Net Zero and DiRAC

Durham HPC Days

May 2024

Alastair Basden

DiRAC / Durham University

DiRAC

- UK tier-1 HPC facility at 4 sites:
  - Cambridge, Durham, Edinburgh, Leicester Universities
  - Head office in University College London
  - Each site with a bespoke design tailored to specific workloads
- Funded by the Science and Technology Facilities Council
  - Serving researchers within the remit of particle physics, astronomy, cosmology, solar, planetary and nuclear

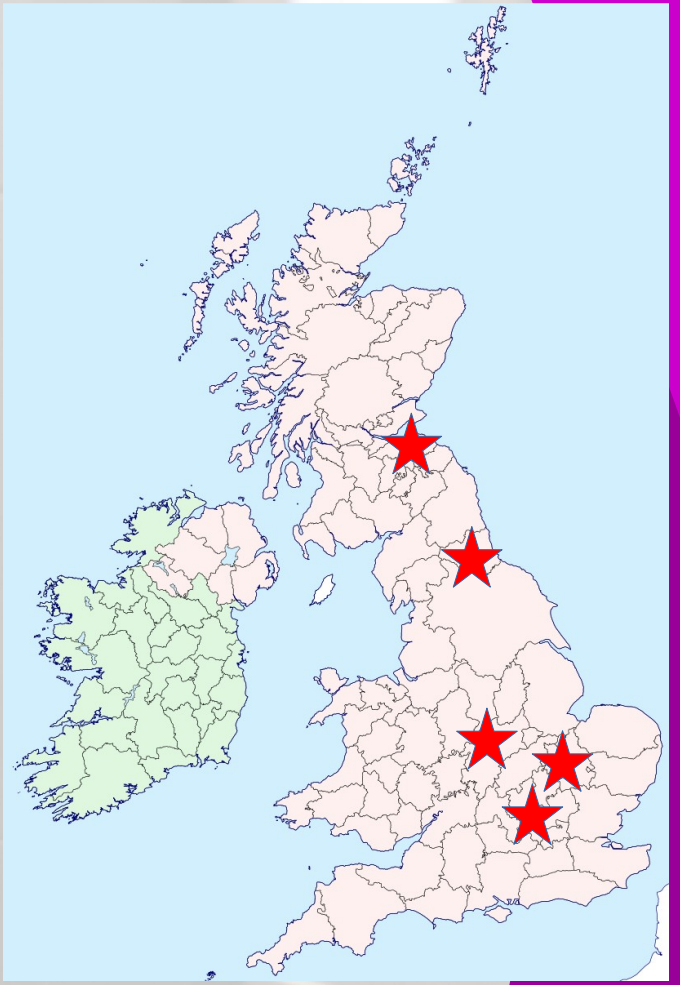

Alastair Basden
DiRAC / Durham University

Net-Zero

- HPC has a problem
  - Significant and growing electricity consumption
  - Associated CO2 emissions
  - Significant embodied (manufacture) CO2
- What are we doing?
- What more can we do?

Alastair Basden
DiRAC / Durham University

Embodied CO2

- Accurate numbers are hard
  - Many manufacturers provide data
  - A lot of guess work
    - Changes with time
- Typically 1200-1800kg CO2 per CPU server

Alastair Basden
DiRAC / Durham University

Science production CO2

- Directly related to power consumption
  - Depends on Carbon Intensity
    - NE currently ~13 gCO2/kWhr (15g year average)
    - GB currently ~141 gCO2/kWhr (155g year average)
  - Typical COSMA8 server ~600W: 5000 kWhr/year
    - 800kg CO2/year (using national value over past year)
    - 1.5-2 years operation for embodied==production
    - 15-20 years using the NE value!
  - As grid greens, need to run for longer
    - Good practice anyway (scarce resources)

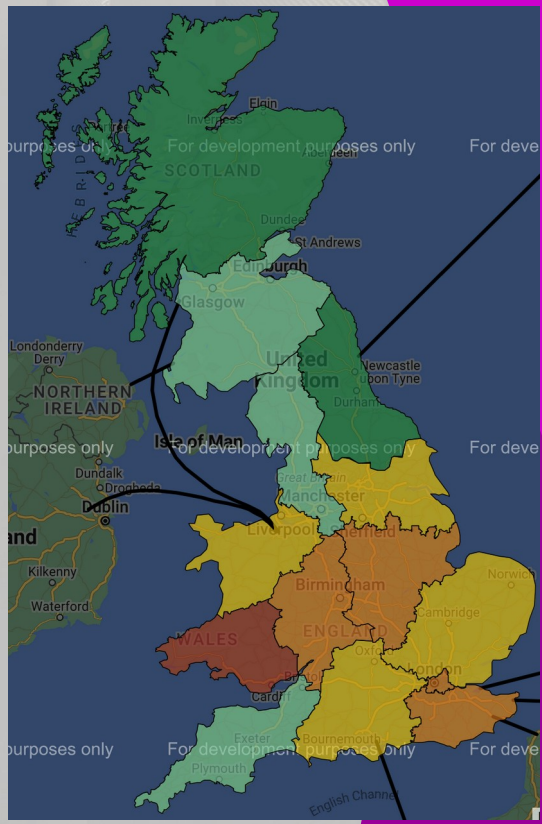

Alastair Basden
DiRAC / Durham University

Transport and decommissioning

- Ignorable to first order

Alastair Basden
DiRAC / Durham University

Awareness

- User awareness can help reduce CO2
  - Brute force vs intelligent parameter searches
  - Selection and efficiency of code
  - Running codes correctly on most appropriate resource
  - Access to job energy statistics
- Within DiRAC: Quarterly emails to users and PIs
  - kWhr, estimated CO2, equivalence
- Admin awareness: Monitoring, consolidation, BIOS, power-off

Resource utilisation

- Key is to keep this high
  - Idle systems use 25-50% electricity
- DiRAC reviews usage on a quarterly basis
  - Emails to underusing PIs
  - Reallocation of resources of projects likely to under use their allocation
    - Based on a study of past usage statistics
- Idle power-off of some systems
  - Lifetime considerations

System design

- Bespoke systems to maximise science output
  - Reduced resource wastage - same science with smaller systems
  - More science per £, or per kWhr
- Consideration of embodied carbon
  - e.g. SSDs are 8x worse than HDDs
    - Swamit Tannu and Prashant J. Nair. 2023. The Dirty Secret of SSDs: Embodied Carbon. SIGENERGY Energy Inform. Rev. 3, 3 (October 2023), 4–9. https://doi.org/10.1145/3630614.3630616

Storage	Energy (10y kwh)	OPEX CO2 UK (10y kg)	OPEX CO2 NE (10y kg)	CAPEX CO2 (kg)	Total CO2 UK (10y kg)	Total CO2 NE (10y kg)
HDD (1TB)	367.9	57.0	5.5	20x2 (5y life)	97.0	45.5
SSD (1TB)	113.8	17.6	1.7	160 (10y life)	177.6 (1.8x)	161.7 (3.6x)

Energy efficient compute

- Bespoke system design
  - Fewer wasted resources
    - Systems matched to the science
- GPU clock frequency settings
- CPU BIOS settings study
- Green500 systems

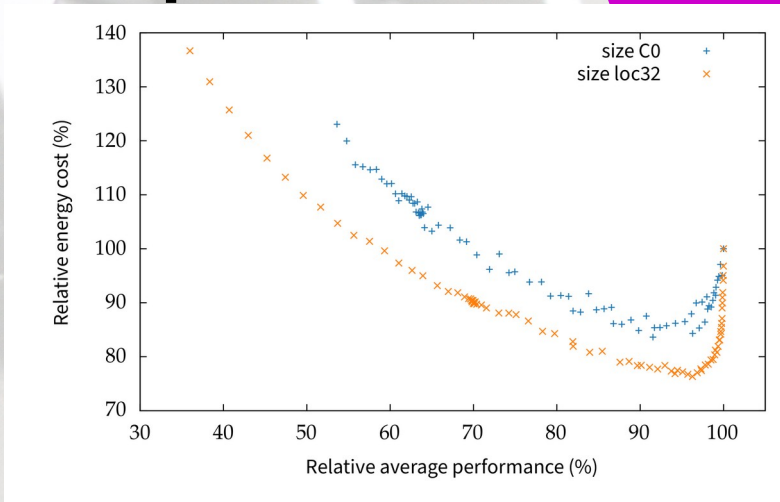

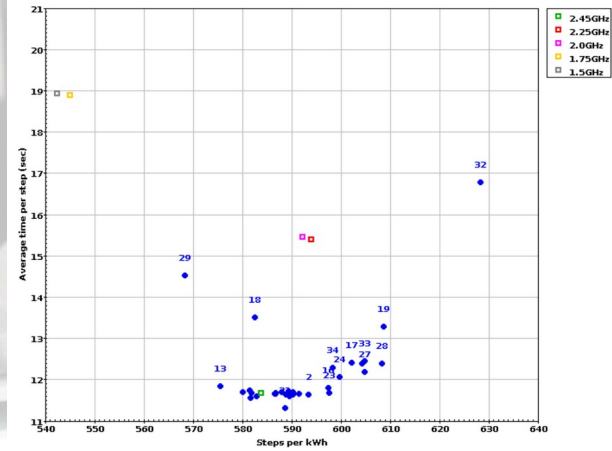

Alastair Basden
DiRAC / Durham University

Energy efficient cooling

- Early Direct Liquid Cooling adoption
- Investigation into Immersion cooling
  - Tank arriving in Durham over the next few months
  - Available to the HPC (admin) community to explore
  - Will be put into production
  - Higher temperature waste heat: reuse/better cooling
  - Lower embodied CO2 (simpler server design)
  - Lower operational power consumption
- Mandated data centre PUE <1.2
  - Free-air coolers

Credit: hypertec.com

Alastair Basden
DiRAC / Durham University

Waste heat reuse

- In the planning stages at some sites
  - Primarily for building heating
  - Expensive and time consuming
    - Why did we not start this 10 years ago?!?
  - Seek to maximise output temperature
- What about the other 6 months?
  - Underground heat storage

Alastair Basden
DiRAC / Durham University

Underground heat storage

- Galleries to Calories scheme in Edinburgh
  - Transfer heat to residential areas
- Mine water heat storage project in Durham
  - Store heat produced in the summer for use in the winter
  - Using old coal mines beneath the site

Alastair Basden
DiRAC / Durham University

Solar panels

- ~£1m deployment at Durham from DiRAC Federation money
  - To demonstrate feasibility of coupling compute with net-zero
  - Lessons learned: 6 month timescales are very challenging - longer-term funding required
  - Requires a lot of good will from Estates
  - Good to have a pre-prepared plan
  - If 5-10% of HPC budget was spent on solar, would eventually be self-sufficient
    - (on average)

Collaboration

- With other UKRI Digital Research Infrastructure projects
  - CEDA HPC-JEEP project
    - What energy data do HPC systems have (for users)
    - Would energy-based charging work?
    - Compute ~80% HPC power usage
  - CEDA IRISCAST project
    - Multi-site snapshot of CO2 production

AI usage study

- Adaptive update of quarterly allocations

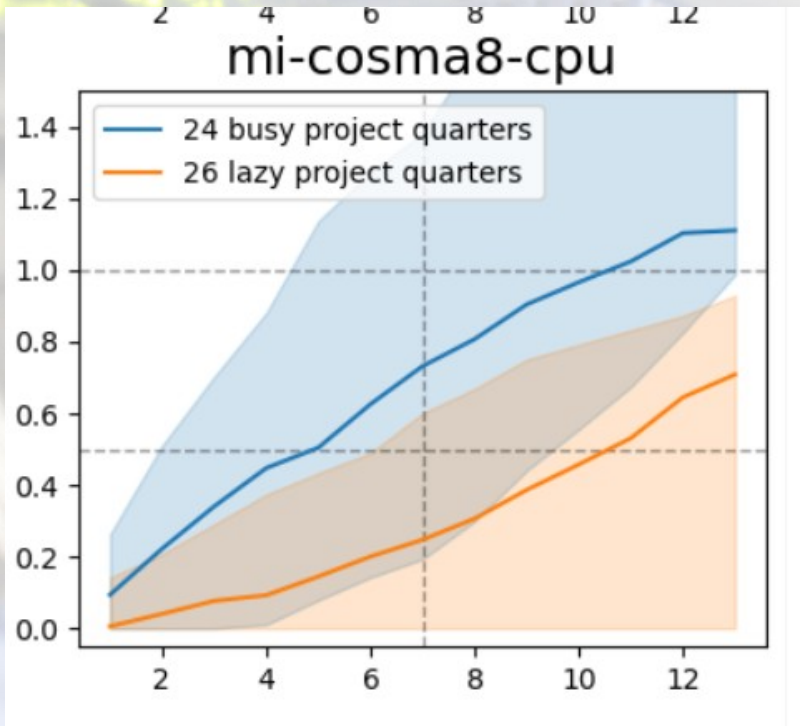

Alastair Basden
DiRAC / Durham University

AI energy consumption study

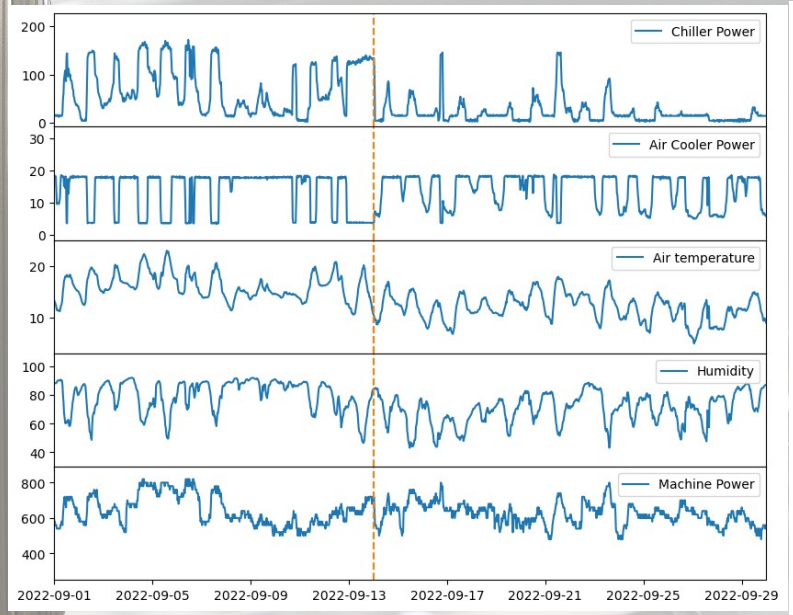

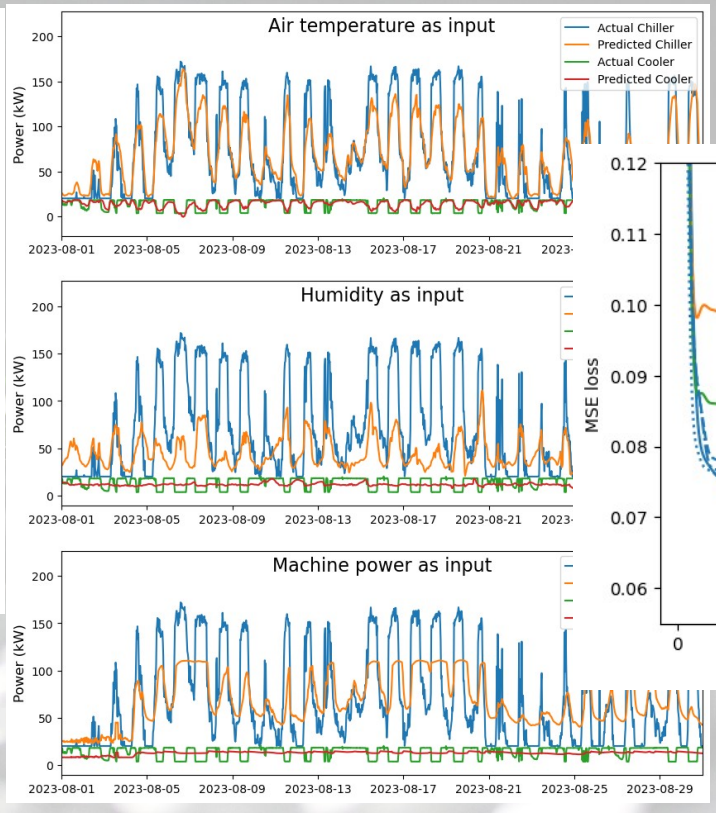

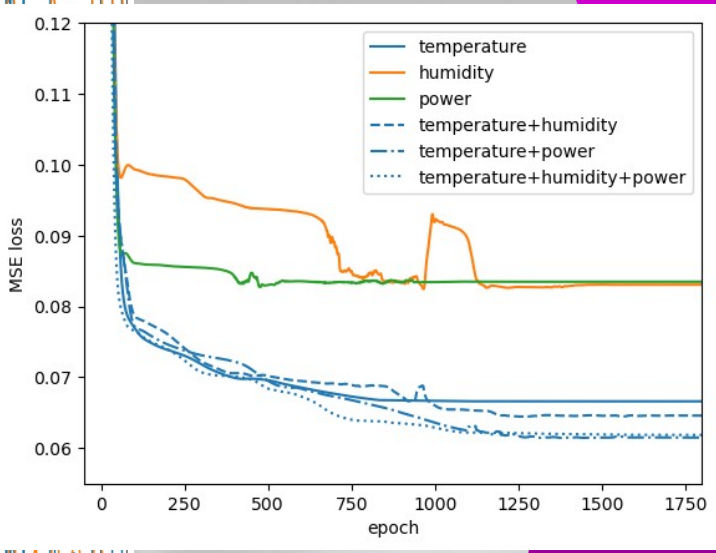

Alastair Basden
DiRAC / Durham University

Free-air cooler / Air temp correlation

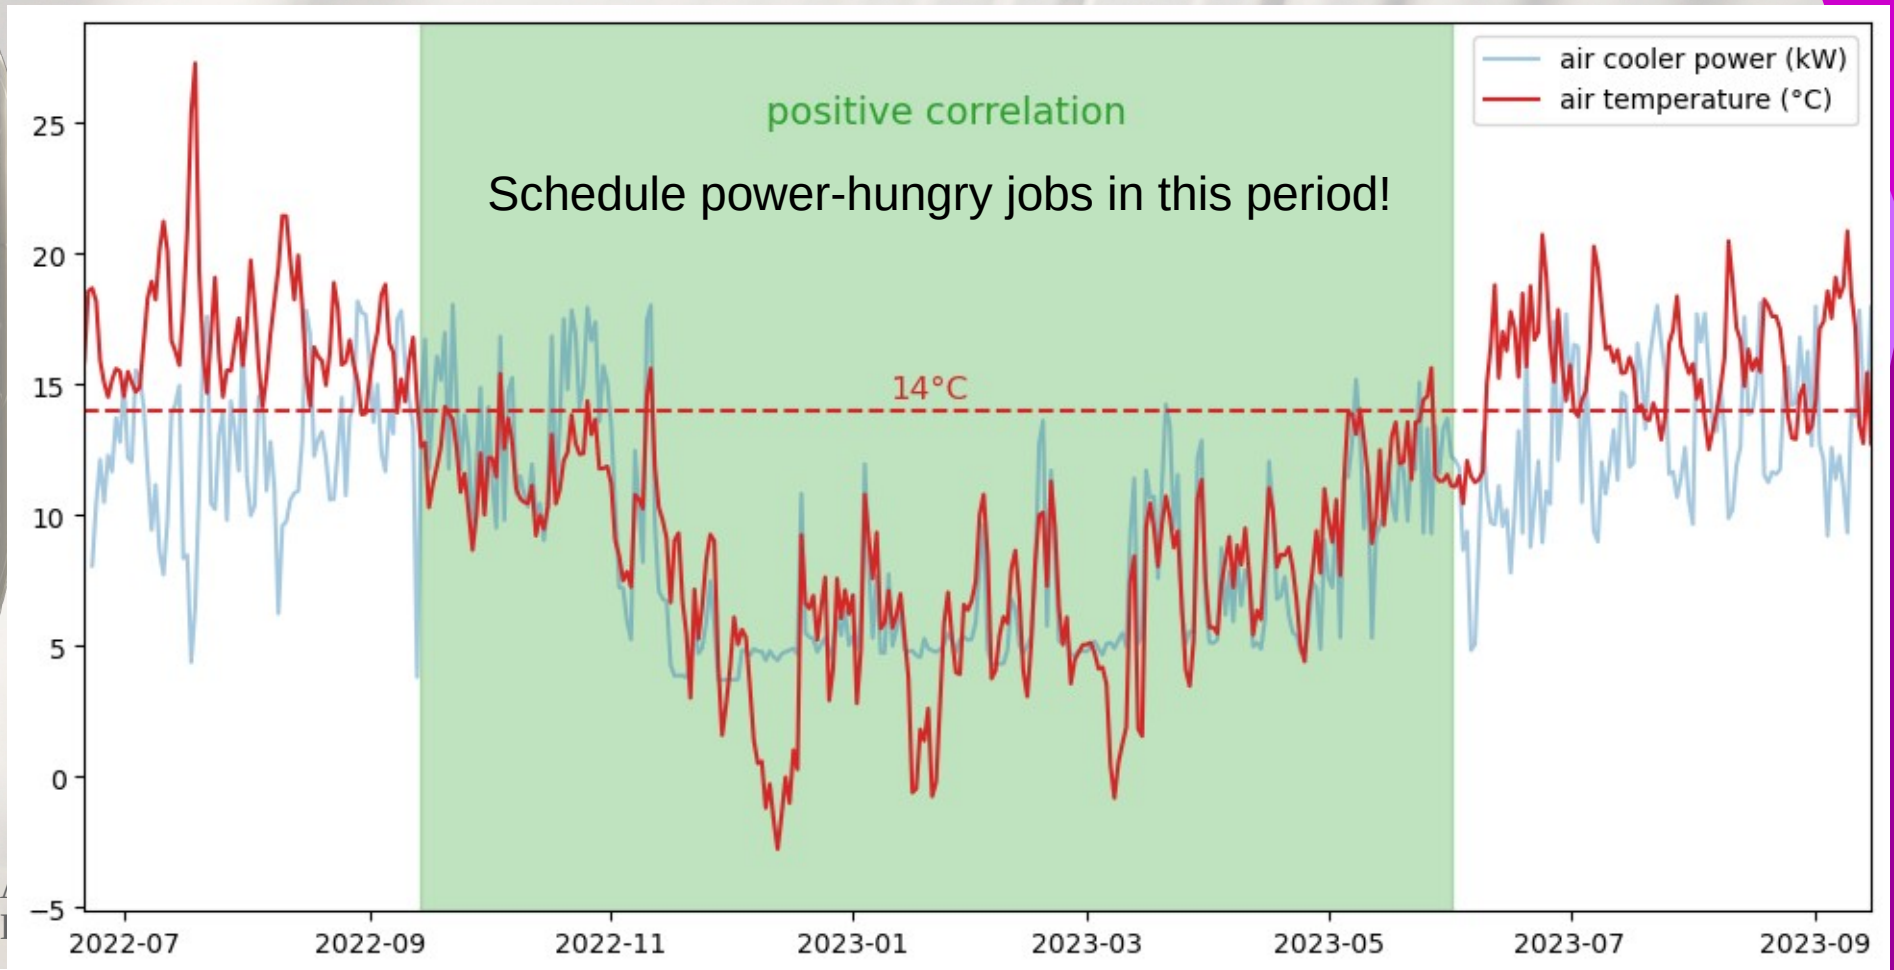

Conclusions

- DiRAC are making small steps
- There is much more to be done
- Key things will be
  - User awareness
  - Code efficiencies
  - Longevity of services