



HAL
open science

Fine-grained methodology to assess environmental impact of a set of digital services

Adrien Berthelot, Eddy Caron, Romain de Laage, Laurent Lefèvre, Alexis Nicolas

► **To cite this version:**

Adrien Berthelot, Eddy Caron, Romain de Laage, Laurent Lefèvre, Alexis Nicolas. Fine-grained methodology to assess environmental impact of a set of digital services. 2024. <hal-04928998>

HAL Id: hal-04928998

<https://hal.science/hal-04928998v1>

Preprint submitted on 4 Feb 2025

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.



HAL Authorization

Fine-grained methodology to assess environmental impact of a set of digital services

Adrien Berthelot^{1,2}, Eddy Caron¹, Romain De Laage³, Laurent Lefèvre¹, and Alexis Nicolas^{2 *}

¹École Normale Supérieure de Lyon, Université Lyon 1, CNRS, Inria, LIP

²Octo Technology

³Université de Neuchâtel

Abstract

The digital sector does not have a neglectable footprint, contributing to global climate change and over-commitment of planetary boundaries. To face environmental challenges, digital services could provide help and even mitigate other sector environmental impact. But, as the assessment of the environmental impacts of digital services remains difficult, digital services are still uncertain solutions. Global assessments of the digital sector already exist but are still often uncompleted, limited to carbon emission, use phase of the equipment and are made at a too large scale to be applicable. As digital services rely on highly mutualized equipment and infrastructure, it is difficult to assess the cost of a single service. In this paper, we propose a methodology to calculate the complete environmental impact of a service. We take a service provider's perspective on how, and with what data from service hosting, to calculate and reduce the service footprint. From a host point of view, we distribute the footprint between the services, allowing to consider precisely the impacts. With such data on the environmental impacts of the hosted services, we propose strategies to prioritize optimization effort and decommissioning policies.

1 Introduction

The Cloud industry has played a major role in establishing the digital world as we know it. The transition to the Cloud is transforming the digital sector from an economy based on the exchange of software and hardware goods to one based on services. As an extension of network infrastructures, Cloud infrastructures have been an essential growth driver for the digital sector. Unfortunately, in view of current environmental issues, this growth raises a number of questions.

Global warming [1], and the overstepping of planetary boundaries [2], are not independent of the phenomenon described above. The Information and Communication Technology (ICT) industry is a significant contributor to global warming [3], while having other important environmental impacts [4], such as rare metal extraction [5] and water consumption [6], to name but a few. Maybe more worrying, this environmental footprint is characterized by strong growth, predicting ever-increasing resource consumption [7]. Moreover, considering the recent trends with the deployment of generative AI as Cloud services with heavy environmental cost [8, 9, 10].

While it's true that Cloud-based digital services

could help reduce the footprint of other sectors of human activity [11], this assertion remains questionable [12]. It is all the more so in a context where the services in question have sometimes been presented as "dematerialization", neglecting on principle their own environmental impact [13]. However, this invisibilization of costs could also be seen as a lack of skills, tools and data to assess the environmental impact of services based on various highly mutualized equipment and infrastructures such as the Cloud [14], making evaluation more complex than for other goods.

Whether to make Cloud services more sustainable by controlling or reducing their impact, or to judge the environmental benefits of a Cloud-based solution, better measurement of environmental impacts is needed. To this end, this article proposes a methodology for calculating impacts on the scale of Cloud services. Whether the services are hosted on a physical instance, e.g. bare metal, or virtual, the methodology enables the impact of hosting to be distributed between services. It also provides the tools to measure the end-to-end impact of each service using hosting information. The method is designed to integrate Life Cycle Assessment (LCA) data to provide an estimate of the various environmental impacts of each service.

These individual assessments of all services then allow the construction of an environmental dashboard enabling informed decisions to control or reduce the

*Cite as: Adrien Berthelot, Eddy Caron, Romain de Laage, Laurent Lefevre, and Alexis Nicolas. 2025. Fine-grained methodology to assess environmental impact of a set of digital services. *Preprint undergoing review*.

service footprint. In this way, services can be prioritized for optimization, downsizing or even closure. Finally, these environmental assessments of services will enable managers to report the environmental impacts of usage to users or customers. As well as helping to comply with relevant standards, this reporting will provide useful data to help judge the environmental benefits or otherwise of Cloud-based solutions.

The rest of this article is organized as follows. Section 2 provides background on sustainability in Cloud, exposing the current limits of environmental assessment of Cloud-based services and presenting why this article could help to reduce the footprint of Cloud activities. Section 3 presents the methodology to assess the environmental impacts of several Cloud-based services. Then, this methodology is applied to a concrete use case, a running service provider, in Section 4. Extending the methodology, following the use case experience, Section 5 highlights 3 methodological contributions are made to the main methodology in order to build better Cloud footprint reduction strategies. Lastly, Section 9 drawn conclusions and guidelines for future research.

2 Background : a long road to sustainability in Cloud computing

Sustainability is a vast and varied subject. It can be assessed in terms of economic, social and environmental objectives [15]. Even when restricted to these environmental aspects, the question of sustainability is broader than that of impact assessment [16]. In fact, the use of a technology has both direct and complex indirect impacts on a number of scales [17]. Taking existing sustainability assessment frameworks for ICTs as a reference [18, 19, 20], this article will discuss sustainability in terms of the direct environmental impacts of Cloud-based services.

This focus, however limited in relation to the previous scales presented, is still insufficiently covered in current work. The various criticisms and current limitations of Cloud-based environmental impact assessment can be summarized as follows:

- **Only considering usage phase :** While the use phase of equipment and infrastructures is obviously studied, the manufacturing phase, and even more so the end-of-life phase, is very rarely taken into account. This is a cause for concern, given that the ICT footprint is less and less determined by the usage phase [21]. The high cost of manufacturing and the complex issue of ICT end-of-life [22] make it essential to factor this cost into environmental impacts.

- **Only considering carbon or electricity :** Related to the previous point, environmental impacts, and even sustainability, are all too often summed up in terms of electricity consumption. If carbon impact is studied at all, it's generally only with regard to the usage phase, i.e. the carbon intensity of the electricity mix used. However, the environmental impact of digital technology, including the Cloud, is much more diverse [4]. This narrow focus runs the risk of confining work on Cloud sustainability to a flawed carbon tunnel vision [23].
- **Only considering large scale :** A less obvious problem, the question of scale is nonetheless decisive in establishing the link between impact assessments and more sustainable Cloud computing. Assessments on the scale of a data-center or infrastructure, based on an inventory of equipment and flows (electricity, water), are disconnected from usage. This lack of user experience is the reason for the choice of service scale proposed in this article.

The definition of service can fluctuate, from the ITU-defined ICT services [24] to the Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) of the Cloud domain [25]. It will therefore be accepted to define a digital service as the interaction of IT equipment and infrastructure to satisfy a need. A digital service is generally divided into 3 parts, as shown in Figure 1.



Figure 1: Three main parts of a digital service

The service scale provides a more comprehensive representation of all the impacts induced by the Cloud infrastructure and its use. This link between Cloud and end-use makes it easier to compare the environmental cost of Cloud-based solutions with each other, or with alternatives. It also enables us to better establish policies for reducing our environmental footprint of Cloud-based services. To better illustrate this contribution, consider 3 common levers for reducing environmental footprint for a provider of Cloud-based services :

- **1. Acting on purchasing:** The simplest lever is to rely on what the market has to offer, and purchase equipment or flows, i.e. electricity, with a lower environmental footprint. It can also be achieved by extending the lifespan of equipment, thereby directly reducing environmental costs. This purchasing policy can even incorporate the

wider issues of equipment recycling and the circular economy [26]. The choice of data-center location has a major influence on the environmental cost of the electricity consumed. Dynamic shifting of workloads between different sites according to the availability of renewable energies [27] straddles the 2nd lever.

- **2. Optimizing:** Doing more with less, or making the most of existing resources. This may involve technical optimization to achieve efficiency gains in Cloud management, as detailed here [28]. This lever also includes gains made through research to reduce the Power Usage Effectiveness (PUE) of sites [29]. It also takes into account efficiency gains on the network [30] and client-side application [31] parts of services. However, with a few exceptions [32, 33], these various approaches remain focused on energy efficiency and its impact on carbon emissions, within the limits mentioned above.

These first two levers represent an already substantial and evolving set of strategies for reducing our environmental footprint. But what if these strategies are insufficient [34], i.e. if they don't enable us to achieve our sustainability objectives? Environmental issues are becoming ever more pressing [35, 36], and there is still considerable uncertainty about the ability of digital services to lighten the net footprint of human activities [37, 38, 39, 40]. It is therefore necessary to consider other possible strategies:

- **3. Scale down** In the context described above, this lever could be indispensable. It would mean closing down services or functionalities, and moving towards digital sufficiency [41]. The question of how, however, remains vast [42, 43]. What should be closed? For which sparing? And to what degree, from a small lowering of quality of service to the closure of one or more services?

This is where the scale of services really becomes important. We need to model the link between usage, resources mobilized and environmental impact in order to organize a coherent reduction in services. In the case of a large-scale environmental cost model, e.g. a datacenter, the consequences of resource reductions are uncertain. The highly mutualized nature of equipment in Cloud-based services makes it nontrivial to consider the consequences on usage of shutting down a piece of equipment.

This is why it is necessary to have a service-wide assessment to better identify the levers for reducing the footprint. We can better understand the footprint of uses to judge whether they are beneficial or not, and whether it is worth preserving them, changing them or close them.

3 General methodology

To calculate the footprint at service level, the methodology uses concepts from LCA, environmental data from LCA and more traditional application monitoring and logging tools. The methodology is run from the point of view of the service manager, from the Cloud where the services are hosted. The core of the methodology is the distribution of the environmental impacts of equipment and infrastructure across the services managed by the host. It is therefore indifferent to the type of hosting, and to the type of environmental data input. What's important is the distribution of the footprint and its influence on the chosen perimeter. First, it is important to present the tools used to supply the data.

3.1 Tools

Life cycle assessment LCA is a global, multi-criteria evaluation method based on the ISO 14040 and 14044 standards [44]. It can be based on complementary standards depending on the sector studied, such as the ITU (International Telecom Union) L1410 standard for ICT goods, networks, and services [24]. It aims to produce an evaluation of the potential environmental impacts of a product or activity. It is global as it considers the entire life cycle of the product(s) studied, using a cradle-to-grave approach. It is also multi-criteria, as it can quantify the impacts by different environmental indicators called *impact categories*. As pictured in Figure 2, an LCA is composed of four interdependent phases. The goal is to, for a defined purpose and perimeter (step 1), account for all the sub-products and elementary flows needed for the study's subject (step 2), and sum their environmental impacts given by life cycle inventory (LCI) data [45] (step 3). Step 4 questions the potential conclusions regarding the initial goals of the study and the uncertainties regarding the hypotheses taken during the previous phases.

Although LCA has only recently been applied in the digital sector [46, 47] compared to other sectors [48, 49], it is widely recognized for its specific qualities. It enables a more comprehensive assessment by taking into account the complete life cycle and the different impact categories, thus avoiding focusing solely on the carbon emissions of the use phase. Moreover, while its use of assumptions is criticized, it nonetheless enables relevant estimates to be made in a context such as that of digital technology, where the industry's lack of transparency [50, 51] could block the study of its damage to the environment. LCA therefore has the best qualities to produce useful quantified estimates of IT products and services footprint. Results that would be reliable

data for decision making or to be re-used, in comparative studies or large scale bottom-up studies. They are many subtype of LCA [52], one target of this article is to use LCA results as input data of environmental footprint and use them in an LCA-inspired framework. This framework will combine the environmental with activity data to produce detailed estimates by services.




Figure 2: Four main stages of LCA

Monitoring activities In order to measure the resource consumption of each service, metrics in terms of CPU usage and RAM usage are calculated at regular intervals. These metrics can be recorded in, for example, a VictoriaMetrics [53] time series database for further processing. The polling frequency of the resources consumption metrics is important. A low value for this interval means greater fidelity in the values collected, as resources consumption can vary greatly from one moment to the next. But a shorter reading frequency also means a higher cost in terms of computing and storage resources [54]. A compromise has to be found between cost and measurement accuracy. In applying the methodology, if the observation period is long enough, a short survey interval is not necessarily necessary. The main issue stays the distribution of the footprint between services. If a longer interval potentially "misses" more CPU or RAM consumption spikes, we can expect, over a long enough period, that these "misses" will be evenly distributed between services.

When retrieving metrics, we must associate them to each service we want to study. The use of containerization (eg. Docker) helped us to achieve this goal thanks to the Linux cgroups technology. Indeed, cgroups can be used to isolate a group of processes and permit to account for utilization of resources of this group of processes [55].

In order to calculate the footprint of a service per user, we need to retrieve information about service usage. All services are accessed through reverse proxies. These components issue access logs containing, among other things, the date and time of access, the IP address of the client accessing the service, its user-agent, the name of the service accessed, the access URL and the amount of transferred data. Analysis of these logs enables us to deduce the number

of accesses, the number of unique visitors and the amount of data transferred. We use GoAccess [56] to perform these analyses. The logs are centralized using Loki [57] and analyzed daily to generate a report with a spread of processing operations. GoAccess considers a unique combination of IP address, User-Agent and date as a single visitor. This number of "visitors" is of course not an exact representation of the number of users, but it does provide a good indicator of affluence. This approximation is useful for public services without authentication. It can, of course, be replaced by more precise data for services with logins and user accounts.

3.2 Distributing the footprint of hosting between services

The methodology calculates the environmental impact of SaaS solutions, which will be referred to as services, managed by a single hosting provider. The evaluation is carried out a posteriori for a defined period of time. Given an hosting capacity H and a set of services S . S can be broken down into several services, s_1, s_2, \dots, s_n , which are hosted on H . H can be a server or a Virtual Machine (VM), or it can be broken down into different parts h_1, h_2, \dots, h_m . A server can host several VMs. The important thing is that the resources represented by H are dedicated solely to S services. For a given period of time, H has an environmental impact I_{Host} . Thus :

$$I_{Host}(S) = \sum_n I_{Host}(s_n) \quad (1)$$

The I_{Host} impact can correspond to one or more environmental impact categories, such as global warming potential (GWP) in kg CO_2 equivalent, or abiotic depletion potential (ADP) in kg Sb equivalent. Using the terminology developed in [21], I_{Host} can be divided into $I_{Host_{capex}}$ and $I_{Host_{opex}}$, as (2). Capex, for capital expenditure, corresponds to a one-time cost, here the impact for manufacturing, transport and end-of-life, it represents the embodied footprint. On the other hand, opex, for operational expenditures, corresponds to the running cost, the usage cost, e.g. the electricity consumption of equipment.

$$I_{Host}(S) = I_{Host_{capex}}(S) + I_{Host_{opex}}(S) \quad (2)$$

Capex of hosting The H footprint will be broken down into 3 resources, CPU, RAM and Storage, plus an extra category "Rest" accounting for what does not fit in other categories, as follows:

$$I_{Host}(S)_{capex} = I_{Host}(S)_{capex}^{CPU} + I_{Host}(S)_{capex}^{RAM} + I_{Host}(S)_{capex}^{Sto} + I_{Host}(S)_{capex}^{Rest} \quad (3)$$

Each s_n service is allocated a share of the footprint of each part. This share is determined by what is known in LCA as an allocation key, such as:

- $a_{CPU}(s_n, h_m)$ determined by the average s_n percentage of CPU usage in h_m , in %, proportional to the total sum of all the s_n average CPU usage in h_m .
- $a_{RAM}(s_n, h_m)$ determined by the average s_n percentage of RAM usage in h_m , in %, proportional to the total sum of all the s_n average RAM usage in h_m .
- $a_{Sto}(s_n, h_m)$ determined by the average s_n storage used in h_m , in %, proportional to the total sum of all the s_n average of used storage of h_m .
- $a_{Rest}(s_n, h_m)$ determined for each service s_n by the average of the three precedent allocation key.

These allocation keys each verify that, for H or one specific h_m :

$$\sum_n a_{CPU}(s_n, h_m) = 1 \quad (4)$$

The “rest” allocation serves to allocate parts of the footprint where usage is not yet measurable. For example, in a physical server, the footprint of the metal case or cabinet. This allocation key could also be used to embodied one resource such as CPU, RAM or Storage in a case where access to one metric is not possible.

The capex footprint of a service is therefore written as :

$$\begin{aligned} I_{Host}(s_n)_{capex} = & \sum_m a_{CPU}(s_n, h_m) \times I_{Host}(s_n, h_m)_{capex}^{CPU} \\ & + a_{RAM}(s_n, h_m) \times I_{Host}(s_n, h_m)_{capex}^{RAM} \\ & + a_{Sto}(s_n, h_m) \times I_{Host}(s_n, h_m)_{capex}^{Sto} \\ & + a_{Rest}(s_n, h_m) \times I_{Host}(s_n, h_m)_{capex}^{Rest} \end{aligned} \quad (5)$$

Opex of hosting The opex footprint of H will be considered here solely as the environmental impact of electricity consumption. Even if this represents the main part of the opex impact, it remains a simplification that leaves out other opex costs, e.g. water consumption [58]. The environmental impact of H can therefore be written as follows:

$$I_{Host_{opex}}(S) = \sum_{s,h} C_s \times EGM_h \times PUE_h \quad (6)$$

Where C is the electricity consumption of service s and EGM , i.e. electricity grid mix, the environmental impact factor associated with electricity consumption in the area of h .

There are several ways of estimating the consumption of hosting H . The low bound is the Idle state and the high bound a 100% load that rises to the TDP of the components. We then try to estimate H 's consumption between the two using this notion of load percentage. So we break down the opex footprint into 2 sub-parts:

$$I_{Host_{opex}}(S) = I_{Host_{opex}}^{Idle}(S) + I_{Host_{opex}}^{Charge}(S) \quad (7)$$

There are several ways of obtaining estimates of these values. They can be obtained by direct measurement [59], benchmarking, manufacturer information, or a mixture [60], as used in Section 4. For each service, the opex footprint is allocated according to 2 allocation keys, such as :

- $a_{Idle}(s_n, h_m)$ equal to $a_{Rest}(s_n, h_m)$
- $a_{Charge}(s_n, h_m)$ equal to $a_{CPU}(s_n, h_m)$ as the CPU remains the main indicator for electricity consumption.

Even if there could be RAM-intensive or disk-intensive workload driving the electricity consumption, CPU activity would also rise in these situations, making it a good indicator not the exclusive cause. So, determining C_s^{Charge} would need more than the CPU usage as input data. The opex footprint of a service s_n is therefore written as :

$$\begin{aligned} I_{Host}(s_n)_{opex} = & \sum_m a_{Idle}(s_n, h_m) \times I_{Host}(s_n, h_m)_{opex}^{CPU} \\ & + a_{Charge}(s_n, h_m) \times I_{Host}(s_n, h_m)_{opex}^{RAM} \end{aligned} \quad (8)$$

3.3 Encompass more to understand better

As shown in Figure 1, now that the impact of data-center hosting has been taken into account, a more accurate view of the service footprint involves considering the footprint of the other two parts. To do this, access logs will be used to assign each service an impact based on its measured usage. The data used are as follows:

- Number of visits per service
- The operating system (OS) associated with these visits
- Volume of data traffic generated per service

Although still using measurable data, this part requires a greater reliance on assumptions, due to the complexity [61] and incompleteness of the metrics characterizing users.

Networks The environmental impact of network usage, i.e. I_{Net} , can be estimated fairly easily via impact models per gigabyte of data transferred. Admittedly, this allocation per gigabyte can be criticized [62, 63, 64], as the cost of transferring a gigabyte is never identical. Nevertheless, this allocation makes it easy to fairly share I_{Net} between different services and requires fewer data and knowledge of the network than other methodologies [65]. The network impact of a service is written as follows:

$$I_{Net}(s_n) = \sum_{type} p_{type} \times d_{n,type} \times IF_{type} \quad (9)$$

Where p_{type} is the percentage of data traffic transferred by the network type and $d_{n,type}$ the data volume associated with the service s_n and network type. The IF_{type} variable is an impact factor, representing the average environmental impact of transferring a gigabyte over the network type. The network type includes its location and category, e.g. whether it's a mobile or fixed network, with mobile networks generally being more costly. It can be difficult to obtain the type of traffic associated with a specific service. Credible assumptions can be made based on the location of visits and the average OS used for these visits. It is reasonable to consider that a majority of traffic coming from mobile OS is related to mobile traffic and that the rest is standard traffic.

End user terminals The impact of terminals, i.e. I_{Term} , is the most difficult to quantify at service level. The e equipment used for a visit has a I_{capex} and I_{opex} footprint that needs to be attributed. The simplistic approach is to take the average duration t_n of a visit for a service s_n and construct an allocation key. An ideal footprint calculation would give, developing the time-based allocation keys, the following equations for **one visit**:

$$\begin{aligned} I_{Term_{capex}}(visit) &= \frac{t}{totalusetime_e} \times I_{capex}(e) \\ I_{Term_{opex}}(visit) &= C_e \times t \times EGM_n \end{aligned} \quad (10)$$

The capex allocation is made using the notion of total use-time of the equipment e . It differs from the lifespan by only taking in account the active use-time. It implies that the equipment e has a total use-time corresponding to only a share of its lifespan. For example, a professional computer may be used 8 hours a day, making its total use-time only a third of its lifespan. This influence greatly $I_{Term_{capex}}$, even more considering 8 hours as an already high average daily use. Equations (10) are also greatly influenced by the visit duration. A notion that may have sense for video streaming service, but less for quantifying

the usage intensity of an online chat where user are connected all day. Moreover, the exact type of equipment e used as end-user terminal can be difficult to detect. The OS of visit can surely identify smartphone users but cannot make the difference between laptop and desktop users, equipment with fairly different footprint. Lastly, it is still hard to measure C_e , the electricity consumption for a visit of time t on e , from a server-side perspective without heavy monitoring of user devices [66]. So, we propose the following simplification for the footprint of terminals of a service s_n :

$$I_{Term}(s_n) = \sum_e p_e \times a_e(t_n) \times (I_{capex}(e) + I_{opex}(e)) \quad (11)$$

Where p_e is the average percentage of use of e . Here $a_e(t_n)$ is a unique time-based allocation for both capex and opex, considering $I_{opex}(e)$ as all the operational impact of e for a same lifespan as used to determined $I_{capex}(e)$. This equation would give an estimate of terminals impact with this abstract average time of visit t_n . More details and discussions are available in Section 5.3.

4 Use case

4.1 Context

Our use case required access to a service provider, giving us access to its access logs and detailed knowledge of its hosting infrastructure. We needed to observe and apply the methodology to heterogeneous services in use by a diverse audience. We chose Picasoft [67], as their managers were able to provide us with this data and own their servers, which enabled us to estimate I_{Host} reliably. Picasoft, is a local service provider with heterogeneous services, for a fine-grained test. We needed such a service provider running several heterogeneous services on the same servers with full access to data during one month.

Picasoft provides freely a set of open-source SaaS solutions to various users described in Table 1. Although, some SaaS solution are providing one specific services, some services are a group of SaaS solution. This is a functional view. For example, Picasoft provides a digital workspace for 3 different associations, including themselves. If both use mainly Nextcloud, it is on different instance, the applications are only accessible to a set of registered users, and they use it alongside other applications. Last example of group, we defined the administration of the server as a service. We included in this service, every application that are needed for all the other service. This includes, the systems metrics use by

administrator and this study, background processes and the online documentation, available as a Wiki.

We conducted a one-month study, from November 17 to December 17, 2023. This period of time allows us to avoid smoothing out the effect of abnormal days or weeks, except certain backup procedures, which were excluded from the study. This period also contains no vacations likely to influence user behavior. The metrics were recorded in a VictoriaMetrics [53] time series database for processing. The polling frequency of the resources consumption metrics has been set to 10 seconds. As explained in section 3.1, we believe that the observation period is long enough for the readings to be balanced across services.

Table 1: Services description

Service name	SaaS solution	Description
podcast	Castopod	Podcast platform
tube	Peertube	Peer-to-peer video platform
team	Mattermost	Online chat
pad	Etherpad	Online text editor
md	CodiMD	Online markdown editor
kanban	Wekan	Project management
voice	Mumble	Voice over IP
drop	Lufi	File sharing service
mobilizon	Mobilizon	Event planning
board	Whiteboard	Virtual whiteboard
paste	Privatebin	Online Pastebin
Other	N/A	Student projects
mail	Postfix	Send email
Picasoft association	Nextcloud (mainly)	Group of services for Picasoft
CET association	Nextcloud (mainly)	Group of services for CET ^a
Solidaire association	Nextcloud (mainly)	Group of services for Solidaire ^b
blog	Plume	Blogging engine
Server administration	Group of services	Wiki documentation, monitoring tools

^a: Local ecological association:

<https://www.compiegne-en-transition.org/fr>

^b: Student association:

<https://solidaires-etudiant-e-s.org/site/syndicats/>

The results of the study will, however, be influenced by the strictly French context and the absence of back-up mechanisms. This will imply a very low carbon footprint for I_{Net} and opex costs, due to

the specificity of the French electricity mix, which is highly decarbonized. For this study, we decided to observe 3 categories of environmental impact:

- Abiotic Depletion Potential (ADP) which contribution to scarcity of minerals and metals resources [68].
- Global Warming Potential (GWP) [69] which represents contribution to climate change.
- Primary Energy (PE) which represents the cumulative energy demand [70], i.e. the total energy footprint. This category is not strictly speaking an environmental impact category, but we consider it is a relevant sustainability criterion, as energy is not unlimited.

We choose these categories for data availability reasons and because they cover the most of ICT footprint [4]. The methodological framework for this case study is summarized in Figure 3. We used environmental impact data from Boavizta [71] for I_{Host} and from Nega-Octet [72] for I_{Net} and I_{Term} . The impact are then distributed between services according to metrics collected by GoAccess and VictoriaMetrics, as explained in section 3.1. GoAccess collect HTTP logs able to describe the number of visit, visits origin and volume of data transferred per services. VictoriaMetrics on its side retrieves data from the dockerstats command to display CPU and RAM use per docker containers on each VM. These data are then aggregated per services to compute a_{CPU} and a_{RAM} , a_{Sto} is directly observed with VictoriaMetrics.

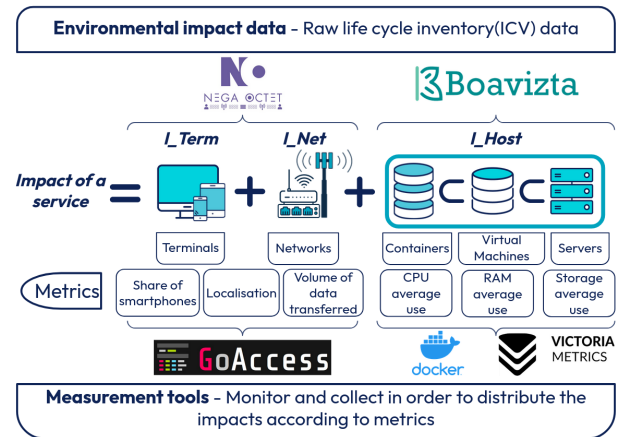


Figure 3: Synthesis of the used methodology, tools and data for the Picasoft case study

4.2 Results

Figure 4 shows the distribution of each impact category between services, while Fig. 5 shows the relationship between services and impact sources for the GWP category only. These two figures give us a

good idea of the quantity and distribution of impacts, allowing some preliminary observations. Firstly, on Figure 4, we can see the relative homogeneity of the cost of services regardless of impact category, i.e. a service is assigned a similar percentage of the footprint in each impact category. Secondly, as expected in Figure 5, the low share of opex and network can be explained by the French context, i.e. decarbonized electricity mix.




Figure 4: Environmental impact distribution of Picasoft services for a month, including hosting and networks impact

While these initial results enable us to report on environmental impacts, they do not in themselves necessarily provide the tools to apply footprint reduction strategies. To this end, we present 3 additions to the basic method in order to better consider the responsibility for impact between departments.

5 Strategies to lower the environmental impact

These strategies are extensions to the calculation methodology. They are intended to assist in the estab-

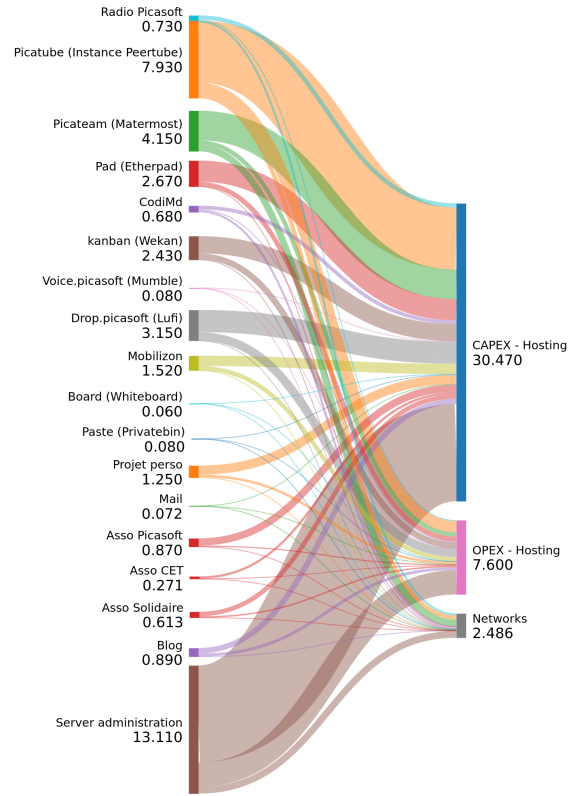


Figure 5: Carbon impact, in kg eq CO₂, of Picasoft services for a month, including hosting and networks impact

lishment of footprint reduction policies. This support is provided by giving different ways of comparing services with each other, for a manager, from a Cloud point of view. Which service should be studied first from an environmental point of view, whether to optimize it, reduce its functionality, limit it or shut it down?

The Picasoft use case will be used to illustrate these contributions, although it is a smaller-scale use case than other SaaS solution deployments.

5.1 Footprint per visitor as an efficiency indicator

The first approach to choosing which service to focus on might naturally be to concentrate on the service with the largest footprint. However, a service may represent a large footprint simply because it's used a lot, leaving little room for optimization. Conversely, we propose to observe in Figure 6 and 7 the footprint of services divided by the number of visitors, in proportion to services with similar footprints. We can then see disparities in terms of efficiency between services of the same order of magnitude in terms of impact.

The footprint divided here is the one illustrated above, including hosting and network. We can see

that, despite a large footprint, some services seem to be more efficient at handling a large number of visits. For example, in the case of server administration, even though some software packages (e.g. Loki or Traefik) involve very few visits at a high cost, this service includes technical documentation in the form of an online wiki. This wiki fulfills its function at a low cost per visit. Conversely, some services appear to be more expensive in this respect. It would be interesting to compare these services with other SaaS solutions meeting similar needs.

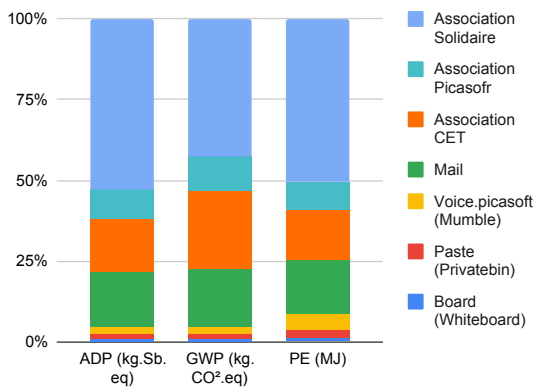


Figure 6: Footprint per visitor of the services covering together less than 5% of the footprint (without end-user terminals)

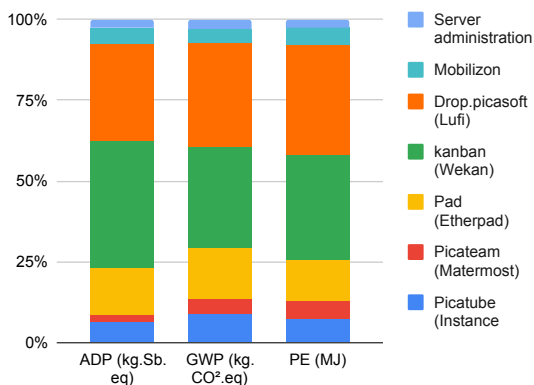


Figure 7: Footprint per visitor of the services covering together more than 80% of the footprint (without end-user terminals)

5.2 The cost of unpredictability

In Section 3.2, we observed that the capex footprint was distributed with allocations based on averages. While averaging makes it easy to share the use of functions such as RAM between services, it can mask problems. Efficiency, in the context of the hosting capex footprint, can be seen as the ability to perform the right sizing, e.g. which VM size for which service? One of the difficulties of this exercise is that

two services may have the same average RAM usage, but a stable need and the other regular peaks of very high consumption. In this situation, the second service pushes for an over-dimensioning of hosting resources. To represent this cost of instability in resource consumption, we have redone the allocations, but using root-mean-square instead of averages. In doing so, we attributed a greater share of the impact to services with abnormal resource utilization times. Table 2 shows the variation in footprint generated by this new distribution of the hosting capex footprint. Services with stable resource utilization are shown in negative and blue. Conversely, services with a larger footprint as a result of the new allocation calculation are shown in yellow.

The yellow part of the table shows the services most influenced by usage, but not necessarily to the same degree. In blue, we see certain services that are stable, such as administration, or with very few users, such as voice. However, even services with few users, such as associations, show strong variations in footprint. We can see that the ADP impact is the most heavily impacted, which is to be expected as this category is mainly influenced by the capex part, including manufacturing impacts.

5.3 The user's cost : where the rest of the footprint lie

While the size of the share of user terminals in the overall digital footprint is open to debate [73, 3], no estimate considers terminals to be a negligible part, even more if we consider multi-criteria assessment [4, 62]. End-user devices are an indispensable component of Cloud-based services. Although their management may be the responsibility of someone other than the service provider, i.e. the user, it would be abnormal not to take into account the footprint of the terminals enabling access to the Cloud-based service. If we want a fair calculation of the environmental impact of a service, we need to allocate to that service a share of the footprint of each terminal used to access the service.

This allocation problem is vast, and includes both methodological and data access limitations. A number of studies have highlighted the uncertainties associated with integrating user terminals into the cost of services [74, 75, 76]. As shown in Section 3.3, $I_Term(S)$ depends on the total duration of use of the terminal, i.e. $I_Term(S)$ depends on both the lifetime of the equipment and its average use by the user. In this way, $I_Term(S)$ depends both on the user's behavior during use of the service and outside of it.

A final limitation to the calculation of terminal footprints is the heterogeneity of services and terminals. Thus, the use of a service can induce significant addi-

Table 2: Footprint variation of CAPEX footprint from distribution by average activity to distribution by RMS activity

Services	Environmental impact variation		
	ADP (kg.Sb.eq)	GWP (kg.CO ₂ .eq)	PE (MJ)
podcast	15.42%	6.35%	6.93%
tube	18.36%	6.49%	7.26%
team	4.76%	1.34%	1.62%
pad	19.45%	10.98%	11.62%
md	-11.86%	-9.21%	-9.55%
kanban	-20.54%	-11.57%	-12.25%
voice	-32.91%	-20.02%	-20.74%
Drop	-41.39%	-15.62%	-17.34%
mobilizon	8.44%	4.68%	4.91%
board	6.83%	1.68%	1.93%
paste	32.31%	20.99%	21.79%
other	9.13%	5.96%	6.13%
mail	26.23%	10.19%	11.22%
Picasoft association	23.86%	8.37%	9.65%
CET association	7.60%	1.93%	2.37%
Solidaire association	-35.76%	-23.47%	-24.21%
blog	6.92%	5.27%	5.37%
Server administration	-9.88%	-4.15%	-4.61%

tional energy consumption on the terminal, the size of this consumption depending on the client-oriented workload of the service and the significant aspect depending on the consumption of the terminal.

The fundamental problem with this heterogeneity, however, remains the allocation criterion, i.e. time. We can allocate 3 hours of terminal time to a video service that has enabled 3 hours of streaming, on the assumption that the terminal is fully mobilized during these 3 hours. However, for an instant messaging service running all day in the background, or an audio podcast service, this time-based allocation loses its meaning. Even knowing how long the user uses the service via his terminal is not enough to allocate a fair share of the footprint to each service. In purely hypothetical terms, it would be necessary to have the same degree of information on terminals as on Cloud servers, which is, in this context, not possible.

However, in Figure 8, we have chosen to represent the cost of terminals according to the different services. We have made the deliberately simplistic assumption of representing each visitor as one minute's use of a personal terminal, irrespective of the service. As a reminder, the term 'visitor' refers to a 'single daily visitor', who may have made several visits to the service during the course of a day, using the same equipment, so this is a very low estimate. Despite the many limitations we have mentioned, we believe it is necessary to represent and consider the impact of terminals, knowing the limited nature of our data. Comparing Figure 5 and Figure 8, it should be noted that even with a minimalist hypothesis on the use of terminals, their carbon impact would be considerable. It's also interesting to note that taking into account the impact of terminals allows us to better appreciate the footprint of certain services that didn't seem to be so costly. The blog service, for example, has so far not been a topic of discussion, yet it generates significant use of terminals.

As with subsections 5.1 and 5.2, the aim here is not to quantify the exact impact, but to add another level of insight into the service footprint. Estimating the size of the terminal footprint between services is important for prioritizing reduction strategies. Even if it's difficult, from the service manager's point of view, to have a direct impact on the terminal part, it's not impossible. The eco-design of services, their ability to be used by the greatest number of terminals, e.g. by proposing lighter versions of services, could make it possible to reduce electricity consumption on the client side of the service[77]. What's more, such efforts could contribute to slowing down the renewal of terminals [78], lengthening their average lifespan, reducing their footprint and ultimately that of the

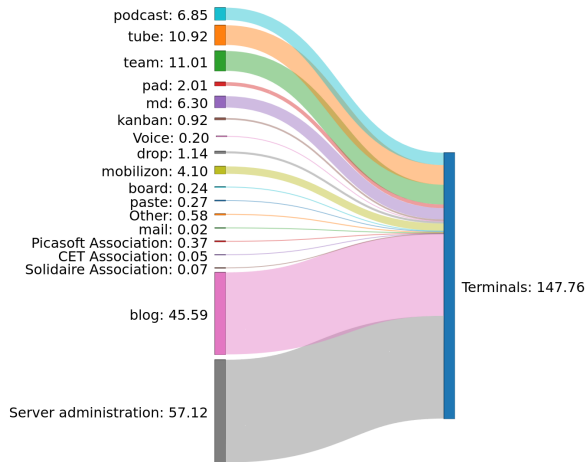


Figure 8: Carbon impact, in kg eq CO₂, of Picasoft service for a month including terminals impacts of a 1mn visit per visitors

associated services.

6 Study Limitations

The main limitation of this paper, apart from the inclusion of terminals discussed above, is the lack of pre-existing methodologies to enable a meaningful comparison. The framing of this methodology, i.e. the multi-criteria evaluation of a set of services hosted on the same resources, remains to the authors' knowledge unique. Other studies and methodologies exist on the notion of service footprints, but they are often limited in terms of impact categories and are not as detailed, focusing on a single service [79]. The authors come to the same conclusions as [79] on the lack of transparent scientific literature on the service footprint, and on the advantages of LCA as an approach for filling these gaps.

One other limitation to the fine allocation of impacts as proposed in the paper is the allocation made for storage. The problem is to allocate to each service a share of the footprint of the storage equipment. This share is of course proportional to the duration of the study and the lifespan of the storage disk, but once again there is the issue of sharing the equipment between services. How can we assess how much storage equipment is used by departments? The naive approach chosen in this article is to look at the average share of disk occupancy in relation to other services. It is assumed that the main criterion for sizing storage, and therefore its footprint, is its capacity in terms of size. However, this criterion can be questioned because the question of access frequency could have been a determining factor. Finally, is the service that 'uses' the disk the most the one that has the largest volume of data stored on it, or the one that performs the most read and write operations?

This paper focuses on the distribution of environmental impacts, but is clearly dependent on both I_{capex} and I_{opex} inputs. While advancements are made to provide a more detailed view of the opex footprint [80, 59], i.e. electricity consumption, the quality of I_{capex} data and its continuity over time leaves much to be desired¹.

7 Discussion

The main research question behind this article is how to calculate the end-to-end footprint of a digital service. Adopting the point of view of a service manager, we have detailed the hosting part and chosen simplified models for the other parts, even incomplete for the terminal part. More detailed approaches for the networks part could have been integrated [81, 65], but they require a level of knowledge of communication infrastructures that is difficult to obtain. Also, apart from specific services [63], the network part tends to have a smaller footprint. In addition, while offering few eco-design levers, i.e. it is not possible to modify the infrastructure but only the use made of it with our limited knowledge of how it works.

Terminals, on the other hand, account for a large proportion of the footprint and could therefore, at the cost of a change in user behavior [76, 82], reveal crucial levers for reducing the footprint of services. This behavior is largely influenced by the integration of the service on the user's terminal and influence back the provider equipments use [83]. In order to gain a better understanding of these levers, it is therefore important to gain a better understanding of or model the use made of terminals in the context of the service under study.

In this article, we have based ourselves on first-hand data: HTTP request logs. These data are reliable, and provide relevant information about the proportion of smartphone users, but do not provide any more information about user sessions per service, e.g. average duration, number of pages visited, etc. On a larger scale, it may be possible to use broader analysis tools, such as Similarweb [84] or HypeStat [85], but at the cost of a loss of precision and reliability [86, 87]. Rather than direct measurement, simulation could be a relevant line of research. Simulating a population of users and their usage, like [88], could enable different implementations of the service to be tested, both on the server side and on the user side. This could help to improve certain service footprint modelling projects such as e-footprint [89].

¹ During the writing of this article, the NegaOctet database project closed without public announcement

8 Future work

Naturally, part of the future work will focus on extending the methodology to include more environmental impact categories, greater precision on equipment consumption and a more detailed calculation of the footprint of terminals. The ideal outcome would be an automatic tool that can be integrated into a control panel. Such a tool would make it possible to continue studying the footprint of services on the basis of a larger number of experiments under real conditions, or at least in simulation.

Apart from this general methodology, it appeared to the authors that the storage function, especially with regard to backup and replication procedures, should be dealt with separately. It would be interesting to study storage as a service, using and adapting our LCA methodology for digital services. This is an active area of research [90, 91] that it would be interesting to observe from the point of view of its environmental cost.

Finally, the services provided by Picasoft are free and open-source equivalents of well-known office suites provided by large companies. It would be interesting to compare the environmental costs of these solutions, which are functionally similar but have radically different implementations and scales.

9 Conclusion

In this article, we presented our methodology for assigning a full environmental cost to a cloud-hosted service. One of the main contributions of our methodology is the distribution of the environmental footprint of hosting across different heterogeneous digital services. This distribution is based on each service's use of the hosting's CPU, RAM and storage resources. In this way, our methodology makes it possible to attribute to each service a share of both the operational and embodied footprint of hosting. In addition, we have proposed methods for including the network and terminal impacts of services. In this way, we have produced a method for calculating the environmental cost of each service, on a full life-cycle and multi-criteria basis, surpassing current evaluation limits.

To illustrate the methodology and demonstrate its applicability, we produced a use case on the digital service provider Picasoft. We were able to observe all the services in real-life conditions and establish an assessment after one month. This assessment includes environmental impacts in three categories, for the supplier as well as for each service. In addition to this ability to report on the impact of each service, the detail of the assessment enabled us to reconsider

the footprint of each service in relation to the others from different angles.

We have therefore proposed 3 new perspectives for reconsidering the footprint of cloud-based services. These perspectives are intended to help prioritize effort between services as part of a footprint reduction policy. Which services should be improved, regulated or closed down, in priority, to meet the cloud service provider's environmental objectives?

The next step and main limit to our current approach remains the footprint of user terminals. How can we include, from the point of view of the service provider, observing and managing from the cloud, this part of the service impact that may be the main one? We have proposed a first approach, but it needs to be improved to adapt to the heterogeneity of services. We may need to move beyond time-based allocation to a more holistic approach. This inclusion of terminals will then enable us to obtain an end-to-end estimate of the environmental cost of each service. The repetition of use-cases such as ours will then enable us to build reliable impact prediction models. In this way, thanks to fine, continuous observation of the environmental impact of services, cloud-based service managers will be able to draw on new, reliable metrics to steer cloud activities towards sustainability.

Acknowledgment

The authors used data and models from the open-source project Boavizta (<https://boavizta.org/en>). The authors would also like to thank the Picasoft association for kindly hosting the use-case for this project. This work was funded by ANRT (CIFRE N° 2021/0576).

References

- [1] H. L. Core Writing Team and J. R. (eds.), *SYR Longer Report*. Geneva, Switzerland: IPCC, 2023. (visited on 03/04/2024).
- [2] K. Richardson *et al.*, "Earth beyond six of nine planetary boundaries," *Science advances*, vol. 9, no. 37, eadh2458, 2023.
- [3] C. Freitag, M. Berners-Lee, K. Widdicks, B. Knowles, G. S. Blair, and A. Friday, "The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations," *Patterns*, vol. 2, no. 9, p. 100340, Sep. 2021, ISSN: 2666-3899. DOI: [10.1016/j.patter.2021.100340](https://doi.org/10.1016/j.patter.2021.100340). (visited on 04/09/2022).

- [4] 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*, Green IT, Dec. 2021. (visited on 12/08/2021).
- [5] G. Pitron, *The Dark Cloud: how the digital world is costing the earth*. Scribe Publications, 2023, 282 pp., ISBN: 978-1-76138-510-0.
- [6] P. Li, J. Yang, M. A. Islam, and S. Ren, *Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models*, Apr. 6, 2023. DOI: [10.48550/arXiv.2304.03271](https://doi.org/10.48550/arXiv.2304.03271). arXiv: [2304.03271](https://arxiv.org/abs/2304.03271) [cs]. (visited on 04/14/2023).
- [7] International Energy Agency. "Electricity 2024 – analysis," IEA. (2024), (visited on 01/25/2024).
- [8] A. de Vries, "The growing energy footprint of artificial intelligence," *Joule*, vol. 7, no. 10, pp. 2191–2194, Oct. 18, 2023, ISSN: 2542-4351. DOI: [10.1016/j.joule.2023.09.004](https://doi.org/10.1016/j.joule.2023.09.004). (visited on 10/24/2023).
- [9] A. Berthelot, E. Caron, M. Jay, and L. Lefèvre, "Estimating the environmental impact of generative-AI services using an LCA-based methodology," in *31ST CIRP CONFERENCE ON LIFE CYCLE ENGINEERING*, Jun. 19, 2024. (visited on 02/07/2024).
- [10] A. S. Luccioni, Y. Jernite, and E. Strubell, *Power hungry processing: Watts driving the cost of AI deployment?* Nov. 28, 2023. DOI: [10.48550/arXiv.2311.16863](https://doi.org/10.48550/arXiv.2311.16863). arXiv: [2311.16863](https://arxiv.org/abs/2311.16863) [cs]. (visited on 12/01/2023).
- [11] IPCC - Working Group III: Mitigation of Climate Change. "Full report 2022, chapter 16.6.2." (2022), (visited on 04/25/2022).
- [12] A. Rasoldier, J. Combaz, A. Girault, K. Marquet, and S. Quinton, "How realistic are claims about the benefits of using digital technologies for GHG emissions mitigation?" In *Eighth Workshop on Computing within Limits 2022*, LIMITS, 245 Main St, Cambridge MA: PubPub, 2022, p. 14.
- [13] B. Fix, "Dematerialization through services: Evaluating the evidence," *BioPhysical Economics and Resource Quality*, vol. 4, no. 2, p. 6, 2019, ISSN: 2366-0120. DOI: [10.1007/s41247-019-0054-y](https://doi.org/10.1007/s41247-019-0054-y). (visited on 03/06/2024).
- [14] R. Itten *et al.*, "Digital transformation—life cycle assessment of digital services, multifunctional devices and cloud computing," *The International Journal of Life Cycle Assessment*, vol. 25, 2020. DOI: [10.1007/s11367-020-01801-0](https://doi.org/10.1007/s11367-020-01801-0).
- [15] U. Nations. "THE 17 GOALS | sustainable development." (), (visited on 03/07/2024).
- [16] A. W. Hjalsted, A. Laurent, M. M. Andersen, K. H. Olsen, M. Ryberg, and M. Hauschild, "Sharing the safe operating space: Exploring ethical allocation principles to operationalize the planetary boundaries and assess absolute sustainability at individual and industrial sector levels," *Journal of Industrial Ecology*, vol. 25, no. 1, pp. 6–19, 2021, ISSN: 1530-9290. DOI: [10.1111/jiec.13050](https://doi.org/10.1111/jiec.13050). (visited on 07/03/2023).
- [17] N. C. Horner, A. Shehabi, and I. L. Azevedo, "Known unknowns: Indirect energy effects of information and communication technology," *Environmental Research Letters*, vol. 11, no. 10, p. 103001, Oct. 5, 2016, Institution: Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States) Publisher: IOP Publishing, ISSN: 1748-9326. DOI: [10.1088/1748-9326/11/10/103001](https://doi.org/10.1088/1748-9326/11/10/103001). (visited on 12/07/2021).
- [18] L. M. Hilty and B. Aebischer, "ICT for sustainability: An emerging research field," in *ICT Innovations for Sustainability*, L. M. Hilty and B. Aebischer, Eds., ser. Advances in Intelligent Systems and Computing, Cham: Springer International Publishing, 2015, pp. 3–36, ISBN: 978-3-319-09228-7. DOI: [10.1007/978-3-319-09228-7_1](https://doi.org/10.1007/978-3-319-09228-7_1).
- [19] A. Lundström and D. Pargman, "Developing a framework for evaluating the sustainability of computing projects," in *Proceedings of the 2017 Workshop on Computing Within Limits*, ser. LIMITS '17, New York, NY, USA: Association for Computing Machinery, Jun. 22, 2017, pp. 111–117, ISBN: 978-1-4503-4950-5. DOI: [10.1145/3080556.3080562](https://doi.org/10.1145/3080556.3080562). (visited on 12/05/2023).
- [20] L. Grimal, I. Di Loreto, N. Burger, and N. Troussier, "Design of an interdisciplinary evaluation method for multi-scaled sustainability of computer-based projects. a workbased on the sustainable computing evaluation framework (SCEF)," *LIMITS Workshop on Computing within Limits*, 2021. DOI: [10.21428/bf6fb269.2ee80cf1](https://doi.org/10.21428/bf6fb269.2ee80cf1). (visited on 04/20/2022).
- [21] U. Gupta *et al.*, "Chasing carbon: The elusive environmental footprint of computing," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, Los Alamitos, CA, USA: IEEE Computer Society, Feb. 1, 2021, pp. 854–867, ISBN: 978-1-66542-235-2. DOI: [10.1109/HPCA51647.2021.00076](https://doi.org/10.1109/HPCA51647.2021.00076). (visited on 02/05/2023).

- [22] V. Forti, C. Baldé, R. Kuehr, and 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, ISBN: 978-92-808-9114-0.
- [23] Jan Konietzko. “Moving beyond carbon tunnel vision with a sustainability data strategy,” www.cognizant.com. Cognizant - Forbes. (Feb. 8, 2022), (visited on 03/07/2024).
- [24] ITU, *ITU I1410 : Methodology for environmental life cycle assessments of information and communication technology goods, networks and services*, Dec. 2014.
- [25] M. Hogan, F. Liu, A. Sokol, and J. Tong, “NIST cloud computing standards roadmap : Version 1.0,” in *NIST Publication 500-291*, Edition: 0 Report Number: NIST SP 500-291v1, Gaithersburg, MD: National Institute of Standards and Technology, 2011, NIST SP 500–291v1. DOI: [10.6028/NIST.SP.500-291v1](https://doi.org/10.6028/NIST.SP.500-291v1). (visited on 03/07/2024).
- [26] R. Bashroush, N. Rteil, R. Kenny, and A. Wynne, “Optimizing server refresh cycles: The case for circular economy with an aging moore’s law,” *IEEE Transactions on Sustainable Computing*, vol. 7, no. 1, pp. 189–200, 2022, Conference Name: IEEE Transactions on Sustainable Computing, ISSN: 2377-3782. DOI: [10.1109/TSUSC.2020.3035234](https://doi.org/10.1109/TSUSC.2020.3035234). (visited on 05/28/2024).
- [27] P. Wiesner, I. Behnke, D. Scheinert, K. Gontarska, and L. Thamsen, “Let’s wait awhile: How temporal workload shifting can reduce carbon emissions in the cloud,” in *Proceedings of the 22nd International Middleware Conference*, ser. Middleware ’21, New York, NY, USA: Association for Computing Machinery, 2021, pp. 260–272, ISBN: 978-1-4503-8534-3. DOI: [10.1145/3464298.3493399](https://doi.org/10.1145/3464298.3493399). (visited on 05/28/2024).
- [28] V. Ostapenco, L. Lefèvre, A.-C. Orgerie, and B. Fichel, “Modeling, evaluating and orchestrating heterogeneous environmental leverages for large scale data centers management,” *International Journal of High Performance Computing Applications*, 2023. (visited on 04/05/2023).
- [29] C. L. Belady and 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, Jul. 2007, pp. 751–755. DOI: [10.1115/IPACK2007-33338](https://doi.org/10.1115/IPACK2007-33338). eprint: https://asmedigitalcollection.asme.org/InterPACK/proceedings-pdf/InterPACK2007/42770/751/2666558/751_1.pdf.
- [30] A. Hounsel, K. Borgolte, P. Schmitt, J. Holland, and N. Feamster, “Comparing the effects of DNS, DoT, and DoH on web performance,” in *Proceedings of The Web Conference 2020*, ser. WWW ’20, New York, NY, USA: Association for Computing Machinery, 2020, pp. 562–572, ISBN: 978-1-4503-7023-3. DOI: [10.1145/336423.3380139](https://doi.org/10.1145/336423.3380139). (visited on 03/08/2024).
- [31] T. Hoßfeld, F. Metzger, and D. Rossi, “Speed index: Relating the industrial standard for user perceived web performance to web QoE,” in *2018 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*, ISSN: 2472-7814, 2018, pp. 1–6. DOI: [10.1109/QoMEX.2018.8463430](https://doi.org/10.1109/QoMEX.2018.8463430). (visited on 03/08/2024).
- [32] V. D. Reddy, B. Setz, G. S. V. R. K. Rao, G. R. Gangadharan, and M. Aiello, “Metrics for sustainable data centers,” *IEEE Transactions on Sustainable Computing*, vol. 2, no. 3, pp. 290–303, 2017, Conference Name: IEEE Transactions on Sustainable Computing, ISSN: 2377-3782. DOI: [10.1109/TSUSC.2017.2701883](https://doi.org/10.1109/TSUSC.2017.2701883). (visited on 05/28/2024).
- [33] M. Vasconcelos, D. Cordeiro, G. da Costa, F. Dufossé, J.-M. Nicod, and V. Rehn-Sonigo, “Optimal sizing of a globally distributed low carbon cloud federation,” in *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, 2023, pp. 203–215. DOI: [10.1109/CCGrid57682.2023.00028](https://doi.org/10.1109/CCGrid57682.2023.00028).
- [34] T. Sukprasert, A. Souza, N. Bashir, D. Irwin, and P. Shenoy, “On the limitations of carbon-aware temporal and spatial workload shifting in the cloud,” in *Proceedings of the Nineteenth European Conference on Computer Systems*, ser. EuroSys ’24, New York, NY, USA: Association for Computing Machinery, 2024, pp. 924–941, ISBN: 9798400704376. DOI: [10.1145/3627703.3650079](https://doi.org/10.1145/3627703.3650079). (visited on 05/28/2024).
- [35] K. Richardson *et al.*, “Earth beyond six of nine planetary boundaries,” *Science Advances*, vol. 9, no. 37, eadh2458, 2023, Publisher: American Association for the Advancement of Science. DOI: [10.1126/sciadv.adh2458](https://doi.org/10.1126/sciadv.adh2458). (visited on 03/08/2024).

- [36] S. Graham *et al.*, “An interdisciplinary framework for navigating social–climatic tipping points,” *People and Nature*, vol. 5, no. 5, pp. 1445–1456, 2023, ISSN: 2575-8314. DOI: [10.1002/pan3.10516](https://doi.org/10.1002/pan3.10516). (visited on 11/07/2023).
- [37] J. C. T. Bieser and L. M. Hilty, “Indirect effects of the digital transformation on environmental sustainability: Methodological challenges in assessing the greenhouse gas abatement potential of ICT,” in *EPiC Series in Computing*, ISSN: 2398-7340, vol. 52, EasyChair, 2018, pp. 68–81. DOI: [10.29007/1x7q](https://doi.org/10.29007/1x7q). (visited on 03/08/2024).
- [38] J. C. T. Bieser, R. Hintemann, L. M. Hilty, and S. Beucker, “A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology,” *Environmental Impact Assessment Review*, vol. 99, p. 107 033, Mar. 2023, ISSN: 0195-9255. DOI: [10.1016/j.eiar.2022.107033](https://doi.org/10.1016/j.eiar.2022.107033). (visited on 05/19/2023).
- [39] V. C. Coroama and F. Mattern, “Digital rebound - why digitalization will not redeem us our environmental sins,” in *Proceedings of the 6th International Conference on ICT for Sustainability, ICT4S 2019, Lappeenranta, Finland, June 10-14, 2019*, Aachen, Germany: CEUR-WS.org, 2019.
- [40] V. C. Coroamă and D. Pargman, “Skill rebound: On an unintended effect of digitalization,” in *Proceedings of the 7th International Conference on ICT for Sustainability*, ser. ICT4S2020, New York, NY, USA: Association for Computing Machinery, 2020, pp. 213–219, ISBN: 978-1-4503-7595-5. DOI: [10.1145/3401335.3401362](https://doi.org/10.1145/3401335.3401362). (visited on 07/06/2022).
- [41] T. Santarius *et al.*, “Digital sufficiency: Conceptual considerations for ICTs on a finite planet,” *Annals of Telecommunications*, May 12, 2022. (visited on 02/13/2023).
- [42] N. Fitzpatrick, T. Parrique, and I. Cosme, “Exploring degrowth policy proposals: A systematic mapping with thematic synthesis,” *Journal of Cleaner Production*, vol. 365, p. 132 764, 2022, ISSN: 0959-6526. DOI: [10.1016/j.jclepro.2022.132764](https://doi.org/10.1016/j.jclepro.2022.132764). (visited on 03/08/2024).
- [43] A. Monnin, E. Bonnet, and D. Landivar, “What the anthropocene does to organizations,” in *35th EGOS Colloquium, Enlightening the Future: The Challenge for Organizations, Sub-theme 67 Critical Organizational Anthropocene Studies*, Jul. 2, 2019.
- [44] W. Klöpffer, “Life cycle assessment,” *Environmental Science and Pollution Research*, vol. 4, no. 4, pp. 223–228, 1997, ISSN: 1614-7499. DOI: [10.1007/BF02986351](https://doi.org/10.1007/BF02986351). (visited on 03/13/2024).
- [45] European Commission. Joint Research Centre. Institute for Environment and Sustainability., *International Reference Life Cycle Data System (ILCD) Handbook :specific guide for Life Cycle Inventory (LCI) data sets*. LU: Publications Office, 2010. (visited on 04/15/2022).
- [46] B. Whitehead *et al.*, “The environmental burden of data centres – a screening LCA methodology,” in *CIBSE ASHRAE Technical Symposium, Imperial College, London UK, 222 Balham High Road, London: CIBSE, May 2012*, p. 17.
- [47] J. Malmodin, D. Lundén, M. Nilsson, and G. Andersson, “LCA of data transmission and IP core networks,” *2012 Electronics Goes Green 2012+*, pp. 1–6, Sep. 2012.
- [48] R. G. Hunt, W. E. Franklin, and 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*, vol. 1, no. 1, pp. 4–7, Mar. 1996, ISSN: 0948-3349, 1614-7502. DOI: [10.1007/BF02978624](https://doi.org/10.1007/BF02978624). (visited on 11/25/2021).
- [49] I. Boustead, “LCA — how it came about: — the beginning in the u.k.,” *The International Journal of Life Cycle Assessment*, vol. 1, no. 3, pp. 147–150, Sep. 1996, ISSN: 1614-7502. DOI: [10.1007/BF02978943](https://doi.org/10.1007/BF02978943). (visited on 12/07/2021).
- [50] T. Billstein, A. Björklund, and T. Rydberg, “Life cycle assessment of network traffic: A review of challenges and possible solutions,” *Sustainability*, vol. 13, no. 20, p. 11 155, Jan. 2021, Number: 20 Publisher: Multidisciplinary Digital Publishing Institute. DOI: [10.3390/su132011155](https://doi.org/10.3390/su132011155). (visited on 02/04/2022).
- [51] S. Pauliuk, G. Majeau-Bettez, C. Mutel, B. Steubing, and K. Stadler, “Lifting industrial ecology modeling to a new level of quality and transparency: A call for more transparent publications and a collaborative open source software framework,” *Journal of Industrial Ecology*, vol. 19, n/a–n/a, Jun. 1, 2015. DOI: [10.1111/jieec.12316](https://doi.org/10.1111/jieec.12316).
- [52] J. B. Guinée, S. Cucurachi, P. J. Henriksson, and R. Heijungs, “Digesting the alphabet soup of LCA,” *The International Journal of Life Cycle Assessment*, vol. 23, no. 7, pp. 1507–1511, 2018, ISSN: 1614-7502. DOI: [10.1007/s11367-018-1478-0](https://doi.org/10.1007/s11367-018-1478-0). (visited on 04/11/2022).
- [53] “Victoria metrics.” (), (visited on 03/13/2024).

- [54] H. J. Syed, A. Gani, R. W. Ahmad, M. K. Khan, and A. I. A. Ahmed, "Cloud monitoring: A review, taxonomy, and open research issues," *Journal of Network and Computer Applications*, vol. 98, pp. 11–26, 2017, ISSN: 1084-8045. DOI: <https://doi.org/10.1016/j.jnca.2017.08.021>.
- [55] K. Kumar and M. Kurhekar, "Economically efficient virtualization over cloud using docker containers," in *2016 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, 2016, pp. 95–100. DOI: [10.1109/CCEM.2016.025](https://doi.org/10.1109/CCEM.2016.025).
- [56] "Goaccess." (), (visited on 03/13/2024).
- [57] "Grafana loki." (), (visited on 03/13/2024).
- [58] D. Mytton, "Data centre water consumption," *npj Clean Water*, vol. 4, no. 1, pp. 1–6, 2021, Publisher: Nature Publishing Group, ISSN: 2059-7037. DOI: [10.1038/s41545-021-00101-w](https://doi.org/10.1038/s41545-021-00101-w). (visited on 03/14/2024).
- [59] M. Jay, V. Ostapenco, L. Lefèvre, D. Trystram, A.-C. Orgerie, and B. Fichel, "An experimental comparison of software-based power meters: Focus on CPU and GPU," in *CCGrid 2023 - 23rd IEEE/ACM international symposium on cluster, cloud and internet computing*, IEEE, May 1, 2023, p. 1. (visited on 03/27/2023).
- [60] "Datavizta," Boavizta association. (), (visited on 03/13/2024).
- [61] I. Bekavac and D. G. Praničević, "Web analytics tools and web metrics tools: An overview and comparative analysis," *Croatian Operational Research Review*, vol. 6, pp. 373–386, 2015.
- [62] L. R. Viana, M. Cheriet, K.-K. Nguyen, D. Marchenko, and J.-F. Boucher, "Sending fewer emails will not save the planet! an approach to make environmental impacts of ICT tangible for canadian end users," *Sustainable Production and Consumption*, vol. 34, pp. 453–466, 2022, ISSN: 2352-5509. DOI: [10.1016/j.spc.2022.09.025](https://doi.org/10.1016/j.spc.2022.09.025). (visited on 12/22/2023).
- [63] G. Guennebaud, A. Bugeau, and A. Dudouit, "Assessing VoD pressure on network power consumption," in *ICT4S - International Conference on Information and Communications Technology for Sustainability*, Rennes, France, Jun. 2023.
- [64] D. Schien, P. J. S. Shabajee, H. B. Akyol, L. Benson, and A. Katsenou, *Help, i shrunk my savings! assessing the carbon reduction potential for video streaming from short-term coding changes*, 2023.
- [65] N. Ullrich, F. M. Piontek, C. Herrmann, A. Saraev, and T. Viere, "Estimating the resource intensity of the internet: A meta-model to account for cloud-based services in LCA," *Procedia CIRP*, The 29th CIRP Conference on Life Cycle Engineering, April 4 – 6, 2022, Leuven, Belgium. Vol. 105, pp. 80–85, 2022, ISSN: 2212-8271. DOI: [10.1016/j.procir.2022.02.014](https://doi.org/10.1016/j.procir.2022.02.014). (visited on 04/03/2023).
- [66] L. Siffre, G. Breuil, A. Noureddine, and R. Pawlak, "Demeter: An architecture for long-term monitoring of software power consumption," in *17th European Conference on Software Architecture*, Istanbul, Turkey, 2023. (visited on 05/29/2024).
- [67] "Picasoft." ().
- [68] L. van Oers, J. B. Guinée, and 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*, vol. 25, no. 2, pp. 294–308, Feb. 1, 2020, ISSN: 1614-7502. DOI: [10.1007/s11367-019-01683-x](https://doi.org/10.1007/s11367-019-01683-x). (visited on 12/15/2023).
- [69] I. P. on Climate Change (IPCC), "The earth's energy budget, climate feedbacks and climate sensitivity," in *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2023, pp. 923–1054. DOI: [10.1017/9781009157896.009](https://doi.org/10.1017/9781009157896.009).
- [70] R. Frischknecht, F. Wyss, S. Büsler Knöpfel, T. Lützkendorf, and M. Balouktsi, "Cumulative energy demand in LCA: The energy harvested approach," *The International Journal of Life Cycle Assessment*, vol. 20, no. 7, pp. 957–969, 2015, ISSN: 1614-7502. DOI: [10.1007/s11367-015-0897-4](https://doi.org/10.1007/s11367-015-0897-4). (visited on 05/03/2024).
- [71] T. Simon, D. Ekchajzer, A. Berthelot, E. Fourboul, S. Rince, and R. Rouvoy, "BoaviztAPI: a bottom-up model to assess the environmental impacts of cloud services," in *HotCarbon'24. Workshop on Sustainable Computer Systems*, Santa Cruz, United States, Jul. 2024.
- [72] "Negaoctet," LCIE Bureau Veritas, APL Datacenter, DDemain and GreenIT.fr. (2023), (visited on 09/29/2023).
- [73] A. Pasek, H. Vaughan, and N. Starosielski, "The world wide web of carbon: Toward a relational footprinting of information and communications technology's climate impacts," *Big Data & Society*, vol. 10, no. 1,

- p. 20 539 517 231 158 994, 2023, ISSN: 2053-9517. DOI: [10.1177/20539517231158994](https://doi.org/10.1177/20539517231158994). (visited on 03/09/2023).
- [74] Y. Arushanyan, "LCA of ICT solutions: Environmental impacts and challenges of assessment," Ph.D. dissertation, KTH Royal Institute of Technology, 2013. (visited on 03/18/2022).
- [75] Y. Arushanyan, E. Ekener-Petersen, and G. Finnveden, "Lessons learned – review of LCAs for ICT products and services," *Computers in Industry*, vol. 65, no. 2, pp. 211–234, 2014, ISSN: 0166-3615. DOI: [10.1016/j.compind.2013.10.003](https://doi.org/10.1016/j.compind.2013.10.003). (visited on 08/28/2023).
- [76] J. Pohl, V. Frick, M. Finkbeiner, and T. Santarius, "Assessing the environmental performance of ICT-based services: Does user behaviour make all the difference?" *Sustainable Production and Consumption*, vol. 31, pp. 828–838, May 1, 2022, ISSN: 2352-5509. DOI: [10.1016/j.spc.2022.04.003](https://doi.org/10.1016/j.spc.2022.04.003). (visited on 07/26/2022).
- [77] K. Yan, X. Zhang, J. Tan, and X. Fu, "Redefining QoS and customizing the power management policy to satisfy individual mobile users," in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2016, pp. 1–12. DOI: [10.1109/MICRO.2016.7783756](https://doi.org/10.1109/MICRO.2016.7783756).
- [78] L. Mosesso, N. Maudet, E. Nano, T. Thibault, and A. Tabard, "Obsolescence paths: Living with aging devices," in *ICT4S 2023 - International Conference on Information and Communications Technology for Sustainability*, Jun. 5, 2023. DOI: [10.1109/ICT4S58814.2023.00011](https://doi.org/10.1109/ICT4S58814.2023.00011). (visited on 02/28/2024).
- [79] D.-G. for Energy (European Commission) *et al.*, *Assessment of the energy footprint of digital actions and services*. LU: Publications Office of the European Union, 2023, ISBN: 978-92-68-02624-3. DOI: [10.2833/478689](https://doi.org/10.2833/478689). (visited on 06/30/2023).
- [80] M. Amaral *et al.*, "Kepler: A framework to calculate the energy consumption of containerized applications," in *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*, ISSN: 2159-6190, 2023, pp. 69–71. DOI: [10.1109/CLOUD60044.2023.00017](https://doi.org/10.1109/CLOUD60044.2023.00017). (visited on 05/28/2024).
- [81] M. Ficher, F. Berthoud, A.-L. Ligozat, P. Sigonneau, M. Wisslé, and B. Tebbani, "Assessing the carbon footprint of the data transmission on a backbone network," in *24th Conference on Innovation in Clouds, Internet and Networks*, Paris, France, 2021. (visited on 12/01/2021).
- [82] A. Nouredine, M. D. Lodeiro, N. Bru, and R. Chbeir, "The impact of green feedback on users' software usage," *IEEE Transactions on Sustainable Computing*, vol. 8, no. 2, p. 280, 2023. DOI: [10.1109/TSUSC.2022.3222631](https://doi.org/10.1109/TSUSC.2022.3222631). (visited on 05/29/2024).
- [83] M. Madon, G. Da Costa, and J.-M. Pierson, "Replay with feedback: How does the performance of HPC system impact user submission behavior?" *Future Generation Computer Systems*, vol. 155, pp. 66–79, 2024, ISSN: 0167-739X. DOI: [10.1016/j.future.2024.01.024](https://doi.org/10.1016/j.future.2024.01.024). (visited on 05/29/2024).
- [84] S, *Similarweb*, 2023. (visited on 10/29/2023).
- [85] H, *Hypestat*, 2023. (visited on 10/29/2023).
- [86] D. Prantl and M. Prantl, "Website traffic measurement and rankings: Competitive intelligence tools examination," *International Journal of Web Information Systems*, vol. 14, no. 4, pp. 423–437, 2018, Publisher: Emerald Publishing Limited, ISSN: 1744-0084. DOI: [10.1108/IJWIS-01-2018-0001](https://doi.org/10.1108/IJWIS-01-2018-0001). (visited on 05/29/2024).
- [87] B. J. Jansen, S.-g. Jung, and J. Salminen, "Measuring user interactions with websites: A comparison of two industry standard analytics approaches using data of 86 websites," *PLOS ONE*, vol. 17, no. 5, e0268212, 2022, Publisher: Public Library of Science, ISSN: 1932-6203. DOI: [10.1371/journal.pone.0268212](https://doi.org/10.1371/journal.pone.0268212). (visited on 05/29/2024).
- [88] M. Madon, G. Da Costa, and J.-M. Pierson, "Characterization of different user behaviors for demand response in data centers," in *EuroPar 2022: Parallel Processing*, J. Cano and P. Trinder, Eds., Cham: Springer International Publishing, 2022, pp. 53–68, ISBN: 978-3-031-12597-3. DOI: [10.1007/978-3-031-12597-3_4](https://doi.org/10.1007/978-3-031-12597-3_4).
- [89] "GitHub - boavizta/e-footprint: A toolkit for exhaustively modeling the environmental impact of digital services." (), (visited on 05/29/2024).
- [90] M. Séguéla, R. Mokadem, and J.-M. Pierson, "Dynamic energy and expenditure aware data replication strategy," in *IEEE International Conference on Cloud Computing Technical Program (CLOUD 2022)*, Barcelona, Spain: IEEE, 2022. (visited on 05/28/2024).
- [91] R. D and G. A, "Survey on data replication in cloud systems," *Web Intelligence*, vol. 22, no. 1, pp. 83–109, 2024, Publisher: IOS Press, ISSN: 2405-6456. DOI: [10.3233/WEB-230087](https://doi.org/10.3233/WEB-230087). (visited on 05/29/2024).