



**Chris Goodyer**

Senior Engineering Manager  
HPC Software

@GoingARM  
22 June 2017

# ARM-supported HPC tools

# Range of SoCs addressing infrastructure

2



**XILINX**  
ZYNQ  
UltraSCALE+

**NXP**  
QorIQ®  
Layerscape  
2080A

**Mellanox**  
TECHNOLOGIES  
BlueField

socionext™  
SC2A11

**apm** applied  
micro®  
X-Gene 3™

**QUALCOMM**  
Centriq 2400

**ALTERA**  
Stratix®10  
FPGA • SoC

**CAVIUM**  
THUNDERX2

One size does not fit all

# Serious ARM HPC deployments starting in 2017

3

Two big announcements about ARM in HPC in Europe



## Bull Atos to Build HPC Prototype for Mont-Blanc Project using Cavium ThunderX2 Processor

January 16, 2017 by [staff](#)

Today the [Mont-Blanc European project](#) announced it has selected Cavium's ThunderX2 ARM server processor to power its new HPC prototype.

The new Mont-Blanc prototype will be built by [Atos](#), the coordinator of phase 3 of Mont-Blanc, using its Bull expertise and products. The platform will leverage the infrastructure of the Bull sequana pre-exascale supercomputer range for network, management, cooling, and power. Atos and Cavium signed an agreement to collaborate to develop this new platform, thus making Mont-Blanc an Alpha-site for ThunderX2.




# GW4

January 17th 2017

## Announcing the **GW4 Tier 2 HPC service, 'Isambard'**: named after Isambard Kingdom Brunel

### System specs:

- Cray CS-400 system
- **10,000+** ARMv8 cores
- HPC optimised software stack
- Technology comparison:
  - x86, KNL, Pascal
- To be installed March-Dec 2017
- £4.7m total project cost over 3 years



I.K.Brunel 1804-1859

Simon McIntosh Smith, [simonm@cs.bris.ac.uk](mailto:simonm@cs.bris.ac.uk),  
[@simonmcs](#)

5

[bristol.ac.uk](http://bristol.ac.uk)

# ARM HPC software ecosystem

# Ecosystem for HPC

---

- List of components needed:
- Linux OS availability
- Compilers
- Libraries
- Debuggers
- Profilers
- Job schedulers

Mix of open source and commercial products and applications...

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
- ARM is on the OpenHPC Technical Steering Committee in order to drive ARM build support

Status: 1.3.1 release out now

- All packages built on ARMv8 for CentOS and SUSE
- ARM-based machines are being used for building and also in the OpenHPC build infrastructure

Functional Areas	Components include
Base OS	RHEL/CentOS 7.1, SLES 12
Administrative Tools	Conman, Ganglia, Lmod, LosF, ORCM, Nagios, pdsh, prun
Provisioning	Warewulf
Resource Mgmt.	SLURM, Munge, Altair PBS Pro*
I/O Services	Lustre client (community version)
Numerical/Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, Mumps
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran)
MPI Families	OpenMPI, MVAPICH2
Development Tools	Autotools (autoconf, automake, libtool), Valgrind, R, SciPy/NumPy
Performance Tools	PAPI, Intel IMB, mpiP, pdtoolkit TAU

# Open source in the ARM HPC ecosystem

7

- Over the past 12 months many more packages and applications have been successfully ported to ARM HPC



# Commercial HPC products simplify the ecosystem

---

8



## Comprehensive

- Comprehensive suite of tools – compiler, libraries, debuggers and profilers



## Performant

- Best in class performance with latest features
- Tuned for a wide range of 64-bit ARMv8-A-based platforms



## Supported

- Commercially supported by ARM



# ARM commercial HPC software portfolio

---

9

## ARM HPC Compilers

COMMERCIALY SUPPORTED  
FORTRAN, C AND C++

## ARM Performance Libraries

BLAS, LAPACK and FFT  
MICRO-ARCHITECTURALLY TUNED



## Allinea Forge (DDT+MAP)

PARALLEL DEBUGGING and PROFILING

## Allinea Performance Reports

PERFORMANCE SUMMARY

# ARM C/C++/Fortran Compiler

10



Commercially supported  
by ARM



Latest features and  
performance optimizations




Optimized OpenMP

- Linux user-space compiler tailored for HPC on ARM
  - Maintained and supported by ARM for a wide range of ARM-based SoCs running leading Linux distributions
  - Based on LLVM, the leading compiler framework
- Latest features go into the commercial releases first
  - Ahead of upstream LLVM by up to an year with latest performance improvement patches
  - SVE support in the assembler, disassembler, intrinsics and autovectorizer
- OpenMP
  - Uses latest open source (now ARM-optimized) LLVM OpenMP runtime
  - Changes pushed back to the community

# OpenMP

## Scalability of LULESH benchmark: GNU libgomp vs ARM-tuned LLVM OpenMP



# ARM Performance Libraries

12

Optimized BLAS, LAPACK and FFT

## Commercial 64-bit ARMv8 math libraries

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



Performance on par  
with best-in-class math libraries

## Best-in-class performance with commercial support

- Tuned by ARM for Cortex-A72, Cortex-A57 and Cortex-A53
- Maintained and supported by ARM for a wide range of ARM-based SoCs
  - Including Cavium ThunderX and ThunderX2 CN99 cores



Commercially Supported  
by ARM

## Silicon partners can provide tuned micro-kernels for their SoCs

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



Validated with  
NAG test suite


# ARM Performance Libraries

13

## Micro-architectural tuning

- ARM cores have a variety of designs, created by both ARM and our partners
- ARM Performance Libraries are creating tailored versions of routines to target these different *micro-architectures*
- It is important to ensure that the correct version is installed on your system

HPE Comanche - Advanced Technology Preview  
DGEMM – 1 thread on Cavium ThunderX2 CN99




# ARM Performance Libraries

14

## Micro-architectural tuning

- ARM cores have a variety of designs, created by both ARM and our partners
- ARM Performance Libraries are creating tailored versions of routines to target these different *micro-architectures*
- It is important to ensure that the correct version is installed on your system

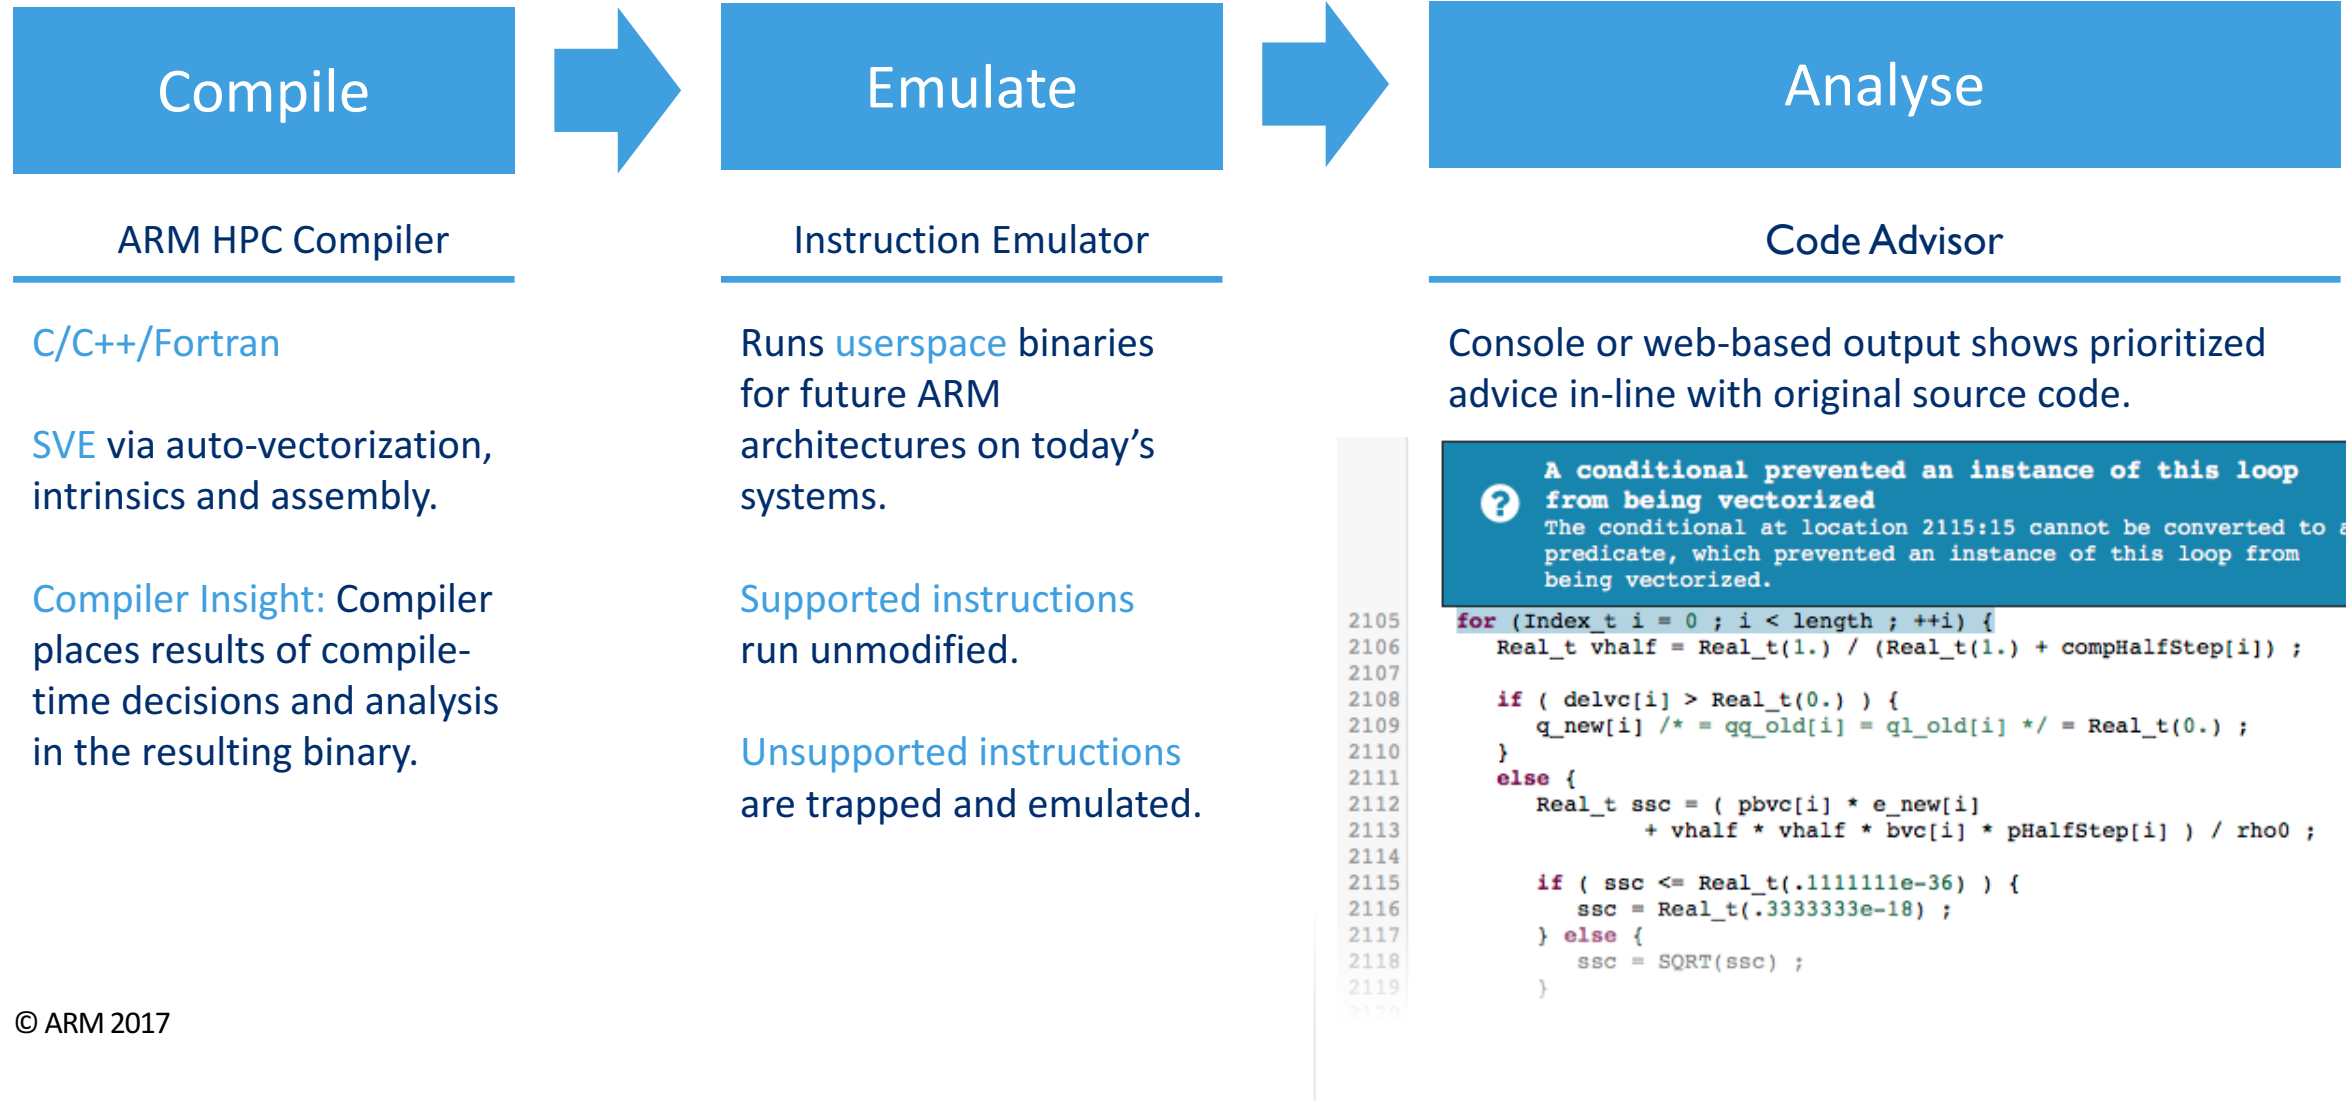
HPE Comanche - Advanced Technology Preview  
DGEMM – 56 threads on Cavium ThunderX2 CN99



# Experimental tools to support SVE

15

With ARM Compiler, Instruction Emulator and Code Advisor



# Introducing the *Compute Library*

16

- Optimized low-level functions for CPU and GPU
  - Most popular CV and ML functions
  - Supports common ML frameworks
- Enable faster deployment of CV and ML
  - Targeting CPU (NEON) and GPU (OpenCL)
  - Significant performance uplift compared to OSS alternatives
- Publicly available now (no fee, MIT license)

CV = Computer Vision  
ML = Machine Learning

## Key Functions categories

Basic arithmetic

Convolutions

Colour manipulation

Feature detection

Neural network

GEMM

Pyramids

Filters

Image reshaping

Mathematical functions



- ARM's ecosystem is built on partnership and choice
  - We work with many organizations to drive hardware design and deliver better software
  - This method enables partners to design different products for different markets
- We license IP at all levels of the stack to help customers be successful
- Our 64-bit server platforms are beginning to see large, main-stream deployments
- Building the software ecosystem and tools is an important part of this story
  - We enhance open source software as well as developing commercially supported options

# ARM

---

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

Copyright © 2017 ARM Limited

[www.arm.com/hpc](http://www.arm.com/hpc)  
[chris.goodyer@arm.com](mailto:chris.goodyer@arm.com)


# ARM Scalable Vector Extensions

Preparing for tomorrow's HPC needs

# Introducing the Scalable Vector Extension (SVE)

20

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

	1	2	3	4
+	5	5	5	5
pred	1	0	1	0
=	6	2	8	4

## Per-lane predication

Operations work on individual lanes under control of a predicate register

for (i = 0; i < n; ++i)				
INDEX i	n-2	n-1	n	n+1
CMPLT n	1	1	0	0

## Predicate-driven loop control and management

Eliminate scalar loop heads and tails by processing partial vectors

	1	2		
+	1	2	0	0
pred	1	1	0	0

## Vector partitioning and software-managed speculation

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

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


## Extended floating-point horizontal reductions

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

# 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



# 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, not media/image processing
- 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](https://developer.arm.com/hpc/resources) from <https://developer.arm.com/hpc/resources>