ARM Transforming HPC

Srinath Vadlamani
HPC Field Application Engineer

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 bring optimized solutions to market quickly.
- 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

arm

Efficient 64-bit architecture for modern workloads

Supported by an extensive commercial and open-source software ecosystem

- Enables real change and choice for HPC market
  - Highly integrated SoC's and efficient cores for greater compute density and power-efficiency
  - Rebalancing of CPU, Memory, and IO
  - Numerous SoC partners for infrastructure and HPC
    - Also addressing cloud solutions
- Provider of commercial and open-source software for HPC
- World-wide partnerships across all industries and markets
  - Over 125B Arm cores shipped in 25yrs; out ships x86 by 20X per year
  - Mobile, embedded, auto, infrastructure, and IoT

arm

Deployments

arm

ThunderX Momentum in HPC Continues to Grow…

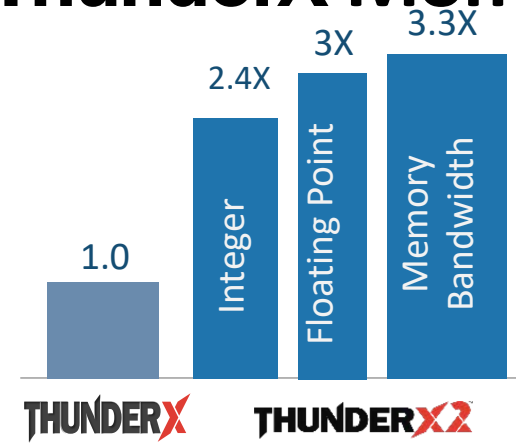

2-3X better HPC performance

Server platforms at World's premier HPC Labs

Significant HPC Engagements

ODM and OEM Platforms

THUNDERX2

Hewlett Packard Enterprise

Density Optimized Rack Mount Server - 2U4N

Atos

X1310- Bull Sequana Compute Blade

CRAY

XC50 Supercomputer

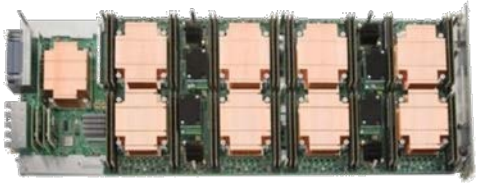

- Chassis with 16 compute blades
- 128 Sockets
- Inter-Aries communcation over backplane

GIGABYTE™

PENGUIN COMPUTING

TundraTM Extreme Scale

OCP Platforms

Inventec

Inventec OCP Server

FOXCONN

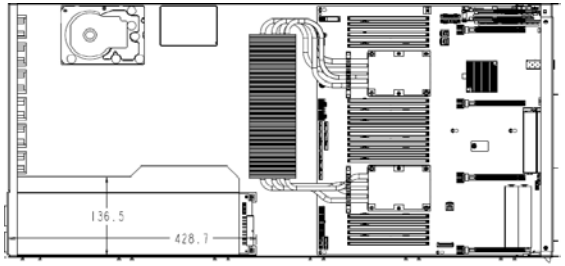

C50 OCP
Project Olympus Server

arm

First purpose-built Arm HPC system by the #1 HPC vendor

HPE Apollo 70 System overview

HPE driving multi-vendor effort to accelerate Arm adoption for HPC

arm
CAVIUM
redhat
Hewlett Packard Enterprise
AMD
SUSE
Mellanox TECHNOLOGIES
RogueWave SOFTWARE

World-class HPE engineering and innovation	- **Leverages proven HPE Apollo 2000 system arch** - **Management support** for HPE provisioning, cluster management and performance software - **Tier 1 vendor services and support** (HPE Pointnext)
Leading Arm technology	- **Cavium ThunderX2 with 32 64-bit Arm v8.1** custom cores for today's HPC workloads - **Greater density and power-efficiency** with highly integrated SoCs - Brings **real change and choice** to HPC market
Density and scalability for HPC clusters	- **Dual socket nodes –** Up to 64 cores per node - **Maximum density** - Up to 80 servers / rack - **High speed interconnects** 100 Gb/s InfiniBand EDR - **Architected to scale** to multi-rack levels

arm

Mellanox ConnectX Family Enables Apollo A70 Interconnect

ConnectX-4 Lx

Best Ethernet Connectivity

On Board ConnectX®-4 Lx - 10G/25GE Ethernet
Smart offloading - Virtual Overlay Network, OVS, RDMA

ConnectX-5

Best HPC and AI Connectivity

100Gb/s InfiniBand and Ethernet
MPI Accelerations engines (Tag Matching, Collectives, Rendezvous protocol)
GPUDirect RDMA, Adaptive Routing, Atomic Operations enhancements
Burst Buffer Offload, NVMe over Fabric

Choice, Flexibility
Highest Performance

HPE Apollo 70

Hewlett Packard Enterprise

arm

ARM® Supercomputing
…Made Possible by Cray

CRAY CATAPULTS ARM-BASED PROCESSORS INTO SUPERCOMPUTING
Cray Adds Arm Processors with Complete Software Stack to the Cray XC50 Supercomputer

Seattle, WA – November 13, 2017 – Global supercomputer leader Cray Inc. (Nasdaq: CRAY) today announced the Company is creating the world's first production-ready, Arm®-based supercomputer with the addition of Cavium (Nasdaq: CAVM) ThunderX2™ processors, based…

CAVIUM

CRAY Programming Environment

arm

Headline results from GW4

University of BRISTOL | Isambard | GW

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

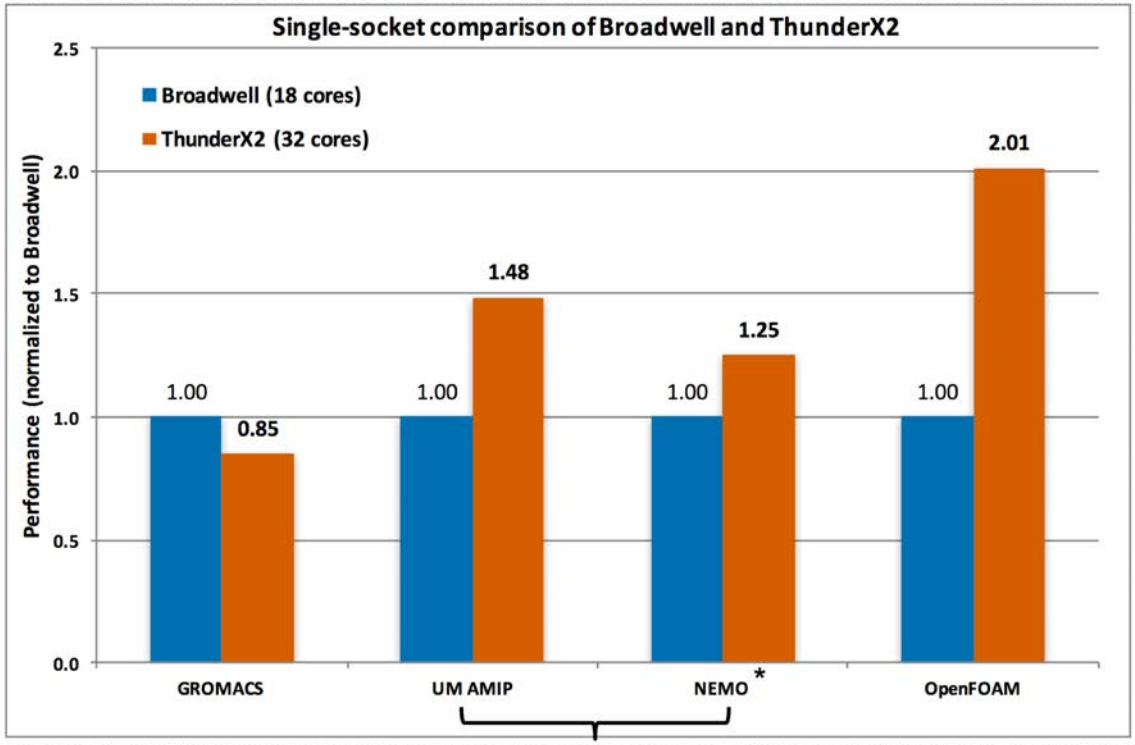

@simonmcs http://gw4.ac.uk/isambard/

Benchmarked by the UK's Met Office

bristol.ac.uk

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

University of BRISTOL

Isambard

GW4

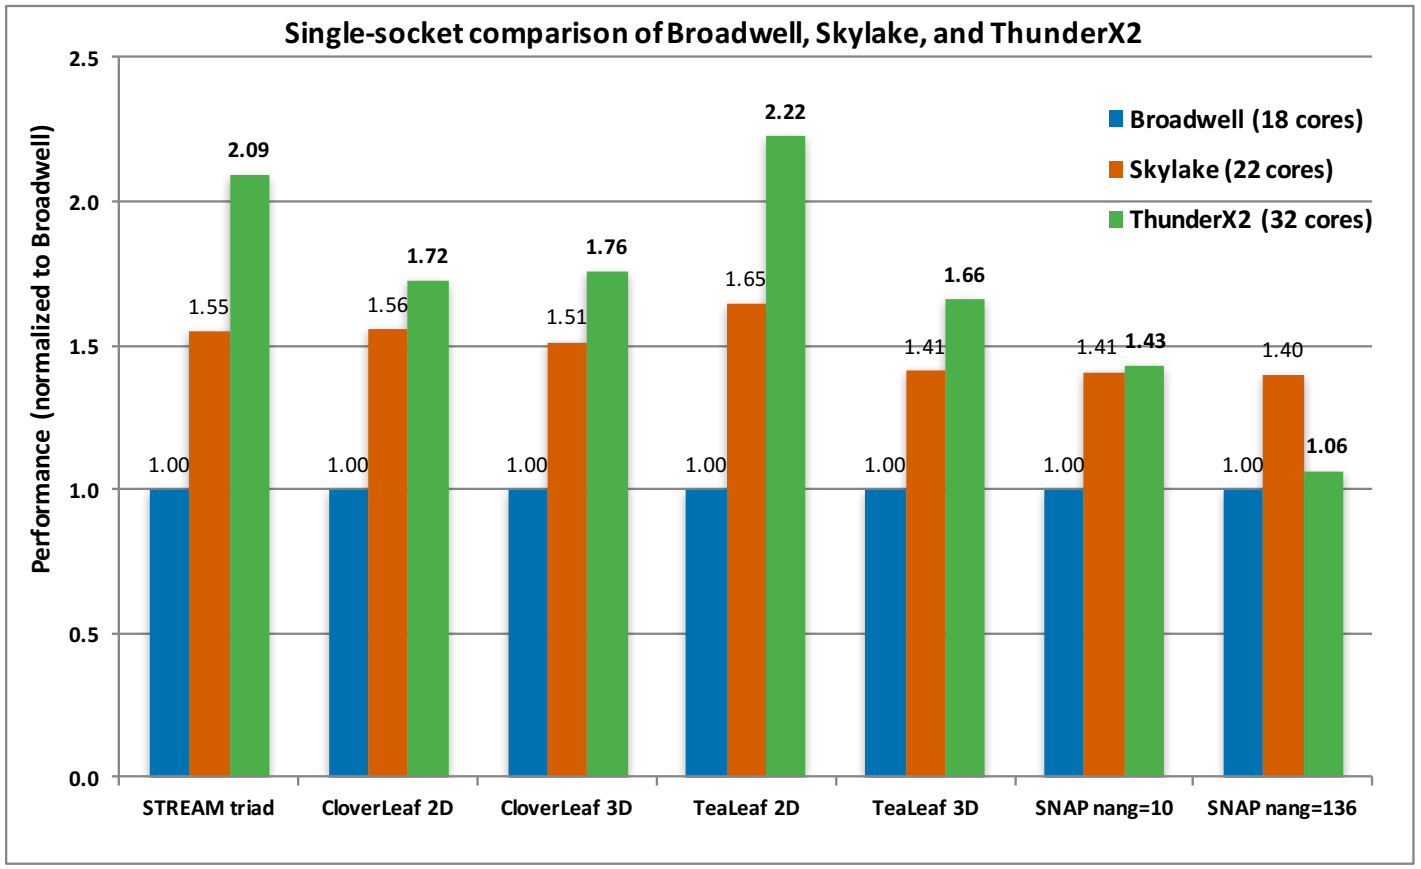

Catalyst UK

Accelerating Arm adoption in the UK

Sites and Target HPC Applications:

- **EPCC**: WRF, OpenFOAM, Rolls Royce Hydra opt, 2 PhD candidates
- **Leicester**: Data-intensive apps, genomics, MOAB Torque, DiRAC collab
- **Bristol**: VASP, CASTEP, Gromacs, CP2K, Unified Model, Hydra, NAMD, Oasis, NEMO, OpenIFS, CASINO, LAMMPS

Typical Cluster for each site:

- 64 x Apollo 70 Compute Nodes (2 racks):
  - Dual socket Cavium 32c, 2.2 GHz
  - 256GB memory (16GB DIMMs)
  - Mellanox IB EDR CX5 Clos
  - 4096+ cores

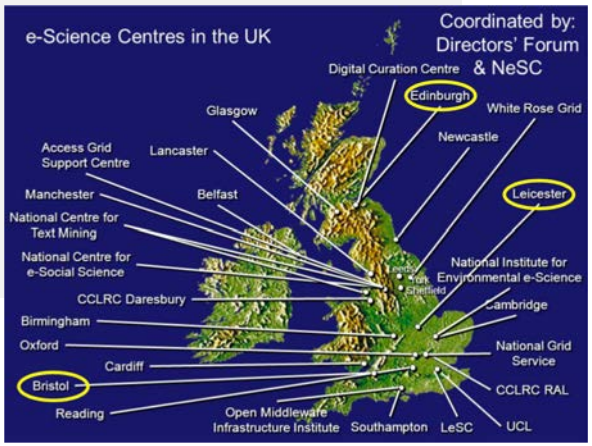

Japan

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...

slides from Fujitsu at ISC'16

arm

Arm Allinea Studio

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

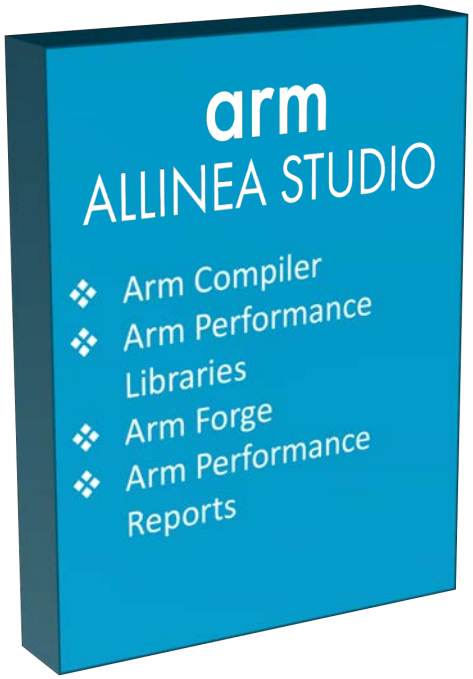

Comprehensive and integrated tool suite for Scientific computing, HPC and Enterprise developers

Cross-platform support covering Arm, x86, Power, and Nvidia

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

arm

Arm DDT parallel debugger

Analyze large-scale MPI, OpenMP, GPGPU applications

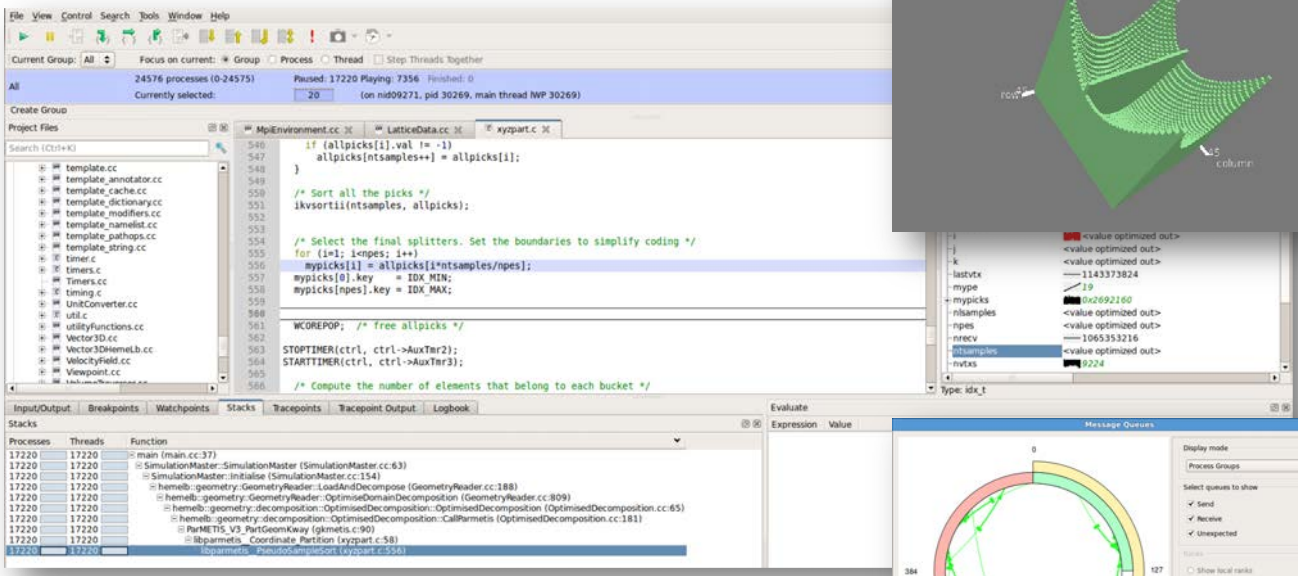

Visualise data structures

Display pending communications

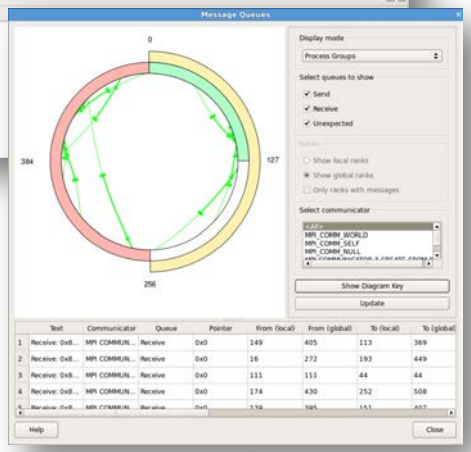

Generate offline reports

arm

Arm MAP lightweight profiler

Visualize **Compute**, **Communication**, **IO** behavior at a glance

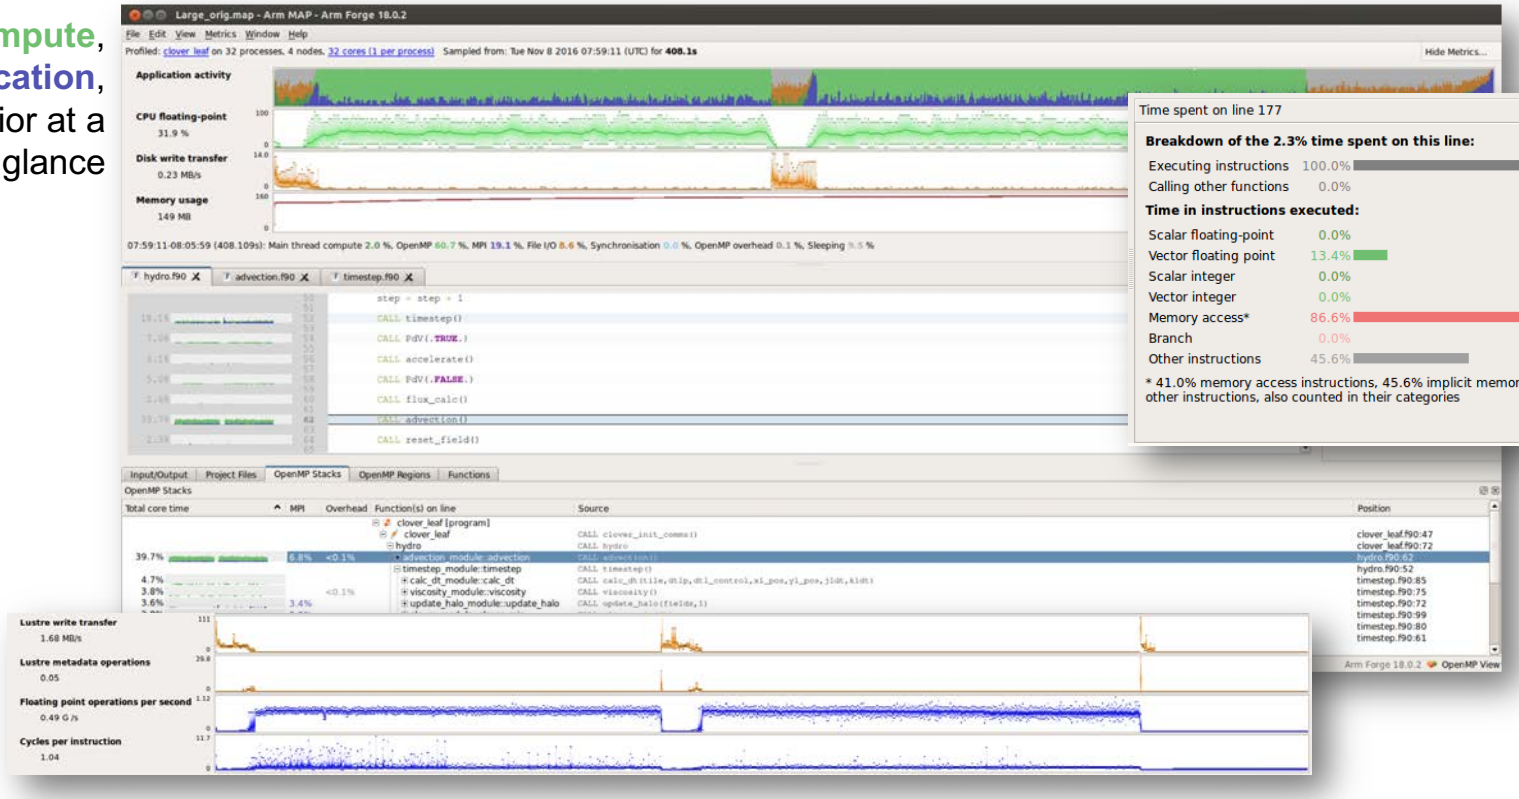

Understand per-line activity

Display advanced metrics

arm

Arm Performance Reports

Understand application efficiency without recompiling

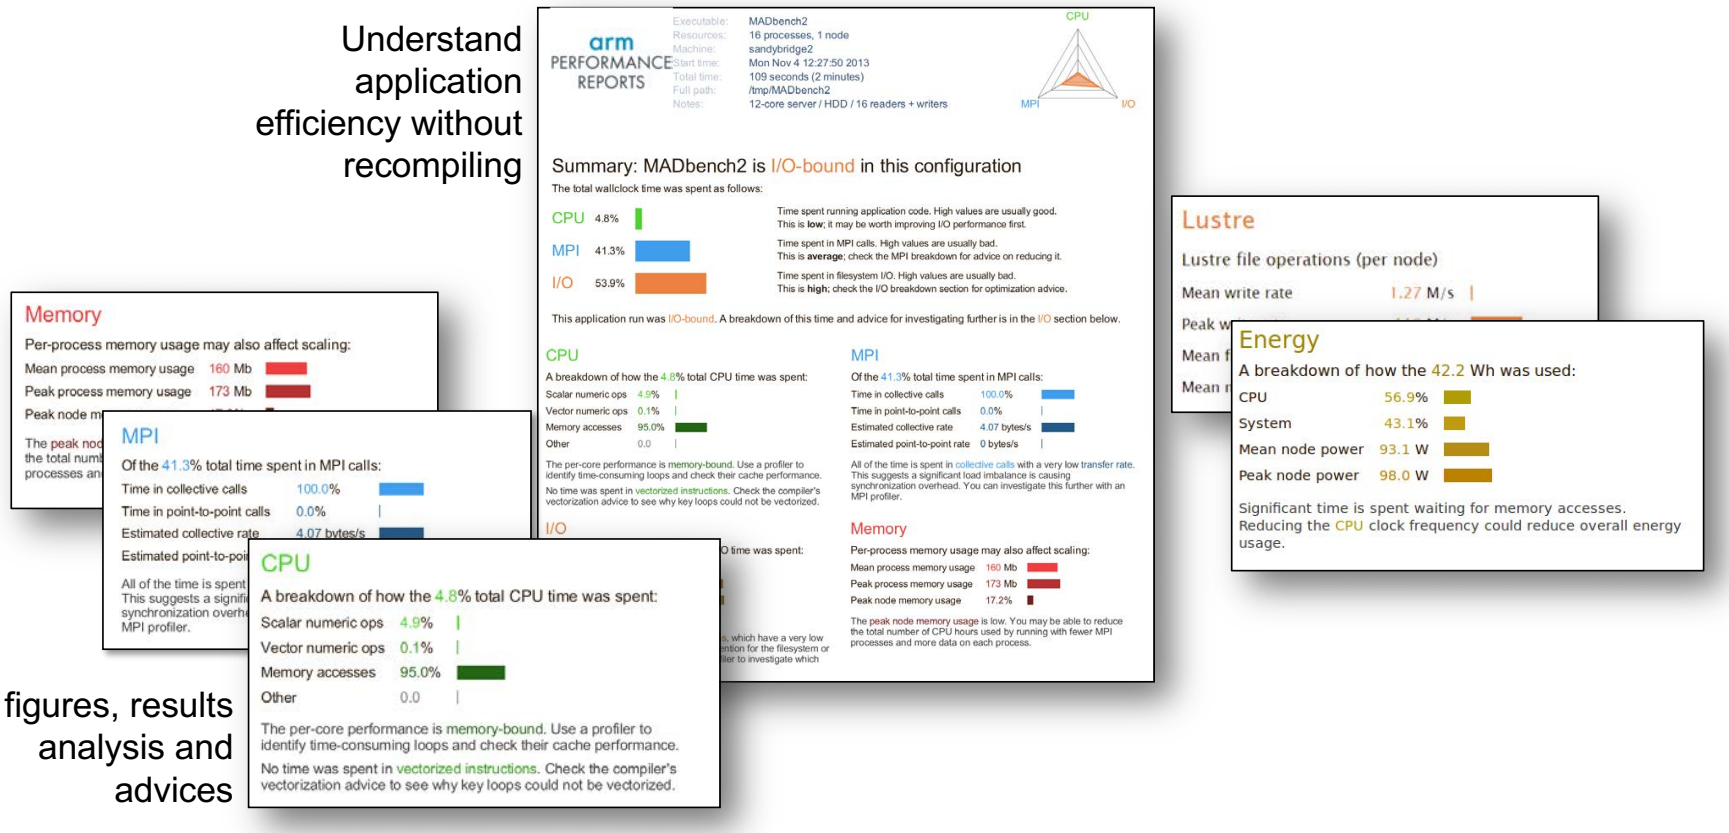

Display advanced metrics

Get figures, results analysis and advices

HPC Infrastructure Tools Group

Worldwide HPC knowledge experts in **cross-platform** enablement and **performance** optimization

Teams:

- Professional Services
- Commercial and open-source SW engineering
- HPC Tools Development, training and support

Offerings:

HPC application modernization, porting, and optimization:

- Identify and resolve scalability bottlenecks.
- Optimize application data movement and vectorization.
- Cross-platform

Performance analysis tools and techniques:

- Deploy and support software tools from Arm and the Arm community.
- Documentation, training, tutorials, and workshops.

Knowledge transfer and Arm adoption:

- Publications, presentations, and community participation.

arm

Useful links

- http://www.goingarm.com
  - User community

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

- https://developer.arm.com/hpc
  - Where to get Arm HPC software with “getting started”
- support-hpc-sw@arm.com

arm

Back up results slides

arm

University of BRISTOL

Isambard

GW4

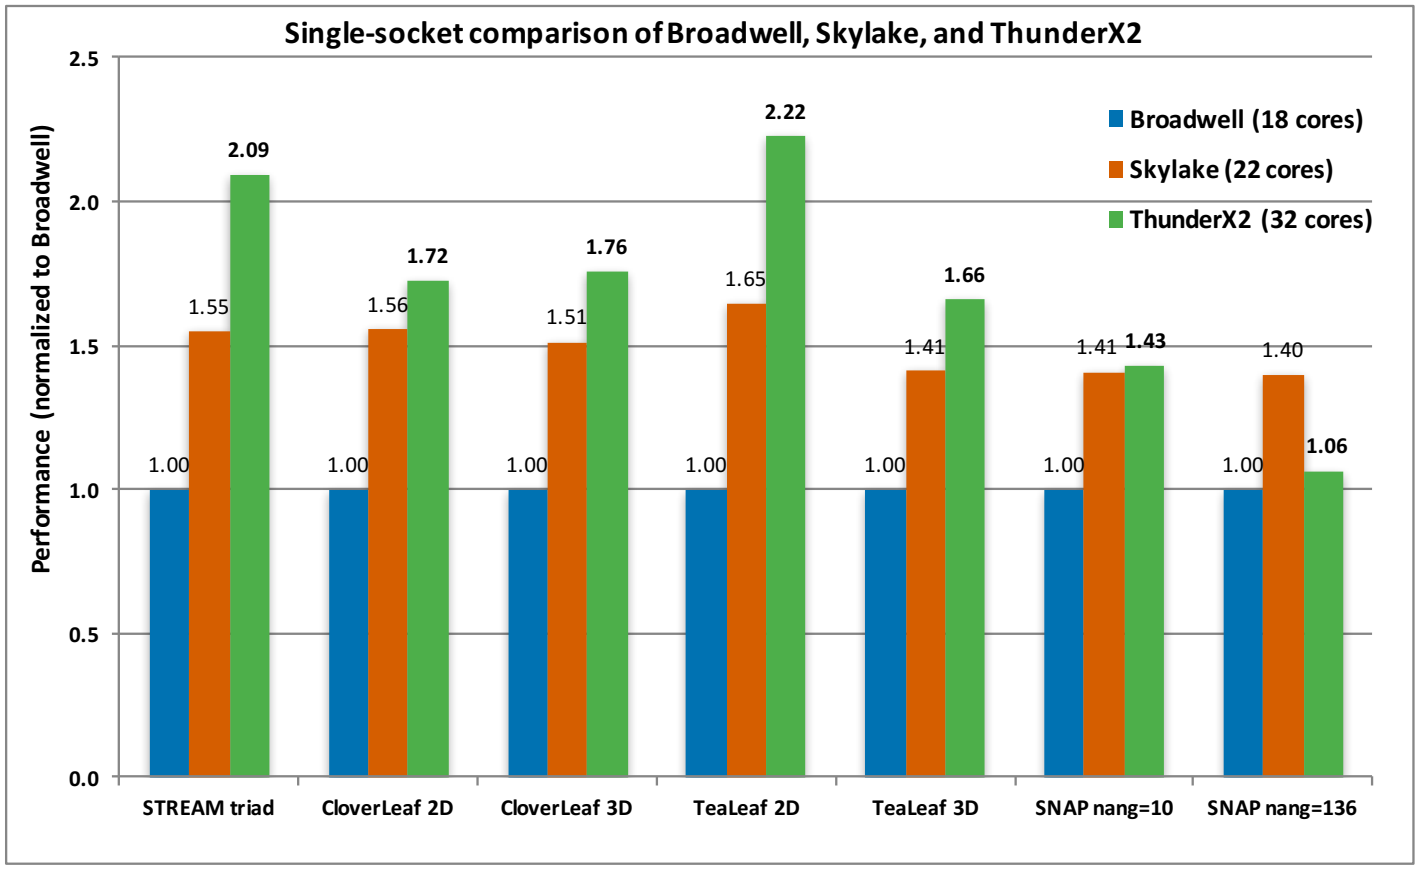

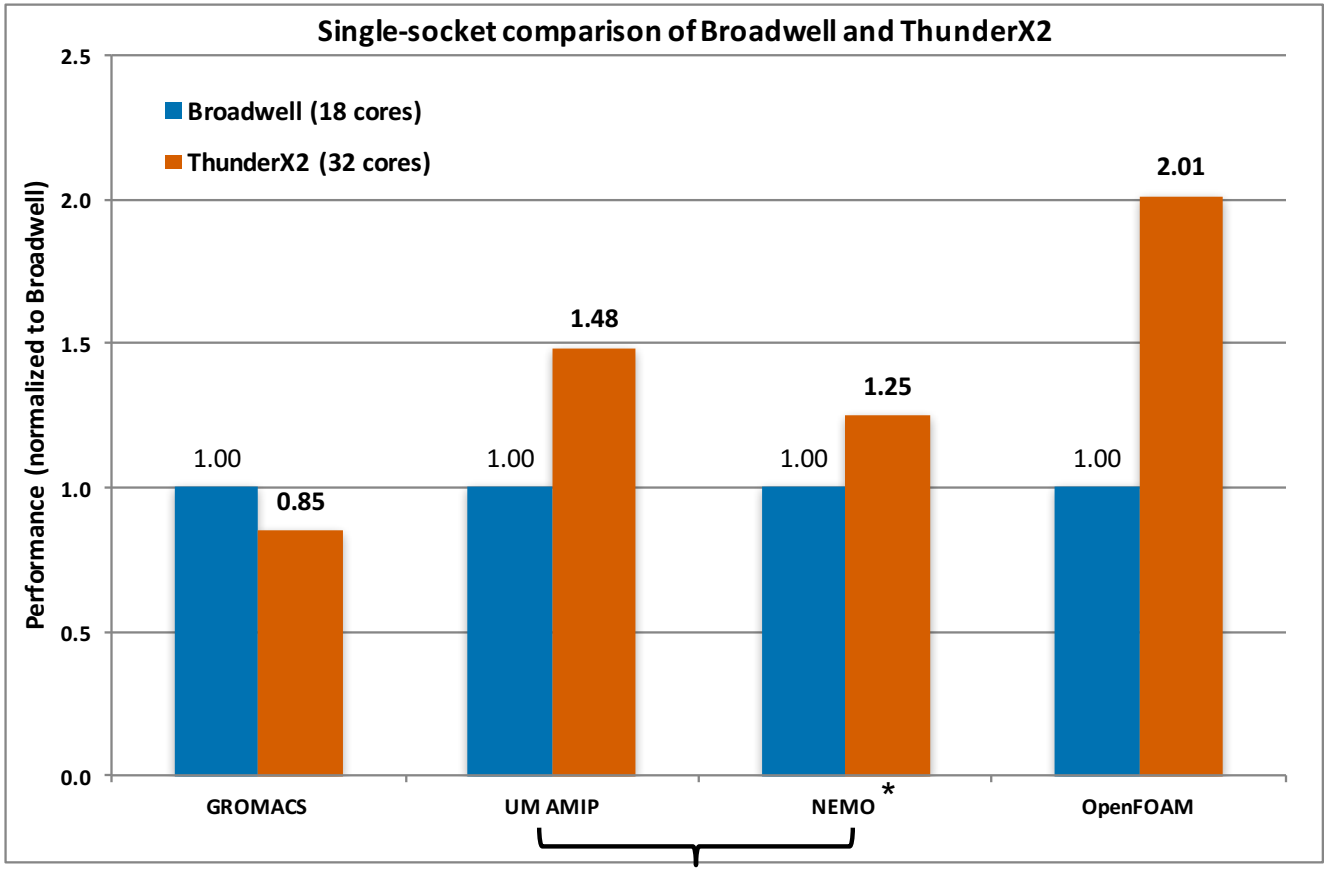

Benchmarked by the UK's Met Office

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

Arm Performance Libraries

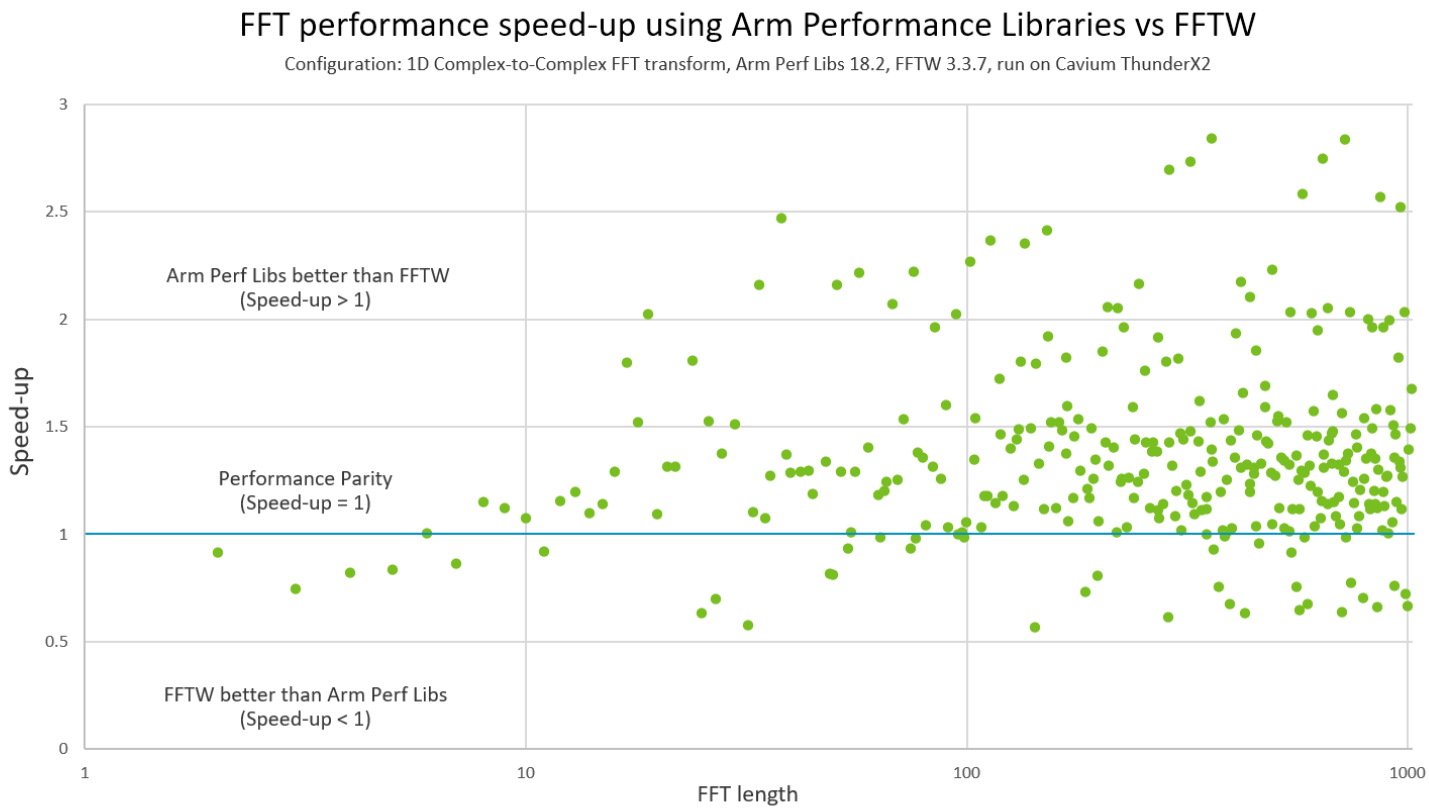

Scalable Vector Extension (SVE)

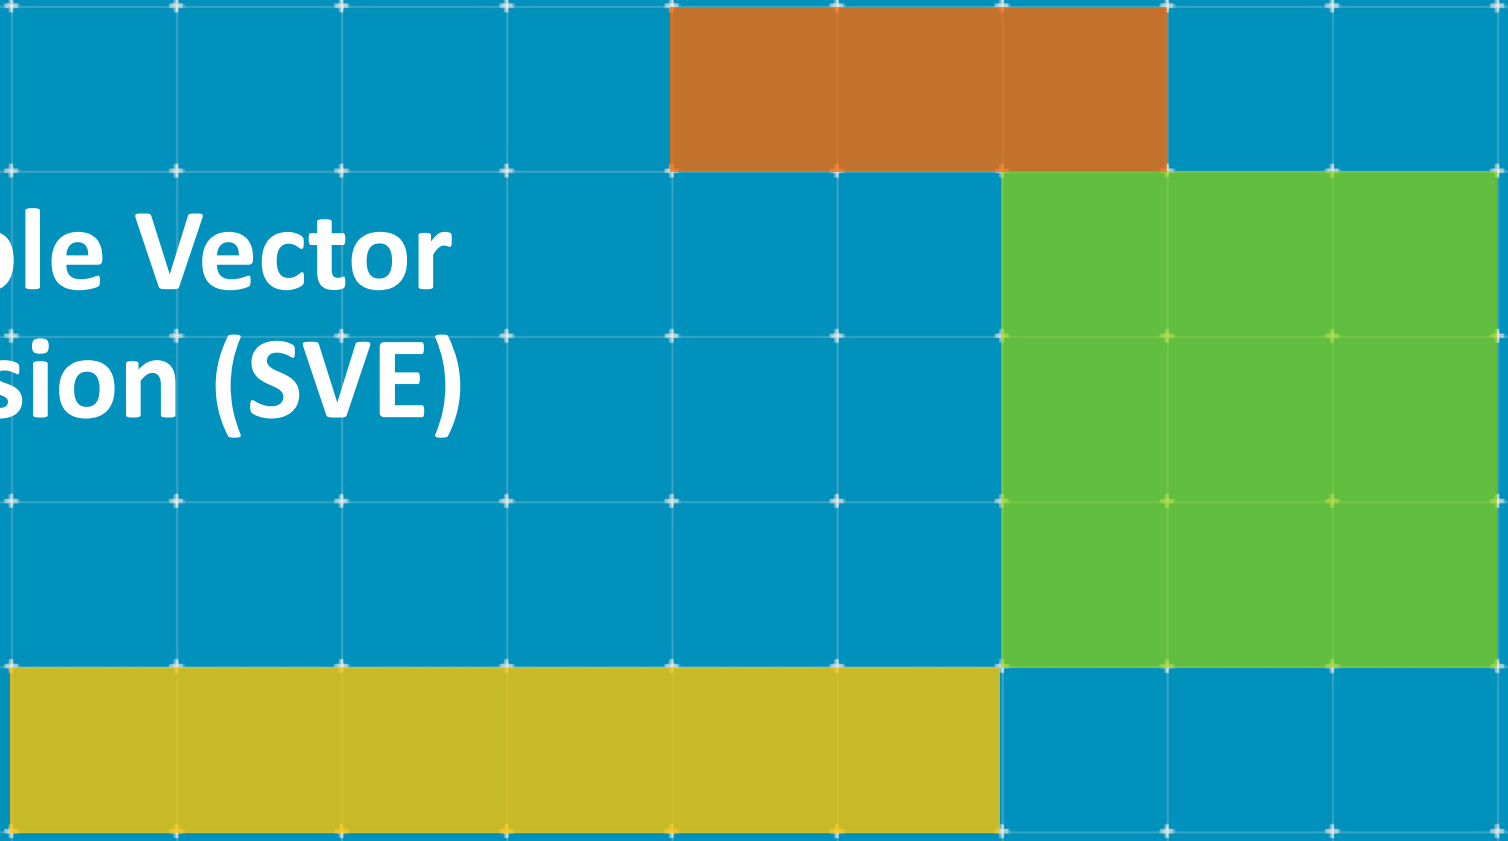

arm

Scalable Vector Extension (SVE)

A vector extension to the Armv8-A architecture with some major new features:

Gather-load and scatter-store

Loads a single register from several non-contiguous memory locations

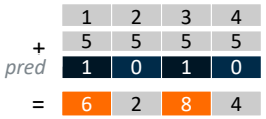

Per-lane predication

Operations work on individual lanes under control of a predicate register

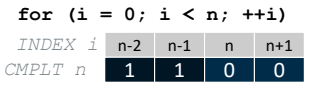

Predicate-driven loop control and management

Eliminate scalar loop heads and tails by processing partial vectors

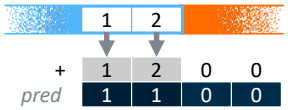

Vector partitioning and software-managed speculation

First Faulting Load instructions allow memory accesses to cross into invalid pages

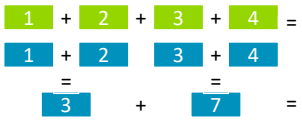

Extended floating-point horizontal reductions

In-order and tree-based reductions trade-off performance and repeatability

arm

What's the vector length?

There is **no** preferred vector length

- Vector Length (VL) is the CPU implementor's choice, from 128 to 2048 bits, in increments of 128
- Adopting a Vector Length Agnostic (VLA) code generation style makes code portable across all possible vector lengths
- VLA is made possible by the per-lane predication, predicate-driven loop control, vector partitioning and software-managed speculation features of SVE
- No need to recompile, or to rewrite hand-coded SVE assembler or C intrinsics

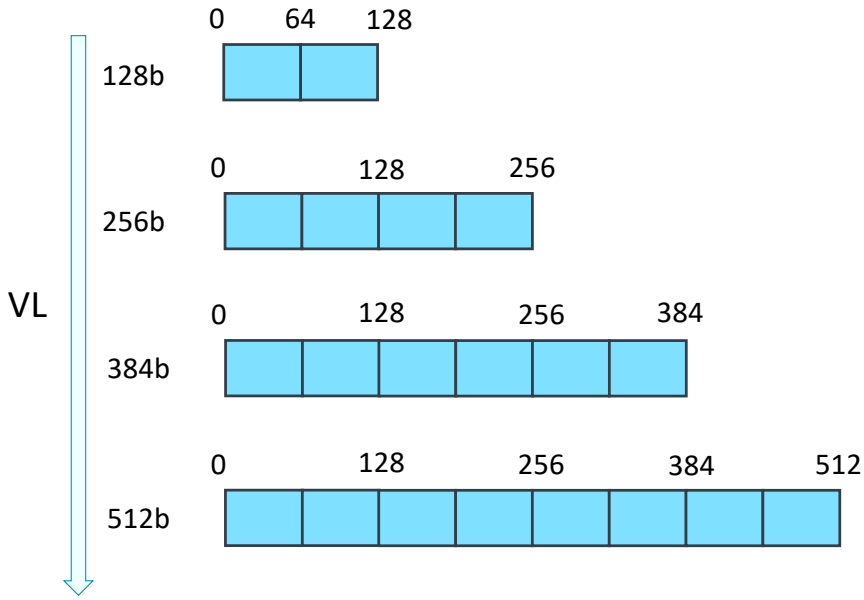

https://developer.arm.com/hpc/resources

arm

More on SVE…

SVE is **not** an extension of Advanced SIMD

SVE is a separate architectural extension with a new set of AArch64 instruction encodings

Does not offer a superset of functionality; focus is HPC scientific workloads, media/image processing, etc.

Further Information

Full ISA Specification due to be released end of March 2017

Numerous worked examples in A sneak peek into SVE and VLA programming from https://developer.arm.com/hpc/resources

Evaluating SVE

with Arm Compiler, Instruction Emulator and Code Advisor

Compile → **Emulate** → **Analyse**

Arm SVE Compiler

C/C++ now, Fortran in 2017

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.

Instruction Emulator

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

Supported instructions run unmodified.

Unsupported instructions are trapped and emulated.

Code Advisor

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