Vrije Universiteit Amsterdam

Bachelor Thesis

Labels, Cards, and Simulation-Based Analysis for Energy Efficiency and Sustainability in Data Centers

Author: Philipp I. Sommerhalter (2640200)

1st supervisor: Prof. dr. ir. Alexandru Iosup
2nd reader: Dr. ir. Animesh Trivedi

A thesis submitted in fulfillment of the requirements for the VU Bachelor of Science degree in Computer Science

September 16, 2022

Abstract

Our planet is rapidly warming, causing harsher extreme weather events, more intense drought, food scarcity, and countless other problems that threaten life on Earth. The need to combat climate change has required all industries to examine and improve their contribution to global warming. To that end, the ICT sector has focused on energy efficiency and sustainability. Data center's (DC), crucial facilities providing global digital services, show the potential for greener operations as they account for 1% to 2% of global energy consumption. A lack of proper energy efficiency metrics has unfortunately stunted the field's ability to use this potential. PUE, a long-used metric for data center efficiency, has been criticized as inadequate. In this thesis, we attempt to improve our analytical view of DC energy efficiency through better metrics, and use these metrics to provide insights into the energy usage of general data center's running business critical workloads in a European context. We do this by performing a short survey of candidate metrics and building an instrument to produce and report such metrics. Our key conceptual contribution is the design of an energy card, of increasing detail as the viewer explores it, used in combination with the OpenDC simulator, a detailed simulator for ICT operations. We also provide unique insights into DC sustainability through our grid analysis tool. Our solution is available at `https://github.com/philippsommer27/opendc-eesr`. The data and results of the experiment can be found at `https://github.com/philippsommer27/experiments-bsc-thesis-2022`

Contents

1

Introduction

Data centers (DC) play a crucial role in an ever increasingly digital world [1], with current trends indicating that they are expected to continue to grow in importance and abundance [2]. These facilities used by academia, governments, businesses and private consumers around the world process and store the 2.5 quintillion bytes that our society generates every day [3]. Data centers are most commonly buildings housing an interconnected network of servers, storage systems, and networking equipment that provides uninterrupted digital services to clients around the world [4]. They feature a vital cooling infrastructure and are supported by key staff that ensure smooth operations of the center. Data centers can range in size and complexity from a cabinet in an office building to *hyperscale* centers operated by the likes of Google and Amazon Web Services, boasting the highest efficiency and computational power [5] [6].

Understanding and improving our information and communications technology (ICT) infrastructure should be a top priority for all stakeholders to ensure sustainable growth for the sector. One way through which experts can analyze, improve, and experiment in cloud infrastructure is with the aid of simulators. Simulation allows for the modelling of the complexity of data centers without the cost and/or risk of deployment on existing systems [7]. Despite its challenges including threats to validity and initial intellectual investment, simulators provide many advantages over other forms of cloud infrastructure exploration such as mathematical analysis and real world experimentation [8]. Data center simulators like OpenDC allow us to relatively quickly explore *'what if'* scenarios with great flexibility and accuracy [7] [9].

Energy efficiency is one of the many aspects of cloud infrastructure requiring urgent investigation and innovation [10]. Energy efficiency can be defined as the ratio between the useful (desired) output of a process to the energy input of that same process, with the optimization goal being maximizing productivity whilst minimizing power needs. [11]. In the context of a data center, this means providing the same (or better) service to the customer while consuming less energy. Data center operators typically prioritize availability and speed to uphold Service Level Agreements (SLA) at the cost of energy use and efficiency. This and other current practices leave much to be desired in terms of energy efficiency. Through data center simulation, we can explore ways to make our ICT infrastructure more energy aware and efficient.

The **European Grid**, is one of the largest synchronous electrical grids in the world. In recent years renewable energy has occupied a larger part of the grid's capacity, marking the continents transition towards carbon neutrality by 2050 [12]. Unfortunately, the volatility of renewable energy production causes strain on the network and leads to a mismatch between demand and supply. Data center have long been poised as proactive grid participants through various interactions [13, 14, 15], with one study in particular [16] claiming that data center worldwide could reduce energy costs by 50% acting as a grid stabilizer. Another way through which data centers can increase their grid 'intelligence' is by using renewable energy production data to modulate utilization. In summary, the smart grid provides many opportunities for sustainable ICT.

1.1 Context

By the turn of the decade, the need for greater sustainability in all industries is more evident than ever. For all stakeholders in cloud computing, there is a shared societal, economic, and intellectual interest in the topic covered in this work.

Experts have long warned of the climate risk associated with rising green house gases in the world's atmosphere [17]. Studies show that there is a 90% to 100% consensus among climate scientists that the increase in Earth's temperature is due to human activity [18]. Many governments and international coalitions have mandated actions to curb the rising average global temperature through directive's such as the European Union's Green Deal [19]. In 2020, it established a goal of a 55% reduction in greenhouse gas emissions by 2030 for all its member states [20]. With the increasing threat of global warming and its disproportionate effect on societies around the world, the major sectors of the global economy face growing pressure to make their activities sustainable and carbon neutral, or even negative [21] [22]. Of these sectors, the Information Technology (IT) and Digital Services industry is growing at an unprecedented rate to meet the growing needs of the world [10]. Lei et al. critically remark that "As demand for data center services rises in the future, scrutiny regarding the climate change impacts of data centers will likely continue" [23]. In fact, tenants are demanding sustainability from their chosen data centers [24]. This presents a financial push for DC designers and operators. The industry has already been proactive in curbing its environmental impact, with initiatives such as the Climate Neutral Data Center Pact', aiming for sector wide climate neutrality by 2030 [25]. To date there few studies have explored the global contribution to greenhouse gases by the ICT industry due to the complexity involved in producing an accurate *life cycle estimate* (LCA) [10].

Data centers consume a non-negligible amount of energy to run. Statistics Netherlands (Centraal Bureau voor de Statistiek) reports that data centers represented 2.8% of the country's electricity consumption in 2020 [26]. Worldwide, they consume upward of 1.5% of the global supply [27]. Google, one of the leading actors in energy efficiency and renewable energy use, reports an energy consumption of 18,571,659 MWh across their entire operation for the fiscal year 2021 [5]. Despite a significant portion of this energy being renewable, the facilities produced 11,371,205 tons of CO2 equivalent greenhouse gases (GHG) for that same year. A figure that, although trending downward, is still equivalent to 2,154,692 gasoline-powered passenger vehicles driven for one year [28]. Many other cloud providers that do not benefit from the advantage of being hyperscale perform worse in terms of sustainability [2]. A recent study [1] aiming to provide a more up to date estimate on data center energy use in the coming decade claims that improvements in energy efficiency, policies, and technologies can absorb the energy consumption of doubling data center compute capacity worldwide [1]. This provides a more optimistic outlook compared to previous assessments that expected a three-fold increase in energy use [29], however, it assumes further progress in energy efficiency. A further analysis conducted by Koot et al. concludes through a simulation model that data center electricity needs would amount to 321 TWh by 2030, up from 286 TWh in 2016 [30]. The varying growth estimates also raise concern about local grid's infrastructure to support the power needs. In 2018, the Dutch Data Center Association warned that the Dutch power grid would need expansion in capacity to meet ICT growth, especially around Amsterdam [31]. Altogether, estimating the energy use of the next decade is an extremely difficult task due to the erratic nature of progress in this field. However, the modest growth of less than 10% between 2010 and 2018 can be attributed in part to far-reaching efforts in the research community; an effort which must continue [32].

The incentives for increased energy efficiency and sustainability are not only societal and environmental; they are also largely economic. Recent geopolitical events, Covid-19 and other factors have caused a drastic surge in energy prices across Europe in the past few years [33]. The price of power in Germany rose by 364% between October 2020 and January 2022, with a peak of 920% December 21st, 2021 [34]. A 2008 study claims that the amortized cost of electricity for a data center can account for up to 15% of the total expenditures of the data center [35]. It is therefore assumed that, provided suitable methods for maintaining or reducing energy costs, data center operators would be incentivized to use them.

The landscape of data center energy efficiency metrics is one which has remained mostly stagnant for over decade, especially at the industry level where it's innovation is most essential. Power usage effectiveness (PUE), a metric proposed by Green Grid in 2006 has remained the de facto metric of energy efficiency in data centers since [36]. It is additionally misrepresented as a holistic metric against the advice of the Green Grid, due to the absence of other metrics [37]. There is a dire need for the multitudes of proposed metrics to be to used in practice, in order to assess their efficacy.

1.2 Problem Statement

Although Dutch data centers boast 86% use of green energy as of 2020 [38], the industry still lacks critical awareness of energy efficiency and sustainability [1]. We decompose this adversity into 3 core problems.

First, **current energy efficiency metrics employed in the data center industry are subpar and incapable of comprehensively capturing the energy efficiency**. For the data centers that do produce better data, it is (often) critically inaccessible to the community [39]. As mentioned in Section 1.1, most data centers solely report PUE as an energy efficiency metric (a more in-depth discussion on PUE's shortcomings can be found in Section 2.4.2). This is due in part to the lack of initiative of data center operators and policy makers to propose new metrics as reporting standards. This causes a twofold problem: data center operators and designers do not have the insight required to implement solutions geared towards ameliorating the energy consumption and efficiency of our ICT infrastructure; policy makers cannot accurately gauge the energy efficiency of data centers and set reasonable targets to increase sustainability. Although there are some instances of sustainability targeting solutions, such as Google's work on temporal load-shifting [40], these are few and far between and not the norm in the industry. Masanet et al. [1] largely attribute this problem to a lack of; [policy making] regarding open data for research and transparency [32, 41], enforced targets on energy efficiency and sustainability (beyond PUE), and research funding.

Second, **there is a lack of an established foundation for communicating data center energy efficiency and sustainability**. A communication barrier between experts and policy makers has historically been a common bottleneck to progress in many fields of science. This hindrance is also present in the context of this work. The use of newer metrics and the need for better efficiency mechanisms cannot be encouraged and enforced by the public (and their representatives) if there is an intellectual barrier. Reports on efficiency and sustainability must be digestible and informative for those who lack technical expertise, but are still affected.

Third, there is an observed **lack of grid awareness by data centers in regards to sustainability**. Many data centers typically claim to use more green energy by percentage than the host country produces. For instance, Dutch data centers report 86% green energy use for 2020 [38], while the Netherlands produced 11.1% green energy in the same year [42]. This is achieved through the purchase of Certificates of Origin, Carbon Certificates, or private energy contracts with renewable energy providers (§2.3). Despite being a step in the right direction, it does not leverage the full potential of data centers to actively modulate their operations based on carbon emissions. Without awareness of the source of the electricity consumed, certificates of origin simply shift carbon emissions to the next industry that did not negotiate green energy deals. The total sum of environmental and social damage remains the same. With grid awareness, data centers can employ various techniques that allow them to be active participants in sustainable energy.

We attempt to address the above problems in part or in full where possible. Due to the wide scope of this topic, we aspire to lay some of the groundwork necessary for the community to answer more questions about the sustainability of our ICT infrastructure.

1.3 Research Questions

In this thesis, we attempt to produce a solution towards a more energy efficient and sustainable global data center infrastructures. To tackle the aforementioned problems we identify three research questions.

1.3.1 Research Question 1

RQ1 How to report the energy efficiency and other sustainability aspects related to running business-critical workloads in general data centers?

Our first research question directs our attention to the problem of improper metrics being used in the industry and a lack of an established foundation for communicating DC energy efficiency and sustainability.

RQ1.1 Which selection of metrics are best suited to broadly and deeply report the energy efficiency and sustainability of a data center?

Striving for innovation in the sustainability domain requires a thorough understanding of the current state. Therefore, the first step to making a data center more energy efficient is to assess its current operations using various metrics. To achieve this, we must study existing metrics related to energy consumption and sustainability, selecting candidates that provide energy-focused insight. This will be achieved by conducting a time-constrained literature survey and selecting those that are most relevant.

RQ1.2 How to effectively communicate the energy efficiency and sustainability of a data center to various stakeholders?

Once the appropriate metrics have been identified, we must establish a suitable medium through which they can be communicated. Seeing as the issue of ICT sustainability is one pertaining to wide range of stakeholders, from students, to technically skilled experts, to policy makers this solution is non-trivial. We must explore ways to report on the sustainability of a data center in an intelligible yet flexible manner.

1.3.2 Research Question 2

RQ2 How to model the sustainability of a general data center running business-critical workloads?

With the understanding provided by **RQ1** we set out to construct a model capable of producing the metrics necessary for sustainability reporting. Modeling this issue provides two clear advantages. First, with current industry practices producing a lack of data and transparency (as discussed in section 1.2) [32], a model provides the greatest approximation opportunity. Second, models allow the possibility of exploring hypothetical scenarios, an essential property in this context, as we aim to demonstrate the potential for increased sustainability through various approaches. Traditional experimentation in the real world, although valuable, involves a capital and time cost too high for such preliminary exploration. Instead, we propose a solution using the OpenDC simulator and building an instrument which infers sustainability metrics from its simulation run.

1.3.3 Research Question 3

RQ3 How can we use our model and solution to draw conclusion on the sustainability of general data centers running business critical workloads?

Having designed and implemented the instrument resulting from **RQ2**, we can analyze the operations of a typical data center under regular operations and modified scenarios. This allows us to quantify the potential for improvement in the industry in terms of sustainability.

1.4 Research Methodology

In addressing the problem statement and its subsequent research questions, we employ various research methodologies.

To answer **RQ1.1**, we employ quantitative research in the form of a shortened literature survey of the field of energy efficiency metrics for data centers [43]. Due to the scope of this work and time constraints a fully comprehensive survey of the matter is infeasible. Instead, we time box our efforts and approach the task by reviewing an existing relevant survey. From there, we synthesize the research terms to be used in a literature search engine[1], and follow the citations to obtain a more in-depth knowledge about each candidate. This methodology is essential in not only finding the most suitable metrics for our case but also in identifying any gaps in the field of research that threaten our goals.

To address **RQ1.2**, we follow a process inspired by the AtLarge design process [44]. This involves an iterative design cycle with multiple rounds of feedback from the relevant community, experts, and stakeholders. Before starting, we carefully analyze the requirements and potential stakeholders.

Answering **RQ2** involves designing and prototyping a system that allows one to investigate the sustainability of the operation of a data center given an energy trace. We use the existing capabilities of OpenDC to produce an energy use scenario, which we then base our analysis on. As such, we approach this solution with compatibility with OpenDC as a priority, but with the aim of not restricting its usage to the simulator. This allows the tool to be used by a wider audience in the community. Calculating the various sustainability metrics involves acquiring additional data, including electricity grid production information and CO2 emission data. For the latter, a small research phase is carried out.

[1] https://scholar.google.com/, https://www.semanticscholar.org/

To address **RQ3**, we designed a series of experiments that analyze the sustainability of cloud computing under various circumstances and assumptions. Through the flexibility implemented in our sustainability analysis tool, we can vary the time frame, geographical location, and assumptions on the proportionality of green energy of our experiments. Our experiments vary along these three axes to detect any sustainability impacts.

Throughout all stages of our work, we impose strong emphasis on open and reproducible science. All implementation follows standard practices and is accompanied by a comprehensive documentation that not only acts as a guide, but explains the theoretical foundation behind the subject at hand. We open-source this document as well in hopes of encouraging the community to expand upon it as research on data center energy efficiency and sustainability grows.

1.5 Thesis Contributions

In answering our research questions, this thesis has produced several contributions.

1. **Conceptual**
   a) We document the finding of a shortened literature survey of the current landscape of energy efficiency and sustainability metrics in data centers.[1]
   b) We design a card, comprised of a selection of energy and sustainability related metrics, offering different levels of information and context depending on the expertise of the user.
   c) We design and research a model to determine the sustainability of a data center's energy usage based on its electrical grid consumption.
2. **Technical**[2]
   a) We implement an instrument **(EESR: Reporting)** to dynamically produce energy efficiency and/or sustainability reports for data centers.[3]
   b) We implement an instrument **(EESR: Grid Analysis)** for configurable analysis of data center sustainability with respect to the synchronous grid of Continental Europe.[3]

[1] https://philipp-sommerhalter.gitbook.io/eesr/

[2] Although we do not contribute directly to the codebase of OpenDC, we consider our work a contribution to the larger OpenDC ecosystem.

[3] https://github.com/philippsommer27/opendc-eesr

3. **Experimental**

    a) We design, implement, and report various experiments that analyze energy sustainability in general data centers running business-critical workloads. We ensure open science by making all data, code, and results accessible. We additionally provide reproducibility guidance for validation and/or inspiration.

4. **Societal**

    a) Presented works to potential adopters and shareholders including SURF[1]. This resulted in an emerging collaboration between SURF and the European Grid Infrastructure[2] (EGI) to leverage our instrument and concepts for use on the ICT infrastructures of Europe. This could lead to high and lasting societal impact.

1.6 FAIR Data and Software

Throughout this work we make an effort to adhere to the FAIR principles of scientific data and software. FAIR, which stands for Findability, Accessibility, Interoperability, and Reuse of digital assets, are a set of guidelines for data sets to ensure open science and reproducibility. This thesis concerns itself with primarily two data sets related to our experiments; the input DC energy trace into EESR and the resulting output data files. We use zenodo, a respected open science platform, to publish our data and relevant code. Zenodo adheres strictly to each fair principle[3]. Additional metadata and instructions are inlcuded in the README of each relevant repository and the EESR documentation.

DOI: 10.5281/zenodo.7084663

Persistent URL: `https://doi.org/10.5281/zenodo.7084663`

License: Creative Commons Attribution 4.0 International[4]

1.7 Plagiarism Declaration

I confirm that this thesis work is my own work, is not copied from any other source (person, Internet, or machine), and has not been submitted elsewhere for assessment.

[1] https://www.surf.nl/en/about-surf
[2] https://www.egi.eu/
[3] https://about.zenodo.org/principles/
[4] https://creativecommons.org/licenses/by/4.0/legalcode

1.8 Thesis Structure

We briefly describe the structure of the rest of this paper. In Chapter 2 the background information pertinent to our work is covered. Chapter 3 is the culmination of our research on energy efficiency and sustainability metrics for data centers. In Chapter 4 the core design of our EESR instrument is explained. The implementation details of this design are described in Chapter 5. We then evaluate our implementation and the various experiments in Chapter 6. Finally, our contribution is summarized and future work is proposed in the concluding chapter (§7).

2

Background

In this chapter, we present a primer on the core topics pertinent to our work. We start by detailing the basics of energy flow and consumption in a general data center (§2.1). We continue by describing the essentials of modelling data centers through discrete event simulation (§2.2). Followed by an explanation of the interactions between data centers and the electrical grid (§2.3), before covering the fundamentals and the use of metrics in our data center context (§2.4). Crucially, we understand the current pitfalls with data center metrics, providing further justification for the efforts (§2.4.2).

2.1 Overview of Energy in General Data Centers

Modern data centers are a complex composite of high-performance servers, reliable networking, and storage solutions with each component defining unique power-draw characteristics. In this section we explicate the main actors in energy source, distribution, and consumption.

The significant energy draw of a data center is in part due to its expected service provision. The Uptime Institute defines an internationally recognized classification standard [45]. Characteristics of the four tiers are presented in table 2.1. Most data centers in Europe belong to Tier III [46], which means an availability of 99.982%. Essential to ensuring this is a robust power system and plentiful compute resources. Its redundancy description, N + 1, means that there is at least one backup component in place. Hyperscale centers, classified as Tier IV, have 2N redundancy exhibiting a fully mirrored support system is in place in case the primary fails. In terms of computing availability, data centers usually over provision resources for peak demand and to handle unexpected bursts of traffic [47]. They do so to uphold Service Level Agreements (SLA) with clients. SLAs guarantee to

Characteristic	**Tier I**	**Tier II**	**Tier III**	**Tier IV**
Site Availability	99.671%	99.749%	99.982%	99.995%
Annual Downtime (hours)	28.8	22.0	1.6	0.4
Redundancy	None	Partial N+1	Full N + 1	2N or 2N + 1

Table 2.1: Relevant characteristics of the Uptime Institute data center classification [49].

the customer various performance targets, such as availability, response time, *greenness*, security, etc., with violations resulting in payouts from the data center. However, over provisioning of resources is energy inefficient. A study of various data centers in 2011 found that the mean utilization of CPUs was 17.76% [48]. A more efficient data center, operated by Google for example, can achieve better numbers of 60% [47].

A DC's constituents can be divided into two main categories: the components responsible for data processing and storage (IT Equipment); and the surrounding support infrastructure (Infrastructure (Non-IT Equipment) and Support Systems). The energy profiles of these categories are discussed in greater detail in the next two subsections.

The summary in the following sections and the basis for figure 2.1 are derived from works [50] [51], [52], and [53].

2.1.1 IT Equipment

The servers, storage solution, and networking of a data center allow it to accomplish its main purpose. Of the IT equipment in a data center, the CPU consumes the most power [51], yet remains the most underutilized resource [47]. This is followed by memory, and then networking. Software wise, data centers usually provision an entire machine through which tenants can host their applications on, or virtualize the machines (with a Hypervisor) and offer the requested amount of compute resources. The latter permits Infrastructure as a Service (IaaS) to be a common business strategy of data centers. Virtualization consolidates the compute capacity into fewer machines, effectively leading to higher IT equipment energy efficiency as either total tenant capacity increases or server count decreases.

2.1.2 Infrastructure (Non-IT Equipment) and Support Systems

As previously mentioned (§2.1), the philosophy of a data center power support system is redundancy. In Marcus J. Ranum's words [54]:

> One person's "paranoia" is another person's "engineering redundancy."

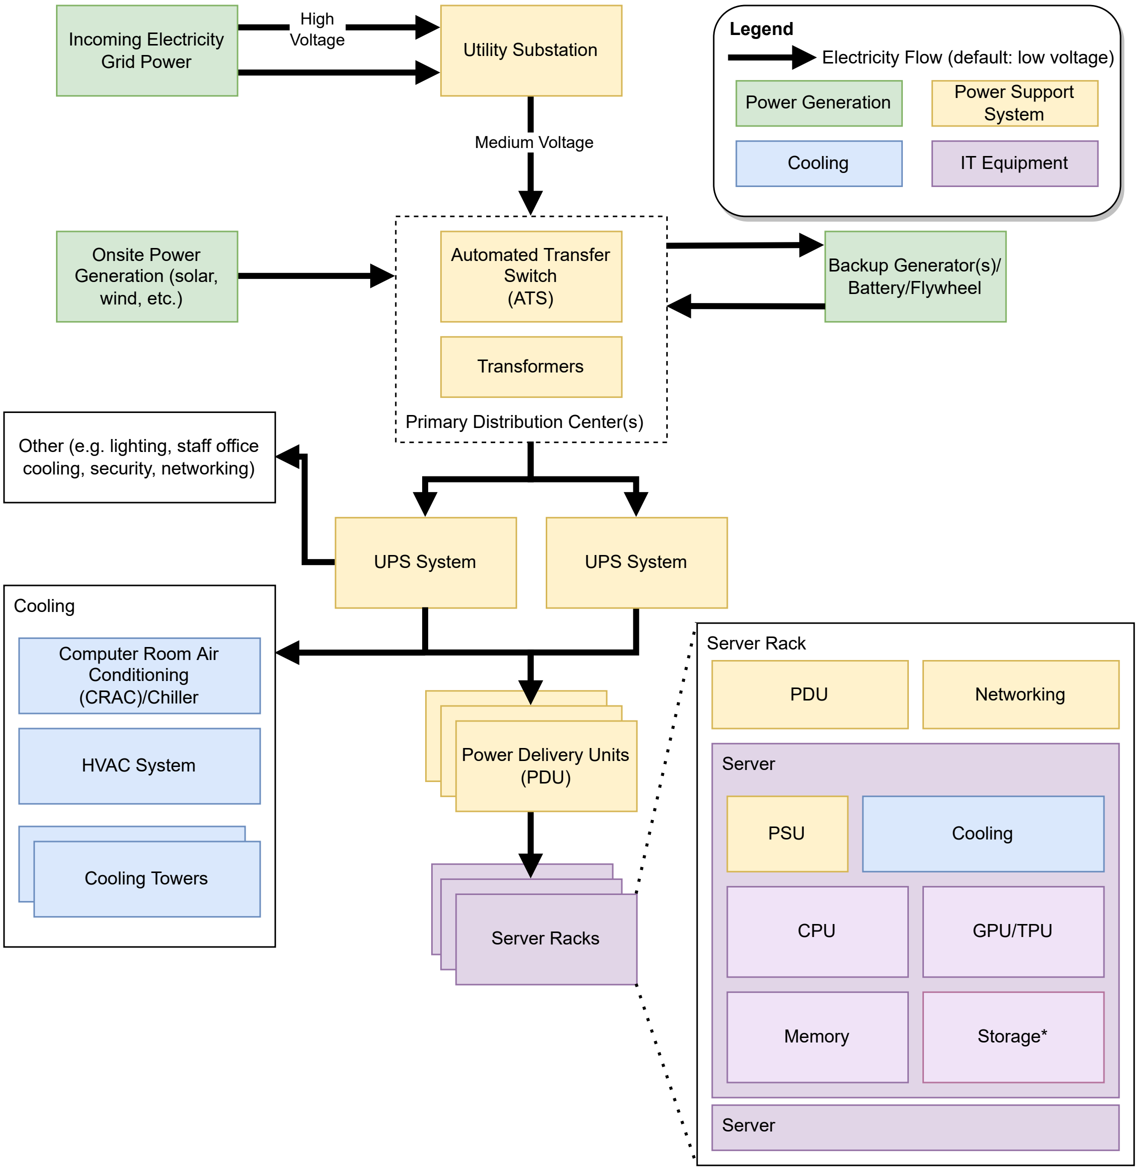

Figure 2.1: Overview of energy flow in a typical data center.

It follows that all of the items described below may appear multiple times in a DC for redundancy and/or capacity requirements. These ancillary systems are tasked with providing the compute components with reliable and sufficient power. Grid power typically enters the site through high voltage AC transmission lines at the utility substation. The substation transforms the electricity into medium voltage, to be distributed to the primary distribution centers. The automated transfer switch (ATS) within a primary distribution center is responsible for selecting the energy source to be used. The possible sources are; the grid (majority), on-site generation (preferred), or the backup power source (diesel generator, battery, or flywheel) in the case of an outage. Although most data centers receive power through the electrical grid, some may have onsite or dedicated alternate energy production facilities (e.g., solar panels, wind farm, geothermal, etc.). [55]. Showcased in Section 2.1.2, a data center will typically feature one or more backup diesel generators in case of a power failure/outage. Low voltage power lines from the distribution center feed into the uninterruptible power supply (UPS) systems. These vital apparatuses contain a battery or flywheel to maintain power flow until backup power can kick in. Note that depending on size and type, the cooling system may be powered through the UPS or the primary distribution center. The same is true for other non-it infrastructure. Nonetheless, data center's floor power distribution unit (PDU) receives its electricity from the UPS, finally powering the server racks. The server racks contain another PDU for its servers, with each server's final power component being the power supply unit (PSU). Here, power is converted into direct current (DC) to be used by the compute components.

At each of these stages, slight power transformation inefficiencies can add up to a significant energy sink. For instance, older UPS systems could lose up to 15% of its input energy due to AC-DC-AC conversion required for storing power in battery [51]. Newer technologies fare better achieving losses of only a 1 to 4%. In order to better monitor energy losses, data center designers and component manufacturers should include sensors in every appliance to measure energy intake, outtake, and loss. This grants operators more precise insight into areas of inefficiencies.

Keeping the server hardware at their optimal operating temperature is a key aspect of data center design. One method is closed-loop cooling, where moving air across the data center has traditionally been used to remove heat. Using a raised floor, cool air can reach the computer room air conditioning (CRAC) unit from below the servers and exit through the top. CRAC units are equipped with a heat exchanger and a pipe to cooling towers which remove excess heat. A more energy efficient alternative to closed-loop cooling is free cooling. This involves forcing outside air in with an HVAC system. The servers themselves are typically cooled with built in fans, or sometimes water cooling. It is clear that data center cooling is no easy task with many considerations to be made; however, its energy draw has historically been significant. The cooling infrastructure typically accounts for about 18% of energy losses.

2.2 Modelling with Discrete Event Simulation

Data center simulation has long been present with multiple simulators available to choose from [56, 57, 58, 59, 60]. However, to date, few efforts have involved assessing sustainability through a data center simulator. OpenDC was built to facilitate investigation of past, present and future data center technologies [9]. It can for instance model emerging serverless and machine learning data center paradigms, providing researchers the tools to guide the next generation of cloud computing. The simulator models the major components found in a data center. From the physical topology, an internal workload and resource scheduling representation consistent with real data centers, to the power systems and energy flows. Most importantly, OpenDC employs a stringent software engineering process which includes open source development, continuous integration (CI) adherence, and careful integration of extensions into the main codebase using the latest and most effective technologies (Kotlin). All of which are core qualities necessary for open and reproducible science.

2.3 Data Centers and The Electrical Grid

The evolution from fragmented national grids into the Synchronous grid of Continental Europe resulted in a much more capable electrical network. As a result, higher energy savings were achieved and more transparent data communication became standard. With it came the formation of the European Network of Transmission System Operators for Electricity (ENTSO-E). This association operates the electricity market and regulates energy cooperation at a continental level. Involving 34 countries, ENTSO-E sees 300 TWh on its network annually [61]. Cross-border grid cooperation allows clients in countries like the Netherlands , low in renewable energy production, to consume green energy from Norway. The continued transformation into a smart grid promises even greater sustainability and stability. Several studies have highlighted the potential for data centers to be direct participants in the market due to their unique energy attributes, allowing them, for example, to assist in load stabilization [62].

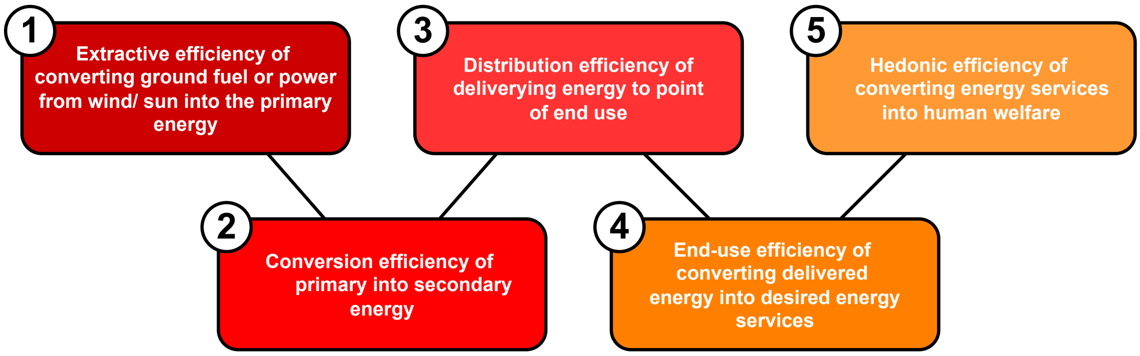

Figure 2.2: Efficiency along Energy Conversion Chains according to Lovins [63].

Many data centers are able to achieve upwards of 85% renewable energy consumption through special long-term contracts with electricity providers called *purchase power agreements* (PPA) [38] [23]. These contracts established far in advance guarantee a certain capacity of electricity from the producer to the consumer.

2.4 Energy Efficiency and Sustainability

In their 2004 paper, Lovins explains the intricacies of energy efficiency as a concept and provides a taxonomic overview [63]. Amongst others, 'incorrect metrics' is listed as a notion lacking "rigorous treatment". According to Lovins, the true efficiency of energy use must viewed as a product chain of efficiency's starting from the source. Depicted in figure 2.2, this chain has as pen-ultimate component the case describing data centers: "End-use efficiency of converting delivered energy into desired energy services". At this stage, the 'service' a data center provides merits precise definition as part of the equation of energy efficiency. Seeing as we are concerned primarily with general data centers, the service would be reliable provisioning of computing power and/or workload completion for clients. In their well-cited[1] article on energy efficiency, Shove proclaims the generalization of the *service* in energy efficiency to be a root for ineffective innovation [64]. That is, when defining, reporting and improving the energy efficiency of *something*, the difference in service must be carefully and fully considered. As an exaggerated example to illustrate this idea: a new data center may run 50% more efficiently due to its novel positioning in the antarctic, but provide no equivalent service as an alternative to its predecessor running a stock market application with low latency requirements.

[1] 200+ citations according to Google Scholars as of 01/08/2022

Not to be confused with energy efficiency, energy conservation refers to the consumption of less energy by doing less. Through the lens of climate change, energy conservation opportunities should also be examined in all industries to see where unnecessary spending occurs. Furthermore, there is a literature gap regarding energy conservation in data centers. In this paper, we focus on energy efficiency and sustainability.

Thus far, we have elaborated on energy efficiency. In this work, we make clear the distinction between it and *energy sustainability*. Energy efficiency is measured purely in terms of the work done. However, it is used to address societal, and environmental goals. To bridge the gap between it and the latter, we delineate energy sustainability: *measuring the environmental sustainability and impact of an energy use.*At this stage in the climate emergency, analyzing human activity through the lens of energy sustainability, as well as efficiency, is imperative to meeting our collective environmental goals.

2.4.1 Metrics

Metrics are the medium of system observation and analysis. Importantly, metrics are used as decision making tools, and when our decisions regarding energy use affects life globally, choosing the right metrics becomes crucial. To be able to innovate, one must first measure performance. Decision making at a data center level and national/international policy setting level can only be based on the information exposed. There are many pitfalls in choosing and using metrics to achieve desired long term goals [65]. In this paper, we only scratch the surface of what may seem like a trivial exercise.

2.4.2 PUE and Other Things Wrong with Data Center Metrics

To understand the need for innovation, we must first dissect the current issues of data center energy efficiency metrics, specifically Power Usage Effectiveness. In this section we summarize the views of various literary works in the community regarding PUE.

$$PUE = \frac{\textit{Total Power Consumption of the Data Center}}{\textit{Energy Consumption of IT equipment}} \tag{2.1}$$

Equation 2.1 succinctly describes how to calculate the PUE of a data center. The metric has a minimum and target value of 1, and an infinite maximum. A PUE score of 1 would mean that 100% of the energy received by a data center is used by IT equipment. The Green Grid specifies various levels of measurement for PUE, ranging from basic to advanced. A higher level provides a more accurate and valuable number. Presented in table 2.2, level 3 allows for a narrow definition of *Energy Consumption of IT equipment*, excluding energy

	1: Basic	2: Intermediate	3: Advanced
IT Equipment Consumption	UPS Outputs	PDU Outputs	IT Equipment Input
Total Data Center Consumption	Utility Inputs	Utility Inputs	Utility Inputs
Measurement Interval	Monthly/Weekly	Daily/Hourly	Continuous

Table 2.2: Green Grid's measurement levels for PUE [37].

inefficiencies of the power delivery system. Furthermore, a measurement interval of at least 15 minutes is encouraged for more actionable insight, as operators can better understand when in the day and during what operations PUE suffers.

The following points particularly become problems when PUE is used to compare data centers.

One of the main issues with PUE is measurement variability, met with a lack of standardization [66]. As table 2.2 highlights, PUE can be measured in various ways, which can impact the final score. Some data centers may achieve a higher PUE by measuring IT equipment use at the rack level as opposed to the individual compute component level [41]. Additionally, a large data center facility can measure total facility usage at the utility box, which would include transformer energy loss, whilst a server closet in an office space would forgo such losses by measuring at a room level.

Often being the only reported metric, PUE also suffers from being used as a 'greenness' metric [41]. Despite the fact that energy usage is usually correlated with carbon emissions, PUE fails to convey a data center that may run on green energy or reuse its energy. Fortunately, nowadays data centers frequently report Green Energy Coefficient (explained in §3.3.1).

Another fairly obvious problem when using PUE as a comparison metric is that it does not account for the varying climates where data centers operate [67]. The cooling requirements of a data center in the tropics will be much greater than those of a data center operating closer to the poles of the earth. There are current plans to build one of the largest data centers in the arctic circle [68]. This data center would be able to achieve a very good PUE score due to the low energy requirements for cooling infrastructure, however it does not mean that the rest of its design and operations are more efficient than another data center.

A seemingly counterproductive property of PUE is that it can sometimes increase when energy efficiency measures are implemented in a data center [36, 66, 69]. Virtualization, a technique described in the next section (§2.4.3), reduces the energy use of IT equipment. If the cooling infrastructure consumption does not scale proportionally to IT equipment use then the PUE value after virtualization will be higher, despite the data center wasting less energy. Another example is an IT equipment refresh of more energy efficient components. Once again, this would in theory increase the energy efficiency of the data center but also increase the PUE value.

Ultimately, PUE has been used more as a marketing tool rather than a genuine metric for measuring energy efficiency [36]. If measured precisely, reported correctly, and interpreted properly PUE can serve as a good metric in a suite of DC energy metrics.

2.4.3 Energy Saving Methods

Fortunately, the ICT industry and the research community are and have been hard at work designing better data centers for energy efficiency. From virtualization methods to cooling systems, there is an opportunity for fine-tuning and innovation in all areas of a data center. We highlight only a fragment of some of the practices, for more details, Kaur et al. survey energy efficiency techniques in their 2015 survey for *ACM CSUR 2015* [70].

Over the years the community has produced numerous viable solutions, many of which are already employed in data centers. In a recent report, Lei et al. summarize some of these strategies for addressing energy use. Included is; server virtualization for IT hardware; adopting more energy efficient servers, storage devices, and power support systems; improved airflow management; raising maximum indoor temperatures to 90 $^\circ$ F from a standard 55 $^\circ$ F [23].

It must be noted that the industry's transition towards more hyperscale data centers is a net positive in terms of energy efficiency and consumption [2], with the worlds best performing hyperscale data centers routinely achieve scores of 1.1 PUE [1]. However, given that the majority (39% as of Q4 2020) of hyperscale centers are based in the United States [71] where growth in the coming decade is expected to be concentrated [72], the need for research on energy efficiency in general data centers is still relevant especially for the European market.

Several works have found the use of a DC power distribution system to be a significant energy saving technique [51] with one in particular claiming that by pairing such a system with onsite solar panels, they can achieve around 17% in energy reduction [73]. Although most current data centers apply a more orthodox AC power system, these findings already deployed at scale provide a promising alternative.

3

Metrics

A lack of well-established performance indicators impedes the industry's evolution towards sustainability. Chapter 3 describes the results of our literature survey on data center metrics. These metrics should be measured and reported by data center operators, and understood and employed by policy makers. In section 2.4.2 we dove into the shortcomings of PUE, the leading data center metrics. The metrics presented in this chapter can act as a replacement or extension to current industry norms. We provide a more extensive overview of each metric in our documentation[1].

Using the right metrics is only half of the solution. DC operators must also adhere to the relevant measurement guidelines and preferably with additional reporting describing how the value was achieved. Additionally, it is recommended that metrics are reported as an annual value, unless specified otherwise.

3.1 Surveying Method

Literature surveys can typically produce an entire high level study. Due to the scope and time frame of this work, we conduct only a short study of DC metrics in order to build our instrument. It must therefor be noted that this survey is by no means exhaustive. We employ a systematic approach to finding relevant works and determining suitability of studied metrics (depicted in figure 3.1). ❶ We start by traversing the metrics detailed in the works of Reddy et al., using it as a well cited basis for our survey [74]. ❷ From it, we select suitable metrics for reporting. ❸ We synthesize search engine keywords, leading us to more studies, and ❹ check the citations for relevant literature. ❺ Repeat steps ❷ - ❹ with new studies found. ❻ Finally we report the metrics in this chapter. This

[1] https://philipp-sommerhalter.gitbook.io/eesr/

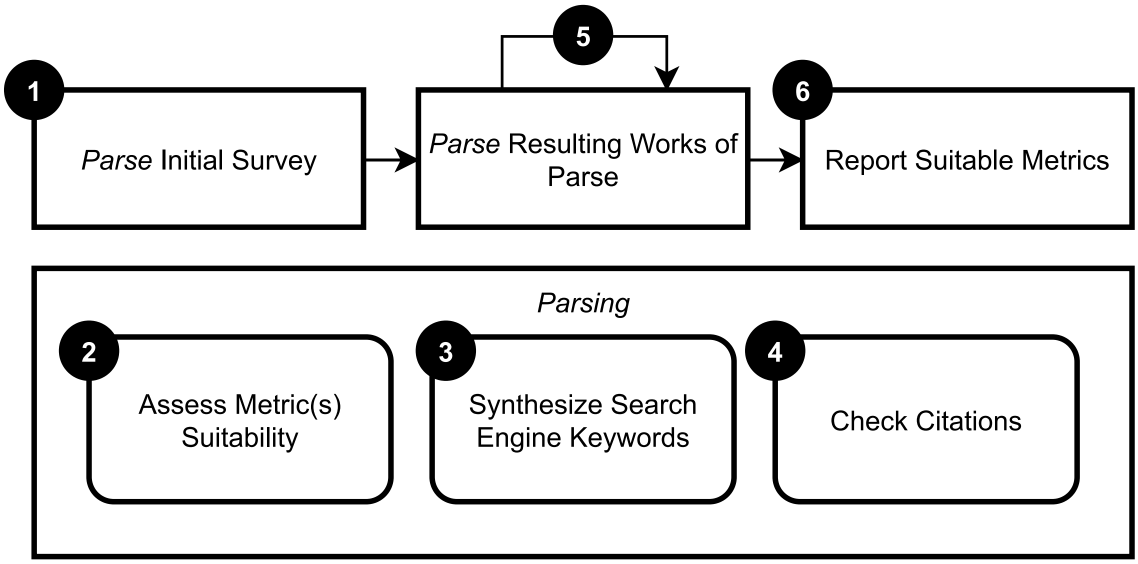

Figure 3.1: Process for surveying suitable metrics.

process has as halting case either: sufficient metrics to implement our instrument or time constraint reached.

3.2 Energy Efficiency Metrics

Having already described PUE (§2.4.2), we report here other energy efficiency metrics for data centers.

3.2.1 TUE

Total-power Usage Effectiveness, proposed by Patterson et al., is an attempt to add a level of depth to measuring data center equipment energy efficiency [69].

To do this, the authors first define *IT Usage Effectiveness* (ITUE). In essence, it is a metric similar to PUE, but for the servers of the facility. It aims to report how much of the energy going into the IT equipment is spent on computing versus cooling, power delivery, etc.

$$ITUE = \frac{Total\ Energy\ into\ the\ IT\ Equipment}{Total\ Energy\ into\ the\ Compute\ Components} \tag{3.1}$$

The total for the numerator should be measured at the power supply level for the computer, whereas the denominator should be measured at the component level (perhaps through software).

The paper then defines TUE as:

$$TUE = ITUE \times PUE \tag{3.2}$$

TUE thus provides a higher precision in the definition of *IT Equipment Energy Use* over PUE. In a sense TUE presents a level 4 measurement of the Green Grid's PUE measurement levels (2.2). This is substantial as servers consume a non-trivial amount of energy that is not directly computing power [ref]. Like many metrics in this chapter, TUE is unfortunately not widely used in practice and remains a conceptual proposition.

3.2.2 DCeP & Proxies

In measuring the efficiency of a data center, the end productivity of each unit of energy is of great interest. The ability to understand whether the current energy draw was used for overhead computation for example, could be extremely valuable in assessing software stack inefficiencies. However, this proves challenging as their is no easy way to quantify 'useful' work in a data center. Nonetheless, in their 2008 white paper, the Green Grid proposes a metric to leverage this concept called *DCeP*.

Data Center energy Productivity (DCeP), is the first in a family of DC productivity (DCP) metrics [75]. With a clear definition of *productivity*, DCeP can be adapted by replacing 'energy' with any optimization goal, such as water usage or land area of facility. The authors of DCeP, emphasize that comparison between data centers is not a priority of the metric. The formula for the metric is:

$$DCeP = \frac{Useful\ Work\ Produced}{Total\ Data\ Center\ Energy\ Consumed\ Producing\ this\ Work} \tag{3.3}$$

The definition of *Useful Work Produced* is then based on how the metric user assigns value to the work of a data center. Although good in concept, this leaves DCeP to be difficult to implement in practice. As such, various proxies have been suggested to serve as productivity indicators [76]. Included in these proxies are; Bits per Kilowatt-hour, Weighted CPU Utilization, Compute Units per Second Trend Curve, Operating System Workload Efficiency, Useful Work Self-Assessment and Reporting. Pros and cons of some of the proxies are featured in [77].

3.3 Energy Sustainability Metrics

In section 2.4 we distinguished energy efficiency from energy sustainability. The following metrics in some way or another quantify the energy sustainability of a data center, with regards to the environment.

3.3.1 GEC

Green Energy Coefficient (GEC) is simply the ratio of green energy consumed by the data center to total energy consumed by the data center (equation 3.4). Although many data center's already loosely report renewable energy percentage, specifying an explicit metric can aid in enforcing rigid measurement guidelines and targets.

$$GEC = \frac{Green\ Energy\ Used\ by\ Data\ Center}{Total\ Data\ Center\ Energy\ Usage} \tag{3.4}$$

A data center relying solely on green energy would yield it the metric's maximum value of 1.0. Despite its simplicity, problems can arise with the metric due to its loose definition and lack of standardization. Although the EU has a taxonomy on what is considered sustainable activity, this does not translate into a clear categorization of energy greenness [78]. Furthermore, a distinction can be made between green energy and renewable energy (see nuclear energy). To increase transparency, our implementation of GEC is calculated with two definitions of green; one for renewable and one for ecologically green (more in section 4.4.3).

Matters become even more convoluted when taking into consideration the purchase and selling of *green certificates*. Various governments and intergovernmental organization recognize these certificates as making an entity's activities 'green'. A data center could therefor report a high GEC score, despite the actual energy used for the facility being 'dirty'. The authors of this paper maintain that the 'greeness' of a data center's electricity consumption is determined by the energy it directly used. More on green certificates can be read in [79, 80]

3.3.2 ERF

We list this metric after consultation with an industry expert on data center efficiency metrics.

Energy Reuse Factor (ERF), captures the sustainability of exporting energy from the data center which would otherwise be rejected to the atmosphere/environment [77]. Most commonly this energy is in the form of heat produced by the cooling infrastructure, whether it be heated water or air. ERF can derived through the following formula:

$$ERF = \frac{Energy\ Reused\ Outside\ Data\ Center}{Total\ Data\ Center\ Energy\ Usage} \tag{3.5}$$

The maximum value of 1.0 represents a 100% reuse of energy where as the minimum, 0.0, means no energy is reused. Difficulties in obtaining the value may arise when considering that exiting energy is usually not in the same form as entering energy.

Various uses for the export energy have been suggested and are readily implemented in data centers around the world [81, 82, 83, 84]. They include connecting to local neighbourhood heating infrastructure or using the excess heat in greenhouses farms. A data center which implements reuse energy not only wastes less energy, but contributes to less emissions that would have resulted from the energy produced which it replaced.

3.3.3 $\mathbf{APC_{ren}}$

Data centers can achieve higher sustainability by adjusting their operations to the tune of clean energy production (e.g. load shifting). DC4Cities proposed in 2014 a metric which captures how much the energy consumption profile of a data center matches the availability of renewable and green energy [85, 86]. Adaptability Power Curve at Renewable Energies ($\mathrm{APC_{ren}}$) is defined through the following two equations:

$$APC_{ren} = 1 - \frac{\sum_{i=1}^{n} |KE_{Reni} - E_{DCi}|}{\sum_{i=1}^{n} E_{Reni}} \tag{3.6}$$

$$K = \frac{\sum_{i=1}^{n} E_{DCi}}{\sum_{i=1}^{n} E_{Reni}} \tag{3.7}$$

Where:

- E_{reni} is the renewable energy production on the grid (or onsite) in kWh,
- E_{DCi} is the total energy consumption of the data center in kWh,
- K serves as an adjustment factor to meaningfully compute E_{reni} against E_{DCi},
- i is the time period, and
- n is the sample size

In essence ($\mathrm{APC_{ren}}$) reports the difference in area between the two curves of DC energy use and renewable energy. As the amount of renewable electricity on the grid is far higher than what a data center consumes, K is needed to equalize their magnitudes. A result of 1.0 (maximum) would mean that the data center matches the production of renewable energy perfectly. The metric has no lower bound, and may have values in the negatives.

3.3.4 CUE

When considering environmental sustainability and climate change, greenhouse gas contribution is a major area of focus. In assessing the sustainability of a data center, we must examine how carbon effective its operations are. Carbon Usage Effectiveness (CUE), another metric by the Green Grid aims to do exactly that [87].

$$CUE = \frac{Total\ CO_2\ Emissions\ Caused\ by\ Data\ Center\ Energy}{IT\ Equipment\ Energy\ Use} \tag{3.8}$$

Defined in equation 3.8, CUE has a unit of carbon dioxide per energy use, with unit prefixes to be determined and properly communicated when reporting. Note that the numerator should included equivalent CO2 emissions from other greenhouse gases. The metric has a target minimum value of 0.0 indicating no carbon emissions associated with DC operations and features no upper bound. CUE alone does not encompass the life cycle assessment of a data center, simply the sustainability of the energy used.

The calculation of *Total CO_2 Emissions Caused by Data Center Energy* is not clearly standardized and can be difficult to do. Each power generation source will incur different carbon emissions and access to this data is often rare or lacking. As such best guest efforts must be made. Calculations should also take into account any emissions from onsite electricity generation.

CUE allows data center operators and relevant stakeholders to see opportunity for workload scheduling based on emissions and procurement of more sustainable energy. Additionally, unlike many metrics mentioned thus far, CUE can be used as a comparison metric between data centers as it is clear that regional differences will affect the value. This can even be useful in determining the construction site of a new data center.

3.4 Honourable Mention

The following metrics are not covered in whole in this paper but are worth investigating for DC energy efficiency and sustainability. The last two items are quasi 'mega-metrics' which attempt to holistically report the energy efficiency and sustainability of a data center in a single number. As mentioned in section 1.1, PUE suffers from being used as a holistic metric. We believe that one metric is not sufficient. Nonetheless, the sub metrics of the two mega-metrics show potential and as such were included here.

- **Water Usage Effectiveness** reports the annual water usage versus energy use of the IT equipment [88]. Seeing as water scarcity is becoming a growing issue, and thatdata centers can employ a large amount of water for their cooling needs, this metric provides a valuable sustainability angle [refs].

- **Corporate Average Data Center Efficiency** combines facility utilization/efficiency and IT utilization/efficiency to highlight areas of potential improvement[89].
- **Datacenter Performance per Energy** combines 4 metrics; ITEU, ITEE, PUE and GEC to provide a single value for the sustainability and efficiency of a data center. The four metrics target four areas of energy improvement. As such stakeholders can determine which to give priority to.

4

Design of Energy Efficiency and Sustainability Reporting (EESR) Instrument

In addressing questions **RQ1.2** and **RQ2** we design the Energy Efficiency and Sustainability Reporting (EESR) instrument.

4.1 Requirements Analysis

In order to ensure that EESR is well designed and suited to the needs of its envisioned users, we first perform a requirements analysis to determine the stakeholders, their use cases, and the functional and non-functional requirements.

4.1.1 Stakeholders

Given the nature of the subject; the direct impact of climate change on the world, all living beings present a relevance to our work. With regards to direct interaction however, we distinguish the following stakeholders to EESR.

S1 **Data center operators** constantly monitor and maintain the facility. They are responsible for ensuring that the center can uphold SLAs, and with tenants becoming increasingly interested in the sustainability of their provider, operators may need various metrics to track efficiency and sustainability performance.

S2 **Data center designers** carry out a significant amount of analysis for the construction (or upgrade) of a data center. Environmental implications and energy efficiency play a growing part in this analysis. With better and more metrics available to them, more sustainable decisions can be made at every stage of the design process, from selecting geographical location to choosing the right scheduling policy.

S3 **Data center customers** and tenants may use sustainability as a decision factor when choosing where to host their services [24]. Through their wallets, these stakeholders indirectly dictate the direction of the industry. Having access to transparent and understandable energy efficiency and sustainability reports for data centers will be essential to them.

S4 **Policy makers** play a fundamental role in the way data centers are designed and operated. In the EU, more and more legislation will be made at a national and continental level, setting strict guidelines for cloud infrastructure energy use. Importantly, the stakeholders making these regulations may lack the technical expertise about the ins and outs data center energy. Furthermore, without adequate data made available to them it is difficult to establish effective and appropriate policies.

S5 **Researchers** give us invaluable knowledge about past present and future of data centers. They help us understand current impacts, projected impacts and ways in which we can improve our design and usage of data centers. Their jobs are made easier and more impactful with better access to data through rich tool sets. Researchers are an indispensable stakeholder in data center sustainability.

S6 **Students** represent the future of data centers. As such they must not only learn about the incredibly complex nature of data centers, but also about ways in which these facilities impact our planet. They are the DC architects, operators, managers, customers and researchers of tomorrow. Knowledge on DC energy efficiency must be made accessible to them if we are to ensure a sustainable ICT path forward.

S7 The **general public** may begin to care more and more about the efficiency and sustainability of the data center that run their daily online services. Just as the average citizen is becoming more aware of the carbon footprint of flying in a plane and what they can do to mitigate it, it is not hard to imagine that the growth in importance and energy use of data centers will result in a similar interest. Like policy makers, the general public needs information to be intelligible to them, with any explanation of new concepts and technicalities being readily available.

4.1.2 Use Cases

The diversity in stakeholders leads to an even more diverse set of use cases for EESR. We generalize these into three core use cases that encompass the possible applications and interactions of EESR. Alongside each use case, involved actors are specified as stakeholders.

U1 DC Energy Efficiency and Sustainability Assessment

Data center operators (**S1**) should be able to assess a given data center's energy characteristics through the lens of efficiency and sustainability. A data center designer (**S2**) or researchers (**S5**) should be able to do the same for a planned or hypothethical DC.

U2 DC Energy Efficiency and Sustainability Reporting

Data center operators (**S1**) should be able to easily report the energy efficiency and sustainability of their centers to; customers (**S3**), policy makers (**S4**), the general public (**S7**) and any other interested stakeholders, in such a way that is understandable and visual coherent with modern standards.

U3 Capabilities Modification and Extension

Researchers (**S5**) and students (**S6**) should be able to fully dissect and extend the system's base capabilities to freely conduct experiments regarding data centers, whilst adapting the instrument to the fast evolving pace of the industry.

4.1.3 Functional Requirements

Based on the use cases envisioned in the previous section, we synthesize the functional requirements of EESR.

FR1 Generate reports pertaining to energy efficiency and sustainability.

The system should allow users to easily produce DC energy reports given a set a validated inputs. Report generation should be largely automatic with as little user interference as possible, encouraging a wider range of possible users.

FR2 Allow report generation to be configurable and shareable.

Despite being easy to operate, the report generation capabilities of the instrument should be highly customizable. In addition to providing various presets, users should be able to provide their own configuration of the report should they wish to do so.

Furthermore, generated reports should allow for easy distribution to any relevant parties.

FR3 Model the energy sustainability of a data center's operation.

EESR must model DC sustainability through various parameters given an energy consumption trace. This model must be as accurate as possible with respect to the lack of accurate modelling data available. The system shall support input not only from OpenDC but also other simulators and real world data centers.

FR4 Produce various sustainability metrics regarding a data center's operation.

The system shall generate a set of specified sustainability metrics, determined in section 3.3, which can be fed into the report generator.

FR5 Enable the exploration of diverse scenarios and their effect on sustainability.

EESR must allow users to explore various 'what if' scenarios given the same input. This allows researchers and the like to examine sustainability under different circumstances, finding potential for improvement.

4.1.4 Non-Functional Requirements

Explicit adherence to a set of non-functional requirements allow us to produce an instrument that not only functionally serves its purpose but is also tailored to the assorted needs of the stakeholders operating or interacting with EESR.

NFR1 Performance

The system should be deliver reports virtually instantaneously, and perform modelling and metric calculation within a reasonable amount of time.

NFR2 Accessibility

Generated reports must meet modern accessibility standards. For instance, ensuring that those with any visual impairments may still interpret them.

NFR3 Usability

Despite its numerous features, EESR must assure usability to a range of audiences through thorough documentation and examples of use.

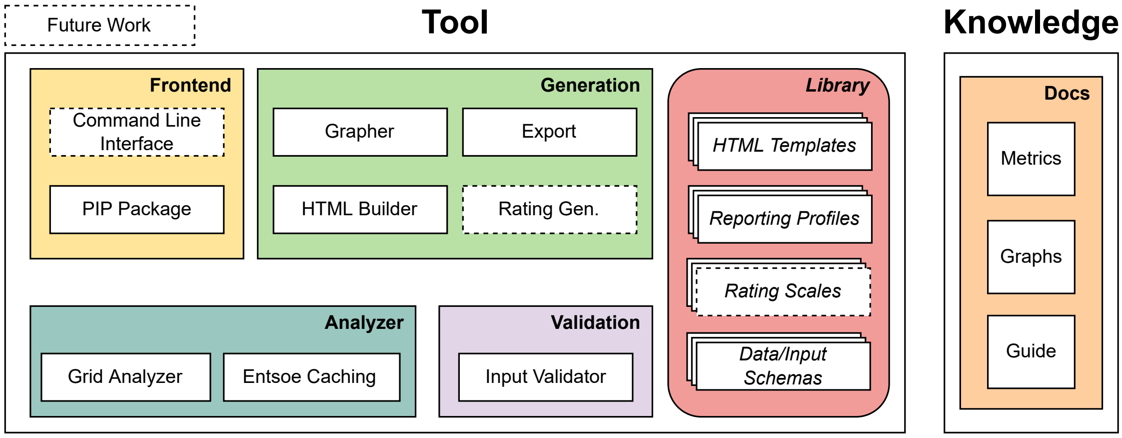

Figure 4.1: Overview of the EESR instrument contribution.

4.2 Overview of EESR

EESR consists primarily of two components; a reporting tool and a DC sustainability modelling system. An interface provides a link between the two components and the user. Accompanying the instrument is an online open source documentation, containing a user guide and a configuration overview. A high level overview of the system and its sub components is visible in figure

4.3 Comprehensive Design of the Reporting Module

In short, the reporting module is nothing more than a custom report generator which has built in functionality for representing various DC efficiency metrics, coupled with a minimal graphing library. As part of the design philosophy of the instrument, most sections of the report can be fully customized with the help of a input format validation.

4.3.1 Design of the Report

Although primarily a creative task, design of the report must involve a systemic approach to be successful. We begin by performing a quick survey of current reporting approaches, taking note of strengths and weakness. Of these inspirations, was Qarnot's 'Carbon Facts' report, which we used a basis for our design [90]. Our design process involved multiple iterations, evolving a concept from paper sketches, to digital prototypes. Each iteration was the result of consultation with various stakeholders and advisors as well as evaluation against our requirements analysis. Figure 4.2 presents the evolution of the report throughout the design phase, from paper and pen sketch to actualization in code.

The final design is a vertical rectangle with four main sections (for *standard* format), positioned from top to bottom. The first section, lists the reported energy efficiency and/or sustainability metrics. Each listed metric implements a mouse tool tip displaying the full name. Additionally, each metric acts as a button linking to an explanation the metric. If using builtin metrics, this button redirects the user to our online documentation. Beside each metric the report may include a lightning bolt icon indicating an energy efficiency metric, or a leaf icon indicating an energy sustainability metric. The second section, displays additional relevant information. This could be additional energy information about the reported data center, such as the total energy usage during the assessment period, or domain specific information e.g. SLA violations or job success rate. The third section holds a graph of the user's choosing providing deeper insight into the reported data center. The graph may provide a breakdown of the various energy sources or energy usage over time. The graph can either be produced by the built in graphing module or be provided externally. There is also an optional graph comment slot under the figure. The final section presents metadata about the source of the data. This includes the assessment period, details about the measurement environment (e.g. simulator, data center, etc.), and lastly a unique UUID generated for each report.

Visually, we employ a simplistic and pleasing look. This is in line with primary purpose of the report; quickly convey the energy efficiency and sustainability of a data center. We omit unnecessary colors and elements. Large sans serif text highlights the most important numbers, whilst smaller text is still visible.

4.3.2 HTML Builder

The base report is generated through HTML. After considering various report generation tools and language libraries, we determined that pure HTML and CSS would allow to build the exact report we designed from the ground up. Furthermore, using HTML allows us to implement all the customization requirements, an option which would be severely limited with other tools. Generation thus uses a provided report *template*, and populates it with the user provided information. This approach permits basic support for user submitted HTML templates, should they wish to have a different format.

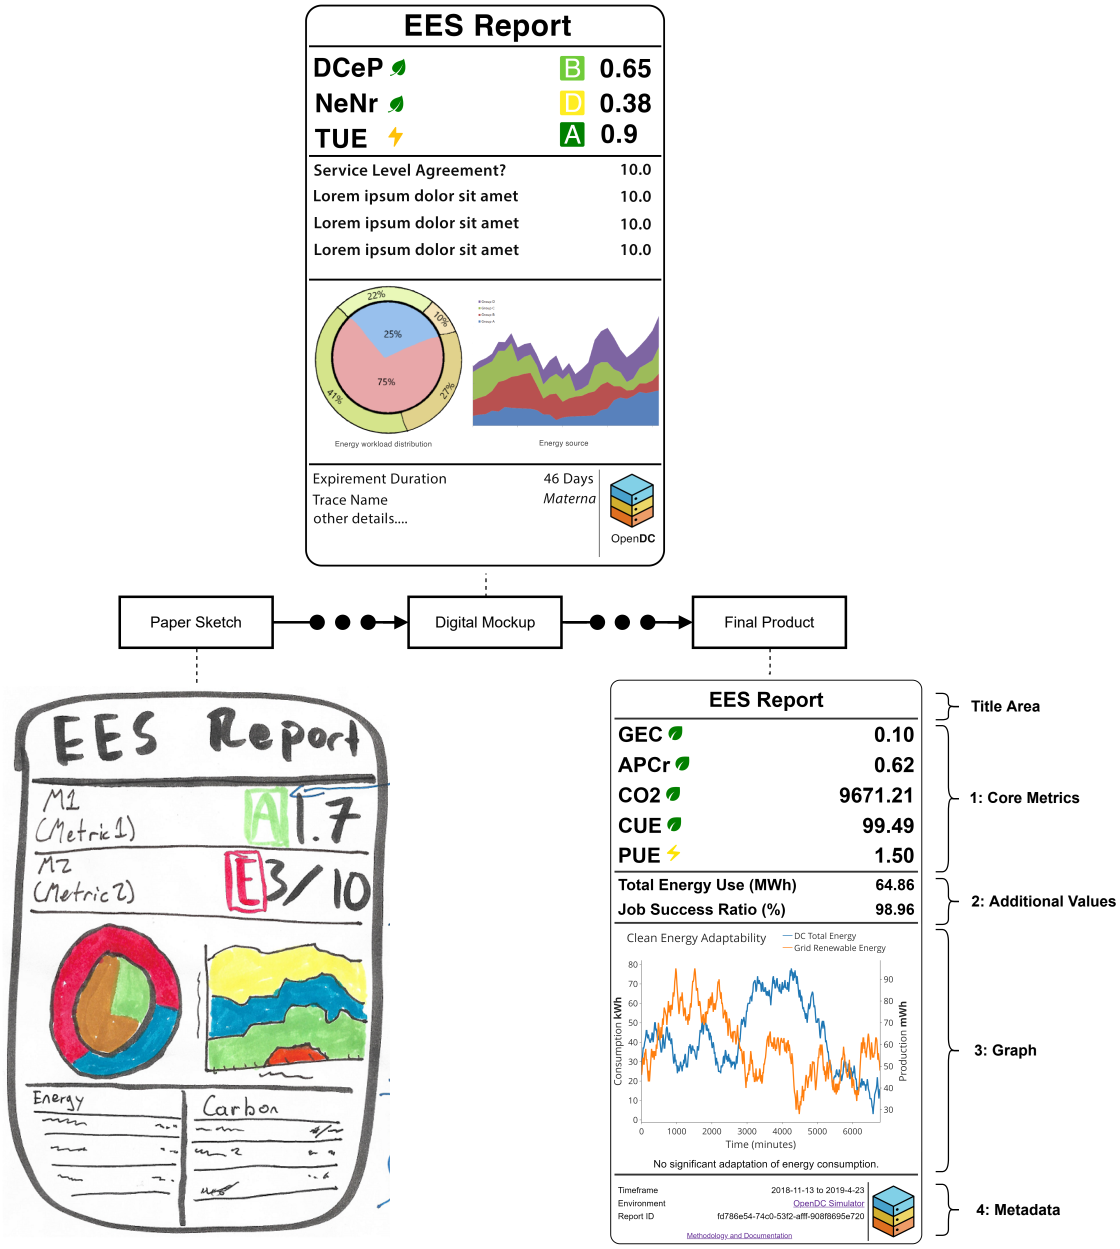

Figure 4.2: Evolution of the EESR report design.

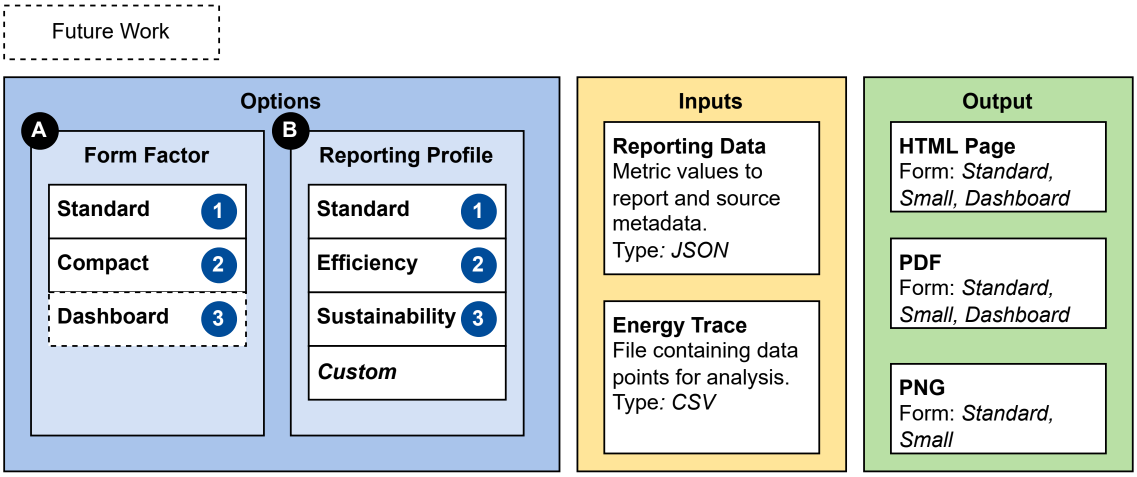

Figure 4.3: Input and configuration of EESR report module.

4.3.3 Configuration

EESR users are presented with a plethora of configuration options, which determine the output of final report. These options are summarized in figure 4.3

4.3.3.1 Form Factor

A First and foremost, users are given the option to choose between two report form factors; 1 *standard* and 2 *compact*. Standard, described above in section 4.3.1 displays everything relevant in a comfortably sized format. Compact, shortens the report to the essentials by removing the third section, and narrowing the overall width. We envision as future work a third form factor: 3 *dashboard*, as an interactive shareable web page displaying multiple graphs, all metrics, and any additional information about the source.

4.3.3.2 Reporting Profiles

B Users can further choose between three builtin *reporting profiles* or supply their own custom profile. Reporting profiles describe which information should be displayed on the generated report. To keep the contents concise and accessible, these profiles compromise completeness by focusing on a specific subject. Reporting profiles only enforce metrics and the graph, but also suggest additional values to be reported. At the time of writing, EESR comes equip with three profiles; 1 *standard profile*, 2 *energy efficiency profile*, and 2 *energy sustainability profile*. These profiles are mere suggestions by the author

	Metrics	**Graph**	**(Suggested) Additional Values**
Standard	TUE 3.2.1, GEC (ren) 3.2.1, CUE 3.3.4	Industry comparison bar chart for DPPE	Total Energy Consumption
Energy	TUE 3.2.1, DCeP 3.2.2	Industry comparison bar chart for TUE	Total Energy Consumption
Sustainability	GEC (green), GEC (ren) 3.3.1, APC_{ren}, Total CO2, CUE 3.3.4	Bar chart of energy sources	Total Energy Consumption

Table 4.1: Breakdown of different reporting profiles provided in EESR.

of the paper, and is subject to change depending on the evolution of the field. Table 4.1 breaks down the characteristics of each profile.

4.3.3.3 Output Formats

Apart from producing an *HTML page*, EESR also supports exporting the report in *PDF* and *PNG* format. Evidently, full report functionality is present only in the HTML format. However, this format is the least shareable as the report will not properly display graphs and other images if it is loaded on a different machine. This is because these assets are stored in the programs data files and can not be inserted into a single HTML file. PDF loses the tool tip functionality in exchange for being fully shareable and maintaining high quality. Crucially, links on the page still work. Finally, exporting as an image provides no interactivity with the report but presents the universally compatible option. However, we expect usage of the image export to be rare.

4.3.4 Validation

Enabling report customization is a schema definition and input validation system. Of the four customization options; form factor, graph, reporting profile, and custom metrics, the latter two have full schema support in EESR. Schemas serve two main purposes. Firstly, they act as a comprehensive documentation for how the instrument expects input. Although one could be perhaps guess the input format from the built in configurations, a schema precisely defines which values are required versus those are optional, their types,

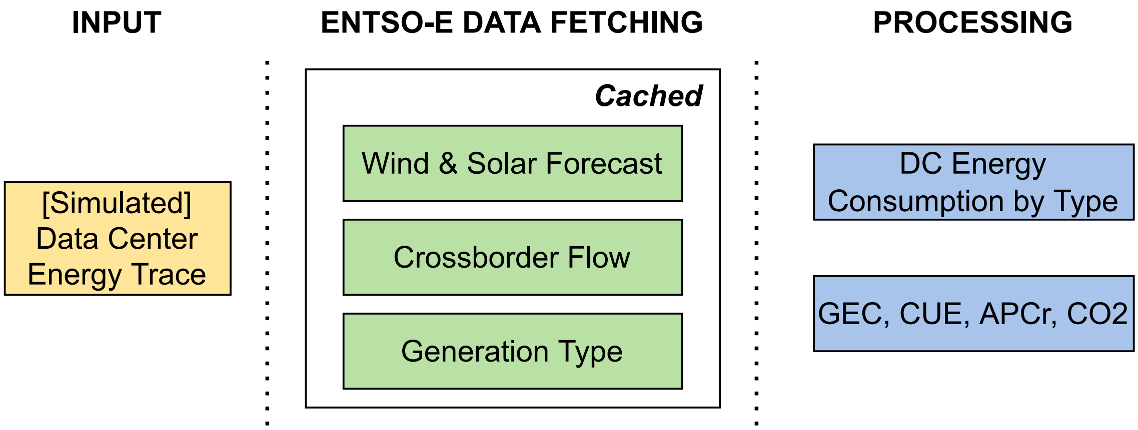

Figure 4.4: Overview of the input and operations of the EESR grid analysis module.

and features descriptions for each input item. As a secondary but equally important purpose, the schemas are used during the validation process of the input. The validation step crucially detects any non-conforming inputs early on in the report generation workflow. This can be achieved with automated tools, providing a significant reduction in implementation effort.

4.4 Comprehensive Design of the Grid Analysis Module

EESR is delivered with an analysis module capable of determining the energy usage of a data center by type of energy. Access to this data opens the door to DC sustainability analysis using different metrics. A general overview of this module is provided in figure 4.4.

As input, the instrument accepts a tabular data structure with three columns: *timestamp*, *it power total*, and *dc power total*. This simple format means that the data can originate from a DC simulation run or a real world DC trace energy trace. From there the system dynamically fetches the relevant energy production data from the ENTSO-E Transparency Platform[1] restful API, after which it performs the necessary calculations. The system also implements a caching system to improve performance and reduce data usage. As output, the instrument produces a summary of metrics in a format ready to be ingested by the reporting module, as well as a full trace of various energy use and sustainability for deeper analysis.

[1] https://transparency.entsoe.eu/

4.4.1 Configuration

The grid analysis module gives users control over two primary dimensions of configuration. First, users can choose to analyze grid data for any period of time between 2016 and the current date, allowing them to explore sustainability under different economic and meteorologic conditions than were present at the time of recording. Secondly, users may determine any country as a data center location that is covered under ENTSO-E (see figure 4.5). With this, experiments can be carried out examining the sustainability difference of deploying the same data center in different countries. Additional parameters are also exposed to the user for tuning.

4.4.2 Modelling DC Energy Sources

National energy production by type can easily be obtained from ENTSO-E. Yet, for a more accurate representation of DC consumption by type we must also consider import energy into the country. One of the outcomes of the Europe electricity market is that countries can purchase electricity from each other. To account for this, when fetching the energy production for a country, EESR grid analysis also fetches cross border electricity flow, and the energy production by type of neighbouring countries. The instrument then adds the additional import energy to the national values, and performs any necessary recalculations (e.g. total energy availability).

We identify two ways of modelling the energy usage by type of a data center's operation. The first, which we name *naive*, assumes a proportional use of electricity by every major industry in a country. In other words, no consumer producer PPA are in effect and every consumer consumes exactly what is on the grid. If the grid's electricity profile is currently comprised of 70% offshore wind electricity, then every industry has an energy consumption profile of 70% offshore wind electricity. The modelling of the naive consumption model is fairly straight forward and can be understood through equation 4.1.

$$DC_i^{type} = DC_i^{total} \times \frac{Grid_i^{type}}{Grid_i^{total}} \tag{4.1}$$

Where:
- DC is the data center,
- $Grid$ is the electricity grid,
- i is the moment in time,
- $type$ is the target energy type, and
- $total$ defines the total electricity consumption (DC)/production (Grid).

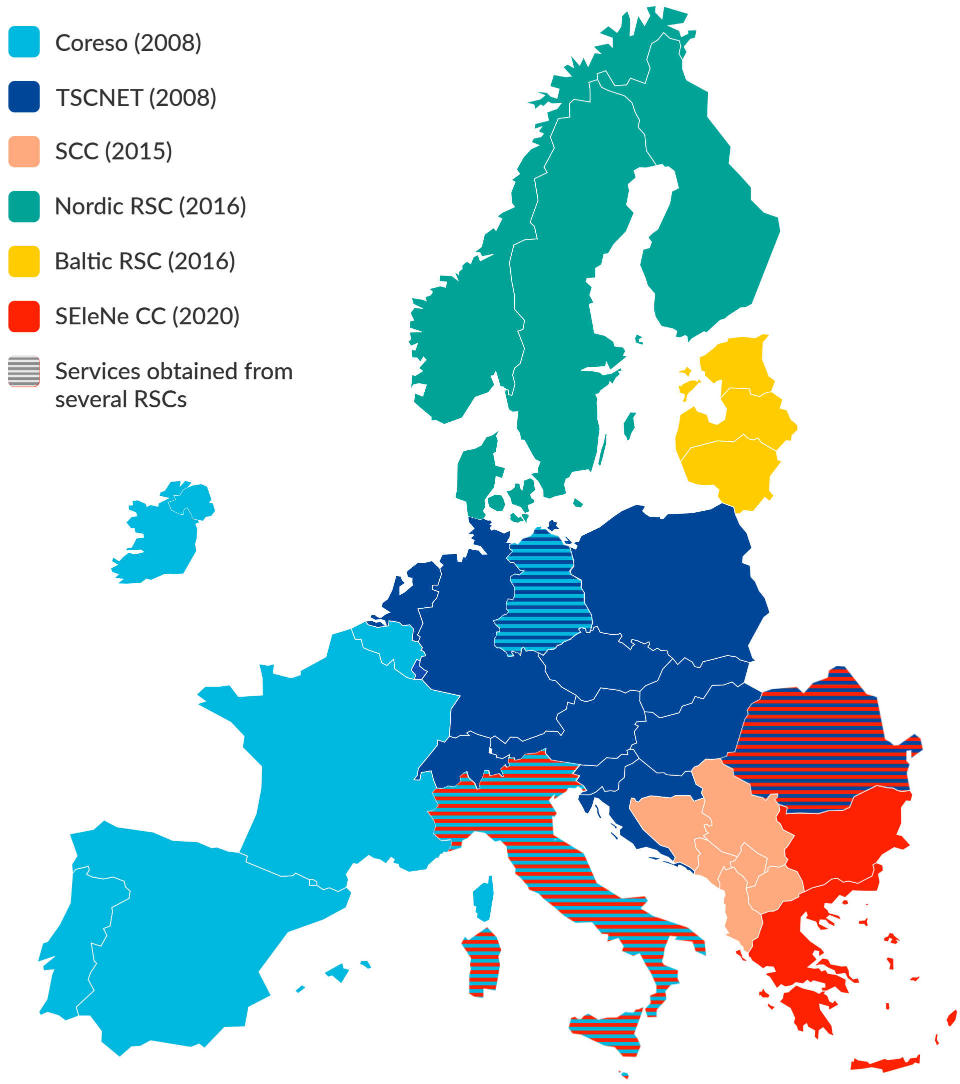

Figure 4.5: ENTSO-E member states, supported in EESR analysis module (figure source: [91])

Such that $Grid_i^{total}$ would mean the total grid electricity production at moment i, and DC_i^{type} means the total electricity consumption by the data center of electricity source *type*. This calculation is then performed for every non-zero electricity type, for every time i.

As an alternative model, a user may choose to calculate DC energy consumption by type using the *greenness assurance* model. Similarly to the naive model, we calculate for each time instance the ratio of the electricity available type to total electricity available. However, we proportionally scale that ratio up or down, depending on whether the energy source is renewable or not, such that the ratio of renewable energy to total energy reaches a GEC level specified by the user, before multiplying it with the data center's total electricity consumption.

In reality, given the prevalence of PPAs, neither of these models mirror what a DC would report as a GEC value. However, the insight is still applicable.

4.4.3 On Renewable vs Green Energy

Many of the sustainability calculations performed, can be defined with either a *renewable* energy target or a *green* energy target. For instance, ignoring its name, calculating GEC could be done with the numerator representing the total green energy or the total renewable consumed. We believe that both derivations are valuable. To that end, we calculate with both definitions in most cases, and allow the user to specify otherwise. Table 4.2 represents our categorization of the different sources provided by ENTSO-E used by EESR based on [92] and [93]. Seeing as there is no universally recognized categorization, we encourage the community to fine tune this taxonomy.

4.4.4 Modelling CO2 usage

Accurately modelling the CO2 emissions of a data center's energy use is extremely difficult. Especially when considering, that the system must support most countries in Europe. Nonetheless, we set out to provide some estimation of CO2 equivalent emissions. We determine for each energy source type, a best case scenario and worst case scenario. There are typically two ways of measuring the global warming potential (GWP) of an activity. Determining the direct quantity of GHG emissions during the activity or, performing a life cycle analysis of the activity. In the case of an activity consuming electricity, a life cycle analysis (LCA) includes everything from extraction of primary energy, to emissions caused by the construction of the power generation facility. The latter method provides a more

Electricity Source	Renewable	Green	gCO2/kWh Best	gCO2/kWh Worst
Mixed	✗	✗	0	0
Biomass	✓	✓	8.5[94]	130[94]
Fossil Brown coal/Lignite	✗	✗	800[94]	1300[94]
Fossil Coal-derived gas	✗	✗	279[95]	849[95]
Fossil Gas	✗	✗	128[96]	434[96]
Fossil Hard coal	✗	✗	600[94]	1050[94]
Fossil Oil	✗	✗	530[94]	900[94]
Fossil Oil shale	✗	✗	385[95]	385[95]
Fossil Peat	✗	✗	126[97]	269[97]
Geothermal	✓	✓	11[98]	47[98]
Hydro Pumped Storage	✓	✓	0	0
Hydro Run-of-river and poundage	✓	✓	2[96]	5[96]
Hydro Water Reservoir	✓	✓	6.1[96]	11[96]
Marine	✓	✓	0	0
Nuclear	✗	✓	5.1[96]	6.4[96]
Other renewable	✓	✓	0	0
Solar	✓	✓	8[96]	83[96]
Waste	✗	✗	376[95]	376[95]
Wind Offshore	✓	✓	12[96]	23[96]
Wind Onshore	✓	✓	7.8[96]	16[96]
Other	✗	✗	0	0

Table 4.2: Categorization and CO2 emissions of different electricity sources used in EESR modelling.

earnest view of the environmental impact of the energy use. As such we aim to provide an LCA value where possible. Due to the scope of this paper the values gathered are not thoroughly audited and simply serve as a starting point for reporting CO2. Unfortunately it is not possible to find a Europe wide LCA value for each energy source. Therefor, we employ a ranked search method for each energy type. When one rank doesn't produce results, we attempt the following rank.

1. LCA GWP value generalized to Europe.
2. LCA GWP value for the Netherlands.
3. LCA GWP value for the World.
4. Non-LCA (instantaneous) GWP emissions value generalized to Europe.
5. Non-LCA (instantaneous) GWP emissions value generalized for the Netherlands.
6. Non-LCA (instantaneous) GWP emissions value generalized for the World.
7. Value of opposite scenario (best/worse)
8. **0.0**

When finding multiple values of the same rank, we prioritize recency and publisher reputation. The findings are detailed in table 4.2.

4.5 Limitations of EESR Design

Models of Energy Consumption by Type

Due to their complex nature the models defined in section 4.4.2 suffer from two major limitations. First, although we take into account incoming cross border electricity flow, we only do this for one hop. This means that we don't calculate the cross border flows of the bordering countries of the primary country and so on. We also don't consider the export of electricity, which could further affect the national ratios of energy type.

CO2 Emissions Collected Data We determine to the best of our ability the CO2 equivalent emissions for each energy source featured in our application. Unfortunately, due to our ranked search method the numbers are not uniform in what they represent. We suspect that this weakens our model significantly.

5

Implementation of EESR

In this chapter, we describe the implementation of our instrument. Highlighting along the way, the major challenges and technical decisions made. The majority of the code base is implemented in Python *3.10.x+*. This choice is rationalized through the language's reputation for data science applications and its large open source community [99]. The use of python further addresses **NFR3**, being considered one of the easier languages to operate for people lacking technical background [100].

We make the EESR open source through GitHub. Usage of EESR is enabled through installation with *pip* with the following command:

```
> pip install opendc-eesr
```

Users can then import the library into any project. The `interface` module should serve most use cases of EESR. Through it, one can perform grid analysis for an input trace and generate reports. For extra control users are welcome to access the individual classes encapsulating the two pillars of the application.

5.1 EESR Reporting

The reporting module's functions can be decomposed into *input validation*, *report generation*, and finally *export*. Report generation further includes the basic graph generator, should users choose to automatically produce figures as opposed to supplying them. Each sub function leverages one or more libraries to simplify implementation and increase codebase quality through community testing and maintenance.

5.1.1 Input Validation using JSON Schema

EESR expects a JSON file adhering to a predefined schema[1] as input for reporting. The schema specifies which *built-in metrics* can be reported, how to provide custom *additional metrics*, *domain* specific information, and finally the required *metadata*. When reporting a built-in metric, the user must only provide the value, as the metric meta-data (*full name*, *documentation link* and *associated icon*) is stored locally. Using the *jsonschema* library[2] the `validator` runs the input against the schema aborting in case of any violations. Last but not least, the program performs a last check guaranteeing that the provided metrics are sufficient for the chosen reporting profile. Defining a custom reporting profile works similarly[3].

5.1.2 HTML Generation

Report generation starts by parsing a bare bones HTML file containing the essential CSS and div elements. The file, acting as blueprint, is then populated using the *domonic* library[4]. Several functions create new HTML elements which are appended to predefined *hooks* in the HTML DOM. As a graphing library we use *Plotly*. Alongside its powerful plotting capabilities, Plotly integrates easily with its sister product *Dash*, enabling the opportunity to implement a dashboard format as future work.

5.2 EESR Grid Analysis

The Pandas library is the main driving force behind our grid analysis module. The entire trace of various metrics and measurements (up to 40 columns and thousands of rows) is stored in a single dataframe. Pandas allows us to comfortably perform column wise arithmetic and append new columns. In table 5.1 an exhaustive list of all the possible resulting columns of an analysis are displayed. As previously mentioned, analysis should be possible for any trace source. However, seeing as development was meant for the OpenDC ecosystem, EESR comes prepackaged with OpenDC trace prepossessing and method for specifically launching OpenDC grid analysis.

[1] `https://github.com/philippsommer27/opendc-eesr/blob/main/eesr/reporting/library/schemas/metrics_input_schema.json`

[2] `https://github.com/python-jsonschema/jsonschema`

[3] `https://github.com/philippsommer27/opendc-eesr/blob/main/eesr/reporting/library/schemas/profile_schema.json`

[4] `https://github.com/byteface/domonic`

Column Name	**Unit**	**Description**
`[Grid Production Type]`	MWh	Amount of energy available on the grid for the specified type.
`grid_total_prod`	MWh	Total energy available on the grid.
`renewable_total`	MWh	Total *renewable* energy available on the grid.
`non_renewable_total`	MWh	Total *non-renewable* energy available on the grid.
`green_total`	MWh	Total *green* energy available on the grid.
`non_green_total`	MWh	Total *non-green* energy available on the grid.
`renewable_perc`	-	Percentage of *renewable* energy available on the grid.
`non_renewable_perc`	-	Percentage of *non-renewable* energy available on the grid.
`green_perc`	-	Percentage of *green* energy available on the grid.
`non_green_perc`	-	Percentage of *non-green* energy available on the grid.
`[DC Production Type]`	kWh	Amount of energy consumed by the DC for the specified type.
`it_energy_total`	kWh	IT equipment total energy consumption.
`dc_energy_total`	kWh	DC total energy consumption .
`dc_renewable_total`	kWh	DC total *renewable* energy consumption.
`dc_non_renewable_total`	kWh	DC total *non-renewable* energy consumption.
`dc_green_total`	kWh	DC total *green* energy consumption.
`dc_non_green_total`	kWh	DC total *non-green* energy consumption.
`total_co2_best`	$kgCO_2eq$	GHG emissions as a result of DC energy consumption, assuming **best case scenario**.
`total_co2_worst`	$kgCO_2eq$	GHG emissions as a result of DC energy consumption, assuming **worst case scenario**.
`cue`	gCO_2eq/kWh	Carbon Usage Effectiveness of the DC.

Table 5.1: Metrics and data points produced by EESR grid analysis.

5.2.1 Grid Awareness Through ENTSO-E

Achieving grid awareness is again greatly aided through use of a community built library: entsoe-py[1]. The library acts as an interface to the public API, conveniently returning a data frame of the requested data. Through it we dynamically request 'Actual Generation per Production Type', 'Cross-Border Physical Flow', and 'Generation Forecasts for Wind and Solar'. Whilst the first two are needed to determine our data center's energy usage profile, the last query type is only used in an auxiliary function of EESR which serves renewable energy predictions for workload shifting and scheduling. We implement the two models designed in section 4.4.2, with the second model unfortunately representing a performance sink due to requiring row by row operation.

5.2.2 Data Gaps and Discrepancy Handling

We recognize the importance of data integrity in producing a model as accurately as possible. Unfortunately, data received from ENTSO-E can often contain gaps or be of varying time units. As such some level of data massaging and filling is implemented at various stages. To highlight this potential threat to validity, we expose every point along the analysis process where data may be modified.

Differing Market Time Units

Market Time Unit is a select length of time for which energy prices are determined. In certain contexts, this is also known as the Imbalance Settlement Period. An MTU of one hour signifies that prices are set for one hour intervals. Complications arise when considering that different bidding zones use different MTUs. Measurement of energy data in ENTSO-E also follows the chosen MTU. Consequently, when processing the import energy the program may encounter disparate time frequencies. By default EESR assumes an MTU of 15 minutes, and attempts to convert data in other time frequencies to 15 minutes. This of course is a loss of potential granularity, as the conversion process divides the total energy production by 2 (for 30 minute MTU) or 4 (for 1 hour MTU).

Gaps in Grid Data

Whether it's measuring infrastructure outages or misalignment due to DST, it can occur that certain queries return a shorter trace than expected. This is problematic when trying to perform column wise arithmetic between two queries, as it would result in the production of null values for a part of the trace. To combat this, the EESR must assume the missing data somehow. In cases where the difference in column lengths is more than 10% of the

[1] `https://github.com/EnergieID/entsoe-py`

expected length, we dismiss the country entirely as an import source. Otherwise, duplicate the last n rows, where n is the difference in column length.

Incomplete Data

Finally, for unknown reasons a grid energy trace may contain null values. When an entire column is null values we simply remove it from the dataframe. However random null values between legitimate numbers are replaced with **0**.

5.2.3 Caching

The caching system for entsoe queries is implemented for various reasons. First, to increase the running time performance of the instrument by eliminating the need to wait for the API to serve our requests. Second, to enable duplicate analysis to be performed in the case of a network outage on either end of the connection. Third, to reduce the load on the ENTSO-E servers. Fourth, to provide a backup for sake of reproducability, should the data ever become unavailable.

Caching is implemented using a dictionary (stored on disk) that maps `CacheEntry`s to a pickle file storing the dataframe. We define a `CacheEntry` as a data class with fields:

- `document kind`: a short string representing the query type,
- `country origin`: the base country for which the request is made,
- `start`: the start time of the query,
- `end`: the end time of the query, and
- `country to`: an optional field indicating the target country in case query is of the cross border flow type.

Every query is first checked for existence in the cache map. If a cache hit occurs, the path to the pickle file is retrieved, and the dataframe is returned to the caller. If there is a cache miss, the data is requested from ENTSO-E. Upon arrival a pickle file name is generated based on a hash of the `CacheEntry`, the dataframe is stored under that name, and an entry is added to the cache map. For testing purposes EESR can be used with caching turned off, which when active will always result in a cache miss.

6

Evaluation

In this section, we evaluate the work produced in this thesis and conduct experiments to gain insight into the energy sustainability of data centers. We present a summary of our findings:

Main Findings

MF1 The current implementation of EESR grid analysis is memory intensive.

MF2 Caching in EESR grid analysis enables a running time improvement of up to 50x and significant savings in data usage.

MF2 EESR can produce reports on energy efficiency and sustainability of the data center.

MF4 EESR can provide deep insights into the energy sustainability of data centers.

MF5 A data center that runs business-critical workloads in the Netherlands in 2020 for three months causes 40.4 tonnes of CO2 equivalent emissions.

MF6 EESR enables analysis of a data center under various 'what if' scenarios.

MF7 The energy sustainability of a data center energy consumption has increased between 2016 and 2020.

6.1 Experimental Setup

In this section, we detail our experimental setup; the inputs, configurations, physical machine, and methodology. All experimental code, data inputs (except data center trace), and results are published as an open source repository[1].

We execute all experiments on a local machine running Windows 10, equipped with an Intel i7 4790k (8 cores), 16 GB of RAM and a gigabit (advertised) fiber optic internet connection. For ease of access and to promote structured implementation, experiments are written and run in a Jupyter Notebook. The Python implementation is CPython (ver 3.10 2), and the Jupyter environment is using IPython version 8.1.1. EESR (ver 0.0.1) is installed through `pip`.

As an input DC energy trace, we simulate an 82 day data center trace in the OpenDC simulator. The trace provided by Solvinity[2], is representative of a general data center that runs business critical workloads in the Netherlands in 2015 and consists of CPU utilization over time. Note that due to confidential information, the DC trace is omitted from our experiment artifact. The DC trace is accompanied by an environment description specifying the topology of the data center. The topology consists of clusters of machines, with specification of the amount of *cores*, core *speed (GHz)*, total *memory (GB)*, number of hosts, *memory (GB)* available to each host, and core count per host for each cluster. The topology is understood by OpenDC as the running environment during the simulation. Simulation employs a linear power model, with a maximum draw of 350 Watts and an idle power use of 200 Watts per machine. We implement a custom output writer which records the total energy consumption of IT equipment of the data center at a frequency of 15 minutes.

As a standard configuration for EESR grid analysis (exceptions are detailed in each experiment), we used a PUE value of 1.59, a figure consistent with the world average as of 2020 [101], to calculate the energy consumption of the entire data center. Furthermore, although the original trace dates from 2015, we simulate the energy trace for the year 2020. The reason is two-fold. First, ENTSO-E only provides complete and usable data from 2016 onward through its web API. Second, most of our assumptions and literature references are from around 2020, to provide relevant information in the thesis. Finally, unless otherwise stated, we specify the Netherlands as the host country. Note that we omit the ENTSO-E api token from the experiment artifact. Individuals can request a token directly from ENTSOE-E [3]

The output of each standard experiment (excluding performance experiments) is a JSON file containing metrics matching the *sustainability reporting profile* (§4.3.3.2), a pickle file

[1] `https://github.com/philippsommer27/experiments-bsc-thesis-2022`

[2] `https://www.solvinity.com/`

[3] `https://transparency.entsoe.eu/content/static_content/download?path=/Static%20content/API-Token-Management.pdf`

encoding the dataframe of the entire trace, and an HTML, PDF, and PNG output file of an EESR report in compact form factor.

6.2 EESR Grid Analysis Performance Evaluation

In section we identified **U1**: the use of EESR by various stakeholders to analyze the sustainability of a real or theoretical data center. As a consequence, non-functional requirement **NFR1** (performance) was determined to ensure that the instrument could be used sensibly in this context. A solution that is too inefficient decreases the chance of adoption and the possible scale of analysis.

We conduct two experiments to investigate the performance the EESR grid analysis module; measuring the impact of caching (running time and network savings) and memory, respectively. The measured experiment is the baseline experiment (§6.3), with reporting disabled.

6.2.1 Memory Impact

Memory impact of performing a grid analysis is assessed using the memory_profiler[1] package. This allows us to observe how much memory is used by the instrument over time. For further insight, we analyze a modified EESR implementation with functions of interest annotated with `@profile`.

The results of the benchmark are visible in 6.1. The `preprocess` method converts the output of the OpenDC simulator into a dataframe intelligible to the grid analysis module. Within the `init` function of the analysis module, the data of the matching grid is obtained for the source country and the neighboring countries and concatenated into the data frame. Finally, the `analyze` function calculates the various sustainability metrics and determines the DC's energy usage by source.

From the moment data pre-processing commences, EESR consumes on average 171.49 megabytes of memory. We can extrapolate that the average memory usage of running a trace spanning an entire year would consume on average 763.3 megabytes of memory. This significant figure exposes a major limitation of our analysis tool. During OpenDC trace pre-processing, the program's memory usage peaks at 208.7 megabytes. We conclude that this is due to the dataframe size being largest during pre-processing as the output file from the OpenDC simulation includes for every time interval the energy usage of each host. During pre-processing the rows are aggregated into total energy usage for all hosts.

[1] `https://github.com/pythonprofilers/memory_profiler`

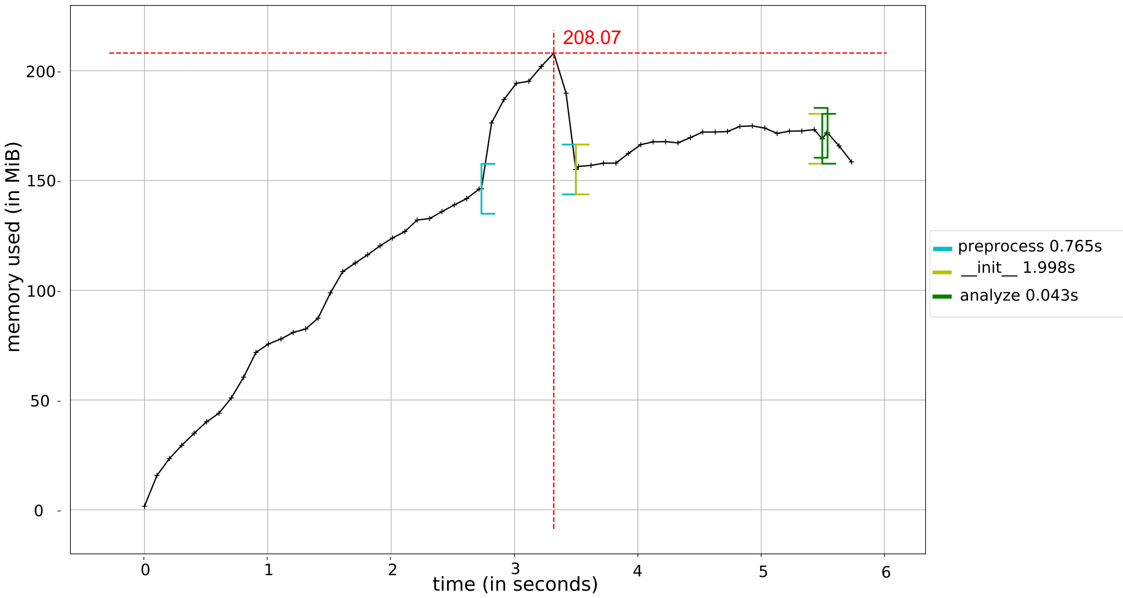

Figure 6.1: Memory profile of running the baseline experiment.

An easy optimization therefor would be to modify our OpenDC writer so that each time interval is already aggregated.

6.2.2 Cache Evaluation

Next, we benchmark the performance improvements of using the cache for ENTSO-E queries. By caching previously fetched queries, the program does not have to request the data from ENTSO-E's API. We envision this scenario to be common among use cases. For instance, stakeholders may have to run an analysis multiple times whilst tweaking configuration, or researches (**S5**) may run different traces of the same time frame with different DC configurations (topology, schedulers, workloads, etc.).

6.2.2.1 Running Time

To observe the time difference between having cached enabled and disabled, python's builtin `timeit` module is used. We perform 4 repetitions of the analysis for each configuration. With caching disabled, analysis takes on average 2 minutes and 22 seconds to complete (with a standard deviation of 2.93 seconds). With caching enabled, the total time is on average 2.8 seconds (with a standard deviation of 197 milliseconds). This is a performance increase of just under 50x.

6.2.2.2 Network Data Savings

To determine the bandwidth savings from using the cache, we monitor the packet transfers that occur during a baseline experiment run using Wireshark. At the time of the benchmark, we record our internet connection status with *speedtest.net*[1]. The results are as follows:

- **Download Speed:** 775.03 Mb/s
- **Upload Speed:** 483.58 Mb/s
- **Latency:** 5 ms

By filtering on the IP address of the API's endpoint we can observe that the total number of packets captured was 41990, with total transfer size of 45.35 megabytes. This bandwidth savings is impactful in two ways. First, it increases the accessibility of the instrument to a wider range of audiences with more limited or costly Internet access. Second, envisioning a scenario of wide adoption of EESR, it reduces the load on ENTSO-E's servers.

6.3 Baseline: Grid Analysis with *naive* Energy Source Model

In this section, we cover our baseline experiment. We examine the simulation of a DC in 2020 in the Netherlands. This demonstrates both the capabilities of our solution as well as gives us insight into the energy sustainability of a general data center. To determine the energy use by type we select the *naive* model which, as explained in Section 4.4.2, assumes a proportional energy use to what is present on the grid. A report of the results is presented in 6.2a.

In summary, the total energy use of the data center amounts to 104.4 megawatt over 82 days. This is equivalent to the annual energy use of 38 households in Amsterdam in 2022[102]. The CUE values indicate that for every kWh of energy consumed, the data center caused on average 615.86 grams of greenhouse gases. This amounted to about 40 tonnes of CO2 equivalent gases expelled into the atmosphere. An APC ren value of 0.48 means that energy consumption matches renewable energy production only somewhat. Finally, 20% of the energy consumed originated from *green* sources, where as 15% originated from renewable sources. This is in line with the production of renewable energy in the Netherlands in 2020 (11.1%) [103], with the largest difference likely due to import of renewable energy.

[1] `https://www.speedtest.net/`

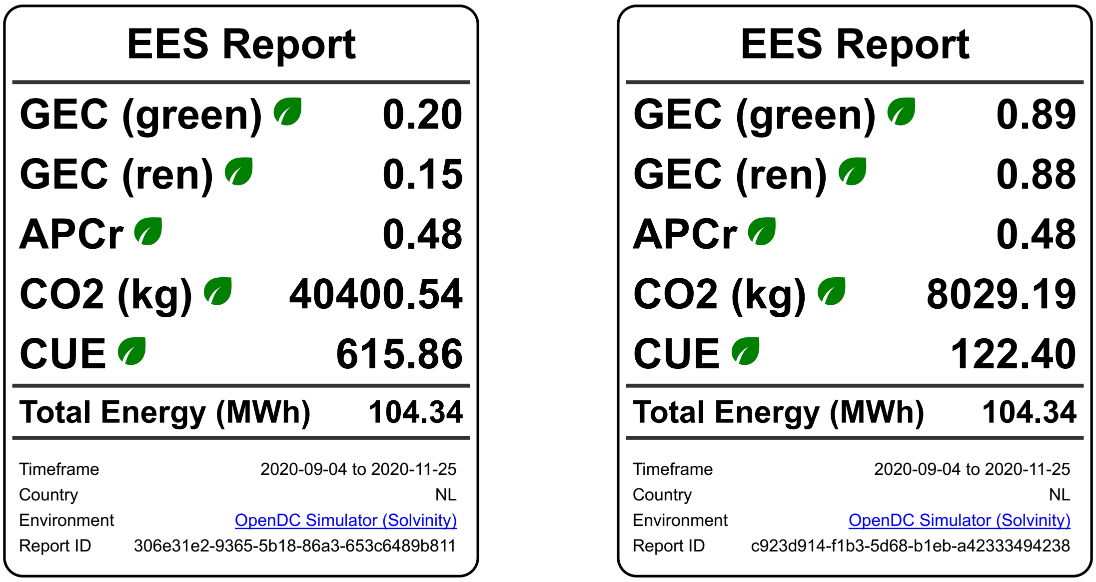

(a) Using *naive* model. **(b)** Using *greenness assurance* model.

Figure 6.2: EESR reports of experiments using different models of energy consumption by source.

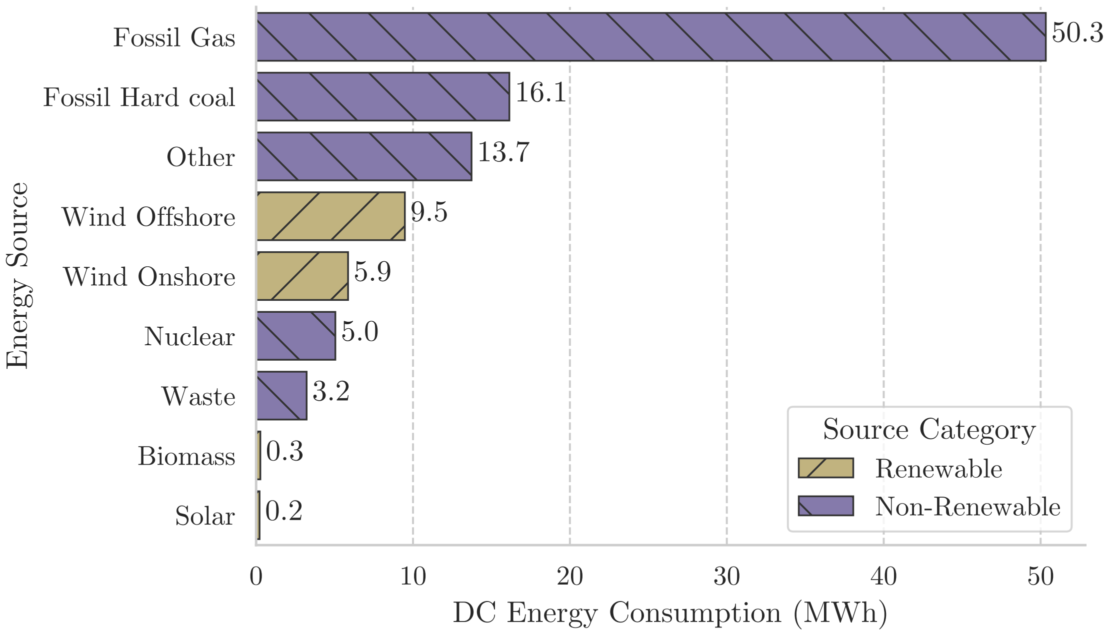

Figure 6.3: Energy consumption of the DC by type *naive*.

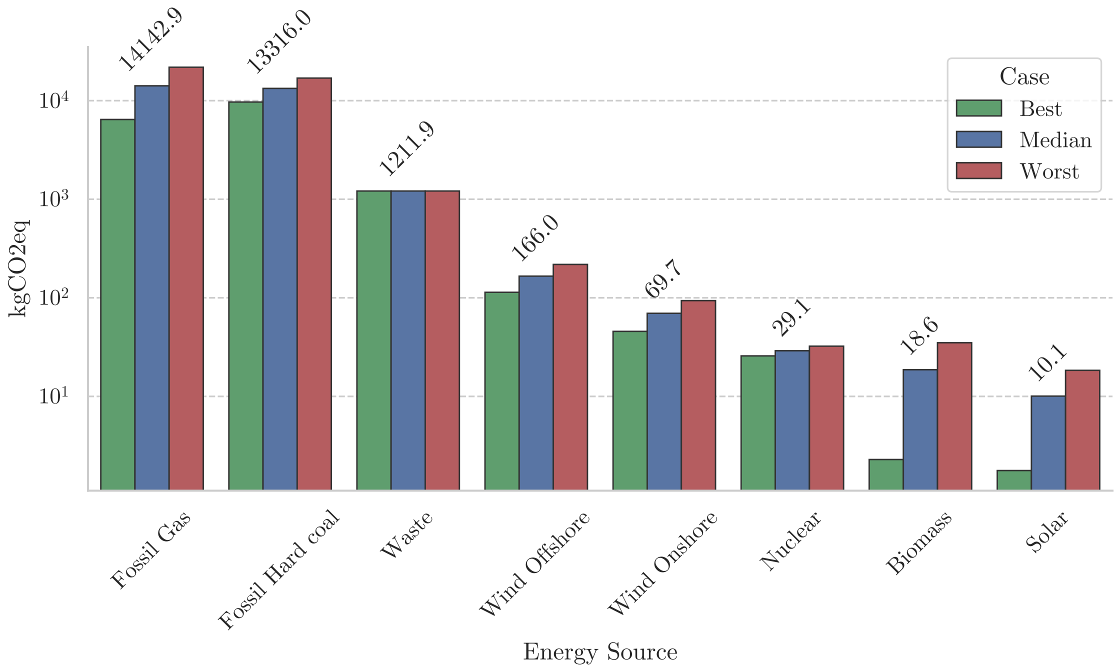

Figure 6.4: Indirect CO2 emissions of the DC through energy use *naive*.

In figure 6.3, we can observe that an overwhelming majority of the energy consumed by the data center was electricity produced with fossil gas. On the x-axis, we plot the energy consumption and on the y-axis, we list the different energy sources. The value of 50.3 MWh is approximately equal to all other energy sources combined. This further explains the low GEC scores reported. Of the major energy sources, only offshore and onshore wind were renewable.

The GHG emissions as a result of the different energy sources are visible in the figure 6.4. The graph shows the total emissions for each source, assuming a best-case, median-case and worse-case CO2 production by type. On the x-axis we list the different energy sources, and on the y-axis we have a logarithmic scale representing total CO2 output. Additionally, for each category, we label the median case emissions. Although they account for three times as much energy use, the CO2 emissions of fossil gas is only 826 more kilograms at a total amount of 14.1 tonnes. As expected, the emissions caused by renewable sources are much lower, with offshore wind causing just 166 kg of CO2.

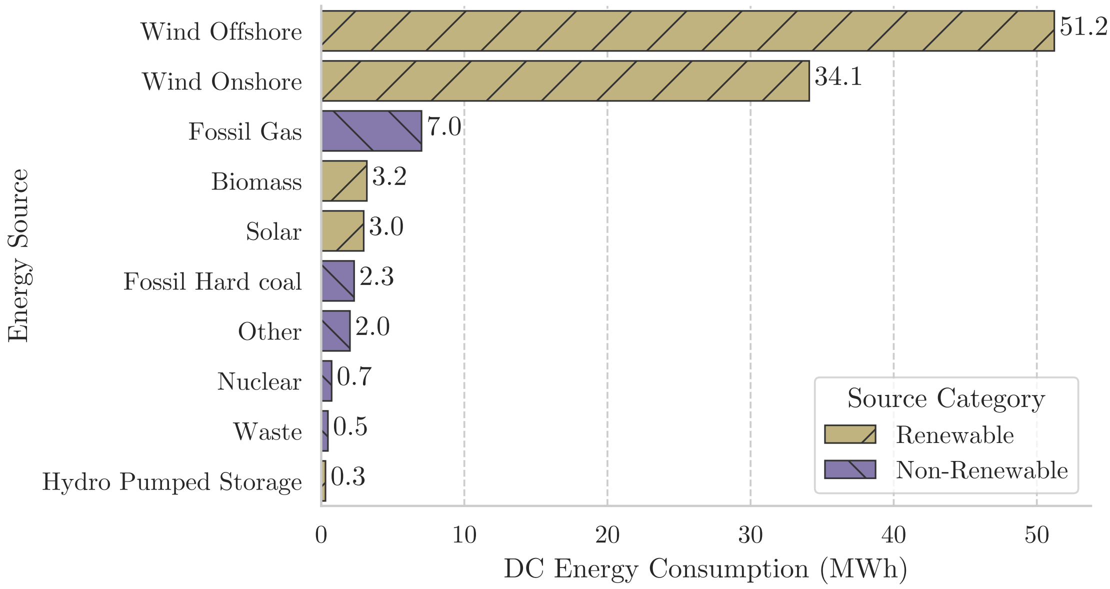

Figure 6.5: Energy consumption of the DC by type *greenness assurance*.

6.4 Grid Analysis with *greenness assurance* Energy Source Model

In comparison to our previous experiment, we now simulate the same data center using the *greenness assurance* (§4.4.2) model with a greenness value of 0.88. This ensures that at every instance, the data center is consuming at least 88% renewable energy. We choose this value based on the claims of the Dutch Data Center association for 2020 [38].

In the report (figure 6.2b), we again see the key metrics calculated by the analysis module. We draw our attention to the significantly reduced total CO2 output of 8029.19 kg. As a result, the CUE score also drops by about 493 grams of CO2 to 122.40.

With the adjusted model, figure 6.5 shows that wind energy now represents the majority of the electricity generated at a combined amount of 85.3 MWh. Renewable sources, biomass and solar energy, represent a larger fraction of the total, placing fourth and fifth, up from eighth and ninth.

In terms of GHG emissions, the top two sources remain fossil gas and fossil hard coal, however, each contributes considerably less by a factor of about 7x. This experiment demonstrates the environmental benefit of ensuring a green energy source for a facility. Through PPA's data centers contribute monetarily to the adoption of renewable energy.

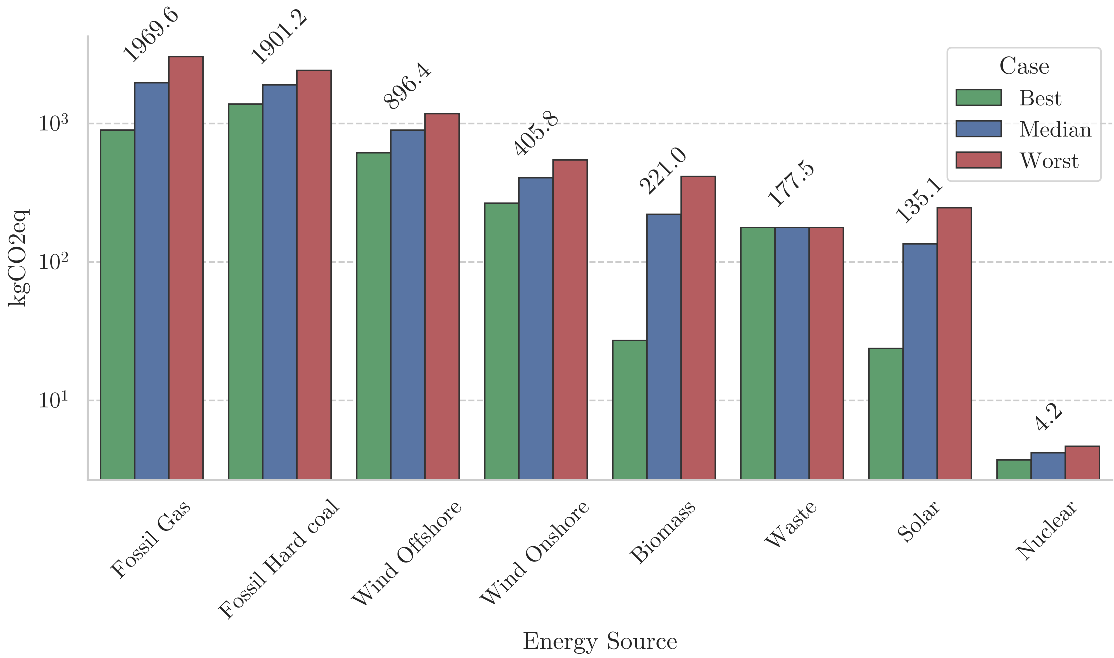

Figure 6.6: Indirect CO2 emissions of the DC through energy use *greenness assurance*.

6.5 Investigation of Seasonal Difference in DC Sustainability

The sustainability of a data center can vary greatly depending on the season due to changing cooling requirements and the availability of wind and solar energy. In this experiment, we quantify this variation. The DC is simulated to start at four different time periods throughout the year. The corresponding months are: January, April, July, and October. The results are presented in figure 6.7. On the y axis of each plot we measure the relevant metric, whilst the x axis represents the start month of the simulation.

The first significant observation is that the DC is more sustainable in winter than in summer. A simulation starting in October achieves a GEC (ren/green) increase in score of 4%. This is likely due to wind farms, which account for the majority of renewable energy in the Netherlands, being more productive during the winter months when the winds are stronger. Counter intuitively, CO2 emissions also increase slightly in winter. This anomaly may be an expose weakness in our model. The production of renewable energy most strongly matches the energy consumption profile of the data center from January to April, as indicated by the maximum value of 0.59. These results hint at the fact that workloads should be concentrated during the winter.

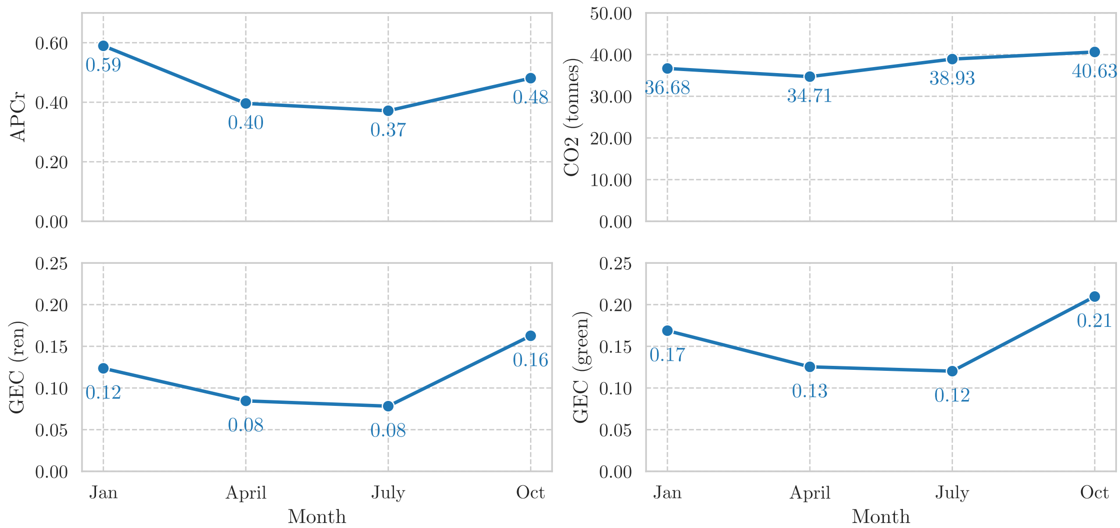

Figure 6.7: Trend plots of various sustainability metrics of DC during different seasons.

6.6 Investigation of Yearly Trends in DC Grid Energy Consumption

Similarly to to the previous experiment, we investigate how the energy sustainability of a data center has evolved over 5 years between 2016 and 2020. The purpose of this experiment is to analyze the direct effect of increasing the adoption of renewable energy on the sustainability of the energy of data centers. For each run, we start the simulation at the same time as the baseline, with only the year changed. We additionally configure the PUE value to match the global average for the year according to the Uptime Institute [101].

The results are illustrated in 6.8. We first notice the significant increase in the GEC metric, especially between 2019 and 2020 where renewable GEC increased by 8% to 15% and green GEC increased by 9% to 20%. These improvements are reflected in the total amount of CO2 released. In 2016, the DC caused 55.93 tonnes of greenhouse gas emissions. By 2020, this number has dropped to just 40.4 tonnes. As expected, the CUE score experiences a similar decline of approximately 28%. A noteworthy observation is that, despite PUE briefly spiking to 1.67 in 2019, from 1.58 in the previous year, increased use of renewable energy meant that CO2 emissions still decreased by 8.71 tonnes. Had this not happened, we could expect total CO2 to increase in 2019 as well due to higher energy consumption. Finally, APC ren remains mostly stable throughout the years increasing by

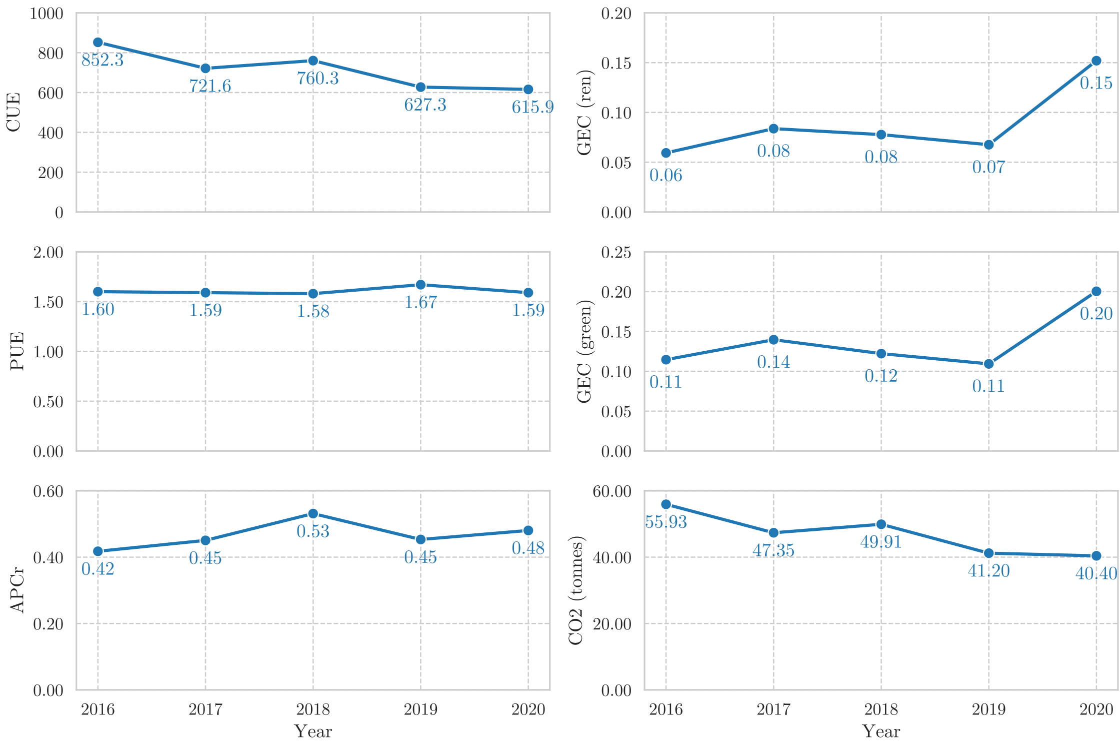

Figure 6.8: Trend plots of sustainability metrics throughout the years 2016-2020.

just 0.06.

6.7 Investigation of DC Sustainability in Different Countries in Europe

As a final experiment, we produce reports of energy sustainability of the same data center workload running in different countries. As mentioned previously, tenants may wish to choose which facility to run their workload based on sustainability factors. This experiments attempts to shed light on the potential differences between some countries in Europe. We choose the following country, describing the choice of each:

1. **Germany** holds the largest number of data centers within its borders, making it a relevant case to study [104].
2. **France** uniquely produces most of its energy through nuclear energy, amongst other things known for its extremely low CO2 emissions. Ranking fourth in the amount of data center's, makes it a valuable option to investigate [104, 105].
3. **Sweden** leads renewable energy production in 2020 with a 60% of its energy coming from various renewable sources [106].
4. **Belgium** ranks low in terms of renewable energy production [106]. We include the country in this context to investigate its carbon impact.

For some countries, ENTSO-E defines sub-grids for different parts of the nation. In the case of Germany and Sweden, where this is true, we specify the control zone that contains the capital city. Note that EESR automatically adjusts for time zones so that the simulation starts at the same time in all countries.

From the generated reports we can extra a few key observations. First, Sweden achieves a high score of renewable and green energy use, which leads to a low CO2 output of 7195.83 kg. Scoring even higher for GEC (green) is France because of its advantageous use of nuclear energy. CO2 and CUE are the lowest among the four countries analyzed. Germany remains an environmentally costly host country choice due to its heavy dependence on fossil coal and gas. A workload running in Germany produces 5.6 tonnes more CO2 than in the Netherlands, totaling 45982.13 kg of greenhouse gases. Of the countries evaluated, France's renewable energy curve most closely follows the data center's energy consumption curve.

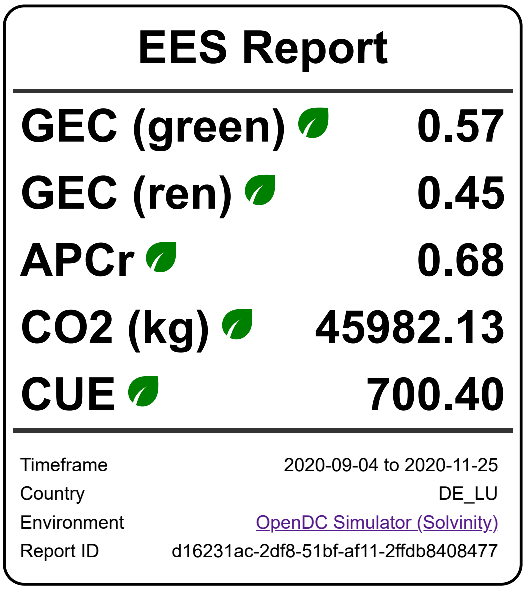

(a) Germany

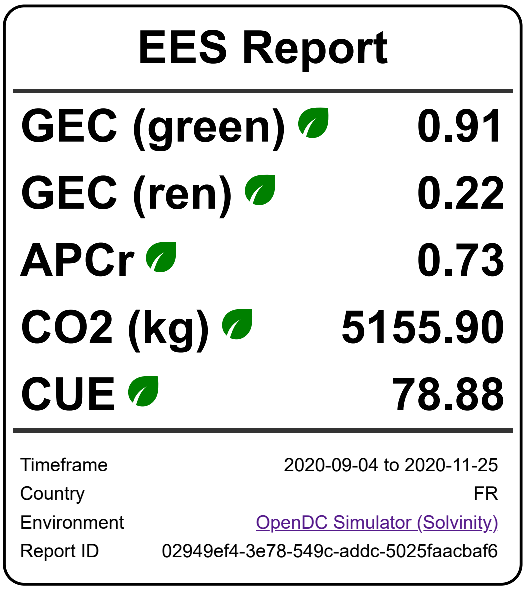

(b) France

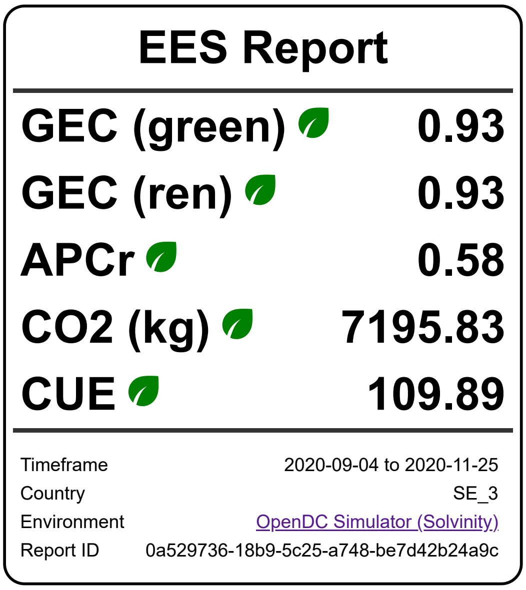

(c) Sweden

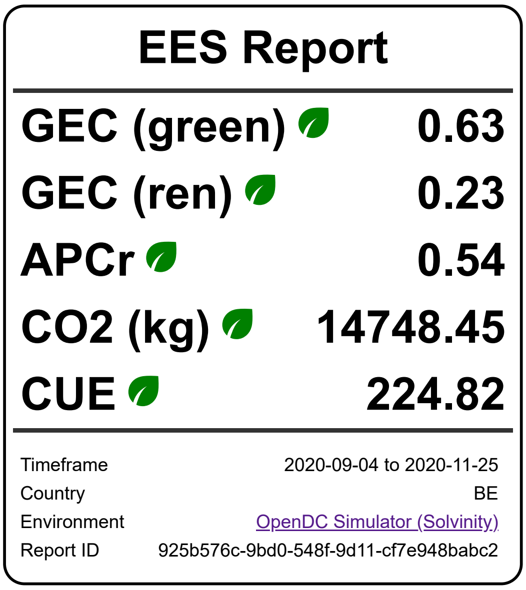

(d) Belgium

Figure 6.9: EESR reports of energy sustainability analysis of various countries.

6.8 Limitations

Input Trace Length. One key limitation of our experiments is the length of the input trace. A time frame of under three months would be considered sub-optimal with regards to measurement recommendations of metrics like CUE, which exclaim that values should be derived from an entire year. We combat this limitation by stating the time frame on our reports and further examining sustainability throughout different stages of the year.

Validation of Model Another significant limitation is that we perform a limited validation of our model. Due to both a lack of similar work available and time constraints, we cannot verify the accuracy of the metrics produced. Nonetheless, the comparison between various experiments proves to be valuable.

Validation of EESR Report A last limitation is the lack of validation of the report generator. Due to the scope of the thesis, we are unable to conduct large-scale surveys to gather feedback on the report. Instead, we validate our reporting solution against the identified requirements in Section 4.1 and obtain feedback from an expert in the field.

7

Conclusion

In this chapter, we conclude our work in this thesis. We set out to answer (a small part of) the big questions regarding the energy sustainability and efficiency of data centers. To achieve this, we detail a set of three research questions in Chapter 1. We now answer these research questions:

7.1 Answering Research Questions

RQ1.1 Which selection of metrics are best suited to broadly and deeply report the energy efficiency and sustainability of a data center?

To address this question, we conduct a short literature survey of data center energy efficiency metrics. The results of this study are presented in Chapter 3. For applicability, we then materialize our research into a set of three reporting profiles detailed in section 4.3.3.2.

RQ1.2 How to effectively communicate the energy efficiency and sustainability of a data center to various stakeholders?

We implement EESR Reporting, a flexible data center energy efficiency reporting tool. To enable ease of use, we automate most of the report generation procedure and ensure that the tool can be easily installed by everyone. Accompanying EESR reporting is an open source knowledge bank of metrics and EESR documentation. This improves our goal of communication, as it breaks down the knowledge barrier.

RQ3 How can we use our model and solution to draw conclusion on the sustainability of general data centers running business critical workloads?

We implement the EESR Grid Analysis module; a configurable DC energy trace analyzer providing insight into the sustainability of a data center's energy consumption. We then leverage our solution to carry out various experiments. From the experiments, we extract multiple main findings about the past and present of data center sustainability.

7.2 Future Work

The possible areas of future work with respect to this thesis are practically limitless given the vast scope of our work. We narrow these avenues to the most valuable and promising prospects:

1. We expect the EESR knowledge bank to expand in detail and with the number of included metrics. Research on energy efficiency metrics is fragmented and only a few surveys gather current work. The knowledge bank attempts to consolidate this knowledge in one accessible place through the contribution of the community.

2. We identify the pressing need to properly validate all aspects of EESR. First, the reporting tool must be presented to various stakeholders for feedback. Second, the numbers produced by the grid analysis module need to be checked against other similar solutions available and/or the real world. This is essential to prove the validity of our work.

3. We see a clear opportunity to improve our assumptions on CO2 emissions. A long-term goal our work attempts to contribute to is to provide a life cycle assessment of data center operations. The key to this mission is accurate greenhouse gas emissions values for our model. Our current data set is generalized to the entire continent and is subpar for some sources. We imagine a version of EESR with the ability to produce an accurate analysis of CO2 emissions based on the specific plants from which electricity originates.

7.3 Envisioned Societal Impact

We envision a long-term societal impact as result of our work (and continued developement). Current and future data centers around Europe could leverage our technology to expand our view of DC energy sustainability and make smarter decisions for our planet. Our research on DC metrics and energy efficiency reporting could be used and expanded to usher the industry into the era of shared information. The ability to assess a facility's energy footprint through the analysis module provides operators and designers the opportunity to rethink DC operations at an energy level. This not only as societal implications for stakeholders but also for climate stability. Naturally, a greener and more energy efficient data center presents a positive economic impact.

More concretely, we see this vision first come to fruition with our ongoing efforts with SURF and EGI. EGI provides data space and computing power for intensive scientific research through many data centers in Europe. As a part of their mission, they aim to increase the sustainability and global warming impact of their services. EESR can play a part in this goal within the next 3 to 5 years.

References

[1] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. **Recalibrating global data center energy-use estimates**. *Science*, **367**(6481):984–986, 2020. 1, 3, 4, 19

[2] IEA. **Digitalisation and Energy**. Technical report, International Energy Agency, Paris, 2017. 1, 3, 19

[3] David Reinsel-John Gantz-John Rydning, J Reinsel, and J Gantz. **The digitization of the world from edge to core**. *Framingham: International Data Corporation*, **16**, 2018. 1

[4] Mauricio Arregoces and Maurizio Portolani. *Data center fundamentals*. Cisco Press, 2003. 1

[5] Google. **Google Environmental Report 2022**. Technical report, Google LLC, 2022. 1, 3

[6] Torsten Hoefler, Ariel Hendel, and Duncan Roweth. **The Convergence of Hyperscale Data Center and High-Performance Computing Networks**. *Computer*, **55**(7):29–37, 2022. 1

[7] Jerry Banks. **Introduction to simulation**. In *Proceedings of the 31st conference on Winter simulation: Simulation—a bridge to the future-Volume 1*, pages 7–13, 1999. 1

[8] J.S. Carson. **Introduction to modeling and simulation**. In *Proceedings of the Winter Simulation Conference, 2005.*, pages 8 pp.–, 2005. 1

[9] Fabian Mastenbroek, Georgios Andreadis, Soufiane Jounaid, Wenchen Lai, Jacob Burley, Jaro Bosch, Erwin Van Eyk, Laurens Versluis, Vincent van Beek, and Alexandru Iosup. **OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters**. In

Laurent Lefèvre, Stacy Patterson, Young Choon Lee, Haiying Shen, Shashikant Ilager, Mohammad Goudarzi, Adel Nadjaran Toosi, and Rajkumar Buyya, editors, *21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, CCGrid 2021, Melbourne, Australia, May 10-13, 2021*, pages 455–464. IEEE, 2021. 1, 15

[10] Lotfi Belkhir and Ahmed Elmeligi. **Assessing ICT global emissions footprint: Trends to 2040 & recommendations**. *Journal of Cleaner Production*, **177**:448–463, 2018. 1, 3

[11] Murray G Patterson. **What is energy efficiency?: Concepts, indicators and methodological issues**. *Energy policy*, **24**(5):377–390, 1996. 2

[12] Eurostat. **Share of energy from renewable sources**, March 2022. `https://ec.europa.eu/eurostat/databrowser/view/nrg_ind_ren/default/table?lang=en`. 2

[13] Ashikur Rahman, Xue Liu, and Fanxin Kong. **A Survey on Geographic Load Balancing Based Data Center Power Management in the Smart Grid Environment**. *IEEE Communications Surveys Tutorials*, **16**(1):214–233, 2014. 2

[14] Hao Wang, Jianwei Huang, Xiaojun Lin, and Hamed Mohsenian-Rad. **Exploring Smart Grid and Data Center Interactions for Electric Power Load Balancing**. **41**(3):89–94, jan 2014. 2

[15] Adam Wierman, Zhenhua Liu, Iris Liu, and Hamed Mohsenian-Rad. **Opportunities and challenges for data center demand response**. In *International Green Computing Conference*, pages 1–10, 2014. 2

[16] Hao Chen, Michael C. Caramanis, and Ayse K. Coskun. **The data center as a grid load stabilizer**. In *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 105–112, 2014. 2

[17] Charles D. Keeling. **The Concentration and Isotopic Abundances of Carbon Dioxide in the Atmosphere**. *Tellus*, **12**(2):200–203, 1960. 2

[18] John Cook, Naomi Oreskes, Peter T Doran, William R L Anderegg, Bart Verheggen, Ed W Maibach, J Stuart Carlton, Stephan Lewandowsky, Andrew G Skuce, Sarah A Green, Dana Nuccitelli, Peter Jacobs, Mark Richardson, Bärbel Winkler, Rob Painting, and Ken

RICE. **Consensus on consensus: a synthesis of consensus estimates on human-caused global warming**. *Environmental Research Letters*, **11**(4):048002, Apr 2016. 2

[19] MARCO SIDDI. **The European Green Deal: Assessing its current state and future implementation**. 2020. 2

[20] DIRECTORATE-GENERAL FOR CLIMATE ACTION EUROPEAN COMMISSION. **Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions**, 2020. https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:52020DC0562. 2

[21] CHARLES D. KEELING, ROBERT B. BACASTOW, ARNOLD E. BAINBRIDGE, CARL A. EKDAHL JR., PETER R. GUENTHER, LEE S. WATERMAN, AND JOHN F. S. CHIN. **Atmospheric carbon dioxide variations at Mauna Loa Observatory, Hawaii**. *Tellus*, **28**(6):538–551, 1976. 2

[22] BARRY S. LEVY AND JONATHAN A. PATZ. **Climate Change, Human Rights, and Social Justice**. *Annals of Global Health*, **81**(3):310–322, 2015. Climate Change, Global Health and Human Rights. 2

[23] NUOA LEI AND ERIC MASANET. *GLOBAL DATA CENTER ENERGY DEMAND AND STRATEGIES TO CONSERVE ENERGY*. 04 2021. 3, 16, 19

[24] RICH MILLER. **The Sustainability Imperative: Green Data Centers and Our Cloudy Future**, Jan 2021. https://datacenterfrontier.com/green-data-center-imperative/. 3, 30

[25] CLIMATE NEUTRAL DATA CENTRE PACT. **Self-Regulatory Initiative Policy Proposal**, 2020. https://www.climateneutraldatacentre.net/wp-content/uploads/2021/06/CNDCP-Policy-Paper_FINAL.pdf. 3

[26] CENTRAAL BUREAU VOOR DE STATISTIEK. **Elektriciteit geleverd aan datacenters, 2017-2020**, Dec 2021. 3

[27] ANDERS ANDRAE. **Should we be concerned about the power consumption of ICT?**, 05 2018. 3

[28] United States Environmental Protection Agency. **Greenhouse Gas Equivalencies Calculator**, 2022.
https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator#results. 3

[29] Tom Bawden. **Global warming: Data centres to consume three times as much energy in next decade, experts warn**, Jan 2016.
https://www.independent.co.uk/climate-change/news/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-expe
html. 3

[30] Martijn Koot and Fons Wijnhoven. **Usage impact on data center electricity needs: A system dynamic forecasting model**. *Applied Energy*, **291**:116798, 2021. 3

[31] NL Times. **Netherlands power grid can't handle tech developments, Data Centers warn**, May 2018.
https://nltimes.nl/2018/05/03/netherlands-power-grid-cant-handle-tech-developments-dat
3

[32] Jonathan Koomey and Eric Masanet. **Does not compute: Avoiding pitfalls assessing the Internet's energy and carbon impacts**. *Joule*, **5**(7):1625–1628, 2021. 3, 4, 6

[33] Madina Khudaykulova, He Yuanqiong, Akmal Khudaykulov, et al. **Economic Consequences and Implications of the Ukraine-Russia War**. *International Journal of Management Science and Business Administration*, **8**(4):44–52, 2022. 4

[34] International Energy Agency. **Evolution of energy prices, Oct 2020-Jan 2022**, 2022.
https://www.iea.org/data-and-statistics/charts/evolution-of-energy-prices-oct-2020-jan-2022. 4

[35] Albert Greenberg, James Hamilton, David A. Maltz, and Parveen Patel. **The Cost of a Cloud: Research Problems in Data Center Networks**. *SIGCOMM Comput. Commun. Rev.*, **39**(1):68–73, Dec 2009. 4

[36] Gemma A. Brady, Nikil Kapur, Jonathan L. Summers, and Harvey M. Thompson. **A case study and critical assessment in calculating power usage effectiveness for a data centre**. *Energy Conversion and Management*, **76**:155–161, 2013. 4, 18, 19

[37] Victor Avelar, Dan Azevedo, Alan French, and Emerson Network Power. **PUE: a comprehensive examination of the metric**. *White paper*, **49**, 2012. 4, 18

[38] Dutch Data Center Association. **State of the Dutch Data Centers 2020**, 2020.
`https://www.dutchdatacenters.nl/en/publications/state-of-the-dutch-data-centers-2020/`. 4, 5, 16, 58

[39] Eric Masanet and Nuoa Lei. **How Much Energy Do Data Centers Really Use?**, March 2020.
`https://www.agci.org/solutions/quarterly-research/2020-03-DataCenters`. 4

[40] Ana Radovanovic. **Our data centers now work harder when the sun shines and wind blows**, 2020.
`https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/`. 4

[41] Nathaniel Horner and Ines Azevedo. **Power usage effectiveness in data centers: overloaded and underachieving**. *The Electricity Journal*, **29**(4):61–69, 2016. 4, 18

[42] Netherlands Environmental Assessment Agency, TNO, Statistics Netherlands (CBS), [Dutch] National Institute for Public Health, and RIVM. **Netherlands Climate and Energy Outlook 2020**, 2020.
`https://www.pbl.nl/sites/default/files/downloads/pbl-2020-netherlands-climate-and-energy-outlook-2020-summary-4299.pdf`. 5

[43] Yair Levy and Timothy Ellis. **A Systems Approach to Conduct an Effective Literature Review in Support of Information Systems Research**. *International Journal of an Emerging Transdiscipline*, **9**, 01 2006. 7

[44] Alexandru Iosup, Laurens Versluis, Animesh Trivedi, Erwin Van Eyk, Lucian Toader, Vincent van Beek, Giulia Frascaria, Ahmed Musaafir, and Sacheendra Talluri. **The AtLarge Vision on the Design of Distributed Systems and Ecosystems**. In *39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019, Dallas, TX, USA, July 7-10, 2019*, pages 1765–1776. IEEE, 2019. 7

[45] Uptime Institute. **Tier Classification System**, 2022. `https://uptimeinstitute.com/tiers`. 11

[46] Research and Markets. **Europe Data Center Market Report 2021-2026: Increasing Demand for High-Density, Mission-Critical Servers, Storage, and Network Infrastructure**, March 2021. `https://www.prnewswire.com/news-releases/europe-data-center-market-report-2021-2026-i` `html`. 11

[47] Xiang Sun, Nirwan Ansari, and Ruopeng Wang. **Optimizing Resource Utilization of a Data Center**. *IEEE Commun. Surv. Tutorials*, **18**(4):2822–2846, 2016. 11, 12

[48] Robert Birke, L. Y. Chen, and Evgenia Smirni. **Report Data Centers in the Wild : A Large Performance Study**. 2012. 12

[49] W Pitt Turner, John H Seader, and Kenneth G Brill. **Industry standard tier classifications define site infrastructure performance**, 2005. 12

[50] Prateek Sharma, Patrick Pegus II, David E. Irwin, Prashant J. Shenoy, John Goodhue, and James Culbert. **Design and Operational Analysis of a Green Data Center**. *IEEE Internet Comput.*, **21**(4):16–24, 2017. 12

[51] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan. *The Datacenter as a Computer: Designing Warehouse-Scale Machines, Third Edition*. Synthesis Lectures on Computer Architecture. Morgan & Claypool Publishers, 2018. 12, 14, 19

[52] Steven Pelley, David Meisner, Pooya Zandevakili, Thomas F. Wenisch, and Jack Underwood. **Power routing: dynamic power provisioning in the data center**. In James C. Hoe and Vikram S. Adve, editors, *Proceedings of the 15th International Conference on Architectural Support for Programming Languages*

and Operating Systems, ASPLOS 2010, Pittsburgh, Pennsylvania, USA, March 13-17, 2010, pages 231–242. ACM, 2010. 12

[53] Philip T. Krein. **Data center challenges and their power electronics**. *CPSS Transactions on Power Electronics and Applications*, **2**(1):39–46, 2017. 12

[54] Matthew J. Hatami. *Oilfield Survival Guide, Volume One: For All Oilfield Situations*. Oilfield Books, LLC, 2018. 12

[55] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. **Data Center Energy Consumption Modeling: A Survey**. *IEEE Communications Surveys & Tutorials*, **18**(1):732–794, 2016. 14

[56] Henri Casanova, Arnaud Giersch, Arnaud Legrand, Martin Quinson, and Frédéric Suter. **Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms**. *Journal of Parallel and Distributed Computing*, **74**(10):2899–2917, June 2014. 15

[57] Dzmitry Kliazovich, Pascal Bouvry, and Samee Ullah Khan. **GreenCloud: a packet-level simulator of energy-aware cloud computing data centers**. *J. Supercomput.*, **62**(3):1263–1283, 2012. 15

[58] Alberto Nuñez, José Luis Vázquez-Poletti, Agustín C. Caminero, Gabriel G. Castañé, Jesús Carretero, and Ignacio Martín Llorente. **iCanCloud: A Flexible and Scalable Cloud Infrastructure Simulator**. *J. Grid Comput.*, **10**(1):185–209, 2012. 15

[59] Rajkumar Buyya, Rajiv Ranjan, and Rodrigo N. Calheiros. **Modeling and simulation of scalable Cloud computing environments and the CloudSim toolkit: Challenges and opportunities**. In Waleed W. Smari and John P. McIntire, editors, *2009 International Conference on High Performance Computing & Simulation, HPCS 2009, Leipzig, Germany, June 21-24, 2009*, pages 1–11. IEEE, 2009. 15

[60] Michael Tighe, Gastón Keller, Michael A. Bauer, and Hanan Lutfiyya. **DCSim: A data centre simulation tool for evaluating dynamic virtualized resource management**. In *8th International Conference on Network and Service Management, CNSM 2012, Las Vegas, NV, USA, October 22-26, 2012*, pages 385–392. IEEE, 2012. 15

[61] Nur IqtiyaniIlham, M. Hasanuzzaman, and M. Hosenuzzaman. **European smart grid prospects, policies, and challenges**. *Renewable and Sustainable Energy Reviews*, **67**:776–790, 2017. 15

[62] Hao Chen, Michael C. Caramanis, and Ayse K. Coskun. **Reducing the data center electricity costs through participation in smart grid programs**. In *International Green Computing Conference*, pages 1–10, 2014. 15

[63] Amory Lovins. **Energy Efficiency, Taxonomic Overview**. *Encyclopedia of Energy*, **2**, 12 2004. 16

[64] Elizabeth Shove. **What is wrong with energy efficiency?** *Building Research & Information*, **46**(7):779–789, 2018. 16

[65] John Hauser and Gerald Katz. **Metrics: you are what you measure!** *European Management Journal*, **16**(5):517–528, 1998. 17

[66] Jumie Yuventi and Roshan Mehdizadeh. **A critical analysis of Power Usage Effectiveness and its use in communicating data center energy consumption**. *Energy and Buildings*, **64**:90–94, 2013. 18

[67] Rădulescu Constanţa Zoie, Rădulescu Delia Mihaela, and Sipică Alexandru. **An analysis of the power usage effectiveness metric in data centers**. In *2017 5th International Symposium on Electrical and Electronics Engineering (ISEEE)*, pages 1–6, 2017. 18

[68] Callum Rivett. **Kolos to build world's largest data centre in Arctic Circle**, May 2020.
`https://technologymagazine.com/cloud-and-cybersecurity/kolos-build-worlds-largest-data-centre-arctic-circle`. 18

[69] Michael K. Patterson, Stephen W. Poole, Chung-Hsing Hsu, Don E. Maxwell, William Tschudi, Henry Coles, David J. Martinez, and Natalie J. Bates. **TUE, a New Energy-Efficiency Metric Applied at ORNL's Jaguar**. In Julian M. Kunkel, Thomas Ludwig, and Hans Werner Meuer, editors, *Supercomputing - 28th International Supercomputing Conference, ISC 2013, Leipzig, Germany, June 16-20, 2013. Proceedings*, **7905** of *Lecture Notes in Computer Science*, pages 372–382. Springer, 2013. 18, 22

[70] Tarandeep Kaur and Inderveer Chana. **Energy Efficiency Techniques in Cloud Computing: A Survey and Taxonomy**. *ACM Comput. Surv.*, **48**(2), oct 2015. 19

[71] Synergy Research Group. **Microsoft, Amazon and Google Account for Over Half of Today's 600 Hyperscale Data Centers**, Jan 2021. 19

[72] Mordor Intelligence. **HYPERSCALE DATACENTER MARKET - GROWTH, TRENDS, COVID-19 IMPACT, AND FORECASTS (2022 - 2027)**, 2021. `https://www.mordorintelligence.com/industry-reports/hyperscale-datacenter-market#`. 19

[73] Keiichi Hirose. **DC powered data center with 200 kW PV panels**. In *2018 International Power Electronics Conference (IPEC-Niigata 2018 -ECCE Asia)*, pages 822–825, 2018. 19

[74] V Dinesh Reddy, Brian Setz, G Subrahmanya VRK Rao, GR Gangadharan, and Marco Aiello. **Metrics for sustainable data centers**. *IEEE Transactions on Sustainable Computing*, **2**(3):290–303, 2017. 21

[75] Dave Anderson, Tahir Cader, Tommy Darby, Nick Gruendler, Rema Hariharan, Anne Holler, Anne Lindberg, Christine Long, Phil Morris, Andy Rawson, Freeman Rawson, Vikram Saletore, Jim Simonelli, Harkeeret Singh, Roger Tipley, Gary Verdun, and John Wallreich. **A Framework for Data Center Energy Productivity**. *White paper*, **13**, 2008. 23

[76] John Haas, Mark Monroe, John Pflueger, Jack Pouchet, Peter Snelling, Andy Rawson, and Freeman Rawson. **Proxy Proposals for Measuring Data Center Productivity**. 23

[77] Energy Star Global Taskforce. **Harmonizing Global Metrics for Data Center Energy Efficiency**, 2014. `https://home.jeita.or.jp/greenit-pc/topics/release/pdf/dppe_e_20140430.pdf`. 23, 24

[78] The European Commission. **EU taxonomy for sustainable activities**, Feb 2022.

https://ec.europa.eu/info/business-economy-euro/banking-and-finance/sustainable-finance/eu-taxonomy-sustainable-activities_en#documents. 24

[79] Anna Darmani, Annika Rickne, Antonio Hidalgo, and Niklas Arvidsson. **When outcomes are the reflection of the analysis criteria: A review of the tradable green certificate assessments**. *Renewable and Sustainable Energy Reviews*, **62**:372–381, 2016. 24

[80] Anna Bergek and Staffan Jacobsson. **Are tradable green certificates a cost-efficient policy driving technical change or a rent-generating machine? Lessons from Sweden 2003–2008**. *Energy Policy*, **38**(3):1255–1271, 2010. Security, Prosperity and Community – Towards a Common European Energy Policy? Special Section with Regular Papers. 24

[81] Marcel Antal, Tudor Cioara, Ionut Anghel, Radoslaw Gorzenski, Radoslaw Januszewski, Ariel Oleksiak, Wojciech Piatek, Claudia Pop, Ioan Salomie, and Wojciech Szeliga. **Reuse of Data Center Waste Heat in Nearby Neighborhoods: A Neural Networks-Based Prediction Model**. *Energies*, **12**(5), 2019. 25

[82] Pei Huang, Benedetta Copertaro, Xingxing Zhang, Jingchun Shen, Isabelle Löfgren, Mats Rönnelid, Jan Fahlen, Dan Andersson, and Mikael Svanfeldt. **A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating**. *Applied Energy*, **258**:114109, 2020. 25

[83] T. Brunschwiler, B. Smith, E. Ruetsche, and B. Michel. **Toward zero-emission data centers through direct reuse of thermal energy**. *IBM Journal of Research and Development*, **53**(3):11:1–11:13, 2009. 25

[84] Hampus Markeby Ljungqvist, Louise Mattsson, Mikael Risberg, and Mattias Vesterlund. **Data center heated greenhouses, a matter for enhanced food self-sufficiency in sub-arctic regions**. *Energy*, **215**:119169, 2021. 25

[85] Smart Cities Cluster. **Cluster Activities Task 3**, Oct 2014. https://www.smartcitiescluster.eu/publications/FP7ProjectCluster-NewEnergyEfficiencyMetricsforDataCentres-task3.pdf. 25

[86] Alexis I Aravanis, Artemis Voulkidis, Jaume Salom, Jacinta Townley, Vasiliki Georgiadou, Ariel Oleksiak, Milagros Rey Porto, Fabrice Roudet, and Theodore Zahariadis. **Metrics for assessing flexibility and sustainability of next generation data centers**. In *2015 IEEE Globecom Workshops (GC Wkshps)*, pages 1–6. IEEE, 2015. 25

[87] Dan Azevedo, Michael Patterson, Jack Pouchet, and Roger Tipley. **Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric**. *White paper*, **32**, Dec 2010. 26

[88] Dan Azevedo, Symantec Christian Belady, and J Pouchet. **Water usage effectiveness (WUE™): A green grid datacenter sustainability metric**. *The Green Grid*, page 32, 2011. 26

[89] James Kaplan, William Forest, and Kindler Noah. **Revolutionizing Data Center Energy Efficiency**, 2014.
https://kuliah.imm.web.id/arsitektur-komputer/Data-center/Revolutionizing_Data_Center_Efficiency.pdf. 27

[90] Qarnot. **Carbon facts: Outil de pilotage de votre consommation énergétique**, 2022.
https://qarnot.com/fr/bas-carbone. 33

[91] ENTSO-E. **Power Regions**, 2022.
https://www.entsoe.eu/regions/. 40

[92] Council of the European Union European Parliament. **Directive (EU) 2018/2001 of the European Parliament and of the Council of 11 December 2018 on the promotion of the use of energy from renewable sources**, Dec 2018. 41

[93] Climate Neutral Data Centre Pact. **Engagements of the Climate Neutral Datacenters Initiative**, 2020.
https://www.climateneutraldatacentre.net/wp-content/uploads/2022/07/20220725__Self-Regulatory-Initiative.pdf. 41

[94] Roberto Turconi, Alessio Boldrin, and Thomas Astrup. **Life cycle assessment (LCA) of electricity generation technologies: Overview, comparability and limitations**. *Renewable and sustainable energy reviews*, **28**:555–565, 2013. 42

[95] Netherlands Enterprise Agency. **The Netherlands: list of fuels and standard CO2 emission factors version of January 2020**, Jan 2020. `https://english.rvo.nl/sites/default/files/2020/03/The-Netherlands-list-of-fuels-version-January-2020.pdf`. 42

[96] United Nations Economic Commission for Europe. **Carbon Neutrality in the UNECE Region: Integrated Life-cycle Assessment of Electricity Sources**, March 2021. `https://unece.org/sed/documents/2021/10/reports/life-cycle-assessment-electricity-generation-options`. 42

[97] Sanni Väisänen et al. **Greenhouse gas emissions from peat and biomass-derived fuels, electricity and heat—Estimation of various production chains by using LCA methodology**. 2014. 42

[98] Annika Eberle, Garvin A Heath, Alberta C Carpenter Petri, and Scott R Nicholson. **Systematic Review of Life Cycle Greenhouse Gas Emissions from Geothermal Electricity**. 2017. 42

[99] KR Srinath. **Python–the fastest growing programming language**. *International Research Journal of Engineering and Technology*, **4**(12):354–357, 2017. 45

[100] AL Sayeth Saabith, MMM Fareez, and T Vinothraj. **Python current trend applications-an overview**. *International Journal of Advance Engineering and Research Development*, **6**(10), 2019. 45

[101] Rhonda Ascierto and Andy Lawrence. **Uptime institute global data center survey 2020**. *Uptime Institute*, **2**, July 2020. 52, 60

[102] James Evans. **Amsterdam Energy Costs**, March 2022. `https://www.citycentreretreat.com/amsterdam-energy-cost.php`. 55

[103] Centraal Bureau voor de Statistiek. **Renewable electricity; production and capacity**, June 2022. `https://opendata.cbs.nl/statline/#/CBS/en/dataset/82610ENG/table?dl=5F0D1`. 55

[104] Statista. **Number of data centers in Europe by country 2021**, August 2021. `https://www.statista.com/statistics/878621/european-data-centers-by-country/`. 62

[105] IEA. **Total energy supply (TES) by source, France 1990-2020**, 2022. https://www.iea.org/countries/france. 62

[106] EUROSTAT. **Share of renewable energy in the EU up to 18.0%**, Jan 2020. 62

Appendix A

Reproducibility

A.1 Abstract

Obligatory

A.2 Artifact check-list (meta-information)

- **Program:** Our main efforts are consolidated into the `EESR` program.
- **Model:** 'naive' and 'greenness assurance' models for EESR grid analysis.
- **Data set:** Available at `https://github.com/philippsommer27/experiments-bsc-thesis-2022/blob/a48dcee006b58b1a73696efccd4de50686d5d628/Input/opendc_out.csv`
- **Run-time environment:** Local Machine
- **Hardware:** i7 4790k, 16GB of RAM, GTX 980, 1 Terabyte of NVME storage.
- **Output:** EESR report and dataframe of energy trace
- **Experiments:** Available at `https://github.com/philippsommer27/experiments-bsc-thesis-2022/blob/a48dcee006b58b1a73696efccd4de50686d5d628/experiments.ipynb`
- **How much disk space required (approximately)?:** 300 MB
- **How much time is needed to prepare workflow (approximately)?:** 15 minutes
- **How much time is needed to complete experiments (approximately)?:** 20 minutes
- **Publicly available?:** Yes
- **Code licenses (if publicly available)?:** Apache License 2.0

A.3 Description

A.3.1 How to access

The following repositories contain all the code and works:

- **Knowledge Bank:** `https://philipp-sommerhalter.gitbook.io/eesr/`
- **EESR Implementation:** `https://github.com/philippsommer27/opendc-eesr`
- **Experiments:** `https://github.com/philippsommer27/experiments-bsc-thesis-2022`
- **Source trace simulated in OpenDC:** `https://github.com/philippsommer27/opendc`

A.3.2 Software dependencies

1. Python
2. Pip

A.4 Installation

Install the main application through pip:

```
> pip install opendc-eesr
```

Clone the repository containing the experiments:

```
> git clone https://github.com/philippsommer27/experiments-bsc-thesis-2022.git
```

A.5 Experiment workflow

Run the two Jupyter Notebooks in the experiment repository. Ensure that you have an ENTSO-E API token.

A.6 Evaluation and expected results

If ran successfully, the *Results* folder will be populated with the outputs of each experiment. Furthermore, the *Plots* folder will contain the relevant plots of each experiment.

A.7 Methodology

Submission, reviewing and badging methodology:

- `https://www.acm.org/publications/policies/artifact-review-badging`
- `http://cTuning.org/ae/submission-20201122.html`
- `http://cTuning.org/ae/reviewing-20201122.html`

...

...