Machine Room and Hardware Theme
Sustainability at the Advance Computing Facility (ACF)

Paul Clark, Adrian Jackson, Alan Simpson, Lorna Smith, Ritchie Somerville, Andy Turner

EPCC, University of Edinburgh

l.smith@epcc.ed.ac.uk

epcc

UKRI EPSRC and NERC

THE UNIVERSITY of EDINBURGH

Hewlett Packard Enterprise

Outline

- Machine Room and Hardware Theme
- Low carbon supply chain – Scope 3 emissions
  - Quantifying emissions from ARCHER2
- Use phase – Scope 2 emissions
  - Quantifying emissions from ARCHER2
    - Comparisons to other emissions sources
    - Scope 3 and Scope 2 balance
  - Reducing Scope 2 emissions
    - Infrastructure, hardware, software, user/application
- Biodiversity
  - Protecting and fostering biodiversity

Machine Room and Hardware

- *We need to bring sustainability to the forefront of the management of our infrastructure*
- *Full and authoritative life cycle analysis of infrastructure will be needed including effective measurement and management of impacts:*
  - *during the use phase*
  - *clear contracts and conditionalities to develop a low carbon supply chain*
  - *and clear end-of-life planning for hardware*

- Overview of key activities at EPCC's Advanced Computing Facility (ACF) and on ARCHER2
- Stimulate discussion and ideas for the following session

Emission sources

- *During the use phase*
  - Scope 2
  - Emissions associated with energy generation to run the system
  - Should also include overheads (e.g. plant energy consumption)
- *Develop a low carbon supply chain and clear end-of-life planning for hardware*
  - Scope 3
  - Emissions associated with manufacture, shipping, disposal etc. of system
  - Ideally, also include emissions associated with construction of data centres
- Scope 1 (not usually important for HPC systems)
  - Direct emissions from the system (e.g. if you had your own energy generation)
- For ARCHER2, we have looked at Scope 2 and Scope 3 emissions

Scope 3 Emissions

Scope 3 (embodied) emissions

Caveat: data sources are often uncertain, so this is a rough estimate

	Compute Nodes	Interconnect	Lustre HDD	Lustre SSD	NFS HDD	
Count	5,860 nodes	768 switches	19,759,200 GB	1,900,800 GB	3,240,000 GB	
$kgCO_2e$ per unit	1,100	200	0.02	0.16	0.02	
$kgCO_2e$ Total	6,400,000	150,000	400,000	300,000	70,000	7,320,000 Total
%	87%	2%	6%	4%	1%	100%*

- Compute nodes dominate by far
  - Used high end estimates for compute node emissions, could be as much as 30% lower
- Even large amounts of storage have a relatively small contribution
- Note that solid state storage has much larger embodied emissions per GB
- Other components (login nodes, cables, CDU, racks) likely contribute smaller amounts and we will try to include these in future work.

Scope 3 (embodied) emissions

- Finding accurate data for scope 3 emissions is challenging – how do we improve this?
- What do these figures mean?
- How do we predict and use these figures to influence future purchasing choices?
- How important are they compared to Scope 2 emissions?
- Understanding the relative importance of Scope 2 and Scope 3 emissions impacts how we deliver services and how long we run them

Dr Elisabetta Boella

Scope 2 Emissions

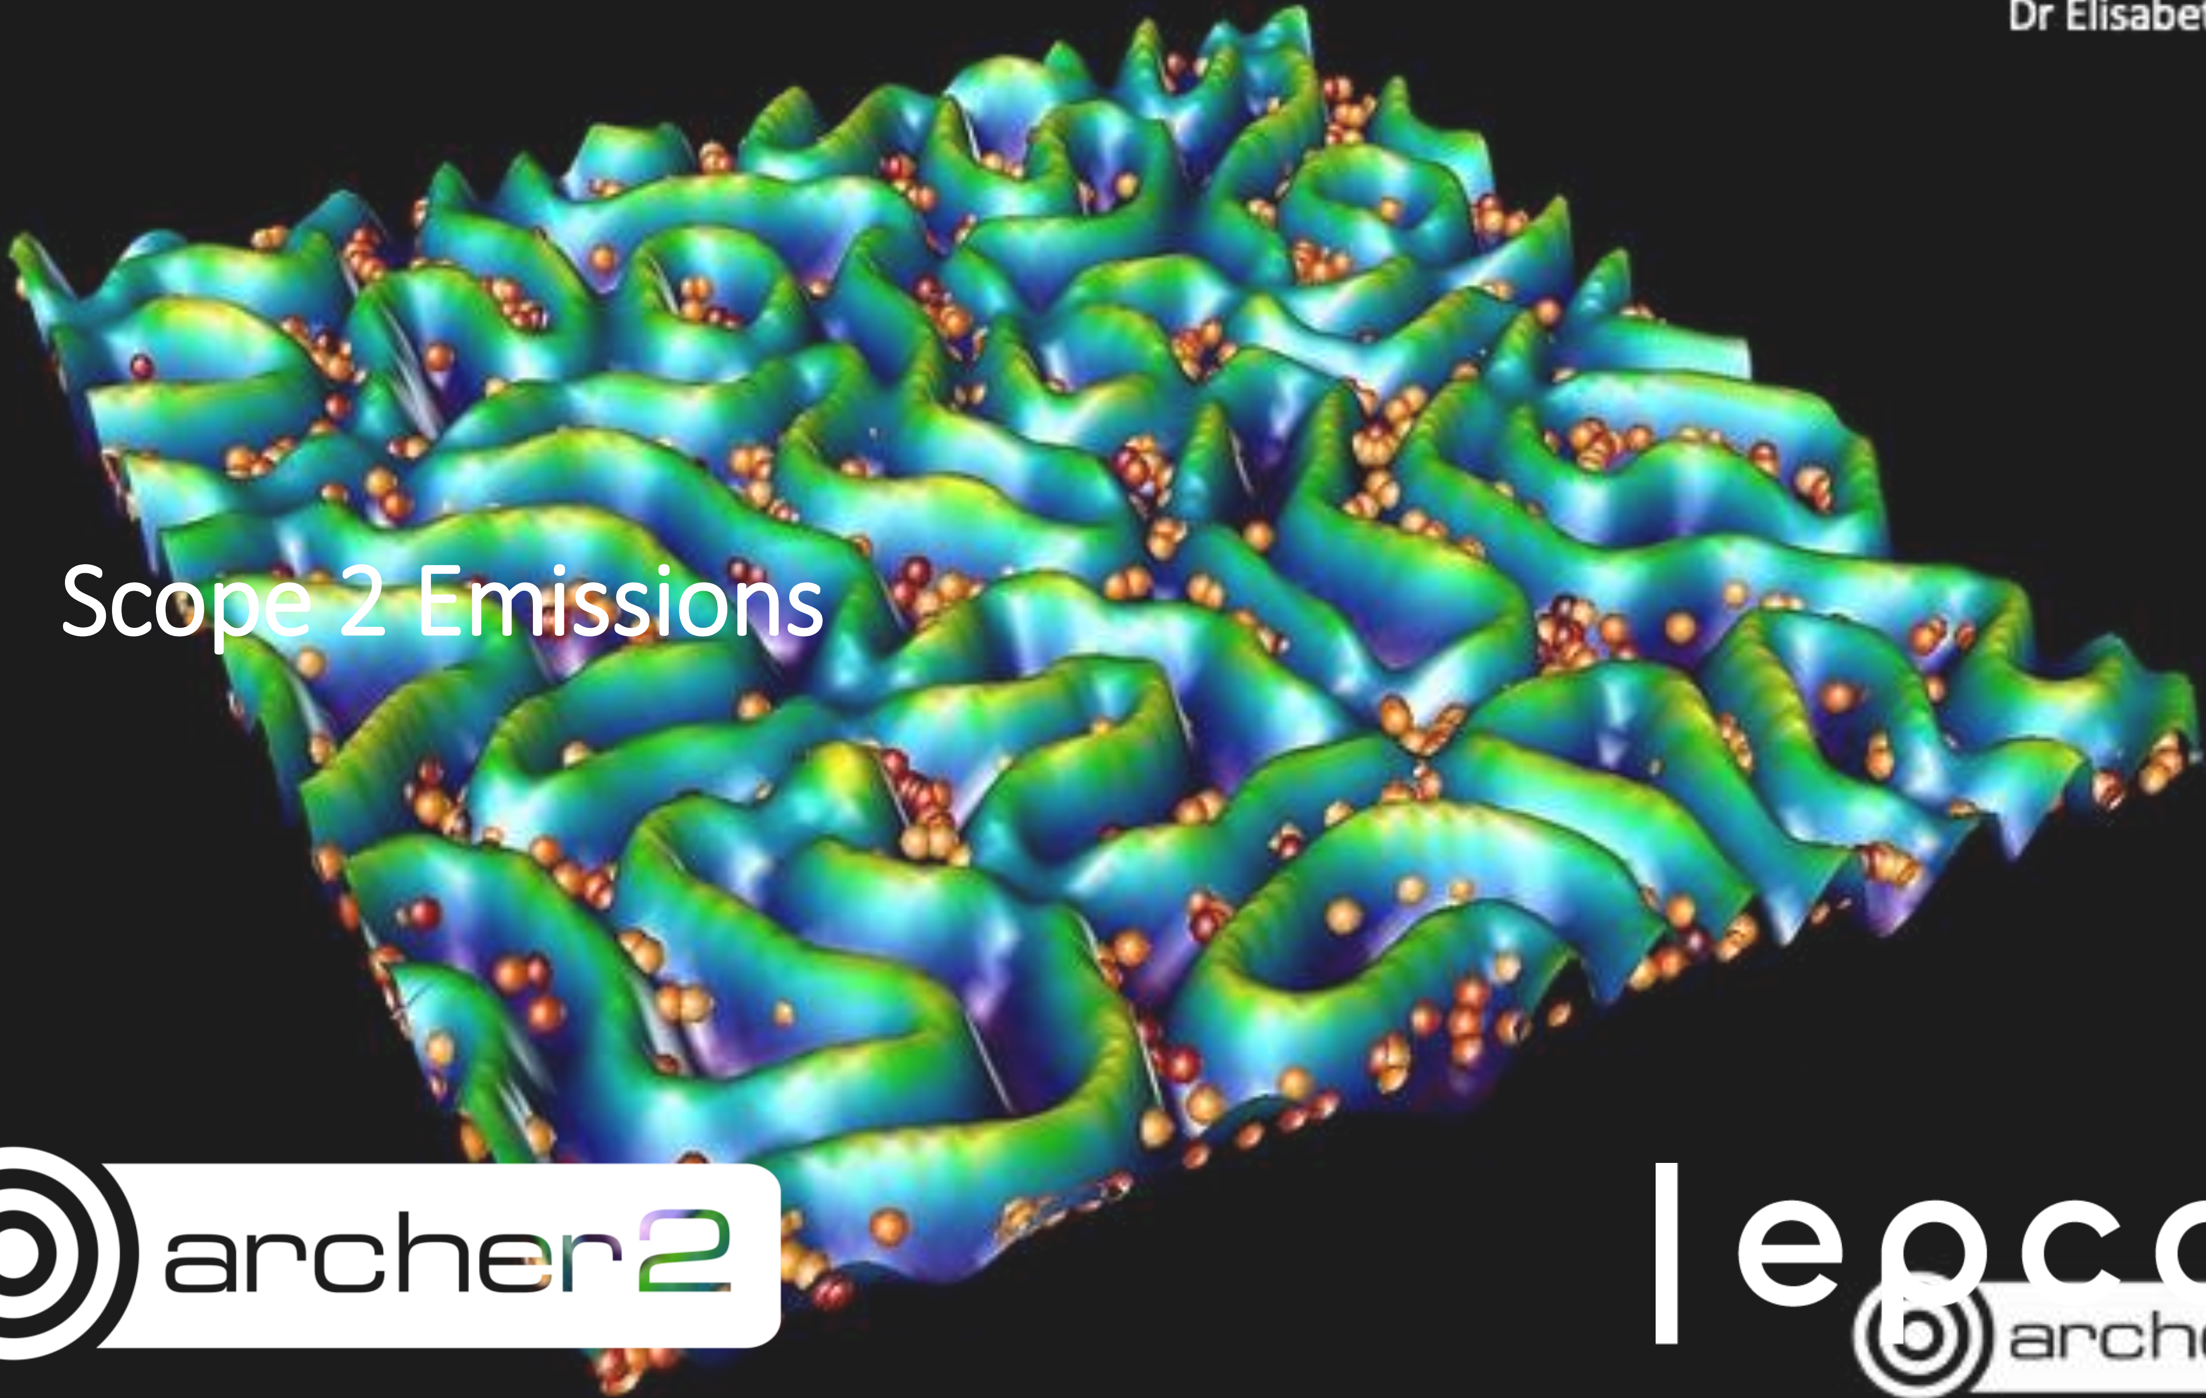

Scope 2 (energy) emissions: power draw

	Compute Nodes	Interconnect	Lustre File System	Home File System	Coolant distribution Units	
Count	5,860 nodes	768 switches	5 File systems	2 File systems	6 CDUs	
Idle (each)	1,300 kW (0.22kW)	170 kW (0.22 kW)	40 kW (8 kW)	16 kW (8 kW)	96 kW (16 kW)	1,600 kW Total
Loaded (each)	2,400 kW (0.41 kW)	240kW (0.24 kW)	40 kW (8 kW)	16 kW (8 kW)	96 kW (16 kW)	2,800 kW Total
%	86%	9%	1%	0%	3%	100%

- Energy use **dominated by compute cabinets**; storage power not important
- Idle power draw of **compute nodes is high** – likely dominated by memory and NIC
    - High system utilisation is critical to both energy and emissions efficiency

Scope 2 emissions scenarios

	Green energy	South Scotland	UK	World
gCo2/kWh	~0	26[2,3]	154[2]	481[4]
Scope 2 Emissions: per annum[1] (kgCO_2)	~0	638,000	3,780,000	11,800,000
Scope 2 Emissions: 6 years (kgCO_2)	~0	3,830,000	22,680,000	70,800,000
Scope 3 Emissions (kgCO2e)	7,320,000	7,320,000	7,320,000	7,320,000
% Scope 2 of total emissions over 6 years	0%	34%	77%	91%

- In low emissions energy scenarios, **emissions are dominated by Scope 3 (embodied) emissions**
- As the energy grid decarbonises, **Scope 3 emissions become more and more important**
- When Scope 3 emissions dominate, **running the system for as long as possible and getting the most out of it gives the best emissions efficiency**
- University of Edinburgh supplied 100% certified renewable energy contract[4]

[1] Assuming 2.8 MW power draw
[2] Median value from 12 months: 1 Apr 2023 – 31 Mar 2024. https://electricityinfo.org/
[3] UK Government guidance is to report using location-specific grid carbon density values: (p48, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/850130/Env-reporting-guidance_inc_SECR_31March.pdf
[3] https://ourworldindata.org/grapher/carbon-intensity-electricity
[4] These types of contracts are largely debunked as a true source of green energy, e.g. https://galleryclimatecoalition.org/green-energy-suppliers/

Scope 2 and Scope 3 emissions

- Understanding the relative importance of Scope 2 and Scope 3 emissions impacts how we deliver services
- Emissions efficiency vs energy efficiency
  - Scope 2 emissions dominate
    - Energy efficiency and emissions efficiency are aligned
    - Reducing energy use reduces cost, improves energy access and improves emissions efficiency
  - Scope 3 emissions dominate
    - Maximising emissions efficiency implies maximising performance (likely maximises energy use)
    - Implications for higher electricity use – is energy efficiency still important?
    - Access to affordable electricity for all is a UN sustainability goal
    - Cost – value for money

Reducing Scope 2 emissions

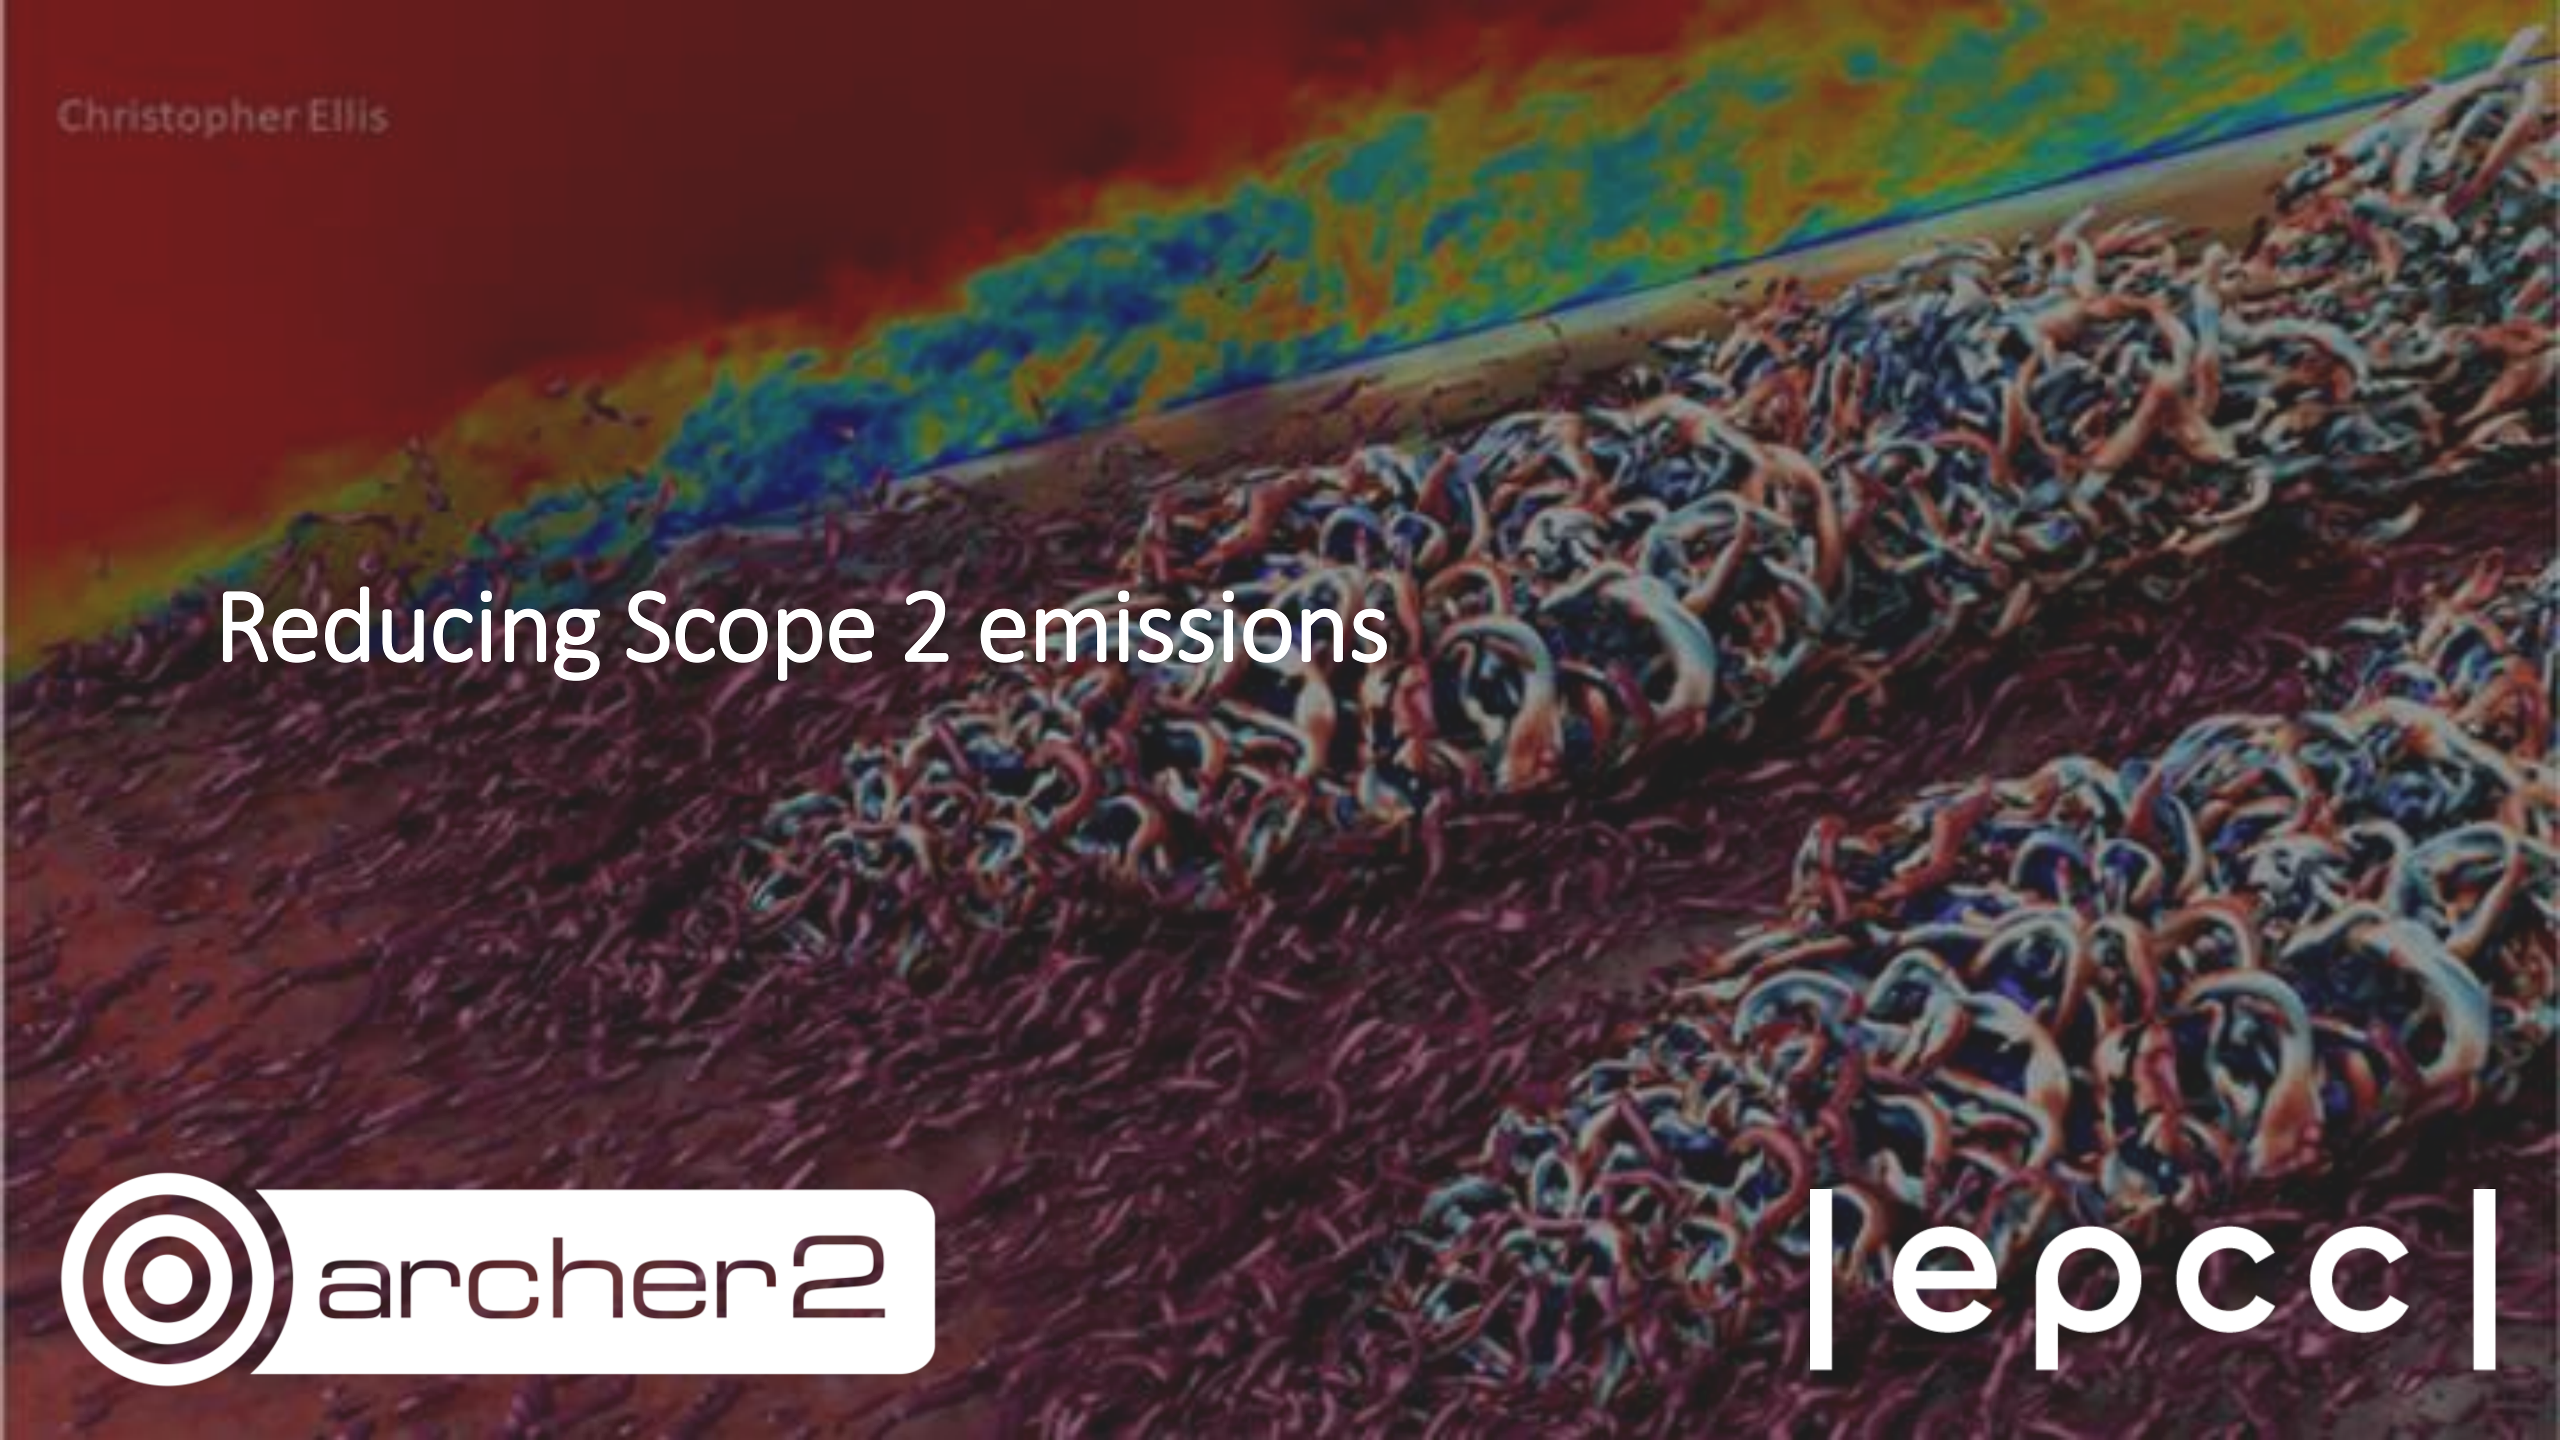

Improving Scope 2 emissions efficiency

- Through the purchase and setup of more energy efficient hardware
  - See earlier questions regarding Scope 3
- Through the use of system software

- HPC nodes often configured to run at highest possible power draw
  - Maximum processer frequency, turbo-boost enabled
  - Gives maximum potential floating point performance
- Many HPC applications are not compute bound
  - Potentially reduce power cap with little impact on performance
  - Improve energy and Scope 2 emissions efficiencies
- Different ways to control power cap
  - Set processor frequency, set node power cap
- Scheduling approaches
  - Link scheduler to emissions forecasts
  - Favour power dense workloads when carbon intensity of energy generation is lowest

Improving Scope 2 emissions efficiency

- Through improvements to the energy efficiency of data centres and their infrastructure

- Cooling accounts for large amounts of the energy consumed by data centres
- Free cooling at the ACF provides a significant reduction in energy consumption
- Next step is to reuse waste heat
  - Reuse for building heating (LUMI, SuperMUC, Isambard3/AI)
- All require infrastructure for transporting heat
- Critical part of net zero for large scale computing

ACF waste heat reuse: ideas

- Connect to nearby University of Edinburgh CHP plant
  - Provides heat to vet school and associated buildings
  - Most technologically straightforward approach
- Geobattery approach
  - Use water flow in abandoned mines underneath ACF
  - Transport waste heat to Edinburgh suburbs

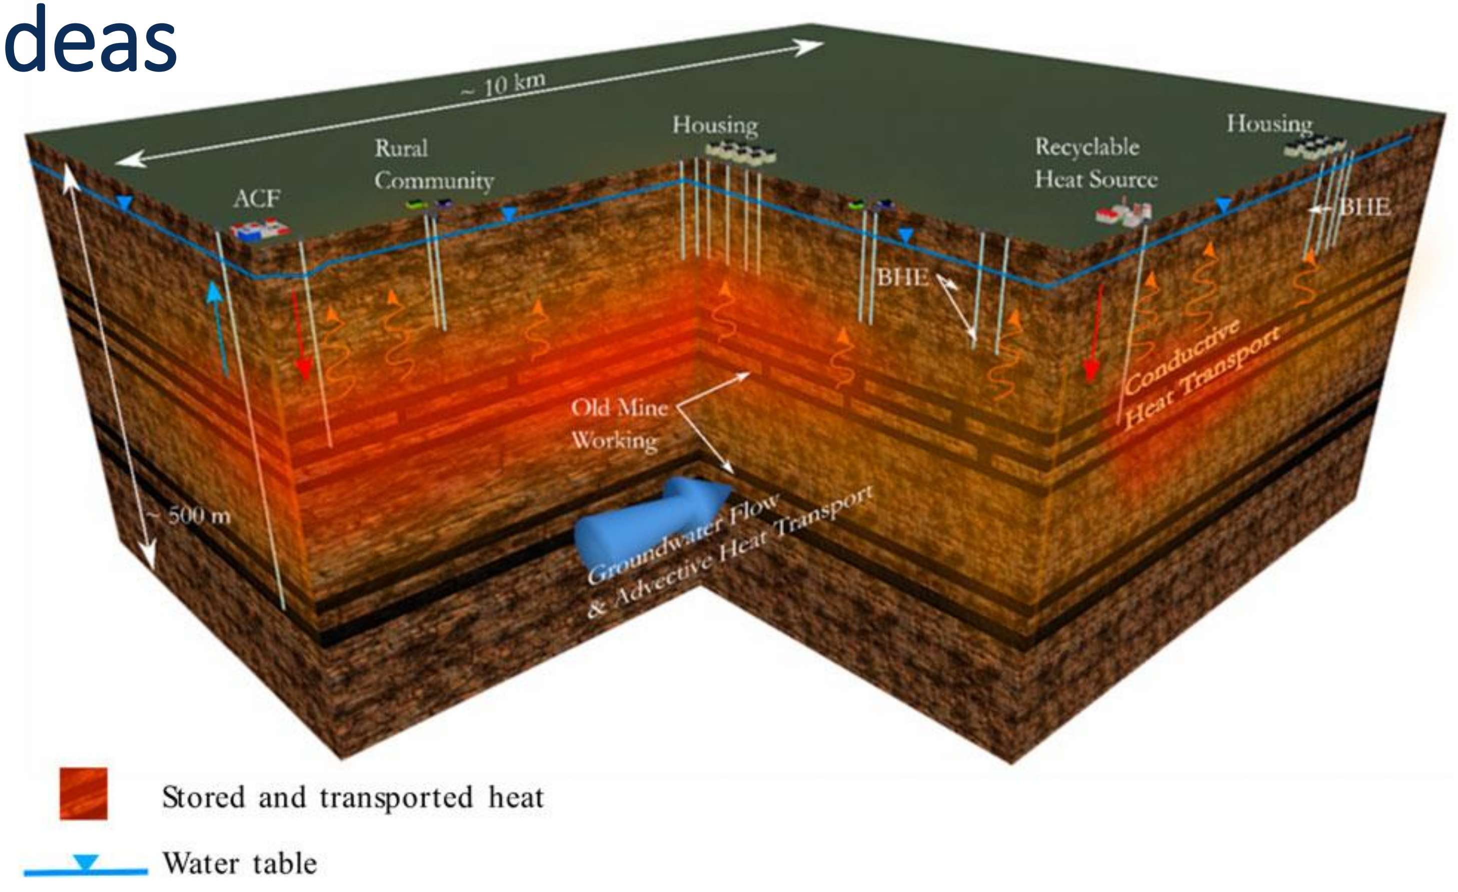

https://www.escubed.org/journals/earth-science-systems-and-society/articles/10.3389/esss.2022.10047/full

Improving Scope 2 emissions efficiency

- Improving emissions through user behaviour
  - Optimising applications, using energy efficient libraries, controlling scheduling

- Important to provide emissions information to users
- Ideally:
  - Compute Scope 2 emissions associated with every job based on energy use and data feed on carbon density from national grid
  - Also include Scope 3 emissions amortised across lifetime of service
- Challenges:
  - Energy reporting per job is not always available and may need to be supplemented with energy use from other components
  - Scope 3 (embodied) emissions data quality is currently poor

- Where is the balance?
  - Infrastructure, hardware, software, user/application

Biodiversity at the Advanced Computing Facility

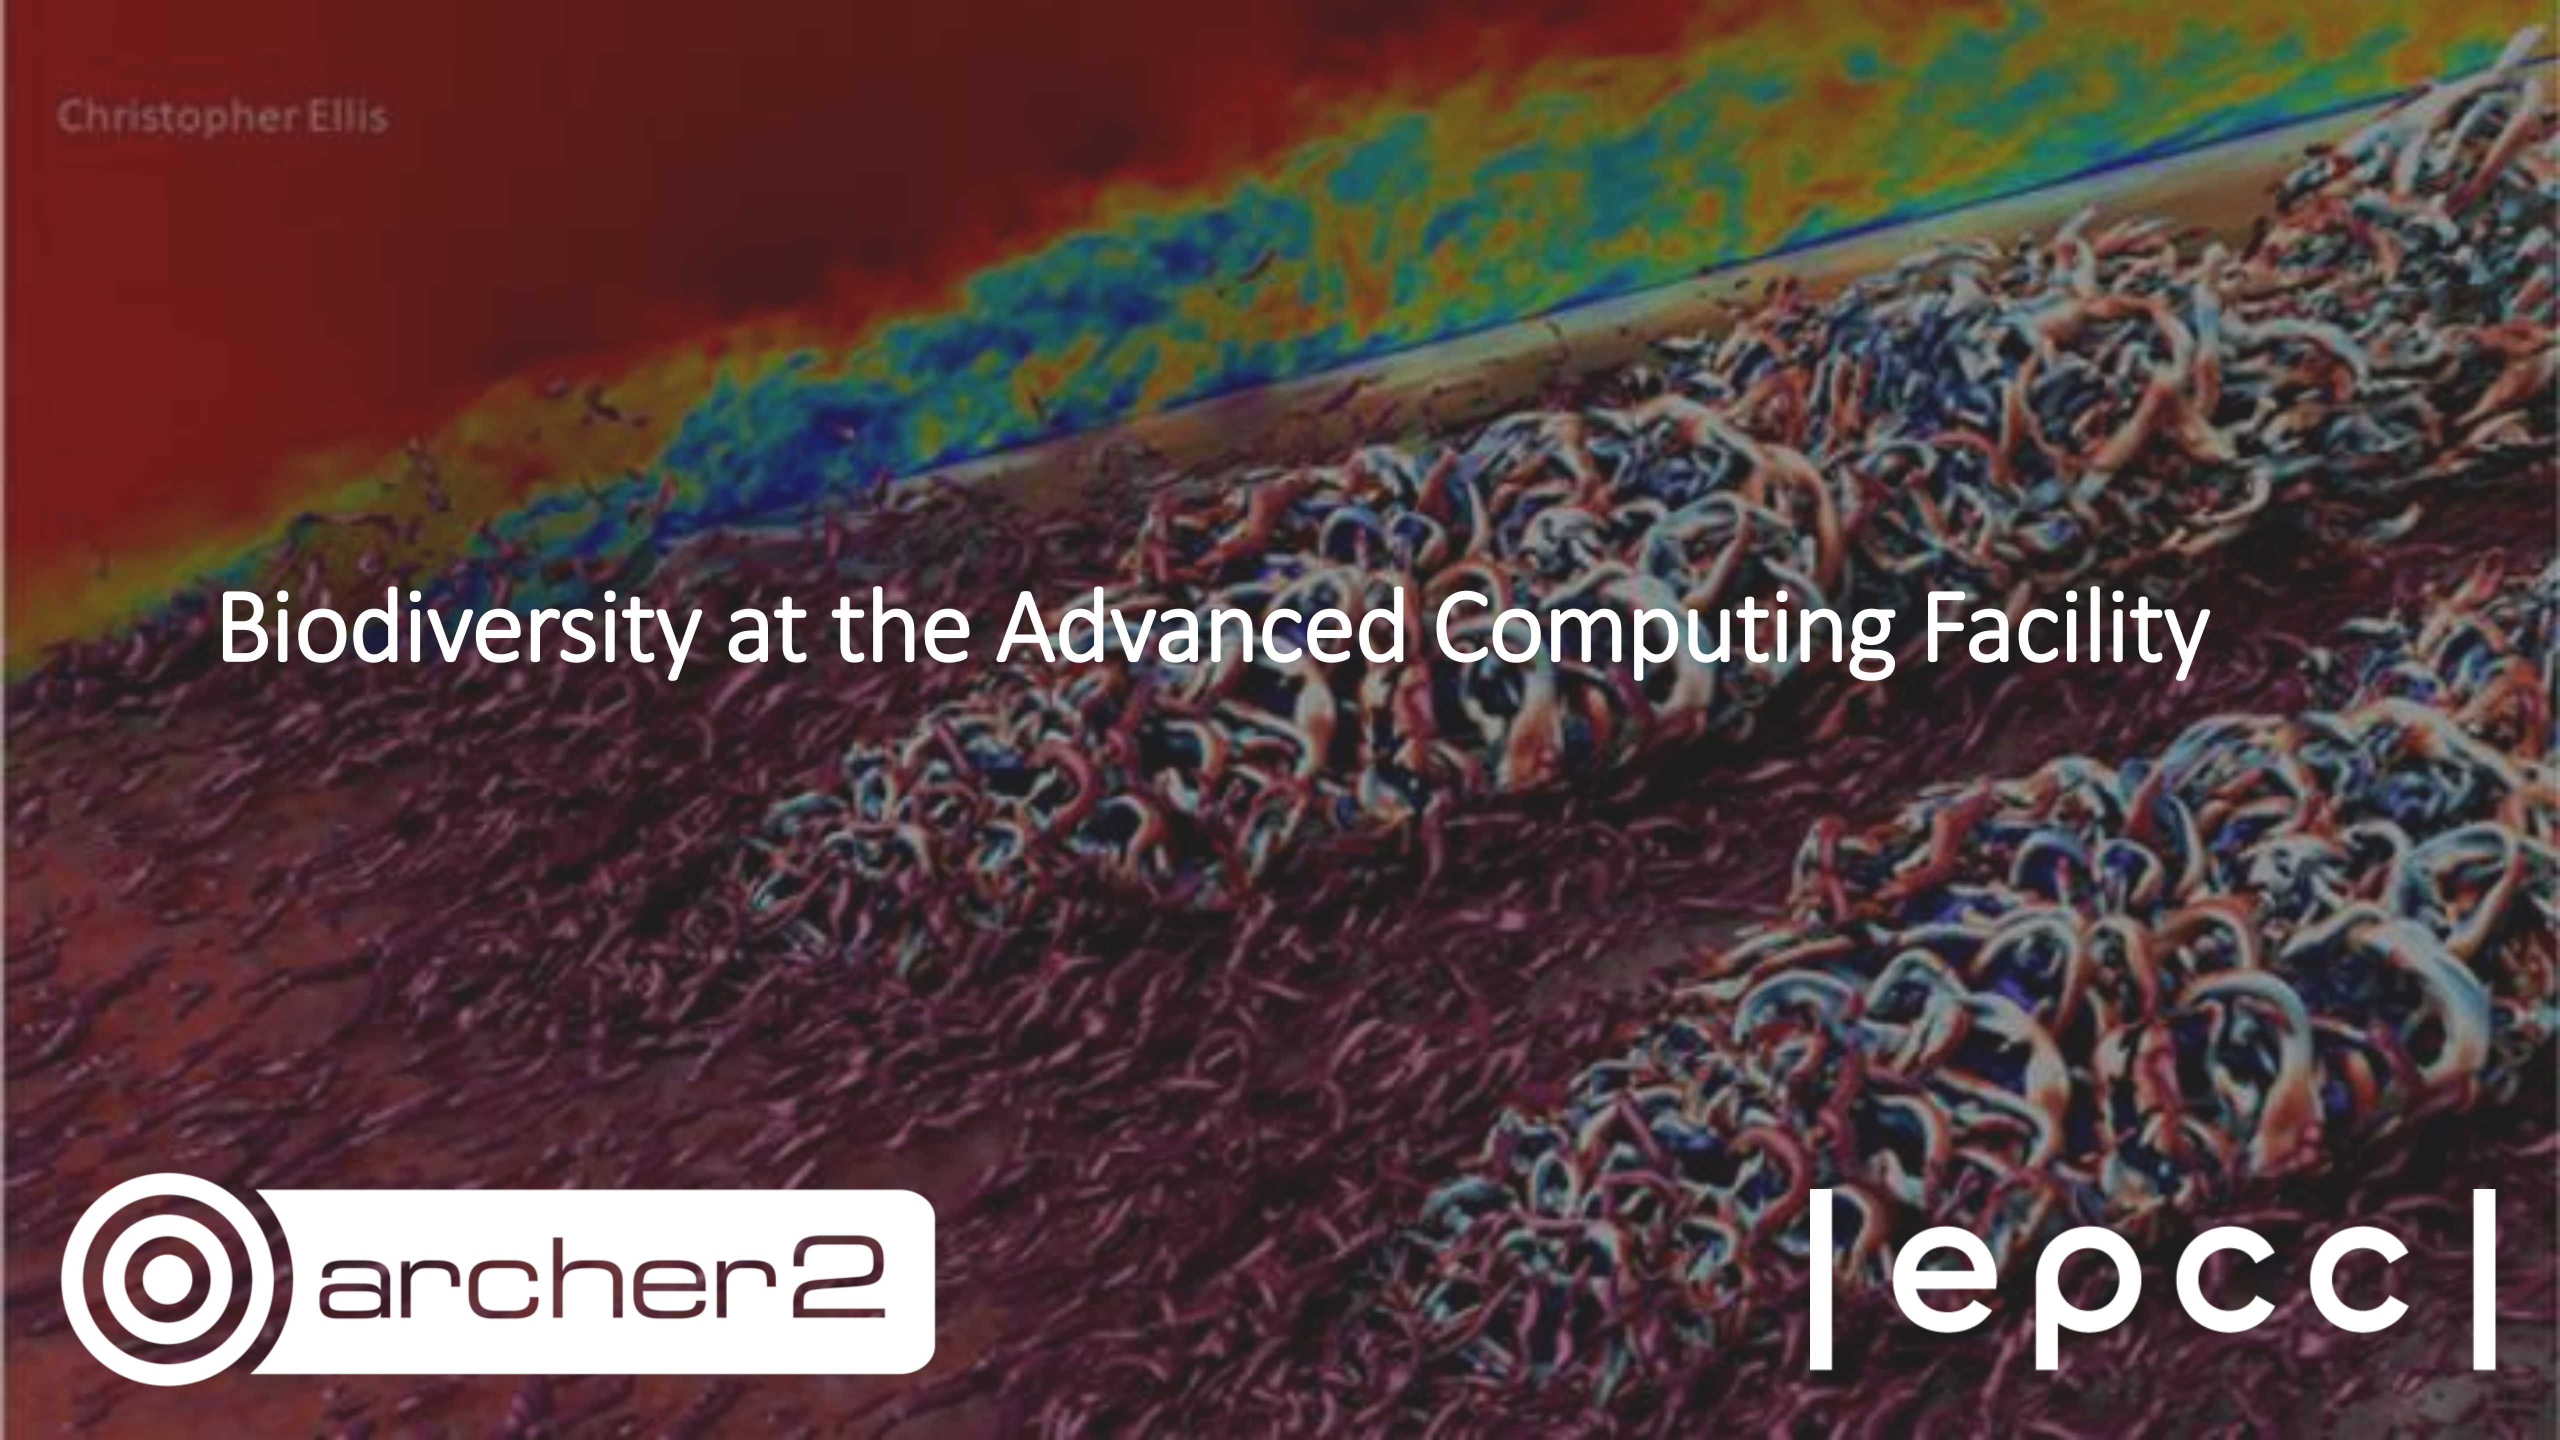

archer2

|epcc|

Data Centres and Biodiversity

- Data centres can have a negative impact on biodiversity
  - Buildings and emissions can result in habitat loss and fragmentation
- ACF
- Construction designed to leave central tree roundel intact
  - This was an original feature of the site
- Planting diverse native species as hedgerows on site
- Creating wildflower meadows to support pollinators
- Creating wildlife corridors through fencing
- Reusing felled wood to enhance habitats available on site
- Creation of dedicated bark pathways to ensure wild ground remains undisturbed by people on site

Summary

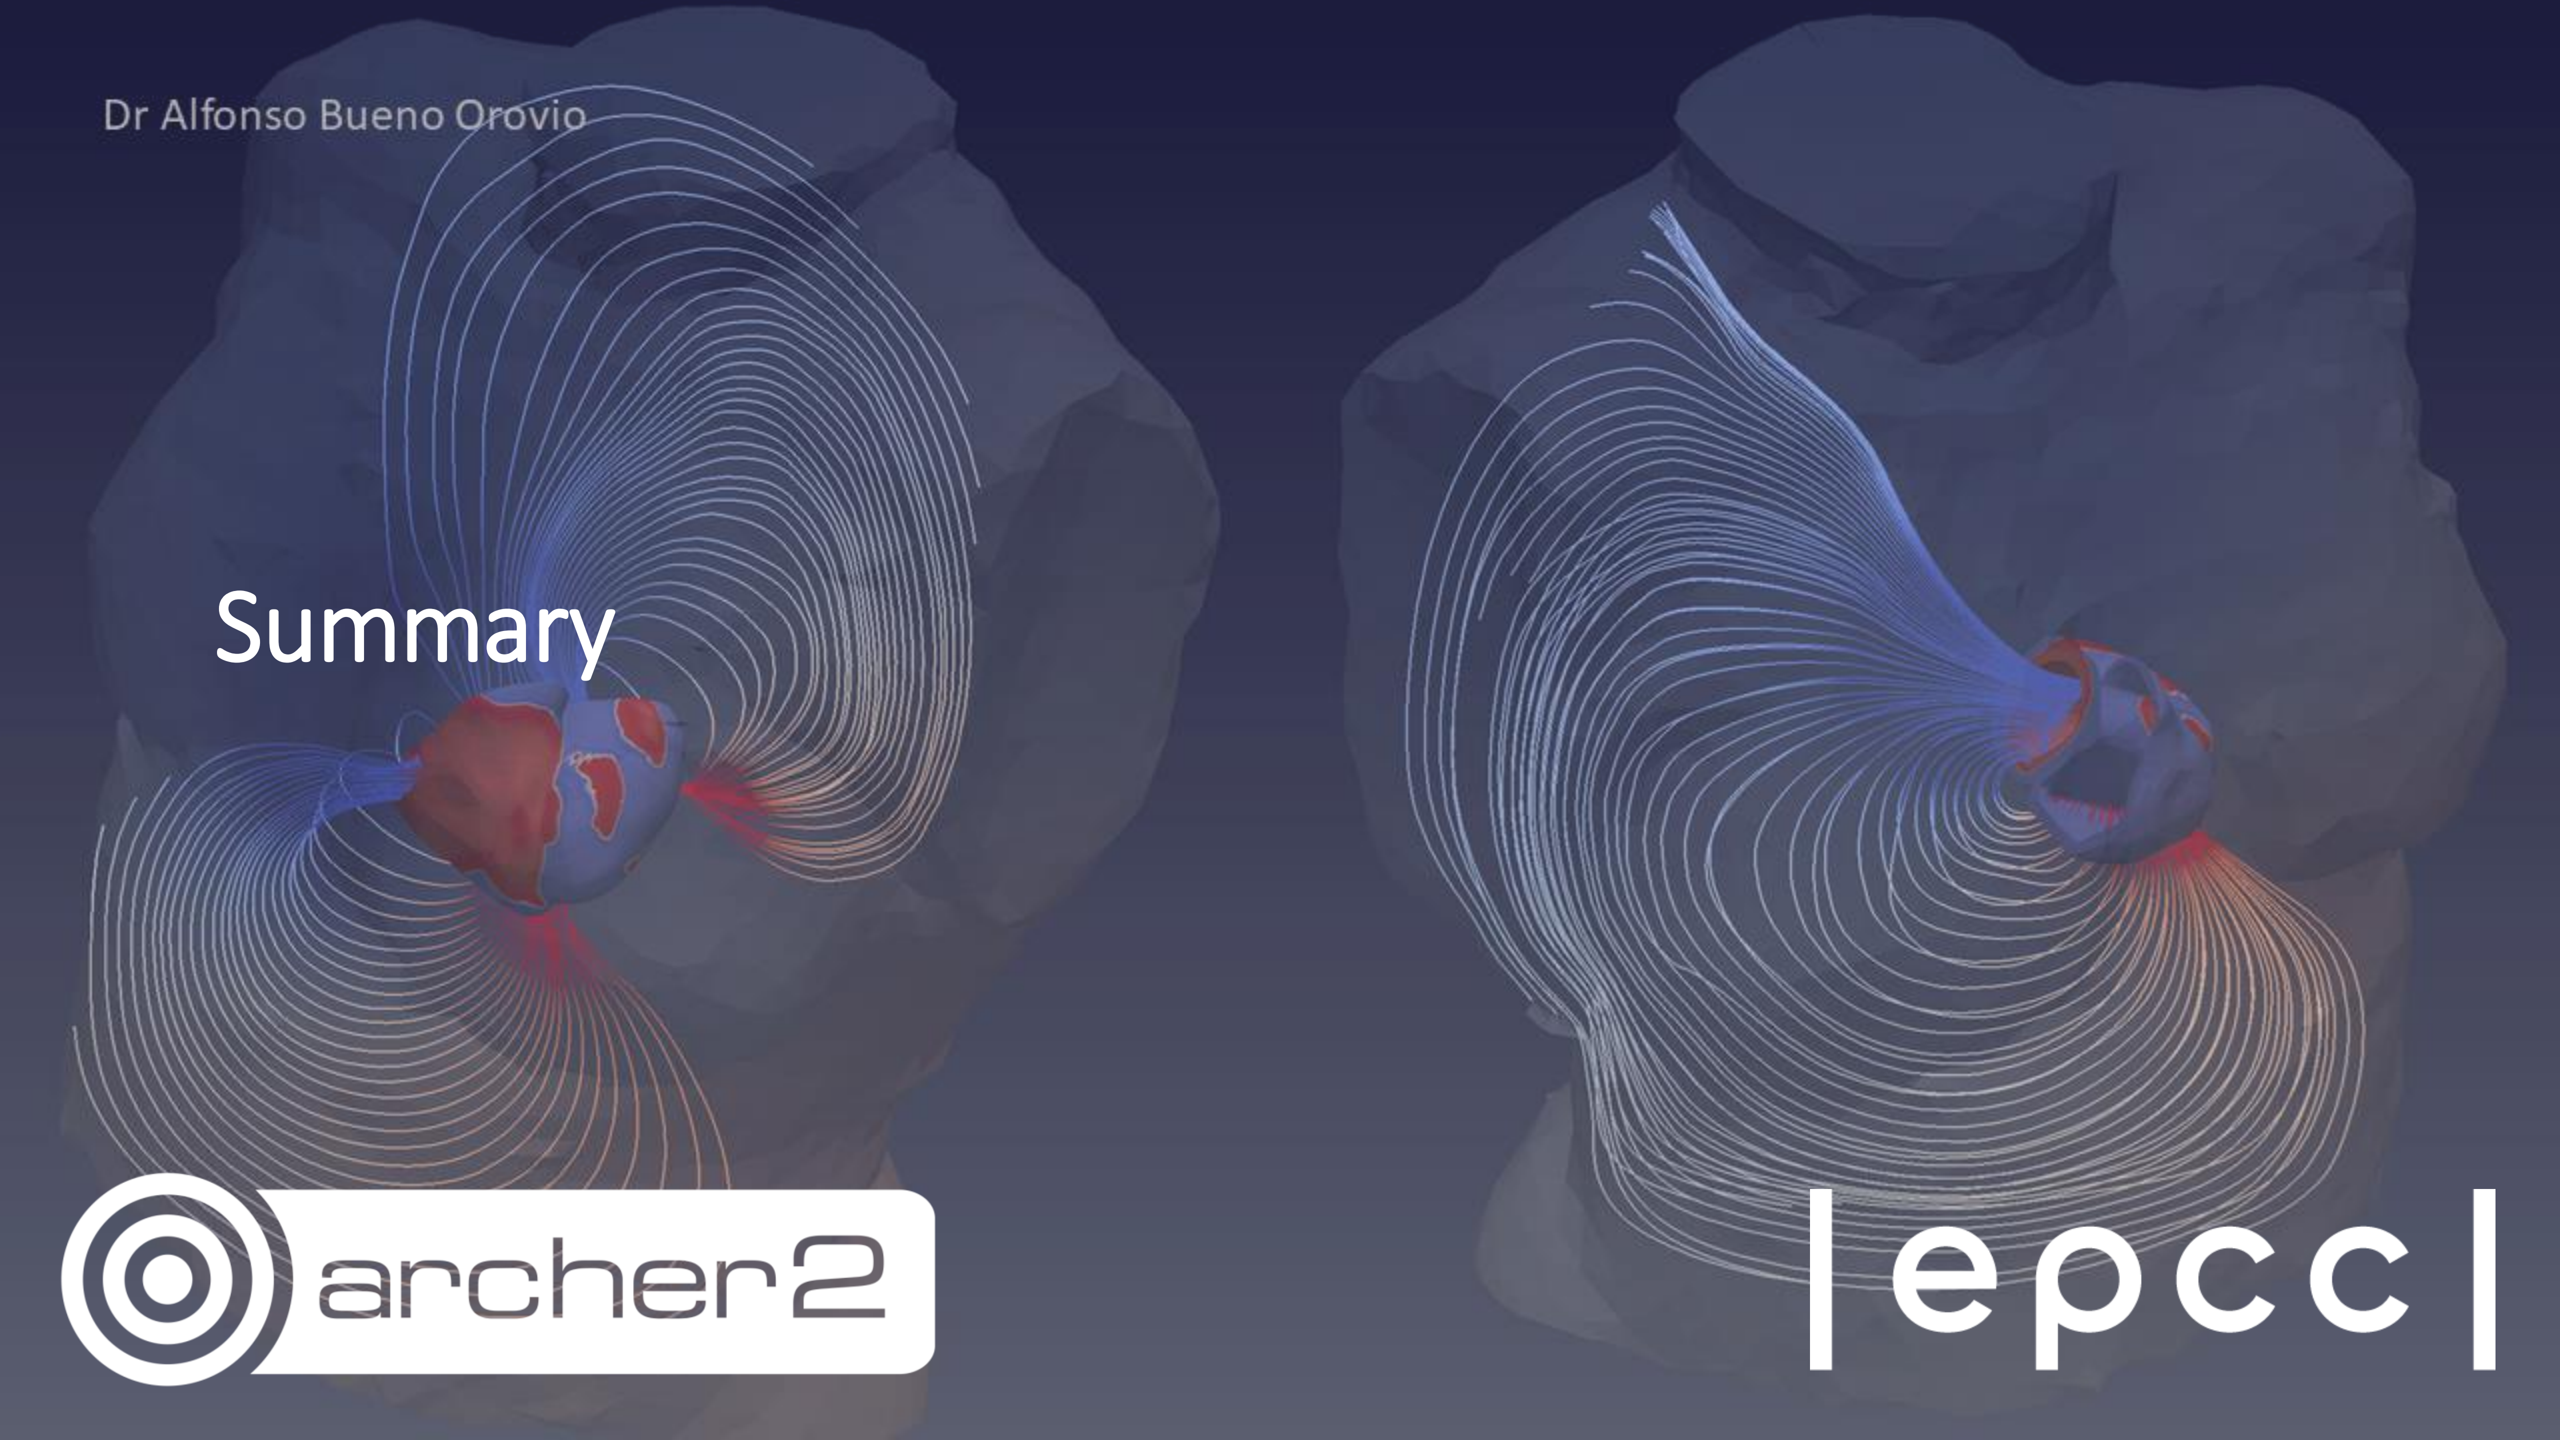

Summary

- Scope 3 emissions
  - How can we control these and plan for these within procurement?
- Scope 2 and Scope 3 emission balance
  - Initial work to understand ARCHER2 emissions
    - Significant amount of lifetime emissions are embodied (Scope 3) due to
      - Low carbon intensity of electricity generation in South Scotland/green energy contract
      - Reduction in ARCHER2 power draw by approx. 0.5 MW
  - Emissions and energy efficiency goals are not always aligned
- Scope 2 emissions
  - Understanding the balance between infrastructure, hardware, software, user/application
- Advanced computing facility
  - Biodiversity core component of site maintenance and development

Rough comparisons

	ARCHER2	FRONTIER	ChatGPT3 (Single training run)	ChatGPT3 (Inference)
Compute Nodes	5,860	9,408	2,640	1,000
Scope 2 Emissions: per annum (kgCO_2)	638,000	38,400,000	552,000* *Unclear how many runs were required	9,000,000^ ^Does not include re-training
Scope 2 Emissions: 6 years (kgCO_2)	3,830,000	230,400,000	< / >	54,000,000
Scope 3 Emissions (kgCO_2e)	7,320,000	8,800,000	6,800,000	2,600,000* * Assume constant query rate

ChatGPT3 estimates from: https://medium.com/@chrispointon/the-carbon-footprint-of-chatgpt-e1bc14e4cc2a

- Scope 3 emissions similar across most large-scale systems
- Scope 2 emissions much larger for GPU systems hosted in USA
- Decarbonising energy grids is critical in GPU era of large scale HPC

Other Comparisons:

Return transatlantic flight:
~1,000 kgCO_2e per passenger

In 2024 it is estimated that globally 9.4Bn passengers on flights[2].
ARCHER2 Scope 2/annum ~= 640 passengers per year.

100g beef: 12.5 kgCO_2e
(100g pulses: 0.04 kgCO_2e)

Roughly 40 kg/year (equivalent of 400 steaks) of red meat eaten per capita in the UK[1]
ARCHER2 Scope 2/annum ~= 51,000 steaks per year (120 peoples' emissions from red meat consumption)

UK house Scope 2 emissions (2,700 kWh/annum[3], electricity only):

S. Scotland: 70 kgCO_2/annum
ARCHER2 Scope 2/annum ~= 9,000 homes
UK: 416 kgCO_2/annum
ARCHER2 Scope 2/annum ~= 1,500 homes
Roughly 30M households in the UK (19,500x ARCHER2)

[1] Source: https://www.nationalfoodstrategy.org/the-report/
[2] Source https://aci.aero/2023/09/27/global-passenger-traffic-expected-to-recover-by-2024-and-reach-9-4-billion-passengers/
[3] Source: https://www.ofgem.gov.uk/average-gas-and-electricity-usage

Decommissioning of Tesseract

- Some HPC systems include decommissioning as part of the contract
- Many resources have financial value when recycled.
- Tesseract was an Extreme Scaling CPU based DiRAC system. Decommissioned in March 2023
- Where possible hardware was recycled
  - Within the ACF and to other centres

Improving Scope 2 emissions efficiency

- HPC nodes often configured to run at highest possible power draw
  - Maximum processer frequency, turbo-boost enabled
  - Gives maximum potential floating point performance
- However, many HPC applications are not compute bound
  - Potentially reduce power cap with little impact on performance
  - Improve energy and Scope 2 emissions efficiencies
- Different ways to control power cap
  - Set processor frequency
    - Advantages: often under user control via scheduler, simple to implement
    - Disadvantages: crude control, particularly on CPU, typically set once for whole run
  - Set node power cap
    - Advantages: fine-grained control, potentially changed multiple times during single run
    - Disadvantages: usually only available to privileged user
- Scheduling approaches
  - Link scheduler to emissions forecasts
  - Favour power dense workloads when carbon intensity of energy generation is lowest