

Isambard:

The World's First Large-Scale Production 64-Bit ARM Supercomputer

Prof Simon McIntosh-Smith

University of Bristol, UK

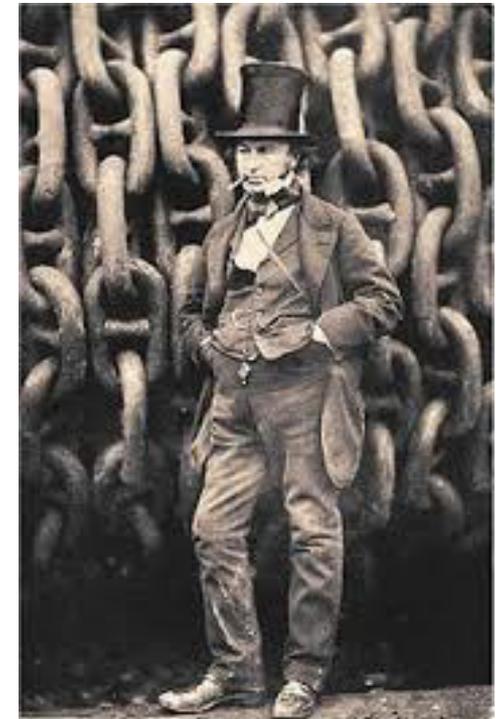
Recent trends in HPC:

- Most of the world's supercomputers are large collections of servers based on commodity processors, typically Intel's x86 CPUs
- **New computer architectures** are just starting to emerge, exploring diverse ways to provide the next jump in performance

Why explore ARM-based supercomputers?

- The architecture development is driven by the *fast-growing mobile space*
- Multiple vendors of ARM-based CPUs:
 - Greater **competition**
 - More **choice**
 - Rapid **innovations**, e.g. in vector instruction set
- **MONT-BLANC** proved the approach is feasible

**'Isambard', a new Tier 2 HPC service from GW4.
Named in honour of Isambard Kingdom Brunel**



I.K.Brunel 1804-1859

Tier 0: international

DoE



**The tiered model of
HPC provision**

Tier 1: national



Tier 2: regional

TIER 2 HPC CENTRES

Edinburgh

Cambridge

University College,
London

Loughborough

Bristol

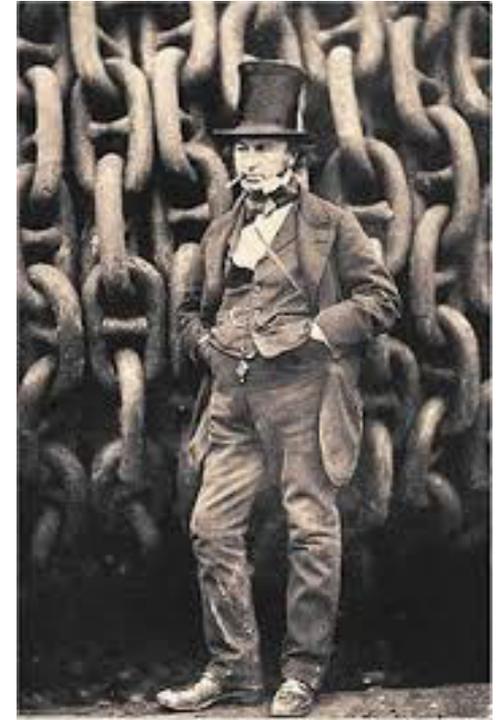
Oxford

Tier 3



Isambard system specification:

- Cray system
- **10,000+** ARMv8 cores
- Cray software tools
 - Compiler, math libs, tools...
- Technology comparison:
 - x86, Xeon Phi, Pascal GPUs
- Phase 1 installed March 2017
- The ARM part arrives early 2018
 - Early access nodes from Sep



I.K.Brunel 1804-1859

Some results:

- My team has had remote access to ARM-based Cavium ThunderX2 early silicon for a few weeks
- We have already managed to compile and run lots of real codes “out of the box”
- Using both GNU and Clang/Flang/LLVM
- Performance already looks very promising!
 - Details will go public later this year

-
- Our focus will be on the top 10 most heavily used codes on Archer as measured in 2017:
 - VASP, CASTEP, GROMACS, CP2K, UM, HYDRA, NAMD, Oasis, SBLI, NEMO
 - Note: 8 of these 10 codes are written in **FORTRAN**
 - Additional codes relevant to project partners:
 - OpenFOAM, OpenIFS, WRF, CASINO, LAMMPS, ...
 - We want to collaborate wherever possible!
 - Need to avoid duplicating effort

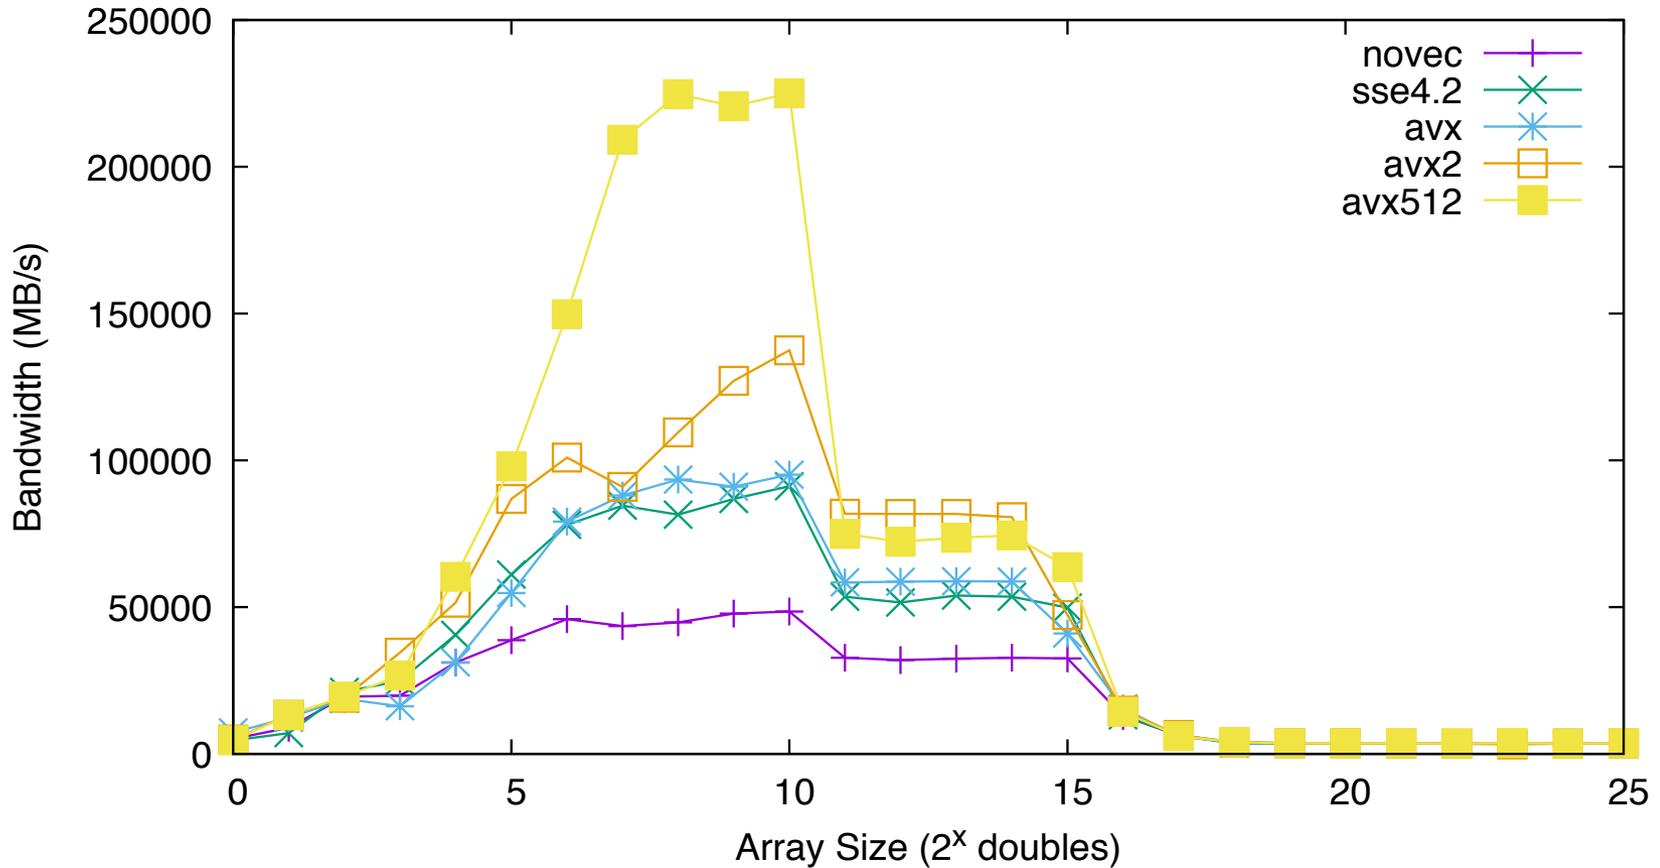
What do we already have working?

- VASP – periodic electronic structure application
- GROMACS – molecular dynamics
- CP2K – quantum chemistry
- OpenFOAM – CFD
- The Unified Model (UM) – weather & climate
- SNAP – deterministic neutral particle transport
- CloverLeaf – hydrodynamics
- TeaLeaf – heat diffusion
- STREAM – memory bandwidth

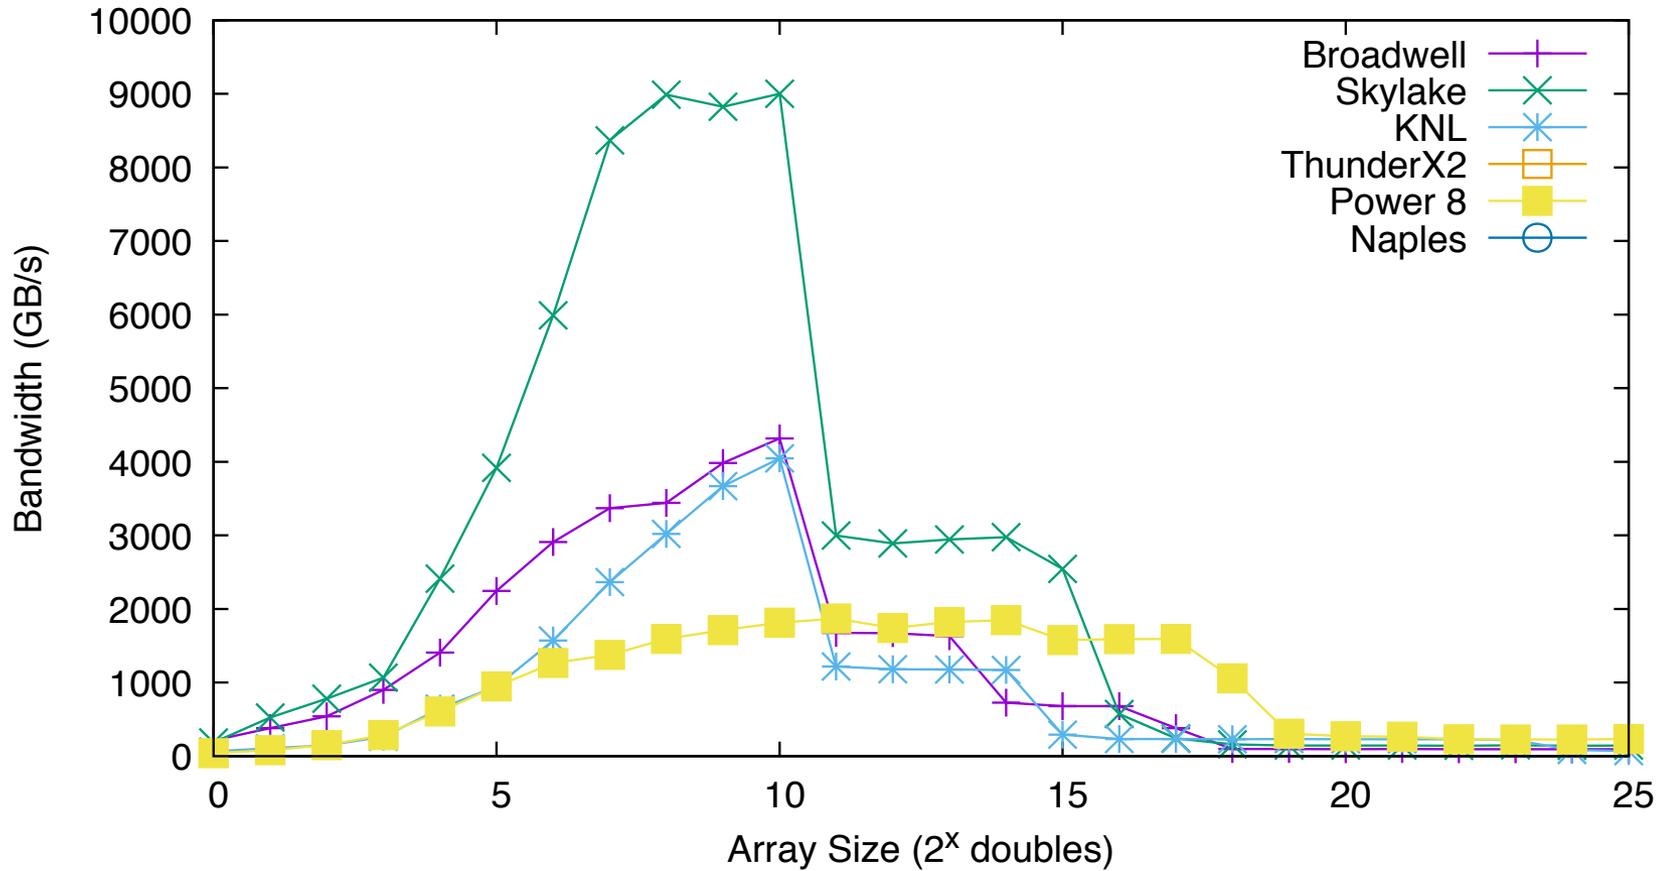
A quick note about vectorisation for NEON / SVE:

- The performance benefit is often from better use of bandwidth, not the FLOPS
- The following graphs are from a version of STREAM modified to measure bandwidths through the cache hierarchy (SC17 poster)

Single core cache bandwidth as measured from a single core - Skylake



Aggregate bandwidth for a single node



Next steps: the first Isambard hackathon

- We have a 2-day intensive session planned for the start of November at Cray's offices in Bristol
- Will focus on the top few codes of our hit list:
 - VASP, CP2K, GROMACS, UM
- Will have access to Cray's ARM compiler
- Have some of the core applications vendors attending the hackathon

Hot off the press: Isambard shortlisted for a Times Higher Education award in the category of “Technological Innovation of the Year”



<http://www.the-awards.co.uk/2017/en/page/shortlisted-page-2017>

Exciting times ahead!

- **In the near term:** greater choice, competition, and innovation, *for production science*
- **In the longer term:** potential for disruptive change to significantly impact scientific software development
- We need to be ready to exploit next generation technologies to deliver new science!

For more information:

- <http://gw4.ac.uk/isambard/>
- <https://www.epsrc.ac.uk/blog/isambardhpc/>
- Twitter: @simonmcs
- Please get in touch to collaborate on optimising applications:
simonm@cs.bris.ac.uk
- Bristol is hiring faculty staff in CS to focus on HPC