



Renewable Energy in Data Centers: the Dilemma of Electrical Grid Dependency and Autonomy Costs

Wedan Emmanuel Gnibga, Anne Blavette, Anne-Cécile Orgerie

► To cite this version:

Wedan Emmanuel Gnibga, Anne Blavette, Anne-Cécile Orgerie. Renewable Energy in Data Centers: the Dilemma of Electrical Grid Dependency and Autonomy Costs. IEEE Transactions on Sustainable Computing, 2023, pp.1-13. 10.1109/TSUSC.2023.3307790 . hal-04189173

HAL Id: hal-04189173

<https://hal.science/hal-04189173>

Submitted on 28 Aug 2023

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Renewable Energy in Data Centers: the Dilemma of Electrical Grid Dependency and Autonomy Costs

Wedan Emmanuel Gribga, Anne Blavette, and Anne-Cécile Orgerie

Abstract—Integrating larger shares of renewables in data centers’ electrical mix is mandatory to reduce their carbon footprint. However, as they are intermittent and fluctuating, renewable energies alone cannot provide a 24/7 supply and should be combined with a secondary source. Finding the optimal infrastructure configuration for both renewable production and financial costs remains difficult. In this paper, we examine three scenarios with on-site renewable energy sources combined respectively with the electrical grid, batteries alone and batteries with hydrogen storage systems. The objectives are first, to size optimally the electric infrastructure using combinations of standard microgrids approaches, secondly to quantify the level of grid utilization when data centers consume/export electricity from/to the grid, to determine the level of effort required from the grid operator, and finally to analyze the cost of 100% autonomy provided by the battery-based configurations and to discuss their economical viability. Our results show that in the grid-dependent mode, 63.1% of the generated electricity has to be injected into the grid and retrieved later. In the autonomous configurations, the cheapest one including hydrogen storage leads to a unit cost significantly more expensive than the electricity supplied from a national power system in many countries.

Index Terms—Sustainable data centers, carbon-neutral computing, renewable resources, energy storage, optimal sizing.

1 INTRODUCTION

In their quest for ever greater performance, the devices that drive our daily lives rely on powerful computing services through the Internet. Hence, the number of worldwide Internet users has doubled, while global Internet traffic has increased by 30% per year since 2010 [1]. This is due to the rapid adoption of social media, video streaming, online gaming, etc. Most of this traffic is handled by data centers that process, manage and store massive amounts of data. Huge efforts are made to increase the capacity of data centers, leading to significant electrical and environmental challenges. In fact, the global data centers electricity use in 2020 was equal to 200-250TWh, or 1% of the world electricity consumption [1]. Fossil (such as natural gas, oil) energies are currently the main sources used to supply the data centers at a global scale [1]. Thus, they represent a source of greenhouse gas (GHG) emissions that “*will not reduce without major concerted political and industrial efforts*” [2].

A promising approach enabling to reduce these GHG emissions and the related pollution is to improve the energy efficiency. It can be achieved by means of virtualisation, load balancing and consolidation. This combination of techniques consists in making the data centers’ resources available to the customers as virtual machines (VMs), distributing them optimally over the minimum physical resources [3], and switching off the idle physical machines [4]. Another technique called Dynamic Voltage and Frequency

Scaling (DVFS) dynamically adjusts the frequency and voltage of processors running VMs. It saves energy when the machines are idle or run less demanding VMs, while meeting the Service Level Agreement (SLA) [5].

Reducing the electricity consumption of data centers is indispensable to reach sustainability. Yet, it does not guarantee by itself carbon neutrality which is an objective for European Union. Indeed, through its Green Deal, the EU Commission committed to reach carbon neutrality by 2050 [6]. Several ways can be adopted to reach carbon neutrality in data centers such as carbon offsetting or using renewable energies use. In this paper, we investigate this latter option.

Supplying data centers with renewable energies, also labelled as ‘green’, is being intensively investigated by large-scale data center operators who have announced that they managed to purchase – through Power Purchase Agreements (PPAs)– and/or to generate enough renewable energy to meet 100% of their operational needs. Such operators are, among others, Google (12 TWh in 2019), Apple (1.7 TWh in 2020) and Facebook (7 TWh in 2020) [1]. However, renewable energies are intermittent and not fully controllable, making it challenging to synchronize their production with the consumption of the data centers. Therefore, it is essential to integrate an additional permanent and more controllable energy source or combination of sources to the electrical infrastructure for the power supply security. This source may be the electrical grid or energy storage systems. However, few studies analyzed the sizing of the renewable energy source and storage equipment in order to precisely meet the electrical demand of data centers. In addition, none of them assess the level of grid utilization when data centers are consuming energy from the grid, nor do they examine the potential effort this may require

- W. E. Gribga and A.-C. Orgerie are with Univ. Rennes, Inria, CNRS, IRISA, Rennes, France. Emails: wedan-emmanuel.gribga@irisa.fr, anne-cecile.orgerie@irisa.fr
- A. Blavette is with Univ. Rennes, ENS Rennes, CNRS, SATIE lab, Rennes, France. Email: anne.blavette@ens-rennes.fr

Manuscript received February 20, 2023, revised May 23, 2023, accepted August 5, 2023.

from the grid operator. Furthermore, some works propose configurations based on renewable energy to cover 100% of data center demand without justifying their economic relevance. This article presents a techno-economic study on data centers that addresses these unresolved concerns. We make the following contributions:

- We provide an optimal sizing of the energy sources (renewable energy source and energy storage devices) in different configurations, autonomous or not, as well as the associated optimal costs.
- We quantify the contribution of the grid (respectively the renewable energy source) to the consumption of the data centers when they are consuming from the grid and analyze the profiles of power exchanged between the data centers and the grid to quantify the power peak that the grid operator has to handle.
- We discuss the economical viability of autonomous datacenters, and provide recommendations to achieve a high penetration of renewable energies in data centers at a lower cost.

The rest of this paper is organized as follows. Section 2 presents recent work using renewable energy in data centers. Section 3 describes the type of data centers considered in this study and their electrical infrastructures. Section 4 describes the component and economic models whereas Section 5 is dedicated to the infrastructure sizing methodology. Section 6 presents the results. Finally, Section 7 summarises the work and analyses the aspects that may have an influence on the results, and improve the work.

2 RELATED WORK

Cloud service providers are taking initiatives to reduce GHG. In January 2021, data center operators and industry associations in Europe made a tangible achievement by launching the Climate Neutral data center Pact [7] which includes a pledge to make data centers climate-neutral by 2030 and has intermediate (2025) targets for power usage effectiveness and carbon-free energy. These efforts result in a rapid penetration of renewable energies in data centers to replace the traditional sources, and may help to reduce or even eliminate the electricity transport and distribution costs, in the case of on-site generation. This section presents some techniques and technologies designed not only to minimize energy consumption, but also to manage it optimally.

2.1 Flexibility of the IT infrastructures

Cloud services fall into two categories. On one hand, inflexible service tasks must be performed as soon as submitted. On the other hand, batch jobs can be delayed provided that the expected quality of service is fulfilled. Therefore, batch jobs may be distributed according to the abundance of renewable energy and may improve the use of renewable energy. Based on this observation, Í. Goiri *et al.* promoted GreenSlot [8], a scheduler that predicts the amount of solar energy of the near future and schedules the workload to maximize the renewable energy consumption while meeting the jobs' deadlines. When the solar generation is not sufficient, it selects the time when the external grid energy is cheap to submit the tasks with respect to the deadlines. L.

Grange *et al.* propose an approach for scheduling batch jobs with due date constraints taking into account the renewable energy's availability to reduce the use of grid-imported energy [9]. Some more recent contributions use machine learning techniques to automatically learn effective job scheduling policies by continually adapting them to the data centers' complex dynamic environment (computing and renewable energy resources) [10], while others extend the scheduling policy exploration to service tasks so as to consume more on-site renewable energy [11].

Several Cloud providers expand their data centers over regions and countries to be close to their customers and for sake of fault tolerance. Each site delivers a local service and may participate in a large ecosystem of data centers via electrical and/or telecommunication interconnections. In non-interconnected or weakly connected electrical networks, such Cloud providers can exploit the temporal variations in on-site productions by routing the load to a data center with available computational resources and renewable electricity. However, for a well interconnected electrical network, energy-aware load migration does not increase the overall renewable energy consumption since it would be consumed in the electrical grid, but allows to qualify this consumption as self-consumption which benefits from economic incentives in numerous countries. For instance, B. Camus *et al.* [12] proposed algorithms that take advantage of this cooperation between data centers to synchronize their on-site photovoltaic production and their consumption by migrating VMs to the most abundant sites and exchanging green energy between sites. Yuan *et al.* [13] formulated a bi-objective optimization problem for distributed green data centers to maximize the profit of service providers and minimize the average task loss possibility of all applications by jointly determining the split of tasks and service rates among multiple data centers.

The state of the art focuses either on proposing algorithms for workload scheduling within a data center or for load balancing among distributed data centers. In this work, we consider a single data center and explore the achievable limits in this case for on-site consumption. We express these limits in terms of electrical grid dependency and autonomy financial costs. To obtain limits independent of specifically tweaked scheduling algorithms, we replay the scheduling policy of real workloads (i.e. from Google and video streaming servers) and do not consider flexing it other than done in actual data centers.

2.2 Flexibility of the electrical infrastructure

As service tasks cannot be synchronized with renewable energy production and since batch jobs can only be shifted to a certain extent, there is a need for flexibility on the power supply side. This is the purpose of integrating Energy Storage Systems (ESSs) to the electrical infrastructure.

ESSs were first used as Uninterruptible Power Supply (UPS), generating power for the data center only when there is grid failure and remaining idle when the grid supply is available. They are sized to the capacity of the data center: a 10MW data center requires a 10MW UPS with several minutes of energy capacity [14]. But in recent years, UPSs also perform other ancillary functions such as load shaving

when the power demand exceeds the power supply limit (i.e. the contract power may be exceeded, thus leading to penalties) [15], [16]. They are exploited as well to provide frequency reserves and voltage control to the grid [17], [18].

ESSs are also adopted to store renewable energy for long period use. Hence, the data centers consume more renewable energy, and become more independent from the grid [19]. For instance, in the green data center prototype Parasol [20] whose power supply is composed of a set of solar panels, a battery bank, and a grid-tie, Í. Goiri et al. experimented a scheduling approach called GreenSwitch that dynamically manages the workload and the energy sources. This combination of supply-side and demand-side mechanisms optimizes the overall cost of electricity, the battery usage, and the quality of service.

The above-mentioned approaches involve an energy mix including potentially non-renewable energy and require the data center to consume electricity from the electrical network. More and more research work focus on autonomous data centers. In 2019, Datazero [21] proposed a data center powered by 100% renewable energy without any contribution from the grid. The renewable power supply is composed of photovoltaic panels and wind turbines and is associated with batteries and an hydrogen storage system [22]. Both the electrical and the computation parts of the data center have independent optimization mechanisms and a negotiation protocol is introduced to match the power consumption with the power generation. In the same direction, N. Lazaar et al. [23] proposed an energy management and sizing strategy for a 1MW standalone green data center powered by a hybrid Tidal and Photovoltaic system associated with a battery and an hydrogen storage system. However, these studies do not provide a cost estimation of such autonomous data centers electrical systems.

Running on 100% renewable energies only targets the scope 1 and 2 of GHG emissions, i.e. the direct emissions and indirect emissions from electricity use [24]. Scope 3 of GHG, including the manufacturing, the outsourced transportation, the post-life phase of each component (servers, routers, etc.), as well as employees commuting and comfort devices, business travels, waste disposal, etc., is usually the largest one in terms of GHG emissions for data centers. This scope is still too little addressed by data center operators, and given the large amount of unknown figures for estimating it, it was considered out of the scope of this paper.

Although all these initiatives seek to capitalize on renewable energy use to reduce pollution in data centers, the optimal sizes of the sources to be used are not investigated. In fact, the survey [25] states about the sizing that *"this problem is not new and has long been studied in the literature for microgrids and similar systems, mostly for coupling renewable energy sources with batteries, hydrogen and other means of storage. However, few works so far have focused on data center applications, especially when considering the specific characteristics of IT load and availability requirements"*. Moreover, the studies using the grid as a source of flexibility do not assess to which extent the data centers are consuming from the grid or any inherent challenge to the grid. Thus, we consider a configuration called grid-dependent where the data center uses only the electrical network in order to analyze these aspects. The survey [25] also found that *"only two works*

consider powering the data center with only renewable energy sources", but none of them tested different configurations in order to evaluate their cost.

In this paper, we propose an economical study of so-called autonomous infrastructures supplying data centers with renewable energies associated to storage devices unconnected to the grid. We consider ESSs used for storing renewable energy, targeting a 100% level of autonomy for datacenters. We consider different configurations (i.e. without storage, and with several types of storage units), and we assess their economical costs for several data center sizes and different workload types. In addition, contrary to previous work, we consider the difference in storage devices' dynamics and their degradation. Hence, to the best of our knowledge, our study is the first providing together 1) indicative optimal sizes of several renewable-energy based electrical systems for various data center scales, 2) a quantification of the grid level of contribution to their consumption (when the data centers consume from the grid), and 3) a quantification of 100% autonomy cost (when the data centers are powered exclusively by the association renewable energy - storage system) for various workload profiles and electrical configurations.

3 SYSTEM ARCHITECTURE

3.1 Data center infrastructure

We consider data centers made of a *task manager*, *compute nodes* (servers) linked together by an *intra-data center network* of homogeneous switches and to the external world by a *telecommunication network*. The task manager is an interface between the data center and its users. It receives some requests from the users and schedules the corresponding jobs on the nodes with available resources (CPU and RAM mainly), by providing the starting and stop time, as well as the required resources. The compute nodes are responsible for executing the jobs received from the manager. Every task instance runs in a virtual machine according to the manager's scheduling. In the data centers of 100 compute nodes and more, the network is a n-ary tree topology with three layers (core, aggregation and edge) [26]. Let n be the number of ports on each switch. The topology can handle up to $\frac{n^3}{4}$ compute nodes and requires $\frac{n^2}{4} + n$ switches. As for small scale data centers (less than 100 compute nodes), they are based on a light tree topology empirically designed and the associated number of switches is precised hereafter.

3.2 Electrical infrastructure

Solar energy is seen today, and particularly in this work, as one of the convenient and safe alternatives to traditional energy sources, due to its sustainability, its attractive ease of installation – no nimbies issue as encountered for wind turbines for example – and its less greedy spatial requirement (most of the time, it is rooftop). Just like its expansion over the last decades, the capacity of photovoltaic (PV) generation increased by 23% in 2020 (representing the second-largest generation growth of all renewable technologies) [27], and by 19% in 2021 (leading the renewable sources) [28]. However, solar energy generation is time and weather related and needs to be associated with a second

more controllable source that may be the electrical grid and/or energy storage devices.

Our goal when examining the different configurations presented in this section is twofold: 1) finding the minimum sizing of the different energy sources in order to avoid load shedding, 2) evaluating the operation cost of these infrastructures in order to analyze their economical viability.

3.2.1 Grid-dependent infrastructure

Cloud operators like Apple have two ways to cover their annual needs at 100% from renewables. On one hand, since 2011, Apple has invested in the creation of renewable energy plants around the world (in the data-center premises or elsewhere) and has reached a cumulative installed power capacity of 1,524 MW (77.2% from solar source, 21.6% from wind turbines, 0.22 % from micro-hydro, and 0.92 % from biogas) in 2020 [29]. In the other hand, as in Nevada, Apple delegates the process of production (not collocated with the data centers), maintenance and management to third party companies. This way of proceeding, known as Power Purchase Agreement (PPA), provided a power capacity of 32MW to Apple data centers in 2020 [29].

Based on this case, we consider a data center with an on-site photovoltaic source and connected to the utility grid, with whom it exchanges some energy, as shown in Figure 1. This electrical system does not include any ESS and uses the grid as a virtual storage in which the surplus of production is injected and retrieved when the renewable production is not sufficient to meet the data center consumption. We designate this infrastructure as "grid-dependent".

Let's name $P_{Gex}(t)$ and $P_{Gim}(t)$ respectively the instantaneous power exported to and imported from the grid at time t . Let's also denote $P_{PV}(t)$ the power generation of the PV plant and $D(t)$ the power consumption of the data center at time t . The electrical infrastructure satisfies the power balance presented in Equation 1.

$$P_{Gim}(t) + P_{PV}(t) = D(t) + P_{Gex}(t) \quad (1)$$

The optimal sizing for this configuration corresponds to the smallest PV plant that allows the annual power generation to cover the data center annual demand.




Figure 1: Grid-dependent data center

3.2.2 Autonomous data center infrastructures

In an autonomous mode, a data center is powered throughout the year by a combination of renewable energy and energy storage sources without resorting to the electrical grid. Figure 2 shows two autonomous infrastructures: one with an electrochemical battery and the second with both a battery and an hydrogen ESS.

Energy Management Strategy of the autonomous infrastructure using only a BESS

In this case, a Battery Energy Storage System (BESS) is associated with the renewable energy source. After powering the data center, any surplus of PV production is used to charge the BESS with an amount of power $P_{ch}(t)$. When the BESS is full, the remaining surplus $P_{curt}(t)$ is curtailed if any. In the case of insufficient production, the battery is discharged with the power $P_{disch}(t)$ in order to meet the load. However, due to physical restrictions, the battery cannot be simultaneously in charge and discharge modes. The instantaneous power exchanged is balanced as in Eq. 2.

$$P_{disch}(t) + P_{PV}(t) = D(t) + P_{ch}(t) + P_{curt}(t) \quad (2)$$

Energy Management Strategy of the autonomous infrastructure using a BESS and a HESS

Batteries are found unreasonably expensive for inter-seasonal storage applications, for which a combination with hydrogen-based storage systems is usually considered. In the third case study, we associate the BESS with a Hydrogen Energy Storage System (HESS) that is cheaper and has a better energy density. A HESS involves mainly an electrolyzer, a tank and a fuel cell. The electrolyzer converts electrical energy into chemical energy by decomposing water into hydrogen (H_2) and oxygen (O_2). The resulting H_2 is pressurized by a compressor (associated with the tank) and injected in the hydrogen tank for storage. The fuel cell carries out the reverse process of electrolysis and uses hydrogen and oxygen or air to generate electrical power.

The BESS plays two roles in this infrastructure: it ensures a daily storage and ensures the power balance for short-term fluctuations. As for the HESS, it is dedicated to inter-seasonal (annual) storage and is subject to the following constraints. To prevent the electrolyzer (respectively fuel cell) from starting repetitively for short slot operations, we impose it to operate for a minimum duration H (Eq. 3a). HESS equipment have slower dynamics compared to the BESS, so they operate with fixed power for a predefined duration H_0 before being able to change (Eq. 3b). The efficiency of the HESS is relatively low when it operates with a low power. Therefore, we impose a power threshold P_0 above which the HESS equipment operates to ensure efficient operation (Eq. 3d). Lower power threshold P_0 was set to 30% of the nominal power like in other work [30]. Finally, we limit the number of on/off switchings per day

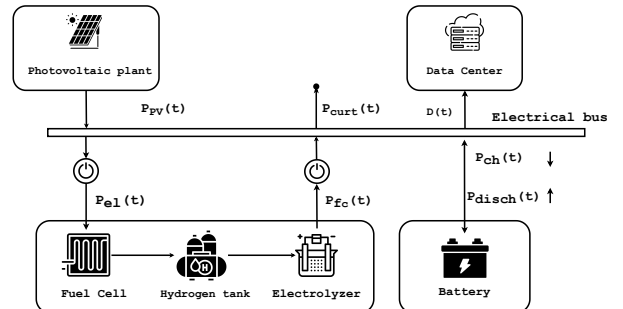


Figure 2: Autonomous data center

(Eq. 3c). Also, the fuel cell and the electrolyzer are not allowed to operate simultaneously (Eq. 3e).

$$P_i(t) \neq 0 \text{ if } \delta_i(t_{on}) = 1 \text{ for } t_{on} \leq t \leq t_{on} + H \quad (3a)$$

$$P_i(t) = P_i(t-1) = P_i(t_0) \neq 0 \text{ if } \delta_i(t) = 1 \text{ for } t_0 + \Delta t \leq t \leq t_0 + H_0 \quad (3b)$$

$$\sum \delta_{switch}^i(t) \leq \lambda_i; \text{ With } t_{k-1} \leq t \leq t_k \text{ and } 1 \leq k \leq 365 \quad (3c)$$

$$P_i(t) \geq P_0 \text{ if } \delta_i(t) = 1 \quad (3d) \quad \delta_{fc}(t) + \delta_{el}(t) \leq 1 \quad (3e)$$

Where $i \in \{fc, el\}$ are the fuel cell and electrolyzer, t_{on} the starting time, $P_i(t)$ the operating power of i , t_0 the first time that the current operating power has been set, λ_i the number of on/off cycles allowed per day for i , t_{k-1} (respectively t_k) the beginning (respectively the end) of day k , $\delta_i(t)$ a binary parameter that equals to 1 when the element i is on at time t and 0 otherwise, and $\delta_{switch}^i(t)$ a binary value that equals to 1 when there is a startup or a switch off and 0 otherwise. $\delta_{switch}^i(t) = |\delta_i(t) - \delta_i(t - \Delta t)|$.

The energy management of the infrastructure is based on heuristics and contains two phases:

Phase 1: The BESS has the priority. It is the default energy management mode and corresponds to when the HESS is off. Hence, when the renewable power generation is in excess, we first charge the battery. When it is full, the electrolyzer can start if and only if the remaining power (if any) is above P_0 . In this situation, the electrolyzer starts and the system switches to Phase 2, otherwise we curtail the rest. In the opposite case where the PV power generation is not sufficient to power the data center, the battery is discharged to meet the power demand. We define a security state-of-charge margin (SoC_{sm}) to avoid too deep discharges of the battery. Once the battery state-of-charge reaches this margin, the fuel cell starts and the system switches to Phase 2.

Phase 2: The HESS has the priority. When the hydrogen storage system is active, the surplus and deficits are absorbed by the HESS and the power balance is ensured, for rapid fluctuations, by the BESS. The HESS operates with a constant power for a time duration of H_0 after which the power may change (for another duration of H_0) according to the data center needs and the PV generation. When the electrolyzer is active and the tank is full, we curtail the excess of power and switch off the electrolyzer. The electrolyzer can also be switched off if the allocated time exceeds H and the battery is not full while the solar system is still generating in excess. As for the fuel cell, we switch it off when the allocated time H is exceeded, the power generation is enough to supply the data center and the state-of-charge of the battery is above the margin (SoC_{sm}). Turning the fuel cell or the electrolyzer off causes the system to switch back to the Phase 1.

Let's name $P_{curt}(t)$ the power curtailed after the battery and/or the tank is full, $P_{fc}(t)$ and $P_{el}(t)$ respectively the operating power of the fuel cell and the electrolyzer. The infrastructure satisfies the power balance in Equation 4.

$$P_{disch}(t) + P_{PV}(t) + P_{fc}(t) = D(t) + P_{ch}(t) + P_{el}(t) + P_{curt}(t) \quad (4)$$

4 ENERGY MODELS

4.1 Energy Consumption of the data centers

Energy consumption is attributable to three main components in data centers: the compute nodes, the intra-data center network and the infrastructure including the cooling

system. The telecommunication network employed outside the data center (i.e. in the Internet) is not included in this study as we focus on the data centers itself.

As showed by Heinrich *et al.* [31], one identifies two terms in the instantaneous power consumption of a compute node (or server): a static term that represents the consumption when the server is turned on but not running any load (it includes the static power consumption of all the server's components like its motherboard, network cards and disks), and a dynamic one that linearly depends on the CPU frequency and the nature of the computational workload. It is true for constant workloads, but remains applicable to non-constant workloads like applications containing several phases of different nature (by independently characterizing each phase) and can be used with Cloud data center traces that provide CPU usage at intervals of time.

After determining the number of ports per switch and the number of switches needed in a data center as described in Section 3.1, we proceed to the identification of corresponding vendor hardware and estimated their consumption using the Cisco Power Calculator [32]. We assume that the switches operate at nominal power. This is slight overestimation since the static consumption of a network switch (without any traffic) is higher than 80% of its nominal power [33]. Thus, the power consumed by the intra-data center network is assumed to be the nominal power of a single switch multiplied by the number of switches in the topology. When the calculated number of ports n does not correspond to any manufactured switch, we used as follows the nearest switches able to satisfy the network topology, while keeping unused links. Let's define n_1 (respectively n_2) so that n_1 (respectively n_2) is the number of hardware ports immediately smaller (respectively greater) than n . For instance $n_1 = 12$, $n_2 = 16$ and $n = 13$ means that no switch of 13 ports exists but we can find switches of 12 and 16 ports in the Cisco catalog. The power consumption when using n ports is evaluated by a linear regression between the two closest existing switches as expressed in Equation 5 where $P(x)$ is the power consumption of a x -ports switch.

$$P(n) = P(n_1) + \frac{P(n_2) - P(n_1)}{n_2 - n_1} \times (n - n_1) \quad (5)$$

As for considering the infrastructure, Power Usage Effectiveness (PUE) is the metric used to account for the energy consumption of the other equipment available in the data center such as the cooling system, the facilities lighting, etc. It represents the ratio between the total facility consumption and the IT power consumption (i.e. consumption of the servers and the switches). Thus, the overall data center's power consumption can be estimated by multiplying the cumulative consumption of the compute nodes and switches by its PUE value.

4.2 Photovoltaic power production

PV systems are designed to convert sunlight into electrical energy through PV cells arranged on a panel. The power generated by a PV panel is the product of its surface, its efficiency of conversion and the solar irradiance [34]. For sake of simplicity, we consider an homogeneous PV plant and assume that the solar irradiance is equally distributed on it. We also neglect the aging effect on the panels performance.

4.3 Battery Energy Storage System

In this study, we consider a Lithium Ion (LI) battery which is the most suitable technology for the type of application considered in the Cloud as it presents the highest cycling rate and energy density. A BESS is represented by its instantaneous state-of-charge (SoC) and its state of health (SOH). Whereas the former designates the amount of energy available in the battery relative to its capacity, the latter represents its state of degradation for having exchanged some energy. Let's consider a BESS whose SoC and capacity at the beginning of the time t of duration Δt are respectively $SoC(t - \Delta t)$ and $C_{Bat}(t - \Delta t)$. Its instantaneous state-of-charge at the end of the time slot is modeled as in Equation 6, neglecting the self-discharge effect compared to the controllable energy flow.

$$\begin{aligned} SoC(t) &= SoC(t - \Delta t) + \frac{\Delta t}{C_{Bat}(t - \Delta t)} \cdot \left(\eta_{ch} \cdot P_{ch}(t) - \frac{P_{disch}(t)}{\eta_{disch}} \right) \\ 0 &\leq P_{ch}(t) \leq P_{max}^{Bat} \text{ and } -P_{min}^{Bat} \leq -P_{disch}(t) \leq 0 \quad (6) \\ SoC_{min} &\leq SoC(t) \leq SoC_{max} \end{aligned}$$

Where, η_{ch} and η_{disch} are respectively the efficiency of charge and of discharge, P_{max}^{Bat} and P_{min}^{Bat} are respectively the nominal power of charge and discharge. We consider a battery with $P_{min}^{Bat} = P_{max}^{Bat}$. The state-of-charge is a normalized value between 0 and 1 and is bound by SoC_{min} and SoC_{max} .

The chemical constituents of a BESS deteriorate as it operates, and that is visible by the decrease of its capacity. In this model, we neglect calendar aging (due to time only) compared to cycling aging (due to operating the battery) which is assumed greater considering the BESS use here. We design a model in which each amount of energy exchanged during charge/discharge operations represent a linear drop in the battery's capacity (Eq. 7). It is assumed that the battery can exchange during its life time N_{cycles} (number of charge/discharge cycles that it can undergo), a total amount of energy that equals to $2 \cdot DOD \cdot N_{cycles} \cdot C_{Bat}(0)$.

$$\begin{aligned} C_{Bat}(t) &= C_{Bat}(t - \Delta t) \\ &- \frac{C_{Bat}(0) - C_{Bat}(t_{EOL}^{Bat})}{2 \cdot DOD \cdot N_{cycles} \cdot C_{Bat}(0)} \cdot \left(P_{ch}(t) \cdot \eta_{ch} - \frac{P_{disch}(t)}{\eta_{disch}} \right) \cdot \Delta t \quad (7) \end{aligned}$$

Where DOD (depth of discharge) refers to the amount of energy flowing in and out of the battery over a cycle. It is represented as a percentage of the total capacity. $C_{Bat}(0)$ and $C_{Bat}(t_{EOL}^{Bat})$ are respectively the initial capacity of the battery and its remaining value at the end of life, t_{EOL}^{Bat} is the time when the battery reaches its end of life.

SOH represents the capacity loss due to the aging process (Eq. 8). It is a normalized value that equals to 1 for a new battery and 0 when the battery lifetime has been reached. In this study, we consider a BESS with the rated energy capacity of a new BESS at the beginning of the simulations but the model is applicable to a second life BESS. Let's name β_{Bat} the percentage of the BESS's capacity remaining at the end of its life. The capacity at end of life is expressed as $C_{Bat}(t_{EOL}) = \beta_{Bat} \cdot C_{Bat}(0)$.

$$\begin{aligned} SOH_{Bat}(t) &= SOH_{Bat}(t - \Delta t) - \frac{C_{Bat}(t - \Delta t) - C_{Bat}(t)}{C_{Bat}(0) - C_{Bat}(t_{EOL})} \quad (8) \\ 0 &\leq SOH(t) \leq 1 \end{aligned}$$

To estimate the spatial size of the BESS, we compute its volume by referencing to a battery used in the real-world

data centers: the EnerSys 12HX505+ model which nominal dimensions are 338x173x2 mm (length x width x height) for a capacity of 1,428 Wh [35]. We assume that the BESS is designed by assembling modules of this reference type.

4.4 Hydrogen Energy Storage System

We choose a Solid Polymer Electrolyte (SPE) electrolyzer and a Polymer Electrolyte Membrane (PEM) fuel cell among the multiple types of electrolyzers and fuel cells, as in previous work from the literature [36]. Under fixed temperature and pressure, and given the hydrogen mass density ρ (33 kWh/kg [22]), the level of hydrogen (LOH) representing the mass of hydrogen in the tank is given by Equation 9.

$$\begin{aligned} LOH(t) &= LOH(t - \Delta t) + \left(\frac{\eta_{el} \cdot P_{el}(t)}{\rho} - \frac{P_{fc}(t)}{\eta_{fc} \cdot \rho} \right) \cdot \Delta t \\ 0 &\leq P_{el}(t) \leq P_{max}^{el}(t) \text{ and } 0 \leq P_{fc}(t) \leq P_{max}^{fc}(t) \quad (9) \\ LOH_{min} &\leq LOH(t) \leq LOH_{max} \end{aligned}$$

Where η_{el} and η_{fc} are respectively the fixed efficiency of the electrolyzer and the fuel cell, $P_{max}^{el}(t)$ and $P_{max}^{fc}(t)$ are their rated power, LOH_{min} and LOH_{max} are the minimum and maximum mass of hydrogen to keep in the tank.

An electrolyzer (respectively a fuel cell) operates for a fixed amount of time N_{unit}^{el} (respectively N_{unit}^{fc}) over its lifetime, at rated operating power. Two types of operation contribute to the reduction of the lifetime: the start/shutdown cycles and the continuous operation. A fuel cell manual [37] states that one start/stop cycle from zero to its nominal power is equivalent to 3 hours of continuous operation. For sake of generality, let's name α_{fc} (respectively α_{el}) the time that a fuel cell (respectively an electrolyzer) loses when it starts and switches off, at rated power. Because the fuel cell (respectively electrolyzer) is not always operating at nominal power, we need to approximate the time lost by starting, operating and shutting down at any power $P_{fc}(t)$ (respectively $P_{el}(t)$).

We consider the actual amount of life reduction to be proportional to the actual power as in [38] and neglect the transient phase in comparison with the time step Δt . Thus, the startup and shutdown are considered instantaneous. Equation 10 models the lifetime reduction during the time slot Δt . The first term is the effect of continuous operation, and the second represents the effect of the switching mode of operation.

$$time_{lost}^i(t) = \frac{P_i(t)}{P_{max}^i(t - \Delta t)} \cdot \left(\Delta t + \delta_{switch}^i(t) \cdot \frac{\alpha_i}{2} \right) ; i \in \{fc, el\} \quad (10)$$

Where $P_{max}^i(t - \Delta t)$ is the rated power of component i at the beginning of the time slot.

The degradation phenomenon results in the increase of the electrolyzer nominal power and the decreases of the fuel cell's life [39]. Let's define β_{el} and β_{fc} as the portion of nominal power remaining at the end of the electrolyzer and the fuel cell life respectively. We consider a linear degradation and neglect the calendar aging over the operational aging (Eq. 11).

$$P_{max}^i(t) = P_{max}^i(t - \Delta t) - \frac{P_{max}^i(0) - P_{max}^i(t_{EOL}^i)}{N_{unit}^i} \cdot time_{lost}^i(t) \quad (11)$$

In this model, $i \in \{fc, el\}$, $P_{max}^i(t)$ is the rated power at time t , t_{EOL}^i is the time when i reaches its end of life, $P_{max}^i(t_{EOL}^i)$ is the rated power of component i at the end

of its life and can be calculated as follows : $P_{\max}^i(t_{EOL}^i) = \beta_i \cdot P_{\max}^i(0)$.

We also define the state of health of the fuel cell (SOH_{fc}) and that of the electrolyzer (SOH_{el}) as normalized parameters that depict the performance loss with aging. Similar to the BESS, SOH_i is equal to 1 for a new component i and gets to 0 at its end of life. We assume that the hydrogen storage system is new at the beginning of the simulation. The SOH at any time t is shown in Equation 12.

$$SOH_i(t) = SOH_i(t - \Delta t) - \frac{|P_{\max}^i(t) - P_{\max}^i(t - \Delta t)|}{|P_{\max}^i(t_{EOL}^i) - P_{\max}^i(0)|} \quad (12)$$

We use the H-TEC PEM Electrolyzer ME450 [40] which dimensions are 13.2x4.0x5.7m for a nominal power of 1MW, as reference to estimate the volume of the electrolyzer. In the same approach, we use FP-100iH fuel cell [41] which dimensions are 5.5x2.2x3.4m for a nominal power of 99 kW as a reference. Let's name P and T the respective absolute pressure (N/m^2) and the temperature (K) in the tank, V the volume (m^3) of the tank, m the mass/LOH (kg) of hydrogen in it, R the specific gas constant (J/kg.K) that equals to 4,116 J/kg.K [42]. The volume of the tank is estimated based on the law of perfect gases ($P.V = m.R.T$). In fact, hydrogen can be considered as an ideal gas over a wide temperature range and even at high pressures [36]. These conditions are met in the tank. The overall volume occupied by the HESS is the sum of the individual volumes of its components.

4.5 Economic modeling

The global cost of an energy source falls into two categories : *costs unaffected by the energy management strategy* and *costs impacted by the management strategy* [43]. For sake of generality, the economic models include a PV plant, a BESS and a HESS. In the case studies where some of these elements are not present, their cost is made equal to zero.

4.5.1 Costs unaffected by energy management strategy

These costs are fixed and the energy management strategy has no impact on them. There are two terms composing them: an *investment cost* and a *fixed operation and maintenance cost*. For the PV plant and the hydrogen tank, we consider that the costs are linked to the calendar aging and that the initial investment is spread over the equipment fixed lifetime. The investment cost $\Delta cost_{inv}(t)$ (\$) incurred by the infrastructure during slot Δt is expressed in Equation 13.

$$\Delta cost_{inv}(t) = \left(\frac{N_{PV} \cdot P_{\max}^{PV} \cdot c_{inv}^{PV}}{N_{unit}^{PV}} + \frac{c_{inv}^{tank} \cdot LOH_{\max}}{N_{unit}^{tank}} \right) \cdot \Delta t \quad (13)$$

Where P_{\max}^{PV} is the nominal power generated by a PV panel, c_{inv}^{PV} (\$/kW) and c_{inv}^{tank} (\$/kg) are respectively the initial investment unit cost in a photovoltaic panel and the tank, N_{unit}^{PV} and N_{unit}^{tank} are respectively the life span of the panel and that of the tank (in units of time).

The fixed operational and maintenance (O&M) costs are periodical fees – such as insurances, property taxes, cleaning and repairing the component, land lease, salaries, etc – paid for the ownership of a component. They are computed by multiplying the fixed O&M unit cost by the nominal power/energy of the considered component. Equation 14

expresses the overall fixed O&M cost $\Delta cost_{O\&Mfix}(t)$ generated within Δt . In this model, the O&M costs of the tank are counted in the O&M of the fuel cell and electrolyzer.

$$\begin{aligned} \Delta cost_{O\&Mfix}(t) = & N_{PV} \cdot P_{\max}^{PV} \cdot c_{O\&Mfix}^{PV} \cdot \Delta t \\ & + c_{O\&Mfix}^{Bat} \cdot C_{Bat}(0) \cdot \Delta t \\ & + \sum_{i \in \{el, fc\}} c_{O\&Mfix}^i \cdot P_{\max}^i(0) \cdot \Delta t \end{aligned} \quad (14)$$

Where for $j \in \{el, fc, PV\}$, $c_{O\&Mfix}^j$ (\$/kW.unit time) is the fixed O&M unit cost in j , $c_{O\&Mfix}^{Bat}$ (\$/kWh.unit time) is the fixed O&M unit cost in the BESS. Generally, the O&M cost is provided on a periodical (annual/monthly) basis, that is divided by the number of time units in the period to obtain the unit cost.

4.5.2 Costs impacted by the energy management strategy

These costs are operating costs and concern the battery, the fuel cell and the electrolyzer. We identify a *degradation cost* and a *variable operation and maintenance cost*. The initial investment of these three components is spread over their lifetimes and each portion of health drop represents a unit of degradation cost. The overall degradation cost $\Delta cost_{degr}(t)$ (\$) generated during Δt is expressed in Equation 15.

$$\begin{aligned} \Delta cost_{degr}(t) = & \frac{c_{inv}^{Bat}}{2 \cdot DOD \cdot N_{cycles}} \cdot \left(\eta_{ch} \cdot P_{ch}(t) + \frac{P_{disch}(t)}{\eta_{disch}} \right) \cdot \Delta t \\ & + \sum_{i \in \{el, fc\}} \frac{c_{inv}^i \cdot P_{\max}^i(0)}{N_{unit}^i} \cdot time_{lost}^i(t) \end{aligned} \quad (15)$$

Where c_{inv}^{Bat} (\$/kWh) is the initial investment unit cost in the battery, c_{inv}^i (\$/kW), for $i \in \{el, fc\}$, is the initial investment unit cost in the fuel cell and the electrolyzer. N_{unit}^i is the lifetime of component i expressed in time units.

The variable O&M costs of an electrical component is the product of its unit variable O&M cost and the amount of energy it consumes or generates. The overall variable O&M cost $\Delta cost_{O\&Mvar}(t)$ (\$) generated during a slot of duration Δt is expressed in Equation 16. The solar panels do not include variable O&M cost.

$$\begin{aligned} \Delta cost_{O\&Mvar}(t) = & c_{O\&Mvar}^{Bat} \cdot \left(\eta_{ch} \cdot P_{ch}(t) + \frac{P_{disch}(t)}{\eta_{disch}} \right) \cdot \Delta t \\ & + \sum_{i \in \{el, fc\}} \delta_i(t) \cdot c_{O\&Mvar}^i \cdot P_i(t) \cdot \Delta t \end{aligned} \quad (16)$$

Where $c_{O\&Mvar}^{Bat}$ (\$/kWh) is the unit variable O&M cost in the battery, $c_{O\&Mvar}^i$ (\$/kW) – $i \in \{el, fc\}$ – is the variable O&M cost in the fuel cell and the electrolyzer.

4.5.3 Total cost function

We assume that an equipment is replaced at the same purchase price when its lifetime is over. However, a dynamic price could be considered to reflect the reality of the market without any prejudice to the validity of the economic methodology nor the models. The overall cost, $cost(t)$ (\$), of the electrical system at time t is the addition of all the costs incurred from $t = 0$ to that time (Eq. 17 with $l \in \{inv, O\&Mfix, O\&Mvar, degr\}$).

$$cost(t) = cost(t - \Delta t) + \sum_l \Delta cost_l(t) = \sum_l \sum_{\tau=0}^t \Delta cost_l(\tau) \quad (17)$$

5 METHODOLOGY

In this paper, we present three case studies where renewable energy covers 100% of data centers' energy needs. This section describes the infrastructures' sizing methods in each case. This methodology is applied to several sizes of data centers characterized by their number of compute nodes.

5.1 Sizing the grid-dependent infrastructure

The goal is to find the number of panels (N_{PV}) and the area occupied by the PV plant, so that the annual energy generation balances the data center energy demand. The number of panels is computed by dividing the annual consumption of the data center by the annual production of a single PV panel (Eq. 18).

$$N_{PV} = \left\lceil \frac{\sum_0^{t_{yr}} D(t)}{\sum_0^{t_{yr}} P_{1PV}(t)} \right\rceil \quad (18)$$

Where t_{yr} is the end of the year and $P_{1PV}(t)$ is the instantaneous power production of one solar panel. Let's name A and A_{total} respectively the surface of a single solar panel and the surface needed to install the solar plant. $A_{total} = N_{PV} \cdot A$. This formula is also applicable to the autonomous case studies.

5.2 Sizing the autonomous infrastructure with a BESS

The objective is to find the number of panels (N_{PV}), the total surface of the solar plant (A_{total}), the capacity of the battery ($C_{Bat}(0)$), its initial SoC ($SoC(0)$) and its nominal power (P_{Bat}^{max}) in order to meet the annual demand. Any eligible BESS should ensure the power supply all along the year (Eq. 19a) and contain at least the same amount of energy at the end of the year as at beginning (Eq. 19b).

$$SoC_{min} \leq \min SoC(t) \leq \max SoC(t) \leq SoC_{max} \quad (19a)$$

$$SoC(t_{yr}) \geq SoC(0) \quad (19b)$$

To do so, let's consider arbitrarily a battery presenting an important capacity C_0 with an initial level of energy SoC_0 that satisfies Equation 19a. The first step is to size the primary source. To size the PV plant, we use a dichotomy search to find the smallest candidate N_{PV} that meets the battery requirements (Eq. 19). The search space range is bounded by N_{PV}^{min} – the PV plant generates the same amount of power as the data center consumption – and N_{PV}^{max} – the generated power is stored into the BESS and is then used to supply the data center – as shown in Equation 20.

$$\underbrace{\left\lceil \frac{\sum_0^{t_{yr}} D(t)}{\sum_0^{t_{yr}} P_{1PV}(t)} \right\rceil}_{N_{PV}^{min}} \leq N_{PV} \leq \underbrace{\left\lceil \frac{\sum_0^{t_{yr}} D(t)}{\eta_{ch} \cdot \eta_{disch} \cdot \sum_0^{t_{yr}} P_{1PV}(t)} \right\rceil}_{N_{PV}^{max}} \quad (20)$$

The second step is to find the smallest $C_{Bat}(0)$ and the exact $SoC(0)$ that ensures the power supply (19a) and $SoC(0) = SoC(t_{yr})$. They can be deduced by homothety (Eq 21).

$$\begin{aligned} C_{Bat}(0) &= \frac{SoC_M - SoC(t_{yr}) + SoC_0 - SoC_m}{SoC_{max} - SoC_{min}} \cdot C_0 \\ SoC(0) &= SoC_{min} + \frac{(SoC_{max} - SoC_{min}) \cdot (SoC_0 - SoC_m)}{SoC_M - SoC(t_{yr}) + SoC_0 - SoC_m} \quad (21) \\ \text{With } SoC_m &= \min_{0 \leq t \leq t_{yr}} SoC(t) \text{ and } SoC_M = \max_{0 \leq t \leq t_{yr}} SoC(t) \end{aligned}$$

The last step consists in choosing a nominal power for the battery according to the following two criteria. First, the battery must satisfy the entire power demand of the data center (Eq. 22a). Secondly, its capacity and its nominal power should be consistent, as a small battery might satisfy the energy requirement but not be able to deliver the required nominal power. We use a Ragone diagram that describes the relationship between the energy density and power density of diverse storage units including LI batteries [44]. Let's define ρ_e^{min} and ρ_e^{max} as respectively the minimum and maximum energy density of a battery, and ρ_p^{min} and ρ_p^{max} its respective minimum and maximum power density. Equation 22 describes the constraints imposed on the BESS nominal power.

$$\max(\max P_{ch}(t), \max P_{disch}(t)) \leq P_{max}^{Bat} \quad (22a)$$

$$\frac{C_{Bat}(0) \cdot \rho_p^{min}}{\rho_e^{max}} \leq P_{max}^{Bat} \leq \frac{C_{Bat}(0) \cdot \rho_p^{max}}{\rho_e^{min}} \quad (22b)$$

Algorithm 1, performing with a complexity of $O(\log_2(N_{PV}^{max} - N_{PV}^{min}))$, summarizes the sizing process.

Algorithm 1: Algorithm for sizing battery and PV

Input: $N_{PV}^{min}, N_{PV}^{max}, C_0, SoC_0$

Output: $N_{PV}, C_{Bat}(0), P_{max}^{Bat}, SoC(0)$

OFFGRID-PV-BESS-SIZING ($N_{PV}^{min}, N_{PV}^{max}$)

Find by dichotomy the smallest N_{PV} satisfying the constraints (Eq. 19) and report SoC_m, SoC_M and $SoC(t_{yr})$

Calculate $SoC(0)$ and $C_{Bat}(0)$ using SoC_m, SoC_M and $SoC(t_{yr})$ as in (Eq. 21)

Choose the smallest P_{max}^{Bat} that meets (Eq. 22)

End Function

return ($N_{PV}, C_{Bat}(0), P_{max}^{Bat}, SoC(0)$)

5.3 Sizing the autonomous infrastructure with a BESS and a HESS

In this case, we aim to find the number of PV panels (N_{PV}), the overall surface of the solar plant (A_{total}), the capacity of the BESS ($C_{Bat}(0)$), its initial SoC ($SoC(0)$), the security margin (SoC_{sm}) and its nominal power (P_{Bat}^{max}). As for the HESS, we need to find the capacity of the tank (LOH_{max}), the initial quantity of hydrogen ($LOH(0)$), the rated power of the electrolyzer (P_{max}^{el}) and that of the fuel cell (P_{max}^{fc}). We used a brute force method to size $C_{Bat}(0)$, $SoC(0)$, SoC_{sm} , LOH_{max} and $LOH(0)$ in order to minimize the overall operation cost (Eq. 17). P_{Bat}^{max} is chosen to satisfy Equation 22, using the Ragone diagram. Lastly, P_{max}^{fc} and P_{max}^{el} are chosen with regard to Equation 23.

$$\max_{i \in \{fc, el\}} P_i(t) \leq \min_{i \in \{fc, el\}} P_{max}^i(t) \quad (23)$$

6 EVALUATION

6.1 Experimentation setup

The models and strategies described above are implemented on top of the SimGrid simulation toolkit [45]. We used a time step of 5 minutes and performed simulations for one and ten years. The input data are reported in Table 1.

The considered servers are based on the Nova cluster installed in Lyon, available on the Grid'5000 experimental platform [50]. They are equipped with 32 Intel Xeon E5-2620 v4 with 8 cores each, 64GB memory, 598GB HDD, 2.3GHz frequency and an Intel Ethernet 10G Ethernet card. We use external wattmeters to make power consumption measurements of these nodes. They consume 79W at idle state and 145W when all the CPUs are 100% loaded.

We use 2 workloads to observe the influence of variable workloads. *Google Cluster Workload Traces 2019* [51] contains real-world Google jobs spread on 9,993 machines. This trace is collected for 1 month (May 2019) and we assume a monthly periodicity to build an annual workload. We consider the jobs as non flexible and simulate them over homogeneous servers described above. We build a power profile by streaming videos on a VLC server on the nodes and measuring their consumption. For these streaming servers, we use the network traffic (from February 2021 to February 2022) provided by the AMS-IX platform [52] as network output and linearly scale their number according to our needs. This hand-made workload is introduced to analyze the impact of the various types of Cloud applications on the costs and dimensions of the autonomous infrastructures, video-streaming being currently one of the most network-hungry type of traffic. Figure 3 shows the annual power consumption of 4,000 Google servers (top) and 4,000 streaming servers (down). One can observe that Google trace is extremely flat (low variations), while the streaming trace presents a higher variability.

We used real rooftop solar power trace collected over one year (2018), in Austin Texas, and available on the Pecan street website [53]. Figure 4 shows the first week power generation. The specifications of the switches are listed in Table 2, based on Cisco equipment. We consider a PUE value of 1.2 in all cases. As a comparison, Google current data centers have a PUE ranging from 1.09 to 1.12 [54].

6.2 Results

6.2.1 Grid-dependent case study

We consider several sizes (number of compute nodes) of data center running the Google workload. Table 3 shows after sizing that we need on average $6.2 m^2$ to $10.2 m^2$ of solar panels to meet the annual energy demand of a server. Hence, a data center of 8,000 servers requires $49,276 m^2$ that

Table 1: Input parameters

PV panels [46] [47]	$P_{max}^{PV} = 175 \text{ W}$	Lifetime = 25yr	$A = 1.305 m^2$
	$c_{O\&Mfix}^{PV} = 18 \text{ \$/kW-y}$	$c_{inv}^{PV} = 1.56 \text{ \$/W}$	
Battery [48] [39]	$SoC_{min} = 0.2$	$SoC_{max} = 0.8$	$DOD = 0.8$
	$N_{cycle} = 2,000$	$\eta_{disch} = 0.9$	$\eta_{ch} = 0.9$
	$c_{O\&Mfix}^{Bat} = 4.4 \text{ \$/kW-yr}$	$c_{inv}^{Bat} = 448 \text{ \$/kWh}$	$\beta_{Bat} = 0.8$
	$c_{O\&Mvar}^{Bat} = 0.5125 \text{ \$/MWh}$		
Fuel cell [48] [39]	Lifetime = 30,000 h	$\eta_{fc} = 0.35$	$\beta_{fc} = 0.9$
	$c_{O\&Mvar}^{fc} = 0.25625 \text{ \$/MWh}$	$\alpha_{fc} = 3 \text{ h}$	$\lambda_{fc} = 6$
	$c_{O\&Mfix}^{fc} = 14.255 \text{ \$/kW-y}$	$c_{inv}^{fc} = 1,320 \text{ \$/kW}$	
Electrolyzer [48] [39]	Lifetime = 30,000 h	$\eta_{el} = 0.35$	$\beta_{el} = 1.1$
	$c_{O\&Mvar}^{el} = 0.25625 \text{ \$/MWh}$	$c_{inv}^{el} = 1,503 \text{ \$/kW}$	$\lambda_{el} = 6$
	$c_{O\&Mfix}^{el} = 14.255 \text{ \$/kW-y}$	$\alpha_{el} = 3 \text{ h}$	
Tank [49] [36]	$T = 288.15 \text{ K}$	$P = 7.10^7 \text{ N/m}^2$	$c_{inv}^{tank} = 333 \text{ \$/kg H2}$
	Lifetime = 30 yr	$LOH_{min} = 0 \text{ kg}$	




Figure 3: Workload power consumption




Figure 4: Solar trace example for the first week of 2018

Table 2: Intra-data center network characteristics

# Servers	# ports	# switches	Power (W)	Switch reference
10	8	4	124	3560CX-8PC-S
50	12	13	200	C8500-12X
100	8	24	124	3560CX-8PC-S
500	13	56	227	3560CX-8PC-S
1000	16	80	282	S-X6516-GE-TX
4000	26	195	395	S-X6516-GE-TX
8000	32	288	470	C6800-32P10G

is equivalent to 12.2 soccer fields of $45m \times 90m$ and may be challenging to acquire. For each size of data center, we compute the self-consumption rate which describes the portion of the data centers' annual consumption that originates from the on-site PV generation (Eq. 24a). The results show that only 35.9% to 37.7% of the power demand is actually satisfied by the on-site PV plant. In other words, 62.3% to 64.1% of the data center actual consumption comes from the utility grid. This is consistent with the theoretical lower bound of 50%, as PV panels do not generated power half of the time on average during the year, and as the Google traces are relatively flat.

$$1 - \frac{\sum_{t=0}^{t_{yr}} P_{Gim}(t)}{\sum_{t=0}^{t_{yr}} D(t)} \quad (24a) \quad \frac{P_{Gim}(t)}{D(t)} \quad (24b)$$

$$\frac{\sum_{t=0}^{t_{yr}} P_{Gex}(t)}{\sum_{t=0}^{t_{yr}} P_{PV}(t)} \quad (24c) \quad \frac{P_{Gex}(t)}{D(t)} \quad (24d)$$

$$\frac{\max_{t \leq t_{yr}} (P_{Gex}(t))}{D_{rated}(t)} \quad (24e) \quad \frac{\sum_{t=0}^{t_{yr}} cost(t)}{\sum_{t=0}^{t_{yr}} D(t)} \quad (24f)$$

We also introduce a power importation ratio as the ratio of the instantaneous power imported from the grid by the instantaneous power consumed by the data center (Eq. 24b), in order to determine the level of implication of the grid. Figure 5 shows the power importation ratio (in orange) for a data center of 4,000 servers. The data center is very often powered entirely by the electrical grid. It is the same for other sizes of data centers.




Figure 5: Injection and importation ratio for 4,000 servers using Google trace (with a zoom for 10 days)

In addition to that, we evaluate the power injection rate that corresponds to the proportion of the annual PV generation injected into the grid (Eq. 24c). It appears that 62.8% to 64.1% of the on-site PV production needs to be virtually stored in the grid as it is not synchronized with the data center consumption. Therefore, it strongly relies on the electrical network as an external energy source.

Power injection ratio is the ratio between the instantaneous power injected into the grid, and the instantaneous power consumption (Eq. 24d). Figure 5 shows that ratio (in blue) for the data center of 4,000 servers. We can observe a high prevalence of injections that can exceed 5 times the instantaneous consumption of the data center. We identify the peak injection ratio as the ratio of the annual peak power injected into the grid, by the nominal power of the data center (Eq. 24e where $D_{rated}(t)$ is the rated consumption of the data center.). The data centers export peak power ranging from 3.4 (340%) to 4.9 (490%) times their nominal power, to the electrical network for the two workload traces respectively. Regarding the penetration of renewable energy in the energy mix, injecting such amount of renewable production into the grid can be positive because it reduces the production from more polluting sources. However, it has a downside due to the non-synchronization of these high injected peaks with the consumption in the grid. So, in order to absorb this large amount of renewable energy, the electrical network operator must bear dispatching costs linked to other power plants, store any excess – thus shifting the issue and responsibility of investing in energy storage equipment from the data center operator to the grid operator – or activate demand-side flexibility means. In other words, the grid operator must handle alone the challenges linked with a potentially high penetration level of renewables.

6.2.2 Autonomous case study using only a BESS

In this configuration, we consider autonomous data centers of different sizes, using only a battery as energy storage mean. The servers first run the Google workload. Table 4 shows an increase of the photovoltaic plant, compared to the grid-dependent infrastructure. As a matter of fact, energy losses occur during the charge and discharge of the battery and they need to be compensated by a higher production. We also observe less than 0.23% (or 2.3% over a 10 year period) degradation of the batteries. However, the lifespan of a battery is 10 years, which means that the batteries are under-utilized in this infrastructure. In terms of volume, the data centers require on average $2.4m^3$ of battery per server.




Figure 6: Impact on the cost of BESS-only infrastructure when increasing the number of streaming servers.

It is important to note that space may be a limiting factor for deploying this infrastructure in large data centers.

We describe the unit operation cost as the ratio between the overall annual cost and the annual energy consumed (Eq. 24f). It is evaluated to 1.115\$/kWh in this configuration. Compared to the European electricity tariffs (for consumers of more than 500MWh/year) in the second half of 2021 [55], this autonomous architecture is equivalent to 5.7 times the average cost and 3.2 times the highest cost (recorded in Denmark), while it is equal to 5.35 times the French tariff. Although the world went through an energy crisis in 2022, such a price is far from being competitive. However, this configuration is theoretical as composed only of expensive batteries, but it provides a basis for exploring more complex infrastructures including hydrogen storage.

Cloud users usually submit applications with various profiles of energy consumption. Therefore, we mix the Google workload with video streaming trace in order to evaluate the effect of the energy profile on the electrical components. We set the percentage of streaming workload from 0 to 100% of the overall load, in each data center. Figure 6 shows the results for a data center of 4,000 servers. Using streaming servers leads to a linear increase of the number of PV (thus their surface and cost), the battery size (thus its cost) and the operation costs. There are two reasons for these observations: streaming applications are more energy consuming and have a higher dynamic range (Fig. 3), therefore they induce more cycles than the Google workload. However, similarly to the Google trace, streaming has a negligible degradation effect on the battery.

6.3 Autonomous case study using a BESS and a HESS

Let's consider data centers of several sizes as before, running the Google traces. In this configuration, the storage system




Figure 7: Comparison of off-grid infrastructures on Google trace

Table 3: Grid-dependent infrastructure results with Google trace

# Servers	$A_{total}(m^2)$	Injection rate (%)	Self-consumption rate (%)	Peak exportation (kW)	Peak importation (kW)	Peak injection ratio
10	108.2	62.8	37.6	11.5	2.2	4.9
50	374.3	64.1	35.9	40	7.1	3.4
100	700.3	64.1	35.9	74.6	13.3	3.6
500	3,552.4	64.1	35.9	378.3	67.5	3.6
1,000	6,524.4	64.1	35.9	695	124.3	3.5
4,000	25,239.6	64.1	35.9	2,688.2	482.1	3.4
8,000	49,276.7	64.1	35.9	5,248.4	942	3.4

Table 4: Autonomous infrastructure using only a BESS results with Google trace

# Servers	$A_{total}(m^2)$	$C_{Bat}(0)$ (MWh)	PV cost (k\$)	BESS cost (k\$)	Operation cost (k\$)	Operation unit cost (\$/kWh)	P_{Bat}^{max} (MW)	Volume	SOH 1 year
100	803.3	22.2	7.5	120.2	127.7	1.115	5.33	252.7	0.997737
500	4,079.2	112.6	38.3	610.3	648.6	1.114	27.1	1,282.4	0.997736
1,000	7,490.8	207.1	70.3	1,121.5	1,191.8	1.115	49.7	2,357.1	0.997737
4,000	28,979.7	801	272.2	4,338.4	4,610.6	1.115	193	9,117.8	0.997737
8,000	56,579.7	1,564	531.5	8,471.2	9002.7	1.115	376	17,803.9	0.997737

is a combination of an HESS and a BESS. We investigate the annual operation cost, the surface occupied by the PV panels and the volume of the storage units in order to conduct a comparative study between this infrastructure and the one including only a BESS. Thus, we compute the ratio between the value of each factor for this infrastructure and that of the previous infrastructure, as presented in Figure 7. The unit operation cost is estimated to 0.450\$/kWh in this configuration, thus 248% cheaper than using only a BESS. Meanwhile, the PV plant has almost doubled in size and the storage devices volume has increased by 1.27. That is due to the low – but representative of most of the current devices – efficiency (35%) of the hydrogen component that we considered. However, the efficiency of new generation fuel cells and electrolyzers can exceed 70% [30], [40], [41] and that will further reduce the cost and the size of the energy sources. As the data center consumption profiles are globally flat and the PV production profiles are obtained by proportionality, the ratios shown in Figure 7 are not significantly influenced by the size of the data center.

Even though this infrastructure is more profitable, it remains less competitive than the average European, French or even Danish electricity tariff. Yet, it is much closer to being competitive in some countries with high tariff (e.g. 0.3203 \$/kWh in Denmark), and it is expected that storage prices will decrease in the coming years. Until it happens, 100% autonomy is not economically attractive and it would be reasonable to combine the storage units with the grid while maintaining a self-consumption above an acceptable threshold.

Let's vary the percentage of streaming from 0 to 100% of the data centers load to quantify the impact of the workload pattern on the equipment sizing and costs. Figure 8a shows a linear increase in the annual operating cost of the electrical infrastructure for 4,000 servers. It has the same overall effect on the size of the battery (capacity and power rating), the tank, the fuel cell and the electrolyzer, as presented in Figure 8b. In this scenario, the battery has an average SOH of 0.937 (or 6.3% of degradation) at the end of the year, with a deviation of 8.3×10^{-4} . Thus, the battery keeps a nearly constant degradation with respect to the streaming share. We also observe that it is better exploited as 63% of its energy potential is exchanged within 10 years (which correspond to the calendar lifespan of the battery).

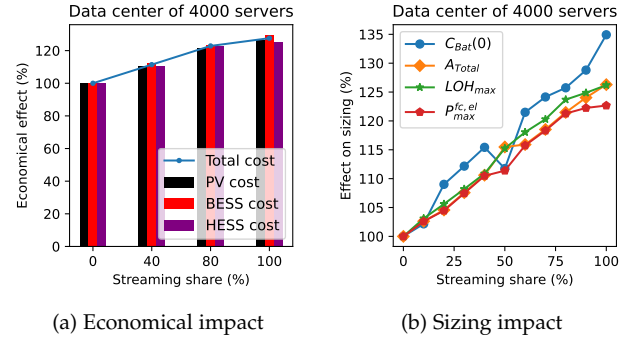


Figure 8: Impact of streaming on the infrastructure.

7 DISCUSSION

This study shows that data centers using the electrical grid with on-site solar energy source are highly dependent on the grid from which they are quasi permanently powered. Moreover, they inject important peaks of power into the grid, which represent a challenge for the grid operator. On the other hand, as expected, combining a BESS and a HESS is the most profitable way to reach a full autonomy. But, it is still costlier than the traditional energy tariffs in European countries. Future work could test intermediate configurations combining the electrical network and storage means, in order to reduce the size of the storage equipment (and thus their cost) while ensuring a high autonomy threshold – that one will progressively increase according to the improvement of storage efficiency and price drop.

Our analysis could be affected by some issues discussed hereafter. First, we consider a given server model and derive its power consumption based on the type of application and the frequency and load of the active processors. This model is extensively used in literature, yet it may be inaccurate for unknown workload. This is why we calibrated it with CPU-intensive workload for Google trace (to get the worst case) and with real streaming measurements for the streaming case. We consider a model based on CPU and memory utilization as these are the metrics provided in the trace. This trace does not include I/O utilization. Other models in the literature [56] can be used to integrate this metric in future work. Also, the base power consumption of our server (79W) represents almost 55% of the peak power consumption (145W). While such a percentage is not unusual

for data center servers, other servers may have different idle-to-peak-power ratio especially when running GPUs. We focused here on a unique server model on which we perform real measurements in order to simplify the analysis, since having heterogeneous servers would make the results more sensitive to workload allocation policy.

The economic model of the battery is relatively conservative: it considers that all expenses, described in [48], must be paid again when the storage is replaced. This would result in a slight overestimation, as some costs (e.g. some installation costs) are paid only for the first time. The solar source may be replaced or combined with wind turbines and drastically impact generation patterns and possibly the conclusions. We focus here on solar energy since it is the most common renewable sources for powering data centers [1]. Indeed, it presents fewer constraints compared to wind generation in terms of required space and nuisance to nearby residents. Yet, the same methodology that was developed here can be employed to conduct a techno-economic study for powering data center through wind sources or for combining solar and wind sources.

Consolidation policies (i.e. allocating the tasks to a minimum number of servers) can reduce the consumption of data centers and thus their electricity bill. Yet, here, the workload comes from a real data center, that apparently does not use consolidation. Modifying the allocation would distance us further from a real data center-based study.

Finally, our study relies on modeling and simulation, which represents an abstraction of the physical environment. However, we carefully selected up-to-date models for their accuracy and practicality. All the models employed for the various physical components (servers, batteries, photovoltaic plants, hydrogen tanks, electrolyzers, etc.) have been validated and used by other work in the literature as shown in Section 4, for purposes different from this study.

8 CONCLUSIONS

In this paper, we conducted a techno-economic study of three types of electric infrastructures (using the external grid, autonomous with battery-based storage only, and autonomous with both battery and hydrogen-based storage) supplying data centers of 10 to 8,000 servers, from 100% renewable energy sources, solar energy here. Our simulations reveal that in the grid-dependent mode, most of the data centers consumption stems from the grid, to which they are therefore highly dependent. We also realize that the data centers regularly inject high peaks of energy into the electrical grid which may be challenging for the grid operator to handle. By assessing the autonomous infrastructures costs, we quantified the additional cost of autonomy, using a combination of BESS and HESS, compared to using PV and network electricity without storage.. In future work, we will consider other renewable energies such as wind that present additional constraints compared to PV electricity in terms of intermittence duration, uncertainty, etc.. We will also consider techniques of load shifting in order to explore their impact on the electricity bill.

ACKNOWLEDGMENTS

This project has received financial support from the CNRS through the MITI interdisciplinary programs. Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations.

REFERENCES

- [1] IEA, "Data Centres and Data Transmission Networks," 2021. [Online]. Available: <https://www.iea.org/reports/data-centres-and-data-transmission-networks>
- [2] C. Freitag, M. Berners-Lee, K. Widdicks, B. Knowles, G. S. Blair, and A. Friday, "The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations," *Patterns*, 2021.
- [3] M. Xu, W. Tian, and R. Buyya, "A survey on load balancing algorithms for virtual machines placement in cloud computing," *Concurrency and Computation: Practice and Experience*, vol. 29, 2017.
- [4] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic Right-Sizing for Power-Proportional Data Centers," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, pp. 1378–1391, 2013.
- [5] C.-M. Wu, R.-S. Chang, and H.-Y. Chan, "A green energy-efficient scheduling algorithm using the DVFS technique for cloud data-centers," *Future Generation Computer Systems*, pp. 141–147, 2014.
- [6] European Commission, "A Clean Planet for all - a European strategic long-term vision for a prosperous, modern, competitive and climate neutral economy," <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52018DC0773>, 2018.
- [7] "Climate Neutral Data Centre Pact," <https://www.climateutraldatacentre.net/>, 2020.
- [8] Íñigo Goiri, M. E. Haque, K. Le, R. Beauchea, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Matching renewable energy supply and demand in green datacenters," *Ad Hoc Networks*, vol. 25, pp. 520–534, 2015.
- [9] L. Grange, G. Da Costa, and P. Stolf, "Green IT scheduling for data center powered with renewable energy," *Future Generation Computer Systems*, vol. 86, pp. 99–120, 2018.
- [10] V. Venkataswamy, J. Grigsby, A. Grimshaw, and Y. Qi, "RARE: Renewable Energy Aware Resource Management in Datacenters," in *Workshop on Job Scheduling Strategies for Parallel Processing*, 2022.
- [11] H. He, H. Shen, Q. Hao, and H. Tian, "Online delay-guaranteed workload scheduling to minimize power cost in cloud data centers using renewable energy," *Journal of Parallel and Distributed Computing*, vol. 159, pp. 51–64, 2022.
- [12] B. Camus, A. Blavette, F. Dufossé, and A.-C. Orgerie, "Harnessing the geographical flexibility of distributed computing clouds for cooperative self-consumption," in *IEEE PES ISGT-Europe*, 2018.
- [13] H. Yuan, J. Bi, M. Zhou, Q. Liu, and A. C. Ammari, "Biobjective task scheduling for distributed green data centers," *IEEE Trans. on Automation Science and Engineering*, vol. 18, pp. 731–742, 2021.
- [14] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE Trans. on Power Systems*, 2018.
- [15] M. Dabbagh, B. Hamdaoui, and A. Rayes, "Peak power shaving for reduced electricity costs in cloud data centers: Opportunities and challenges," *IEEE Network*, vol. 34, no. 3, pp. 148–153, 2020.
- [16] X. Peng, T. Bhattacharya, T. Cao, J. Mao, T. Tekreeti, and X. Qin, "Exploiting renewable energy and UPS systems to reduce power consumption in data centers," *Big Data Research*, vol. 27, 2022.
- [17] D. Al Kez, A. M. Foley, F. W. Ahmed, M. O'Malley, and S. Mueeen, "Potential of data centers for fast frequency response services in synchronously isolated power systems," *Renewable and Sustainable Energy Reviews*, vol. 151, 2021.
- [18] P. Peng, M. Chen, Y. Li, Y. Zhao, N. Li, H. Liu, B. Wang, Y. Sun, and X. Tang, "Research on Energy Storage Type of Uninterruptible Power Supply Technology in Internet Data Center," in *Int. Conf. on Power and Energy Systems*, 2022, pp. 553–558.
- [19] Y. Li, A.-C. Orgerie, and J.-M. Menaud, "Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a Cloud data center," in *Euromicro Int. Conf. on Parallel, Distributed, and Network-Based Processing*, 2017.
- [20] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, "Parasol and greenswitch: Managing datacenters powered by renewable energy," *ACM SIGPLAN Notices*, vol. 48, no. 4, pp. 51–64, 2013.

- [21] J.-M. Pierson *et al.*, "Datazero: Datacenter with zero emission and robust management using renewable energy," *IEEE Access*, vol. 7, pp. 209–230, 2019.
- [22] M. Haddad, G. Da Costa, J.-M. Nicod, M.-C. Péra, J.-M. Pierson, V. Rehn-Sonigo, P. Stolf, and C. Varnier, "Combined IT and power supply infrastructure sizing for standalone green data centers," *Sustainable Computing: Informatics and Systems*, vol. 30, 2021.
- [23] N. Lazaar, M. Barakat, M. Hafiane, J. Sabor, and H. Gualous, "Modeling and control of a hydrogen-based green data center," *Electric Power Systems Research*, vol. 199, 2021.
- [24] G. Radonjić and S. Tompa, "Carbon footprint calculation in telecommunications companies—the importance and relevance of scope 3 greenhouse gases emissions," *Renewable and Sustainable Energy Reviews*, vol. 98, pp. 361–375, 2018.
- [25] G. Rostirolla *et al.*, "A survey of challenges and solutions for the integration of renewable energy in datacenters," *Renewable and Sustainable Energy Reviews*, 2022.
- [26] T. Wang, Z. Su, Y. Xia, and M. Hamdi, "Rethinking the Data Center Networking: Architecture, Network Protocols, and Resource Sharing," *IEEE Access*, vol. 2, pp. 1481–1496, 2014.
- [27] P. Bojek, "Solar PV," IEA report, <https://www.iea.org/reports/solar-pv>, 2021.
- [28] IRENA, "Renewable Capacity Statistics," <https://www.irena.org/publications/2022/Apr/Renewable-Capacity-Statistics-2022>, 2022.
- [29] Apple, "Apple Environmental Progress Report," https://www.apple.com/environment/pdf/Apple_Environmental_Progress_Report_2021.pdf, 2021.
- [30] P. Marocco, D. Ferrero, A. Lanzini, and M. Santarelli, "Optimal design of stand-alone solutions based on RES + hydrogen storage feeding off-grid communities," *Ener. Conv. and Management*, 2021.
- [31] F. C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarié, S. Hunold, A.-C. Orgerie, and M. Quinson, "Predicting the energy-consumption of MPI applications at scale using only a single node," in *IEEE Int. Conf. on Cluster Computing*, 2017.
- [32] CISCO, "Cisco's power consumption calculator," <https://cpc.cloudapps.cisco.com/cpc/>, accessed April 2023.
- [33] A. Adelin, P. Owezarski, and T. Gayraud, "On the impact of monitoring router energy consumption for greening the internet," in *IEEE/ACM Int. Conf. on Grid Computing*, 2010, pp. 298–304.
- [34] M. Mackie, H. Sun, and J. Jiang, "Reinforcement Learning Based Load Balancing for Geographically Distributed Data centres," *IEEE PES Innovative Smart Grid Technologies Europe*, 2021.
- [35] EnerSys, "Battery range summary," <https://www.enersys.com/493bb4/globalassets/documents/product-documentation/datasafe/hx-plus/emea/fr-hx-rs-009-07-2019.pdf>, 2019.
- [36] B. Li, "Sizing and operation of multi-energy hydrogen-based microgrids," Ph.D. dissertation, Univ. Bourgogne, 2018.
- [37] Ballard Power Systems, "Nexa Power Module User's Manual," <http://faculty.stust.edu.tw/~wcchang/MAN5100078.pdf>, 2003.
- [38] P. Garcia, J. P. Torreglosa, L. M. Fernandez, F. Jurado, R. Langella, and A. Testa, "Energy management system based on technoeconomic optimization for microgrids," *Electric Power Systems Research*, vol. 131, pp. 49–59, 2016.
- [39] D. Morin, Y. Stevenin, C. Grolleau, and P. Brault, "Evaluation of performance improvement by model predictive control in a renewable energy system with hydrogen storage," *International Journal of Hydrogen Energy*, vol. 43, no. 45, 2018.
- [40] "H-TEC PEM Electrolyzer ME450," Data sheet, <https://www.h-tec.com/en/products/detail/h-tec-pem-electrolyser-me450/me450/>, 2022.
- [41] Fuji Electric Global, "Fuel Cells specifications," Data sheet, <https://www.fujielectric.com/products/fuelcell/spec.html>, 2022.
- [42] "Fluid properties and conversions," in *Handbook of Valves and Actuators*. Butterworth-Heinemann, 2007, pp. 433–483.
- [43] P. Marocco, D. Ferrero, E. Martelli, M. Santarelli, and A. Lanzini, "An MILP approach for the optimal design of renewable battery-hydrogen energy systems for off-grid insular communities," *Ener. Conversion and Management*, vol. 245, 2021.
- [44] Y. Zhang, "Performance and ageing quantification of electrochemical energy storage elements for aeronautical usage," Ph.D. dissertation, Université de Bordeaux, 2019.
- [45] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2899–2917, 2014.
- [46] V. Ramasamy, D. Feldman, J. Desai, and R. Margolis, "US solar photovoltaic system and energy storage cost benchmarks: Q1 2021," National Renewable Energy Lab., Tech. Rep., 2021.
- [47] SolarWorld, "SolarWorld Sunmodule SW 165/175/185 mono," Technical data sheet, <https://documents.unboundsolar.com/legacy/pdfs/module%20pdf%20folder/SW165175185.pdf>, 2006.
- [48] K. Mongird, V. Viswanathan, J. Alam, C. Vartanian, V. Sprenkle, and R. Baxter, "2020 grid energy storage technology cost and performance assessment," *Energy*, 2020.
- [49] "Materials-Based Hydrogen Storage," Office of Energy Efficiency and Renewable Energy, <https://www.energy.gov/eere/fuelcells/materials-based-hydrogen-storage>, 2020.
- [50] D. Balouek *et al.*, "Adding virtualization capabilities to the Grid'5000 testbed," in *Cloud Computing and Services Science*. Springer, 2013, vol. 367, pp. 3–20.
- [51] J. Wilkes, "Google cluster-usage traces v3," <https://github.com/google/cluster-data/blob/master/ClusterData2019.md>, Google, Tech. Rep., 2020.
- [52] "Amsterdam exchange point," <https://www.ams-ix.net/ams/documentation/total-stats>, 2021–2022.
- [53] Pecan Street, "Pecan street dataport," <https://www.pecanstreet.org/dataport/>, 2018.
- [54] Google, "Google data centers efficiency," <https://www.google.com/about/datacenters/efficiency>, accessed 2023.
- [55] Eurostat, "Electricity price statistics," https://ec.europa.eu/eurostat/databrowser/view/NRG_PC_205_custom_3477477/default/table?lang=en, 2021.
- [56] R. Ferreira da Silva, H. Casanova, A.-C. Orgerie, R. Tanaka, E. Deelman, and F. Suter, "Characterizing, Modeling, and Accurately Simulating Power and Energy Consumption of I/O-intensive Scientific Workflows," *J. of computational science*, 2020.



Wedan Emmanuel Gribga received the engineering degree in embedded electronic systems and systems control in 2020, at the National School of Applied Sciences in Marrakech, Morocco. He is currently a Ph.D student in Computer Science at the University of Rennes 1. His research interests focus on energy efficiency, energy management optimization, Cloud Computing and distributed systems.



Anne Blavette received her Ph.D. degree in Electrical Engineering from University College Cork, Ireland in 2013. She was the recipient of a European Marie Skłodowska-Curie Fellowship to work at the SATIE lab, Ecole Normale Supérieure of Rennes, where she has been a tenured full-time researcher since 2015. Her research interests include the optimised energy management in power systems.



Anne-Cécile Orgerie received her Ph.D. degree in Computer Science from Ecole Normale Supérieure de Lyon (France) in September 2011. She has been a full time researcher at CNRS in the IRISA laboratory (Rennes, France) since October 2012. Her research interests focus on energy efficiency, Cloud Computing, Smart Grids and distributed systems.