



Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale

Jinghan Sun*
UIUC
js39@illinois.edu

Zibo Gong*
UIUC
zibog2@illinois.edu

Anup Agarwal
CMU
anupa@andrew.cmu.edu

Shadi Noghiabi
Microsoft Research
shadi.noghiabi@gmail.com

Ranveer Chandra
Microsoft Research
ranveer@microsoft.com

Marc Snir
UIUC
snir@illinois.edu

Jian Huang
UIUC
jianh@illinois.edu

ABSTRACT

Modular data centers (MDCs) that can be placed right at the energy farms and powered mostly by renewable energy, is a flexible and effective approach to lowering the carbon footprint of data centers. However, the main challenge of using renewable energy is the high variability of power produced, which implies large volatility in powering computing resources at MDCs, and degraded application performance due to the task evictions and migrations. This causes challenges for platform operators to decide the MDC deployment.

To this end, we present SkyBox, a framework that employs a learning-based approach for platform operators to explore the efficient use of renewable energy with MDC deployment across geographical regions. SkyBox is driven by the insights based on our study of real-world power traces from a variety of renewable energy farms – the predictable production of renewable energy and the complementary nature of energy production patterns across different renewable energy sources and locations. With these insights, SkyBox uses the coefficient of variation metric to select the qualified renewable farms, it can identify a set of farms with complementary energy production patterns with a subgraph identification algorithm. After that, SkyBox enables smart workload placement and migrations to further tolerate the power variability. Our experiments with real power traces and datacenter workloads show that SkyBox has the lowest carbon emissions compared with existing approaches. SkyBox also minimizes the negative impact of the power

variability on cloud applications, enabling it an effective solution of utilizing renewable energy for modern data centers.

CCS CONCEPTS

• **Hardware** → **Renewable energy**; • **Information systems** → **Data centers**; • **Computing methodologies** → **Planning and scheduling**.

KEYWORDS

Renewable energy, modular data center, cloud computing, workload scheduling

ACM Reference Format:

Jinghan Sun, Zibo Gong, Anup Agarwal, Shadi Noghiabi, Ranveer Chandra, Marc Snir, and Jian Huang. 2024. Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale. In *ACM Symposium on Cloud Computing (SoCC '24)*, November 20–22, 2024, Redmond, WA, USA. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3698038.3698544>

1 INTRODUCTION

Data centers today consume more than 2% of total U.S. power [1] and emit even more carbon than the aviation industry [79]. The power consumption of data centers is projected to further increase in the near future [78]. As a result, major datacenter vendors have to purchase carbon credits to temporarily offset the carbon impact of data centers [46, 71, 95]. However, this can only temporarily mitigate the carbon impact of data centers. To curb the carbon footprint of data centers at scale, a promising long-term solution is to rely more on renewable energy sources (e.g., solar and wind) as opposed to non-renewable sources (e.g., oil and gas). This is feasible, especially considering the recently reduced cost of renewable energy [2, 56]. However, the key challenge of utilizing renewable energy is their variability across time and space. For instance, solar power production varies across time and geographical locations, and wind power production changes frequently based on the direction and speed of wind.

To manage the variability of renewable energy production, two common approaches have been explored. The first

*Co-primary authors



This work is licensed under a Creative Commons Attribution International 4.0 License.

SoCC '24, November 20–22, 2024, Redmond, WA, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1286-9/24/11.

<https://doi.org/10.1145/3698038.3698544>




Figure 1: The target system architecture of SkyBox. It facilitates the deployment of renewable energy-based modular data centers across multiple geographical regions at scale.

approach is to transmit energy (a mix of renewable and non-renewable energy) via transmission lines to data centers. This approach incurs significant monetary cost, additional carbon footprint, and increased complexity of power management, due to the uncontrollable mix of renewable and non-renewable energy sources [2]. Specifically, about half of the cost is due to the power transmission and distribution [37]. As an alternative, batteries are used to store power and supply at a later time. However, they are expensive for large-scale datacenter deployment and minuscule in scale [77]. For instance, the battery capacity in the US is only 0.4% of the overall solar and wind capacity [20, 34].

To address these issues, both academic and industry community [23, 36, 39, 72] proposed to move data centers closer to renewable energy farms and powering data centers using renewable energy. In this paper, we focus on modular data centers (MDCs), and co-locate MDCs with renewable energy farms as shown in Figure 1. This is because modular data centers have low construction cost, and offer great mobility across different regions [16, 39, 43]. The renewable energy-based modular data centers rMDCs can alleviate the heavy use of batteries and power transmission lines, and have a great potential to achieve zero emissions. Most recently, we have seen that datacenter vendors have been planning to deploy MDCs to meet the increasing demands on cloud computing while achieving the carbon-free goal [8, 36, 45, 72].

However, as we deploy modular data centers, the power variations in renewable energy farms create challenges. First, due to the power variability of renewable energy sources, the data center resources can be powered on or off as renewable power increases or decreases. This can cause applications to fail, interrupt, or migrate, slowing down their execution. Second, as we scale the deployment of renewable-based data centers across multiple regions, it is non-trivial to consider the volatility. It creates a large exploration space for identifying an optimized solution, as we have multiple geographically distributed sites with varying power over time. Third, the dynamic nature of the volatility patterns further complicates the challenges and creates the need for dynamic solutions.

To this end, we develop SkyBox, a framework that uses a holistic and learning-based approach to enabling the efficient

use of renewable energy with rMDC deployment at scale. With SkyBox, we aim to answer three research questions: (1) where to deploy rMDCs across geographical regions? (2) how to maximize the efficiency of renewable energy sources across multiple rMDCs? and (3) how to enable smart application placement and migrations across rMDCs to minimize the performance impact of the power variability of renewables. By answering them, we wish SkyBox will facilitate the deployment and operation of rMDCs at scale.

We drive the design of SkyBox with three main observations based on our study of the distribution and production patterns of energy sources generated at more than 500 energy farms in total. We first observe that not all renewable farms are good candidates to place rMDCs as some farms consistently do not have enough stable power output. Second, the power production can be complementary across multiple geographically distributed energy farms, and their aggregate power is more stable than each individual site. Third, while the variability is high, renewable energy production is predictable for a reasonable prediction horizon (i.e., 3–24 hours). This offers sufficient time for datacenter operators to identify power changes in advance and migrate workloads hours ahead for tolerating the power variability.

SkyBox proposes three main techniques leveraging the aforementioned insights. (1) Site pruning for reducing the exploration space when deploying rMDCs. It quantifies the coefficient of variation of the power production of renewable energy farms based on their historical traces. This enables datacenter operators to identify the sites that have relatively stable power supply and capacity. (2) Subgraph identification for pinpointing rMDC locations that have stable aggregated power production. Within the same subgraph, the power production pattern of each renewable energy farm is complementary. Therefore, the subgraph will deliver stable aggregated power production as a whole. (3) Smart virtual machine (VM) placement and migration for minimizing the impact of power variability, by developing a Mixed-Integer Program (MIP) model for enabling optimized VM placement and migration across rMDCs.

Specifically, as building rMDCs at every renewable farm is neither practical nor needed (e.g., Europe alone has over 550 renewable farms), we locate a subset of renewable farms to deploy rMDCs (site pruning). We first analyze the historical power traces to select farms that exhibit low power variability and complementary power production patterns. After the site pruning, we could identify a small set of locations for rMDCs, where each rMDC is represented by a time series of “compute capacity” instead of a constant compute capacity. Therefore, we can capture and represent the dynamic changes of its co-located renewable energy farm.

Second, to tolerate the power variability, we run applications on this fleet of geographically distributed rMDCs as

a single resource with stable aggregated power. We partition the possible sites for rMDC deployment into independent subgraphs (subgraph identification). Based on these subgraphs, SkyBox can help cloud platform vendors decide the locations for rMDC deployment. With the SkyBox framework, cloud vendors can also periodically re-identify the subgraphs to ensure each subgraph has stable aggregated power production, and migrate cloud applications across rMDCs if needed based on their power availability.

Third, as cloud platform operators manage a diverse set of cloud applications, SkyBox assigns each cloud application or VM to a subgraph and then to a rMDC within the subgraph (VM placement and migration). We make smart decisions with our learning-based approaches to minimize the impact of interruptions caused by the power variability. SkyBox formulates the VM placement and migration problem into an optimization problem using the Mixed-Integer Program (MIP) model [38], as it can make optimized decisions for a large set of VMs and rMDCs with a global perspective. SkyBox takes VM properties (e.g., regular or evictable) into consideration to match each VM with its best-fit rMDC, and can effectively handle dynamic power supply patterns by leveraging the predictability of renewable power supply.

We implement SkyBox framework by enabling the replay of the power variations for each rMDC with real-world power traces from renewable energy farms, and the replay of the VM execution and migration across rMDCs with real-world VM traces from cloud platforms [9]. Our evaluation shows that, with careful selection of rMDC sites, and proper placement of VMs in rMDCs, SkyBox can reduce the total carbon footprint by 46% with low VM migration frequency, in comparison with conventional datacenter deployment approaches. Therefore, we believe that SkyBox can complement current data centers to meet the compute demand in a more sustainable manner.

As modular data centers have lower construction costs compared to conventional data centers, in combination with the minimal use of batteries, SkyBox helps datacenter vendors identify the rMDCs deployment that incurs low embodied carbon footprint. With its minimized usage of the power grid, it also has a minimal operational carbon footprint. As SkyBox utilizes the aggregated power production of a few stable energy farms, it minimizes the overprovisioning of power and compute resources, which further reduces both the embodied and operational carbon emission. Overall, we make the following contributions in the paper:

- We conduct a characterization study of renewable energy with power traces from real-world energy farms (§2).
- We design a site-pruning technique using historical traces to locate viable renewable sites for rMDCs (§4.2).




Figure 2: Power variation in renewable energy farms.

- We propose a subgraph identification technique that can identify small complementary subgroups of energy farms based on the prediction of their power variability (§4.3).
- We develop optimization techniques for smart VM placement and migration to minimize the negative impact of power variability on VM performance and maximize the efficiency of rMDCs (§4.4).
- We implement SkyBox framework to facilitate the deployment of rMDCs at scale, and conduct a detailed evaluation using real-world renewable energy and VM traces (§5).


2 CHARACTERIZATION OF RENEWABLE ENERGY

To motivate the design of SkyBox, we first characterize the variability of renewable energy sources. The production of renewable energy depends on various factors. Without loss of generality, we quantify the variability of two main renewable energy sources, solar and wind, by analyzing two representative years-long datasets: (1) EMHIRES dataset [44, 55] with traces from 555 sites in Europe and (2) ELIA dataset [32] from 25 sites in Belgium that includes per-site power production traces and forecasts.

As expected, we observe significant variation across space (different farm locations), time (times of the day and seasons), and power resources (solar and wind). Figure 2 shows a 4-day sample of solar production, normalized to the maximal energy production capacity. Solar energy follows a periodic diurnal pattern, but days overcast with heavy clouds can significantly reduce the peak production (3.5% vs. 77% in the following day), and days with variable cloud patterns cause spiky energy production. And different energy sources exhibit different patterns. The wind energy production for the same duration exhibits sharp peaks and valleys (depending on the weather conditions), but rarely drops to zero.

2.1 Complementary Variability Patterns

Despite the large variability in a single renewable energy farm, we observe that these variability patterns are often complementary among different renewable energy farms across a geographical region. This complementary pattern



(a) Complementary power production of three sites. (b) Duration of the complementary patterns.

Figure 3: Reducing the variability in renewable energy production by aggregating multiple sites.

can be generated by using different energy sources (e.g., wind vs. solar), and geographical locations with different impacts of micro-climates and weather (e.g., same solar source but in two different locations, one of which is covered by clouds).

Figure 3a demonstrates such an example from the EMHIRES dataset. To gauge the variability of power produced over time, we use the coefficient of variation (CoV) of renewable power produced at different timestamps as the metric. CoV is the standard deviation divided by the mean, so a higher CoV implies more variability. In Figure 3a, by combining the energy sources from one solar farm (Solar_1) and one wind farm (Wind_1) we deliver more stable energy than the single solar farm (5.6× lower CoV). Adding another wind farm (Wind_2) further reduces the variability and decreases the value of CoV by an additional 1.5×. Thus, we reach that:

Observation 1: Power variability in renewables can be masked by selecting a subgroup of complementary sites.

We further analyze how long the complementary sites can remain complementary. We define a group of sites as complementary sites if their aggregated power production is stable ($CoV < 0.4$) for at least one week. Our results show that the complementary sites remain stable for 14.4 days on average and up to 44 days, as shown in Figure 3b. For the example shown in Figure 3a, the three renewable sites can maintain a complementary pattern for 30 days. Therefore, our analysis suggests that we can re-identify complementary sites periodically but not frequently.

Observation 2: The complementary sites could preserve the complementarity for 2 weeks on average and up to 6 weeks.

2.2 Predictability of Renewables

The main cause of power variability is the weather conditions, which can be predicted accurately a few hours ahead. We show the power forecasts provided in the ELIA dataset [32] (based on the weather forecasts) in Figure 4. The predictions for near-future power production are accurate enough to capture important trends. The mean absolute percentage error (MAPE) for the next 3-hour predictions is 8.5–9.0%, for day-ahead predictions is 18–25%, and for week-ahead predictions is 44%–75%. This predictability provides insights into

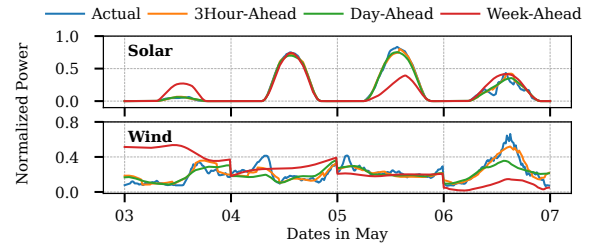


Figure 4: Energy prediction of solar and wind in near (e.g., 3-hour-ahead) and far-away future (week-ahead).

how the power resources will change, and when will workload migrations be needed (see the detailed discussion in §4.4). As the power changes can be predicted within several hours to one day ahead, we have sufficient time to migrate applications to further tolerate the power variability.

Observation 3: Renewable energy is predictable for a reasonable prediction horizon (3–24 hours) in the future.

3 BACKGROUND OF MODULAR DATA CENTERS

To facilitate our discussion on rMDC deployment, we now present the technical background of modular data centers.

3.1 Modular Data Centers

Modular data centers (MDCs) attract much attention from datacenter vendors, as they have low construction cost and installation time [51, 89, 99]. Many cloud vendors have also deployed MDCs in production [53, 73, 86]. An MDC organizes the server racks, cooling system, power supply, and batteries in one or more containers, which provides convenience for installation and shipping. In Table 1, we list the configurations and costs of an example MDC FusionModule-2000 [94]. A typical MDC consists of 10 racks (each rack with 15 servers), 3 battery cabinets with a default 15-minute backup time, and 3 cooling containers. Among all the components, server racks take a major portion of the embodied carbon (i.e., carbon emission of manufacturing, construction, and shipping), installation cost, and building footprint. The battery cabinets have less installation cost and building footprint, as they are usually used for emergency backups with a small capacity. It is worth noting that the MDC construction is much less affected by its location compared to traditional site-built data centers, as the MDCs are often organized into standard, prefabricated containers off-site [51, 89, 99].

3.2 Renewable Energy-based MDCs

Modular data centers are well-suited to collocate with renewable energy farms for their flexibility and low construction costs. Major datacenter providers have recently invested in

Table 1: A summary of the core components in a modular data center, and their cost and characteristics.

Modular DC	Server Rack	Cooling	Battery
Power (<i>kW</i>)	150	35	-
Embodied Footprint (<i>tCO₂e</i>)	88.7	3.7	5.5
Cost (\$)	450K	10.9K	46.9K
Footprint (<i>m²</i>)	20.6	5.8	5.8
Capacity	10 racks	-	15 mins

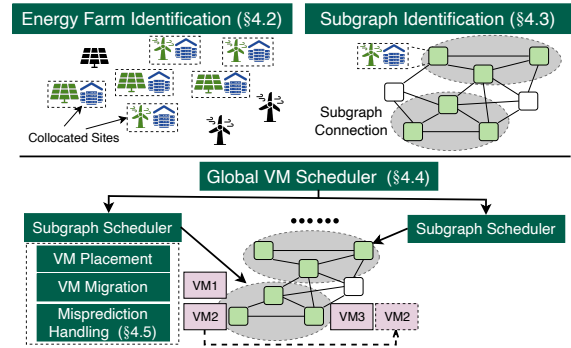
building rMDCs that can be easily deployed alongside renewable energy farms [52, 69]. However, the key challenge with collocation is the variability of the renewable energy production (see §2). To cope with this, current studies usually leverage batteries and/or power grids. We discuss their pros and cons as follows.

Batteries. Batteries can tolerate the volatility of renewable power supply [2, 68, 96] by storing the surplus renewable energy and discharging it when the renewable power is insufficient. However, deploying batteries in rMDCs faces two critical challenges. First, the current rMDC battery capacity is insufficient to tolerate the renewable energy production variation (up to several hours). Second, increasing the battery size will increase the cost and building footprint. Considering the commercial-scale battery (\$1250/*kWh* [77]) with a one-hour backup time (for servers in Table 1), the hardware and installation cost accounts for >30% of the total rMDC cost, and the building footprint reaches >40% of the total footprint. Therefore, it is less practical to fully rely on batteries to tolerate long-lasting renewable power supply fluctuations.

Power grid. The power grid serves as a backup energy source for rMDCs, as it is a more stable power supply. Drawing energy from the power grid helps the rMDC keep servers running when the renewable power supply drops. However, the grid mainly supplies power from non-renewable energy sources today, which have significantly larger carbon intensity [48]. Therefore, extensive power grid use will increase the operational carbon footprint of the rMDC. In our work, we aim to minimize the use of power grid in rMDCs.

4 DESIGN AND IMPLEMENTATION

The goal of SkyBox is to best utilize renewable energy when collocating MDCs with renewable energy farms while minimizing the negative impact on application performance. SkyBox has a comprehensive consideration of different energy sources. SkyBox uses renewable energy as its primary source and the power grid (non-renewable) as a backup when renewable energy becomes insufficient. Since the power grid usage is a major source of operation carbon footprint, SkyBox maximizes the use of renewable energy for rMDCs and minimizes the use of power grid and batteries to reduce carbon emissions when renewable energy is insufficient.

**Figure 5: System overview of SkyBox.**

4.1 SkyBox Overview

SkyBox takes a holistic and learning-based approach. It has four core components as shown in Figure 5.

- **Identification of renewable energy farms:** To decide where to deploy rMDCs, SkyBox uses the stability of power supply as the key metric. By collocating rMDCs with stable renewable energy farms, it ensures a reliable power supply, and also reduces embodied carbon as it can overprovision fewer servers (§4.2).
- **Subgraph identification of complementary sites:** Our study reveals that a small group of renewable energy farms with complementary patterns can produce stable aggregated power supply (§2.1). SkyBox identifies a set of subgraphs from the selected rMDCs, where each of them has complementary renewable energy sources. With the subgraph candidates, we can select ones with the stablest power supply (§4.3).
- **VM placement and migration:** Since each individual rMDC of a stable subgraph may still face power variability, SkyBox performs optimized VM placement and migration within the subgraph. Utilizing the predictability of renewable energy production (see §2.2) and VM characteristics (e.g., lifetime), SkyBox places VMs on stabler rMDC and performs VM migrations from rMDCs with insufficient power to those with excessive power. The optimized VM placement and migration minimizes the usage of power grid and reduces the operational carbon footprint (§4.4).
- **Misprediction handler:** Predictions of the future power supply and VM lifetime may incur misprediction. SkyBox gracefully handles the mispredictions and it incurs minimal overhead (§4.5).

4.2 Identification of Energy Farms for rMDCs

SkyBox first decides which renewable energy farm to collocate with each rMDC. There are a large number of renewable farms, and not all of them are required or suitable for building




Figure 6: Improved power stability and reduced embodied carbon footprint (lower is better).

a rMDC. For example, there are over 500 farms in Europe currently [44, 55], so it is unnecessary to deploy rMDCs in each farm site. Typically many factors govern which locations are best for deploying data centers, they include the proximity to users, operating cost, geography, and others [15, 40, 92]. However, the stability and complementarity of power production of each renewable farm are playing the most crucial role in determining the rMDC sites [2, 45]. This is because the available computing resources of rMDCs are determined by the fluctuated power production, which in turn affects the services of cloud applications. SkyBox focuses on this technical aspect, with the goal of assisting cloud platform vendors decide the rMDC locations.

SkyBox prefers renewable energy farms with stable power supply. This is for two major reasons. First, collocating rMDCs with stable farms guarantees higher resource availability, which reduces the power grid usage and incurs less VM outages. Second, the collocation scheme helps reduce the embodied carbon footprint and the demand for batteries. This is because rMDC with a higher power supply fluctuation has to overprovision additional server capacity to utilize the peak renewable energy production. rMDCs with more stable power supply require less server capacity and battery capacity, therefore mitigating the embodied carbon footprint and construction cost.

SkyBox uses the coefficient of variation (CoV) of the power production as the key metric to rank the farms. Figure 6 (left) shows the CoV of the three selected energy farms is much less (51%) than that of the selected energy farms. In comparison with rMDCs supplied with stable power sources (e.g., power grid), Figure 6 (right) shows the additional embodied carbon cost of building rMDCs of these selected farm sites is much less than that of unselected farms (see the detailed procedure of calculating carbon emissions in §5.1). This is because these rMDCs need fewer servers for the overprovisioning for tolerating the power variability of collocated energy farms. Although SkyBox causes embodied carbon footprint for constructing rMDCs, it significantly reduces the total carbon emissions by reducing the operational carbon footprint. We will discuss it in details in §5.

The site identification of SkyBox selects rMDCs mainly based on the stability of renewable energy production since this is the main concern of using renewable energy. Note

that SkyBox is flexible to integrate constraints specified by cloud vendors (e.g., location constraints) in the site selection procedure. Furthermore, even if we do not select rMDCs with the stablest power supply, SkyBox can still reduce the total carbon footprint with its subgraph identification and optimized VM placement and migration (see §4.3 and §4.4).

4.3 Complementary Subgraph Identification

In §4.2, we identify a set of renewable energy farms with stable power supply to collocate with rMDCs. However, three major challenges remain. First, managing these rMDCs across multiple geographical regions will be complicated. It brings challenges to datacenter operators, especially considering the diverse power production patterns at different farms. Second, the selected farms individually cannot constantly generate stable power production over time. Third, as cloud providers place their VM workloads across all the rMDCs, the decision space could be extremely large. An ill-judged decision could cause the inefficiency of VMs and even unexpected VM outages due to the unstable power supply.

Key idea: To address those challenges, we identify disjoint subgraphs from these rMDCs. Each subgraph has a subset of renewable energy farms. To build a subgraph, we use three metrics: (1) *the CoV of the aggregated power production*; (2) *the minimum power production of the farms*; and (3) *the distance between sites*. The first and second metrics are based on the historical power production of each farm. They guarantee that we select the subgraphs with sufficient and stable aggregated power supply. The third metric ensures the selected subgraphs have low communication overheads. In practice, cloud vendors make sure their data centers are at least 100 miles apart to guarantee physical isolation. They should also be not too far away from each other to ensure low network overhead (e.g., less than 1,000 miles between European data centers of a major cloud vendor [54, 74]). Therefore, we set the upper bound limit for the distance between sites in a subgraph as 500 miles by default, so that the rMDC deployment will fall within this range. Note that this parameter is configurable. SkyBox can identify more complementary subgraphs with a larger upper limit, which offers the flexibility for cloud vendors to decide their exploration space.

We enumerate all possible subgraphs and rank them using the aforementioned three metrics. We iterate through the set of subgraphs in the rank order and select the subgraph that has no intersection with the previously selected subgraph. After that, we have a set of subgraphs that have complementary patterns.

Key study results: Exhaustively enumerating all the subgraph candidates can lead to non-trivial overhead, since the number of subgraph candidates grows combinatorially as we




Figure 7: The impact of different subgraph size K . It shows the stable aggregated power and its average CoV.

increase the number of farm sites and the size of subgraphs. Ideally, we wish to minimize the size of subgraphs, while retaining the complementary benefits, which can simplify the power management and task scheduling across data centers. **Small subgraphs are sufficient for rMDC deployment:** Identifying subgraphs from a large set of farm sites is a non-trivial problem. The number of subgraph choices grows combinatorially as we increase the number of sites and vary the size of subgraphs. Ideally, we wish to keep the size of subgraphs as small as possible, while retaining the complementary benefits. With a small subgraph, we can simplify the power management as well as the task scheduling across data centers. According to our study of the 555 farms (the EMHIRES dataset [44, 55]), we find that small subgraphs with 3 sites have complementary power patterns, and can provide a sustained power capacity. We show the aggregated power and average CoV of the power with different subgraph size K in Figure 7. In the best case (Ideal in Figure 7), all farm sites are grouped into a single subgraph, it has the highest stable aggregated power and the lowest variability (CoV). As we vary the size K in Figure 7, the stable aggregated power of the subgraphs with size $K = 3$ can reach to the 80–90% of the ideal stable power, and its CoV is also comparable to the ideal case. As we further increase the subgraph size, we will enlarge the exploration space. However, it does not bring more significant benefits. Therefore, we set $K = 3$ as the subgraph size in SkyBox. After the subgraph identification, we observe that the identified subgraph of SkyBox consists of (1) three wind farms or (2) a mix of solar and wind farms. SkyBox does not group three solar farms together into a subgraph, since this often has higher variability.

Subgraph identification is a dynamic problem: As the power patterns of a renewable site may change over time. Our study shows that subgraphs usually retain their complementarity for two to six weeks. Therefore, SkyBox can re-identify subgraphs that have complementary power patterns every two weeks or longer for stable aggregated power supply. Note that the subgraph re-identification procedure only defines how to group rMDCs into subgraphs, it will not incur any additional VM migration or suspension.

As each subgraph has stable aggregated power supplies from renewables, cloud vendors can deploy rMDCs close

to renewable energy farm sites in these subgraphs. Note that cloud vendors can select one or more subgraphs as the candidates for rMDCs deployment, or deploy rMDCs at the farm sites covered in all the subgraphs. After the rMDC deployment, we do not need to move the deployed rMDCs frequently (even upon subgraph re-identification). Instead, SkyBox suggests migrating cloud applications across deployed rMDCs for lower cost. And SkyBox develops a set of algorithms for cloud VM placement and migration (§4.4) to minimize its impact on the application performance.

SkyBox is still applicable if any regulations enforce certain constraints on certain rMDCs that should not form a subgraph for VM migrations. SkyBox may identify fewer subgraph candidates, but this does not affect its feasibility. SkyBox is a framework that helps cloud vendors explore and decide the deployment plans of modular data centers.

4.4 VM Placement and Migration

After identifying the subgraphs, we now place VMs across these subgraphs and their rMDCs. A simple approach is heuristic-based VM placement, such as greedily placing VMs onto the rMDCs in the order of their resource availability. However, heuristic-based approaches are often suboptimal, especially when the power supply changes over time, and different VMs have different priorities and demands for computing resources. Therefore, it is difficult for heuristic-based approaches to identify the best rMDC for each VM.

SkyBox formulates the VM placement and migration problem into an optimization problem using the Mixed-Integer Program (MIP) model [38]. SkyBox employs the MIP model for three reasons. First, the MIP model can make optimized decisions for a large set of VMs and rMDCs with a global perspective. Second, it can take VM properties (e.g., regular or evictable) into consideration to match each VM with its best-fit rMDC. Third, the MIP model can effectively handle the dynamic power supply patterns of rMDCs by leveraging the predictability of renewable power supply to find optimized solutions in advance.

SkyBox employs a hierarchical approach for its VM placement and migration policies. It takes three steps to make decisions: (1) VM placement among subgraphs, (2) VM placement within each subgraph, and (3) VM migration across rMDCs within a subgraph.

VM placement among subgraphs: SkyBox first relies on heuristic-based approaches for VM placement across subgraphs. This is because the aggregated power supply of each subgraph is relatively stable. We empirically find that the heuristic-based approach can handle such a scenario well. Specifically, SkyBox employs the best-fit placement algorithm. We maintain a list of subgraphs ordered by the amount

Table 2: Input constants of our MIP model.

Symbol	Interpretation
M, D, N, \mathbb{K}_n, T	A set of M regular VMs, D delay-insensitive VMs, N subgraphs, K rMDCs for subgraph n , time interval T
$Power_{mt}$	The power draw of VM m at time t
Mem_m	The memory size of VM m
$Lifetime_m$	The predicted lifetime of VM m
RS_{nkt}	Renewable power supplied to the k th rMDCs of subgraph n at time t
$Power_{Migr}$	Power consumption of migration (per GB VM states)
CI_{NR}, CI_R	Carbon Intensity of non-renewable energy and renewable energy [2]

of excess power. When a VM arrives, we place it on the subgraph with the highest amount of resources available. If the subgraph has no excess power, we remove it from the list.

$$\text{O1 Carbon: } \min \sum_{n \in \mathbb{N}, k \in \mathbb{K}_n, t \in T} NR_{nkt} \cdot CI_{NR} + RU_{nkt} \cdot CI_R \quad (1)$$

$$\text{O2 Uptime: } \max \sum_{m \in \mathbb{D}} \frac{Lifetime_m}{Lifetime_m + Downtime_m}, \quad \text{s.t.}$$

$$\text{C1 Power: } \forall n \in \mathbb{N}, k \in \mathbb{K}_n, t \in T :$$

$$RU_{nkt} \leq RS_{nkt}$$

$$Consum_{nkt} = NR_{nkt} + RU_{nkt} = \sum_{m \in \text{MUD}} X_{mnt} \cdot Power_{mt}$$

Formula 1: The key objectives and model constraint.

VM placement among rMDCs in a subgraph: After identifying the suitable subgraph, SkyBox decides where to place the VM among rMDCs within each subgraph. Since the renewable power supply for each individual rMDC is less stable than the aggregated power supply of each subgraph, the VM placement problem across rMDCs is more dynamic and complex. Therefore, we do not use a heuristic-based approach, but instead formulate it into an optimization problem using the MIP model. We define the input constants and variables of MIP model in Table 2 and Table 3 and explain them below. **MIP input:** The MIP model takes the following inputs: (1) VM resource configuration, including its memory size (Mem_m), its VM category (regular or evictable¹), and its power consumption ($Power_{mt}$), which can be estimated by the number of vCPUs and its utilization [57]; (2) Estimated VM lifetime ($Lifetime_m$), based on the insights from prior studies [12, 25, 50] that VM lifetime can be estimated with high accuracy using VM properties (e.g., VM type and hosted application type); (3) Current and predicted power supply of rMDCs (RS_{nkt}), with the high predictability of renewable power supply (see §2.2). SkyBox has the strong robustness of handling mispredictions, as discussed in §4.5 and §5.7.

¹Cloud platforms usually offer evictable VMs that run at a much lower price and priority than regular VMs. They can be evicted if needed [5, 70].

Table 3: Variables used in our MIP model.

Symbol	Domain	Interpretation
X_{mnkt}	$\{0, 1\}$	Whether VM m is powered up in the k th rMDC of the subgraph n at time t
$M_{mnk_1k_2t}$	$\{0, 1\}$	Whether VM m is migrated from the rMDC k_1 to k_2 of the subgraph n at time t
NR_{nkt}	$\mathbb{R}_{\geq 0}$	Non-renewable power used by the k th rMDCs of the subgraph n at time t
RU_{nkt}	$\mathbb{R}_{\geq 0}$	Renewable power used by the k th rMDCs of the subgraph n at time t
$Consum_{nkt}$	$\mathbb{R}_{\geq 0}$	Power consumption in the k th rMDCs of the subgraph n at time t
$Downtime_m$	$\mathbb{R}_{\geq 0}$	The actual downtime of VM m

These MIP inputs are essential in the MIP model, as they inform the model about the VM energy consumption and renewable energy supply, which helps the MIP model limit the number of VMs that can run in rMDCs.

MIP output: For VM placement within subgraphs, the MIP model denotes placement decisions as X_{mnkt} (see Table 3). It sets $X_{mnkt} = 1$, if VM m is powered on at k th rMDC in subgraph n at time t . The MIP model also decides the non-renewable energy usage from the power grid. We use NR_{nkt} to represent the amount of non-renewable power supply to the k th rMDC in subgraph n at time t through the grid.

MIP objectives: SkyBox aims to minimize the total carbon footprint and maximize the VM uptime. To achieve these goals, we define two objectives in our MIP model. The first objective (O1 Carbon in Formula 1) is defined to minimize the non-renewable energy usage of rMDCs over the course of time, since it is the primary cause of the operational carbon footprint (i.e., the carbon intensity of non-renewable energy CI_{NR} is significantly greater than that of renewable energy CI_R). The second objective maximizes the uptime percentage for evictable VMs (O2 Uptime). This objective is defined to make sure that the evictable VMs will also have reasonable good performance when the renewable energy supply is unstable. As for regular VMs, SkyBox guarantees that they will never be actively powered off due to the insufficient renewable power supply.

MIP variables and constraints: We list one of the key constraints of our MIP model in Formula 1. The constraint enforces that the total power consumption of all running VMs ($Consum_{nkt}$) on each rMDC cannot exceed the total power supply (C1 Power) that includes both renewable energy production and the non-renewable energy from the power grid. **VM migration across rMDCs in a subgraph:** If an rMDC experiences a decrease of renewable power supply, SkyBox can migrate some of its VMs to another rMDC within the same subgraph that has excessive power supply. Therefore, we integrate a power model to quantify the overhead of VM migration, where its power consumption is proportional

to the size of VM states migrated [64]. Larger VMs incur higher overhead to the total power consumption and are more sensitive to migrations. Even though the VM live migration introduces extra power consumption, it only incurs a millisecond-level of VM downtime [63, 83]. Given the low migration frequency of SkyBox (0.015 times per VM per hour on average, see Figure 14), it will not incur much performance overhead to the VMs, especially for the VMs running batch processing and throughput-oriented workloads.

Additional MIP output: VM migration is represented by the symbol $M_{mnk_1k_2t}$ (Table 3) in our MIP model. $M_{mnk_1k_2t}$ denotes whether a VM m should be migrated from k_1 th rMDC to k_2 th rMDC of subgraph n at time t . Using the same objectives (O1 and O2), the MIP model will identify an optimized VM migration plan for each rMDC.

Additional MIP variables and constraints: The MIP model uses $M_{mnk_1k_2t}$ and input constant $Power_{Migr}$ to estimate the migration overhead of a VM m . For each VM, its migration overhead counts toward the total power consumption of both migration source and target rMDC (C1' Power of Formula 2). Following the prior constraint (C1 Power), the total consumption cannot exceed the total power supply.

C1' Power: $\forall n \in \mathbb{N}, k \in \mathbb{K}_n, t \in \mathbb{T}$:

$$Consum_{nkt} = \sum_{m \in \text{MUD}} [X_{mnkt} \cdot Power_{mt} + Power_{Migr} \cdot Mem_m \cdot (\sum_{k_1 \in \mathbb{K}_n} M_{mnk_1kt} + \sum_{k_2 \in \mathbb{K}_n} M_{mnk_2kt})] \quad (2)$$

Formula 2: The additional key constraint of the model.

Sources of benefits: We expect that our MIP model can make optimized decisions by exploiting the following four sources of benefits: (1) VMs with higher migration overheads can be placed at rMDCs with more stable predicted renewable power supplies. (2) VMs with longer lifetimes can be scheduled to rMDCs with longer predicted renewable energy availability. The MIP model prefer to place VMs with shorter lifetimes to less stable farms, such as the solar farms. (3) We prefer to choose VMs with lower migration overheads to migrate when possible. And (4) we use non-renewable energy from the power grid for rMDCs with unstable renewable power supplies.

4.5 Misprediction Handling

Our MIP model relies on the prediction of renewable power production and VM lifetimes, which may be slightly inaccurate. For instance, a 3-hour-ahead prediction has 9% error on average. To detect such mispredictions at runtime, SkyBox compares the real renewable energy production at each rMDC against its predicted value. Similarly, we also obtain the prediction errors for the VM lifetime prediction. In




Figure 8: The misprediction handling mechanisms.

Table 4: Misprediction handling for different scenarios.

ID	VM Lifetime	Renewable Power	Possible Exceptions
①	Over-predict	Under-predict	N/A
②	Over-predict	Over-predict	Power Deficiency
③	Under-predict	Under-predict	Power Deficiency, Extra VM Lifetime
④	Under-predict	Over-predict	Power Deficiency, Extra VM Lifetime

SkyBox, minor prediction errors are tolerable. While the actual power supply of an rMDC may be lower than predicted, it may still exceed the actual power consumption. And large mispredictions may cause extra power grid usage or evictable VM outages. Therefore, we design a misprediction handling mechanism to minimize their impact.

We categorize the misprediction into two types: under-prediction and over-prediction, where an under-predicted value is smaller than its actual value, and vice versa. For example, an under-predicted power production means more actual power is produced than predicted, and an over-predicted VM lifetime means the VM's actual lifetime is shorter than predicted. We present different misprediction scenarios in Table 4, and the misprediction handling workflow in Figure 8.

In ①, the actual VM lifetime is shorter than predicted, and the actual power supply is larger than predicted. Based on the predicted values, the MIP model may place fewer VMs on the rMDC than it can actually support, and these VMs may run shorter than its predicted lifetime. In this case, the power consumption never exceeds the power supply. The rMDC will have excessive power, and it can serve as the migration target of other rMDCs. Thus, no specific handling is needed.

In ②, both the VM lifetime and the renewable power supply are over-predicted. Since the power supply is over-predicted, our MIP model may place more VMs on the rMDC than it can support, resulting in a power deficiency at the rMDC (i.e., the power supply at an rMDC cannot meet its total power consumption). To handle this, we employ three techniques: *VM migration*, *evictable VM shutdown*, and *power transmission from the grid*. First, we prioritize VM migration since it introduces minimal VM downtime and carbon overhead. To perform migration, we select another rMDC which has the most excessive power supply in the same subgraph as the migration target. If such an rMDC exists, we will migrate

VMs to this target rMDC until it cannot host more VMs, or the power consumption of the source rMDC falls below the power supply. We prioritize the VMs with less amount of VM states to migrate, as they incur lower migration overheads. If the power consumption still exceeds the power supply, we will power off evictable VMs until their VM uptime percentage is below a predefined threshold (empirically set to 90%). Finally, SkyBox will use the power grid if still necessary.

In ③, the actual VM lifetime is longer than its prediction. The MIP model might place VMs onto the rMDC that does not have a sufficient power supply for its entire lifetime. This problem can be handled with similar mechanisms in ②. Since the MIP model receives the under-predicted VM lifetime as its input, it cannot make placement and migration decisions for the rest of their actual lifetimes. To handle this issue, SkyBox adds an offset to the VM lifetime prediction when we detect that it is under-predicted. This offset is adjusted using the average difference between the predicted and the actual lifetime of completed VMs. SkyBox periodically adjusts its VM lifetime prediction on an hourly basis, which is aligned with the decision interval of its MIP model.

In ④, it may incur similar problems caused by an over-predicted power supply in ② and an under-predicted VM lifetime in ③. We handle them in the same manner as we have discussed above.

SkyBox can tolerate spikes in VM utilization, as we can employ the same misprediction handling mechanism as discussed above. This is because the impact of an under-predicted VM utilization on SkyBox is similar to an over-predicted renewable power supply and vice versa.

4.6 SkyBox Implementation

We develop SkyBox framework by enabling the replay of the power variations for each rMDC with real-world power traces from renewable energy farms, and the replay of the VM execution and migration across rMDCs with real-world VM traces from data centers [9]. It takes renewable power traces and VM traces as its inputs and replays the power supply in each rMDC and the VM events (i.e., VM deployment, migration, or shutdown). At runtime, SkyBox updates the power production at each rMDC, triggers VM placements, and decides VM migration. The decision interval of SkyBox is configurable. We use one hour by default to make it aligned with the granularity of the real-world power traces.

We use Gurobi [49] to implement our MIP model for VM placement and migration, the MIP model uses a 3-hour-ahead prediction for its inputs such as the future power supply. The migration decisions of each subgraph are executed concurrently. SkyBox is not limited to power production at a coarse time scale, it can also work with a fine time scale. The only

reason we use hour-scale power production data is due to the time granularity of the open-sourced power traces.

We also developed a prototype system on a six-node cluster using the OpenStack cloud framework that live-migrates VMs and executes real applications in each VM. Our prototype system emulates the WAN setup with six nodes from two CloudLab clusters [31] (one from UMass and one from Clemson). We use PowerTOP [30] to track the power consumption on each server, which is used to limit the total power consumption and estimate the carbon footprint. The prototype system helps validate the accuracy of our SkyBox testbed and measure the performance impact of VM migrations on real applications. Our experimental results with various compute- and memory-intensive applications (e.g., graph computing with GraphBIG and Redis benchmarks) show that our testbed has high accuracy (> 97% correctness in measuring VM uptime) and the VM migrations do not introduce much performance overhead (less than 8%).

5 SKYBOX EVALUATION

Our evaluation shows that: (1) SkyBox reduces the total carbon footprint by 46% (§5.2) with reduced monetary cost (§5.3), and minimizes its impact on VMs, compared with baseline rMDC placement policies (§5.4); (2) It remains performant with large-scale rMDCs deployment, different battery capacities, and various VM workloads (§5.5); (3) It minimizes the use of power grid for reduced carbon emissions (§5.6); and (4) it shows strong robustness to the mispredictions of renewable power production and VM lifetime (§5.7).

5.1 Experimental Setup

We use the EMHIRES dataset [44, 55] to obtain the power production data of renewable farms. We select a representative eight-week period of the dataset, which includes different time periods when the variability of renewable energy production is high (e.g., Week 4 and 7 in Figure 9) and low (e.g., Week 1 and 8). The variability of the eight-week period is also centered around the yearly average value. Similar to the prior work [41, 42], we take a constant portion of the power generated from each renewable energy farm, such that its maximum power capacity equals to the maximum power consumption of rMDC servers. For each rMDC, we keep its original 15-minute battery capacity.

We use the VM workloads from the Azure Cloud Dataset [9]. The dataset has various VM properties, including the VM configurations (e.g., memory size), VM lifetime, and its category (regular or evictable). The VM lifetimes in the dataset range from a few hours to a few weeks. We inject prediction errors to the VM lifetime prediction to evaluate the robustness of SkyBox (see §5.7). Their arrival rates maintain a high power utilization (90%) over datacenter servers [41, 42]. By




Figure 9: Carbon footprint with different rMDC placements in eight weeks (Emb.: Embodied; Op.: Operational).

Table 5: Key parameters in the evaluation of SkyBox.

Category	Component	Value	Lifetime
Cost [13, 16, 77]	Server	\$3,000 per server	4 years
	Battery	\$1,250 per kWh	10 years
	Power Trans.	\$300K per km	20 years
	Construction (Cooling, electricity, etc.)	\$10 per watt	20 years
Carbon Intensity [48]	Solar	41 gCO ₂ eq per kWh	-
	Wind	11 gCO ₂ eq per kWh	-
	Brown	700 gCO ₂ eq per kWh	-
Embodied Footprint [33, 35, 82]	Server	591 kgCO ₂ eq per server	4 years
	Battery	146 kgCO ₂ eq per kWh	10 years
	Cooling Facility	50 kgCO ₂ eq per m ²	20 years

default, the VM dataset has 10% evictable VMs. We vary the proportion of evictable VMs in our sensitivity analysis (§5.5). **rMDC setup:** We first obtain the top six rMDCs based on the CoV of the power supply and group them into two stablest subgraphs (see §4.2 and §4.3). Each rMDC follows the configurations shown in Table 1. We further evaluate SkyBox with an increasing number of rMDCs in §5.5.

Carbon footprint model: We quantify the total carbon footprint with three categories: amortized embodied carbon, operational carbon from the power grid, and operational carbon from renewable energy. The embodied carbon footprint is measured with servers, batteries, and cooling facilities, and it is amortized over their lifetime. We list their details in Table 5. The operational carbon footprint is measured by the product of total energy consumption (from wind, solar, or power grid) and the carbon intensity (see Table 5).

Baseline policies: We compare SkyBox against baseline rMDC placement policies. They differ in the way that rMDCs are placed with multiple energy sources (i.e., renewable energy, battery, and power grid). For all baseline policies, rMDCs prioritize to use renewable energy and use the power grid as a backup. Note that we do not compare non-renewable-based data centers, as they incur a high operational carbon footprint. We summarize the baseline policies as follows:

- **Centr-Global** deploys a centralized rMDC at the geometric center of all the selected renewable energy farms and connects it to all the farms with power transmission lines.
- **Centr-Graph** is similar to Centr-Global, but it deploys a centralized rMDC per subgraph. They are ideal policies as they can aggregate all the renewable energy into a rMDC.




Figure 10: Comparison of the cumulative carbon footprint with different rMDC placement policies.

- **Distr-Grid** co-locates the rMDC with each renewable energy farm. Each rMDC handles its VM workloads independently without migrating VMs.
- **Distr-Battery** adds extra batteries based on Distr-Grid. The battery charges excessive renewable power and discharges when renewable power is insufficient. We set the extra battery capacity of each rMDC such that it can sustain one hour of server operation [58]. We further vary its battery capacity in the sensitivity analysis (§5.5).

SkyBox is similar to Distr-Grid but live-migrates VMs within each subgraph from rMDCs having insufficient power to those with excessive power supply using the MIP model. We also compare SkyBox with its variants to show the benefit of each of its component: **SkyBox-NoSI** uses a random site selection policy; **SkyBox-NoSG** uses random subgraph identification; and **SkyBox-BestEffort** uses the best-effort VM placement policy instead of the MIP model. SkyBox-BestEffort prioritizes migrating VMs to the rMDC with the most excessive power production, suspends evictable VMs if needed, and uses the power grid as the last choice.

5.2 Carbon Footprint Reduction

We first evaluate SkyBox in reducing the carbon footprint. We show the cumulative carbon footprint of different rMDC placement policies over eight weeks in Figure 10. Compared to Distr-Grid and Distr-Battery, SkyBox produces 46% and 39% less total carbon footprint, respectively. This is because (1) SkyBox groups rMDCs into subgraphs, which share a more stable aggregated power supply than each individual




Figure 11: Breakdown of carbon footprint of subgraphs.

Table 6: Breakdown of the amortized monetary cost of different data center placement policies over one year.

Category	Centr-Global	Centr-Graph	Distr-Grid	Distr-Battery	SkyBox
Servers	\$562K	\$614K	\$675K	\$675K	\$675K
Battery	\$23K	\$25K	\$28K	\$141K	\$28K
Power Trans.	\$62M	\$61M	\$45K	\$45K	\$45K
Construction	\$375K	\$410K	\$450K	\$450K	\$450K
Total	~\$63.0M	~\$62.0M	~\$1.2M	~\$1.3M	~\$1.2M

rMDC; and (2) it performs effective VM migration within each subgraph with the predictions of the power supply.

We also compare SkyBox with the two centralized rMDC placement policies: Centr-Global and Centr-Graph. Since Centr-Global can aggregate all renewable energy production into a centralized data center, it has the most stable power supply. Compared to this ideal placement policy, SkyBox has slightly higher operational carbon, as it incurs extra migration overhead. SkyBox also produces slightly more embodied carbon, as it needs to provision more servers to make use of the peak power supply. Compared to Centr-Graph, SkyBox has almost the same (0.1% more) total carbon footprint. This is because Centr-Graph statically connects rMDCs with the renewable power farms in each subgraph. While these two baseline policies have a similar carbon footprint as that of SkyBox, their monetary costs are several times higher than that of SkyBox (see §5.3), making them less attractive.

SkyBox achieves more carbon footprint reduction, when the power production is unstable. As shown in Figure 9, SkyBox obtains more benefits during the fifth to the seventh week, when the CoV of energy production is high (i.e., 51%–64% reduction compared to Distr-Grid). This is because the baselines require more power grid usage when renewable production is unstable, causing a high operational carbon footprint. In contrast, SkyBox can perform optimized VM placement and migration, which significantly reduces the operational carbon footprint caused by the use of the power grid, as shown in the carbon breakdown in Figure 11.

As the power grid becomes more renewable in the future (e.g., its carbon intensity is projected to decrease by 20 ~ 30% in 2050 [3]), SkyBox can further reduce its total carbon footprint. SkyBox still produces 38% and 30% less carbon

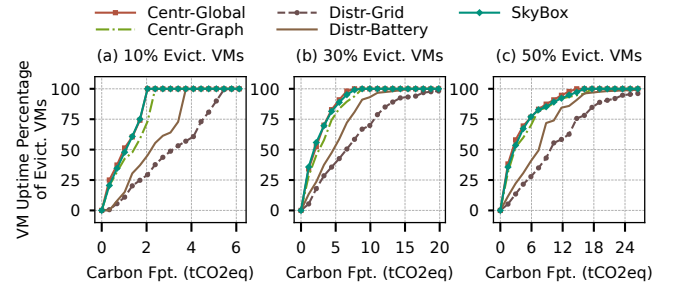


Figure 12: Efficiency of rMDCs with different percentages of evictable VMs (quantified with evictable VM uptime).

than the baseline placement policies Distr-Grid and Distr-Battery. This is because the operation carbon footprint from the grid is still the bottleneck of the total footprint.

5.3 Monetary Cost Reduction

We evaluate the monetary cost of SkyBox amortized over one year, and compare it with the baselines. With the monetary cost breakdown and the expected lifetime for each rMDC component in Table 5, we show our study results in Table 6.

SkyBox and Distr-Grid have the minimum monetary cost among all the rMDC placement policies. Distr-Battery has 13% additional costs due to the extra battery capacity. Centr-Global and Centr-Graph have lower server costs since they need fewer deployed servers than SkyBox to reach the same computing capacity. As the centralized data centers aggregate the power supply from multiple renewable energy farms, which generates a smoothed power supply curve, allowing Centr-Global and Centr-Graph to provision fewer servers to make use of their peak power supply. However, Centr-Global and Centr-Graph have to connect all renewable energy farms to the centralized rMDC via long power transmission lines (with an average distance of around 700 km), causing higher cost (62× more than SkyBox). Thus, these solutions are less practical in terms of cost efficiency. In contrast, SkyBox migrates VM workloads upon insufficient power supply rather than transmitting renewable energy from other renewable energy sites. It helps minimize the power transmission cost.

5.4 VM Uptime Improvement

SkyBox improves the overall efficiency of rMDCs as reflected by the evictable VM uptime. Figure 12 presents the trade-off between evictable VM uptime and their operational carbon footprint, with different settings of the evictable VMs (from 10% to 50%). With the same operational carbon footprint, SkyBox achieves a similar efficiency as Centr-Global and Centr-Graph. This is because they have relatively stable




Figure 13: Total carbon footprint with an increasing number of rMDC sites.

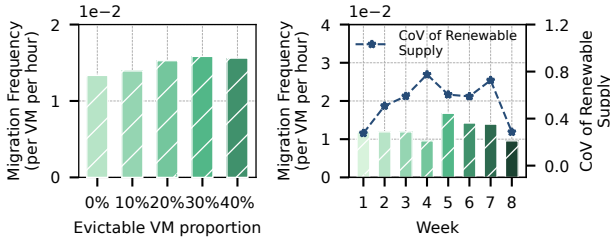


Figure 14: VM migration frequency of SkyBox.

renewable energy supplies, which minimizes the negative impact (i.e., VM shutdown) on evictable VMs.

SkyBox always delivers better efficiency than Distr-Grid and Distr-Battery. For instance, Figure 12 (left) shows that, when the carbon footprint is 2.0 tCO₂eq, SkyBox ensures the evictable VMs have 100% uptime, while Distr-Grid and Distr-Battery only have 29% and 45% uptime, respectively. This is because SkyBox can timely migrate VMs from the rMDCs that have insufficient renewable power, and keep these VMs operational. But Distr-Grid and Distr-Battery have to suspend these VMs; otherwise, they have to use the power grid to fill the gap, which results in an increased operational carbon footprint.

SkyBox does not rely on evictable VMs to gain benefits. With different proportions of evictable VMs (10%-50%), even when the uptime percentage of all evictable VMs keeps at 100%, SkyBox still effectively reduces the carbon footprint, because of the stable renewable power offered in subgraphs.

5.5 Sensitivity Analysis

Varying the number of rMDCs: We show the scalability of SkyBox by increasing the number of rMDCs from 6 to 54. As shown in Figure 13, the total carbon footprint of SkyBox is close to ideal, and SkyBox consistently outperforms Distr-Grid and Distr-Battery as we increase the number of rMDCs. We also observe that SkyBox has even greater carbon reduction (from 45% to 56%) with more rMDCs. This is because a larger pool of candidate subgraphs (e.g., 20 candidates with 6 sites vs. 24.8K candidates with 54 sites) increases the opportunities for SkyBox to identify more stable subgraphs.

Varying migration frequency: We evaluate SkyBox under different migration frequencies, where the frequency is




Figure 15: Total carbon footprint as we vary the battery capacity in modular data centers.




Figure 16: Total carbon cost with various evictable VMs.

measured by the number of migration events per VM per hour of runtime. Figure 14 (left) shows that under different evictable VM proportions, SkyBox obtains consistent carbon footprint reduction with low migration frequency (around 0.015). When the variation of renewable energy production is high (e.g., Week 4 in Figure 14), SkyBox avoids migrating VMs with long lifetimes to reduce the VM migration overheads. Consider that live migration has low power consumption, only sub-second downtime to VMs, and low impact on application performance [4, 24, 64], SkyBox does not incur much migration overhead. We show that SkyBox delivers constant migration frequency over time (Figure 14, right), under the setting of 0% evictable VMs.

Varying battery capacity: We present the benefit of SkyBox using different battery capacities (from a 15-minute to a 4-hour backup time). As shown in Figure 15, the increased battery capacities help offset more operational carbon footprint for SkyBox. SkyBox without extra battery can bring more benefits than Distr-Battery with a large battery capacity. This is because Distr-Battery has limited opportunity to keep the batteries charged upon insufficient power supply, while SkyBox can migrate VMs to another rMDC.

Varying evictable VM proportion: In Figure 16, we change the proportion of the evictable VMs from 0% to 40%. The total carbon footprint decreases for all the placement policies, as evictable VMs offer more flexibility in VM scheduling. SkyBox consistently outperforms Distr-Grid and Distr-Battery. Even with 0% evictable VMs (i.e., all VMs are regular), SkyBox outperforms Distr-Grid and Distr-Battery, with a reduction of the total carbon footprint of 44% and 35% respectively. SkyBox also incurs only 2% and 14% extra carbon compared to the Centr-Graph and Centr-Global, respectively.




Figure 17: Benefits of different SkyBox components.

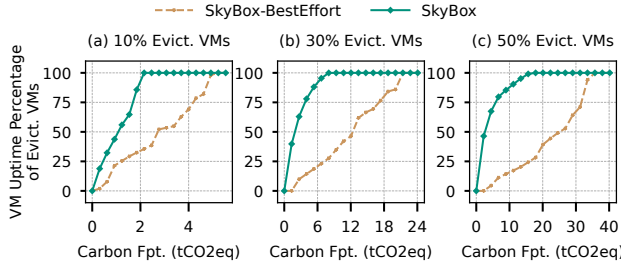


Figure 18: Comparison of the VM placement/migration policy in SkyBox vs. the best-effort approach.

5.6 Benefit Breakdown of SkyBox

We now evaluate the benefit of each SkyBox component. We compare SkyBox with its variants, each replacing one SkyBox component with a baseline policy (e.g., MIP Model vs. Best-Effort VM allocation). We first show the comparison of the total carbon footprint in Figure 17. By identifying stabler renewable energy farms, SkyBox reduces the total carbon footprint by 53%, compared to the random site selection scheme (SkyBox-NoSI). With subgraph identification, SkyBox groups multiple rMDCs to further achieve stable aggregated power production. It helps SkyBox to reduce 37% carbon footprint compared with the random subgraph identification (SkyBox-NoSG). Furthermore, SkyBox can find a more optimized VM placement and migration plan using our MIP model, leading to a 39% less carbon footprint than the best-effort VM placement policy (SkyBox-BestEffort).

To further evaluate how the MIP model of SkyBox optimizes VM placement and migration, we compare SkyBox against SkyBox-BestEffort. In Figure 18, with the same carbon footprint, SkyBox achieves higher VM uptime for evictable VMs. Their gap is larger with more evictable VMs. This is because the scheduling of a larger set of VMs across rMDCs forms a more complicated decision space, it is harder for the heuristic-based approach (i.e., best-effort) to find optimized VM placement and migration plans. When the MIP model of SkyBox makes decisions, it takes the future power supply changes and different VM properties (e.g., memory size and priority) into account and considers the global and long-term impact. Therefore, it can make better migration plans.




Figure 19: The impact of different misprediction ratios of power supply and VM lifetime in SkyBox.

5.7 Resilience to Mispredictions

We now demonstrate the robustness of SkyBox in handling the mispredictions of power supply and VM lifetime. In Figure 19, we vary the maximum misprediction ratios of the renewable power production and the VM lifetime from -50% to 50%, and present the percentage of increased carbon footprint against SkyBox with accurate predictions.

The mispredictions of power production incur a small increase in the total carbon footprint (0.6%-9.7%). And in practice, power production usually has a low prediction error. For instance, the misprediction ratio of a 3-hour-ahead power production prediction is typically $\pm 20\%$ (9% on average, see §2.2), where SkyBox has less than 3.7% carbon footprint overhead. As for VM lifetime mispredictions, their impact is trivial (up to 3%), showing that SkyBox can tolerate inaccurate predictions (up to 50%) of VM lifetime.

It is worth noting that SkyBox with mispredictions still outperforms Distr-Grid and Distr-Battery with accurate predictions (see Figure 10 and 19). SkyBox obtains at least 33% and 42% carbon footprint reduction compared with Distr-Grid and Distr-Battery, respectively. This is because SkyBox can effectively perform VM migration upon mispredictions, as discussed in §4.5.

6 RELATED WORK

Renewable energy for data centers. Using renewable energy to power data centers has been investigated in prior studies [39, 41–43, 61, 62, 81, 92, 106]. They benefited the recent development of containerized, mobile, truck, and edge data centers [18, 29, 97] that enable platform operators to move data centers closer to renewable energy farms with low cost. Most recently, a majority of popular cloud providers have been planning to use renewable energy in their data centers [2, 8, 45, 71], and discussed how to build carbon-aware data centers with renewable energy [2]. Our work SkyBox shares the same goal with these prior studies, but proposes a new approach to utilizing renewable energy for modular data centers. SkyBox employs a holistic and learning-based approach to formulate a different problem and explores different design tradeoffs.

A majority of prior studies focused on a single data center for optimizing its scheduling decisions. They did not explore the complementary energy patterns across multiple sites [39, 41]. In these prior studies, they either targeted only one specific type of application (e.g., Hadoop) [42], or one type of carbon emission (i.e., neglecting embodied emission) [92]. Our work proposes a holistic approach, and discusses different design tradeoffs, such as the tradeoff between the VM performance and the carbon footprint, distributed rMDCs vs. centralized data center, and the tradeoff between embodied and operation carbon footprint with different rMDC deployment settings.

Energy and carbon efficiency of data centers. To improve the energy efficiency of data centers, a variety of computer architecture and systems techniques have been developed [27, 28, 58, 59, 85, 93, 104], including the dynamic voltage and frequency scaling, and the power management in different server components (e.g., CPU cores, caches, and memory). Researchers also developed many power-aware scheduling policies across the entire systems stack [4, 17, 19, 21, 23, 40–42, 50, 67, 98, 100, 101, 105], such as workload scheduling and VM placement. Recently, more studies have focused on improving the carbon efficiency [14, 47, 48, 80, 91, 102, 103]. These studies are orthogonal to our work. SkyBox focuses on exploring the efficiency of renewable energy with modular data centers, by addressing the variability challenge of using the renewable energy.

Energy-based elasticity and reliability. To provide reliable and elastic computing for cloud services, prior studies proposed power-aware workload migration policies across data centers, depending on the energy cost and availability [11, 22, 65, 66]. And they investigated the use of VM migration [84, 88, 90], replication [26, 60], and checkpointing [75, 76] to manage faults and power outages. Different from these prior works, SkyBox targets the power variability of renewables, and develop new algorithms for VM placement and migrations across geo-distributed modular data centers. Modern data centers employed resource harvesting techniques, such as Spot VMs and Harvest VMs [5–7, 10, 70, 87] to address the variability in computing resource demand by adaptively reclaiming spare resources. While our goal is to optimize the use of unstable renewable energy, these resource harvesting approaches can be applied to further utilize the renewable energy. We wish to explore the resource harvesting techniques in rMDCs in the future work.

7 CONCLUSION

In this work, we develop SkyBox, a framework that employs a holistic and learning-based approach for deploying geo-distributed modular data centers at scale. It tackles the power variability of renewables with the insights derived from our

study of real-world power traces. SkyBox presents a set of techniques to help cloud vendors identify the suitable sites for deploying modular data centers. It also develops a set of algorithms to facilitate VM placement and migration across deployed data centers. Our experiments show the capabilities of SkyBox in helping cloud vendors deploy modular data centers in collocation with renewable energy farms.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and our shepherd Prashant Shenoy for their insightful comments and feedback. We also thank the members in the Systems Platform Research Group at University of Illinois Urbana-Champaign for discussing and proofreading this paper. They include Daixuan Li, Shaobo Li, Benjamin Reidys, Haoyang Zhang, and Yiqi Liu. This work was partially supported by NSF grant CCF-1919044.

REFERENCES

- [1] “How much energy do data centers consume?” <https://www.techtarget.com/searchdatacenter/tip/How-much-energy-do-data-centers-consume>, 2022.
- [2] B. Acun, B. Lee, F. Kazhamiaka, K. Maeng, M. Chakkaravarthy, U. Gupta, D. Brooks, and C.-J. Wu, “Carbon explorer: A holistic framework for designing carbon aware datacenters,” in *Proceedings of the 28th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’23)*, Vancouver, Canada, 2023.
- [3] U. E. I. Administration, “Eia projects that renewable generation will supply 44” <https://www.eia.gov/todayinenergy/detail.php?id=51698>, 2022.
- [4] S. Akoush, R. Sohan, A. Rice, A. W. Moore, and A. Hopper, “Free lunch: Exploiting renewable energy for computing,” in *13th Workshop on Hot Topics in Operating Systems (HotOS XIII)*. Napa, CA: USENIX Association, May 2011.
- [5] Amazon, “Spot Instances,” <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html>, 2020.
- [6] P. Ambati, Í. Goiri, F. Frujeri, A. Gun, K. Wang, B. Dolan, B. Corell, S. Pasupuleti, T. Moscibroda, S. Elnikety *et al.*, “Providing slos for resource-harvesting vms in cloud platforms,” in *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI ’20)*, 2020, pp. 735–751.
- [7] A. AWS, “Burstable performance instances,” <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances.html>, 2020.
- [8] A. AWS, “Why moving to the cloud should be part of your sustainability strategy,” <https://aws.amazon.com/blogs/publicsector/why-moving-cloud-part-of-sustainability-strategy/>, 2021.
- [9] Azure, “Microsoft azure traces,” <https://github.com/Azure/AzurePublicDataset/tree/master>, 2022.
- [10] M. Azure, “Introducing B-Series, our new burstable VM size,” <https://azure.microsoft.com/en-us/blog/introducing-b-series-our-new-burstable-vm-size/>, 2017.
- [11] S. Bahrami, V. W. Wong, and J. Huang, “Data center demand response in deregulated electricity markets,” *IEEE Transactions on Smart Grid*, vol. 10, no. 3, pp. 2820–2832, 2018.

- [12] H. Barbalho, P. Kovaleski, B. Li, L. Marshall, M. Molinaro, A. Pan, E. Cortez, M. Leao, H. Patwari, Z. Tang *et al.*, “Virtual machine allocation with lifetime predictions,” *Proceedings of Machine Learning and Systems*, vol. 5, 2023.
- [13] L. A. Barroso, U. Hölzle, and P. Ranganathan, *The datacenter as a computer: Designing warehouse-scale machines*. Springer Nature, 2019.
- [14] N. Bashir, T. Guo, M. Hajiesmaili, D. Irwin, P. Shenoy, R. Sitaraman, A. Souza, and A. Wierman, “Enabling sustainable clouds: The case for virtualizing the energy system,” in *Proceedings of the ACM Symposium on Cloud Computing*, 2021, pp. 350–358.
- [15] J. L. Berral, Í. Goiri, T. D. Nguyen, R. Gavaldà, J. Torres, and R. Bianchini, “Building green cloud services at low cost,” in *2014 IEEE 34th International Conference on Distributed Computing Systems*. IEEE, 2014, pp. 449–460.
- [16] J. L. Berral, Á. Goiri, T. D. Nguyen, R. Gavaldà, J. Torres, and R. Bianchini, “Building green cloud services at low cost,” in *Proceedings of the 2014 IEEE 34th International Conference on Distributed Computing Systems (ICDCS'14)*, 2014.
- [17] J. L. Berral, Í. Goiri, R. Nou, F. Julià, J. Guitart, R. Gavaldà, and J. Torres, “Towards energy-aware scheduling in data centers using machine learning,” in *Proceedings of the 1st International Conference on energy-Efficient Computing and Networking*, 2010, pp. 215–224.
- [18] F. Board, “Truck data center: The core component of the networked truck,” <https://www.fleetboard.info/digital-solutions/hardware/truck-data-center/#/>, 2022.
- [19] G. C. Buttazzo, G. Lipari, M. Caccamo, and L. Abeni, “Elastic scheduling for flexible workload management,” *IEEE Transactions on Computers*, vol. 51, no. 3, pp. 289–302, 2002.
- [20] U. o. M. Center for Sustainable Systems, “US Grid Energy Storage Factsheet,” <http://css.umich.edu/factsheets/us-grid-energy-storage-factsheet>, 2020.
- [21] A. Chatzipapas, D. Peditakis, C. Rotsos, V. Mancuso, J. Crowcroft, and A. Moore, “Challenge: Resolving data center power bill disputes: The energy-performance trade-offs of consolidation,” in *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems*, 2015, pp. 89–94.
- [22] H. Chen, M. C. Caramanis, and A. K. Coskun, “The data center as a grid load stabilizer,” in *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2014, pp. 105–112.
- [23] A. A. Chien, R. Wolski, and F. Yang, “The zero-carbon cloud: High-value, dispatchable demand for renewable power generators,” *The Electricity Journal*, vol. 28, no. 8, pp. 110–118, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1040619015001931>
- [24] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, “Live migration of virtual machines,” in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation-Volume 2*, 2005, pp. 273–286.
- [25] E. Cortez, A. Bonde, A. Muzio, M. Russinovich, M. Fontoura, and R. Bianchini, “Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms,” in *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17)*, 2017, pp. 153–167.
- [26] B. Cully, G. Lefebvre, D. Meyer, M. Feeley, N. Hutchinson, and A. Warfield, “Remus: High availability via asynchronous virtual machine replication,” in *Proceedings of the 5th USENIX symposium on networked systems design and implementation*. San Francisco, 2008, pp. 161–174.
- [27] H. David, C. Fallin, E. Gorbato, U. R. Hanebutte, and O. Mutlu, “Memory power management via dynamic voltage/frequency scaling,” in *Proceedings of the 8th ACM international conference on Autonomic computing*, 2011, pp. 31–40.
- [28] H. David, E. Gorbato, U. R. Hanebutte, R. Khanna, and C. Le, “Rapl: Memory power estimation and capping,” in *2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED)*. IEEE, 2010, pp. 189–194.
- [29] Delta, “All-in-one containerized data center solutions,” <https://www.deltapowersolutions.com/en/mcis/data-center-solutions-containerized-datacenter.php>, 2022.
- [30] R. H. Documentation, “Powertop,” https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/7/html/power_management_guide/powertop#PowerTOP, 2024.
- [31] D. Duplyakin, R. Ricci, A. Maricq, G. Wong, J. Duerig, E. Eide, L. Stoller, M. Hibler, D. Johnson, K. Webb, A. Akella, K. Wang, G. Ricart, L. Landweber, C. Elliott, M. Zink, E. Cecchet, S. Kar, and P. Mishra, “The design and operation of CloudLab,” in *Proceedings of the USENIX Annual Technical Conference (ATC)*, Jul. 2019, pp. 1–14. [Online]. Available: <https://www.flux.utah.edu/paper/duplyakin-atc19>
- [32] ELIA, “Solar and Wind Generation,” <https://www.elia.be/en/grid-data/power-generation>, 2021.
- [33] E. Emilsson and L. Dahllöf, “Lithium-ion vehicle battery production-status 2019 on energy use, co2 emissions, use of metals, products environmental footprint, and recycling,” 2019.
- [34] U. Energy Information Administration, “Electricity explained: Electricity generation, capacity, and sales in the united states,” <https://www.eia.gov/energyexplained/electricity/electricity-in-the-us-generation-capacity-and-sales.php>, 2020.
- [35] H. P. Enterprise, “Hpe product carbon footprint – hpe proliant dl580 gen10 server data sheet,” https://www.hpe.com/psnow/doc/a50005153enw.pdf?jumpid=in_pdfviewer-psnow, 2023.
- [36] Facebook, “Facebook’s u.s. renewable energy impact study,” <https://www.rti.org/publication/facebook-us-renewable-energy/fulltext.pdf>, 2021.
- [37] R. L. Fares and C. W. King, “Trends in transmission, distribution, and administration costs for us investor-owned electric utilities,” *Energy Policy*, vol. 105, pp. 354–362, 2017.
- [38] FrontlineSolvers, “Mixed-integer programming (mip) problems,” <https://www.solver.com/integer-constraint-programming>, 2022.
- [39] I. n. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, “Parasol and greenswitch: Managing datacenters powered by renewable energy,” in *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'13)*, Houston, Texas, USA, 2013.
- [40] I. n. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, “Parasol and greenswitch: Managing datacenters powered by renewable energy,” vol. 41, no. 1, p. 51–64, 2013.
- [41] Í. Goiri, K. Le, M. E. Haque, R. Beauchea, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, “Greenslot: scheduling energy consumption in green datacenters,” in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011, pp. 1–11.
- [42] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, “Greenhadoop: leveraging green energy in data-processing frameworks,” in *Proceedings of the 7th ACM european conference on Computer Systems*, 2012, pp. 57–70.
- [43] Á. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, “Designing and managing data centers powered by renewable energy,” *IEEE Micro*, vol. 34, no. 3, 2014.
- [44] I. Gonzalez Aparicio, T. Huld, F. Careri, F. Monforti, and A. Zucker, “EMHRES dataset Part II: Solar power generation,” *European Meteorological derived High resolution RES generation time series for present and future scenarios. Part II: PV generation using the PVGIS model*, 2017.

- [45] Google, “Our third decade of climate action: Realizing a carbon-free future,” <https://blog.google/outreach-initiatives/sustainability/our-third-decade-climate-action-realizing-carbon-free-future/>, 2020.
- [46] Google, “Cloud sustainability,” <https://cloud.google.com/sustainability>, 2021.
- [47] U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu, “Act: Designing sustainable computer systems with an architectural carbon modeling tool,” in *Proceedings of the 49th Annual International Symposium on Computer Architecture*, 2022, pp. 784–799.
- [48] U. Gupta, Y. G. Kim, S. Lee, J. Tse, H.-H. S. Lee, G.-Y. Wei, D. Brooks, and C.-J. Wu, “Chasing carbon: The elusive environmental footprint of computing,” in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2021, pp. 854–867.
- [49] L. Gurobi Optimization, “Gurobi optimizer reference manual,” 2018.
- [50] O. Hadary, L. Marshall, I. Menache, A. Pan, E. E. Greeff, D. Dion, S. Dorminey, S. Joshi, Y. Chen, M. Russinovich *et al.*, “Protean: {VM} allocation service at scale,” in *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI '20)*, 2020, pp. 845–861.
- [51] J. Hamilton, “Architecture for modular data centers,” *arXiv preprint cs/0612110*, 2006.
- [52] IBM, “Prefabricated modular data center—Add data center capacity where and when you need it,” <https://www.ibm.com/downloads/cas/YXKNDQAO>, 2014.
- [53] IBM, “Ibm cloud prefabricated modular data center.” <http://ftpmirror.your.org/pub/misc/ftp.software.ibm.com/common/ssi/ecm/bu/en/buo03045gben/global-technology-services-resiliency-services-bu-sales-guide-buo03045gben-20170801.pdf>, 2017.
- [54] D. Infra, “Microsoft azure—Zs data center locations: Regions and availability zones,” <https://dgtlinfra.com/microsoft-azure-data-center-locations/>, 2024.
- [55] G. A. Iratxe, A. Zucker, F. Careri, F. Monforti, T. Huld, and J. Badger, “EMHIRES dataset Part I: Wind power generation,” 2016.
- [56] JustEnergy, “Renewable Energy vs. Traditional Energy: The Real Cost,” <https://justenergy.com/blog/renewable-energy-vs-traditional-energy-the-real-cost/>, 2021.
- [57] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya, “Virtual machine power metering and provisioning,” in *Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 39–50.
- [58] V. Kontorinis, L. E. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, D. M. Tullsen, and T. S. Rosing, “Managing distributed ups energy for effective power capping in data centers,” in *2012 39th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2012, pp. 488–499.
- [59] A. Kumbhare, R. Azimi, I. Manousakis, A. Bonde, F. Frujeri, N. Mahalingam, P. Misra, S. A. Javadi, B. Schroeder, M. Fontoura *et al.*, “Prediction-based power oversubscription in cloud platforms,” *arXiv preprint arXiv:2010.15388*, 2020.
- [60] H. A. Lagar-Cavilla, J. A. Whitney, A. M. Scannell, P. Patchin, S. M. Rumble, E. De Lara, M. Brudno, and M. Satyanarayanan, “Snowflock: rapid virtual machine cloning for cloud computing,” in *Proceedings of the 4th ACM European conference on Computer systems*, 2009, pp. 1–12.
- [61] C. Li, Y. Hu, R. Zhou, M. Liu, L. Liu, J. Yuan, and T. Li, “Enabling datacenter servers to scale out economically and sustainably,” in *Proceedings of the 46th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'13)*, Davis, California, 2013.
- [62] C. Li, A. Qouneh, and T. Li, “Is witch: Coordinating and optimizing renewable energy powered server clusters,” in *Proceedings of the 39th Annual International Symposium on Computer Architecture (ISCA'12)*, Portland, Oregon, 2012.
- [63] H. Liu, H. Jin, X. Liao, C. Yu, and C.-Z. Xu, “Live virtual machine migration via asynchronous replication and state synchronization,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 22, no. 12, pp. 1986–1999, 2011.
- [64] H. Liu, C.-Z. Xu, H. Jin, J. Gong, and X. Liao, “Performance and energy modeling for live migration of virtual machines,” in *Proceedings of the 20th International Symposium on High Performance Distributed Computing*, ser. HPDC '11. New York, NY, USA: Association for Computing Machinery, 2011.
- [65] Z. Liu, I. Liu, S. Low, and A. Wierman, “Pricing data center demand response,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 42, no. 1, pp. 111–123, 2014.
- [66] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, “Data center demand response: Avoiding the coincident peak via workload shifting and local generation,” *Performance Evaluation*, vol. 70, no. 10, pp. 770–791, 2013.
- [67] K. Mahajan, M. Chowdhury, A. Akella, and S. Chawla, “Dynamic query {Re-Planning} using {QOOP},” in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, 2018, pp. 253–267.
- [68] S. McCluer and J.-F. Christin, “Comparing data center batteries, flywheels, and ultracapacitors,” *White paper*, vol. 65, p. 202, 2008.
- [69] Microsoft, “Introducing the microsoft azure modular datacenter,” <https://azure.microsoft.com/en-us/blog/introducing-the-microsoft-azure-modular-datacenter/>, 2020.
- [70] Microsoft, “Use Azure Spot Virtual Machines,” <https://docs.microsoft.com/en-us/azure/virtual-machines/spot-vms>, 2020.
- [71] Microsoft, “Azure sustainability,” <https://azure.microsoft.com/en-in/global-infrastructure/sustainability/>, 2021.
- [72] Microsoft, “Microsoft announces intent to build a new datacenter region in finland, accelerating sustainable digital transformation and enabling large scale carbon-free district heating,” <https://news.microsoft.com/europe/2022/03/17/microsoft-announces-intent-to-build-a-new-datacenter-region-in-finland-accelerating-sustainable-digital-transformation-and-enabling-large-scale-carbon-free-district-heating/>, 2022.
- [73] Microsoft, “Modular datacenter overview,” <https://learn.microsoft.com/en-us/azure-stack/mdc/mdc-overview>, 2022.
- [74] Microsoft, “Microsoft datacenters,” <https://datacenters.microsoft.com/globe/explore/>, 2024.
- [75] J. Mohan, A. Phanishayee, and V. Chidambaram, “Checkfreq: Frequent, fine-grained DNN checkpointing,” in *19th USENIX Conference on File and Storage Technologies (FAST '21)*, 2021, pp. 203–216.
- [76] A. Moody, G. Bronevetsky, K. Mohror, and B. R. De Supinski, “Design, modeling, and evaluation of a scalable multi-level checkpointing system,” in *SC'10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2010, pp. 1–11.
- [77] NREL, “Commercial battery storage,” https://atb.nrel.gov/electricity/2022/commercial_battery_storage, 2022.
- [78] D. of Energy, “Clean energy resources to meet data center electricity demand,” <https://www.energy.gov/policy/articles/clean-energy-resources-meet-data-center-electricity-demand>, 2024.
- [79] T. S. project, “Lean ICT – TOWARDS DIGITAL SOBRIETY,” Tech. Rep., 03 2019.
- [80] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care *et al.*, “Carbon-aware computing for datacenters,” *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2022.
- [81] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramaniam, “Carbon-aware energy capacity planning for datacenters,” in *2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of*

- Computer and Telecommunication Systems*. IEEE, 2012, pp. 391–400.
- [82] B. X. Rodriguez, “Embodied carbon of heating, ventilation, air conditioning and refrigerants (hvac+ r) systems,” Ph.D. dissertation, 2019.
- [83] A. Ruprecht, D. Jones, D. Shiraev, G. Harmon, M. Spivak, M. Krebs, M. Baker-Harvey, and T. Sanderson, “Vm live migration at scale,” *ACM SIGPLAN Notices*, vol. 53, no. 3, pp. 45–56, 2018.
- [84] A. Ruprecht, D. Jones, D. Shiraev, G. Harmon, M. Spivak, M. Krebs, M. Baker-Harvey, and T. Sanderson, “Vm live migration at scale,” *ACM SIGPLAN Notices*, vol. 53, no. 3, pp. 45–56, 2018.
- [85] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, “Green ai,” *Commun. ACM*, vol. 63, no. 12, p. 54–63, Nov. 2020. [Online]. Available: <https://doi.org/10.1145/3381831>
- [86] A. W. Services, “Aws announces aws modular data center for u.s. department of defense joint warfighting cloud capability,” <https://aws.amazon.com/blogs/publicsector/announcing-aws-modular-data-center-u-s-department-defense-joint-warfighting-cloud-capability/>, 2023.
- [87] P. Sharma, A. Ali-Eldin, and P. Shenoy, “Resource deflation: A new approach for transient resource reclamation,” in *Proceedings of the Fourteenth EuroSys Conference 2019*, 2019, pp. 1–17.
- [88] P. Sharma, S. Lee, T. Guo, D. Irwin, and P. Shenoy, “Spotcheck: Designing a derivative iaas cloud on the spot market,” in *Proceedings of the Tenth European Conference on Computer Systems*, 2015, pp. 1–15.
- [89] J. Shuja, A. Gani, S. Shamshirband, R. W. Ahmad, and K. Bilal, “Sustainable cloud data centers: a survey of enabling techniques and technologies,” *Renewable and Sustainable Energy Reviews*, vol. 62, pp. 195–214, 2016.
- [90] R. Singh, D. Irwin, P. Shenoy, and K. K. Ramakrishnan, “Yank: Enabling green data centers to pull the plug,” in *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, 2013, pp. 143–155.
- [91] A. Souza, N. Bashir, J. Murillo, W. Hanafy, Q. Liang, D. Irwin, and P. Shenoy, “Ecovisor: A virtual energy system for carbon-efficient applications,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 2023, pp. 252–265.
- [92] C. Stewart and K. Shen, “Some joules are more precious than others: Managing renewable energy in the datacenter,” in *Proceedings of the workshop on power aware computing and systems*. IEEE, 2009, pp. 15–19.
- [93] C. Subramanian, A. Vasani, and A. Sivasubramanian, “Reducing data center power with server consolidation: Approximation and evaluation,” in *2010 International Conference on High Performance Computing*. IEEE, 2010, pp. 1–10.
- [94] H. Technologies, “Huawei fusionmodule2000 smart modular data center solution,” <https://e.huawei.com/en/material/networkenergy/dcenergy/450821c30cb5499eb5e2fe39f09942b9>, 2021.
- [95] The Climate Pledge, “See which companies and organizations have committed to net zero carbon by 2040,” <https://www.theclimatepledge.com/us/en/Signatories>, 2021.
- [96] C. C. Thompson, P. K. Oikonomou, A. H. Etemadi, and V. J. Sorger, “Optimization of data center battery storage investments for microgrid cost savings, emissions reduction, and reliability enhancement,” *IEEE Transactions on Industry Applications*, vol. 52, no. 3, pp. 2053–2060, 2016.
- [97] K. Trailer, “Mobile data center,” <https://www.kytrailer.com/Specialty-Trailers/Category/1/Commercial/Type/6/Mobile-Data-Center>, 2022.
- [98] A. Verma, L. Pedrosa, M. R. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, “Large-scale cluster management at Google with Borg,” in *Proceedings of the European Conference on Computer Systems (EuroSys)*, Bordeaux, France, 2015.
- [99] K. V. Vishwanath, A. Greenberg, and D. A. Reed, “Modular data centers: how to design them?” in *Proceedings of the 1st ACM workshop on Large-Scale system and application performance*, 2009, pp. 3–10.
- [100] G. Von Laszewski, L. Wang, A. J. Younge, and X. He, “Power-aware scheduling of virtual machines in dvfs-enabled clusters,” in *2009 IEEE International Conference on Cluster Computing and Workshops*. IEEE, 2009, pp. 1–10.
- [101] C. Wang, B. Urgaonkar, Q. Wang, G. Kesidis, and A. Sivasubramanian, “Data center power cost optimization via workload modulation,” in *2013 IEEE/ACM 6th International Conference on Utility and Cloud Computing*. IEEE, 2013, pp. 260–263.
- [102] J. Wang, U. Gupta, and A. Sriraman, “Peeling back the carbon curtain: Carbon optimization challenges in cloud computing,” in *Proceedings of the 2nd Workshop on Sustainable Computer Systems*, 2023, pp. 1–7.
- [103] P. Wiesner, I. Behnke, D. Scheinert, K. Gontarska, and L. Thamsen, “Let’s wait awhile: How temporal workload shifting can reduce carbon emissions in the cloud,” in *Proceedings of the 22nd International Middleware Conference*, 2021, pp. 260–272.
- [104] C.-M. Wu, R.-S. Chang, and H.-Y. Chan, “A green energy-efficient scheduling algorithm using the dvfs technique for cloud datacenters,” *Future Generation Computer Systems*, vol. 37, pp. 141–147, 2014.
- [105] C. Zhang, A. Kumbhare, I. Manousakis, D. Zhang, P. Misra, R. Assis, K. Woolcock, N. Mahalingam, B. Warrior, D. Gauthier, L. Kunnath, S. Solomon, O. Morales, M. Fontoura, and R. Bianchini, “Flex: High-availability datacenters with zero reserved power,” in *Proceedings of the International Symposium on Computer Architecture (ISCA’21)*, June 2021.
- [106] K. Zhang, P. Wang, N. Gu, and T. D. Nguyen, “Greendrl: managing green datacenters using deep reinforcement learning,” in *Proceedings of the 13th Symposium on Cloud Computing*, 2022, pp. 445–460.