

CATS: The Climate Aware Task Scheduler

Colin Sauze and Abhishek Dasgupta

And the rest of the CATS team:

Sadie Bartholomew, Andrew Walker, Loïc Lannelongue, Thibault Lestang, Tony Greenberg, Lincoln Colling, Adam Ward and Carlos Martinez




What?

- Time shifting computing to reduce carbon emissions from electricity.
- Focus on small to mid sized HPC and HTC systems.
 - Tend to be less loaded, near 100% load prevents time shifting.

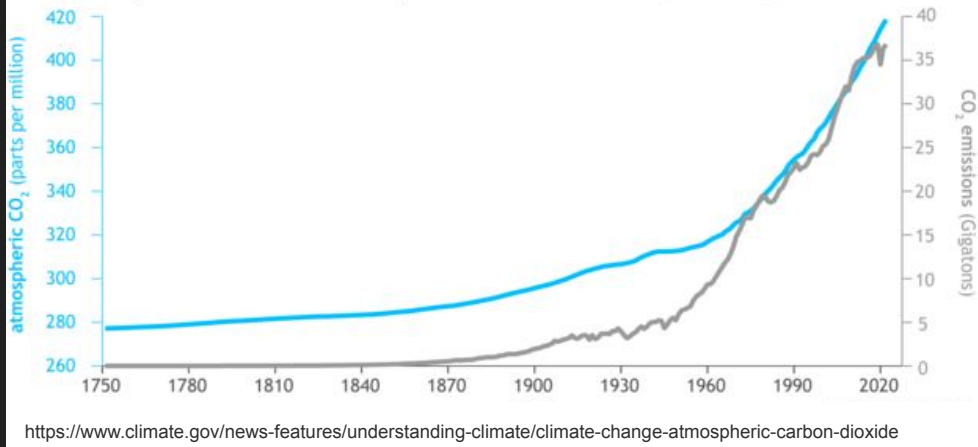


Why?

CARBON DIOXIDE OVER 800,000 YEARS




Global atmospheric carbon dioxide compared to annual emissions (1751-2022)




Global temperature change

Relative to average of 1971-2000 [°C]



DAILY SEA SURFACE TEMPERATURE 60°S–60°N

Data: ERA5 1979–2023 • Credit: C3S/ECMWF



<https://showyourstripes.info/c>

<https://climate.copernicus.eu/2023-track-become-warmest-year-after-record-october>

Our Motivation



Image credits: from IT Crowd (Channel 4),
linked from <https://i.imgur.com/208mpa.jpg>

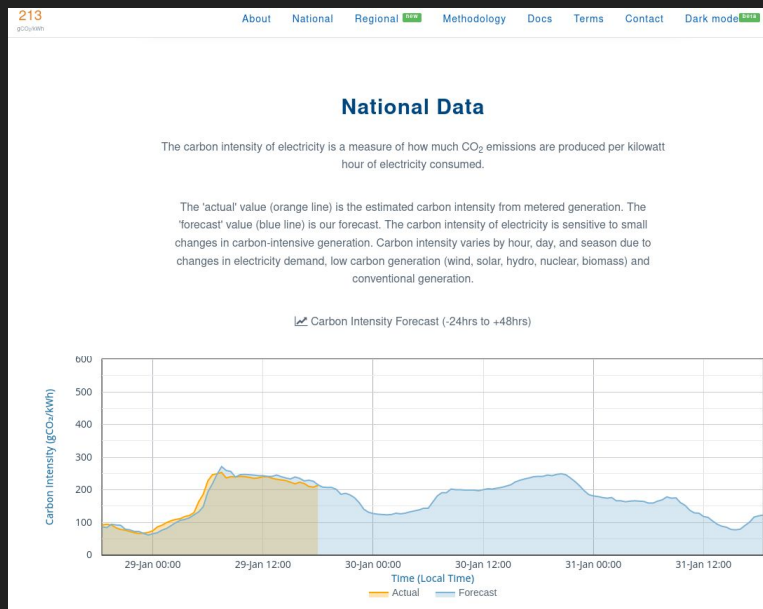
Many of us need to do computationally or data intensive science ...
... but can we do this and not set the world on fire?

How?

- Shift compute to times when electricity has lower carbon intensity
- Carbon Intensity of electricity in the UK is very variable over time
 - Windy and/or sunny = lower carbon
 - Between 0 and 400 g/kWh
 - EU average 251 g/kWh in 2022 (1)
- Regional variations
 - Scotland usually lowest
 - Lots of wind power, some hydro
 - Southern England or London highest
 - Still highly dependent on gas, some wind and solar, not much wind power
- Regions are interconnected
 - But capacity is limited
 - Interconnections to Denmark, France, Ireland, Belgium, the Netherlands and Norway

Carbon Intensity API

- Carbonintensity.org.uk provides regionalised forecasts
 - 30 minute time window
 - 48 hours into the future
 - JSON and XML APIs



Carbon Intensity - Regional ^{beta}

[GET /regional](#) [docs](#)
gets current carbon intensity for GB regions

[GET /regional/england](#) [docs](#)
gets current carbon intensity for England

[GET /regional/scotland](#) [docs](#)
gets current carbon intensity for Scotland

[GET /regional/wales](#) [docs](#)
gets current carbon intensity for Wales

[GET /regional/postcode/{postcode}](#) [docs](#)
gets current carbon intensity for specified outward postcode

[GET /regional/regionid/{regionid}](#) [docs](#)
gets current carbon intensity for specified region

[GET /regional/intensity/{from}/fw24h](#) [docs](#)
gets carbon intensity for 24h after specified datetime for GB regions

[GET /regional/intensity/{from}/fw24h/postcode/{postcode}](#) [docs](#)
gets carbon intensity for 24h after specified datetime for specified outward postcode

[GET /regional/intensity/{from}/fw24h/regionid/{regionid}](#) [docs](#)
gets carbon intensity for 24h after specified datetime for specified region

[GET /regional/intensity/{from}/fw48h](#) [docs](#)
gets carbon intensity for 48h after specified datetime for GB regions

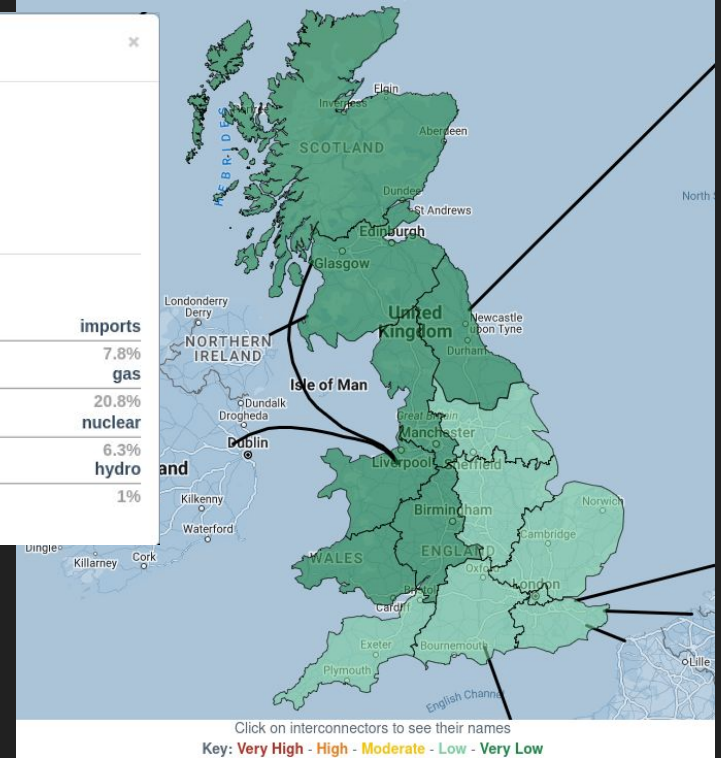

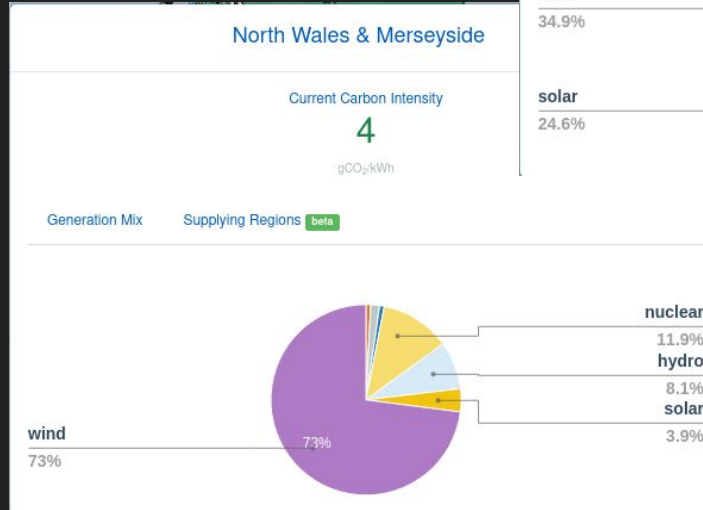
[GET /regional/intensity/{from}/fw48h/postcode/{postcode}](#) [docs](#)
gets carbon intensity for 48h after specified datetime for specified outward postcode

[GET /regional/intensity/{from}/fw48h/regionid/{regionid}](#) [docs](#)
gets carbon intensity for 48h after specified datetime for specified region

[GET /regional/intensity/{from}/pt24h](#) [docs](#)
gets carbon intensity for 24h before specified datetime for GB regions



An example, October 6th 2023

A very windy and somewhat sunny day

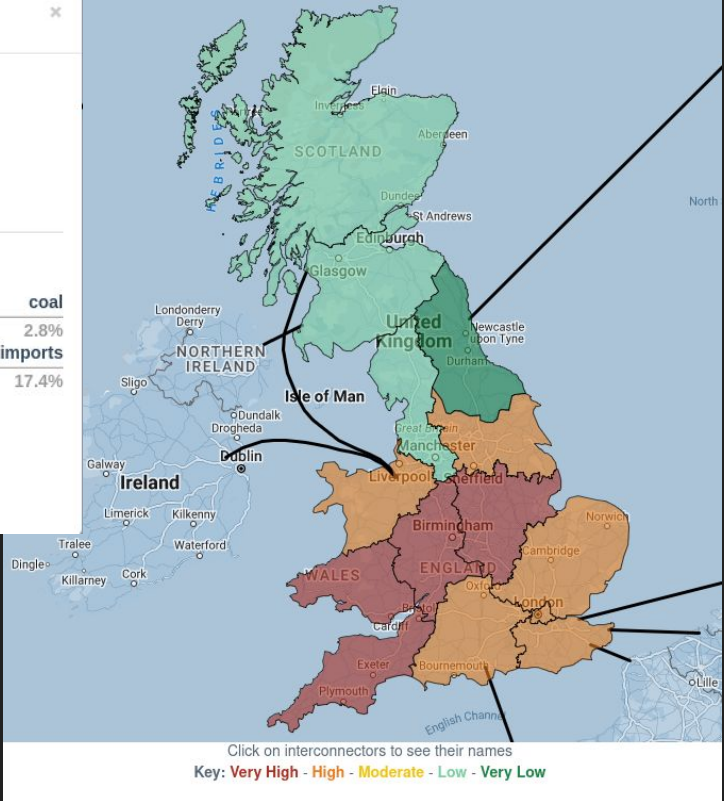


October 9th 2023

58.75 times more carbon dioxide than Oct 6 per kWh in North Wales!



3.14 times more in South England




How much could we save by timeshifting?

- Fictional HPC:
 - 64 core AMD EPYC 7773X (Milan) CPUs
 - 10 nodes, 2 CPUs per node, 20 CPUs total, 1280 cores
 - Fully loaded CPU = 255W, Idle CPU = 37.5W (from <https://www.phoronix.com/review/amd-epyc-7773x-linux/9>)
 - Idle saving = 217.5W per CPU
 - Cluster idle vs peak = 4.35kW
- Time shifting reduces grid intensity from 200g/kWh to 50g/kWh = 150g/kWh reduction
- 12 hour job using all cores
- $12\text{h} * 4.35\text{kW} = 52.2\text{kWh}$
- $52.2\text{kWh} * 0.15\text{kg} = 7.83\text{kg}$
- Comparable to driving an average car (150g/km) 50km (7.5kg)

Why you should timeshift your compute!

- Time shifting your job = not driving 50 km to work
- How many of us don't drive to work for environmental reasons?
- How many employers discourage driving to work for this reason?
- Let's time shift our compute too!

Savings vary across the world



(a) *Flexible Start* optimization for Dense 201.

RESEARCH ARTICLE OPEN ACCESS

Measuring the Carbon Intensity of AI in Cloud Instances

Authors: Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz,

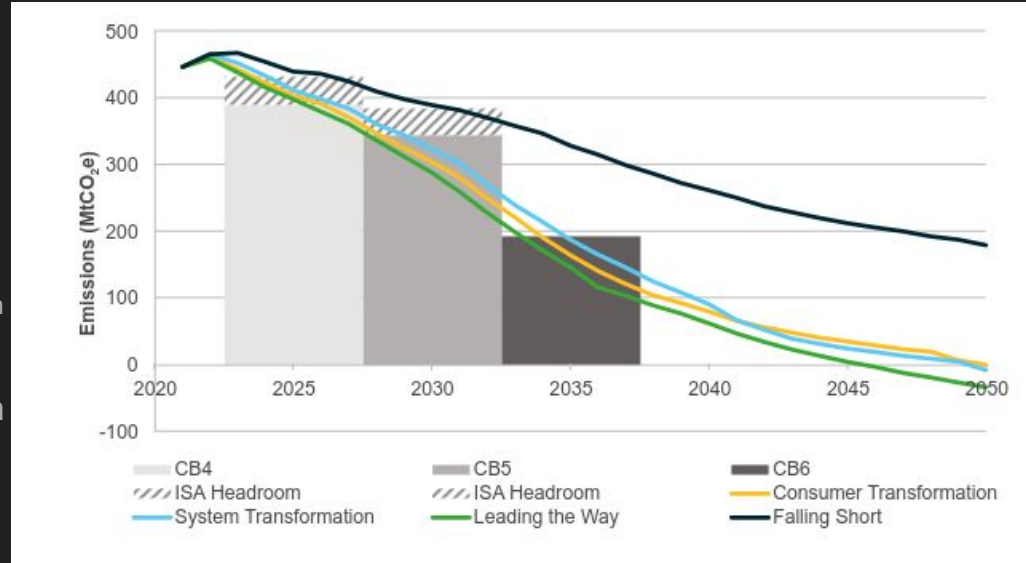
Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCarlo, Will Buchanan [Authors Info & Claims](#)

FACCT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency • June 2022 • Pages 1877–1894 • <https://doi.org/10.1145/3531146.3533234>

<https://doi.org/10.1145/3531146.3533234>

But won't the grid be net zero soon anyway?

- Maybe?
 - UK Grid 495 gCO₂/kWh in 2008, 155 gCO₂/kWh in 2023
 - Planned to be net-zero by 2035
- If we can do something now, then why wait?
- Getting zero emission grid on cold, dark, calm winter days will be difficult.
 - Expensive energy storage, synthetic fuels or carbon capture needed
 - Reducing peak demand will help
- Increasing number of days with overproduction
- Time shifting useful if you generate your own power too (e.g. rooftop solar)
- Variable rate electricity tariffs, save money as well as carbon.




From the National Grid Future Energy Scenarios in 5
<https://www.nationalgrideso.com/document/283056/download>

Introducing the climate-aware task scheduler

1. Users submit a runtime and postcode

2. CATS figures out the best time to start the job



TADAA, you've reduced your carbon footprint by **70%**

CATS proof of concept

- Built in a day at the Software Sustainability Institute's Collaborations Workshop Hackday 2023
 - Won 1st prize!
- Python script targeting the classical UNIX 'at' scheduler
- Intended as proof of concept for users who do 'small scale' computing (e.g. a few hours on a workstation overnight)



Limitations of CATS

- Won't be able to do much on systems at/near 100% load
- Relies on user specifying the job length correctly
- Another user might run during the time we're avoiding
- Only works in the UK right now
 - Other countries don't have freely available regional carbon intensity data or carbon intensity forecasts
 - Links to APIs to enable this are welcome (<https://github.com/GreenScheduler/cats/issues/22>)
- Not the only thing you can/should do to reduce the climate impact of your HPC
 - Lots of emissions from scope 3 (manufacturing)
 - Cooling
 - Storage
 - Networks

Using CATS

```
python -m cats -d <job_duration> --loc <postcode>
```

- Postcode is a proxy for location, can also be specified in a config.yml file.
- Will return the time to start a job in an At friendly format (yyyymmddhhMM) to stdout
 - Additional info returned in JSON format to stderr
- Pipe a command to At to run it and substitute the -t option for the output from CATS, for example to run ls:

```
ls | at -t `python -m cats -d 5 --loc OX1`
```


Estimating Carbon Used by a Job

We can get estimates on how much a job used


Providing we supply:

- PUE (Power Use Efficiency, the amount of power that went on compute vs cooling)
- CPU/GPU type and Thermal Design Power
- Job run time

```
location: "EH8"
api: "carbonintensity.org.uk"
PUE: 1.20 # > 1
partitions:
  CPU_partition:
    type: CPU # CPU or GPU
    model: "Xeon Gold 6142"
    TDP: 9.4 # Thermal Design Power in W/core
  GPU_partition:
    type: GPU
    model: "NVIDIA A100-SXM-80GB GPUs"
    TDP: 300
    CPU_model: "AMD EPYC 7763"
    TDP_CPU: 4.4
```

Demo Video

<https://github.com/GreenScheduler/cats/blob/main/cats.gif?raw=true>



Next Steps

- Version 1.0 release this month
 - Cleaning up the command line options
- Slurm Integration
 - Simplest method: use sbatch to offset start time
 - Green queues
 - Integrating carbon accounting as a Slurm plugin (will need rewrite in C)
 - Funding from the Software Sustainability Institute for a few months of developer time
- We plan to test CATS on Slurm on some real HPCs this summer

Any questions?

CATS Github: <https://github.com/GreenScheduler/cats>

Get in contact:

Colin Sauze - colin.sauze@noc.ac.uk

Abhishek Dasgupta - abhishek.dasgupta@dtc.ox.ac.uk

