5. Green Software Metrics

Sustainable Software Engineering
CS4575

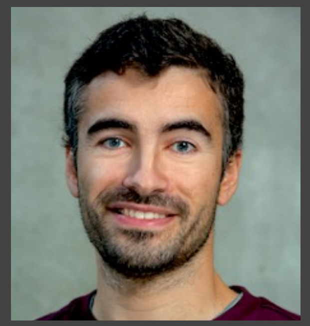

Luís Cruz
L.Cruz@tudelft.nl

Carolin Brandt
C.E.Brandt@tudelft.nl

Enrique Barba Roque
E.BarbaRoque@tudelft.nl

Bitcoin example

- **1 bitcoin** transaction is equivalent to more than **1.5 million VISA** transactions.

- Day-to-day metrics are easy to grasp
    - If we say **8 gigajoules**, it’s a bit more difficult to understand.
- These numbers keep changing (check it here: https://www.statista.com/statistics/881541/bitcoin-energy-consumption-transaction-comparison-visa/)

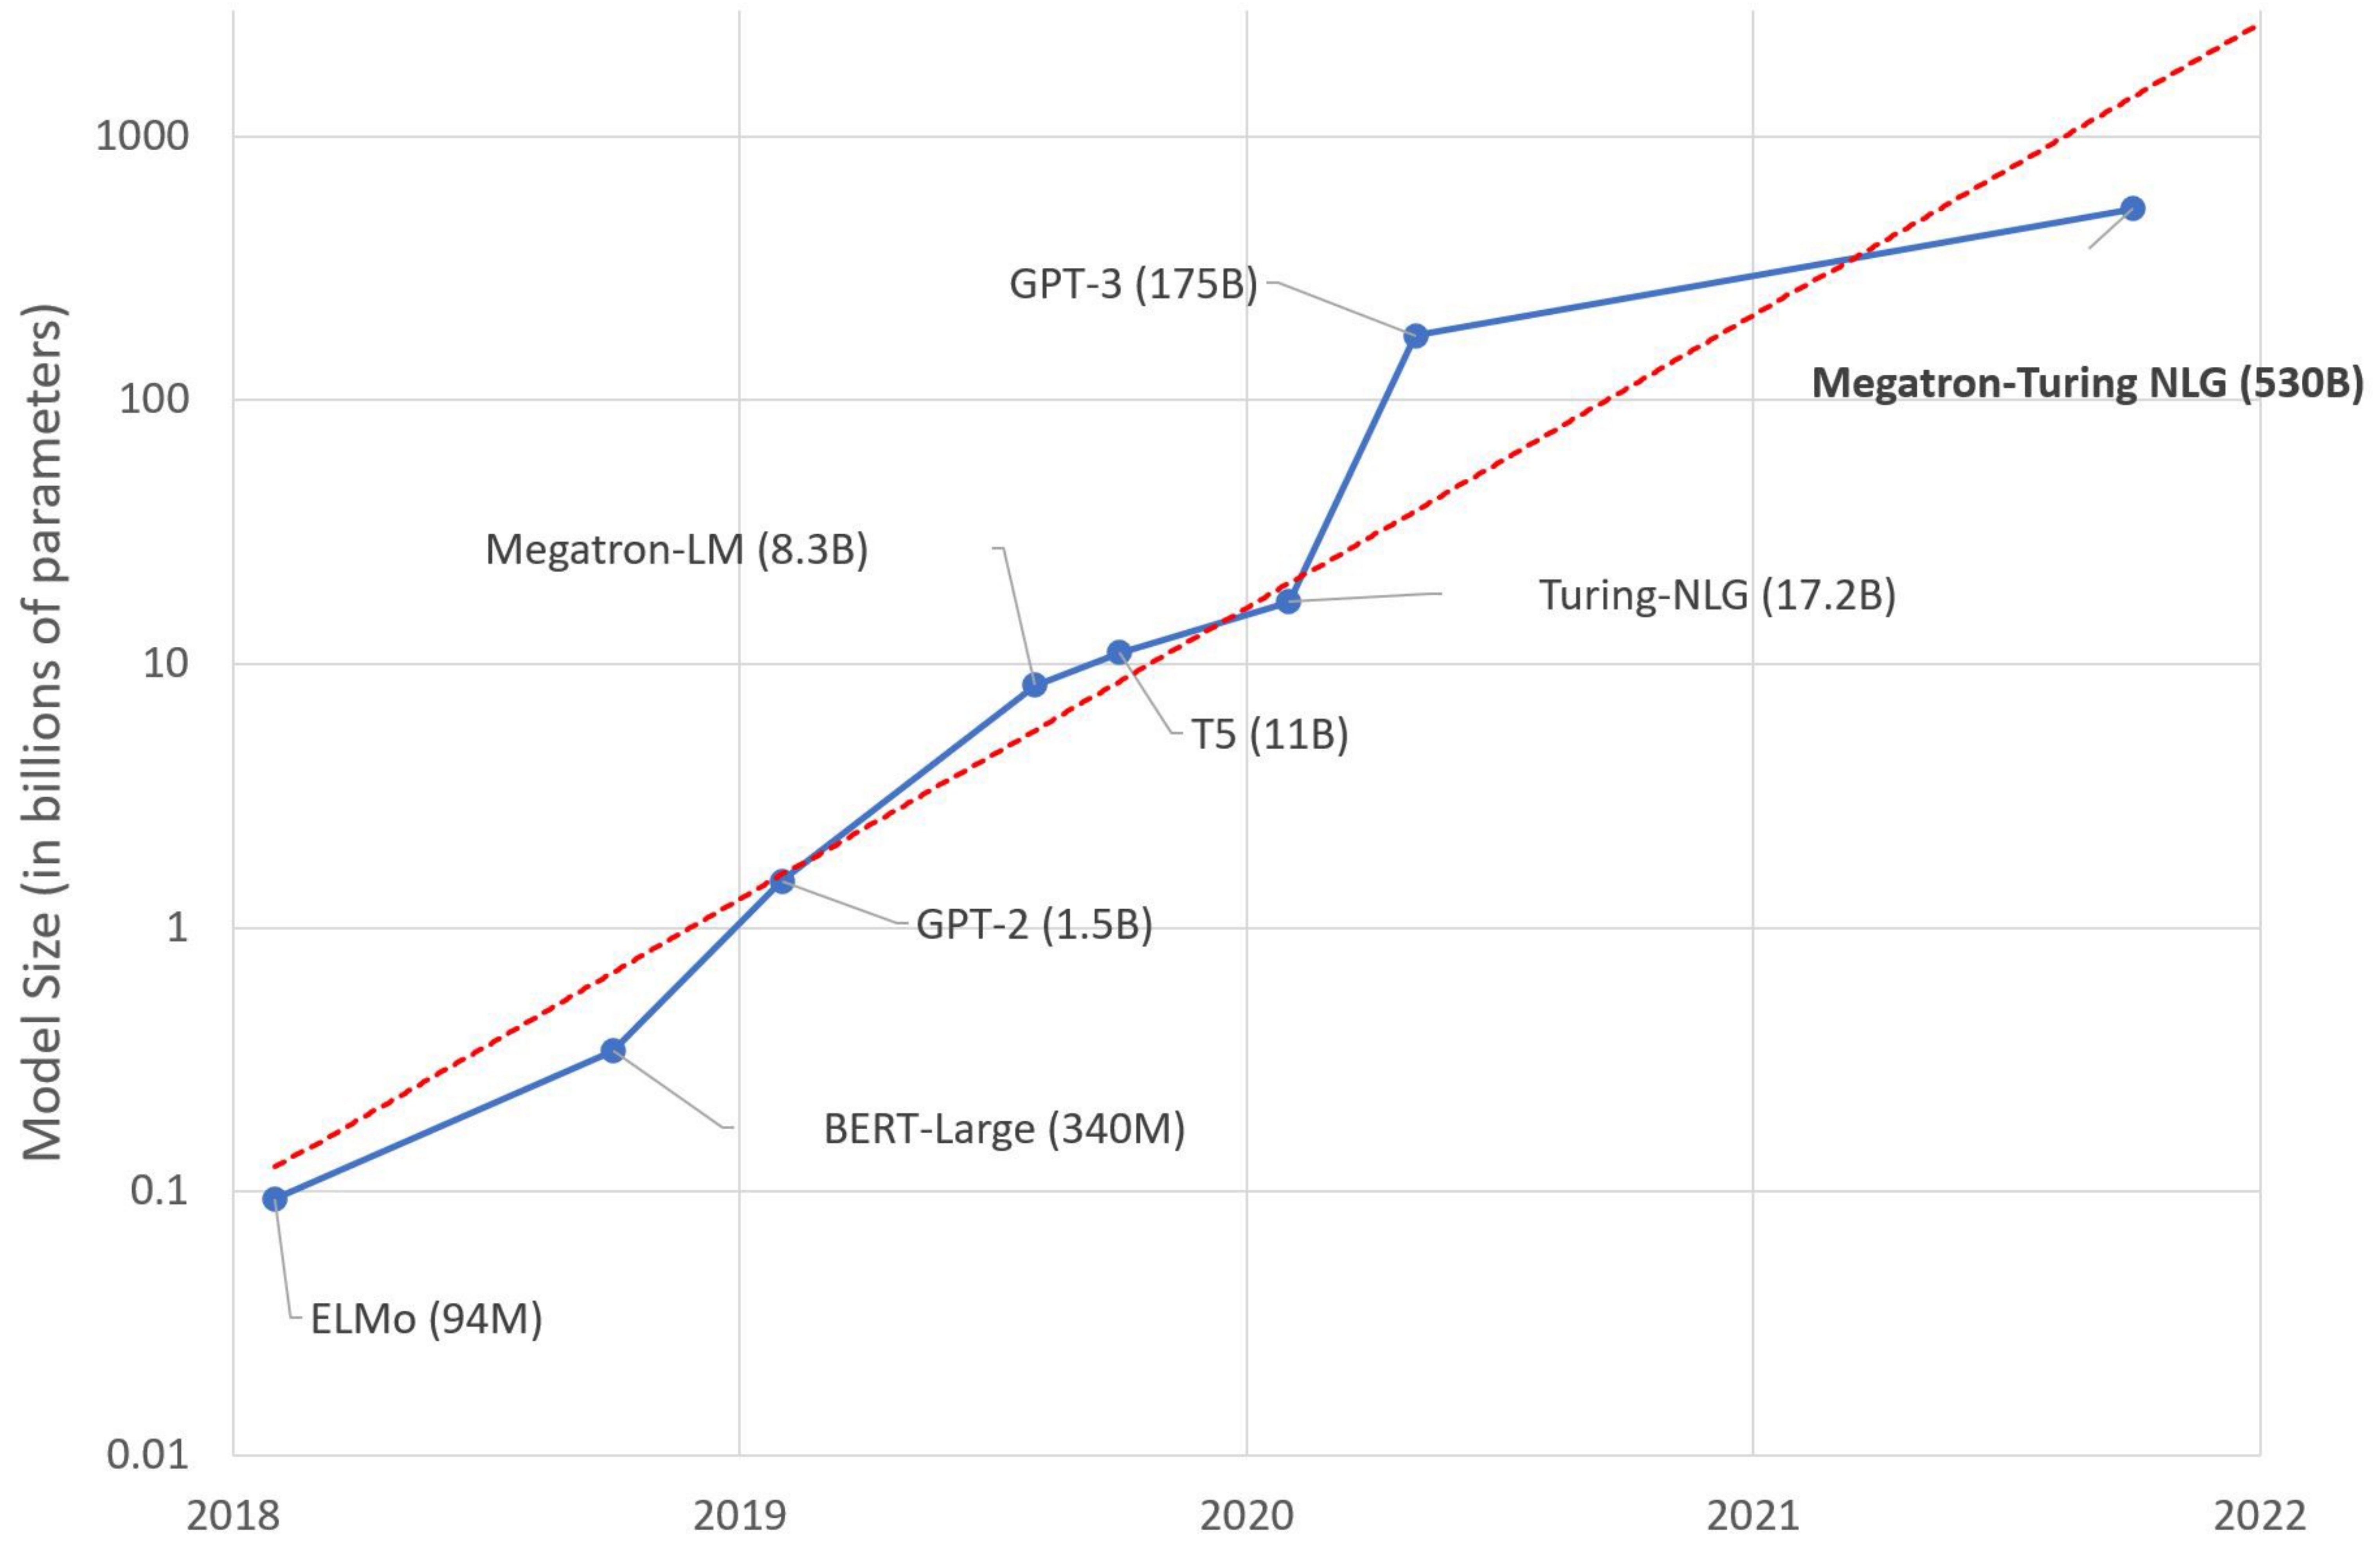
Model Size (in billions of parameters)
1000
100
10
1
0.1
0.01
2018
2019
2020
2021
2022
GPT-3 (175B)
Megatron-Turing NLG (530B)
Megatron-LM (8.3B)
Turing-NLG (17.2B)
T5 (11B)
GPT-2 (1.5B)
BERT-Large (340M)
ELMo (94M)

(Electrical) Energy

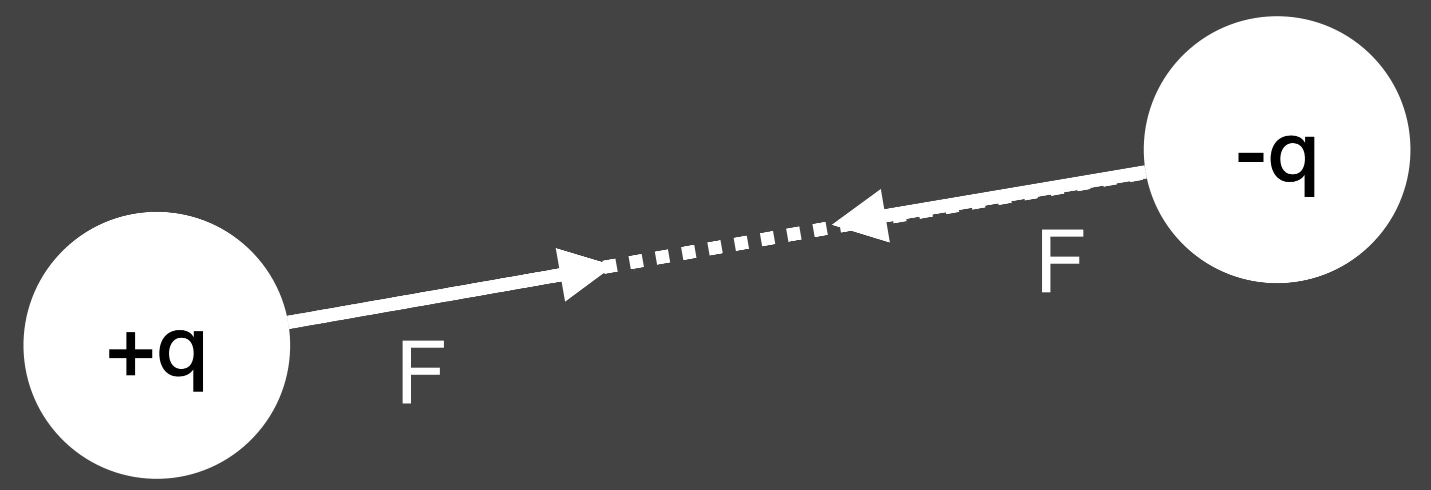

- **Work** required to move charged particles.
  - Same concept but different perspective when talking about **thermal**, **mechanical**, or **nuclear** energy.
- Most common units:
  - **joule (J)** - recommended; scientific communications; metric from the International System of Units
  - **kilowatt-hour (kWh)** - more common, e.g., used for household electricity consumption

Power

- Amount of **work** being done **per unit of time**.
  - Commonly measured in **watts (W)**.

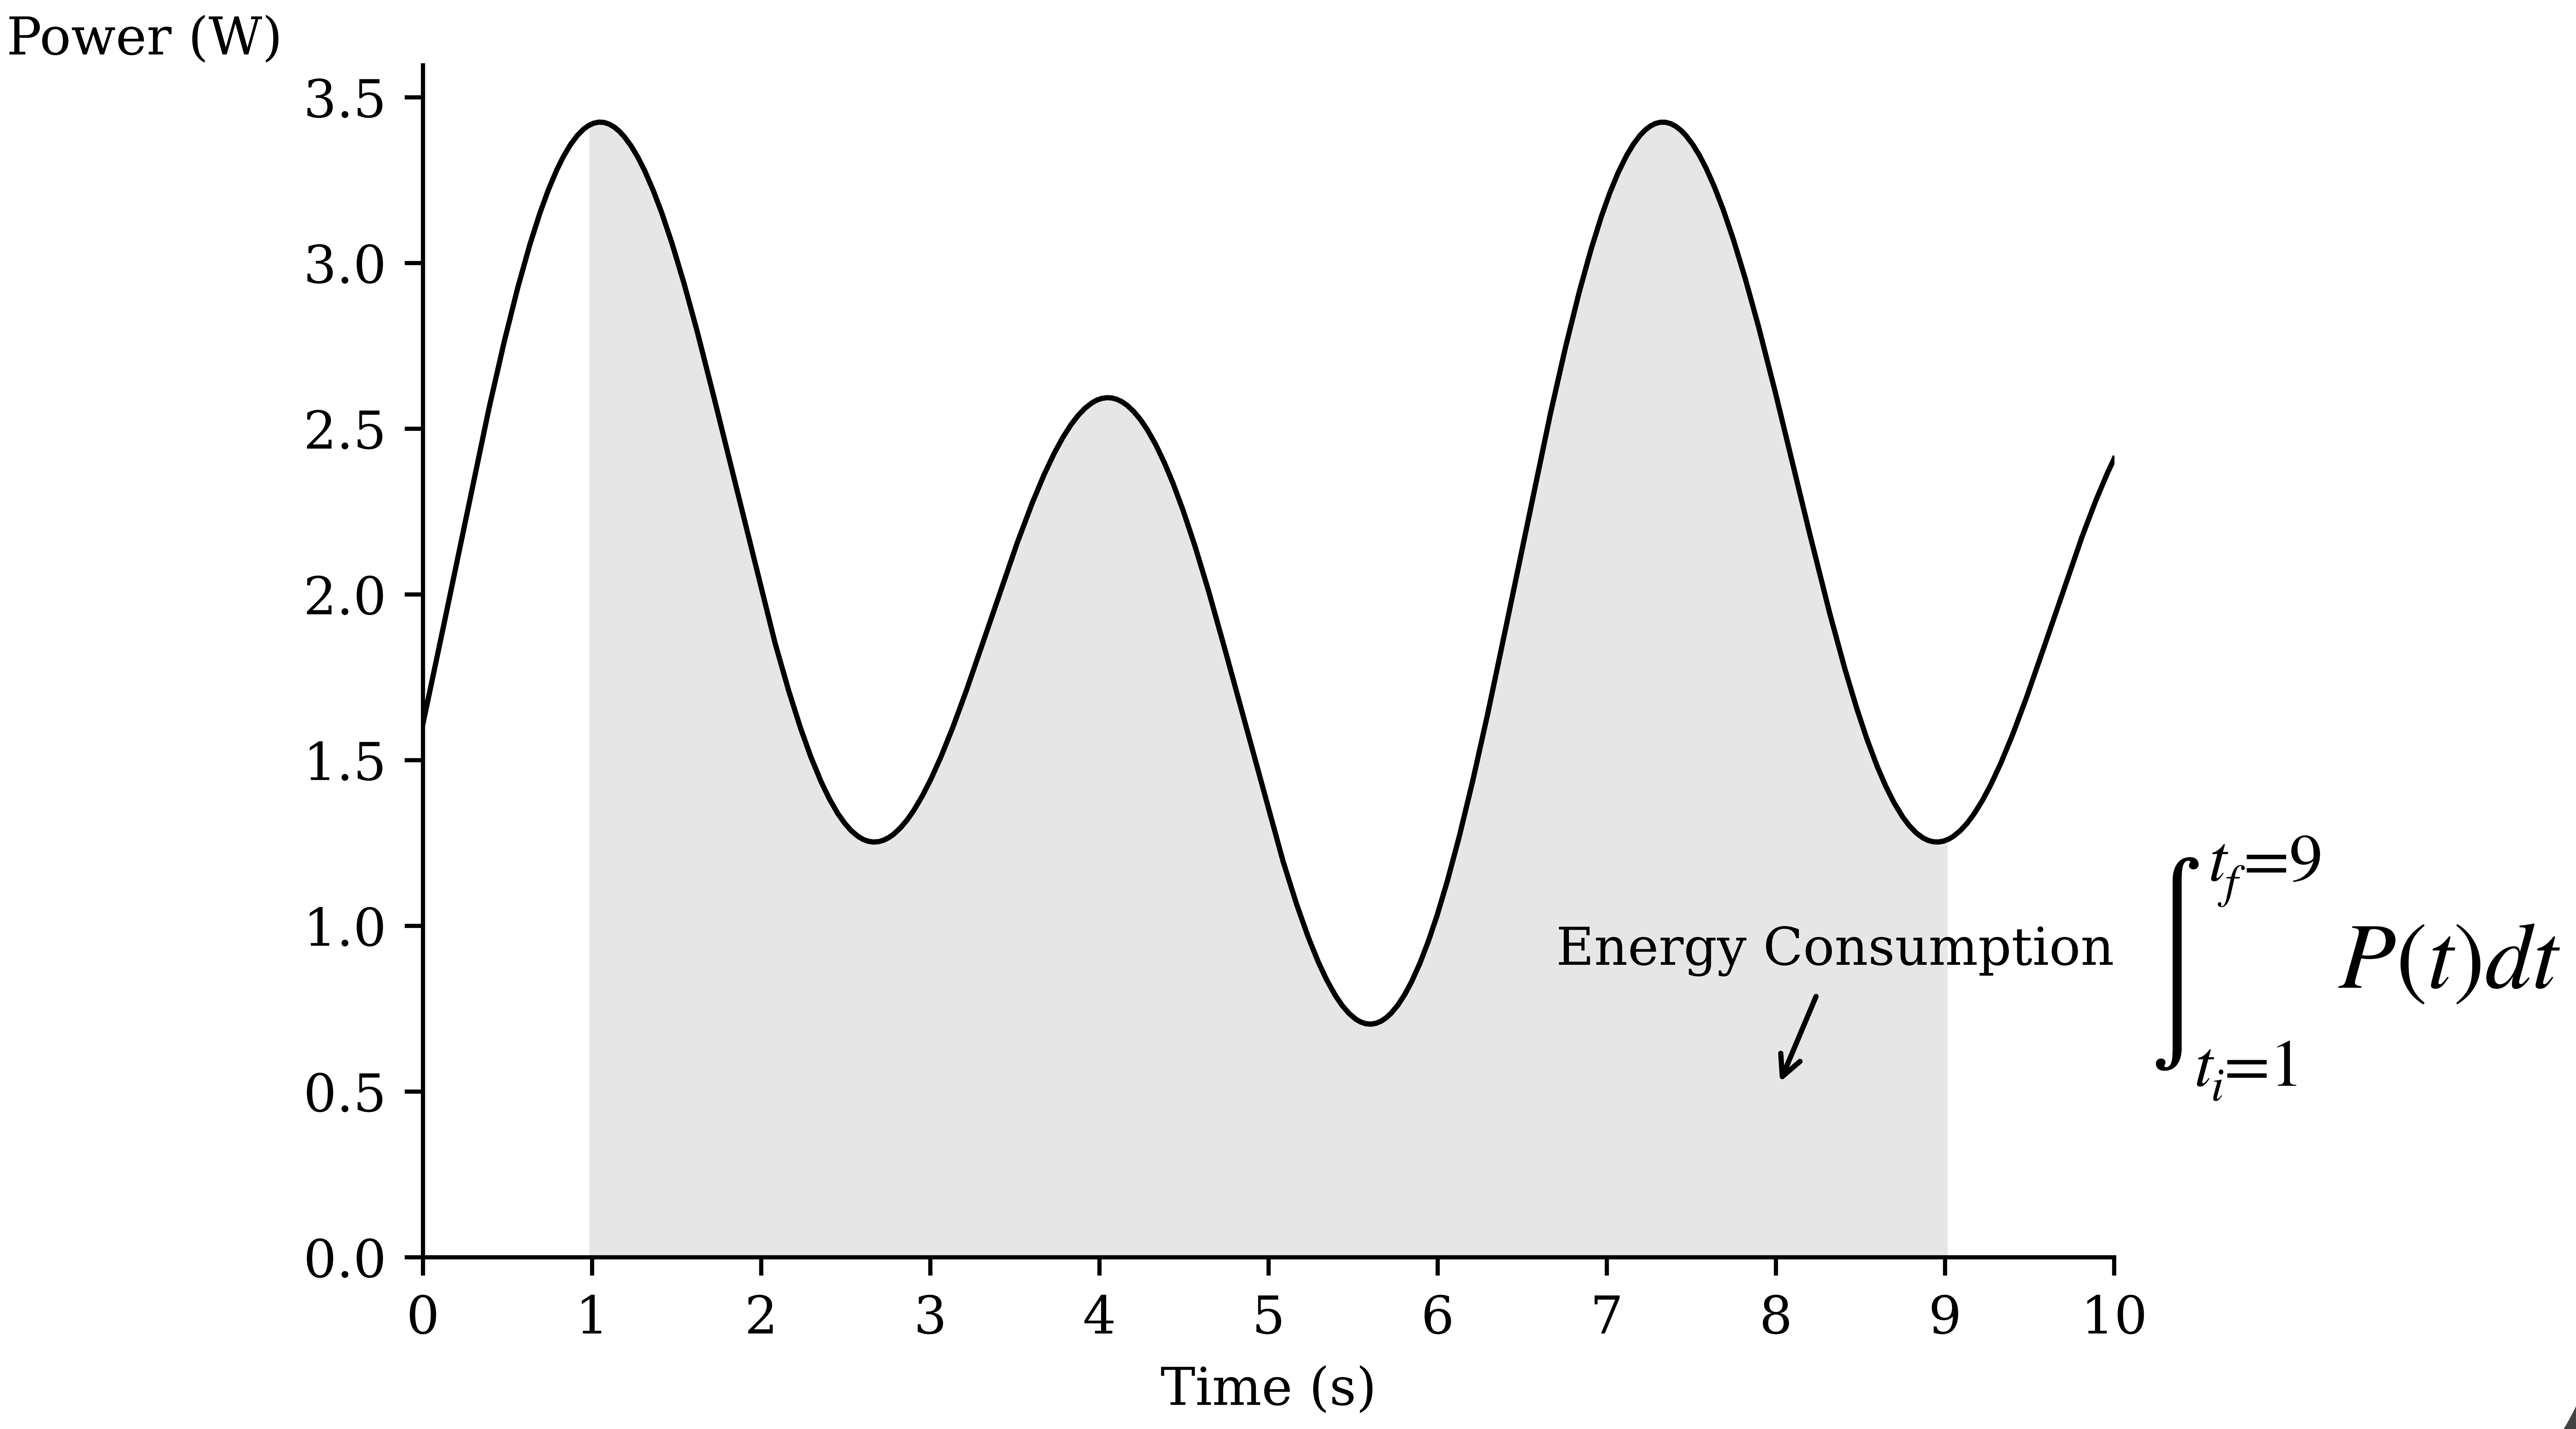

Power (W)
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
0
1
2
3
4
5
6
7
8
9
10
Time (s)
Energy Consumption
$\int_{t_i=1}^{t_f=9} P(t)dt$

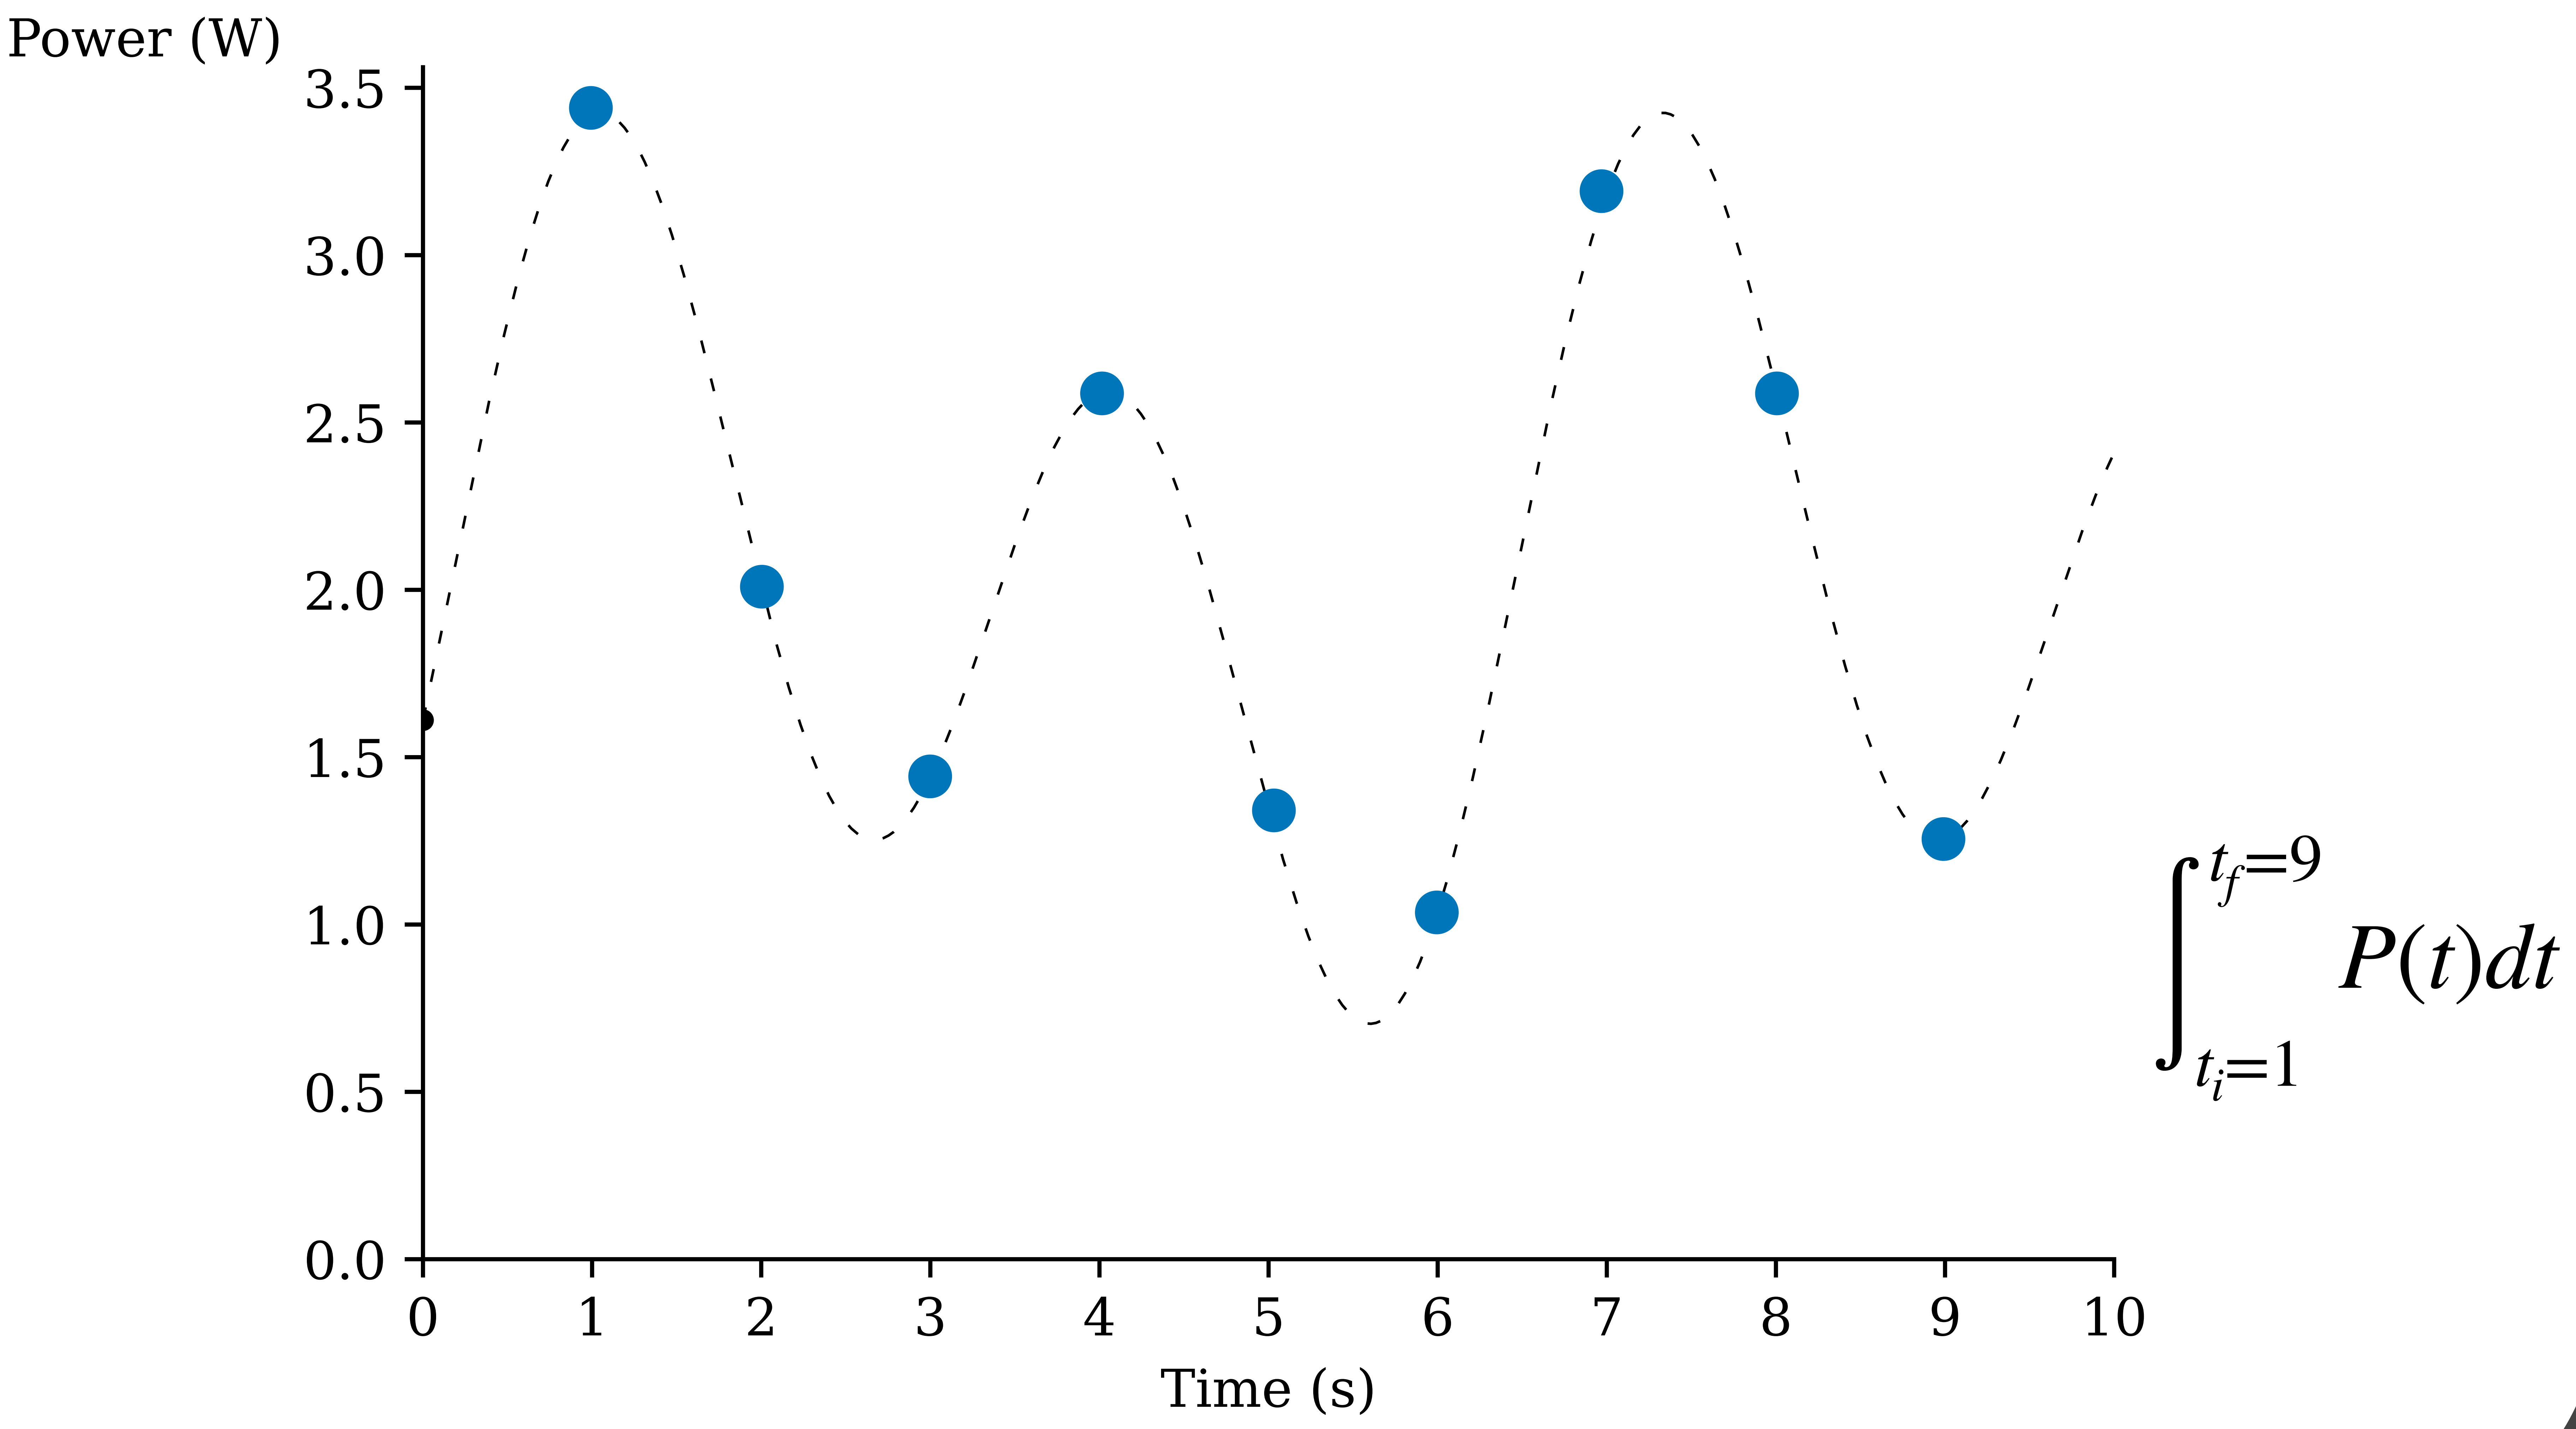
Power (W)
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
0
1
2
3
4
5
6
7
8
9
10
Time (s)
$\int_{t_i=1}^{t_f=9} P(t)dt$

$$A = \frac{a+b}{2}h \quad = \frac{P_{t_n} + P_{t_{n+1}}}{2} \cdot \Delta t \; = \frac{3.5\text{W} + 2.0\text{W}}{2} \cdot 1\text{s} = 2.75\text{J}$$

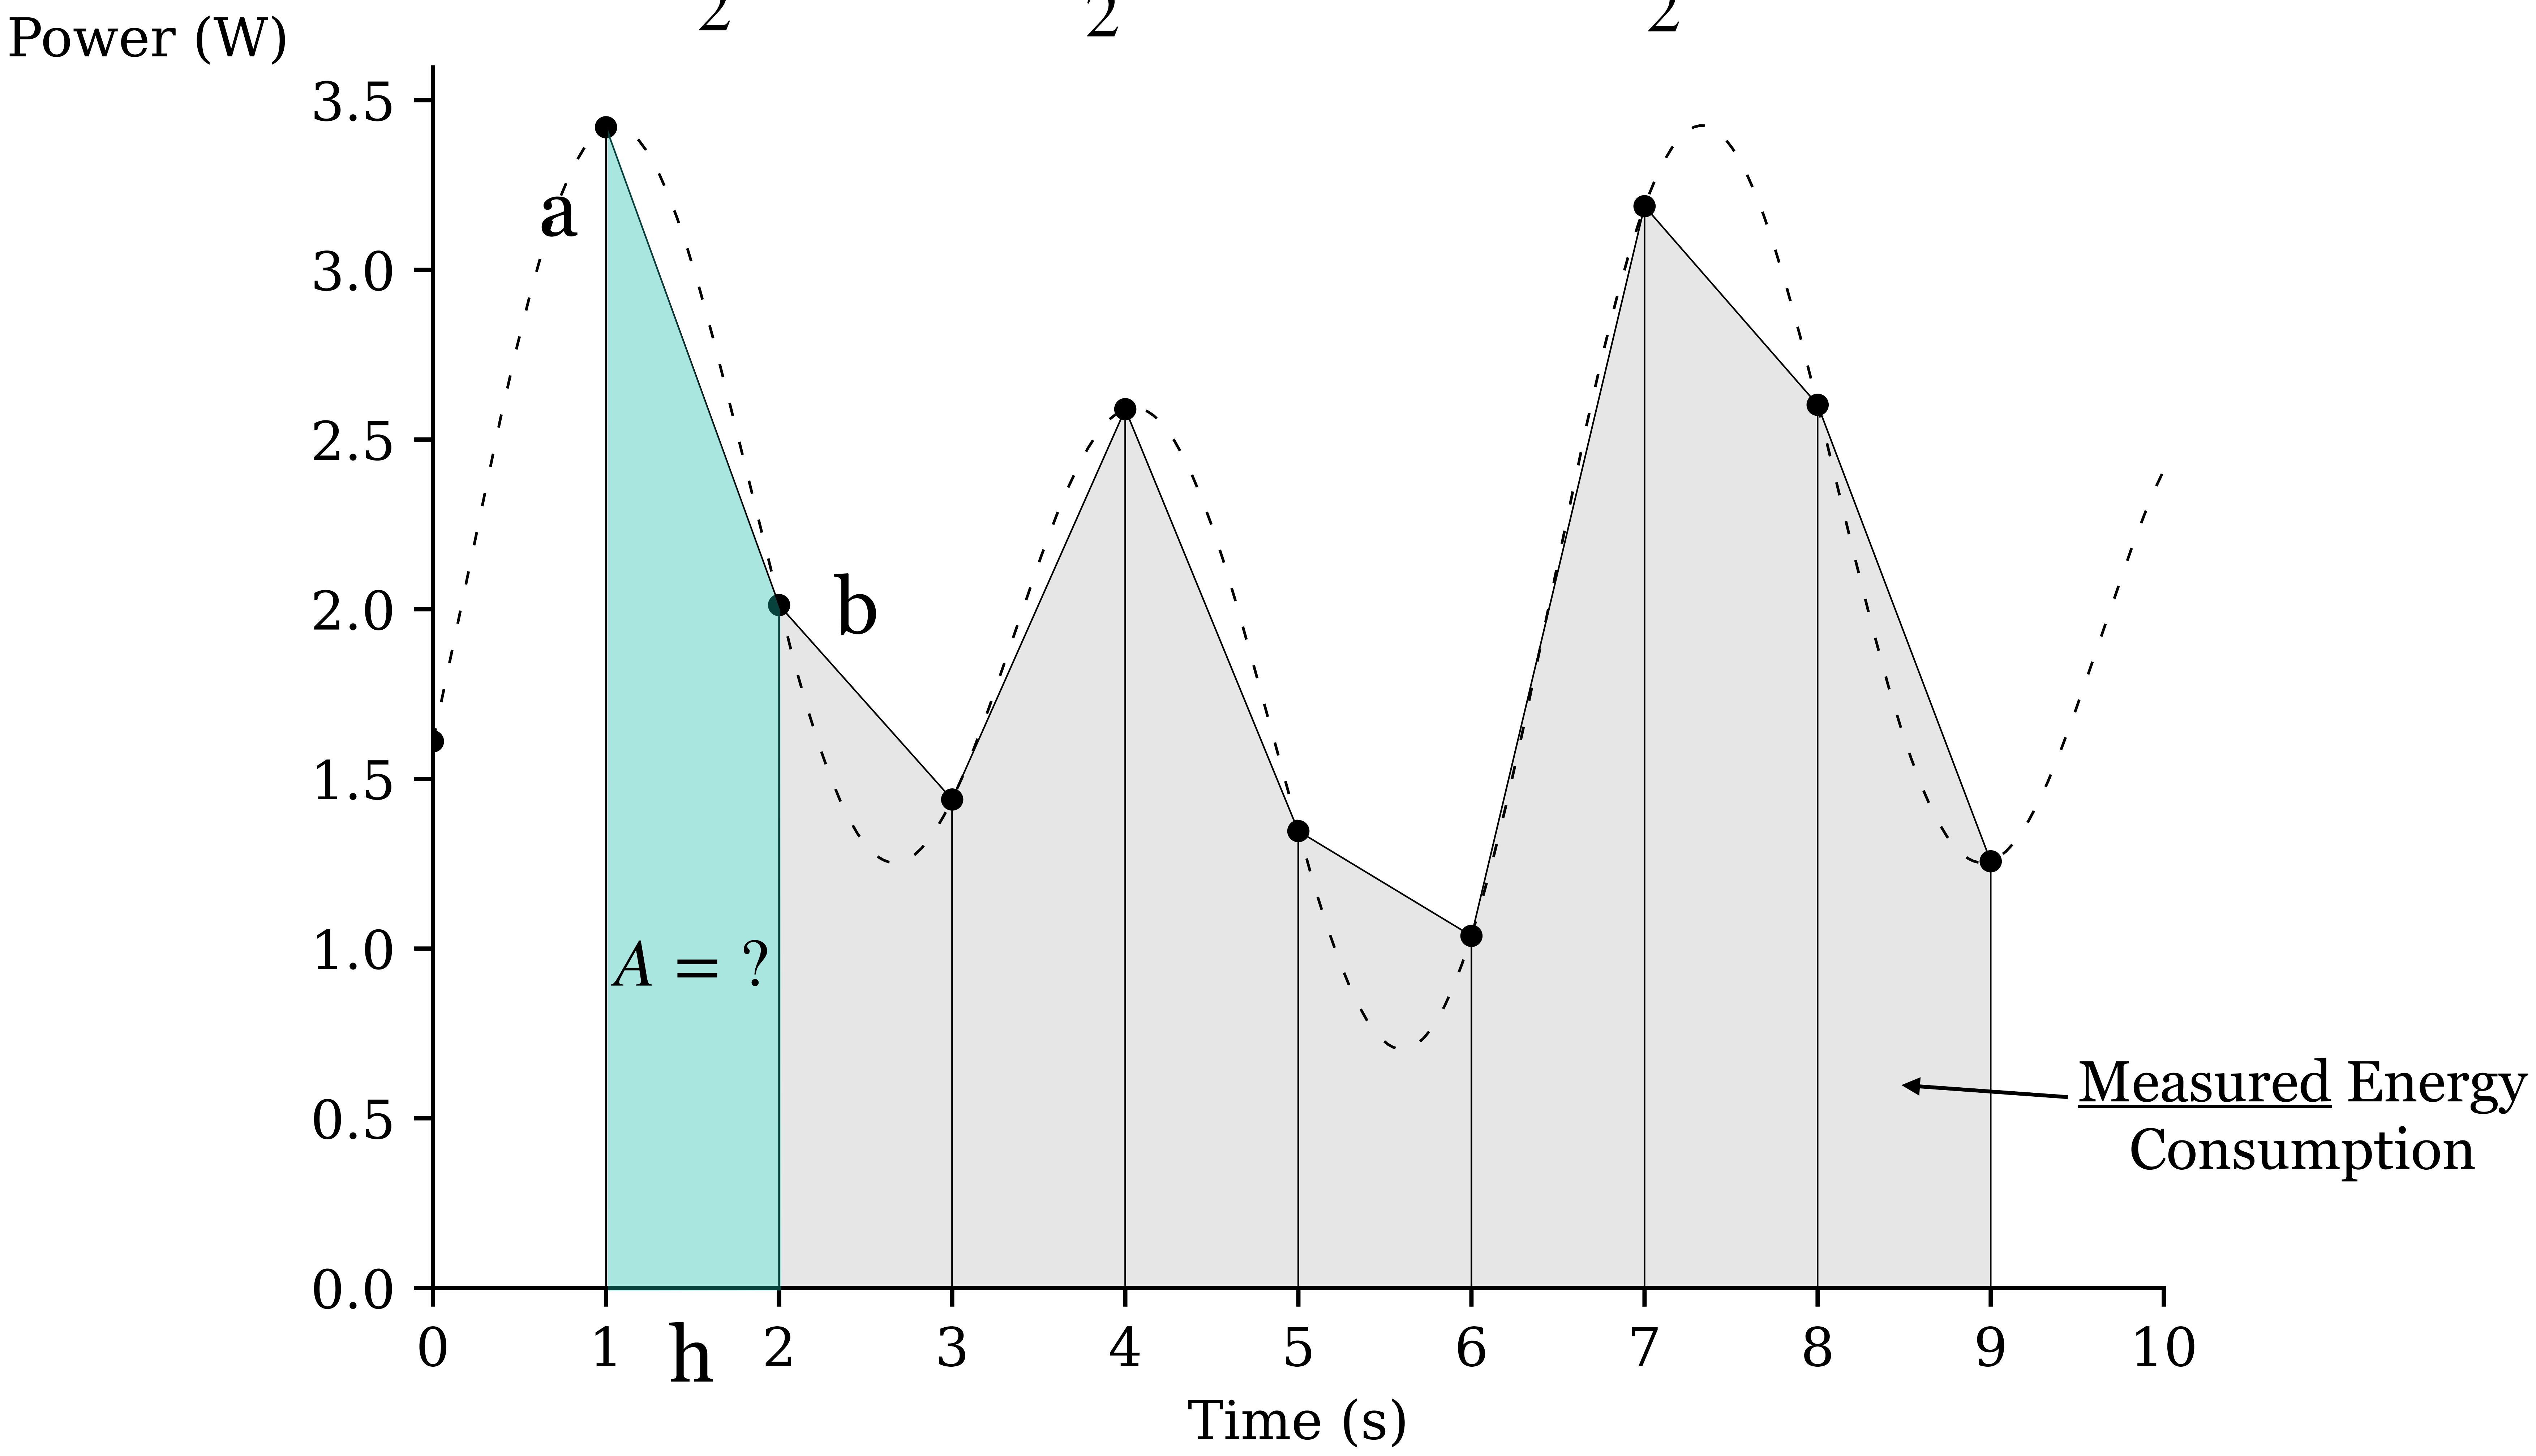

Trapezoid Rule

$$\int_{t_0}^{t_n} P(t)dt \approx \frac{\Delta t}{2}[P(t_0) + 2P(t_1) + 2P(t_2) + \ldots + 2P(t_{n-1}) + P(t_n)]$$

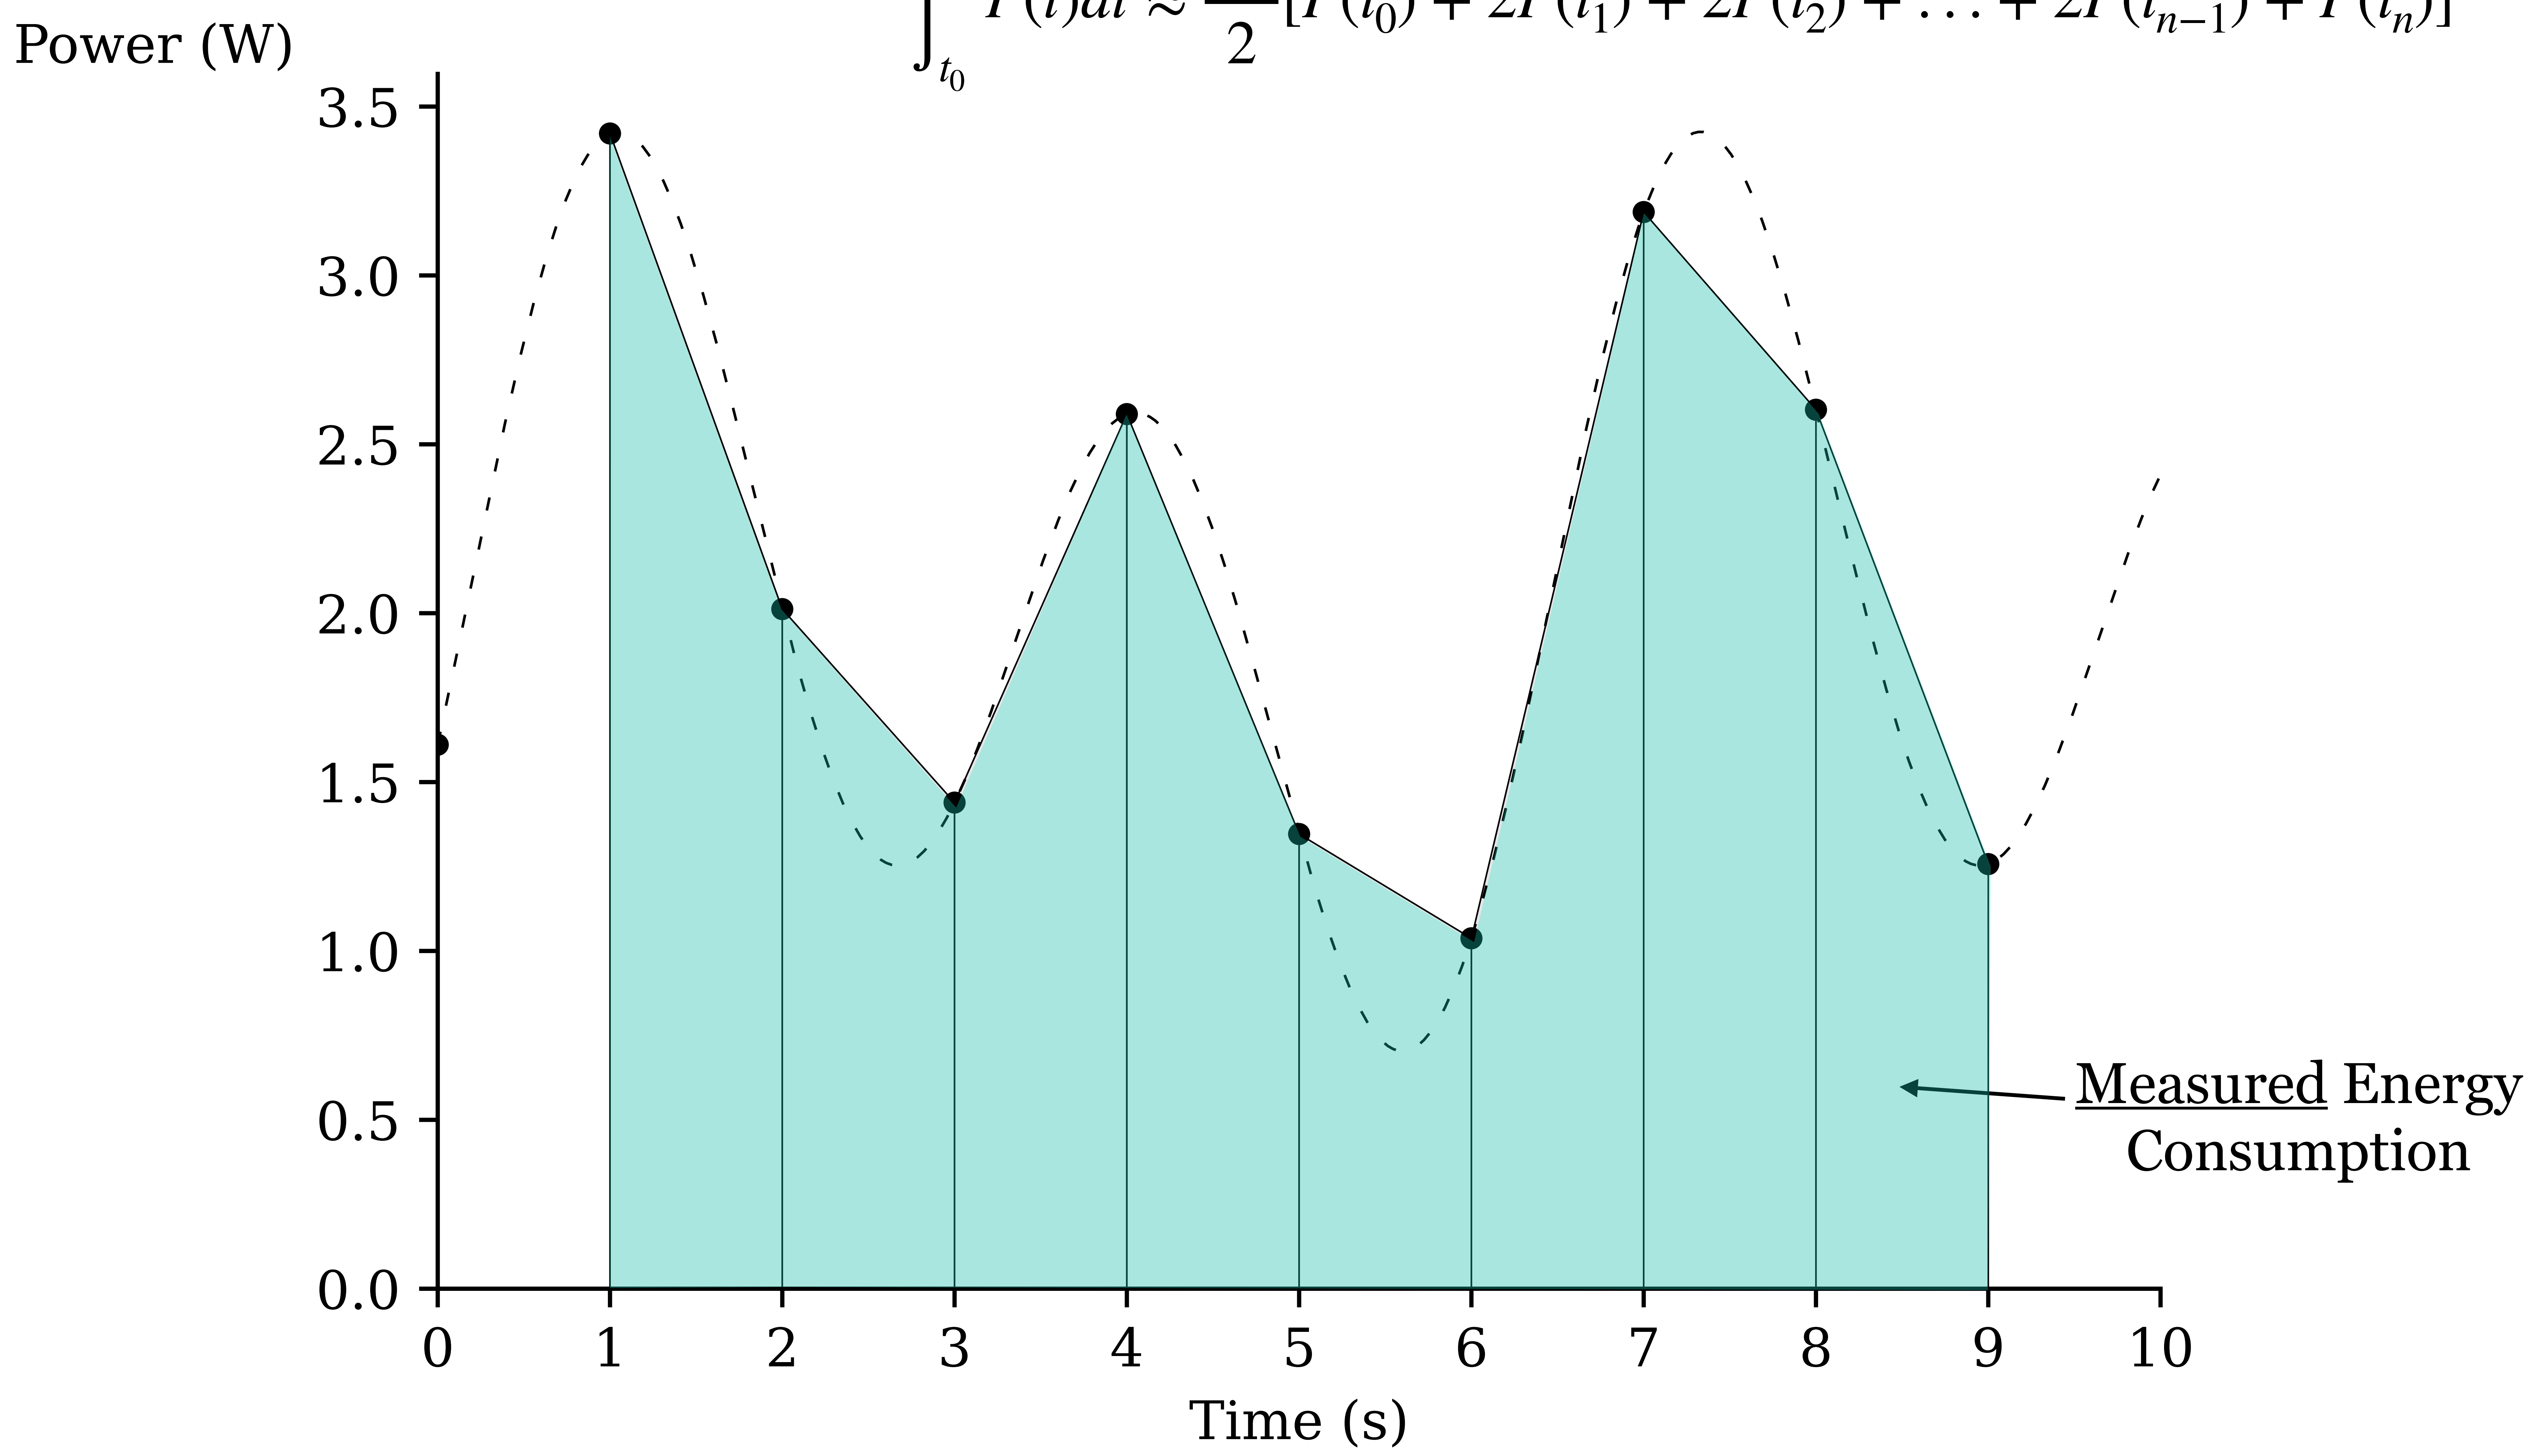

⚠️ Sometimes you cannot assume that the sampling interval (Δt) is always the same. ⚠️

Trapezoid Rule in Python

```python
import numpy as np
energy_consumption = np.trapz(power_sample, timestamps)
```

Average power

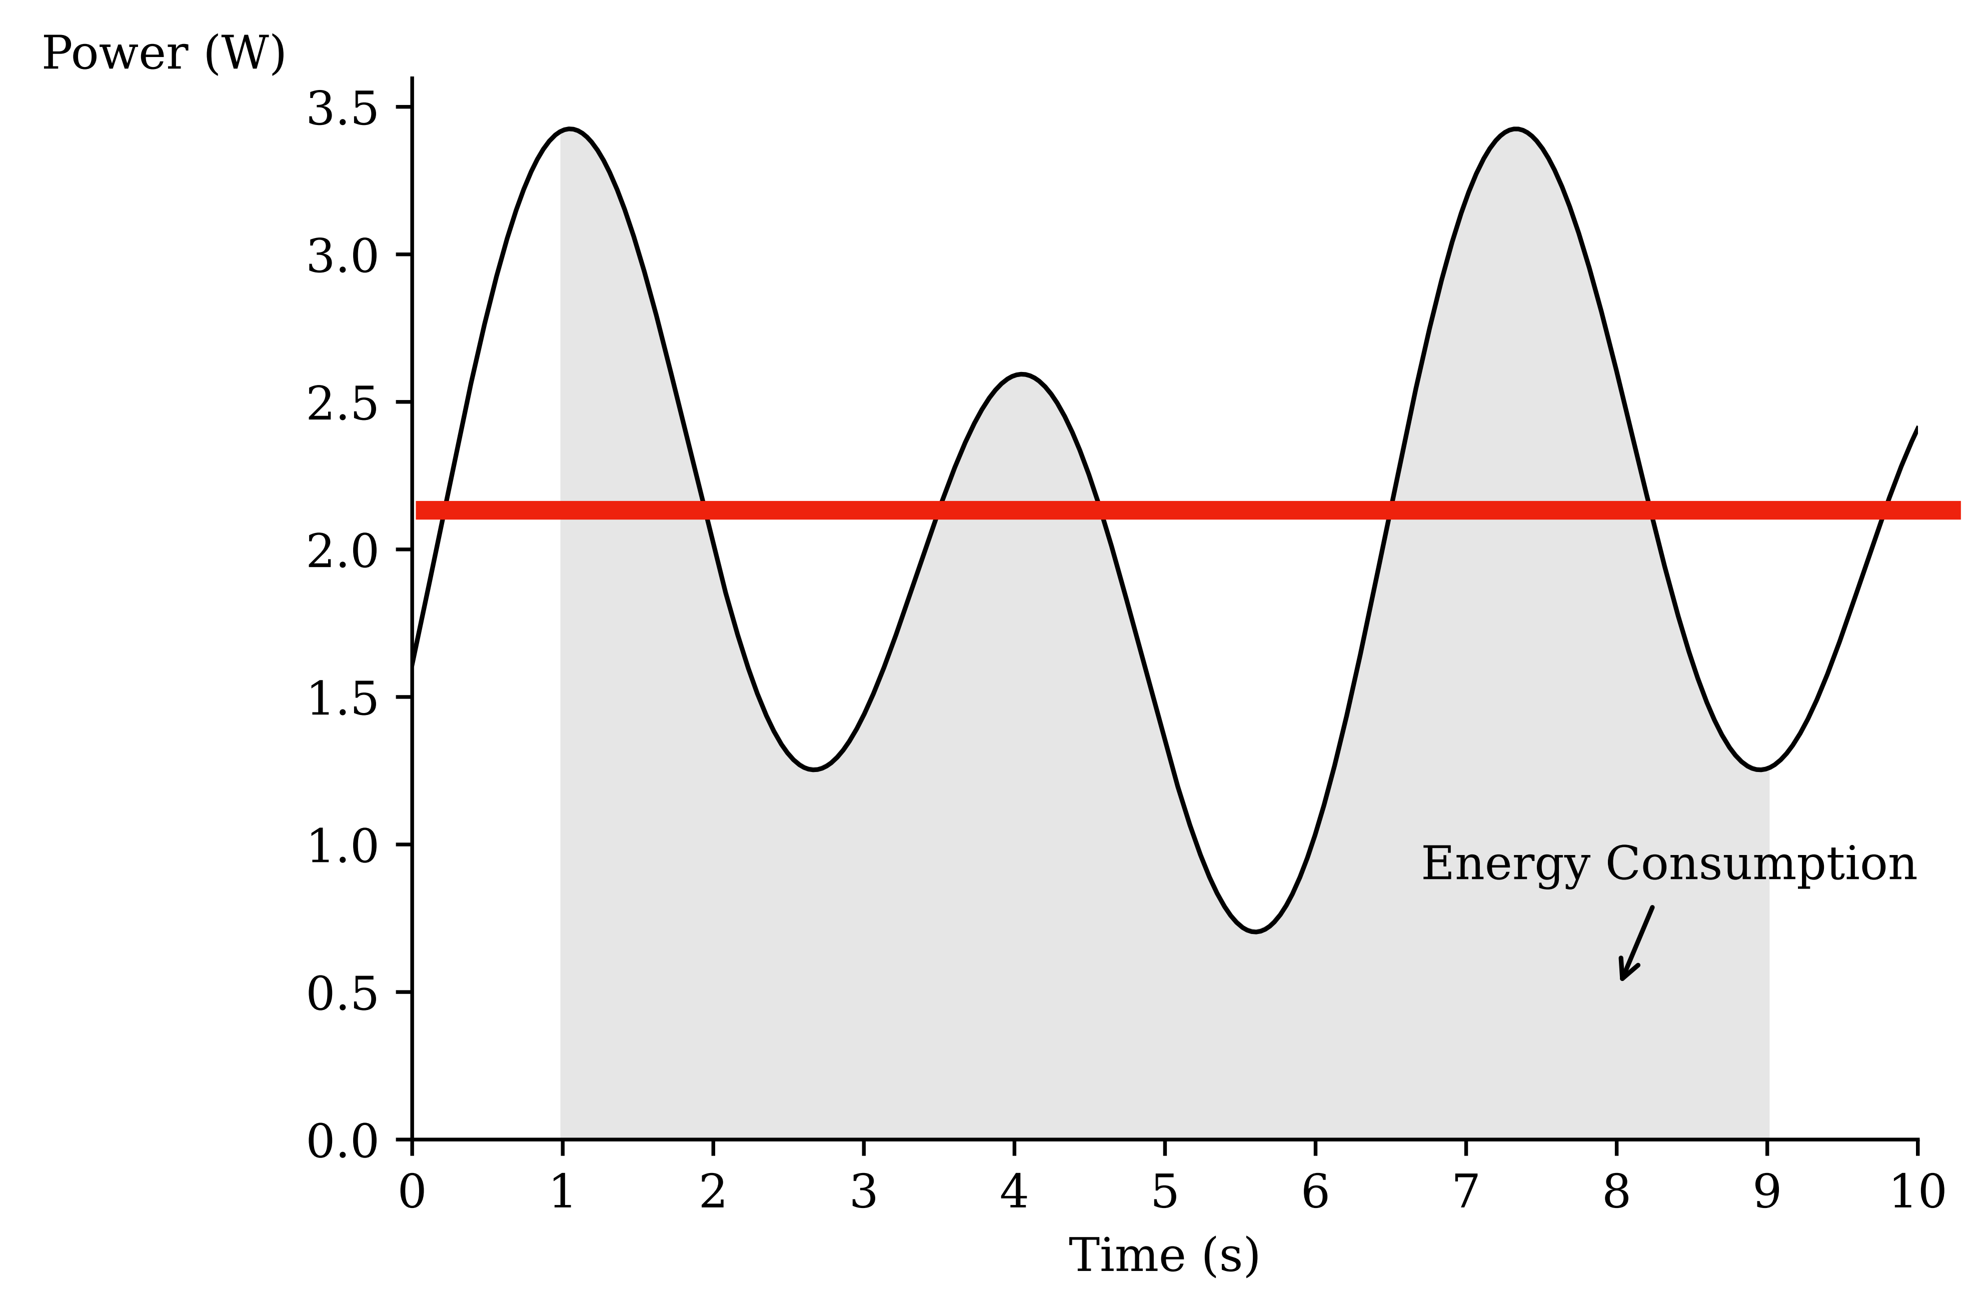

$$Energy = P_{avg} \cdot \Delta t$$

- **Easy to convert** to energy consumption
  - Simply multiply by the **elapsed time**.
- (This is another reason to **always collect time data** along with energy metrics.)

Power or **Energy**?

- **Average power consumption makes sense** when we report the consumption of a **continuous use case**. E.g., reading an ebook in your computer.
- **Energy consumption** makes sense in **one-off use cases**. E.g., energy consumption of a bitcoin transaction.

Learning activity

- Pair up with one colleague and discuss potential software use cases where one should test energy efficiency.
  - Choose one use case where **energy consumption** is the best metric to discuss/test energy efficiency.
  - Choose one use case where **average power consumption** is the best metric to discuss/test energy efficiency.

Energy Delay Product (EDP)

- In some cases, to achieve less energy consumption, one simply runs the software on a low power mode of the CPU.
    - E.g., setting the CPU at a low frequency will make execution slow but more energy-efficient.
- Energy consumption metric that **penalizes slow runs**

$$EDP = E \times t = \Delta P \times t^2$$

- Gives more importance to application runtime, with the goal of making both **low energy and fast** runtime applications.

?

The typical notebook battery has between 2,000 and 6,000 milliamp hours (mAh)

mAh

This is not energy or power.
It is a unit of electric charge.

Electric charge

- International System of Units (SI): **Coulomb (C)**.
- 1 electron has $1.602176634 \times 10^{-19}$ coulombs. Moving the electron around the electric field requires work (energy consumption).
- **mAh** is the most common metric to specify the capacity of batteries.
- **1 mAh = 3.6C**
- To compute the actual energy of a battery we need to factor in **voltage**:

$$Energy = Voltage \times Charge$$

- E.g., for a battery with a capacity of 1000mAh:

$$1000\text{mAh} \times 3.8\text{V} = 3800\text{mWh} = 3.8\text{Wh} = 3.8 \times 3600\text{J} = 13680\text{J}$$

Why do we use **charge units** for batteries?

- There is a continuous change of voltage throughout a **charge/discharge cycle**.
- E.g., it can start with 4.5 V at a “100%” capacity and from to 3.0V at 5% capacity.
- Most devices use voltage to compute their **battery level percentage**.

International System of Units (SI)

- Energy: **Joule**
- Power: **Watt**
- Charge: **Coulomb**
- (Time: **second**)

SI Units are difficult to grasp

- Whenever talking to a general audience use relative units:
    - Compare to the other well-known things:
        - Yearly **household** energy consumption
        - Yearly **country** electrical energy consumption. (e.g., https://ccaf.io/cbeci/index/comparisons)
        - **Driving kms** with a standard car
        - Percentage of a normal **battery** charge cycle.
    - Compare to other software artefacts/usecases:
        - E.g, percentage of **Version A over Version B**.

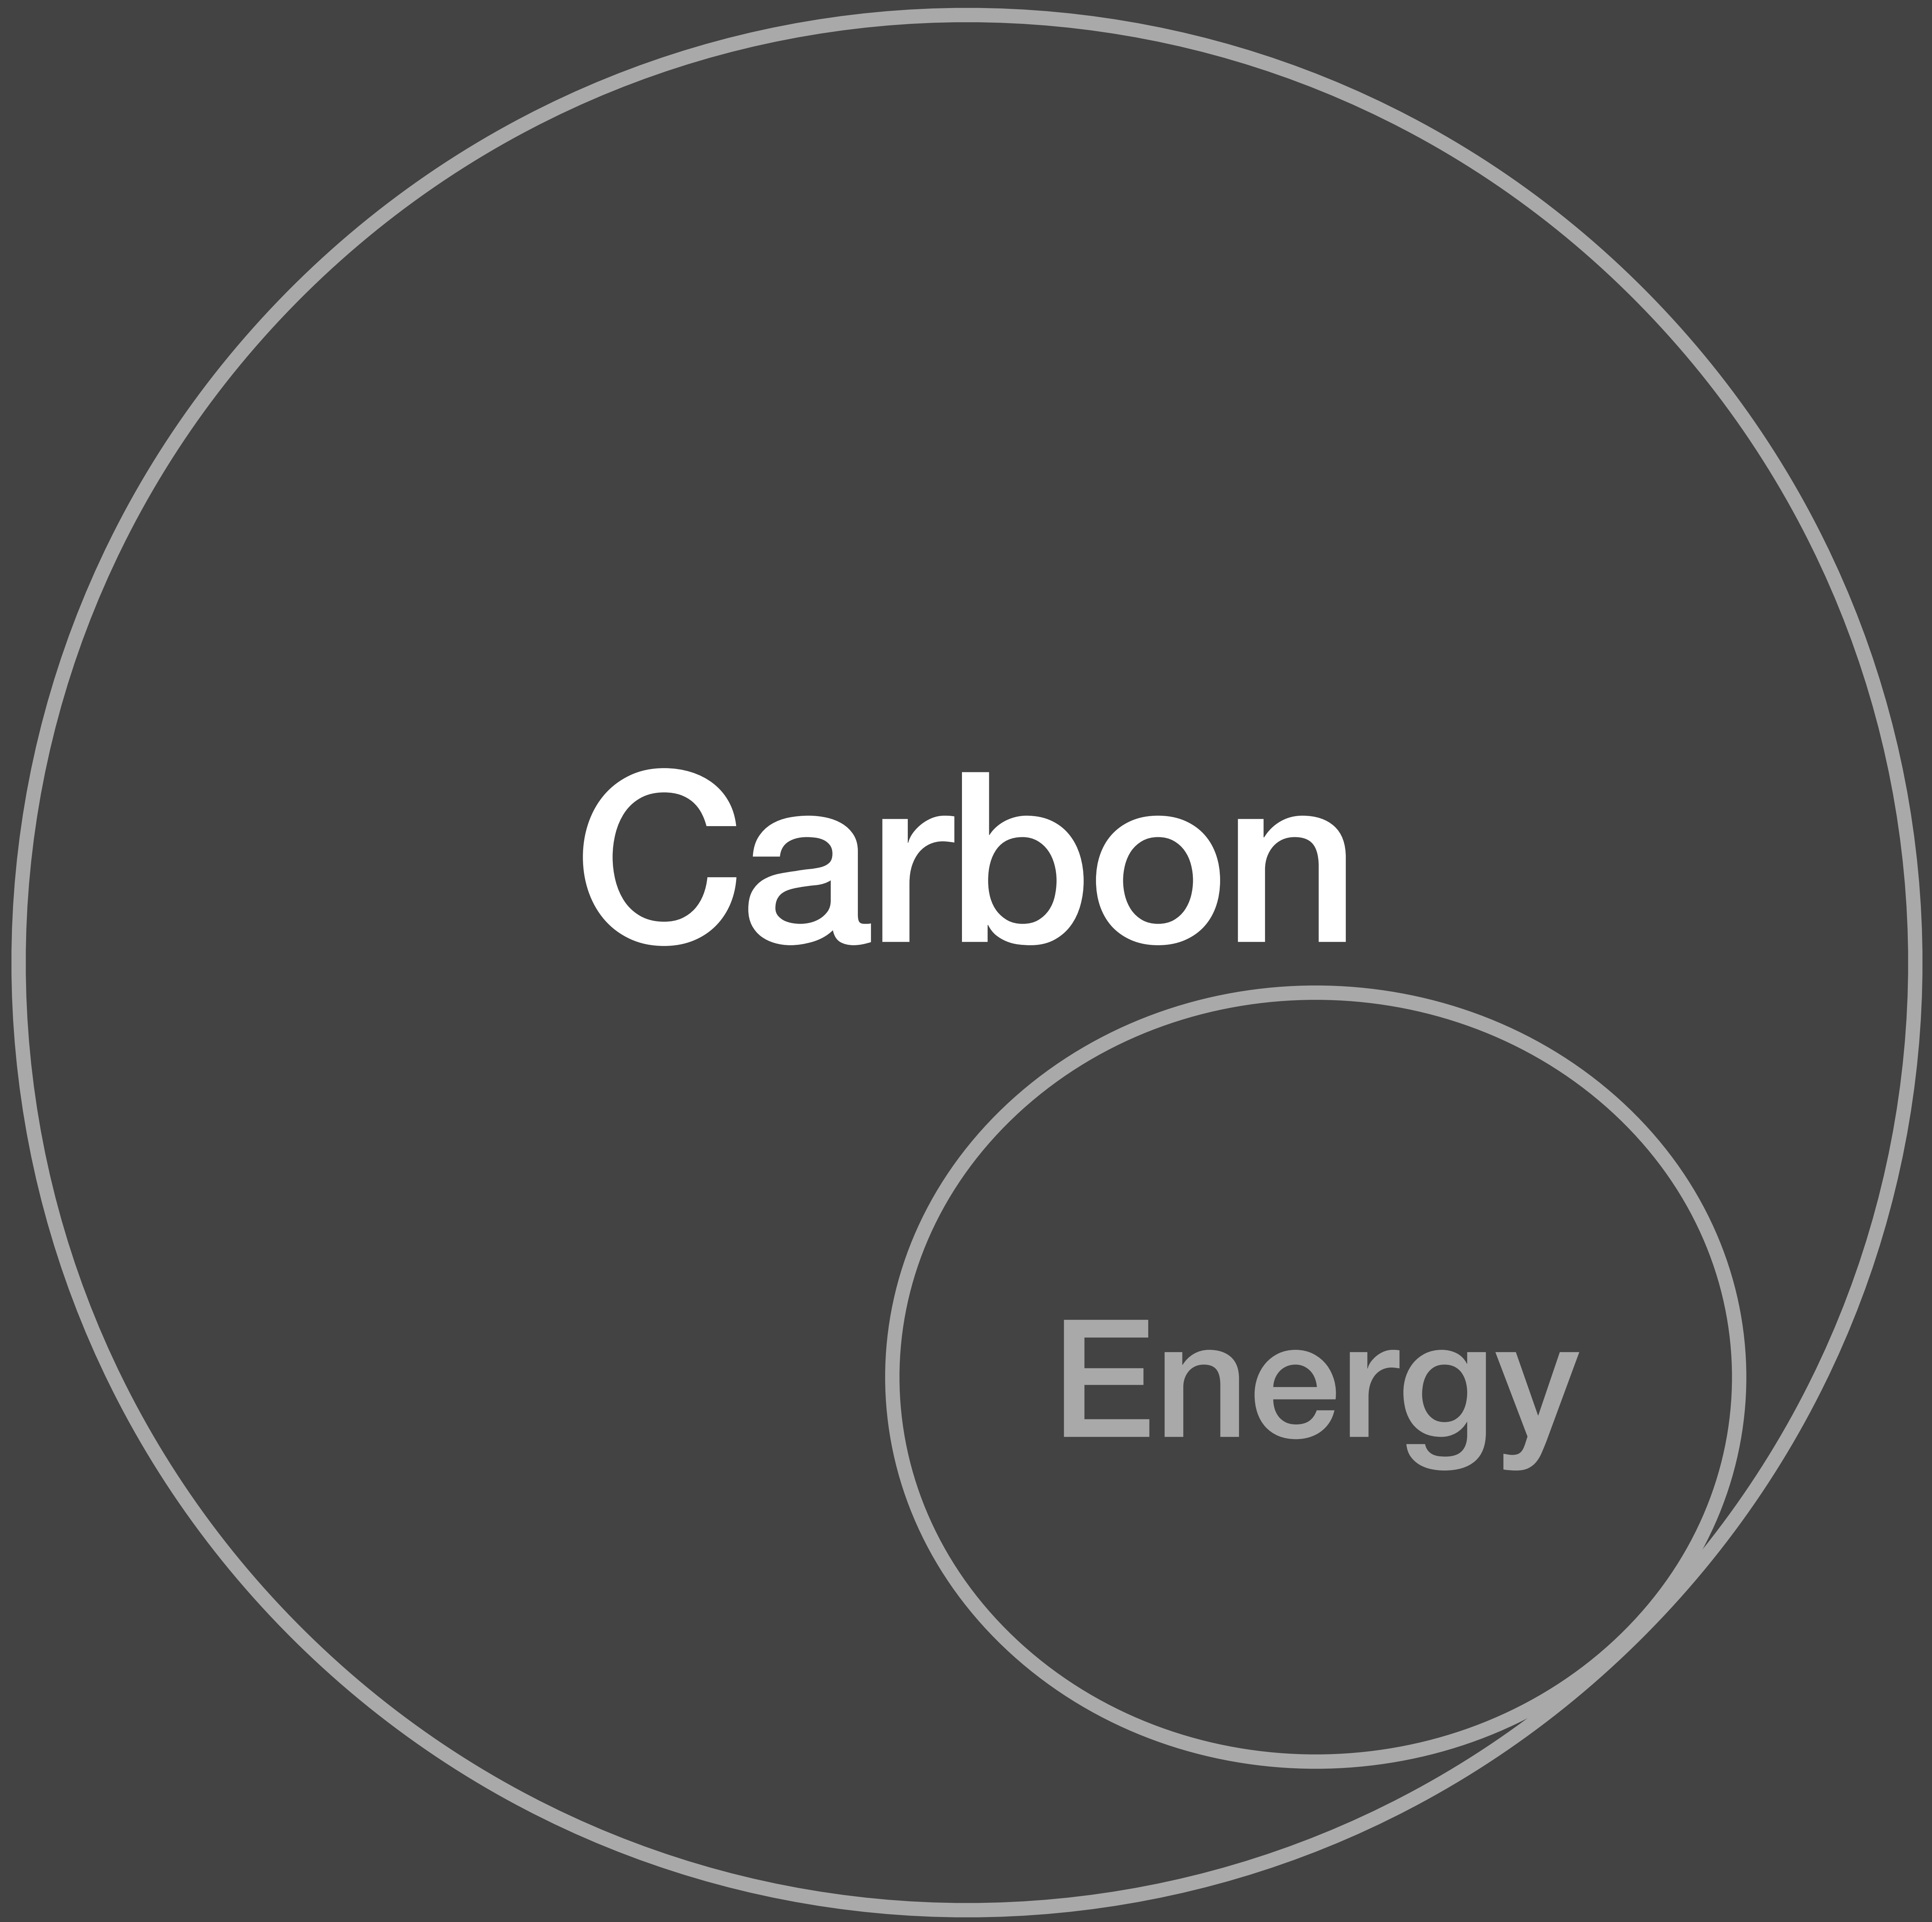
Carbon
Energy

Carbon

- **Greenhouse gas (GHG)** emissions **by human activities** are the main root of the **Global Warming**.
- There are many GHGs but they have different impacts on global warming
  - The most harmful: **Carbon Dioxide** and **Methane**.
- The **Kyoto protocol,** signed in 1997, defines **7** main GHGs.
  The other five: Nitrous Oxide (N2O), Hydrofluorocarbons (HFCs), Perfluorocarbons (PFCs), Sulphur Hexafluoride (SF6), and Nitrogen Trifluoride (NF3). https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:Greenhouse_gas_(GHG)
- All these GHGs have a different impact in the atmosphere.
  How can we report GHG emissions in a single unit?

Carbon Dioxide Equivalent (CO2-eq)

- A weight function that combines all gas emissions into their carbon dioxide equivalent
- 1kg of **Methane** (CH4) is estimated to be **21 times** more harmful than 1kg of **Carbon Dioxide** (CO2).
- The weight function relies on the **estimation of the impact** of GHGs over a period of **100 years** when compared to **carbon dioxide**. Aka 100-global-warming potential (**100-GWP**).

$$CO_2eq = \sum_{g \in GHG} (GWP_g \cdot m_g)$$

- Co2-eq is expressed in mass – e.g., **kg**CO_2eq

Greenhouse Gas	100-GWP
Carbon dioxide CO_2	1
Methane CH_4	21
Nitrous oxide N_2O	310
Sulphur hexafluoride SF_6	23900

Carbon Dioxide Equivalent (CO2-eq)

$$CO_2eq = \sum_{g \in GHG} (GWP_g \cdot m_g)$$

Greenhouse Gas	100-GWP
Carbon dioxide CO_2	1
Methane CH_4	21
Nitrous oxide N_2O	310
Sulphur hexafluoride SF_6	23900

- **As an example, imagine that to run our software system our electricity provider emits 1000Kg of CO2, 20Kg of CH4, 5Kg of N2O, 0Kg of the remaining GHG.**

$$\begin{aligned} CO_2eq &= GWP_{CO_2} \cdot m_{CO_2} + GWP_{CH_4} \cdot m_{CH_4} + GWP_{N_2O} \cdot m_{N_2O} \\ &= 1 \times 1000 + 21 \times 20 + 310 \times 5 \\ &= 2670 \text{kg} CO_2eq \end{aligned}$$

Note:

- **100-GWP** is only an estimation;
- different sources reveal different estimations;
- there is also the 20-GWP and the 500-GWP.

I use this source: Foster et al. (2017) **Changes in Atmospheric Constituents and in Radiative Forcing**
https://archive.ipcc.ch/pdf/assessment-report/ar4/wg1/ar4-wg1-chapter2.pdf

Carbon credits (quick detour)

- Strategy used to **regulate allowed emissions** and to make carbon emission rights **tradable**.
- Each entity (e.g., company/country) has a budget of carbon credits.
- Entities can **buy carbon credits** from other entities when they are **over budget**.
  - In the case of companies, carbon credits can only be bought from **GHG mitigation projects**.
- 1 carbon credit = 1 tonne CO_2-eq
- **Consequence**: the price of carbon credits is rising and carbon trading is starting to be interesting for investors.

When should we use Carbon vs Energy?

- Energy/Power is more useful at the **software usecase** level.
- Carbon is more useful at the **infrastructure** level (e.g., datacenter) or at the **project** level (e.g., the impact of developing a full software project).
- Choose your metrics wisely ;)

- Carbon emissions ✅
- How do we go from **energy** consumption to **carbon** consumption?

Carbon intensity

- How much **carbon** is emitted **per kWh** of electricity consumed.
- The common unit: gCO_2eq/kWh 🤷‍♂️
- E.g., **gas-based power plants** emit **less carbon** than **coal-based plants**.
- The power grid is a mix between different sources of electricity – different locations have different carbon intensity.

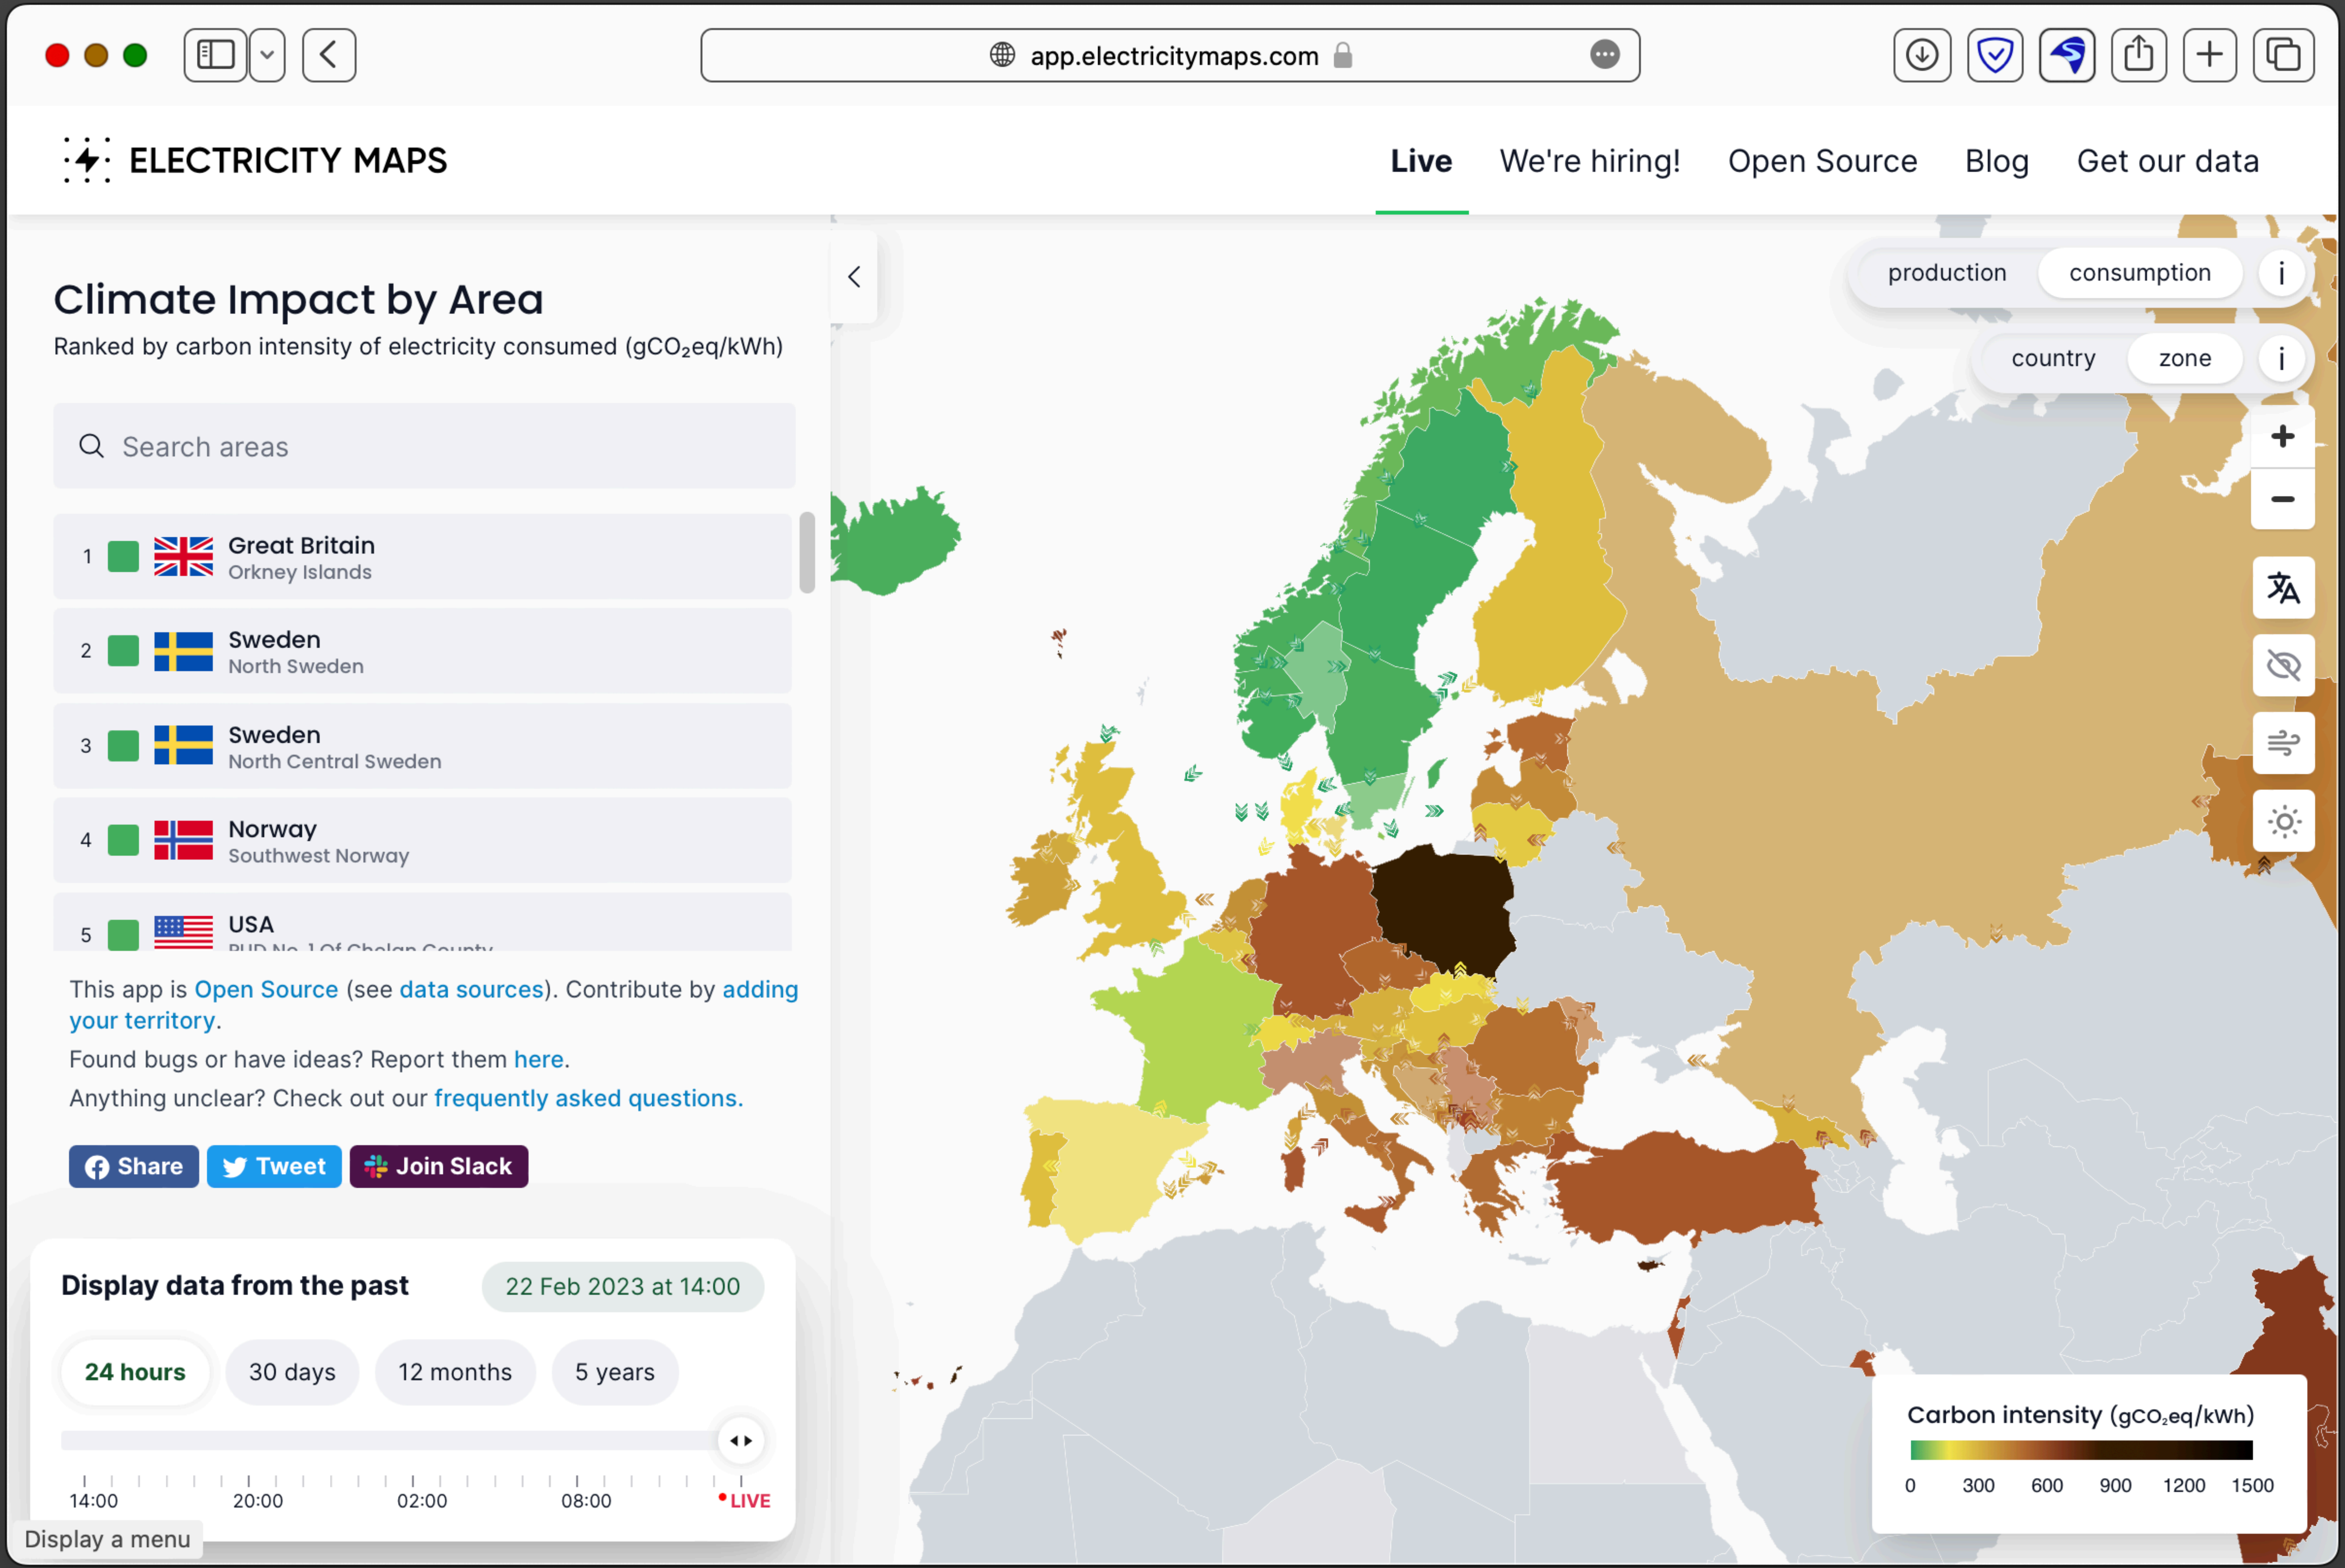

https://app.electricitymaps.com/map

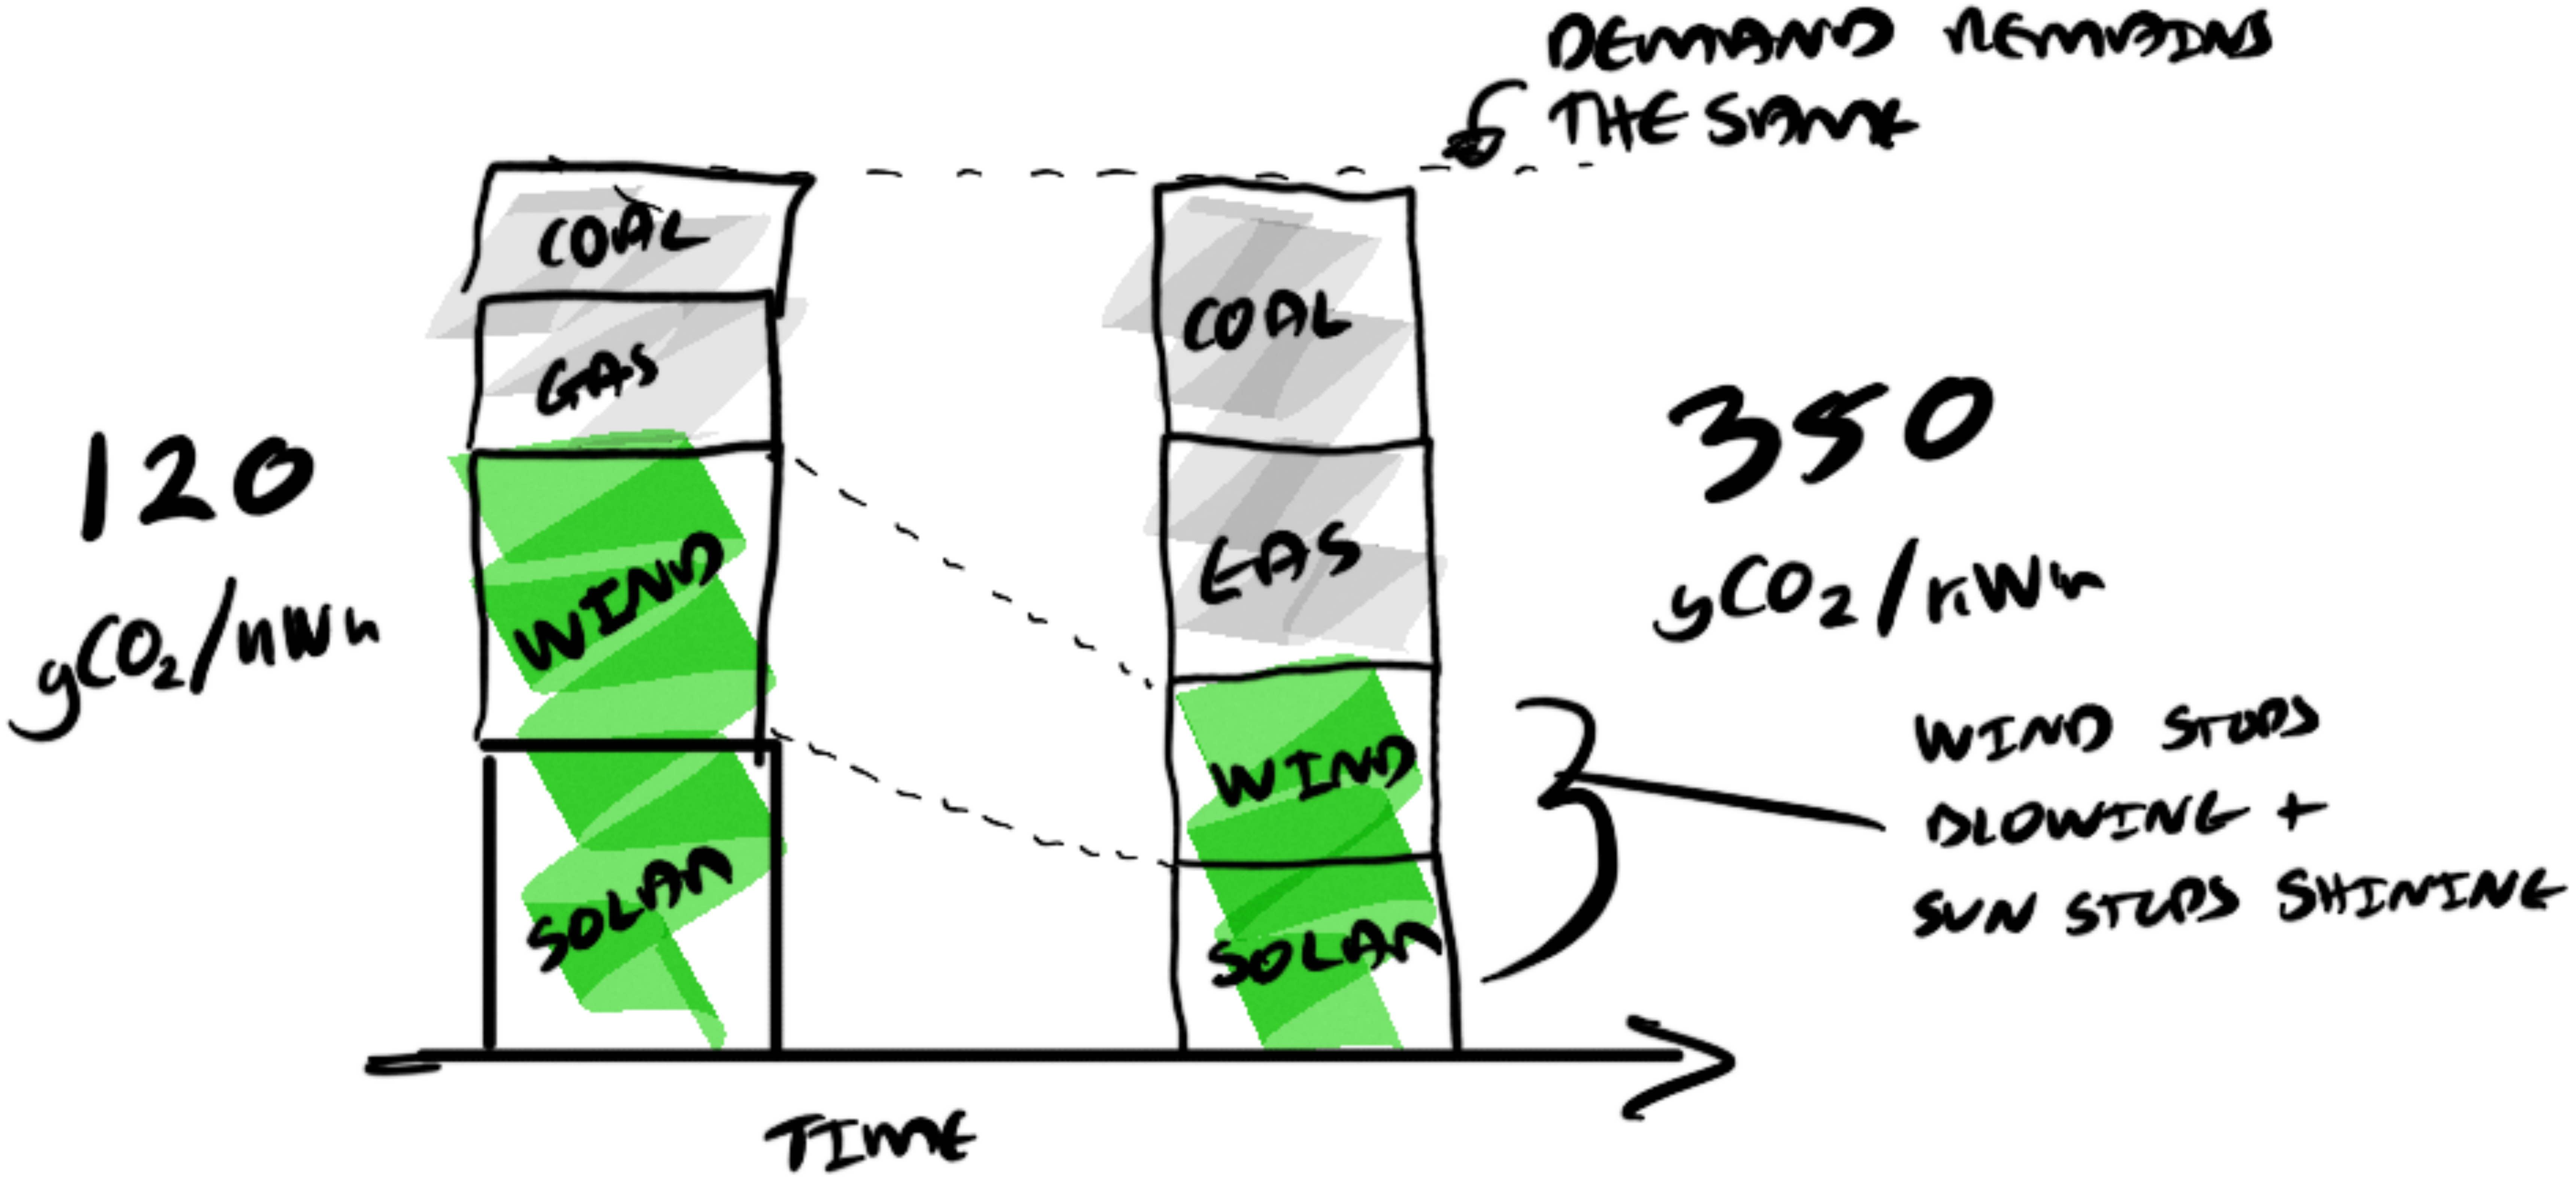

By Asim Hussain: https://principles.green/principles/carbon-intensity/

📝 Reducing software energy consumption can help reduce the carbon intensity. Why?

?

⚠️

One would expect zero carbon intensity from solar-panels or wind farms, but that’s not the case.

Marginal Power Plant

- Renewable-based power plants cannot adapt to demand.
- When demand is higher than the existing power in the electricity grid, we need a power plant that is able to scale up to that demand.
  - This is usually done by **fossil-based power plants**. They are called the **marginal power plants**.
- The problem is that marginal power plants do not scale down to zero.
  - There is always a minimum carbon that needs to be emitted, even if there is a lot of renewable energy in the grid.

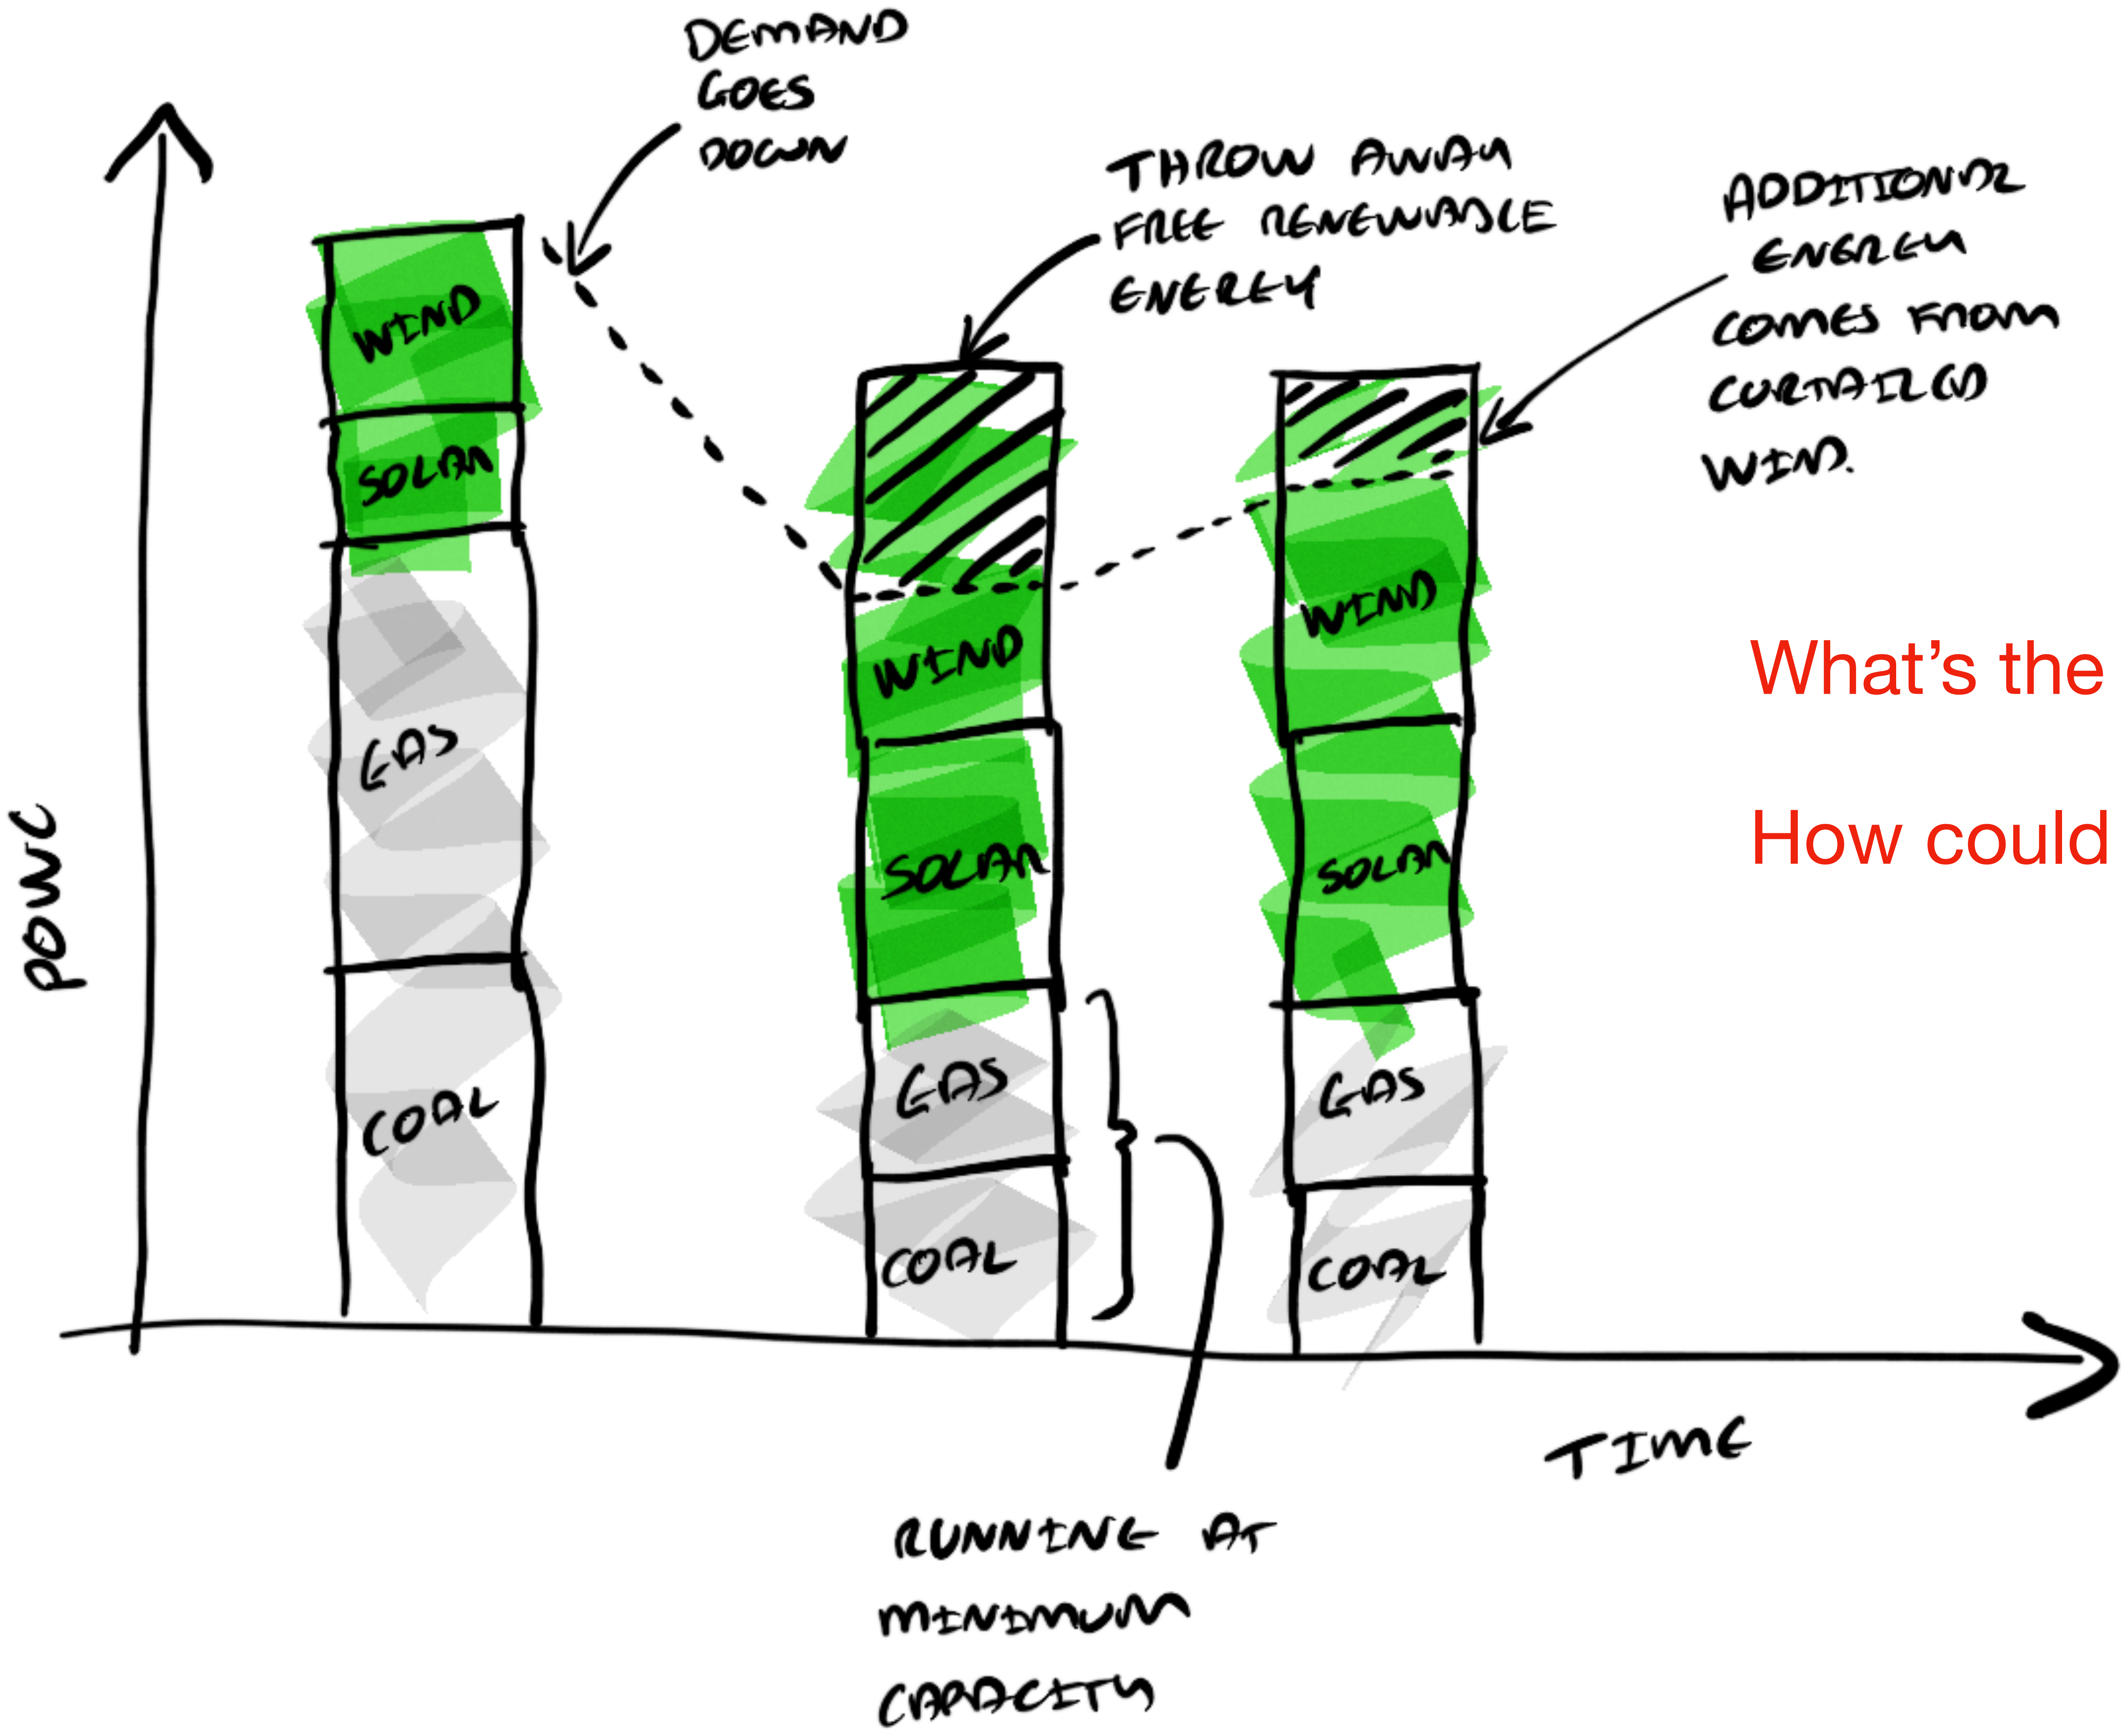

What’s the problem here?

How could we fix this?

By Asim Hussain: https://principles.green/principles/carbon-intensity/

Marginal Carbon Intensity

- Increase or decrease in carbon emissions in the electrical grid, in **response to an infinitesimal** increase/decrease in power demand/supply.

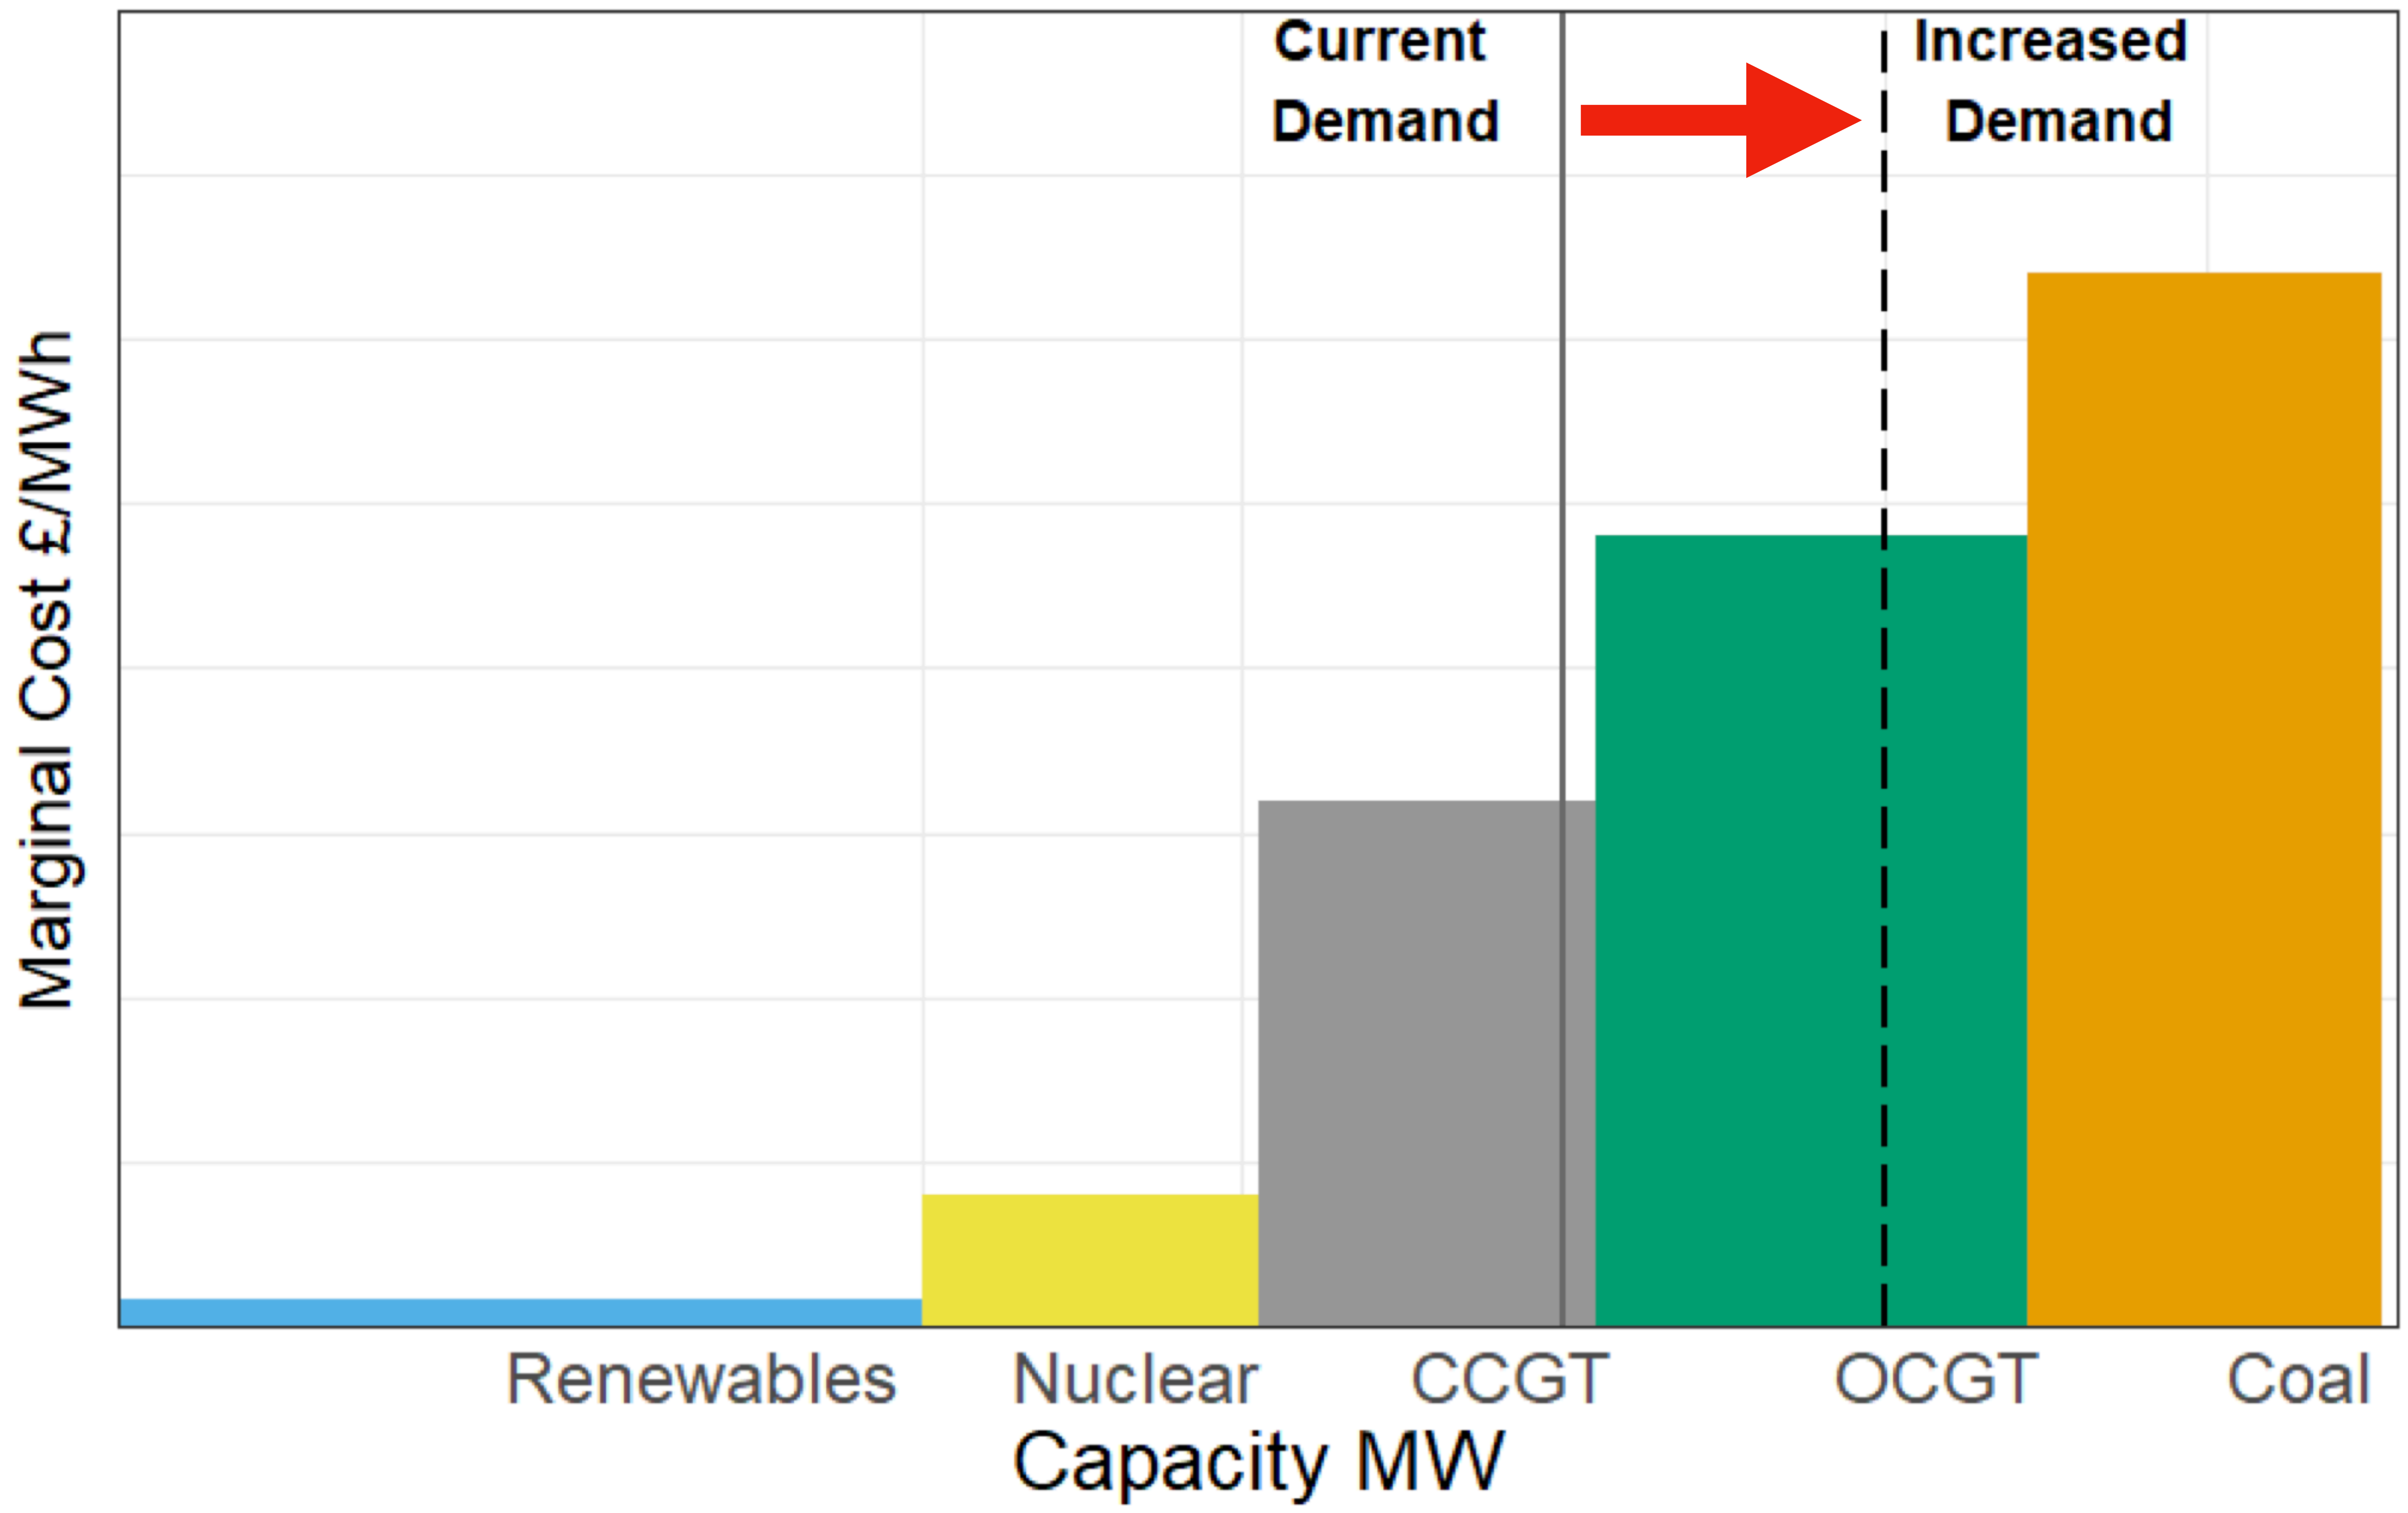

From: “Literature Review: On the effectiveness of a Marginal Carbon Intensity Signal ”

Why is marginal carbon intensity relevant for software?

- **Tip:** consider a task scheduler in a datacenter.

Recap

?

- Power
- Energy
- Average Power
- Energy Delay Product
- Electric charge (battery capacity)
- Carbon dioxide equivalent (carbon emissions)
    - 100-global-warming potential
- Carbon Intensity
- Marginal Carbon Intensity

Further Reading

?

- Blog post on energy units: https://luiscruz.github.io/2023/05/13/energy-units.html