Working towards Net-Zero in DiRAC

Alastair Basden

DiRAC/Durham

Durham University
Institute for Computational Cosmology

DiRAC
High Performance Computing Facility

ExCALIBUR

Importance of Net-Zero

- 2024: 1.54 degree rise
- Current projection is 3.1 degrees by 2100
  - Moderately catastrophic

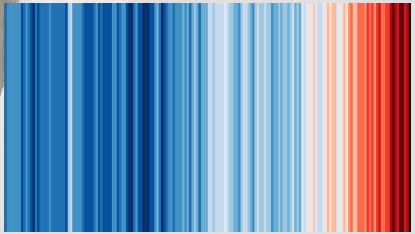

showyourstripes.info

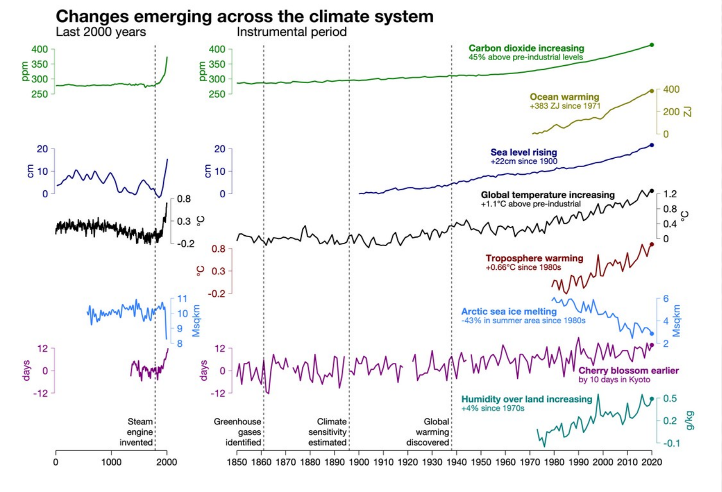

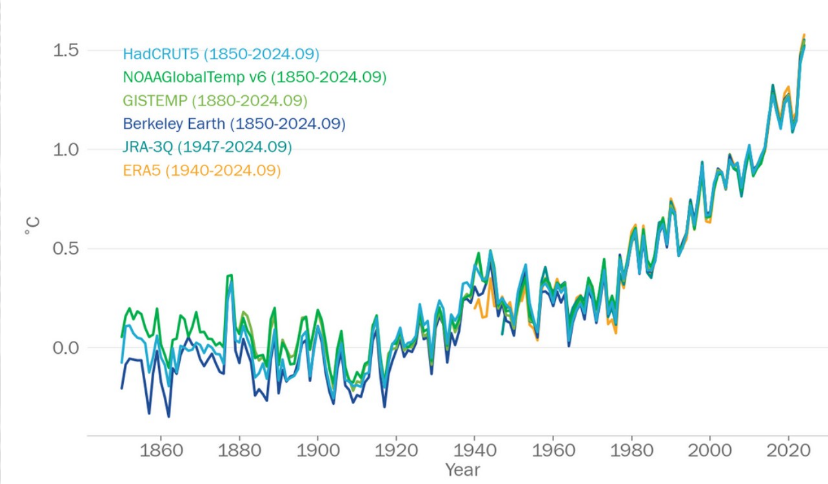

Alastair Basden
DiRAC / Durham University

www.climate-lab-book.ac.uk

climate-lab-book.ac.uk

Net Zero: What is our goal?

- Max science per CO_2?
- Zero CO_2?
- Somewhere in between?

Alastair Basden
DiRAC / Durham University

The HPC context

- HPC is a large energy user:
  - COSMA8 - ~300kW
    - 70T CO2 / year (assuming regional carbon intensity)
    - 1 flight getting nearly half way across the Atlantic
      - 50 people going to SC24
  - TURSA - ~400kW
  - DIaL3 - ~170kW
  - CSD3 - ~
- The electricity grid is getting greener
  - Embodied and flying has a massive impact!

Alastair Basden
DiRAC / Durham University

Regional variation

- carbonintensity.org.uk
- What does this mean?
- How interconnected is the grid?
  - Should we use regional or national values?
  - Probably somewhere in between
    - Regions are connection but transmission losses and bandwidth
    - Carbon accounting will likely be based on average UK value

East England: 107 g CO_2/kWhr
East Midlands: 209
England: 133
GB: 125
London: 123
North East England: 23
North Scotland: 31
North Wales & Merseyside: 75
North West England: 47.
Scotland: 38
South East England: 133
South England: 190
South Scotland: 25
South Wales: 253
South West England: 244
Wales: 180
West Midlands: 126
Yorkshire: 143
g CO_2/kWhr (year Dec 2023-Nov 2024)

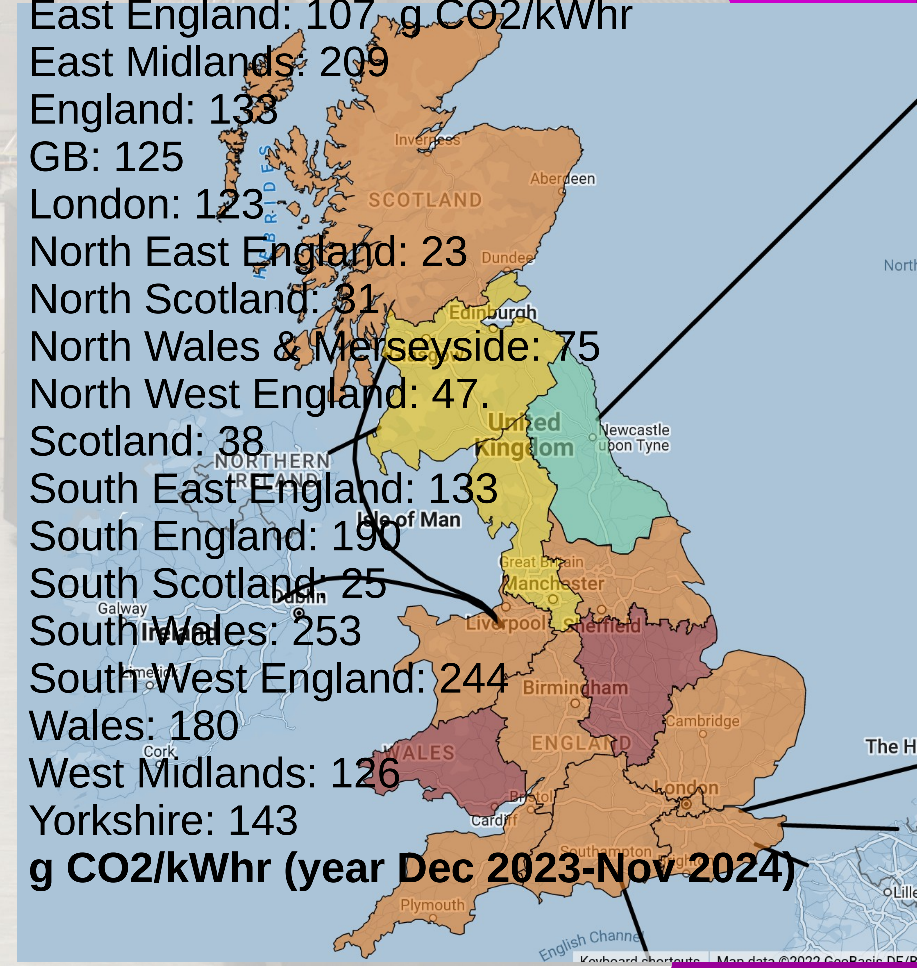

Alastair Basden
DiRAC / Durham University

Embodied CO2

- A key consideration:
  - 850kg CO2 for a COSMA8 compute node
    - 450T for COSMA8 (compute only)
    - As given by Dell: Manufacture and logistics
    - With electricity @23g CO2/kWhr for North East:
      - 7 years of operation to equal emboddied CO2 (Cost ~£10k, electricity cost ~£7k)
      - Closer to 10 years considering other system components
  - Some servers closer to 2T:
    - Close to 20 years to balance CO2 sources
- We should run our systems for longer
  - Stewardship of other resources
- And encourage manufacturers to reduce CO2
  - Include requirements within procurements
    - At least for reporting
    - Of embodied, **not** whole life cycle (13x CO2 intensity assumed)

DiRAC steps

- DiRAC has considered Net-Zero from instantiation
  - 2012 - PUE mandated to be below 1.2
    - Standard was 1.6 at the time
    - Evaporative or dry-air cooling
  - DLC on all DiRAC services since 2020
  - Recent audit to help develop a formal DiRAC Net Zero strategy
  - Continue to evolve to a more sustainable facility…
  - Repurposing old systems: COSMA5, COSMA6

Alastair Basden
DiRAC / Durham University

System energy data

- All DiRAC sites capture energy information
  - Whole system
  - Per-node energy
  - Storage energy
  - Cooling and UPS energy
  - Some capture network fabric energy

User energy reporting

- Energy usage per job:
  - Captured on some systems
  - User searchable
- Quarterly energy usage email on some systems

Alastair Basden
DiRAC / Durham University

Investigations into down-clocking

- Small reduction in performance (10%) for larger reduction in energy usage (30%)
- And BIOS settings:
  - Multi-dimensional parameter space
  - Code dependent
  - Sweet spot improving performance while reducing power

Caveat: embodied CO_2

- Accepting 10% longer run times could mean:
  - Less science (is this acceptable? Are we too generous with RAC allocations?)
  - Science extended over a longer period
    - i.e. into the purchase of a new system
    - In which case, the gain is minimal due to high embodied CO_2
  - Better efficiency or ways of working
    - i.e. the time wasn’t actually required
  - Could we reduce the system size by 10% at procurement and run at full speed?
    - This would yield a REAL CO_2 saving*

*Depending on assumptions related to CO_2

Alastair Basden
DiRAC / Durham University

Bespoke design

- DiRAC bespoke design principles help reduce CO_2
  - Less physical infrastructure required
    - Lower embodied CO_2
    - (driven by capital budget, rather than CO_2 concerns)
  - Higher science per kWhr
  - A single large generic system delivering the same capabilities would have significantly increased CO_2
  - Focus on capability workloads

Solar panel installation

- ~£1m solar panels installed in Durham in 2023
  - Funded by DiRAC Federation
  - 85% of energy dedicated to COSMA

Alastair Basden
DiRAC / Durham University

Heat storage

- We do not currently use waste heat within DiRAC
  - There are some low hanging fruits/demonstrators
  - But requires Estates involvement
  - Heat is low grade
    - Investigations into raising the temperature
      - e.g Immersion cooling

Future heat reuse and storage

- Investigation into waste heat reuse
  - EPCC: Calories to Galleries project
    - Transfer heat to inhabited areas using natural water flows
    - Concept stage
  - Durham: ICHS project
    - Using flooded mine workings to store heat in summer
      - Allowing reuse in winter (driving heat pumps)
    - Drilling of test shafts may start this summer
    - Prototype feasibility study stage

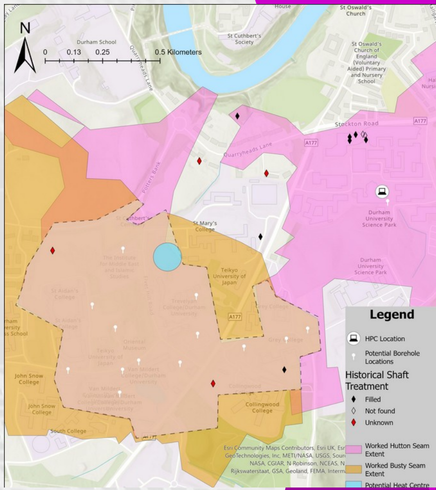

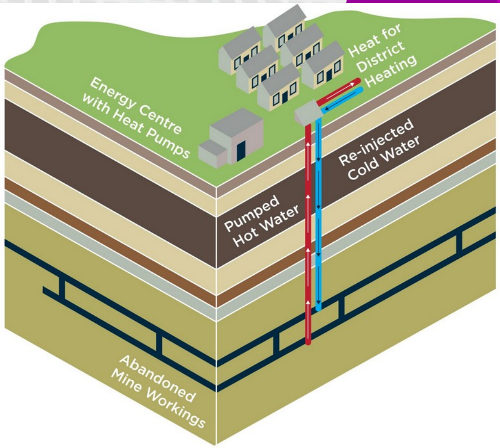

Alastair Basden
DiRAC / Durham University

Immersion cooling prototype

- Funded to install an immersion tank and servers
  - Made available to UK DTP professionals to visit and explore
  - Funding will allow replacement of old COSMA5 system
    - Compressing ~300 old nodes into 8 new nodes.

Alastair Basden
DiRAC / Durham University

Node power-off

- Less well used Slurm partitions allow node power-off at Durham
  - Multi-year data looking at node failure rates
  - COSMA5, DINE
  - Others to be included eventually
- Note: No need to do this for a well used system

Node power reduction

- Should we look at reducing node power when carbon intensity is high?
  - Potentially - depends on embodied CO_2 calculation
    - Study to be performed

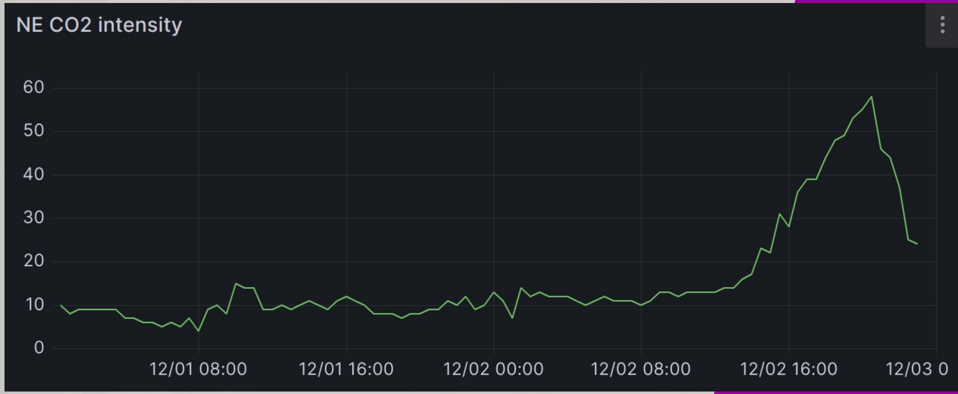

Alastair Basden
DiRAC / Durham University

RSE work

- Large improvements in code performance
  - Significantly increased science per CO2
    - Though this may not reduce CO2 itself!
- Examples:
  - GRID, Prompi, SWIFT

Alastair Basden
DiRAC / Durham University

DRI Net-Zero studies

- IRISCAST
  - Carbon cost audit
- HPC-JEEP
  - Per-job energy statistics
  - Typically, 80% of system energy is compute
    - Rest is a balance of fabric and storage
    - (Excluding PUE)

Future work

- Grid-scale storage technologies
  - Cut demand from grid at times of peak CO2
- Ongoing work to improve monitoring
- Investigation into charging by kWhr
  - Note: UKRI CO2 accounting
- Investigation into load reduction at peak CO2 intensity