



HAL
open science

Energy-related Impact of Redefining Self-consumption for Distributed Edge Datacenters

Wedan Emmanuel Gnibga, Anne Blavette, Anne-Cécile Orgerie

► To cite this version:

Wedan Emmanuel Gnibga, Anne Blavette, Anne-Cécile Orgerie. Energy-related Impact of Redefining Self-consumption for Distributed Edge Datacenters. IGSC 2024 - 15th International Green and Sustainable Computing Conference, Nov 2024, Austin, United States. pp.1-7. hal-04770489

HAL Id: hal-04770489

<https://hal.science/hal-04770489v1>

Submitted on 6 Nov 2024

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

Energy-related Impact of Redefining Self-consumption for Distributed Edge Datacenters

Wedan Emmanuel Gnibga*, Anne Blavette†, Anne-Cécile Orgerie*

*Univ. Rennes, Inria, CNRS, IRISA, Rennes, France, Email: {wedan-emmanuel.gnibga, anne-cecile.orgerie}@irisa.fr

†Univ. Rennes, ENS Rennes, CNRS, SATIE lab, Rennes, France, Email: anne.blavette@ens-rennes.fr

Abstract—Self-consumption is developing worldwide to increase renewable electricity consumption and reduce electricity bills. It can be carried out individually or collectively (grouping several entities), but is generally restricted geographically. In datacenters, one may use load shifting to benefit from the self-consumption tariffs, at the cost of increasing energy consumption. An alternative would be to extend the self-consumption rules to wider perimeters. This paper proposes a comparative study on several aspects influencing the collective self-consumption (CSC) of an Edge infrastructure, including spatial load shifting, temporal load shifting and extending the current rules to encompass wider geographical boundaries. Spatial shifting under the current CSC scheme is found more cost-effective (3.9% of cost reduction) with a negligible increase in energy consumption (0.19%), compared to the revised definition of collective self-consumption which leads to 3.7% cost reduction and no increase in energy consumption. Moreover, allowing up to 10% of the user tasks to be shifted in time can further increase the self-consumption rate by 1.3%.

Index Terms—Collective self-consumption, distributed systems, load shifting, carbon footprint.

I. INTRODUCTION

Electricity users have transformed from mere consumers to key players in the transition to modern and less carbon-intensive grids. Worldwide, individuals and businesses set up their own renewable energy systems (generally PV) to self-consume their production or even to sell any excess power they generate to other consumers. To further incentivize this practice, participants can organize themselves into Energy Communities (ECs) [1], which are supported by specific legislation. Frieden et al. [2] gave an overview of regulations in major European countries that set geographical boundaries of ECs and provide specific guidelines for their organization and their operation. Under certain conditions, which may vary from one country to another, the energy consumption of energy communities may (partly) be qualified as collective self-consumption. Under this scheme, electricity is considered as self-consumed when it is produced and consumed within the boundaries of the energy community at the same instant, and is paid at a lower tariff or is even free. In other words, collective self-consumption is an incentive that allows energy customers to synchronize their global electricity production with their consumption, benefiting the power grid.

As the existing rules restrict the geographical radius, distributed datacenters (DCs) such as Edge infrastructures equipped with photovoltaic power plants cannot benefit from synchronizing their global power generation and demand.

Consequently, numerous research propose “follow-the-sun” approaches that aim to shift the DC load to where (spatial load shifting) and when (temporal load shifting) the panels generate abundantly. These practices have two objectives: 1) consume more “green” energy on-site to claim sustainability credits (i.e. increase the self-consumption and “reduce” the carbon impact from the perspective of the DCs operator), and 2) reduce electricity bills as the self-produced electricity is free once the PV investment is done. However, generating electricity at one location and consuming it at another within the same uncongested grid leads to the same global carbon impact (although with some increased network losses), and therefore shifting the load to the specific generation site may be unnecessary. Instead, it may increase the overall carbon footprint at the global scale, as more energy is needed to carry out this operation. We believe that an alternative to spatial load shifting would be to extend the rules of (collective) self-consumption allowing distributed DCs to benefit from more favorable tariffs by coordinating themselves, thereby achieving a synchronization at a large scale and cutting down the overall carbon impact. However, key questions emerge from that hypothesis: will redefining CSC unlock more financial advantages? Does it deliver significantly greater energy savings compared to the traditional load shifting? Does it substantially influence the collective self-consumption? To answer these questions, we propose a comparative economic and environmental analysis of two definitions of collective self-consumption in an Edge computing infrastructure. This infrastructure is made of small-scale DCs named here Edge-DCs. Our main contributions are to:

- formulate an extended definition of CSC where energy produced at one Edge-DC and consumed at another will be considered as self-consumed. This would be accompanied by a financial compensation to the grid operator for using the grid. Then, we compare the resulting electrical and economical gains to those of the current CSC rules.
- we quantify the cost-effectiveness and self-consumption benefits of spatial and temporal load shifting within existing and extended CSC frameworks, and their impact of Edge-DCs energy consumption.

In the following, Section III presents our methodology, including the collective self-consumption frameworks and the load shifting strategies, then Section IV presents the results and discussions, and Section V presents our conclusions.

II. RELATED WORK

Today, many research projects rely on spatial load shifting through virtual machine migration. For instance, the authors of [3], [4] base on that principle to dynamically route the load to the DCs where renewable energy is the most readily available. While this inter-site load exchange leads to consume free electricity and thus, to make financial profits, it requires an increase in energy consumption to perform such migrations [5]. In addition to spatial shifting, one may delay some of the load (with the clients approval) to times when renewable energy is most abundant [6]. This is referred to as temporal shifting. The Lawrence Berkeley National Laboratory (LBNL) has quantified the benefits of this spatial-temporal flexibility [7]. They show that up to 27% of the DCs' power load can be moved, potentially outweighing some of the drawbacks associated with not participating in CSC schemes.

The European Commission [8] believes that extending the definition of collective self-consumption to encompass broader distances would promote better coordination between energy consumption and production. In France, the decree [9] organizing CSC limits it to a perimeter of 2 km with an aggregated power generation capacity below 3 MW. However, although local CSC presents both local and global advantages, an extended CSC would retain the global advantage (i.e. production/consumption balancing) and could therefore be incentivized. Hence, distant entities (such as PV-powered DCs) could benefit from more favorable tariffs. The carbon-neutrality objectives are indeed driving DC operators to move towards powering their facilities with renewable energy sources, often installed on-site. Therefore, an extended collective self-consumption could open doors to align the geographically distributed production and load, to claim sustainability credits and to make profit.

III. METHODOLOGY

A. System description

The system considered in this study is a nation-wide distributed Edge infrastructure made of 260 Edge-DCs, as presented in [10]. They are all connected to the grid. Two types of DCs are identified according to their electrical infrastructure: 1) those with a solar power source installed on-site, 2) those with no other source but the grid. It is worth noting that the electrical system does not include energy storage equipment as our objective is to examine and compare strategies enabling the DCs load to be synchronized with the PV generation, that would help to reduce the need for storing energy. Self-consumption is usually incentivized by electricity selling prices being lower than purchasing prices. To benefit from it, we have investigated three solutions: 1) **spatial load shifting**: this approach aims to maximize self-consumption, as per the current definition, by strategically shifting workloads among sites through the telecommunication networks. 2) **Extended self-consumption definition**: we formulate an extended definition of collective self-consumption (CSC) where energy injected into the grid by one site and

simultaneously consumed by another, regardless of distance, is considered as self-consumed and paid only for a network usage fee if it is not consumed and produced on the same site. 3) **Temporal load shifting**: we investigate how strategically shifting energy consumption patterns across time can further improve the CSC and reduce the costs. We use a distributed resource management approach. When coordination is necessary, such as for exchanging workload during spatial load shifting, we employ P2P negotiations.

B. Improving the collective self-consumption

1) Spatial load shifting:

In this study, we aim to establish an upper limit for the potential cost reductions and overall on-site energy consumption achievable through spatial load shifting. Therefore, we assume that there are no geographical restrictions on where the shifted load can be relocated within a given national power system (i.e. no congestion due to consumption or production peaks). However, we do not account for the energy consumption of the telecommunication network that might be substantial in a context of intensive VMs migration. Indeed, the telecommunication network might not belong to the Edge-DCs' operator and thus, cannot be considered within the same entity for the self-consumption strategy.

Let's call N_i^A the number of virtual machines (VMs) running on site i during the current time slot t , $R_{k,i}$ (respectively $P_{k,i}(t)$) the IT resources (CPU and RAM) requested (respectively power demand at time t) by VM k ($\forall k \in [0, N_i^A]$) and $R_{\max,i}$ the maximum IT resources available on site i . Let's assume that site i has no (sufficient) on-site renewable power generation, and there are other DCs with surplus of PV power generation. Let j be any of these DCs with surplus of on-site power generation. Spatial load shifting objective is to maximize on-site consumption (Equation 1a) at j , as the energy produced and consumed simultaneously at the DC j is free. Hence, at each time step t and for each task k running on i , the bilateral negotiations consist in finding a site j with enough free resources (Eq. 1b) and sufficient surplus of PV generation to power k (Eq. 1c), in the set of sites equipped with on-site power plant (\mathcal{N}_{PV}). If no suitable site can be identified for load relocation, the process terminates. Similarly, if site i is equipped with a power plant and its on-site generation becomes sufficient to meet its remaining load after a certain amount of spatial load shifting, it stops negotiating (Eq. 1d).

$$\min_{P_{load,j}(t)} P_{PV,j}(t) - P_{load,j}(t) \quad (1a)$$

$$R_{k,i} + \sum_{l \leq N_j^A} R_{l,j} \leq R_{\max,j} ; j \in \mathcal{N}_{PV} ; i \neq j \quad (1b)$$

$$P_{k,i}(t) \leq P_{PV,j}(t) - P_{load,j}(t) \quad (1c)$$

$$P_{PV,i}(t) \leq P_{load,i}(t) \quad (1d)$$

Where $P_{PV,j}(t)$ and $P_{load,j}(t)$ are the power generated by the power plant and the power demand of the site j respectively at t . $R_{l,j}$ are the resources used by the task l ($l \leq N_j^A$).

2) Extended collective self-consumption definition:

Unlike spatial load shifting which implies a supposedly increased consumption due to VM migration, another strategy

is studied here in the form of an extension of the collective self-consumption definition. Under this extended definition of collective self-consumption (CSC) we propose in this study, any amount of energy produced and consumed simultaneously is considered as self-consumed, regardless of the distance between the consumption and production sites. In essence, we relax the spatial restriction and impose a strict temporal constraint. However, the total amount of power generation remains capped at 3MW, as per the current French regulation. This approach eliminates the need for physical load migration to benefit from "free" energy, once the PV panels are paid. However, it is assumed here that the DCs operator must pay financial compensation to the grid operator for utilizing the electrical network for the amount of consumed energy that does not match with on-site production.

Equation 2 describes the rules adopted to extend the collective self-consumption definition. For the sake of generality, let's consider $P_{PV,i}(t)$ as the power generation on a site i . Note that $P_{PV,i}(t) = 0$ for sites with no power plant. During each decision-making interval, a designated coordinator collects information on the surplus power generation (Eq. 2a) and power deficit (Eq. 2b) experienced by all DCs. This coordinator, based on a predefined allocation strategy, then allocates the accumulated surplus to sites facing an energy shortfall. Defining the allocation for each site is necessary to distinguish between the energy assumed to be bought from the grid, and this bought from the group of Edge-DCs. We employ a prorate allocation scheme (Eq. 2c) as introduced in [1]. Hence, to each site is attributed a coefficient $\epsilon_i \in [0, 1]$ representing the portion of the collective surplus that should be allocated to that site. The power allocated to a site i does not exceed the amount needed to balance the deficit (Eq. 2d).

$$P_{surp,i}(t) = \max\{0 ; P_{PV,i}(t) - P_{load,i}(t)\} \quad (2a)$$

$$P_{def,i}(t) = \max\{0 ; P_{load,i}(t) - P_{PV,i}(t)\} \quad (2b)$$

$$\epsilon_i = [P_{def,i}(t)] / [\sum_{j=0}^N P_{def,i}(t)] ; \sum \epsilon_i \leq 1 \quad (2c)$$

$$P_{com,p,i}(t) = \min\{P_{def,i}(t) ; \epsilon_i \cdot \sum_{j=0}^N P_{surp,j}(t)\} \quad (2d)$$

$$P_{com,s,i}(t) = \frac{P_{surp,i}(t)}{\sum_{j=0}^N P_{surp,j}(t)} \cdot \sum_{j=0}^N P_{com,p,j}(t) \quad (2e)$$

Where $P_{surp,i}(t)$ and $P_{def,i}(t)$ are respectively the surplus and deficit on site i at time t , $P_{com,p,i}(t)$ and $P_{com,s,i}(t)$ are amount of power considered as imported and exported within the perimeter of the DC respectively.

If the total generated power exceeds the collective demand, the rest will be considered as purchased by the energy provider. The coordinator always makes sure that, $\sum_{i=0}^N P_{com,s,i}(t) = \sum_{i=0}^N P_{com,p,i}(t)$ where N denotes the number of sites. This ensures that all the power that is considered as exported to the community is equal to the amount that is considered as imported from the community, thus preventing any non-physical double allocation of energy (i.e. exchanged within the community and exchanged with the energy supplier).

3) Temporal load shifting:

There are two types of tasks submitted at each DC. Some tasks (called service workload) are urgent and need to be processed

right away. Others (called batch workload) are flexible and can be scheduled for later as long as they meet the client's quality of service (QoS) requirements. This flexibility in batch workloads creates an opportunity to capitalize on the PV plants generation. Temporal load shifting consists in postponing certain tasks until periods when the local power plant produces a surplus of electricity, thus improving the synchronization between production and consumption. It comprehends two phases as represented in Figure 1.

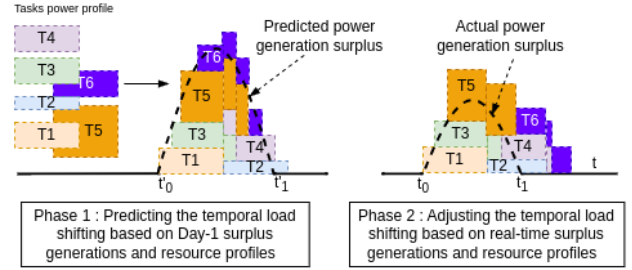


Figure 1: Temporal load shifting representation. E.g. the actual production (Phase 2) falls short compared to the predictions (Phase 1), so the task T6 is postponed, and finally deployed when there is no foreseeable surplus of PV production.

In the first stage, we use information about solar power generation and IT resource availability from the previous day to identify periods when there is surplus of generation. We then progressively (at each instant of the current day) predict which batch tasks can be delayed to these times. The strategy aims to use up all the excess solar power generation. In this work, we rely on a day-ahead persistence forecast model [11] to predict IT resource availability and on-site power surplus. This forecast model is a simple method that assumes that the future will be the same as the present. In other words, we assume that whatever is happening today will happen tomorrow. Let i be a DC, d be the current day, $\vec{P}_{surp,i}(d)$ and $\vec{R}_i(d)$ be respectively vectors of the predicted surplus of power production and the resources availability on day d . As a site's actual resources available at the beginning of days $d-1$ (previous day) and d (current day) might differ, we adjust the entries $\vec{R}_i(t) \in \vec{R}_i(d)$ by $\tilde{R}_i(t) = R_i(t-24h) + R_i(t_{start,d}) - R_i(t_{start,d}-24h)$ where $t_{start,d}$ is the time when the day d starts and $R_i(t_{start,d})$ is the actual resources available at that time. t'_0 and t'_1 are respectively, the instants when site i is predicted to start and stop generating surplus power. Let $\vec{P}_{surp,i}(d-1)$ be the actual surpluses of power generation on day $d-1$. The persistence forecasting model yields $\vec{P}_{surp,i}(d) = \vec{P}_{surp,i}(d-1)$.

The second phase fine-tunes the load-shifting decisions made in the first phase, considering real-time conditions. In fact, predictions might not be perfect as weather conditions and user requests actually change from a day to another. Sometimes, tasks might be scheduled for times with less solar power generation than expected, or when resources are not available. In these cases, the second phase might postpone those tasks again to the next time slot. In contrast, if a day turns out to be sunnier than predicted, there might be

more excess solar power generation on a given instant, and potentially occurring earlier or lasting longer ($t_0 \neq t'_0$ and $t_1 \neq t'_1$). Hence, we reschedule the previously delayed tasks. This process is illustrated in Phase 2 of Figure 1.

C. Per-site cost and energy management models

This section describes the constraints of energy management and the costs at each DC for a given time t , when all the tasks are deployed. For sake of generality, let's consider a site i with an on-site power plant (on the sites with no PV, the power generation is 0), under the extended collective self-consumption definition. The energy transactions are subject to the power balance represented in Equation 3a. A site cannot both purchase energy from the grid and sell energy back to the grid simultaneously as shown in Equation 3b. Similarly, Equation 3c prevents a site from consuming and injecting energy simultaneously from/into the DCs energy community. Each site can only import (Eq. 3d) or export (Eq. 3e) power up to a certain limit set in a contract with the energy provider. In this work, we consider the same power limit for imports and exports, regardless of the allocation of this power (within the DC community or to/from the energy supplier).

The energy cost in the site i is formalized as in Equation 3f. It is composed of 1) energy supplier subscription fees of the considered site, 2) grid import costs, 3) the PV plant's total cost of ownership (TCO) estimated using the model proposed in [12], and 4) fees for electrical network usage (only in the extended CSC definition scheme).

$$\begin{aligned} P_{g,p,i}(t) + P_{com,p,i}(t) + P_{PV,i}(t) \\ = P_{load,i}(t) + P_{g,s,i}(t) + P_{com,s,i}(t) \end{aligned} \quad (3a)$$

$$P_{g,p,i}(t) \cdot P_{g,s,i}(t) = 0 \quad (3b)$$

$$P_{com,p,i}(t) \cdot P_{com,s,i}(t) = 0 \quad (3c)$$

$$P_{com,p,i}(t) + P_{g,p,i}(t) \leq P_{subs,i} \quad (3d)$$

$$P_{com,s,i}(t) + P_{g,s,i}(t) \leq P_{subs,i} \quad (3e)$$

$$\begin{aligned} cost(t) = & c_{g,subs} + c_{g,p}(t) \cdot P_{g,p,i}(t) \cdot \Delta t + TCO_{PV,i}(t) \cdot \Delta t \\ & - c_{g,s}(t) \cdot P_{g,s,i}(t) \cdot \Delta t + \alpha \cdot P_{com,p,i}(t) \cdot \Delta t \end{aligned} \quad (3f)$$

Where $TCO_{PV,i}(t)$ is the instantaneous PV total cost of ownership, $c_{g,p}$, $c_{g,s}$ and α (€/kWh) are respectively the grid purchase and sale tariffs from the energy supplier, and the electrical network usage fees (applicable to the extended definition of collective self-consumption), $c_{g,subs}$ (€) and $P_{subs,i}$ (kW) are respectively the instantaneous subscription fees (calculated by dividing the monthly fees by the number of time slots in a month) and the subscription power capacity. We consider a discrete time step of duration Δt .

IV. CASE STUDY

A. Input data and performance metrics

To perform the analysis, we ran simulations on the Sim-Grid toolkit [13]. We ran simulations for forty (40) days, setting the time step to thirty (30) minutes that matches the interval between data readings from the French smart meters called Linky [14]. We focus on the results from the thirty (30) last days (a month), during which the system has reached the steady state in the simulation and allows to analyze a typical

billing period. We adopt the classical DCs and workload power profiles described in [10]. IT resource power demand has two main components: 1) static power demand that refers to the power consumed when a server is turned on, 2) dynamic power demand that represents the additional power consumed when the server is actively processing tasks.

To simulate real-world conditions and the local spatial variability of solar power generation, we used an annual rooftop solar power data from Austin, Texas collected in 2018 [15]. We split the data into subsets of weekly generation. So every week, each site with a solar power plant is randomly assigned a subset of data corresponding to the generation of 125 solar panels. That ensures that each site experiences a diverse range of solar production throughout the simulation. The study employs fixed electricity tariffs from energy provider EDF for 12 kVA non-residential customers in 2023. These tariffs concern both purchasing power [16] ($c_{g,subs} = 21.88$ €/month where 5.5% of this tariff is paid for VAT, $c_{g,p} = 0.211$ €/kWh where 20% represents VAT) and selling excess electricity back to the grid [17] ($c_{g,s} = 0.13$ €/kWh). The grid operator is assumed here to receive compensation fee of $\alpha = 0.05$ €/kWh.

In this study, we evaluate five monthly metrics: the total cost $Cost_T$ –including the electrical bills and the solar power plant's total cost of ownership (TCO)– as shown in Equation 4a, the aggregated energy consumption E_{load} (Eq. 4b), the amount of energy consumed outside the perimeter of the Edge infrastructure $E_{g,p}$ (Eq. 4c), the monthly peak $P_{g,s,max}$ of power generated and non consumed by the sites (Eq. 4d) and the collective self-consumption rate SC_r that represents the complement of grid contribution to the power consumption.

$$\begin{aligned} Cost_T = & \sum_i \sum_t c_{g,subs} + c_{g,p}(t) \cdot P_{g,p,i}(t) \cdot \Delta t + TCO_{PV,i}(t) \cdot \Delta t \\ & - \sum_i \sum_t c_{g,s}(t) \cdot P_{g,s,i}(t) \cdot \Delta t + \alpha \cdot P_{com,p,i}(t) \cdot \Delta t \end{aligned} \quad (4a)$$

$$E_{load} = \sum_i \sum_t P_{load,i}(t) \cdot \Delta t \quad (4b) \quad E_{g,p} = \sum_i \sum_t P_{g,p,i}(t) \cdot \Delta t \quad (4c)$$

$$P_{g,s,max} = \max_t \sum_i P_{g,s,i}(t) \quad (4d) \quad SC_r = 100 \cdot (1 - \frac{\sum_i \sum_t P_{g,p,i}(t)}{\sum_i \sum_t P_{load,i}(t)}) \quad (4e)$$

B. Results of the CSC frameworks

Let's first consider the following three scenarios: 1) the system adheres to the current French CSC framework without any spatial or temporal load shifting. This baseline serves as a reference point to assess the effectiveness of optimizations introduced in other scenarios, 2) the system adheres to current CSC framework and spatial load shifting is processed among the sites, 3) we apply the extended CSC definition. Figure 2 shows the power demand patterns for a typical day, along with the portion supplied by the grid. A summary of the monthly metrics is presented in the top part of Table I.

We observe in Figure 2 that under the current French CSC rules, spatial load shifting leads to peaks in power demand during daytime, with a negligible amplitude increase of 1.6% compared to the average power demand. Also, the energy overhead due to VM migration, supposedly leading to an increase in energy consumption, is partially compensated during the non generating periods, due to the resources consolidation policy we employ (similar to [10]), which periodically switches off the unused resources. Overall, spatial

Table I: Monthly comparison of the strategies.

| Scenarios | Energy consumption E_{load} (MWh) | Energy imports $E_{g,p}$ (MWh) | Energy injection $E_{g,s}$ (MWh) | Peak injection $P_{g,s,max}$ (MW) | CSC_r (%) | Cost(€) |
|--|-------------------------------------|--------------------------------|----------------------------------|-----------------------------------|-------------|---------|
| Baseline (French CSC) | 678 | 561 | 168 | 1.222 | 17.29 | 102947 |
| Spatial load shifting (French CSC) | 679.3 | 510 | 116 | 0.94 | 24.94 | 98913 |
| Extended CSC | 678 | 442 | 49 | 0.728 | 34.85 | 99090 |
| Current CSC: spatial + temporal shifting (5% tasks) | 679.6 | 509.8 | 115.6 | 0.939 | 24.97 | 98966 |
| Current CSC: spatial + temporal shifting (10% tasks) | 680 | 510.2 | 115.6 | 0.942 | 24.96 | 99024 |
| Extended CSC: temporal shifting (5% tasks) | 678.5 | 438 | 44 | 0.694 | 35.49 | 98897 |
| Extended CSC: temporal shifting (10% tasks) | 679.8 | 434 | 39 | 0.652 | 36.16 | 98840 |

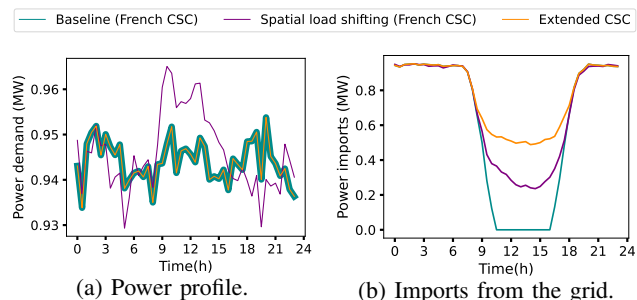


Figure 2: Performance of the extended CSC compared to the current definition (with and without spatial load shifting).

load shifting increases the energy consumption by 1.3MWh (or 0.19%) over one month, compared to the baseline which is negligible. However, as stated earlier, we do not account for the energy consumption of the telecommunication network that might be substantial. In contrast, spatial load shifting reduces daytime power demands from outside the perimeter of the Edge infrastructure. In particular, it reduces the peak of power injected into the grid by 23.1% (282kW). This results in a monthly reduction of the energy imports from the grid by 51MWh (or 9.1%), represents a collective self-consumption rate improvement by 7.65% and reduces costs by of 3,9%.

In the extended CSC definition, the sites consume the same amount of energy as in the baseline, hence the superposition of both profiles in Figure 2. Moreover, energy consumption outside of the infrastructure perimeter is reduced by 21.2% (119MWh) while the Edge infrastructure reduces and sometimes eliminates its reliance on additional sources of power during daytime. The self-consumption rate is thus improved by 17.6% compared to the baseline. This framework reduces the peaks of power injected and not immediately consumed by 40.6% (495kW). It is noteworthy that this does not change the actual physical power flow in the grid. The same amount of physical power is drawn and exported from/into the grid as in the baseline scenario, but accounting practices makes the difference. However, this metric can provide some legibility to the DC operator on its global impact on the network. The extended CSC definition reduces the monthly costs by 3.7%.

In summary, extending the self-consumption perimeter to accommodate distant sites benefits to the Edge-DCs operator in terms of cost reduction. In the other hand, spatial load shifting achieves similar financial benefits with a negligible increase in energy consumption while also reducing congestion and line losses. However, spatial load shifting suffers from a lower load/PV production synchronization potential compared

to extending the CSC perimeters: 1) the DCs' resources are limited, which bounds the amount of load that can be migrated. Moreover, Edge applications require proximity to users for low latency, restricting their migration options. 2) the telecommunication networks have limited bandwidth, which also bounds the amount of migrated load. Intensive load shifting may require to redesign more powerful telecommunication network, thus increasing the power demand on the network side, both for its operation and its manufacturing.

C. Association of the spatial and temporal load shiftings in the current CSC scheme

Now, we analyze the effect of associating temporal and spatial load shifting under the current CSC framework. We assume that a portion of the tasks can be delayed (batch tasks). However, all tasks must be completed within a 24-hour deadline. We simulate scenarios where different shares of tasks (successively 0%, 5% and 10%) can be delayed. Figure 3 shows the results for a typical day of the month and the monthly results are summarized in the middle part of Table I. We observe that this combination has no significant effect on the monthly energy and economical metrics. However, it further increases the energy consumption, especially due to the fact that more resources are required to operate the delayed tasks during the daytime. That results in a negligible increase in the costs (<0.1%). Hence, this combination may not be attractive from a financial perspective, and maybe also from an environmental perspective (more energy consumed but from a higher share of renewables).

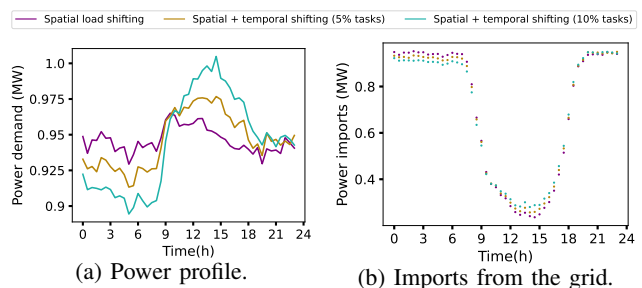


Figure 3: Current CSC with spatial and temporal load shifting.

D. Temporal load shifting in the extended CSC definition

In this section, we perform temporal load shifting under the extended CSC definition. Figure 4 illustrates the power demand and grid import profiles (outside of the Edge infrastructure perimeter) on a typical day of the month, for three scenarios. Respectively 0% (i.e. no temporal load shifting),

5% and 10% of user tasks can be postponed. We observe that from 0AM to 9PM, the power demand is reduced when we perform temporal load shifting, slightly reducing the power demand from the grid. During the daytime, the power demand increases as the result of the additional load previously delayed. However, it does not impact the power demand from the grid, as the PV plants generate enough to power the load. In the afternoon (after 5PM), the overall power demand drawn from outside the sites perimeter is higher than when no temporal load shifting is considered. That may be due to 2 reasons: 1) some long tasks deployed during daytime were not terminated, 2) some tasks could not be started during the daytime due to some sites resources unavailability or low power generation compared to what was predicted.

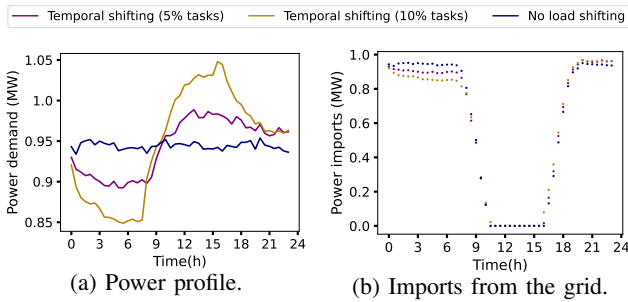


Figure 4: Extended CSC with temporal load shifting.

The monthly results are summarized in the bottom part of Table I. Temporal load shifting increases the energy consumption (0.5MWh and 2.8MWh respectively for 5% and 10% batch tasks which represents an increase of less than 1% compared to the baseline). This is solely due to the fact that shifting tasks changes the overall resource requirements. Consequently, more servers need to be powered on during the day, leading to higher static power demand. We observe peaks in power demand during daytime, which coincides with periods of high solar generation. In contrast, there is a decrease in power request early in the day. This pattern reflects the strategy of postponing tasks from mornings to sunnier periods, effectively leveraging renewable energy. Increasing the number of tasks shifted further amplifies this observation. Hence, the strategy leads to the decrease of power imports from the grid (outside of the infrastructure perimeter) by 0.9% and 1.81% respectively (for 5% and 10% of batch tasks) compared to when no load shifting is processed. This represents an increase of 0.64% and 1.31% of the CSC rate. The strategy also reduces the costs by 0.19% and 0.25% compared to no temporal load shifting, and the energy injections by 10.2% and 20.41%.

In summary, performing temporal load shifting in the proposed extended CSC scheme increases slightly the collective self-consumption rate and reduces the peaks of power injected out of the DCs perimeter. This scenario is even more cost effective than the case where spatial load shifting is performed in the current CSC scheme (Section IV-B).

E. Discussion

Our comparative analysis has not considered the benefits of co-localizing production and consumption on line losses and congestion reduction. However, one of the advantages of CSC over a given geographical perimeter is indeed to reduce line losses. Hence, comparing the energy consumption increase generated by VMs migration to the line losses induced by the extended self-consumption definition would be insightful. However, such losses are highly dependent on the electrical network impedance topology and associated consumption/production patterns. Hence, it is necessary to carry out case-by-case studies on a large number of representative examples, which is out of the scope of this paper.

Given the low power flows by each site (i.e. a few kW), it may be assumed that the excess PV energy will be physically absorbed by other consumers located close to the DC, thus reducing line losses. In this case, depending on the impedance of the lines/transformers and on the levels of local DCs consumption/production, spatial load shifting may result in either more or less losses. Regarding congestion, the low power involved per site leads to think that their contribution to power congestion reduction would be negligible.

Also, we do not consider the energy consumed by the telecommunication network for VM migrations as it does not impact the DC operator's revenue. Yet, it make spatial load shifting even less sustainable as this energy is consumed from an energy mix including non-renewable sources. Moreover, the extended collective self-consumption can easily be applied to large-scale DCs by relaxing the limitation of the aggregated power capacity (3MW in the French current regulation), and extensively to multi-factory industries.

V. CONCLUSIONS

Collective self-consumption aims to balance energy use and potentially generate profit for local producers and consumers within the same grid. The extension of the CSC area is envisaged as a promising means to further synchronize production and consumption. In this paper, we conduct a comparative analysis on the impact of several strategies and two definitions of collective self-consumption. The results show that spatial shifting reduces costs by 3.9% with a negligible increase in energy consumption (0.19%). It is closely followed by the extended CSC approach that reduces the costs by 3.7% with no increase in energy consumption. Moreover, allowing up to 10% of the tasks to be delayed can further increase profits by 0.25% in the extended scheme (becoming more cost-effective than spatial shifting under the current scheme), while being environmentally detrimental in the current scheme.

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] A. D. Mustika, R. Rigo-Mariani, V. Debusschere, and A. Pachurka, "A two-stage management strategy for the optimal operation and billing in an energy community with collective self-consumption," *Applied Energy*, vol. 310, 2022.
- [2] D. Frieden, A. Tuerk, J. Roberts, S. D'Herbemont, A. F. Gubina, and B. Komel, "Overview of emerging regulatory frameworks on collective self-consumption and energy communities in Europe," in *International Conference on the European Energy Market (EEM)*, pp. 1–6, 2019.
- [3] Y. Cao, F. Cao, Y. Wang, J. Wang, L. Wu, and Z. Ding, "Managing data center cluster as non-wire alternative: A case in balancing market," *Applied Energy*, vol. 360, 2024.
- [4] M. Zakarya, A. A. Khan, M. R. C. Qazani, H. Ali, M. Al-Bahri, A. U. R. Khan, A. Ali, and R. Khan, "Sustainable computing across datacenters: A review of enabling models and techniques," *Computer Science Review*, vol. 52, p. 100620, 2024.
- [5] B. Camus, A. Blavette, F. Dufossé, and A.-C. Orgerie, "Self-Consumption Optimization of Renewable Energy Production in Distributed Clouds," in *IEEE International Conference on Cluster Computing*, pp. 370–380, 2018.
- [6] I. Riepin, T. Brown, and V. Zavala, "Spatio-temporal load shifting for truly clean computing," *arXiv preprint arXiv:2405.00036*, 2024.
- [7] G. Ghatikar, V. Ganti, N. Matsun, and M. A. Piette, "Demand response opportunities and enabling technologies for data centers: Findings from field studies," tech. rep., Lawrence Berkeley National Lab., 2012.
- [8] P. J. Coulon, "Opinion of the European Economic and Social Committee on individual and collective energy self-consumption as a factor in the fight for the green and energy transition, and for economic and social balance." Official Journal of the European Union, <http://data.europa.eu/eli/C/2024/873/oj>, accessed 24-04-2024.
- [9] "Arrêté du 21 novembre 2019 fixant le critère de proximité géographique de l'autoconsommation collective étendue." <https://www.legifrance.gouv.fr/loda/id/JORFTEXT000039417566/>, 2020. Accessed: 2024-04-10.
- [10] W. E. Gnibga, A. Blavette, and A.-C. Orgerie, "Latency, Energy and Carbon Aware Collaborative Resource Allocation with Consolidation and QoS Degradation Strategies in Edge Computing," in *IEEE International Conference on Parallel and Distributed Systems*, 2023.
- [11] Y. Wang, D. Millstein, A. D. Mills, S. Jeong, and A. Ancell, "The cost of day-ahead solar forecasting errors in the United States," *Solar Energy*, vol. 231, pp. 846–856, 2022.
- [12] W. E. Gnibga, A. Blavette, and A.-C. Orgerie, "Renewable energy in data centers: the dilemma of electrical grid dependency and autonomy costs," *IEEE Transactions on Sustainable Computing*, 2023.
- [13] 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.
- [14] Enedis, "Tout savoir sur le compteur électrique Linky." <https://particulier.edf.fr/fr/accueil/gestion-contrat/compteur/compteur-linky.html#fn3>. Accessed:10-09-2024.
- [15] Pecan Street, "Pecan street dataport." <https://www.pecanstreet.org/dataport/>, 2018. Accessed: 2024-04-10.
- [16] EDF, "Fiche descriptive Offre de fourniture d'électricité au Tarif Bleu ." https://www.edf.fr/sites/entreprise/files/2023-02/Fiche_CRE_Tarif_Bleu_Non_residentiel_Fevrier_2023.pdf. Accessed: 2024-04-10.
- [17] EDF, "Tarif rachat photovoltaïque: découvrez les prix à jour en 2023." <https://mypower.engie.fr/conseils/couts-panneaux-solaires/tarif-achat-photovoltaique-rachat-electricite.html>. Accessed: 2024-04-10.