



arm

Arm Developments in HPC

CIUK 2018, Manchester

Dr. Oliver Perks (olly.perks@arm.com)

12th December 2018

70%


of the world's population
uses Arm technology



arm

HPC Processors and Deployments

History of Arm in HPC: A Busy Decade




2011 Calxeda

- 32-bit ARmv7-A – Cortex A9



2011-2015
Mont-Blanc 1

- 32-bit Armv7-A
- Cortex A15
- First Arm HPC system




2014
AMD Opteron
A1100

- 64-bit Armv8-A
- Cortex A57
- 4-8 Cores



2015
Cavium
ThunderX

- 64-bit Armv8-A
- 48 Cores



2017
(Cavium)
Marvell
ThunderX 2

- 64-bit Armv8-A
- 32 Cores



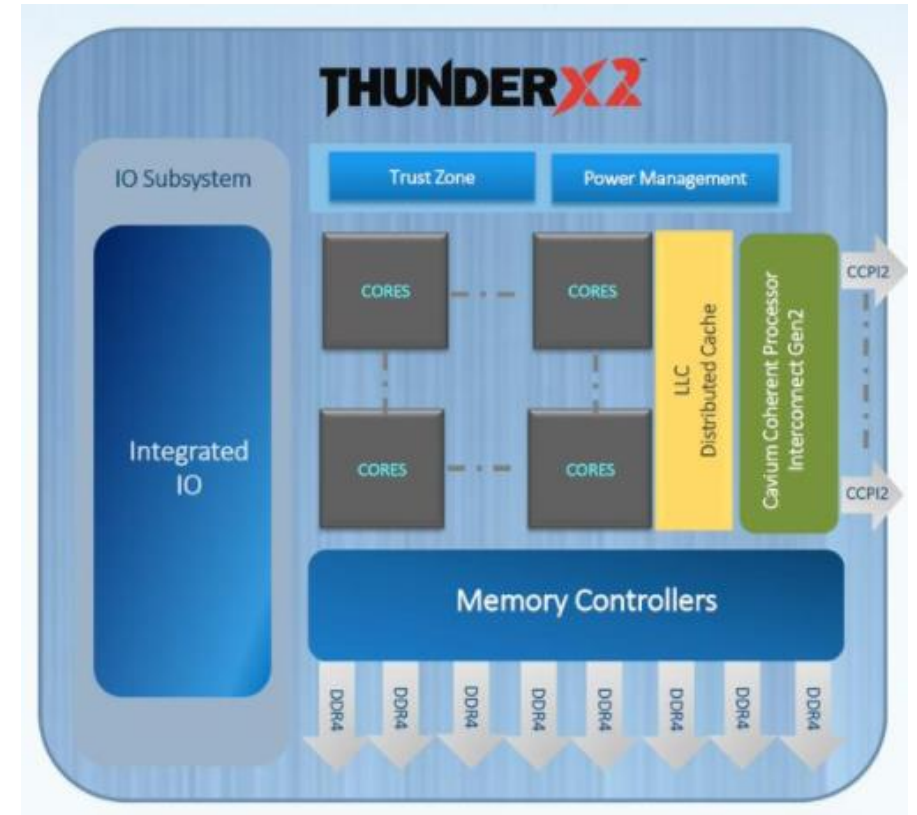
2019
Fujitsu A64FX

- First Arm chip with SVE vectorisation
- 48 Cores



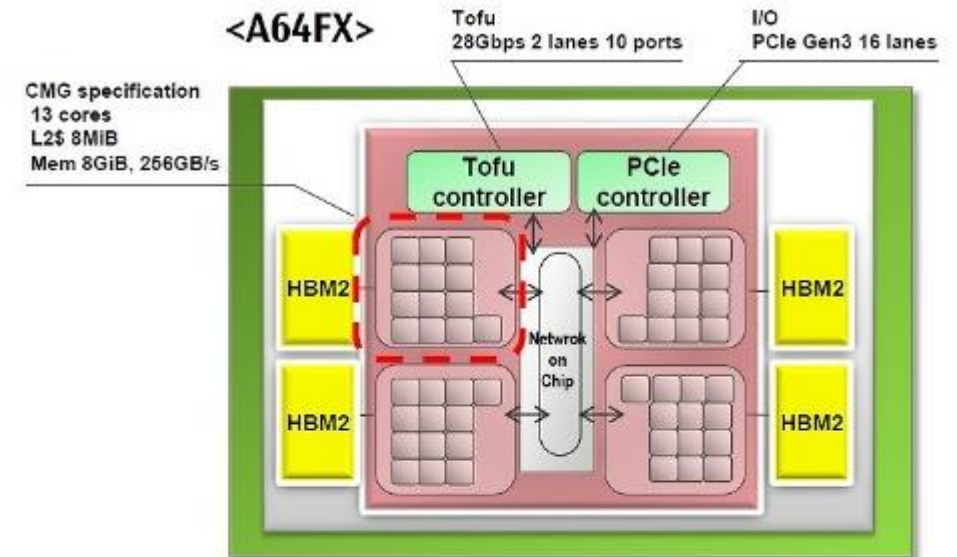
Marvell ThunderX2 CN99XX

- Marvell's next generation 64-bit Arm processor
 - Taken from Broadcom Vulcan
- 32 cores @ 2.2 GHz (other SKUs available)
 - 4 Way SMT (up to 256 threads / node)
 - Fully out of order execution
 - **8 DDR4 Memory channels (~250 GB/s Dual socket)**
 - Vs 6 on Skylake
- Available in dual SoC configurations
 - CCPI2 interconnect
 - **180-200w / socket**
- No SVE vectorisation
 - 128-bit NEON vectorisation



Fujitsu A64FX


- Chip designed for RIKEN POST-K machine
- 48 core 64-bit Armv8 processor
 - + 4 dedicated OS cores
- With SVE vectorisation
 - 512 bit vector length
- 32 GB HBM2
 - No DDR
 - 1 TB/s bandwidth
- TOFU 3 interconnect



	A64FX (Post-K)	SPARC64 Xlfx (PRIMEHPC FX100)
ISA (Base)	Armv8.2-A	SPARC-V9
ISA (Extension)	SVE	HPC-ACE2
Process Node	7nm	20nm
Peak Performance	>2.7TFLOPS	1.1TFLOPS
SIMD	512-bit	256-bit
# of Cores	48+4	32+2
Memory	HBM2	HMC
Memory Peak B/W	1024GB/s	240GB/s x2 (in/out)

Deployments: Isambard @ GW4

- **10,752** Armv8 cores (168 x 2 x 32)
 - **Cavium ThunderX2 32core 2.1GHz**
- Cray XC50 'Scout' form factor
- High-speed **Aries** interconnect
- Cray HPC optimised software stack
 - CCE, Cray MPI, Cray LibSci, CrayPat, ...
 - OS: CLE (Cray Linux Environment)
- Phase 2 (the Arm part):
 - **Accepted Nov 9th!**



Deployments: Catalyst

Arriving now!



Fulhame Catalyst system at EPCC

- Deployments to accelerate the growth of the Arm **HPC** ecosystem
- Each machine will have:
 - 64 HPE Apollo 70 nodes
 - Dual 32-core Marvell ThunderX2 processors
 - 4096 cores per system each
 - 256GB of memory / node
 - Mellanox InfiniBand interconnects
- OS: SUSE Linux Enterprise Server for HPC



Bristol: VASP, CASTEP, Gromacs, CP2K, Unified Model, NAMD, Oasis, NEMO, OpenIFS, CASINO, LAMMPS



THE UNIVERSITY of EDINBURGH

EPCC: WRF, OpenFOAM, Two PhD candidates



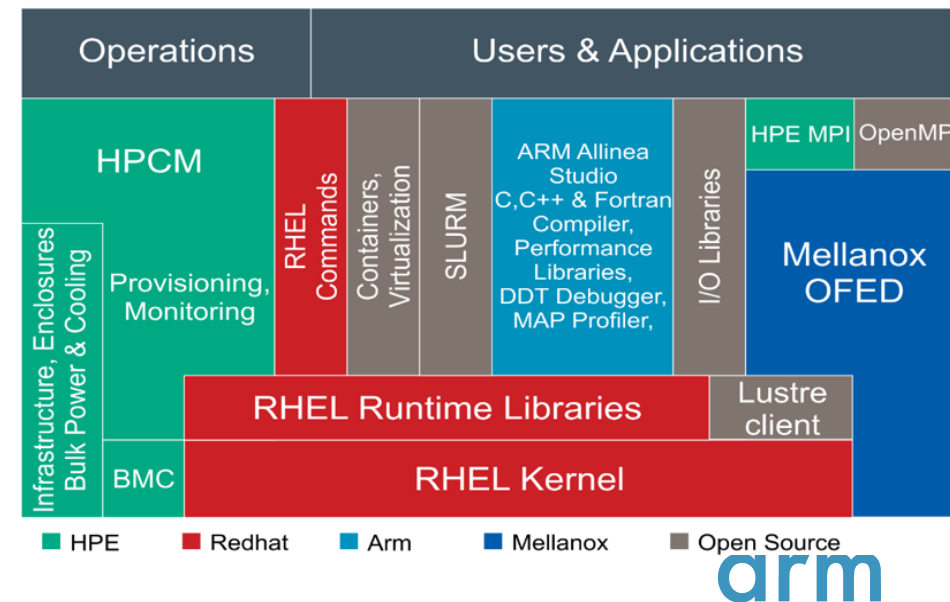
UNIVERSITY OF LEICESTER

Leicester: Data-intensive apps, genomics, MOAB Torque, DiRAC collaboration

Deployments: HPE Astra at Sandia


Mapping performance to real-world mission applications

- HPE Apollo 70
- #204 on Top500
 - 1.5 PFLOPs Rmax (2.0 PFLOPs Rpeak)
- Marvell ThunderX2 processors
 - 28-core @ 2.0 Ghz
 - 332 TB aggregate memory capacity
 - 885 TB/s of aggregate memory bandwidth
- 2592 HPE Apollo 70 nodes
 - 145, 152 cores
- Mellanox EDR InfiniBand
- OS: RedHat



The Cloud

Open access to server class Arm



AWS Graviton Processor



AWS Launches First Arm Cloud Instances

By Doug Black

November 28, 2018



c1.large.arm

With 96 physical Arm cores, this server is anything but a lightweight - and it comes with 128 GB of RAM for just \$0.50/hr. Nice!

arm

Software
Ecosystem



– Easy HPC stack deployment on Arm

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

Arm and partners actively involved:

- Arm is a silver member of OpenHPC
- Linaro is on Technical Steering Committee
- Arm-based machines in the OpenHPC build infrastructure

Status: 1.3.4 release out now

- Packages built on Armv8-A for CentOS and SUSE
- <https://developer.arm.com/hpc/hpc-software/openhpc>

Functional Areas	Components include
Base OS	CentOS 7.4, SLES 12 SP3
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, Hypra, 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, pdttoolkit, TAU, Scalasca, Score-P, SIONLib

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



Best-in-class performance



Commercially Supported
by Arm



Validated with
NAG test suite

Commercial 64-bit ArmV8-A math Libraries

- Commonly used low-level maths routines – BLAS, LAPACK and FFT
- Optimised maths intrinsics
- Validated with NAG's test suite, a de facto standard

Best-in-class performance with commercial support

- Tuned by Arm for specific core - like TX2 and Cortex-A72
- Maintained and supported by Arm for wide range of Arm-based SoCs

Silicon partners can provide tuned micro kernels for their SoCs

- Partners can contribute directly through open source routes
- Parallel tuning within our library increases overall application performance

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 petaflop applications)

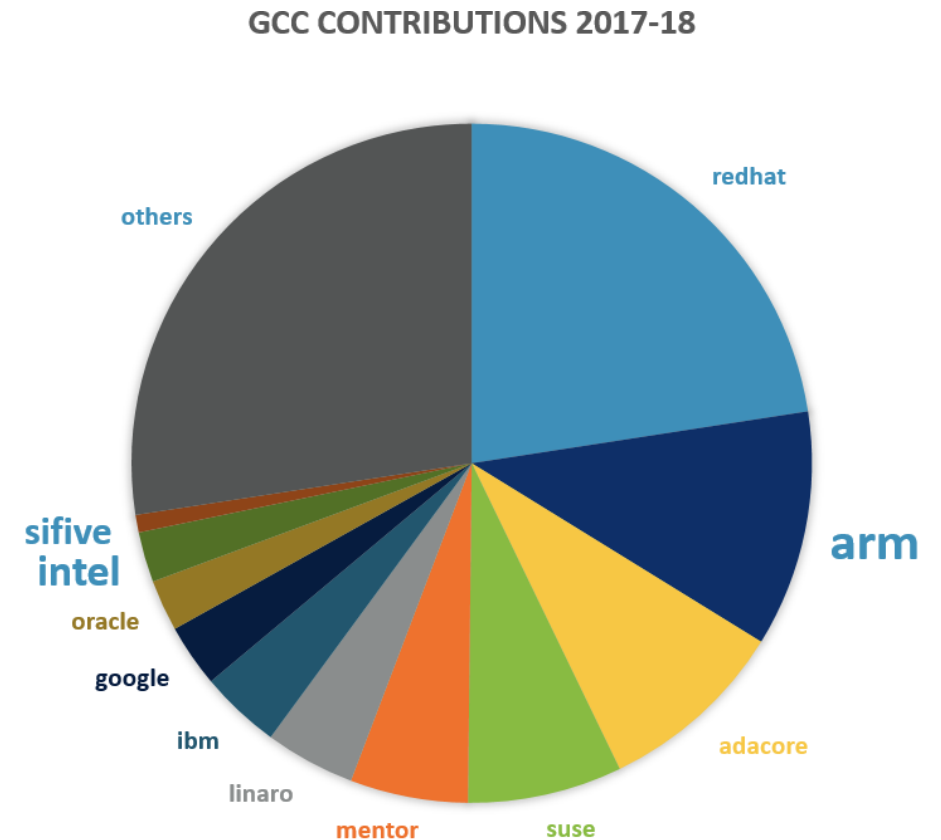
Easy to use by everyone

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

GCC is a major compiler in the Arm ecosystem

Arm the second largest contributor to the GCC project


- On Arm, GCC is a first class compiler alongside commercial compilers.
 - GCC ships with Arm Compiler for HPC
- SVE support since GCC 8.0
- NEON, LSE, and others well supported



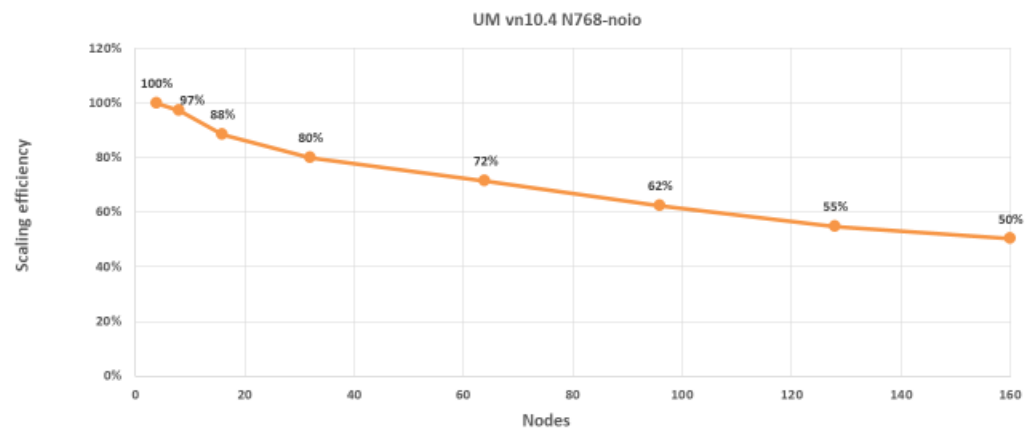
arm

Application Performance

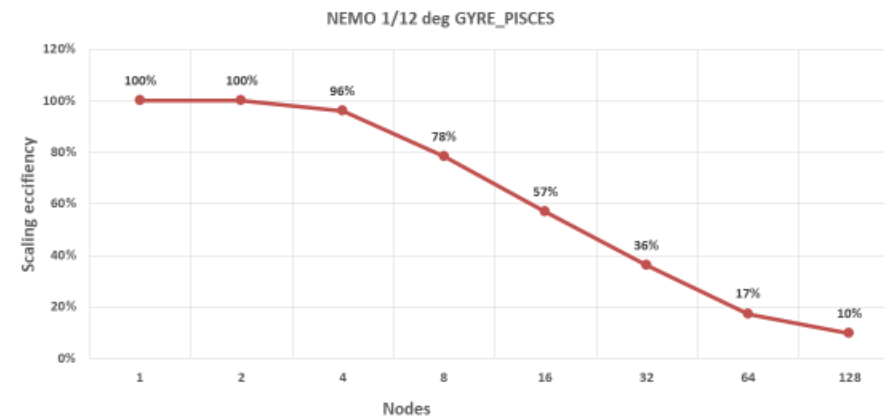
Single node performance results



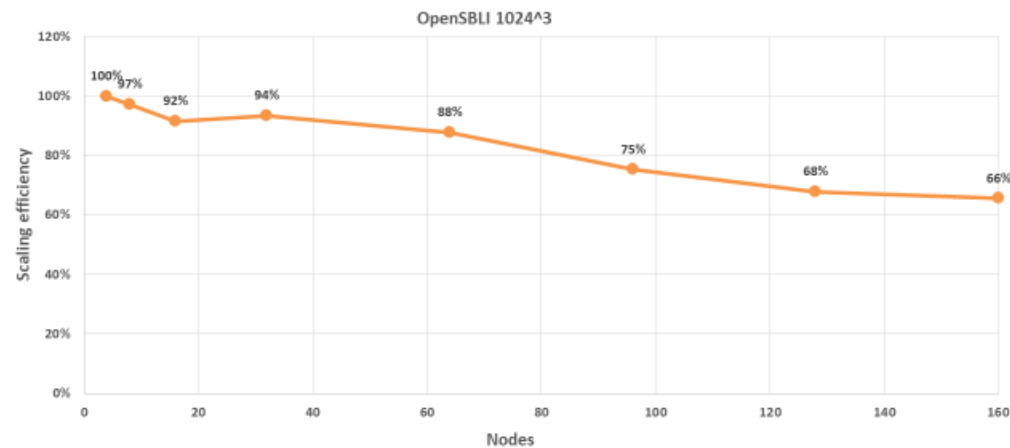
UM scalability, up to 10,240 cores




NEMO scalability, up to 8,192 cores



OpenSBLI scalability, up to 10,240 cores



GROMACS scalability, up to 8,192 cores



ArmPL: DGEMM performance

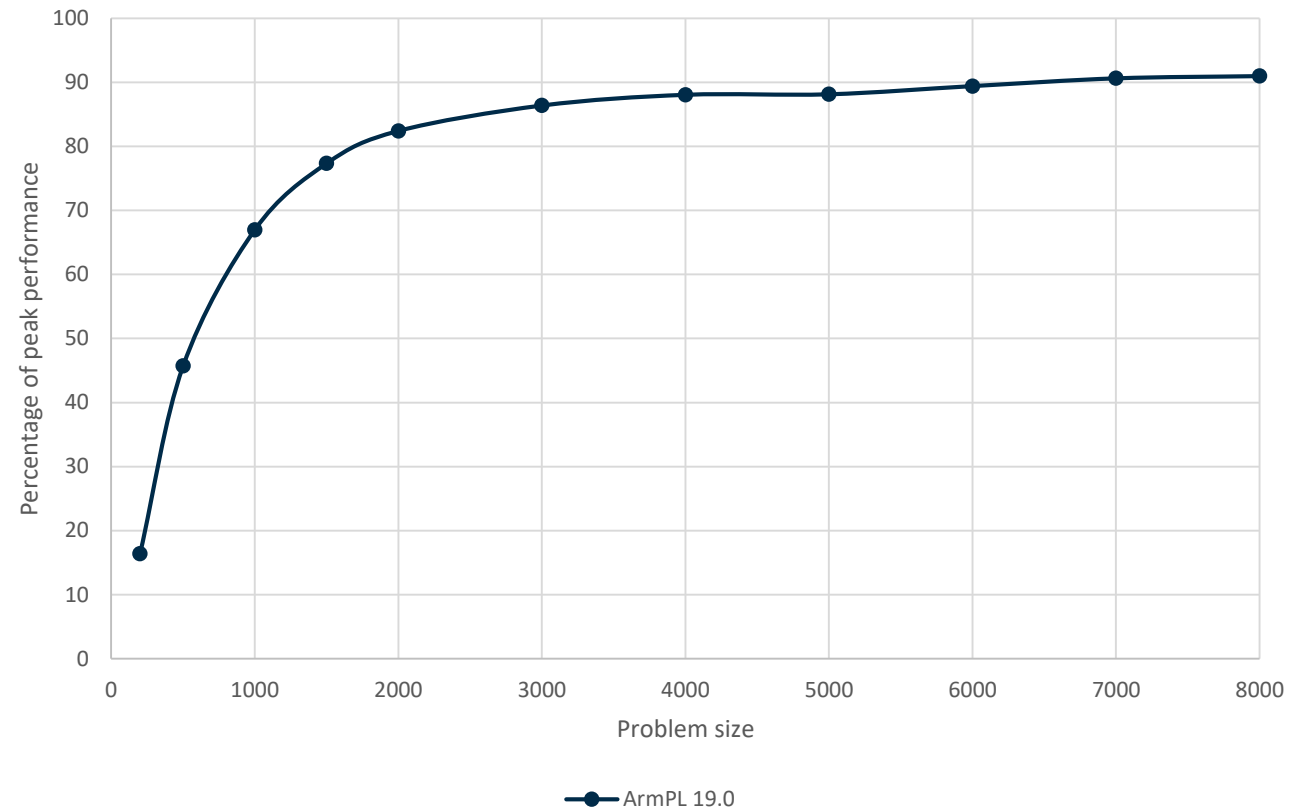
Excellent serial and parallel performance

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 (not shown)

Parallel performance at 90% of peak
Improved parallel performance for small problem sizes

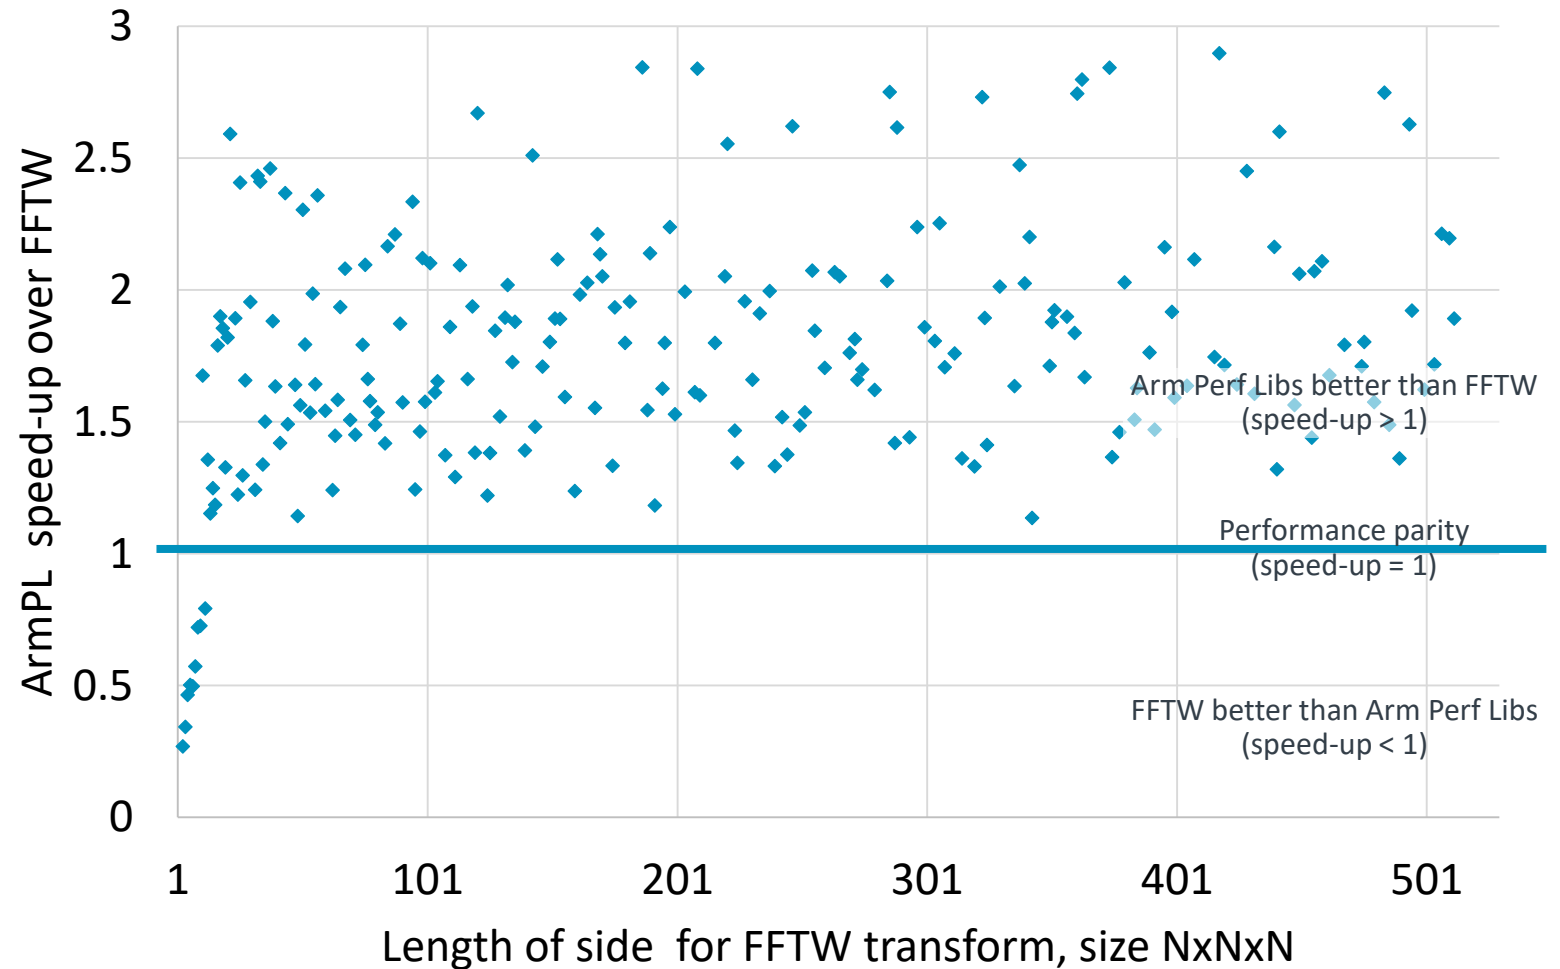
Arm PL 19.0 DGEMM parallel performance improvement for Cavium ThunderX2 using 56 threads



ArmPL 19.0 FFT 3-d complex-to-complex vs FFTW 3.3.7

- FFT timings on CTX2
 - 1-d complex-complex double precision
 - Higher means new Arm Performance Libraries FFT implementation better than FFTW 3.3.7
- Results show:
 - Average about 30% faster than FFTW
 - Lots of this comes from better usage of vectors on CTX2
 - Cases where we are still slower are:
 - Very small (no major issue)
 - Powers of primes
 - 143, 187 and 198 times tables

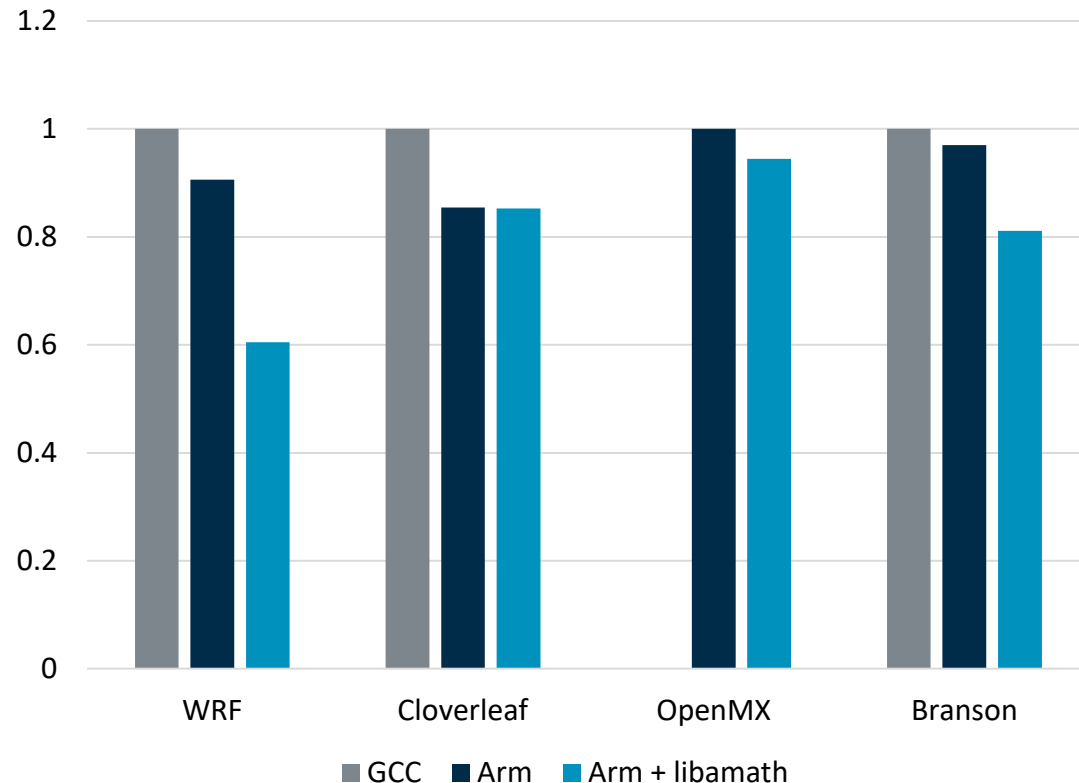
Arm Performance Libraries 19.0 vs FFTW 3.3.7
Complex-to-complex double precision 3-d transforms



Math Routine Performance

Aarch64 optimised Glibc maths routines

Normalised runtime



Arm PL provides libamath

- With Arm PL (*-armpl*).
 - Algorithmically better performance than standard library calls
 - No loss of accuracy
 - single and double precision implementations of: `exp()`, `pow()`, and `log()`
 - single precision implementations of: `sin()`, `cos()`, `sincos()`, `tan()`
- ...more to come.*

arm

Scalable Vector Extension (SVE)

What makes it a Scalable Vector Extension?



	1	2	3	4
+	5	5	5	5
<i>pred</i>	1	0	1	0
=	6	2	8	4

SVE does not have a fixed vector length

- Vector Length is hardware implementation choice of 128 to 2048 bits.

SVE is not an extension of AArch64 128-bit Advanced SIMD (*aka* Neon)

- A completely new set of instructions that can scale to future generations of vector processor.
- Initial focus is HPC scientific supercomputers and machine learning

SVE begins to remove barriers to auto-vectorization of high-level languages

- Per-lane predication on vector registers
- Compilers are often unable to vectorize code due to control and data dependencies.
- Software-managed speculative vectorization allows more code to be vectorized

```
for (i = 0; i < n; ++i)
  INDEX i
  CMPLT n
```

n-2	n-1	n	n+1
1	1	0	0

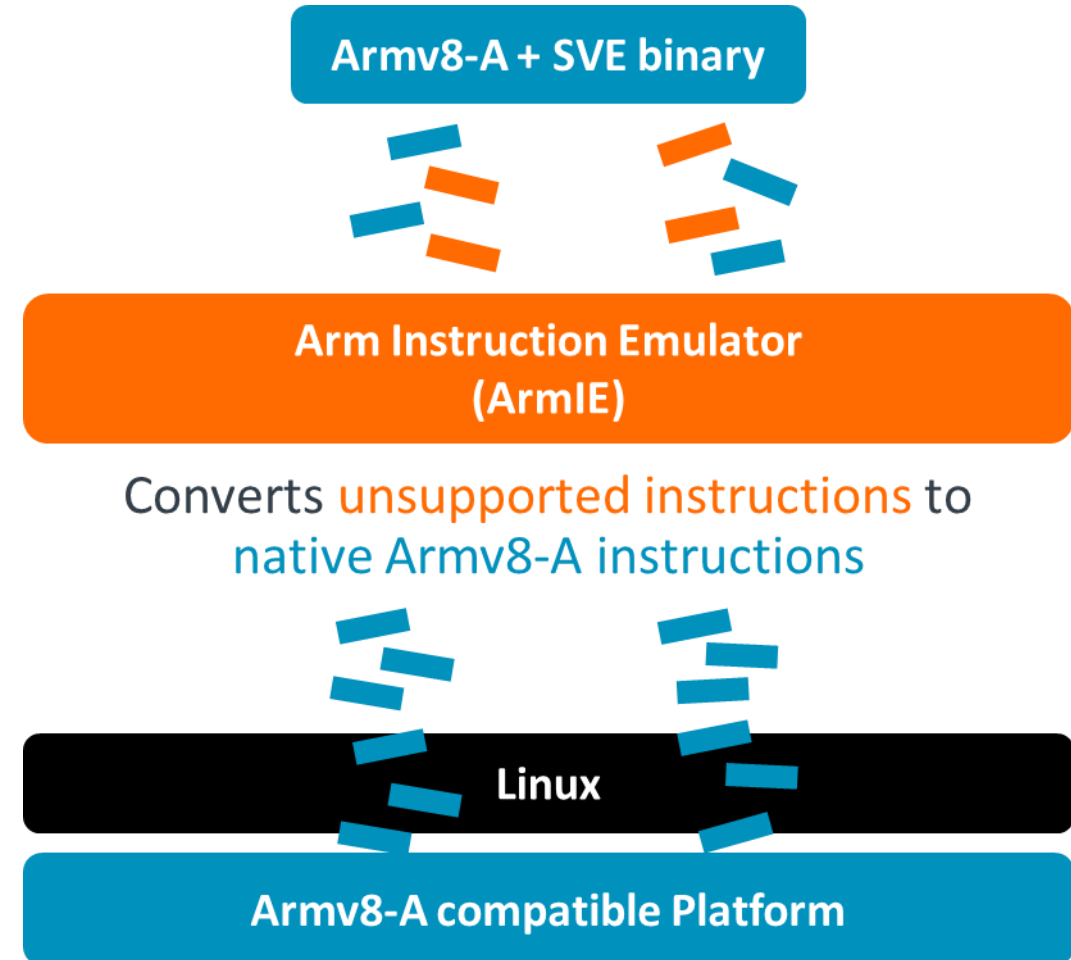
1	2	0	0
+	1	2	0
<i>pred</i>	1	1	0

1	+	2	+	3	+	4	=
1	+	2	+	3	+	4	=
3	+	7	=				

Running SVE without SVE Hardware

Arm Instruction Emulator: ArmIE

- Compilers can generate SVE
 - Arm Compiler, GNU, Fujitsu, etc.
- However no hardware to actually run it on
 - A64FX will be the first SVE hardware
- Need a way of understanding behaviour
 - Both compiler and application
 - Gem5 is great for hardware simulation, but slow
- ArmIE lets you test SVE vectorisation
 - With real applications
 - Emulate different vector widths



arm

Going Forward

The Future of Arm in HPC

What's next?

Processors

- By more vendors
 - Marvell, Ampere, Amazon, HiSilicon, Fujitsu
- Targeting different market segments
- All built on the Arm ecosystem
- Supported by the tools

Deployments

- Large scale and small scale deployments
- Increased exposure to the architecture
- More applications and libraries ported to Arm
 - Including ISVs
- Increased community

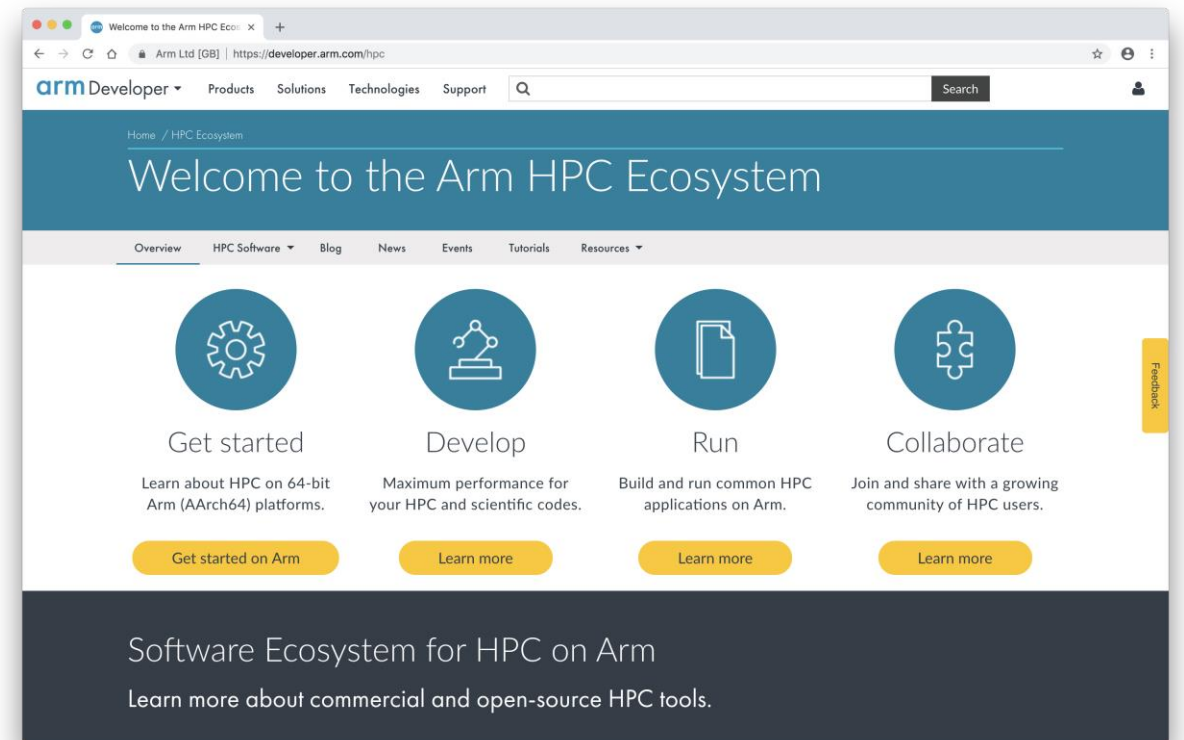
Commitment from Arm

- Neoverse – IP roadmap for silicon vendors
- Investment in software ecosystem
 - E.g. F18
- Support for customers
 - Applications
 - Software
 - Performance

Arm HPC Ecosystem website: www.arm.com/hpc

Get involved

- News, events, blogs, webinars, etc.
- Guides for porting HPC applications
- Quick-start guides for tools
- Links to community collaboration sites
- Arm HPC Users Group (AHUG)



arm

Come and visit us Booth #37

The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

www.arm.com/company/policies/trademarks

