arm Research

Arm's role in co-design for the next generation of HPC platforms

Filippo Spiga
Software and Large Scale Systems

What it is Co-design?

Abstract: Preparations for Exascale computing have led to the realization that future computing environments will be significantly different from those that provide Petascale capabilities. This change is driven by energy constraints, which is compelling architects to design systems that will require a significant re-thinking of how algorithms are developed and implemented. **Co-design has been proposed as a methodology for scientific application, software and hardware communities to work together.** This chapter gives an overview of co-design and discusses the early application of this methodology to High Performance Computing.

[…] **The co-design strategy is based on developing partnerships with computer vendors and application scientists and engaging them in a highly collaborative and iterative design process well before a given system is available for commercial use.** The process is built around identifying leading edge, high-impact scientific applications and providing concrete optimization targets rather than focusing on speeds and feeds (FLOPs and bandwidth) and percent of peak.

"On the Role of Co-design in High Performance Computing", *Transition of HPC Towards Exascale Computing* (2013)

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

Based on the #1 architecture and #1 ecosystem with more than 100 billion chips shipped to date

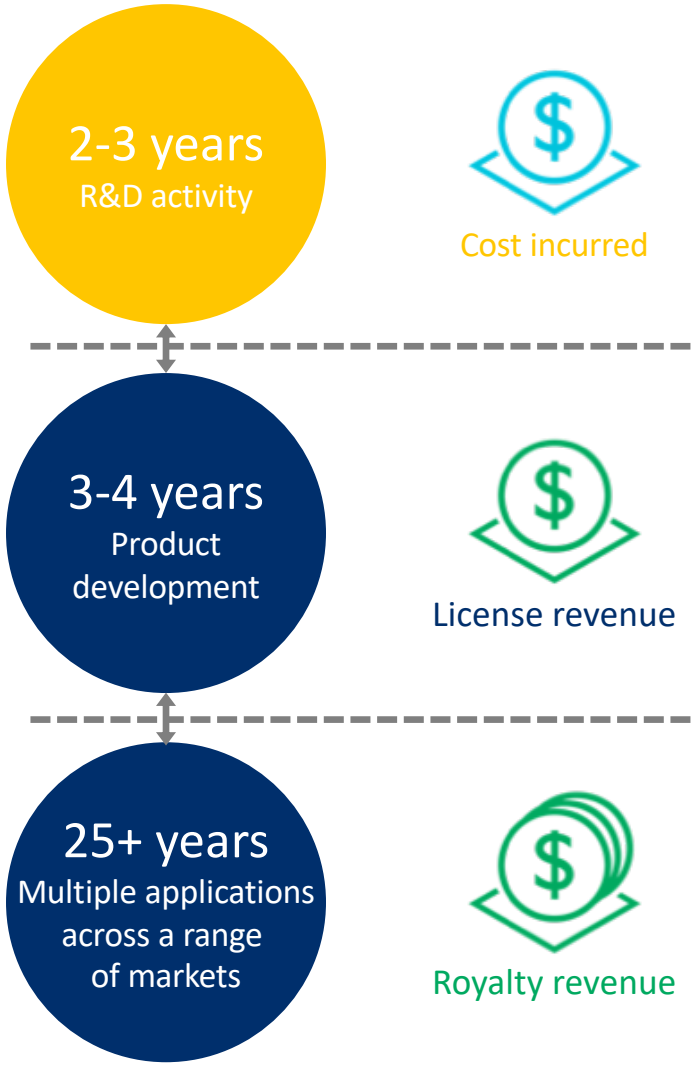

arm Research

Why Arm matters in HPC?

Power Efficiency

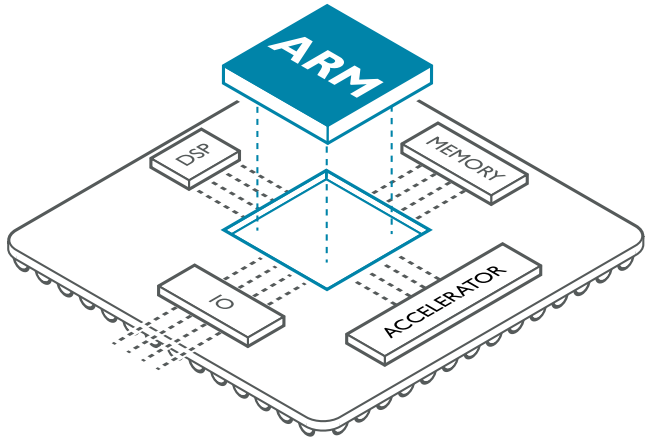

Choice

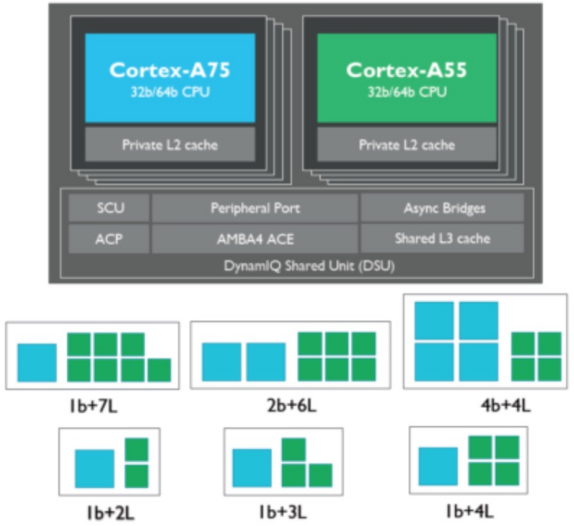

Customization

arm Research

Arm HPC strategy

Mission

Enable the world's first Arm supercomputer(s)

Strategy

Enablement + Co-Design + Partnership

Building Blocks

Enablement

- **Address gaps** in computational capability and data movement within Architecture
- **Seed the software ecosystem** with open source support for Armv8 and SVE libraries, tools, and optimized workloads
- Provide **world class tools** for compilation, analysis, and debug at large scale.

Co-Design

- Work with key end-customers in DoE, DoD, RIKEN, and EU to **design balanced architecture**, uArchitecture and SoCs based on real-world workloads, not benchmarks.
- Develop **simulation and modelling tools** to support co-design development with end-customers, partners, and academia.

Partnership

- Work with Architecture partners to quickly bring **optimized solutions to market**.
- Work with ATG & uArchitecture design teams to **steer future designs** to be more relevant for HPC, HPDA, and ML
- Work with key ISVs to enable mid-market

European HPC initiatives adopting Arm for HPC

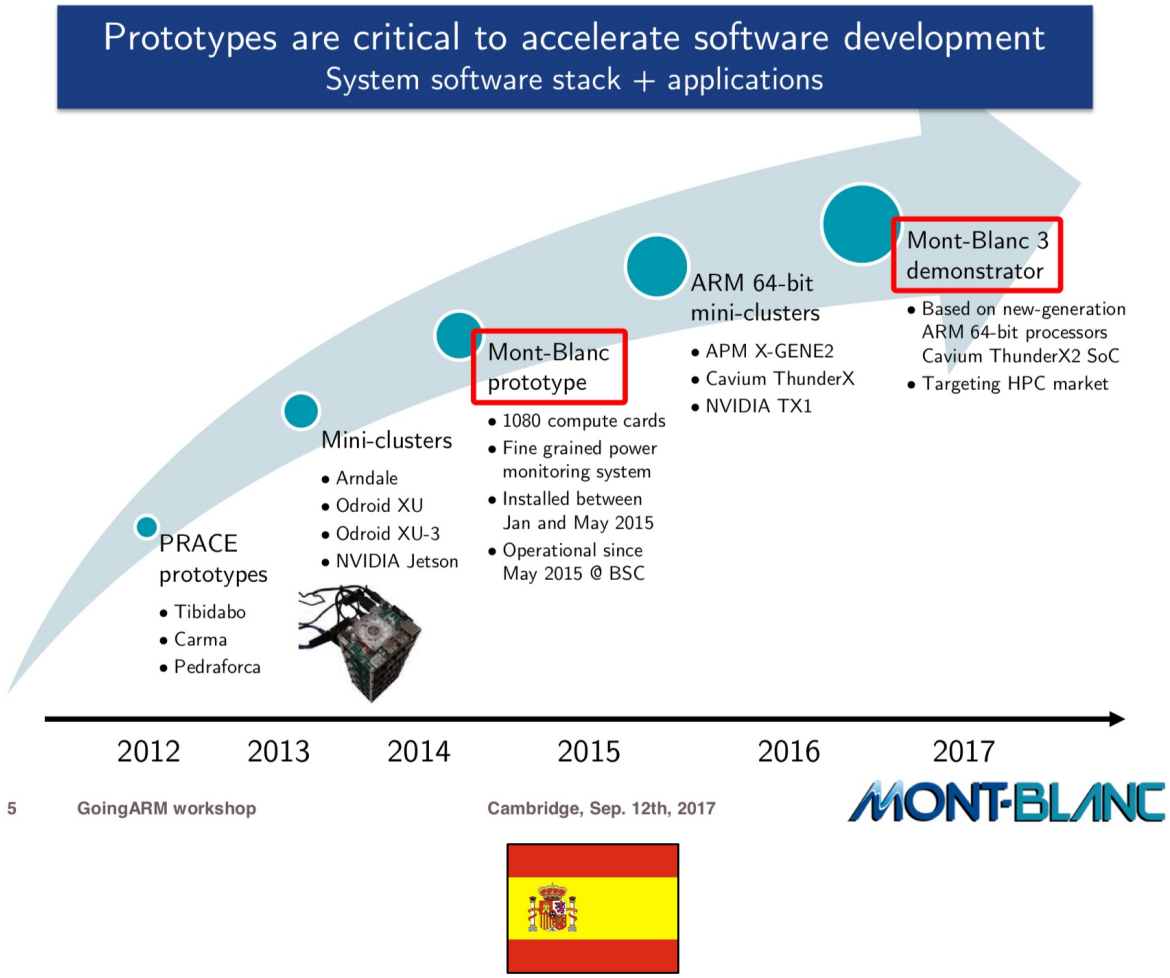

University of BRISTOL Isambard GW4

Isambard system specification:

- Cray system
- **10,000+** ARMv8 cores
- Cray software tools
  - Compiler, math libs, tools…
- Technology comparison:
  - x86, Xeon Phi, Pascal GPUs
- Phase 1 installed March 2017
- The ARM part arrives early 2018
  - Early access nodes from Sep

I.K.Brunel 1804-1859

Simon McIntosh Smith, simonm@cs.bris.ac.uk, @simonmcs http://gw4.ac.uk/isambard/ 6 bristol.ac.uk

arm Research

Other WW HPC initiatives adopting Arm for HPC

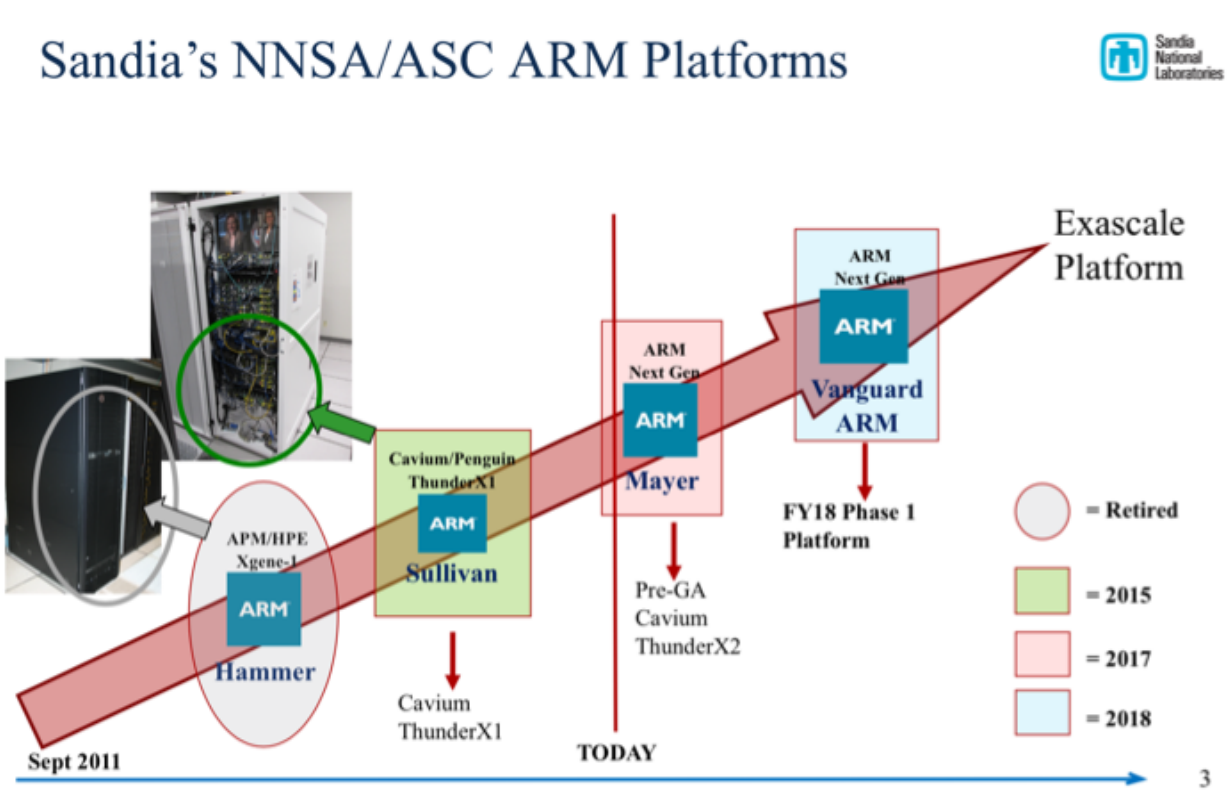

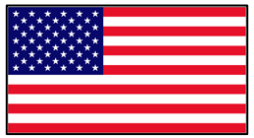

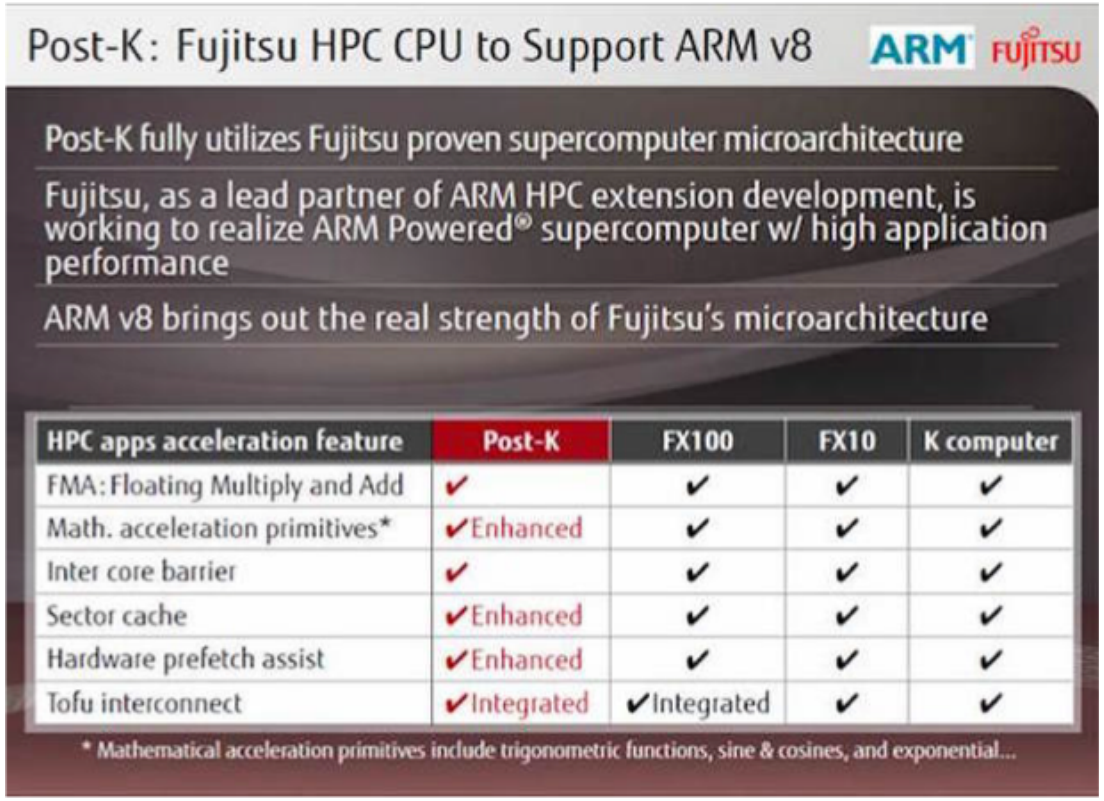

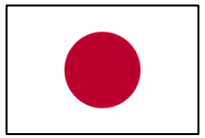

arm Research

Co-design with Arm

Analysis of applications to devise the most efficient solutions

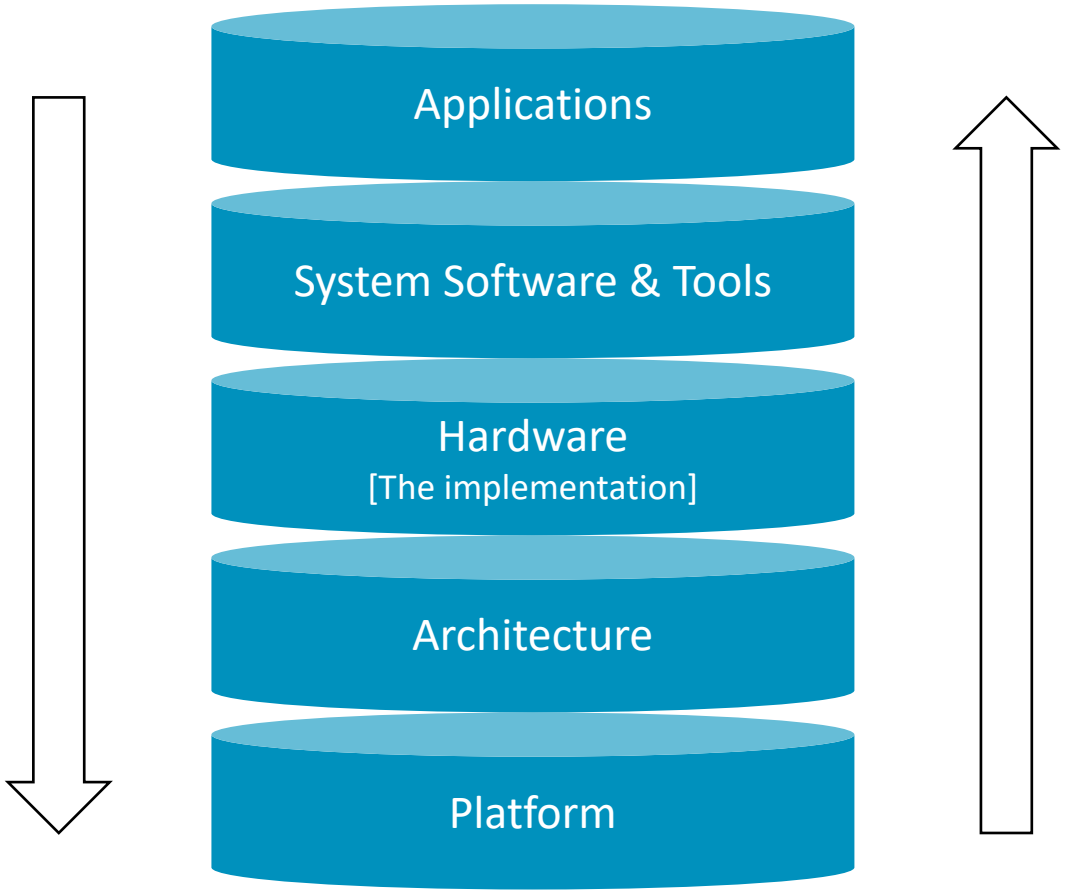

Issues and opportunities to exploit

arm Research

HPC Applications

OpenFOAM

LAMMPS

QUANTUMESPRESSO

NAMD
Scalable Molecular Dynamics

EPOCH

Geant 4

GROMACS FAST. FLEXIBLE. FREE.

WRF

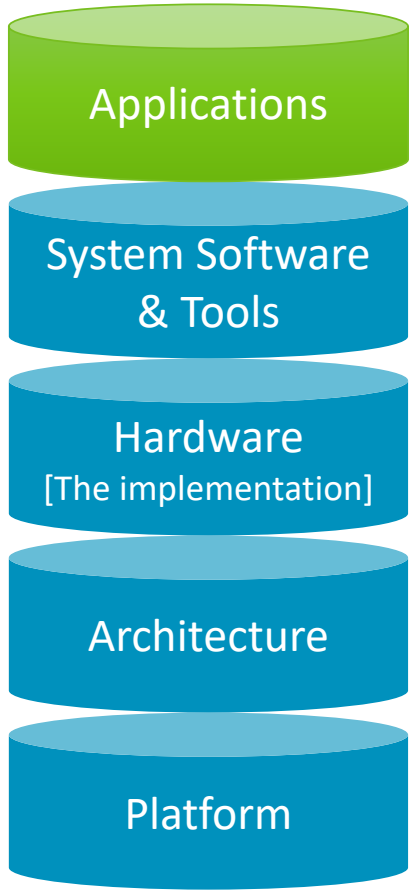

- Initial focus on successfully building with both **Arm** and **GCC** compilers across a broad front.
- Often only modest changes to environment variables, build scripts and architecture files are needed

Public wiki & repo = https://gitlab.com/arm-hpc/packages/wikis/home

Arm Allinea Studio

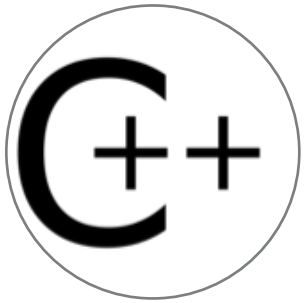

C/C++ Compiler

- C++ 14 support
- OpenMP 4.5 without offloading
- SVE ready
- Compiler Optimization report

Fortran Compiler

- Fortran 2003 support
- Partial Fortran 2008 support
- OpenMP 3.1
- SVE ready

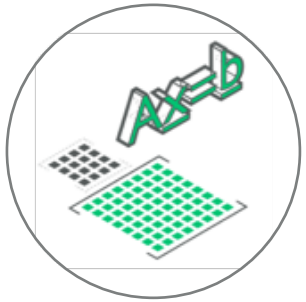

Performance Libraries

- Optimized math libraries
- BLAS, LAPACK and FFT
- Threaded parallelism with OpenMP
- GEMM tuning for ThunderX2 and Qualcomm Falkor
- FFT optimizations
- Batched BLAS

Forge (DDT and MAP)

- Profile, Tune and Debug
- Scalable debugging with Arm Allinea DDT
- Parallel Profiling with Arm Allinea MAP
- Arm hardware performance counters support
- Interop with Arm Compiler and Libraries

Performance Reports

- Analyze your application
- Memory, MPI, Threads, I/O, CPU metrics
- Arm hardware performance counters support
- Interop with Arm Compiler and Libraries

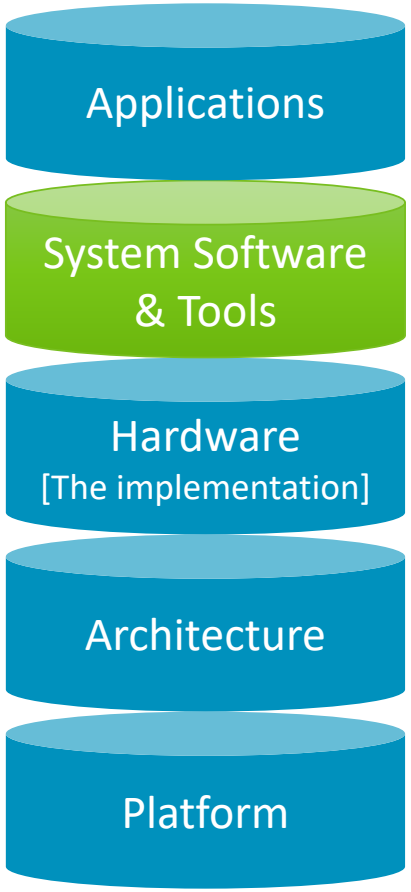

Website: https://www.arm.com/products/development-tools/hpc-tools/allinea-studio

arm Research

Arm Compiler - Building on LLVM, Clang and Flang projects

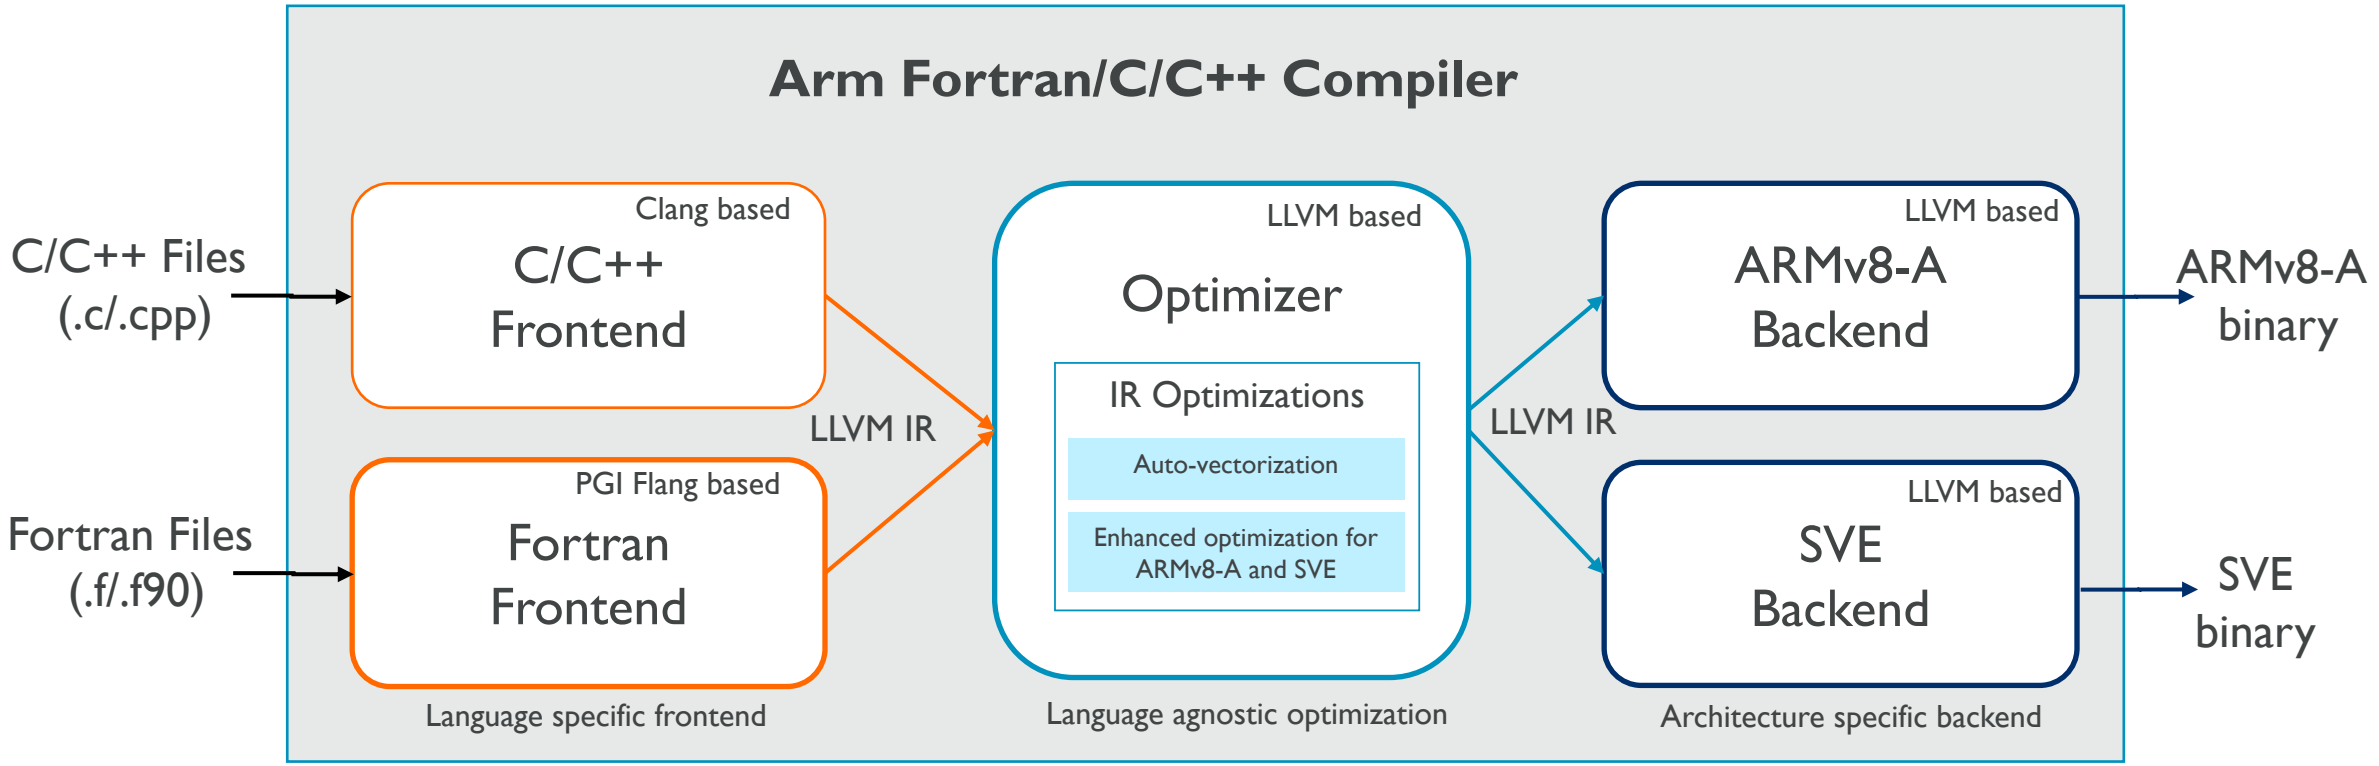

arm Research

System Software & Tools roadmap

Fortran Compiler
- Improvements in debugging
- Increased Fortran 2008 support
- Improved OpenMP 4.5 support

All compilers
- Improvements in optimization report

More features in compilers

- Application specific tuning and optimization
- For Cavium ThunderX2 and other server-class Arm-based platforms

More optimizations for current hardware

- SVE enabled Performance Libraries
- Application specific tuning and optimization in Compilers and Libraries for SVE

Getting ready for SVE-based future hardware

arm Research

The hardware

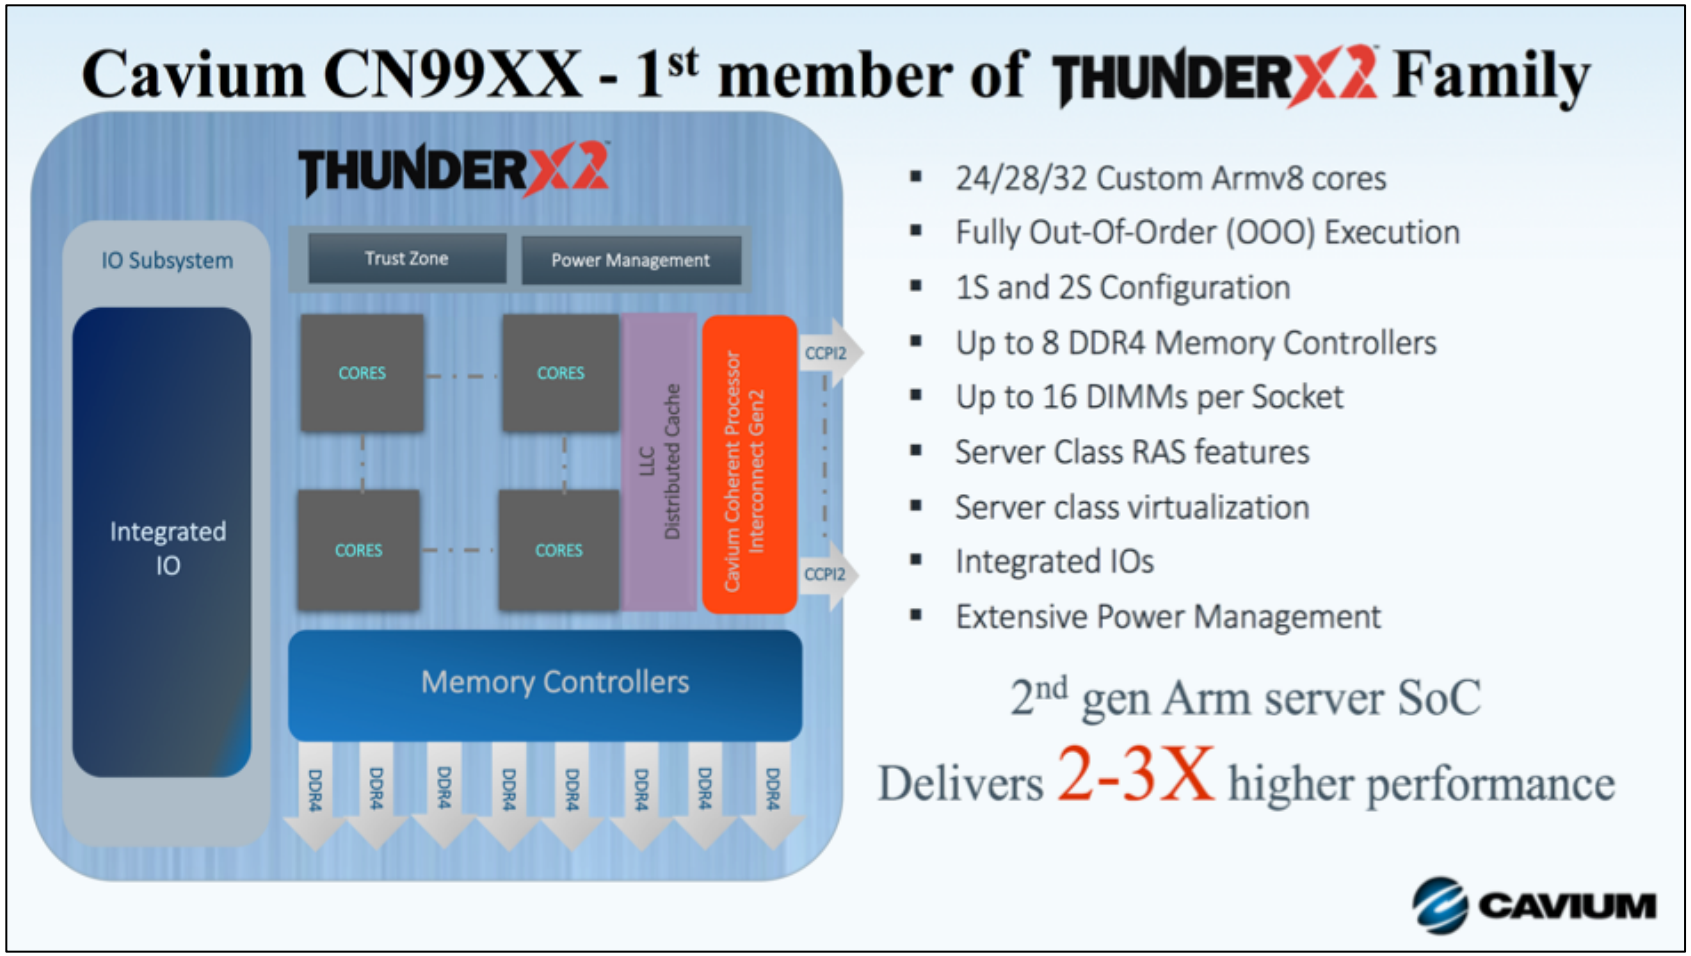

2016

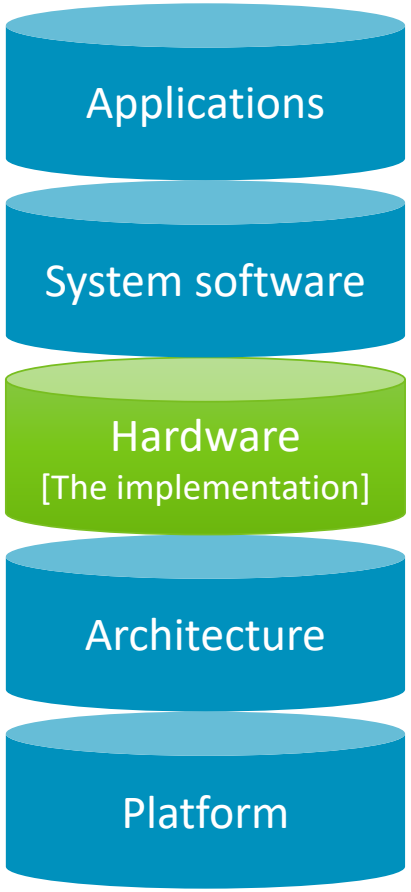

arm Research

The hardware

Qualcomm Centriq™ 2400 Processor

Accelerating Innovation in the Data Center

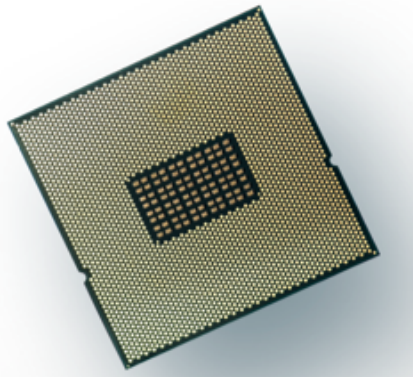

Performance
48-core SoC design with optimized performance for throughput data center workloads

Power
World's first and only 10nm processor for server-class performance and leading power efficiency

Total Cost of Ownership
Significant hardware CapEx and OpEx savings

Ecosystem
Robust Arm64-based ecosystem across infrastructure software and core operating system components

Security
Reliable security for the modern datacenter founded on immutable root of trust and TrustZone architecture

10

2017

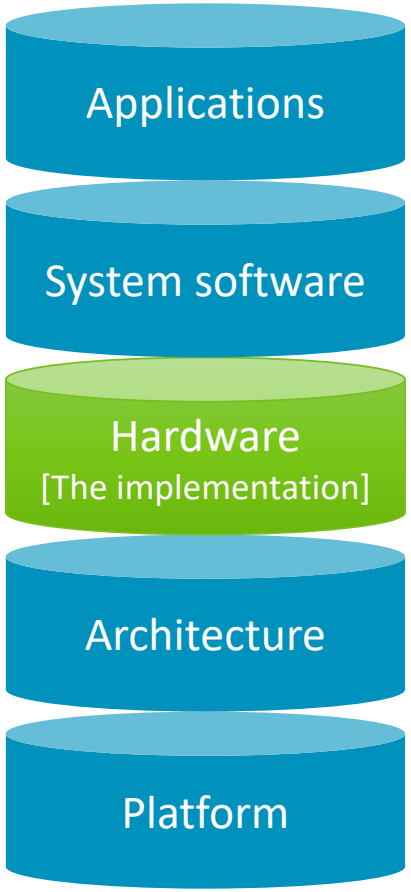

arm Research

DGEMM performance with Arm tools 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

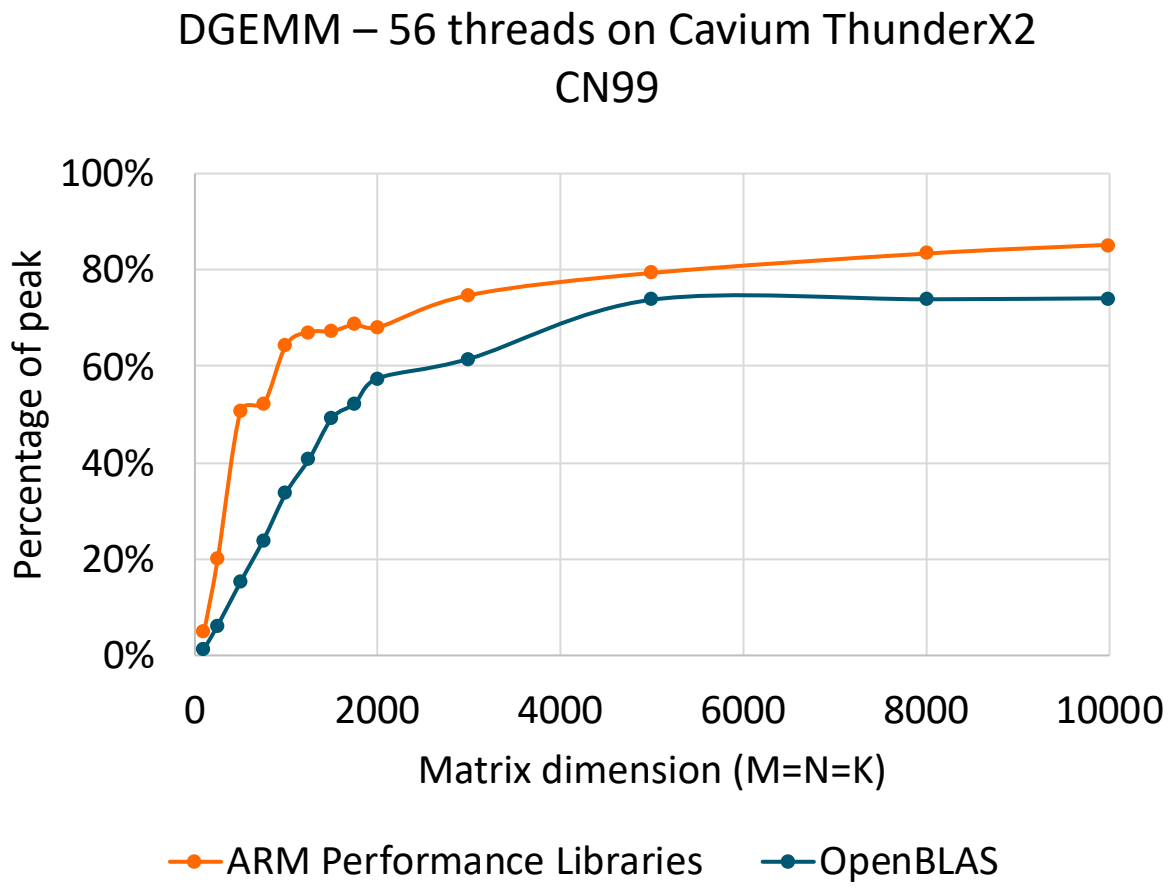

arm Research

Results from Isambard on ThunderX2

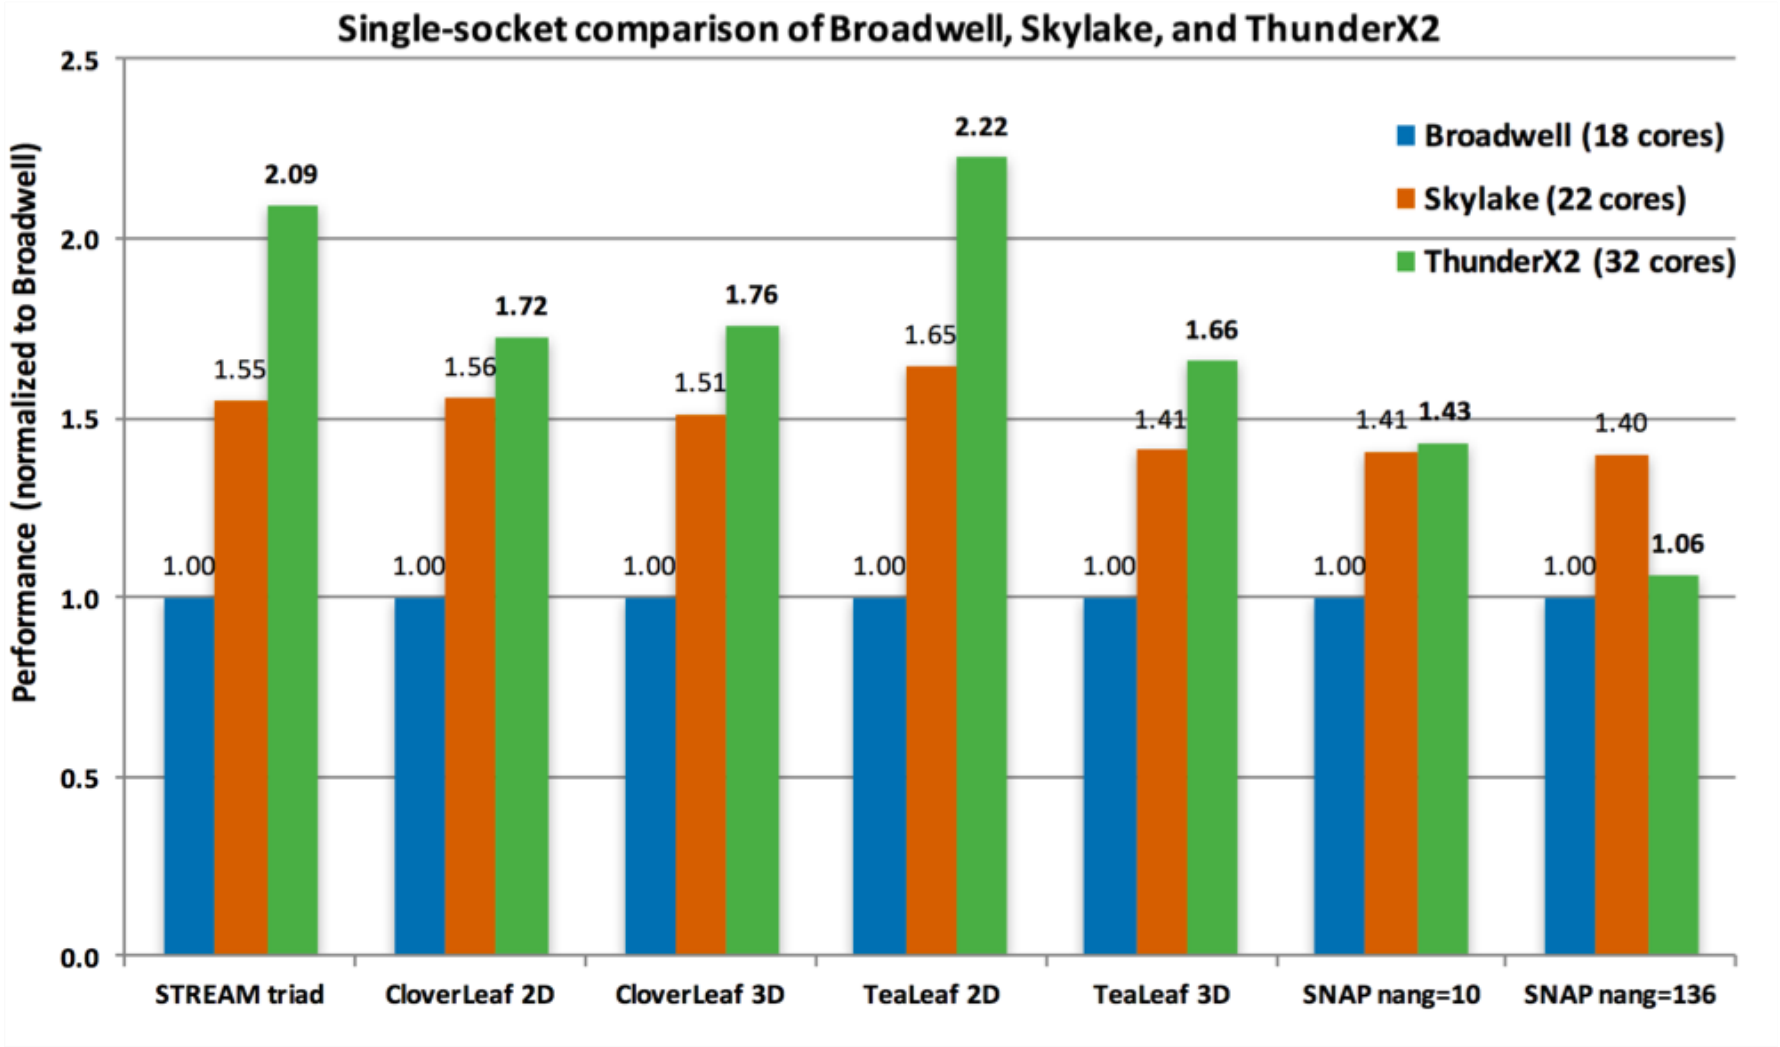

Source: http://gw4.ac.uk/isambard

arm Research

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
- Common-sense says you need high vector utilisation to achieve significant speedups
  - Compilers often unable to vectorize due to intra-vector data & control dependencies
  - Begins to address some of the traditional barriers to auto-vectorization (e.g. control flow)

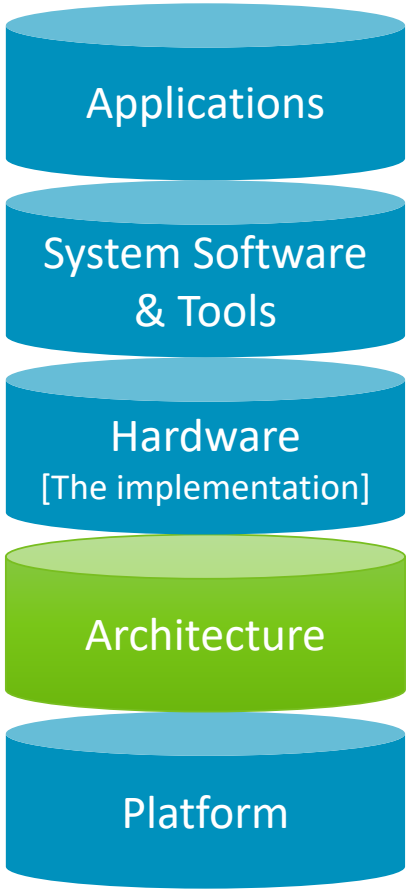

arm Research

Scalable Vector Extension (SVE)

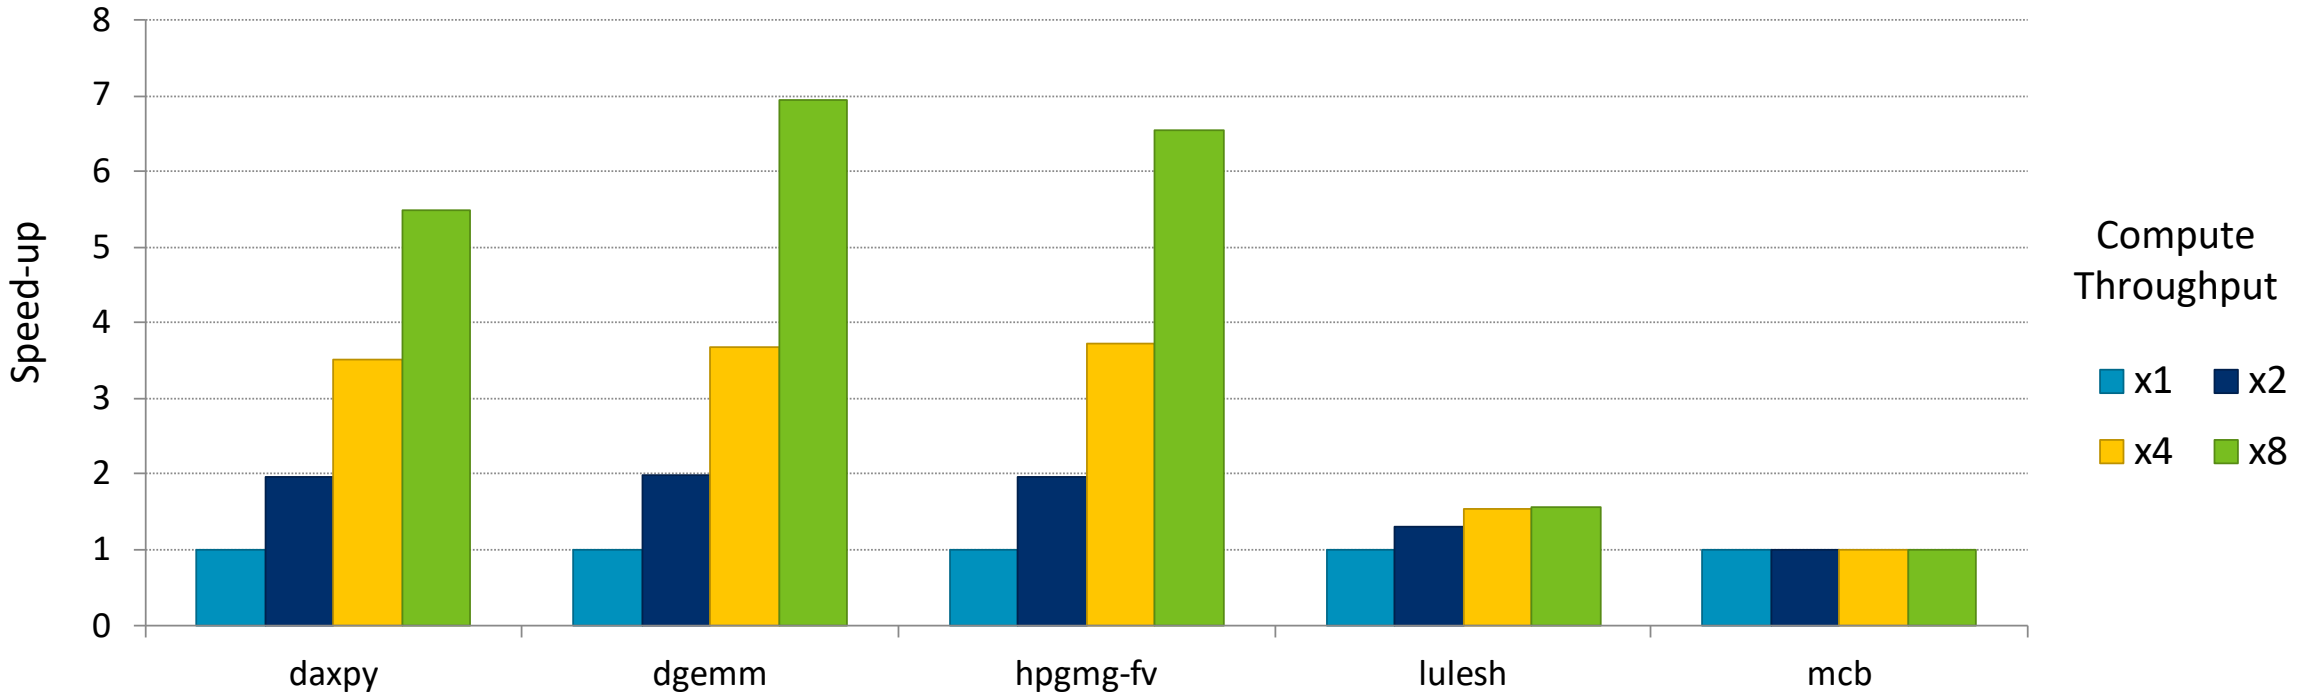

HPCAFE-2017 "*Energy-efficient HPC in Mont-Blanc and beyond: an ARM hardware and software perspective*", Roxana Rusitoru

arm Research

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.

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.

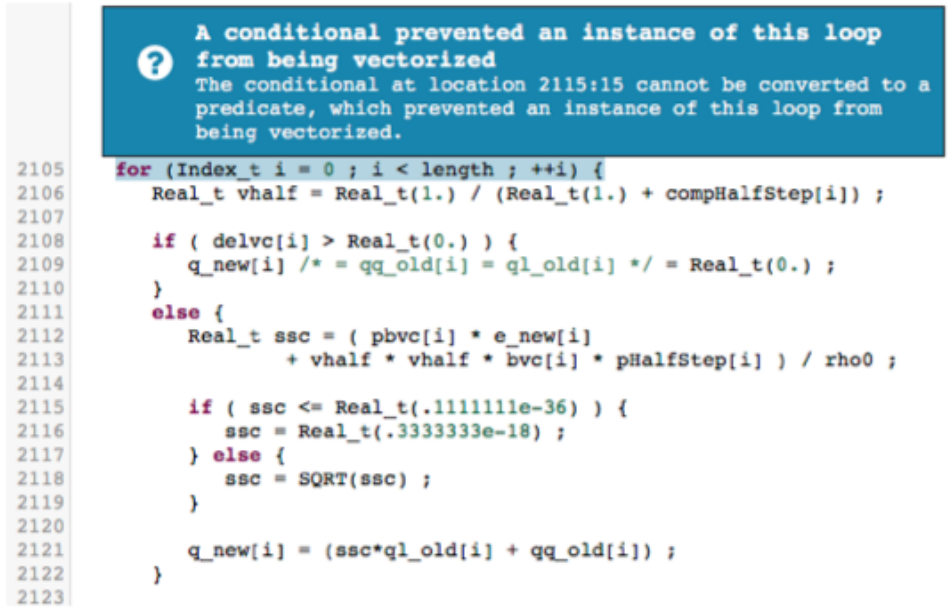

arm Research

Jülich efforts exploring SVE for scientific applications

Objective: understand if apps cat benefit from SVE, assess quality and readiness of tools

- Various Arm-based SoC (Huawei Taishan)
- Several applications of interest: QE, KKRnano, GRID, BQCD
- Results on MiniKKR show Arm-based SoC (no SVE) similar performance figure versus x86
- Estimate performance using instruction/branch counting (*dynamic*) and critical path analysis (*static*)

JÜLICH FORSCHUNGSZENTRUM

"**Early Experience with ARM SVE**", presented at SC'17 Arm SVE User Meeting by D. Pleiter (JSC)
http://www.goingarm.com/slides/2017/SVE_SC17/GoingArm_SVE_SC17_Arm_Dirk.pdf

"**Exploring SVE for scientific applications**", presented at HiPEAC'18 by S. Nassyr (JSC)
http://www.goingarm.com/slides/2018/HiPEAC2018/julich_hipeac_goingarm_2018.pdf

Platforms

- **HPE Apollo 70**
  https://news.hpe.com/hpe-helps-businesses-capitalize-on-hpc-and-ai-applications-with-new-high-density-compute-and-storage/

- **CRAY XC50**
  http://investors.cray.com/phoenix.zhtml?c=98390&p=irol-newsArticle&ID=2316352

- **ATOS Bull Sequana x1310**
  https://atos.net/en/2017/press-release/general-press-releases_2017_06_19/atos-expands-range-supercomputers-include-arm-processors-new-bull-sequana-x1310

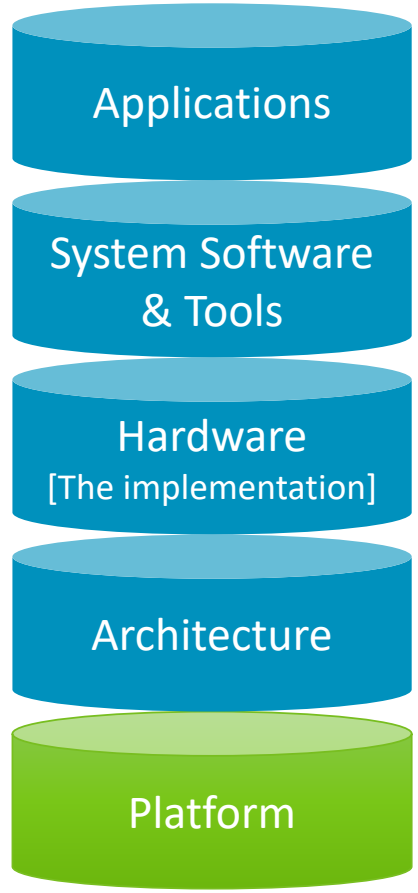

arm Research

Building communities

Outside the people we collaborate with, various complementary Arm HPC communities already exist:

- **Arm HPC User Group** (SC) and **GoingArm** (ISC/ArmRS/HiPEAC)
- **Arm HPC Google Group** (https://groups.google.com/forum/arm-hpc)
- **Arm HPC GitLab pages** (https://gitlab.com/arm-hpc/)

Our app work is **engaging with code owners and users** to get suitable test cases, to get Arm support built in, and including helping them make testing part of their development processes

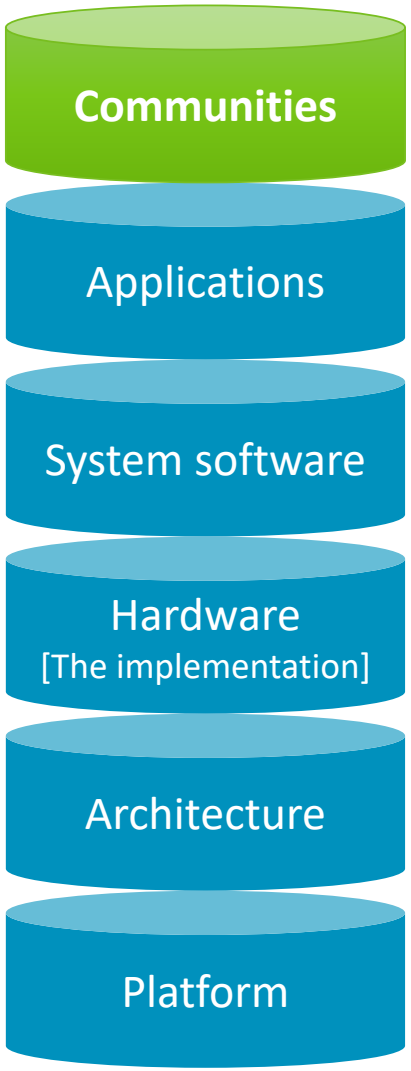

arm Research