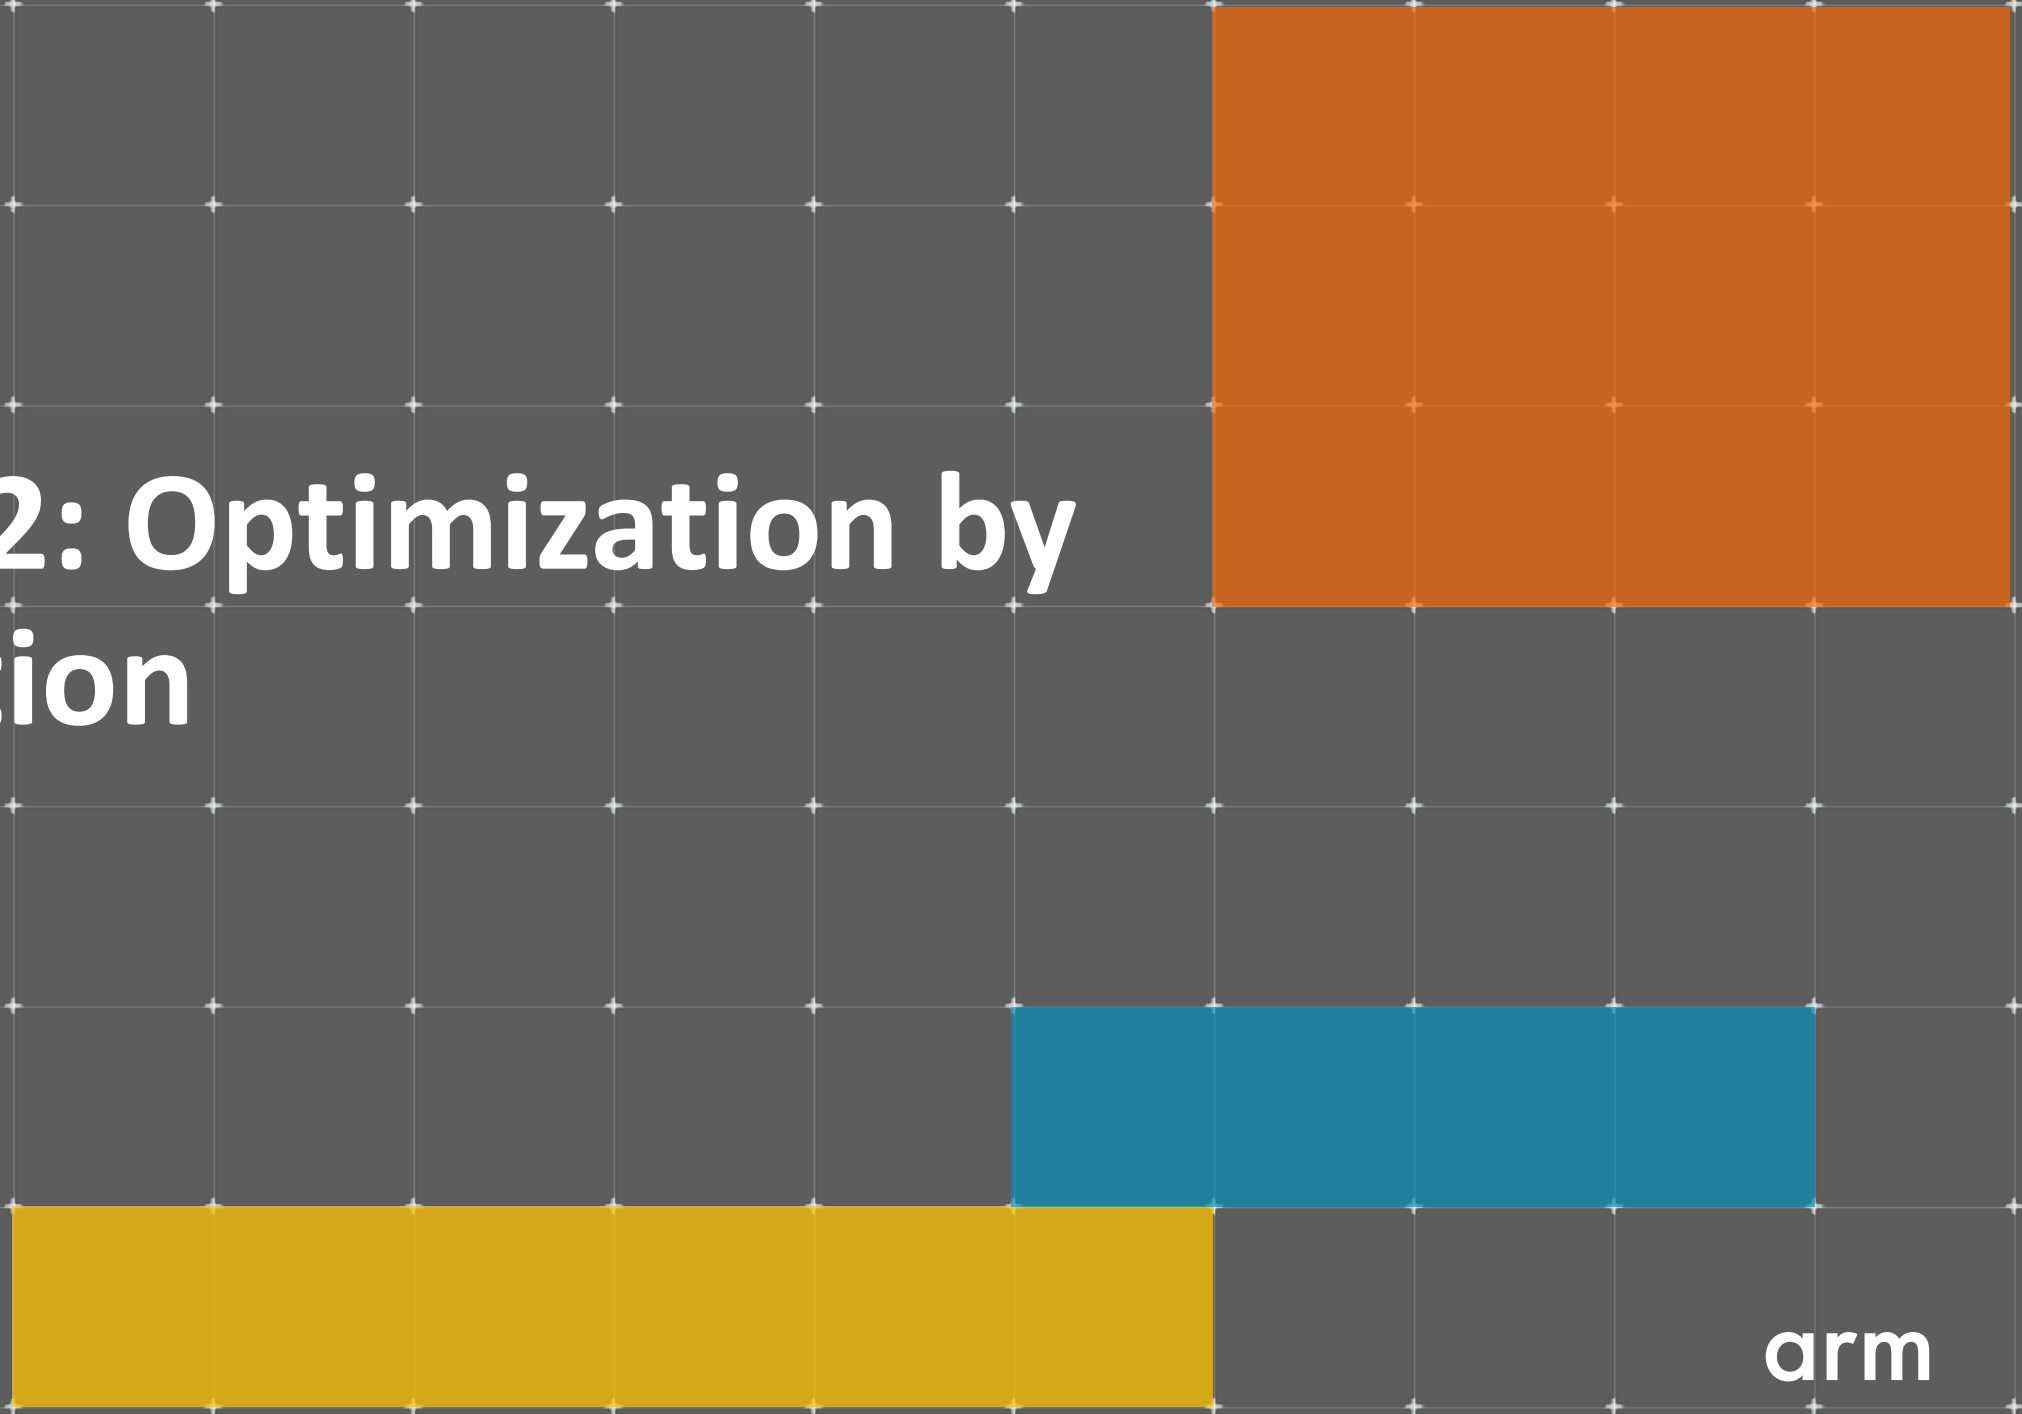

arm

Identifying and Resolving Performance Issues

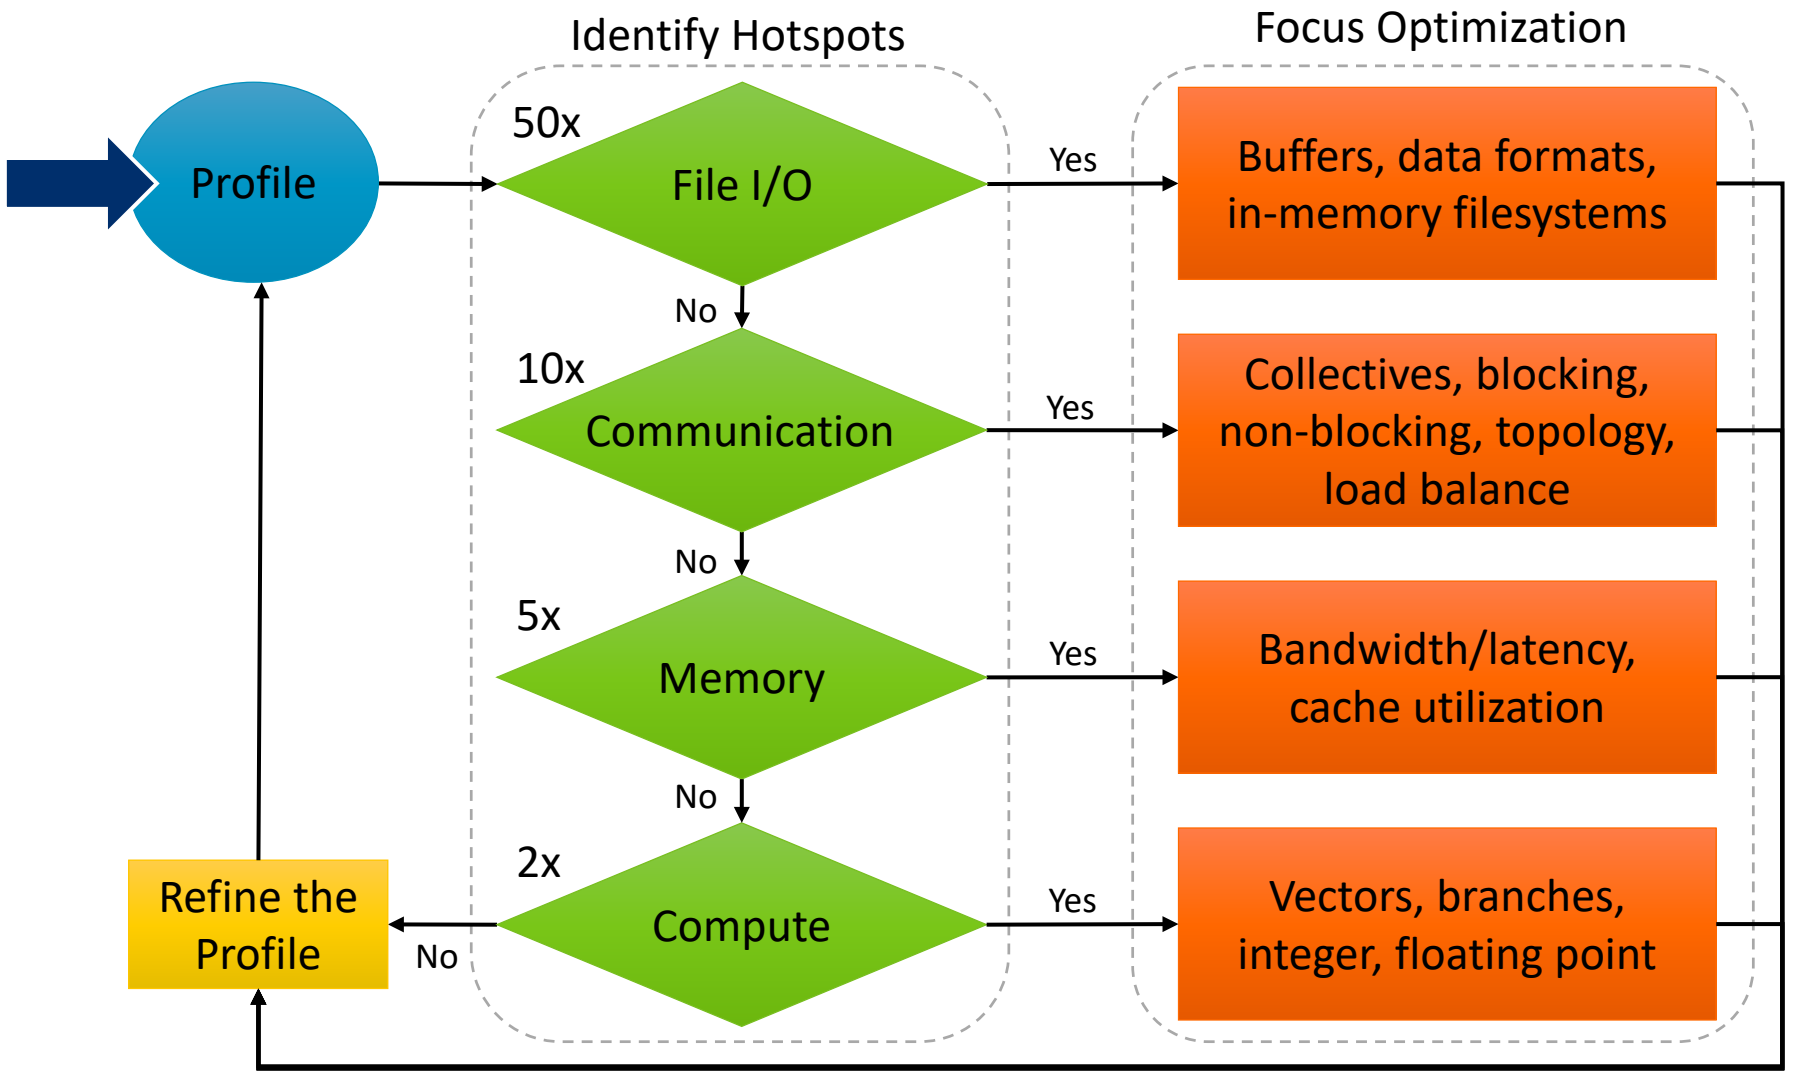

arm

VI-HPS and the tools ecosystem

See the http://www.vi-hps.org/tools/ for an excellent view of the tools ecosystem.

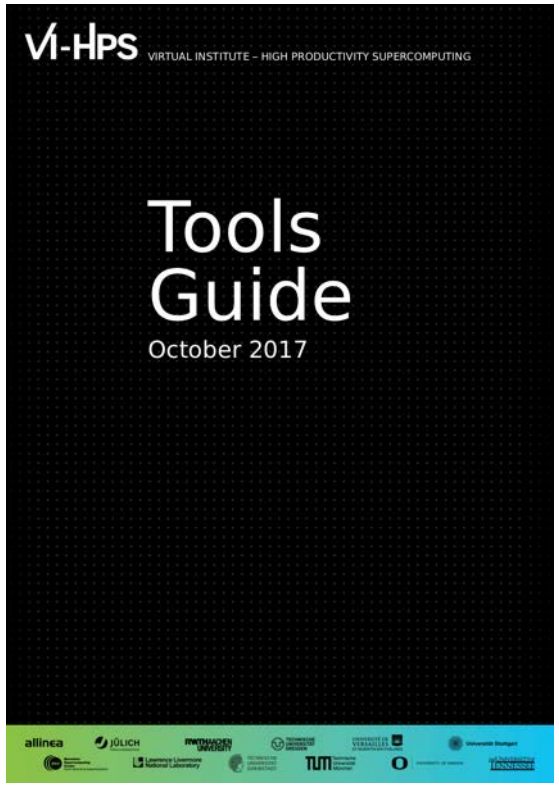

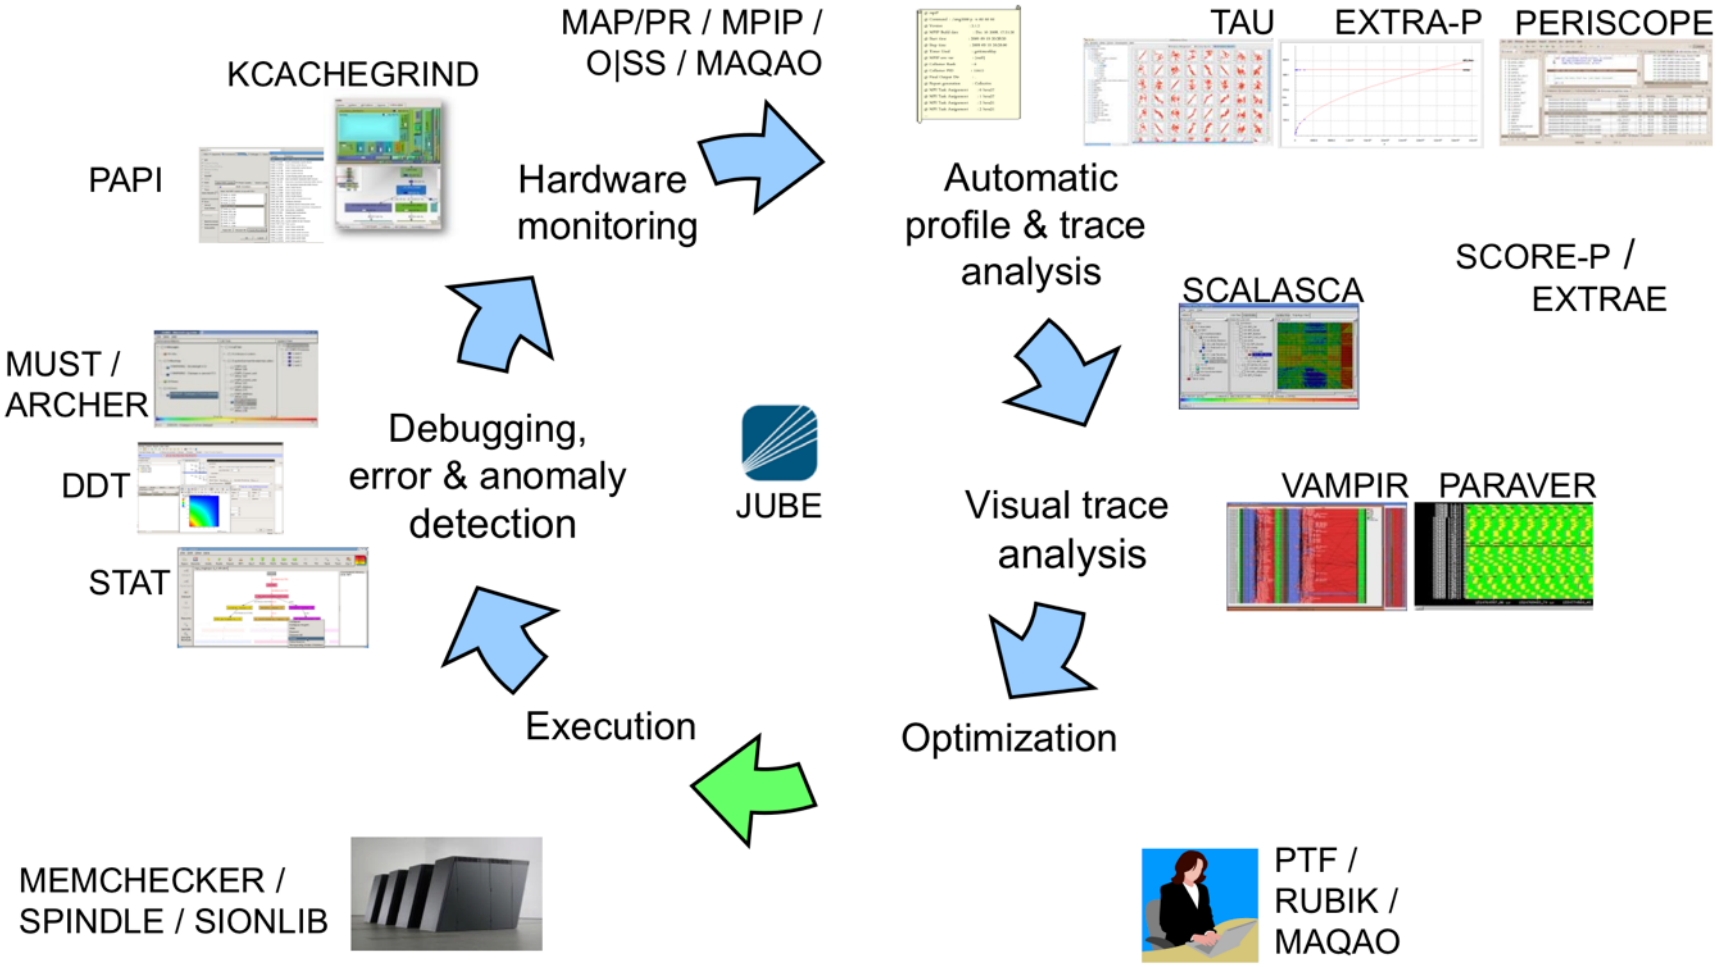

arm

Arm Forge Professional

A cross-platform toolkit for debugging and profiling

Commercially supported by Arm

Fully Scalable

Very user-friendly

The de-facto standard for HPC development

- Available on the vast majority of the Top500 machines in the world
- Fully supported by Arm on x86, IBM Power, Nvidia 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 petaflopic applications)

Easy to use by everyone

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

arm

Run and ensure application correctness

Scalable tool for interactive and automated debugging

- Run with the representative workload you started with
- Ensure application correctness with **Arm Forge Professional**
- Integrate Arm Forge to your CI workflows for automated & non-interactive debugging

Examples:

```
$> ddt -offline mpirun -n 48 ./example
$> ddt mpirun -n 48 ./example
```

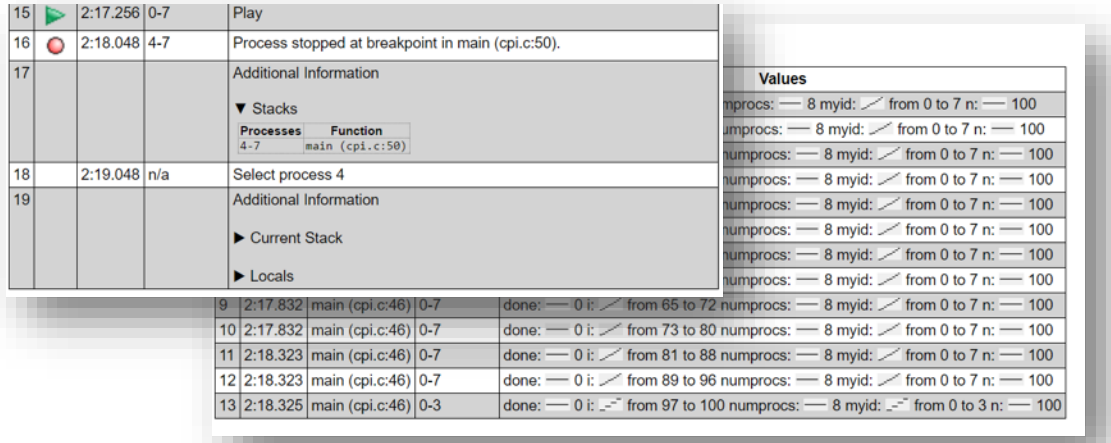

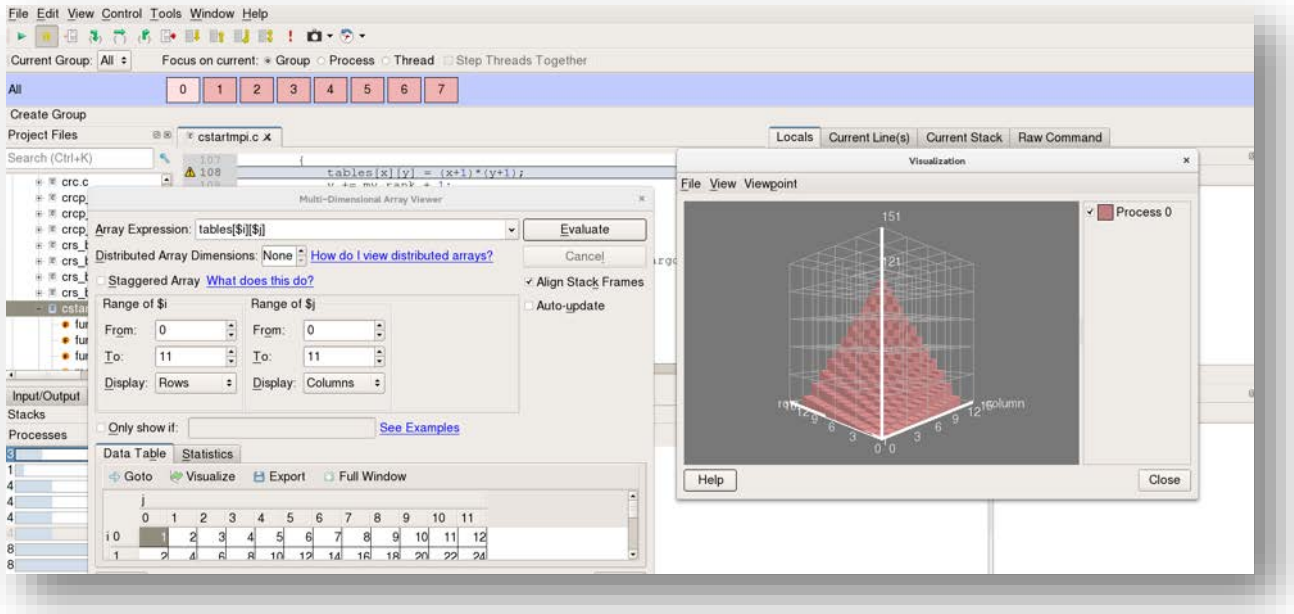

Optimize the application for Arm

Identify bottlenecks and rewrite some code for better performance

- Run with the representative workload you started with
- Measure all performance aspects with **Arm Forge Professional**

Examples:

```
$> map -profile mpirun –n 48 ./example
```

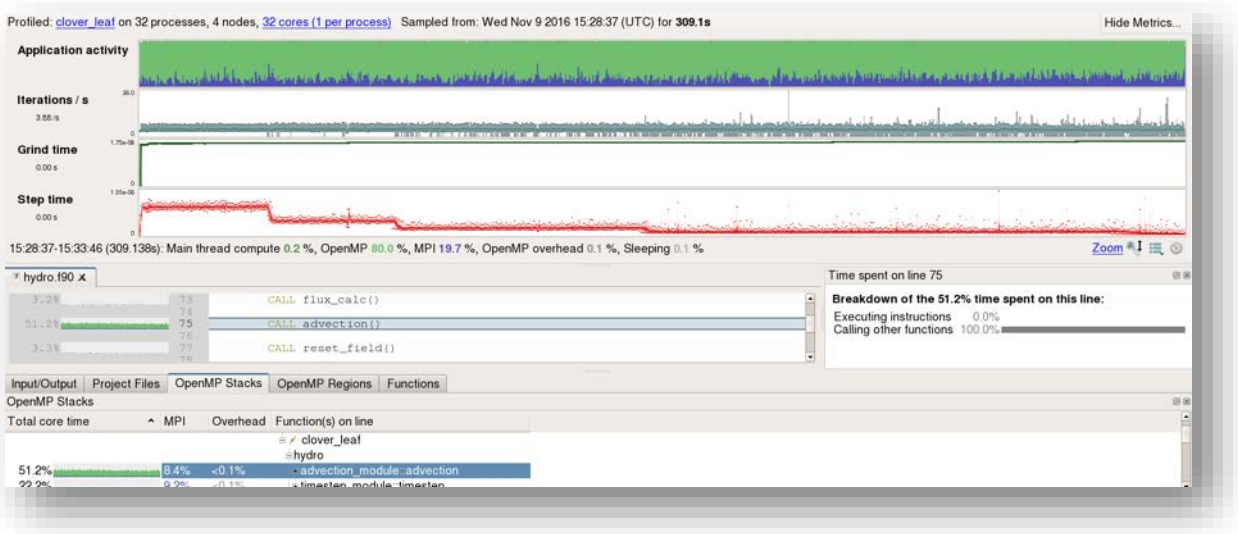

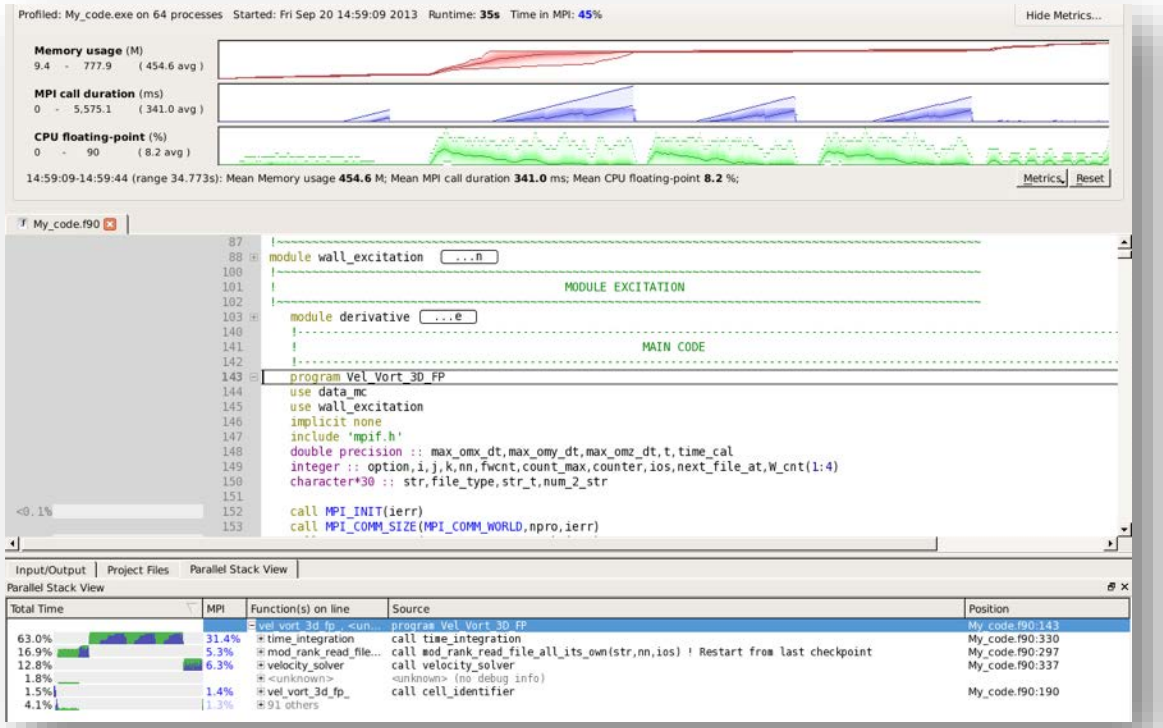

arm

Performance metrics in MAP

- MAP has a set of available metrics
  - Designed to support generic case of performance profiling
  - Presented along with call stack timelines
- Time classification
  - Based on call stacks – MPI, OpenMP, I/O, Synchronization
- Specific metrics
  - MPI call and message rates (P2P and collective bandwidth)
  - I/O data rates (POSIX or Lustre)
  - Energy data (IPMI or RAPL for Intel)
- Instruction information (hardware counters)
  - X86 – instruction breakdown + PAPI
  - Aarch64 – Perf metric for hardware counters

Custom metrics interface

- MAP supports the development of user metrics
- We provide a custom metric interface
  - API for safe calls to common functions
- Let’s you develop your own metrics of interest
  - Link to application metrics (units / s, error values)
  - Link to libraries (specialist communication or I/O)
  - System metrics (custom energy monitors)
- Integrates directly into MAP and Performance Reports
  - XML files for aggregation methods
- Need to consider overheads and thread safety

Custom metric – MUSCLE2 & LU error terms

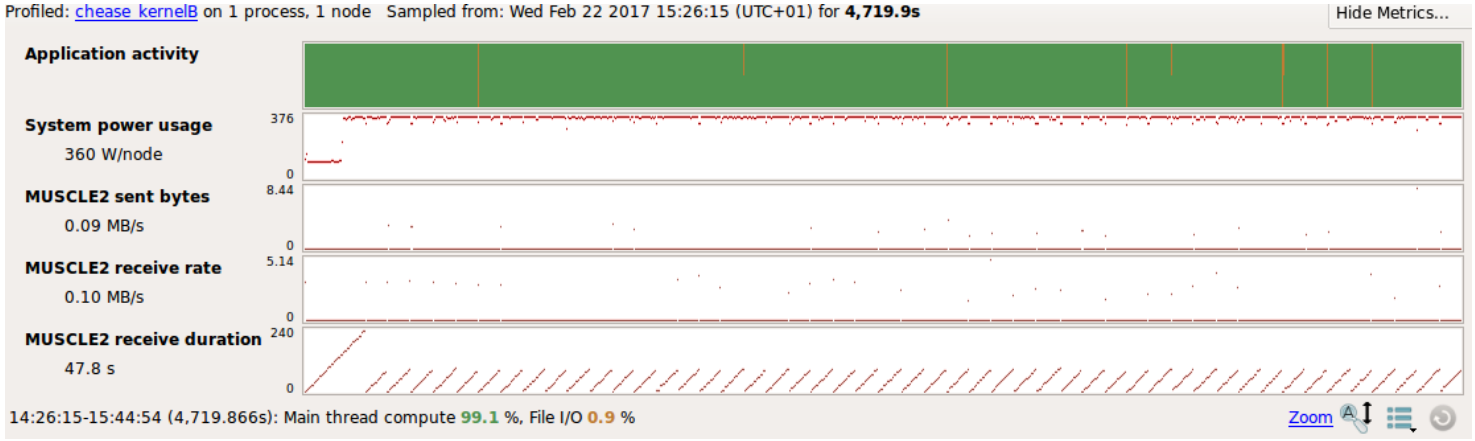

Instrumentation of MUSCLE2 library

Record communication volumes and times

Data collected along with ‘normal’ MAP metrics

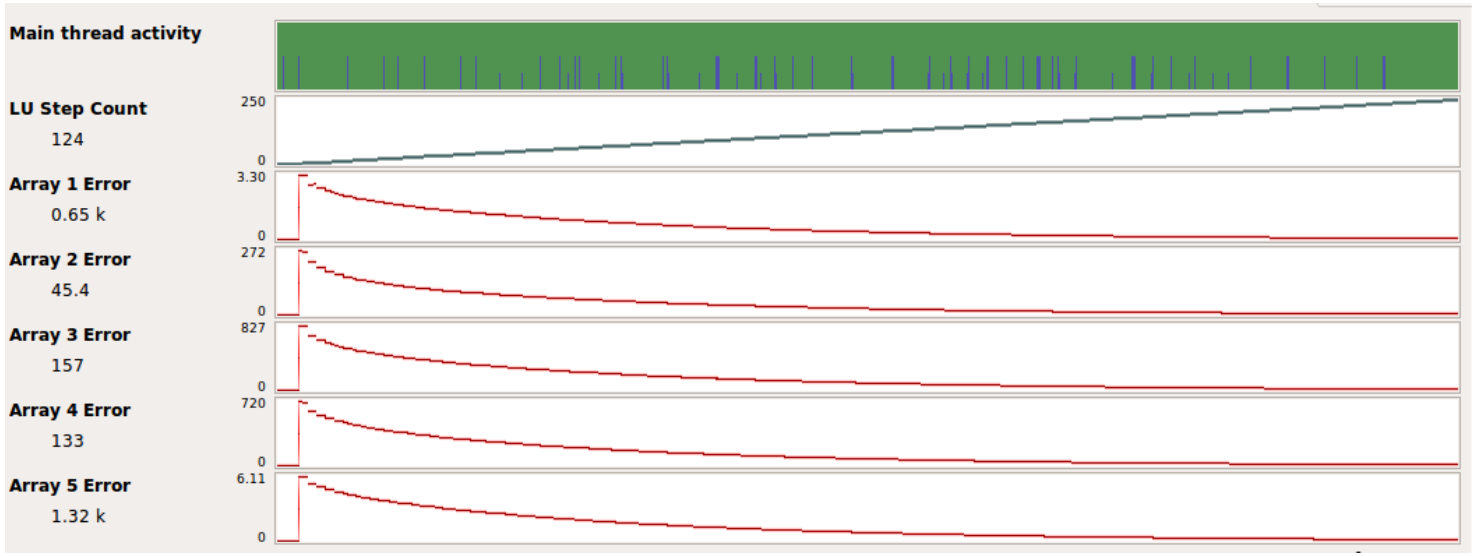

- Instrumentation of NPB LU application
- Record error terms of solve
- Plot over time and step count for optimisation

arm

JSON Export

Export map profile data to JSON file

- Command line or GUI
- Provides meta data + samples

```
operks@eslogin001:/work/y14/y14/operks/CloverLeaf_ref> map --export=clover.json clover_leaf_6p_1n_4t_2017-02-09_12-18.map
Loading MAP file clover_leaf_6p_1n_4t_2017-02-09_12-18.map...
...done
Collecting samples...
...done
Calculating...
...done
Collecting samples...
...done
Calculating...
...done
MAP generated clover.json
```


```
{
  "info": {
    "command_line": "aprun -n 6 -N 6 -S 3 -d 4 ./clover_leaf",
    "create_version": "7.0",
    "machine": "mom5",
    "metrics": {
      "memory_per_node": {
        "max": 67658141696,
        "mean": 67658141696,
        "min": 67658141696,
        "sum": 67658141696,
        "var": 0
      },
      "num_cores_per_node": {
        "max": 48,
        "mean": 48,
        "min": 48,
        "sum": 288,
        "var": 0
      },
      "num_omp_threads_per_process": {
        "max": 3,
        "mean": 3,
        "min": 3,
        "sum": 18,
        "var": 0
      },
```


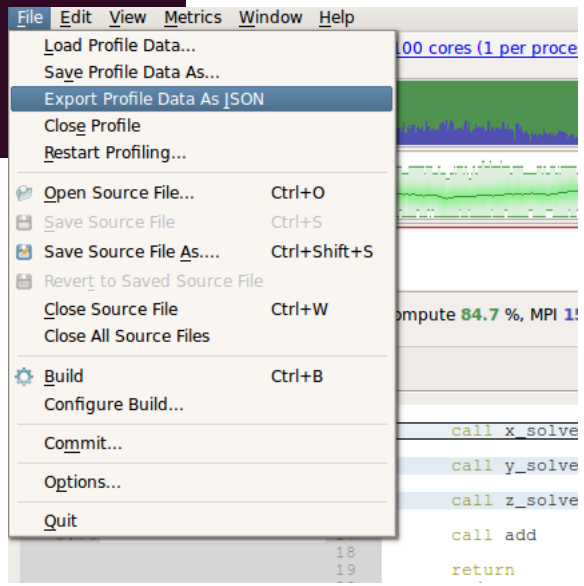

arm

Quantum Collisions Success Story

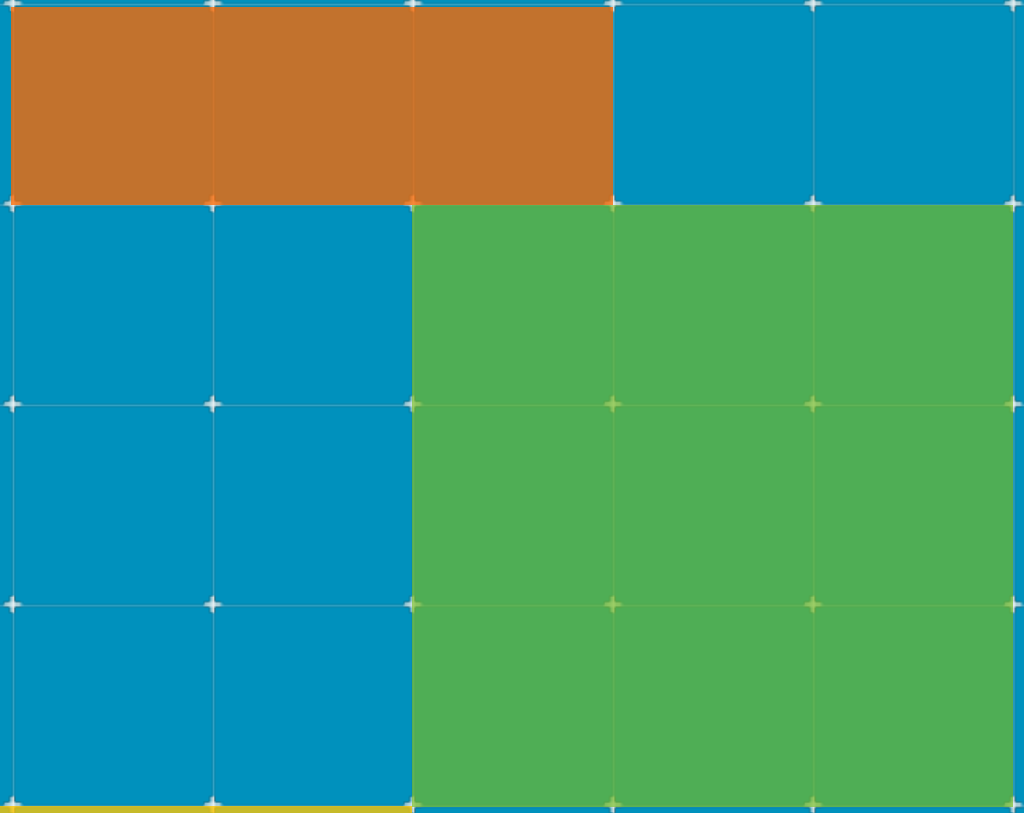

arm

CCC and the ORNL GPU Hackathon @ Pawsey

Quantum collisions in atomic and molecular physics

CCC: Quantum mechanics

- Fusion energy
- Laser science
- Lighting industry
- Medical imaging / therapy
- Astrophysics

Igor Bray, Head of Physics and Astronomy, and the Theoretical Physics Group, in the Faculty of Science and Engineering, at Curtin University

Curtin University

arm

Initial Profile

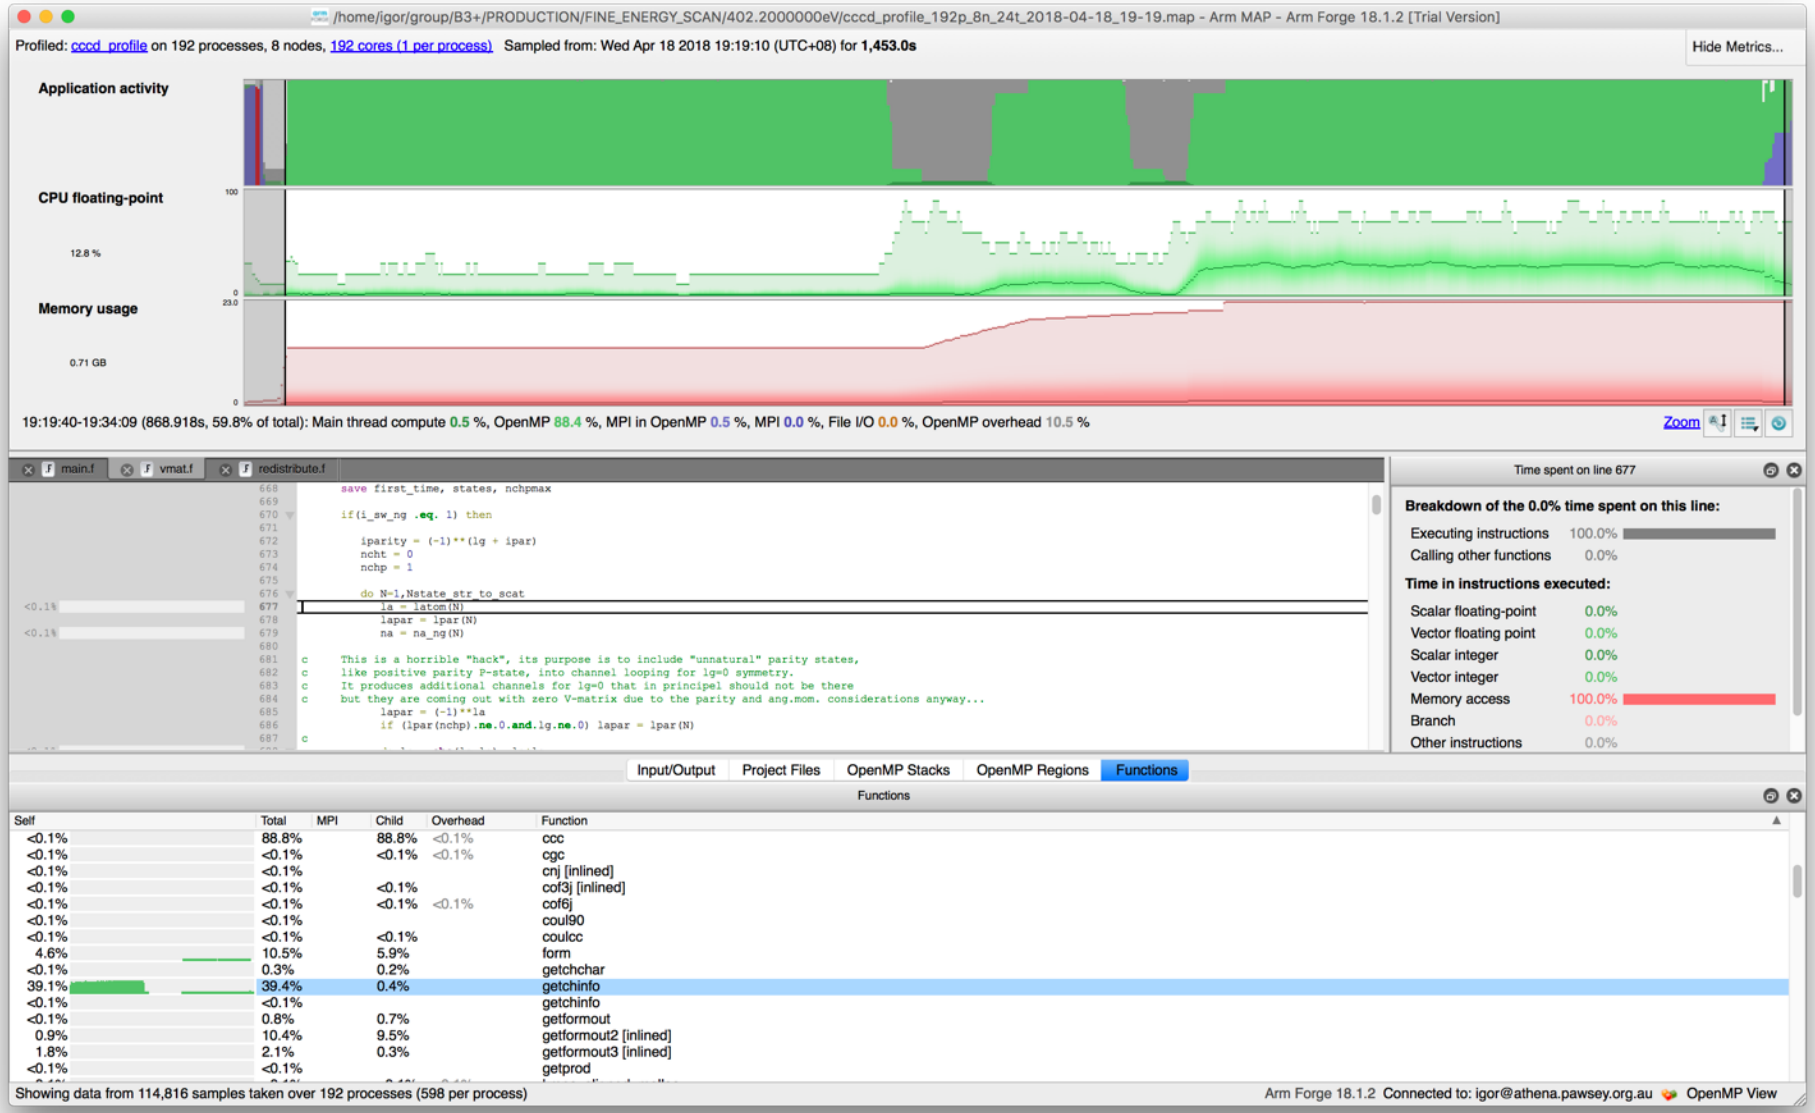

Load balancer is imbalanced?

Before:

0	8	0	-10	199	329	492	1.21	13530	0	89	-1	**91%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff
1	8	0	-7	591	573	872	1.97	45150	0	350	0	**80%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff
2	8	0	-16	894	762	1153	2.28	77028	0	607	1	**86%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff
3	8	0	-24	916	886	1331	2.05	99681	0	766	2	**91%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

Initial Profile

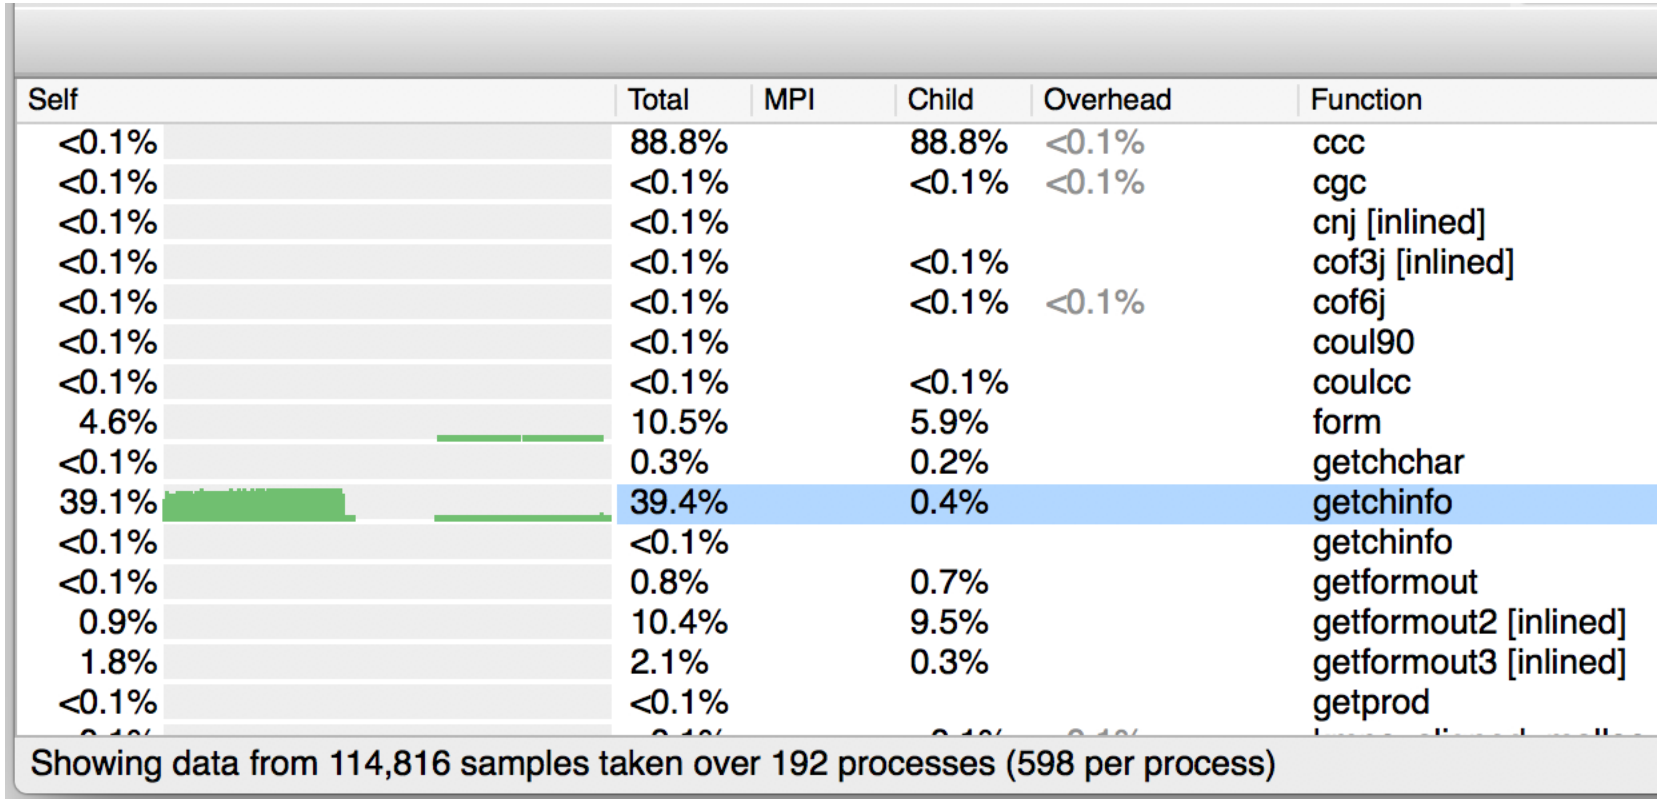

Self	Total	MPI	Child	Overhead	Function
<0.1%	88.8%		88.8%	<0.1%	ccc
<0.1%	<0.1%		<0.1%	<0.1%	cgc
<0.1%	<0.1%				cnj [inlined]
<0.1%	<0.1%		<0.1%		cof3j [inlined]
<0.1%	<0.1%		<0.1%	<0.1%	cof6j
<0.1%	<0.1%				coul90
<0.1%	<0.1%		<0.1%		coulcc
4.6%	10.5%		5.9%		form
<0.1%	0.3%		0.2%		getchchar
39.1%	39.4%		0.4%		getchinfo
<0.1%	<0.1%				getchinfo
<0.1%	0.8%		0.7%		getformout
0.9%	10.4%		9.5%		getformout2 [inlined]
1.8%	2.1%		0.3%		getformout3 [inlined]
<0.1%	<0.1%				getprod

Showing data from 114,816 samples taken over 192 processes (598 per process)

Surprise! Didn’t expect that.

arm

Results and Final Profile

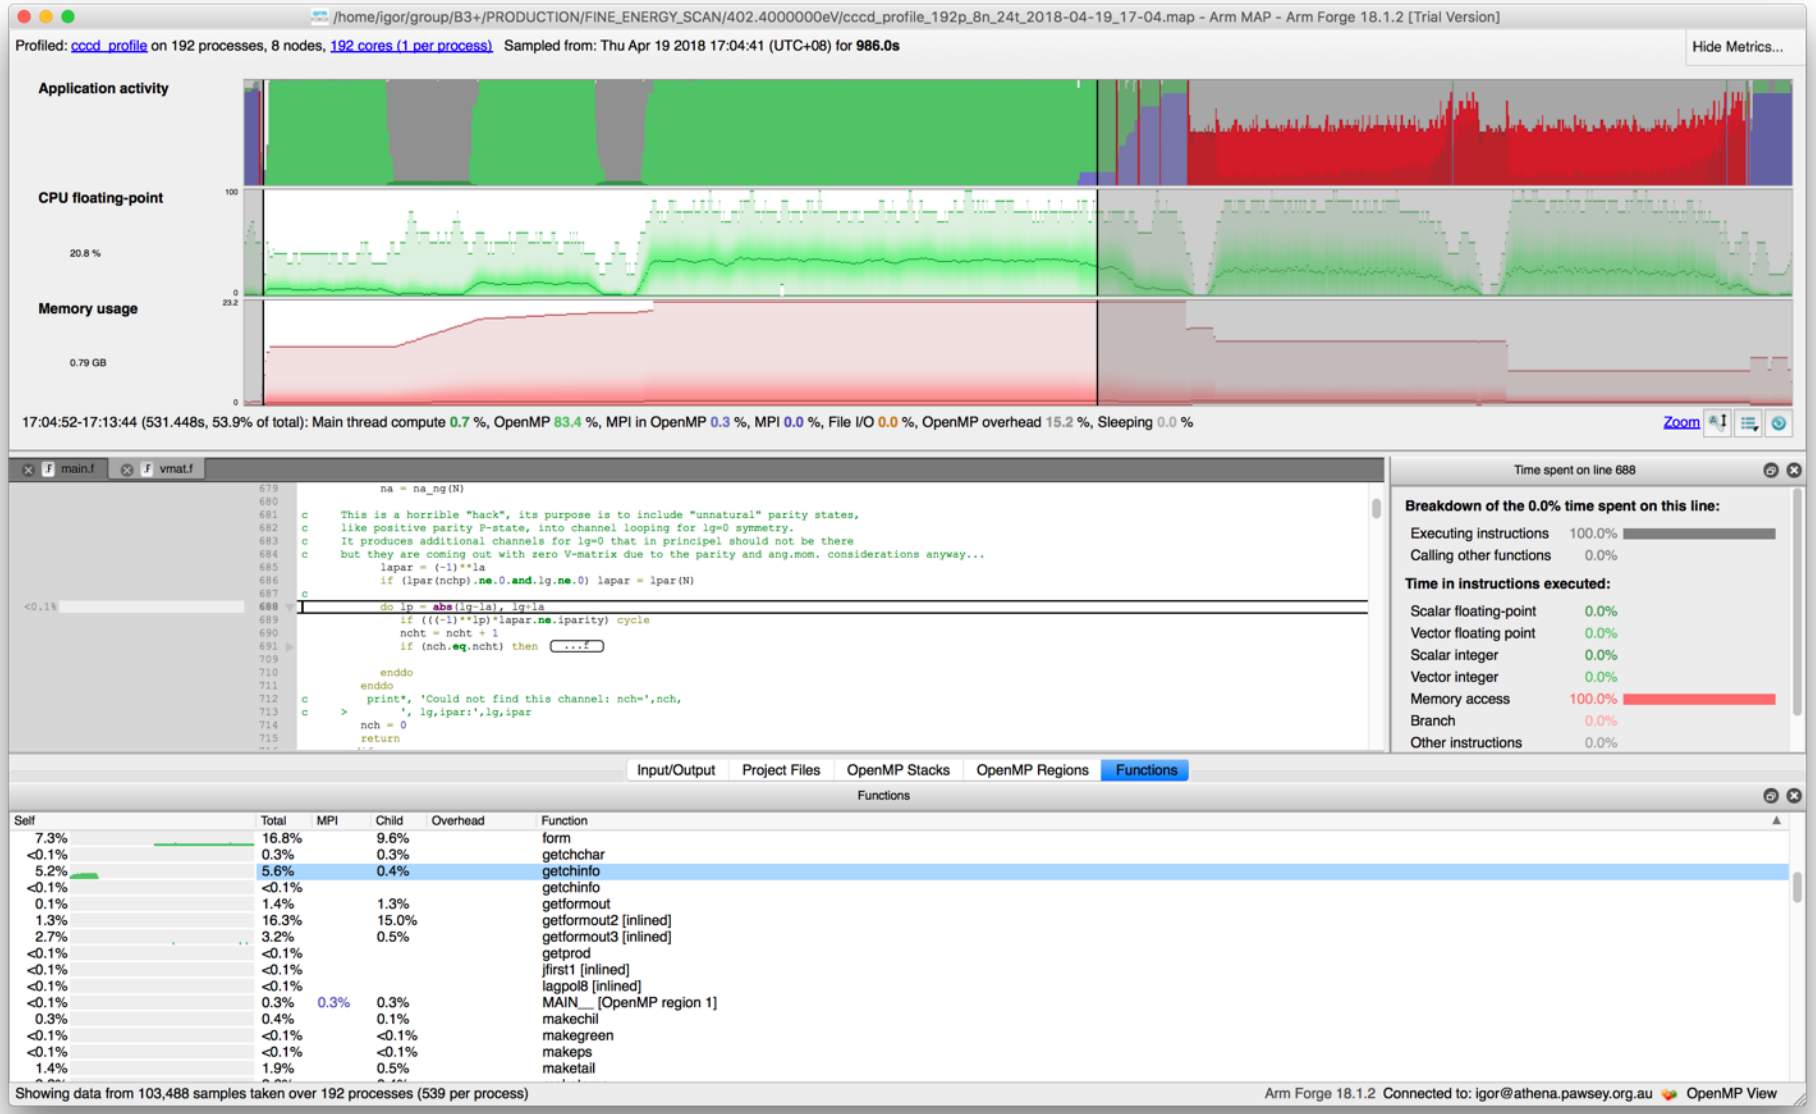

Results and Final Profile

4.6%	10.5%	5.9%	form
<0.1%	0.3%	0.2%	getchchar
39.1%	39.4%	0.4%	getchinfo
<0.1%	<0.1%		getchinfo
<0.1%	0.8%	0.7%	getformout

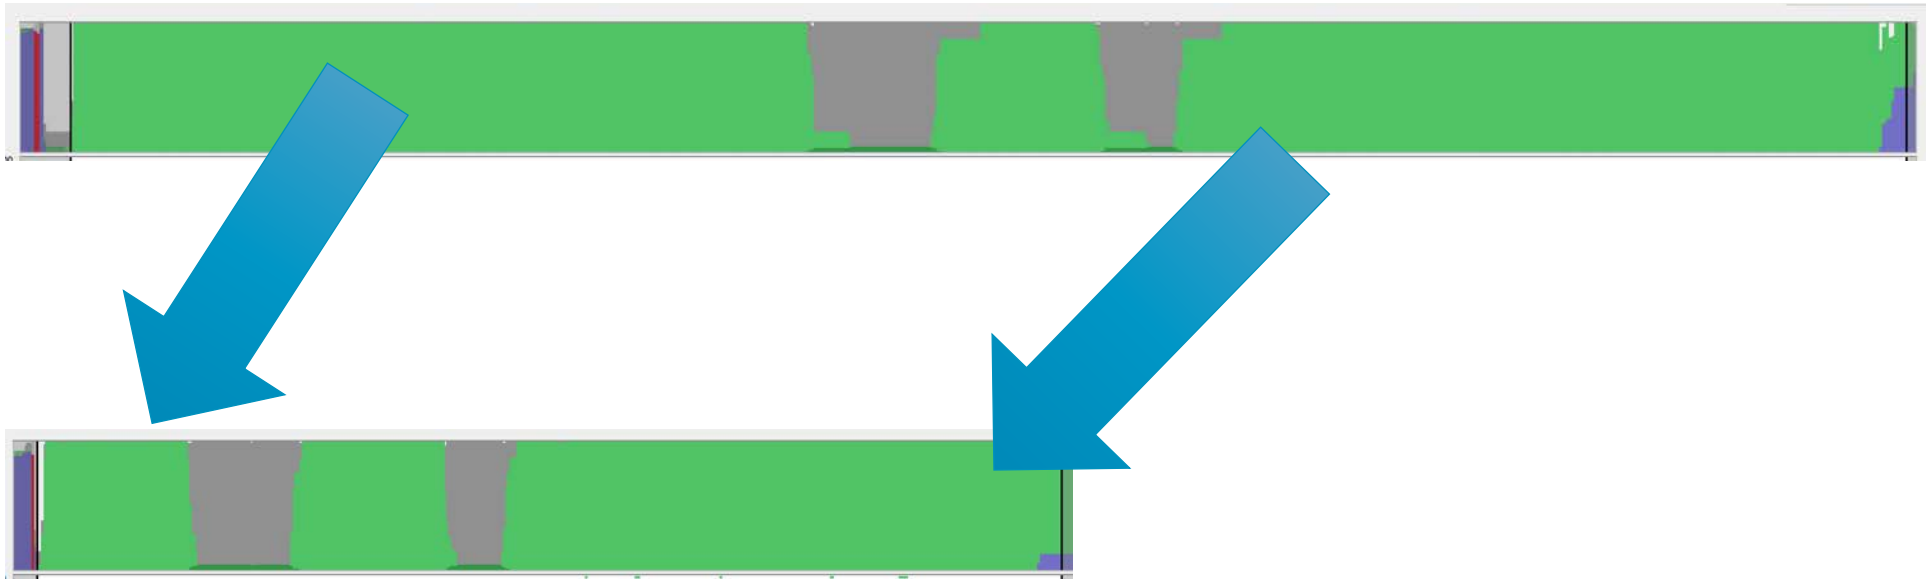

Self	Total	MPI	Child	Overhead	Function
7.3%	16.8%		9.6%		form
<0.1%	0.3%		0.3%		getchchar
5.2%	5.6%		0.4%		getchinfo
<0.1%	<0.1%				getchinfo
0.1%	1.4%		1.3%		getformout

arm

Balanced load balancer

Before:

0 8 0 -10 199 329 492 1.21 13530 0 89 -1 **91%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

1 8 0 -7 591 573 872 1.97 45150 0 350 0 **80%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

2 8 0 -16 894 762 1153 2.28 77028 0 607 1 **86%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

3 8 0 -24 916 886 1331 2.05 99681 0 766 2 **91%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

After:

0 8 0 -10 174 329 492 1.06 13530 0 85 -1 **93%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

1 8 0 -11 415 577 872 1.40 43956 0 340 0 **97%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

2 8 0 -11 616 757 1153 1.55 79003 0 592 1 **97%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

3 8 0 -12 667 874 1331 1.46 105111 0 734 2 **96%** LG,node,ipar,inc,vt,i1,i2,tperi,nch,naps,mt,prev LG,eff

arm

Wrap Up

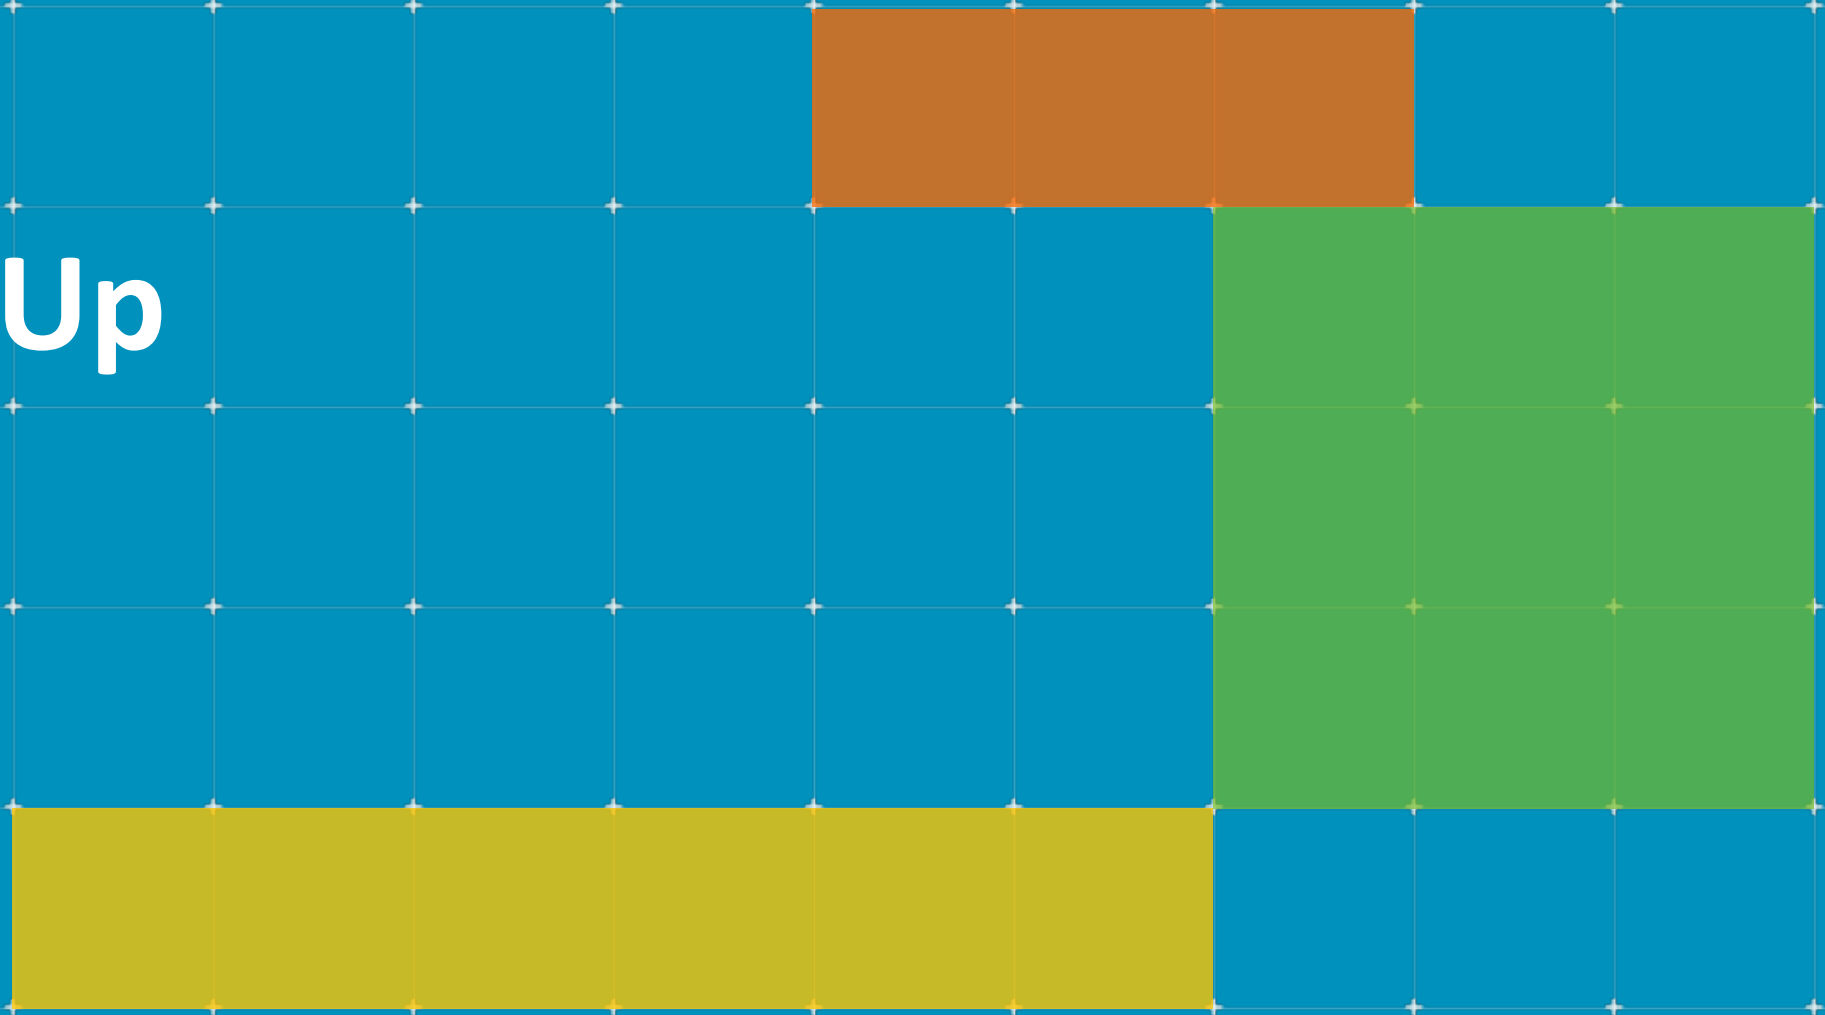

arm

Takeaways

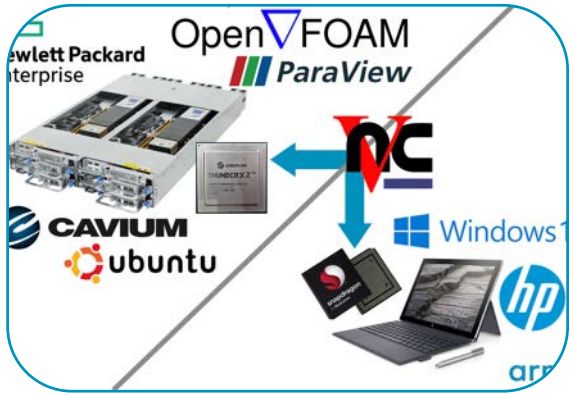

Applications ecosystem is mature and growing

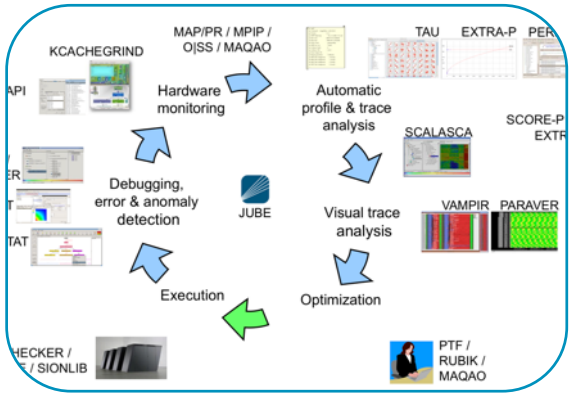

Tools ecosystem is mature and growing

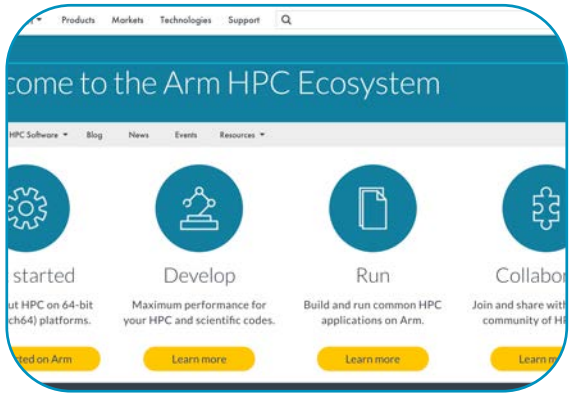

arm.com/hpc for porting and tuning training and resources

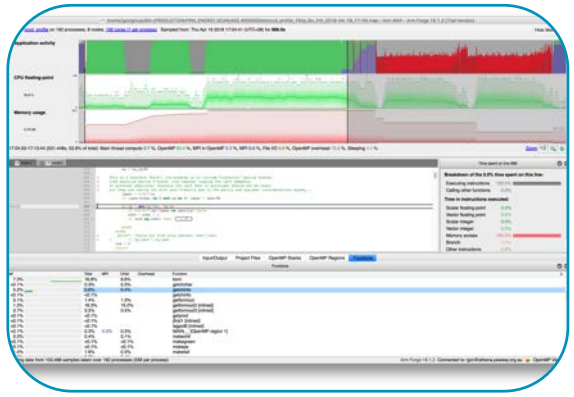

For cross-platform performance improvements, use Arm Forge

arm

Thank You!
Danke!
Merci!
谢谢!
ありがとう!
Gracias!
Kiitos!
감사합니다
धन्यवाद