Transparency first: assessing (and reporting) the Software Carbon Intensity of large-scale AI systems

Ivano Malavolta

i.malavolta@vu.nl

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP

VU

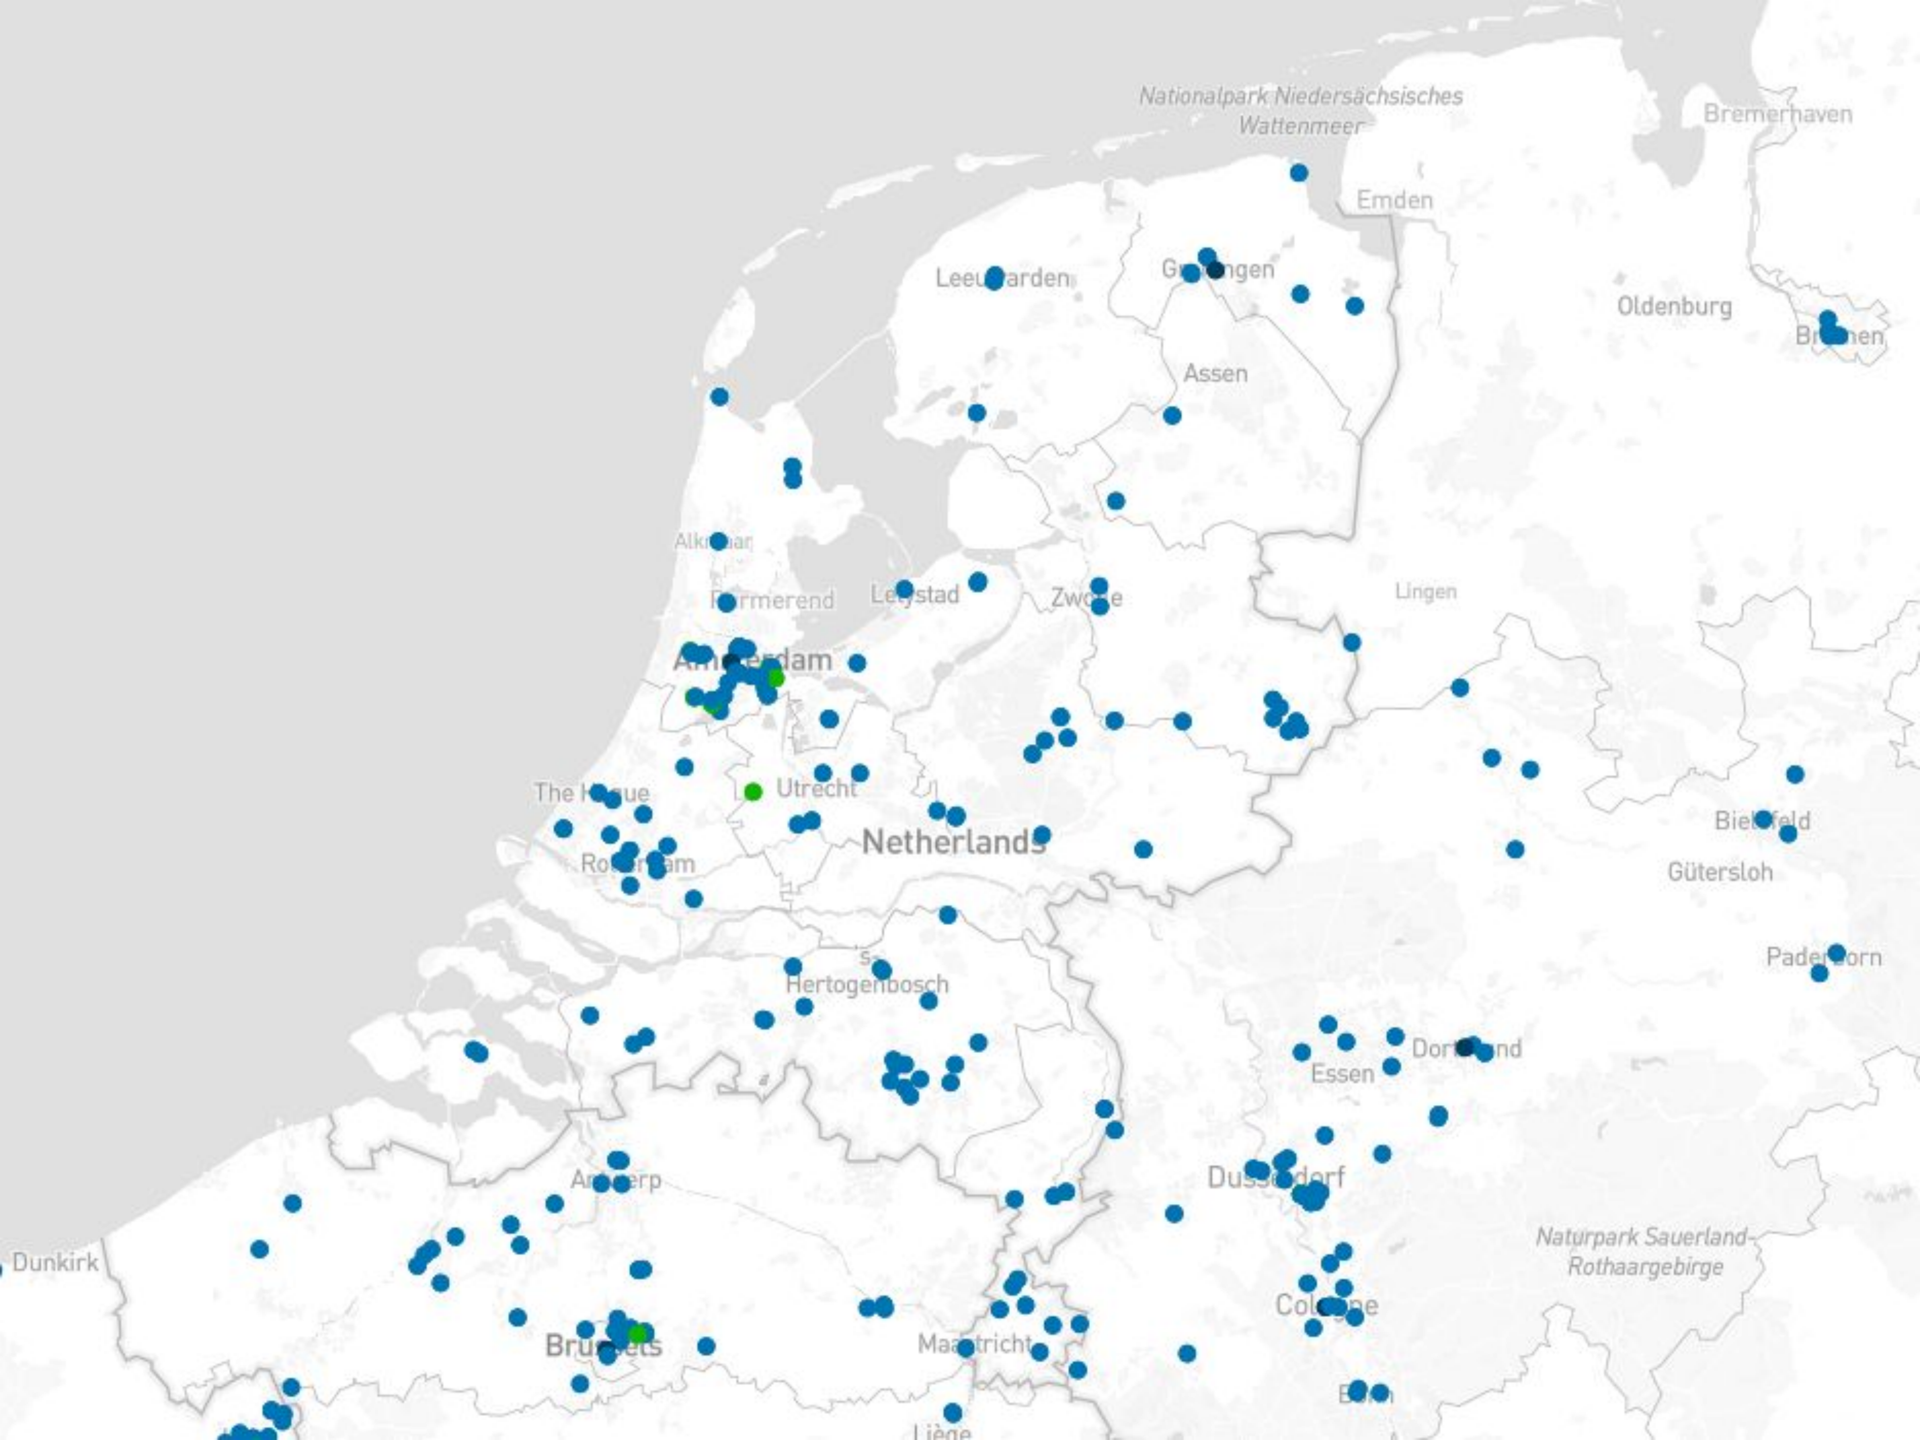

Nationalpark Niedersächsisches
Wattenmeer
Bremerhaven
Emden
Oldenburg
Assen
Lingen
Utrecht
Netherlands
Gütersloh
Hertogenbosch
Essen
Naturpark Sauerland-
Rothaargebirge
Dunkirk
Liège

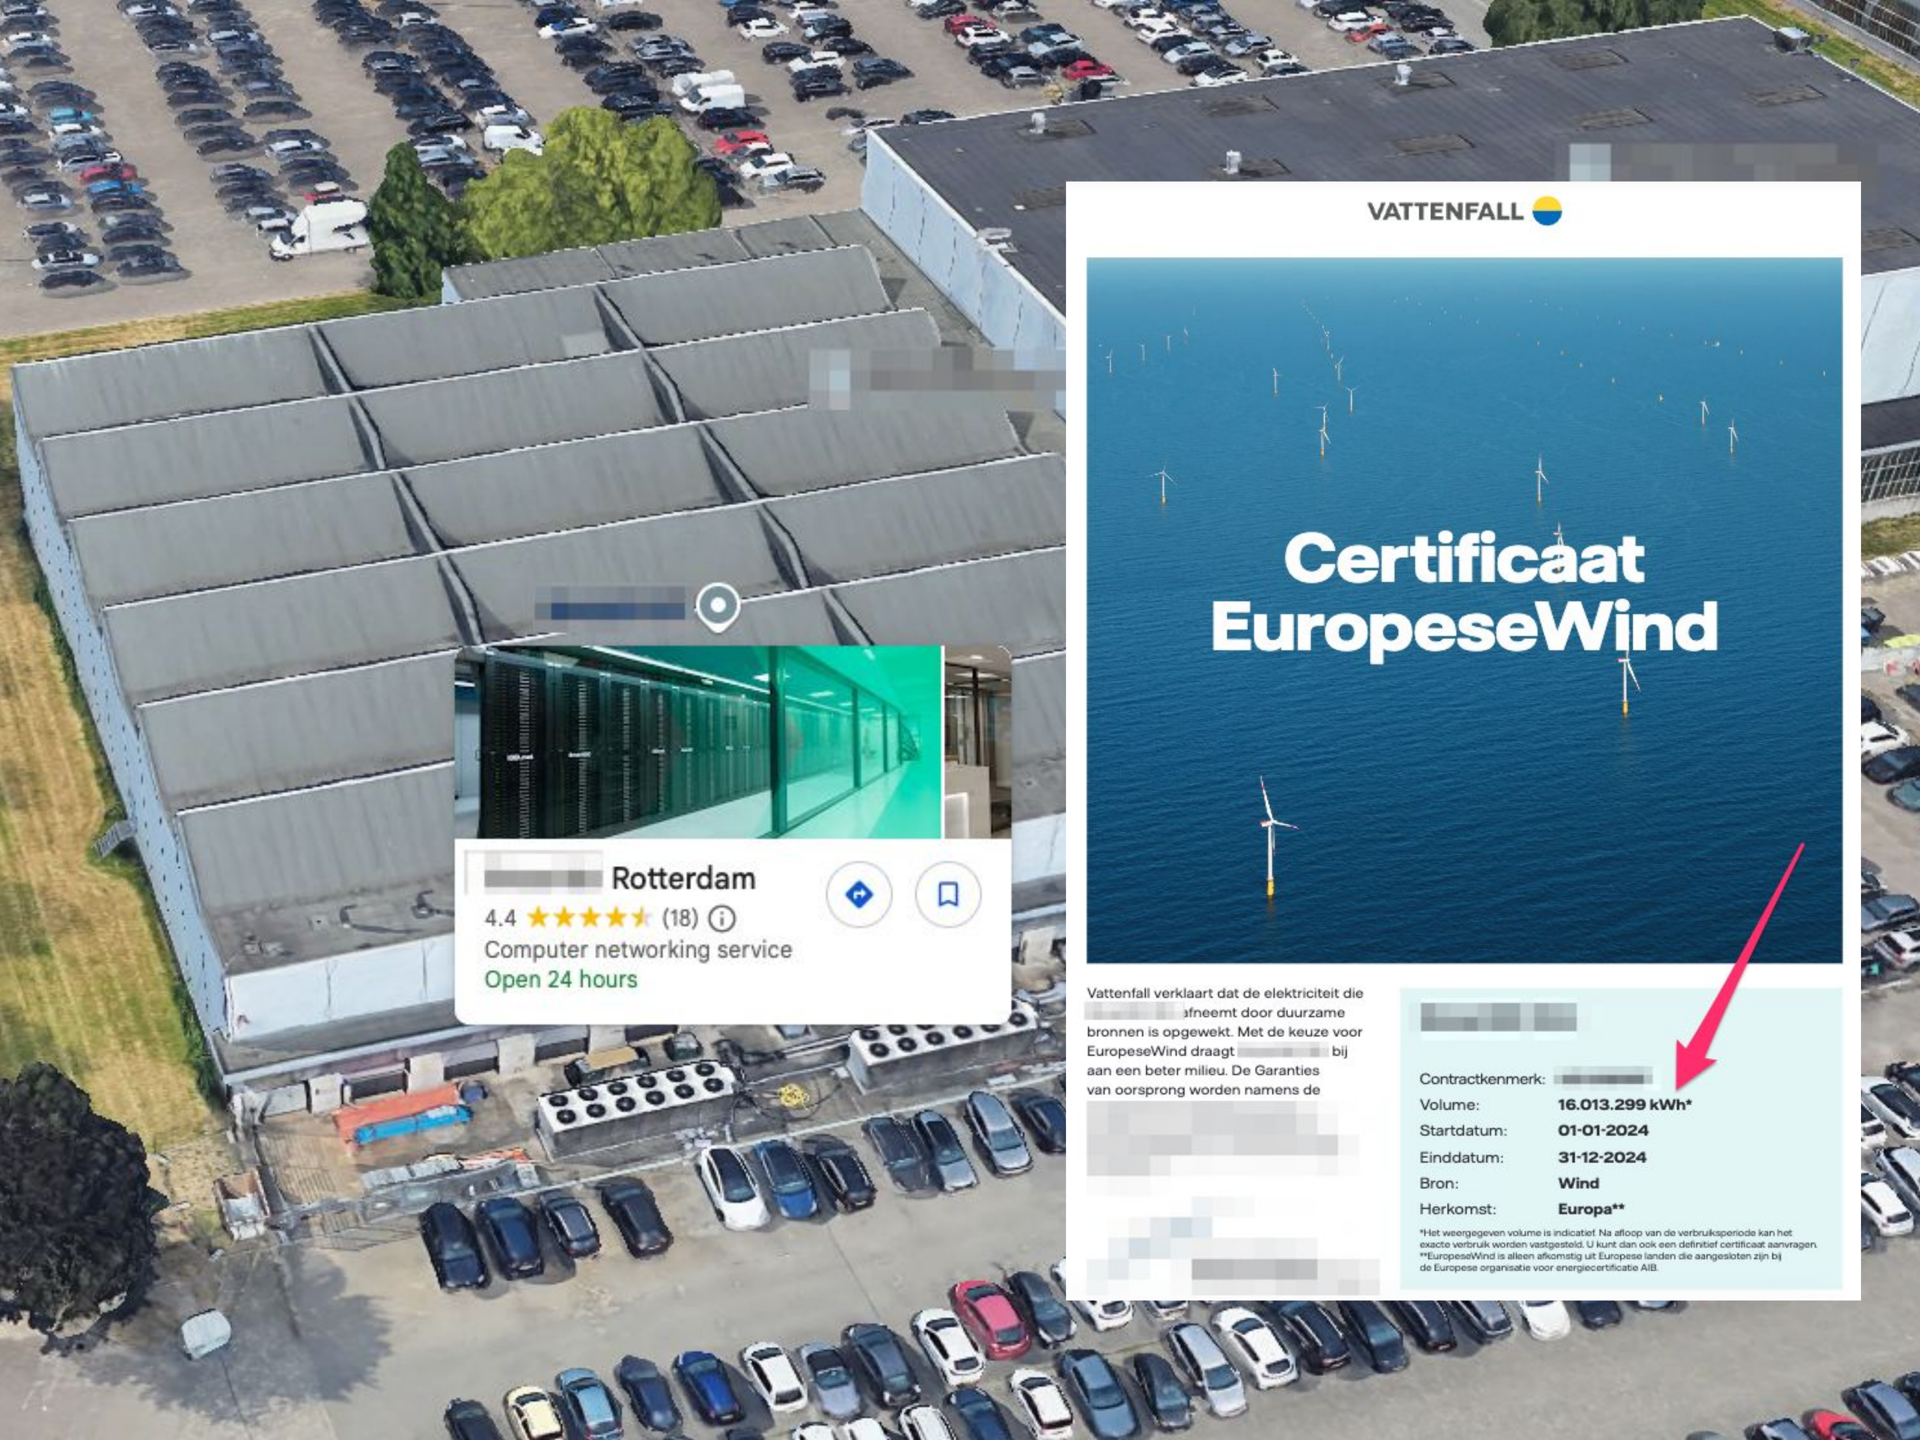

Rotterdam
4.4 (18)
Computer networking service
Open 24 hours
VATTENFALL
Certificaat
EuropeseWind
Vattenfall verklaart dat de elektriciteit die
afneemt door duurzame
bronnen is opgewekt. Met de keuze voor
EuropeseWind draagt bij
aan een beter milieu. De Garanties
van oorsprong worden namens de
Contractkenmerk:
Volume: 16.013.299 kWh*
Startdatum: 01-01-2024
Einddatum: 31-12-2024
Bron: Wind
Herkomst: Europa**
*Het weergegeven volume is indicatief. Na afloop van de verbruiksperiode kan het exacte verbruik worden vastgesteld. U kunt dan ook een definitief certificaat aanvragen.
**EuropeseWind is alleen afkomstig uit Europese landen die aangesloten zijn bij de Europese organisatie voor energiecertificatie AIB.

Let’s do some math

Total energy consumption
16,013,299 kWh → ~16 GWh

Average household consumption per capita
1.6MWh → 0.0016 GWh

1 datacenter == ~10k people

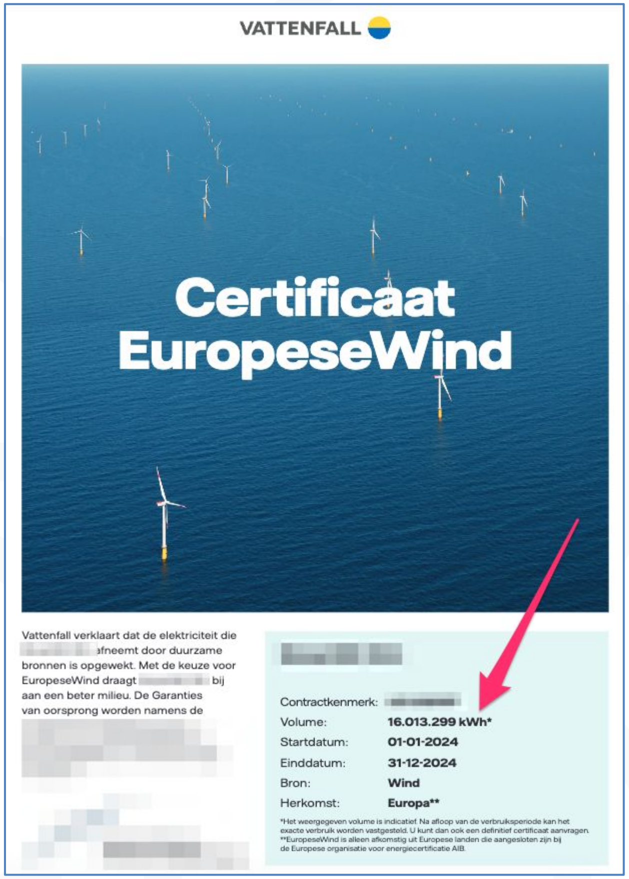

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Electricity_and_heat_statistics#Consumption_of_electricity_per_capita_in_the_household_sector

The S2 group & the Green Lab

Software Carbon Intensity

International Standard

ISO/IEC 21031:2024

Information technology — Software Carbon Intensity (SCI) specification

Edition 1
2024-03

Reference number
ISO/IEC 21031:2024

© ISO 2024

The S2 group & the Green Lab

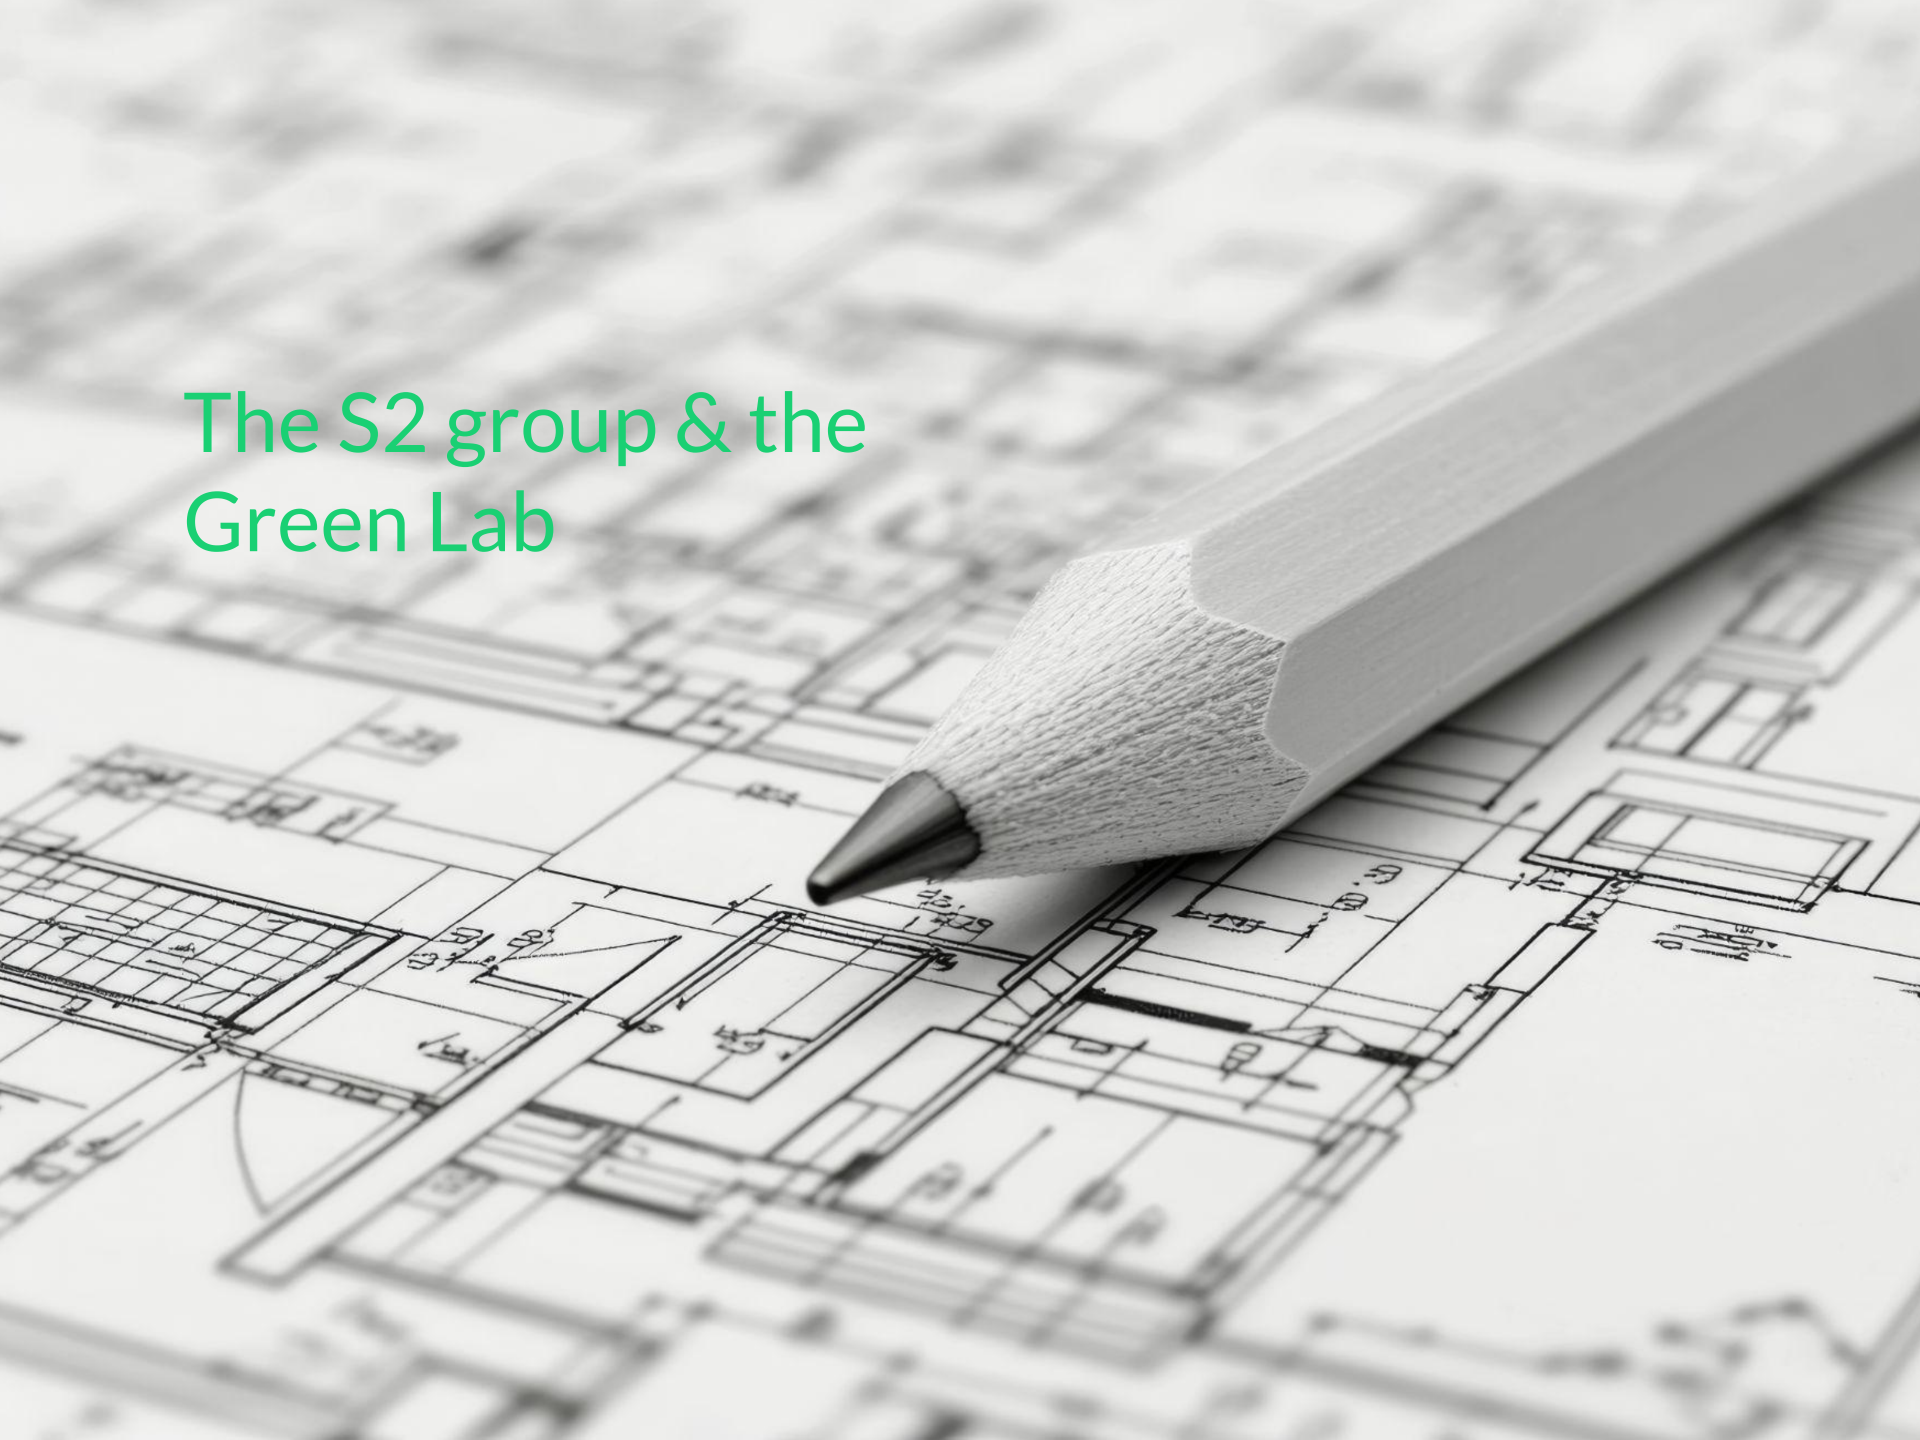

The S2 group @VU Amstedam

I/O Magazine > December 2020

Group portrait

Sustainability in software engineering

By Bennie Mols Images Ivar Pel

GROUP PASSPORT

RESEARCH FIELD

Software and Sustainability: creating software engineering knowledge that makes software better, smarter, and more sustainable; Software architecture; Software design and modelling; Software quality assessment.

EMPLOYEES

1 Professor
9 assistant and associate professors
6 postdocs and PhD students
1 junior lecturer
2 research assistants.

INSTITUTION

Vrije Universiteit Amsterdam

FACILITIES

Green Lab: a lab for experiments with the optimisation of the energy efficiency and other quality properties of software deployed on servers and mobile devices like smartphones.

WEBSITES

http://patricialago.nl
http://s2group.cs.vu.nl

12

13

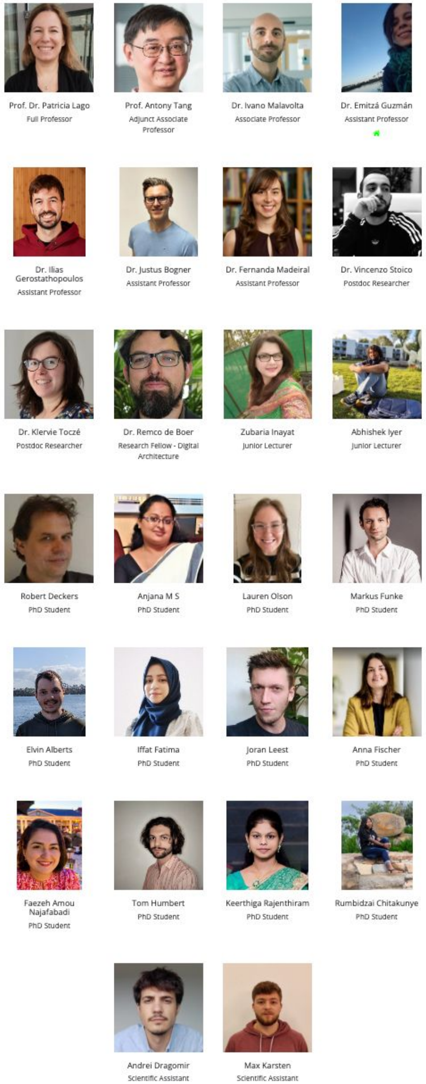

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

[I/O Magazine, Issue 3, Dec. 2020]

Master track: Software Engineering and Green IT

OVERVIEW COMPUTER SCIENCE: JOINT DEGREE WITH UVA

LANGUAGE OF INSTRUCTION	English
DURATION	2 years
TUITION FEE	Tuition fees
APPLICATION DEADLINE	September intake 1 April for EU/EEA and non-EU/EEA students* 1 June for holders of a Dutch bachelor's degree (with a Dutch or EU/EEA nationality). * EU/EEA students (exept for Dutch students) with an international degree who do not need housing services through VU Amsterdam can still apply until 1 June. February intake 1 November for non-EU/EEA students 1 December for EU/EEA students
START DATE	1 September and 1 February
STUDY TYPE	Full-time
SPECIALIZATIONS	- Big Data Engineering - Computer Systems Security - Foundations of Computing and Concurrency - Internet and Web Technology - Parallel Computing Systems - Software Engineering and Green IT

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

The Green Lab

https://arxiv.org/abs/2407.05689

A MASTER COURSE

Students measure real software products

A PLATFORM

Our infrastructure for experimenting on software

- energy efficiency
- performance
- ...

A COLLABORATION PLATFORM

Industry-driven experiments

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Green Lab infrastructure

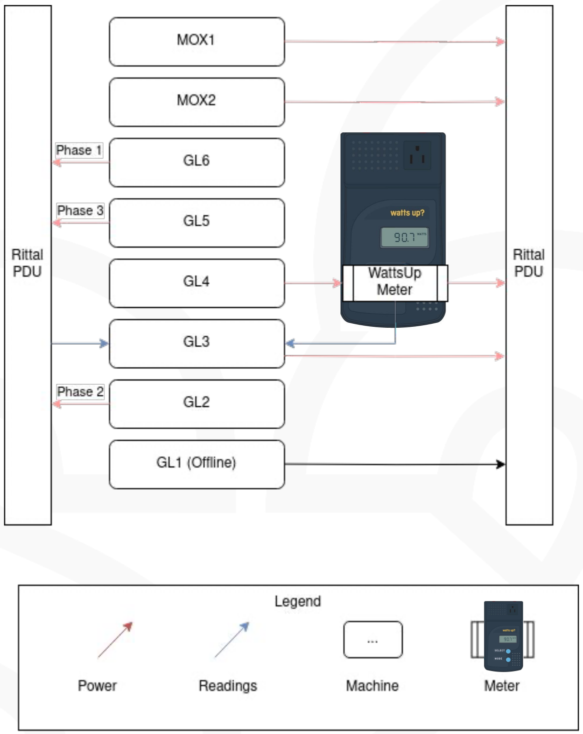

https://arxiv.org/abs/2407.05689

ID	HD	RAM	CPU (Intel Xeon)	Operating System
MOX1	36Tb	196Gb	Silver 4112@2.60GHz	Debian 9
MOX2	36Tb	384Gb	Silver 4208@2.10GHz	Debian 11
GL6	36Tb	384Gb	Silver 4208@2.10GHz	Ubuntu 20.04
GL5	36Tb	384Gb	Silver 4208@2.10GHz	Ubuntu 22.04
GL4	1Tb	16Gb	E5335@2.00GHz	Ubuntu 22.04
GL2	1Tb	32Gb	E3-1231@3.40GHz	Ubuntu 22.04
GL3	126Gb	8Gb	E5345@2.33GHz	Ubuntu 22.04

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Software tools

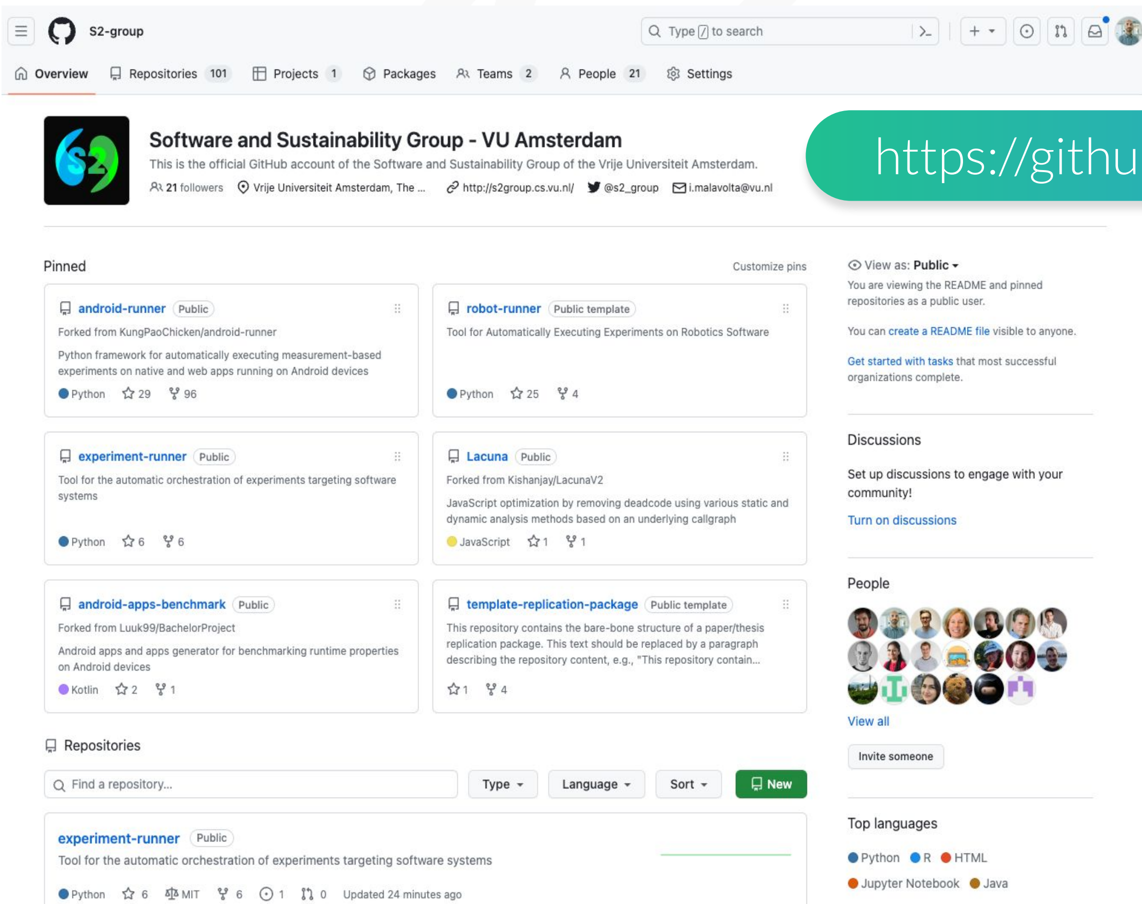

https://github.com/S2-group

SOFTWARE AND
SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Experiment Runner

A framework to **automatically** execute measurement-based experiments

https://github.com/S2-group/experiment-runner

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Anatomy of an experimental run

Guldner et al. Development and evaluation of a reference measurement model for assessing the resource and energy efficiency of software products and components-Green Software Measurement Model (GSMM). Future Generation Computer Systems, 155, pp. 402-418, 2024.

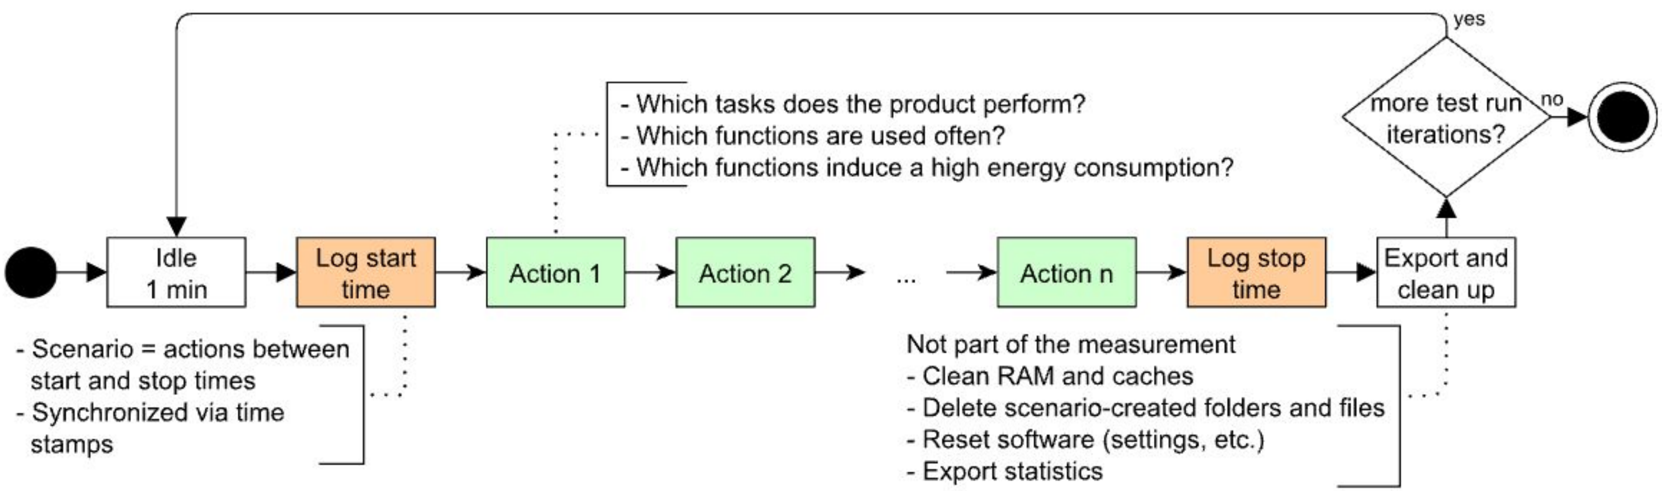

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Runners as a learning platform

Ivano Malavolta, Vincenzo Stoico, Patricia Lago. **Ten Years of Teaching Empirical Software Engineering in the context of Energy-efficient Software**, 2024. https://arxiv.org/abs/2407.05689

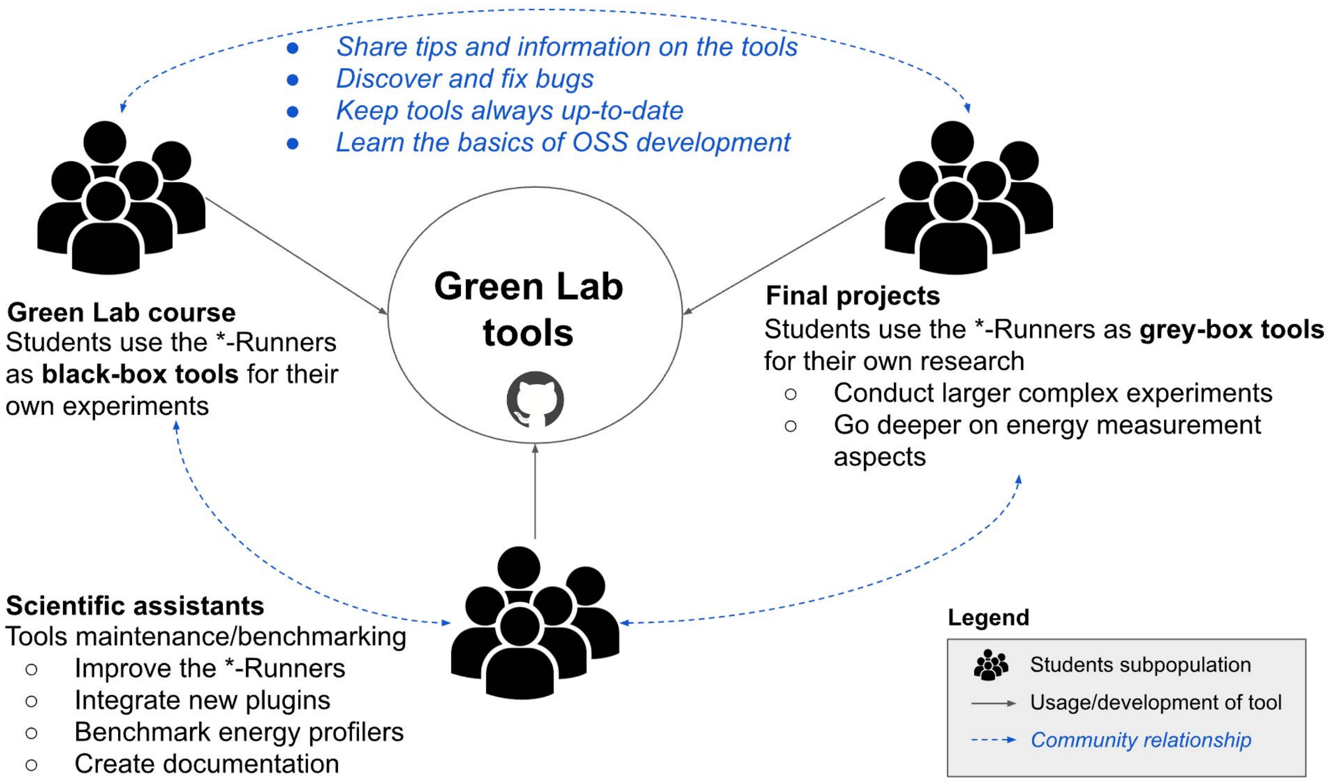

Software Carbon Intensity

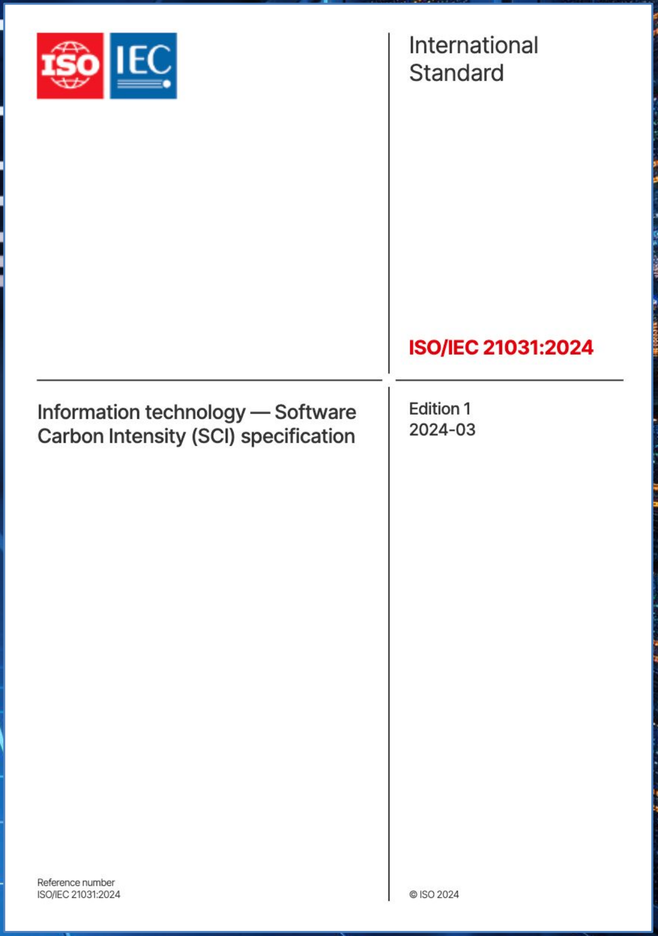
International Standard

ISO/IEC 21031:2024

Information technology — Software Carbon Intensity (SCI) specification

Edition 1
2024-03

Reference number
ISO/IEC 21031:2024

Basic concepts: Reducing carbon emissions

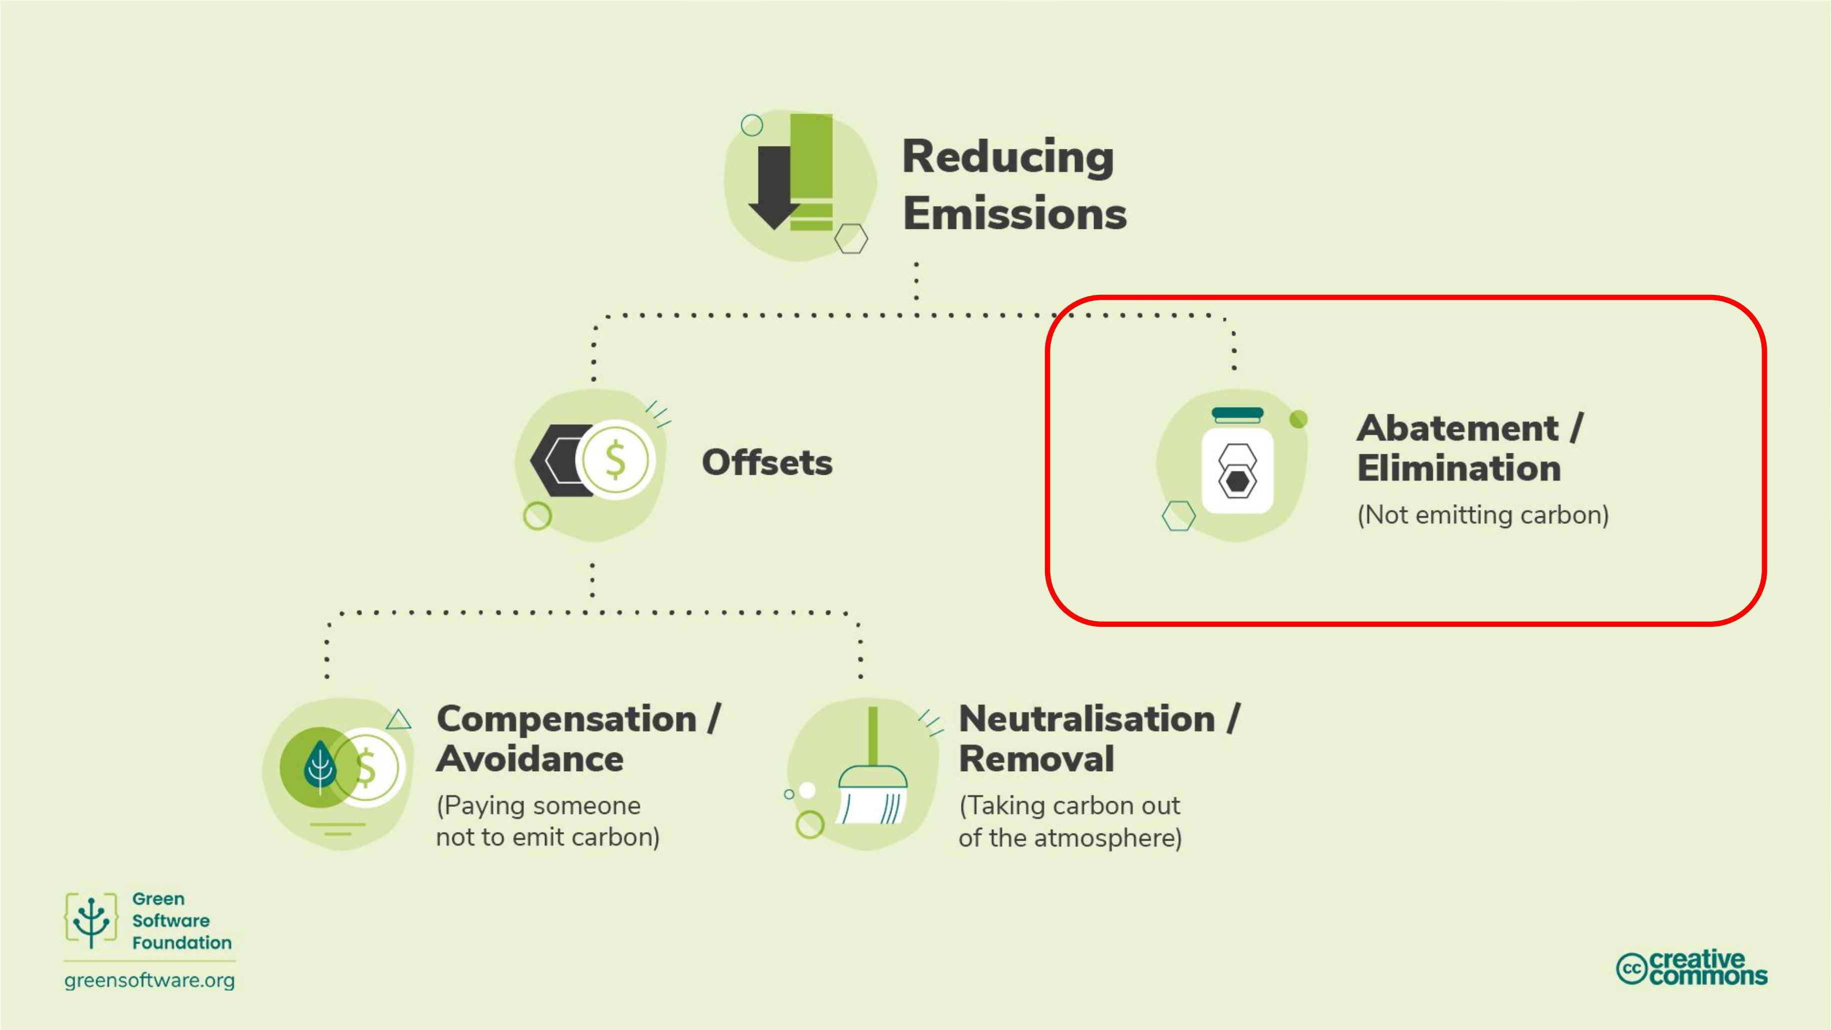

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Basic concepts: Strategies for greener software

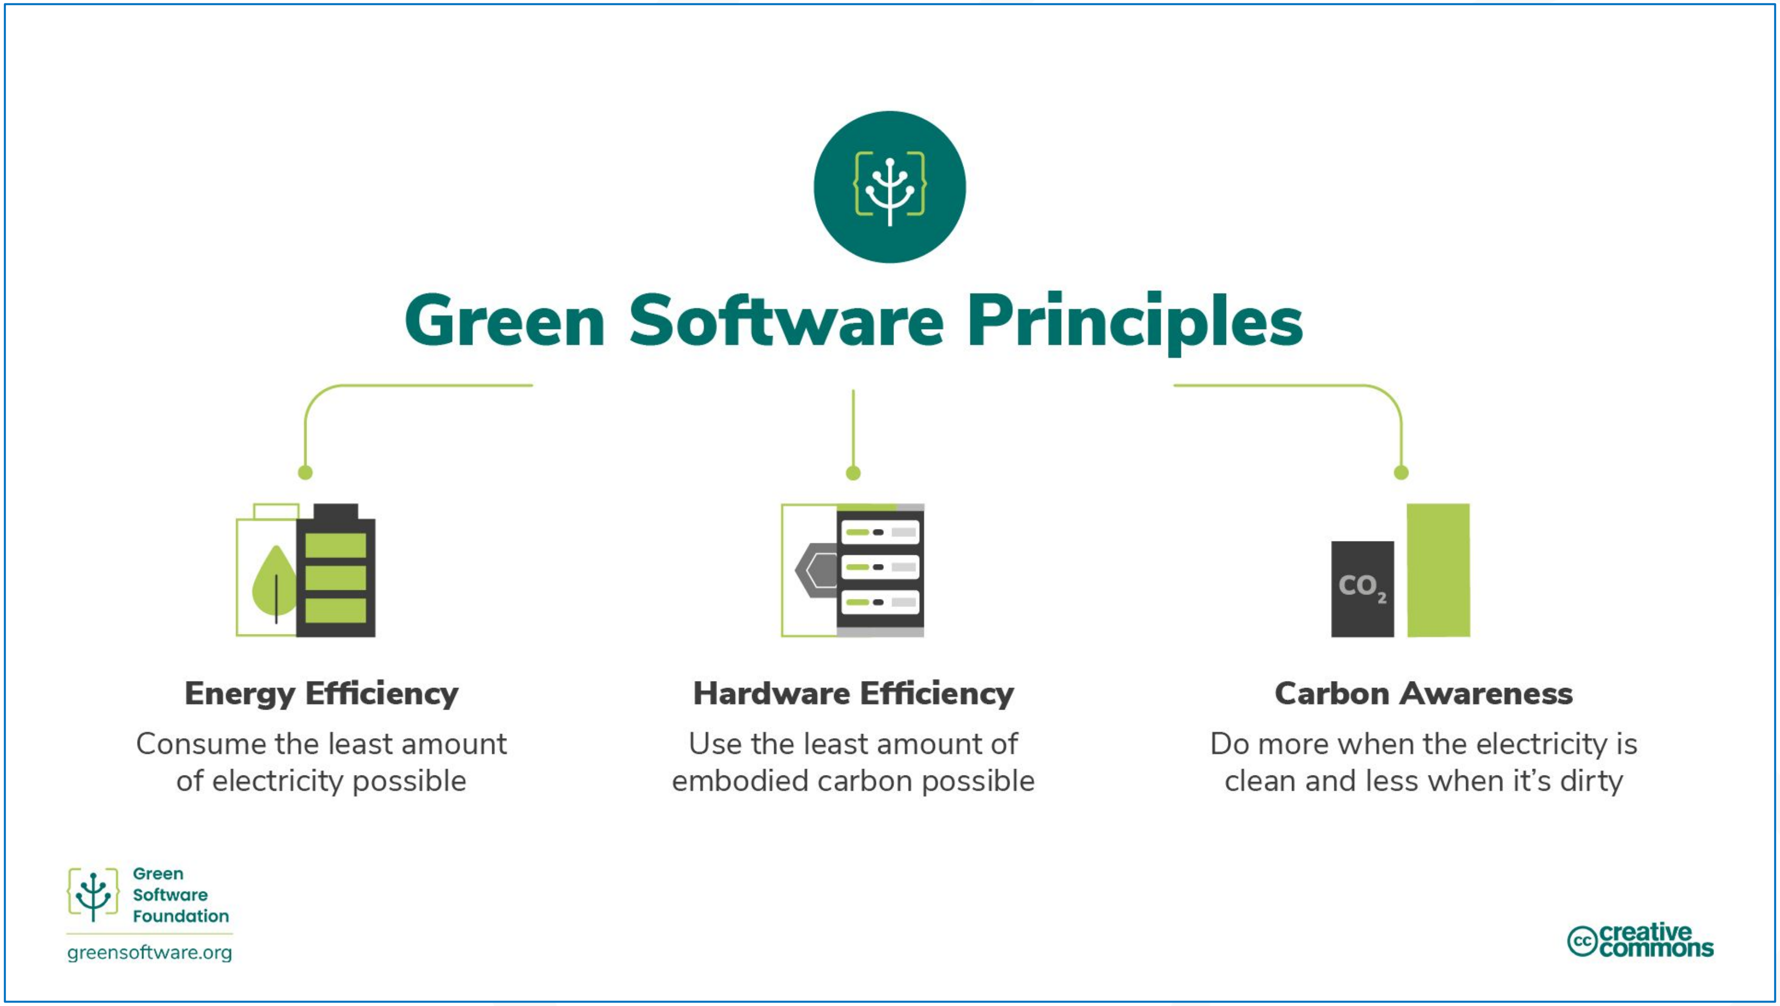

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Basic concepts: Carbon intensity

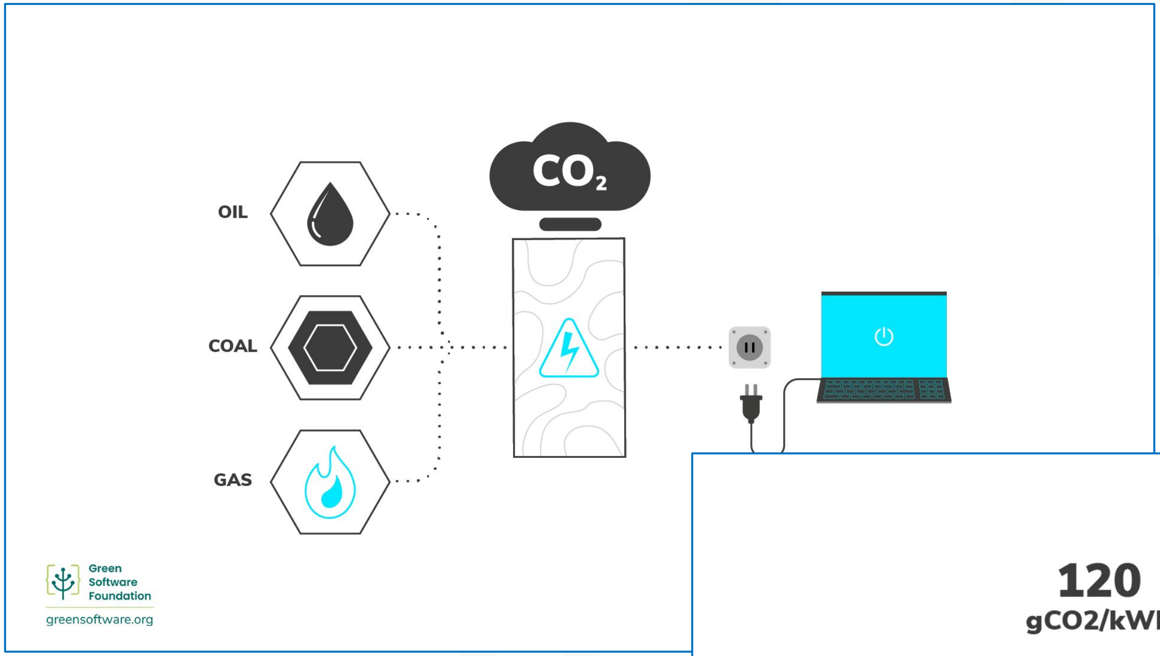

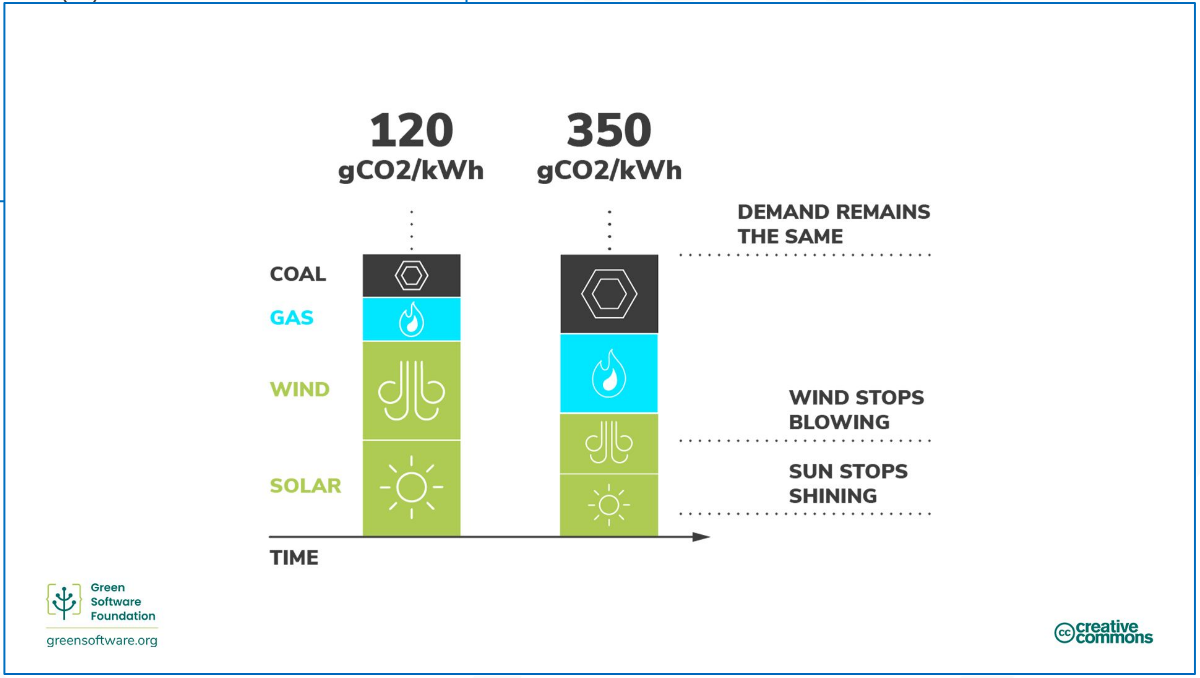

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Basic concepts: Carbon intensity

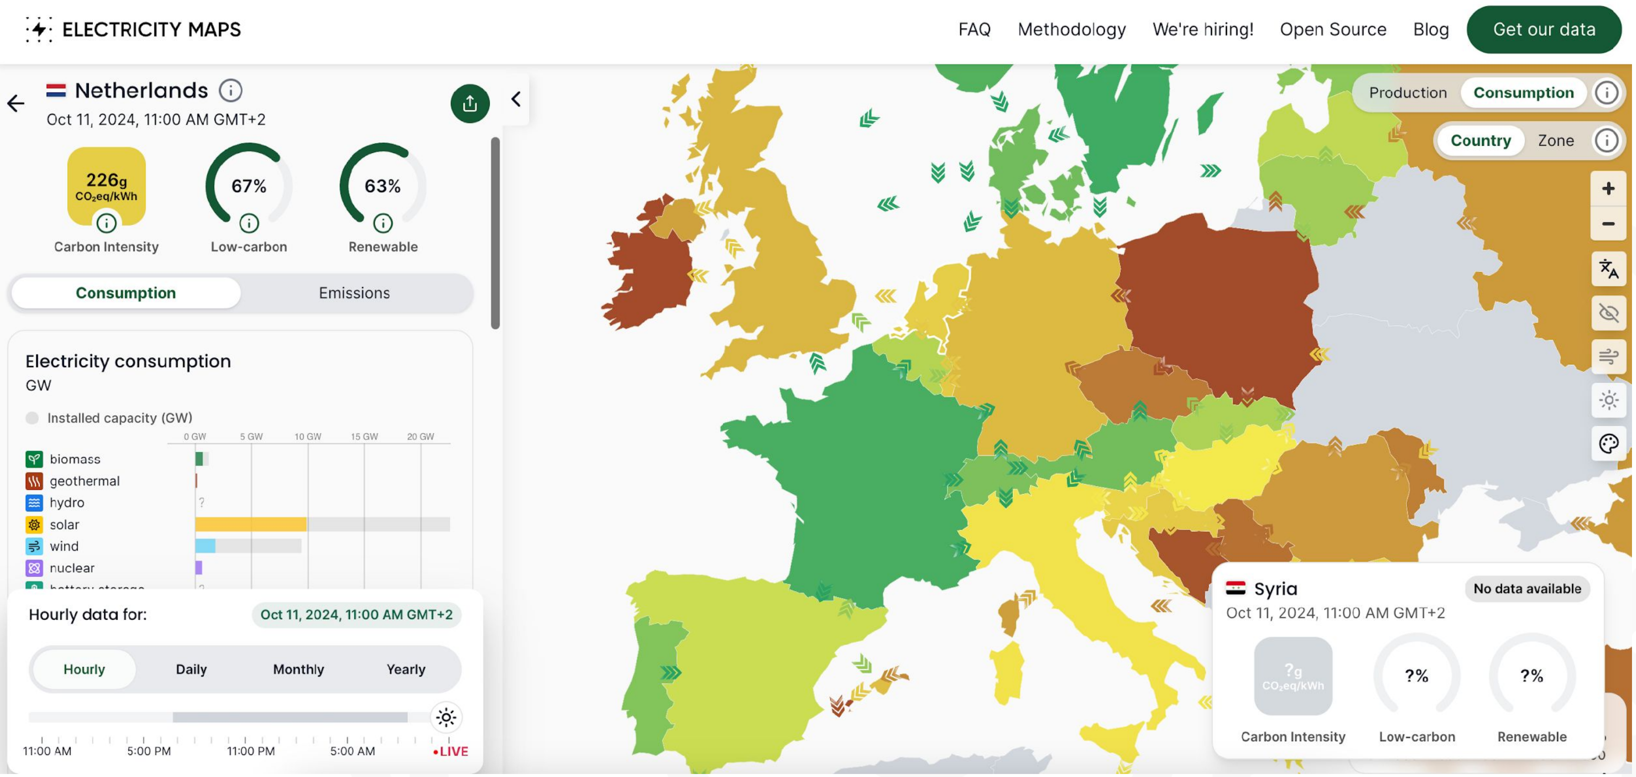

https://app.electricitymaps.com/zone/NL
https://www.electricitymaps.com/methodology

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Basic concepts: embodied emissions

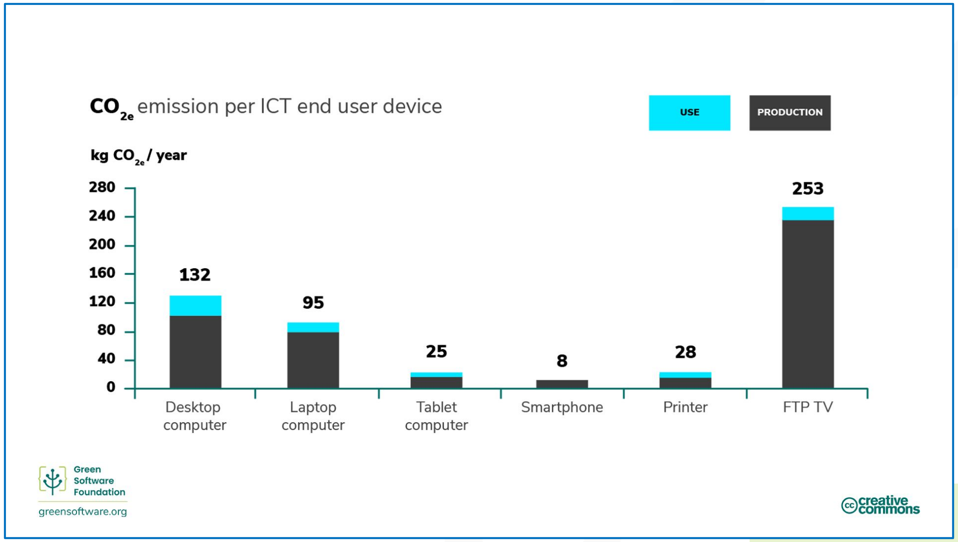

Carbon amortization + utilization

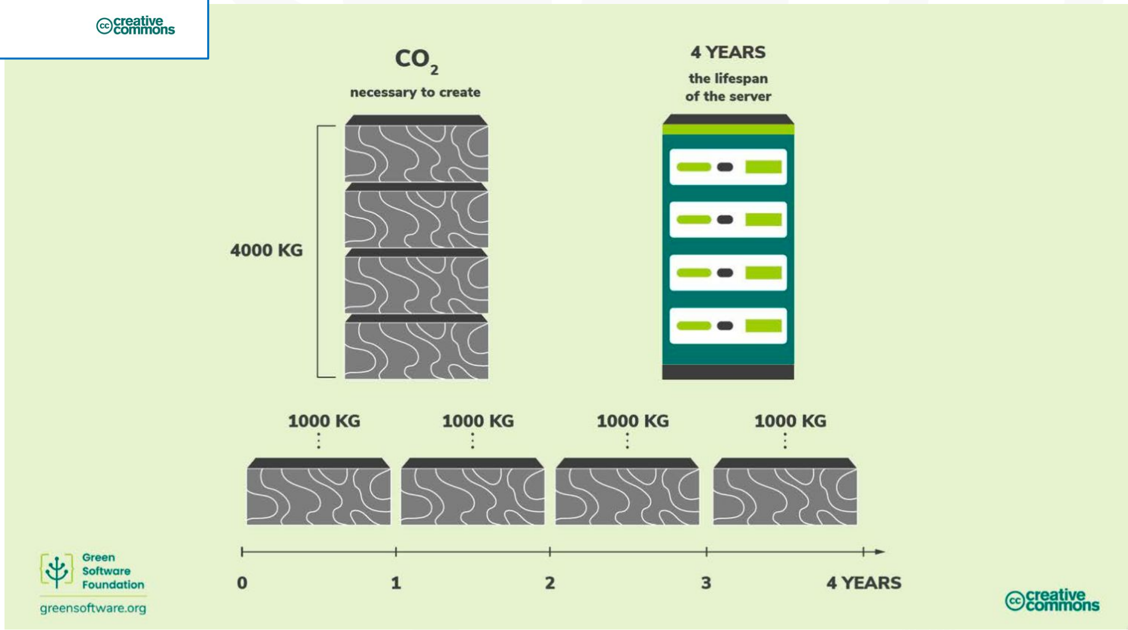

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

The Software Carbon Intensity formula

Carbon emitted per kWh of energy, gCO2/kWh

Carbon emitted through the hardware that the software is running on

$$SCI = ((E * I) + M) \text{ per } R$$

Energy consumed by software in kWh

Functional Unit; this is how software scales, for example per user or per device

Reporting in research

A CO2 EMISSION CALCULATION

In order to calculate the CO_2-emissions of this experiment, the power consumption of the Raspberry Pi and the Laptop must first be determined and also the average CO_2-emissions in grams per kWh in the country where the experiment is being carried out. The power consumption of the Lenovo - ThinkPad T14s Gen 2 is **6.3 Watt**[14] in idle mode. For the Raspberry Pi it is **1.9 Watt**[15] in idle mode. Finally the average emission per kWh for the Netherlands is **268.48 g per kWh**[16] in the year 2023. Now the CO_2-emissions are calculated according to the SCI guide [17]. The formula for calculating the Software Carbon Intensity (SCI) is given by:

$$SCI = (E \times I) + \frac{M}{R} \tag{3}$$

- **SCI**: Software Carbon Intensity, which quantifies the carbon emissions associated with the use of software over a given time period.
- **E** (Energy consumption): Measured in kilowatt hours (kWh), it represents the total energy consumption. This includes:
  - Total energy consumption of the experiment converted to kWh.
  - Laptop idle energy usage times the total runtime of the experiment in kWh.
  - Raspberry Pi idle energy usage times the total runtime of the experiment in kWh.
- **I** (Emissions factors): Represents the carbon emissions per kWh of energy consumption for the Netherlands.
- (**M** (Embodied emissions data): Refers to the emissions produced during the manufacture and transportation of hardware (e.g., servers, laptops, and mobile devices) used to run the software. **IGNORED** for this calculation because not applicable.)
- (**R** (Reference or lifespan): Represents the reference amount, such as the useful lifespan or total usage period of the hardware. It normalizes the embodied emissions over the time period. **IGNORED** for this calculation because not applicable.)

The result is calculated automatically by the R script co2_calculation.R[18] using the run_table.csv data and the formula Equation 3 where the variables M and R are ignored because they are not applicable for this experiment. The result for this experiment is a Software Carbon Intensity (SCI) of **245.27 g CO2e**.

[14] https://static.lenovo.com/ww/docs/regulatory/eco-declaration/eco-thinkpad-t14s-gen-2-amd.pdf
[15] https://www.pidramble.com/wiki/benchmarks/power-consumption
[16] https://www.statista.com/statistics/1290441/carbon-intensity-power-sector-netherlands/
[17] https://sci-guide.greensoftware.foundation/
[18] https://github.com/royderegt/experiment-runner/blob/master/r-scripts/co2_calculation.R

Running example

Goal: to calculate the SCI of a software application called "MovieRecommender" running on a Google Cloud VM

Scaling factor: #API requests
Average monthly requests: 20k
Region: US-East*
VM config: e2-standard-4

E2 standard | E2 high-mem | E2 high-cpu | E2 shared-core

E2 standard machine types have 4 GB of system memory per vCPU.

Machine types	vCPUs	Memory (GB)	Max number of Persistent Disk (PDs)†	Max total PD size (TiB)	Local SSD	Maximum egress bandwidth (Gbps)‡
`e2-standard-2`	2	8	128	257	No	4
`e2-standard-4`	4	16	128	257	No	8
`e2-standard-8`	8	32	128	257	No	16
`e2-standard-16`	16	64	128	257	No	16
`e2-standard-32`	32	128	128	257	No	16

† Persistent Disk and Hyperdisk usage is charged separately from machine pricing.
‡ Maximum egress bandwidth cannot exceed the number given. Actual See Network bandwidth.

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Software boundary

Software boundary = all supporting infrastructure and systems that significantly contribute to the software's operation

Examples:
- compute resources
- storage
- networking equipment
- memory
- monitoring
- idle machines
- logging
- scanning
- build and deploy pipelines
- testing
- …

If evaluating a complex system, compute the SCI of one component at a time, then aggregate

MovieRecommender

Boundary:
- all services of the application
- the infrastructure for running them (vCPUs and server only)

We are leaving out: memory, disk , racks, cooling water resources, CI/CD pipelines, etc.

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Functional unit (R)

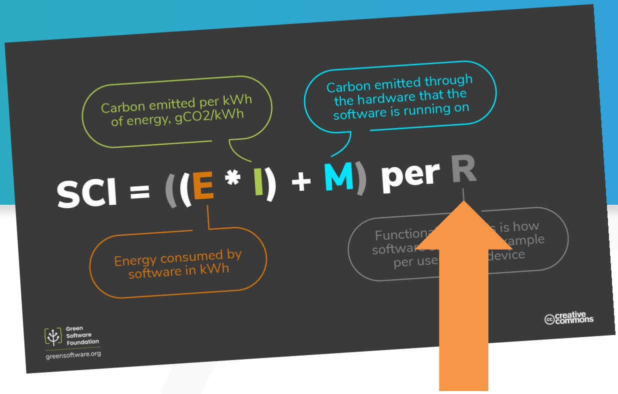

Functional unit defines how your software scales

<u>The goal of the SCI is to quantify how much carbon is emitted per one unit of R</u>

Examples:

- #API call/request
- Benchmarks executed
- #users
- Minute/time unit
- Devices
- Physical site
- Data volume
- Batch/Scheduled Jobs
- Database transactions

MovieRecommender | Functional unit: one API call

It must be **consistent** across all the components in the software boundary

Energy (E)

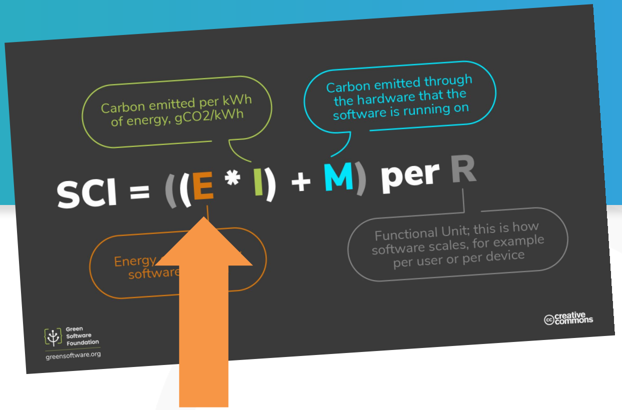

E is the energy consumed by a software system for a functional unit of work R

Unit of measure: kWh

Four main strategies to compute E:

1. Tool-based [**MEASURED**]
2. API-based techniques [**CALCULATED**]
3. Performance engineering techniques [**CALCULATED**]
4. Public datasets [**CALCULATED**]

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

E -- Tool based

In tool-based approaches you integrate your software with tools that are directly measuring the energy consumed for each functional unit R.

Examples of useful tools:

- PyJoular
- energibridge
- Scaphandre
- any RAPL wrapper
  but check your SCI boundary!
- etc.

Extensive list here

This is what we primarily do in the Green Lab course

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

E -- API-based techniques

You integrate your software with APIs that provide at runtime energy values for VM instances, based on CPU usage, instance type, location and duration of usage, etc.

Examples of available APIs:

- Climatiq
- CloudCarbonFootprint

Available APIs typically report Energy Carbon Intensity (E * I), not only E!
The SCI specification calls it Operational emissions O.

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Example of Climatiq request and response

```
curl --request POST \
--url https://api.climatiq.io/compute/v1/azure/instance \
--header "Authorization: Bearer $CLIMATIQ_API_KEY" \
--data '{
    "region": "uk_west",
    "instance": "h8",
    "duration": 24,
    "duration_unit": "h"
}'
```

https://www.climatiq.io/docs/api-reference/computing#vm-instance

((E * I) for memory

((E * I) for CPU

```
{
    "total_co2e": 0.7436,
    "total_co2e_unit": "kg",
    "memory_estimate": {
        "co2e": 0.1382,
        "co2e_unit": "kg",
        "co2e_calculation_method": "ar5",
        "co2e_calculation_origin": "source",
        "emission_factor": {
            "name": "Electricity supplied from grid",
            "activity_id": "electricity-supply_grid-source_supplier_
            "id": "efac3e9c-89c4-4ac0-9af6-d2bb428302d8",
            "access_type": "public",
            "source": "BEIS",
            "source_dataset": "Greenhouse gas reporting: conversion
            "year": 2024,
            "region": "GB",
            "category": "Electricity",
            "source_lca_activity": "electricity_generation",
            "data_quality_flags": []
        },
        "constituent_gases": {
            "co2e_total": 0.1382,
            "co2e_other": null,
            "co2": 0.1368,
            "ch4": 0.00002143,
            "n2o": 0.000003071
        },
        "activity_data": {
            "activity_value": 0.6675,
            "activity_unit": "kWh"
```

```
    "cpu_estimate": {
        "co2e": 0.1065,
        "co2e_unit": "kg",
        "co2e_calculation_method": "ar5",
        "co2e_calculation_origin": "source",
        "emission_factor": {
            "name": "Electricity supplied from grid",
            "activity_id": "electricity-supply_grid-source_supplier_
            "id": "efac3e9c-89c4-4ac0-9af6-d2bb428302d8",
            "access_type": "public",
            "source": "BEIS",
            "source_dataset": "Greenhouse gas reporting: conversion
            "year": 2024,
            "region": "GB",
            "category": "Electricity",
            "source_lca_activity": "electricity_generation",
            "data_quality_flags": []
        },
        "constituent_gases": {
            "co2e_total": 0.1065,
            "co2e_other": null,
            "co2": 0.1054,
            "ch4": 0.00001651,
            "n2o": 0.000002366
        },
        "activity_data": {
            "activity_value": 0.5143,
            "activity_unit": "kWh"
        },
```

E -- Performance engineering techniques

These techniques are generally about:

1) identifying which hardware components to consider in each machine (at least, CPU, GPU, memory)
2) collecting data from tech specs/datasets about the average power consumption P_i of each component i and summing up all of them into P
3) the application of the usual $E = P * t$ formula, where t is the total time window considered in the SCI calculation

Points of attention:

- Which hardware components to consider
  It depends on your SCI boundaries
- The sources you use for collecting the P_i values
  Example: Intel Xeon Platinum 8270 datasheet
- How to include utilization in your P (e.g., average CPU utilization)
  Example: see DEF formula

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

E -- Public datasets

You use public sources and references for energy estimates for computing resources

Examples of available datasets:

- Boavizta Cloud dataset
- Boavizta servers dataset
- Climatiq open data explorer
- Spec.org average Watts for CPUs

As an alternative, Cloud Jewels coefficients (created by the Sustainable Task Force at Etsy):

- `2.10 Wh per vCPUh [Server]`
- `0.89 Wh per TBh for HDD storage [Storage]`
- `1.52 Wh per TBh for SSD storage [Storage]`

Cloud Jewels coefficients are (very rough) fixed estimates, which can be used as final resort. See here for the details about the used method.

Example of calculation via a public dataset

MovieRecommender

Energy Carbon intensity (E * I) = (0.0013*24*30)*1000 = **936 gC02e**

https://www.climatiq.io/data/explorer

GCP (us-east-4) CPU LCA Activity: `use_phase` 0.0013 $kgCO_2e$/CPU-hour	CCF	2021	Virginia, US	NumberOver Time

Emission intensity for usage of a single CPU in kg CO2e per hour for the Google Cloud Platform data centers in the given location using the assumptions and grid emissions factors available in the source as of date accessed. It is assumed that vCPU utilization is 50%. The source does not clarify if the kgCO2e value is calculated using either IPCC Fourth Assessment Report (AR4) or IPCC Fifth Assessment Report (AR5) methodologies.

ACTIVITY ID	`cpu-provider_gcp-region_us_east_4`
ID	`523aa3b8-81d2-468b-bcf2-846f48d86266`
SOURCE	CCF
YEAR	2021
YEAR RELEASED	2021
REGION	Virginia, US (US-VA)
SECTOR	Information and Communication
CATEGORY	Cloud Computing - CPU
UNIT TYPE(S)	Number Over Time
EMISSION FACTORS	CO_2e 0.001331 kg/CPU-hour Data Quality
CO_2e CALCULATION METHOD	**Method applied:** AR4 **Methods supported:** AR4 **Origin:** Source

Energy carbon intensity (I)

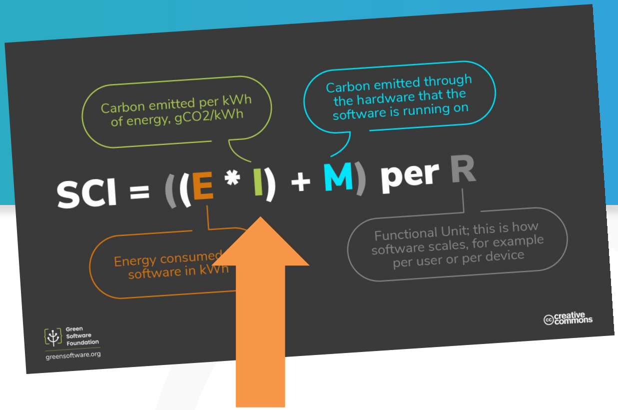

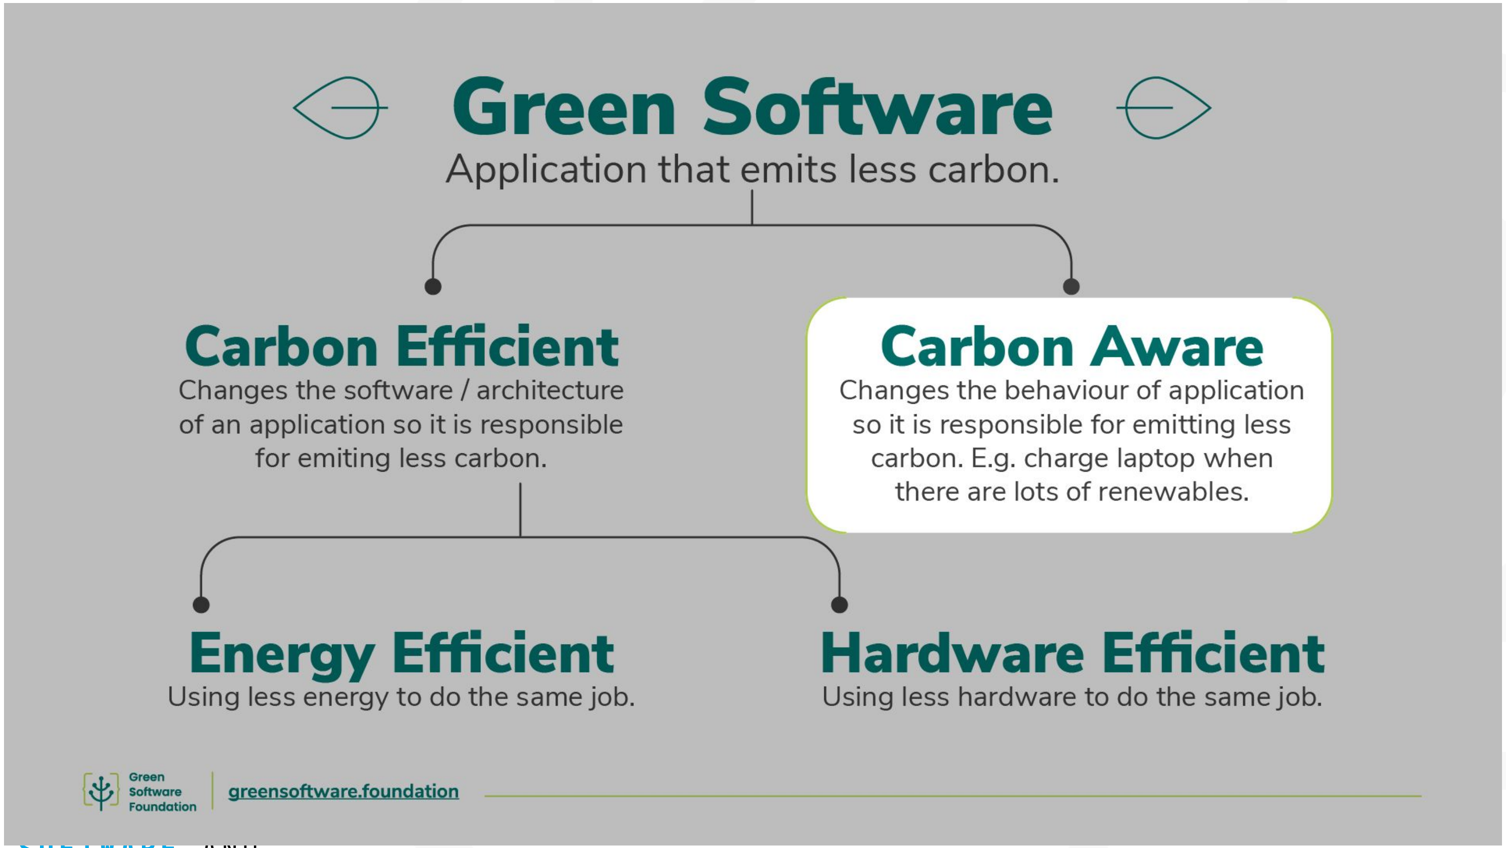

https://github.com/Green-Software-Foundation/carbon-aware-sdk

Energy carbon intensity (I)

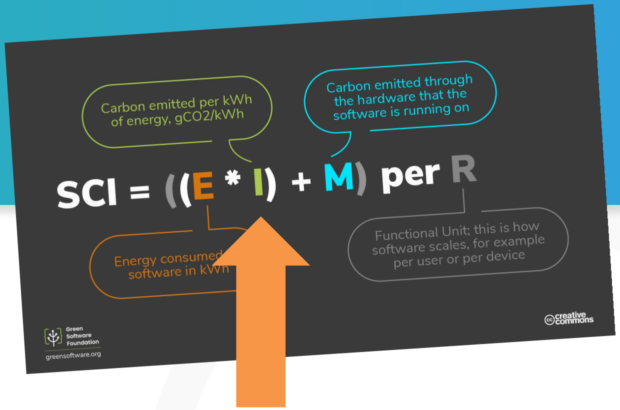

or better:
Location-Based Marginal Carbon Intensity

Carbon intensity = the measure of the greenhouse gas emissions associated with producing electricity

It is expressed in **gCO_2eq/kWh** - grams of carbon dioxide equivalents emitted per kilowatt hour of consumed electricity

No single method to calculate this, it depends primarily on the region where energy is consumed

Energy carbon intensity (I)

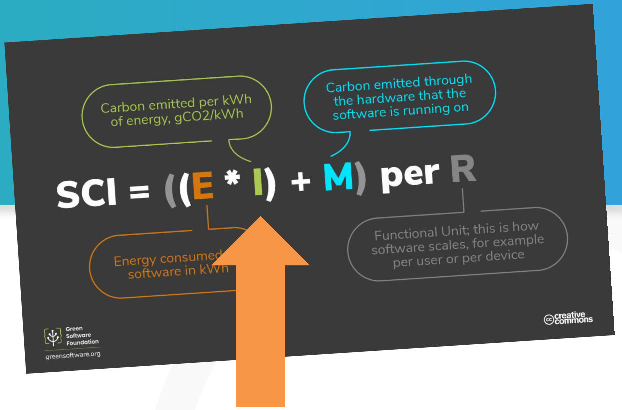

Two main strategies to follow to compute it:

- **API-based techniques**: your software uses real-time APIs and act based on the levels of carbon intensity
  - This is the preferred method
  - Examples: https://codecarbon.io
    https://app.electricitymaps.com
    https://github.com/Green-Software-Foundation/carbon-aware-sdk

- **Lookup open datasets**: you use historical databases
  - Example: https://ourworldindata.org/grapher/carbon-intensity-electricity

Embodied emissions (M)

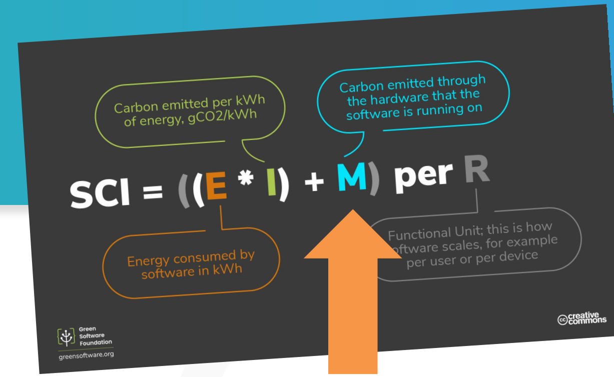

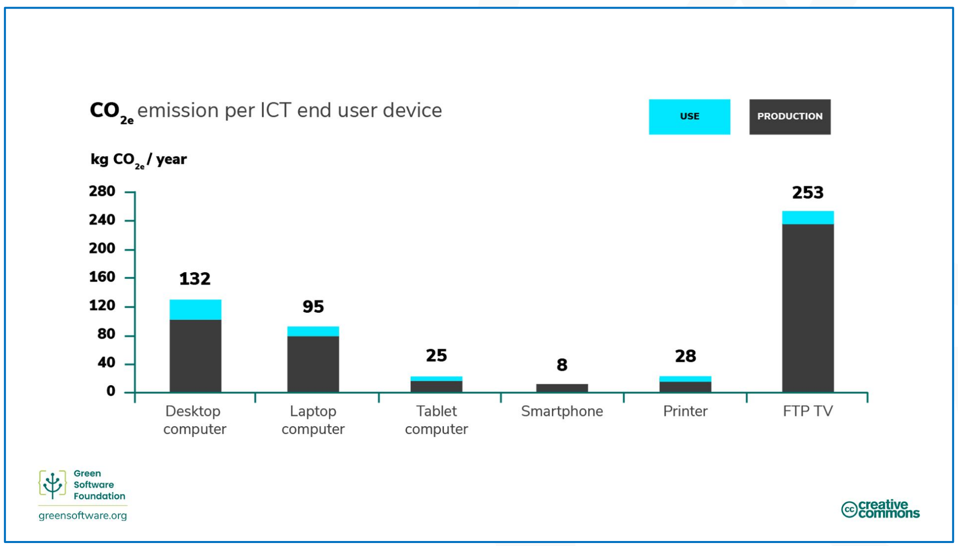

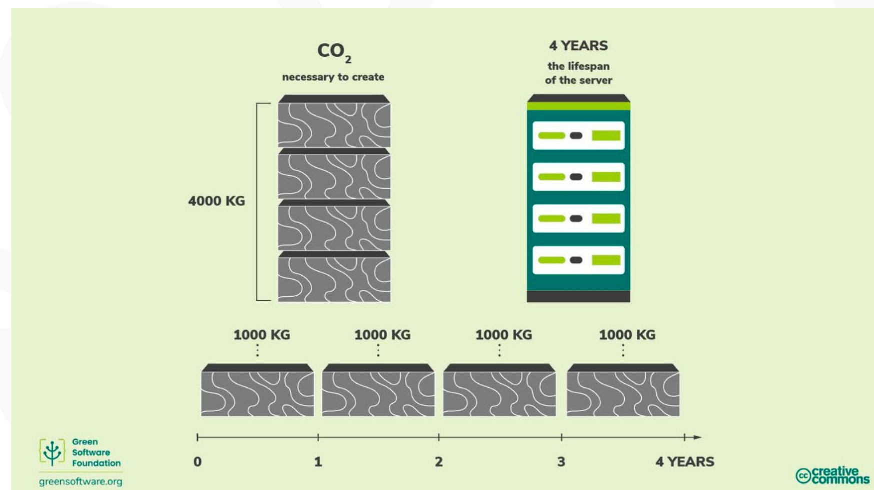

Carbon amortization

Embodied emissions (M)

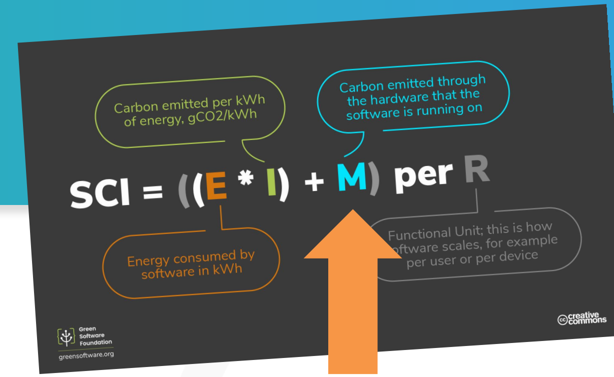

M = the <u>fraction</u> of the total embodied emissions of the device is allocated to the software in grams of carbon (gCO2eq)

$$M = TE * TS * RS$$

where:

- **TE = Total Embodied Emissions** = the sum of Life Cycle Assessment (LCA) emissions for all hardware components
- **TS = Time-share** = the share of the total life span of the hardware reserved for use by the software
- **RS = Resource-share** = the share of the total available resources of the hardware reserved for use by the software

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Embodied emissions (M)

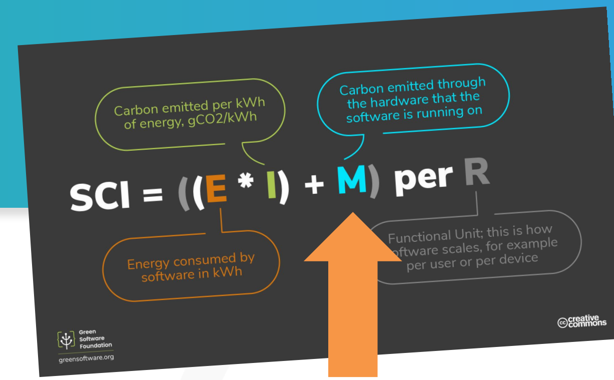

Time | Resources

$$M = TE * TS * RS$$

$$M = TE * (TiR/EL) * (RR/ToR)$$

where:

- TiR = **Time Reserved** = the length of time the hardware is reserved for use by the software in hours/days/months/years
- EL = **Expected Lifespan** = the anticipated time that the equipment will be installed in hours/days/months/years
- RR = **Resources Reserved** = the amount of resources reserved for use by the software
- ToR = **Total Resources** = the total amount of resources available

Final SCI – Example

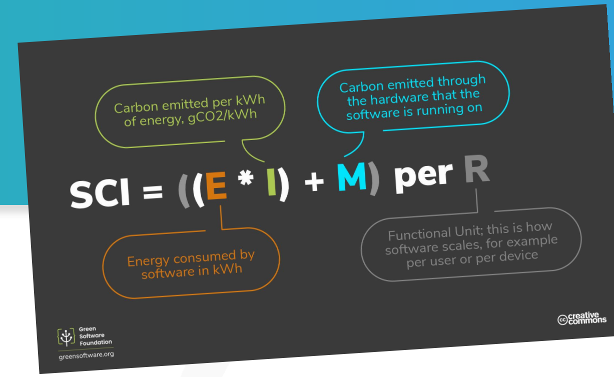

Here we compute the final value of SCI for the system under analysis

Time horizon = 1 month

R = 20,000 API requests

O = (E * I) = 936 gC0$_2$e

M = TE * (TiR/EL) * (RR/ToR)

TE = 1,230.3 kgCO$_2$e = 1,230,300 gCO$_2$e

TiR = 1 month

EL = 4 years = 48 months

RR = 4 vCPUs

ToR = 32 vCPUs

M = 1,230,300 * (1/48) * 4/32) = 3,203.90625 gCO$_2$e

SCI = (936 + 3,203.9) gC0$_2$e per month

= (936 + 3,203.9) / 20000 = 0.206995 gC0$_2$e per API call

SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP
VRIJE UNIVERSITEIT AMSTERDAM

Wrap-up

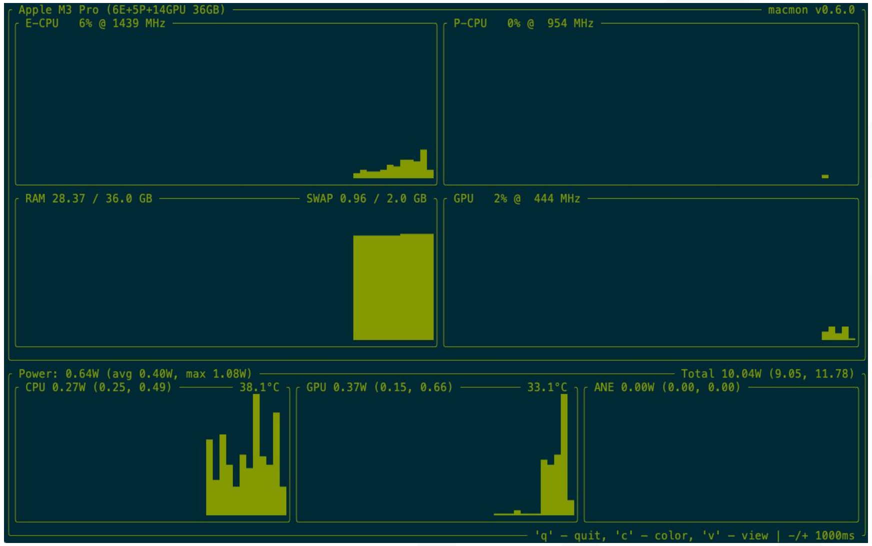
Apple M3 Pro (6E+5P+14GPU 36GB)
macmon v0.6.0
E-CPU 6% @ 1439 MHz
P-CPU 0% @ 954 MHz
RAM 28.37 / 36.0 GB
SWAP 0.96 / 2.0 GB
GPU 2% @ 444 MHz
Power: 0.64W (avg 0.40W, max 1.08W)
Total 10.04W (9.05, 11.78)
CPU 0.27W (0.25, 0.49)
38.1°C
GPU 0.37W (0.15, 0.66)
33.1°C
ANE 0.00W (0.00, 0.00)
'q' – quit, 'c' – color, 'v' – view | -/+ 1000ms

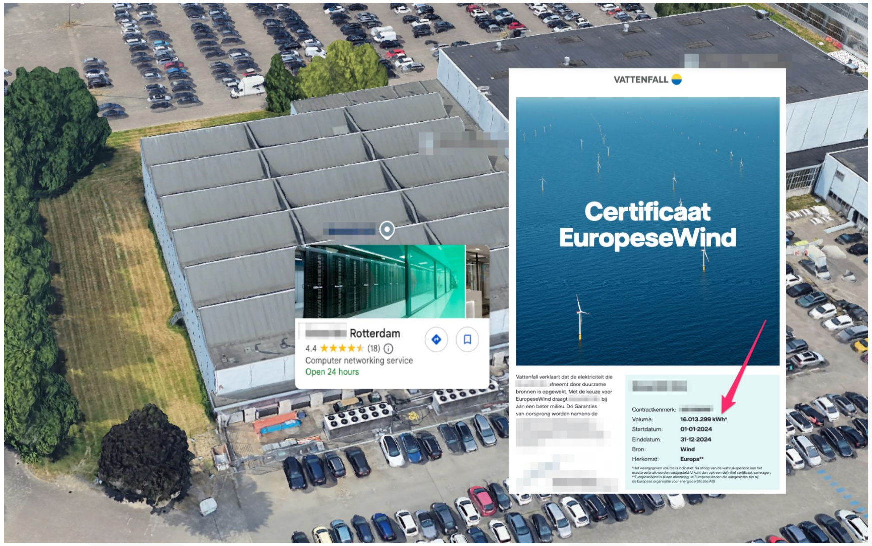
VATTENFALL
Certificaat EuropeseWind
Rotterdam
Computer networking service
Open 24 hours


The Green Lab
https://arxiv.org/abs/2407.05689
A MASTER COURSE
Students measure real software products
A PLATFORM
Our infrastructure for experimenting on software
energy efficiency
performance
...
COMPUTER SCIENCE: SOFTWARE ENGINEERING & GREEN IT
Software for a Sustainable Digital Society
A COLLABORATION PLATFORM
Industry-driven experiments
SOFTWARE AND SUSTAINABILITY
S2 RESEARCH GROUP

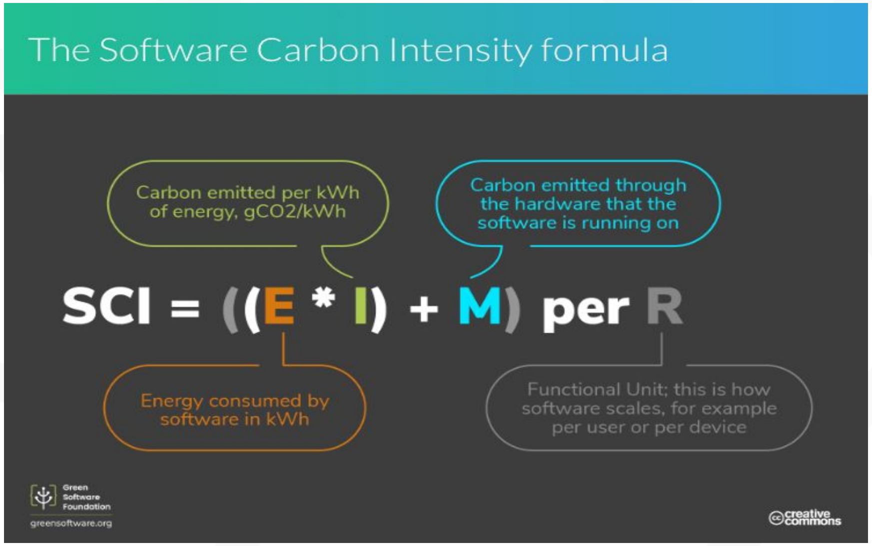
The Software Carbon Intensity formula
Carbon emitted per kWh of energy, gCO2/kWh
Carbon emitted through the hardware that the software is running on
SCI = ((E * I) + M) per R
Energy consumed by software in kWh
Functional Unit; this is how software scales, for example per user or per device
Green Software Foundation
greensoftware.org
creative commons

References

- Fibonacci implementations: https://realpython.com/fibonacci-sequence-python/#using-iteration-and-a-python-function
- Basics about green software

  https://learn.greensoftware.foundation
- Software Carbon Intensity Specification

  https://www.iso.org/standard/86612.html
- SCI official guide

  https://sci-guide.greensoftware.foundation
- SCI case studies (for inspiration)

  https://sci-guide.greensoftware.foundation/CaseStudies
- Example of usage of SCI in a scientific study:

  - YouTube video: https://www.youtube.com/watch?v=BOHKSK8GFYg

  - Study: https://hotcarbon.org/assets/2024/pdf/hotcarbon24-final109.pdf