Prof Simon McIntosh-Smith
Isambard PI
University of Bristol /
GW4 Alliance

Scaling Results From Isambard: the First Generation of Arm-based Supercomputers

University of Bristol

EPSRC

GW4

Isambard system specification

- **10,752** Armv8 cores (168n x 2s x 32c)
  - **Cavium ThunderX2 32core 2.1→2.5GHz**
- Cray XC50 'Scout' form factor
- High-speed **Aries** interconnect
- Cray HPC optimised software stack
  - CCE, Cray MPI, math libraries, CrayPAT, …
- **Phase 2 (the Arm part):**
  - Delivered Oct 22nd, handed over Oct 29th
  - Accepted Nov 9th
  - Upgrade to final B2 TX2 silicon, firmware, CPE completed March 15th 2019

University of BRISTOL

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

GW4


CRAY
Isambard
EPSRC
Pioneering research
and skills
GW4
UNIVERSITY OF
BATH
University of
BRISTOL
CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD
UNIVERSITY OF
EXETER
Met Office

Cavium ThunderX2, a seriously beefy CPU

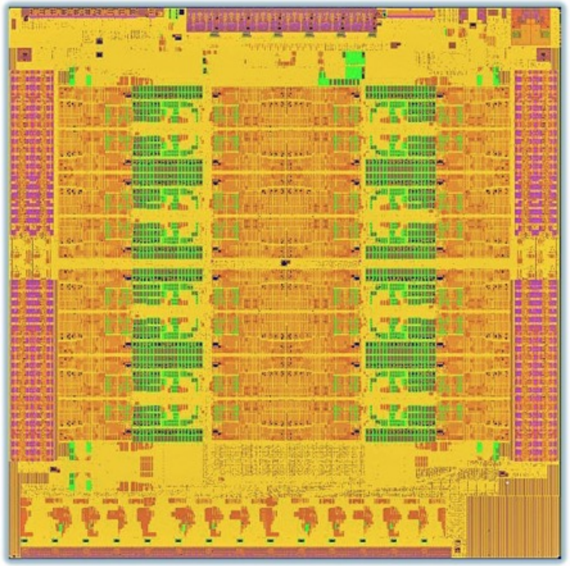

- 32 cores at up to 2.5GHz
- Each core is 4-way superscalar, Out-of-Order
- 32KB L1, 256KB L2 per core
- Shared 32MB L3
- Dual 128-bit wide NEON vectors
  - Compared to Skylake's 512-bit vectors, and Broadwell's 256-bit vectors
- 8 channels of 2666MHz DDR4
  - Compared to 6 channels on Skylake, 4 channels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

Recap of Single Node results from CUG 2018

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

SKL 20c Intel Skylake Gold 6148, **$3,078** each
TX2 32c Cavium ThunderX2, **$1,795 each** (near top-bin)

Previous single node performance results

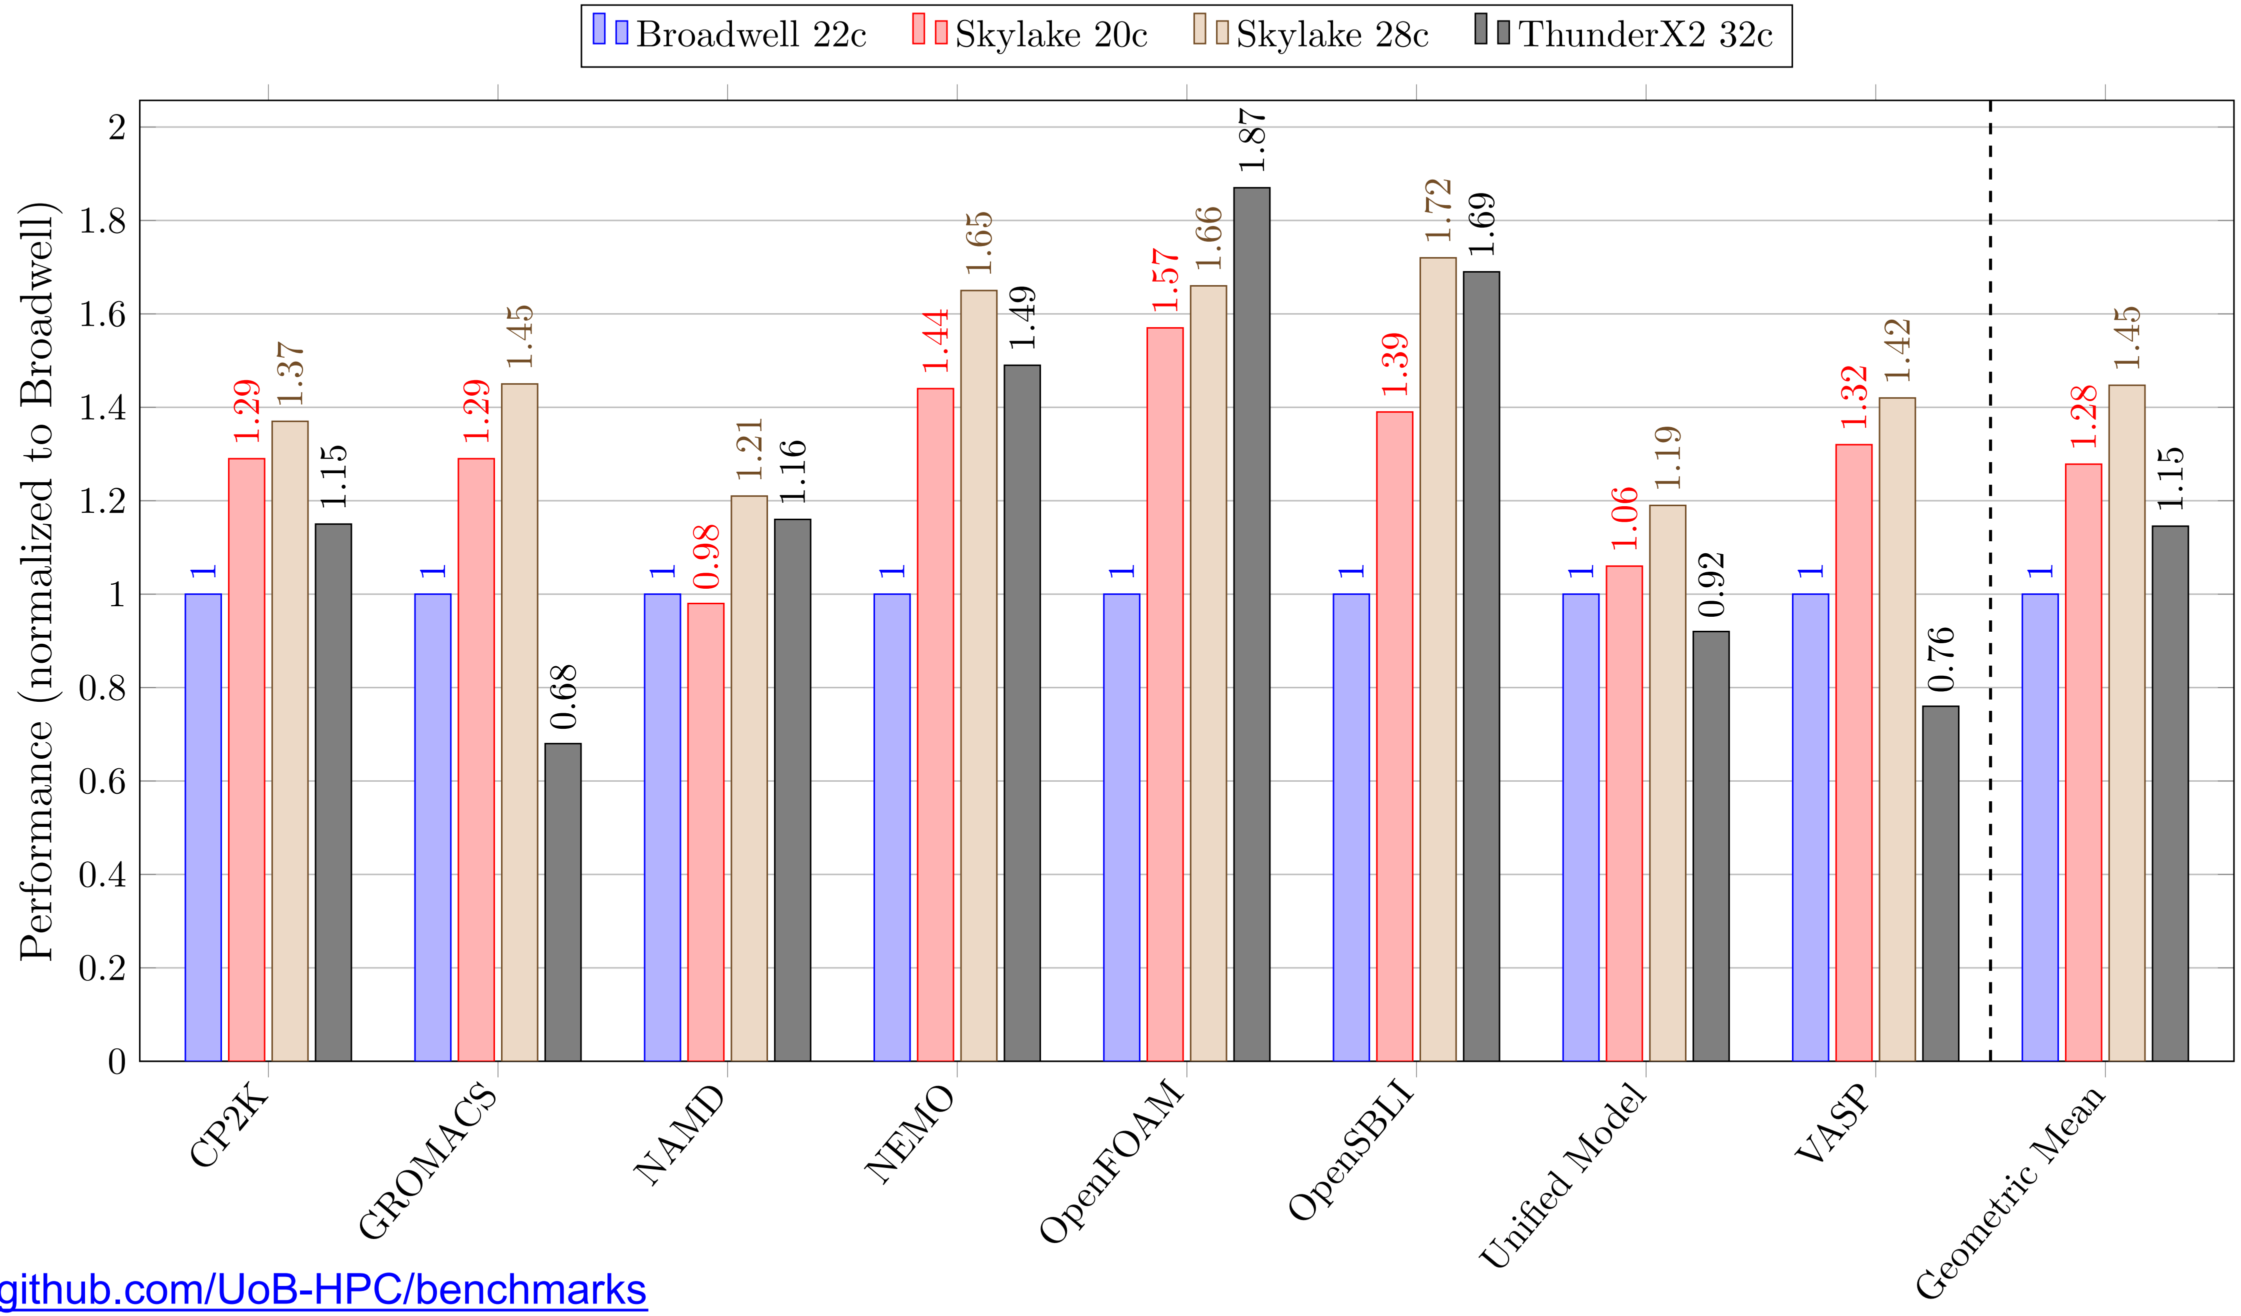

https://github.com/UoB-HPC/benchmarks

Scalability comparisons

- We've plotted results using 'Scaling (parallel) efficiency'
- We've compared against two x86-based XC50 systems:
  - Horizon using Intel Skylake Gold 6148 20-core CPUs at 2.4GHz
  - Swan using Intel Skylake Platinum 8176 28-core CPUs at 2.1GHz
  - Could only go up to 64 nodes on these systems, though we could have gone up to 164 on Isambard
- All the results are for **strong scaling**, except SNAP
- All of these systems use the same interconnect (Aries) and the same O/S and MPI library, so this is a good test of whether Arm-based ThunderX2 scales as well as x86

CloverLeaf scaling – relative performance

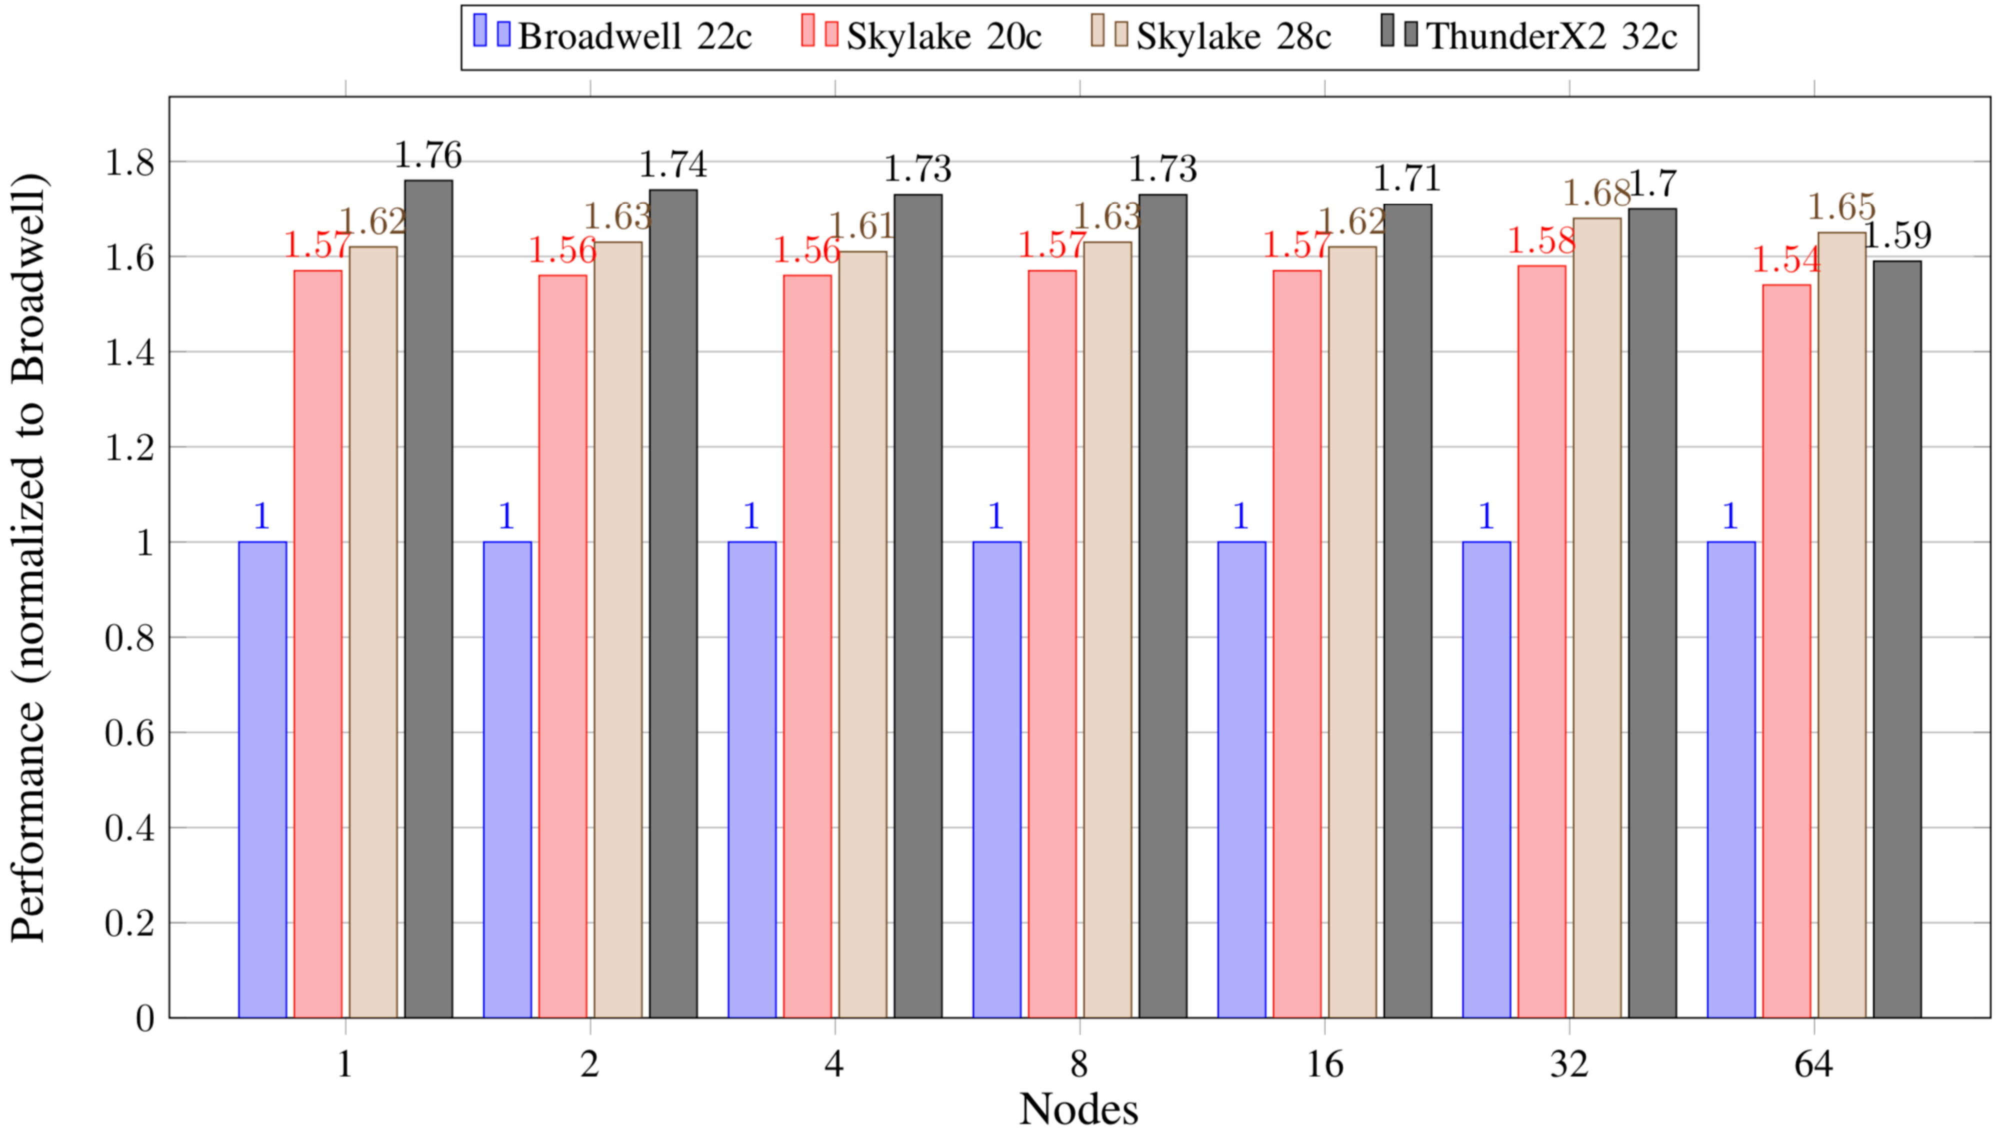

CloverLeaf scaling – parallel efficiency

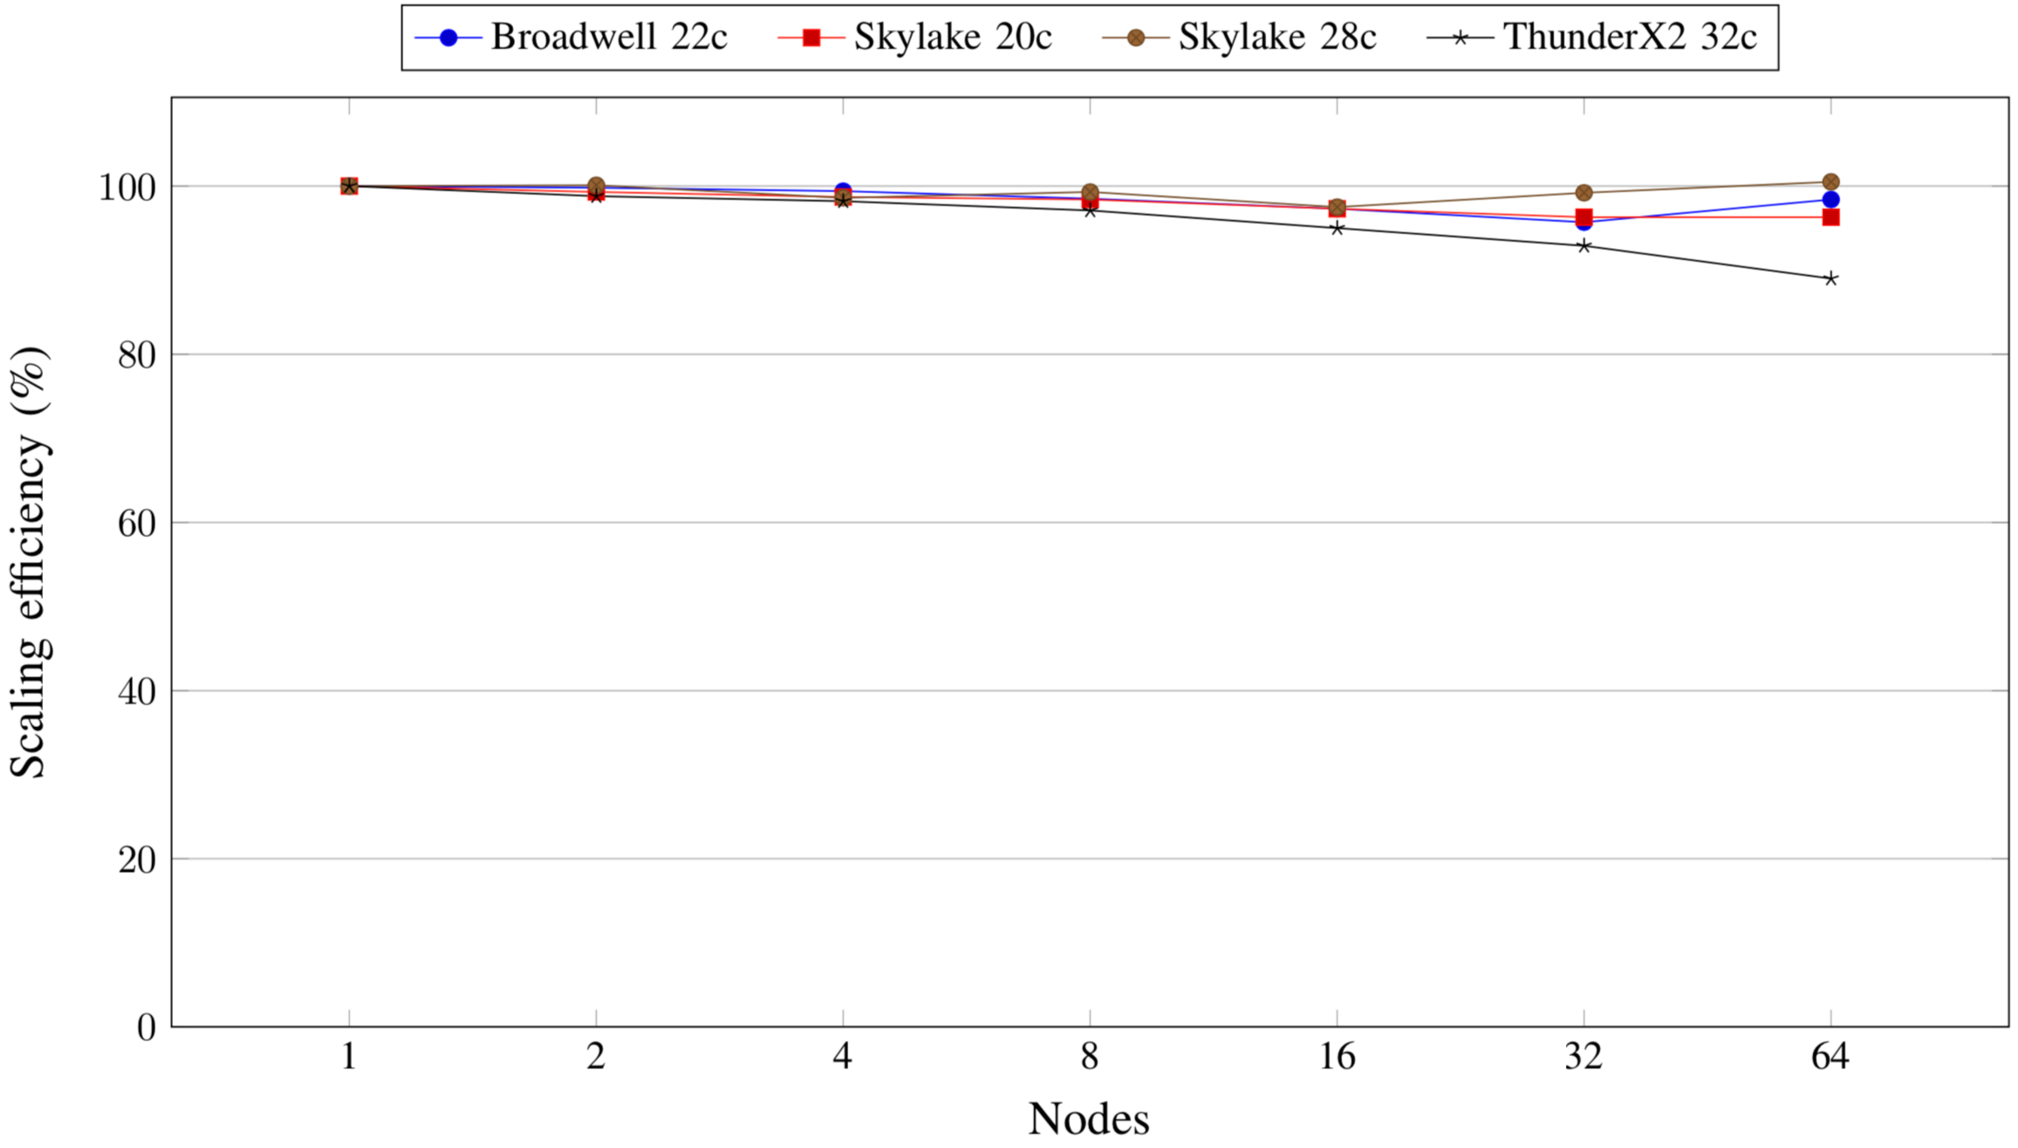

TeaLeaf scaling – relative performance

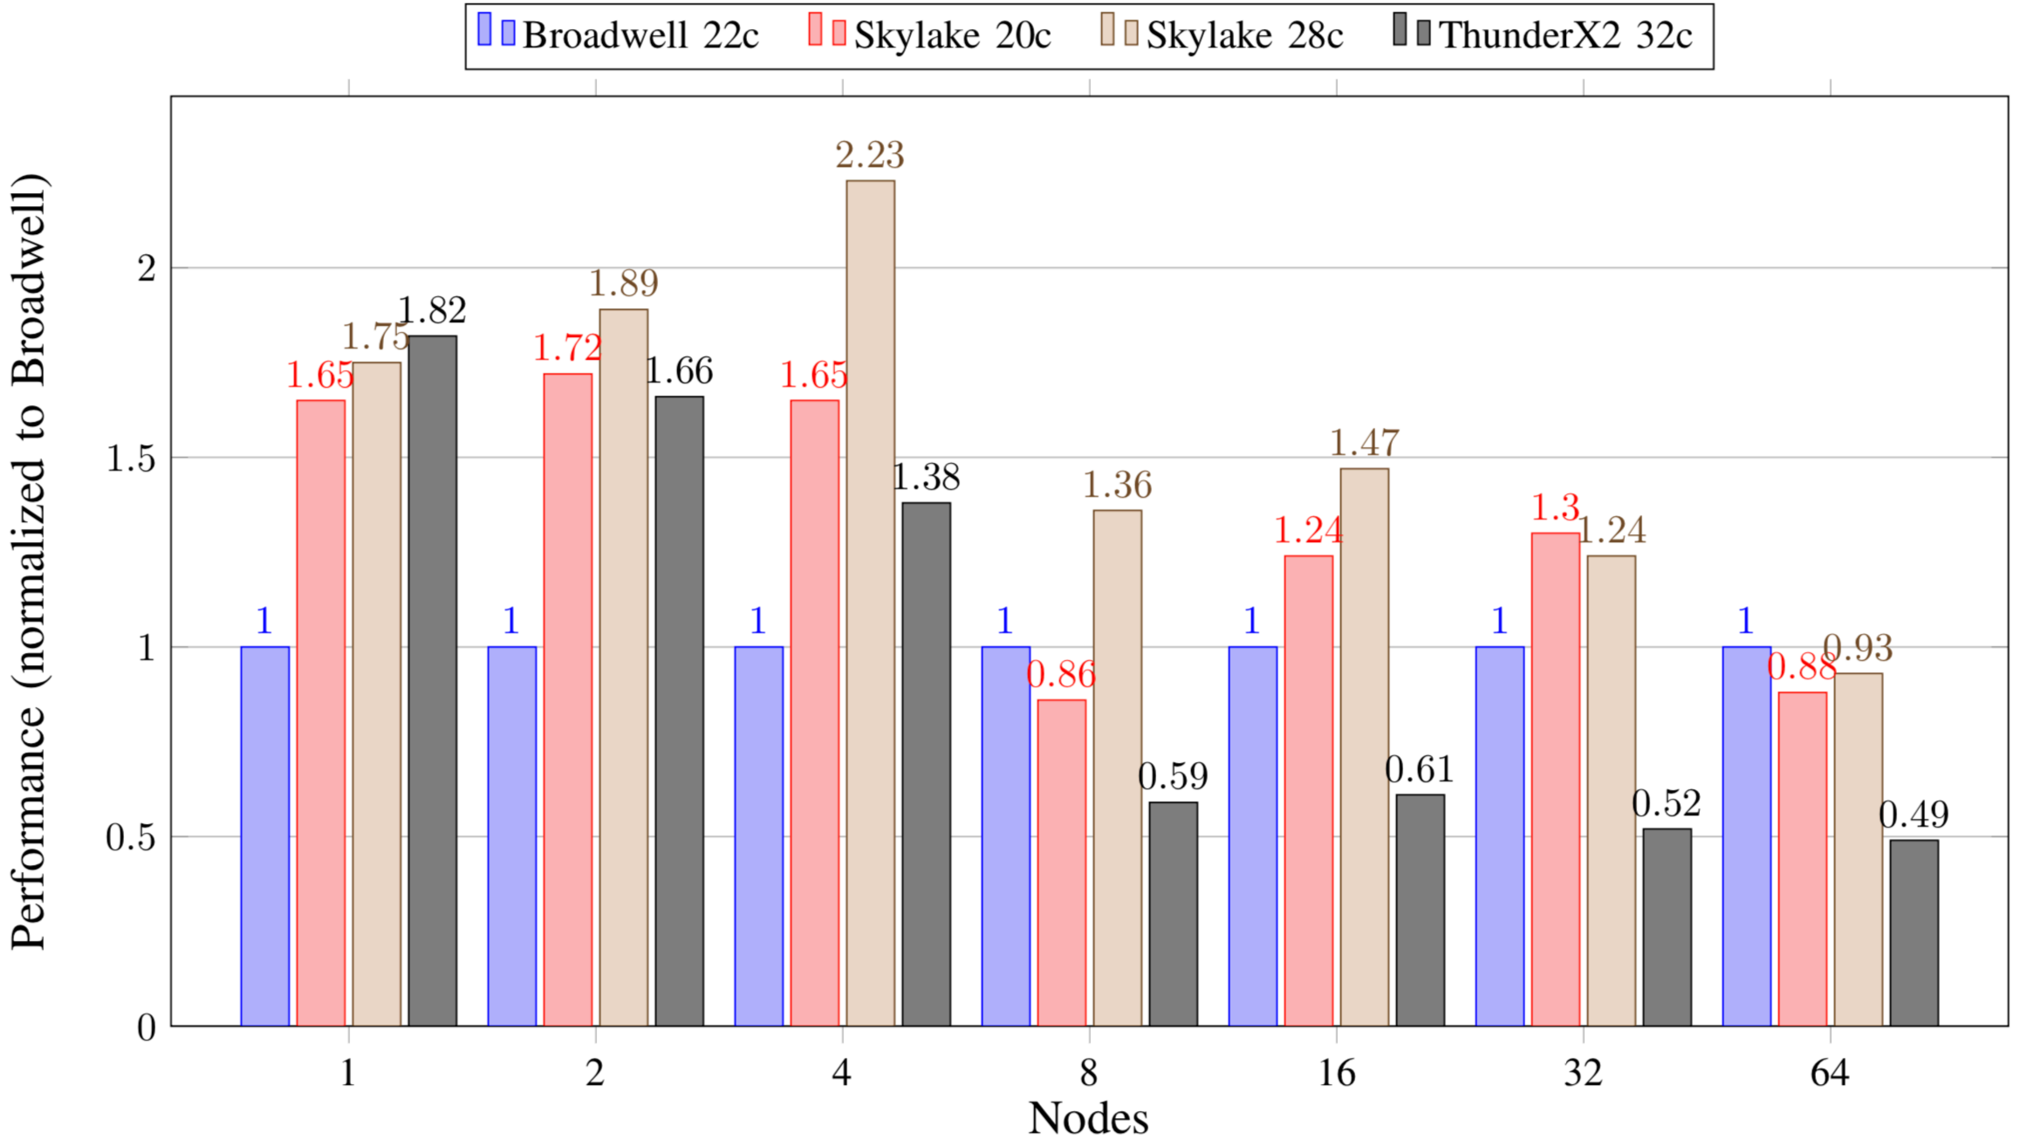

TeaLeaf scaling – parallel efficiency

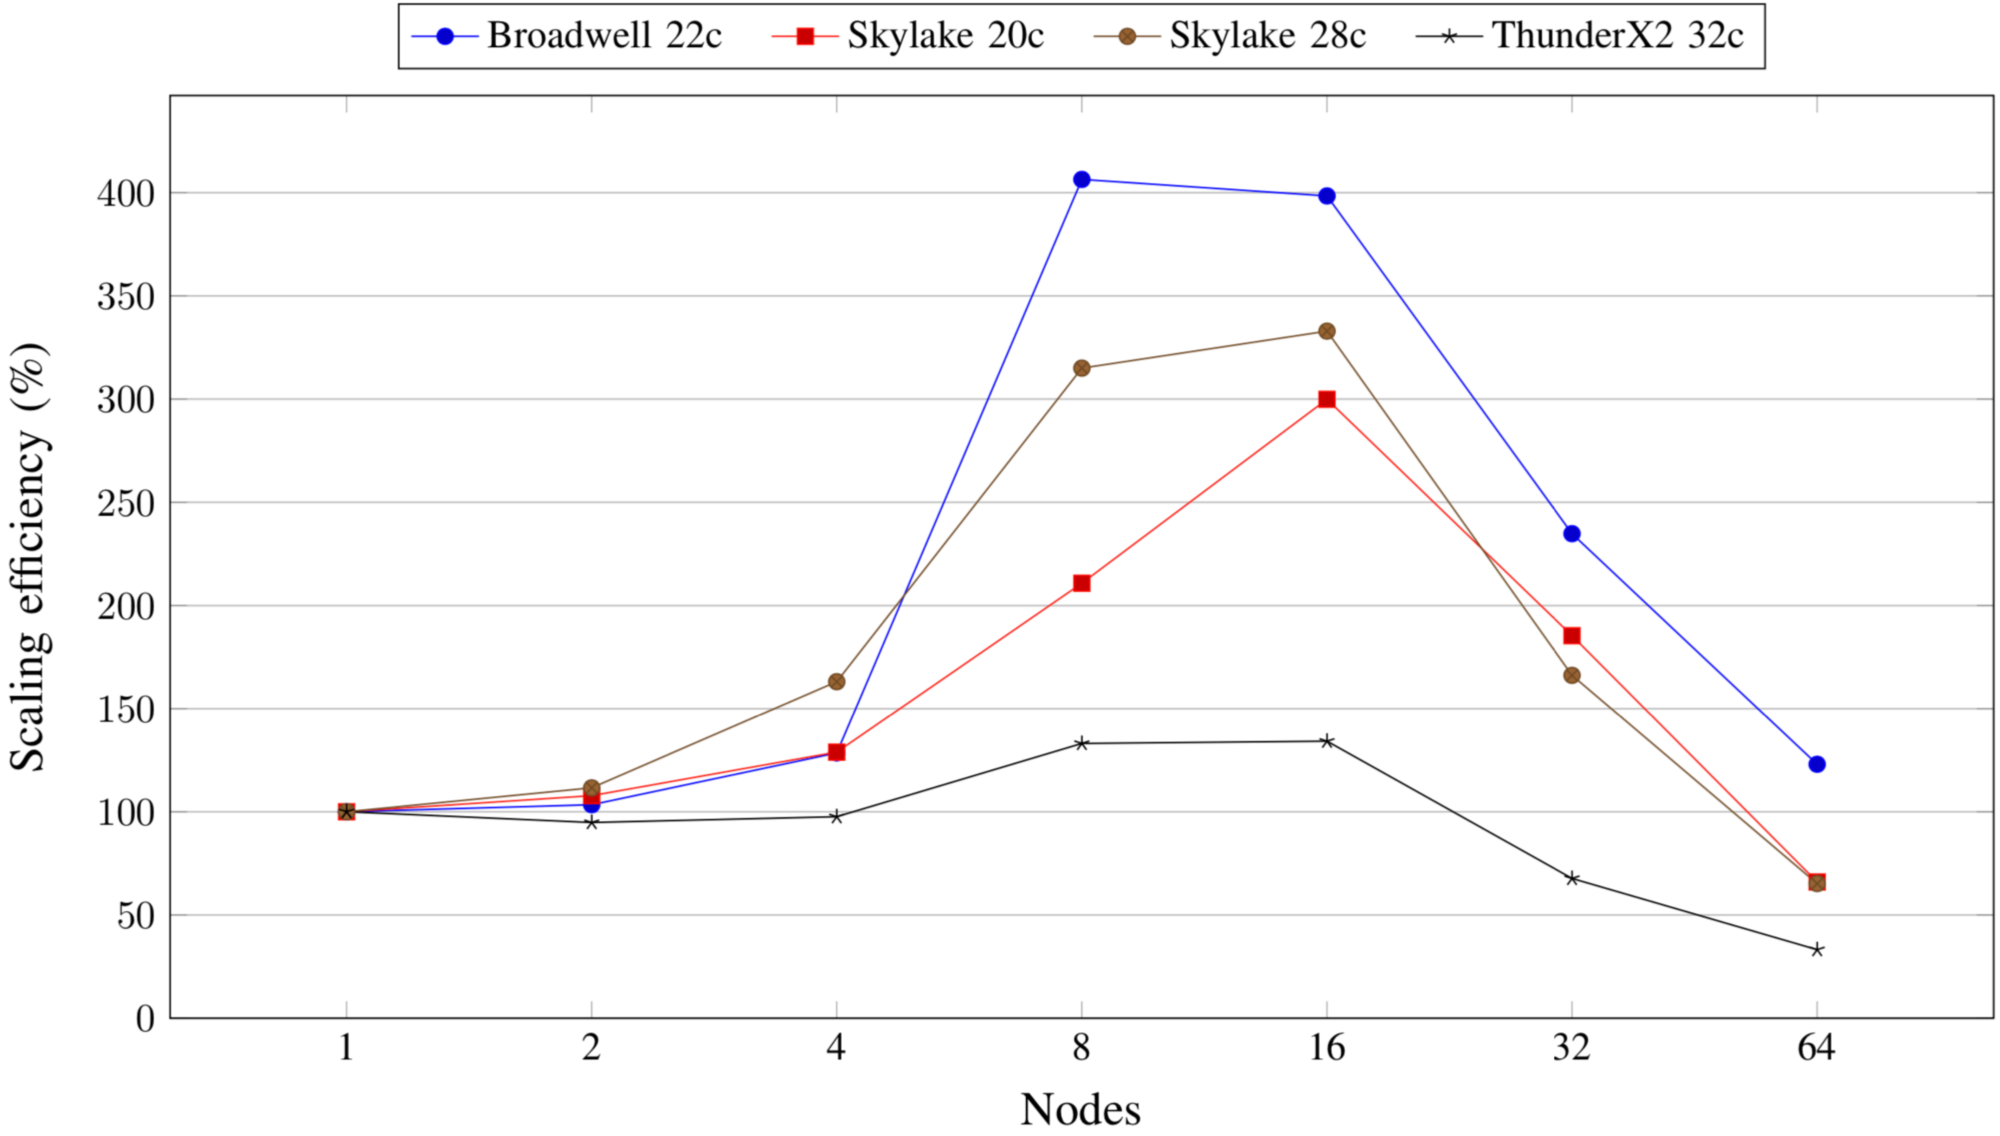

SNAP scaling – relative performance

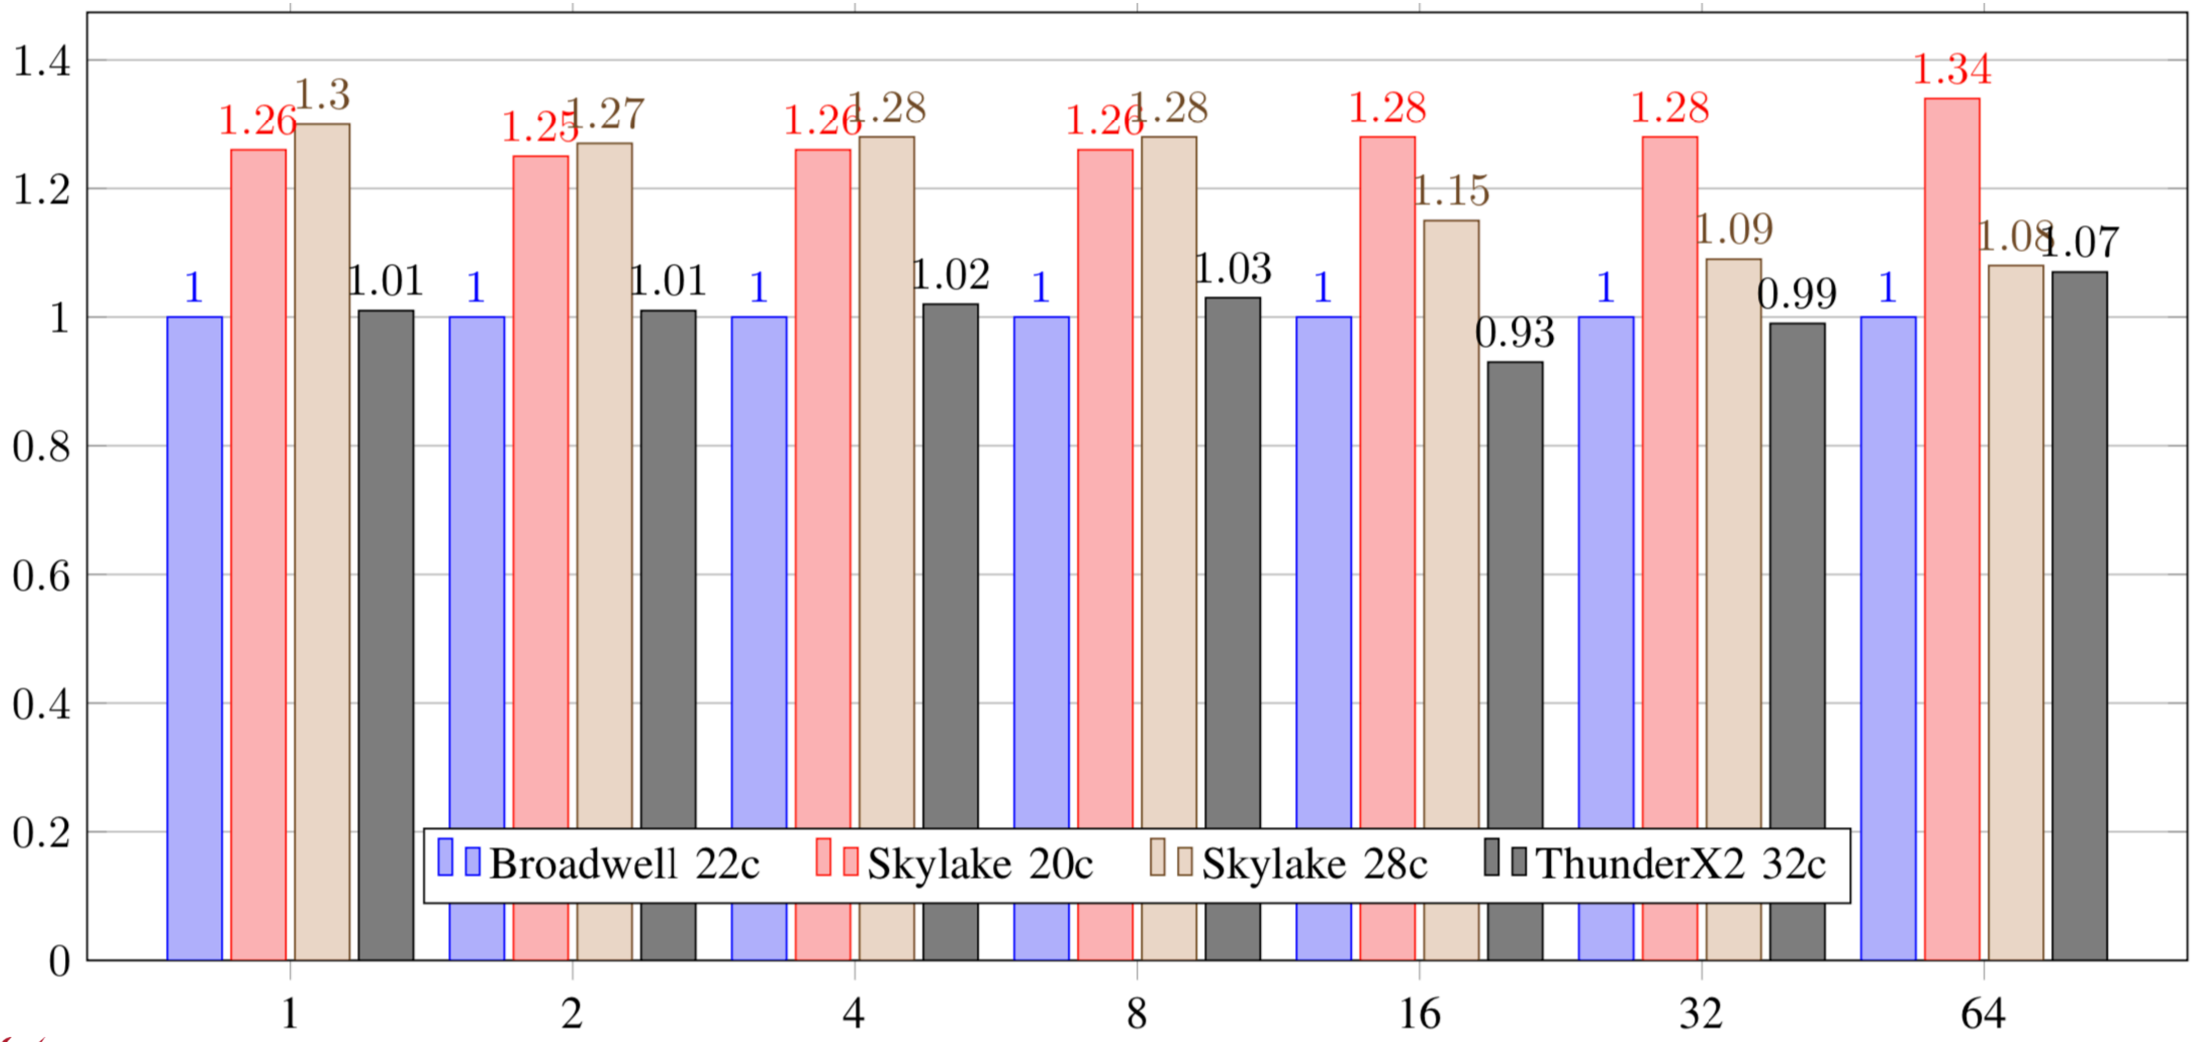

SNAP scaling – parallel efficiency

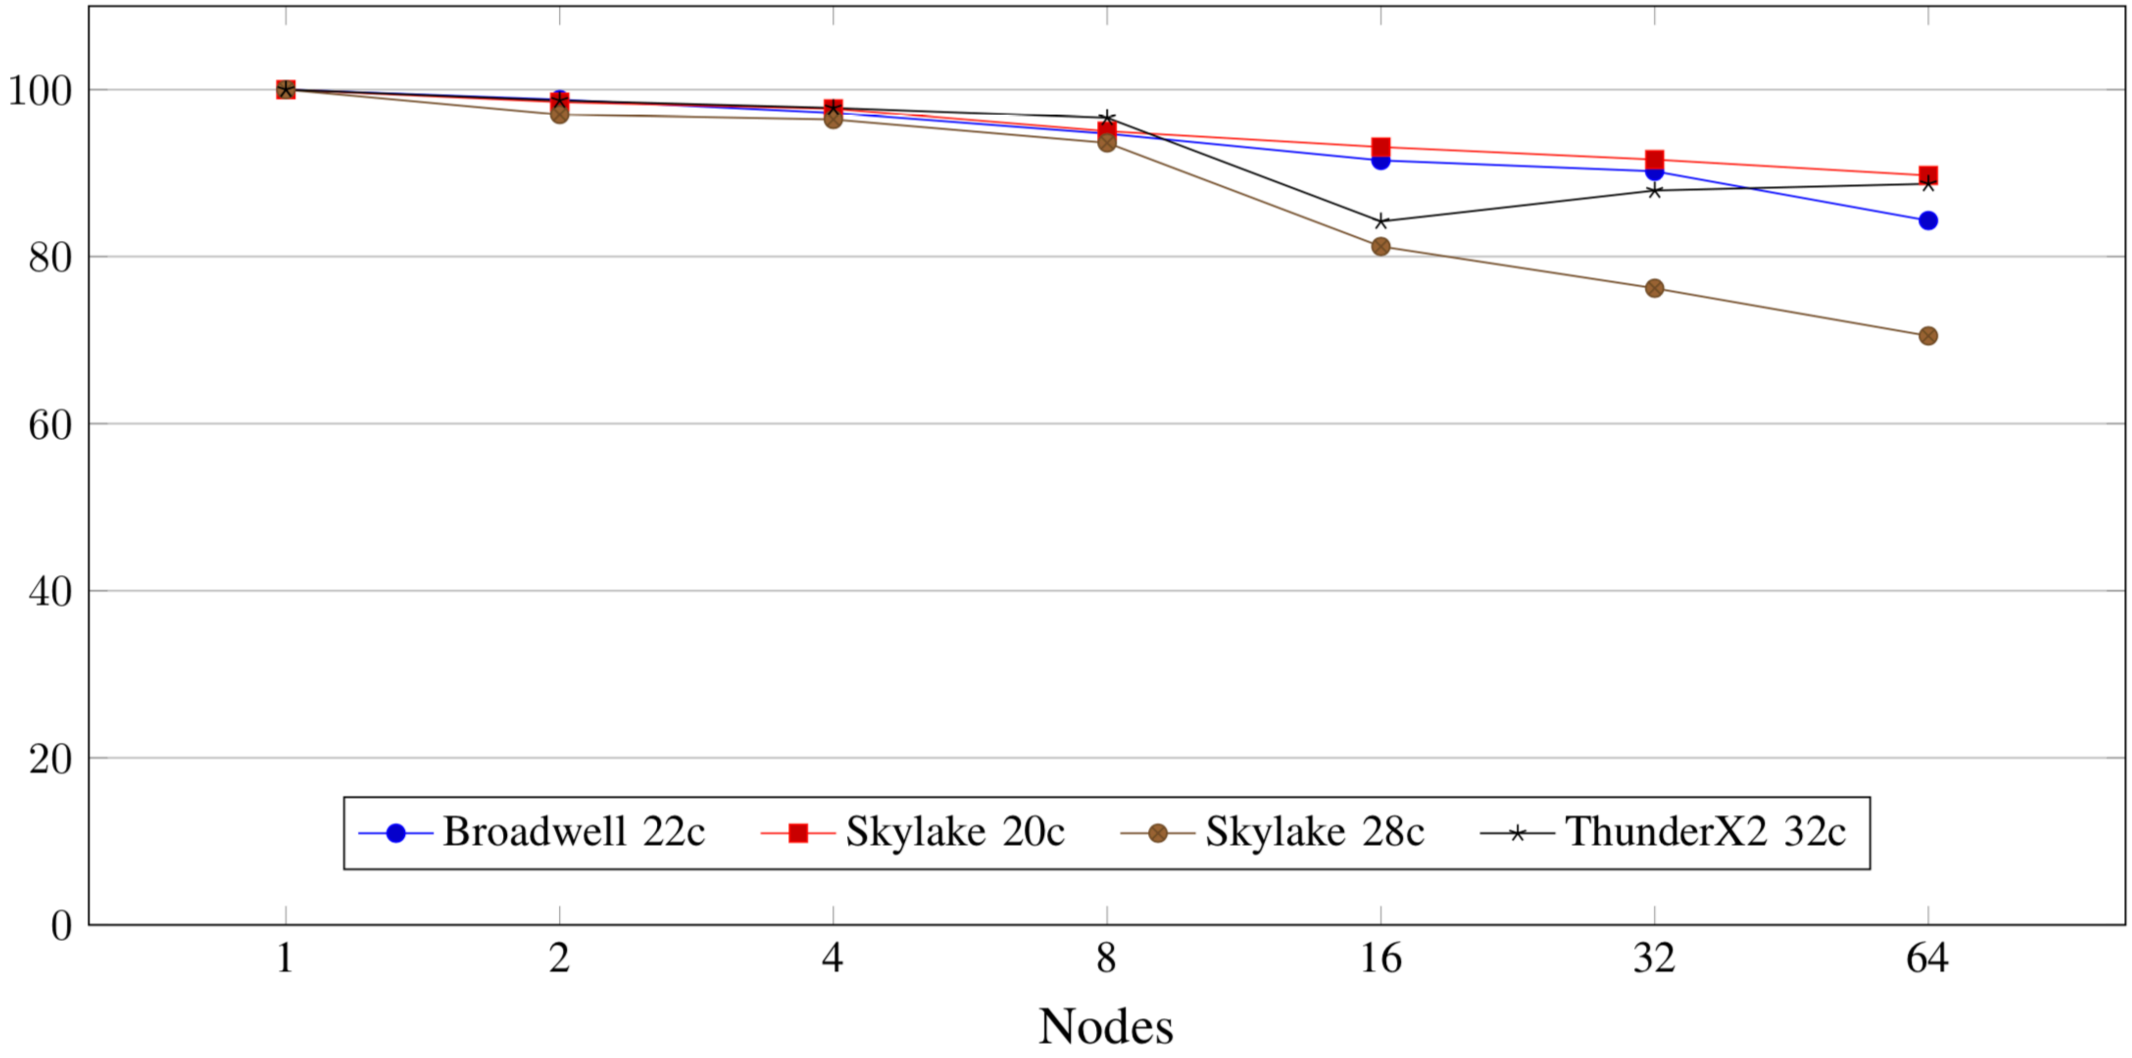

University of BRISTOL

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

GW4

GROMACS scaling – relative performance

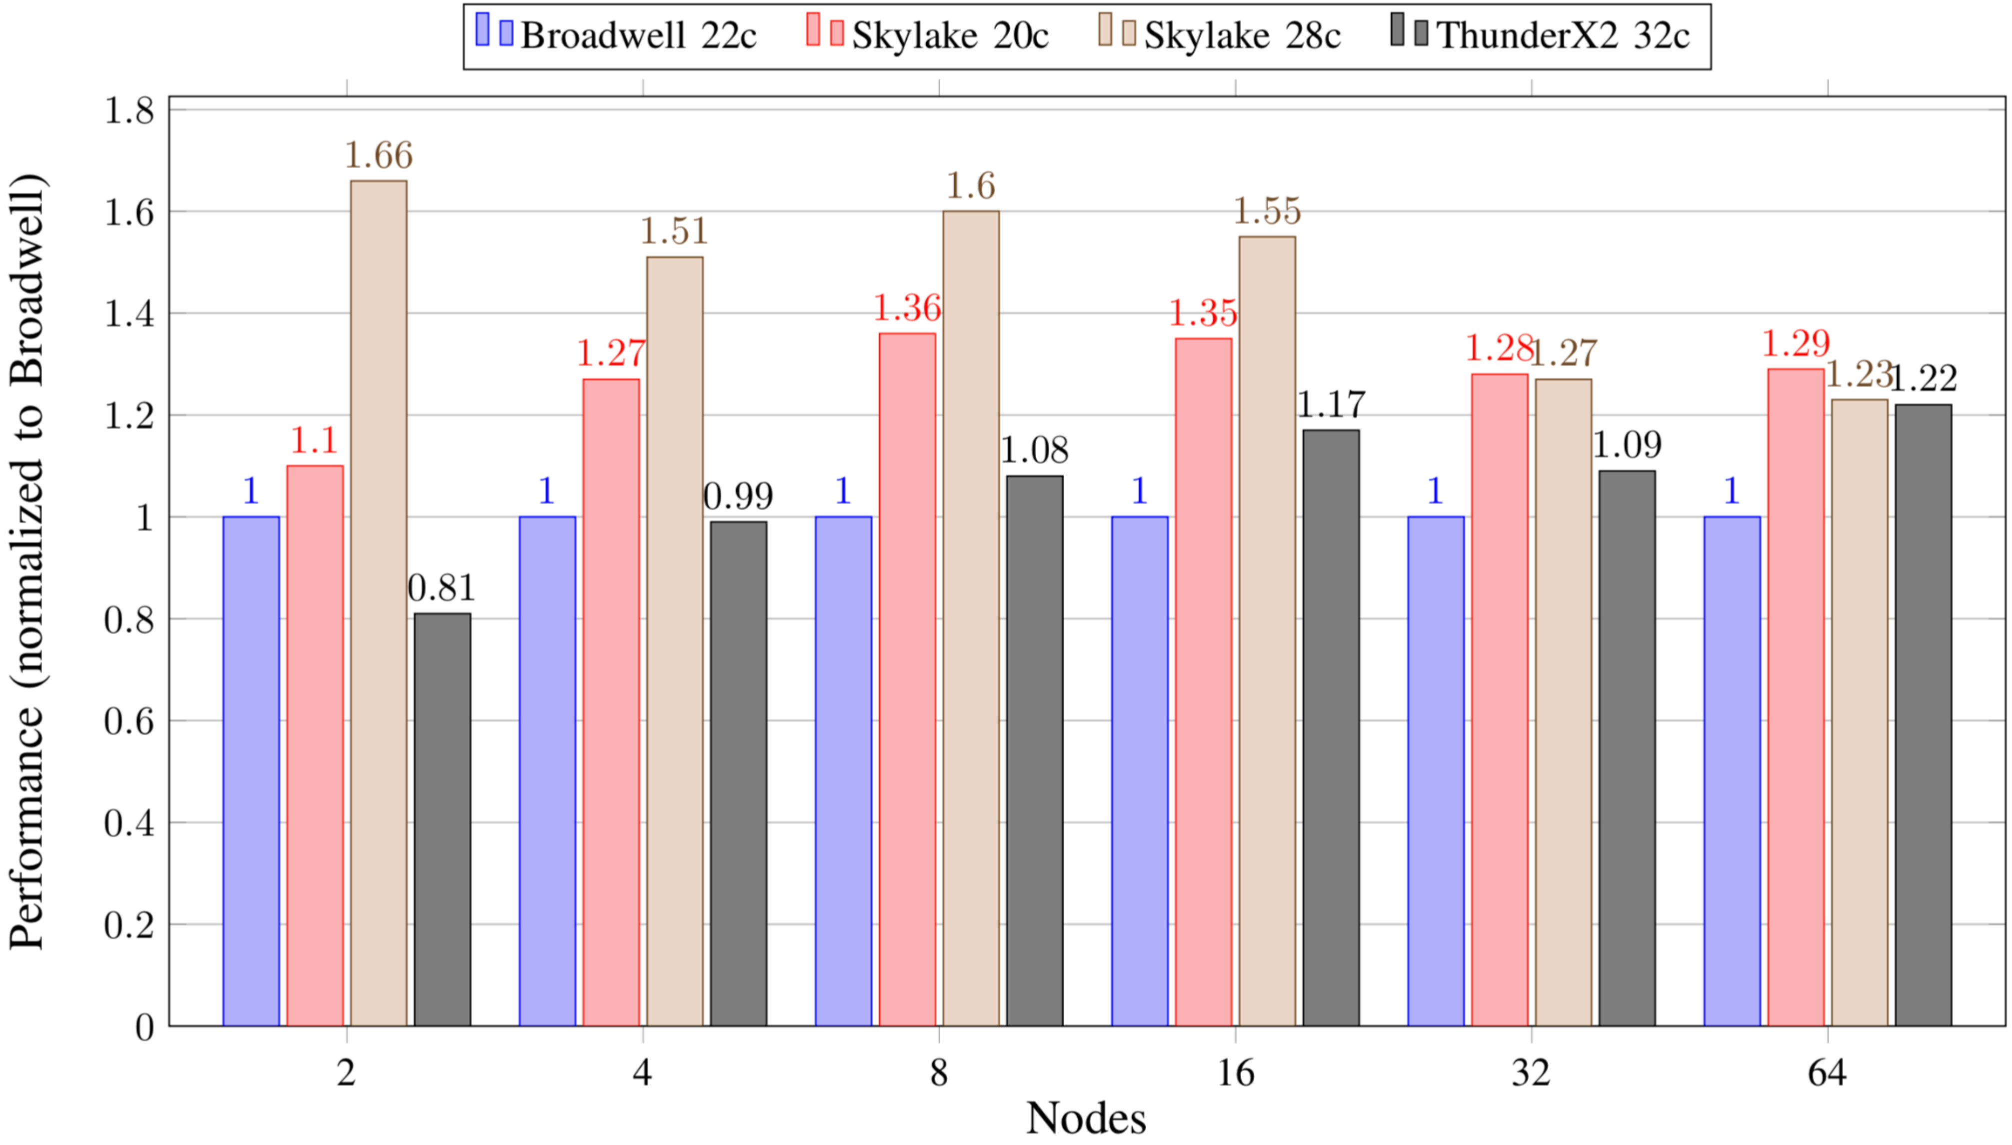

GROMACS scaling – parallel efficiency

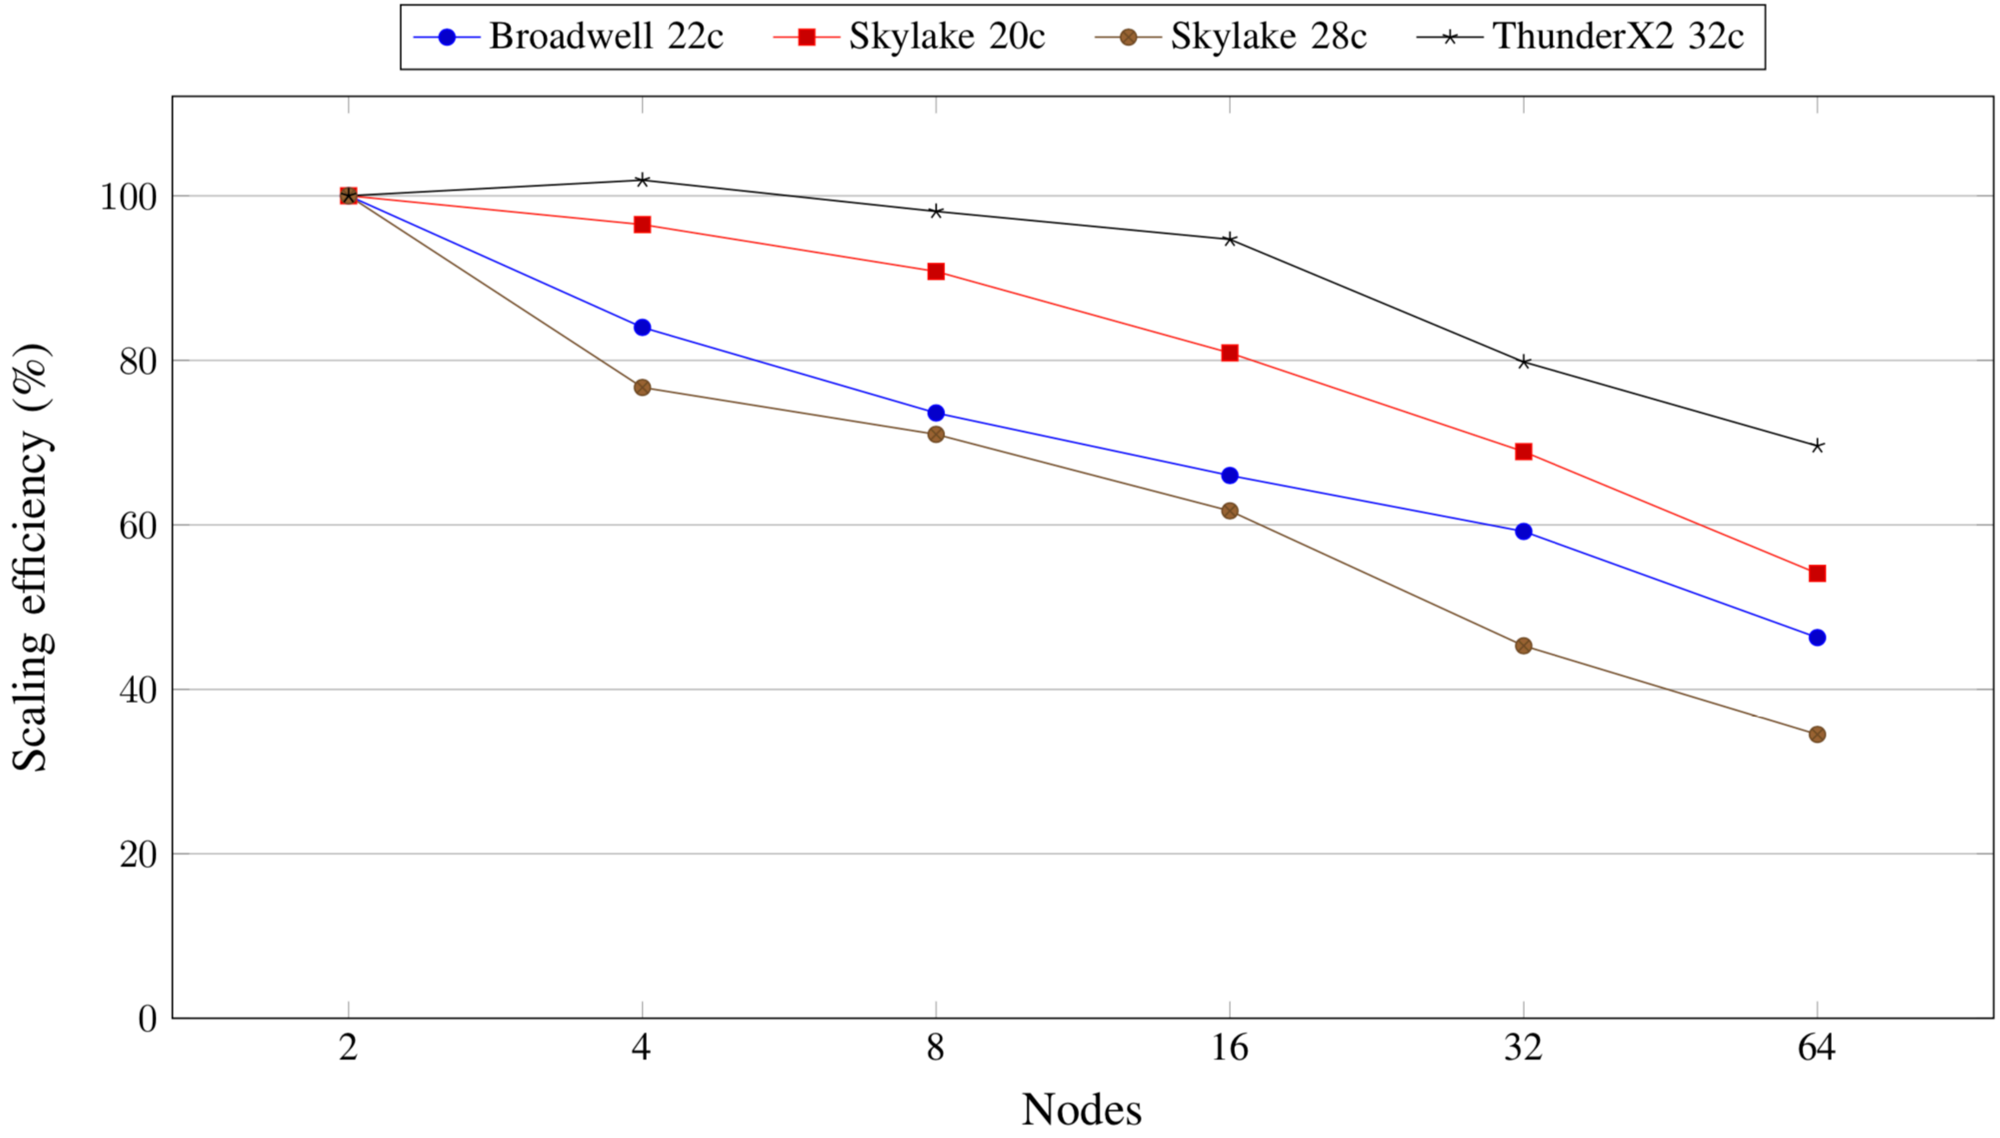

NEMO scaling

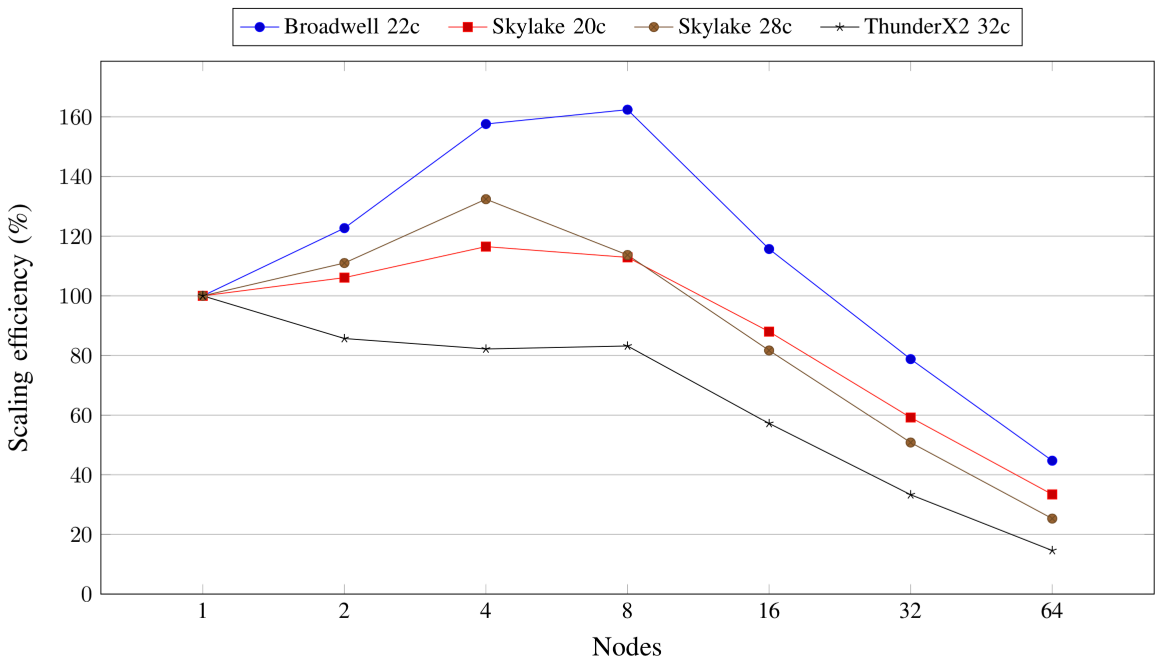

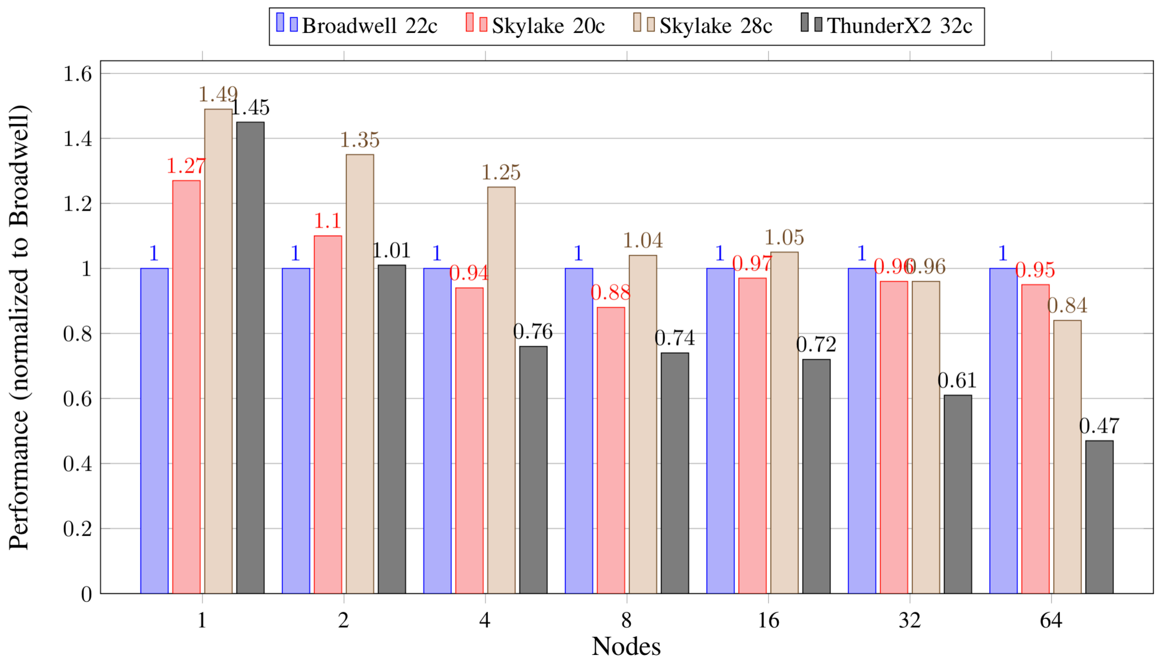

Parallel efficiency **Relative performance**

OpenFOAM scaling

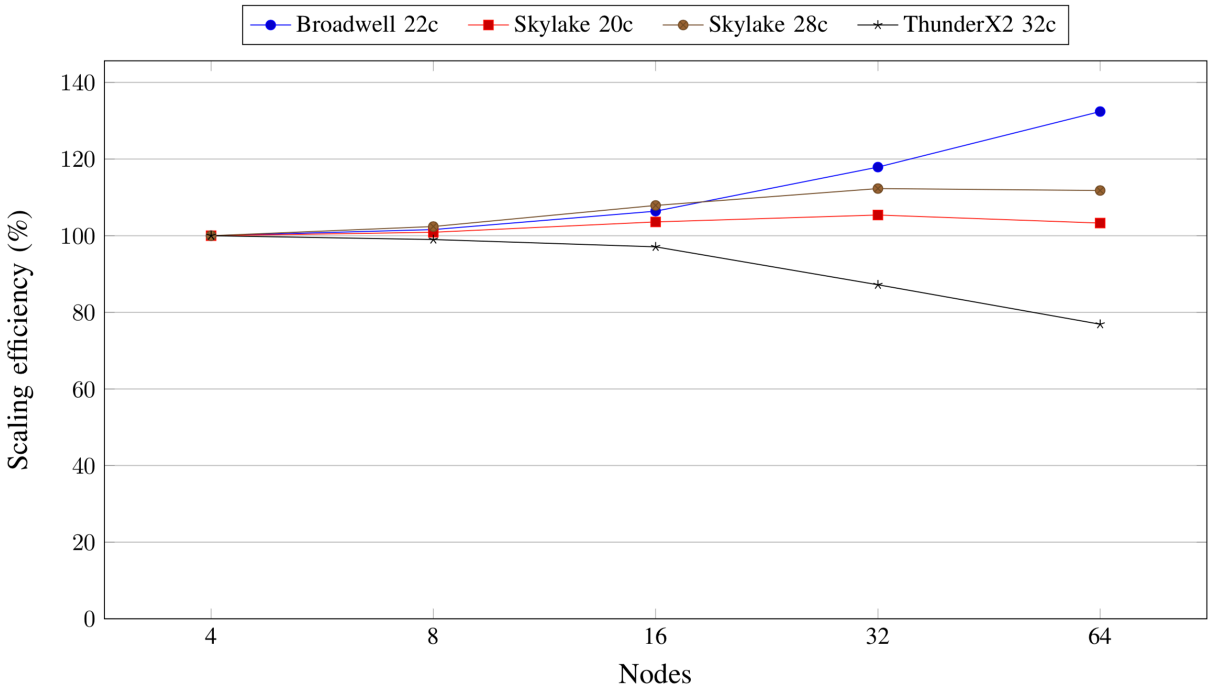

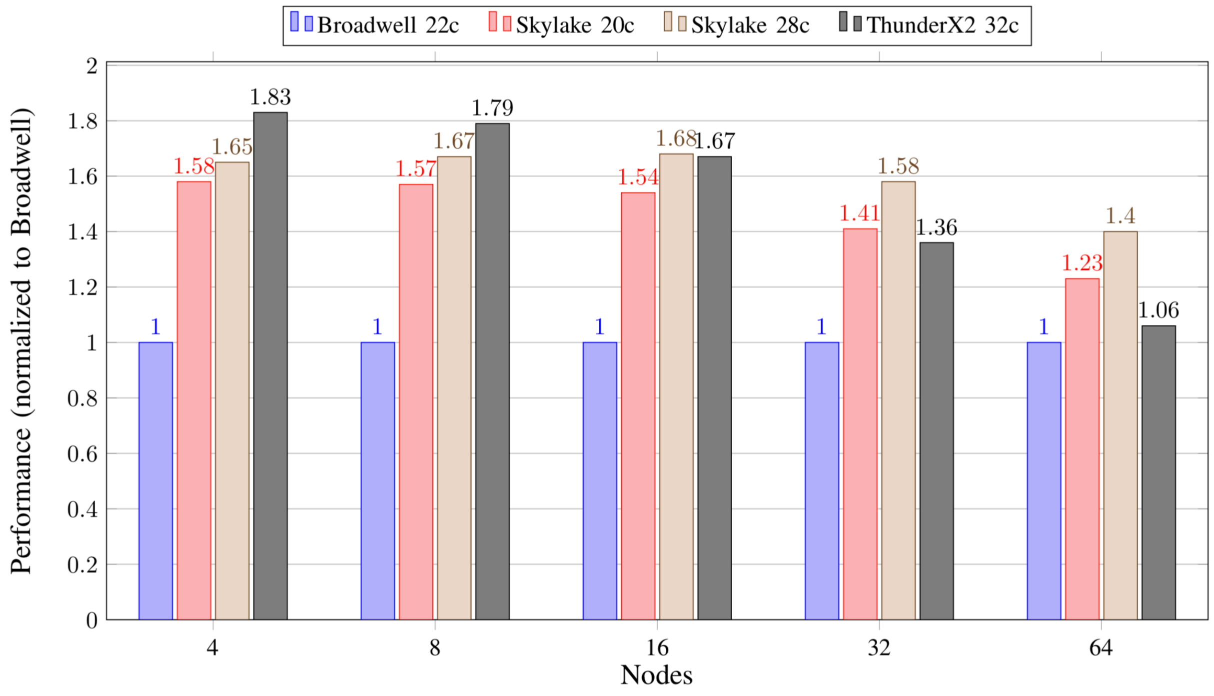

Parallel efficiency

Relative performance

University of BRISTOL

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

GW4

OpenSBLI scaling

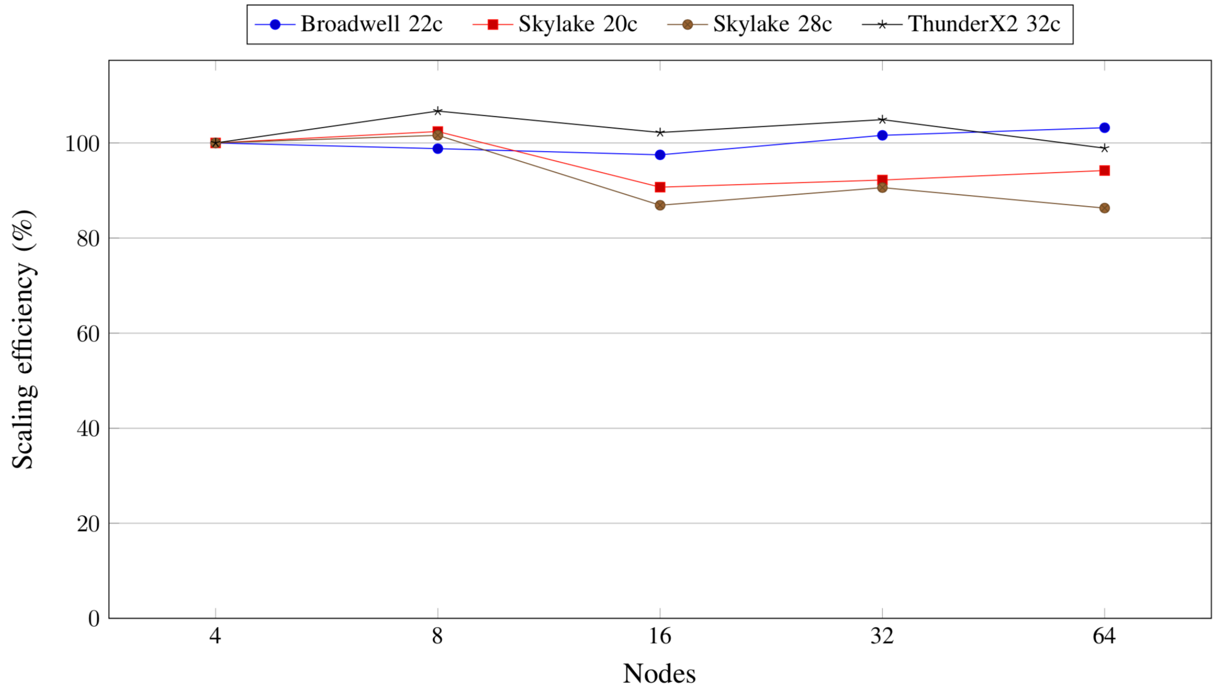

Parallel efficiency

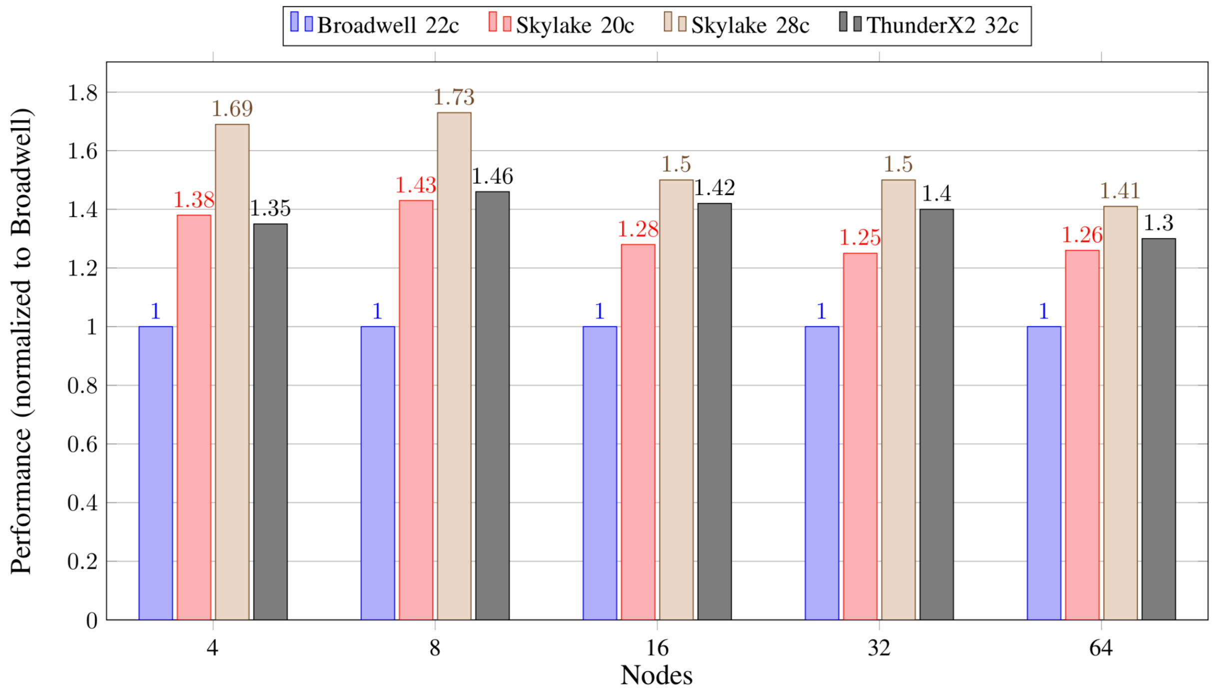

Relative performance

University of BRISTOL

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

GW4

VASP scaling

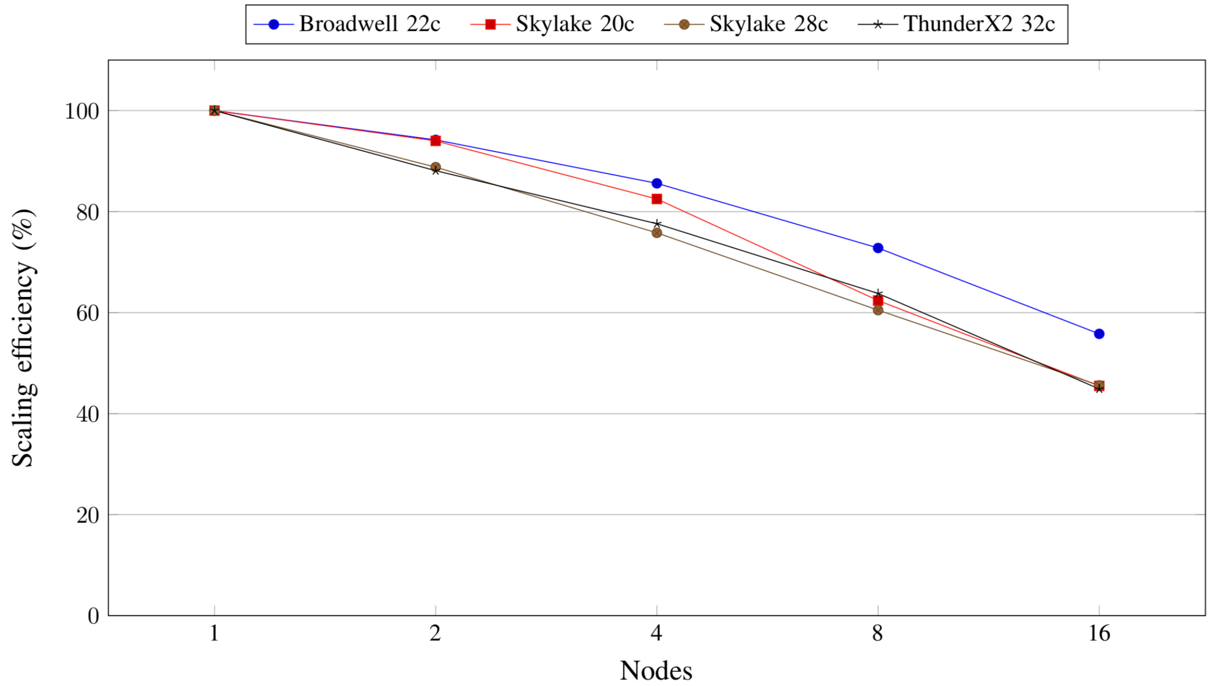

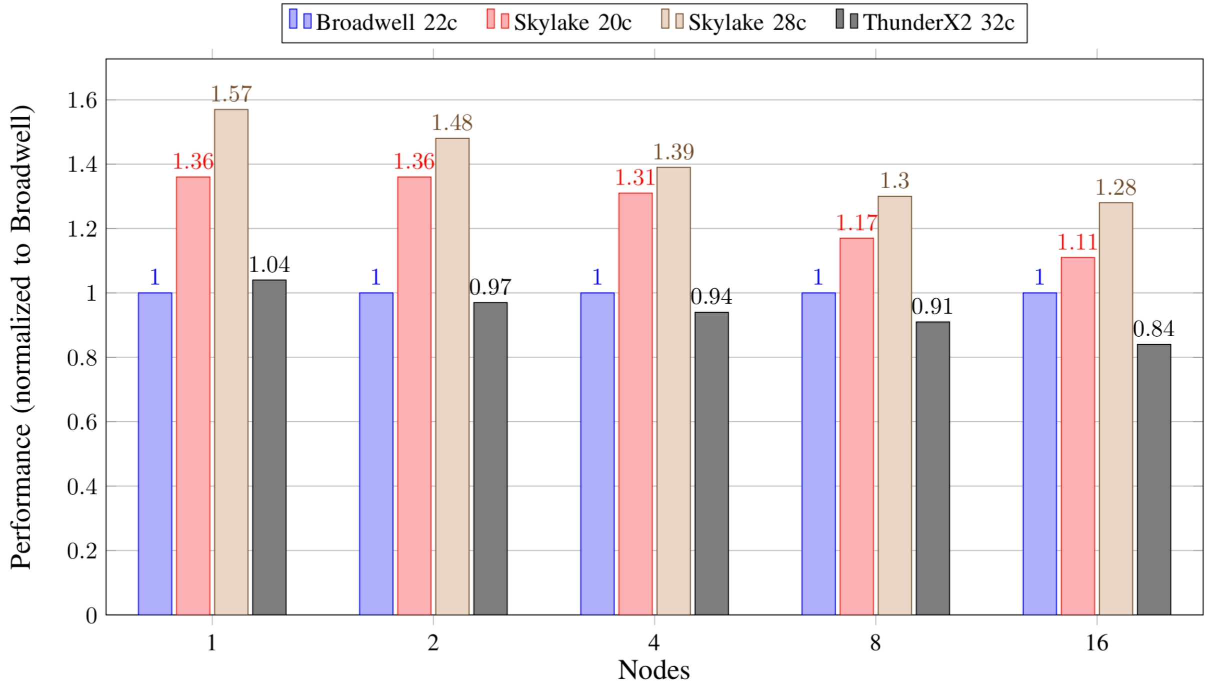

Parallel efficiency

Relative performance

University of BRISTOL

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

GW4

Which compilers were best in each case?

Benchmark	ThunderX2	Broadwell	Skylake
STREAM	GCC 8.2	Intel 2019	CCE 8.7
CloverLeaf	CCE 8.7	Intel 2019	Intel 2019
TeaLeaf	CCE 8.7	Intel 2019	Intel 2019
SNAP	CCE 8.7	Intel 2019	Intel 2019
GROMACS	GCC 8.2	GCC 8.2	GCC 8.2
NEMO	CCE 8.7	CCE 8.7	CCE 8.7
OpenFOAM	GCC 7.3	GCC 7.3	GCC 7.3
OpenSBLI	CCE 8.7	Intel 2019	CCE 8.7
VASP	GCC 7.3	Intel 2019	Intel 2019

Isambard scaling summary

- **Arm-based systems appear to scale just as well as x86 ones**
- **For certain codes that were compute-bound at low scale, these became network bound at 'real' scale, levelling the playing field**
- **We're seeing a minor issue with scaling in two cases, appears to be related to MPI collectives** – investigations are underway
- **The software stack has been robust, reliable and high-quality** (both the commercial and open source parts)
- Now have evidence that **Arm-based systems are real alternatives for HPC,** reintroducing much needed ***competition*** to the market

The Bristol HPC team doing this work

James Price

Andrei Poenaru

Tom Deakin

Also thanks go to:

- The Isambard project members: the GW4 Alliance, the Met Office, Arm, Marvell and Cray
- Cray for access to the Swan and Horizon x86 systems
- EPSRC for funding the project

University of BRISTOL

EPSRC

GW4

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

Backup

Comparison of compilers on Arm

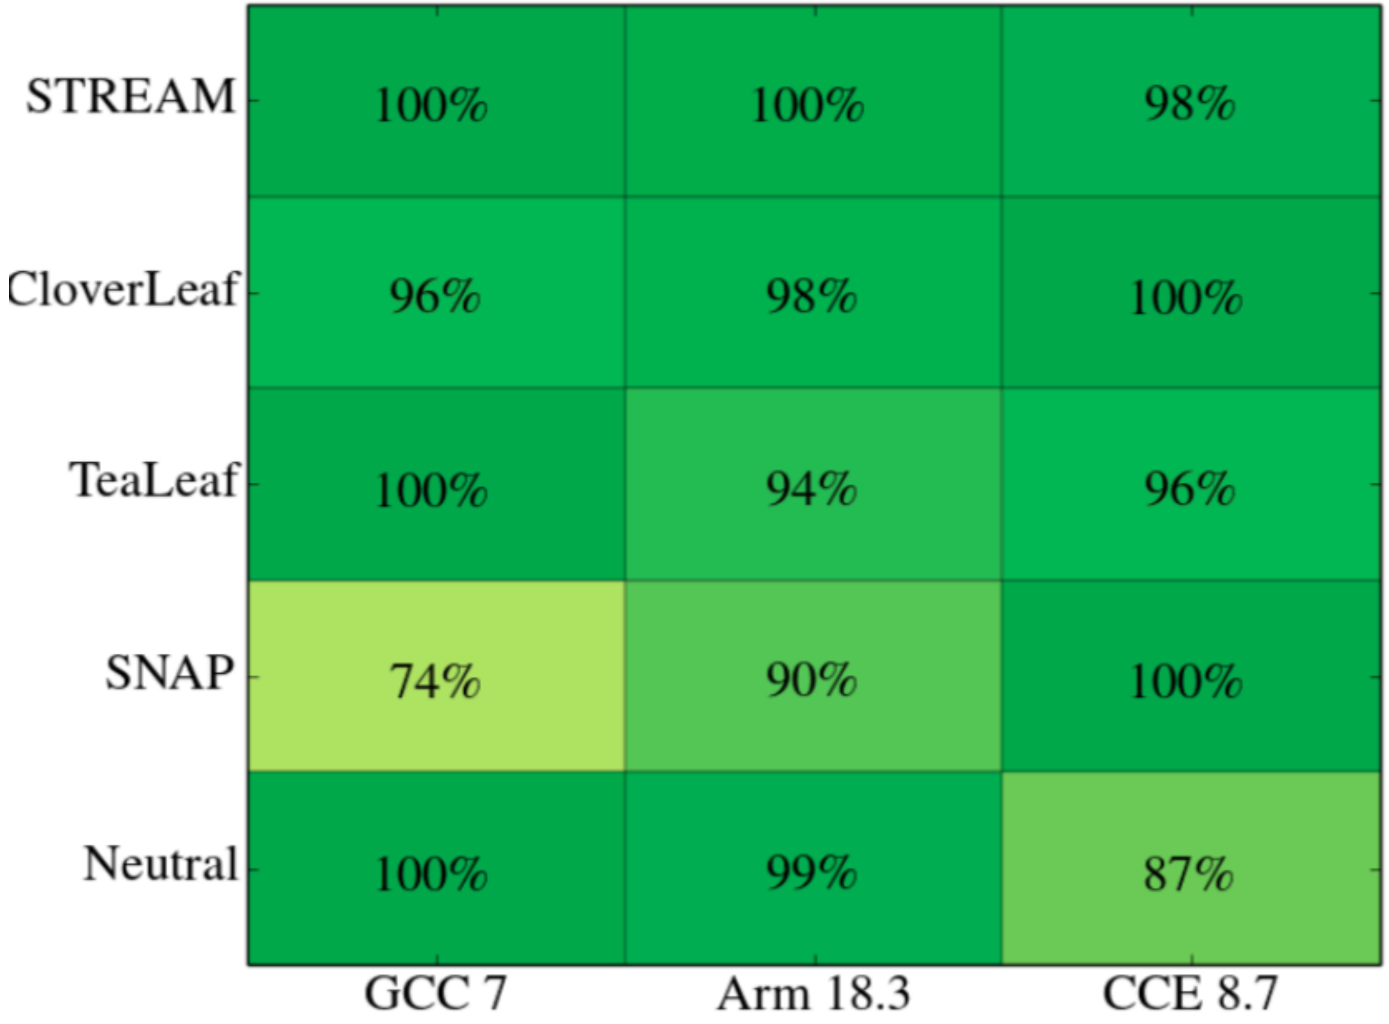

University of BRISTOL

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

GW4