

Prof Simon McIntosh-Smith

HPC research group

University of Bristol

@simonmcs

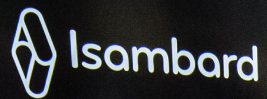


Status and performance of two Arm supercomputers: the Cray Isambard Tier-2 service and HPE Catalyst

Isambard system specification

- 10,752 Armv8 cores (168n x 2s x 32c)
 - Marvell ThunderX2 32 core 2.1→2.5GHz
 - 256 GB RAM per node, bandwidth >240 GB/s
- 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):**
 - Installed in November 2018, accepted in 1 week!
 - Upgraded silicon (to B2), firmware and stack Mar19
 - Now have 185 registered users, 63 are external
 - **PRODUCTION SERVICE opened to all users May 2019**
 - First Arm-based production service in the world!





EPSRC

Pioneering research
and skills

GW4



UNIVERSITY OF
BATH



University of
BRISTOL

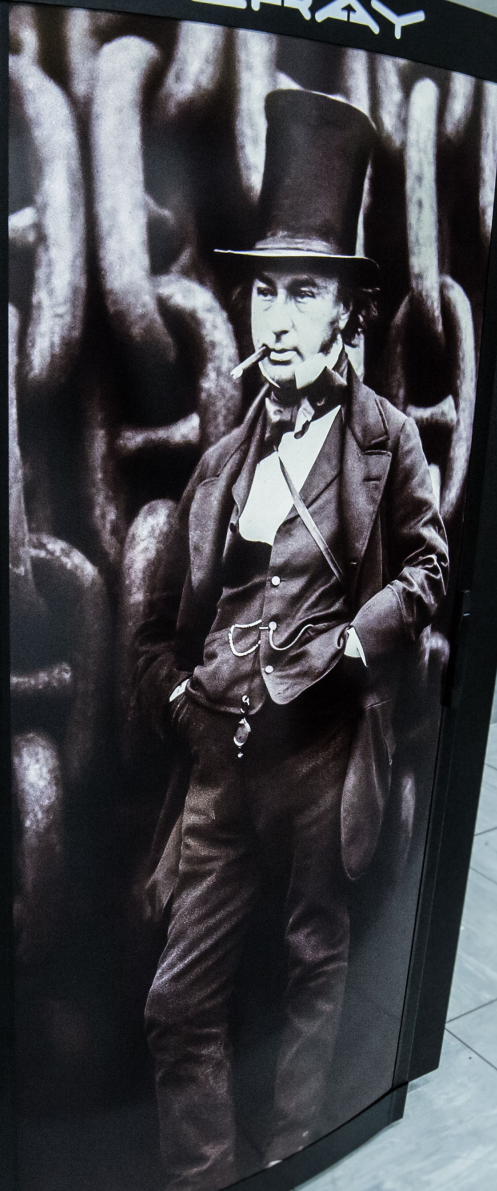
CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD

UNIVERSITY OF
EXETER



arm

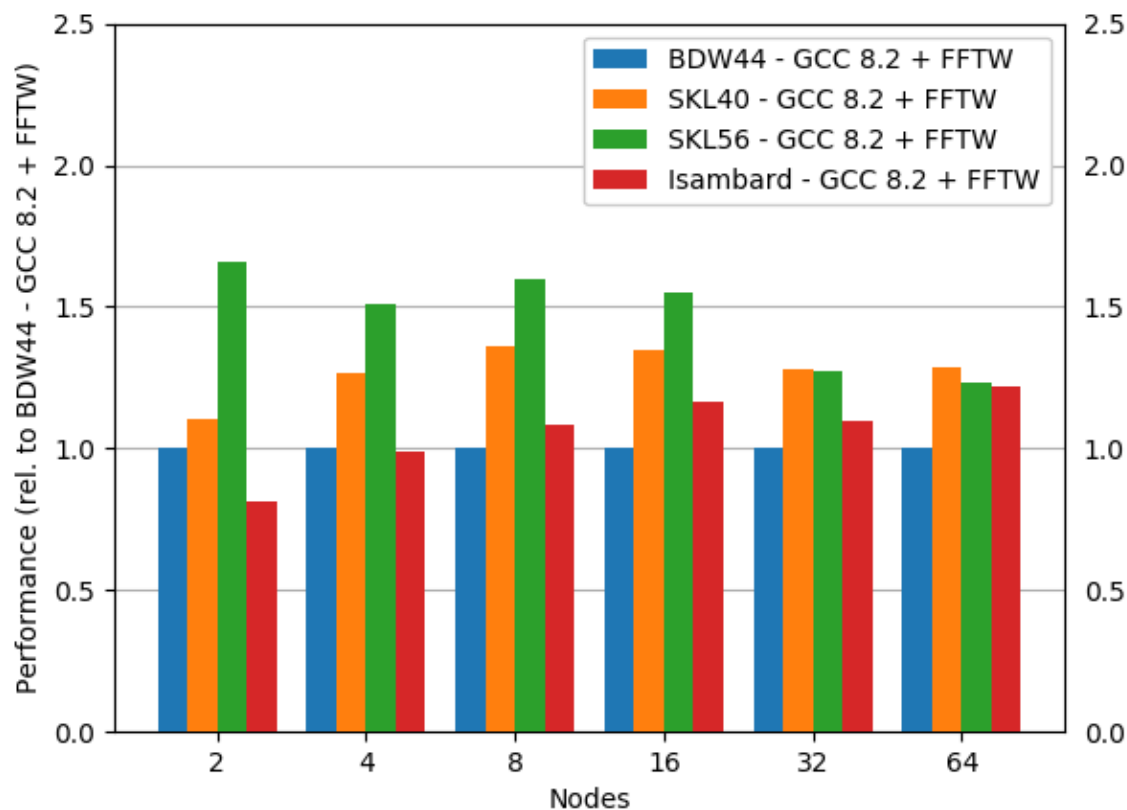
CRAY



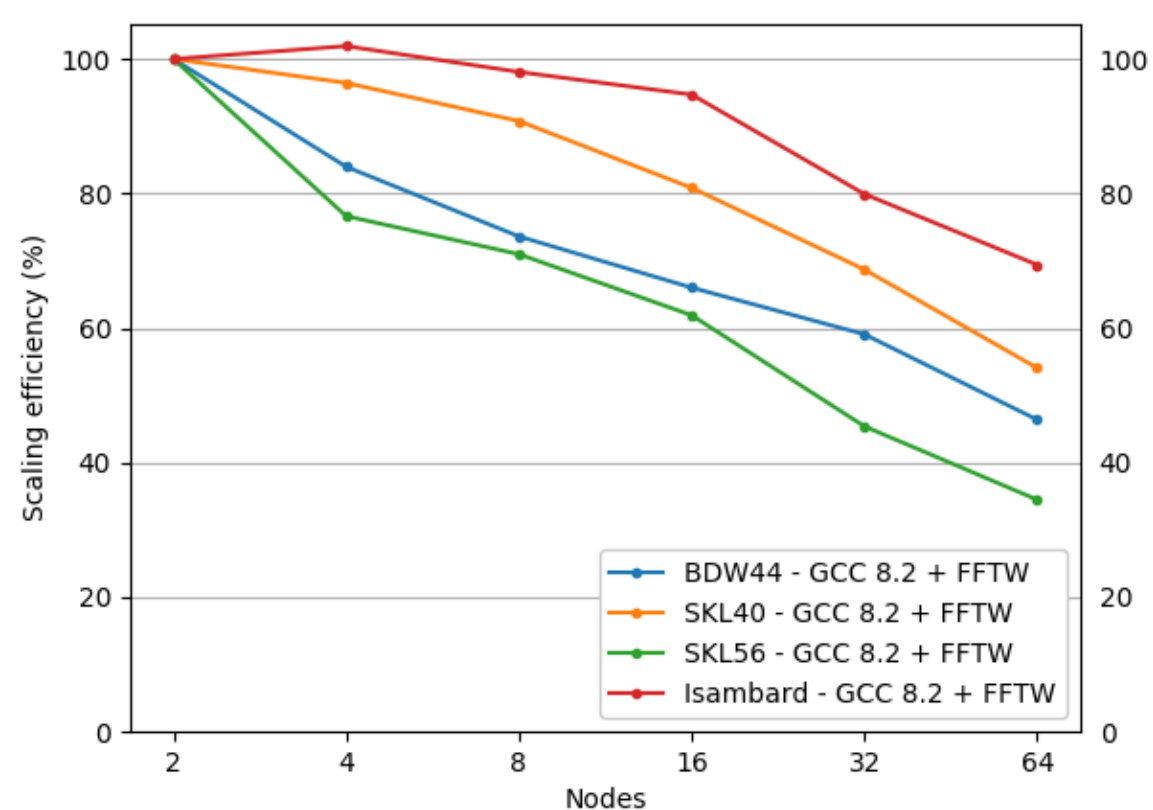
Scalability comparisons between x86 and Arm

- We explore relative performance and scaling efficiency
- We've compared against three x86-based XC50 systems:
 - Swan using Intel Broadwell E5-2699 22-core CPUs at 2.2GHz
 - Horizon using Intel Skylake Gold 6148 20-core CPUs at 2.4GHz
 - Swan using Intel Skylake Platinum 8176 28-core CPUs at 2.1GHz
- All the results are for **strong scaling**, up to 64 nodes
- All of these systems use the same interconnect (Cray's 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

GROMACS (42 million atoms, ARCHER benchmark)

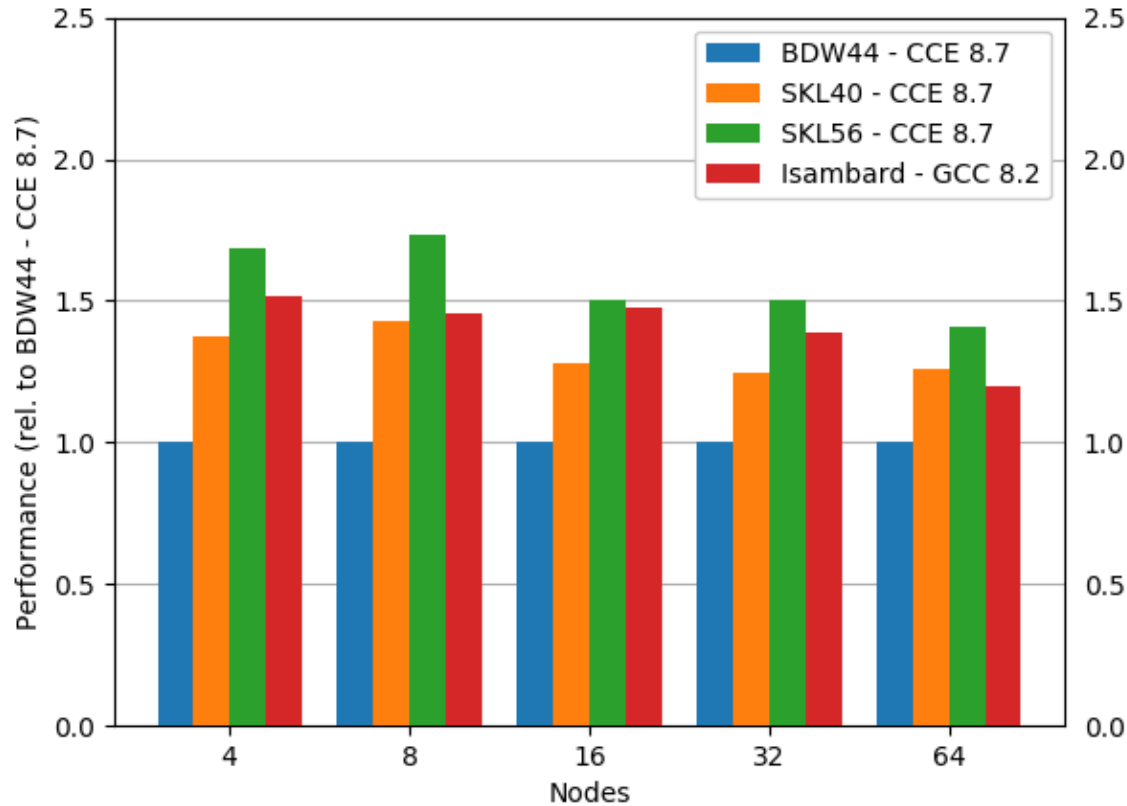


Relative performance

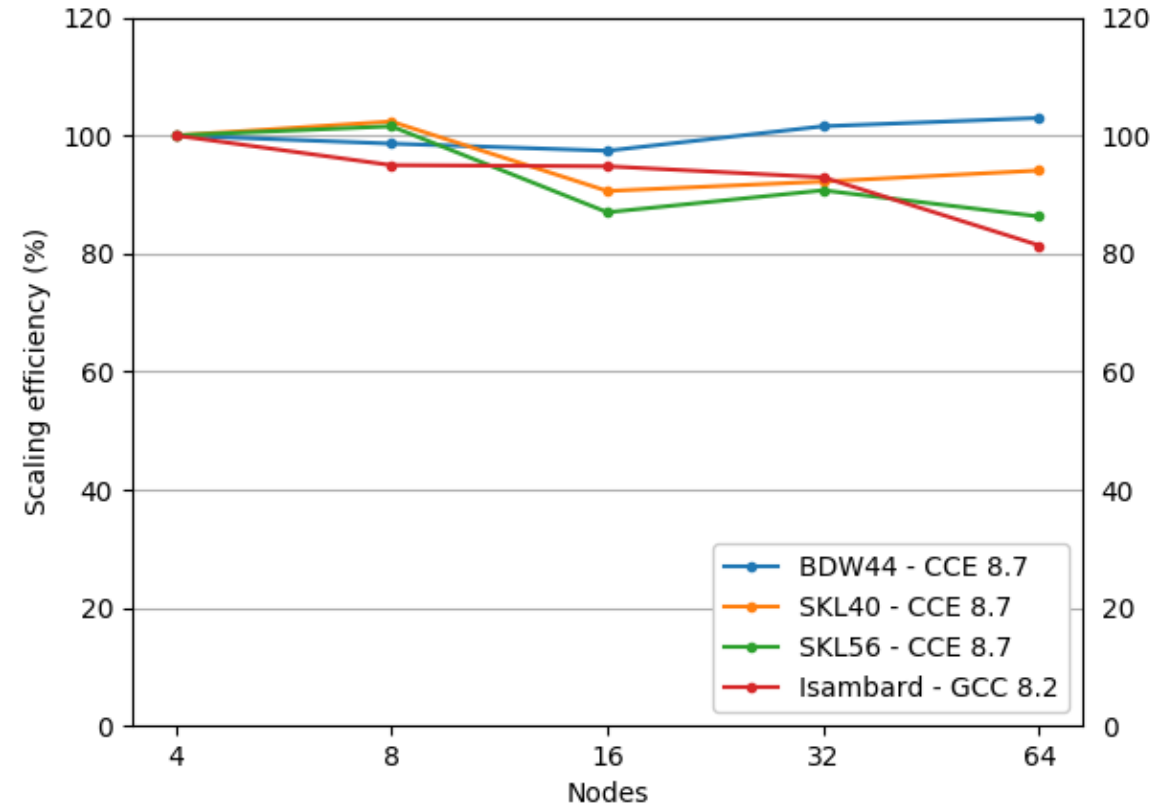


Parallel efficiency

OpenSBLI (1024^3, ARCHER benchmark)

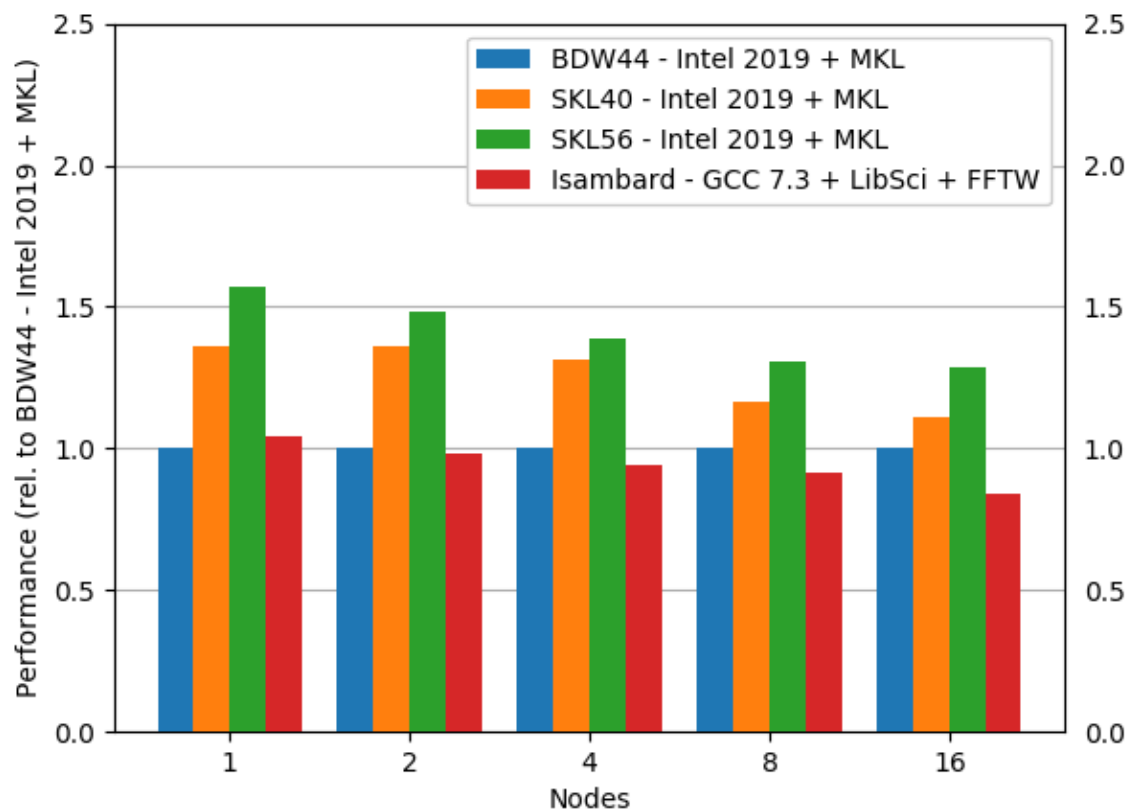


Relative performance

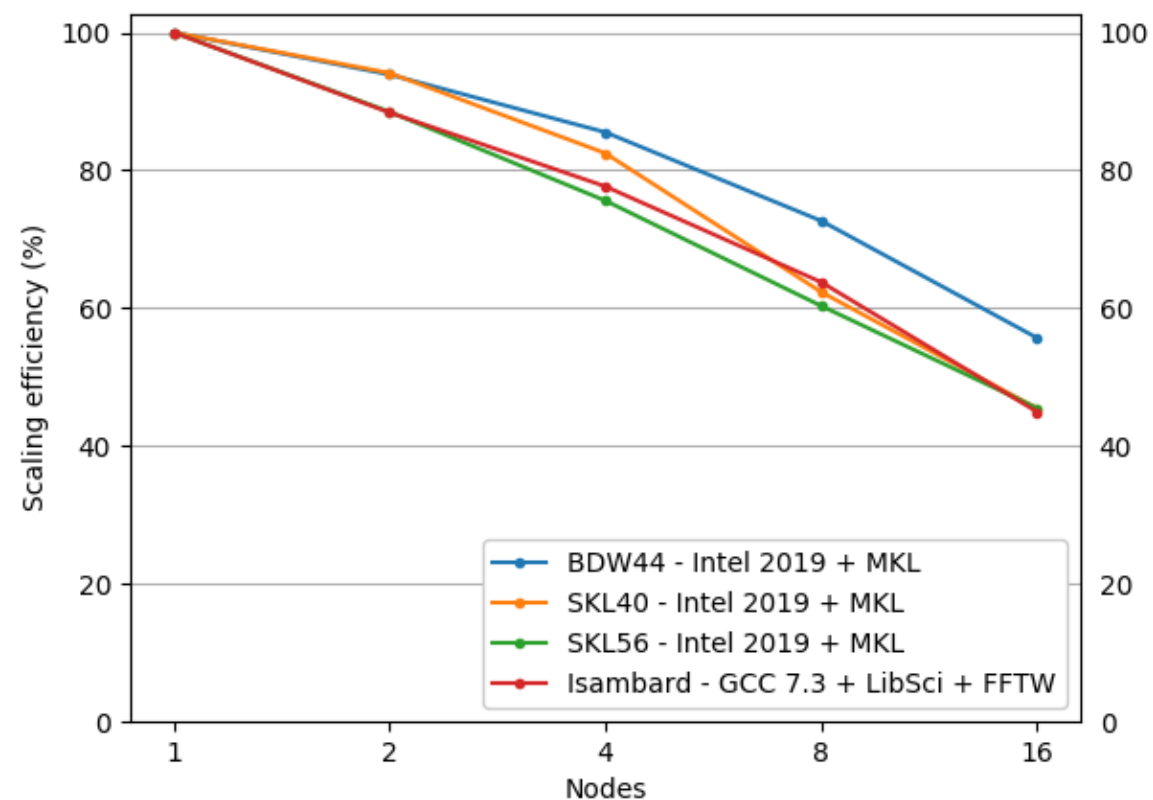


Parallel efficiency

VASP (PdO, 1392 atoms)



Relative performance



Parallel efficiency

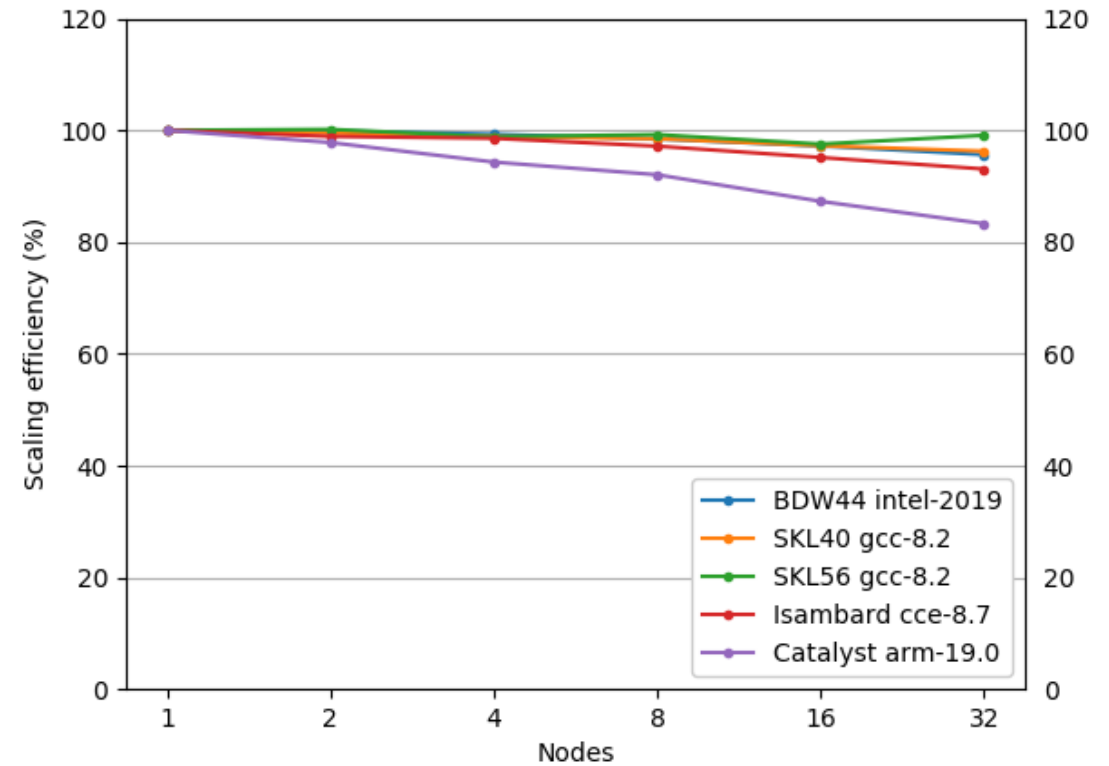
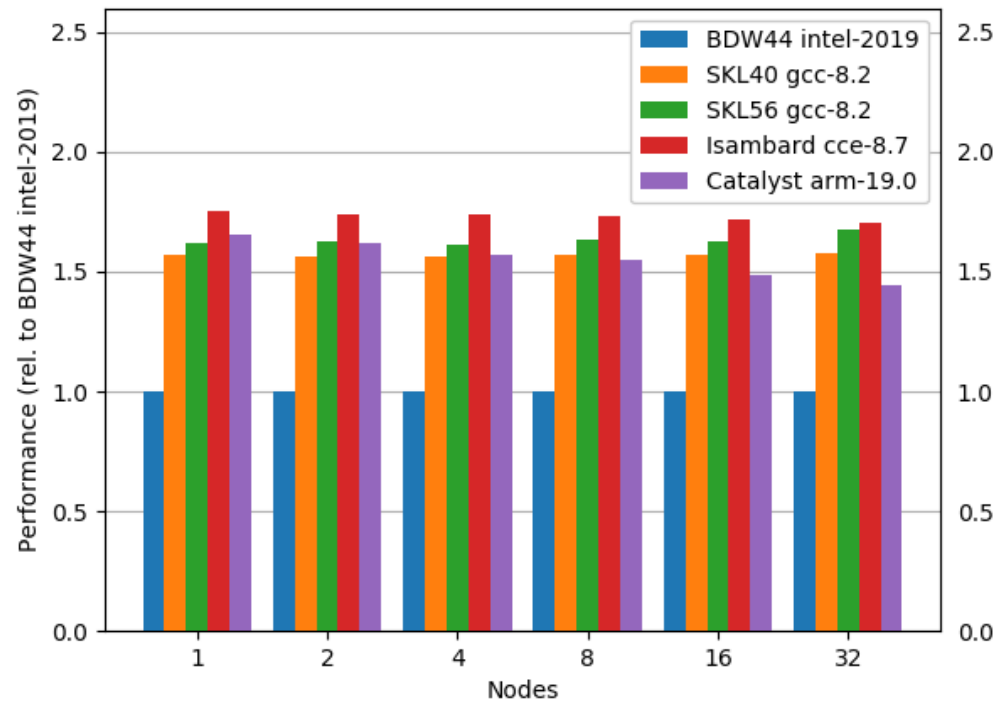
Isambard scaling summary

- **Arm-based systems appear to scale just as well as x86 ones**
- **For certain codes that were compute- or bandwidth-bound at low scale, these became network bound at 'real' scale, levelling the playing field**
- **We're seeing an issue with scaling in a few cases; appears to be in the MPI implementation** (currently investigating with Cray)

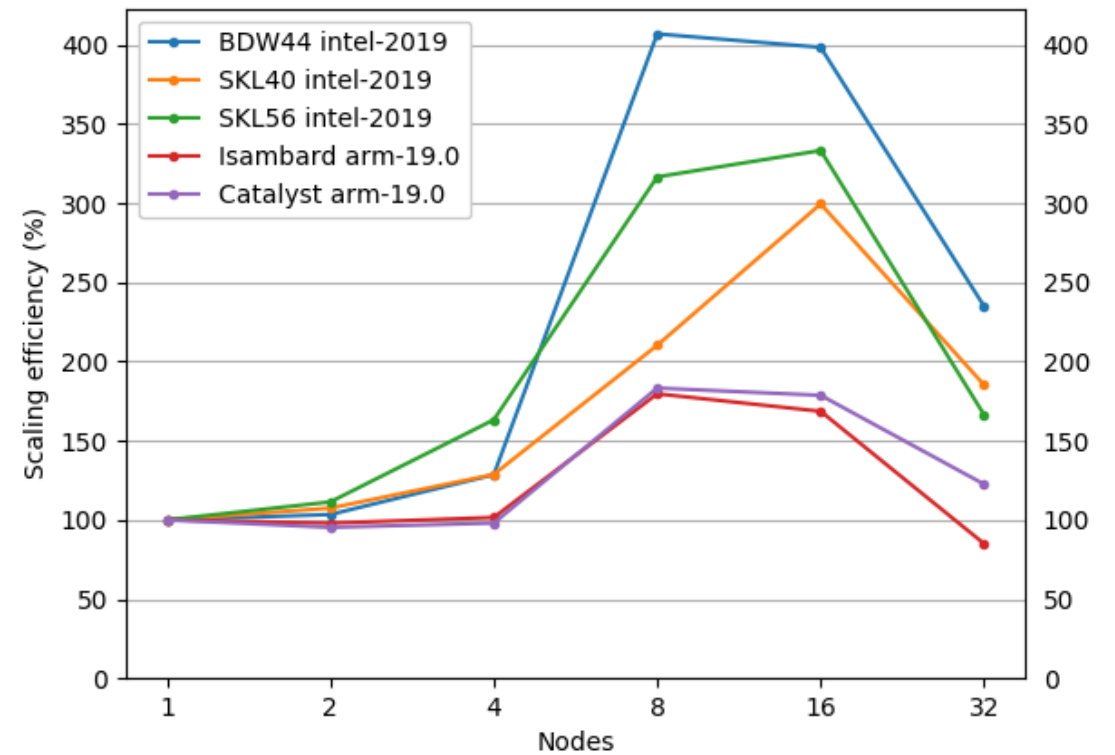
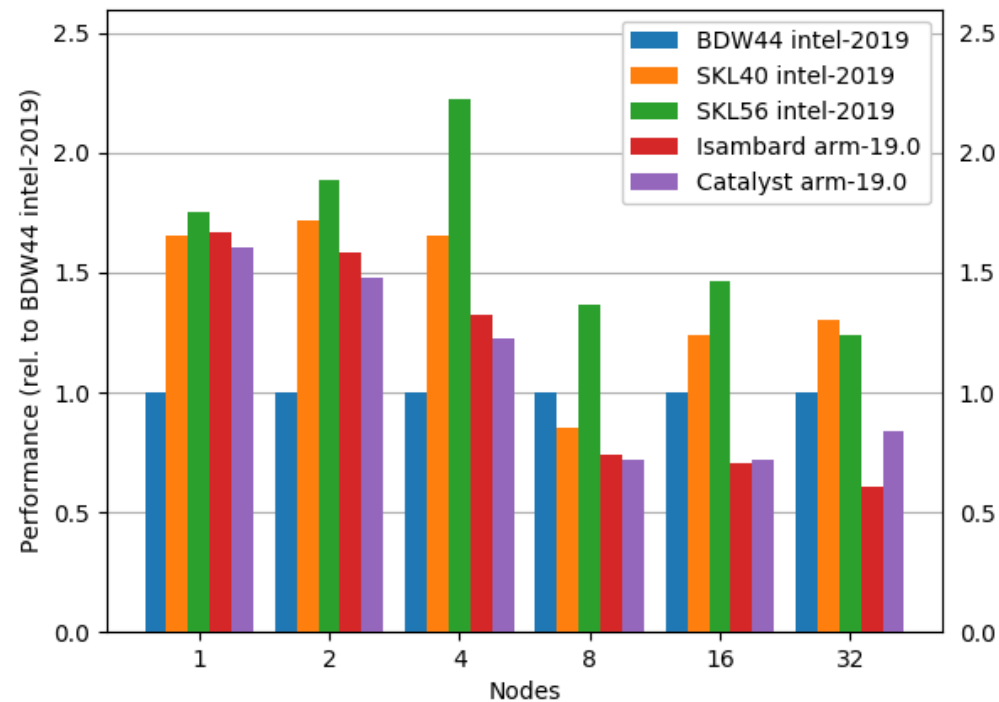
Comparing between multiple Arm-based supercomputers

- **Bristol** is one of the few sites in the world with multiple different Arm-based supercomputers
 - Added an HPE Catalyst system in 1Q2019
 - 64-node Apollo70 system, 4,096 cores, ThunderX2 CPUs
- **Isambard** and **Catalyst** together enable us to compare across:
 - **Networks:** Cray Aries vs Mellanox IB
 - **Software stacks:** open source vs Cray
- The following charts have never been shown in public before...

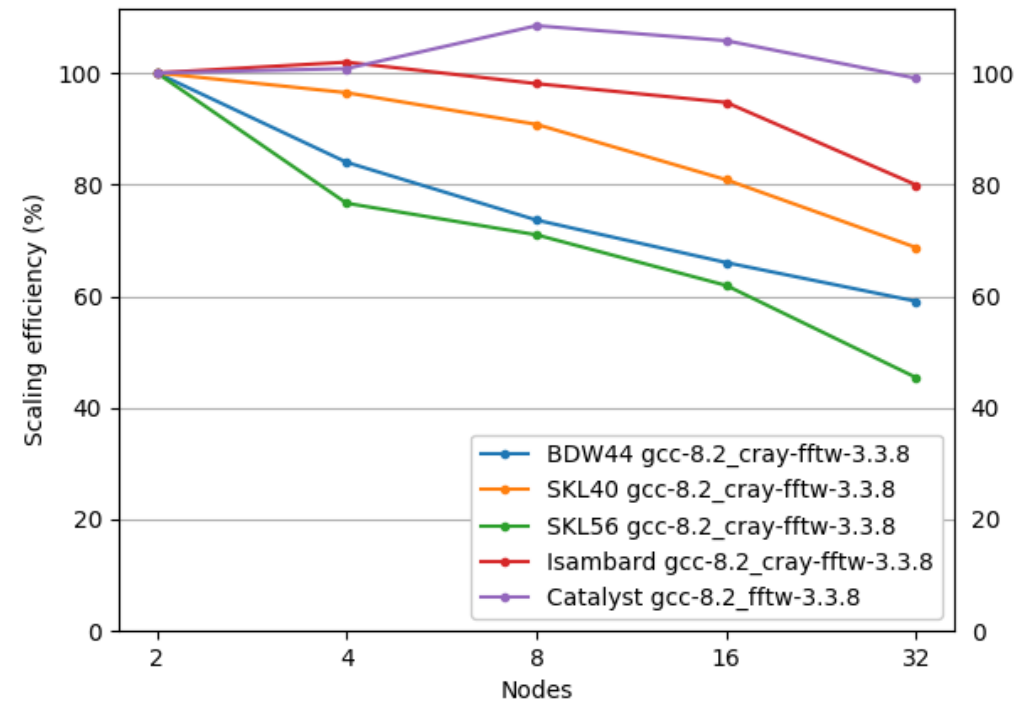
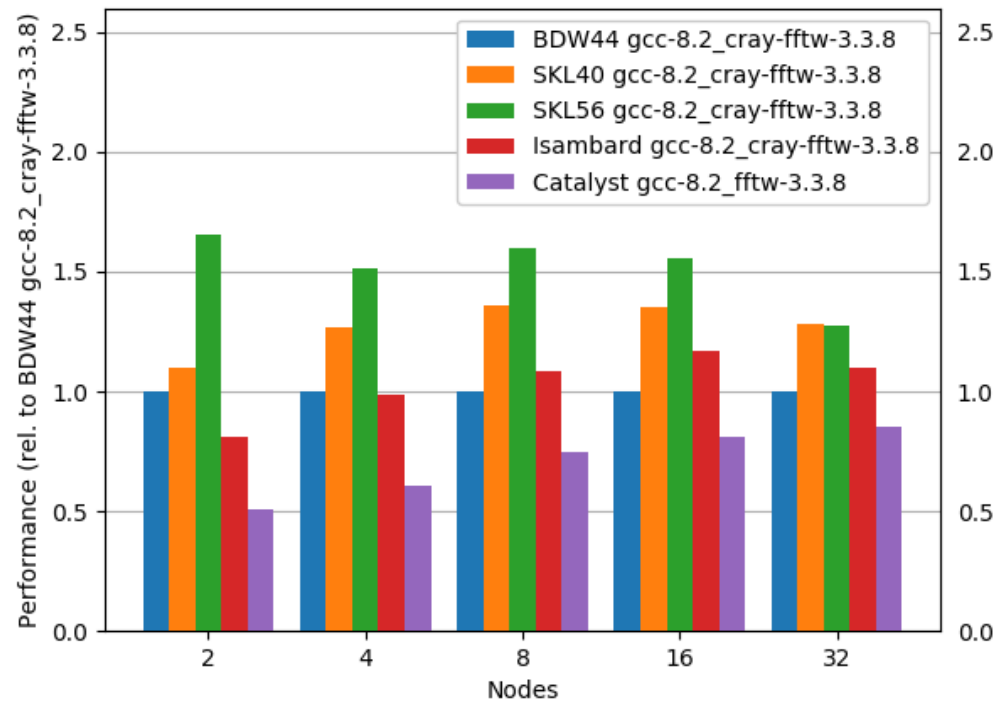
CloverLeaf – system comparison



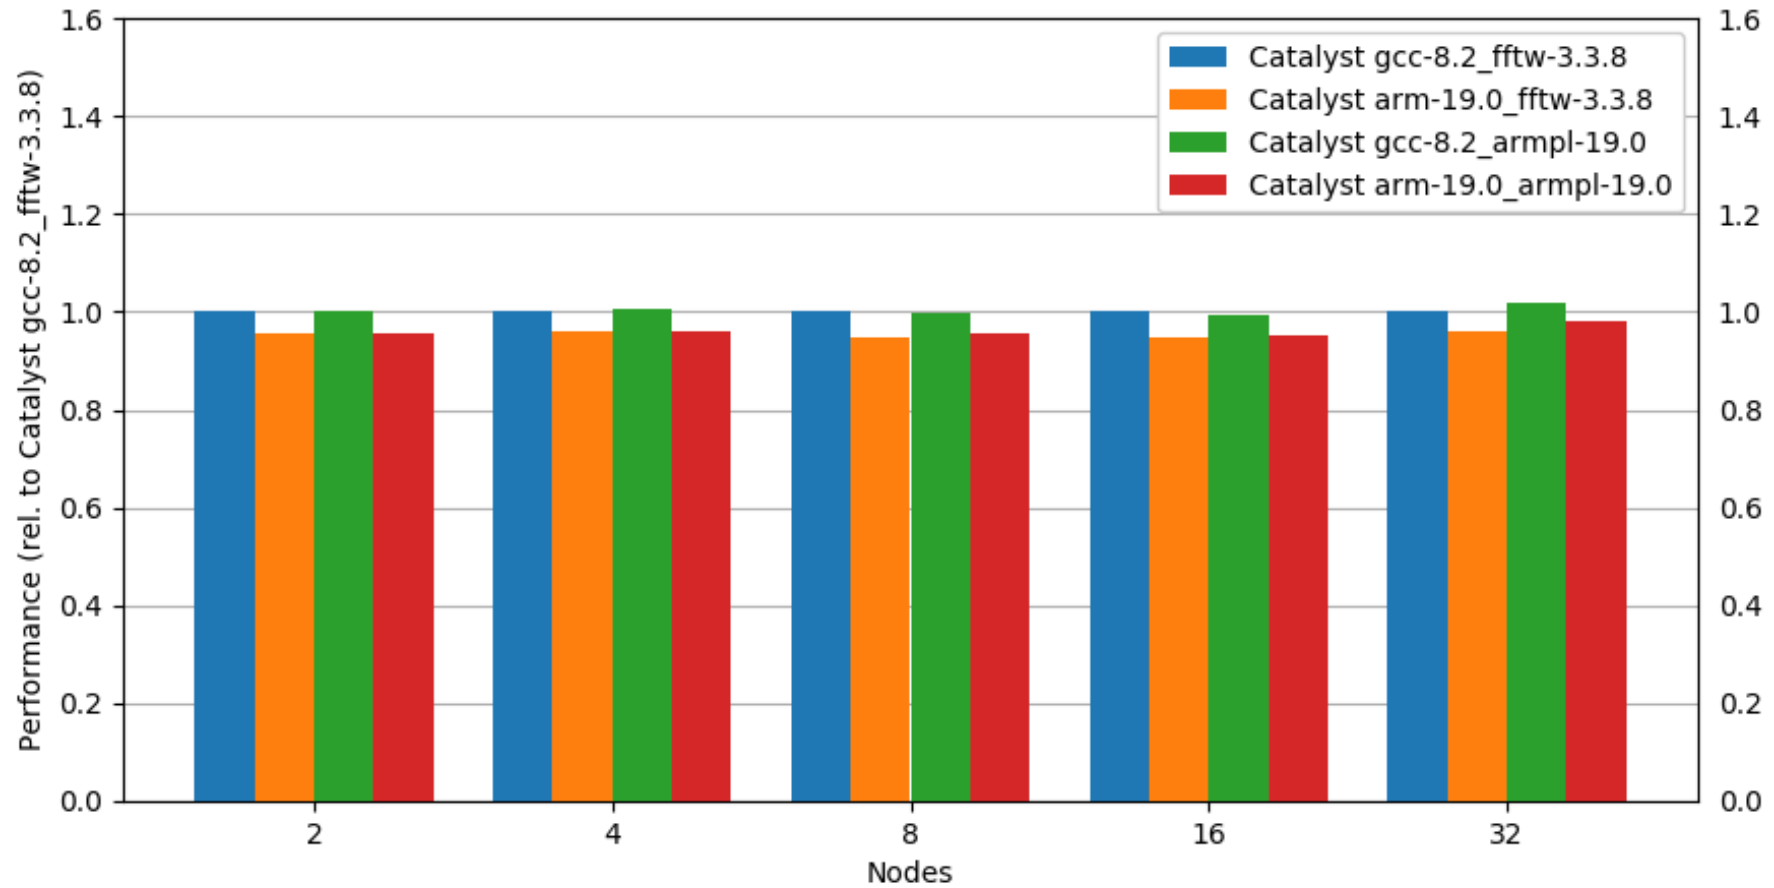
TeaLeaf – system comparison



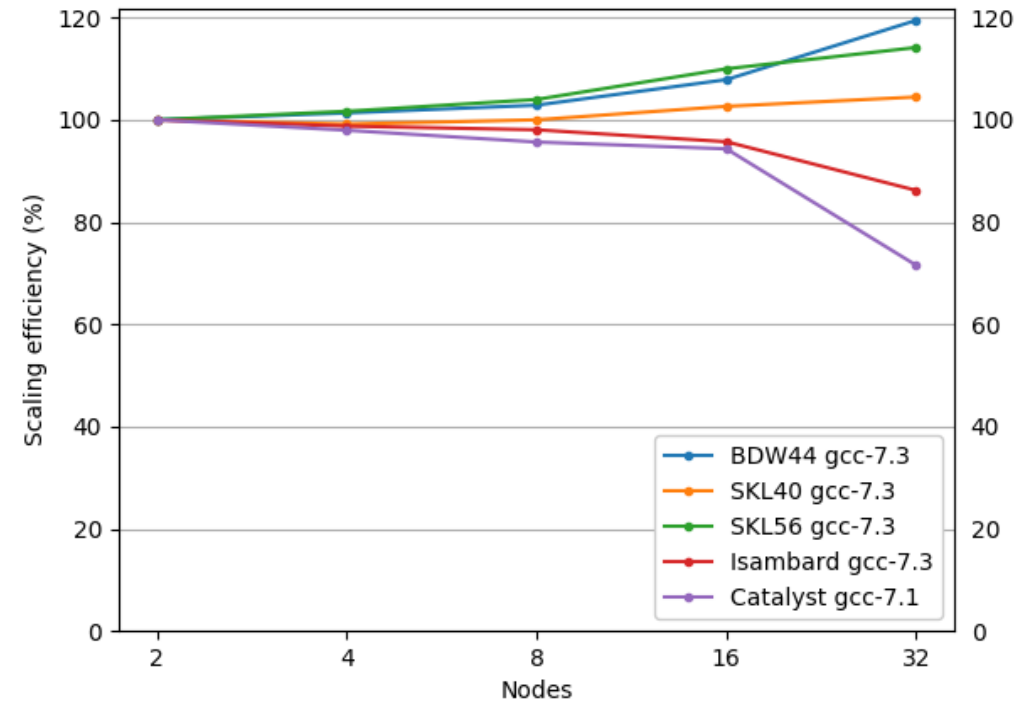
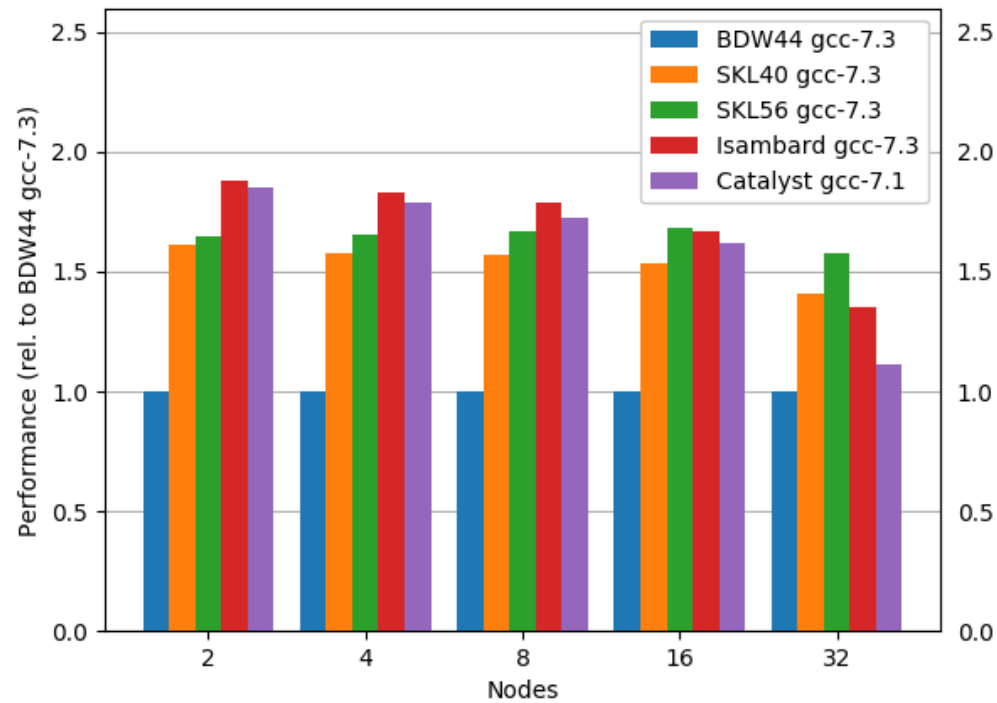
GROMACS – system comparison



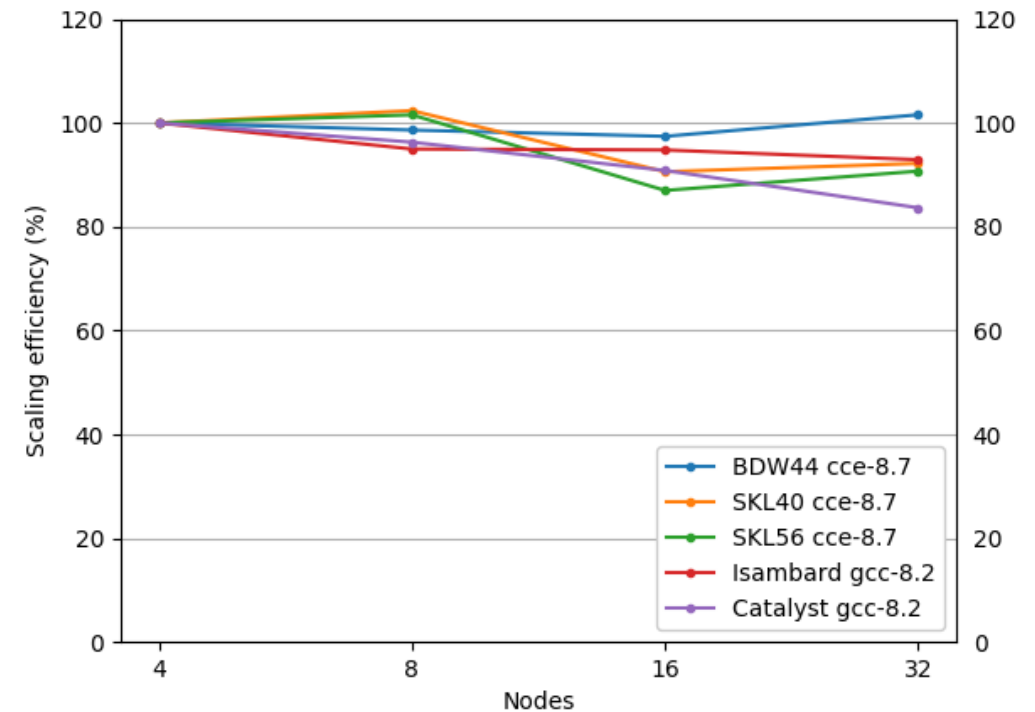
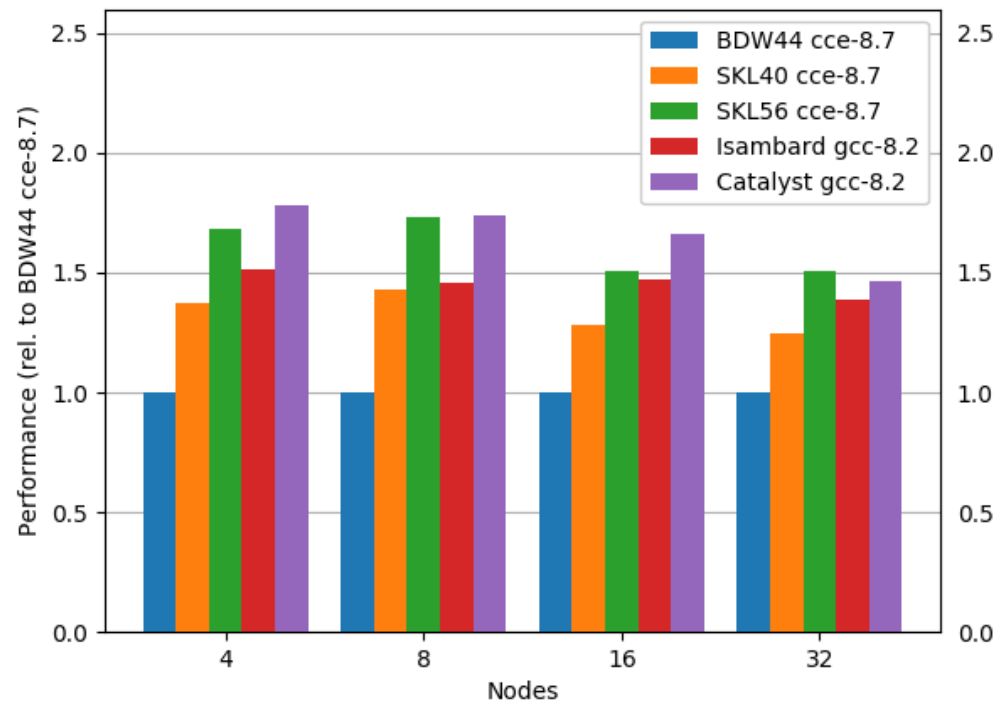
GROMACS - toolchains



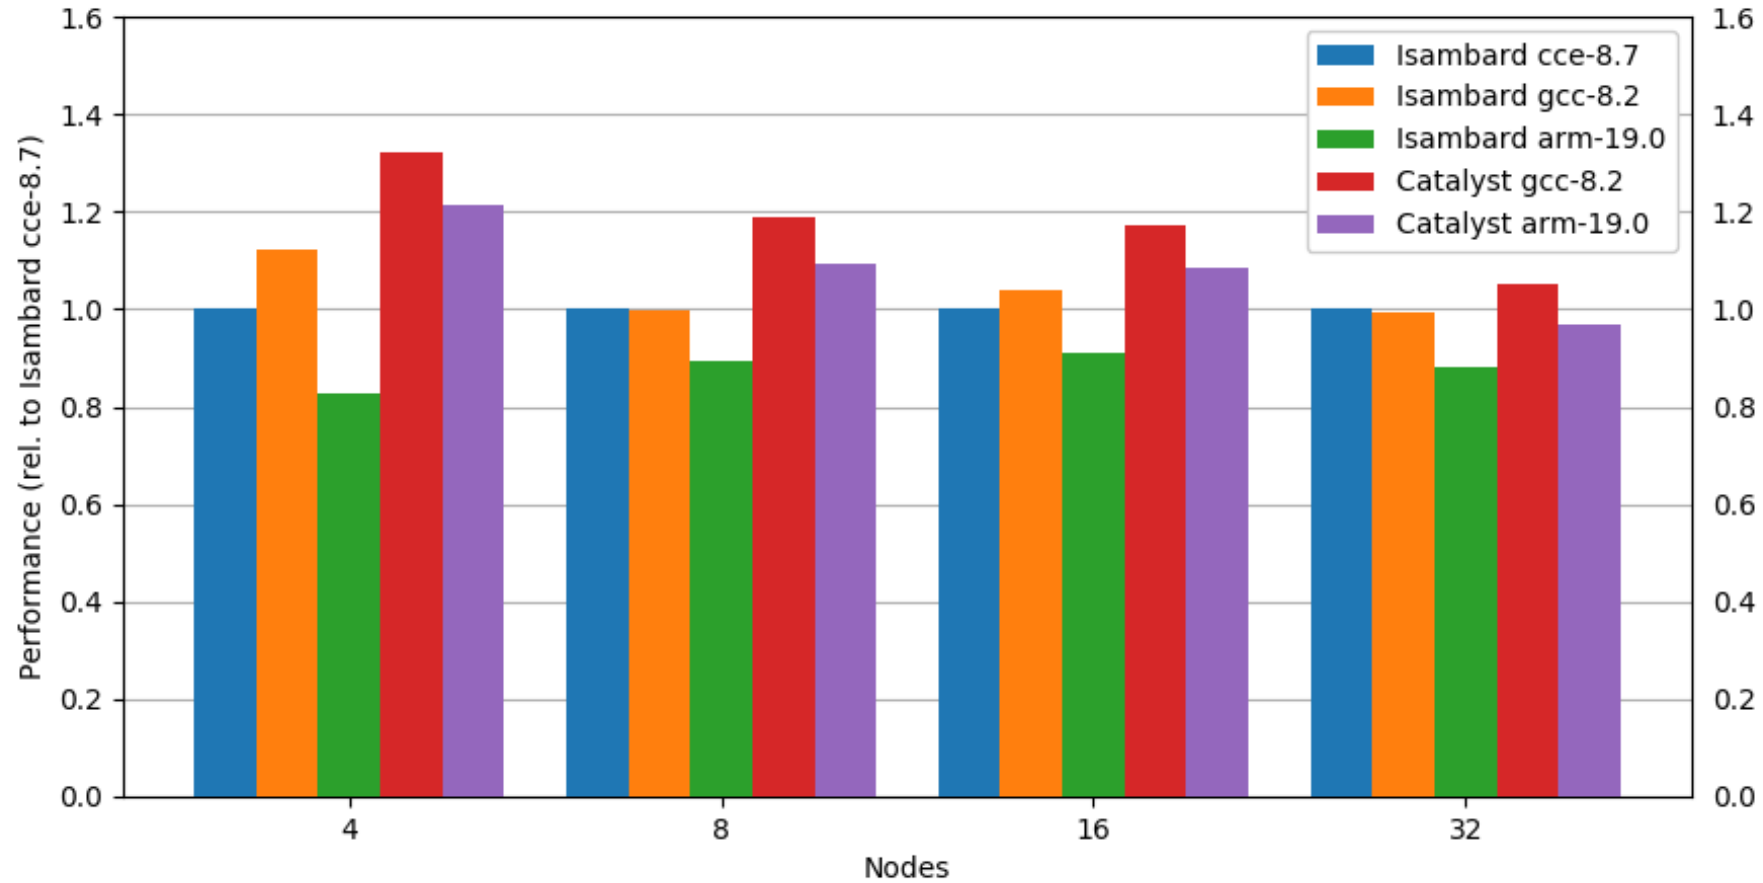
OpenFOAM – system comparison



OpenSBLI – system comparison



OpenSBLI - toolchains



Conclusions

- Arm-based supercomputers are now in production, doing real science
- Available from multiple vendors
- **Solid, robust software toolchains from multiple vendors**
 - Both open source and proprietary
- Arm-based systems scale just as well as x86 ones
- Now have evidence that **Arm-based systems are real alternatives for HPC**, reintroducing much needed competition to the market

For more information

Scaling Results From the First Generation of Arm-based Supercomputers

S. McIntosh-Smith, J. Price, A. Poenaru and T. Deakin, CUG 2019, Montreal

<http://uob-hpc.github.io/2019/06/07/CUG19.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>