Sustainable Reality
25th September 2024

UNIVERSITY of York

ucisa Awards 2024 Sustainable Digital Project or Initiative Winner

Viking - One year on

Emma Barnes & Richard Fuller

alcesflight EcoDataCenter Sunspeed PHYSICAL IT SERVICES MADE SIMPLE paloalto NETWORKS AMD CORNELIS NETWORKS Arelion NVIDIA DELL

Introduction

- UoY is a research-intensive Russell Group University
- Research IT team about 6 years old
- Group manages infrastructure used for research and provides individual support
- Fully integrated into IT Services
- Research Software Engineers also sit within this team

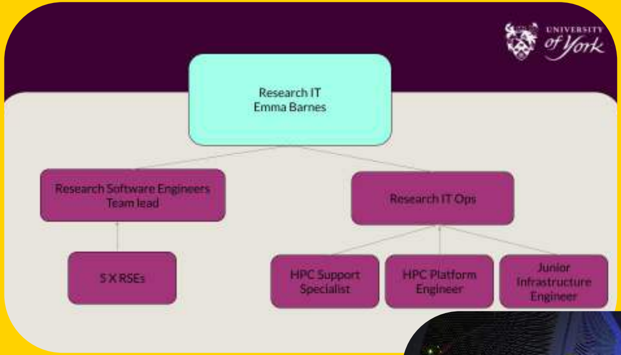

UoY Journey to Sustainable Research IT

Viking Service

A facility that caters for the research performed at the University.

Over almost 5 years

1400 Registered users

350 Separate projects

>200 research outputs (papers, articles, datasets, software etc.)

At least **£20 million** in grant funding where Viking has contributed

University first: using an external service provider (Alces!) to run our HPC facility

University first: using an external data centre to host the hardware

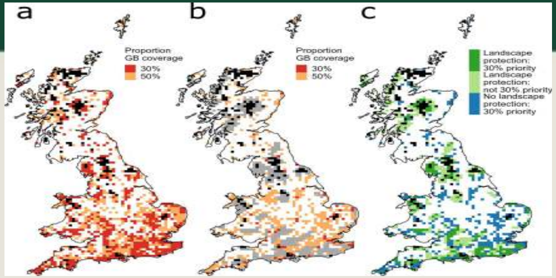

Translating area-based conservation pledges into efficient biodiversity protection outcomes *Cunningham, C.A., Crick, H.Q.P., Morecroft, M.D. et al. Commun Biol* ***4,*** *1043 (2021). https://doi.org/10.1038/s42003-021-02590-4*

Photon quantum entanglement in the MeV regime and its application in PET imaging

Watts, D.P., Bordes, J., Brown, J.R. et al. Nat Commun ***12,*** *2646 (2021). https://doi.org/10.1038/s41467-021-22907-5*

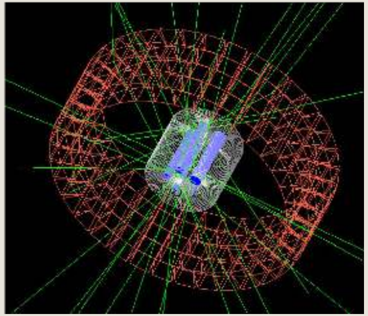

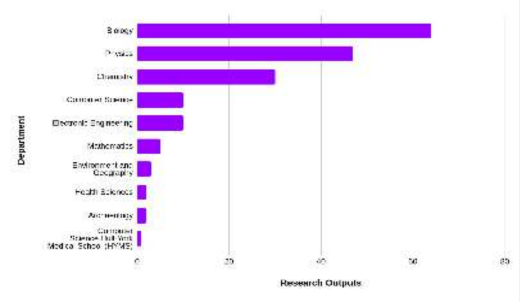

Eleanor Joan Green,
Department of Archeology

Viking 2.0

Our need will be the real creator

(or, necessity is the mother of invention)

- Electricity has become more expensive
- Budget has not increased
- Same capital buys you more nodes (and GPUs!)
- Can't host on campus
- UK data centres not inspiring
- Cloud too expensive given our usage (and capital preference)

Coupled with...

UNIVERSITY of York

- It's batch processing - surely it can go anywhere?
- We do climate research
- GPU usage - trends, CryoEM, AI/ML...
- Don't forget about water

The idea

"Can't we host it somewhere cold?"


Falun, Sweden
60.6° N

Developing the solution

UNIVERSITY of York

- Alces had already done some overseas HPC
- Recommended EcoDataCenter
- Who were inspiring and competent, in a second language
- Site visit...

Benefit Highlights

Sustainability

~100% renewable energy

Better than net-zero

Constructed from wood

Heat re-used

PUE of 1.15

Comprehensive reporting

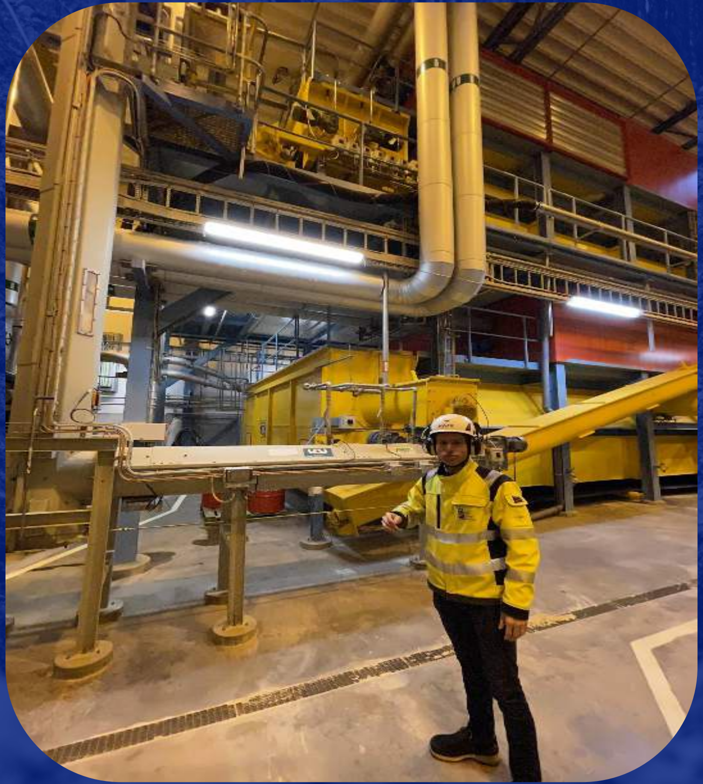

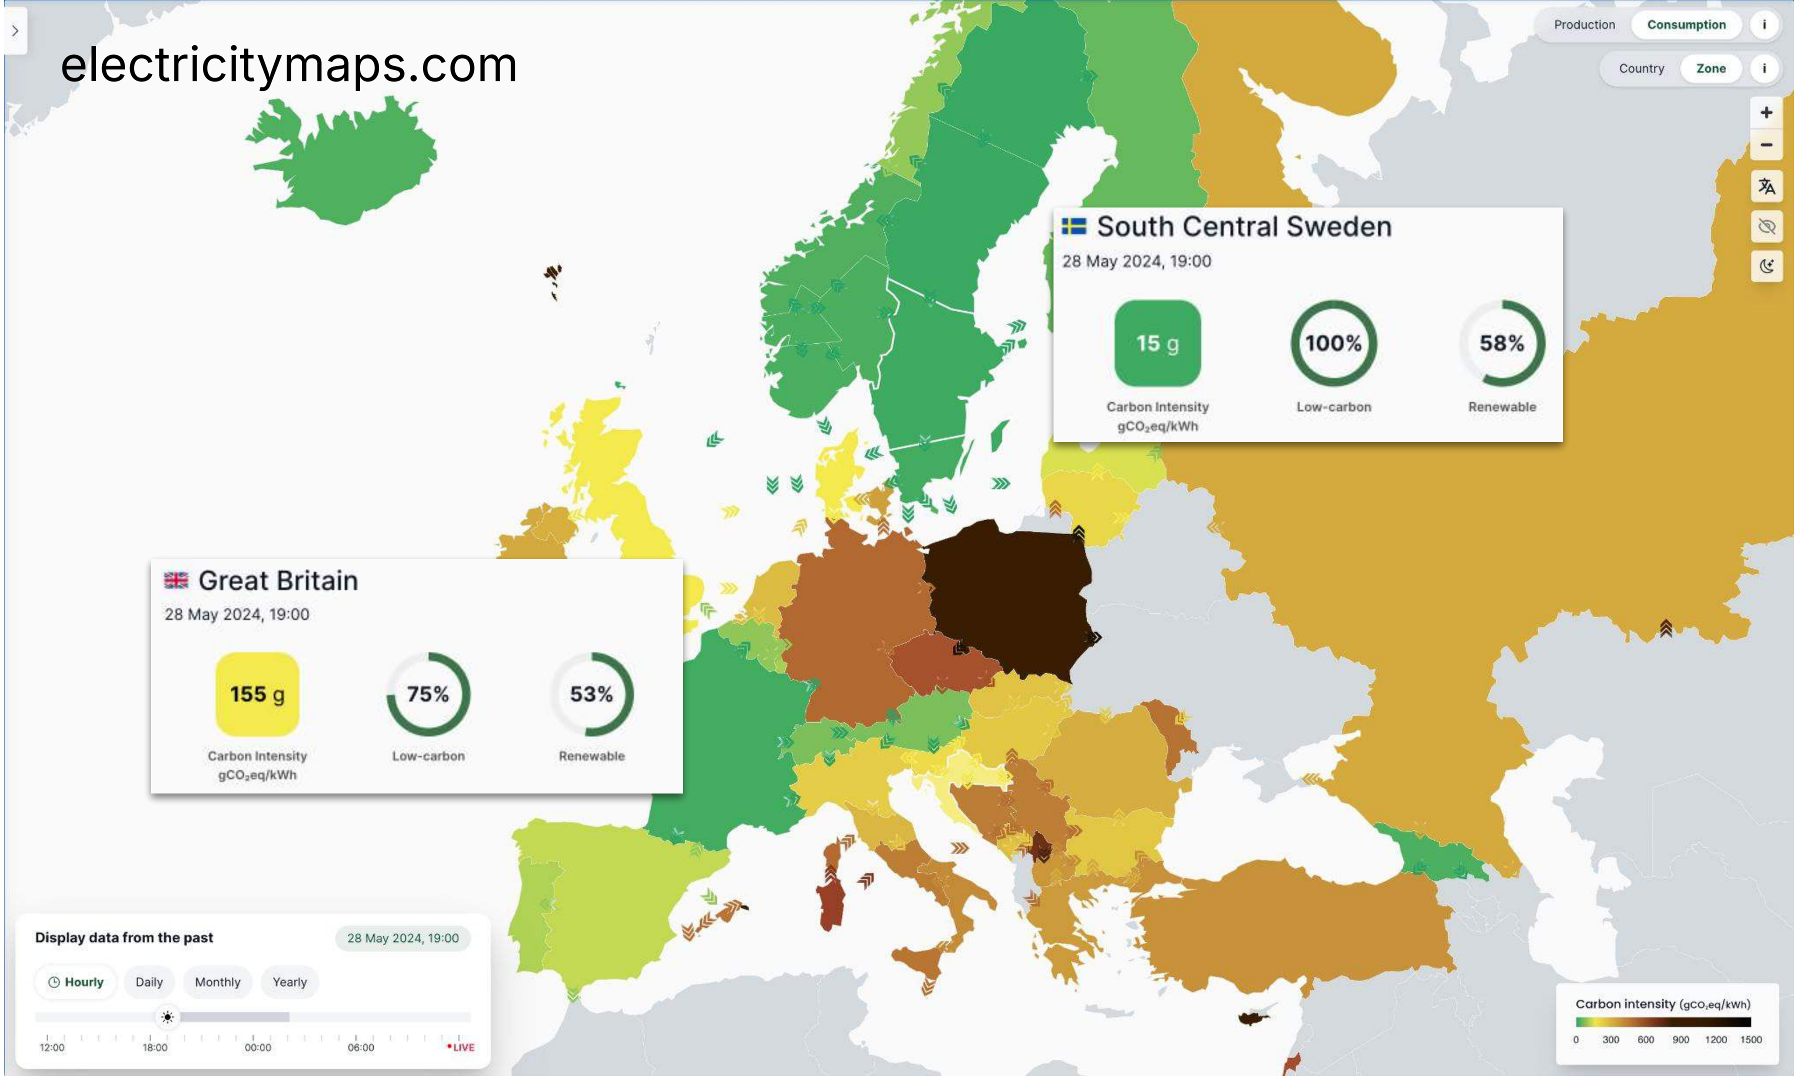
electricitymaps.com
Production
Consumption
Country
Zone
South Central Sweden
28 May 2024, 19:00
15 g
Carbon Intensity gCO₂eq/kWh
100%
Low-carbon
58%
Renewable
Great Britain
28 May 2024, 19:00
155 g
Carbon Intensity gCO₂eq/kWh
75%
Low-carbon
53%
Renewable
Display data from the past
28 May 2024, 19:00
Hourly
Daily
Monthly
Yearly
12:00
18:00
00:00
06:00
LIVE
Carbon intensity (gCO₂eq/kWh)
0 300 600 900 1200 1500

Cost

Paying for power less than what we paid in the UK in 2018

Fixed price for five years

Price went down(!) with a new wind farm

From EcoDataCenter talk at the Viking launch

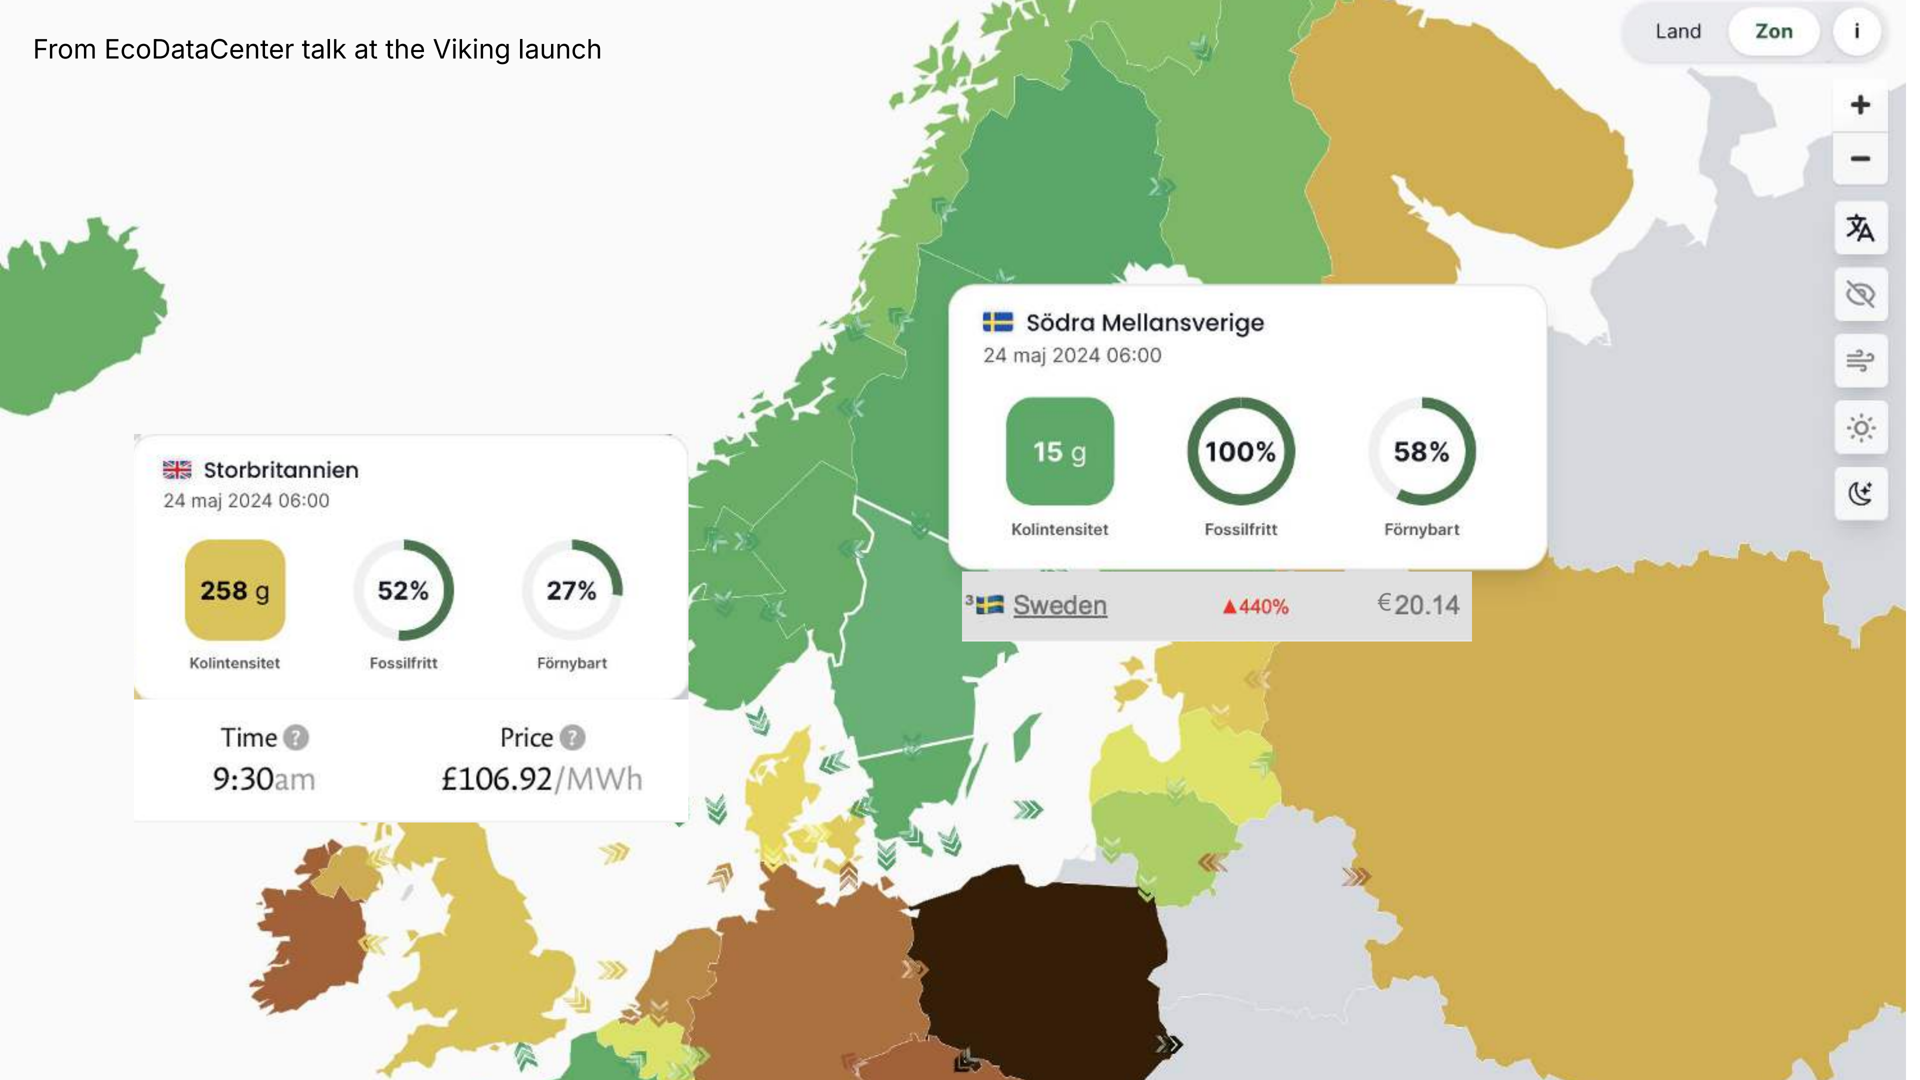

UNIVERSITY of York

Reduce on-premise DCs

Moving all of our batch processing GPUs to Sweden

Future of compute is "Sweden or cloud"

Designed for HPC

They want our heat, density is good

Air, water and immersion cooling

Security and flexibility

- ISO 9001, 14001, 27001 and validated for EN 50600
- Anything from a rack to your own building
- Straightforward to grow

Facts and Figures

UNIVERSITY of York

Viking Service

	Viking2	Viking1
Compute node only CPU cores	12,864	7000
Total standard compute nodes	134	170
Compute node generation	AMD EPYC3 7643	Intel Xeon 6138
Memory per compute node	512 GB	127 nodes 192GB
		33 nodes 384 GB
High memory node	2x 2TB	2 x 768 GB
High memory node	1x 4 TB	2 x 1.5 TB
GPUs	48 A40	8 V100s
	12 H100	
Scratch (PB)	1.5	2.5
Warm storage (PB)	2.0	
Usable NVME storage (TB)	215	48
Interconnect type	100Gb OPA	100Gb Mellanox

Viking Service

- Soft launch October 2023,
- ~3600 applications installed
- Reached 75% utilisation
- New docs
  https://vikingdocs.york.ac.uk
- Working on:
  - OpenOnDemand
  - Partition updates to improve usage

Viking Service

UNIVERSITY of York

- Soft launch October 2023,
- ~3600 applications installed
- Reached 75% utilisation
- New docs https://vikingdocs.york.ac.uk
- Working on:
    - OpenOnDemand
    - Partition updates to improve usage

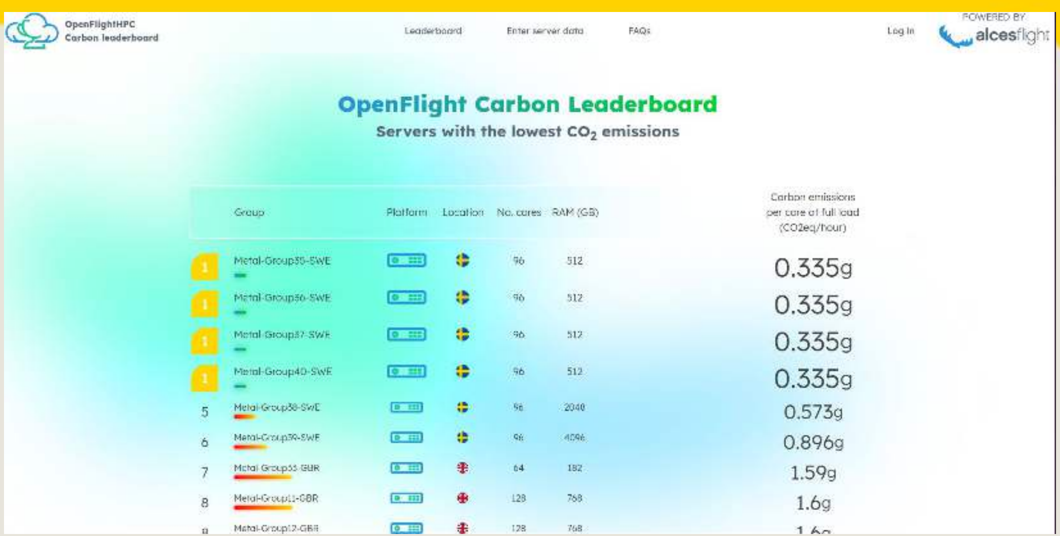

https://leaderboard.openflighthpc.org/leaderboard/grouped

Facts and figures

Average power usage per month

Viking1 ~ 53910 kWhr

Viking2 ~ 82291 kWhr

University of York
April 2024

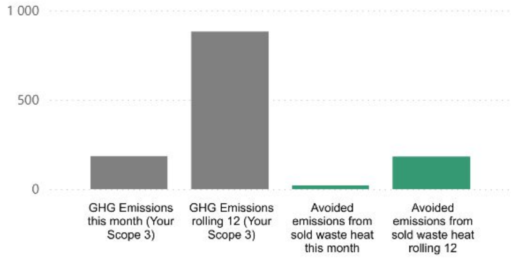

Customer Climate and Water Report

EcoDataCenter

Total GHG emissions this month (E-liabilities)

	Month (kg CO2e)	12 Months (kg CO2e)
EcoDataCenter's Scope 1 emissions	185	882
GHG Emissions From Fuels	185	882
Refrigerant leaks	0	0
EcoDataCenter's Scope 2 emissions	0	0
Hydro Power 75% (generation)	0	0
Wind Power 25% (generation)	0	0
EcoDataCenter's Scope 3 emissions	1 213	7 226
Purchased goods and services – construction materials	28	127
Capital goods – datacenter installations	275	1 178
Upstream emissions of purchased fuels	42	199
Upstream emissions of purchased electricity (life cycle emissions)	732	4 779
Transmission and distribution losses	37	163
Upstream transportation and distribution	8	51
Emissions from waste generated in operations	31	248
Business travel	21	166
Employee commuting	39	315

Customer's Scope 3 GHG emissions

Allocated GHG emissions

These are EcoDataCenter's scope 1 and 2 emissions.

These emissions should be reported as your Scope 3, category 8. Upstream leased assets' emissions according to CSRD.

Current month: 185 kg CO2e **12 months: 882 kg CO2e**

Customer's Scope 4 Avoided emissions

Avoided emissions from waste heat sold to third party

The avoided emissions are from waste heat sold to parties outside of our system boundaries and might be accounted by others too. The waste heat is replacing need of district heating

Current month: 18 kg CO2e **12 months: 181 kg CO2e**

Monthly Customer Energy Consumption

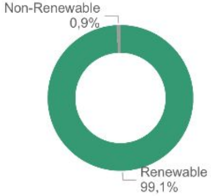

Monthly Customer water usage

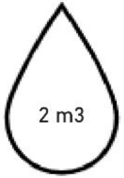

Customer Energy Consumption

Total power consumption:	82 291 kWh
Equipment power consumption:	77 285 kWh
Shared power consumption:	5 006 kWh
Renewable:	81 569 kWh
Non-Renewable:	722 kWh
Renewable %:	99,1 %
Non-renewable %:	0,9 %

Customer Water Consumption

Allocated water use for cooling current month	2 m3
Allocated water use for cooling - rolling 12 months	5 m3

Facts and figures

Carbon usage per year

Viking1 ~ 134 Tonnes CO2

Viking2 ~ 2.4 Tonnes CO2

98% reduction in Carbon!

OpenFlightHPC
Carbon leaderboard

Leaderboard Enter server data FAQs Log in

POWERED BY alcesflight

OpenFlight Carbon Leaderboard

Servers with the lowest CO_2 emissions

	Group	Platform	Location	No. cores	RAM (GB)	Carbon emissions per core at full load (CO2eq/hour)
1	Metal-Group35-SWE			96	512	0.335g
1	Metal-Group36-SWE			96	512	0.335g
1	Metal-Group37-SWE			96	512	0.335g
1	Metal-Group40-SWE			96	512	0.335g
5	Metal-Group38-SWE			96	2048	0.573g
6	Metal-Group39-SWE			96	4096	0.896g
7	Metal-Group33-GBR			64	182	1.59g
8	Metal-Group11-GBR			128	768	1.6g
8	Metal-Group12-GBR			128	768	1.6g

https://leaderboard.openflighthpc.org/leaderboard/grouped

Feedback from funders ("Outstanding")

"The university will help off-set energy consumed by the microscope/filter through investment in renewable energy and hydroelectricity"

Highlights so far…

We won an award!

UCISA Awards 2024 – Sustainable Digital Project or Initiative Award

https://www.ucisa.ac.uk/Events/2024/March/UCISA24/Event-Other-Info-List/UCISA-2024-Awards

UNIVERSITY of York

Viking Launch event and press interest!

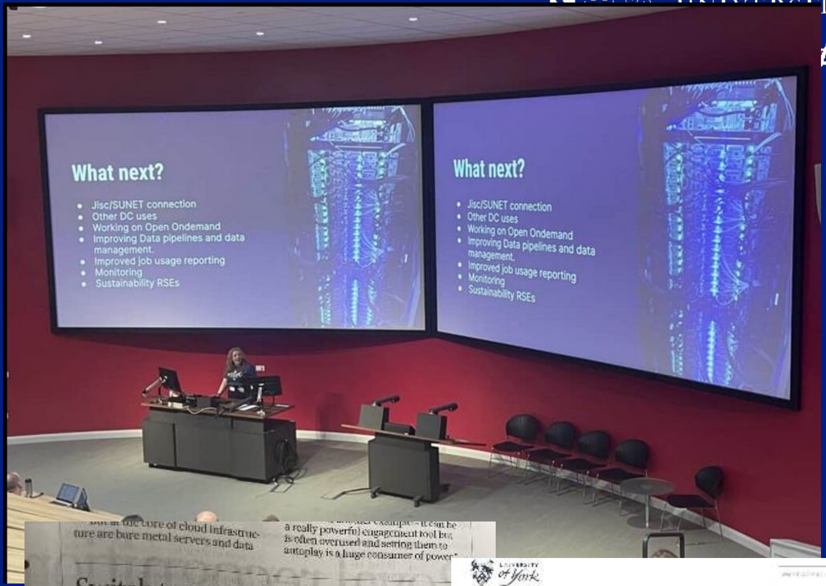

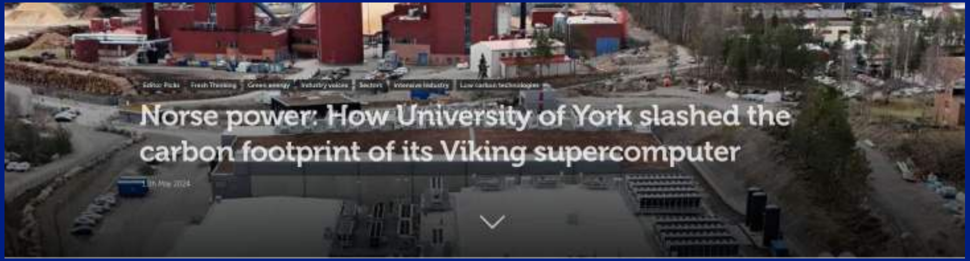

ture are bare metal servers and data

a really powerful engagement tool but is often overused and setting them to autoplay is a huge consumer of power."

Switch to renewable data centre

THE University of York has located a new high performance computer facility dubbed "Viking 2" at a renewable data centre in Sweden.

EcoDataCenter, the Swedish data centre which is now home to Viking 2, runs on 100 per cent renewable energy: 75 per cent from hydropower and 25 per cent from wind power. It is also the world's first data centre to be made from wood.

It is almost double the size of the previous facility making it a powerful tool for supporting the University's world-leading research.

Richard Fuller, assistant director of IT University of York, said: "Given the university's strong commitment to net-zero, and the fact we will be doing climate research on this facility, it was extremely pleasing to find a partner so focused on environmental sustainability, and with the capability and maturity to deliver. So much time is spent talking about sustainability, and it's exciting to be taking positive steps, and to be leading on sustainable research computing facilities."

The university was recently awarded the Sustainable Digital Project or Initiative Award at the ... with the move to EcoDataCenter forming a key part of the submission.

Sign up to become a Wildlife Friend

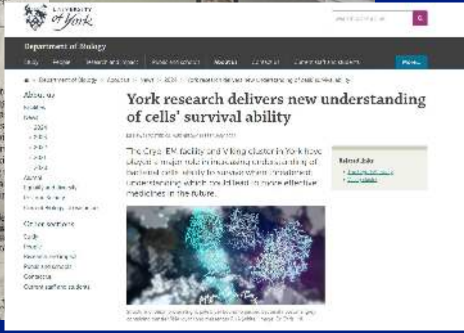

https://energyadvicehub.org/york-university-reduce-carbon-footprint-supercomputer/

Ancient DNA from archaeological grape seeds reveals the history of winemaking in Chianti, Italy

- 100s of GB of raw DNA sequencing data are generated with Illumina technology
- Viking processes the data using software developed for ancient DNA
- Ancient samples are compared to large reference databases to create figures and evaluate statistical significance

Oya Inanli

PhD student in the MSCA-funded ChemArch Doctoral Network

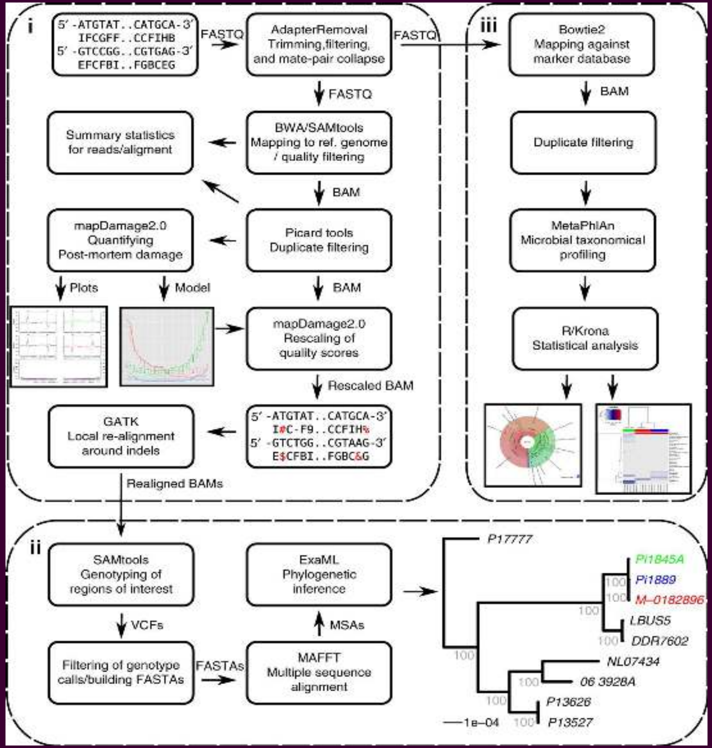

The automated PALEOMIX bioinformatic pipeline for ancient DNA developed by Schubert *et al.* (2014). *Nature Protocols* 9, 1056-1082.

Ancient DNA from archaeological grape seeds reveals the history of winemaking in Chianti, Italy

- 100s of GB of raw DNA sequencing data are generated with Illumina technology
- Viking processes the data using software developed for ancient DNA
- Ancient samples are compared to large reference databases to create figures and evaluate statistical significance

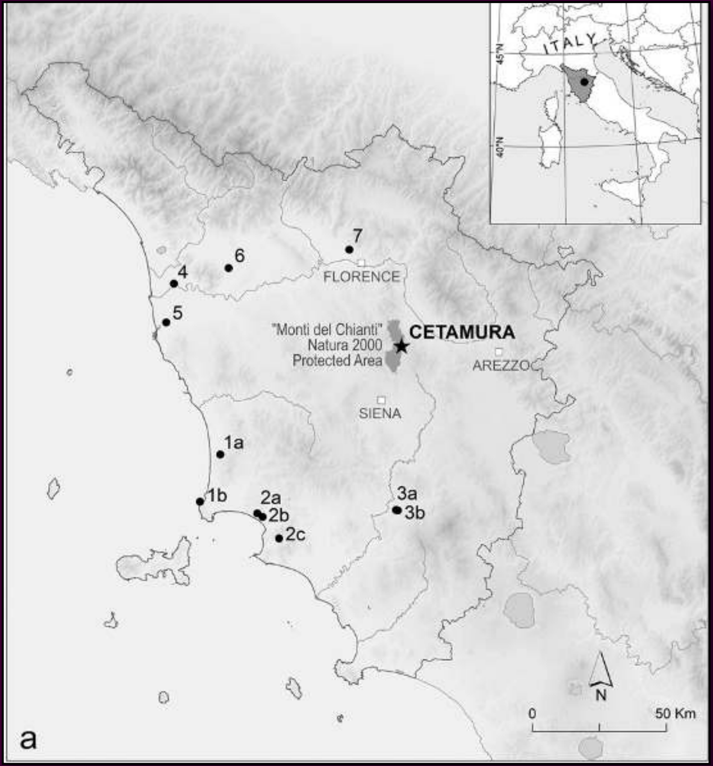

The automated PALEOMIX bioinformatic pipeline for ancient DNA developed by Schubert *et al.* (2014). *Nature Protocols* 9, 1056-1082.

Oya Inanli

PhD student in the MSCA-funded ChemArch Doctoral Network

Cryo-EM: illuminating new biology one molecule at a time

Case study: structural studies of bacterial ribosomes

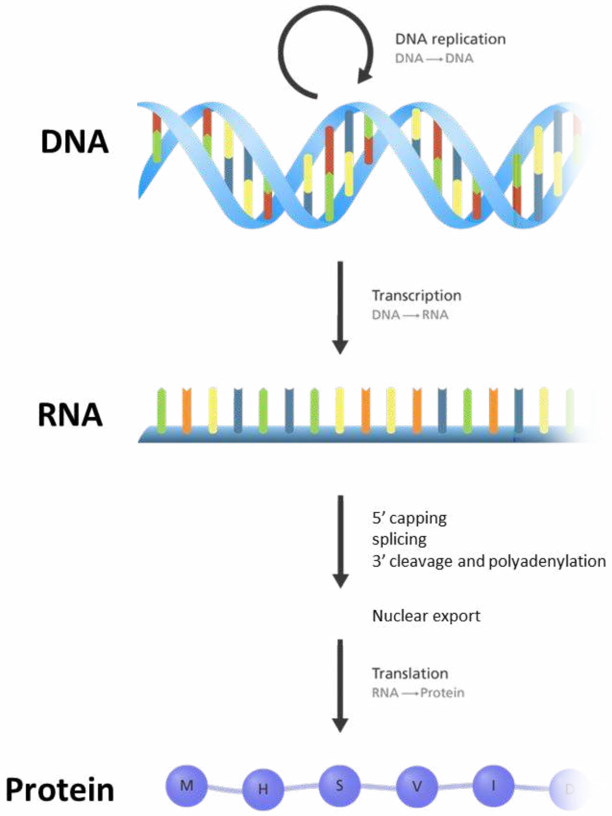

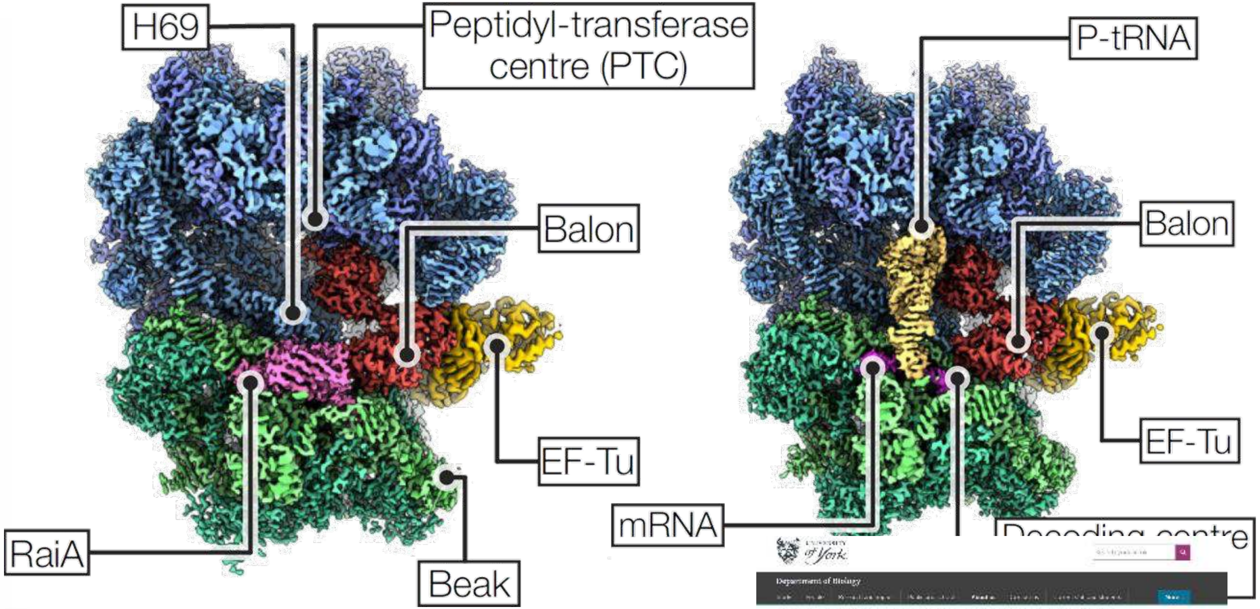

Article | Open access | Published: 14 February 2024

A new family of bacterial ribosome hibernation factors

Karla Helena-Bueno, Mariia Yu. Rybak, Chinenye L. Ekemezie, Rudi Sullivan, Charlotte R. Brown, Charlotte Dingwall, Arnaud Baslé, Claudia Schneider, James P. R. Connolly, James N. Blaza, Bálint Csörgő, Patrick J. Moynihan, Matthieu G. Gagnon, Chris H. Hill & Sergey V. Melnikov

Nature **626**, 1125–1132 (2024) | Cite this article

19k Accesses | **1** Citations | **144** Altmetric | Metrics

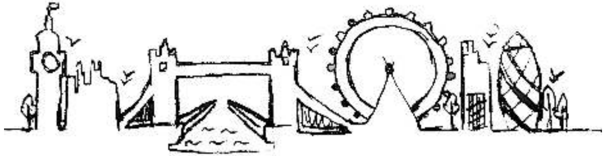

Generations of London English

Devyani Sharma

Kathleen McCarthy

Paul Kerswill

Elisa Passoni

Andy Gibson

Joe Pearce

Sam Hellmuth

Sue Fox

Naturalistic data (adolescents and adults)

Hire PDRA-1	Materials prep	Fieldwork, coding, forced-alignment	Analysis: Phonetic, syntactic, social	Dissemination

Experimental data (adolescents and adults)

	Materials prep	Laboratory experiments	Analysis: Perception, cognition, control, recall	Dissemination

Naturalistic and experimental data (children)

Hire PDRA-2	Materials prep	Hire UG assistants	School visit	Analysis	School visit	Analysis	School visit	Analysis

Corpus creation

	Materials prep	Hire PGRA-3	Transcription, data cleaning, corpus construction	UKDS upload, public samples on website

March 2023 — February 2026

      

Person-specific automatic speaker recognition

Understanding the behaviour of individuals for applications of ASR

£1,012,570

pasr.york.ac.uk

2022-25

Netherlands Forensic Institute
Ministry of Justice and Security

OxfordWaveResearch

UNIVERSITY of York

E·S·R·C
ECONOMIC & SOCIAL RESEARCH COUNCIL

Forensic Voice Comparison

Unknown Speaker

(e.g. incoming call to a call centre/ threatening phone call/ bugged car)

vs.

Known Speaker

(e.g. call from verified account holder/ police interview)

Forced alignment

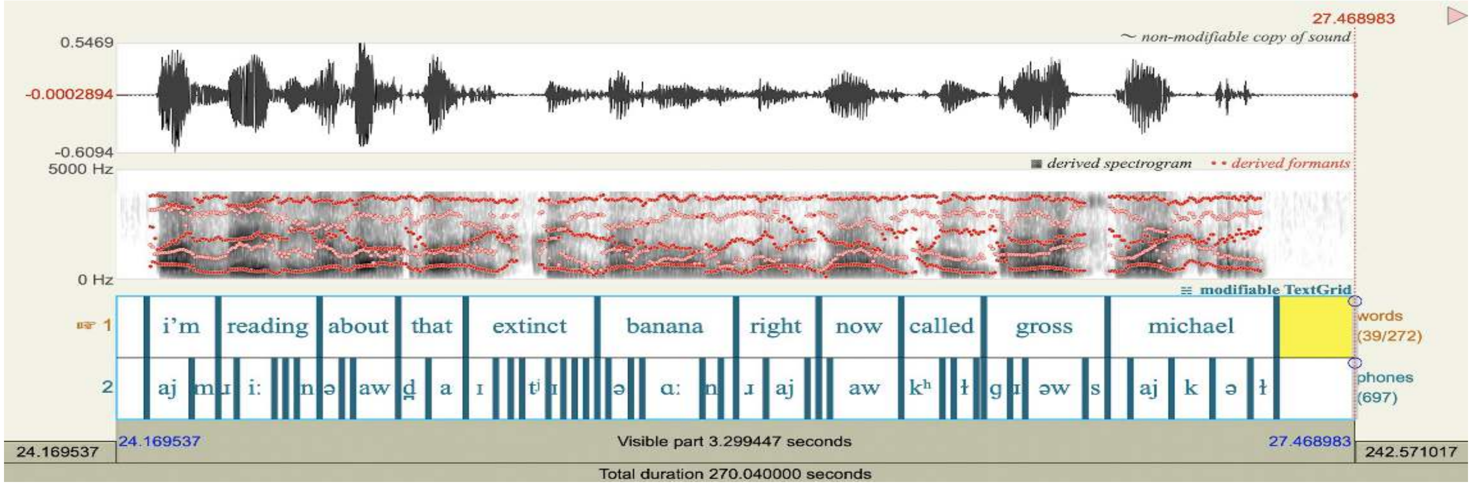

Crystal Structure Prediction

- Use computational methods to find stable crystal structures (or materials) with ‘good’ properties.
  - Good? Magnetic memory, semiconductor devices, battery & solar technology etc.
  - Stable? Want them to stay as they are for long enough to be useful.

- Isn’t there some formula for this?
  - No…but
  - DFT is correct, but it’s costly, though parellisable.
  - Using ML to speed up discovery

Scott Donaldson
Supervised by Matt Probert
School of Physics, Engineering and Technology

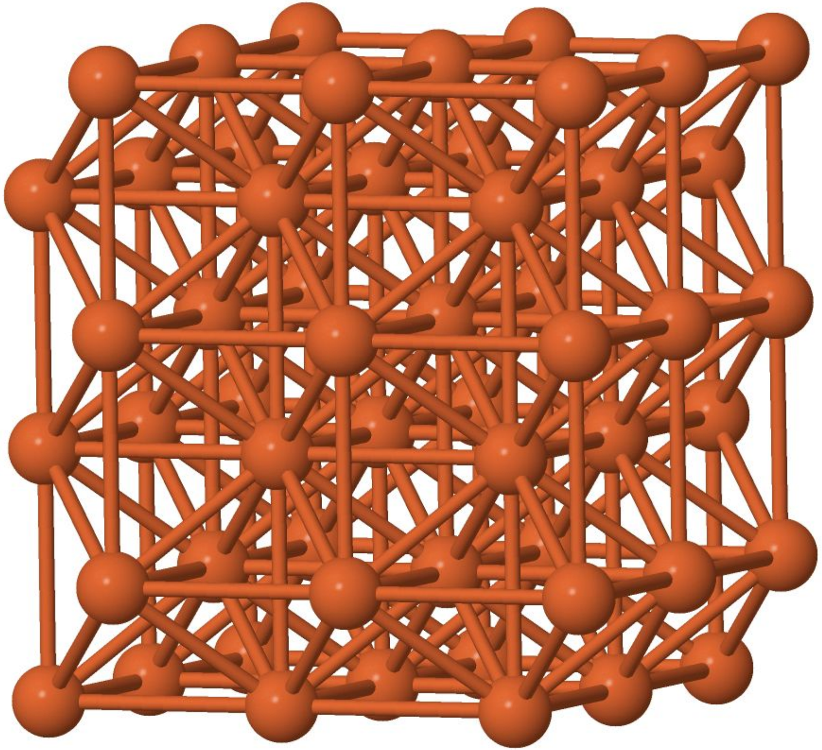

Computational Cost and Crystal Structure Prediction

- How long does this take with ML?

- About 100,000 structures.
- About 30 seconds each.
- Over 34 days on one core...

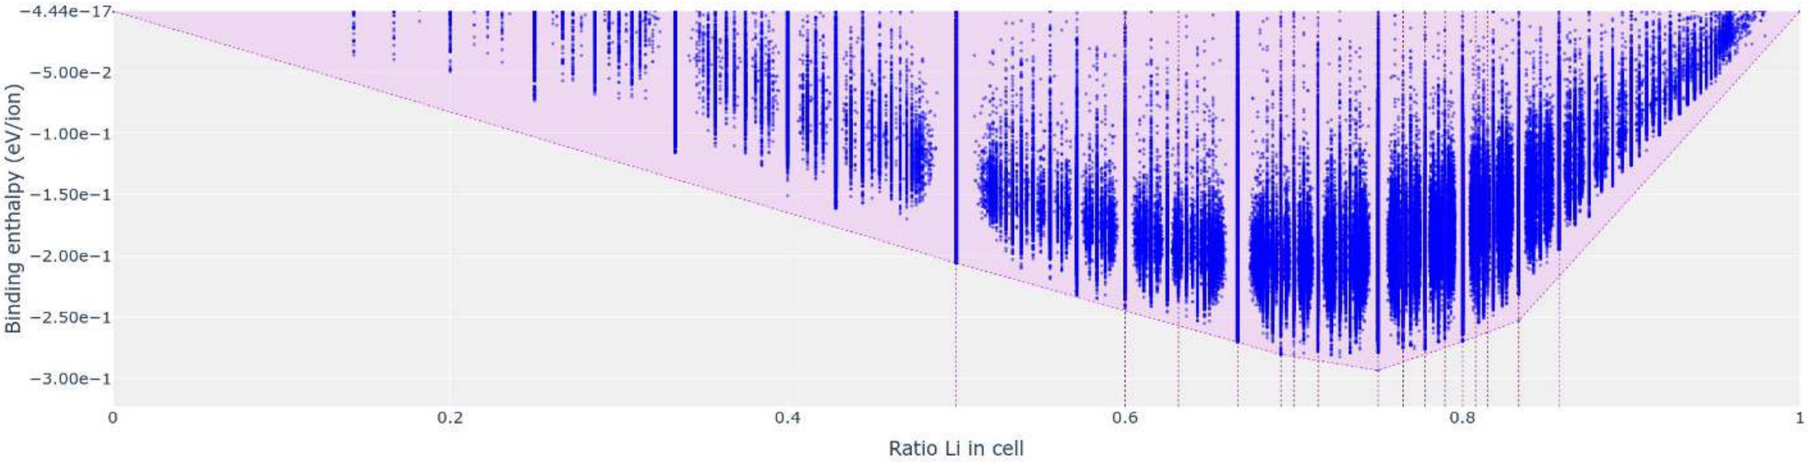

What next and where are we now?

What next?

- Increase network bandwidth
- Other DC uses
- Open OnDemand
- Improving data pipelines and data management
- Improved job usage reporting
- Monitoring
- Sustainability RSEs

Contact emma.barnes@york.ac.uk and richard.fuller@york.ac.uk