

The HPC Hardware and Development Lab @Durham

Durham University

Alastair Basden, Billy Carlyle, Mark Lovell, Paul Walker

COSMA

Alastair Basden
DiRAC / Durham University



DiRAC
High Performance
Computing Facility

What is it?

- A collection of new and novel HPC hardware
- Accessible to all UK-based academic researchers
 - So that they don't have to get their own
 - Single login to access all components
 - Shared file system for data
 - >350 users signed up in Oct 2026
 - 80% from outside Durham
- To facilitate code porting, debugging and testing
 - Code profiling, performance analysis
 - Advising future system purchases



- The HPC Hardware Lab is part of the COSMA facility:
 - Est. 2001
 - Includes the DiRAC Memory Intensive HPC service
 - COSMA7, COSMA8
 - primarily covering STFC science
 - Alongside the 3 other DiRAC systems
 - The Institute for Computational Cosmology dedicated computing system
 - COSMA5
 - The HPC Hardware and Development Lab



DiRAC

- Established 2009
- Provides HPC to the STFC theory community
 - Particle Physics, Astrophysics, Cosmology, Solar System and Planetary Science and Nuclear Physics
- Three services:
 - Extreme Scaling: Edinburgh (TURSA)
 - Data Intensive: Cambridge and Leicester (CSD3 and DiaL)
 - Memory Intensive: Durham (COSMA)
- Co-designed and tailored for specific workloads
 - Bespoke systems for the science being carried out
- The Hardware Lab offers key insights for the DiRAC co-design process
 - Instrumental for TURSA, COSMA8

HPC Hardware Lab @Durham

- Mission: Provide access to the latest HPC hardware to researchers and DRI/RTP professionals from across the UK
 - For code testing, performance tuning and debugging
 - To advise on purchase of future technologies
 - To allow informed decisions to be made whenever funding appears
- Accessible to all
 - Single authentication system to lots of novel hardware
- An emphasis on production/going live
 - No shelf-ware!

History



- Came together almost by accident
 - 2019: Intel provide a 56-core 6TB Cascade Lake system
 - For testing Non-volatile DIMM performance (Apache Pass)
 - 2019: University funding for the DINE cluster
 - 16 node BlueField DPU test system
 - 2019: ExCALIBUR announced, with a H&ES component
 - Hardware and Enabling Software: \$4.5m
 - Perfect timing for the hardware lab

Hardware Lab Components

- DINE: Durham Intelligent Networking Environment
- DINE2: Durham Integrated Next-gen Environment
- CPU compute
 - X86, ARM
- GPU compute
 - NVIDIA, AMD, Intel
- Composability
- Networking
- Storage laboratory
- Environmental
 - Solar panel installation
 - Immersion cooling
 - Heat storage
 - Logging and awareness
- Quantum

DINE

- Durham Intelligent Network Environment
 - A 24 node (initially 16) system for investigation of networking technologies
 - And other things
 - 32 cores and 512GB RAM per node, AMD ROME
- Has hosted (since 2019):
 - BlueField-1
 - BlueField-2 (host-separated and embedded modes)
 - Rockport 100G 6D Torus Ethernet network
- Now used for production workloads
 - Charging usage by kWhr
- Funded by Durham, DiRAC, ExCALIBUR



DINE-2

- Durham Integrated Next-gen Environment
 - 8 node Intel Sapphire Rapids system
 - 64 cores, 2TB RAM per node
 - Currently hosts a CerIO composable PCIe-5 fabric
 - 12x NVIDIA GPUs, assignable in any number to any host
 - 400G Ethernet test network (now relocated)
- Used for production workloads
 - Also charges for usage by kWhr
- Funded by DiRAC, IRIS and SKA


CPU compute

- Providing users with access to cutting edge CPU technologies:
- X86:
 - AMD: Turin, Genoa, Bergamo, Milan, Milan-X, Rome
 - Milan-X has 768GB Cache
 - Intel: Sierra Forest, Emerald Rapids, Sapphire Rapids, Ice Lake, Cascade Lake, Sky Lake
 - Cascade Lake has 6TB RAM (Apache Pass)
- ARM:
 - NVIDIA: Grace
 - Exciting new system due shortly
- RISC-V: Tenstorrent QuietBox with Blackhole cards
 - High VFM inferencing
- Funded by OEMs, Intel, AMD, Dell, DiRAC, ExCALIBUR



GPU Compute

- Access to small numbers of new (and not-so new) GPUs
- AMD MI300X
 - 8x MI300X, cache coherent global address space
- AMD MI300A
 - Integrated CPU-GPU (APU)
- AMD MI210, MI100, MI50
- NVIDIA RTX PRO 6000
- NVIDIA H100 (Grace-hopper system)
 - PCIe version also coming soon, X86
- NVIDIA A100, A30, V100
- Intel Ponte Vecchio
- Direct ssh (including Jupyter) and queue-based access
- Funded by Dell, AMD, Intel, DiRAC, ExCALIBUR, IRIS



Credit: AMD


Dual access model

- Many of the systems are single nodes
 - Aim to facilitate rapid developer access
 - And also dedicated profiling/benchmarking
- Slurm-by-night / SSH-by-day:
 - Free-for-all for testing and debugging
 - Slurm queue allowing 6-hour jobs at night
 - To facilitate single user testing and profiling


Composability

- Infrastructure-as-a-service
 - The ultimate goal for commercial cloud-type systems
 - Is it relevant for HPC?
 - How does performance suffer?
 - Is it stable?
 - What are the use cases?
- 2 composable systems
 - And a Gen-Z test system
 - Defunct: technology transferred to CXL

Liquid composable system





- Installed 2021
- A PCIe4-based composable fabric: 4 lanes to each server
- 3 A100 GPUs shared between 4 servers
 - Including a login node
 - GPUs per node can be changed in a few clicks
- 12TB RAM shared between these servers
 - Can be changed and reconfigured
- Positives: It works, can allow high RAM/GPU configs
- Negatives:
 - Bandwidth is shared
 - GPUs sometimes fail requiring a full stack reboot
 - RAM/kernel issues (and no Rocky9 support yet)
 - Rack-scale limitations
 - Bottlenecks
- ExCALIBUR funded





CerIO composable system

- Installed 2024
- PCIe-5-based fabric
 - No central switch: Uses a flit-based torus topology
 - Full data centre scalability
 - 200Gb/s to nodes
 - 300Gb/s inter-card bandwidth
- 8 compute nodes, 8x A30 GPUs, 16TB RAM
- Will allow networking and composability within a single fabric
 - Exciting for future compute designs
- IRIS/SKA/ExCALIBUR/DiRAC



Network fabrics

- Rockport 6D Torus network
 - A “switchless” fabric for 100G Ethernet
 - Trained on DINE in 2021
 - Installed on COSMA7 in 2022
 - 224 nodes (half the cluster) replaced IB
 - Allows direct comparison of fabrics
 - At full HPC problem-size scale
 - Works well
 - Performance comparable to InfiniBand
 - For real workloads
 - Handles congestion well
 - ExCALIBUR/DiRAC funded
- 400/800G Ethernet fabric
- InfiniBand



Storage sub-lab

- Various different storage technologies
 - Most make it into production
- High-performance scratch Lustre (NVMe)
- DAOS (NVMe)
- Ceph (Object and bulk storage)
- StorJ private cloud
- VAST (NVMe)
- Globus (data transfer)
- Lustre (efficient bulk parallel storage)
- ZFS (Bulk storage)
- Tape: Atempo Miria
- NVMe RAID test systems
- Funded by DiRAC/IRIS/SKA/ExCALIBUR





The Net-Zero lab

- All HPC-related things relevant for Net-Zero

The Net-Zero lab

- HPC is a huge energy user
 - COSMA ~1MW at peak
 - (~5GWh/year, ~120T CO₂, 70 people crossing the Atlantic and back)
 - Responsibility to keep this as low as possible
- 2023: Installation of ~£1m solar panels
 - Funded by DiRAC
 - Investigation into the interplay between supply and demand






Cooling sub-lab

- A sequence of technologies:
 - Hot aisle (2006?)
 - Passive cooled rear doors (pre 2010)
 - Active cooled rear doors (2018)
 - Direct liquid cooling (2020)
 - Immersion cooling (2024-5)
 - As a national object-of-study
 - Support for visits to Durham for operators to learn this technology
 - Reduced operational and embodied CO2
 - EPSRC DRI funding
- Free air coolers: 12kW to cool 800kW load
 - Dependent on external temperature



Mine water heat storage

- 
- HPC produces a lot of heat
 - We can heat buildings with this in the winter
 - What can we in the summer?
 - Store it underground for extraction in the winter
 - Data centre is sitting on old, flooded mine workings
 - Ongoing project to investigate feasibility
 - In particular, how fast does the water flow?
 - EPSRC funded



Active power capping

- When the local grid Carbon Intensity rises:
 - Power cap nodes to a lower value
 - Many codes unaffected



Energy awareness


- Quarterly reporting to users around energy usage
 - Ability to query energy usage for each submitted job
- Monitoring of power on a node, rack, room and system scale
- Power-down of unused nodes
 - On quieter systems

Embodied carbon

- We keep our embodied carbon low
 - The AI boom has helped with this
- We need to keep systems in production for ~10 years to balance embodied vs operational CO2

Quantum computing

- No quantum compute in Durham (yet)
- However, the hardware lab administered national access to:
 - DWAVE quantum annealer
 - QuEra neutral atom quantum computer
- ExCALIBUR funded



Considerations / Lessons learned

- Verifying access
- Not all successes
- Length-of-interest varies
- Single node systems introduce access complications
- Production rather than retirement
 - Driver issues for older kit
- Ongoing funding is challenging
- Researcher buy-in is important
 - Need a reason, need to feel invested
 - SHAREing project helping to facilitate this

Accessing the hardware lab

- Sign up on SAFE:
 - safe.epcc.ed.ac.uk/dirac
- Apply to join an appropriate project code
 - cosma.readthedocs.io
- Visits arranged upon request

Hardware Lab Outputs

- Key outputs from the hardware lab are:
 - Up to date knowledge of performance on new technologies
 - Experience profiling and optimising codes
 - Code preparation for future systems
 - Training on new technologies and tools
 - User awareness
 - Input into future system design

Conclusion

- The Durham HPC Hardware Laboratory
 - Accessible for UK researchers
 - Single login
 - HPC environment
 - Cutting edge technologies
 - Let us know if there is something of particular interest
 - Always happy to host!
 - A steady funding line would be a good thing!