RISC-V: POWERING THE FUTURE OF HIGH PERFORMANCE COMPUTING?

Nick Brown (EPCC)

n.brown@epcc.ed.ac.uk

epcc

THE UNIVERSITY OF EDINBURGH

Why for HPC?

1. **Openness** where anyone can freely take the standard and build compatible, but specialised, solutions
   - Means that hardware can be driven by HPC requirements and market demand.

2. **Community driven** resulting in participants ranging from hardware vendors to software developers able to shape the standard
   - HPC community can get involved and have a say in the evolution of RISC-V to ensure it meets our needs.

3. **Modularity** the standard can be used in a flexible manner enabling development of bespoke hardware that targeted at a specific range of applications
   - Specialisation to deliver improved performance and energy efficiency for HPC

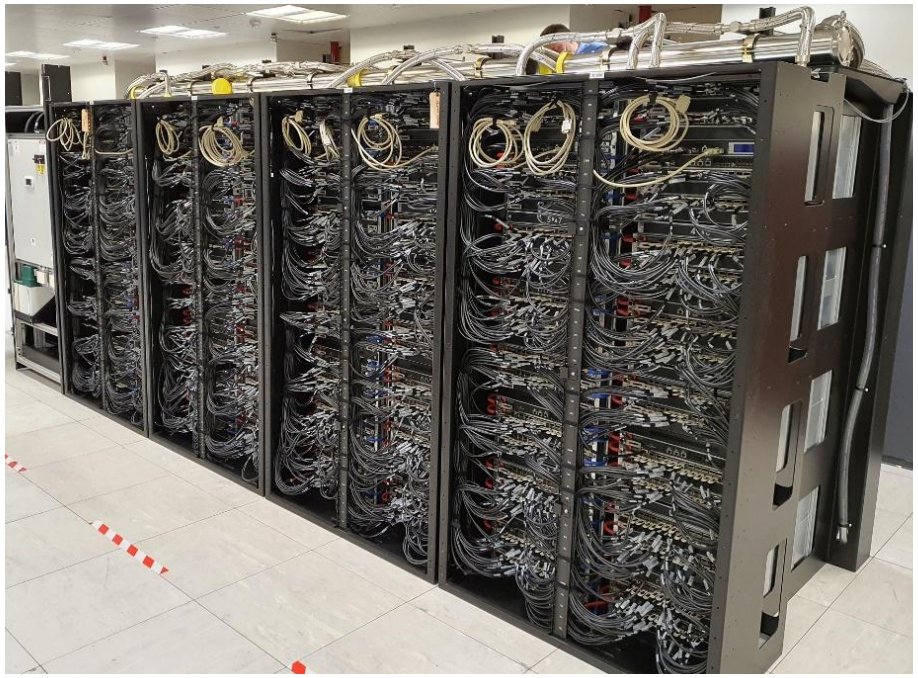

Already a (reasonable) ecosystem

- We set up a RISC-V HPC lab for free access
  - Initially embedded SoCs which are not realistic for HPC, but enable people to kick the tyres
  - https://riscv.epcc.ed.ac.uk

- Generally, a lot of the libraries and frameworks you would expect are either available already or easy to build
  - Although often not optimised specifically for RISC-V
- Surveyed applications on ARCHER2, a Cray-EX
  - 27.7% (35% when ignoring unknown category) of runtime is by applications that run on RISC-V
  - 55% of users can run their applications on RISC-V

Application	Percentage ARCHER2 usage	Number of users	Supports RISC-V?
VASP	18.4%	153	No
Met Office UM	8.9%	37	No
GROMACS	6.8%	115	Yes
Python	4.2%	64	Yes
LAMMPS	4.1%	45	Yes
OpenFOAM	4.1%	49	Yes
CP2K	4.0%	48	Yes
Nektar++	2.7%	14	No
CASTEP	2.5%	46	Yes
NEMO	1.8%	19	No
SENGA	1.3%	5	No
CESM	0.7%	10	Yes
Quantum Espresso	0.7%	46	Yes
NAMD	0.5%	3	Yes
WRF	0.1%	4	Yes
Other	12.4%	121	No
Unknown	26.8%	-	N/A

epcc | THE UNIVERSITY OF EDINBURGH

64-core SG2042: A commodity RV HPC CPU?

- First commodity available high-core count RISC-V CPU
  - Includes 128-bit RVV 0.7.1
  - Core for core is faster than SpacemiT K1 which itself provides RVA22 + RVV 1.0
  - NASA's parallel benchmark suite kernels
    - Thanks to E4 compute engineering for access to their dual socket RISC-V system!

Hardware	Mop/s	% performance of SG2042
Sophon SG2042	472.97	-
VisionFive V2	121.62	25.7%
VisionFive V1	39.31	8.3%
SiFive U740	49.03	10.4%
All Winner D1	47.71	10.1%
Banana Pi F3	146.83	31.0%
Milk-V Jupyter	158.64	33.5%

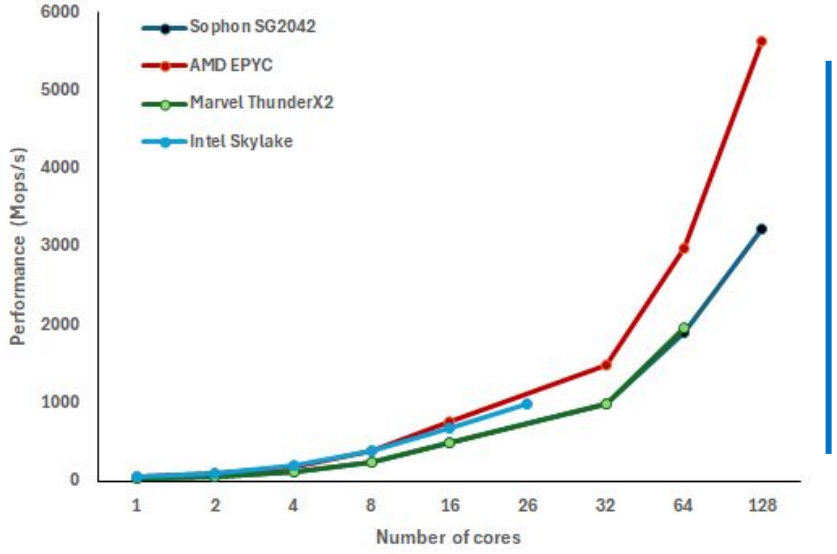

Embarrassingly Parallel
Compute bound

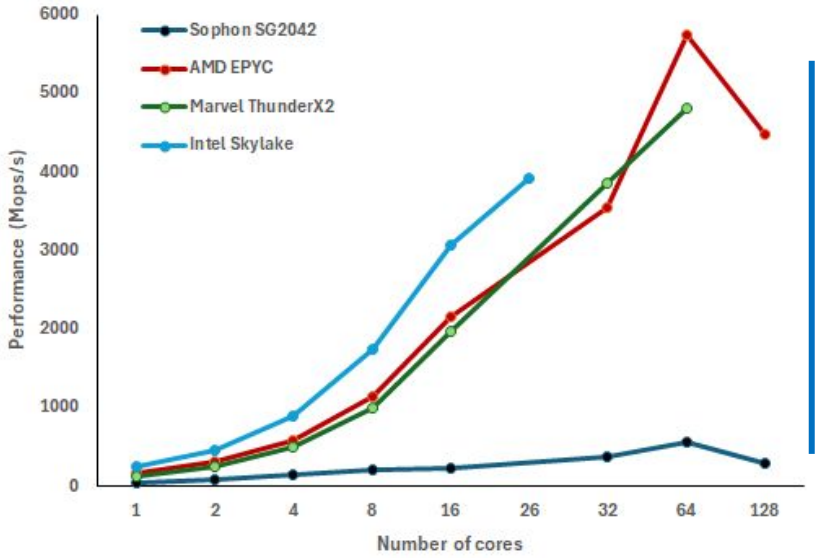

Integer Sort
Memory latency bound

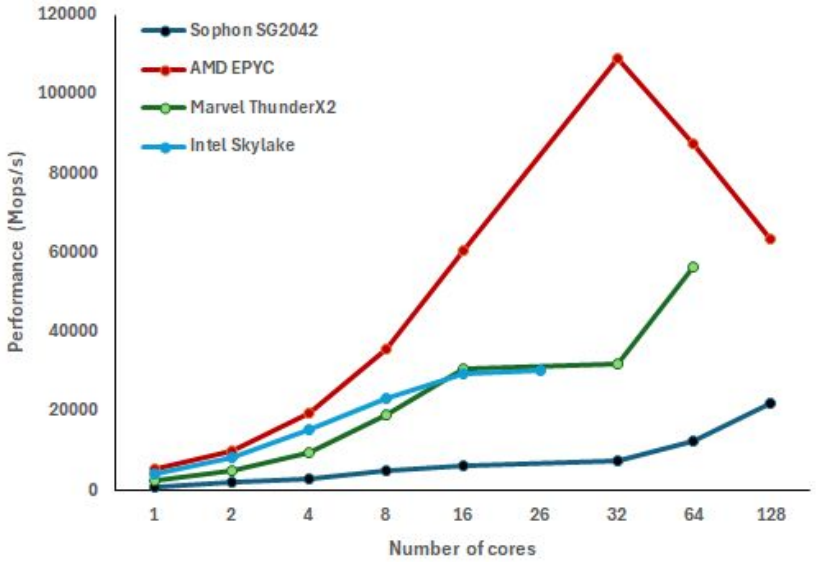

Multi Grid (MG)
Memory bandwidth bound

Missing parts of the ecosystem that we need to address

- ***Software tooling***
  - We need better support for profilers and debuggers
    - Not just the tooling, but also ensuring performance events specification that is being worked on currently suits HPC
  - Cluster management tools are needed (often supplied by vendors)
- ***Scaling out***
  - ***Filesystems:*** support for BeeGFS, NFS but not Lustre or GPFS
  - ***Networking***: Open Fabrics Interface (OFI) ported to RISC-V, so theoretically high-performance networking stacks should run but these require maturing & vendor support
  - ***GPUs***: Are ubiquitous in HPC, people are experimenting with driving Nvidia and AMD GPUs from RISC-V, but work is still early

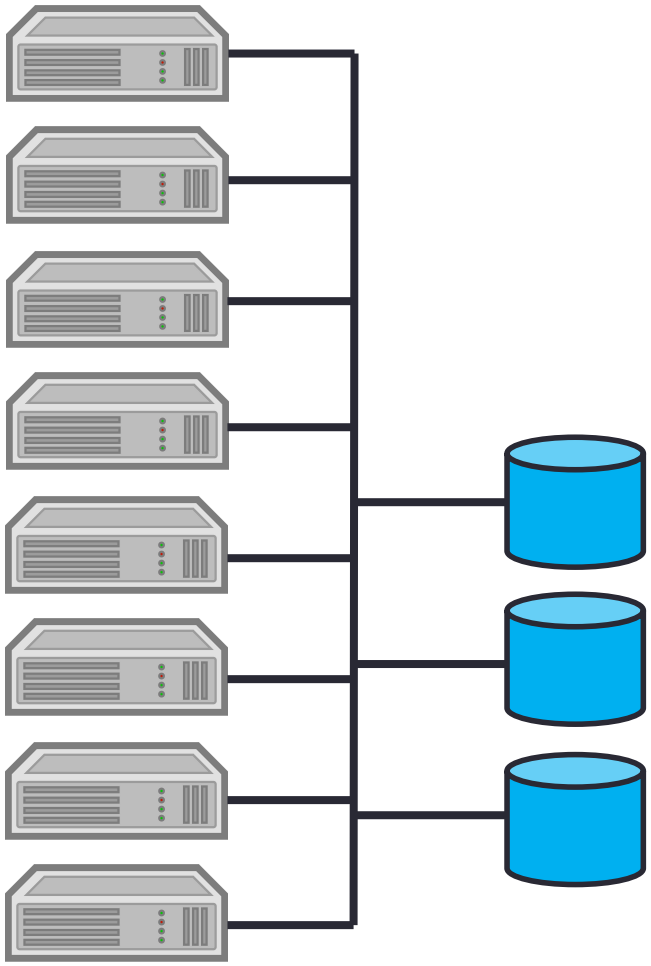

RISC-V accelerators

- Potentially, RISC-V accelerators are an easier adoption proposition for the HPC community
  - Can be fit into existing systems, don't need to worry about other aspects such as networking, filesystems etc

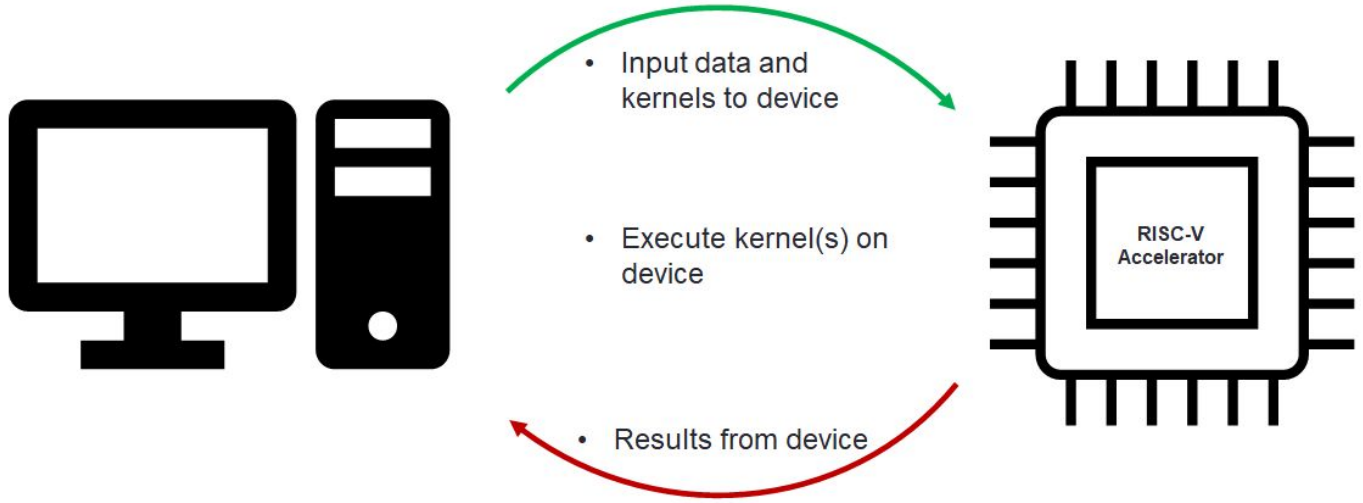

- HPC is currently dominated by GPUs
  - But these are power hungry, and vendors are more and more focussing on lower precision than FP64
  - So having other options will be very beneficial
    - Currently codes are redesigned for the (GPU) accelerator where possible, with RISC-V we can work the other way where accelerators are designed for the needs of codes
    - The flexibility and choice enabled by RISC-V is key here

A whole load of RISC-V accelerators for HPC….

Accelerating HPC on the Tenstorrent Wormhole

- Tenstorrent have committed to an *open first* approach with their software ecosystem
  - Which makes it much more accessible to directly program their architecture via their Metalium framework

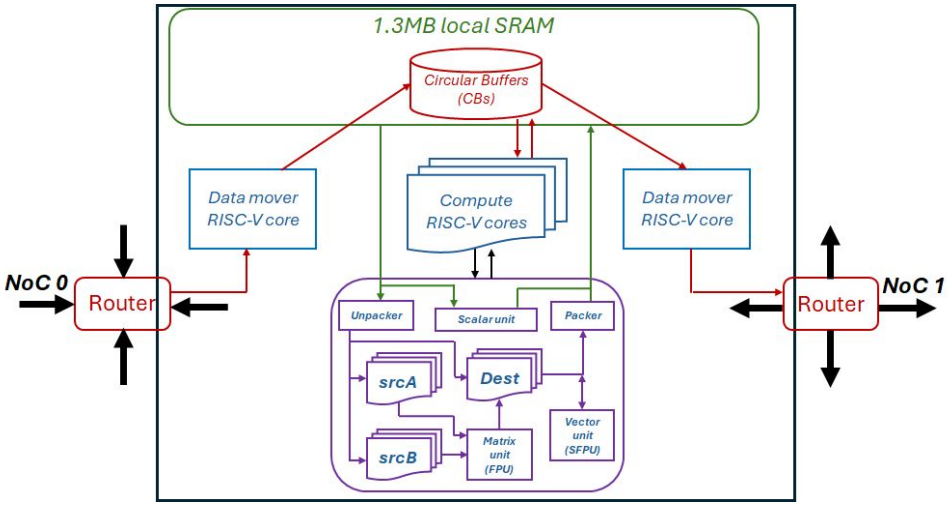

- Stencil based algorithms are ubiquitous in HPC
  - On the Grayskull over 108 Tensix cores performance is comparable to a 24-core Xeon Platinum (Cascade Lake) but at 5 times less energy

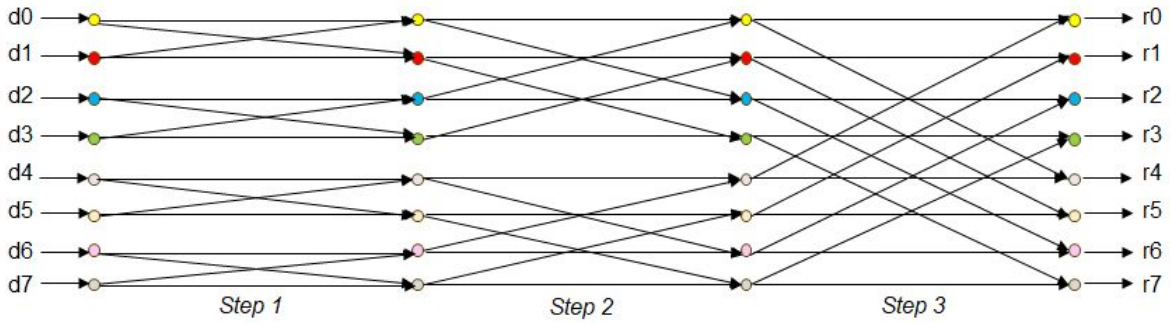

- FFTs were voted the most important numerical algorithm in our lifetimes
  - Used extensively in HPC!
  - 2D FFT is around half the performance of a 24-core Xeon, but at over 5 times less power draw
  - And the Wormhole card is much cheaper too!

HPC community engagement

- The HPC community is showing interest
  - HPC SIG organises workshops, panels and collaborative BoFs at key HPC conferences, with attendance growing
    - Our next workshop is at ISC next month, with a packed programme https://riscv.epcc.ed.ac.uk/community/workshops/isc25-workshop/
    - Call for papers is open for the RISC-V for HPC workshop at SC in November https://riscv.epcc.ed.ac.uk/community/workshops/sc25-workshop/
- Recent HPC presence on the RISC-V website
  - https://riscv.org/industries/risc-v-hpc/

Conclusions and action points….

- Flexibility provide by RISC-V is a clear benefit to HPC and progress is being made
- The next generation of high performance RISC-V CPUs has potential to form a key alternative for HPC machines
  - But “everything else” must be ready including applications, libraries tooling etc.

- The HPC SIG has identified several high priority areas to address based around gaps that currently exist in the ecosystem and we welcome comments/feedback
  - *https://edin.ac/4mahlV5*
- But also, critical to engage more closely with specific task group activities
  - Matrix extension task groups, once standardised ensure these can be exploited by HPC
  - Performance events specification TG

epcc

THE UNIVERSITY OF EDINBURGH