



HAL
open science

BoaviztAPI: a bottom-up model to assess the environmental impacts of cloud services

Thibault Simon, David Ekchajzer, Adrien Berthelot, Eric Fourboul, Samuel Rince,
Romain Rouvoy

► **To cite this version:**

Thibault Simon, David Ekchajzer, Adrien Berthelot, Eric Fourboul, Samuel Rince, et al.. BoaviztAPI: a bottom-up model to assess the environmental impacts of cloud services. HotCarbon'24. Workshop on Sustainable Computer Systems, Jul 2024, Santa Cruz, United States. <hal-04621947v1>

HAL Id: hal-04621947

<https://hal.science/hal-04621947v1>

Submitted on 26 Jun 2024 (v1), last revised 6 May 2025 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

BoaviztAPI: A Bottom-Up Model to Assess the Environmental Impacts of Cloud Services

Thibault Simon
Orange Labs, Univ. Lille, Inria,
CNRS, UMR 9189 CRISTAL
Rennes, France
thibault.simon@inria.fr

David Ekchajzer
Hubblo, Université Paris-Saclay,
Univ. Evry, IMT-BS, LITEM,
Evry-Courcouronnes, France
david.ekchajzer@universite-paris-
saclay.fr

Adrien Berthelot
ENS de Lyon, OCTO Technology,
Inria
Lyon, France
adrien.berthelot@ens-lyon.fr

Eric Fourboul
Hubblo, Boavizta
Montpellier, France
eric@hubblo.org

Samuel Rince
Alygne, Boavizta
Paris, France
samuel@sento.re

Romain Rouvoy
Univ. Lille, Inria,
CNRS, UMR 9189 CRISTAL
Lille, France
romain.rouvoy@univ-lille.fr

ABSTRACT

Awareness surrounding the environmental impacts of the digital industry has led numerous professionals to incorporate these considerations into their work. However, the conceptualization of environmental impacts has often been narrowed to the scope of carbon footprint. This limitation can be attributed to various technical and data accessibility constraints, hindering a comprehensive evaluation, including a multi-criteria analysis over the entire life cycle of digital technologies.

In response to these limitations, we introduce a comprehensive bottom-up evaluation method suitable for servers and cloud instances, employing a life cycle thinking approach. We start by modeling the life cycle impacts of a server based on its hardware configuration. Then, we aggregate these with the impacts of its technical and physical environment to define the impacts of a cloud platform. We finally model the impacts of a cloud instance as a portion of the cloud platform.

The proposed approach has been implemented as an open-source toolkit and published as an API. This initiative aims to provide developers and researchers with a tool for conducting environmental evaluations of their infrastructure based on open data and open methodologies, enhancing their ability to explore the environmental materiality of ICT products, services, and infrastructures.

CCS CONCEPTS

• **General and reference** → *Evaluation*; • **Hardware** → **Impact on the environment**; • **Social and professional topics** → *Sustainability*; • **Information systems** → **Data centers**.

KEYWORDS

Cloud, Environmental Impact, Life Cycle Assessment, Bottom-up, Multicriteria, Open-data

1 INTRODUCTION

Six out of the nine planetary boundaries introduced by Rockstrom have been exceeded [1]. Regarding the climate change boundary, Freitag *et al.* [2] estimate that in 2020 between 2.1% and 3.9% of worldwide greenhouse gas emissions could be attributed to the

digital sector. Furthermore, the sector also contributes to the overshooting of other planetary boundaries, as shown for Europe in [3]. The hosting infrastructure is estimated to contribute between 18 and 45% to the sector’s carbon footprint [2, 4–6]. As data centers experience substantial growth [7], their associated environmental impacts are likely to increase, further hindering the path toward the sector’s sustainability. This trend is especially pronounced for cloud infrastructures: according to CISCO, the global cloud data centers traffic has been growing at a 25% *compound annual growth rate* (CAGR) between 2016 and 2021 [8].

In this context, major cloud providers, such as AWS [9], Azure [10], and GCP [11], have introduced tools enabling customers to assess the carbon footprint associated with their cloud usage. However, the underlying methodologies adopted by these actors have not been subjected to rigorous transparency procedures, which presents a significant threat to ensuring the consistency of results between tools. This lack of transparency results in inconsistencies that hinder the comparability of results, stemming from variations in the scope considered and allocation methods employed.

Table 1 summarizes the key differences between the three major cloud providers carbon assessment methodologies. One can observe that some parameters—denoted as (✓)—are only partially accounted for, or using non-standardized approaches. Others, such as third-party services, can be accounted for in some tools but not in others.

To address these limitations, the open source project *Cloud Carbon Footprint* (CCF) [12] introduced a provider-agnostic approach. However, its calculation remains constrained by several factors. For the *resource extraction & manufacturing* phase, impacts are systematically estimated based on the Dell R740 Life Cycle Assessment [13], regardless of the hardware components used. Power is modeled based on benchmarks that measure power at the machine level [14], which CCF associates with *cloud instances* based on CPU architecture. Not only does the CPU architecture not seem to be a good proxy for estimating the power consumption of a CPU, but even less so for the whole machine. For example, a storage server’s primary source of power consumption is typically the hard drives. A more specific model for each component seems to be required.

All of these methodologies to assess the environmental impact of cloud infrastructures only address the carbon footprint dimension.

	Azure	GCP	AWS	CCF	This paper
Lifecycle phase					
Resource extraction & manufacturing	✓	✓	✗	✓	✓
Distribution	?	?	?	✓	(✓)
Usage	✓	✓	✓	✓	✓
End of life	✓	✗	✗	(✓)	✗
Included resources					
Buildings	✗	✓	✗	✗	✓
IT equipments	✓	✓	✓	(✓)	(✓)
Technical & building environment	✓	✓	?	✗	✓
Employee commuting	✗	✓	✗	✗	✗
Coolant leaks	?	✗	?	✗	✗
Impact of IP traffic	✓	✗	?	✗	✗
Fossil fuel energy source on site	?	✓	✓	✗	✗
Unallocated resources (IDLE)	✓	✓	?	✗	✗
Third-party services (overhead)	✓	✓	?	✗	✗
Electricity method					
T&D lost	?	✗	?	(✓)	(✓)
Manufacture of energy infrastructures	?	✓	?	?	✓
Location-based	(✓)	✓	✗	(✓)	✓
Market-based	✓	✓	✓	(✓)	✗

Table 1: GHG emissions accounting perimeter of cloud services

However, these infrastructures also have important environmental impacts on other dimensions, such as metal extraction [3, 15, 16]. This *carbon tunnel vision* [17] obscures environmental issues that could become critical and potentially leaves the door open to deleterious transfer [18, 19] from one environmental impact to another.

The ICT sector must contribute to the path toward sustainable development actively. This entails not only monitoring and mitigating environmental footprints across various impact categories, but also considering different life cycle phases rather than focusing solely on energy consumption. Developing a comprehensive understanding of the environmental implications of cloud computing equipment and services, including servers and instances, is essential for informed decision-making. However, there is still a lack of open standards to facilitate comparable and verifiable results that are provider-agnostic. In light of such limitations, this paper proposes a bottom-up approach for the modeling of the environmental impacts of cloud infrastructures, implemented as an open-source API.¹

2 METHODOLOGICAL BACKGROUND

Based on the following principles, this paper presents an approach that could be used to meet the above-mentioned requirements.

¹<https://github.com/Boavizta/boaviztapi>




Figure 1: Modeling approach used to estimate the environmental impacts of a cloud platform and instances

Attributional Life Cycle Assessment (A-LCA). Defined by the ISO 14044 [20], LCA is a method both employed in the industrial and academic context [21, 22]. It is used to estimate the potential environmental impacts of a product, a process, a project, or, in our context, a service. The methodology presented in this paper follows its main principles:

- The most important phases of the life cycle of servers and cloud services are considered: raw material acquisition, manufacturing (both are included in what is called *embodied impacts* in the rest of the paper) and usage. It should be noted that distribution to end-users and end-of-life impacts are excluded, as data on e-waste collection and recycling rates are scarce [3, 23].
- A multi-criteria approach is followed, providing impact factors for three environmental impact categories, as defined by [24]:
 - *Abiotic Resource Depletion of minerals and metals (ADP)*, which assesses the use of minerals and fossil raw materials. This indicator is a model for assessing the contribution of mineral and metal extraction to their progressive scarcity, using antimony grams equivalent as a metric [25].
 - *Primary Energy (PE)*, which includes all energy, direct and indirect, used in any phase of the life cycle. This represents the total cumulative energy demand of the assessed system [26].
 - *Global Warming Potential (GWP)*, which evaluates the effects on global warming. This well-known indicator is also a model linking greenhouse gas emissions to global warming. It is expressed in grams of CO_2 equivalent [27].

Modeling approaches. As illustrated in Figure 1, the environmental impacts of a *cloud platform* are mostly modeled using a *bottom-up* approach, whereby the impacts of each resource required to fulfill the service it provides are aggregated. Unlike a top-down approach, this allows the proportion of the contribution of each resource to the global impacts to be identified. The impacts of the technical and building environment are allocated on the *cloud platform* using a top-down approach.

Open data. This methodology leverages openly accessible data to ensure that the results are transparent, reproducible, and verifiable. Two types of open data are used:

- **Market & technical data:** characteristics of components, devices, and cloud instances available on the market.
- **Impact factors:** which convert physical quantities into environmental impacts. Those factors are extracted from publicly available life cycle assessments.

3 SERVER MODEL

Following a bottom-up approach (cf. Figure 1), the assessment of the environmental impacts of *cloud instances* starts with a component-level assessment of servers. Their impacts are then aggregated and integrated with the technical and building environment to constitute a *cloud platform*, which is subsequently allocated into individual *cloud instances*. In this section, we define the foundational layer of the model—the estimation of a server’s environmental impact—and further validate it against the LCA of a Dell R740 [13]. While similar bottom-up methods already exist [28], our approach aims to deliver: a more detailed and comprehensive calculation of the embodied footprint, multi-criteria impacts, and a calculation of the usage impact specific to our use case.

Variables denoted as \mathcal{F} represent *impact factors*, encapsulating environmental impacts associated with a specific quantity for any given *impact criteria* defined in Section 2—i.e., ADP, PE, and GWP. These impact factors are, thus, expressed in units of impact, such as kgCO_2e , per quantity, such as Wh of electricity. Variables designated as \mathcal{I} are environmental impacts, expressed solely in the impact unit.

3.1 Embodied impact of a physical server

The impact assessment relies on the life cycle modeling of hardware components, encompassing their most important phases: raw material acquisition, and manufacturing, collectively termed as *embodied impacts* (cf. Section 2). In all subsequent equations, embodied impact factors \mathcal{F}^e are amortized over the life cycle expectancy \mathcal{D} (in hours) to obtain the embodied impact for an hour of usage.

3.1.1 CPU. For most electronic components, the primary source of impacts lies in the process of engraving semi-conductors [29]. Consequently, their impacts directly depend on both their die size and the engraving technology employed.

For CPUs, the die size in mm^2 is multiplied by the corresponding impact factor \mathcal{F}_{cpu}^{die} with a base impact \mathcal{I}_{cpu}^{base} (containing packaging, heatsink socket, and transportation) to obtain an environmental impact factor per hour of usage:

$$\mathcal{F}_{cpu}^e = \frac{die_{cpu} \times \mathcal{F}_{cpu}^{die} + \mathcal{I}_{cpu}^{base}}{\mathcal{D}} \quad (1)$$

For both of these impact constants, we propose default values developed in Table 2 over the three considered impact categories in Table 2, based on a 14 nm engraving process.

We retrieve CPU die sizes from the TechPowerUp CPU specs database [30]. Crowd-sourced characteristics for more than 1750 CPU models are available within the API.²

Using Equation 1 for two CPU units with a die size of 6.94 cm^2 , the non-amortized embodied carbon footprint is estimated at 45.62




Figure 2: Bottom-up modeling of the manufacturing environmental impacts of a Dell R740 defined in [13] using our approach

$\text{kgCO}_2\text{e/h}$, while it accounts for approximately $46.76 \text{ kgCO}_2\text{e}$ for 1 hour in Dell’s LCA.

3.1.2 NAND memory. The impact factors per hour of both SSD and RAM sticks are obtained with the Equation 2, albeit utilizing distinct impact factors for their respective density and capacity.

$$\forall flash \in \{ssd, ram\} : \mathcal{F}_{flash}^e = \frac{\frac{capacity}{density} \times \mathcal{F}_{flash}^{die} + \mathcal{I}_{flash}^{base}}{\mathcal{D}} \quad (2)$$

For the twelve 32GB sticks using a density of 1.79 GB/mm^2 from [31], we obtain $534.60 \text{ kgCO}_2\text{e/h}$ using Equation 2, while the Dell LCA reports $553.33 \text{ kgCO}_2\text{e}$ for 1 hour.

The server combines multiple SSD disks. Considering a density of 19 GB/cm^2 extracted from [32], we obtain $52.65 \text{ kgCO}_2\text{e/h}$ for the 400 GB disk with Equation 2, and $3,607.77 \text{ kgCO}_2\text{e/h}$ for the 8 3.84 TB ones, versus $64.1 \text{ kgCO}_2\text{e}$ and $3,373.5 \text{ kgCO}_2\text{e}$ for the R740, respectively. More up-to-date crawled SSD densities are available within our model.³

3.1.3 Others. Apart from CPU, memory, and storage, other components are required for a server which we categorized as *others*.

The *Power Supply Unit* (PSU)’s embodied impact is estimated at $72.71 \text{ kgCO}_2\text{e}$ using the impact factor \mathcal{F}_{psu}^e and the R740 PSU weight of $2,992 \text{ kg}$ [13]. The rest of the component’s impact factor is static in our approach and computed as follows:

$$\mathcal{F}_{others}^e = \frac{\mathcal{I}_{motherboard}^e + \mathcal{I}_{psu}^e + \mathcal{I}_{assembly}^e + \mathcal{I}_{casing}^e}{\mathcal{D}} \quad (3)$$

By using the open source factors provided in Table 2, the remaining server’s components account for $295.49 \text{ kgCO}_2\text{e/h}$, compared to $207.07 \text{ kgCO}_2\text{e}$ reported in the R740 LCA.

3.1.4 Total. The total embodied impact factor of a server per hour of usage is finally calculated as the sum of its components’ impacts as follows:

$$\mathcal{F}_{server}^e = \mathcal{F}_{cpu}^e + \mathcal{F}_{ram}^e + \mathcal{F}_{storage}^e + \mathcal{F}_{others}^e \quad (4)$$

²https://github.com/Boavizta/boaviztapi/blob/main/boaviztapi/data/crowdsourcing/cpu_specs.csv

³https://github.com/Boavizta/boaviztapi/blob/main/boaviztapi/data/crowdsourcing/ssd_manufacture.csv

Constant	Impact	Value	Unit
\mathcal{F}_{cpu}^{die}	GWP	1.97	kgCO ₂ e/mm ²
	ADP	5.87e−07	kgSbeq/mm ²
	PE	2.65e+01	MJ/mm ²
\mathcal{I}_{cpu}^{base}	GWP	9.14	kgCO ₂ e
	ADP	2.04e−02	kgSbeq
	PE	156.43	MJ
$\mathcal{F}_{flash}^{die}$	GWP	2.20	kgCO ₂ e/cm ²
	ADP	6.30e−05	kgSbeq/cm ²
	PE	2.73e+01	MJ/cm ²
\mathcal{I}_{ram}^{base}	GWP	5.22	kgCO ₂ e
	ADP	1.69e−03	kgSbeq
	PE	74.00	MJ
\mathcal{I}_{ssd}^{base}	GWP	6.34	kgCO ₂ e
	ADP	5.63e−04	kgSbeq
	PE	73.98	MJ
\mathcal{I}_{hdd}^e	GWP	3.11e+01	kgCO ₂ e
	ADP	2.50e−04	kgSbeq
	PE	2.76e+02	MJ
\mathcal{F}_{psu}^e	GWP	2.43e+01	kgCO ₂ e/kg
	ADP	8.30e−03	kgSbeq/kg
	PE	3.52e+02	MJ/kg
$\mathcal{I}_{motherboard}^e$	GWP	6.61e+01	kgCO ₂ e
	ADP	3.69e−03	kgSbeq
	PE	8.36e+02	MJ
$\mathcal{I}_{assembly}^e$	GWP	6.68	kgCO ₂ e
	ADP	1.41e−06	kgSbeq
	PE	6.86e+01	MJ
\mathcal{I}_{rack}^e	GWP	1.50e+02	kgCO ₂ e
	ADP	2.02e−02	kgSbeq
	PE	2.20e+03	MJ

Table 2: Impact constants, extracted from [31]

This results in a total life cycle value—i.e. non-amortized—of 4, 536.13 kgCO₂e/h against a 4, 244.76 kgCO₂e baseline reported in the Dell R740 LCA [13]. However, without transparent factors, modeling methodology, and allocation choices provided, this 6.42% variation cannot be further detailed, emphasizing the need for open methodologies.

The comprehensive evaluation of environmental impacts across various categories is depicted in Figure 2, emphasizing the necessity of adopting holistic methodologies in impact assessments to achieve thorough evaluations of ICT impacts. These impacts vary across different impact categories, highlighting the need to avoid carbon-centric perspectives to reveal potential shifts in impact distribution among various categories.

3.2 Usage impact of a physical server

The total energy consumption of a server \mathcal{E}_{server} , in Wh, is calculated as the sum of its components (C) power consumption \mathcal{P} , in W, over a given duration \mathcal{T} , in hours:

$$\mathcal{E}_{server} = \sum_{c \in C} (\mathcal{P}_c \times \mathcal{T}) \quad (5)$$

To obtain the associated environmental impact factor for one hour of usage \mathcal{F}_{server}^u , this consumption is multiplied by the impact factor \mathcal{F}_{em} representing the electricity mix—i.e., the environmental impacts associated to the production and transport of energy.

However, contrarily to embodied impacts, the power drawn by components is not static: it depends on the workload executed. While it can be physically measured, such measurements are often not available, especially in cloud instances. As such, we propose an alternative approach to estimate the component’s power draw for a given workload.

3.2.1 CPU power. For a given CPU workload $w \in [0\%, 100\%]$ which represents the ratio of time the CPU is actively processing tasks to the total time observed—expressed as a percentage—the power consumption \mathcal{P}_{cpu} is estimated using:

$$\mathcal{P}_{cpu}(w) = a \times \log(b \times (w + c)) + d \quad (6)$$

The parameters a , b , c , and d are determined for a logarithmic regression model fitted on sample data (w, \mathcal{P}_{cpu}) collected from empirical measurements. We selected a logarithmic regression model based on the power consumption curves of CPUs for AWS cloud instances as described in the work of [33]. Here, we propose another way to determine them inspired from [34], using the *Thermal Design Power* (TDP) value of a given CPU and a generic power consumption profile derived from [33] that expresses the power draw as a ratio of the TDP. By fitting a power consumption profile \mathcal{P}_{cpu} using power measurements sampled at 0%, 10%, 50% and 100% workloads, which correspond to TDP ratios of 0.12, 0.32, 0.75, and 1.02, respectively, we can describe the workload as a continuous variable, enhancing the estimation of CPU power draw. Further details on the modeling of components’ power consumption can be found in the project’s documentation.⁴

3.2.2 NAND memory power. While the power consumption of SSD and RAM also depends on their workloads, defining such workload can be tedious, particularly for SSDs. Furthermore, we mostly observed consistent power consumption profiles and thus used constant power consumption values for both idle and active states.

For a given model, the TechPowerUp SSD specs database [35] can be used to estimate the idle and active (averaged) power consumption values. For RAM banks, we use the averaged values from [33]: 0.19 W/GB and 0.54 W/GB in idle and active states, respectively. By default, RAM banks are constantly considered in an active state.

3.2.3 Power of other components. The power consumption of the rest of the server is considered as overhead and thus set as a factor f of the sum of CPU, RAM, and SSD energy consumption. This ratio is an average that counts in the power consumption of the motherboard, fans, and PSU(s). Its power draw is estimated as follows, with a factor value by default at $f = 0.2$:

$$\mathcal{P}_{other}(w) = f \times (\mathcal{P}_{cpu} + \mathcal{P}_{ram} + \mathcal{P}_{ssd}) \quad (7)$$

Note that we do not include graphic or AI accelerators (GPUs, TPUs, etc.) in this study, their power consumption must be modeled separately and will be addressed in future work.

⁴<https://doc.api.boavizta.org/Explanations/components/cpu/>

4 CLOUD MODEL

According to Figure 1, the subsequent stage of the modeling process is to analyze the impacts of *cloud platforms*, which will be used in the assessment of *cloud instances*.

4.1 Embodied & usage impacts of a cloud platform

We define a *cloud platform* as the aggregation of a cluster of servers and their technical and building environment required to provide cloud services. A *cloud platform* offers a pool of resources assigned to cloud services: vCPU, vRAM, storage, and shared resources—power supply, motherboard, casing and technical environment.

To account for the embodied impacts of the technical and building environment, we refer to the study published by ARCEP [36], which reports on embodied impact factors per m^2 of a server room. The technical environment includes the building, generators, chillers, inverters, and batteries, as well as a wide range of equipment such as electrical and network cables, lighting, fuel oil storage tanks, etc. Knowing the electricity consumption per m^2 and the *Power Usage Effectiveness* (PUE) used in [36], we infer the embodied impacts of the technical environment \mathcal{F}_{DC}^e as $2.61e-02 \text{ kgCO}_2e/\text{kWh}$, $6.61e-01 \text{ MJ}/\text{kWh}$, and $1.17e-06 \text{ kgSbeq}/\text{kWh}$ for respectively GWP, PE and ADP criteria. As such, the embodied impact of the technical and building environment can be estimated for a given energy consumption \mathcal{E} as follows:

$$\mathcal{I}_{env}^e(\mathcal{E}) = \mathcal{E} \times \mathcal{F}_{DC}^e \quad (8)$$

In order to account for the usage impacts of the technical and building environment, the PUE of the cloud infrastructure is applied to the usage impacts of the server. The PUE is defined as the ratio of electricity consumed by the facility to the electricity consumed by the IT equipment [37]. The usage impacts are therefore defined as follows for a given energy consumption \mathcal{E} , using the electricity mix \mathcal{F}_{em} :

$$\mathcal{I}_{env}^u(\mathcal{E}) = \mathcal{E} \times (PUE - 1) \times \mathcal{F}_{em} \quad (9)$$

4.2 Embodied & usage impacts of a Cloud instance

A *cloud instance* is modeled as a part of a *cloud platform*. Its impacts encompass both a share of the technical and building environment impacts, \mathcal{I}_{env} , and a share of each of the servers' component impact factor \mathcal{F}_r . This share is computed for each component using the quantity assigned to the instance Q_r^u over the total available resources on the platform Q_r^u . Such quantity Q materializes as vCPUs for CPUs and GBs for RAM and disks. For the *others* resources, which are not explicitly assigned to a *cloud instance*—i.e. motherboard, PSU, assembly, and casing (cf. Equation 3)—we have chosen an allocation based on vCPU, although this may depend on the type of *cloud instance*.

$$\mathcal{F}_{inst} = \mathcal{I}_{env}(\mathcal{E}_{inst}) + \sum_{r \in \{cpu, ram, ssd, others\}} \frac{Q_r^u(instance) \times \mathcal{F}_r}{Q_r^u(platform)} \quad (10)$$

Equation 10 allows to obtain the instance embodied and usage impacts for an hour, each share is multiplied by the *cloud platform* total impacts for the resource (\mathcal{F}_r). This assumes that the server




Figure 3: GWP impact of a cloud instance hosted in France with 4 vCPUs, 8 GB of vRAM and 80 GB of SSD storage used for a year, calculated using Equation 10

is used continuously over its lifetime. To estimate the share of technical and building environment associated with the instance, Equation 8 is used with the instance energy consumption \mathcal{E}_{inst} . For one-hour usage, this energy consumption is estimated using Equation 5, using the ratio $\frac{Q_r^u(instance)}{Q_r^u(platform)}$ to obtain for each component the share to allocate to the instance.

To illustrate the aforementioned approach we consider a *cloud platform* comprising a single server with technical characteristics and environmental impacts defined in Section 3, with a lifespan of 5 years. We consider that the infrastructure is hosted in France where \mathcal{F}_{em} is equal to $0.098 \text{ kgCO}_2e/\text{kWh}$ [38]. The PUE is set arbitrarily to 1.5. CPU power is estimated at 104.75 W using Equation 6 and the open factors within the APL.⁵ Each SSD disk consumes 5.7 W .

The impact of a year of usage for a *cloud instance* with 4 vCPUs, 8 GB of RAM, and 80 GB of SSD storage are estimated using Equation 10 in Figure 3. One can observe that contrary to the server in Figure 2, the impact of SSDs is smaller due to a small share of disks reserved. The technical and building environment substantially increase the total impacts, both embodied and usage.

4.3 Limits & future work

As exposed in Fig. 1, our approach does not cover the entire perimeter of a cloud infrastructure. This includes third-party services hosted in the *cloud platform* that serve multiple customers. These can be technical services, such as the control plane, or customer services, such as billing. In addition, the usage impacts of non-server IT equipment, such as network equipment, are not taken into account. This also minimizes the impacts of the technical and building environment, as we do not consider their consumption when applying the PUE. Finally, we do not consider servers in idle states. This assumes that all servers are occupied, which is not always the case [39]. They are part of the cloud infrastructure and allow for rapid and important scaling. These exclusions mean that certain parts of the infrastructure remain unallocated to customers.

In addition, our allocation strategy poses some issues. It assumes that no resource is overcommitted (when a resource is used by two or more cloud instances at the same time). If a resource is overcommitted [40], all the instances sharing a physical resource would double account its impacts. These examples underscore the

⁵<https://doc.api.boavizta.org/Explanations/usage/power/>

significant margins of error associated with ICT environmental assessments, which remain inadequately quantified and documented, even though they have a significant impact on the final assessment results. In future work, we aim to enhance the tracing of uncertainties throughout the modeling process, employing approaches such as those suggested by [41].

To enhance the modeling of energy consumption, our future work includes Energizta, a collaborative project aimed at collecting and reporting open data on server energy usage.⁶ Certain services, such as Artificial Intelligence, Machine Learning, cryptomining, and *High Performance Computing* (HPC), heavily rely on GPUs, which will be incorporated in future work.⁷

Finally, the choice of a vCPU-based allocation for the *others* component and the technical and building environment remains arbitrary. This choice will be challenged in future work to better represent the link between these resources and the scaling of *cloud instance*.

5 CONCLUSION

This paper introduces our bottom-up methodology for assessing the environmental impacts of servers and *cloud instance* solutions based on crowd-sourced data. As an open-source project, we emphasize the transparency and reliability of our bottom-up methodology. While a top-down approach would allow for a broader scope, the bottom-up approach allows us to pinpoint which part of the *cloud instance* is responsible for the most impacts, empowering stakeholders to identify actionable reduction levers. It has been notably utilized by companies, such as Orange [42],⁸ Sweep⁹, and Sami,¹⁰ as well as for conducting GHG assessments in France. The research community has also used it as a basis to estimate and reduce various ICT aspects' environmental footprint, such as Kubernetes scheduling, AI, infrastructure management... [43–49]. Furthermore, the methodology can be enhanced in specific areas, benefiting from the expertise of each contributor in each segment of the digital value chain. We believe that this collaborative approach to evaluation will allow us to gradually and rigorously expand the scope of our evaluation methodology while maintaining confidence in the tools.

6 ACKNOWLEDGEMENTS

We extend our gratitude to Pierre Rust for his insightful comments and assistance in the development of this methodology. Additionally, we wish to express our sincere appreciation to the open-source contributors not previously mentioned in this paper, namely: Benoit Petit, Olivier de Meringo, Bruno Thomas, Jacob Valdemar, Benjamin Le Bigot, Benjamin Davy, Christophe Saugé, Simon Shillaker, Valentin Chaud, Luke Kuzmish, Olivier Delcroix, and Seb Solere. We welcome contributions and encourage readers to contact the authors with any challenges they encounter.

This work is partly supported by the French government through the *Agence Nationale de la Recherche* (ANR) under the France 2030 program, including funding from the CARECLOUD (ANR-23-PECL-0003) and DISTILLER (ANR-21-CE25-0022) grants.

⁶<https://boavizta.github.io/Energizta/>

⁷<https://github.com/Boavizta/boaviztapi/issues/65>

⁸<https://orange.fr>

⁹<https://sweep.net>

¹⁰<https://sami.eco>

REFERENCES

- [1] K. Richardson, W. Steffen, W. Lucht, J. Bendtsen, S. E. Cornell, J. F. Donges, M. Drüke, I. Fetzer, G. Bala, W. von Bloh, G. Feulner, S. Fiedler, D. Gerten, T. Gleeson, M. Hofmann, W. Huiskamp, M. Kumm, C. Mohan, D. Nogués-Bravo, S. Petri, M. Porkka, S. Rahmstorf, S. Schaphoff, K. Thonicke, A. Tobian, V. Virkki, L. Wang-Erlandsson, L. Weber, J. Rockström, Earth beyond six of nine planetary boundaries, *Science Advances* 9 (37) (2023) eadh2458. doi:10.1126/sciadv.adh2458.
- [2] C. Freitag, M. Berners-Lee, K. Widdicks, B. Knowles, A. Friday, The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations, *Patterns* 2 (9) (2021) 100340. doi:10.1016/j.patter.2021.100340.
- [3] Bordage, F., de Montenay, L., Benqassem, S., Delmas-Orgelet, J., Domon, F., Prunel, D., Vateau, C. et Lees Perasso, E., Digital technologies in Europe: an environmental life cycle approach, [Online; accessed 18. Mar. 2024]. URL <https://www.greenit.fr/wp-content/uploads/2021/12/EU-Study-LCA-7-DEC-EN.pdf>
- [4] L. Belkhir, A. Elmelig, Assessing ict global emissions footprint: Trends to 2040 & recommendations, *Journal of Cleaner Production* 177 (2018) 448–463. doi:10.1016/j.jclepro.2017.12.239.
- [5] D. Lundén, J. Malmödin, P. Bergmark, N. Lövehagen, Electricity consumption and operational carbon emissions of european telecom network operators, *Sustainability* 14 (5) (2022).
- [6] A. Pasek, H. Vaughan, N. Starosielski, The world wide web of carbon: Toward a relational footprinting of information and communications technology's climate impacts, *Big Data & Society* 10 (1) (2023) 20539517231158994. doi:10.1177/20539517231158994.
- [7] Electricity 2024 – analysis (2024). URL <https://www.iea.org/reports/electricity-2024>
- [8] C. V. Networking, Cisco global cloud index: Forecast and methodology, 2016–2021, White paper. Cisco Public, San Jose 1 (2016). URL https://virtualization.network/Resources/Whitepapers/0b75cf2e-0c53-4891-918e-b542a5d364c5_white-paper-c11-738085.pdf
- [9] Understanding your carbon emission estimations - AWS Billing, [Online; accessed 18. Mar. 2024] (2024). URL <https://docs.aws.amazon.com/awscaccountbilling/latest/aboutv2/ccft-estimation.html>
- [10] A new approach for Scope 3 emissions transparency, [Online; accessed 18. Mar. 2024] (2021). URL https://download.microsoft.com/download/7/2/8/72830831-5d64-4f5c-9f51-e6e38ab1dd55/Microsoft_Scope_3_Emissions.pdf
- [11] Carbon Footprint reporting methodology, [Online; accessed 18. Mar. 2024] (2024). URL <https://cloud.google.com/carbon-footprint/docs/methodology>
- [12] Overview | Cloud Carbon Footprint, [Online; accessed 18. Mar. 2024] (Feb. 2024). URL <https://www.cloudcarbonfootprint.org/docs>
- [13] Busa, Andreas and Hegeman, Malcolm, Life cycle assessment of dell r740 (2019). URL https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf
- [14] All Published SPEC SPECpower_ssj2008 Results, [Online; accessed 8. May 2024] (May 2024). URL https://www.spec.org/power_ssj2008/results/power_ssj2008.html
- [15] B. K. Sovacool, The precarious political economy of cobalt: Balancing prosperity, poverty, and brutality in artisanal and industrial mining in the democratic republic of the congo, *The Extractive Industries and Society* 6 (3) (2019) 915–939. doi:10.1016/j.exis.2019.05.018.
- [16] G. Pitron, The Dark Cloud: how the digital world is costing the earth, Scribe Publications, 2023.
- [17] Jan Konietzko, Moving beyond carbon tunnel vision with a sustainability data strategy (2022). URL <https://www.cognizant.com/us/en/insights/insights-blog/moving-beyond-carbon-tunnel-vision-with-a-sustainability-data-strategy-codex7121>
- [18] R. L. Lewison, A. F. Johnson, J. Gan, R. Pelc, K. Westfall, M. Helvey, Accounting for unintended consequences of resource policy: Connecting research that addresses displacement of environmental impacts, *Conservation Letters* 12 (3) (2019) e12628. doi:10.1111/conl.12628.
- [19] M. Amidpour, M. H. Khoshgoftar Manesh, Chapter 11 - environmental impacts consideration, in: M. Amidpour, M. H. Khoshgoftar Manesh (Eds.), *Cogeneration and Polygeneration Systems*, Academic Press, 2021, pp. 195–213. doi:10.1016/B978-0-12-817249-0.00011-2.
- [20] W. Klöpffer, Life cycle assessment, *Environmental Science and Pollution Research* 4 (4) (1997) 223–228. doi:10.1007/BF02986351.
- [21] I. Boustead, LCA – how it came about: – the beginning in the u.k., *The International Journal of Life Cycle Assessment* 1 (3) (1996) 147–150. doi:10.1007/BF02978943.
- [22] R. G. Hunt, W. E. Franklin, R. G. Hunt, LCA – how it came about: – personal reflections on the origin and the development of LCA in the USA, *The International Journal of Life Cycle Assessment* 1 (1) (1996) 4–7. doi:10.1007/BF02978624.
- [23] V. Forti, C. Baldé, R. Kuehr, G. Bel, The Global E-waste Monitor 2020. Quantities, flows, and the circular economy potential, United Nations University

- (UNU)/United Nations Institute for Training and Research (UNITAR) – co-hosted SCYCLE Programme, International Telecommunication Union (ITU) & International Solid Waste Association (ISWA), 2020.
- [24] S. Fazio, F. Biganzoli, V. De Laurentiis, L. Zampori, S. Sala, E. Diaconu, Supporting information to the characterisation factors of recommended of life cycle impact assessment methods (2018). doi : 10.2760/002447.
- [25] L. van Oers, J. B. Guinée, R. Heijungs, Abiotic resource depletion potentials (ADPs) for elements revisited—updating ultimate reserve estimates and introducing time series for production data, *The International Journal of Life Cycle Assessment* 25 (2) (2020) 294–308. doi : 10.1007/s11367-019-01683-x.
- [26] R. Frischknecht, F. Wyss, S. Büsser Knöpfel, T. Lützkendorf, M. Balouktsi, Cumulative energy demand in LCA: the energy harvested approach, *The International Journal of Life Cycle Assessment* 20 (7) (2015) 957–969. doi : 10.1007/s11367-015-0897-4.
- [27] T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, P. M. Midgley, *Climate change 2013: The physical science basis*, Tech. rep., Intergovernmental Panel on Climate Change, Working Group I Contribution to the IPCC Fifth Assessment Report (AR5)(Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.), 1535 pp. (2013).
- [28] U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, C.-J. Wu, ACT: designing sustainable computer systems with an architectural carbon modeling tool, in: *Proceedings of the 49th Annual International Symposium on Computer Architecture, ISCA '22*, Association for Computing Machinery, 2022, pp. 784–799. doi : 10.1145/3470496.3527408.
- [29] T. Pirson, T. P. Delhaye, A. G. Pip, G. Le Brun, D. Bol, The Environmental Footprint of IC Production: Review, Analysis, and Lessons From Historical Trends, *IEEE Trans. Semicond. Manuf. PP* (99) (2023) 1. doi : 10.1109/TSM.2022.3228311.
- [30] TechPowerUp, Cpu specs database, [Online; accessed 18. Mar. 2024] (2024). URL <https://www.techpowerup.com/cpu-specs/>
- [31] J. Gröger, R. Liu, L. Stobbe, J. Druschke, N. Richter, Green Cloud Computing: lebenszyklusbasierte Datenerhebung zu Umweltwirkungen des Cloud Computing: Abschlussbericht, Umweltbundesamt, 2021. URL https://www.umweltbundesamt.de/sites/default/files/medien/5750/publikationen/2021-06-17_texte_94-2021_green-cloud-computing.pdf
- [32] J. Choe, Memory technology 2021: Trends & challenges, in: *International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*, 2021, pp. 111–115. doi : 10.1109/SISPAD54002.2021.9592547.
- [33] B. Davy, Building an aws ec2 carbon emissions dataset, [Online; accessed 22. Jun. 2024] (2021). URL <https://medium.com/teads-engineering/building-an-aws-ec2-carbon-emissions-dataset-3f0fd76c98ac>
- [34] A. Noureddine, A. Bourdon, R. Rouvoy, L. Seinturier, Runtime monitoring of software energy hotspots, in: *ASE, ACM*, 2012, pp. 160–169.
- [35] TechPowerUp, Ssd specs database, [Online; accessed 18. Mar. 2024] (2024). URL <https://www.techpowerup.com/ssd-specs/>
- [36] Y. Aiouch, A. Chanoine, L. Corbet, P. Drapeau, L. Ollion, V. Vigneron, Evaluation de l'impact environnemental du numérique en France et analyse prospective - La librairie ADEME (2022). URL <https://librairie.ademe.fr/consommer-autrement/5226-evaluation-de-l-impact-environnemental-du-numerique-en-france-et-analyse-prospective.html>
- [37] C. L. Belady, C. G. Malone, Metrics and an infrastructure model to evaluate data center efficiency, in: *Proceedings of the ASME 2007 InterPACK Conference collocated with the ASME/JSME 2007 Thermal Engineering Heat Transfer Summer Conference*, Vol. ASME 2007 InterPACK Conference, Volume 1, 2007, pp. 751–755. doi : 10.1115/IPACK2007-33338.
- [38] N. Scarlat, M. Prussi, M. Padella, Quantification of the carbon intensity of electricity produced and used in Europe, *Appl. Energy* 305 (2022) 117901. doi : 10.1016/j.apenergy.2021.117901.
- [39] J. W. Pierre Delforge, America's data centers consuming and wasting growing amounts of energy, Tech. rep., NRDC (2015). URL <https://www.nrdc.org/resources/americas-data-centers-consuming-and-wasting-growing-amounts-energy>
- [40] P. Jacquet, T. Ledoux, R. Rouvoy, ScroogeVM: Boosting cloud resource utilization with dynamic oversubscription, *IEEE Transactions on Sustainable Computing* (2024) 1–13Publisher: IEEE. doi : 10.1109/TSUSC.2024.3369333.
- [41] É. Guégain, T. Simon, A. Rahier, R. Rouvoy, Managing Uncertainties in ICT Services Life Cycle Assessment using Fuzzy Logic, in: *ICT4S'24 - 10th International Conference on ICT for Sustainability*, IEEE, Stockholm, Sweden, 2024.
- [42] T. Simon, P. Rust, R. Rouvoy, J. Penhoat, Uncovering the environmental impact of software life cycle, in: *2023 International Conference on ICT for Sustainability (ICT4S)*, 2023, pp. 176–187. doi : 10.1109/ICT4S58814.2023.00026.
- [43] J. V. Andreasen, Carbon Efficient Karpenter: Optimizing Kubernetes Cluster Autoscaling for Carbon Efficiency, Master's thesis, Aarhus University (2024). URL <https://github.com/JacobValdemar/carbon-efficient-karpenter-thesis>
- [44] M. Proske, Ecodesign concepts and environmental assessments addressing product obsolescence, Ph.D. thesis, TU Berlin, <https://depositonce.tu-berlin.de/bitstreams/677ac517-6c03-41c2-b285-5a1449771df3/download> (2023).
- [45] M. Cortes-Cornax, P. Lago, C. Roncancio, A case study of environmental considerations and opportunities in cyber physical systems, *IET Software* 17 (4) (2023) 424–434. doi : 10.1049/sfw2.12130.
- [46] B. Everman, T. Villwock, D. Chen, N. Soto, O. Zhang, Z. Zong, Evaluating the carbon impact of large language models at the inference stage, in: *2023 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, 2023, pp. 150–157. doi : 10.1109/IPCCC59175.2023.10253886.
- [47] A. Berthelot, E. Caron, M. Jay, L. Lefèvre, Estimating the environmental impact of generative-ai services using an lca-based methodology, *Procedia CIRP* 122 (2024) 707–712, 31st CIRP Conference on Life Cycle Engineering. doi : 10.1016/j.procir.2024.01.098.
- [48] N. Patel, A. Mamandipoor, D. Quinn, M. Alian, Xfm: Accelerated software-defined far memory, in: *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '23*, Association for Computing Machinery, New York, NY, USA, 2023, p. 769–783. doi : 10.1145/3613424.3623776.
- [49] C. Morand, A.-L. Ligozat, A. Névéol, MLCA: a tool for Machine Learning Life Cycle Assessment, in: *ICT4S'24 - 10th International Conference on ICT for Sustainability*, IEEE, Stockholm, Sweden, 2024.