SC'17 PANEL, "THE ARM SOFTWARE ECOSYSTEM: ARE WE THERE YET?"

- Moderator
  - CJ Newburn, NVIDIA
- Arm SW tools providers
  - Eric Van Hensbergen, Arm
  - Larry Kaplan, Cray
  - Shinji Sumimoto, Fujitsu
- Data center reps
  - Simon McIntosh-Smith, Bristol/Isambard
  - Matt McLean, U Michigan
  - Veronica Vergara, ORNL
- Material from other contributors from the audience
  - Trent D'Hooge, LLNL
  - Kevin Pedretti, Sandia
  - Jay Kruemcke, SuSE

LINE UP

- Welcome, overview – CJ Newburn, NVIDIA
  - Systems, package readiness, wiki
- Panelist presentations
  - Arm SW tools providers
    - Eric Van Hensbergen, Arm
    - Larry Kaplan, Cray
    - Shinji Sumimoto, Fujitsu
  - Data center reps
    - Simon McIntosh-Smith, Bristol/Isambard
    - Matt McLean, U Michigan
    - Veronica Vergara, ORNL
- Experience of the SW Ecosystem with current deployments - audience
- Questions from and discussion with the audience
- Wrap up

SYSTEMS - EARLY

Data Center Site:	BSC	BSC	BSC	Sandia	Sandia
System name:	Mont-Blanc prototype	Thunder	JetsonTX1	Hammer	Sullivan
Total nodes	~1000	5	15	48	37
Login nodes	5	1	1	1	1
Compute nodes	~1000	5	14	47	36
Compute nodes with GPUs	1 per node	0	14		
CPU type, # cores	Samsung Exynos Dual	ThunderX 1 48	JetsonTX1, 4	XGene-1	ThunderX 1
GPU type (if any)	MALI T604		NVIDIA Maxwell		
CPUs per node	2	96	4	1	2
GPUs per node	1		1		
DRAM memory per node	4GB	128GB	4GB	64GB	64GB
SSD per node	64GB	64GB	16GB		
Interconnect	1 GbE	40 GbE	1 GbE	10GbE	10GbE
OEM	Bull	E4	Self made	HPE	Penguin
Delivered	2015	2015	2015	2015	2015
OS(es)	Ubuntu 14.04	Ubuntu 14.04	Ubuntu 14.04	RHEL 7	RHEL 7
Testbed or prod	Test bed / Porting / Scalability study	Test bed / Porting	Test bed / Porting / Power monitoring	Test Bed / Porting	Test Bed / Porting

SYSTEMS - CURRENT/FUTURE

Data Center Site:	LLNL	ORNL	ORNL	Sandia	Sandia	Sandia	U Mich	U Mich	UK Bristol	UK Bristol
System name:		Envoy	Wombat	Mayer	Comanche-T	Comanche-X	Cavium-ThunderX (Hadoop)	Cavium-ThunderX (Hadoop)	Isambard early access	Isambard
Total nodes	68	16	16	33	2	4	41	120+	8	160
Login nodes	2		1	1	2	4	4	10-Jun	1	2
Compute nodes		16	11	32			37	110+	7	158
Compute nodes with GPUs	64	0	4						0	0
CPU type, # cores	ThunderX 2 28	ThunderX1 (pre-prod) 48	ThunderX2 (pre-prod) 28	ThunderX 2 28	ThunderX2 28	ThunderX2 28	ThunderX 48	ThunderX	ThunderX 2 28	ThunderX 2 32
GPU type (if any)	AMD		AMD							
CPUs per node	2	2	2	2	2	2	2	2	2	2
GPUs per node	2		2							
DRAM memory per node	256 GB	128GB	256GB	256GB	256GB	256GB	512GB	512GB	256GB	256GB
SSD per node	256GB		480GB				6	6		
Interconnect	EDR IB	10GbE	EDR IB	IB	10GbE	10GbE	40GbE	40GbE	10GbE	Ares
OEM	HPE	Cray	HPE	HPE	HPE	NDA Info	Gigabyte	Gigabyte	Cray	Cray
Delivered	2017	2017	2017	2017	2017	2017	2017	2019	2017	2018
OS(es)	TOSS 3 / RHEL 7 based	Cray, based on SLES 12	RHEL 7	RHEL 7	RHEL 7	SLES	RHEL 7	RHEL 7	SLES	SLES
Testbed or prod	targeting prod	Testbed	Testbed	Test Bed / Porting	Test Bed / Porting	Test Bed / Porting	Development Platform	prod!	Test Bed / Porting	prod

PACKAGE READINESS

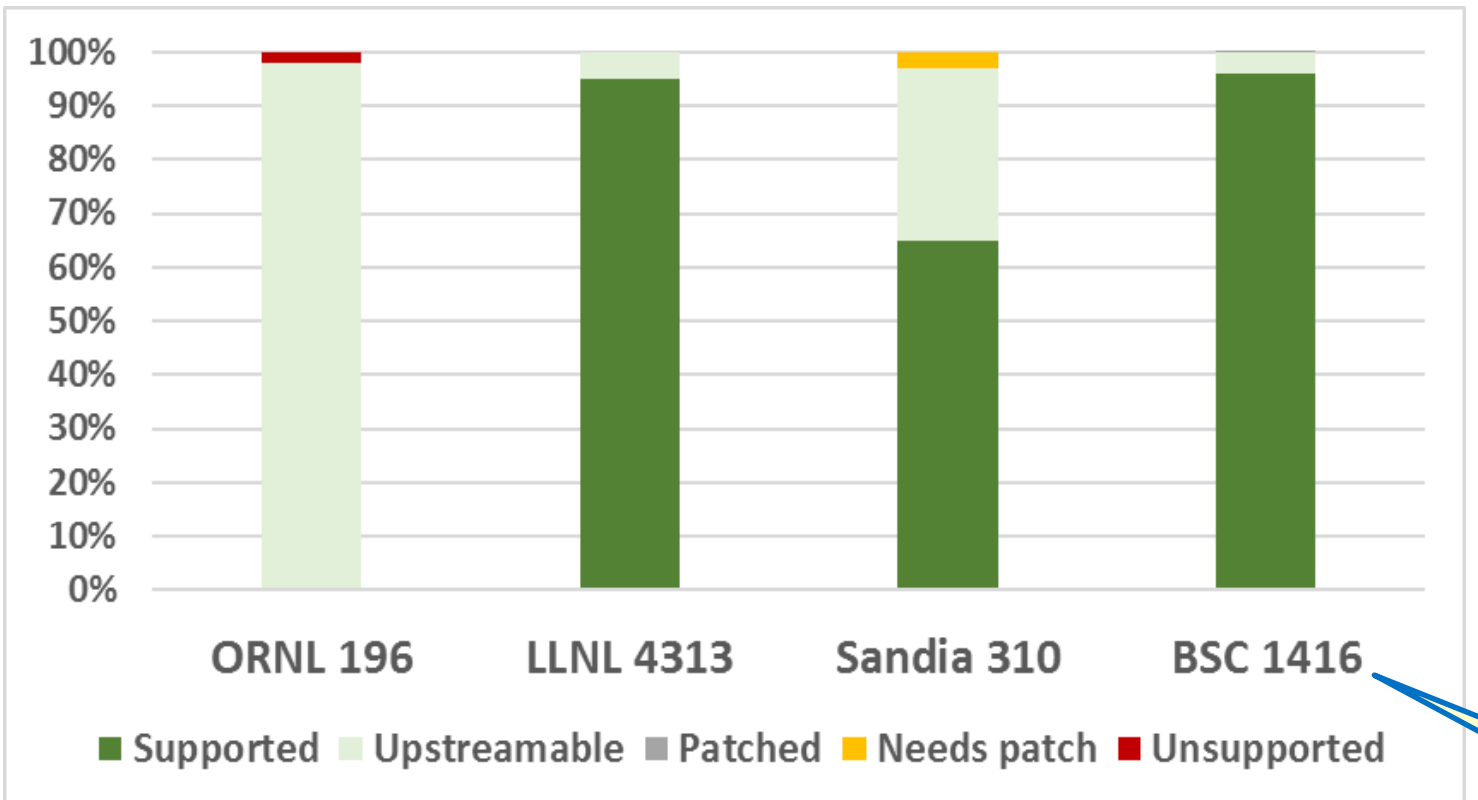

Supported - commercial package officially supports 64-bit ARMv8, e.g. from RHEL, EPEL

Upstreamable - open-source project builds from upstreamed source on Arm with no modifications (apart from tweaking compiler options for performance);

Patched - detailed modification steps for the package are available

Needs Patch - patches are required, but are not documented in detail

Unsupported - haven't made it work yet

pkgs of interest

BUGS

- Compilers @ Sandia
  - Open 5
  - Resolved 2
- In general issues are fairly minor and are expected to be resolved quickly.
- Vendors are engaging and resolving issues quickly.

GitLab Projects Groups Snippets Help Sign in / Register

Arm HPC Users Gr... packages Wiki Home

Home

Last edited by **arm-hpc packages pipeline** about 14 hours ago

Page history

This Wiki exists to capture and link to information regarding the packages considered critical for HPC.

Download the summary Excel Spreadsheet

Please make any modifications you like to the individual package pages. Especially desirable contributions are:

- Elaborating on details of what the package is, where they are sparse;
- Mentioning yourself if you are actively working on it or have some

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

ARM® at Cray

Larry Kaplan
Chief Software Architect

SC17 Denver, CO hpc connects

ARM is Coming to the Cray XC50

Thunder X2 Processor Daughter Card

2-Socket Nodes

8 DIMMs per Socket

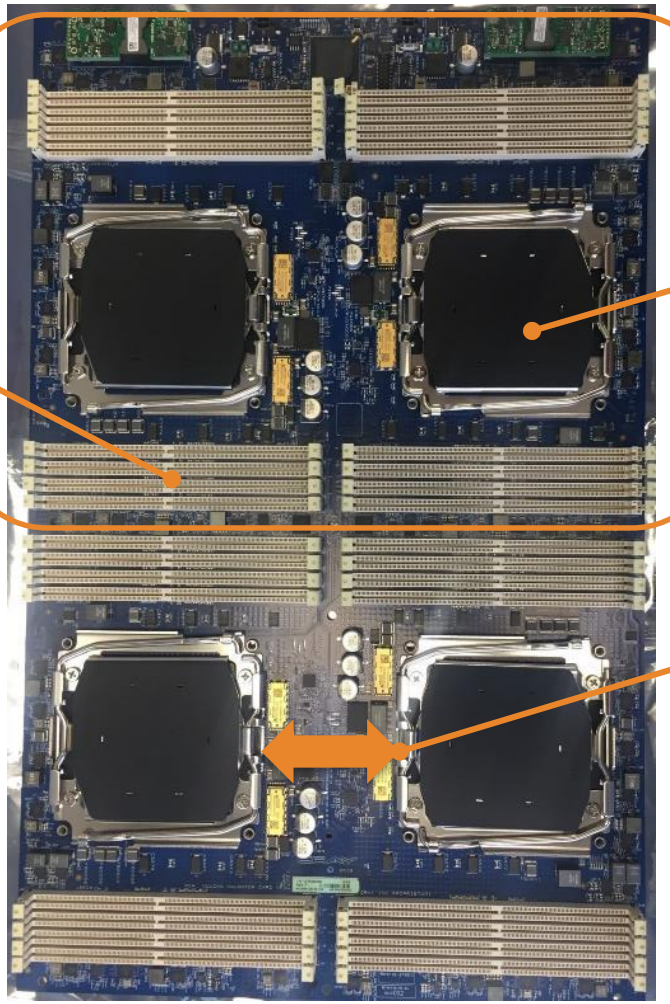

High memory and inter-chip link bandwidth permit a flat (uniform) 300 GB/s local memory system!

CAVIUM

ThunderX2 Processor
32 Cores
2.1 GHz base frequency + boost
537 Gflops
32 MB L3 Cache

Inter-chip Interconnect
75 GB/s/dir
24 lanes @
25 Gbps

Cray Collaboration Evident in ARM SVE

Feature	Benefit
Scalable vector length (VL)	Increased parallelism while allowing implementation choice of VL
VL agnostic (VLA) programming	Supports a programming paradigm of write-once, run-anywhere scalable vector code
Gather-load & Scatter-store	Enables vectorization of complex data structures with non-linear access patterns
Per-lane predication	Enables vectorization of complex, nested control code containing side effects and avoidance of loop heads and tails (particularly for VLA)
Predicate-driven loop control and management	Reduces vectorization overhead relative to scalar code
Vector partitioning and SW managed speculation	Permits vectorization of uncounted loops with data-dependent exits
Extended integer and floating-point horizontal reductions	Allows vectorization of more types of reducible loop-carried dependencies
Scalarized intra-vector sub-loops	Supports vectorization of loops containing complex loop-carried dependencies

Nigel Stephens – https://community.arm.com/processors/b/blog/posts/technology-update-the-scalable-vector-extension-sve-for-the-armv8-a-architecture

Programming Environment for Cray Systems

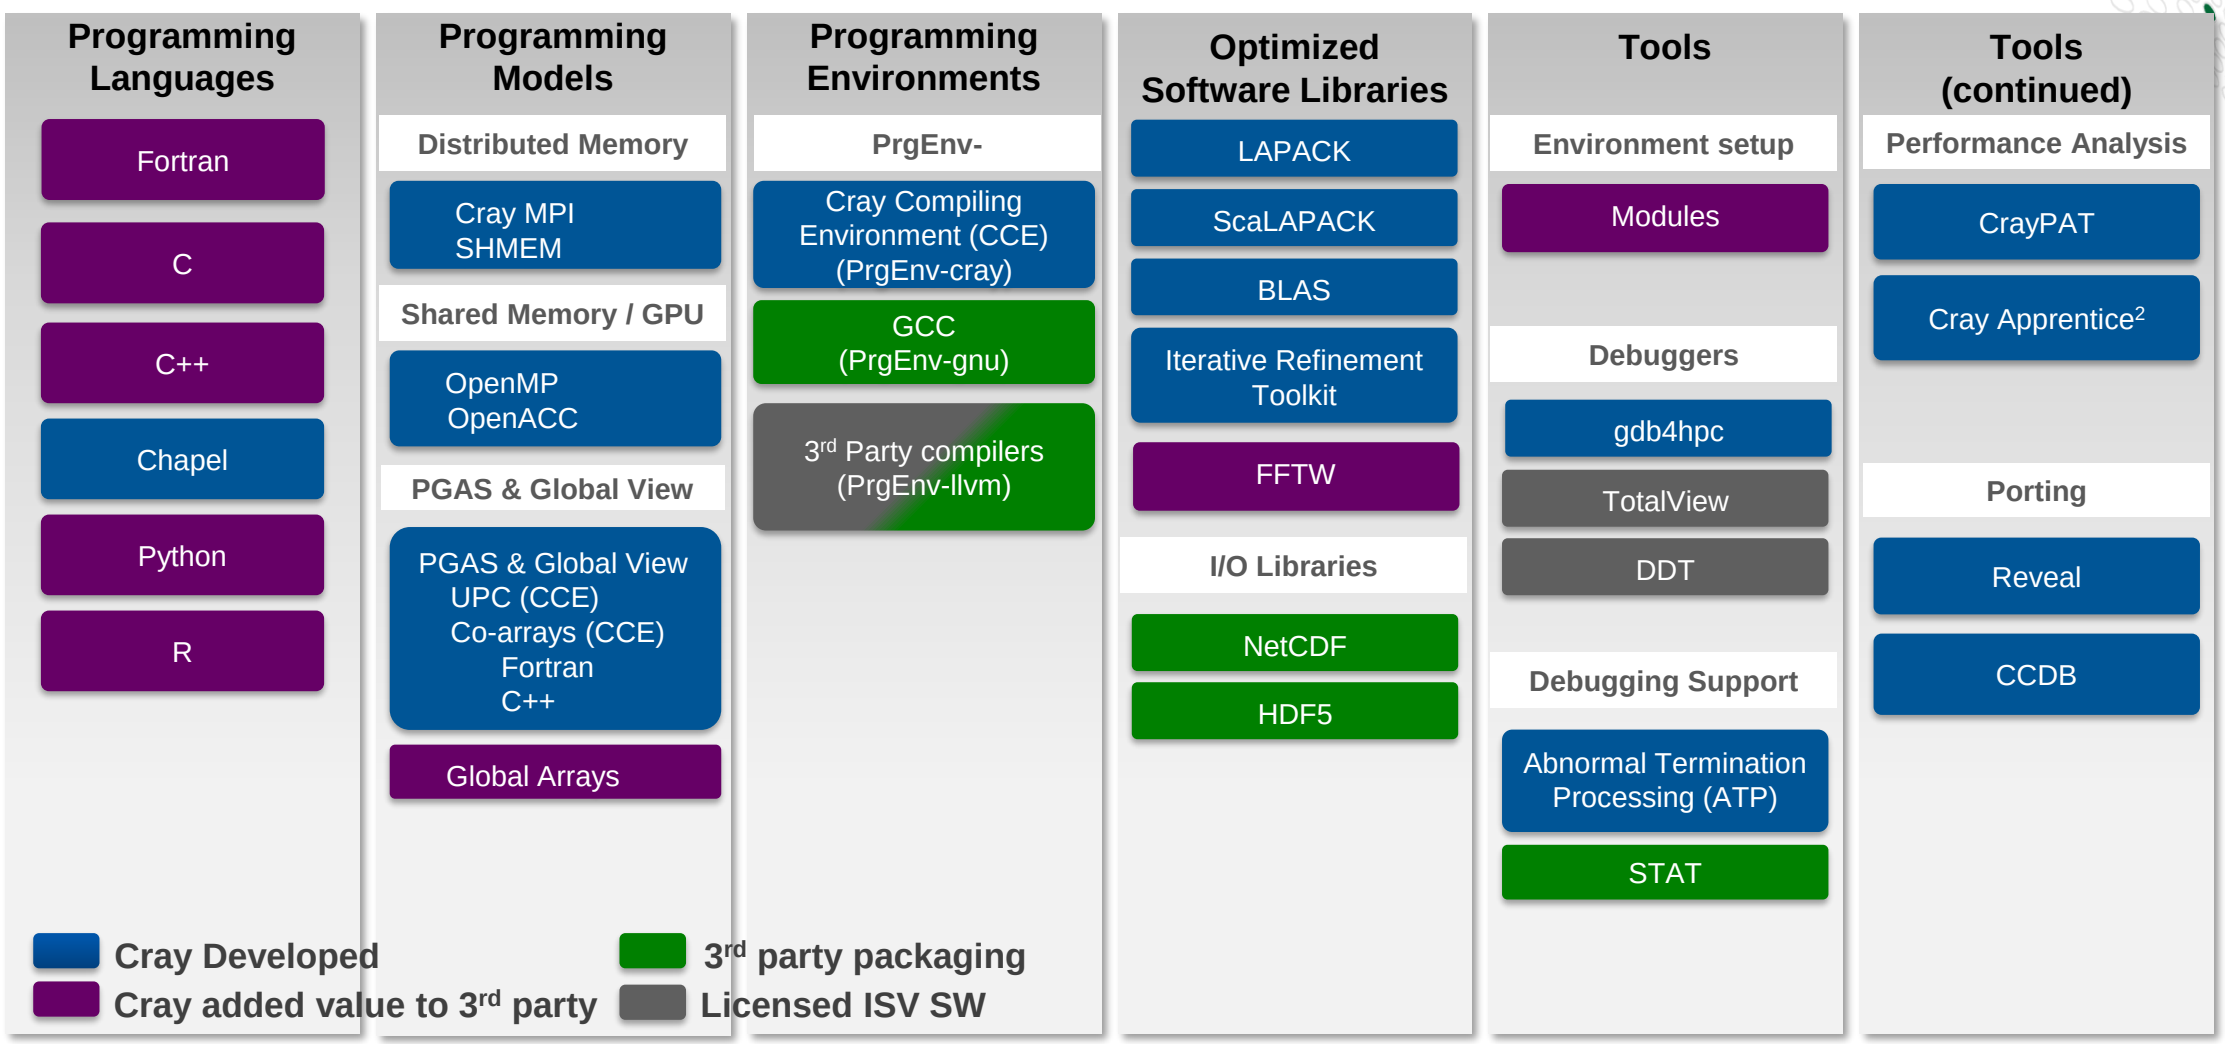

CCE vs. Alternative ARM Compilers on 135 HPC Benchmarks

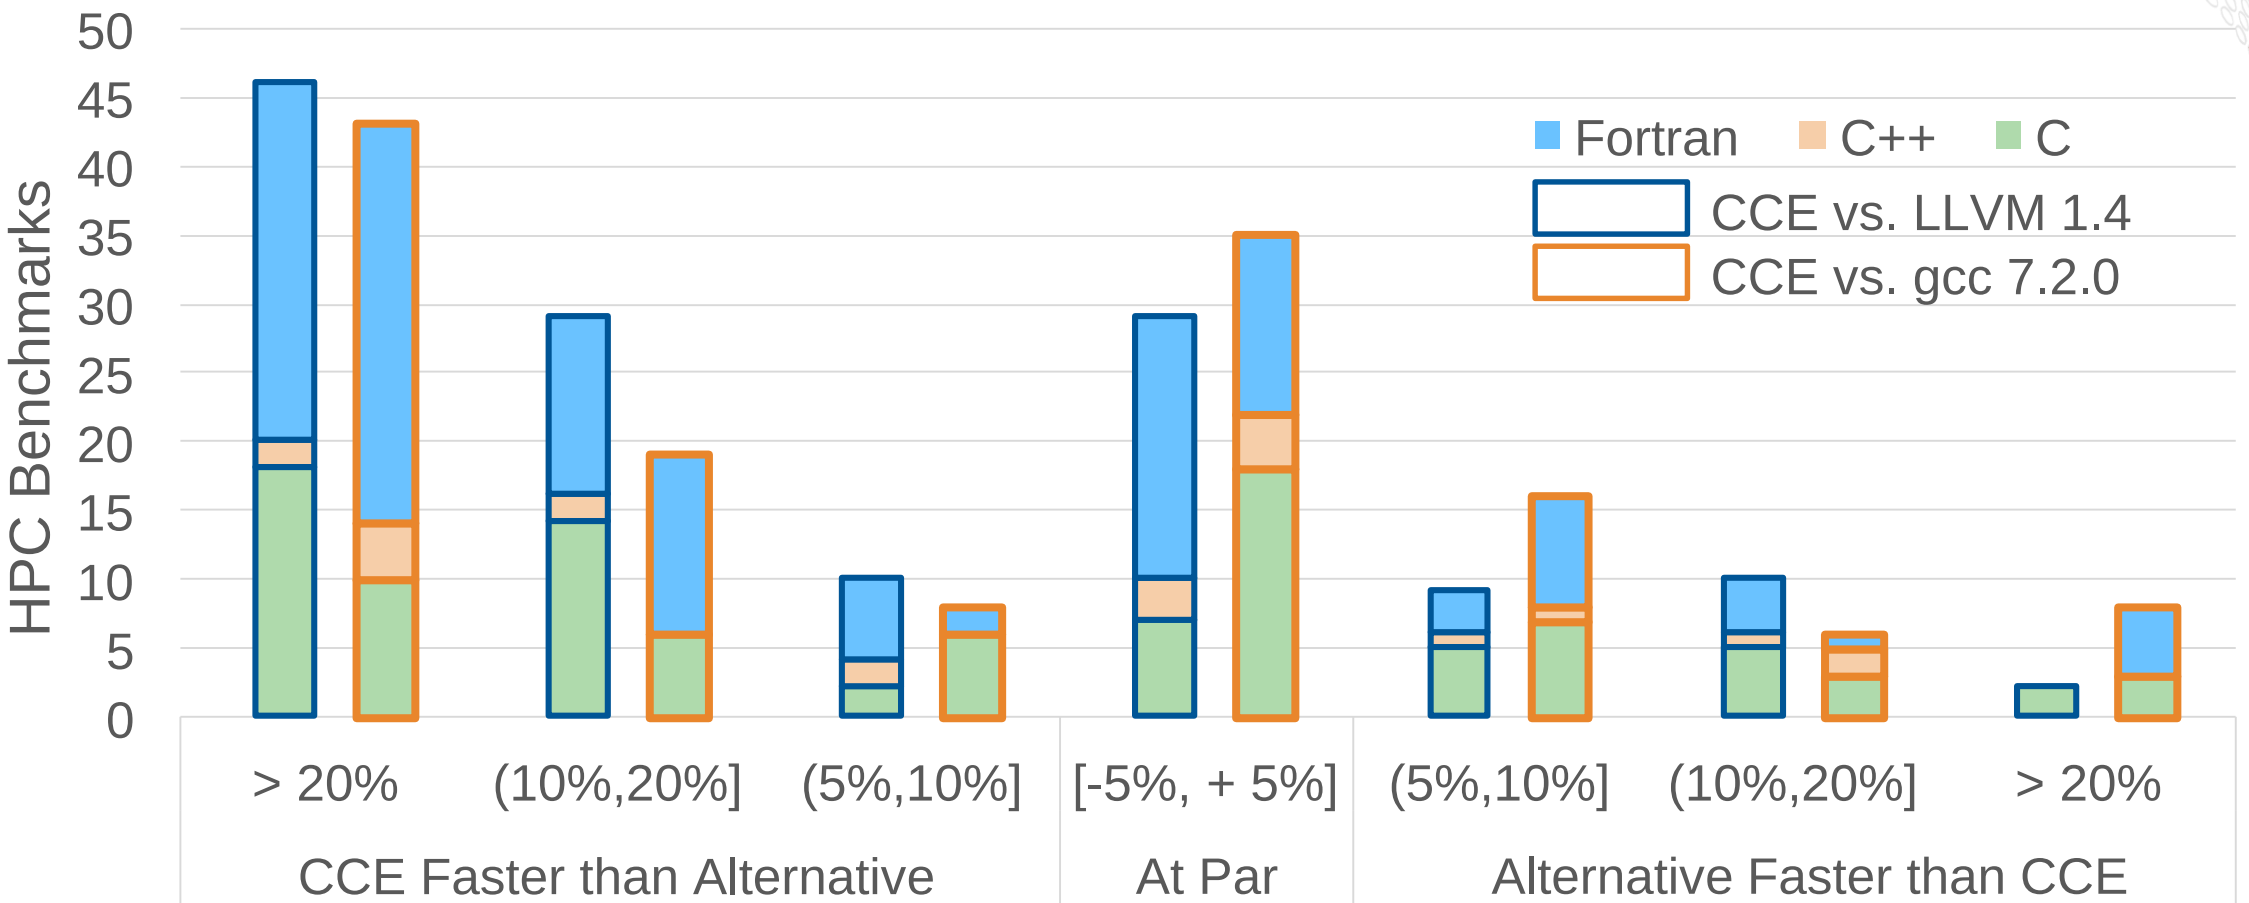

Supercomputing 2017

ARM Ecosystem Readiness

Eric Van Hensbergen
<eric.vanhensbergen@arm.com>

Supercomputing 2017

Are we there yet?
Arm Ecosystem Panel

Arm Allinea Studio

Built for developers to achieve best performance on Arm with minimal effort

arm
ALLINEA STUDIO

- Arm Compiler
- Arm Performance Libraries
- Arm Forge
- Arm Performance Reports

Comprehensive and integrated tool suite for Scientific computing, HPC and Enterprise developers

Seamless end-to-end workflow from getting started to advanced optimization of your workloads

Commercially supported by Arm engineers

Frequent releases with continuous performance improvements

Ready for current and future generations of server-class Arm-based platforms

Available for a wide-variety of Arm-based server-class platforms

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 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
- Batch 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

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

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

Migrating High Performance to 64-bit Arm

An end-to-end workflow to get the most of your modern Arm-based platform

Understand application behaviour now

Set a reference for future work

- Choose a representative test cases with known results
- Analyse performance on existing hardware (e.g. x86)
  with **Arm Performance Reports**
- Test scaling and note compiler flags

Example:

```
$> perf-report mpirun -n 16 mmult.exe
```

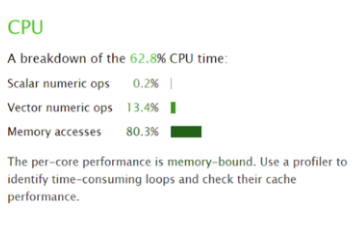

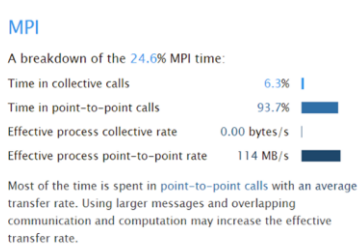

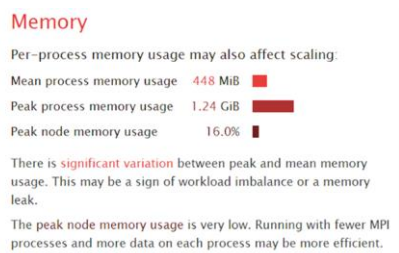

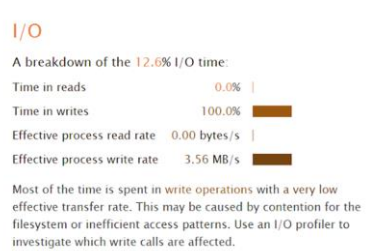

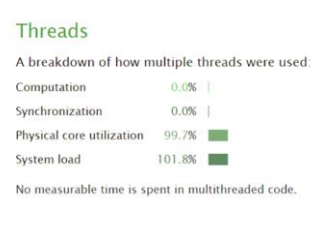

arm

Compile and link your application on Arm

Application porting should be effortless

- Modify the Makefile/installation scripts to ensure compilation for aarch64 happens
- Compile the code with the **Arm C/C++/Fortran Compiler**
- Link the code with the **Arm Performance Libraries**

Examples:

```
$> armclang -c –I/path/armpl/include example.c -o example.o
$> armclang example.o -L/path/armpl/lib -larmpl_lp64 -o example.exe -lflang -lflangrti -lm
```

Note: Most codes are compiled with “mpicc”. The MPI library needs to be compiled first!

arm

Run and ensure application correctness

Combination of debugging and re-compilation

- Run with the representative workload you started with
- Verify application output, tweak compilation optimisations with **Arm C/C++/Fortran Compiler**
- Ensure application correctness with **Arm Forge Professional**

Examples:

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

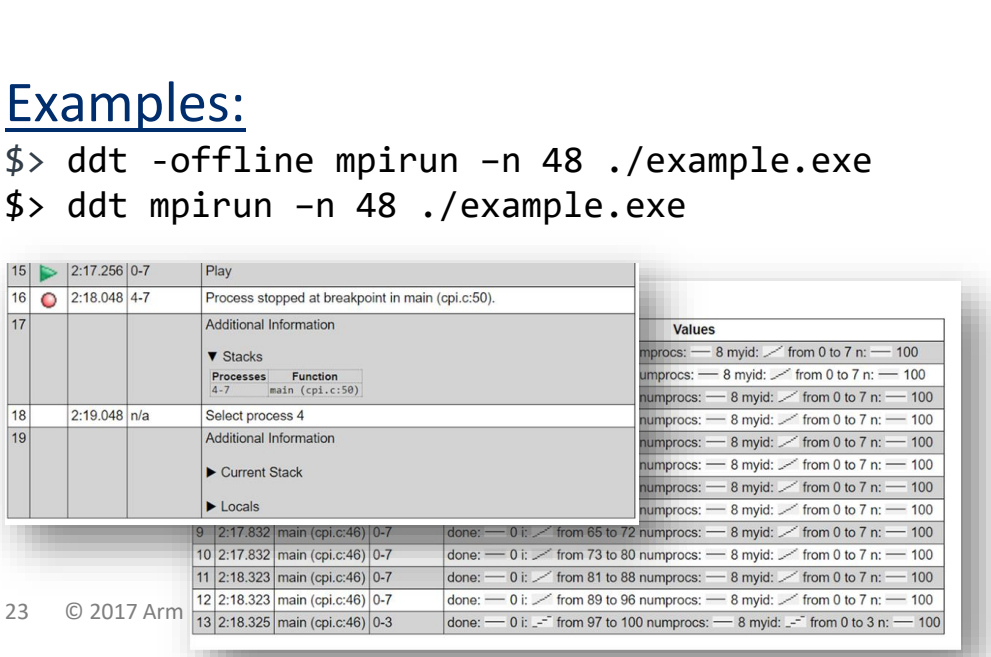

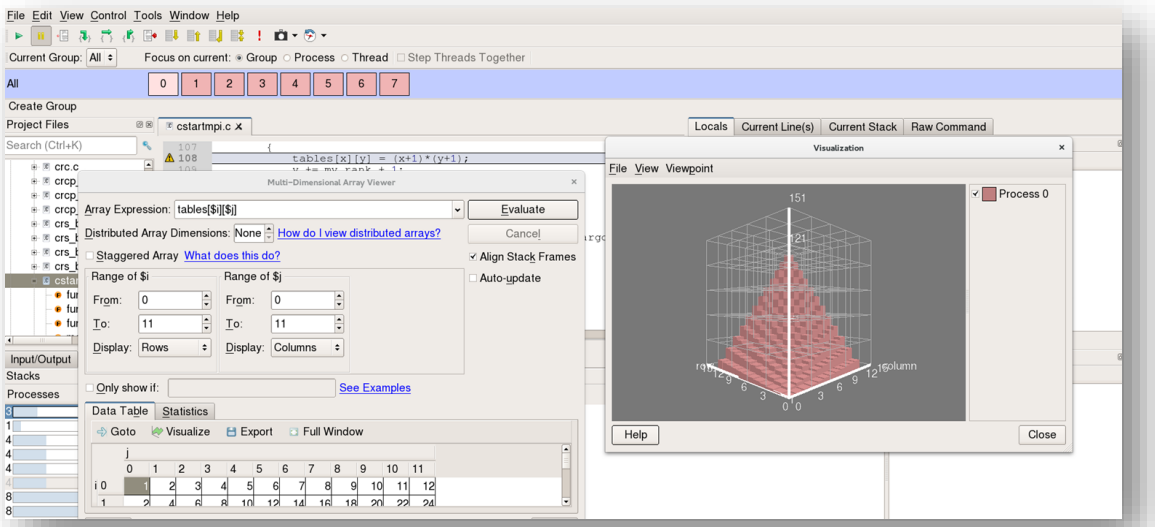

Optimise the application for aarch64

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.exe
```

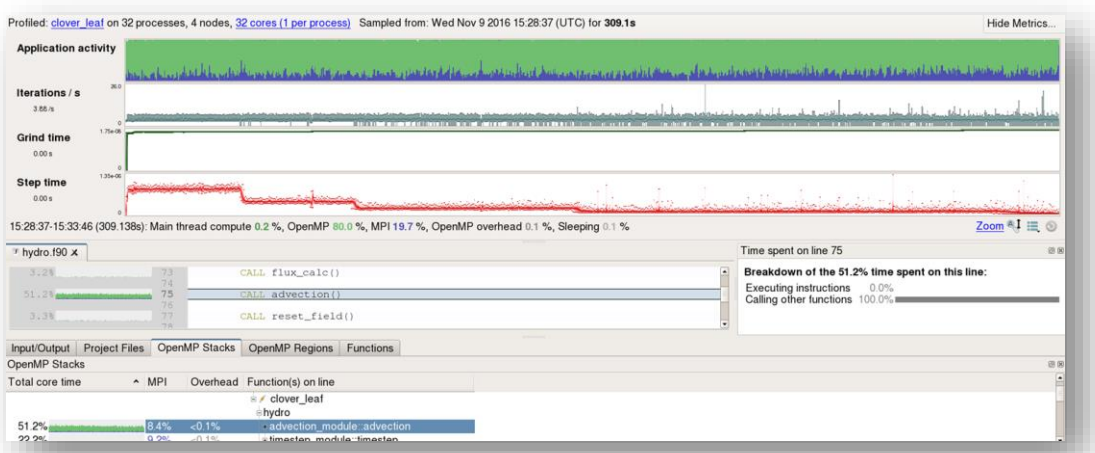

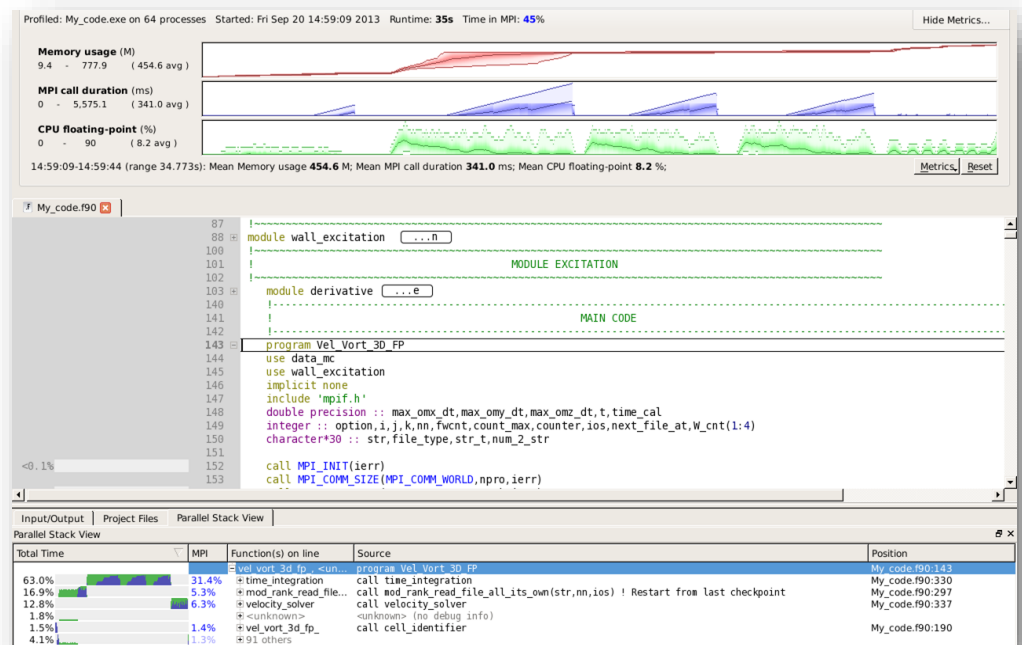

Demonstrate performance gains on Arm

A before/after comparison is a double edge sword… but also our most powerful ally

Before

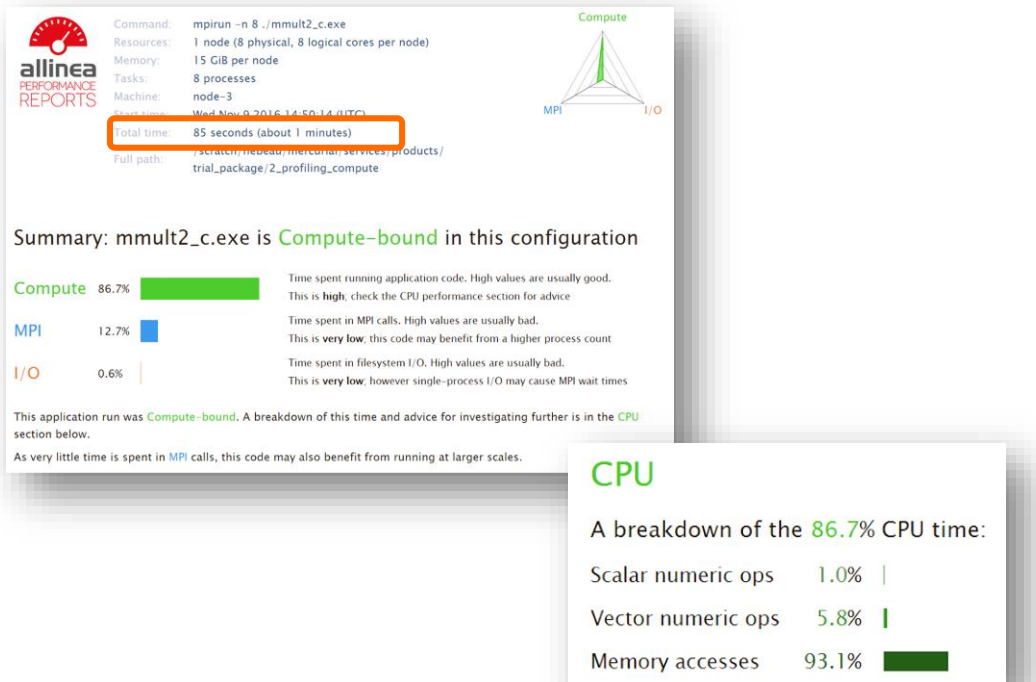

After

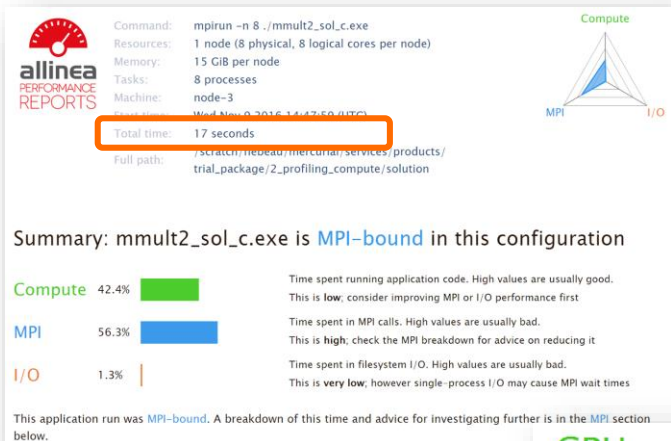

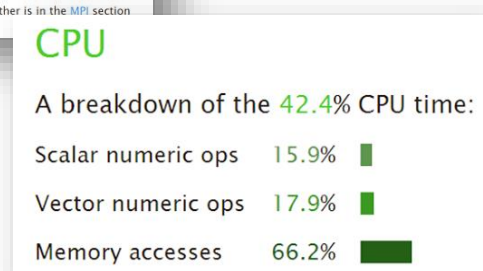

arm

Scalable Vector Extension (SVE)

There is no preferred vector length

- Vector Length (VL) is hardware choice, from 128 to 2048 bits, in increments of 128
- *Vector Length Agnostic* (VLA) programming adjusts dynamically to the available VL
- No need to recompile, or to rewrite hand-coded SVE assembler or C intrinsics

SVE is not an extension of Advanced SIMD

- A separate architectural extension with a new set of A64 instruction encodings
- Focus is HPC scientific workloads, not media/image processing

Amdahl says you need high vector utilisation to achieve significant speedups

- Compilers often unable to vectorize due to intra-vector data & control dependencies
- SVE also begins to address some of the traditional barriers to auto-vectorization

Evaluating SVE

Compile

Arm Compiler

C/C++/Fortran code

SVE via auto-vectorization, intrinsics and assembly.

Compiler Insight: Compiler places results of compile-time decisions and analysis in the resulting binary.

Supplied with SVE Performance Libraries.

Emulate

Arm Instruction Emulator

Runs userspace binaries for future Arm architectures on today's systems.

Supported instructions run unmodified.

Unsupported instructions are trapped and emulated.

Analyse

Arm Code Advisor

Console or web-based output shows prioritized advice in-line with original source code.

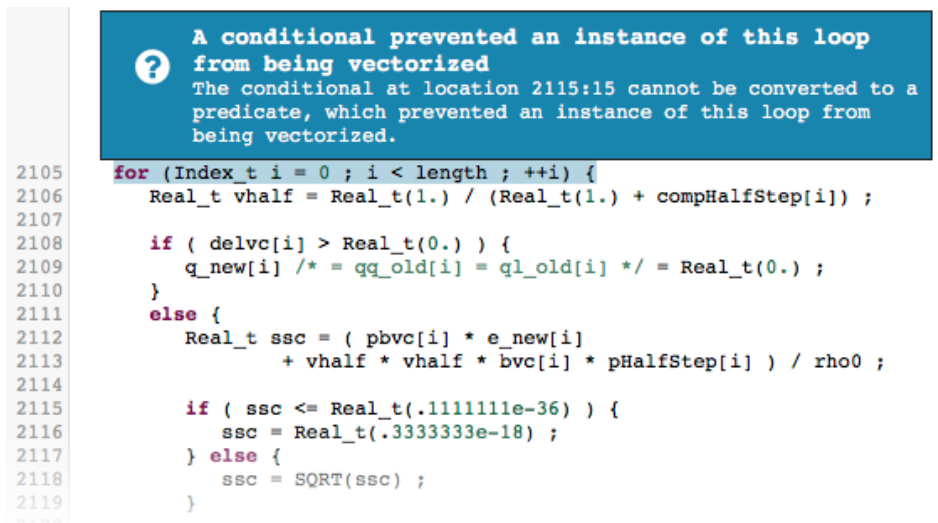

Arm Instruction Emulator (Beta)

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

- 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

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
```

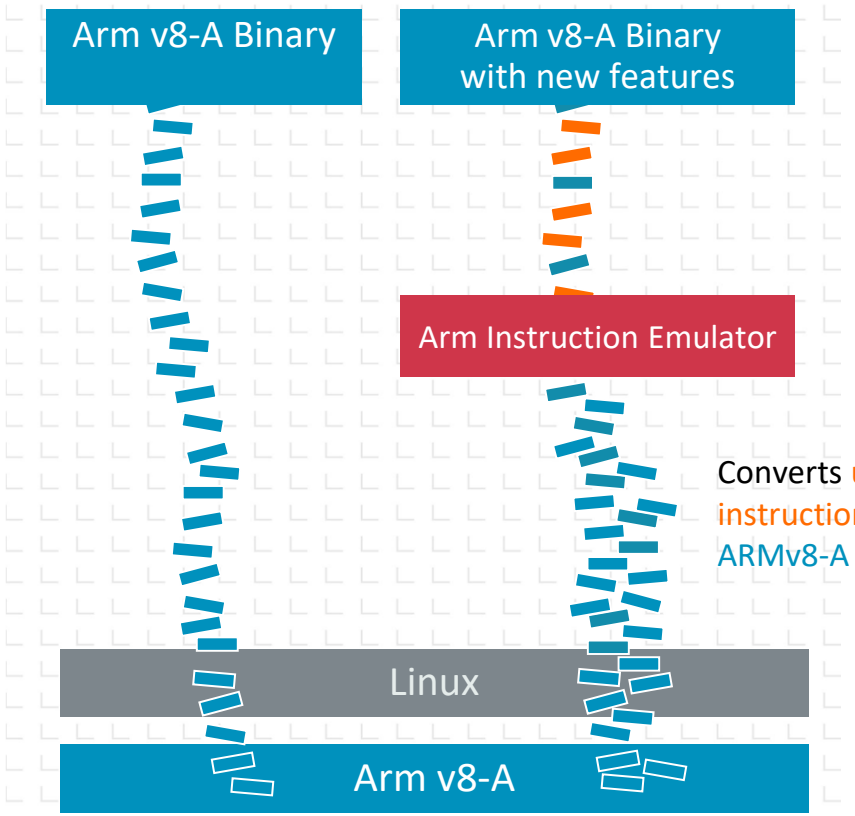

arm

Arm Code Advisor (Beta)

Actionable insights based on static and run time information

Performance Advice

- Compiler vectorization hints.
- Compilation flags advice.

Insights from compilation and runtime

- Compiler Insights are embedded into the application binary by the ARM Compilers.

Extensible Architecture

- Users can write plugins to add their own analysis information.
- Data accessible via web-browser, command-line, and REST API to support new user interfaces.

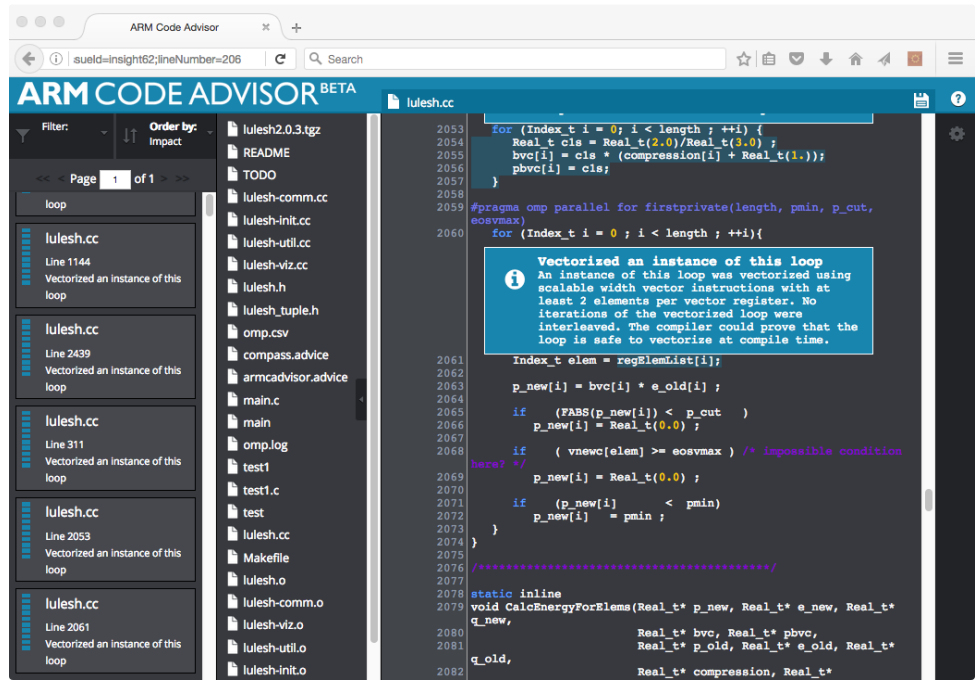

– Easy HPC stack deployment on Arm

OpenHPC is a community effort to provide a common, verified set of open source packages for HPC deployments

Arm's participation:

- Silver member of OpenHPC

Status: 1.3.3 release out now

- Packages built on Armv8-A for CentOS and SLES
- Arm-based machines in the OpenHPC build infrastructure

Functional Areas	Components include
Base OS	CentOS 7.3, SLES 12 SP2
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagios, pdsh, pdsh-mod-slurm, prun, EasyBuild, ClusterShell, mrsh, Genders, Shine, test-suite
Provisioning	Warewulf
Resource Mgmt.	SLURM, Munge
I/O Services	Lustre client (community version)
Numerical/Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, SuperLU_Dist,Mumps, OpenBLAS, Scalapack, SLEPc, PLASMA, ptScotch
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran), LLVM
MPI Families	OpenMPI, MPICH
Development Tools	Autotools (autoconf, automake, libtool), Cmake, Valgrind,R, SciPy/NumPy, hwloc
Performance Tools	PAPI, IMB, pdtoolkit, TAU, Scalasca, Score-P, SIONLib

arm

gitlab.com/arm-hpc

Community site with useful resources on HPC packages on Arm

Status of various HPC software packages on Arm

Packages in the 'application' category

Package	Last Modified	BuildMaturity	CompilesARMCompiler	CompilesGCC	NEONOptimized
openfoam	2017-08-02	NeedsPatch	Yes	Yes	-
openfoamplus	2017-08-02	NeedsPatch	Yes	Yes	-
picard	2017-07-10	-	-	-	-
quantum-espresso	2017-10-19	NeedsPatch	Yes	Yes	-

Recipes to build packages with GCC and Arm Compiler

Build instructions

Downloading and unpack the packages

```
wget http://www.qe-forge.org/gf/download/frsrelease/240/1075/qe-6.1.tar.gz
wget http://www.qe-forge.org/gf/download/frsrelease/240/1073/qe-6.1-test-suite.tar.gz

# Unpack tar file of src
tar zxf qe-6.1.tar.gz
cd qe-6.1
```

Compiler configuration

```
F77=armflang
```


arm

20 YEAR MY HPC EXPERIENCE

PC cluster history shows…..

- 1997: I joined RWCP project to develop open source PC cluster system software called SCore, and developed communication libraries called PM/PMv2 and SCore Clusters.
- 1990's: Proprietary World in Super Computing.
  - No open source HPC ecosystem
  - 1994: PC cluster was developed, called Beowulf clusters, as a poor-man's super computer.
- 2000's: Commercial Use started in Japan
  - Open standard based systems became important
- Now: Intel based PC cluster is dominant in HPC.

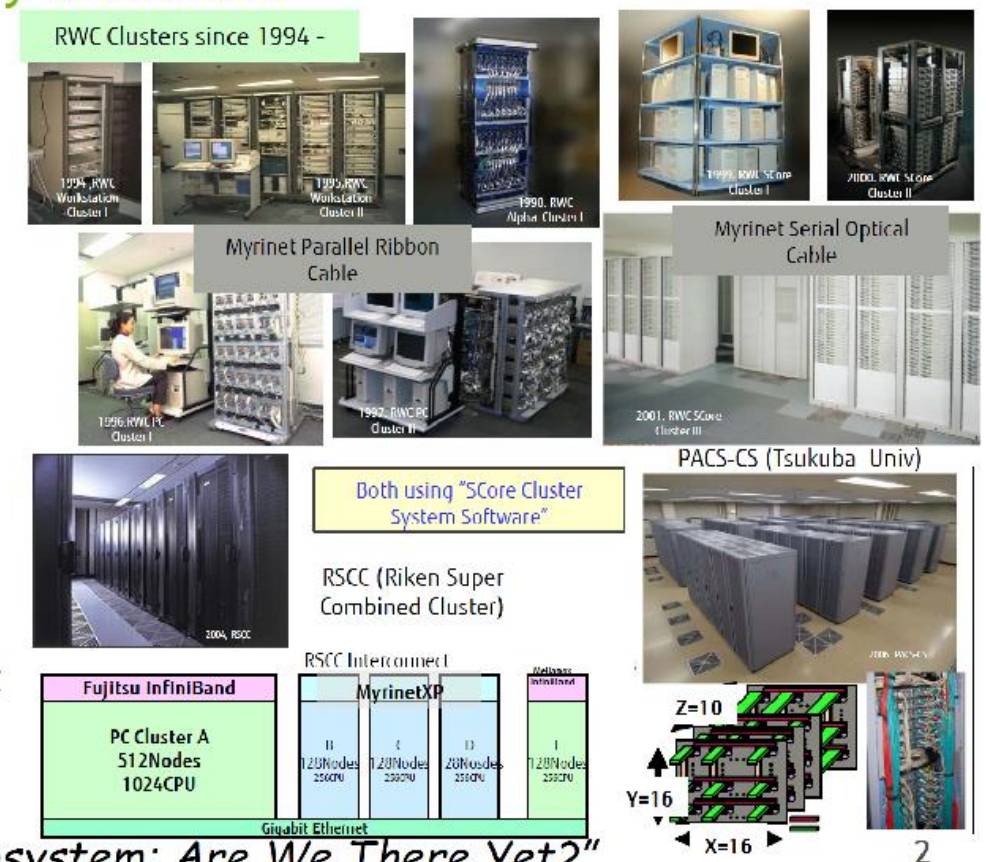

SC'17 Panel, "*The Arm Software Ecosystem: Are We There Yet?*"

INSIGHTS

Should be More Open Standard Based System

- Key learnings
  - More open standard system will be widely used
  - Software ecosystem is a critical issue every time
- Fujitsu decided to use Arm ISA for the next supercomputer to realize more open standard system.
  - Becoming SVE Lead Partner
  - RIKEN and Fujitsu are developing Arm based super computer called Post-K

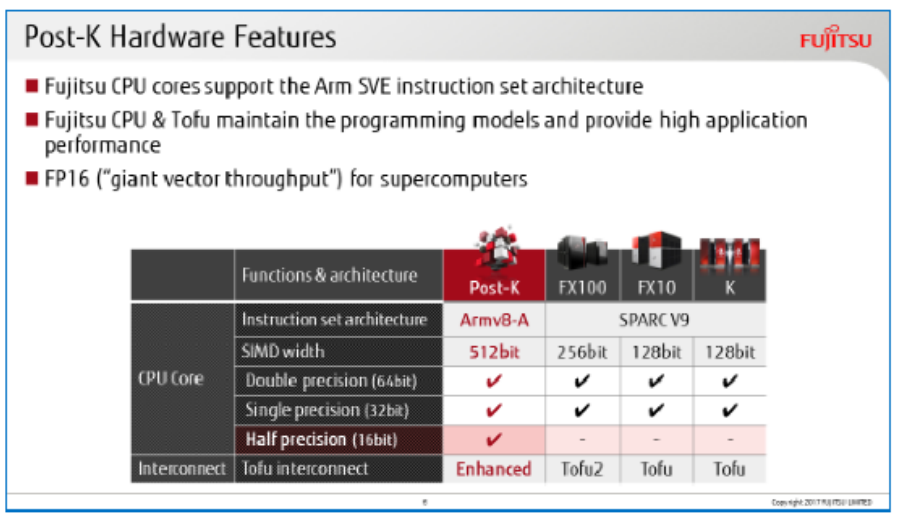

SC'17 Panel, *"The Arm Software Ecosystem: Are We There Yet?"*

Post-K Software Stack

FUJITSU

- Valuable feedbacks through "co-design" from application R&D teams

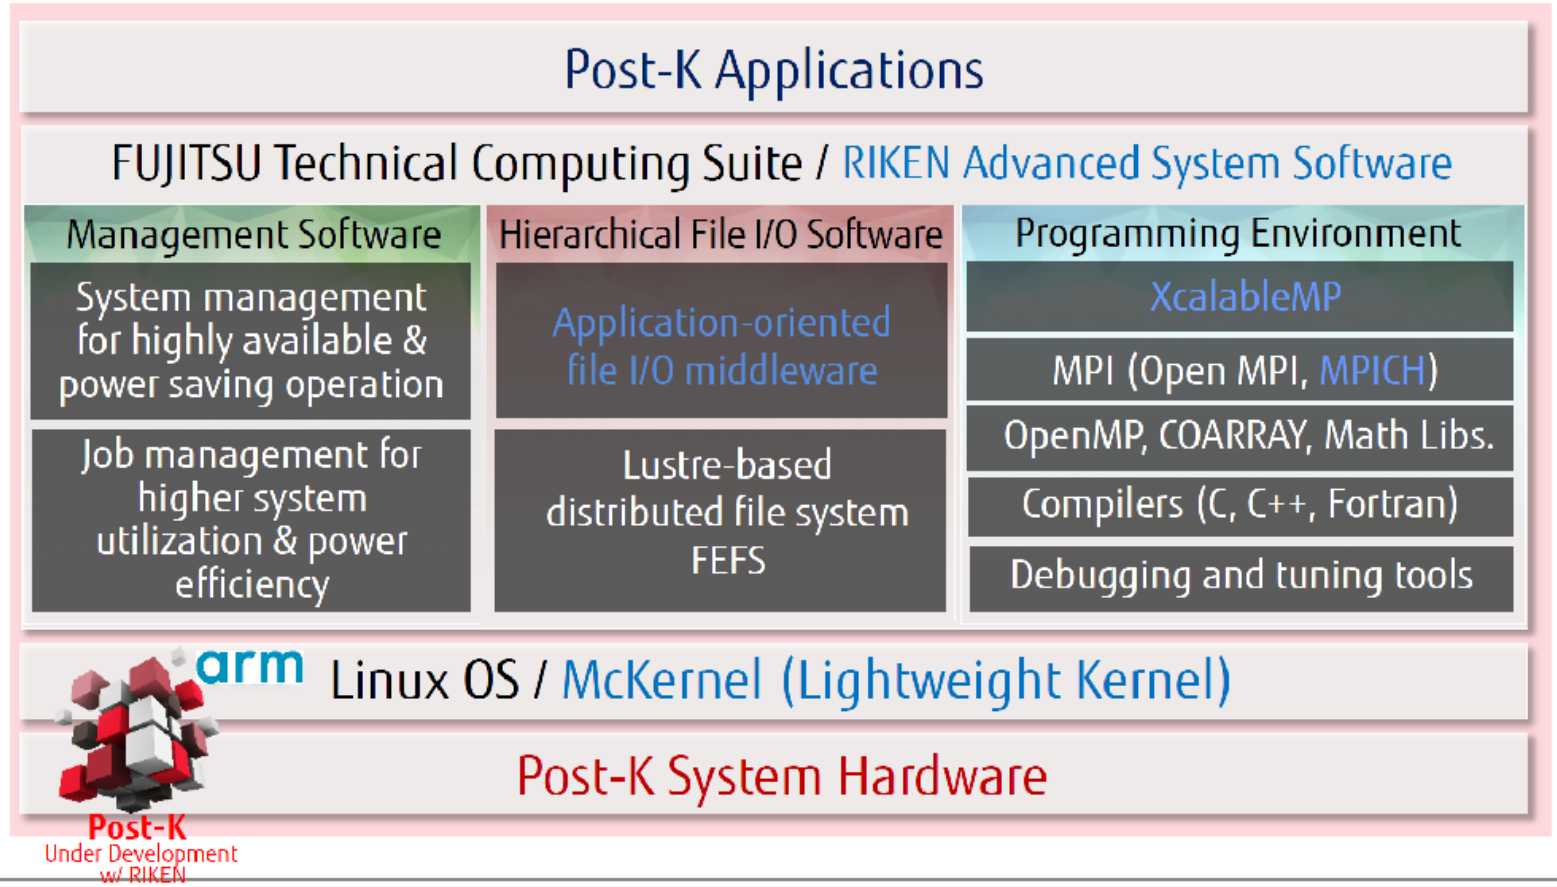

SVE Enabled FUJITSU Compiler for Arm

FUJITSU

- Targets 512bit-wide vectorization as well as Vector-length-agnostic
  - Fixed-vector-length facilitates optimizations such as constant folding
- Language Standard Support
  - Fully supported : Fortran 2008, C11, C++14, OpenMP 4.5
  - Partially supported : Fortran 2015, C++1z, OpenMP 5.0
- Inherits options/features of K computer, PRIMEHPC FX10 and FX100
- Supports Arm C Language Extensions (ACLE) for SVE
  - ACLE allow programmers to use SVE instructions as C intrinsic functions

```
// C intrinsics in ACLE for SVE
svfloat64_t z0 = svld1_f64(p0, &x[i]);
svfloat64_t z1 = svld1_f64(p0, &y[i]);
svfloat64_t z2 = svadd_f64_x(p0, z0, z1);
svst1_f64(p0, &z[i], z2);
```

```
// SVE assembler
ld1d    z1.d, p0/z, [x19, x3, lsl #3]
ld1d    z0.d, p0/z, [x20, x3, lsl #3]
fadd    z1.d, p0/m, z1.d, z0.d
st1d    z1.d, p0, [x21, x3, lsl #3]
```

FUJITSU

What is going on to build up Arm HPC ecosystem?

- Fujitsu's Perspective and Activities

Strategy for building up Arm HPC Ecosystem

FUJITSU

- Background: Not 20 years ago
  - Huge HPC ecosystem has been developed on IA(Intel) based system.
  - Therefore, easy migration of ecosystem from the IA's is very important.
- Two Strategies:
  - Keeping binary compatibility among Arm based systems including Linux distribution image.
  - Keeping source level portability between IA and Arm
- Actions with related community:
  - Keeping Arm Binary Level Compatibility: Arm and SIG HPC in Linaro(SBSA,SBBR)
  - Building and Keeping Source Level Portability: OpenHPC

arm

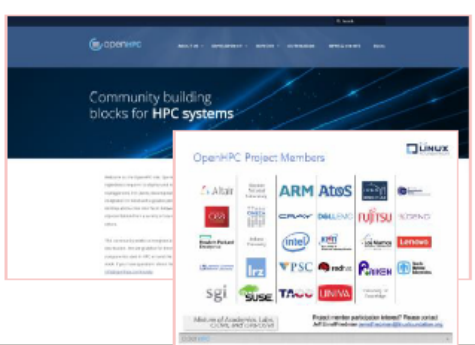

Fujitsu's Arm HPC Development Plan

FUJITSU

- Arm HPC software development plan with Arm HPC community

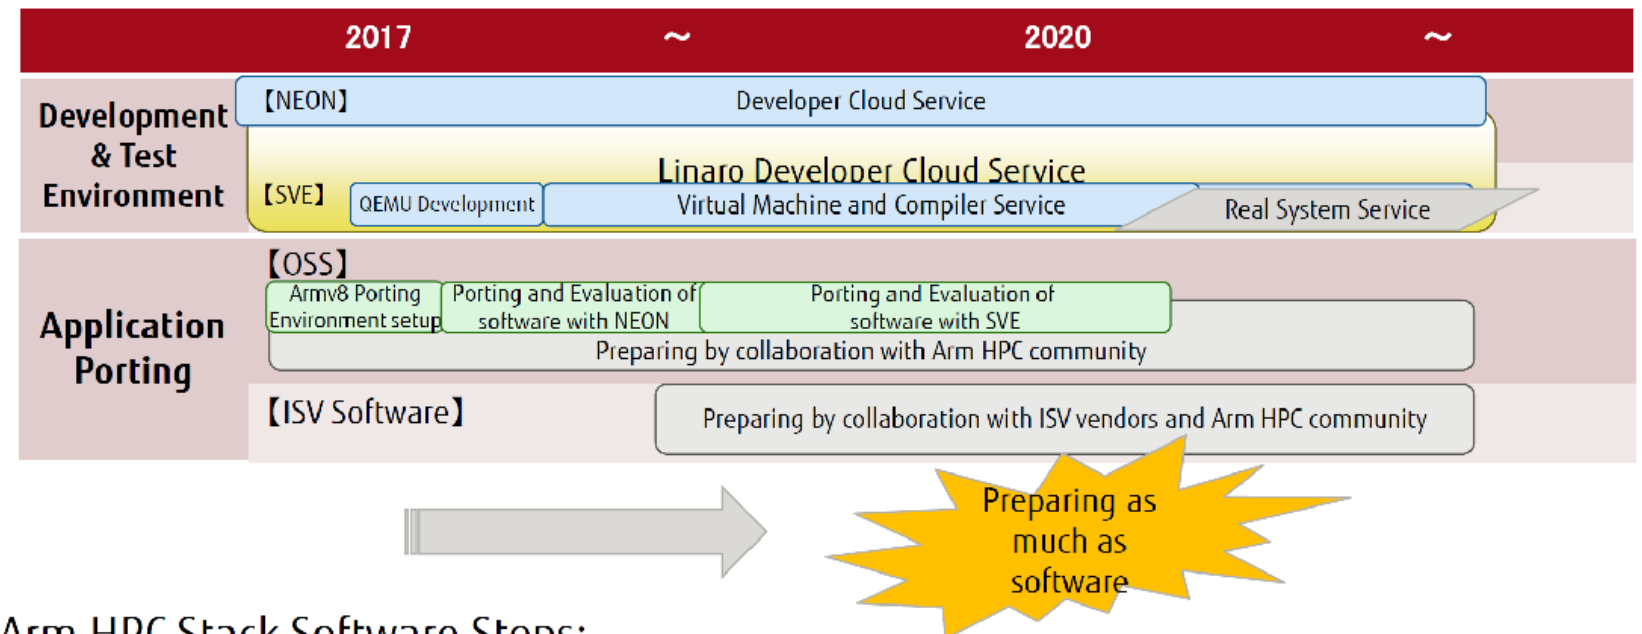

- Preparing Arm HPC Stack Software Steps:
  - Developing&Test Environment: Now NEON. Next SVE using Linaro Developer Cloud
    - For SVE, VM based developer service first. Next real system service
  - Application Porting: Now targeting OSSs. Next ISVs with Arm HPC community

Activities of Arm HPC User Group https://gitlab.com/arm-hpc

FUJITSU

- Fujitsu plans to contribute our experiences to Arm HPC User Group
  - Now comparing our results to the Arm HPC User Group site results.

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

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

shaping tomorrow with you

Post-K Target Application:

IDC HPC User Forum 2016@ Austin: from talk of Project Leader Prof. Ishikawa

FUJITSU

Target Applications' Characteristics

AICS

	Target Application		
	Program	Brief description	Co-design
①	GENESIS	MD for proteins	Collective comm. (all-to-all), Floating point perf (FPP)
②	Genomon	Genome processing (Genome alignment)	File I/O, Integer Perf.
③	GAMERA	Earthquake simulator (FEM in unstructured & structured grid)	Comm., Memory bandwidth
④	NICAM+LETK	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)	Comm., Memory bandwidth, File I/O, SIMD
⑤	NTChem	molecular electronic (structure calculation)	Collective comm. (all-to-all, allreduce), FPP, SIMD,
⑥	FFB	Large Eddy Simulation (unstructured grid)	Comm., Memory bandwidth,
⑦	RSDFT	an ab-initio program (density functional theory)	Collective comm. (bcast), FFP
⑧	Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)	Comm., Memory bandwidth, SIMD
⑨	CCS-QCD	Lattice QCD simulation (structured grid Monte Carlo)	Comm., Memory bandwidth, Collective comm. (allreduce)

RIKEN

5

Keeping Arm Binary Level Compatibility

FUJITSU

- Operating System Level Binary Portability
  - Two Specification defined by Arm, Linaro etc.
    - SBSA(Server Base System Architecture)
    - SBBR(Server Base Boot Requirements)
  - Armv8 distribution, such as RedHat, SUSE, can be used without modification
- System Software Level Binary Portability
  - Linaro is building system software stack for Arm HPC.
- Application binary portability with different SIMD width
  - Provided by Scalable Vector Extension(SVE) Specification

Execution Binary Portability

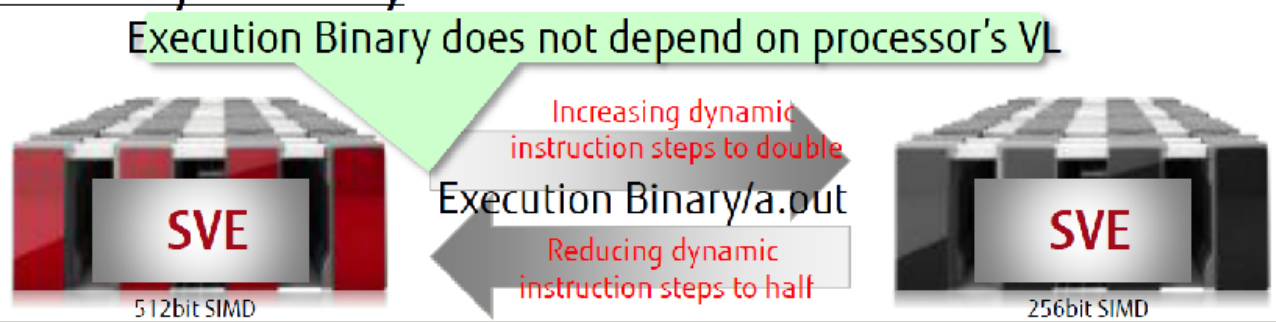

Towards building up Arm HPC Ecosystem

FUJITSU

- Goal:
  - Horizontal multi vendor collaboration to build up commodity platform
- Building Steps:
  1. Preparing Arm HPC system software standards
  2. Building Arm HPC Market
  3. Expanding Arm HPC Market as an HPC Commodity Platform
- The First Step: Preparing Arm HPC system software standards
  - Having a relationship with Arm HPC community is very important!

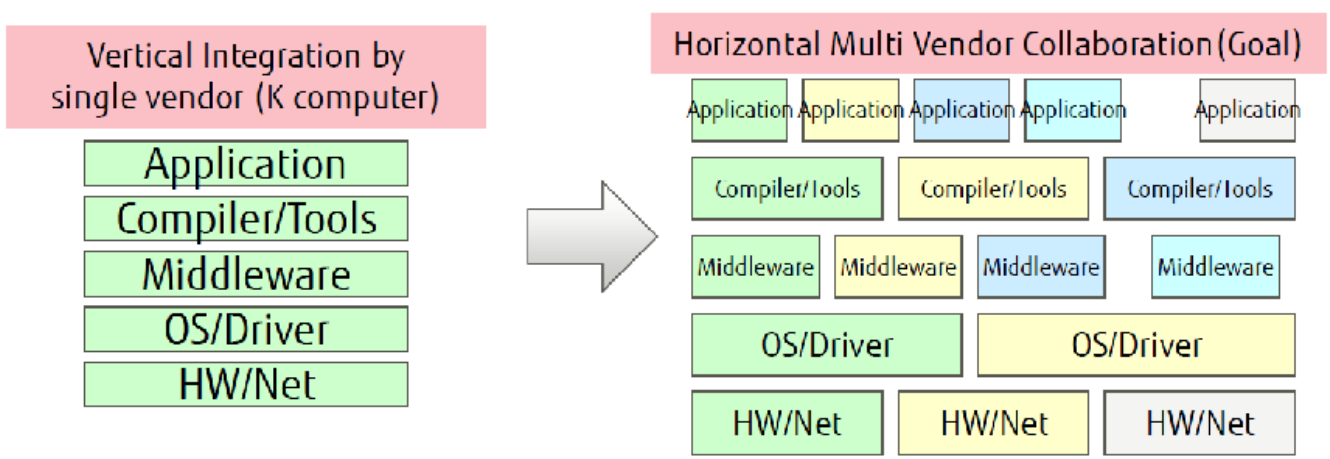

University of BRISTOL | Isambard | GW4

SC'17 Panel
"The Arm Software Ecosystem: Are We There Yet?"

Prof Simon McIntosh-Smith

University of Bristol, UK

University of BRISTOL | Isambard | GW4

'Isambard', a new Tier 2 HPC service from GW4.

Named in honour of Isambard Kingdom Brunel

University of BRISTOL

UNIVERSITY OF BATH

UNIVERSITY OF EXETER

CARDIFF UNIVERSITY
PRIFYSGOL CAERDYDD

Met Office

EPSRC

CRAY THE SUPERCOMPUTER COMPANY

ARM

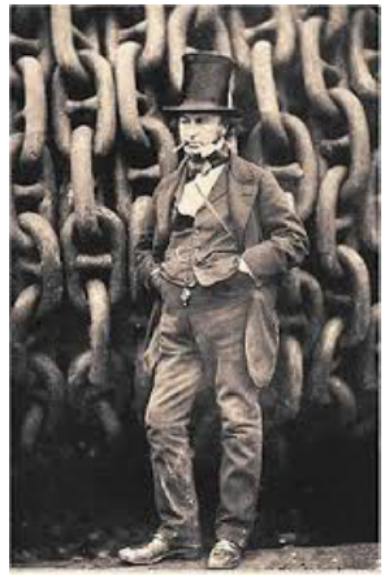

I.K.Brunel 1804-1859

University of BRISTOL | Isambard | GW4

Isambard system specification (red = new info):

- Cray "Scout" system – XC50 series
  - Aries interconnect
- **10,000+** Armv8 cores
  - Cavium ThunderX2 processors
  - 2x 32core @ 2.1GHz per node
- Cray software tools
- Technology comparison:
  - x86, Xeon Phi, Pascal GPUs
- Phase 1 installed March 2017
- The Arm part arrives Q2 2018

I.K.Brunel 1804-1859

University of BRISTOL | Isambard | GW4

- The Isambard project's focus will be on the top 10 most heavily used codes on Archer in 2017:
  - VASP, CASTEP, GROMACS, CP2K, UM, HYDRA, NAMD, Oasis, SBLI, NEMO
  - Note: 8 of these 10 codes are written in **FORTRAN**
- Additional important codes for project partners:
  - OpenFOAM, OpenIFS, WRF, CASINO, LAMMPS, …
- We want to collaborate wherever possible!
  - Accelerate the adoption of Arm in HPC

University of BRISTOL | Isambard | GW4

- We received our 8 node test system 2 weeks ago
- Ran our first hackathon 1 week ago
- So far, every application and mini-app we've tried, has compiled, run correctly, and performed well, out of the box
- That includes the Met Office's production climate/weather code, the UM
  - Millions of lines of Fortran and many complex dependencies
- Cray's first native version of CCE for Arm, 8.6.4, already performing well
- GCC 7.x and Arm Clang/Flang 18.x in good shape
- Math libraries such as OpenBLAS seem reasonable

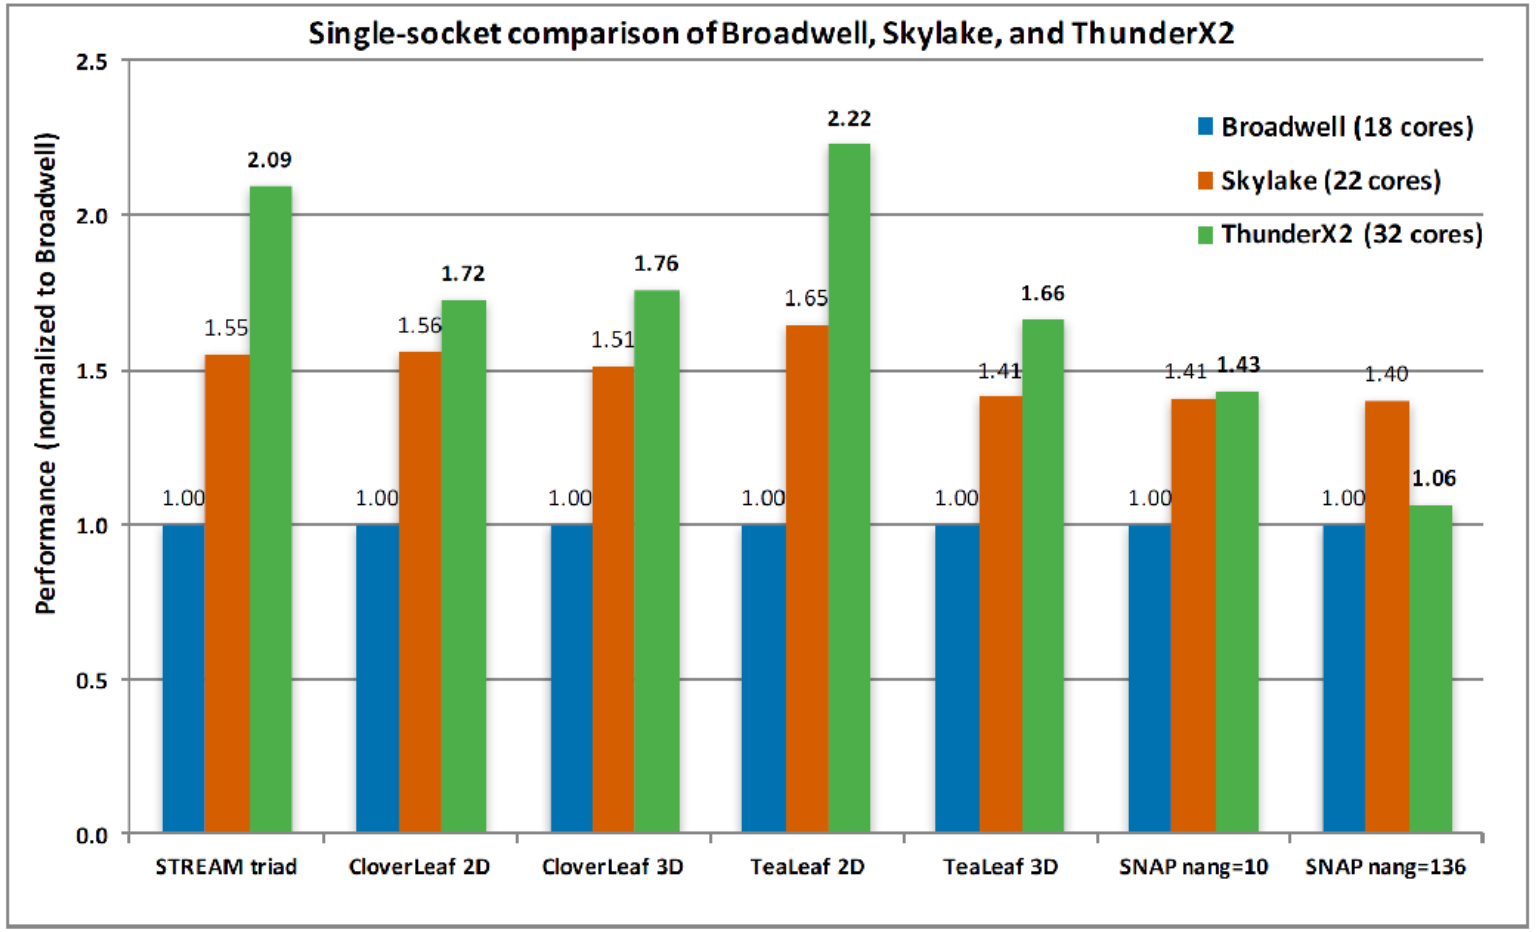
Single-socket comparison of Broadwell, Skylake, and ThunderX2
Performance (normalized to Broadwell)
2.5
2.0
1.5
1.0
0.5
0.0
Broadwell (18 cores)
Skylake (22 cores)
ThunderX2 (32 cores)
1.00
1.55
2.09
1.00
1.56
1.72
1.00
1.51
1.76
1.00
1.65
2.22
1.00
1.41
1.66
1.00
1.41
1.43
1.00
1.40
1.06
STREAM triad
CloverLeaf 2D
CloverLeaf 3D
TeaLeaf 2D
TeaLeaf 3D
SNAP nang=10
SNAP nang=136

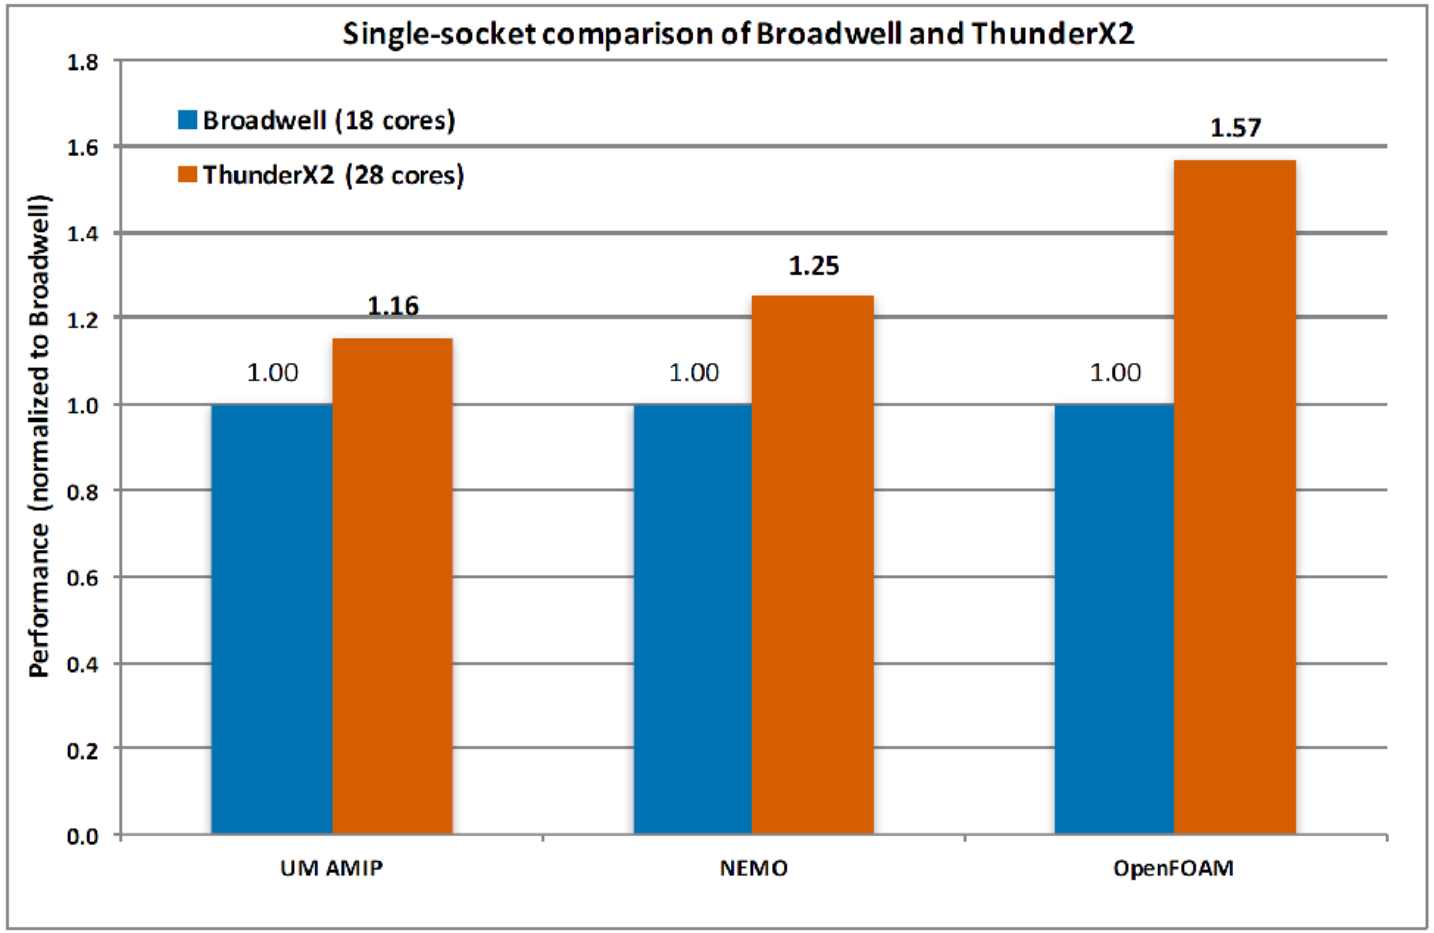
Single-socket comparison of Broadwell and ThunderX2
Performance (normalized to Broadwell)
1.8
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
Broadwell (18 cores)
ThunderX2 (28 cores)
1.00
1.16
1.00
1.25
1.00
1.57
UM AMIP
NEMO
OpenFOAM

Details for benchmarks comparisons:

- ThunderX2 systems
  - 28c, 2.0GHz, 2400MHz DDR4, A1 silicon, SLES 12 SP3
  - 32x, 2.5GHz, 2667MHz DDR4, A1 silicon,
  - Mixture of Cray CCE 8.6.4, GCC 7.2 and Arm Clang 18.0
  - Alpha release (pre-production) hardware, beta release software
- Broadwell system
  - 18c, 2.1GHz, 2400MHz DDR4, Xeon E5-2695 v4
  - Intel 2017 compiler except for UM and NEMO, which used CCE 8.5.8
- Skylake system
  - 22c, 2.1GHz, 2667MHz DDR4, Xeon
  - Intel 2018 compiler

University of BRISTOL | Isambard | GW4

For more information:

- http://gw4.ac.uk/isambard/
- https://github.com/UoB-HPC/GW4-Isambard
- Twitter: @simonmcs
- Email: simonm@cs.bris.ac.uk

ARC-TS ADVANCED RESEARCH COMPUTING TECHNOLOGY SERVICES UNIVERSITY OF MICHIGAN

Advanced Research Computing — Technology Services provides access to and support for advanced computing resources. ARC-TS facilitates new and more powerful approaches to research challenges in fields ranging from physics to linguistics, and from engineering to medicine.

Matt McLean
mattmc@umich.edu

- Big Data Systems Administrator @ ARC-TS
- HPC admin
- UNIX admin
- MFan!

The DSI

In 2015, U-M announced the Data Science Initiative, investing $100 million over five years to enhance opportunities for student and faculty researchers across the university to tap into the enormous potential of big data. The initiative included the goal of expanding U-M's research computing capacity.

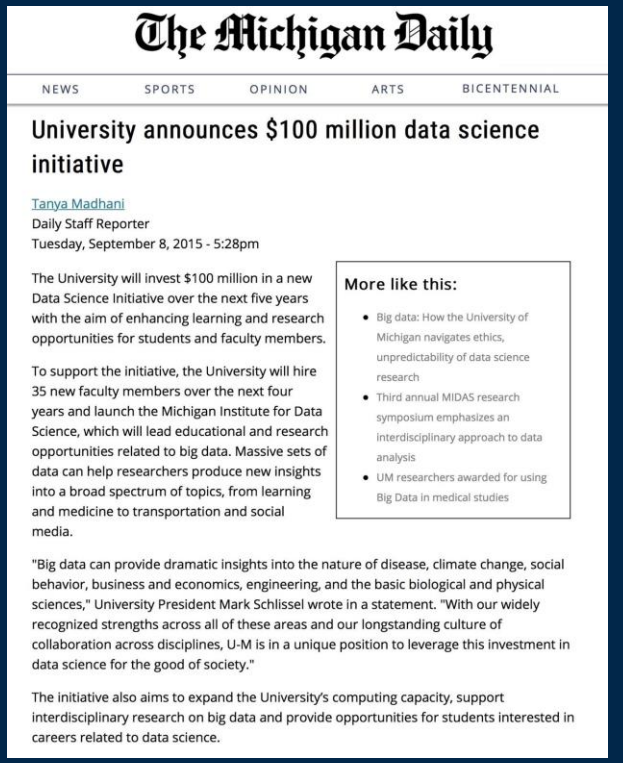

The Michigan Daily

NEWS SPORTS OPINION ARTS BICENTENNIAL

University announces $100 million data science initiative

Tanya Madhani
Daily Staff Reporter
Tuesday, September 8, 2015 - 5:28pm

The University will invest $100 million in a new Data Science Initiative over the next five years with the aim of enhancing learning and research opportunities for students and faculty members.

To support the initiative, the University will hire 35 new faculty members over the next four years and launch the Michigan Institute for Data Science, which will lead educational and research opportunities related to big data. Massive sets of data can help researchers produce new insights into a broad spectrum of topics, from learning and medicine to transportation and social media.

More like this:

- Big data: How the University of Michigan navigates ethics, unpredictability of data science research
- Third annual MIDAS research symposium emphasizes an interdisciplinary approach to data analysis
- UM researchers awarded for using Big Data in medical studies

"Big data can provide dramatic insights into the nature of disease, climate change, social behavior, business and economics, engineering, and the basic biological and physical sciences," University President Mark Schlissel wrote in a statement. "With our widely recognized strengths across all of these areas and our longstanding culture of collaboration across disciplines, U-M is in a unique position to leverage this investment in data science for the good of society."

The initiative also aims to expand the University's computing capacity, support interdisciplinary research on big data and provide opportunities for students interested in careers related to data science.

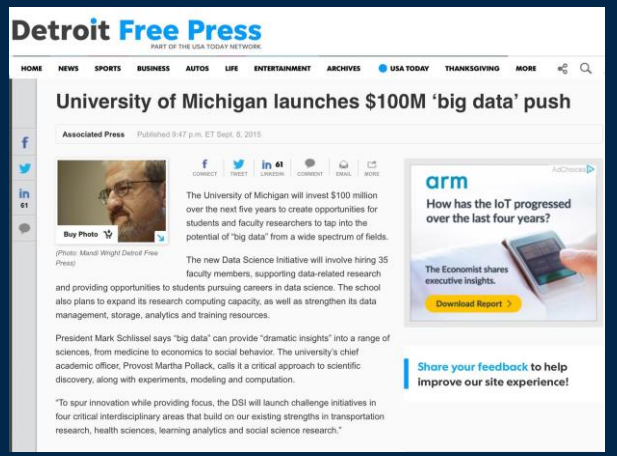

Detroit Free Press
PART OF THE USA TODAY NETWORK

HOME NEWS SPORTS BUSINESS AUTOS LIFE ENTERTAINMENT ARCHIVES USA TODAY THANKSGIVING MORE

University of Michigan launches $100M 'big data' push

Associated Press Published 9:47 p.m. ET Sept. 8, 2015

(Photo: Mandi Wright Detroit Free Press)

The University of Michigan will invest $100 million over the next five years to create opportunities for students and faculty researchers to tap into the potential of "big data" from a wide spectrum of fields.

The new Data Science Initiative will involve hiring 35 faculty members, supporting data-related research and providing opportunities to students pursuing careers in data science. The school also plans to expand its research computing capacity, as well as strengthen its data management, storage, analytics and training resources.

President Mark Schlissel says "big data" can provide "dramatic insights" into a range of sciences, from medicine to economics to social behavior. The university's chief academic officer, Provost Martha Pollack, calls it a critical approach to scientific discovery, along with experiments, modeling and computation.

"To spur innovation while providing focus, the DSI will launch challenge initiatives in four critical interdisciplinary areas that build on our existing strengths in transportation research, health sciences, learning analytics and social science research."

Share your feedback to help improve our site experience!

ARC-TS, in furtherance of that goal, set out to design and implement a big data Hadoop cluster for use by researchers across campus.

UNIVERSITY OF MICHIGAN

Cavium Partnership

In late 2016, Cavium, Inc. (NASDAQ: CAVM) agreed to gift U-M 100+ ThunderX servers valued at $3.5 million to build a large scale Hadoop cluster.

The gift is structured in three phases, and by year three the cluster will have over 10,000 cores, 55TB of memory, and 9PB of storage.

Because of this gift, ARC-TS will be able to provide a large scale Hadoop ecosystem **for free** to any researcher at U-M.

MARKETSINSIDER PORTFOLIO

MARKETS STOCKS INDICES COMMODITIES CURRENCIES MUTUAL FUNDS ETFS BONDS NEWS

University of Michigan Partners with Cavium on Big Data Computing Platform for U-M Researchers

PRESS RELEASE PR Newswire

Nov. 7, 2017, 12:13 PM

SHARE

SAN JOSE, Calif., Nov. 7, 2017 /PRNewswire/ -- Cavium™, Inc. (NASDAQ: CAVM), a leading provider of semiconductor products that enable secure and intelligent processing for enterprise, data center, wired and wireless networking and the University of Michigan today announced a new partnership that will position U-M as a leader in data-intensive scientific research by creating a powerful Big Data computing cluster using dual socket servers powered by Cavium's ThunderX® ARMv8-A workload optimized processors.

The University RECORD

News for Faculty and Staff

November 7, 2017

U-M, Cavium partner on big data research computing platform

By Dan Meisler

Advanced Research Computing

Topic: Information Technology

Like 64 Tweet Share print

A new partnership between the University of Michigan and Cavium Inc., a San Jose-based provider of semiconductor products, will create a powerful new big data computing cluster available to all U-M researchers.

The $3.5 million ThunderX computing cluster will enable U-M researchers to, for example, process massive amounts of data generated by remote sensors in distributed manufacturing environments, or by test fleets of automated and connected vehicles.

The cluster will run the Hortonworks Data Platform providing Spark, Hadoop MapReduce, and other tools for large-scale data processing.

+ MORE INFORMATION

- Cavium ThunderX Cluster at

CAVIUM

UNIVERSITY OF MICHIGAN

Hortonworks HDP

Java OpenJDK on ARM is ready today.

While the Hadoop ecosystem is written in Java, some components still have x86_64 dependencies or optimizations.

In mid 2017, ARC-TS partnered with HortonWorks, a leading provider of curated and tested Hadoop Ecosystem components to develop a version of its signature Hadoop software bundle HDP for use on ThunderX.

UNIVERSITY OF MICHIGAN

At U-M, the core of Hadoop is successfully running on ARM

What we have running today

- HDFS - HA
- YARN - HA
- Hive 2 - HA
- Spark 2
- Mapreduce
- Pig

```
[root@cavium-dn0037 ~]# hdfs dfs -df -h
Filesystem                Size    Used  Available  Use%
hdfs://cavium-hadoop  2.9 P  16.2 T      2.7 P    1%
```

```
[root@cavium-dn0037 ~]# /usr/hdp/2.6.2.0-205/spark2/bin/spark-submit --version
Welcome to
      ____              __
     / __/__  ___ _____/ /__
    _\ \/ _ \/ _ `/ __/  '_/
   /___/ .__/\_,_/_/ /_/\_\   version 2.1.1.2.6.2.0-205
      /_/

Using Scala version 2.11.8, OpenJDK 64-Bit Server VM, 1.8.0_121
```

```
[root@cavium-dn0037 ~]# hadoop version
Hadoop 2.7.3.2.6.2.0-205
Subversion https://github.com/hortonworks/hadoop-release.git -r fd80ec475
Compiled by root on 2017-09-08T01:52Z
Compiled with protoc 2.5.0
From source with checksum 90b73c4c185645c1f47b61f942230
```

```
[root@cavium-dn0037 ~]# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 7.3 (Maipo)
[root@cavium-dn0037 ~]# uname -io
aarch64 GNU/Linux
```

UNIVERSITY OF MICHIGAN

ARM Software Ecosystem: Are we there yet?

Early Experiences at ORNL

Verónica G. Vergara Larrea

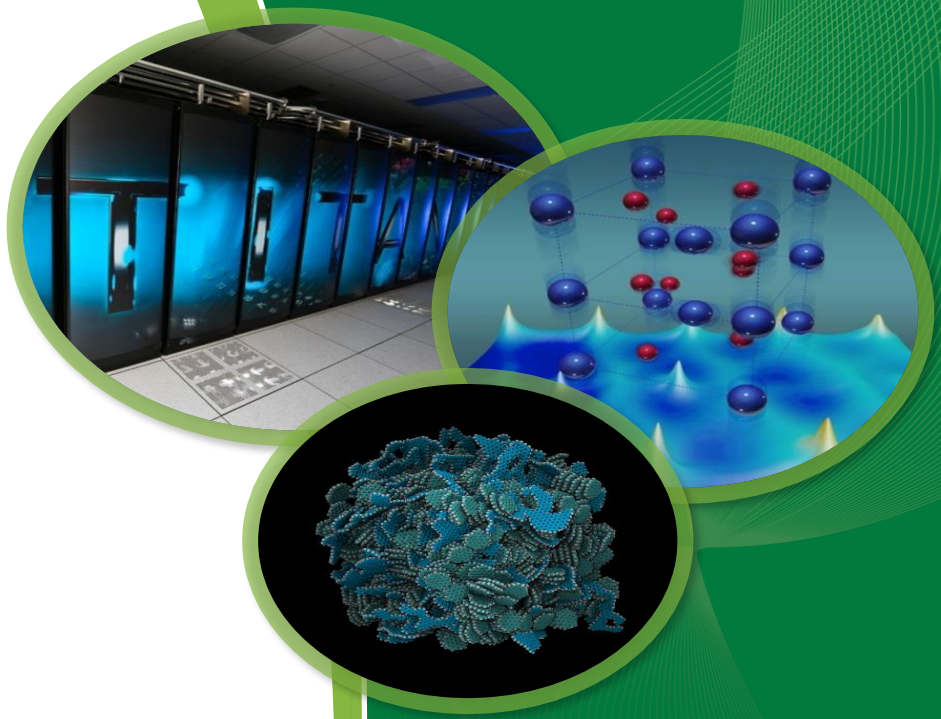

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Experiences at ORNL

- ARM systems at ORNL
  - Current systems:
    - **ARM1:** 16 node Cray Envoy system
      - Two 48-core ThunderX1 processors per node
    - 2 node HPE test system (arrived last week)
      - Two 48-core ThunderX2 processors per node (pre-production)
  - Next system:
    - **Wombat:** 16 node HPE system (1 login and 15 compute nodes)
      - Two 28-core ThunderX2 processors per node (pre-production)
      - Four nodes have two AMD GPUs each
- Critical parts of the Arm SW ecosystem for you
  - Compilers that support latest language standards and programming models
  - ...

Experiences at ORNL

  - Optimized math and I/O libraries
  - Tools for profiling and debugging
  - Drivers that support attached accelerators
  - Support for parallel file systems
- Biggest successes
  - Able to get Spack setup and working on the compute nodes
  - Built 192 libraries and dependencies, as well as Python packages
  - Able to run several acceptance applications built with the GCC compiler on Arm1
    - minisweep, LAMMPS, GTC, and more.
- Key concerns
  - Cross-compiling remains a challenge
  - Cray's Libsci does not provide C bindings
  - No Lustre client support for ARM
    - ORNL is leading this effort

Insights / General Thoughts

- Key lessons
  - Where there's greatest hope
    - For the most part, it has been relatively easy to build libraries with GCC
    - No need for cross-compiling on our next ARM system
  - What the biggest problems and surprises are likely to be
    - Applications are not tuned for the architecture
      - Sure, applications run but are they running fast? Next TODO on Wombat
    - Experimenting/evaluating AMD GPUs
- Where there's lack of consensus
  - Homogeneous vs heterogeneous ARM based systems
  - Which workloads are best suited for the architecture?
- What the path is to greater agreement
  - Functionality and performance evaluations will be key
  - Convince vendors and open source developers to start providing ARM builds
- Aspirations and inspirations
  - Run **all** the acceptance applications!
  - Do a full evaluation comparing both AMD and NVIDIA GPUs

SC'17 PANEL, "THE ARM SOFTWARE ECOSYSTEM: ARE WE THERE YET?"

Trent D'Hooge, LLNL

EXPERIENCE

- Existing and future deployments
  - Applied Micro, ThunderX1, ThunderX2
- Biggest successes
  - RedHat progress in bringing aarch64 software stack in sync with x86_64
  - Additional packages building and passing basic tests
- Key concerns
  - Users will find many issues as we scale up. Critical that fixes are pushed upstream for everyone to benefit from.
  - Proprietary software and lessons learned not shared with the community
  - GPU support from Nvidia and AMD. ROCm and Cudatoolkit

INSIGHTS

- Success will be:
  - Aarch64 clusters running in production
  - Viable option for future COTS clusters and software is not in the TCO equation
- Concerns
  - Dedication to HPC from ARM manufactures
    - Intel is spending a lot of time and money to get into ARM's core business, why does ARM want to go into Intel's? Where is the money in HPC?

SC'17 PANEL, "THE ARM SOFTWARE ECOSYSTEM: ARE WE THERE YET?"

Kevin Pedretti, Sandia

Compiler Dashboard

Early ThunderX2 Hardware, Single node

Sandia National Laboratories

Workload	GCC 7.1	Vendor A	Vendor B
STREAM			
GUPS			
MiniFE			
Pennant			
Sweep3D			
LULESH			
OpenMPI 2.1.2			
Kokkos Kernels			
Trilinos			
NaluCFD			

Things are working surprisingly well

There are some issues, but being rapidly addressed and resolved

Performance is looking good.
Excellent on memory bandwidth, on par for compute. Should get significantly better with GA hardware and software tuning.

Key: Fastest | Middle | Slowest

Results from Si Hammond @ Sandia

Compiler Dashboard

Early ThunderX2 Hardware, Single node

Sandia National Laboratories

Workload	GCC 7.1	Vendor A	Vendor B
STREAM			
GUPS			
MiniFE			
Pennant			
Sweep3D			
LULESH			
OpenMPI 2.1.2			
Kokkos Kernels			
Trilinos			
NaluCFD			

Key: Fastest | Middle | Slowest

Results from Si Hammond @ Sandia

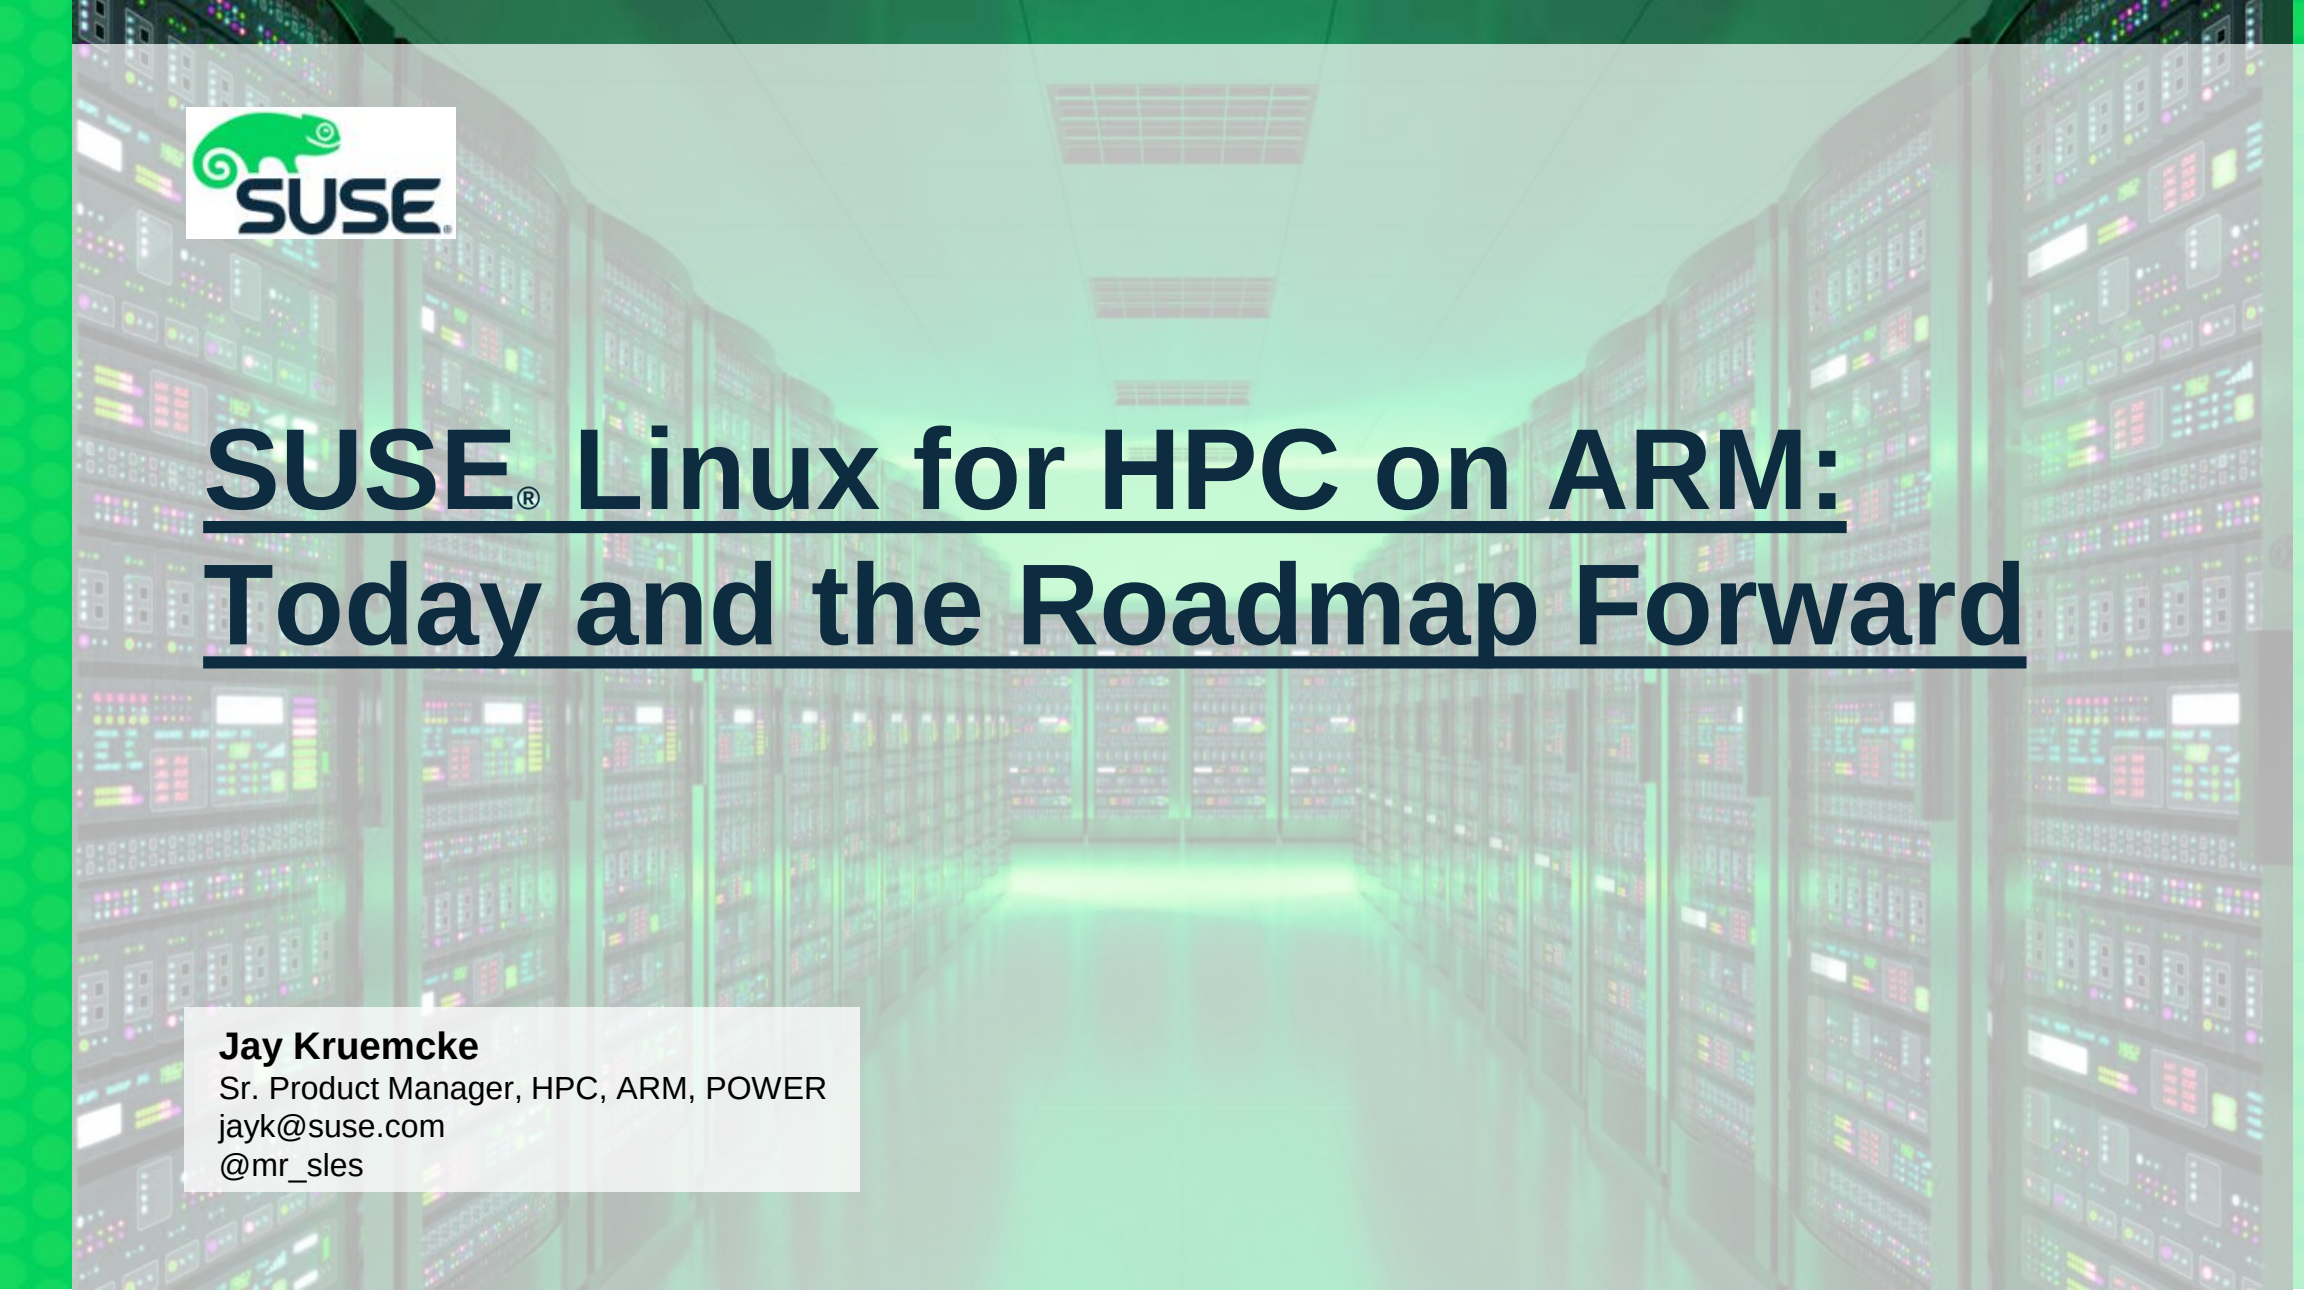

SUSE® Linux for HPC on ARM: Today and the Roadmap Forward

Jay Kruemcke
Sr. Product Manager, HPC, ARM, POWER
jayk@suse.com
@mr_sles

SUSE Linux Enterprise Server ARM Roadmap*

Offering commercial Linux support for ARM AArch64 since November 2016

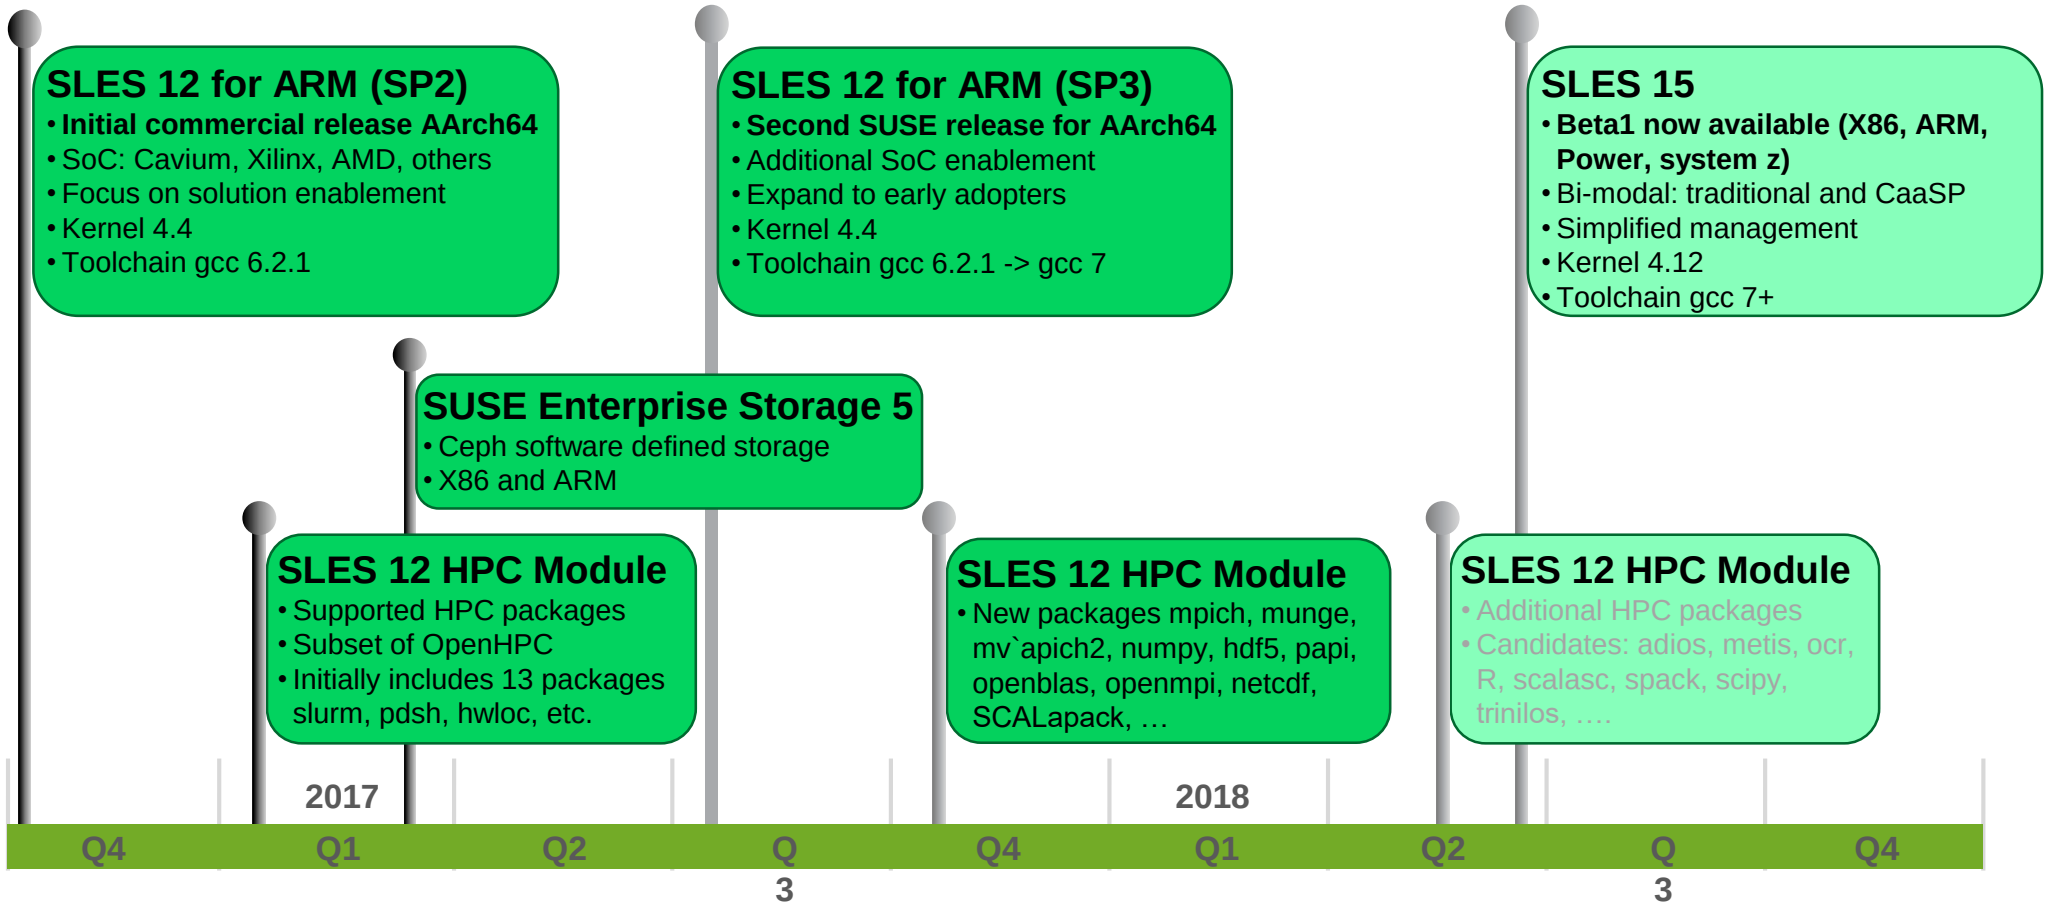

*All statements regarding future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only

SUSE Linux Enterprise 12 for ARM

SP3 includes enablement for these SoC vendors and chips*

- Advanced Micro Devices (AMD) – Opteron A1100
- Applied Micro - X-Gene 1, X-Gene 2
- Cavium - ThunderX
- **Cavium - ThunderX2 CN99 (SP3)**
- **Cavium - Octeon TX (SP3)**
- **HiSilicon - Hi1616 (SP3)**
- **Marvell - Armada 7K/8K (SP3)**
- **New Co / MACOM / Applied Micro - X-Gene 3 (SP3)**
- NXP - QorIQ LS2085A / LS2045A, LS2080A / LS2040A
- **NXP - QorIQ LS1043A (SP3)**
- **Qualcomm - Centriq 2400 (SP3)**
- **Rockchip - RK3399 (SP3)**
- Xilinx – Zynq UltraScale+ MPSoC

- Upstream kernel version: 4.4
- KVM with libvirt
- GCC 6.2.1 -> GCC 7 (3Q17)
- HPC module

Raspberry Pi 3 Model B images are also available with free, one year self-support
http://tinyurl.com/slespi

** Please check with your specific hardware vendor. Due to the rapidly evolving availability of ARM System on a Chip hardware, not all platforms have undergone the same degree of hardware testing*

SUSE Linux Enterprise HPC Module

Simplifying access to supported HPC software

- All packages supported by SUSE
  - Support included in the SLES Subscription
- Available for X86 and ARM platforms beginning with SLES 12 SP2
- Includes high demand packages for HPC workloads – more to come
- Flexible release schedule. Releases are independent of Service Pack schedule

Package		*HPC Module 1Q17*	*HPC Module 4Q17*
conman		0.2.7	**0.2.8**
cpuid	(X86 only)	20151017	**20170122**
fftw			**3.3.6**
hdf5			**1.10.1**
hwloc		1.11.5	
lua-filesystem		1.6.3	
lua-lmod		6.5.11	**7.6.1**
lua-luaterm		0.7	
lua-luaposix		33.2.1	
memkind	(X86 only)	1.1.0	
mpiP			**3.4.1**
mrsh		2.12	
munge		0.5.12	
mvapich2			**2.2**
netcdf			**4.4.1.1**
netcdf-cxx			**4.3.0**
netcdf-fortran			**4.4.4**
numpy			**1.13.3**
openblas			**0.2.20**
openmpi			**1.10.7**
papi			**5.5.1**
pdsh		2.31	**2.33**
petsc			**3.7.6**
phdf5			**1.10.1**
powerman		2.3.24	
prun		1.0	
rasdaemon		0.5.7	
ScaLAPACK			**2.0.2**
slurm		16.05.8	**17.02.09**

Other ARM related SUSE Products

SUSE Enterprise Storage

Available for X86 and ARM since 1Q17

SUSE Manager

SUSE OpenStack Cloud

Technology Demonstration

PANELIST QUESTIONS

- Have we fully enumerated the *de facto* dependences on a given SW ecosystem? What are they? (Refer to efforts like SPACK, CERN, OpenHPC, ORNL)
- What are the hardest technical challenges in porting to Arm? Examples include assembly code, memory ordering, lack of maturity wrt package versioning, interoperability.
- What fraction of apps that we care about will not run into those problems? Consider the starting apps at Fujitsu, Isambard, CERN, and the acceptance tests at ORNL.
- Which will be the last apps and frameworks to get ported to Arm, for technical reasons, and why?
- How much of an overlap can we expect between HPC and server, and different usages like simulation, deep learning and analytics, as we move to Arm?