Bootstrapping a HPC Ecosystem

Eric Van Hensbergen

Fellow – Senior Director of HPC

Software and Large Scale Systems Research

Teratech Forum

June 19, 2018

ARM computing is everywhere

#1

shipping GPU in the world is Mali

6.6Bn

ARM-based embedded chips shipped in 2016

> 5Bn

people using ARM-based mobile phones

>100Bn

ARM-based chips to date

The ARM business model

Global leader in the development of semiconductor IP

- R&D outsourcing for semiconductor companies

Innovative business model

- Upfront license fee – flexible licensing models
- Ongoing royalties – typically based on a percentage of chip price
- Technology reused across multiple applications

Create and transform markets

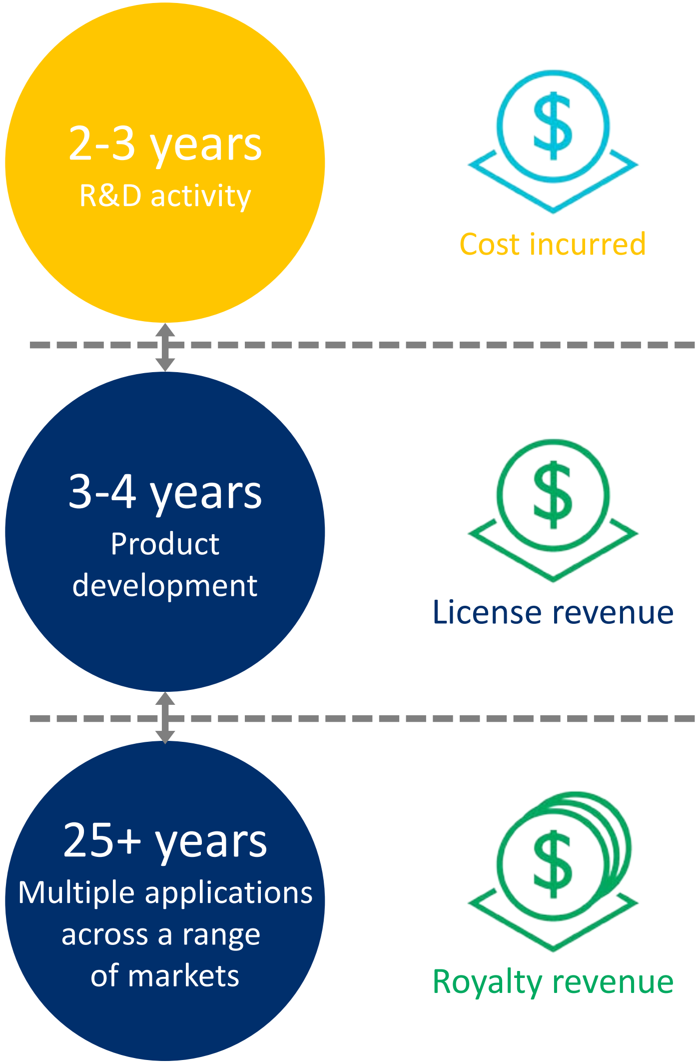

arm

A Brief History of ~~ARM~~ Arm

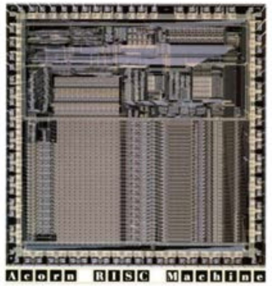

1985 ARM1
3µ 6k gates
7mm x 7mm

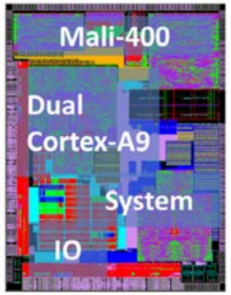

Cortex-A9 SOC
40nm 100M gates
7.4mm x 6.9mm

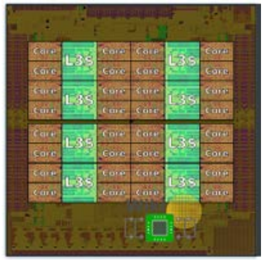

Cavium ThunderX2 SoC
32 Cores / 128 Threads
16nm

1985

1990

2005

2018

Early Days: ARM11MP

- Introduced in 2005
- Implementation in Cortex-A9
- Snoop control unit controlled coherence
  - Connects between 1 and 4 cores
  - Initiate L2 AXI memory accesses
  - arbitrate between Cortex-A9 processors requesting L2 accesses
  - manage ACP accesses

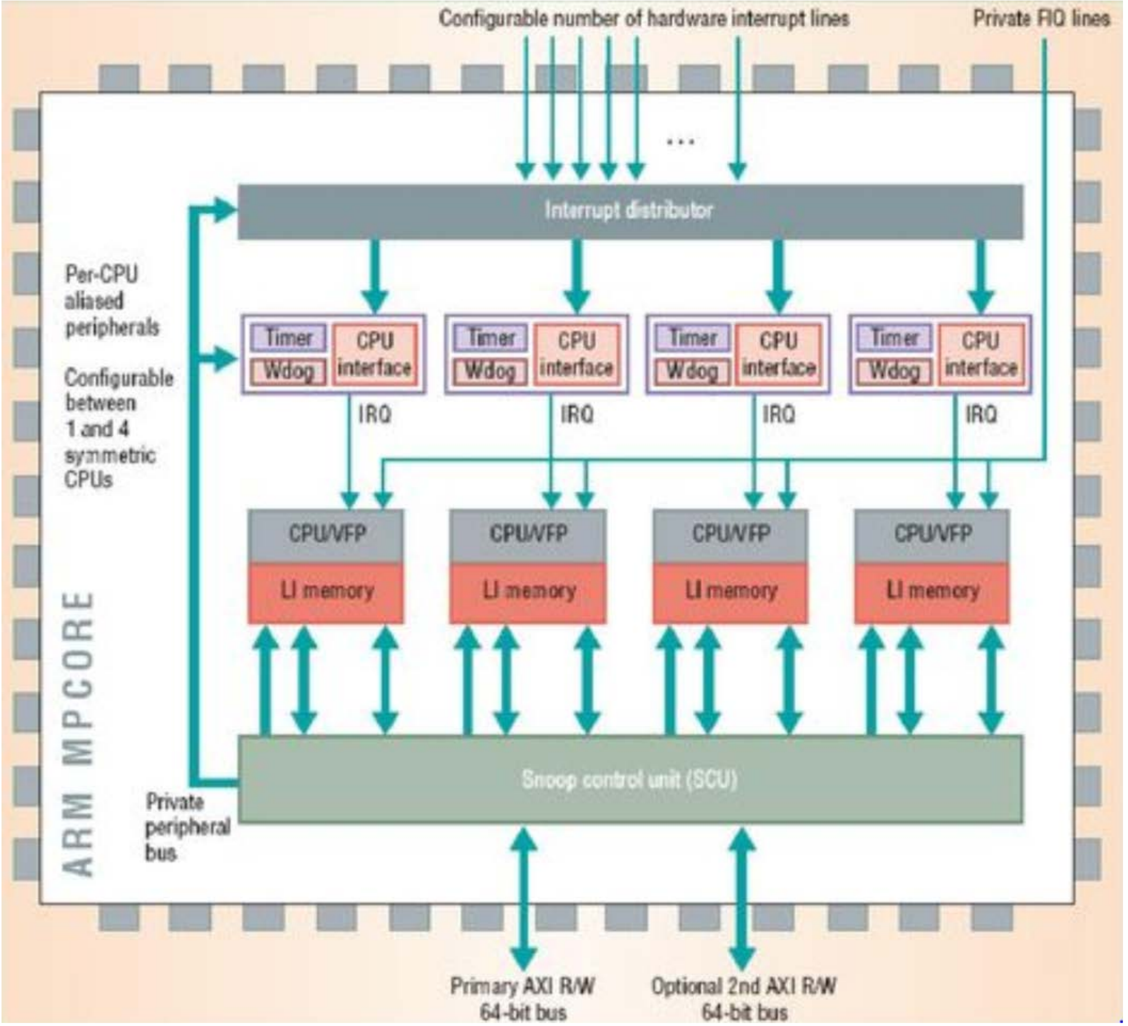

2011-2012: Tegra2 & SECO Boards

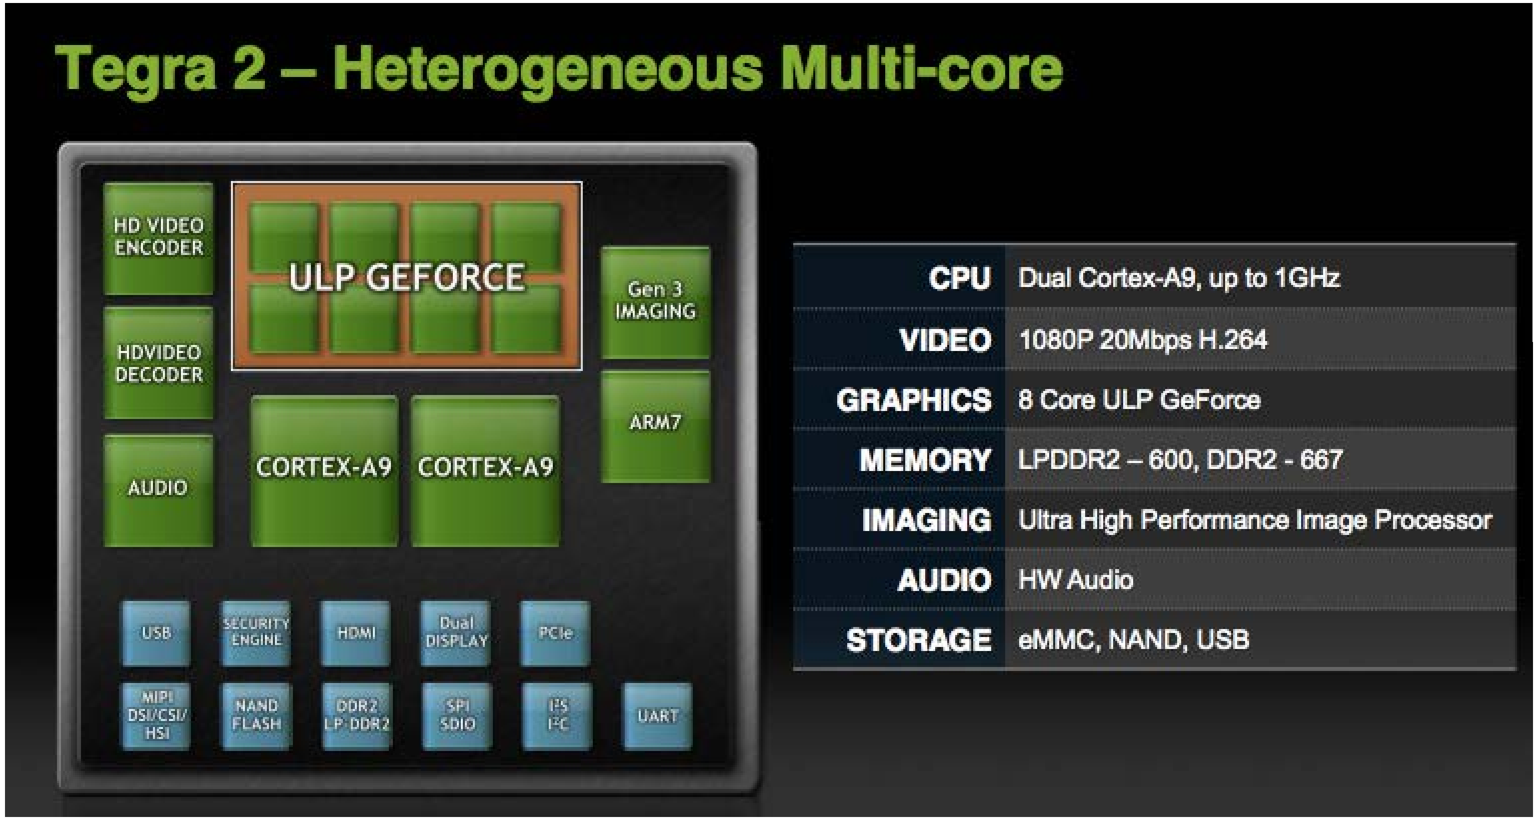

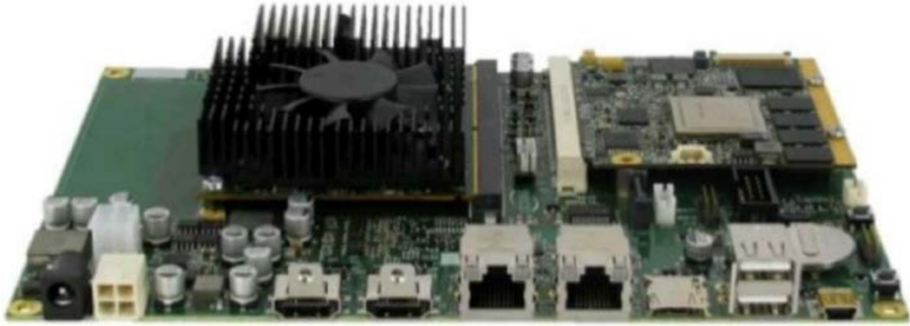

arm

2013- Montblanc Prototype 2014

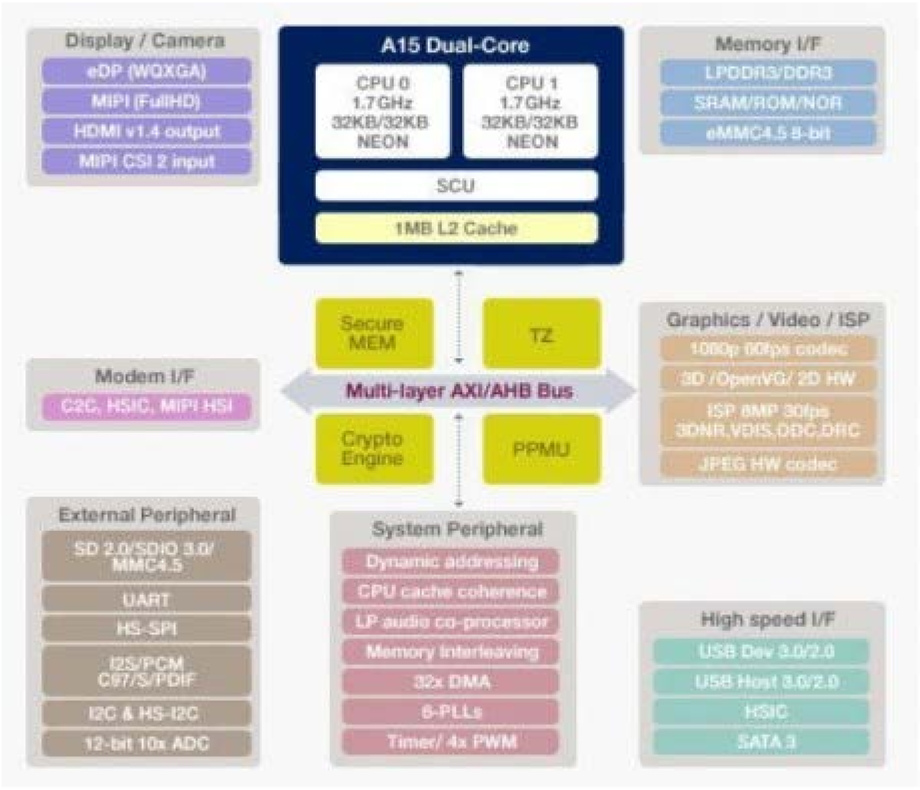

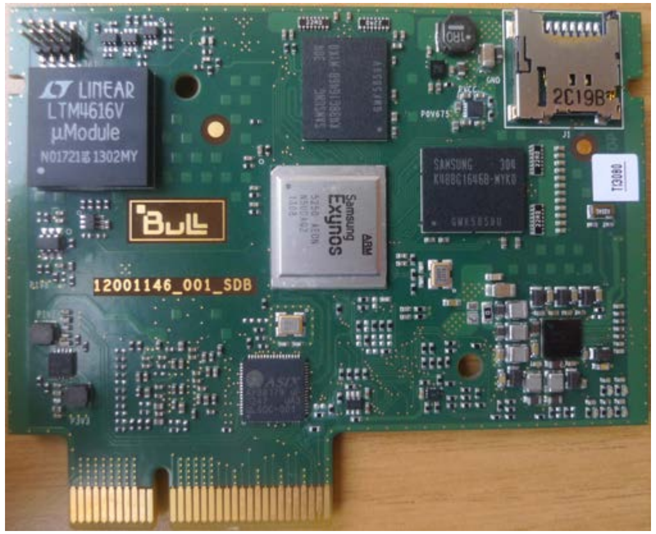

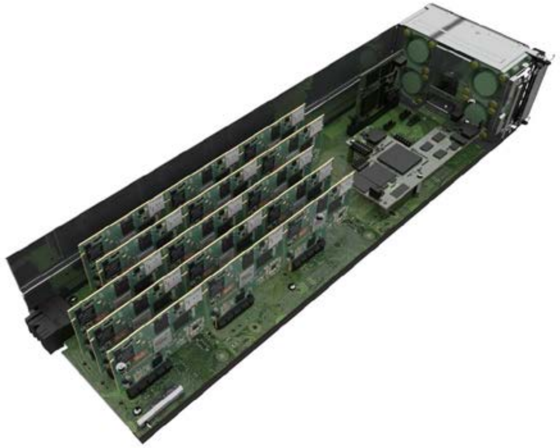

arm

Early Montblanc Stack

- Custom Linux kernels
  - No distro support
  - Limited driver support
- Open-source stack
  - No commercially supported tools for HPC
  - MPI unoptimized
- Academic runtime stacks to support heterogenous systems

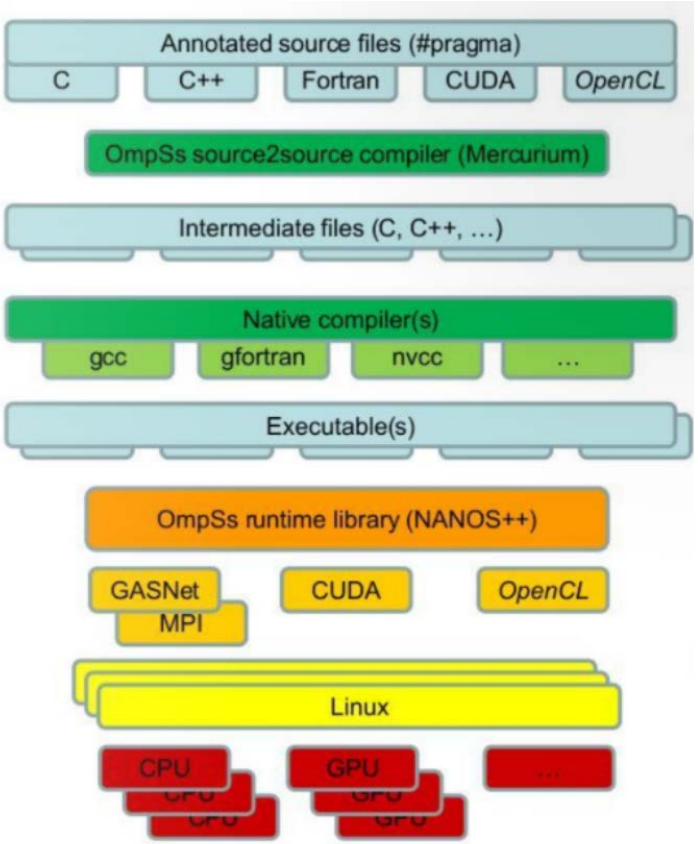

2013-2014 Prototype

Opportunities

- Extremely low-power, good track towards byte/flop ratio
- Most software "just worked" after compilation with Arm tools
- High level of integration on SoC limited need for complicated motherboards
- Low-power nodes & small form-factor allowed very dense packaging at rack level

Challenges

- Software porting to GPU cumbersome, particularly w/OpenCL
- Mobile SoC had extraneous devices (video/audio drivers, etc.), but missing high-performance PCI
- Resulting USB Networking less than performant for HPC
- Exiting platforms had limited memory capacity
- 32-bit, low single thread and low-FLOPs
- Embedded GPUs were disappointing on performance and programmability

November 2014: Co-Design w/DoE Fast-Forward 2 Program

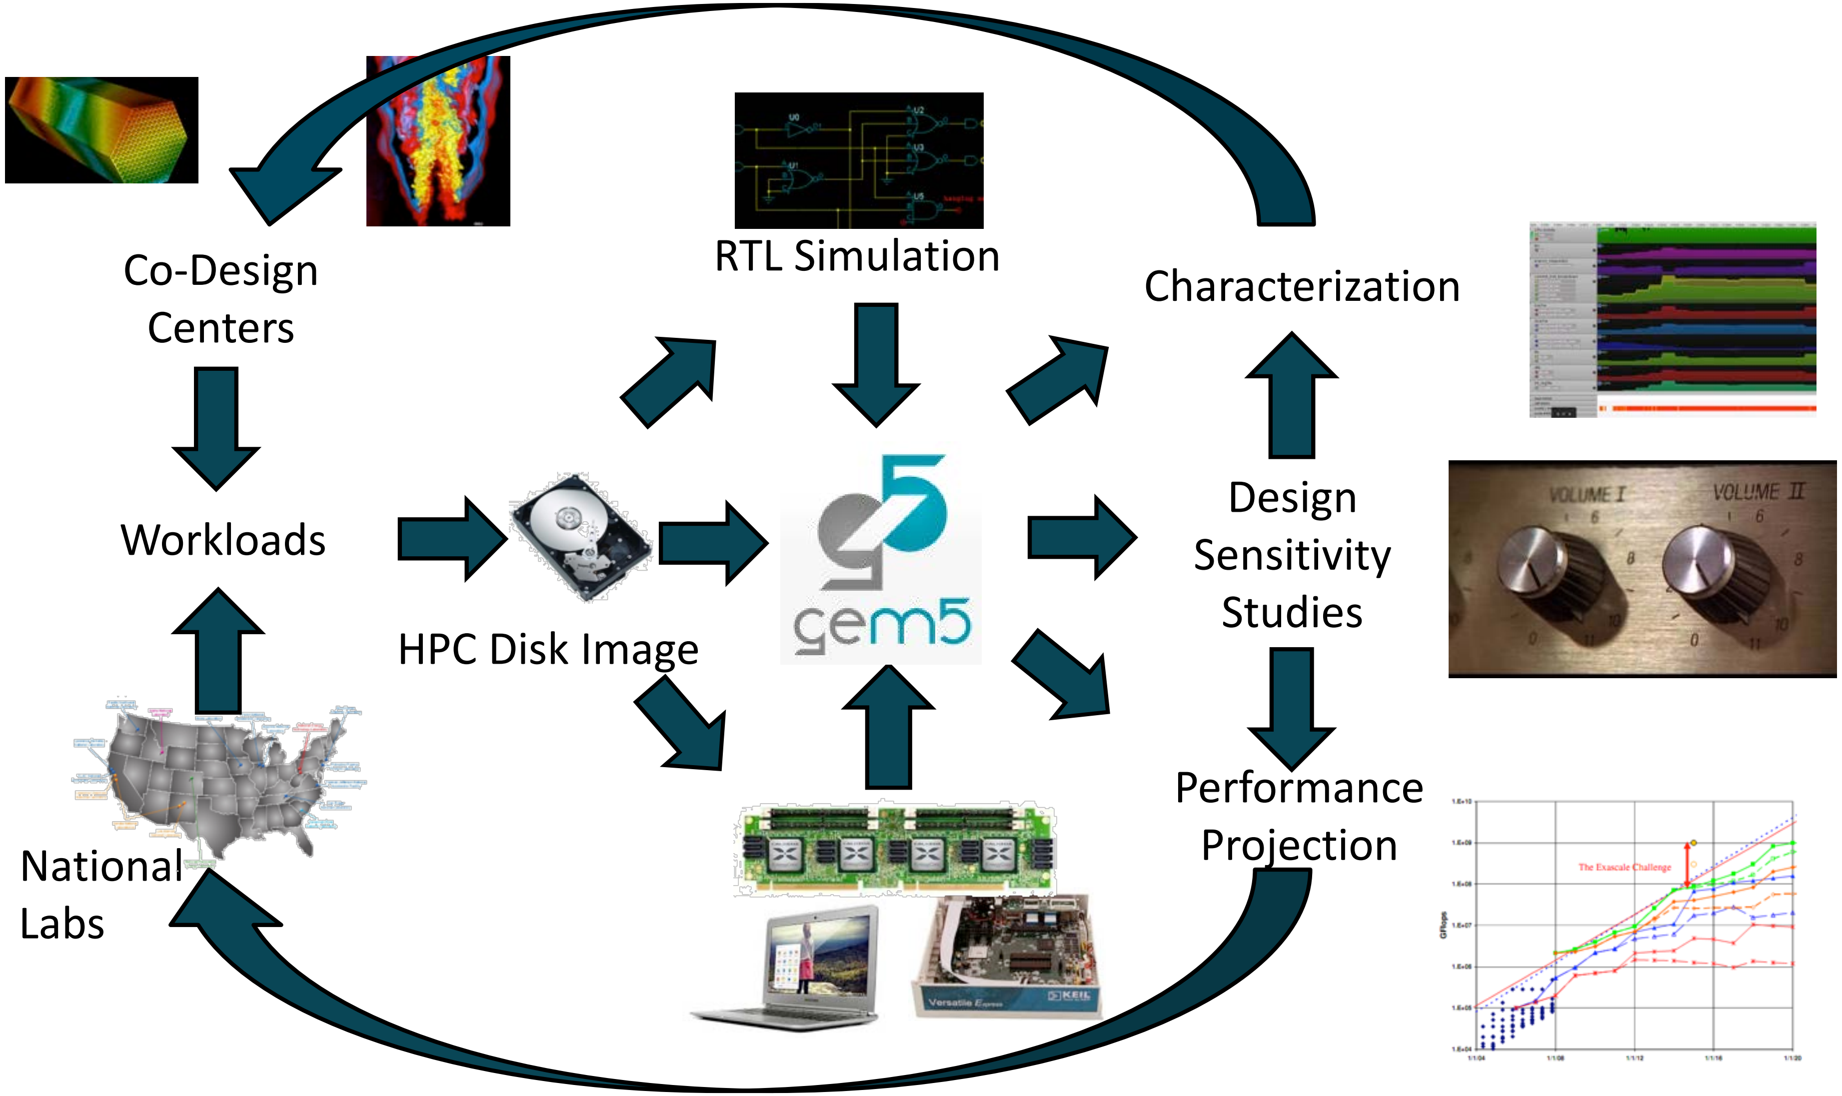

arm

Montblanc: Other Mobile Testbed Platforms

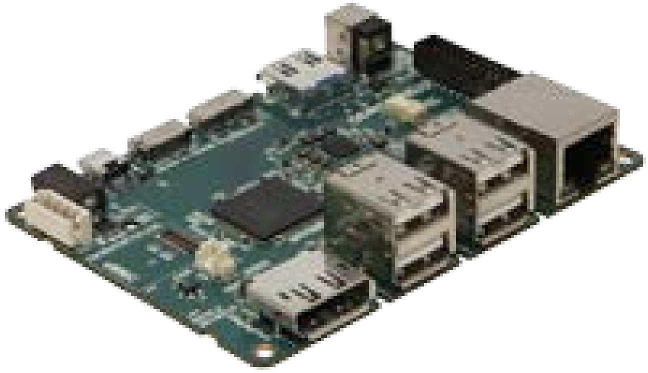

Arm big.LITTLE performance evaluation (mini-apps)

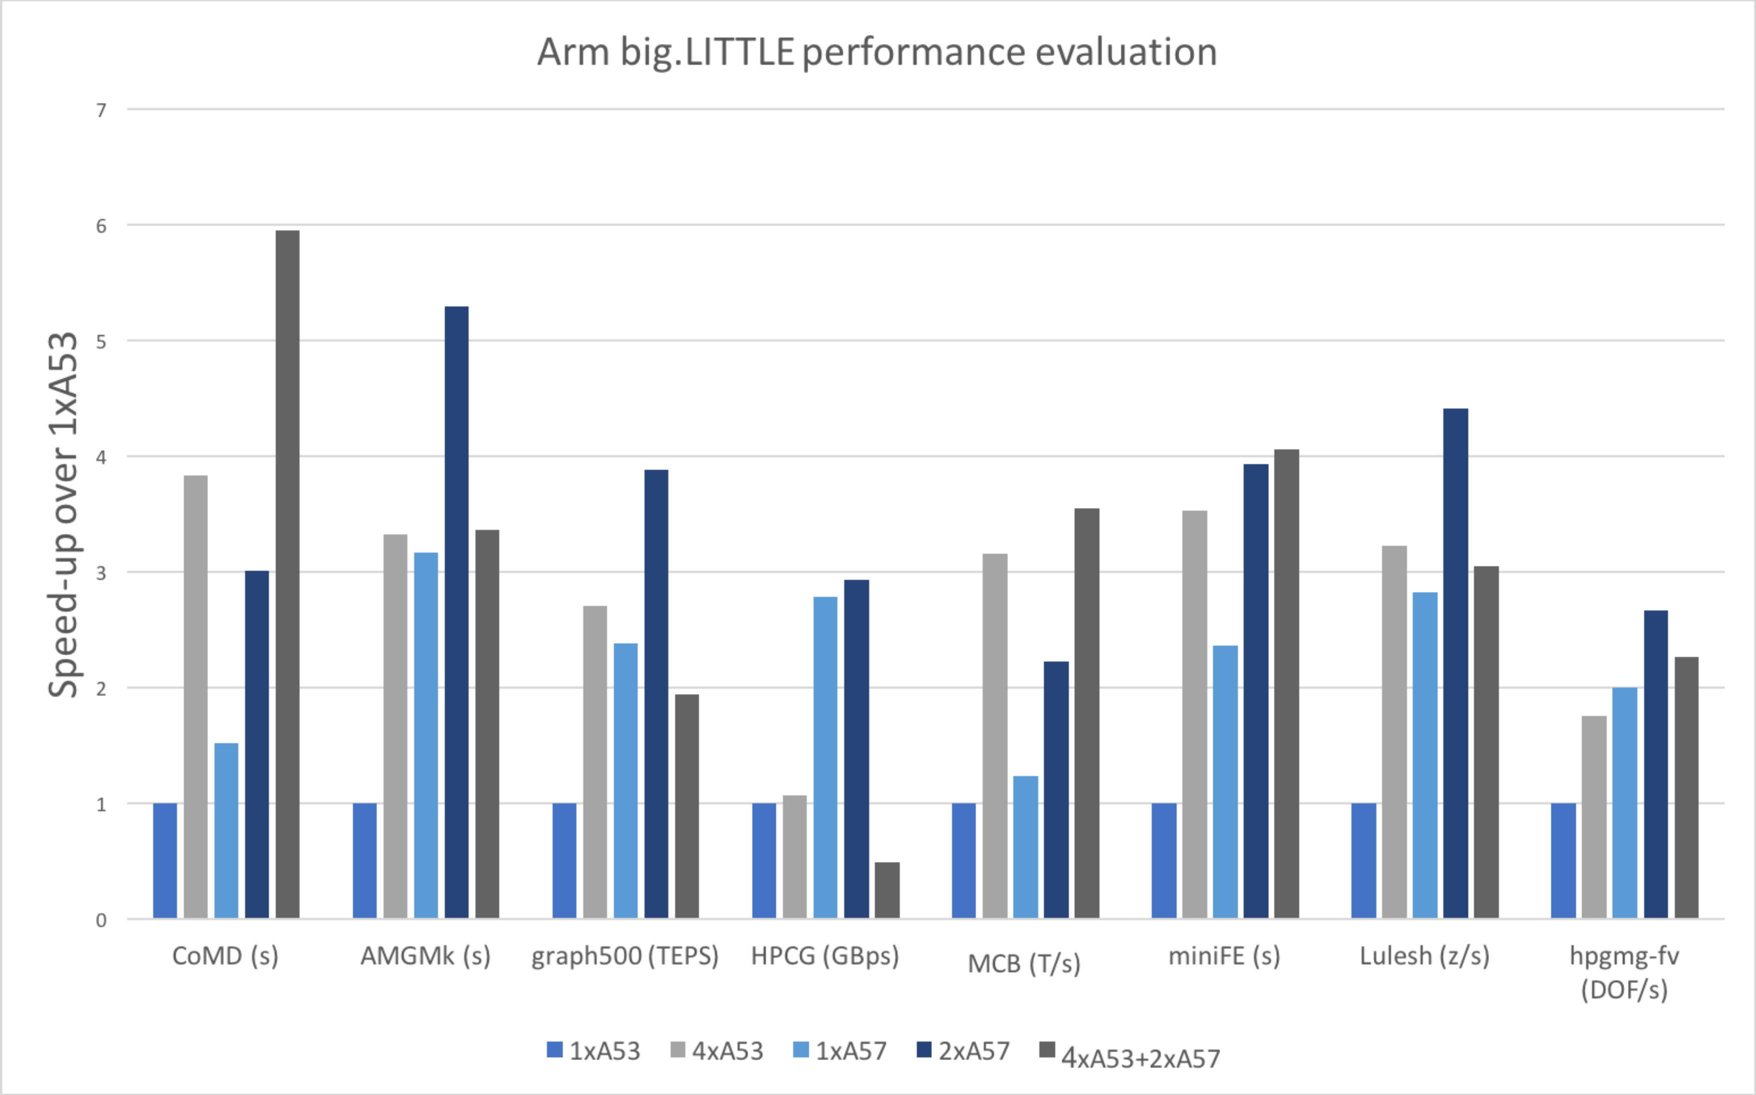

64-bit Arm Servers

APM XGENE-1

X-Gene1 (Storm)

- First generation ARMv8 CPU
- 8 Cores 2.4GHz
- 8MB L3 cache
- 4 DDR memory channels
- PCIE Gen3
- SATA Gen3
- 1/10G interfaces
- Integrated NIC
- 40nm TSMC
- System management
- Security, networking and scale-out accelerators

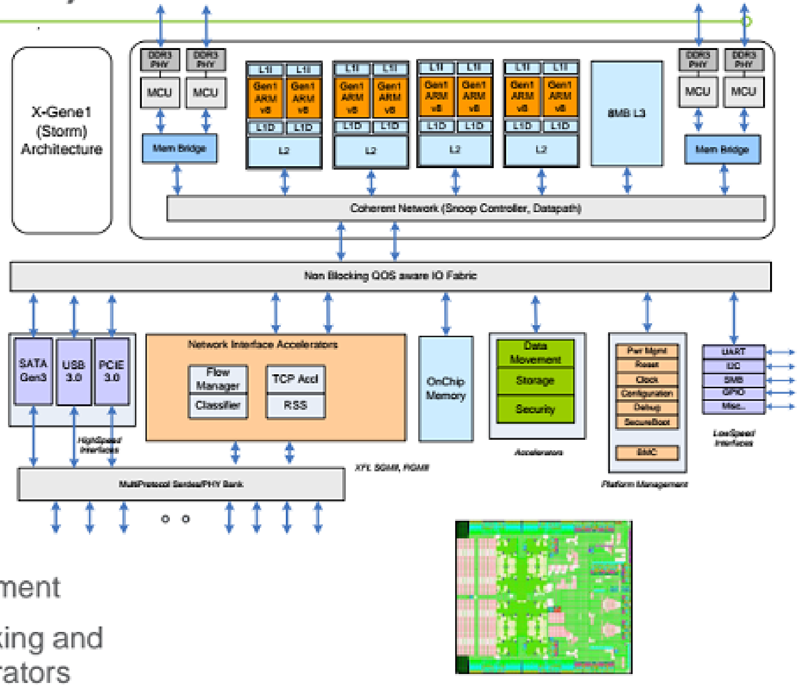

Cavium Thunder X1

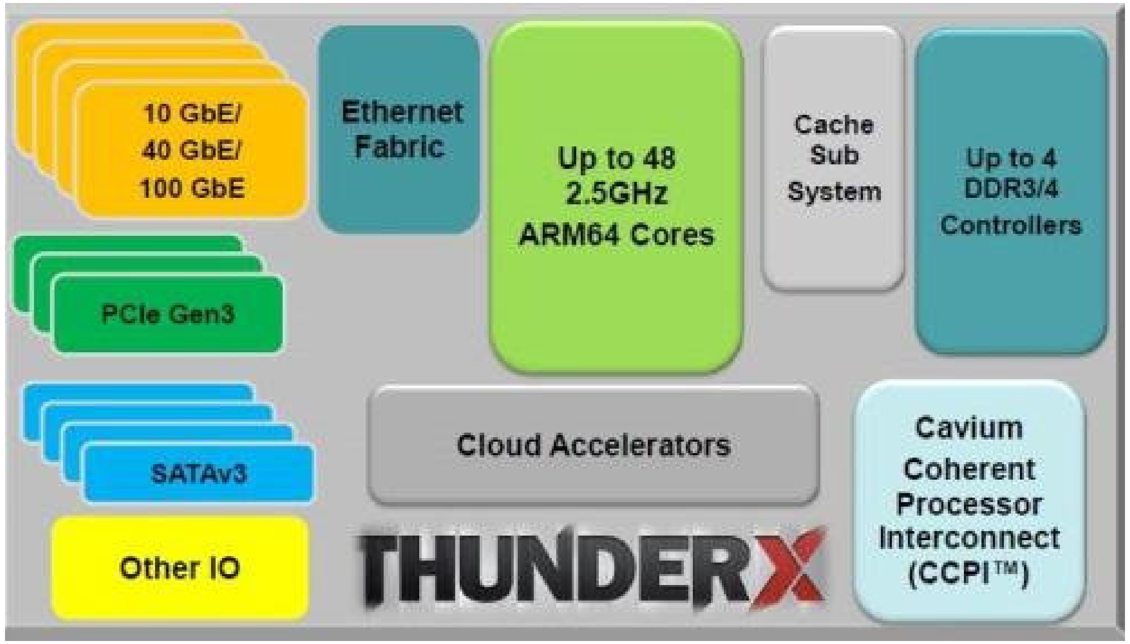

arm

2014-2016 Arm Initial Server Hardware

Opportunities

- 64-bit hardware operated at high frequencies, increasing single thread performance and overall FLOP count
- Server class hardware provided reasonable memory capacities and peripherals
- High core counts on TX1 proved effective at data centric workloads

Challenges

- Driver support sorely lacking for high performance networking and GPUs
- Single thread performance still lacking, ultimately limiting effectiveness of platforms
- Larger core-counts revealed interesting behavior within Linux and poor NUMA configuration support in firmware
- Commercial availability of hardware led to calls for commercial supported software stacks and OS distributions
- FLOP support lacking (Cavium ThunderX1 and X-Gene not initially targeting HPC)

June 2015: Arm announces Arm Performance Library

Enable the wide variety of ARM cores available today without adding complexity to the software ecosystem.

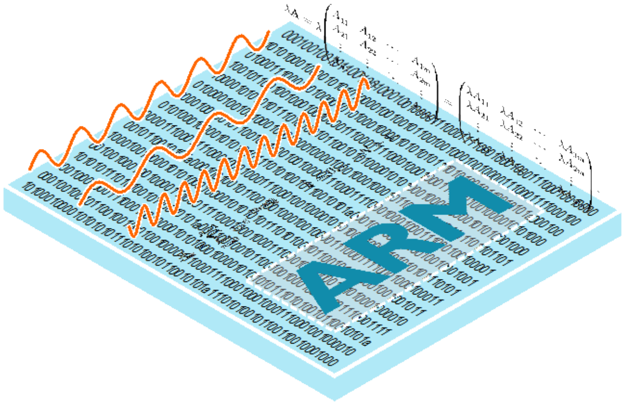

- Commercially supported 64-bit ARMv8 vendor math libraries for scientific computing.
- Built and validated using technology from the Numerical Algorithms Group (NAG).
- ARM silicon partners provide us with tuned kernels.

Capabilities:
➔ **BLAS**
➔ **LAPACK**
➔ **FFT**

Tuned for:
- Cortex-A53, A57, A72
- Applied Micro X-Gene®
- Cavium® ThunderX

arm

2016: Scalable Vector Extension

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

arm

June 2016: Japan Announces Arm based Post-K

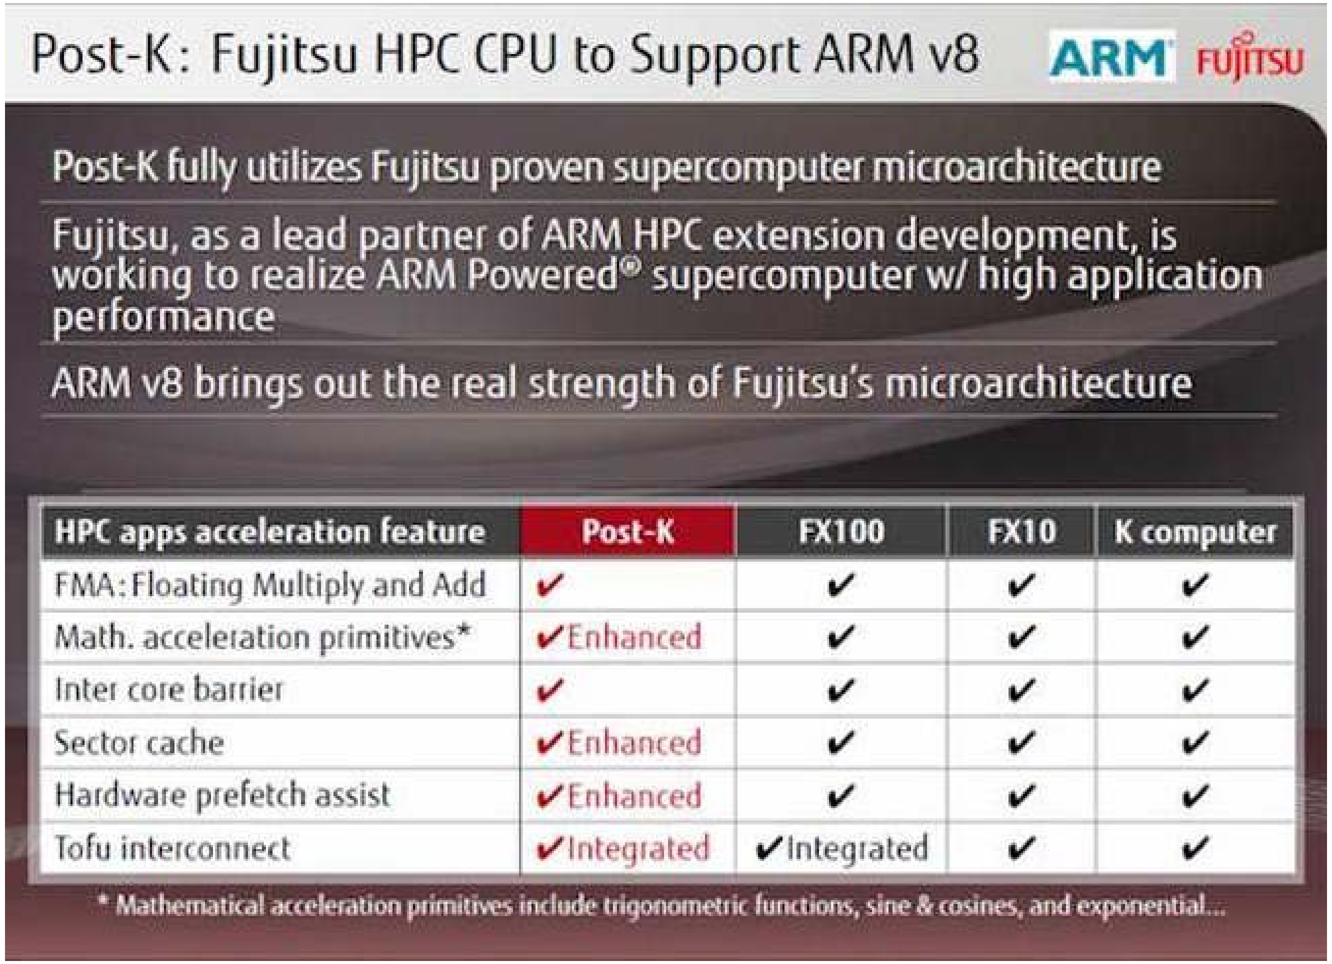

Post-K: Fujitsu HPC CPU to Support ARM v8

ARM FUJITSU

Post-K fully utilizes Fujitsu proven supercomputer microarchitecture

Fujitsu, as a lead partner of ARM HPC extension development, is working to realize ARM Powered® supercomputer w/ high application performance

ARM v8 brings out the real strength of Fujitsu's microarchitecture

HPC apps acceleration feature	Post-K	FX100	FX10	K computer
FMA: Floating Multiply and Add	✔	✔	✔	✔
Math. acceleration primitives*	✔Enhanced	✔	✔	✔
Inter core barrier	✔	✔	✔	✔
Sector cache	✔Enhanced	✔	✔	✔
Hardware prefetch assist	✔Enhanced	✔	✔	✔
Tofu interconnect	✔Integrated	✔Integrated	✔	✔

* Mathematical acceleration primitives include trigonometric functions, sine & cosines, and exponential...

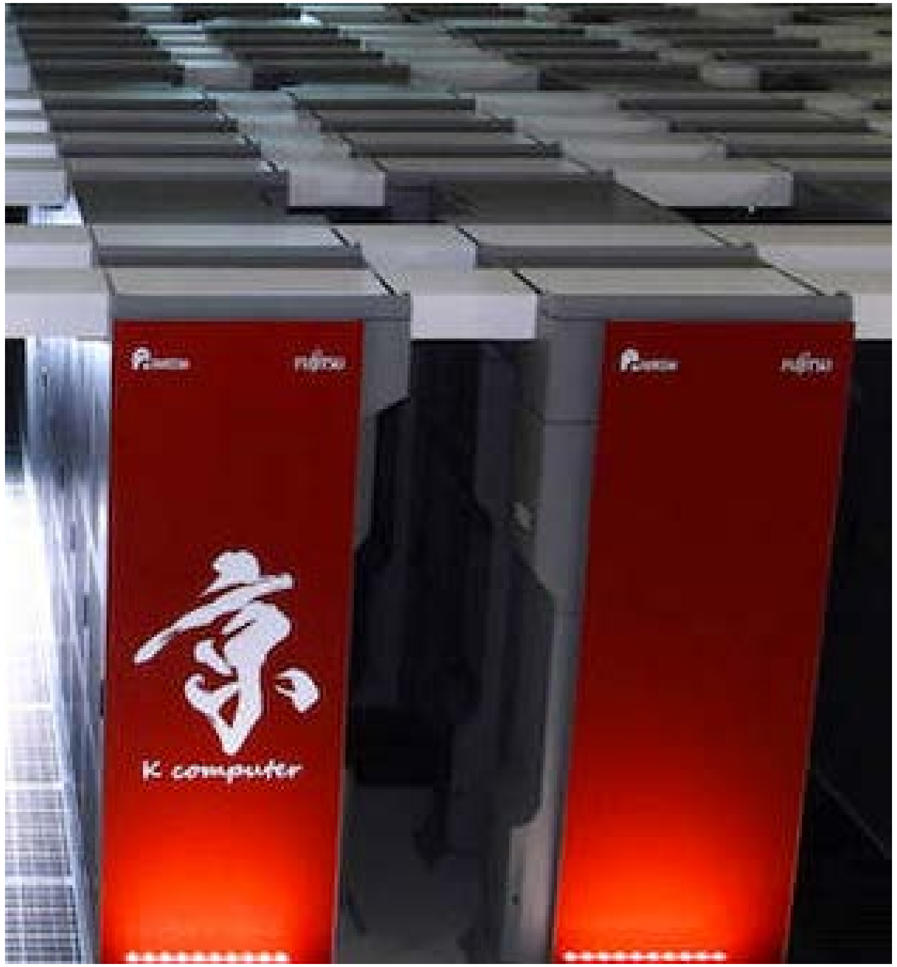

arm

openHPC – 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
- Packages built on Armv8-A for CentOS and SLES
- Arm-based machines in the OpenHPC build infrastructure
- Technical Preview in 2016, Full Release in 2017

Functional Areas	Components include
Base OS	CentOS, SLES
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

September 2016: Linaro Starts HPC-SIG

HPC Special Interest Group

The Linaro HPC SIG is a collaboration building on the work from open source projects

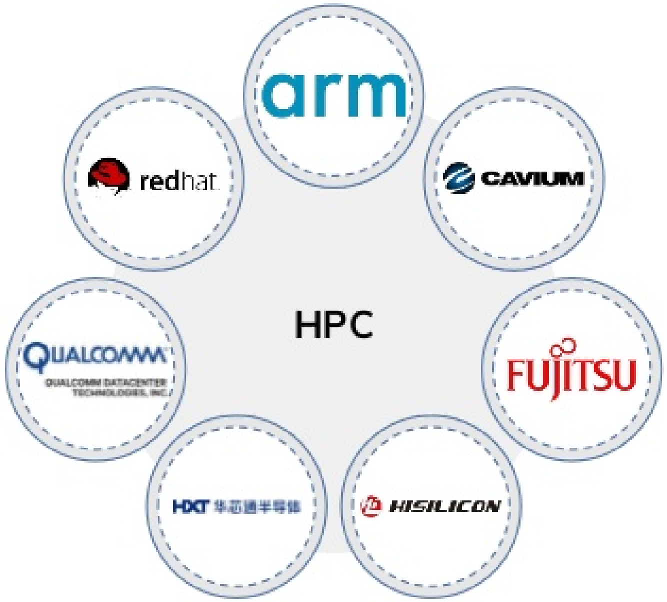

Driving enterprise-class, open-source HPC development on Arm

Identify and adopt standards to make HPC deployment on Arm a commercial imperative. Develop real-world use cases that reap the benefits of Arm while ensuring interoperability, modularization, orchestration

Lower deployment & management barriers

Leverage the Linaro Developer Cloud and other services to rapidly develop cost-effective Cloud-based HPC development frameworks and generate reference implementations for commercial prototyping and deployment

Enable the data-driven economy

Machine learning and Deep learning are both critical to the future of HPC, specifically as the path toward exascale computing. Driving engineering in HPDA and Machine learning algorithms will aid the success of organizations to fully capitalize on these technologies

Member-driven with Advisory Board

Members determine work to be completed by engineering resources while the advisory board provides subject matter expertise on HPC requirements and guidance on the ongoing HPC SIG strategic direction and roadmap

HPC

Linaro Background

Technical Steering Committee (TSC)

Office of the CTO (OCTO)

Working Groups

- Kernel Consolidation
- Toolchain
- Graphics
- Power Management

Segment Groups

- LEG (15 companies)
- LNG (12 companies)

Optimization, consolidation and innovation

Platform Engineering

- LAVA
  Validation & Benchmarking
- **Builds & Baselines**
  OE/Yocto, Android, Fedora, Ubuntu
- QA
  Continuous Integration for Linaro output

Validation, benchmarking and release management

Landing Teams
working under NDA

ARM
BROADCOM
FUJITSU
HISILICON
LG Electronics
ST ERICSSON
SAMSUNG

SoC support and optimization

arm

2017: AToS, HPE and Cray Announce Products

AToS Sequana (ISC 2017)

HPE Apollo 70 (SC 2017)

Cray XC50 (SC 2017)

arm

Cavium CN99XX - 1st member of THUNDERX2 Family

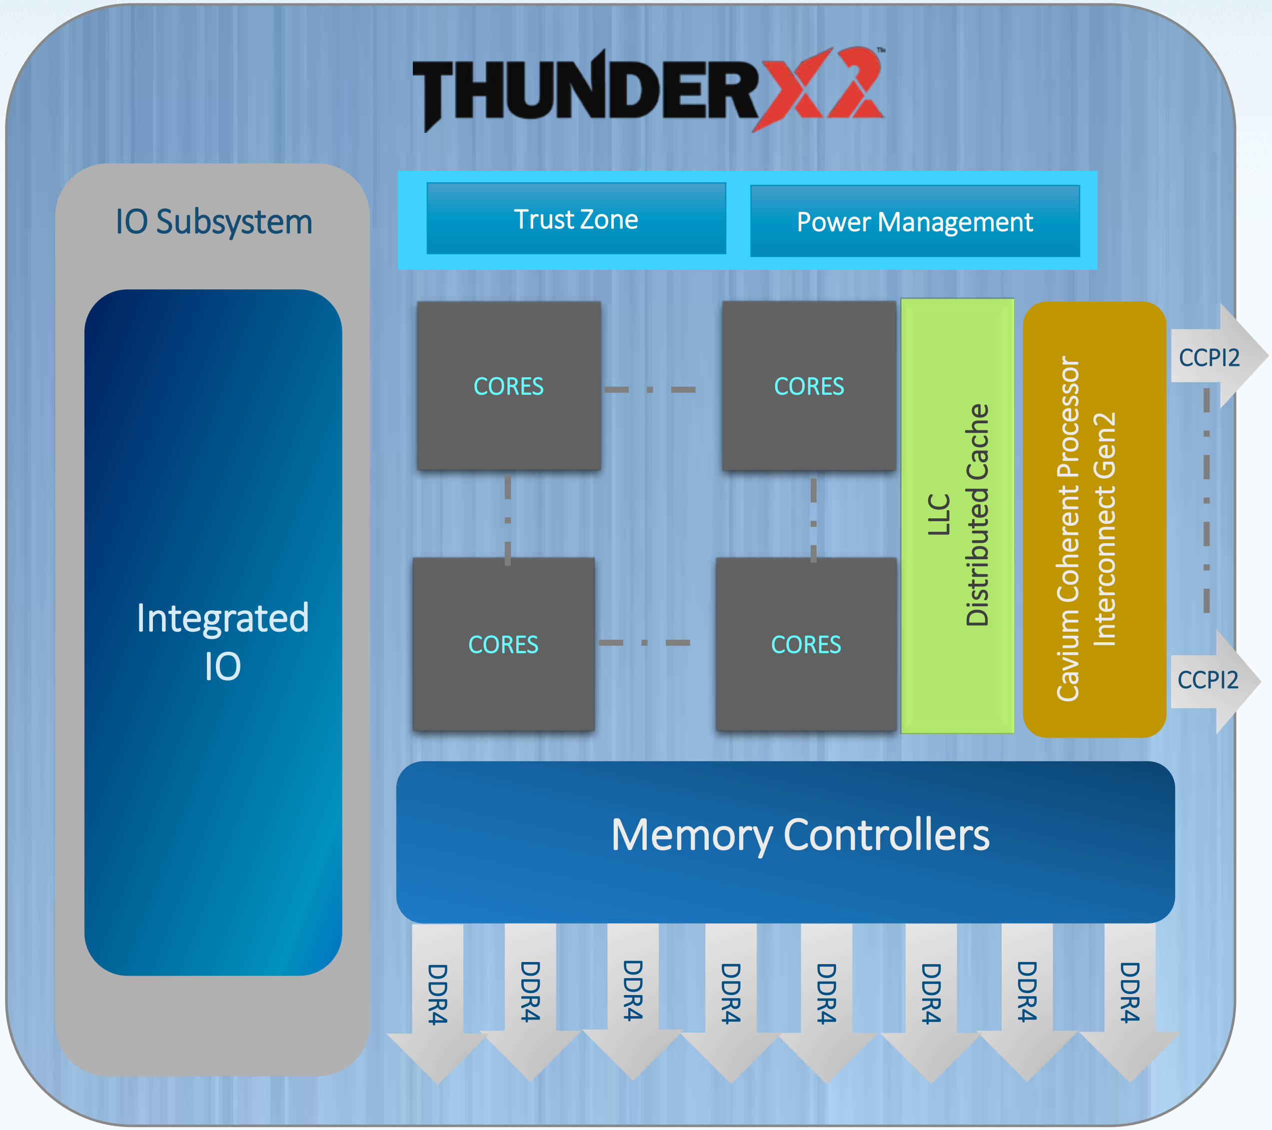

- 24/28/32 Custom ARMv8 cores
- Fully Out-Of-Order (OOO) Execution
- 1S and 2S Configuration
- Up to 8 DDR4 Memory Controllers
- Up to 16 DIMMs per Socket
- Server Class RAS features
- Server class virtualization
- Integrated IOs
- Extensive Power Management

2nd gen ARM server SoC

Delivers 2-3X higher performance

CAVIUM

Headline results from GW4

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 @ >2GHz per node
- Cray software tools
- Technology comparison:
  - x86, Xeon Phi, Pascal GPUs
- Phase 1 installed March 2017
- The Arm part arrives early 2018

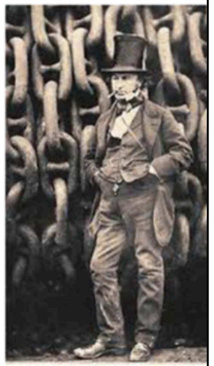

I.K.Brunel 1804-18

@simonmcs http://gw4.ac.uk/isambard/ 5 bristol.ac

University of BRISTOL | Isambard | GW4

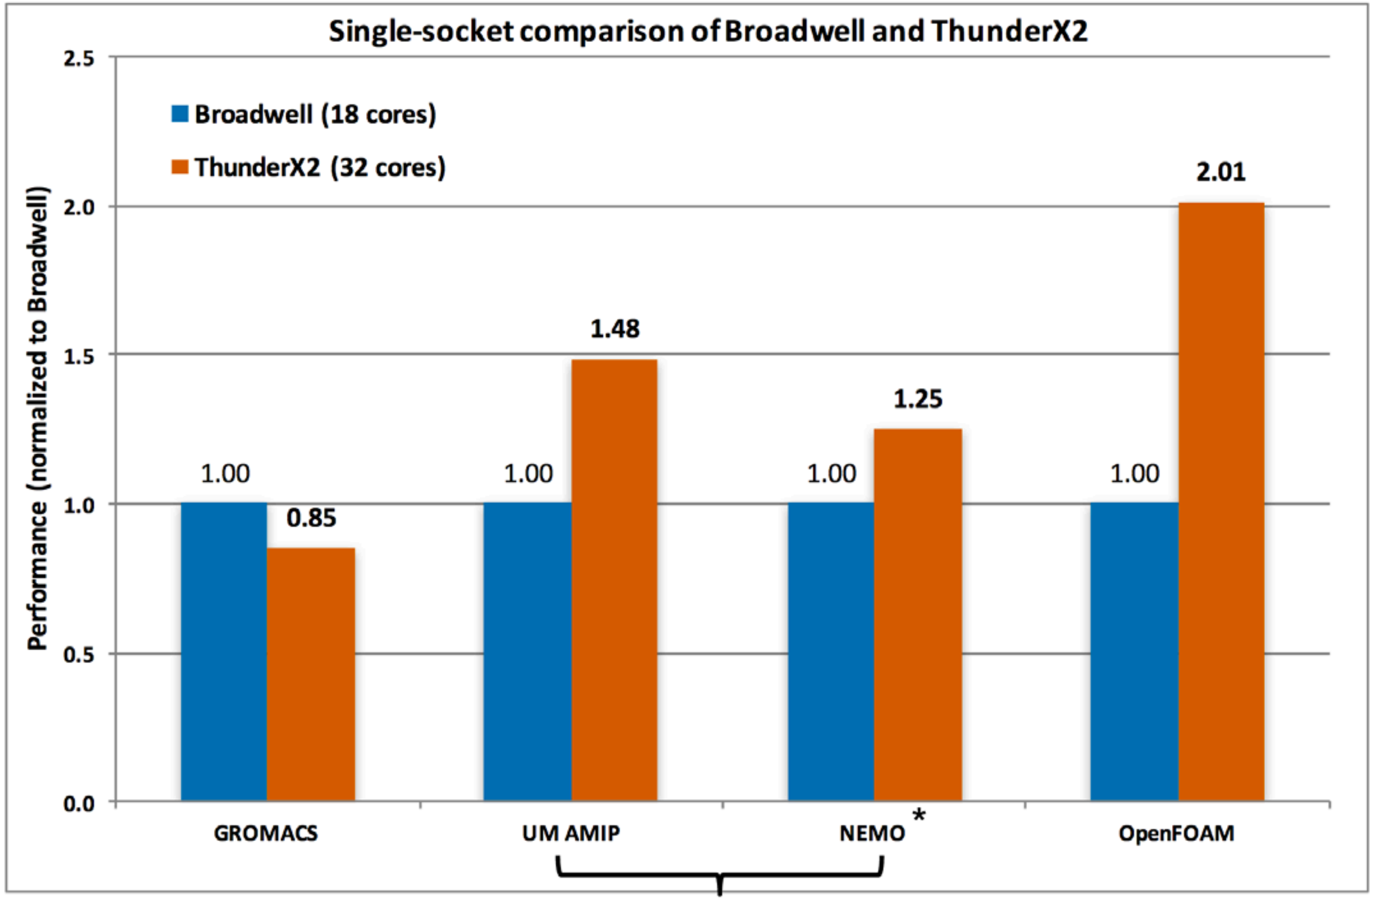

Benchmarked by the UK's Met Office

@simonmcs http://gw4.ac.uk/isambard/ bristol.ac.uk

* = NEMO runs from a 28 core, 2.0GHz TX2

Dibona: ThunderX2 based system

The Mont-Blanc 3 demonstrator

→ **Codename: "Dibona"**

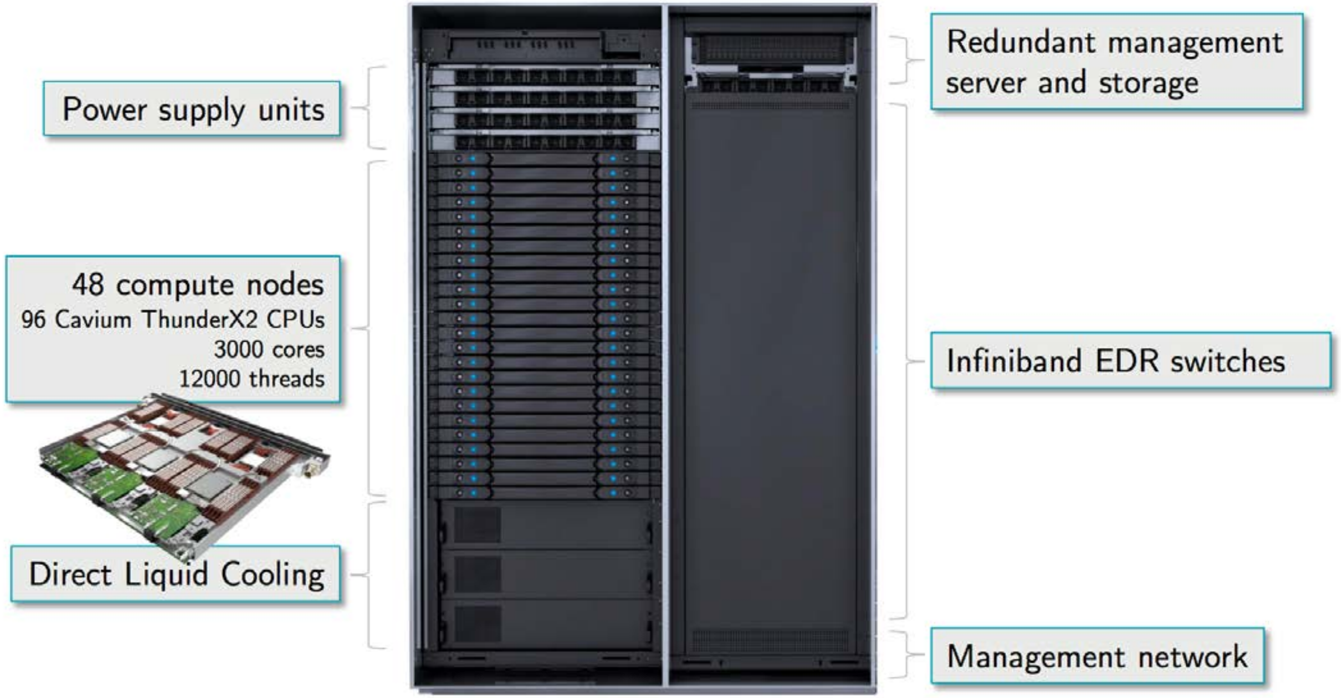

7 GoingARM workshop Frankfurt, June 22nd, 2017 MONT-BLANC

arm

June 2017: Arm Announces Flang Support

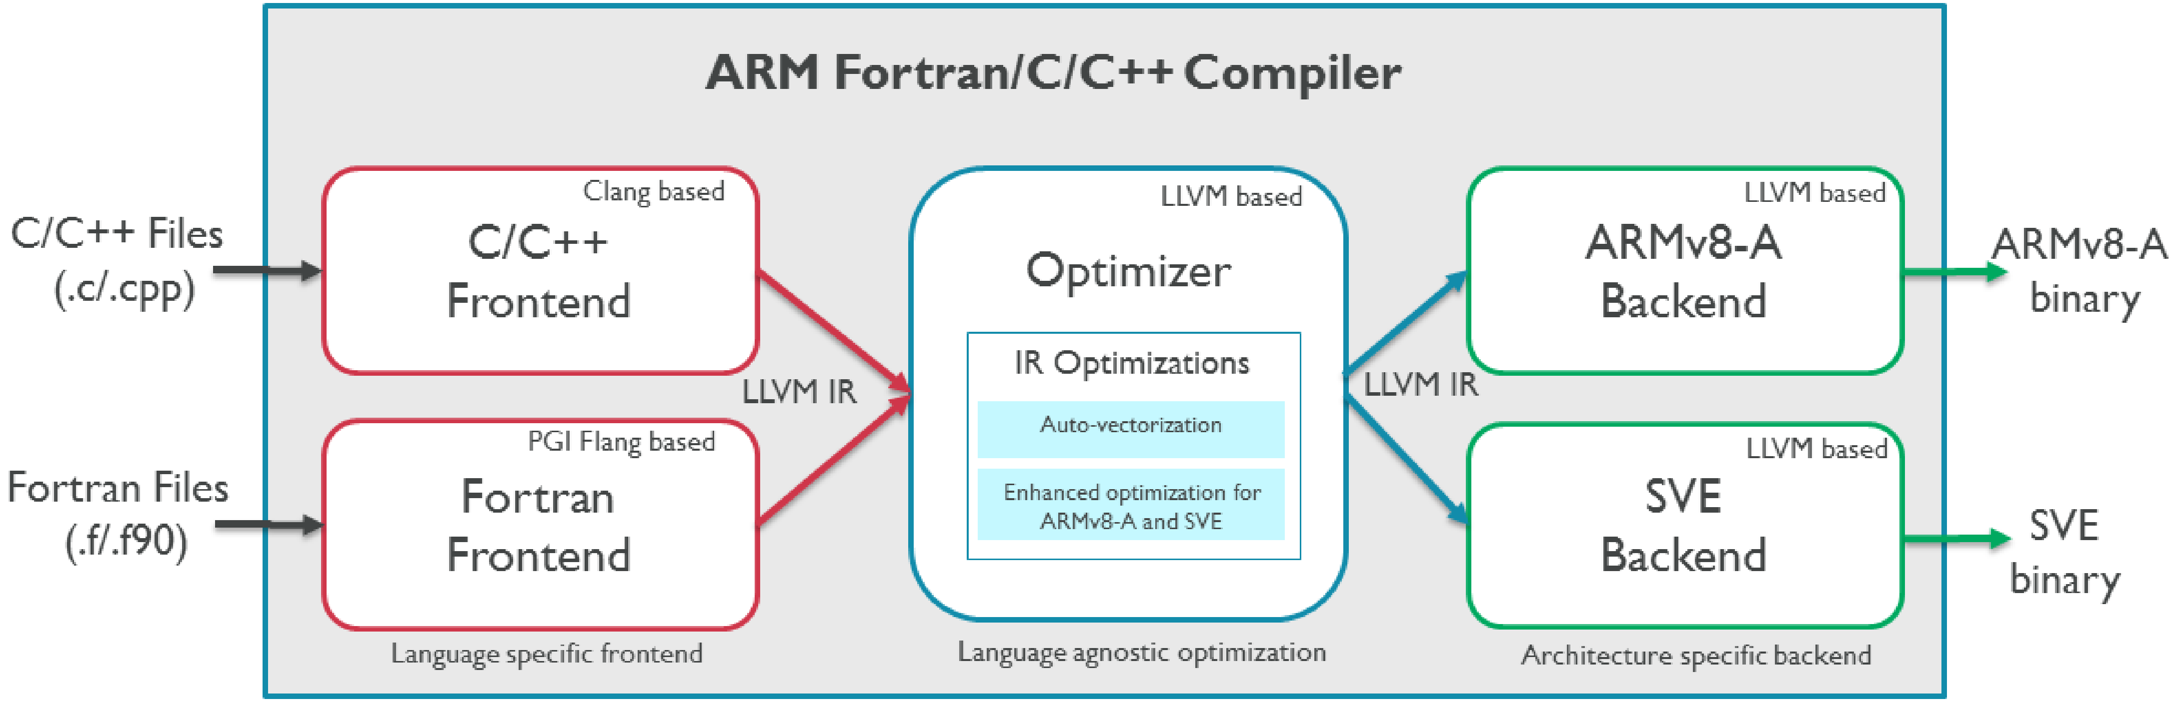

2017: Arm HPC ecosystem

Porting to Arm

Arm is engaging directly with partners and HPC scientific code developers to support porting and optimisation of common HPC libraries, tools and applications.

Initial focus on successfully building with both **Arm** and **GCC** compilers across a broad front.

Often only modest changes to environment variables, build scripts and architecture files are needed

Degree of commonality between codes

OpenFOAM

NAMD
Scalable Molecular Dynamics

QUANTUMESPRESSO

GROMACS
FAST. FLEXIBLE. FREE.

Geant 4

EPOCH

WRF

arm

2017: Montblanc Software Stack

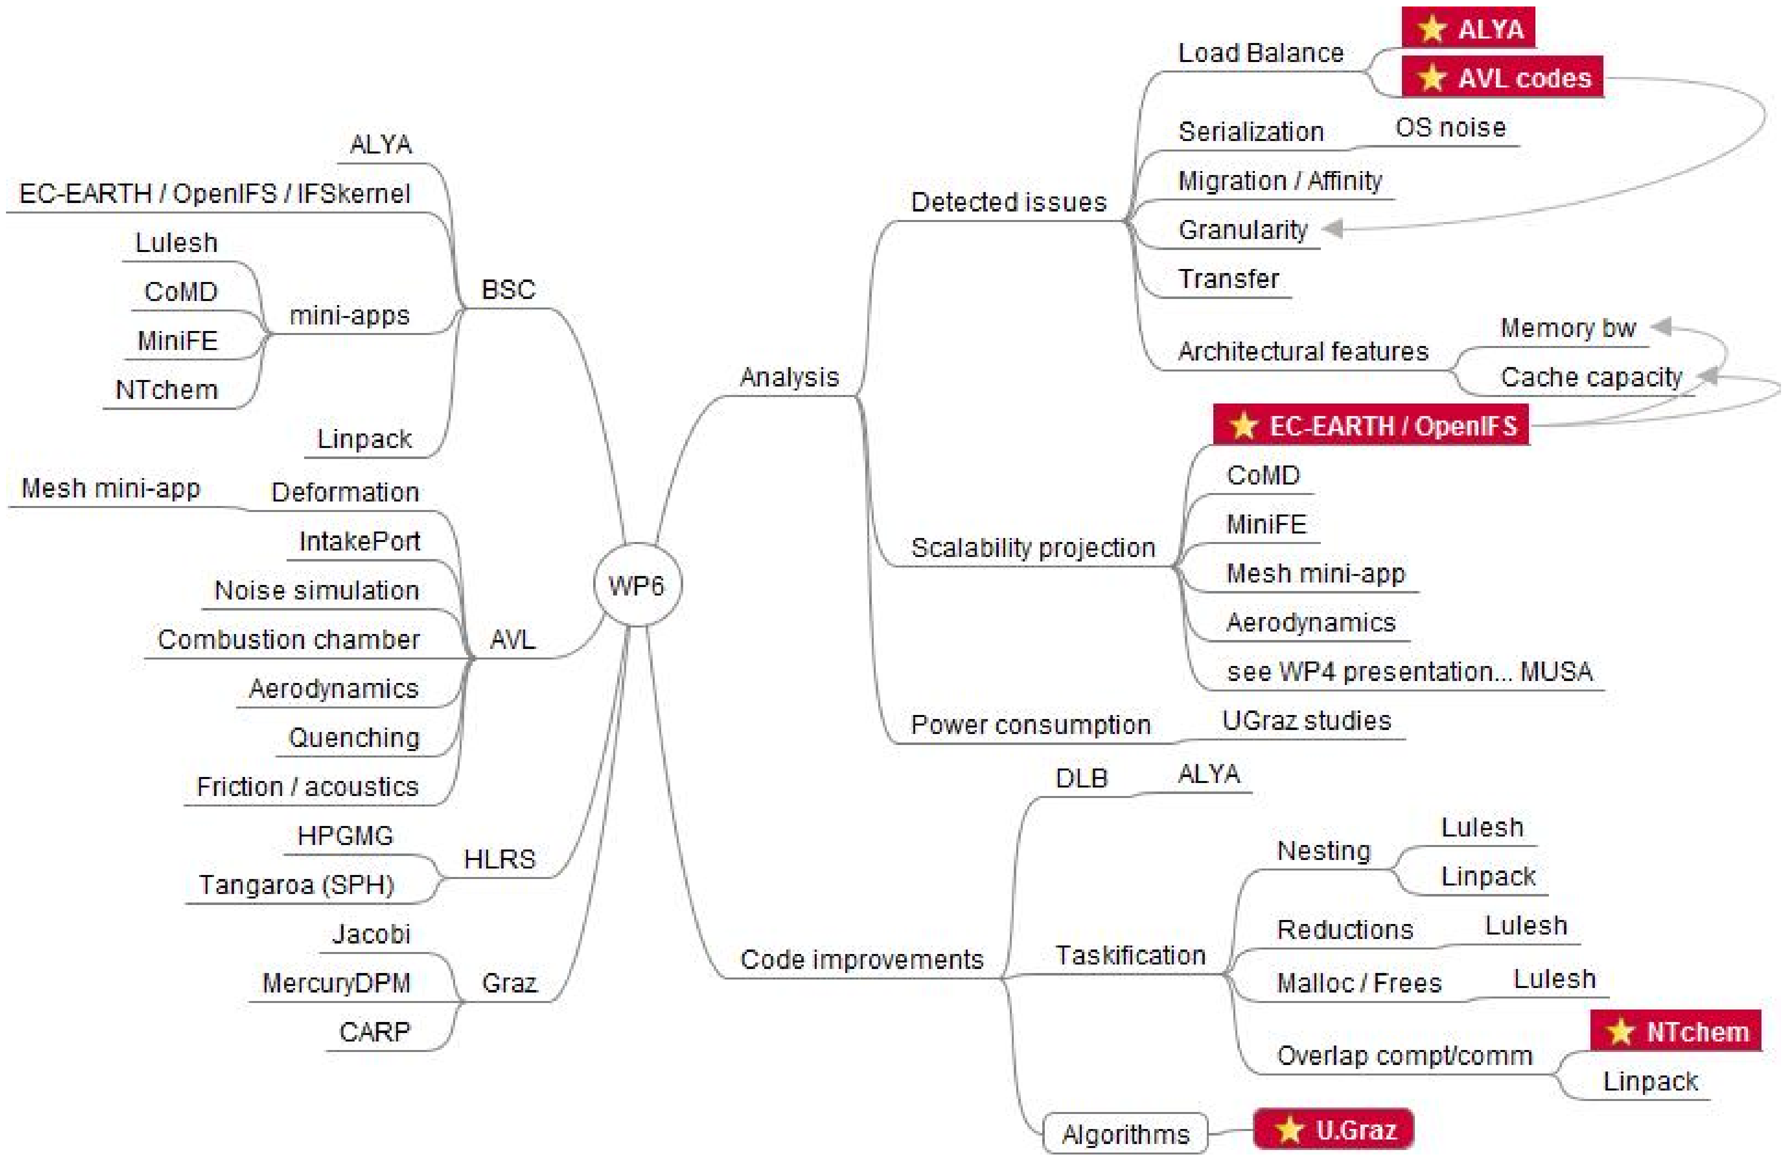

June 2017: gitlab.com/arm-hpc

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

Open Source Compiler Highlights in 2017

Performance Improvements

- GCC SPEC CPU2006:
  - INT/FP+4% comparing to GCC-7
  - Selected Glibc scalar math functions significantly optimized (wrf +60%)
  - GCC loop vectorizer enhanced (hmmer +30%)
- Glibc single-thread optimization up to 25%-150% improvement in benchmark
- LLVM: About 1% performance improvement for Armv7-A

Architecture enablement

- Continuously adding new architecture support to GCC and LLVM
  - Include Cavium ThunderX2 and Qualcomm Falkor
- Armv8.3-A released in GCC-7
- Armv8.4-A and SVE are targeted to be included in the GCC-8 release

Work in progress

- Resolving GCC vectorizer regression
- Vect-math library for Arm/AArch64, enable more vectorization opportunities
- Further loop optimizations in GCC
- LLVM GlobalISel pass enablement
- GDB Fortran enhancement

November 2017: Arm Allinea Studio

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

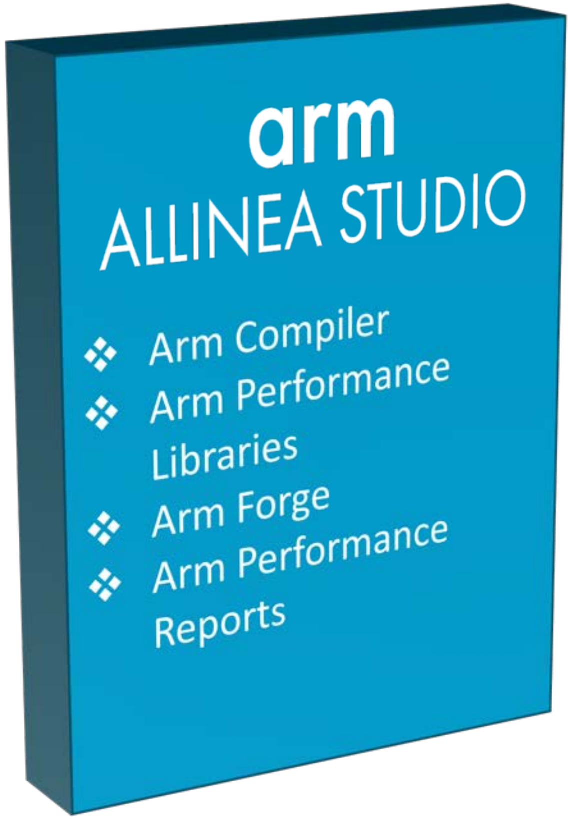

Comprehensive and integrated tool suite for Scientific computing, HPC and Enterprise developers

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

2017-2018 2nd Generation Servers

Opportunities

- Arrival of ThunderX2 and Qualcomm Centriq provided ample memory capacity and world class memory bandwidth
- Second generation servers had better PCI and more driver support
- ThunderX2 provided high single thread performance and FLOP rates
- Commercial support for software and OS distros readily available

Challenges

- Consolidation within the silicon market pressuring Arm server chip companies
- Limited initial hardware availability made porting of non-open-source code difficult
- Facing chicken and egg problem with ISV community
- Vector widths still limited to 128-bit NEON, SVE support coming in 3rd generation

arm

Momentum Timeline of commercial HPC tools from Arm

Continued commitment to high performance computing

SC15
- Launch of **Performance Libraries -** Optimized BLAS, LAPACK and FFT

SC16
- Launch of **C/C++ Compiler -** Linux user space compiler

2016 Dec
- Acquisition of **Allinea,** the leading HPC tools provider

ISC17
- Beta availability of **Fortran Compiler**

SC17
- Launch of **Arm Allinea Studio**
- A comprehensive and integrated tools suite for Scientific Computing, HPC and Enterprise

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.

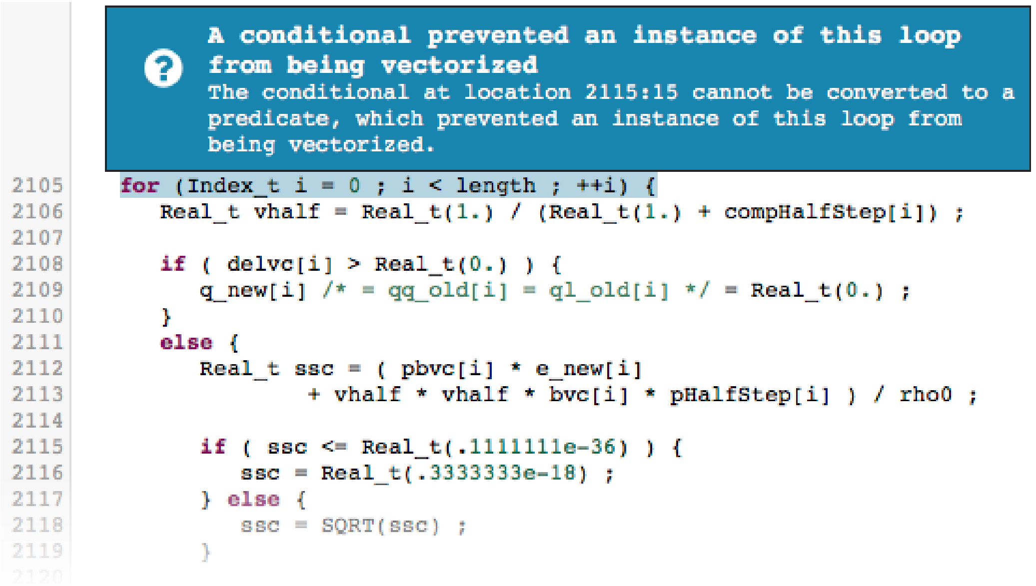

Challenge: Core Count Versus Memory Bandwidth

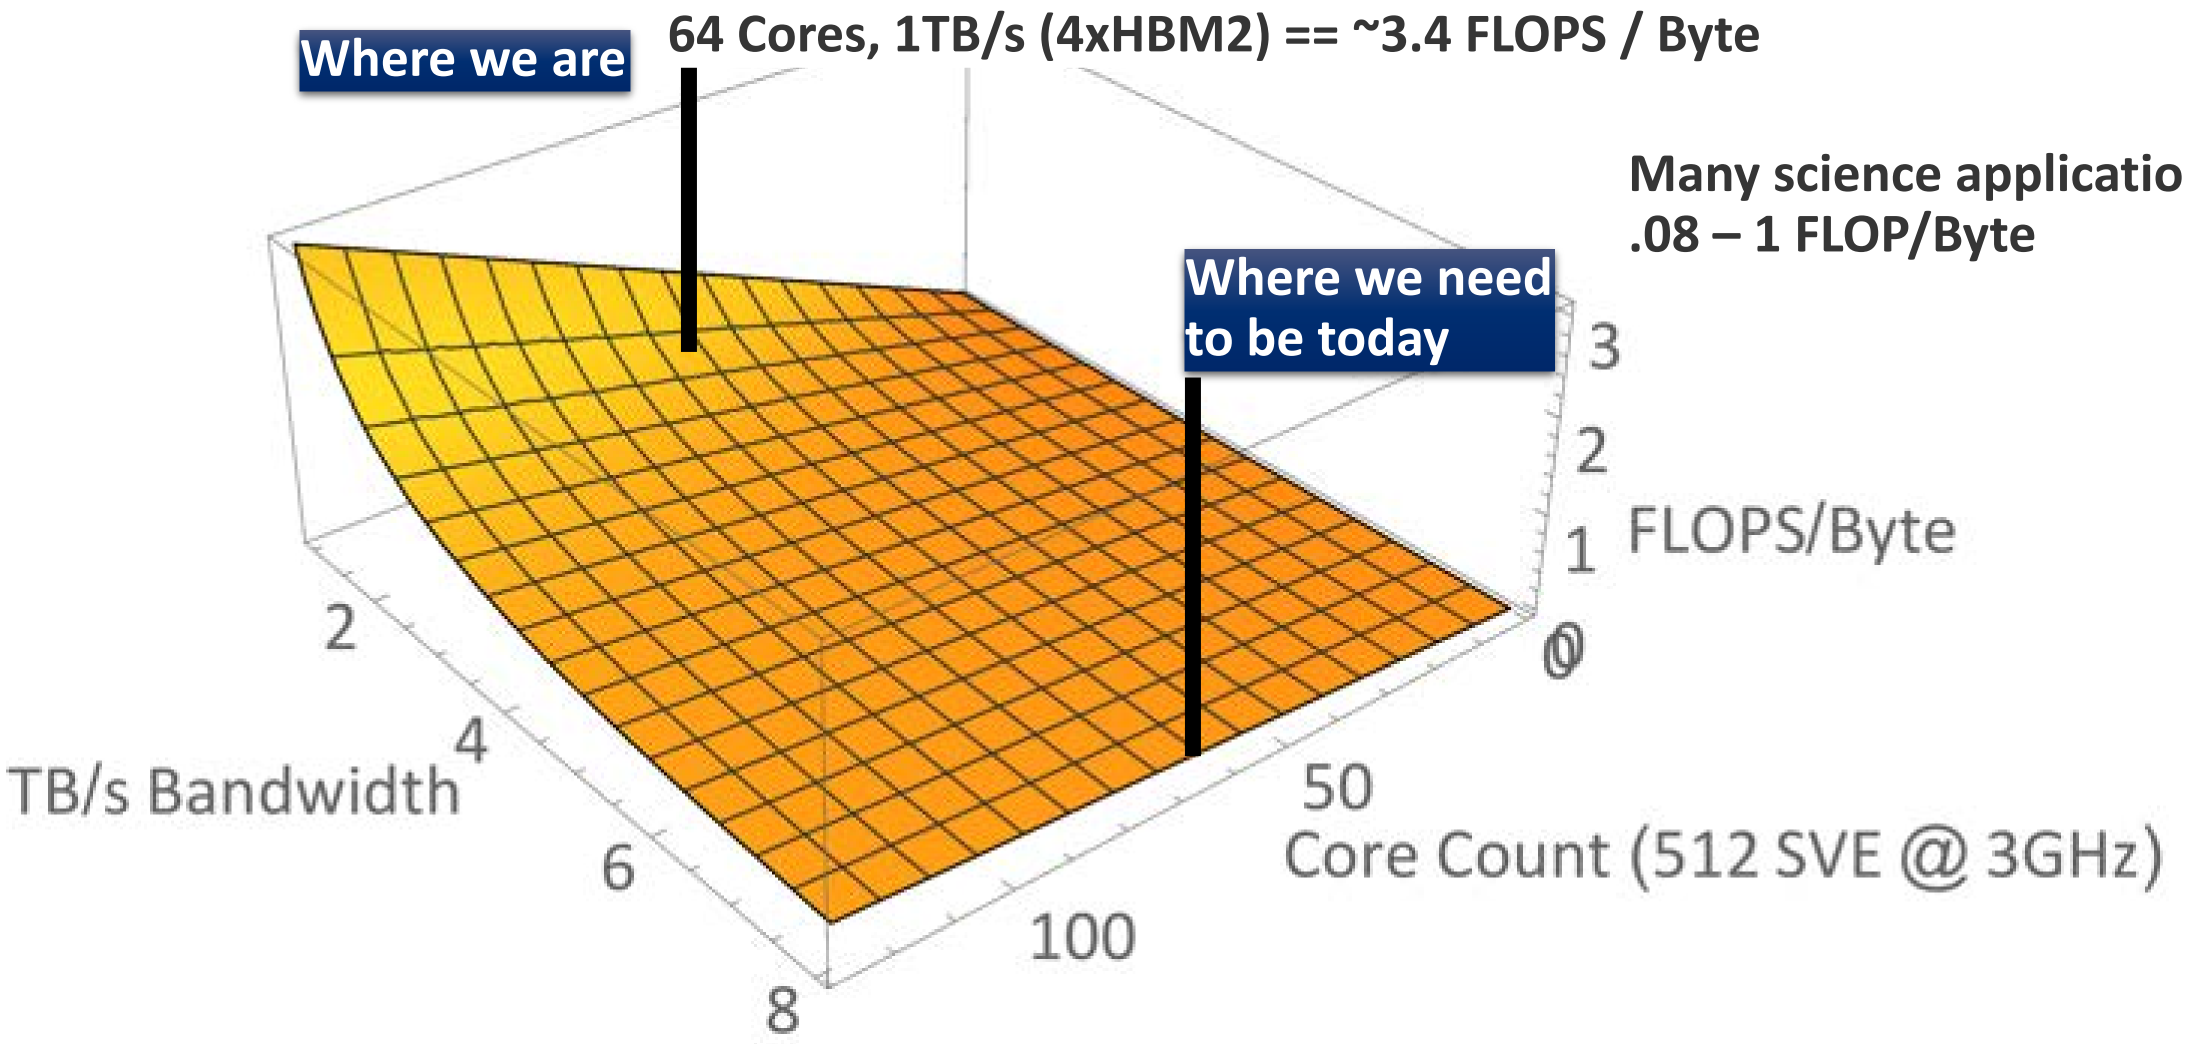

Many science applications
.08 – 1 FLOP/Byte

arm

Research: Gather-Scatter & Data Reorg Near-Memory

The Sparse Data Reduction Engine (SPiDRE)

Standard Gather Scatter

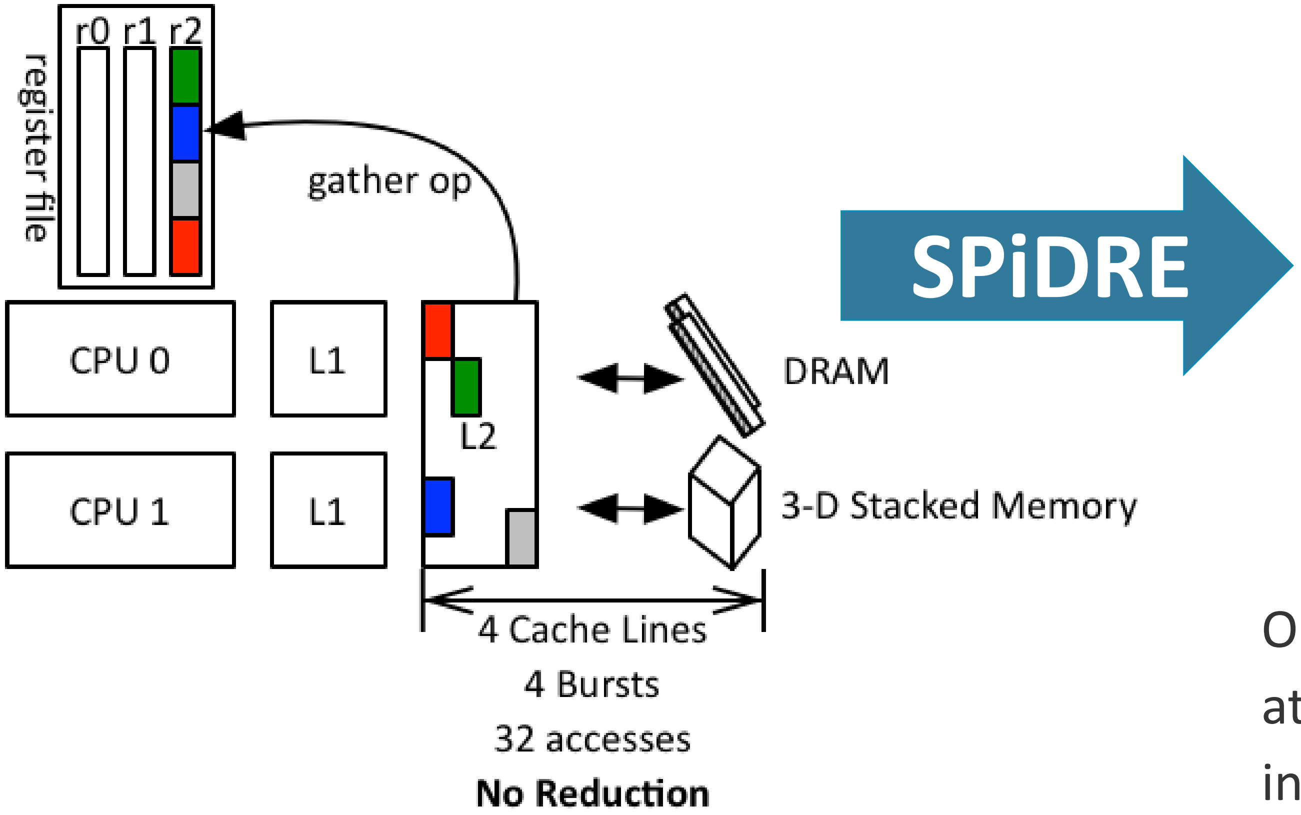

Near-Memory Gather-Scatter

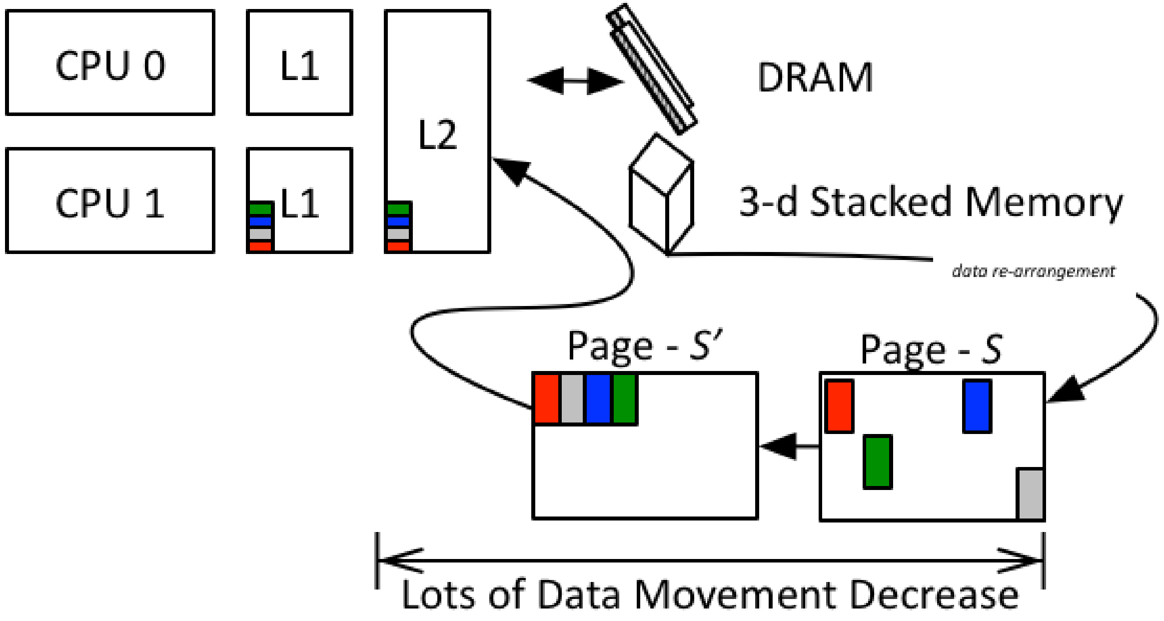

Opportunity for increased throughput
at memory controller via higher bandwidth
interface to memory than driven by cores

arm

Challenge: On-SoC Scaling & Accelerator Coupling

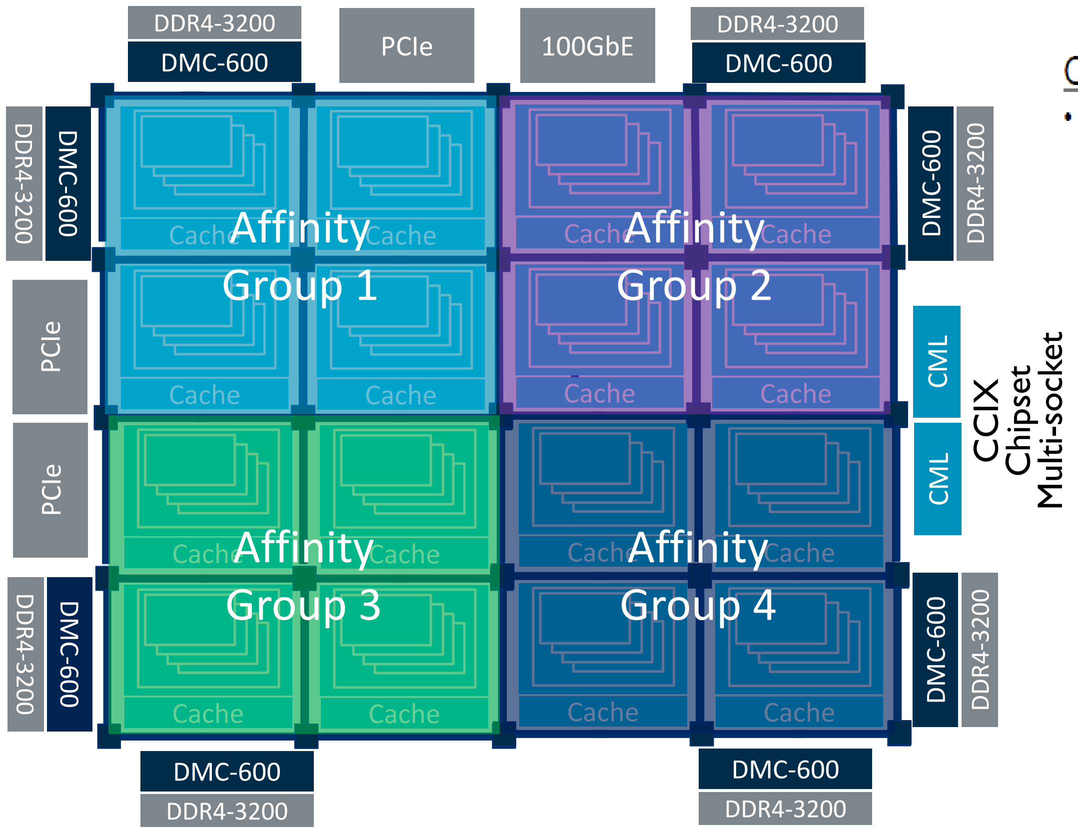

Create affinity working groups, aka on-chip NUMA

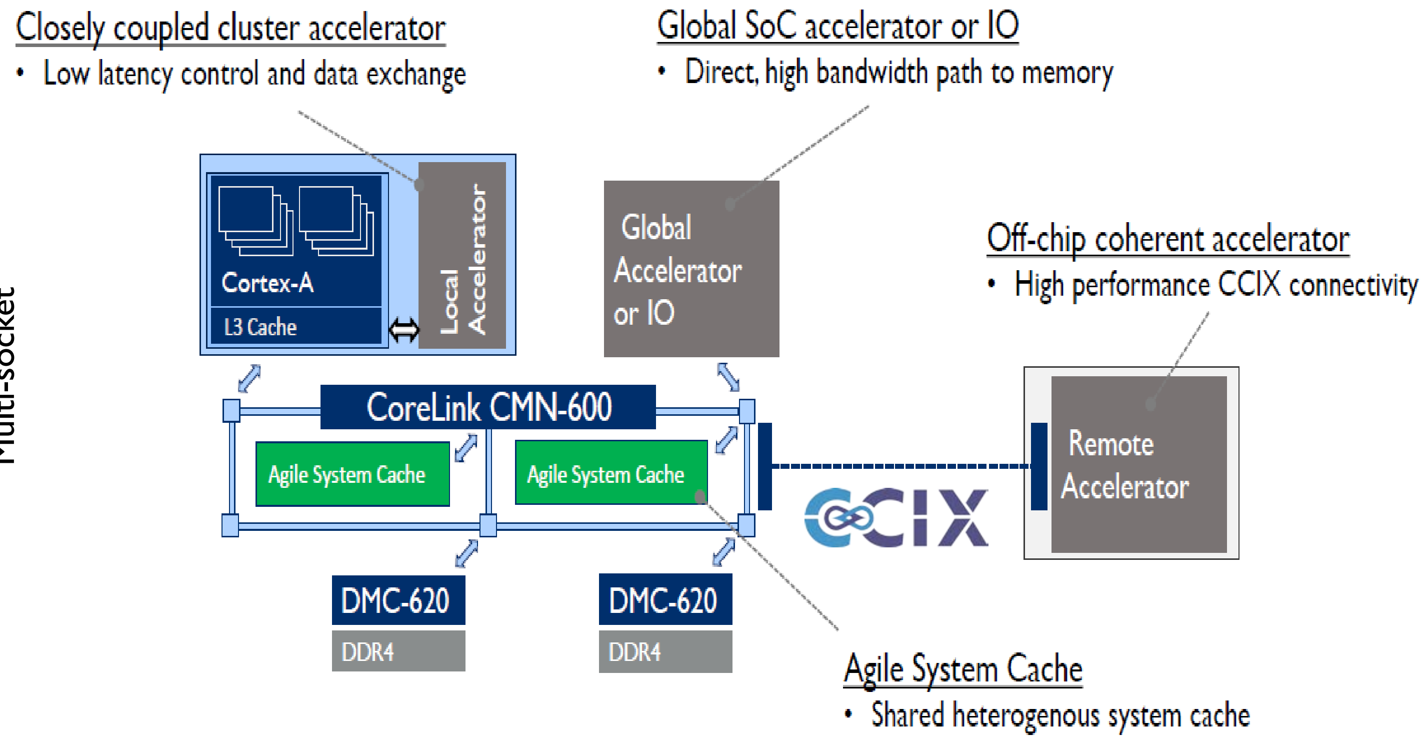

Flexible coherent acceleration options (local, global, disaggregated)

arm

Hardware Software Stack Overview

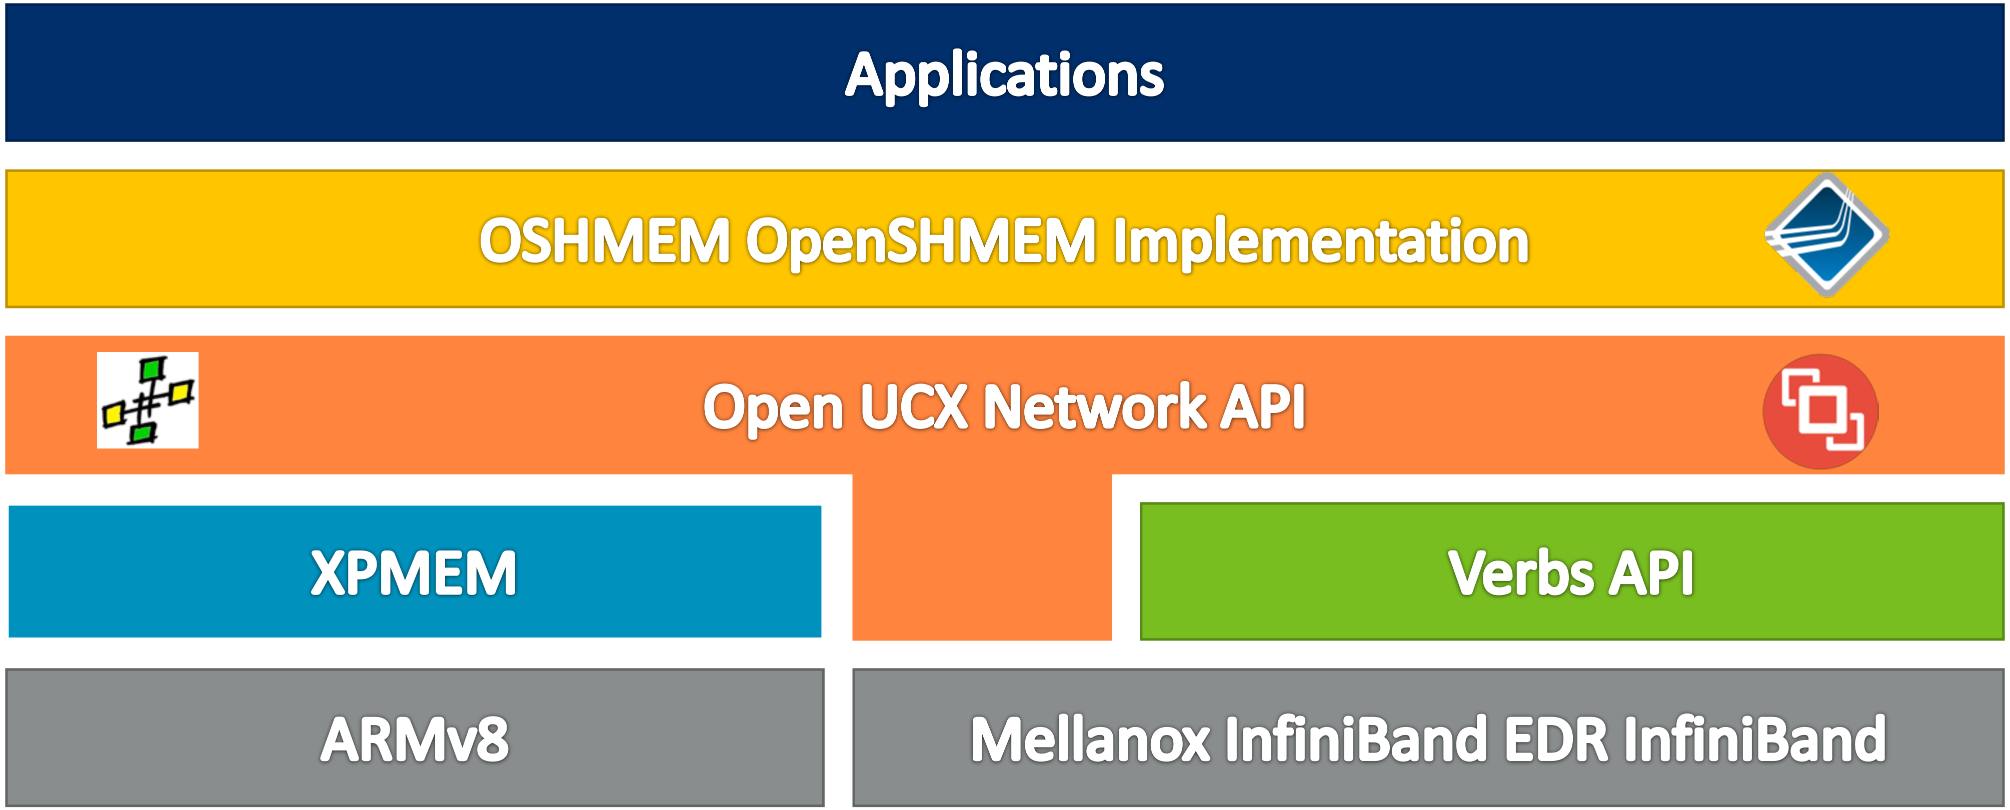

Arm Research

- Continued SVE evolution including enhanced support for Machine Learning, Graph Analytics, and Big Data
- On node scaling improvements including evaluating optimizations for task based parallelism
- Architectural extensions supporting tighter coupling and increased efficiency for off-node interconnects
- Data movement and reorganization optimizations through platform, microarchitectural and architectural techniques
- Continued application analysis and optimizations with an eye towards SVE and other architectural roadmap improvements
- gem5 SVE support
- Support EPI in whatever way possible

Arm Products

- Allinea Studio
  - Continuous improvement of Allinea toolchain with compilers, performance libraries, and microarchitectural analysis tools
- Open Source Enablement
  - OpenHPC involvement
  - GCC and LLVM upstream contributions
- SVE Exploration Tools
  - ARMie DynamoRIO SVE Support
  - Early evaluation version of Arm commercial tools and libraries supporting SVE
  - Arm code advisor vectorization analysis framework

arm

One More Thing....

http://www.arm.com/hpc

