


Dr. James Price

University of Bristol /
GW4 Alliance



Isambard: tales from the world's first Arm-based production supercomputer

'Isambard' is a new UK Tier 2 HPC service from GW4



Isambard Kingdom Brunel
1804-1859



The tiered model of HPC provision

Tier 0: international



Tier 1: national



Tier 2: regional

TIER 2 HPC CENTRES

Edinburgh	Cambridge	UCL
Loughborough	Bristol	Oxford


Tier 3



Why explore Arm-based supercomputers?

- Need to ensure that there is sufficient competition for CPUs in HPC procurements
- The architecture development is driven by the *fast-growing mobile space*
- Multiple vendors of Arm-based CPUs:
 - Greater **competition**
 - More **choice**
 - Exciting **innovations**, e.g. in vector instruction set

CRAY

 Isambard

EPSRC
Pioneering research
and skills

GW4

 UNIVERSITY OF
BATH


 University of
BRISTOL

CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD

UNIVERSITY OF
EXETER


Met Office


arm



Isambard system specification


- **10,752** Armv8 cores (168 x 2 x 32)
 - **Cavium ThunderX2 32 core @ 2.1GHz**
 - **256 GB DDR4 memory per node**
 - **500 TB Lustre filesystem**
- Cray XC50 'Scout' form factor
- High-speed **Aries** interconnect
- Cray HPC optimised software stack
 - **CCE, Cray MPI, math libraries, CrayPAT, ...**
- Technology comparison (Phase 1):
 - **x86, Xeon Phi, POWER9, Pascal/Volta GPUs**
- Phase 2 accepted 9th November 2018
- 25% of machine time allocated to EPSRC users via RAP





Cavium ThunderX2, a seriously beefy CPU

- 32 cores at up to 2.5GHz
- Each core is 4-way superscalar, Out-of-Order
- 32KB L1, 256KB L2 per core
- Shared 32MB L3
- Dual 128-bit wide NEON vectors
 - Compared to Skylake's 512-bit vectors, and Broadwell's 256-bit vectors
- 8 channels of 2666MHz DDR4
 - Compared to 6 channels on Skylake, 4 channels on Broadwell
 - AMD's EPYC also has 8 channels




Benchmarking platforms

Processor	Cores	Clock speed GHz	TDP Watts	FP64 TFLOP/s	Bandwidth GB/s
Broadwell	2 × 22	2.2	145	1.55	154
Skylake Gold	2 × 20	2.4	150	3.07	256
Skylake Platinum	2 × 28	2.1	165	3.76	256
ThunderX2	2 × 32	2.2	175	1.13	320

- BDW 22c** Intel Broadwell E5-2699 v4, **\$4,115** each (near top-bin)
- SKL 20c** Intel Skylake Gold 6148, **\$3,078** each
- SKL 28c** Intel Skylake Platinum 8176, **\$8,719** each (near top-bin)
- TX2 32c** Cavium ThunderX2, **\$1,795 each** (near top-bin)

Key architectural comparisons (node-level, dual socket)



Isambard's core mission: deploying Arm in production HPC


Starting by porting/benchmarking/optimizing codes from the top 10 most heavily used on ARCHER:

- **VASP, CASTEP, GROMACS, CP2K, UM, NAMD, Oasis, SBLI, NEMO**
- Most of these codes are written in FORTRAN

Additional important codes for project partners:

- **OpenFOAM, OpenIFS, WRF, CASINO, LAMMPS, ...**

Single-node performance for top applications on ARCHER



Single-node performance summary

- Performance is competitive with contemporary Intel processors
 - ThunderX2 is **faster** when codes are dominated by memory bandwidth
 - ThunderX2 is **slower** when FLOP/s and L1 cache bandwidth is critical
- Next-gen Arm CPUs will increase FLOP/s + cache bandwidth
 - Introduction of SVE will allow vector width of up to 2048-bits
 - Fujitsu A64FX chip unveiled recently with 512-bit SVE
 - Expecting 512-bits to be a common choice for server chips

Arm software ecosystem


- Three mature compiler suites:
 - GNU (gcc, g++, gfortran)
 - Arm HPC Compilers (armclang, armclang++, armflang)
 - Cray Compiling Environment (CCE)
- Three mature sets of math libraries:
 - OpenBLAS + FFTW
 - Arm Performance Libraries (BLAS, LAPACK, FFT)
 - Cray LibSci + Cray FFTW
- Multiple performance analysis and debugging tools:
 - Arm Forge (MAP + DDT, formerly Allinea)
 - CrayPAT / perftools, CCDB, gdb4hpc, etc

Benchmark	ThunderX2	Broadwell	Skylake
STREAM	Arm 18.3	Intel 18	CCE 8.7
CloverLeaf	CCE 8.7	Intel 18	Intel 18
TeaLeaf	CCE 8.7	GCC 7	Intel 18
SNAP	CCE 8.6	Intel 18	Intel 18
Neutral	GCC 8	Intel 18	GCC 7
CP2K	GCC 8	GCC 7	GCC 7
GROMACS	GCC 8	GCC 7	GCC 7
NAMD	Arm 18.2	GCC 7	GCC 7
NEMO	CCE 8.7	CCE 8.7	CCE 8.7
OpenFOAM	GCC 7	GCC 7	GCC 7
OpenSBLI	CCE 8.7	Intel 18	CCE 8.7
UM	CCE 8.6	CCE 8.5	CCE 8.7
VASP	GCC 7.2	Intel 18	Intel 18

Comparison of compilers on Arm

Exact same issues on x86


	GCC	Arm	CCE
STREAM	97%	100%	99%
CloverLeaf	92%	95%	100%
TeaLeaf	99%	95%	100%
SNAP	74%	87%	100%
Neutral	100%	94%	85%
CP2K	100%	BUILD	CRASH
GROMACS	100%	91%	CRASH
NAMD	83%	100%	BUILD
NEMO	-	-	100%
OpenFOAM	100%	97%	BUILD
OpenSBLI	-	-	100%
Unified Model	84%	72%	100%



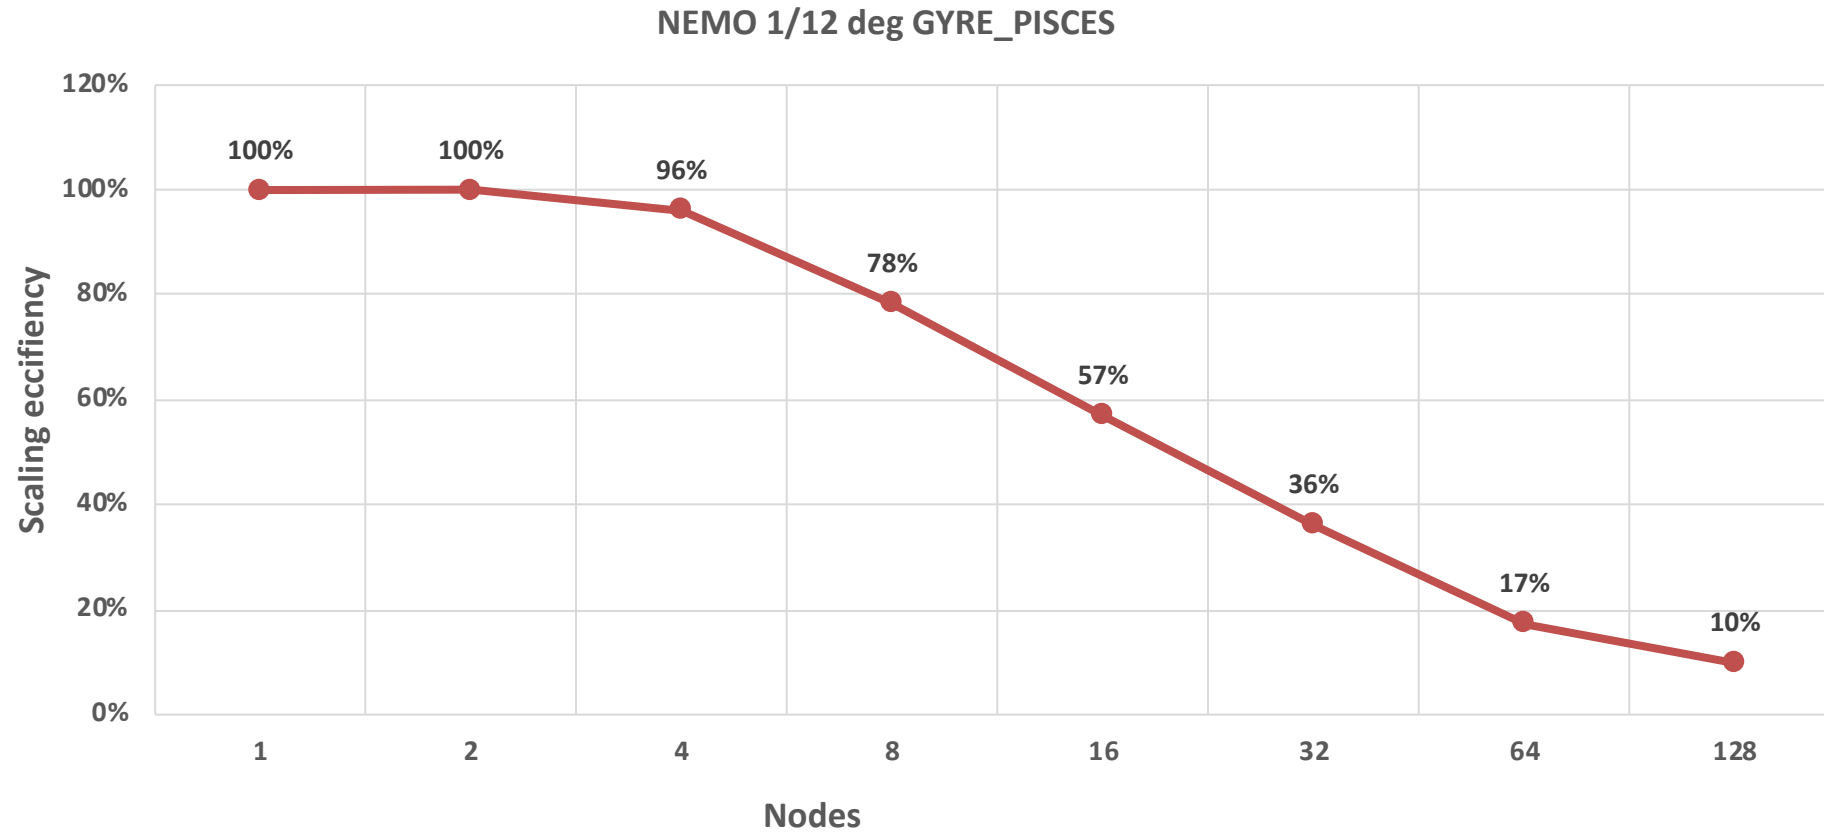
Scalability comparisons

- We've run some of the same applications at scale (up to **160 nodes / 10,240 cores**) on Isambard to test scalability for ThunderX2 with the Aries interconnect
- Results are plotted as '**Scaling efficiency**' versus one or two nodes
- These are early results, generated quickly in the first few days after acceptance with little time to tune scaling etc. We expect the results to improve even further as we continue to work on them


UM scalability, up to 10,240 cores



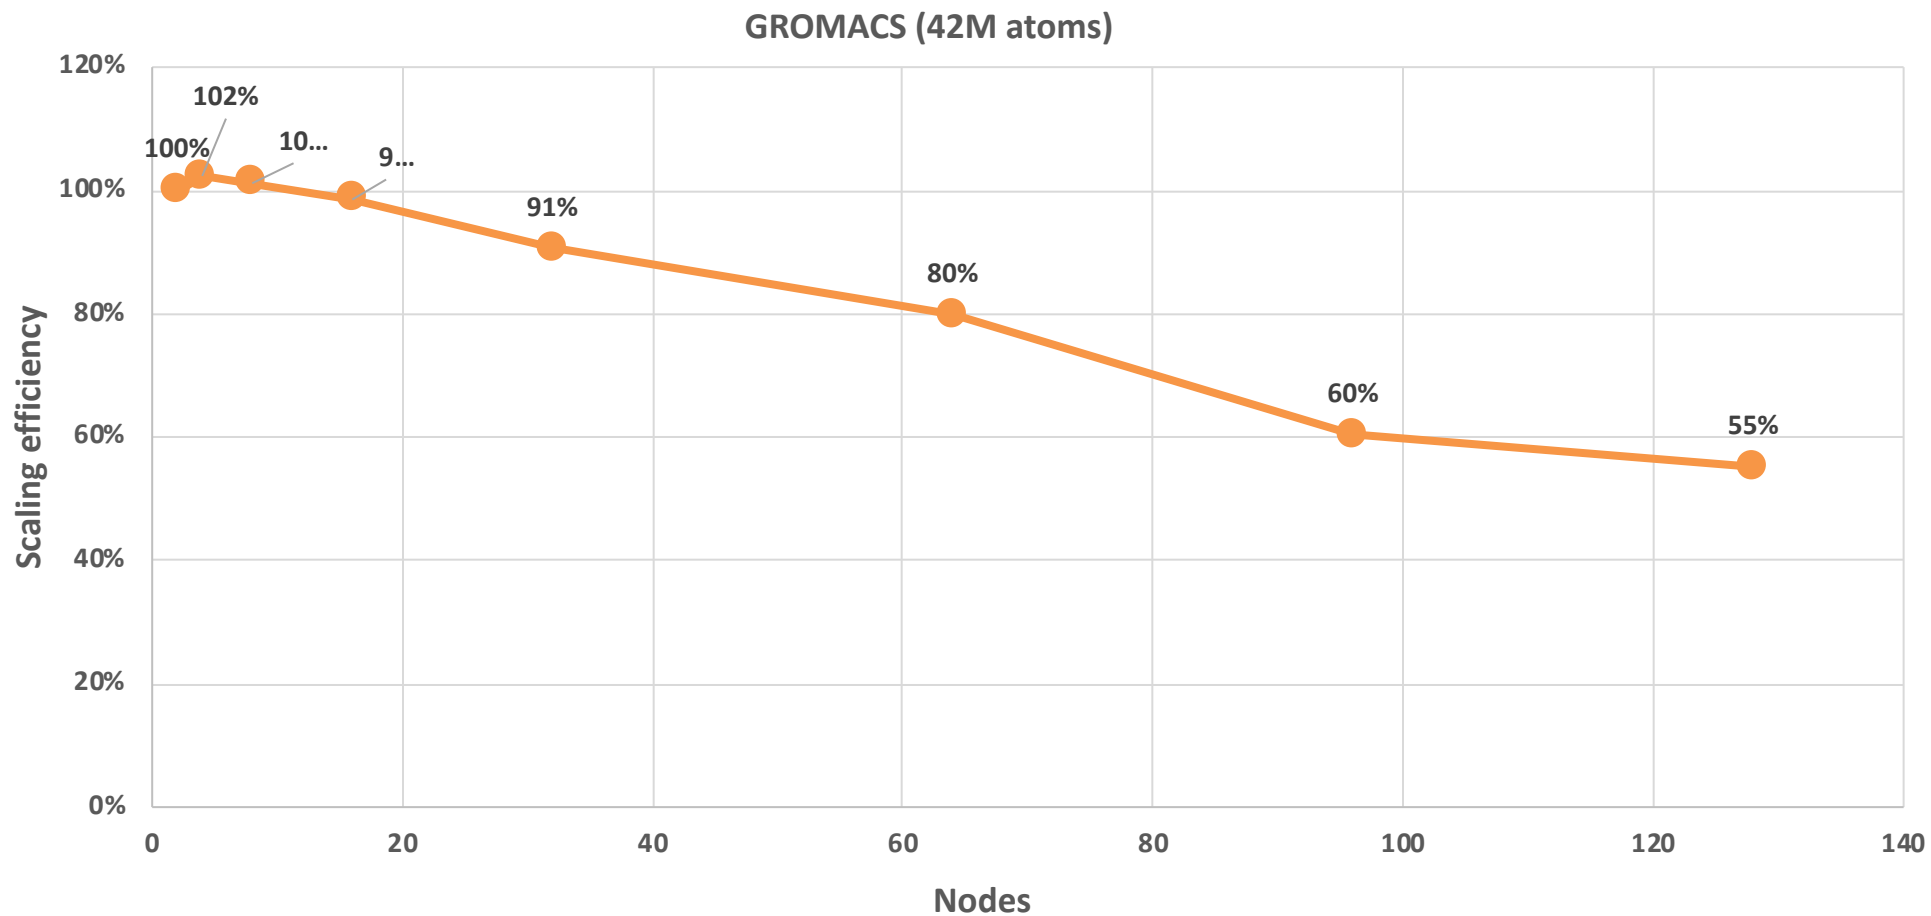
NEMO scalability, up to 8,192 cores



OpenSBLI scalability, up to 10,240 cores



GROMACS scalability, up to 8,192 cores



Conclusions

- Results show **ThunderX2 performance is competitive with current high-end server CPUs**, while **performance per dollar is compelling**
- **The software tools ecosystem is already in good shape**
- The Isambard system is now being configured for production service, and will soon be available to researchers across the UK
- The signs are that **Arm-based systems are now real alternatives for HPC**, reintroducing much needed competition to the market

For more information

Comparative Benchmarking of the First Generation of HPC-Optimised Arm Processors on Isambard

S. McIntosh-Smith, J. Price, T. Deakin and A. Poenaru, CUG 2018, Stockholm

<http://uob-hpc.github.io/2018/05/23/CUG18.html>

Bristol HPC group:

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

Isambard:

<http://gw4.ac.uk/isambard/>

Build and run scripts:

<https://github.com/UoB-HPC/benchmarks>