Andrei Poenaru

University of Bristol HPC Research Group

Isambard 2: Adding SVE to the World's First Production ARMv8 Supercomputer

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4

Introduction to the GW4 Isambard 2 supercomputer

- **Isambard 2** is a £4.1M EPSRC project, run by a consortium of the GW4 Alliance, the Met Office, HPE/Cray, Fujitsu and Arm, to deliver a Tier-2 HPC service to researchers across the UK and around the world
- Funded in late 2019, Isambard 2 builds on Isambard 1's achievements as the world's first Arm64-based production supercomputer
- **Isambard 1** has been a huge success, proving for the first time that Arm works for supercomputing in production environments

Isambard 2 production system

- **21,504** Armv8 cores (336n x 2s x 32c)
  - **Marvell ThunderX2 32 core @2.5GHz**
- Cray XC50 ‘Scout’ form factor
- High-speed **Aries** interconnect
- Cray HPC optimised software stack
  - Compilers, math libraries, CrayPAT, …
  - Also comes with all the open source software toolchains: GNU, Clang/LLVM etc.
- Multi-Architecture Comparison System
- Hosted for the Consortium by the Met Office in Exeter

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4

Isambard 2's A64FX Apollo80 system

- Isambard 2 includes A64FX CPUs from Fujitsu
  - 72 nodes connected with 100 Gbps IB
    - 3,456 cores, 72 TB/s memory bandwidth, 202 TFLOP/s 64-bit
  - Comes with a Cray software stack
    - CCE, GNU, Arm Compiler
  - Fujitsu compiler coming soon

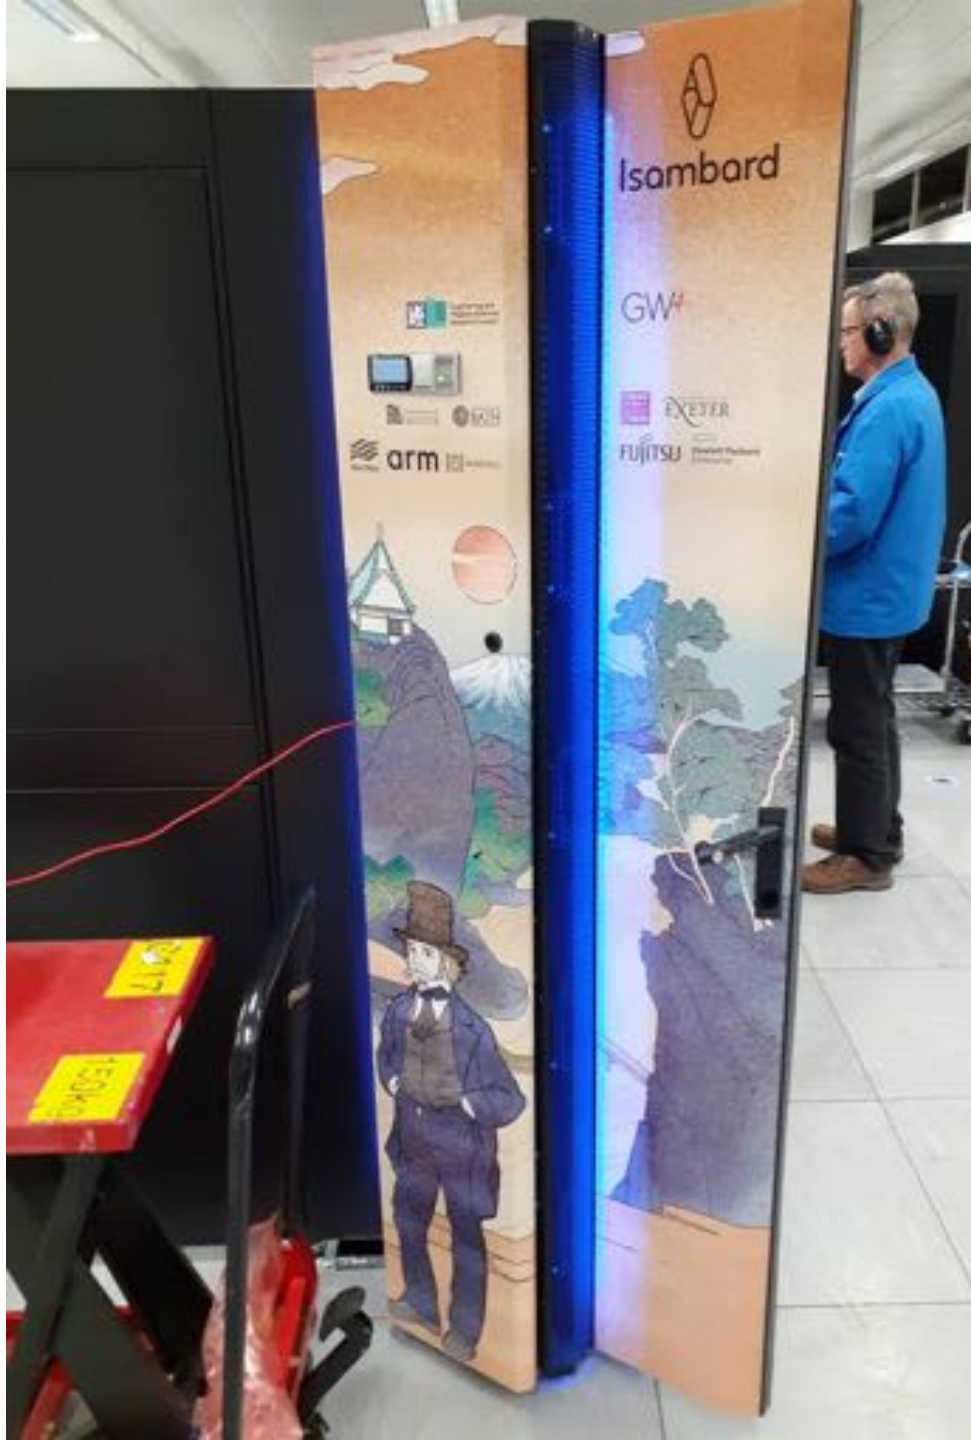

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

Fujitsu's A64FX

http://www.fujitsu.com/global/about/resources/news/press-releases/2018/0822-02.html

https://github.com/fujitsu/A64FX

- 48 cores, 1.8 – 2.2GHz
  - 4 CMGs
- >2.7 TFLOP/s double precision
- 2x 512-bit vector pipelines per core
  - ARMv8.3-A + SVE
- 1 TByte/s main memory bandwidth
  - 4 stacks of HBM2
- ~170 Watts
- High speed interconnect
- 8.7B transistors, 7 nm

FUJITSU

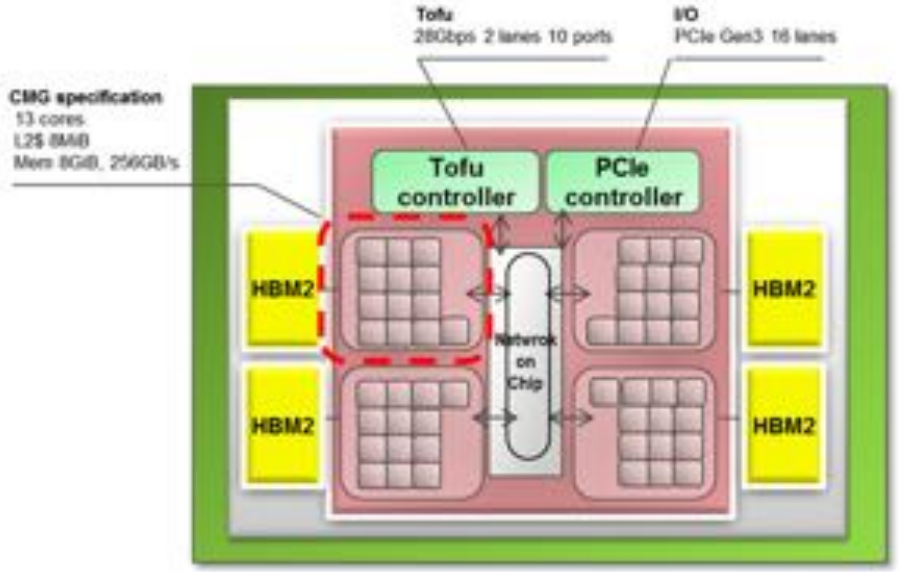

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4

A64FX benchmarking

Platform	Cores	Clock Speed	Peak FLOP/s (d.p.)	Peak memory BW
A64FX	48	2.2 GHz	3.4 TFLOP/s	1,024 GB/s
ThunderX2	2x32	2.5 GHz	1.3 TFLOP/s	320 GB/s
Skylake	2x28	2.1 GHz	2.8 TFLOP/s	256 GB/s
Cascade Lake	2x20	2.1 GHz	2.0 TFLOP/s	282 GB/s

Best compilers and optimization flags used in every case

- For A64FX this was usually the Fujitsu compiler
- For ThunderX2 this was usually the Cray compiler
- For SKL and CLX this was usually the Intel compiler

Where available, flat MPI, flat OpenMP and hybrid MPI+OpenMP were all tested, with the best result reported.

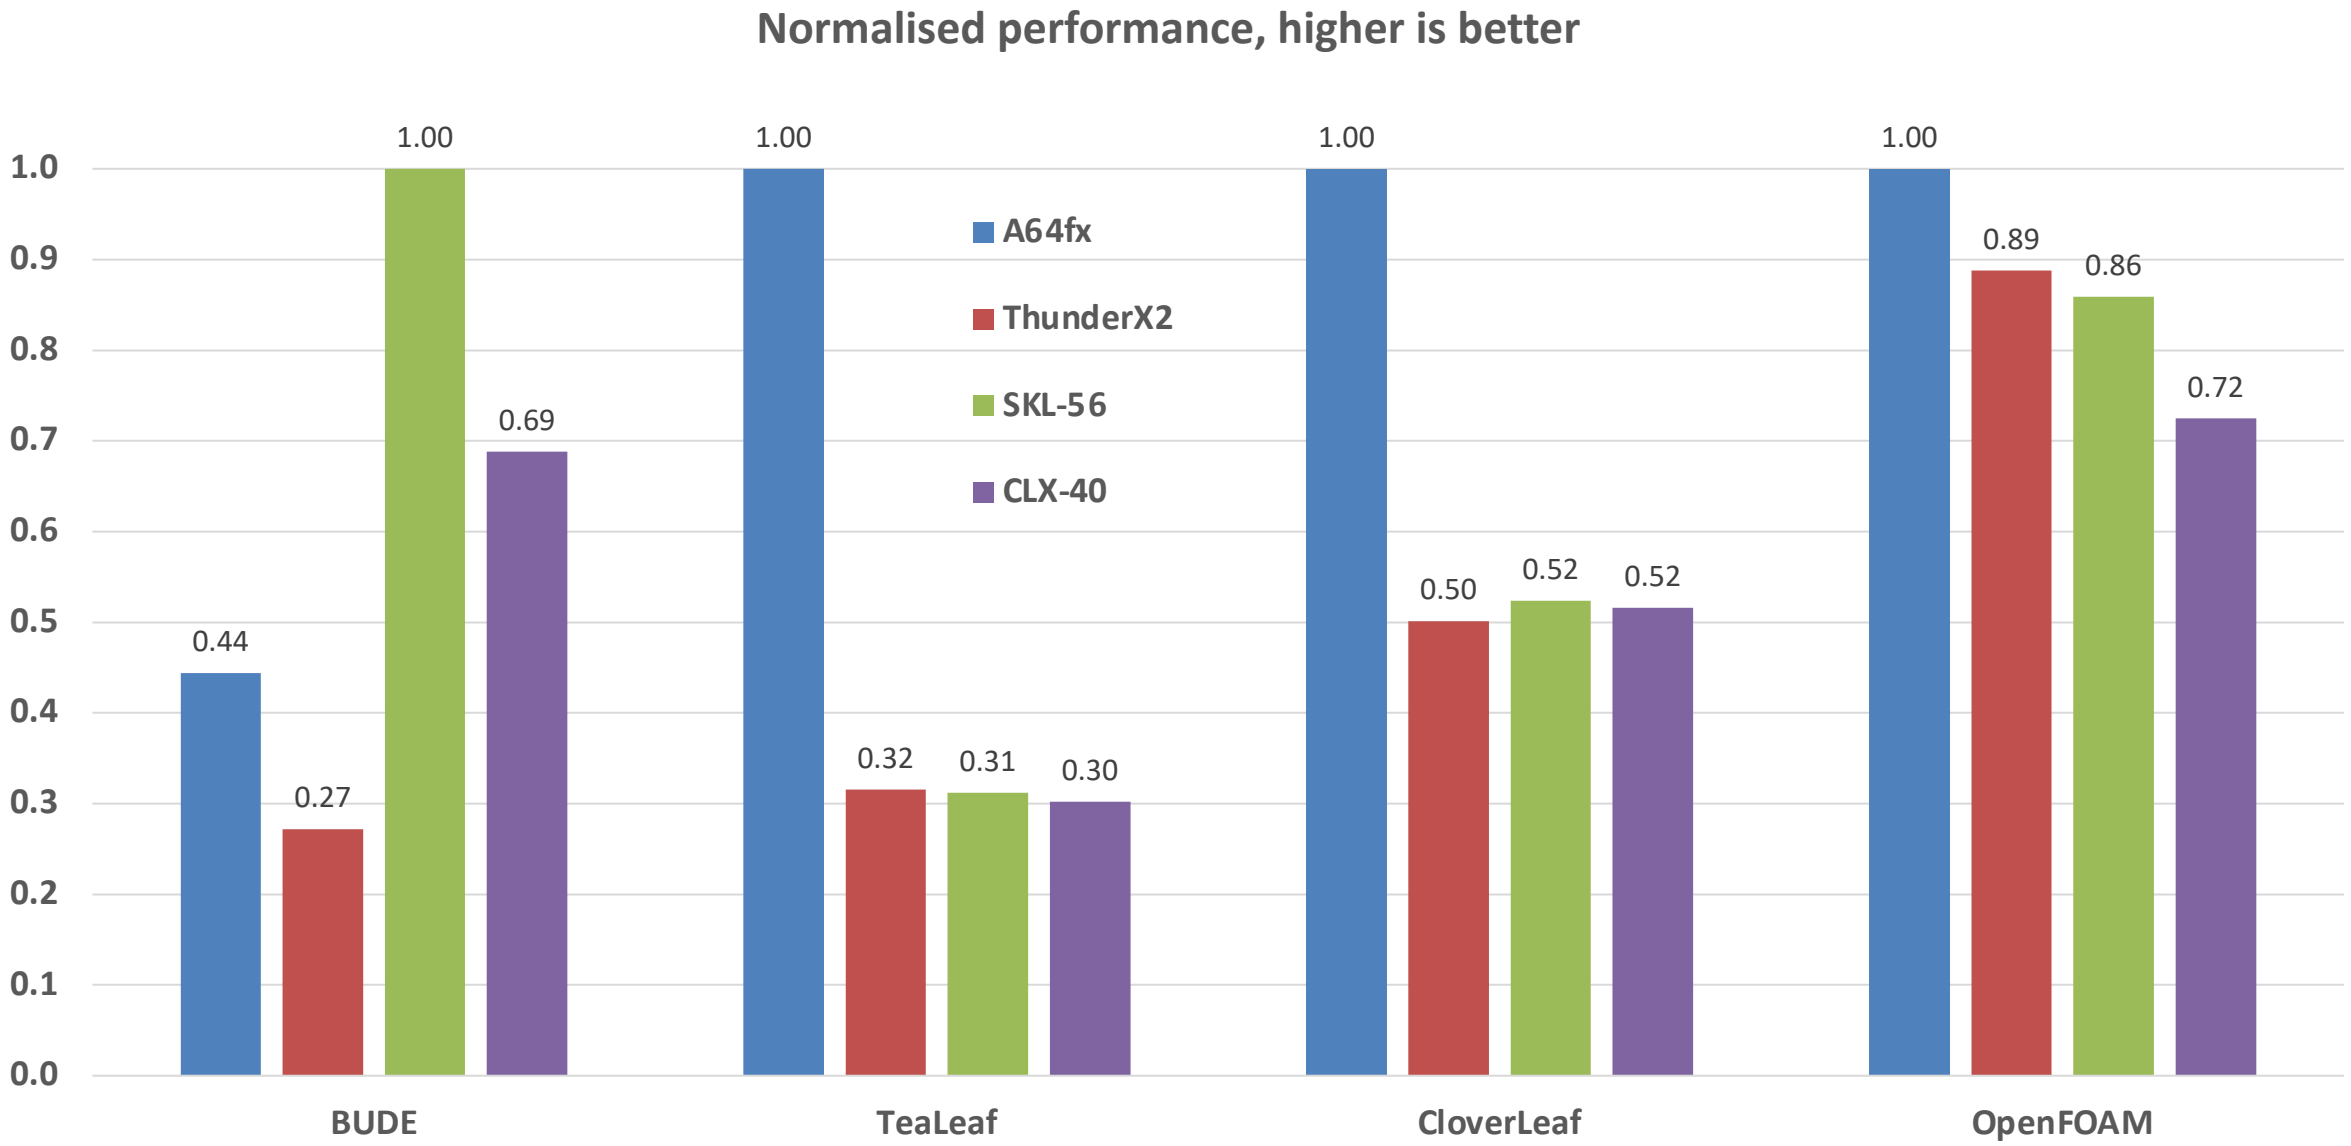

Performance normalised to the fastest processor for each application

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4

BUDE performance across programming models, normalised to OpenMP

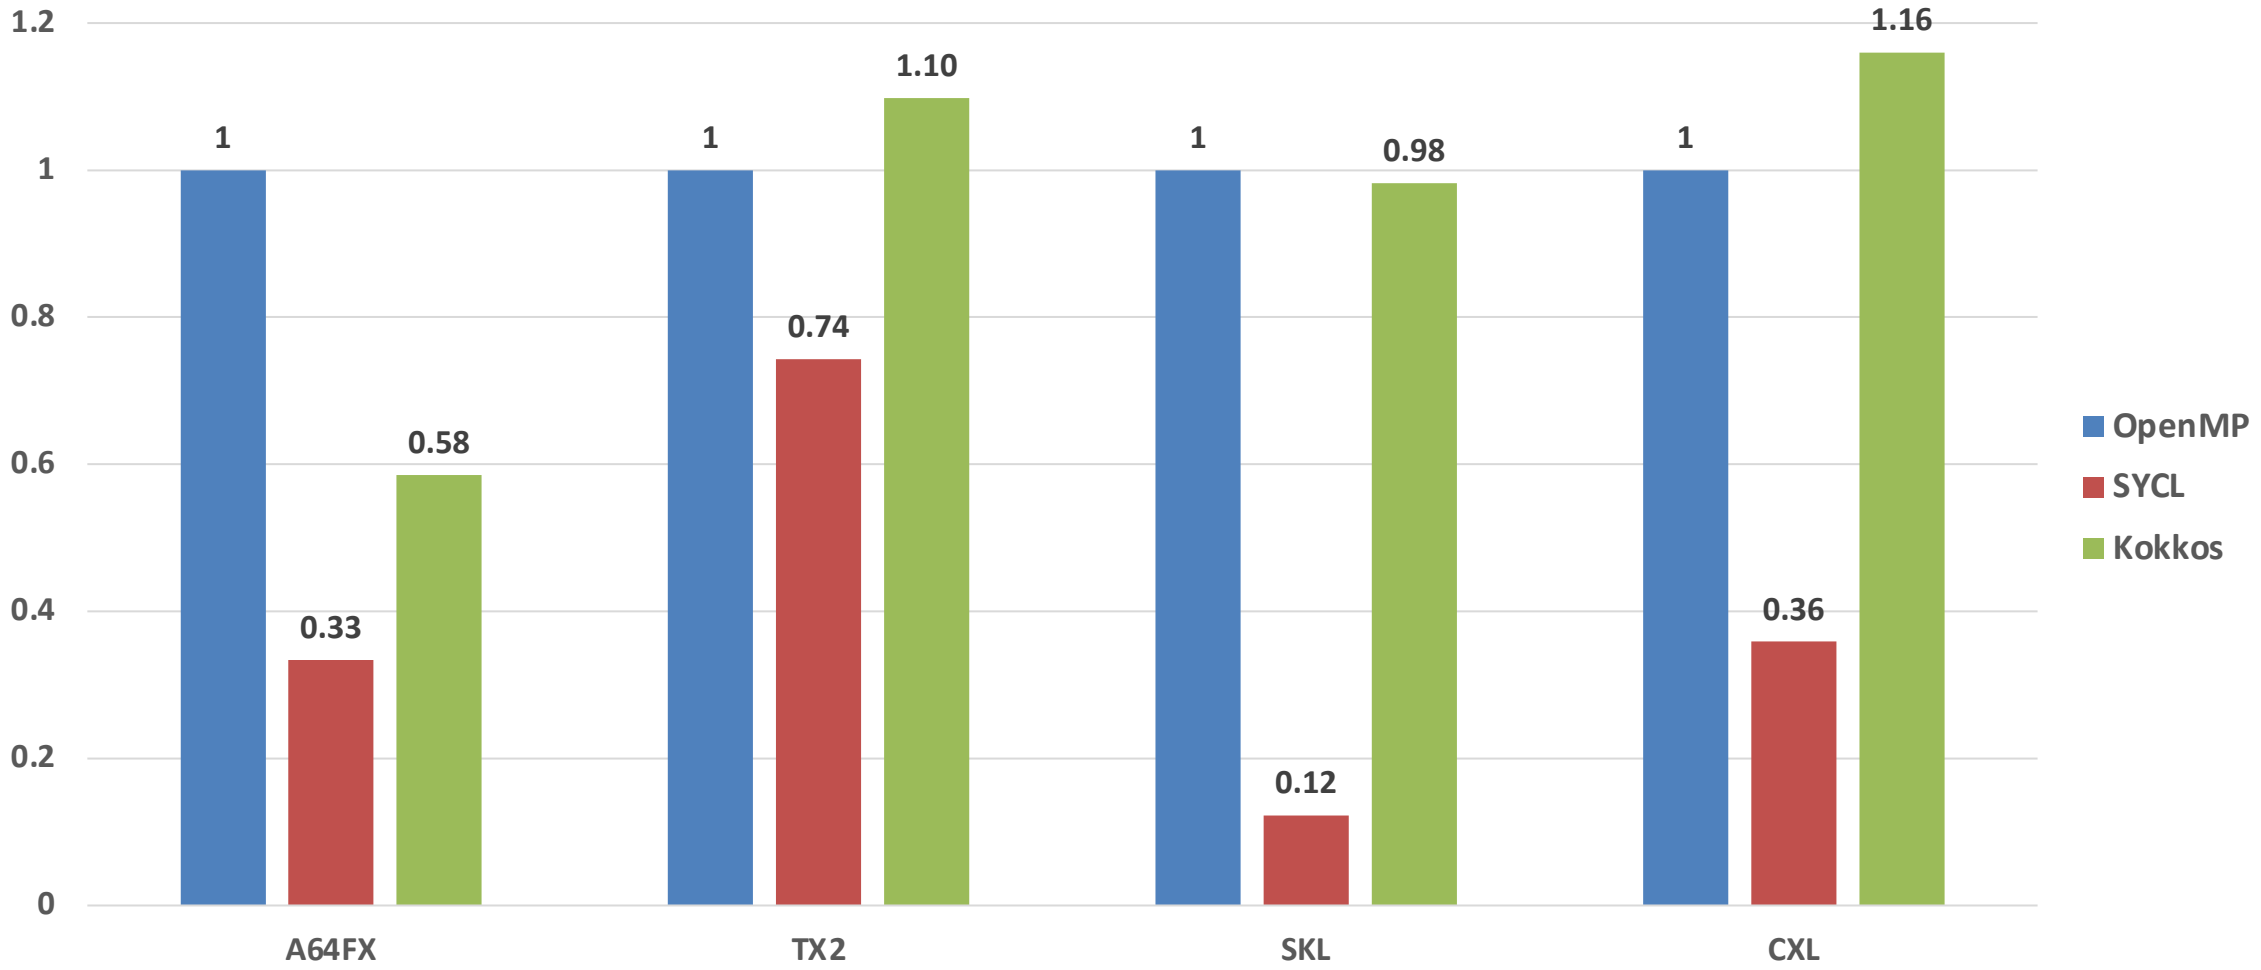

Early results: SYCL and Kokkos not yet optimised for A64FX

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4

Early Experience with Workloads on Isambard 2

- Everything works out-of-the-box!
  - The same experience we had with ThunderX2 in Isambard 1
  - No specific programming model or language needed
- Compiler support and libraries are already available
  - Cray, Arm, Fujitsu support A64FX
  - GCC support in 11 (experimental)
- Optimised libraries and higher-level frameworks are continuously being improved
- More full scale benchmarks soon, e.g. GROMACS or NAMD

Working with the Fujitsu A64FX

- No difference from working with other general-purpose CPUs
- 4 CMGs (NUMA nodes)
  - Core binding is particularly important
  - Some applications benefit from running 4 MPI ranks/node
- Out-of-Order architecture benefits from software pipelining and optimised instruction scheduling
  - Use a compiler with a good cost model
- There is a configurable “sector cache”…

It's easy to apply for time on Isambard

- Please contact the Isambard PI, Prof. Simon McIntosh-Smith simonm@cs.bris.ac.uk, who will help you determine if Isambard will work for you. If it will, applying for an account is quick and easy.

- Small amounts of pump-priming time are available for free, to try porting, optimizing for Arm etc.

- Larger amounts of time for real science runs can be applied for via the regular EPSRC "Access to HPC" calls, or via some CCPs.

Summary

- A64FX already looks very promising, beating cutting-edge dual-socket nodes in most tests so far
- Easy to use – in most cases running unmodified flat MPI, or hybrid MPI+OpenMP
- Performance is similar to GPUs, but with a significantly lower barrier to entry in ease of use
- Isambard 2 makes most of the major technologies available in one place, enabling rigorous comparative benchmarking

Thank you

- https://uob-hpc.github.io

- https://github.com/UoB-HPC/benchmarks
- https://github.com/UoB-HPC/performance-portability

University of BRISTOL

UKRI Engineering and Physical Sciences Research Council

GW4