Calculating User-Centric Carbon Footprints for HPC

Christian Wassermann[1], Mario Bielert[2], Gert Vanberg[1], Daniel Hackenberg[2], Christian Terboven[1], and Matthias S. Müller[1]

[1] IT Center, RWTH Aachen University, Germany
[2] ZIH, CIDS, TU Dresden, Germany

Sustainable HPC State of the Practice Workshop 2024 in Kobe, Japan

IT Center | RWTH AACHEN UNIVERSITY

Increasing Accessibility of Carbon Emissions in HPC

Increasing Accessibility of Carbon Emissions in HPC

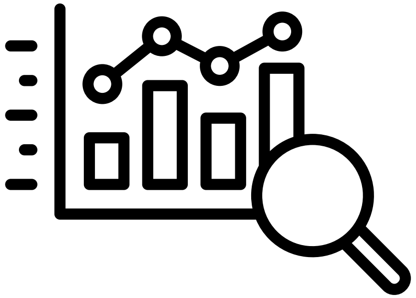

1. Measure

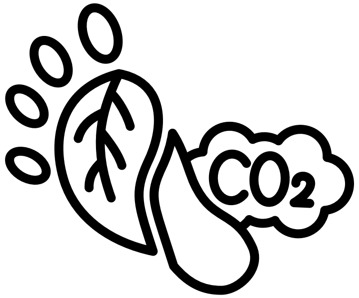

2. Optimize

Increasing Accessibility of Carbon Emissions in HPC

1. Measure

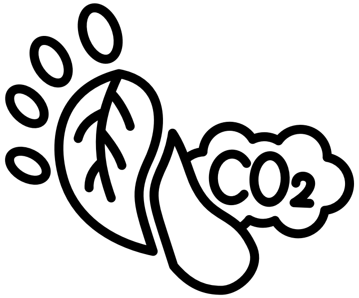

2. Optimize

Yearly power usage effectiveness (PUE)

or

Yearly carbon usage effectiveness (CUE)

[1]

IT Center | RWTH AACHEN UNIVERSITY

Increasing Accessibility of Carbon Emissions in HPC

1. Measure

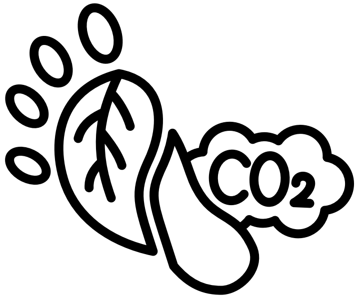

2. Optimize

Yearly power usage effectiveness (PUE)

or

Yearly carbon usage effectiveness (CUE)

Data Gap

[1]

Our Key Contributions

C1: **Carbon footprinting methodology** for HPC operators **and** HPC users

C2: Definition of **three carbon accounting levels** to facilitate comparability

C3: **Two case studies** analyzing the **variability of calculation parameters**

Agenda

- Carbon Footprint Formula
- Carbon Accounting Levels
- Case Studies at RWTH Aachen & TU Dresden
- Summary & Outlook

Agenda

- **Carbon Footprint Formula**
- Carbon Accounting Levels
- Case Studies at RWTH Aachen & TU Dresden
- Summary & Outlook

Carbon Footprint Formula

$$Carbon\ Footprint_{Start,\ End} = \overline{\qquad\qquad\qquad\qquad}$$

- $Carbon\ Footprint_{Start,\ End}$: Carbon emissions in $\frac{gCO_2e}{coreh}$

Carbon Footprint Formula

$$Carbon\ Footprint_{Start,\ End} = \frac{IT\ Energy * PUE\ \ *}{}$$

- $Carbon\ Footprint_{Start,\ End}$: Carbon emissions in $\frac{gCO_2e}{coreh}$
- $IT\ Energy$: Energy consumption of cluster (subcomponent)
- PUE: Interim PUE (iPUE) for energy overhead, e.g., of cooling

Carbon Footprint Formula

$$Carbon\ Footprint_{Start,\ End} = \frac{IT\ Energy * PUE\ * \langle \overrightarrow{Energy\ Mix},\ \overrightarrow{Emission\ Factors} \rangle}{}$$

- $Carbon\ Footprint_{Start,\ End}$: Carbon emissions in $\frac{gCO_2e}{coreh}$
- $IT\ Energy$: Energy consumption of cluster (subcomponent)
- PUE: Interim PUE (iPUE) for energy overhead, e.g., of cooling
- $\overrightarrow{Energy\ Mix}$: Portions of energy sources for electricity generation
- $\overrightarrow{Emission\ Factors}$: Carbon intensities for life-cycle emissions of energy sources
  - Fixed coefficients derived from large studies, e.g., by Intergovernmental Panel on Climate Change (IPCC) [3]

Carbon Footprint Formula

$$Carbon\ Footprint_{Start,\ End} = \frac{IT\ Energy * PUE\ * \langle \overrightarrow{Energy\ Mix},\ \overrightarrow{Emission\ Factors} \rangle}{Core\ Hours}$$

- $Carbon\ Footprint_{Start,\ End}$: Carbon emissions in $\frac{gCO_2e}{coreh}$
- $IT\ Energy$: Energy consumption of cluster (subcomponent)
- PUE: Interim PUE (iPUE) for energy overhead, e.g., of cooling
- $\overrightarrow{Energy\ Mix}$: Portions of energy sources for electricity generation
- $\overrightarrow{Emission\ Factors}$: Carbon intensities for life-cycle emissions of energy sources
  - Fixed coefficients derived from large studies, e.g., by Intergovernmental Panel on Climate Change (IPCC) [3]
- $Core\ Hours$: Available core hours
  - With one core hour corresponding to one physical CPU core being used for one hour

Agenda

- Carbon Footprint Formula
- **Carbon Accounting Levels**
- Case Studies at RWTH Aachen & TU Dresden
- Summary & Outlook

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary					
intermediate					
advanced					

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power				
intermediate	continuous integrated energy				
advanced	continuous integrated energy				

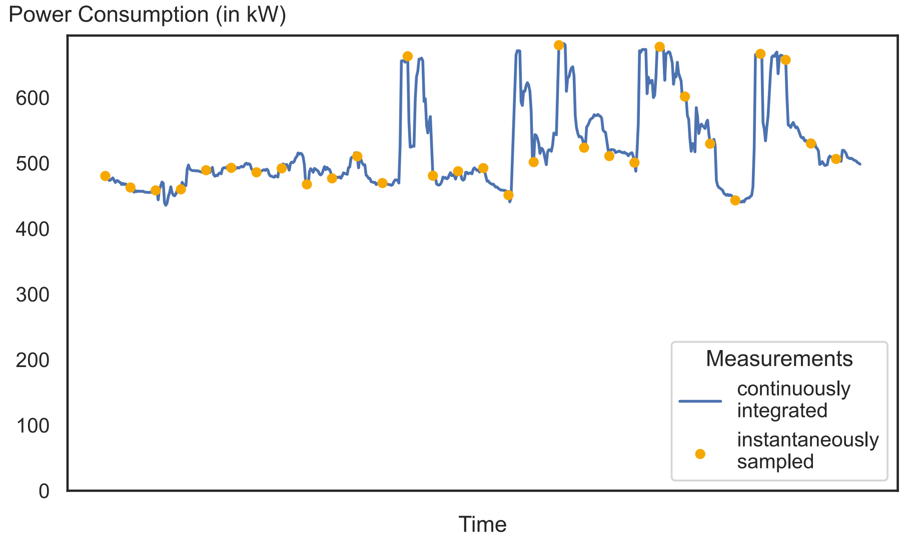

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only			
intermediate	continuous integrated energy	aggregate, compute-only			
advanced	continuous integrated energy	separate, all contributing			

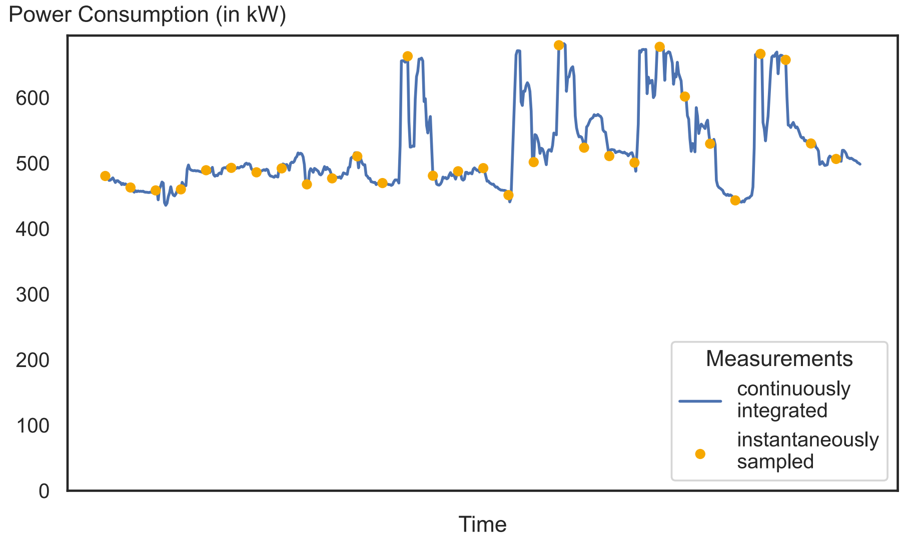

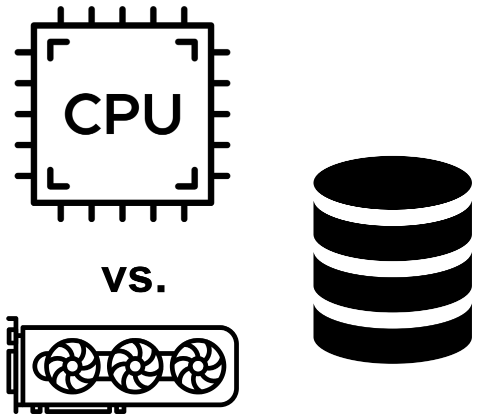

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated		
intermediate	continuous integrated energy	aggregate, compute-only	measured		
advanced	continuous integrated energy	separate, all contributing	measured		

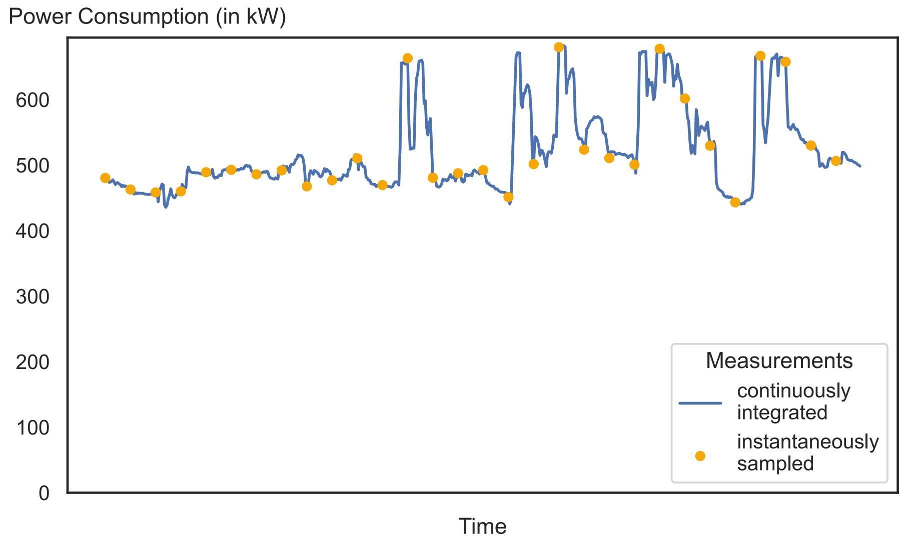

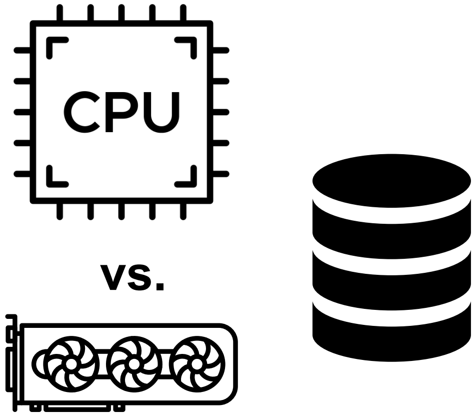

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	
advanced	continuous integrated energy	separate, all contributing	measured	usable	

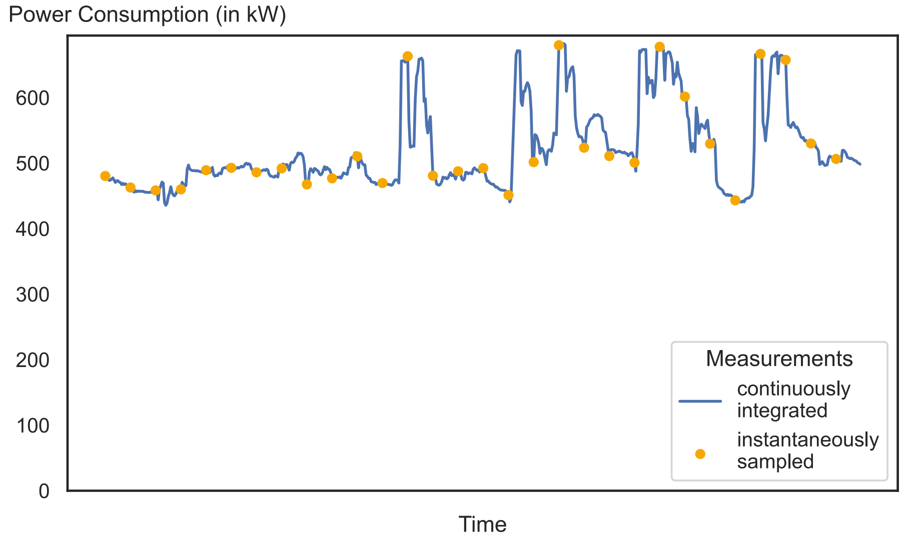

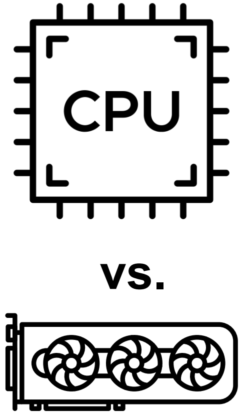

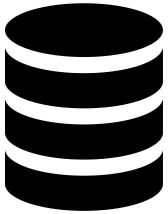

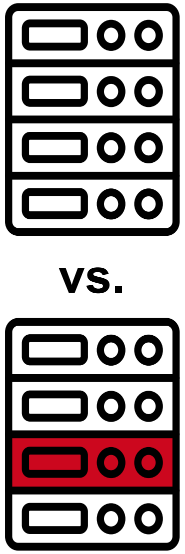

Christian Wassermann, Mario Bielert, Gert Vanberg, Daniel Hackenberg, Christian Terboven, Matthias S. Müller

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	yearly
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	monthly
advanced	continuous integrated energy	separate, all contributing	measured	usable	hourly

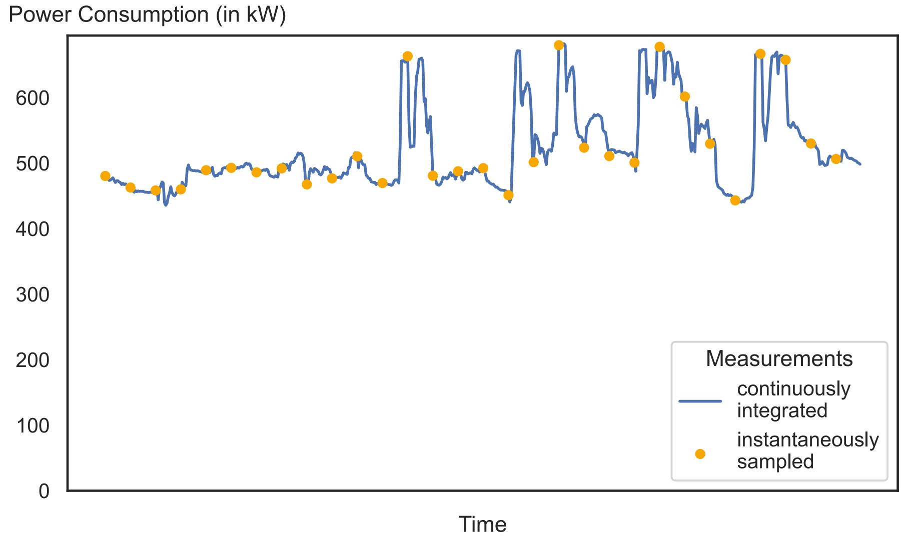

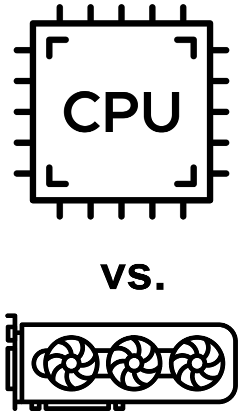

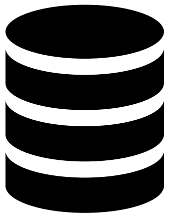

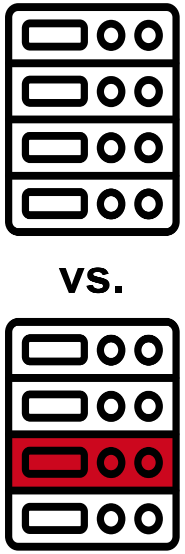

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	yearly
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	monthly
advanced	continuous integrated energy	separate, all contributing	measured	usable	hourly

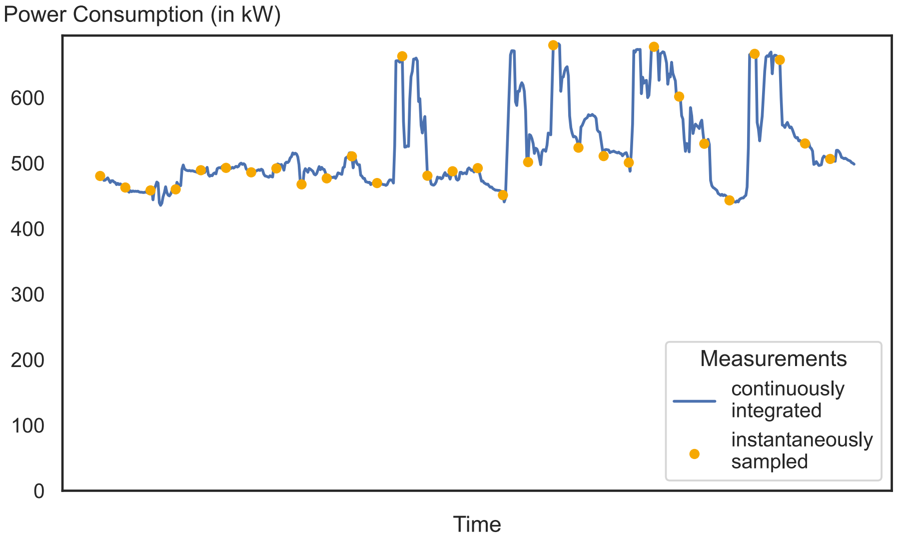

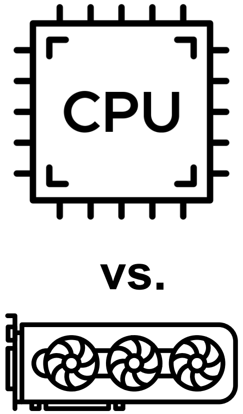

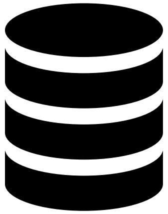

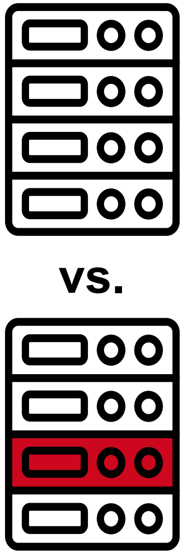

- The **elementary** level **has a low barrier to entry**

Calculating User-Centric Carbon Footprints for HPC
Christian Wassermann, Mario Bielert, Gert Vanberg, Daniel Hackenberg, Christian Terboven, Matthias S. Müller
Sustainable HPC State of the Practice Workshop 2024 | 24.09.2024

IT Center RWTH AACHEN UNIVERSITY

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	yearly
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	monthly
advanced	continuous integrated energy	separate, all contributing	measured	usable	hourly

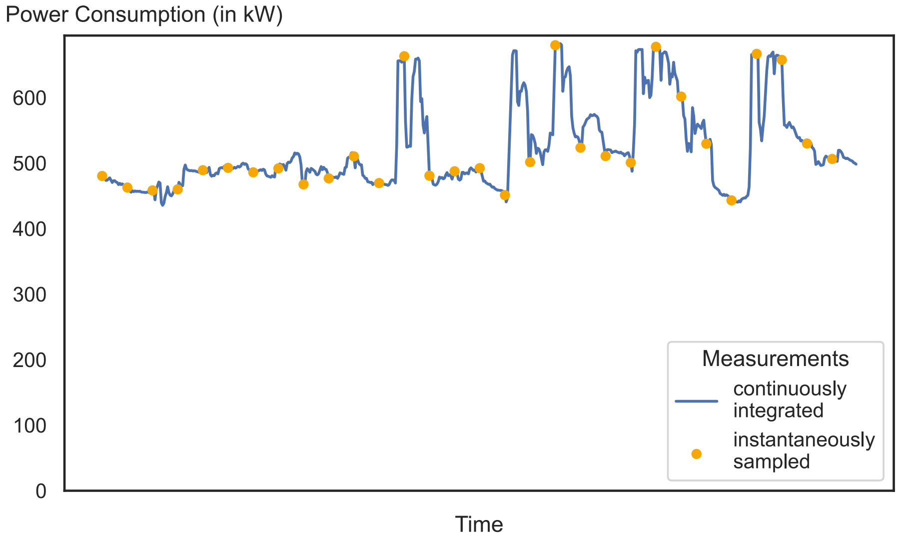

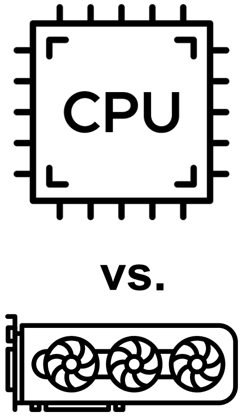

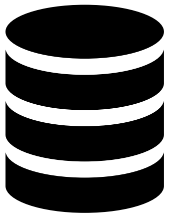

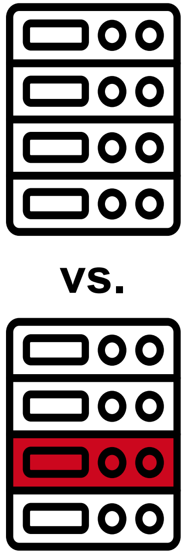

- The **elementary** level **has a low barrier to entry**
- The **intermediate** level **requires measurements** instead of estimates

Carbon Accounting Levels

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	yearly
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	monthly
advanced	continuous integrated energy	separate, all contributing	measured	usable	hourly

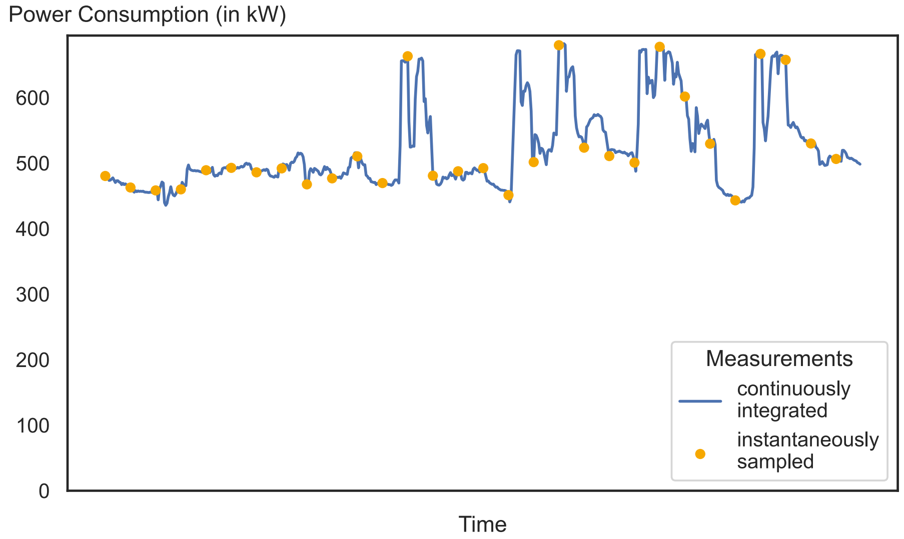

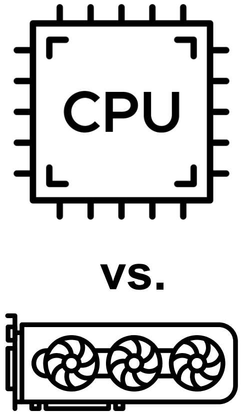

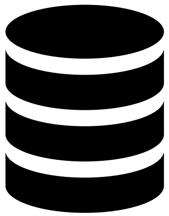

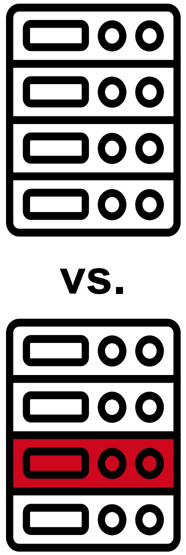

- The **elementary** level **has a low barrier to entry**
- The **intermediate** level **requires measurements** instead of estimates
- The **advanced** level **refines the granularity** of the calculated carbon footprints

Agenda

- Carbon Footprint Formula
- Carbon Accounting Levels
- **Case Studies at RWTH Aachen & TU Dresden**
- Summary & Outlook

Case Studies at RWTH Aachen & TU Dresden – Overview

- Cluster data analyzed:
  - **CLAIX-2018** at RWTH Aachen:
    - In operation since 2018-12 → **One year of data starting 2022-10-01**
    - 1243 CPU-only nodes with 2 x 24 core Intel Skylake
    - 54 GPU-enabled nodes with 2 x NVIDIA V100
  - **Barnard** at TU Dresden:
    - In operation since 2023-10 → **Eight month of data starting 2023-10-01**
    - 630 CPU-only nodes with 2 x 52 core Intel Sapphire Rapids
    - Free cooling
  - **Time resolution: 15 min** (after aggregation)

- Available **core hours derived from SLURM**'s `sacctmgr` node events

- **German energy mix** data of German federal network agency:
  - **Time resolution: 15 min**
  - Two energy mixes considered:
    - **Green**: Market-based method considering **renewable energy certificates (RECs)**
    - **Gray**: Location-based method considering **all energy sources**

[2]

Case Studies at RWTH Aachen & TU Dresden – Cluster Energy Consumption & iPUE

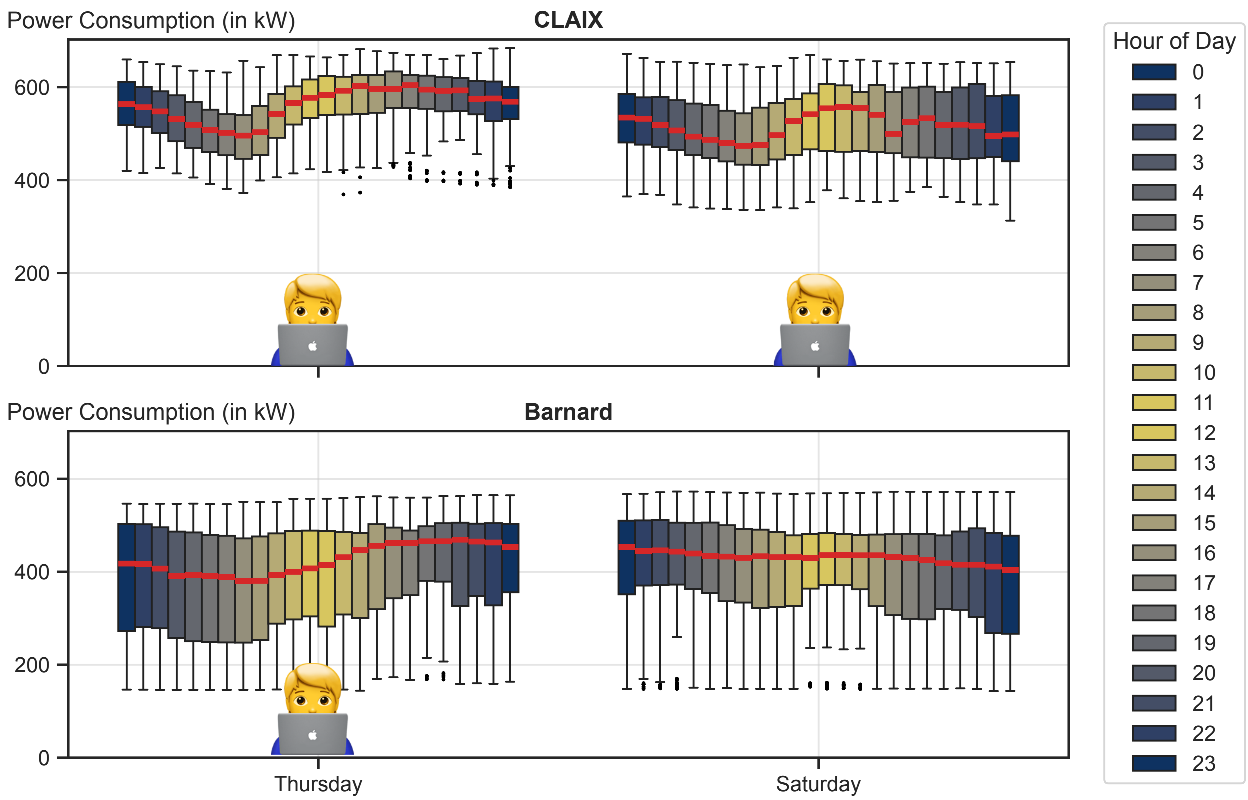

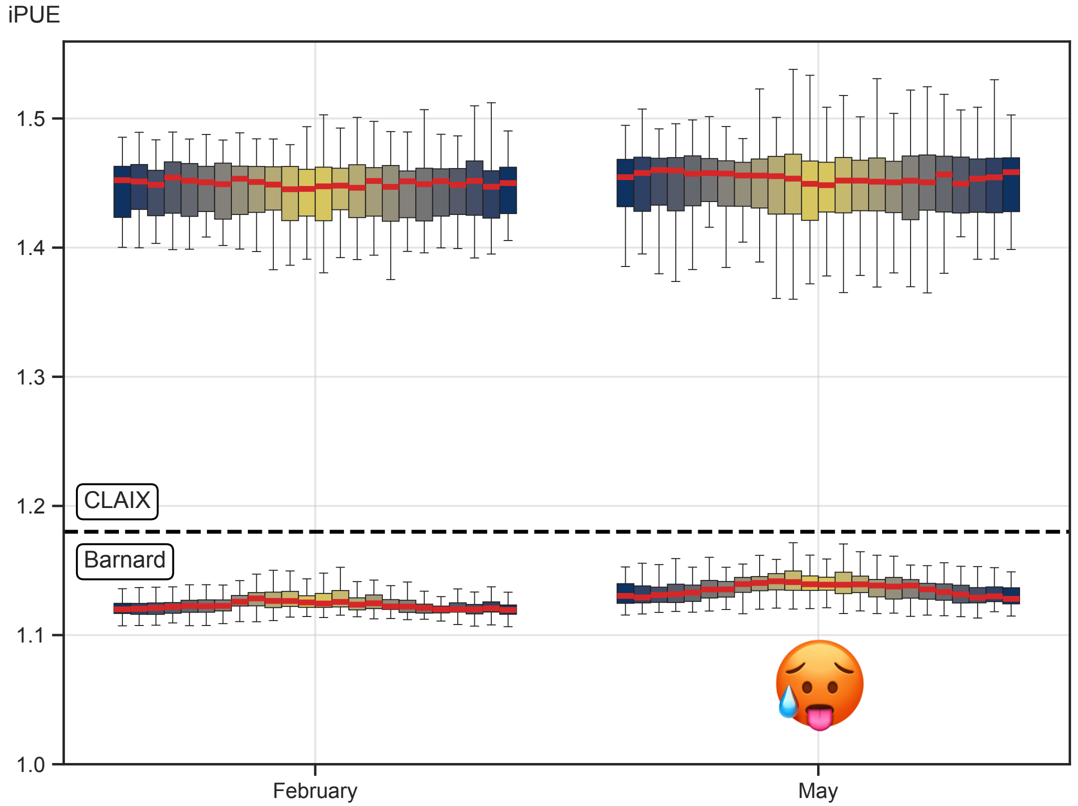

- Increasing power consumption during working hours
- Decreasing power consumption otherwise

- CLAIX without free cooling: no clear trends
- Barnard with free cooling:
  - Warmer months coincide with higher iPUEs
  - Daytime coincides with higher iPUEs

Case Studies at RWTH Aachen & TU Dresden – Electricity Generation

- Gray 6 times more emissions than Green

Energy Mix	Carbon Intensity
Gray	316.95 gCO_2e/kWh
Green	51.28 gCO_2e/kWh

- Qualitatively similar hourly trends:
  - In "sunny" months, low carbon around noon
  - Otherwise, fluctuations without clear trend
- But: Gray with proportionally larger changes

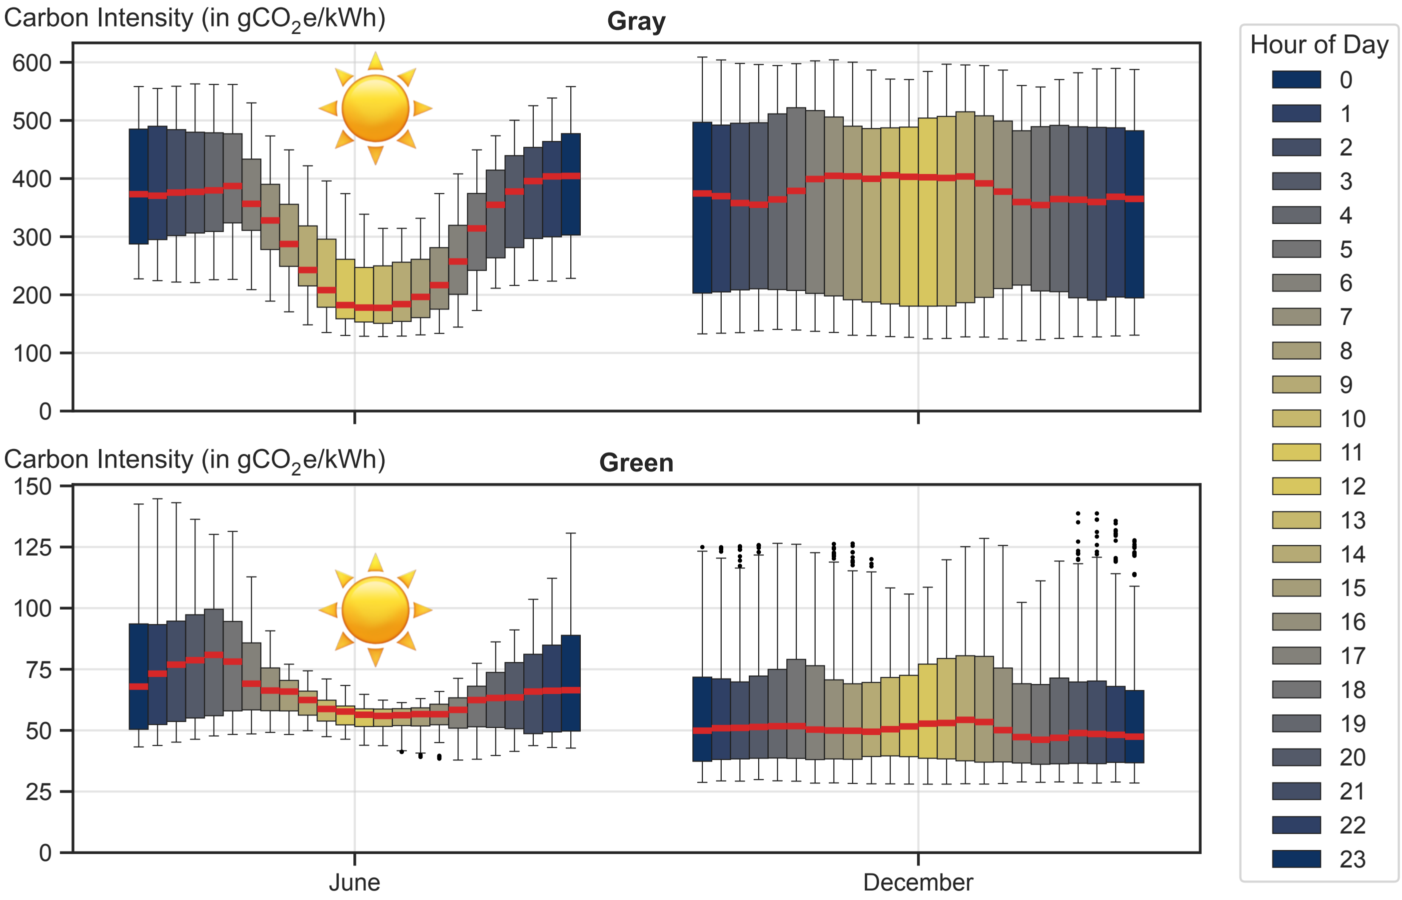

IT Center
RWTH AACHEN UNIVERSITY

Case Studies at RWTH Aachen & TU Dresden – Carbon Accounting Levels

- Energy mix: Gray >> Green
- Level: elementary < intermediate
- Benefits of Barnard show in
  - lower iPUE values, and
  - lower energy per core hour

  resulting in lower carbon footprints

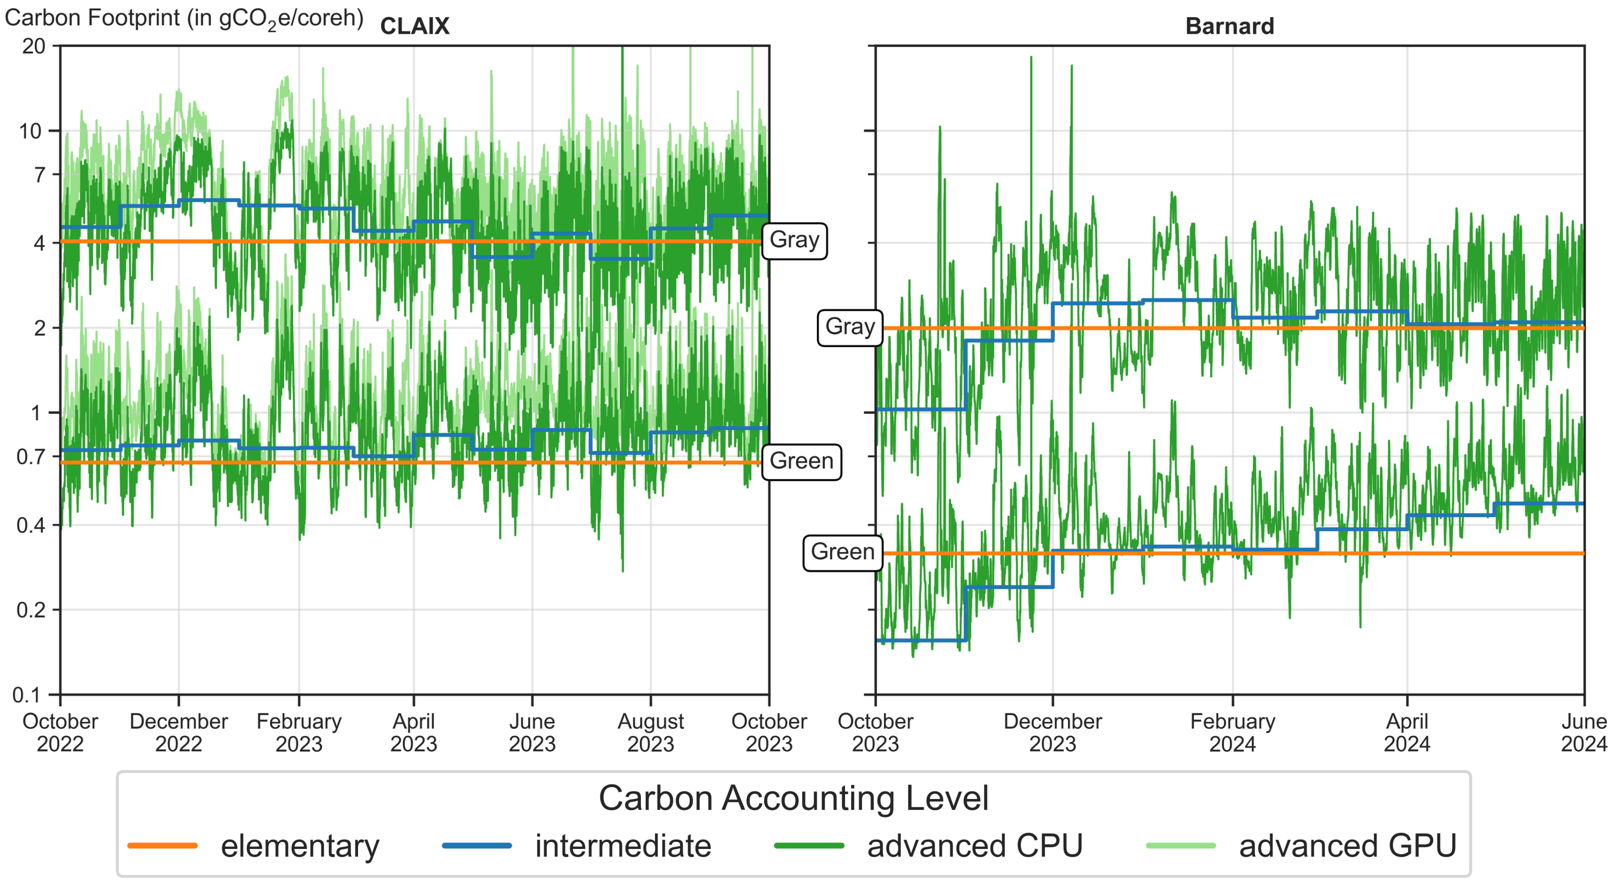

Agenda

- Carbon Footprint Formula
- Carbon Accounting Levels
- Case Studies at RWTH Aachen & TU Dresden
- **Summary & Outlook**

Summary & Outlook

- Carbon calculation formula:

$$Carbon\ Footprint_{Start,\ End} = \frac{IT\ Energy * PUE\ * \langle \overrightarrow{Energy\ Mix}, \overrightarrow{Emission\ Factors} \rangle}{Core\ Hours}$$

- Carbon accounting levels:

Level	Measurements	Balancing Boundary	PUE	Core Hours	Resolution
elementary	instantaneously sampled power	aggregate, compute-only	estimated	installed	yearly
intermediate	continuous integrated energy	aggregate, compute-only	measured	usable	monthly
advanced	continuous integrated energy	separate, all contributing	measured	usable	hourly

- Case studies show time-sensitivity of carbon footprints depending on the cluster load, iPUE, and energy mix

Outlook

Thank you for your attention!

- Incorporate information of the embodied carbon
- Involvement of HPC users: Reporting of their carbon footprints + Incentivising flexible job scheduling

References

- [1] L. Lannelongue, J. Grealey, and M. Inouye, ‘Green Algorithms: Quantifying the Carbon Footprint of Computation’, Advanced Science, vol. 8, no. 12, 2021, doi: 10.1002/advs.202100707.
- [2] https://smard.de
- [3] O. Edenhofer et al., ‘Climate Change 2014: Mitigation of Climate Change’, IPCC, Nov. 2018. [Online]. Available: https://www.ipcc.ch/site/assets/uploads/2018/02/ipcc_wg3_ar5_full.pdf
- [4] T. Gibon, Á. J. Hahn Menacho, and Mélanie Guiton, ‘Carbon Neutrality in the UNECE Region: Integrated Life-cycle Assessment of Electricity Sources’, UNECE, Apr. 2022. [Online]. Available: https://unece.org/sites/default/files/2022-04/LCA_3_FINAL%20March%202022.pdf
- [5] B. Knowles, ACM TechBrief: Computing and Climate Change. 2021.

- CC by popcornarts https://thenounproject.com/icon/storage-disc-2636532/
- CC by Misha Petrishchev https://thenounproject.com/icon/gpu-1132941/
- CC by Zoya Khan https://thenounproject.com/icon/data-center-6386087/
- CC by yogi rista https://thenounproject.com/icon/earth-6619582/
- CC by littlestar23 https://thenounproject.com/icon/measure-6414403/
- CC by krisna arga muria https://thenounproject.com/icon/carbon-footprint-7133151/
- CC by Gary Raftery https://www.bpb.de/175089
- Flaticon License by Good Ware https://www.flaticon.com/free-icon/cpu-tower_2432574
- https://commons.wikimedia.org//wikiFile:Flag_of_Europe.svg

Backup Slides

Units in Carbon Footprint Formula

$$Carbon\ Footprint_{Start,\ End} = \frac{IT\ Energy * PUE\ * \langle \overrightarrow{Energy\ Mix}, \overrightarrow{Emission\ Factors} \rangle}{Core\ Hours}$$

$$= \frac{[kWh] * [unitless] * \langle [unitless], [\frac{gCO_2e}{kWh}] \rangle}{[coreh]}$$

$$= \frac{[kWh] * [unitless] * [\frac{gCO_2e}{kWh}]}{[coreh]}$$

$$= [\frac{gCO_2e}{coreh}]$$

Carbon Intensity Studies

- Studies on life-cycle greenhouse gas **emissions of energy sources**
  - **Worldwide scope**: 2014 by IPCC (= Intergovernmental Panel on Climate Change)
  - **EU focus**: 2022 by UNECE (= United Nations Economic Commission for Europe)

IPCC

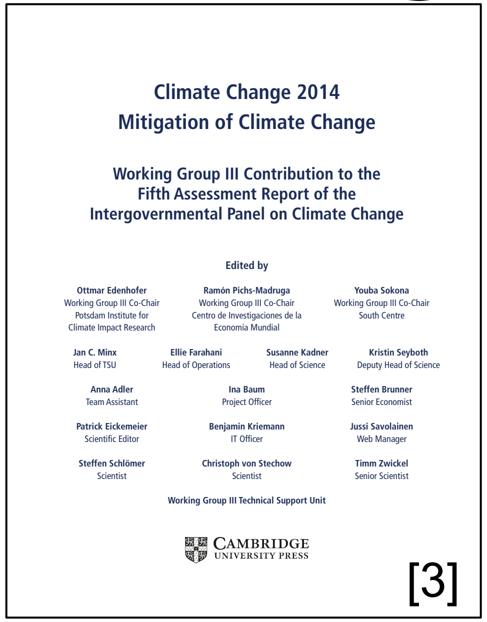
Climate Change 2014
Mitigation of Climate Change

Working Group III Contribution to the
Fifth Assessment Report of the
Intergovernmental Panel on Climate Change

Edited by

Ottmar Edenhofer
Working Group III Co-Chair
Potsdam Institute for
Climate Impact Research

Ramón Pichs-Madruga
Working Group III Co-Chair
Centro de Investigaciones de la
Economía Mundial

Youba Sokona
Working Group III Co-Chair
South Centre

Jan C. Minx
Head of TSU

Ellie Farahani
Head of Operations

Susanne Kadner
Head of Science

Kristin Seyboth
Deputy Head of Science

Anna Adler
Team Assistant

Ina Baum
Project Officer

Steffen Brunner
Senior Economist

Patrick Eickemeier
Scientific Editor

Benjamin Kriemann
IT Officer

Jussi Savolainen
Web Manager

Steffen Schlömer
Scientist

Christoph von Stechow
Scientist

Timm Zwickel
Senior Scientist

Working Group III Technical Support Unit

CAMBRIDGE
UNIVERSITY PRESS

[3]

Carbon Intensities (in gCO_2e/kWh)

IPCC	Energy Source	UNECE
11	Wind Onshore	12
12	Wind Offshore	13.5
24	Hydropower	11
44.5	Solar	21
230	Biomass	
134	Other Renewable	
12	Nuclear	5.1
490	Fossil Gas	430
820	Hard Coal	1000

UNECE

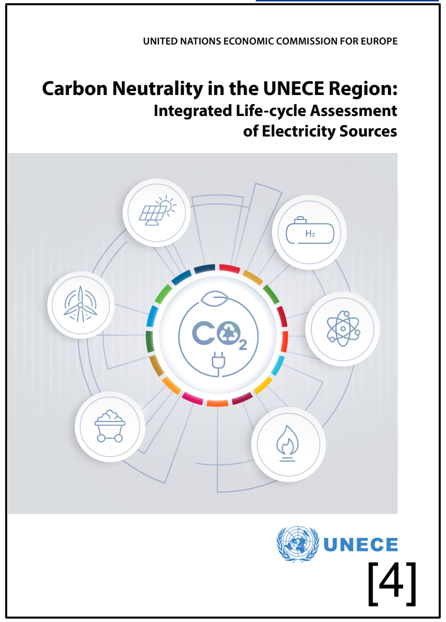

[4]

Choice of Advanced Time Resolution

- Note: ECDF plots
- With coarser time resolution:
  - Less low carbon time frames
  - Less high carbon time frames
- Visual assessment:
  - Indistinguishable: 15 min & 1 h
  - Close match: 6 h
- ⚠️ Some energy mixes are only published every 1 h

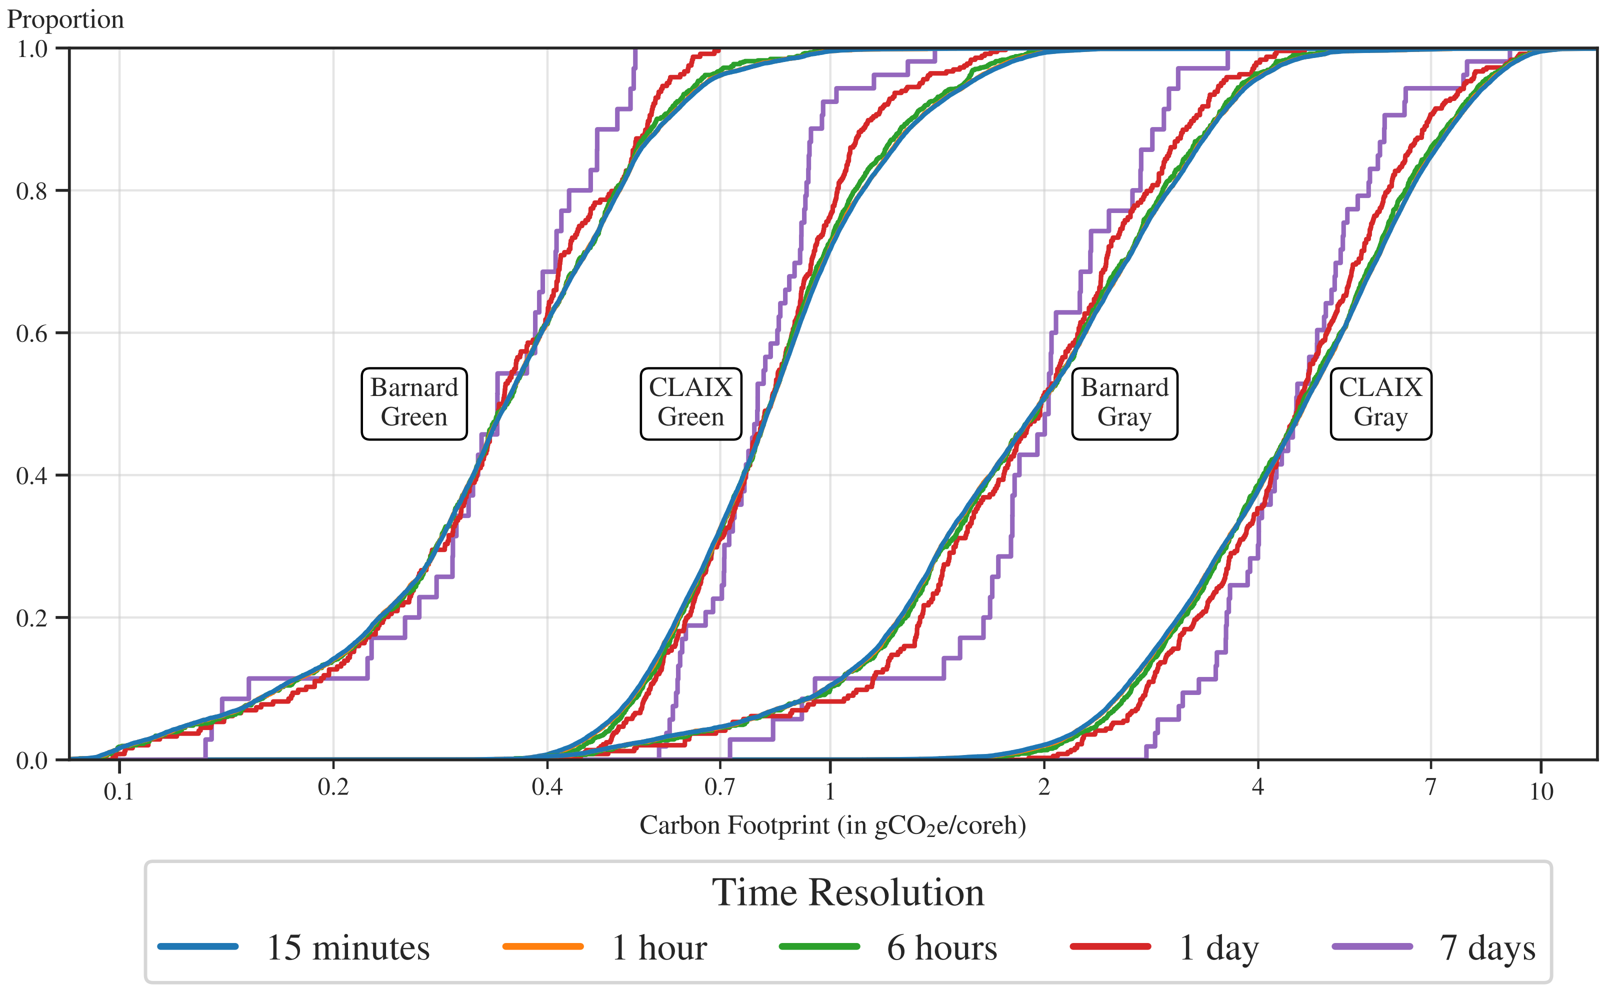

Discussion of Core Hours

- “Core Hour” corresponds to usage of one physical for one hour
  - Hyper-threads are typically not scheduled individually
  - GPUs are typically accounted via equivalent core hours
    - For example, 96 CPU cores and 4 GPUs → 24 core hours per GPU hour

- Choice between usable, used, and billed core hours
  1. “Usable” decouples the carbon footprint from idle resources → **incentivise shifting to lower load periods**
  2. “Usable” assigns the carbon footprint of idle resources to the HPC operator
  3. “Usable” can cause > 100% carbon emissions when multiplied by job billing due to overallocation of memory / GPUs
  4. ”Used” requires **time-consuming processing of privacy-related job data**
  5. “Used” does not attribute overallocation of memory / GPUs fairly → similar issue as 3.
  6. “Billed” fixes 5. but also entails 4.