

# Enabling Sustainable Clouds: The Case for Virtualizing the Energy System

Noman Bashir<sup>1</sup>, Tian Guo<sup>2</sup>, Mohammad Hajiesmaili<sup>1</sup>, David Irwin<sup>1</sup>

Prashant Shenoy<sup>1</sup>, Ramesh Sitaraman<sup>1</sup>, Abel Souza<sup>1</sup>, and Adam Wierman<sup>3\*</sup>

<sup>1</sup>University of Massachusetts Amherst <sup>2</sup>Worcester Polytechnic Institute <sup>3</sup>California Institute of Technology

## ABSTRACT

Cloud platforms' growing energy demand and carbon emissions are raising concern about their environmental sustainability. The current approach to enabling sustainable clouds focuses on improving energy-efficiency and purchasing carbon offsets. These approaches have limits: many cloud data centers already operate near peak efficiency, and carbon offsets cannot scale to near zero carbon where there is little carbon left to offset. Instead, enabling sustainable clouds will require applications to adapt to when and where unreliable low-carbon energy is available. Applications cannot do this today because their energy use and carbon emissions are not visible to them, as the energy system provides the rigid abstraction of a continuous, reliable energy supply. This vision paper instead advocates for a "carbon first" approach to cloud design that elevates *carbon-efficiency* to a first-class metric. To do so, we argue that cloud platforms should *virtualize the energy system* by exposing visibility into, and software-defined control of, it to applications, enabling them to define their own abstractions for managing energy and carbon emissions based on their own requirements.

## CCS CONCEPTS

• **Computer systems organization** → **Cloud computing; Special purpose systems.**

## KEYWORDS

Carbon-efficiency, virtualization, cloud computing, edge

\* Author list in alphabetical order by last name.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org). SoCC '21, November 1–4, 2021, Seattle, WA, USA  
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8638-8/21/11...\$15.00

<https://doi.org/10.1145/3472883.3487009>

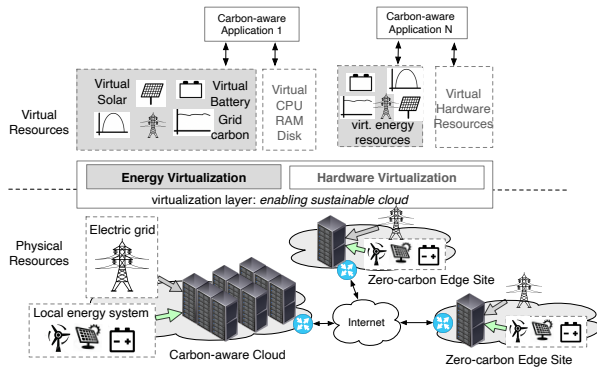
## ACM Reference Format:

Noman Bashir<sup>1</sup>, Tian Guo<sup>2</sup>, Mohammad Hajiesmaili<sup>1</sup>, David Irwin<sup>1</sup> and Prashant Shenoy<sup>1</sup>, Ramesh Sitaraman<sup>1</sup>, Abel Souza<sup>1</sup>, and Adam Wierman<sup>3</sup>. 2021. Enabling Sustainable Clouds: The Case for Virtualizing the Energy System. In *ACM Symposium on Cloud Computing (SoCC '21)*, November 1–4, 2021, Seattle, WA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3472883.3487009>

## 1 INTRODUCTION

Cloud platforms offer many advantages over building out and managing a private infrastructure, including low upfront costs, pay-as-you-go pricing, and rapid scalability. These advantages, along with an explosion of data and data-driven workloads [50, 54, 61], has fueled exponential growth in cloud capacity, which has been doubling roughly every four years for more than a decade [26]. This growth has been accelerating recently due to increasing demand for AI and machine learning (ML) applications. For example, recent estimates suggest the computation required to train state-of-the-art AI/ML models, e.g., for facial recognition, has been doubling every 3.4 months for a decade, which is significantly faster than Moore's Law [1]. If these trends continue, we can expect the cloud's exponential growth to persist for the foreseeable future, and possibly even accelerate further.

Cloud platforms have long had a strong financial incentive to optimize energy-efficiency to lower their operating expenses, which are massive for hyper-scale cloud data centers. These optimizations have been quite successful: despite the end of Dennard scaling [14], the cloud's energy demand grew much more slowly than expected over the past decade [44]. The success was largely due to industry's intense focus on reducing data centers' power usage effectiveness (PUE) to near 1 by aggressively optimizing energy-efficiency *en masse* across hardware, software, and cooling systems. However, there are few significant remaining opportunities left to further optimize energy-efficiency, as many cloud data centers already operate near peak efficiency [46]. As a result, cloud platforms can no longer rely on improving their energy-efficiency to mitigate their energy growth. *Thus, moving forward, the cloud's continued exponential growth is likely to translate directly into exponentially rising energy demand.*



**Figure 1: Our vision of a sustainable cloud-edge infrastructure that virtualizes the energy system, enabling applications to optimize their own energy and carbon.**

As cloud platforms’ energy demand grows, there is increasing concern about its environmental sustainability. Indeed, researchers are gradually shifting their focus from reducing operating costs by minimizing energy usage to reducing environmental impact by minimizing carbon emissions. For example, there has been recent work highlighting deep learning’s carbon emissions [13, 51] and advocating for the design of zero-carbon clouds [17]. Many prominent technology companies have also recognized the problem and set ambitious goals to reduce their carbon footprint, primarily by purchasing carbon offsets [2, 10, 21, 35, 48]. However, carbon offsets are only a transitional mechanism that cannot scale to near zero carbon, since, at that point, there is little carbon left to offset. Of course, the pronouncements above also coincide with a burgeoning financial incentive to “go green,” as solar energy is already the cheapest form of electricity in recorded history, and its cost is expected to continue to decline [22].

This recent emphasis on sustainability recognizes that the cloud’s rising energy consumption is not actually the problem: rather, *the problem is the carbon footprint of that energy consumption and its negative impact on the environment*. However, optimizing data centers for *carbon-efficiency*, i.e., the work done per kilogram of carbon emitted, differs substantially from optimizing them for *energy-efficiency*, i.e., the work done per joule of energy consumed. To illustrate, a data center could be highly energy-inefficient, e.g.,  $PUE \gg 2$ , but also highly carbon-efficient if it is powered entirely by *green* energy from co-located renewables with zero carbon emissions. Likewise, a data center could be highly energy-efficient, e.g.,  $PUE \sim 1$ , but also highly carbon-inefficient if it is powered by *brown* energy from grid generators burning fossil fuels. While there has been decades of research on optimizing cloud energy-efficiency, there has been little research on optimizing cloud carbon-efficiency. To address the problem, we advocate for a “carbon first” approach to cloud design that re-focuses research on optimizing carbon-efficiency by elevating it to a first-class metric.

A distinguishing characteristic of low-carbon energy is that it is *unreliable*: it is not always available at any single location all the time, but instead varies widely, and not entirely predictably, both temporally and geographically. Thus, *enabling sustainable clouds that are carbon-efficient will require applications to adapt to when and where low-carbon energy is available*. Unfortunately, applications cannot do this today because their energy and carbon are not visible to them, as the energy system provides the rigid abstraction of a continuous, reliable energy supply. To address the problem, this vision paper instead argues for *virtualizing the energy system* by exposing visibility into, and software-defined control of, it to cloud applications, providing them the flexibility to define their own abstractions for managing energy and carbon emissions based on their own requirements.

Figure 1 illustrates our vision of a sustainable cloud-edge infrastructure that virtualizes the energy system, enabling applications to optimize their energy usage and carbon emissions. Our approach effectively extends the end-to-end principle to the energy system [41], and to some extent takes prior approaches that advocate delegating resource management to applications to their logical conclusion [20, 27].

## 2 UNSUSTAINABLE APPROACHES

Ironically, current approaches to enabling sustainable clouds, which focus on improving energy-efficiency and purchasing carbon offsets, are themselves not sustainable. That is, as we discuss below, these approaches are not sufficient to eliminate the cloud’s carbon emissions, and it is not necessarily clear how much they contribute to lowering them, especially over long periods. Thus, developing new approaches that unequivocally and directly reduce carbon emissions is critical. **Energy-Efficiency.** Improving the energy-efficiency of computer systems, in general, has been an active research area for nearly three decades [57], and of data centers, in particular, for at least two decades [16]. Over this period, researchers have developed myriad techniques that have pushed computational energy-efficiency to near its physical limits. These techniques run the gamut from hardware mechanisms, such as dynamic voltage and frequency scaling (DVFS) [57] and power capping [29, 37], to software policies, such as consolidating workload to shutdown idle servers [16], to facility and cooling system optimizations, such as leveraging “free” cooling [25] and optimizing power delivery [38]. Improvements in the latter are captured by PUE, and have been an intense focus of data center operators over the past decade, largely because PUEs were initially high ( $\gg 2$ ) and thus there was significant room for improvement.

At this point, however, state-of-the-art data centers have PUEs within 6-10% of optimal, i.e.,  $PUE \sim 1$  [8], leaving little room for further improvement. Hardware and software systems similarly operate near their peak efficiency and thus

also cannot be substantially improved. As a result, *we cannot expect to significantly reduce the cloud's energy consumption and thus its carbon emissions by further improving energy-efficiency*. In fact, the efficiency improvements over the past decade has not stopped the cloud's energy growth, but only slowed it down relative to initial estimates [32]. This may be due, in part, to Jevon's Paradox, which observes that increasing resource efficiency often increases, rather than decreases, that resource's consumption [12, 49]. This paradox is well-known in energy economics, and occurs because increasing energy's efficiency also lowers its cost, which can increase demand, such that it more than offsets the decrease in usage from improving efficiency. *Hence, while there are many benefits to improving energy-efficiency, reducing energy usage, and thus carbon emissions, is not always one of them.*

**Carbon Offsets.** The other common approach to reducing cloud carbon emissions is by purchasing carbon offsets, which represent a quantifiable reduction in carbon emissions. While there are many types of carbon offsets, technology companies often purchase them by either subsidizing renewable energy projects or purchasing renewable energy credits [18]. This renewable energy is not necessarily located near cloud data centers, and typically does not directly power them. However, when accounting for carbon emissions, these carbon offsets are treated as location-agnostic by assuming each kilowatt-hour (kWh) of carbon-free renewable energy that is purchased displaces a kWh of carbon-intense electricity consumed from data centers' local grid. The typical accounting period for carbon offsets under greenhouse gas (GHG) protocol is one calendar year [9]. Thus, a "carbon neutral" company purchases enough carbon offsets each year to balance its carbon emissions. With this approach, companies can also be "carbon negative" by purchasing more carbon offsets than their carbon emissions [48].

Technology companies have led in the adoption of carbon offsets, and many have used them to achieve annualized location-agnostic carbon-neutrality [2, 10, 35]. Yet, these and other companies' operations are still responsible for a significant amount of direct carbon emissions. To address the problem, Google recently announced that it aims to be "carbon free" by 2030, in part, by piloting a new form of carbon offset, called Time-based Energy Attribute Certificates (T-EACs), which have an hourly location-specific accounting regime [55]. T-EACs recognize that grid carbon emissions vary over time and by location, based on the mix of generators used to satisfy demand, and thus incentivize using energy when and where carbon emissions are low. As a result, T-EACs more directly offset companies' direct carbon emissions. Unfortunately, carbon offsets, including T-EACs, are only effective as a transitional mechanism, since we must ultimately reduce *absolute* global carbon emissions to near zero, where there is little carbon left to offset. *To reach zero carbon,*

*we must move beyond carbon offsets and focus on changing operations to always run on zero-carbon energy, e.g., from solar, wind, hydro, nuclear, and geothermal.* Carbon offsets may also disincentive, and actually delay, these operational changes by providing a means for reducing carbon emissions for less than it would cost to make such changes [34].

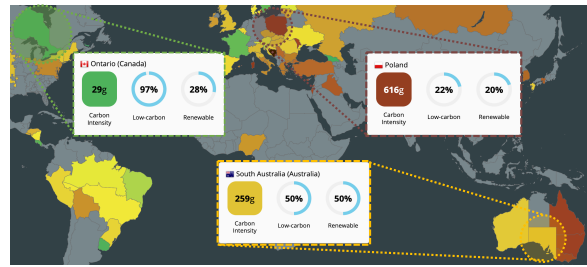
### 3 A CARBON FIRST APPROACH

To address the limitations of current approaches, we advocate for a "carbon first" approach that elevates *carbon-efficiency* to a first-class metric in cloud design. Carbon-efficiency is a measure of computational work done per kilogram of carbon, and other greenhouse gas (GHG), emissions. We focus specifically on Scope 2 emissions from using electricity [9], which represent the vast majority of cloud platforms' carbon emissions. Unlike energy-efficiency, carbon-efficiency optimizations are not bound by Jevon's Paradox because carbon is not a resource cloud platforms consume, but, rather, an energy by-product, which they can eliminate by operating when and where zero-carbon energy is available.

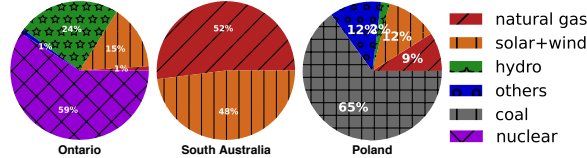
Of course, to optimize any metric, we must be able to both measure and control it. Currently, cloud platforms and users have little visibility into their energy usage and carbon emissions, and little-to-no control over them. The cloud's energy system delivers power from two primary sources: the electric grid and an increasingly rich local energy system, which may consist of on-site renewables, such as solar, and batteries. These two power sources present different trade-offs. The grid provides the convenient abstraction of a reliable power source with a non-zero carbon footprint, while the local energy system is an unreliable source of clean power. In §3.1 and §3.2, we discuss the potential for optimizing carbon-efficiency by exposing visibility and control of energy and carbon emissions from both sources. §4 then makes the case for exposing this visibility and control directly to cloud applications by virtualizing the energy system.

#### 3.1 Electric Grid: Visibility and Control

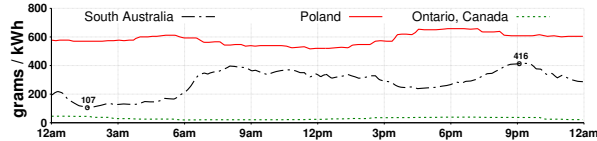
The electric grid presents electrically-powered devices with the abstraction of a reliable supply of power on demand (up to some maximum power). The grid exposes a hardware interface to this abstraction in the form of standardized electrical sockets, which implicitly encode information about power's output, i.e., voltage and frequency, via their form factor. This simple abstraction has been amazingly successful for over 100 years, since the grid's inception, at supporting countless uses of electrical energy, most of which were unimaginable a century ago. However, this simplicity is increasingly becoming a barrier to innovation, as it exposes no interface for receiving information about grid energy's characteristics, i.e., its energy sources and carbon emissions, and no means for



(a) Grid carbon emissions vary by 21× across regions.



(b) Varying mix of generators in different regions



(c) Hourly carbon emissions vary by 4× over a single day.

**Figure 2: Grid carbon emissions vary spatially (a) and temporally (c) based on the mix of active generators (b).**

controlling these characteristics. This lack of visibility and control prevents systems from optimizing carbon-efficiency. **Exposing Visibility.** The grid's energy comes from a large mix of generators that have a wide range of carbon emissions. For example, hydroelectric power plants have zero carbon emissions, while carbon emissions from thermal generators vary widely based on their fuel type, capacity, and real-time power output. The grid varies its set of active generators, and their real-time power output, over time to match the energy demand of its connected devices. Since the grid exercises little-to-no control over the devices that connect to its socket interface, it also has little-to-no control over variations in its energy demand. *These variations in demand and generation cause wide variations in grid energy's carbon emissions.*

Recently, new services, such as electricityMap [5], have emerged that estimate grid carbon emissions by collecting and analyzing real-time data on grid operations published by grid balancing authorities, such as Independent System Operators (ISOs) and Regional Transmission Organizations (RTOs). This data reveals that generator operations exhibit an order of magnitude variation in their carbon emissions across space and time. Such spatial and temporal variations present an untapped opportunity for reducing the carbon emissions of cloud platforms and applications.

**Variations Across Space.** Figure 2(a) shows the average carbon intensity (in g-CO<sub>2</sub>/kWh) of different world regions from electricityMap, where darker colors indicate higher carbon intensity. The variations are significant with Ontario, which

relies on nuclear and hydro power, having 21× lower carbon intensity than Poland, which relies on thermal generators that burn coal. Figure 2(b) shows the different generator mixes in Ontario, S. Australia, and Poland from left to right.

**Variations Across Time.** Figure 2(c) similarly shows that grid energy's carbon intensity also varies significantly over time at any single location. On this day, S. Australia's carbon intensity varied by 4× (between 2am and 9pm), largely because solar and wind energy account for 48% of the energy generation in S. Australia. By contrast, Poland and Ontario's carbon intensity are much less variable because they predominantly rely on coal and nuclear, respectively.

In some sense, electricityMap and similar services supplement the grid's socket interface with a new interface that enables devices to retrieve energy information. Thus, *we can use these services to also expose grid carbon emissions to cloud platforms and applications.* ElectricityMap already exposes a real-time API that software could integrate programmatically. Google has recently taken steps in this direction by using electricityMap data to provide users carbon intensity estimates for cloud regions to inform their choices [7]. Interestingly, Google and other cloud platforms do not yet expose similar visibility into cloud applications' fine-grained power usage even though such information is typically visible to low-level software. System software, such as RAPL [37], is also capable of attributing power to individual cores, which can be used as input to models capable of attributing power to individual processes, containers, and virtual machines (VMs) on a server. Thus, while Google's current carbon data is useful for coarse-grained region selection, users cannot yet track their own fine-grained application-specific energy usage, and thus carbon emissions. *We argue that exposing such fine-grained data is necessary for optimizing the carbon emissions of cloud platforms and applications.*

**Exposing Control.** Exposing the visibility above would enable direct monitoring of grid energy's carbon emissions, but is useless without a means to control these emissions. Such control is possible either by changing energy's supply, i.e., the grid's mix of generators, or its demand, i.e., consumption. Unfortunately, consumers cannot directly influence grid energy's carbon emissions by altering the mix of generators supplying electricity. *Consumers, including cloud platforms, can, however, vary their energy consumption, and hence carbon emissions, by shifting workloads in either time—to a low-carbon period—or space—to a low-carbon region.*

Cloud platforms can shift energy usage across time by controlling either computation via job scheduling, i.e., running jobs during low-carbon periods, or energy via battery scheduling, i.e., charging and discharging batteries during low- and high-carbon periods, respectively. The latter is important for interactive services that cannot simply shift their computation in time. Google has begun experimenting internally

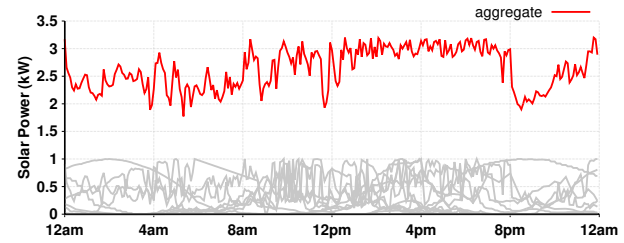


with such carbon-aware time-shifting in its data centers [40]. Cloud platforms can similarly shift energy usage across space by moving computation, e.g., jobs and requests, to regions with low carbon intensity. As others have noted [11, 53], moving computation, which is equivalent to moving tiny amounts of energy in the form of the bits that encode information, is many orders of magnitude more efficient and less costly than moving energy. Thus, *cloud platforms should focus on moving computation to energy, rather than moving energy to computation*. Cloud platforms' geographically distributed infrastructure enables such carbon-aware movement.

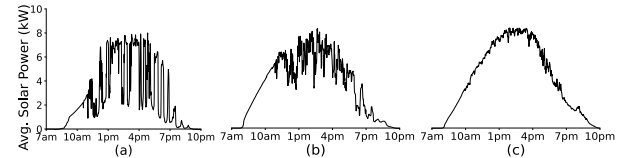
Unfortunately, cloud platforms do not expose any means for controlling energy usage (and thus carbon emissions), such as power capping [45], even though these mechanisms are also provided by low-level software, including RAPL [37], and can be applied to individual processes, containers, and VMs. There has been substantial prior work using such control to perform time and space shifting but mostly for optimizing energy costs, which, unlike carbon emissions, are visible, e.g., via job scheduling [23, 24], battery scheduling [36, 56], and request routing [39]. Some of these approaches may also be useful in optimizing cloud carbon-efficiency. However, a key difference with prior work is that we eventually must reduce carbon emissions to zero. This will likely require capping carbon emissions at some point, and then progressively lowering the cap. Such caps will break the grid's abstraction of reliable energy and force cloud platforms and applications to handle it. In contrast, prior work generally focuses on minimizing cost assuming reliable energy, and thus does not address problems with unreliable energy, which may force applications to make compromises in their performance and reliability to limit carbon emissions. As we discuss in §3.2, techniques designed for transient computing [42, 47], which focus on designing reliable applications on unreliable servers, may be useful in handling this unreliability.

### 3.2 Local Energy: Visibility and Control

Cloud platforms can also draw energy from their local energy system, which may include both substantial battery capacity and co-located renewables, such as solar or wind energy [31]. Unfortunately, hyper-scale cloud data centers are generally too power-dense for renewables to power a significant fraction of their operations. For example, the average solar power density across the earth's surface is  $\sim +55\text{W/m}^2$  [15], assuming ideal solar efficiency at the Shockley-Queisser limit, and only  $\sim +33\text{W/m}^2$  at current solar cell efficiencies. Thus, assuming average solar density, making a 100MW data center net-zero using co-located solar would require covering an area of nearly 450 acres (or  $\sim 1.35\text{km} \times 1.35\text{km}$  area). While such a large co-located solar farm may be possible [31], there are many issues that can prevent it, such as lack of available



**Figure 3: Estimated individual and aggregate solar across 16 AWS regions distributed across the earth.**



**Figure 4: Solar power aggregated across (a) 1, (b) 7, and (c) 15 sites, from left to right, within a 50-mile radius.**

land, unsuitable terrain, and a large distance to a grid connection. As a result, we expect that most hyper-scale data centers will require some grid energy, and thus will need to rely, in part, on their local grid to reduce carbon intensity.

The characteristics above, though, motivate more widely distributing computation to lower its power density to match renewables, rather than aggregating it at a single location. Hence, we envision sustainable clouds will include many smaller, more widely distributed, micro data centers that are largely self-powered, primarily using solar and batteries [24]. This approach, while more sustainable, also has some financial benefits. As mentioned earlier, moving computation is cheaper than moving energy, and the difference in cost will only grow as solar prices continue to decline. The cost disparity is not apparent now because the grid already exists, and thus represents a sunk cost. Cloud platforms have also already started building out such a distributed infrastructure, albeit for a different reason: to support edge computing near end-users that can deliver low-latency services [28]. *Cloud platforms should leverage this opportunity to also improve their sustainability*. Prior work has already demonstrated that renewable-powered edge sites are viable [24].

Distributing multiple such edge sites across large and small geographical regions is also beneficial because it can reduce renewable energy's unreliability. Consider that, over large geographic regions, solar energy is highly reliable: in the extreme, the earth's surface receives near constant solar radiation. This effect is also present when applied to just a few sites, as shown in Figure 3, which plots estimated solar output across 16 AWS regions spread across the world (gray), and their aggregate (red). The figure shows that aggregate solar output is much less volatile than any individual site. Solar volatility also decreases when aggregating across many sites in a small region, as shown in Figure 4. This figure shows aggregate solar energy, from left to right, across 1, 7, and 15

sites (each with 8kW capacity) over a cloudy day within a 50-mile radius. As before, the aggregate is much less volatile than a single site. Thus, *while renewable energy is not always available at one location, it is likely available somewhere.*

**Exposing Visibility.** Self-powered edge data centers have complete visibility into their rich local energy system, including its real-time solar (or wind) generation, grid power consumption, battery charge level, and energy usage of computing infrastructure. These sub-systems generally include fine-grained energy monitoring that report power consumption at high resolutions, e.g., every second or minute, using embedded sensors or, for servers, hardware counters. This energy-related information is typically available via programmatic APIs exposed by different hardware components, e.g., battery charge controller or solar inverter, and should be exposed and integrated into management software to provide applications visibility into the energy system.

**Exposing Control.** Although local generation from renewable sources is intermittent and variable, edge data centers have full control over their local energy system, enabling them to dynamically coordinate power flow between computing infrastructure and the grid, solar, and batteries [24]. Hyper-scale data centers with co-located renewables have similar control, but as mentioned above, are more likely to rely on grid energy. Specifically, data centers can control: when to use, store, or net meter solar energy; when and how to charge batteries, e.g., from the grid or solar, and when to discharge them; and if and when to use grid energy and how much. Renewable energy research has shown that such control decisions can directly optimize carbon emissions [30].

However, a key challenge for renewable-powered infrastructure is its unreliability: if grid energy is not available, e.g., due to carbon capping, it may require throttling or shutting down servers due to a lack of energy. Relying more on unreliable renewable energy will further break the grid's reliability abstraction. Instead, carbon-efficient applications will need to be designed to handle new renewable energy dynamics. These dynamics are reminiscent of those exhibited by spot/preemptible cloud servers [3, 4, 6]. While spot/preemptible servers are cheap, cloud platforms may revoke them at any time. These revocation "failures" are expected, and prior work has designed many techniques to mitigate their impact in various applications, primarily by judiciously checkpointing in-memory state [42, 43, 47, 52, 59, 60].

An important difference between spot/preemptible dynamics and renewable dynamics is that, with the former, cloud platforms unilaterally choose which servers to revoke without warning, while, with the latter, systems have a wider range of choices to satisfy energy drops by controlling either their computation, e.g., selecting servers to

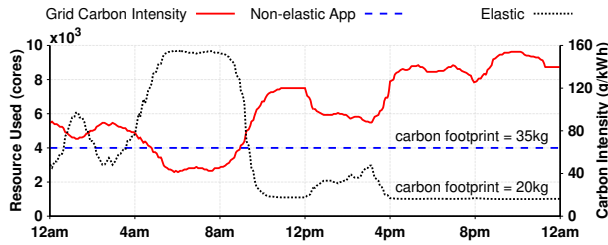
throttle/shutdown, or their energy system, e.g., discharging batteries. Of course, under an energy shortage, applications might also choose to move computation to another site without a shortage. *Different applications will make different decisions based on their own specific requirements.*

## 4 VIRTUALIZING THE ENERGY SYSTEM

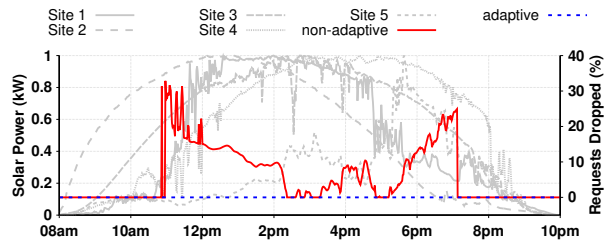
The previous section argues that optimizing carbon-efficiency requires being able to measure and control it. There is really no technical barrier to doing so, as §3 outlines a rich set of existing mechanisms for exposing visibility and control of energy usage and carbon emissions to cloud platforms. Unfortunately, simply exposing visibility and control is not enough, since cloud platforms are not in a position to exercise this control to optimize carbon-efficiency because they have little visibility into the applications that run on them. Applications have a wide range of characteristics and performance requirements that affect how they may choose to optimize carbon-efficiency. For example, a batch job, such as training a large distributed ML model, might choose to handle a renewable energy shortage (or high-carbon grid energy) by capping its power usage if it can tolerate some delay. In contrast, an interactive service might handle a similar situation by either discharging a battery, re-routing requests to a site with ample clean energy, or some combination.

Thus, we argue that enabling sustainable clouds requires *virtualizing the energy system* to expose the visibility and control from §3 directly to applications. We envision an application's virtual energy system providing a virtual solar array, which supplies a configurable share of a physical solar array's real-time power, a virtual battery, which offers a configurable share of a physical battery's capacity, and a grid connection, which provides access to a configurable amount of carbon-intensive energy. An exogenous policy would determine each application's share of variable solar power and battery capacity at each site. For example, public cloud platforms might sell solar and battery shares at each site for some price independently of hardware resources. Cloud platforms could also directly incentivize carbon-efficiency by setting per-user carbon caps or placing an explicit price on carbon.

Critically, virtualizing the energy system would enable *individual applications* to measure, control, and thus optimize their own carbon-efficiency. Applications would have full control over their virtual energy system via an application-level API, enabling them to explicitly regulate how much grid power they use, their battery charging and discharging, e.g., from solar or grid, and their energy consumption by setting per-VM (or container) power caps. This API would also enable applications to access energy information, such as their real-time energy usage, grid carbon emissions, and solar generation, or receive asynchronous notifications when



**Figure 5: Resource and carbon usage for an elastic and inelastic ML job. The carbon-aware elastic job finishes at the same time, and reduces carbon emissions by 45%.**



**Figure 6: Solar power across multiple edge sites (gray), and requests dropped by adaptive and non-adaptive serverless edge application (blue and red). The non-adaptive app drops 32% of requests due to lack of energy.**

significant changes occur, such as a large decrease in solar power due to a passing cloud. Applications would then be responsible for matching their allocated energy’s supply to demand to optimize their carbon-efficiency. That is, applications would be allocated computing resources independently of energy, and would use their virtual energy system to explicitly allocate their available power across their resources, e.g., by setting power caps or deferring/moving computations. Our approach essentially applies the end-to-end principle to the energy system [41], similar to how exokernels applied it to managing hardware [20].

As with exokernels, virtualizing the energy system can enable the development of new application-specific abstractions for managing unreliable low-carbon energy from within the software stack, without requiring access to privileged parts of the system, e.g., the energy system hardware and grid. Importantly, providing applications access to a low-level abstraction for managing their energy and carbon usage *does not* require them to use it directly. Instead, most applications will interact with libraries that provide application-specific abstractions for managing energy and carbon, just as exokernels and Mesos [27] enable applications to use different library OSes and cloud frameworks.

The system would aggregate and map many applications’ virtual controls over their virtual energy system onto the physical system, e.g., by enforcing power caps and regulating battery charging/discharging, and strictly enforce the constraint that an application’s power usage does not exceed the power its virtual energy system is supplying. For

example, if an application’s virtual solar power drops, and it does not respond by lowering power caps to match the drop, the system might terminate VMs to satisfy the constraint, similar to an OS’s Out-of-Memory killer.

To demonstrate the benefits of such flexibility, Figure 5 depicts the execution over time of elastic and inelastic distributed ML training, along with grid energy’s variable carbon-intensity. In this case, we use MLperf [33] power-performance results to model ImageNet [19]. The inelastic variant uses the same resources, e.g., GPUs, throughout the day, while the elastic variant has an autoscaler that uses a sigmoid function to dynamically scale resources up and down during low and high carbon periods, respectively. The experiment shows that by scaling resource usage in response to variations in grid carbon intensity, the elastic carbon-aware variant lowers overall carbon emissions 45% without any effect on job completion time.

Figure 6 then depicts the performance of a serverless edge application that runs on edge sites powered by renewable energy. Requests are generated based on the latency-sensitive serving tasks in a Google workload trace (first day of cell a) [58]. The gray lines depict solar output at each site, while the blue and red lines depict the percentage of requests dropped by an adaptive and non-adaptive application. For the non-adaptive application, we evenly and statically divide the workload across sites, while the adaptive application routes workload to sites based on their energy availability. We report the drop rate only when the aggregate power across sites is enough to serve the load. The figure shows that the non-adaptive service drops up to 32% of requests due to lack of available energy, even though energy is available at other sites, while the adaptive service never drops requests and has a zero carbon footprint.

## 5 CONCLUSIONS

This paper advocates for a “carbon first” approach to cloud design that elevates carbon-efficiency to a first-class metric. Carbon-efficiency has generally been ignored by computer systems researchers in the past because it requires tighter integration with, and visibility into, an often opaque energy system that provides the convenient abstraction of reliable energy on demand. However, the environmental cost of maintaining this simple abstraction has become too high, as it masks energy’s unreliability and carbon emissions from applications. Thus, we argue for virtualizing the energy system to expose the visibility and control applications need to optimize carbon emissions based on their own requirements.

**Acknowledgements.** We thank our shepherd Sangeetha Abdu Jyothi and the reviewers for their comments, electricityMap for the carbon data, and AWS. This work is funded by NSF grants 2105494, 2045641, and VMware.

## REFERENCES

- [1] 2018. OpenAI Blog, AI and Compute. <https://openai.com/blog/ai-and-compute/>.
- [2] 2019. Reuters, Amazon Vows to be Carbon Neutral by 2040, buying 100,000 Electric Vans. <https://www.reuters.com/article/us-amazon-environment/amazon-vows-to-be-carbon-neutral-by-2040-buying-100000-electric-vans-idUSKBN1W41ZV>.
- [3] 2020. Amazon EC2 Spot Instances. <https://aws.amazon.com/ec2/spot/>.
- [4] 2020. Azure Spot Virtual Machines. <https://azure.microsoft.com/en-us/pricing/spot/>.
- [5] 2020. Electricity Map. <https://www.electricitymap.org/map>.
- [6] 2020. Google Preemptible Virtual Machines. <https://cloud.google.com/preemptible-vms>.
- [7] 2021. Carbon free energy for Google Cloud regions. <https://cloud.google.com/sustainability/region-carbon>.
- [8] 2021. Google Data Centers Efficiency. [google.com/about/datacenters/efficiency/](https://google.com/about/datacenters/efficiency/).
- [9] 2021. Greenhouse Gas Protocol. <https://ghgprotocol.org/>.
- [10] Nicola Acutt. 2018. Radius: Stories at the Edge, Achieving Carbon Neutrality. <https://www.vmware.com/radius/achieving-carbon-neutrality/>.
- [11] A. Agarwal, J. Sun, S. Noghabi, S. Iyengar, A. Badam, R. Chandra, S. Seshan, and S. Kalyanaraman. 2021. Virtual Battery: Redesigning Cloud Computing for Renewable Energy. In *HotNets*.
- [12] B. Alcott. [n.d.]. Jevons' Paradox. *Ecological Economics* 54, 1 ([n. d.]), 9–21.
- [13] E. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *ACM FAccT*.
- [14] Mark Bohr. 2007. A 30 Year Retrospective on Dennard's MOSFET Scaling Paper. *IEEE Solid-State Circuits Society Newsletter* 12, 1 (Winter 2007), 11–13.
- [15] Greg Bothun. 2020. Basics of Solar Energy. <http://zebu.uoregon.edu/disted/ph162/14.html>.
- [16] Jeffrey S Chase, Darrell C Anderson, Prachi N Thakar, Amin M Vahdat, and Ronald P Doyle. 2001. Managing Energy and Server Resources in Hosting Centers. *ACM SIGOPS Operating Systems Review* 35, 5, 103–116.
- [17] A. Chien. 2021. Driving the Cloud to True Zero Carbon. *CACM* 64, 2 (February 2021).
- [18] Alyssa Daniels. 2020. Environmental Leader, Google Signs PPA for 140MW from Solar Farm in Texas. <https://www.environmentalleader.com/2020/09/google-candela-texas-solar-ppa/>.
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [20] Dawson R. Engler, M. Frans Kaashoek, and James O'Toole. 1995. Exokernel: An Operating System Architecture for Application-Level Resource Management. In *ACM Symposium on Operating System Principles (SOSP)*, Vol. 29. 251–266.
- [21] Darrell Etherington. 2020. TechCrunch, Google Claims Net Zero Carbon Footprint over its Entire Lifetime, Aims to only use Carbon-Free Energy by 2030. <https://techcrunch.com/2020/09/14/google-claims-net-zero-carbon-footprint-over-its-entire-lifetime-aims-to-only-use-carbon-free-energy-by-2030/>.
- [22] S. Evans. 2020. CarbonBrief, Solar is now 'cheapest electricity in history,' confirms IEA. <https://www.carbonbrief.org/solar-is-now-cheapest-electricity-in-history-confirms-iea>.
- [23] Inigo Goiri, Ryan Beauchea, Kien Le, Thu D. Nguyen, Md. E. Haque, Jordi Guitart, Jordi Torres, and Ricardo Bianchini. 2011. GreenSlot: Scheduling Energy Consumption in Green Datacenters. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. Seattle, Washington, 449–471.
- [24] Inigo Goiri, William Katsak, Kien Le, Thu D. Nguyen, and Ricardo Bianchini. 2013. Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy. In *ACM Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Vol. 48. Houston, Texas, 51–64.
- [25] I. Goiri, T. Nguyen, and R. Bianchini. 2015. CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters. In *ASPLOS*, Vol. 50. 253–265.
- [26] Synergy Research Group. 2020. *Hyperscale Data Center Count Reaches 541 in Mid-2020; Another 176 in the Pipeline*. Technical Report. Synergy Research Group, Reno, NV (United States). <https://www.srgresearch.com/articles/hyperscale-data-center-count-reaches-541-mid-2020-another-176-pipeline>
- [27] Benjamin Hindman, Andy Konwinski, Matei Zaharia, Ali Ghodsi, Anthony Joseph, Randy Katz, Scott Shenker, and Ion Stoica. 2011. Mesos: A Platform for Fine-grained Resource Sharing in the Data Center. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. Boston, Massachusetts, 295–308.
- [28] Jeremy Hsu. 2019. How YouTube led to Google's cloud-gaming service: The tech that made YouTube work everywhere promises to do the same for games. *IEEE Spectrum* 56, 09 (2019), 9–10.
- [29] C. Imes, H. Zhang, K. Zhao, and H. Hoffman. 2019. CoPPER: Soft Real-Time Application Performance Using Hardware Power Capping. In *International Conference on Autonomic Computing (ICAC)*. 31–41.
- [30] Rishikesh Jha, Stephen Lee, Srinivasan Iyengar, Mohammad H. Hajiesmaili, David Irwin, and Prashant Shenoy. 2020. Emission-aware Energy Storage Scheduling for a Greener Grid. In *ACM International Conference on Future Energy Systems (e-Energy)*. 363–373.
- [31] Penny Jones. 2012. DataCenterDynamics, Apple Confirms Solar Farm at Maiden Data Center. <https://www.datacenterdynamics.com/en/news/apple-confirms-solar-farm-at-maiden-data-center/>.
- [32] J. Koomey. 2011. Growth in Data Center Electricity Use 2005 to 2010. <https://www.koomey.com/research.html>.
- [33] MLPerf. 2021. MLPerf™ v1.0 Inference Closed-Power ResNet-v1.5 server/offline (entries 1.0-{70,72,73,74}). <https://mlcommons.org/en/inference-datacenter-10/>
- [34] E. Niiler. 2020. Wired, Do Carbon Offsets Really Work? It Depends on the Details. <https://www.wired.com/story/do-carbon-offsets-really-work-it-depends-on-the-details/>.
- [35] Kevin O'Sullivan. 2020. The Irish Times, Facebook Commits to Net-Zero Carbon Emissions by 2030. <https://www.irishtimes.com/news/environment/facebook-commits-to-net-zero-carbon-emissions-by-2030-1.4354701>.
- [36] Darshan S. Palasamudram, Ramesh K. Sitaraman, Bhuvan Urgaonkar, and Rahul Urgaonkar. 2012. Using Batteries to Reduce Power Costs of Internet-Scale Distributed Networks. In *ACM Symposium on Cloud Computing (SoCC)*. 1–14.
- [37] Srinivas Pandravadu. 2014. Running Average Power Limit. <https://01.org/blogs/2014/running-average-power-limit-%E2%80%93rapl>.
- [38] S. Pelley, D. Meisner, P. Zandevakili, T. Wenisch, and J. Underwood. 2010. Power Routing: Dynamic Power Provisioning in the Data Center. In *ASPLOS*, Vol. 38. 231–242.
- [39] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. 2009. Cutting the Electric Bill for Internet-Scale Systems. In *SIGCOMM*. 123–134.
- [40] A. Radovanovic. 2020. Google Blog, Our data centers now work harder when the sun shines and wind blows. <https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/>.



- [41] Jerome H. Saltzer, David P. Reed, and David D. Clark. 1984. End-To-End Arguments in System Design. *ACM Transactions on Computer Systems* 2, 4 (November 1984), 277–288.
- [42] Prateek Sharma, Tian Guo, Xin He, David Irwin, and Prashant Shenoy. 2016. Flint: Batch-Interactive Data-Intensive Processing for Transient Servers. In *ACM European Conference on Computer Systems (EuroSys)*. London, United Kingdom, 1–15.
- [43] Prateek Sharma, Stephen Lee, Tian Guo, David Irwin, and Prashant Shenoy. 2015. SpotCheck: Designing a Derivative Cloud on the Spot Market. In *ACM European Conference on Computer Systems (EuroSys)*. Bordeaux, France, 1–15.
- [44] Arman Shehabi, Sarah Josephine Smith, Dale A. Sartor, Richard E. Brown, Magnus Herrlin, Jonathan G. Koomey, Eric R. Masanet, Nathaniel Horner, Ines Lima Azevedo, and William Linter. 2016. *United States Data Center Energy Usage Report*. Technical Report LBNL-1005775. Lawrence Berkeley National Lab (LBL).
- [45] Kai Shen, Arrvindh Shriraman, Sandhya Dwarkadas, Xiao Zhang, and Zhuang Chen. 2013. Power Containers: An OS Facility for Fine-grained Power and Energy Management on Multicore Servers. In *ACM Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 65–76.
- [46] Prashant Shenoy and Thomas Wensich. 2015. *NSF Workshop on Sustainable Data Centers*. Technical Report. National Science Foundation.
- [47] Rahul Singh, David Irwin, Prashant Shenoy, and K.K. Ramakrishnan. 2013. Yank: Enabling Green Data Centers to Pull the Plug. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 143–156.
- [48] Brad Smith. 2020. Official Microsoft Blog, Microsoft will be Carbon Negative by 2030. <https://blogs.microsoft.com/blog/2020/01/16/microsoft-will-be-carbon-negative-by-2030/>.
- [49] S. Sorrell. 2009. Jevons' Paradox Revisited: The Evidence for Backfire from Improved Energy Efficiency. 37, 4 (2009), 1456–1469.
- [50] Rick Stevens, Valerie Taylor, Jeff Nichols, Arthur Barney Maccabe, Katherine Yelick, and David Brown. 2020. *AI for Science*. Technical Report. Argonne National Lab.(ANL), Argonne, IL (United States).
- [51] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and Policy Considerations for Modern Deep Learning Research. In *AAAI Conference on Artificial Intelligence (AAAI)*. 13693–13696.
- [52] Supreeth Subramanya, Tian Guo, Prateek Sharma, David Irwin, and Prashant Shenoy. 2015. SpotOn: A Batch Computing Service for the Spot Market. In *Proceedings of the Sixth ACM Symposium on Cloud Computing (SoCC)*. Kohala Coast, Hawai'i, 1–13.
- [53] J. Switzer, R. McGuinness, P. Pannuto, G. Porter, A. Schulman, and B. Raghavan. 2021. *TerraWatt: Sustaining Sustainable Computing of Containers in Containers*. Technical Report. arXiv:2102.06614 <https://arxiv.org/abs/2102.06614>
- [54] Domenico Talia. 2013. Clouds for scalable big data analytics. *Computer* 46, 5 (2013), 98–101.
- [55] Maud Texier. 2021. Google Cloud: A timely New Approach to Certifying Clean Energy. <https://cloud.google.com/blog/topics/sustainability/t-eacs-offer-new-approach-to-certifying-clean-energy>.
- [56] R. Uргаonkar, B. Uргаonkar, M. Neely, and A. Sivasubramaniam. 2011. Optimal Power Cost Management Using Stored Energy in Data Centers. In *SIGMETRICS*. 221–232.
- [57] M. Weiser, B. Welch, A. Demers, and S. Shenker. 1994. Scheduling for Reduced CPU Energy. In *OSDI*.
- [58] J. Wilkes. 2020. *Google Cluster-Usage Traces v3*. Technical Report. Google Inc. Posted at <https://github.com/google/cluster-data/blob/master/ClusterData2019.md>.
- [59] Ying Yan, Yanjie Gao, Yang Chen, Zhongxin Guo, Bole Chen, and Thomas Moscibroda. 2016. TR-Spark: Transient Computing for Big Data Analytics. In *ACM Symposium on Cloud Computing (SoCC)*. Santa Clara, California, 484–496.
- [60] Youngseok Yang, Geon-Woo Kim, Won Wook Song, Yunseong Lee, Andrew Chung, Zhengping Qian, Brian Cho, and Byung-Gon Chun. 2017. Pado: A Data Processing Engine for Harnessing Transient Resources in Datacenters. In *ACM European Conference on Computer Systems (EuroSys)*. Belgrade, Serbia, 575–588.
- [61] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *USENIX Symposium on Networked System Design and Implementation (NSDI)*. 15–28.