

# Establishing the HPC Hardware Lab @Durham

Durham University

Alastair Basden, Peter Draper, Mark Lovell, Fawada Qaisir, Richard Regan, Paul Walker

Alastair Basden  
DiRAC / Durham University



Durham  
University

Institute for Computational  
Cosmology

**DiRAC**  
High Performance  
Computing Facility

ExCALIBUR  
10



- HPC Hardware Lab hosted by:
  - the COSmology MAchine
    - Est. 2001
    - The DiRAC Memory Intensive service
      - primarily covering STFC science
      - Alongside the 3 other DiRAC systems
  - 3 generations currently in operation
  - Newest being COSMA8
    - ~70k cores, 0.5PB RAM, 20PB storage, 26PB tape
    - Bespoke design appropriate for main workloads
    - Tier-1 national service



# 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

# HPC Hardware Lab @Durham

- Mission: Provide access to the latest HPC hardware to users 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
  - Followed by multiple H&ES awards, and from other sources

# Establishing the Lab

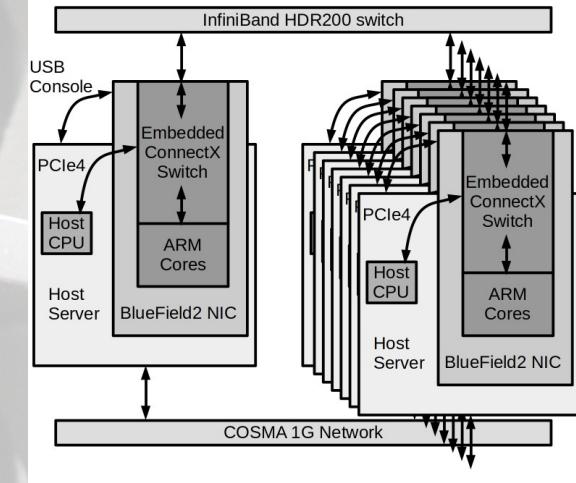
- Took a while to realise we had a Lab!
  - Hardware running alongside COSMA
  - Piggybacking on account creation, home space etc
    - Easy to make accessible to users
  - Eventually realised we had something that people wanted!
    - Demand continues to grow

# Hardware Lab Components

- Two test clusters:
  - 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
- Investigating feasibility of charging by kWhr from April
- Funded by Durham, DiRAC, ExCALIBUR
- A small cluster to accept hardware is useful!



# DINE-2

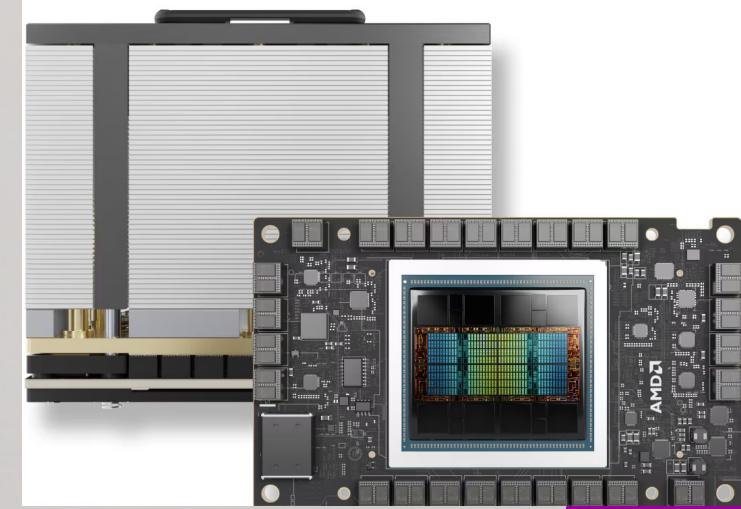
- Durham Integrated Next-gen Environment
  - 8 node Intel Sapphire Rapids system
    - 64 cores, 2TB RAM per node
  - Currently hosts a CerIO composable PCIe-5 fabric
    - 8x NVIDIA GPUs, assignable in any number to any host
  - 400G Ethernet fabric (800G switch)
- Also part of the kWhr charging/allocation study
- Funded by DiRAC, IRIS and SKA
- Lack of PCIe slots

# CPU compute

- Providing users with access to cutting edge CPU technologies:
- Coming soon: AMD Turin
- AMD Genoa and Bergamo
- NVIDIA Grace
- Intel Emerald Rapids, Sapphire Rapids, Ice Lake
- AMD Milan-X (extreme cache version: 768MB L3 cache)
- AMD Milan, Rome
- Intel Cascade Lake (with Apache Pass RAM, 6TB)
- Funded by OEMs, Intel, AMD, Dell, DiRAC, ExCALIBUR

# GPU Compute

- Access to small numbers of latest (and not-so latest) GPUs
- AMD MI300X
  - 8x MI300X, cache coherent global address space
- AMD MI300A
  - 4x MI300A, APU, global RAM
- AMD MI210, MI100, MI50
  - MI50 were not well used - ensure there is demand!
- 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
  - Difficulties arising from single node systems
- Funded by Dell, AMD, Intel, DiRAC, ExCALIBUR, IRIS
- Coming soon: NVIDIA RTX6000



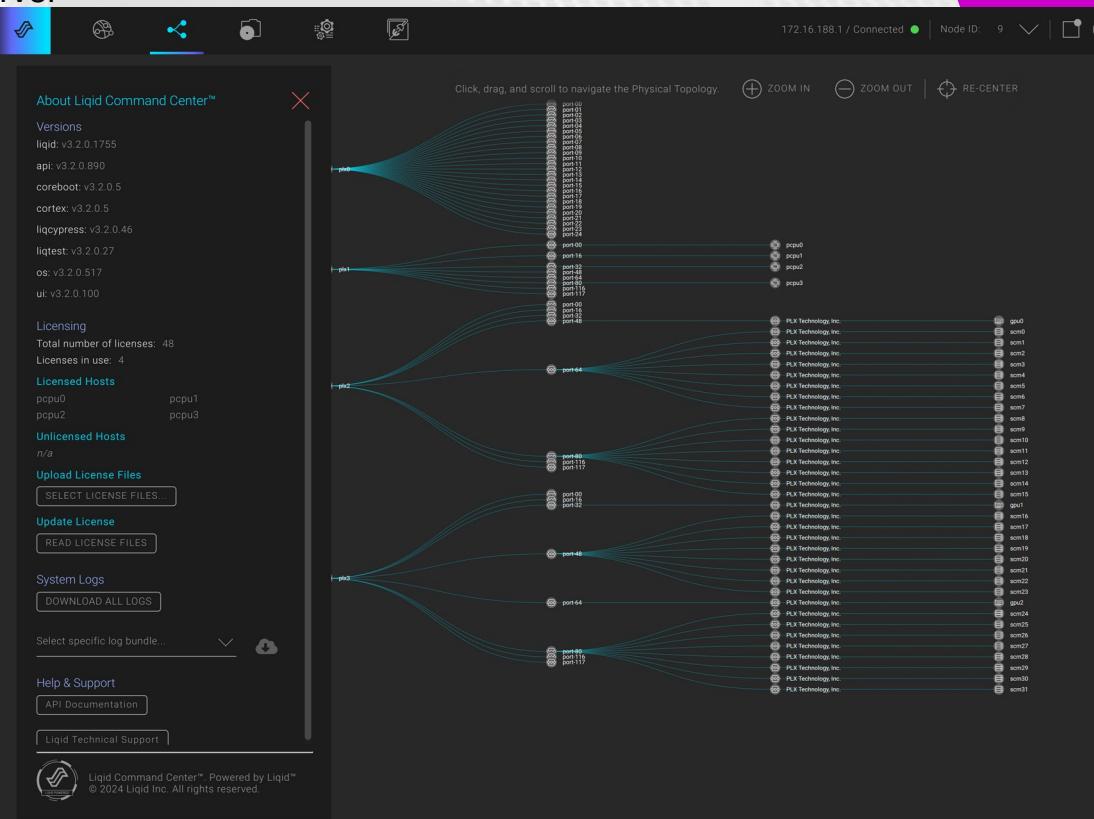
Credit: AMD

# 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
    - Key learning point: Use higher TRLs
      - We are not suited for hardware development or early kit

# 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



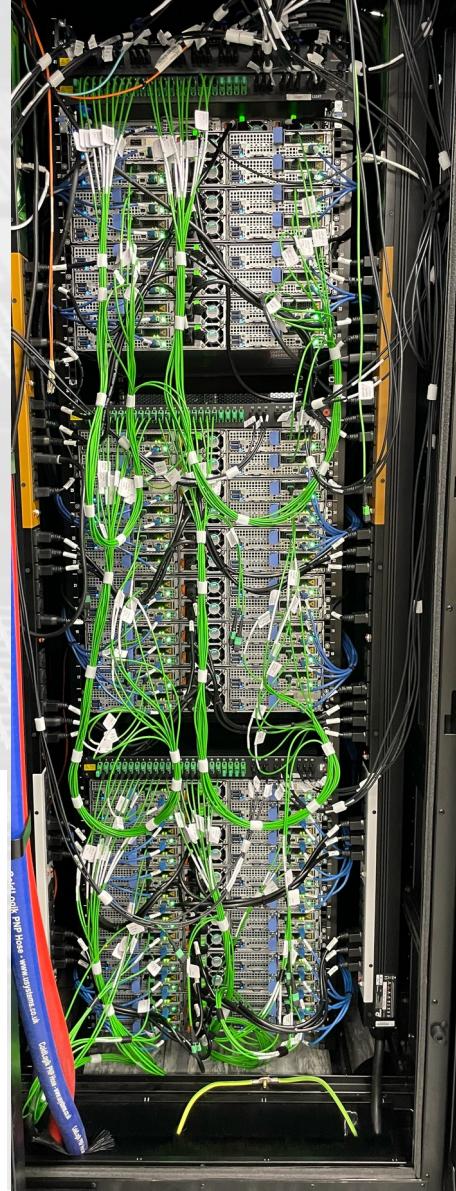
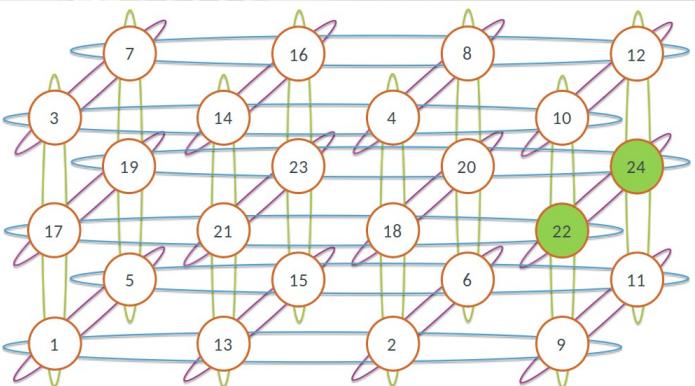
# CerIO composable system

- Installed 2024
- PCIe5-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
- **Initially didn't work! - not always possible to rush into production**
  - Now stable (firmware updates, co-design)
  - Expansion pending
  - Some downsides (reboots, etc)
- IRIS/SKA/ExCALIBUR/DiRAC



# Network fabrics

- Rockport 6D Torus network
  - A “switchless” fabric for 100G Ethernet
  - Trailed 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
- 800G Ethernet fabric
  - 400G to nodes
  - On the DINE2 cluster



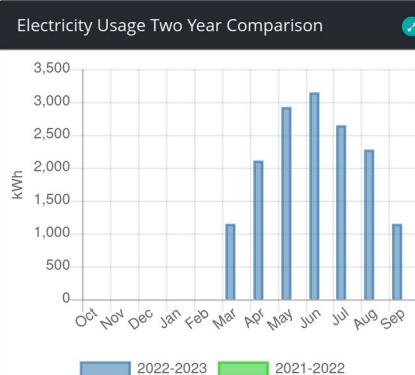
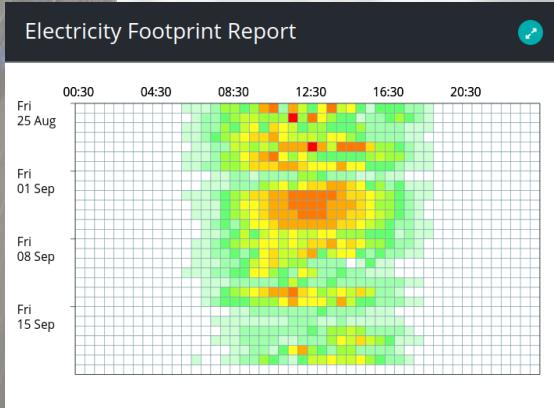
# Storage sub-lab

- Various different storage technologies
- High-performance scratch Lustre (NVMe)
- DAOS (NVMe)
- Ceph
- StorJ private cloud
- VAST (NVMe)
- Globus (data transfer)
- Lustre (efficient bulk storage)
- Tape: Atempo Miria
- NVMe RAID test systems
- Funded by DiRAC/IRIS/SKA/ExCALIBUR



# Environmental-related

- HPC is a huge energy user
  - COSMA ~1MW at peak
    - (~5GWh/year, ~120T CO<sub>2</sub>, 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
    - And how to engage with Estates



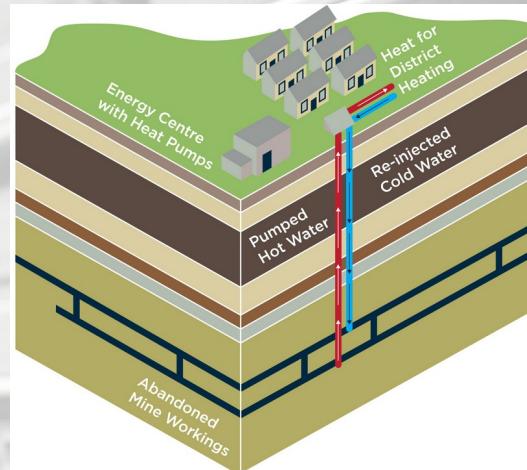
# Cooling sub-lab

- A sequence of technologies:
  - Hot aisle (2006?)
  - Passive cooled rear doors (pre 2010)
  - Active cooled rear doors (2018)
  - Direct liquid cooling (2020)
  - Immersion cooling (2025)
    - As a national object-of-study
    - Support for visits to Durham for operators to learn this technology
    - Reduced operational and embodied CO<sub>2</sub>
- Free air coolers: 12kW to cool 800kW load
  - Dependent on external temperature

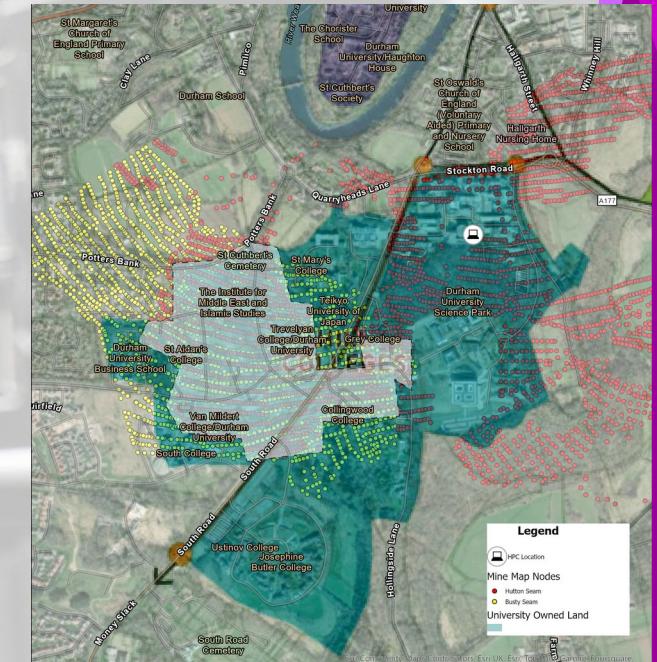


# Mine water heat storage

- The ICHS project
  - Storing waste heat in flooded mines
    - For reuse during the winter

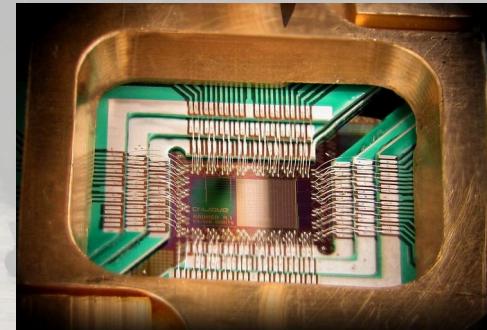


Alastair Basden  
DiRAC / Durham University



# Quantum computing

- No quantum compute in Durham (yet)
- However, the hardware lab administers national access to:
  - DWAVE quantum annealer (expired)
  - QuEra neutral atom quantum computer
- Quantum for HPC acceleration
  - This is becoming feasible
- ExCALIBUR funded



# Operational costs

- Generally swallowed up in other things
  - Electricity
  - Installation
  - Maintenance
  - Updates, improvements, etc
  - There is currently no ideal UKRI method

# 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

# Future plans\*

- Bespoke AI accelerators
- Lower spec GPUs
- Turin CPU system
- CXL composable systems
  - Large shared-memory fabrics
- Ultra-Ethernet fabric

\* dependent on funding!

# Accessing the hardware lab

- Sign up on SAFE:
  - [safe.epcc.ed.ac.uk/dirac](http://safe.epcc.ed.ac.uk/dirac)
- Apply to join an appropriate project code: (durham.readthedocs.io)
  - do009: General purpose
  - do015: Cerio compstable system
  - do016: NVIDIA GPUs
  - do017: Intel GPUs
  - do018: AMD GPUs
- We are not too prescriptive to users: no agreements, SLAs, etc, low barriers to use
  - We like feedback, but don't push too heavily for it

# Conclusion

- The Durham HPC Hardware Laboratory
  - Accessible for UK researchers
    - Account creation on SAFE
    - Single login
    - HPC environment
  - Cutting edge technologies
  - Let us know if there is something of particular interest
    - Always happy to host!