Heidelberg Institute for Theoretical Studies

GSI ITEE-Palaver

EcoFreq: compute with cleaner energy via carbon-aware power scaling

Alexey Kozlov

Staff scientist

CME group, HITS gGmbH

alexey.kozlov@h-its.org

May 24, 2022

Top500 power trend

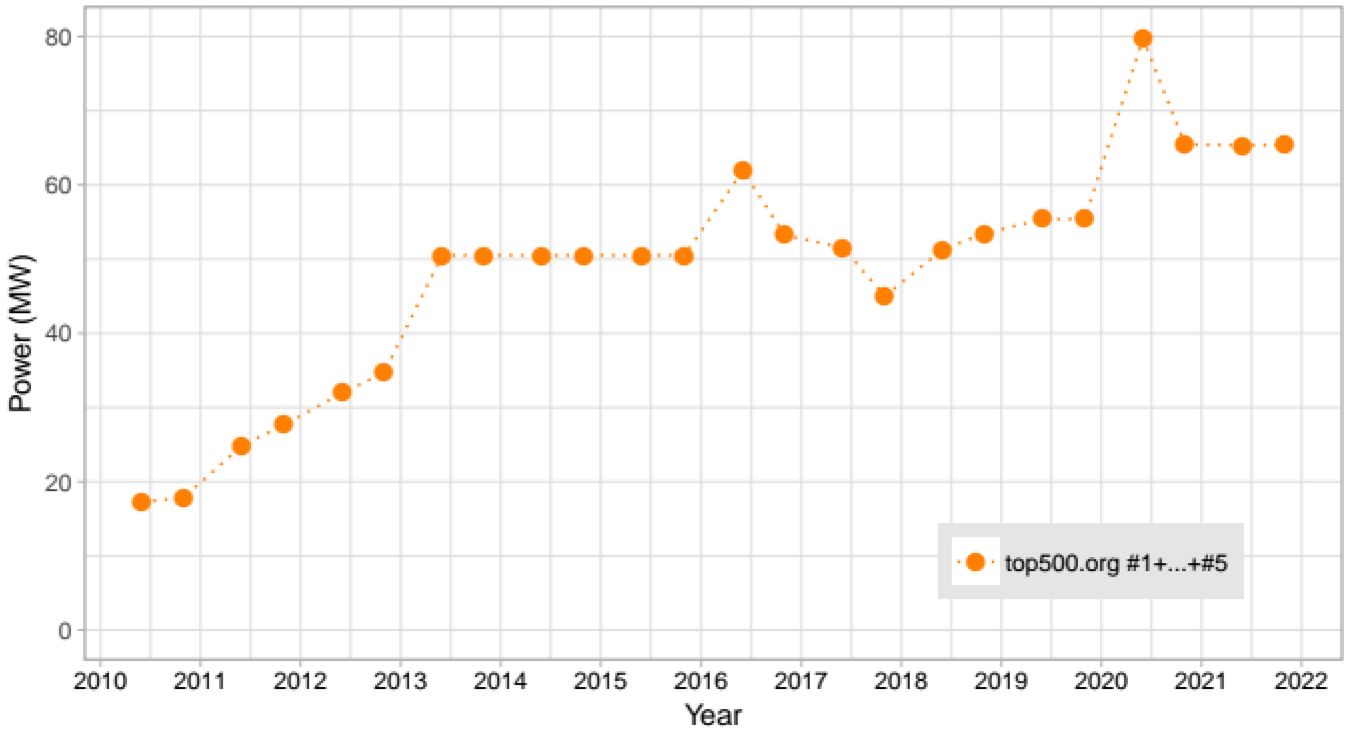

Faster hardware+software != energy savings

Efficient vs. Green

- Energy efficiency → relative improvement
  - More FLOPS/W, FLOPS/$, FLOPS/rack, ...
  - Battery life, range, ...
- GreenIT → absolute impact
  - Negative “side effects”
  - Environment, climate, society
- 0% < overlap < 100%

Electricity sources

Technology	Risks		
	Climate	Environmental	Political
Fossil (coal, natural gas)	–	–	–
Nuclear	+	+ / –	–
Renewables	+	+	+

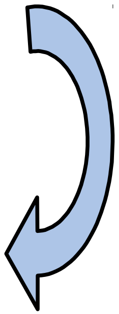

Energy mix DE (2021)

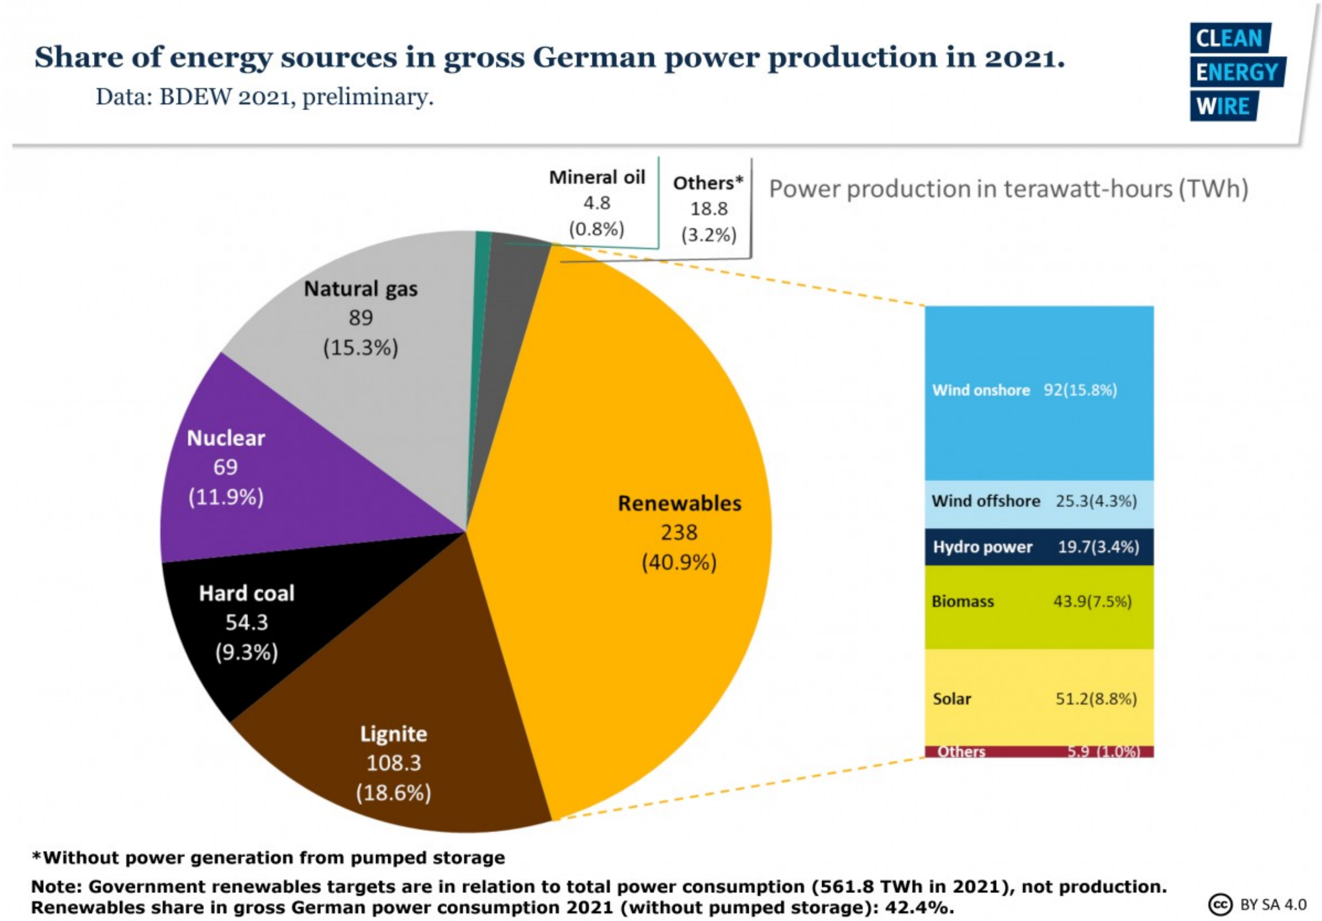

2014 IPCC, Global warming potential of s

Life cycle CO_2 equivalent (including albedo effect) from selected electricity supply technologies.[2][3] Arranged by decreasing median (gCO_2eq/kWh) values.

Technology	Min.	Median	Max.
Currently commercially available technologies			
Coal – PC	740	820	910
Biomass – Cofiring with coal	620	740	890
Gas – combined cycle	410	490	650
Biomass – Dedicated	130	230	420
Solar PV – Utility scale	18	48	180
Solar PV – rooftop	26	41	60
Geothermal	6.0	38	79
Concentrated solar power	8.8	27	63
Hydropower	1.0	24	2200[1]
Wind Offshore	8.0	12	35
Nuclear	3.7	12	110
Wind Onshore	7.0	11	56

Avg. carbon intensity:
~350 g CO_2eq/kWh

Real-time energy mix

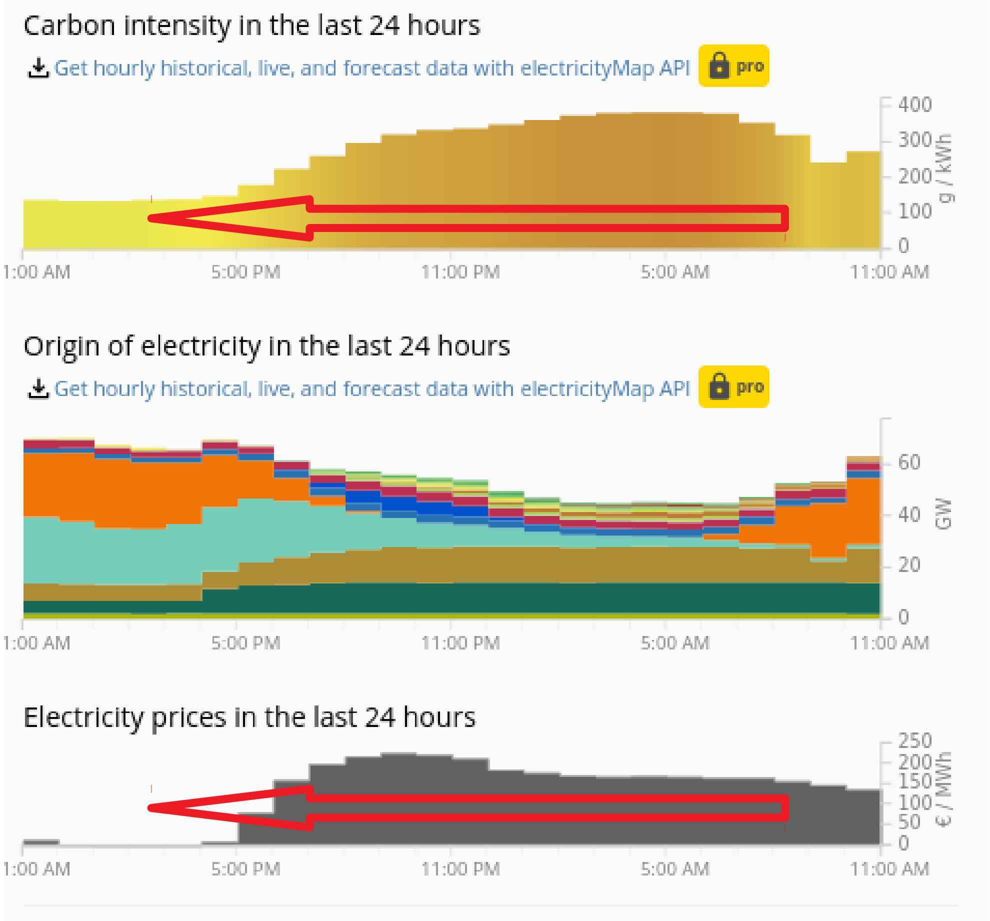

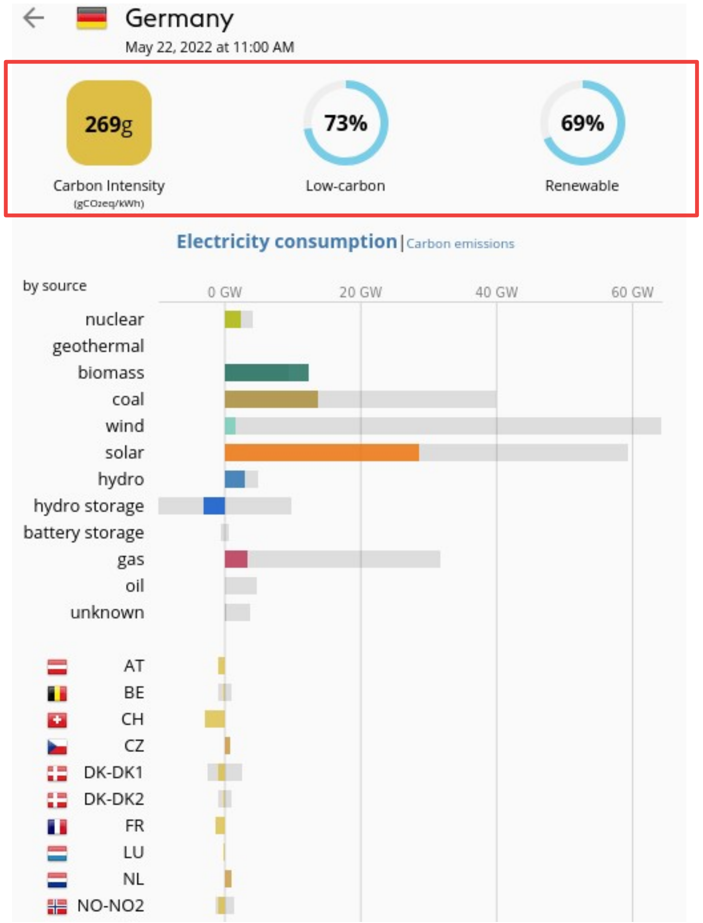

Compute with cleaner energy

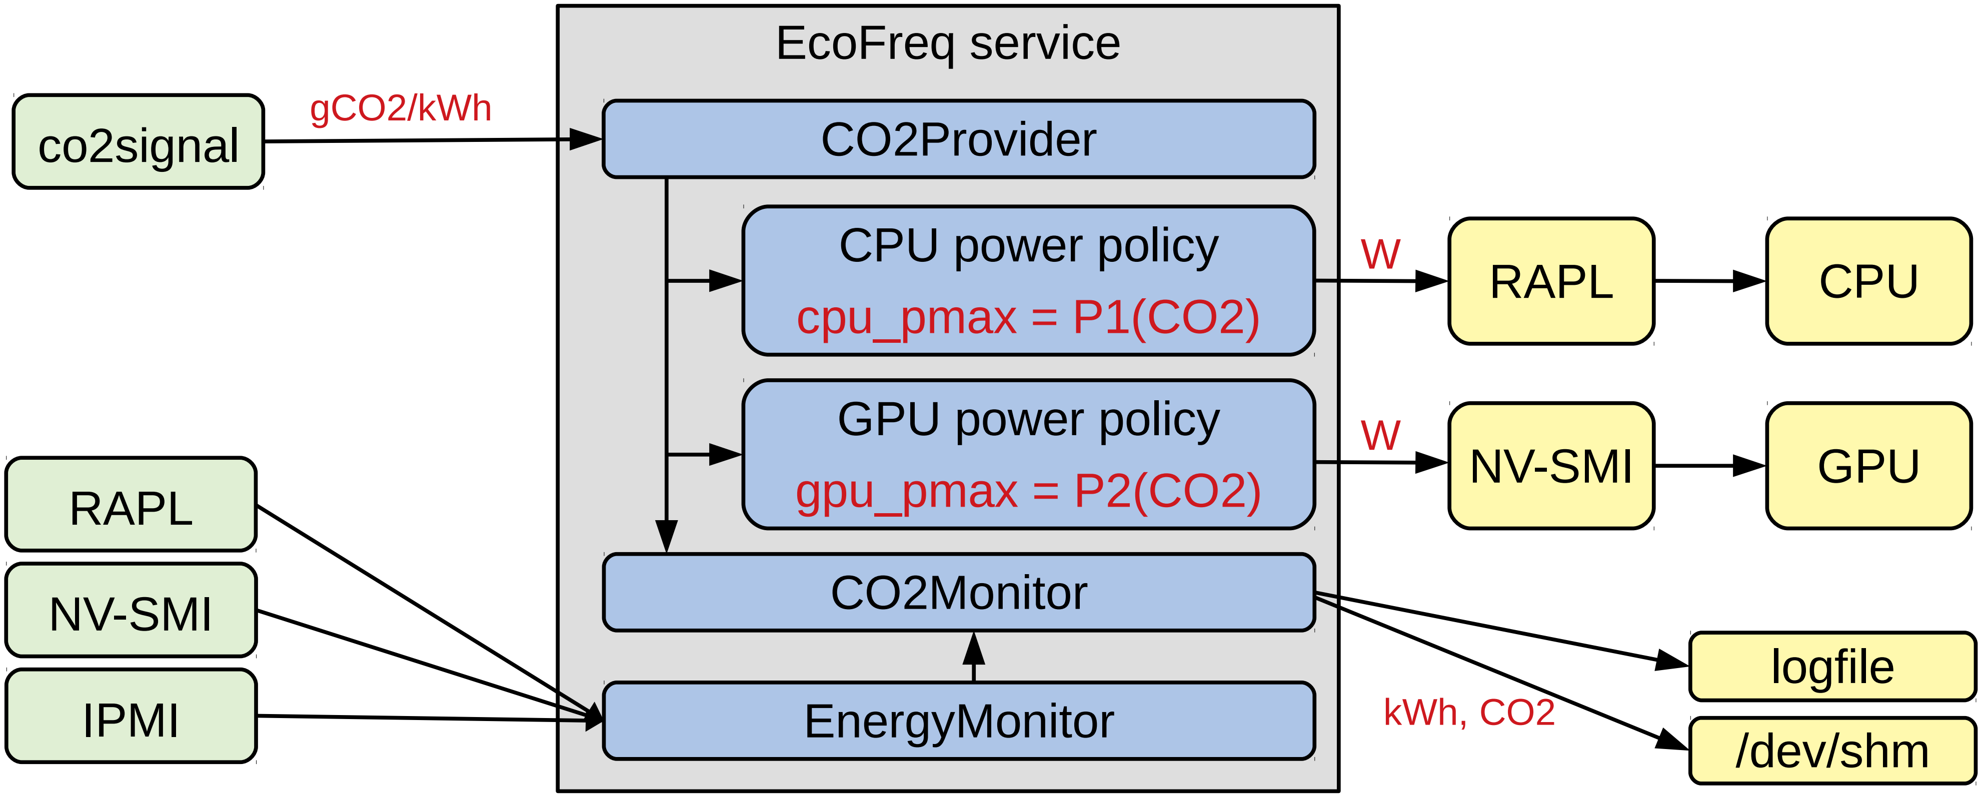

- Proof-of-concept implementation: https://github.com/amkozlov/eco-freq

Input / optimization target

- Life-cycle CO_2 → default
- Marginal CO_2
- Renewable / low-carbon %
- Electricity cost

Carbon intensity data

- co2signal / electricityMap
  - 50+ countries worldwide
  - Real-time carbon intensity: free
  - Marginal, prices, historical data: paid service
- ENTSO-E
  - Europe/EU
  - Generation, load, prices – history and forecast
  - Free

Power scaling: techniques

- RAPL / DVFS
  - Dynamic power / frequency limits → ~50% - 100% TDP
  - Supported by most CPUs/GPUs (Intel, AMD, NVIDIA)
- Utilization capping
  - e.g. cpulimit: https://github.com/opsengine/cpulimit
- Adaptive parallelization
  - Adjust # threads / MPI ranks

Power scaling: advantages

- Transparent to the workload
  - No profiling, recompilation etc.
  - Long jobs are fine (no interruption / restart)
- Also works without job queue / scheduler
- No generation forecast needed
  - But can be used if available

Power scaling policy

- Linear (default so far)
  - Pmax ~ 1/CI, where CI = current carbon intensity
- Static / manual
  - e.g. 80% by default → usually best energy efficiency
  - Increase if: solar/wind surplus, low price, urgent job...
- Hardware- and workload-aware → future work

Evaluation metrics

- CO_2-to-solution (cf. energy-to-solution)
  - Job perspective $\rightarrow \Sigma CO_2$ over program runtime
- CO_2-per-hour (-month/-year)
  - System-wide perspective
- CO_2-to-solution ~ CO_2-per-hour iff
  - Utilization $\rightarrow$ 0%
  - Idle power = 0

} rather unrealistic!

EcoFreq: Evaluation

- Linear policy, no HW/SW profiling, no forecast

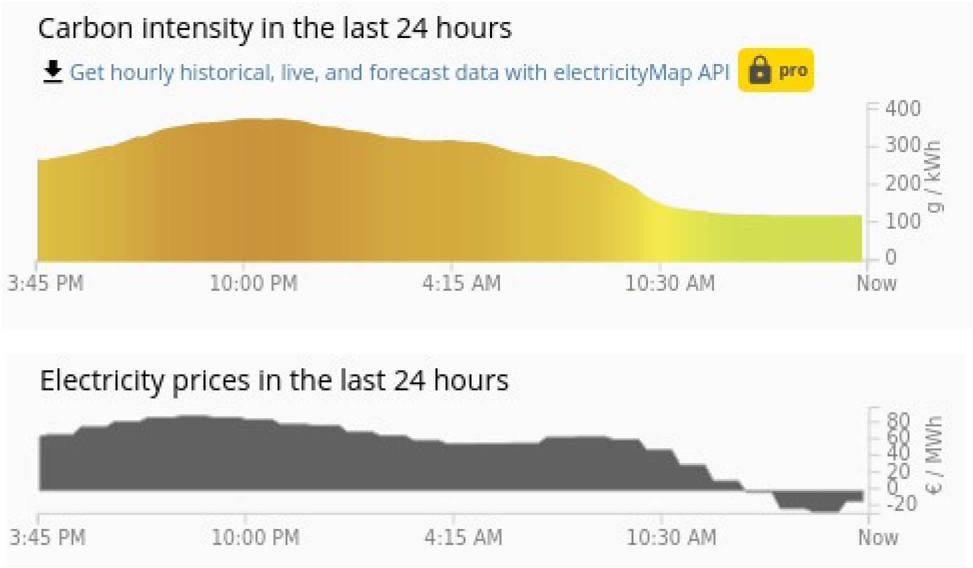

Germany, June 11-12, 2021

RAxML-NG v1.0.2 @ Intel Xeon Platinum 8260, 48C			
	Baseline	EcoFreq	Diff. %
Time [s]	58502	62257	+6.4 %
Energy [kWh]	9.533	9.082	-4.7 %
CO_2-to-solution [g]	2590	2307	-10.9 %
CO_2-per-hour [g] (100% utilization)	153	133	-13.1 %

EcoFreq: Demo

```
$ sudo ./ecofreq.py

#Timestamp              gCO2/kWh    CPU_Pmax [W]   GPU_Pmax [W]   SYS_Pavg [W]   Energy [J]     CO2 [g]
2021-06-11T23:14:18          380         223.000             NA        398.650   358785.000      37.874
…
2021-06-12T09:44:48          262         275.750             NA        529.639   476675.000      34.632
…
2021-06-12T13:44:59          133         330.000             NA        594.097   534687.500      19.778
```

```
$ ./ecostat.py
EcoStat v0.0.1

Loading data from log file: /var/log/ecofreq.log

Time interval:                 2021-05-18 - 2021-06-11
Duration active:               23 days, 23:15:57
Duration inactive:             17:06:18
CO2 intensity range [g/kWh]:   108 - 449
CO2 intensity mean [g/kWh]:    284
Energy consumed [J]:           657629645.0
Energy consumed [kWh]:         182.675
CO2 emitted [kg]:              51.701283
```

```
$ ./ecorun.py -p linear raxml-ng

[...]


time_s:       882.708
pwr_avg_w:    553.422
energy_j:     488510.0
energy_kwh:   0.136
co2_g:        51.112
```

Next steps

- Hardware- and workload-aware policies
- SLURM plugin
- Evaluation on a larger system → GSI?

Discussion