

Combining Renewable Solar and Open Air Cooling for Greening Internet-Scale Distributed Networks

Vani Gupta

University of Massachusetts, Amherst
vani@cs.umass.edu

Prashant Shenoy

University of Massachusetts, Amherst
shenoy@cs.umass.edu

Ramesh K. Sitaraman

University of Massachusetts, Amherst
Akamai Technologies
ramesh@cs.umass.edu

ABSTRACT

The widespread adoption and popularity of Internet-scale Distributed Networks (IDNs) has led to an explosive growth in the infrastructure of these networks. Unfortunately, this growth has also led to a rapid increase in energy consumption with its accompanying environmental impact. Therefore, energy efficiency is a key consideration in operating and designing these power-hungry networks. In this paper, we study the greening potential of combining two contrasting sources of renewable energy, namely solar energy and Open Air Cooling (OAC). OAC involves the use of outside air to cool data centers if the weather outside is cold and dry enough. Therefore OAC is likely to be abundant in colder weather and at night-time. In contrast, solar energy is correlated with sunny weather and day-time. Given their contrasting natures, we study whether synthesizing these two renewable sources of energy can yield complementary benefits. Given the intermittent nature of renewable energy, we use batteries and load shifting to facilitate the use of green energy and study trade-offs in brown energy reduction based on key parameters like battery size, number of solar panels, and radius of load movement. We do a detailed cost analysis, including amortized cost savings as well as a break-even analysis for different energy prices. Our results look encouraging and we find that we can significantly reduce brown energy consumption by about 55% to 59% just by combining the two technologies. We can increase our savings further to between 60% to 65% by adding load movement within a radius of 5000kms, and to between 73% to 89% by adding batteries.

CCS CONCEPTS

• **Computer systems organization** → **Cloud computing**; • **Hardware** → **Power and energy**; **Renewable energy**.

KEYWORDS

internet-scale distributed networks, energy-efficient data centers, renewable energy, open air cooling

ACM Reference Format:

Vani Gupta, Prashant Shenoy, and Ramesh K. Sitaraman. 2019. Combining Renewable Solar and Open Air Cooling for Greening Internet-Scale Distributed Networks. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems (e-Energy '19)*, June 25–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3307772.3328307>

1 INTRODUCTION

Internet-scale Distributed Networks (IDNs) are massive geographically distributed networks of inter-connected data centers housing hundreds of thousands of servers. Content Delivery Networks (CDNs) are examples of such networks, and they deliver most of the Internet traffic content today, e.g. streaming media, web applications, social networking content, web objects and other downloadable content. Given the immense size of these networks, they consume massive amounts of energy incurring energy bills that run into tens of millions of dollars annually [41]. Growth in data center electricity usage slowed down from 2005 to 2010 as compared to the previous five years from 2000 to 2005 due to the economic slowdown, virtualization, and other efficient data center practices [28]. However, regardless of that, electricity usage of data centers in the US still grew by 36% from 2005 to 2010 totaling about 2% of total US electricity use in 2010 [28].

Given the energy costs and its environmental impact, major IT companies like Google, Facebook, Apple have all committed to greener practices and renewable sources of energy. Google matched 100% of the 2017 electricity consumption of their global operations with renewable energy purchases [23]. Facebook has committed to reducing its greenhouse gas emissions by 75% and powering its global operations with 100% renewable energy by the end of 2020 [18]. In a 2018 press release Apple has stated that they are globally powered by 100% renewable energy [3].

Significant research has been done on making data centers energy-efficient. Part of this work is focused on reducing energy consumption itself [30] [33] [11] [43]. Other work has focused on utilizing renewable energy via local load scheduling, geographical load balancing and data center provisioning and site selection [21] [20] [32] [31] [19] [6]. However, although cooling energy is a major portion of the energy consumed by a data center, efficiency in data center cooling has received much less attention in comparison. In this paper, we take a more comprehensive view of energy consumption in a data center and consider not only energy to power servers, but also energy used for cooling. We study the greening potential of synthesizing two contrasting sources of renewable energy: solar energy and a renewable form of cooling known as open air cooling (OAC). We note that solar energy is more abundant in sunny locations and during day-time. In contrast, OAC is available when the

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.
e-Energy '19, June 25–28, 2019, Phoenix, AZ, USA
© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6671-7/19/06...\$15.00
<https://doi.org/10.1145/3307772.3328307>

weather outside is cold and dry enough. Therefore, OAC is available in colder locations and during night-time. We evaluate if the contrasting nature of these two technologies yields complementary benefits. Given renewables are intermittent in general, and the renewables we have chosen to study are complementary in time and space, we use batteries and load shifting for smoothing the supply of green energy. We study the greening potential of combining these two technologies against two yardsticks: reduction in brown energy and cost effectiveness. To realistically evaluate the greening potential, we use an extensive real-world load trace from Akamai, one of the leading CDN providers in the world [38].

Our work is most applicable to IDNs like CDNs that have a global deployment of servers and replication of services and content. CDNs have an extensive network of servers scattered all over the globe so they can be proximal to their end-users. In addition, their content is replicated widely so that it can be served reliably and with low latency to end-users. We take advantage of both these defining CDN features in designing our solution as they allow us to move load between data centers. This load movement can affect latency, and so in our solution we consider different radii of load movement and incorporate it as a variable parameter into our analysis.

Contributions: To the best of our knowledge, our solution is novel as it synthesizes two renewable technologies, solar energy and OAC, and evaluates their greening potential in the context of an IDN, with large-scale real-world load traces. Specifically, our contributions include:

- *Synthesizing solar energy and OAC as contrasting and complementary technologies:* Motivated by the contrasting and complementary nature of solar energy and OAC, we use a simple greedy algorithm that enables us to use solar energy and OAC efficiently. A *net-zero year* (nzy) data center produces as much energy from renewables in a year as it needs to entirely offset its brown energy consumption in that year. Just by introducing OAC alone to the mix of half the number of panels it takes to be net-zero year, we show that we can go from 34% reduction to about 54.9% brown energy reduction. With panels needed to be net-zero year, we can go from 41.5% to about 59.4% savings. We see even higher savings by employing both batteries and load movement. We incorporate several key parameters that can be used to model trade-offs while evaluating energy efficiency. Some of these parameters include radius of load movement, battery capacity, number of solar panels installed, battery cost and lifetime, solar panel cost and lifetime, and energy prices.

- *Evaluation using an extensive real-world trace:* We evaluate the greening potential of solar energy and OAC extensive load traces from Akamai [38]. The traces consist of information on from 724 global data center locations including 100,592 servers deployed all over the world. We also use year-long weather data for OAC from over 650 locations. In addition, we use a year's worth of PVWatts solar data. Using this data, we simulate our solution for a whole year, parallelizing our runs by week to reduce the time of running. We then evaluate our solution against several metrics measuring total brown energy reduction, peak reduction, cost savings and a break-even analysis. We vary battery capacity as a function of the average day's load in a data center.

- *Brown energy reduction evaluation:* We evaluated how well the mix of solar energy and OAC reduces brown energy consumption.

Energy companies often charge their customers for both the energy consumed and the peak energy drawn. As part of this analysis, we studied two metrics: 1) total brown energy reduction and 2) peak energy reduction. For brown energy reduction, we studied how our results vary with addition of OAC to solar energy, with the addition of load movement, and also with the addition of batteries.

Allowing a radius of 5000kms with the combination of solar energy and OAC, we can increase our savings to 60.3% or panels0.5 and to about 65% with net-zero year panels. Our results show that with a battery capacity of half the average day's load at each data center, we can significantly increase the reduction in brown energy to over 73% for panels0.5 and over 89% with net-zero year panels, without moving any load. For percentage peak reduction, we see a reduction between 10% and up to 40% depending upon the number of panels installed, the battery capacity and radius of load movement. Fixing the radius of load movement to 1000kms, and varying battery capacity and panels as shown above, we can achieve a reduction of about 11% in the worst case to about 26% with greater battery capacity and larger number of panels.

- *Amortized cost analysis:* We evaluated the cost saving potential of our solution given investment in different combinations of battery capacities and number of panels. We calculated yearly cost savings based on yearly savings in brown energy consumption costs and yearly amortized expenditure for batteries and panels. We find significant cost savings for moderate and high energy prices, ranging from 9.9% all the way to 60.3% based on different parameter values. Even for low price for energy, if we do not use batteries and have 0.5nzy panels, we see cost savings from 22% to 41%. However, with 0.5avgdayload batteries and 0.5nzy panels, savings drop to between 3% to about 8.4%, and for other combinations of panels and batteries we incur a loss in the case of low price of energy. With the prices of batteries and solar panels on the decline, we believe the results for lower energy prices should also improve in the future.

- *Break-even analysis:* With a higher price of energy, for half the panels it takes to be nzy, we see break-even periods as low as 6 years. For a moderate and low energy prices, we can achieve break-even periods of 8.9 years and between 14.9 years respectively. Again, with the cost of solar panels and batteries declining, these numbers should improve in the future.

- *Cost Analysis based on future projections:* Given the price of solar panels and batteries is falling, and the price of energy over the long run is increasing, we evaluated our solution using projected prices of batteries, panels, and energy. We found dramatic increases in brown energy reduction and break-even periods even for the projected lower end price of energy. Even for the low price of energy, for which we incurred a loss in certain cases with current prices, we see cost savings of 23.9% to 55.9%.

Roadmap: We present the background in Section 2 and our algorithm in Section 3. We present our experimental methodology and empirical results in Sections 4 and 5 respectively. We end with related work and conclusions in Sections 6 and 7 respectively.

2 BACKGROUND

Internet-Scale Distributed Networks: An IDN provides modern Internet services via its network of servers housed in a large number of data centers spread all over the globe. An example of an IDN is

a Content Delivery Network (CDN) that serves content to clients on the web reliably and with low latency. The three main entities in a CDN system include the content providers, the CDN provider, and the end users [9]. Content providers interested in distributing their content to end-users contract with CDN providers so they can use the CDN’s infrastructure to help distribute their content transparently, reliably, and in a timely fashion. Content is replicated by the CDN’s distribution systems to edge servers located in a diverse set of geographically distributed locations. On receiving a user request, the request routing system assigns the user to the appropriate nearby server to ensure low response times. Therefore, the two defining characteristics of a CDN are *global deployment* of servers and *replication* of services. Both these features work in conjunction with each other to provide services that are proximal to end-users. We use these two features to enable us to move load between data centers, although with a possible increase in latency. We move load in a constrained manner by restricting the radius of load movement, and study the greening potential of solar energy and OAC with radius of load movement as a variable parameter.

Data centers require massive amounts of energy to run and maintain servers and other supporting equipment. The bulk of the energy consumed by a data center comes from powering its servers and for cooling them. About 56% energy is used to power servers and about 30% is used for cooling and the rest 14% is used for power conditioning and networking equipment [40]. In this paper, we refer to the energy used to power servers as ‘*server energy*’ and the energy used for cooling as ‘*cooling energy*’.

Open Air Cooling (OAC): Data centers need cooling to keep server and other equipment at recommended operating temperatures. Depending upon weather conditions existing in the data center location, air from outside can be brought into the data center to cool servers. Stated simply, OAC involves the use of outside air to cool data centers. Broadly, there are two flavors of OAC: air-side and water-side. Air-side uses outside air for cooling, and water-side use water as a cooling medium circulating through cooling towers. Another version of air-side free cooling is evaporative cooling where outside air in conjunction with evaporating water is used for cooling. Given cooling energy accounts for a significant portion of the energy consumption, renewable cooling using outside air can be considered to be virtually ‘free’ in comparison to using HVAC chillers - and is therefore sometimes also referred to as ‘free cooling’. Given OAC uses outside air for cooling, the availability of OAC depends on the temperature and humidity conditions existing outside. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) has defined different classes of data centers based on the temperature and humidity ranges that they can tolerate [26]. Classes are labeled A1 through A4, with most restrictive to least restrictive. We assume our data centers belong to the A1 class, which requires the smallest range for temperature and humidity, and represents more commonly deployed basic equipment today. We use the existing weather conditions outside and the ASHRAE requirements for A1 class of data centers to determine whether OAC is available or not. Using the most restrictive class also gives us the ability to do the worst case analysis with respect to OAC availability. A1 ranges for temperature and humidity are listed in Table 1.

Class	Dry-Bulb Temp (°C)	Humidity Range	Max Dew Point (°C)
A1	15 to 32	20% to 80%	17

Table 1: ASHRAE’s allowable ranges for temperature and humidity for A1 class of data centers.

Modeling Server and Cooling Energy: The main source of energy consumption in a data center is the energy used to power and cool servers. To model server energy, we use the linear model of energy consumption [5]. Energy to power servers is determined using normalized load λ (where $0 \leq \lambda \leq 1$), where λ is the load on the server as a fraction of server capacity. Idle servers also consume approximately 60% of energy. Hence, we use λ equals $P_{idle} + (P_{peak} - P_{idle})\lambda$, as the power consumed by a server serving normalized load λ , where P_{idle} is the power consumed by an idle server, and P_{peak} is the power consumed by the server under peak load. To determine the total energy consumed by the cluster of servers in a data center, we assume that we can consolidate load between servers, and shutdown the remaining servers that are idle to conserve energy [30]. Power Usage Effectiveness (PUE) is the measure of how efficiently a data center uses energy. It is the ratio of the total energy used by the data center (including energy used by the IT equipment, cooling energy, and other overhead) and the amount of energy used by the IT equipment. We use the average PUE of 1.8 [34] when determining cooling energy consumed by the data center.

Geographical Variations in Solar Energy and OAC Availability: We see variations in solar output and OAC based on factors like temperature, season, time of day, northern or southern hemisphere location, climate, weather conditions [24] [25]. Therefore, using renewables efficiently involves handling the variations in and availability of renewable output. In this paper, we use a combination of load movement and battery storage to mitigate the problem of intermittent availability of solar energy and OAC. Given the geographical diversity of data center locations and replicated content and services, we use load shifting to take advantage of renewables. We vary radii of load movement to control latency. To enable us to store excess solar energy, we assume that batteries are available at all data center locations. We vary battery capacity installed at a data center location as a function of the average day’s load for that data center. Our analysis can also be modified to include net-metering. However, given net-metering is not consistently available globally, we make a simplifying assumption that all data center locations employ batteries to store excess solar energy.

Net-zero Data center: A ‘net-zero energy’ data center is designed and managed in a manner that uses on-site renewables to entirely offset the use of any non-renewable energy from the grid [4]. Therefore, given a period of time, a net-zero data center produces at least as much on-site renewable energy as it consumes. With this definition, a ‘net-zero year’ data center is net-zero over the period of a year. For a data center, we define the *net-zero year solar panels* as the number of solar panels needed by the data center to be net-zero year. For our experiments, we vary the number of solar panels installed at a data center as a function of its *net-zero year solar panels*.

Variable and Value	Notation
battery capacity = x *(avg day's load)	bcapx
num solar panels = x *(net-zero year number of panels)	panelsx or xnzy
radius of load movement = xkms	r=x

Table 2: Parameters values and related notation used to refer to them in the paper

Metrics for evaluating proposed solution: To evaluate the combined greening potential of solar energy and OAC, we measure reductions in both *energy consumption* and *cost*. We use reduction in total brown energy consumption and reduction in peak energy drawn from the grid to determine how effective the combination of solar energy and OAC is in greening IDNs. We use amortized cost savings and a break-even analysis to evaluate how effective the algorithm is with respect to cost.

Parameter Values and Related Notation: In this paper, we study our algorithm by varying parameters like battery capacity and number of solar panels. We vary battery capacity installed at a data center as a function of the average day's load for that data center. We consider three different fractions: 0, 0.5*(average day's load), and 1*(average day load). We vary the number of solar panels as a function of the net-zero number of panels for a data center. We consider two fractions: 0.5*(net-zero year number of panels), and 1*(net-zero year number of panels). In addition to these, we also vary the radius of load movement and use a notation $r=x$ to mean that a maximum radius of load movement of x kms was used in our simulation. It is cumbersome to refer these cases using their full descriptive text for battery capacity and panels as listed above. Therefore, we use a shorter notation and list the mapping of the full text to its notation in Table 2. For example, to refer to a case in which we employ a battery capacity of 0.5*(average day's load) and install 0.5*(net-zero year number of panels), in our plots and empirical results we use a notation *bcap0.5 and panels0.5 (or 0.5nzy)*.

Problem Statement: IDNs consume massive amounts of energy. The bulk of the energy consumed by data centers consists of energy used to power servers and to cool them [40]. One way IDNs can be made greener is by replacing brown energy consumption by energy generated from renewable sources. Solar energy is correlated with sunny weather and day-time. In contrast, OAC is more abundant in colder weather and night-time. In this paper, we study the potential of using two *contrasting* and *complementary* sources of renewable energy (namely solar energy and OAC) in their ability to reduce brown energy consumption in IDNs in a cost effective fashion. Given the intermittent nature of renewable energy, in general, and the complementary nature of these two specific sources, we use batteries and load movement as facilitators for smoothing supply of green energy. Specifically, in this paper we try to study two aspects:

- *The potential for replacing brown energy with a combination of solar energy and OAC in IDNs.*
- *The cost effectiveness of combining these two contrasting sources of renewable energy in our IDN setting.*

3 ENERGY-AWARE LOAD SCHEDULING ALGORITHM

We describe our greedy heuristic algorithm in the following paragraphs. We assume that we have the ability to cool load using OAC

Parameter	Value
Loss %	14
System Capacity	0.275 kW
Module Type	Standard
Timeframe	Hourly
Azimuth	180 deg for northern hemisphere and 0 for southern
Tilt	Absolute value of latitude
Dataset	'TMY2' for US Locations and 'Intl' for others

Table 3: Parameters for PVWatts Data

as long as the weather conditions outside permit us to do so. We also assume we have the on-site solar panels at each data center location. Further, we assume that we have batteries available locally to store excess solar energy. Finally, we assume we can leverage redundancy and data replication in IDNs by moving load to locations where there is more renewable energy available.

Our algorithm works as follows. If OAC is available, we use that for cooling data centers. If solar energy is being generated by locally installed solar panels, we use that to meet local energy demand, including cooling energy if OAC is not available. For remaining server and cooling load, we use locally installed batteries. If any load is left over, we try to shift it to other locations with surplus green energy, constrained by a maximum radius of load movement. We do load shifting in two iterations. In the first iteration, we move load to locations that have both surplus solar energy and OAC. In the second iteration, we move load to locations that have surplus solar energy and no OAC. This allows us to use solar energy from data centers that did not get selected in the first iteration. For both iterations, load shifting is constrained to remain within a maximum radius of load movement to control latency. Finally, for any remaining load, we draw energy from the grid. We store any unused solar energy in batteries for future use. The pseudo-code for the algorithm is listed in Appendix A.

4 EXPERIMENTAL METHODOLOGY

We performed our experiments on a month-long Akamai trace. This extensive trace has a granularity of 5 minutes and consists of information on 100,592 servers in 724 global data center locations from around the world. The data set consists of information for fields like load, requests, bytes, number of servers, server capacity, latitude, longitude, city, state, and country.

Our solar data set was acquired from the PVWatts [37] website. It consists of a year-long dataset for solar energy generation at a granularity of one hour. We assume that the power rating of a solar panel ranges from 200 watts to 350 watts [14] and take an average value of 275 watts as the power rating per panel. We list values of parameters used for PVWatts solar data in Table 5. For any other required parameters, we used the default values listed in the PVWatts version 5 manual [12].

For determining OAC availability we used a year-long weather dataset for the year 2012 from the National Oceanic and Atmospheric Administration (NOAA). This global dataset contains several metrics, including hourly dry-bulb temperature and dew point. Given that the location of our data centers, we mapped which weather station was closest and used its weather data as being representative of the weather at the data center location. Given the

NOAA has a vast network of weather stations, we could map most of our data centers to weather stations within 10kms. For most of the remaining data centers, we could map a weather station within 40kms.

Weather data used for OAC and solar data had a granularity of one hour. However, the load trace has a granularity of 5 minutes. We therefore assumed that the weather and solar output do not change much during the hour, and use the hours value for each of the 5-minute timeslots that fall within the hour. Additionally, our weather data and solar energy data was year-long, however, the Akamai load trace was month-long. To simplify, we assumed that the load trace pattern repeats throughout the year. However, our algorithm does not fundamentally depend upon or exploit the fact that the load pattern repeats throughout the year. Therefore, it would also be applicable to a yearly load trace in which the load pattern is different for each month.

We analyzed our metrics by varying several parameters. For a given data center, we varied battery capacity as a function of the average day’s load, and considered battery capacities of zero, half of the average day’s load, and a full average day’s load. For each data center, we determined the number of solar panels we need to be net-zero year i.e the number of panels needed to produce enough solar energy to cover the total energy needs of the data center for a year. For our experiments, we varied the number of panels from half of the net-zero year number of panels to a full net-zero year number of panels. Given the size of our datasets, running our algorithm sequentially would have been computationally expensive. Therefore, we parallelized our algorithm by week and in order to do a worst case analysis, we assumed a starting battery charge of zero at the beginning of each week.

5 EMPIRICAL RESULTS

We evaluated the greening potential of solar energy and OAC in the context of both brown energy reduction and cost effectiveness. We analyzed several metrics, namely brown energy reduction, peak reduction, cost savings, and break-even points. We describe our findings related to these metrics in the paragraphs below.

5.1 Brown Energy Reduction

Brown energy reduction is calculated by taking the average of percentage reduction in brown energy across all the data centers of the IDN for the year. Our results show the following:

- *Combining solar energy and OAC yields significant benefits:*

Figure 1 shows the brown energy reduction we can achieve with the combination of solar energy and OAC by different months of the year. Solar energy output is higher in the summer months when there is plenty of sunshine. Therefore, we see the reduction in brown energy peak in the summer months when we use solar energy alone. In contrast, OAC is more abundant when the outside weather is cold and dry enough. Therefore savings from OAC are higher in the winter months and dip in the summer months. Combining these two technologies, we can achieve a much higher savings of between 49.7% to about 60% throughout all the months of the year as shown by the green line. Figure 2 shows how our yearly average percent savings increase when we combine solar energy with OAC. As seen by comparing the left two bars of Figure 2 (a) and (b), just by




Figure 1: Plot shows how solar energy and OAC combine to yield higher savings across various months of the year for panels0.5 and r=0.

introducing OAC alone to the mix of 0.5nzy panels, we can go from 34% reduction to about 55% average brown energy reduction. With nzy panels, we can go from 41.5% to about 59.4% savings.

- *Load movement leads to more savings:* As seen in Figure 2 (a and b), savings increase with increasing r. For r=5000kms, we can increase our average reduction from 54.9% to 60% for panels=0.5nzy and from 59.4% to 65% for nzy panels.

- *Batteries help significantly:* As seen by the leftmost bars in Figure 2 (c and d), in the absence of batteries, doubling the number of solar panels increases savings from 34% to 41.5% for the solar energy only scenario and from 54.9% to only about 59.4% for the combination for solar energy and OAC. Without batteries, instantaneous solar energy produced is wasted. However, as shown by the bars to the right in Figure 2 (c), by employing batteries with bcap0.5, we can significantly increase the reduction in brown energy to over 48% for panels0.5 and over 74.9% for nzy panels with solar energy alone. For the combination of solar and OAC in Figure 2 (d), we can increase savings to 73% for 0.5 net-zero year panels and over 89% with net-zero year number of panels.

- *Diminishing returns with increase in battery capacity:* Reduction in brown energy increases with larger battery capacity, however, we see diminishing returns. Figure 2 (d) shows the jump in savings from 0 battery capacity to 0.5 is dramatic – from 54% to 73% for 0.5nzy panels. However the jump from 0.5 to 1 is not that large – 73.2% to 73.7%. For a larger number of solar panels (shown by the red bars in Figure 2 (d)), the same diminishing returns with batteries are observed and we see a jump in reduction from 59% to 89% to 91% as we increase the battery capacity from 0 to 0.5 to 1. This trend is also preserved for the solar energy only scenario as we can see from Figure 2 (c).

- *Application-specific parameter values:* We can achieve similar gains in brown energy reduction with different sets of parameter values. These parameter values could be chosen based on the specific needs of applications, e.g. we may choose to not move load for latency sensitive applications, whereas for latency tolerant applications, we may choose to move load and save on battery costs. As an example, suppose we would like to achieve approximately 70% reduction in brown energy consumption. We can achieve this in

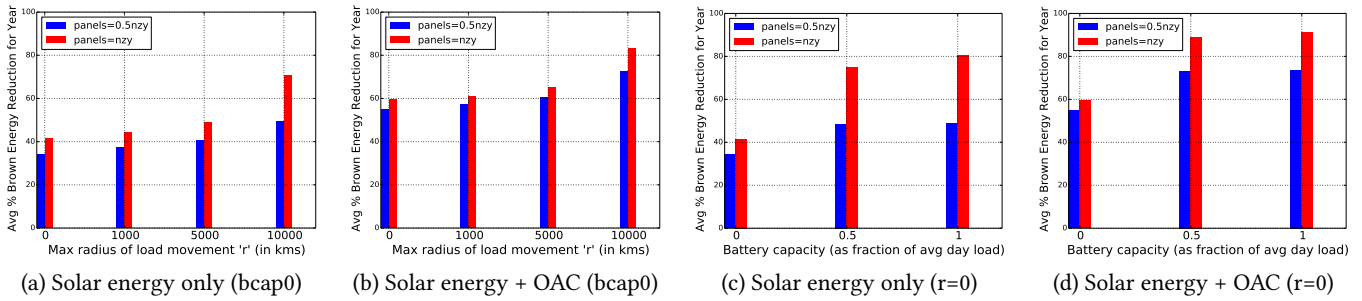




Figure 2: We see a significant increase in brown energy reduction as we move from solar energy only (a & c) to solar energy + OAC (b & d). Increasing r (a & b) yields larger savings. Increasing battery capacity (c & d) helps but shows diminishing returns.

two different ways using different combinations of load movement, battery capacity, and solar panels. The two ways from the above plots are: From from Figure 2 (b), bcap0 panels0.5 and r=10,000 and from Figure 2 (d), panels0.5 bcap0.5 with r=0. The former scenario is better suited for applications that can tolerate latency, where as the latter can be employed in case of latency-sensitive applications though with an added expenditure for batteries.

months. However, for a place like Las Vegas (see lowest blue line corresponding to panels0.5 bcap0 in Figure 3(b)), where solar energy is available for most of the year, we get a curve that dips in the summer months, when OAC is not as abundant. These shapes change with the addition of load movement and batteries to the mix, as both of those alter the basic assumptions about locational variations of OAC and solar. Also, locations that are mostly high in solar energy output (e.g. Las Vegas which is ranked as the third highest city in the United States based on percentage annual sunshine [35]), have an advantage over locations that are excellent for OAC year round (e.g. Anchorage where the highest average year round temperature is 19 °C and the average dew point is -2 °C [47]). Solar output can be used for meeting both server energy demand, as well as for cooling purposes. However, OAC can only be used for cooling. From the plots, with sufficiently high number of solar panels and battery size, we can nearly see a high reduction in brown energy consumption year round. For Anchorage, however, in the summer months we see a dip in brown energy reduction due to lesser solar energy availability. The curves also show diminishing returns when battery capacity is increased successively from zero, to half of the average day's load, to a full average day's load.



(a) Anchorage for r=1000kms



(b) Las Vegas for r=1000kms

Figure 3: Figure showing reduction in brown energy across different months for Anchorage and Las Vegas

- *Location Based Results:* Trade-offs for specific locations vary significantly depending on the local availability of solar energy and OAC and their interplay. For a place like Anchorage (see lowest blue line corresponding to panels0.5 bcap0 in Figure 3(a)), where OAC is available for most of the year, the shape of the curve depends on the availability of solar energy, which peak in the summer

5.2 Peak Reduction

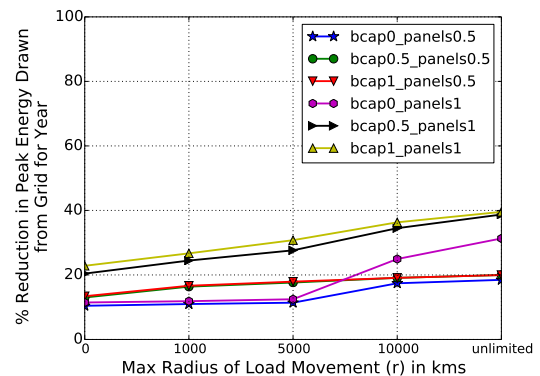


Figure 4: Plot showing significant gains in peak reduction. Increasing solar panels, battery capacity and r result in higher reductions.

This metric measures the average percentage peak reduction for peak energy drawn from the grid for the year. We first determine

Resource	Parameter	Value
Battery	Price/kWh lifetime	\$190/kWh 10 yrs
Solar Panels	Price/Wac lifetime	\$2.1/Wac 25 yrs

Table 4: Price and lifetime for batteries and solar panels. Cost for commercial solar panels and lithium-ion batteries was used.

the maximum energy drawn for a data center for the year for the original load trace. We then determine the maximum energy drawn for the new load incorporating solar panels, OAC and load movement (for $r > 0$) under the greedy heuristic algorithm. We then calculate the percentage reduction for each data center based on the above values and finally average them. Our results are shown in Figure 4.

- *Significant reduction in peak energy:* As shown in Figure 4, we can see an overall reduction between 10% and up to 40% depending upon the number of panels installed, the battery capacity and radius of load movement. Fixing the radius of load movement to 1000kms, and varying battery capacity and panels as shown above, we can achieve a reduction of about 11% in the worst case to about 26% with greater battery capacity and larger number of panels. With a larger radius of load movement, we can see significantly higher percentages of reduction. As an example, with a $r=10,000$ kms we can see a decrease of over 35% with $bcap1$ and nzy panels.

5.3 Cost Analysis

In this section, we evaluate how well the combination of solar energy and OAC performs with respect to cost savings. To this end, we consider the following aspects: 1) Yearly amortized cost savings and 2) Break-even analysis. We describe these in detail in the following paragraphs.

With the battery and solar cost and lifetime parameters [7] [44] [36] [13] listed in Table 4, we studied cost savings and break-even periods under three different prices of energy from low, to moderate, to high. The following three scenarios were analyzed:

- **Low Price - 7¢/kWh:** This is closer to the industrial price of electricity in the US [1] and is the lower end price for our analysis.
- **Moderate Price - 12¢/kWh:** This is based on a blended value of 12¢/kWh midway between our low and high cost values of 7¢/kWh and 17¢/kWh.
- **High Price -17¢/kWh:** This is on the higher end of the non-household energy prices found in countries in Europe [17].

5.3.1 Yearly Amortized Cost Savings. We calculate *original yearly cost* of brown energy drawn from the grid for the original trace. We then calculate the *new yearly cost* of brown energy for the new reduced load after incorporating solar panels, OAC and load movement (for $r > 0$) under the greedy heuristic algorithm. To account for the yearly cost of panels and batteries, we calculate expense for panels and batteries and amortize the price over their lifetime to determine the *yearly amortized cost* for these investments. We then add the *yearly amortized cost* to the *new yearly cost*. Finally, we find the percentage reduction in cost using

the *original yearly cost* and *new yearly cost* calculated above. The results for the metric are discussed below.

- *Cost savings are directly proportional to the price of energy:* From Figure 5 we see higher savings in cost as we move from a low to a moderate to a high energy price. With a higher per unit energy price, every unit of brown energy drawn from the grid that is replaced with green energy reduces a larger amount from the operational cost.

- *Significant cost savings for moderate and high energy prices:* As seen in Figure 5, significant cost savings can be achieved for moderate and high energy prices (plots b and c). Savings range from 9.9% to 60.3% based on different parameter values. With moderate energy prices, for $bcap0.5$ and $panels0.5$, we can see a savings of about 32% without any load movement. For the higher price and same battery size and panels, savings are much higher at 44.4%.

- *Savings in some cases with low energy prices:* From Figure 5 (a), we see that with lower energy prices, we can yield cost savings if we employ fewer number of panels ($0.5nzy$) coupled with either no batteries or batteries with a smaller capacity of $bcap0.5$. With $panels0.5$ and $bcap0$, we see savings ranging from 22% to 41% depending on r . With $panels0.5$ and $bcap0.5$, we see a savings of 3% to about 8.4% depending on r . For other combinations of panels and battery capacities, we incur a loss. However, with prices of solar panel installation and batteries on the decline, we expect these cost savings in this case to improve going forward.

- *Middle ground provisioning:* As seen from the green line in subplots of Figure 5, $bcap0.5$ and $panels0.5$ yields no losses for the low energy price and yields significant savings for the higher energy price. This coupled with the fact that $bcap0.5$ and $panels0.5$ yields significant average percent brown energy reduction, (73% for 0.5 net-zero year panels and over 89% with net-zero year number of panels), makes it a good middle ground for achieving both objectives of reducing brown energy consumption and saving on cost.

- *Sensitivity of metric in inversely proportional to energy price:* Generally speaking, this metric is more sensitive to change in parameters (i.e. battery capacity and number of panels) with lower energy prices, as compared to higher energy prices. Observing Figure 5, we see that the lines successively span out less as we go from low to moderate to high prices. For the lower energy price for $r=0$, the savings range from 22% to -48%. For the moderate energy price, savings range from about 35.8% to about 10%. Finally, for the higher energy price, savings range from 46% to about 29%. Therefore, decisions to switch between different battery capacities and number of panels have a greater effect on cost savings when prices are low, as compared to when they are higher.

5.3.2 Break-even Analysis. In this section, we look at the number of years it takes to break even on the expenditure made towards batteries and solar panels. We determine brown energy cost for the year for the original trace and for the new trace after our algorithm has been run. We calculate the difference of these two to get cost savings for the year. We then find the capital expenditure incurred on batteries and solar panels across the IDN, and divide it by the savings for the year to get the number of years it would take to recover the cost.

Figure 7 gives an idea of the break-even period across different combinations of battery capacity and panels. Figure 6 focuses on




Figure 5: Plots show significant amortized savings for moderate and high energy prices. For the lower energy price, we see losses for higher battery capacity and larger number of panels. However, even for the lower energy price, we see significant savings without batteries, and we can see some savings with bcap0.5.




Figure 6: Plot shows a decrease in the number of years to break even as the price of energy goes up (for $r=0$).

$r=0$ and the combination of panels and battery capacity for which the number of break-even years are the lowest:

- *Break-even period is inversely proportional to energy price:* Figure 6 shows that for half the nzy panels and a low energy price, we see a break-even period of about 14.9 years. This falls to 8.7 years for the moderate price, and 6 years for the higher price of energy. The same trend is observed for all combinations of panels and capacities as seen in Figure 7. Therefore, the higher the price of energy, the lower the number of years to break even. This is because for every unit of brown energy reduced, we get larger savings when we multiply it with the higher unit cost of energy versus a lower unit cost of energy.

- *Finding a middle ground:* The break-even period is very similar for 1) bcap0 and panels0.5; and 2) bcap0.5 and panels0.5. For the higher energy price and with bcap0 and panels0.5, it takes between about 4.6 to 6.1 years to break even depending upon the values of r . With bcap0.5 and panels0.5, it takes about the same number of years (between 6.7 to 6.3) to break-even. This trend is also observed for lower and moderate energy prices as well. Therefore, from an overall solution standpoint considering bcap0.5 is useful in brown

Parameter	Cost (constant 2017 dollars)
Lower Electricity Cost Projection (¢/kWh)	7.98
Moderate Electricity Cost Projection (¢/kWh)	13.67
Higher Electricity Cost Projection (¢/kWh)	19.36
Solar Panel Cost (\$/Wac)	1.30
Battery Cost (\$/kWh)	70

Table 5: Projected Electricity, Solar Panel and Battery Costs

energy reduction and cost savings, bcap0.5 and panels0.5 emerges as the preferred option between 1 and 2.

5.4 Cost Analysis with Future Projections

Given the price of solar panels and batteries is on the decline, and the price of energy is on the rise, we evaluated our algorithm for 2030 price projections of electricity, solar panels, and batteries. For electricity prices, we used the projected average US electricity price in 2030 [15], we then calculated the current ratio of the average price across all sectors to the current industrial price of electricity [1] to determine the industrial electricity price for 2030. We then used the percentage increase in price to scale up our low, moderate and high prices used in the paper. We used the SunShot study targets for installed solar panel cost in \$/Watt in the beyond 2020 [16] as well as their 2030 targets [45], in conjunction with the current commercial solar panel per watt rates [36] to determine the installed cost of commercial panels in 2030. We used the Bloomberg New Energy Finance (BNEF) projection for the lithium-ion battery cost in 2030 [7]. Table 5 shows the projected values we used (in constant 2017 dollars). As a simplifying assumption we assumed that the lifetime of batteries and solar panels remains the same as the current values uses. If the lifetime were to increase in the future, that would yield even higher cost savings.

With the projected values of parameters discussed above, we re-looked at how well the algorithm performs with respect to: 1) yearly amortized cost savings for our algorithm and 2) break-even analysis. Our findings are discussed below:




Figure 7: The break even period is inversely proportional to the price of energy. With a moderate amount of battery capacity and panels, we can achieve close to the lowest break even periods compared to others.

5.4.1 Yearly Amortized Cost Savings with Future Cost Projections:

The results for this metric are discussed below.

- *Dramatic increase in cost savings:* As seen in Figure 8, cost savings showed a dramatic increase across the board for all combinations of parameters. Figure 8 (a) shows that for the lower price of energy, range from 23.9% to 55.9%. None of the combinations of parameters result in a loss, like we saw with current prices. From Figure 8 (b) shows that with moderate energy prices, we can see savings of 38.6% to 68.9%. With the future higher energy price, we see even higher savings ranging from 44.7% to 77.06%.

5.4.2 Break-even Analysis with Future Cost Projections:

The results for this metric are discussed below:

- *Dramatic decrease in number of years to break even:* We see a huge decrease in the number of years it takes to break even with the projected prices. From Figures 6 and 9, we can see that for bcap0, panels0.5 and $r=0$, for the new low price, the number of years it takes to break even falls from 14.9 years to 8.08 years. For the moderate price it falls from 8.7 to 4.71, and from 6.1 to 3.33 for the high price. We see the similar trend for bcap0.5, panels0.5 and $r=0$ where the number of years are reduced by approximately half between the current and projected costs. In addition, we see from Figure 9, that the break even years with bcap0.5 are marginally less than without batteries. Given the decrease in the prices of batteries and solar panels, and the higher energy cost, for 0.5nzy panels in the future it would in fact take marginally less time break even if we employ a battery capacity of bcap0.5, than if we do not have any batteries at all.

5.5 Discussion

Our analysis shows that combining solar energy and OAC can significantly reduce brown energy consumption in IDNs. Load movement and batteries can yield further savings. We find that savings due to load movement are most pronounced over larger distances where the the night-day difference is apparent. Therefore applications that are not latency sensitive have the most to gain from load movement. Batteries with a capacity of half of the average day’s load can significantly increase savings. We also see that batteries not only increase savings, but are also cost effective with moderate and high energy prices. Therefore in locations where energy prices are

moderate to high, deploying batteries with solar panels is beneficial. With lower energy prices we can achieve cost savings in certain cases. With future projected prices of solar panels, batteries and energy, we find dramatic increases in cost savings and break even periods for all prices.

6 RELATEDWORK

Given energy efficiency is important for sustainability, significant work has been done in the area of data centers energy management. Part of this work has focused on reducing energy at the server level. Work includes shutting off servers during off-peak times and switching between high and low power states to prevent wear and tear [43] [30] [33] [11]. Allocation of energy between user applications taking into account user priorities and the lifetime of the battery has also been studied [48]. Prior work has also looked at OS level power management by real-time monitoring of the CPU to keep it utilized to a certain percentage [39].

Separately, another part of prior work has focused on energy-efficiency at the data center level. Job scheduling to maximize solar energy usage without violating user deadlines has been studied [21] [22]. Prior work has looked at using solar energy and wind energy prediction to increase green energy usage and cut down canceled jobs [2]. There has been work on job migration between two sets of servers (one powered by energy from the grid and another by wind energy) with the goal of maximizing wind energy usage [29]. Prior work has also looked at energy capacity planning finding the best ratio of renewables given a location and workload or given carbon footprint goals [8] [42]. Given cooling accounts for a large portion of data center energy consumption, work has also been done on use of cooling technologies in modular data centers [27] and on unified management of data centers depending upon renewable availability, cooling efficiency, workload fluctuations, and price of energy [10]. Although the above work provides excellent solutions for data center energy management, it is not targeted towards a network-level setting, which is the focus of this paper.

There has been significant prior on network-level energy management as well. Studies have investigated the use of load balancing using the ‘follow the renewables’ approach to almost entirely power their data centers using a renewable mix of wind and solar energy




Figure 8: Future projection plots show dramatic increases in amortized savings for moderate and high energy prices. For the lower energy price scenario for $r=0$, we see a savings of 23.9% to 55.9% with no losses for any combination. This is an improvement from the current price scenario.

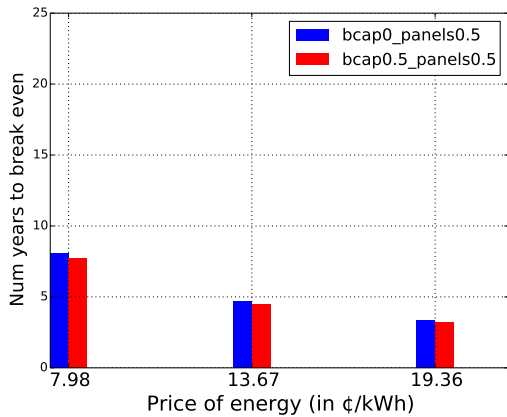


Figure 9: Plot shows a significant decrease in the number of years to break-even with future cost projections (for $r=0$).

[31] [32]. Prior work has also studied user request routing for greening data centers [46]. Solutions have been proposed for dispatching requests to data centers in a way that maximizes renewable energy and stays within a budget [49]. Work has been done to assign users to data centers based on the three-way mix of latency, price of electricity, and carbon footprint [19]. Prior work has also looked into site selection for green data centers using a follow-the-renewables approach [6]. However, none of these studies explicitly consider a combination of solar energy and open air cooling as part of their renewable mix. Most of them do not evaluate their solution on as extensive real-world, global trace as we have done in our paper. These studies also do not explicitly consider the impact of varying storage capacities on their outcomes. Efficient provisioning of solar panels for net-zero IDNs based on geographical solar energy availability has been previously studied [25]. However, this work is for offline panel provisioning. In contrast, we do not focus on solar panel provisioning, and instead we assume that solar panels are installed at every data center location. Existing work has also looked at geographical load movement to study the potential of open air cooling for serving the cooling energy needs of IDNs [24]. However, in this paper, we study the *combined* potential of solar

energy and OAC for net-zero IDNs, considering both *server energy* and *cooling energy* while determining data center energy demand.

7 CONCLUSIONS

In this paper, we studied the greening potential of solar energy in conjunction with OAC given their contrasting natures. To that end, we implemented a simple greedy heuristic and evaluated it on an extensive Akamai load trace. We considered several metrics broadly analyzing brown energy reduction and cost effectiveness of employing a combination of solar energy and OAC in IDNs. We found that just by introducing OAC alone to the mix of 0.5nzy panels, brown energy reduction increases from 34% to about 54.9%. With nzy panels, we can go from 41.5% to about 59.4% savings. We can increase our savings further to between 60% to 65% by adding load movement within a radius of 5000kms. With batteries and $r=0$, we are able to significantly reduce brown energy consumption by 73% (for 0.5nzy panels) and over 89% (for nzy panels). We could also achieve peak energy reduction of about 10% to 40%. Therefore the combination of solar energy and OAC enables significant brown energy savings. Our cost analysis showed that for moderate to higher prices of energy we can achieve significant cost savings from 9.9% to 60.3%. For low energy prices, we found that we can still achieve between 22% to 41% savings with panels0.5 and bcap0. For bcap0.5 panels0.5, we see small savings of between 3% to 8.4%. In other cases with a low energy price, we incurred a loss. With a higher price of energy, we could observe break-even periods as low as 6 to 8.7 years. With energy prices on the rise and solar and battery prices declining, we re-looked at the potential under projected prices. We saw dramatic increases in cost savings, with savings between 23.9% to 55.9% even for the lower projected energy price. With $r=0$ and panels0.5, the number of break-even years reduced significantly by roughly 45% for bcap0 and by roughly 50% for bcap0.5. Overall, we showed that the combination of solar energy and OAC has significant greening potential for IDNs.

Acknowledgements: We thank the anonymous reviewers for their valuable comments. This research was supported by NSF grants 1763834, 1802523, 1836752, 1413998, and 1763617.

REFERENCES

- [1] U.S. Energy Information Administration. 2018. *Industrial Price of Energy US*. https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_6_a.
- [2] B. Aksanli, J. Venkatesh, L. Zhang, and T. Rosing. 2011. Utilizing Green Energy Prediction to Schedule Mixed Batch and Service Jobs in Data Centers. In *Proceedings of the 4th Workshop on Power-Aware Computing and Systems (HotPower '11)*. ACM, New York, NY, USA, Article 5, 5 pages. <https://doi.org/10.1145/2039252.2039257>
- [3] Apple Newsroom 2018. *Apple Now Globally Powered by 100 Percent Renewable Energy*. <https://apple.co/2lGaMX4>.
- [4] M. Arlitt, C. Bash, S. Blagodurov, Y. Chen, T. Christian, D. Gmach, C. Hyser, N. Kumari, Z. Liu, M. Marwah, A. McReynolds, C. Patel, A. Shah, Z. Wang, and R. Zhou. 2012. Towards the Design and Operation of Net-Zero Energy Data Centers. In *13th InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. 552–561. <https://doi.org/10.1109/ITHERM.2012.6231479>
- [5] L. Barroso and U. Holzle. 2007. The Case for Energy-Proportional Computing. *IEEE Computer* (2007).
- [6] J. L. Berral, Í Goiri, T. D. Nguyen, R. Gavalda, J. Torres, and R. Bianchini. 2014. Building Green Cloud Services at Low Cost. In *IEEE 34th International Conference on Distributed Computing Systems, ICDCS 2014, Madrid, Spain, June 30 - July 3, 2014*. 449–460.
- [7] Bloomberg NEF 2018. *BNEF New Energy Outlook 2018*. <https://bit.ly/2Or3NDE>.
- [8] M. Brown and J. Renau. 2011. ReRack: Power Simulation for Data Centers with Renewable Energy Generation. *SIGMETRICS Perform. Eval. Rev.* 39, 3 (Dec. 2011), 77–81. <https://doi.org/10.1145/2160803.2160865>
- [9] R. Buyya, M. Pathan, and A. Vakali. 2008. *Content Delivery Networks* (1st ed.). Springer Publishing Company, Incorporated.
- [10] T. Chen, X. Wang, and G. B. Giannakis. 2016. Cooling-Aware Energy and Workload Management in Data Centers via Stochastic Optimization. *IEEE Journal of Selected Topics in Signal Processing* 10, 2 (March 2016), 402–415. <https://doi.org/10.1109/JSTSP.2015.2500189>
- [11] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam. 2005. Managing Server Energy and Operational Costs in Hosting Centers. In *Proceedings of the 2005 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '05)*. ACM, New York, NY, USA, 303–314. <https://doi.org/10.1145/1064212.1064253>
- [12] A. P. Dobos. 2014. *PVWatts Version 5 Manual*. <http://www.nrel.gov/docs/fy14osti/62641.pdf>.
- [13] Energy Sage 2018. *How Long Will My Solar Panels Last?* <https://bit.ly/2Khr3k>.
- [14] Energy Sage 2019. *Evaluating Solar Panel Efficiency*. <https://www.energysage.com/solar/buyers-guide/solar-panel-efficiency/>.
- [15] L. Schwartz et al. 2017. *Electricity End Uses, Energy Efficiency, and Distributed Energy Resources Baseline*. <https://bit.ly/2DlGqci>.
- [16] Woodhouse et al. 2016. *The Role of Advancements in Solar Photovoltaic Efficiency, Reliability, and Costs*. <https://www.nrel.gov/docs/fy16osti/65872.pdf>.
- [17] Eurostat 2018. 'Non-household Price of Energy Germany'. https://ec.europa.eu/eurostat/statistics-explained/index.php/Electricity_price_statistics.
- [18] Facebook Newsroom 2018. *On Our Way to Lower Emissions and 100% Renewable Energy*. <https://bit.ly/2Pdlzew>.
- [19] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav. 2012. It's Not Easy Being Green. *SIGCOMM Comput. Commun. Rev.* 42, 4 (Aug. 2012), 211–222. <https://doi.org/10.1145/2377677.2377719>
- [20] Í Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini. 2013. Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy. *SIGARCH Comput. Archit. News* 41, 1 (March 2013), 51–64. <https://doi.org/10.1145/2490301.2451123>
- [21] Í Goiri, K. Le, Md. E. Haque, R. Beauchea, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini. 2011. GreenSlot: Scheduling Energy Consumption in Green Datacenters. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC '11)*. ACM, New York, NY, USA, Article 20, 11 pages. <https://doi.org/10.1145/2063384.2063411>
- [22] Í Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini. 2012. GreenHadoop: Leveraging Green Energy in Data-processing Frameworks. In *Proceedings of the 7th ACM European Conference on Computer Systems (EuroSys '12)*. ACM, New York, NY, USA, 57–70. <https://doi.org/10.1145/2168836.2168843>
- [23] Google 2018. *Google Environmental Report 2018*. <https://bit.ly/2C99nYV>.
- [24] V. Gupta, S. Lee, P. Shenoy, R. K. Sitaraman, and R. Urgaonkar. 2016. How to Cool Internet-scale Distributed Networks on the Cheap. In *Proceedings of the Seventh International Conference on Future Energy Systems (e-Energy '16)*. ACM, New York, NY, USA, Article 9, 12 pages. <https://doi.org/10.1145/2934328.2934337>
- [25] V. Gupta, P. Shenoy, and R. K. Sitaraman. 2018. Efficient Solar Provisioning for Net-Zero Internet-Scale Distributed Networks. In *2018 10th International Conference on Communication Systems Networks (COMSNETS)*. 372–379. <https://doi.org/10.1109/COMSNETS.2018.8328221>
- [26] T. Harvey, M. Patterson, and J. Bean. 2012. Updated Air-side Free Cooling Maps: The Impact of ASHRAE 2011 Allowable Ranges. *The Green Grid* (2012). <https://bit.ly/2TM7Wf5>.
- [27] R. Khalid, A. P. Wemhoff, and Y. Joshi. 2017. Energy and Exergy Analysis of Modular Data Centers. *IEEE Transactions on Components, Packaging and Manufacturing Technology PP*, 99 (2017), 1–13. <https://doi.org/10.1109/TCPMT.2017.2705055>
- [28] J. Koomey. 2010. Growth in Data Center Electricity Use 2005 to 2010. (2010). <http://bit.ly/30x8FhN>.
- [29] C. Li, A. Qouneh, and T. Li. 2012. iSwitch: Coordinating and optimizing renewable energy powered server clusters. In *2012 39th Annual International Symposium on Computer Architecture (ISCA)*. 512–523. <https://doi.org/10.1109/ISCA.2012.6237044>
- [30] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. 2011. Dynamic Right-Sizing for Power-Proportional Data Centers. In *2011 Proceedings IEEE INFOCOM*. 1098–1106. <https://doi.org/10.1109/INFCOM.2011.5934885>
- [31] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. Andrew. 2011. Geographical Load Balancing with Renewables. *SIGMETRICS Performance Evaluation Review* 39, 3 (2011), 62–66. <http://dblp.uni-trier.de/db/journals/sigmetrics/sigmetrics39.html#LiuLWLA11>
- [32] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. Andrew. 2011. Greening Geographical Load Balancing. In *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '11)*. ACM, New York, NY, USA, 233–244. <https://doi.org/10.1145/1993744.1993767>
- [33] V. Mathew, R. K. Sitaraman, and P. Shenoy. 2013. Energy-Efficient Content Delivery Networks Using Cluster Shutdown. In *International Green Computing Conference, IGCC 2013, Arlington, VA, USA, June 27-29, 2013, Proceedings*. 1–10. <https://doi.org/10.1109/IGCC.2013.6604510>
- [34] R. Miller. 2011. *The Average PUE is 1.8: Uptime Institute*. <http://bit.ly/1eEqN9d>.
- [35] NOAA Sunniest Cities 2004. *NOAA Ranking of Cities Based on Percentage Possible Sunshine*. <https://bit.ly/2S9aal>.
- [36] NREL 2018. *U.S. Solar Photovoltaic System Cost Benchmark: Q1 2018*. <https://www.nrel.gov/docs/fy19osti/72399.pdf>.
- [37] NREL PVWatts V5 API 2014. *PVWatts V5*. <https://developer.nrel.gov/docs/solar/pvwatts-v5/>.
- [38] E. Nygren, R. K. Sitaraman, and J. Sun. 2010. The Akamai Network: A Platform for High-Performance Internet Applications. *ACM SIGOPS Operating Systems Review* 44, 3 (2010), 2–19.
- [39] V. Pallipadi and A. Starikovskiy. 2006. The Ondemand Governor: Past, Present and Future. In *Proceedings of Linux Symposium, vol. 2, pp. 223-238*. <https://doi.org/pub/linux/kernel/people/lenb/acpi/doc/OLS2006-ondemand-paper.pdf>
- [40] S. Pelley, D. Meisner, T. Wenisch, and J. VanGilder. 2009. Understanding and Abstracting Total Data Center Power. In *Workshop on Energy-Efficient Design*.
- [41] A. Qureshi. 2010. Power-Demand Routing in Massive Geo-Distributed Systems. (2010).
- [42] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramaniam. 2012. Carbon-Aware Energy Capacity Planning for Datacenters. In *Proceedings of the 2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '12)*. IEEE Computer Society, Washington, DC, USA, 391–400. <https://doi.org/10.1109/MASCOTS.2012.51>
- [43] N. Sharma, S. Barker, D. Irwin, and P. Shenoy. 2011. Blink: Managing Server Clusters on Intermittent Power. *SIGARCH Comput. Archit. News* 39, 1 (March 2011), 185–198. <https://doi.org/10.1145/1961295.1950389>
- [44] K. Smith, A. Saxon, M. Keyser, B. Lundstrom, Ziwei Cao, and A. Roc. 2017. Life prediction model for grid-connected Li-ion battery energy storage system. In *2017 American Control Conference (ACC)*. 4062–4068. <https://doi.org/10.23919/ACC.2017.7963578>
- [45] Solar Energies Technology Office, U.S. Department of Energy 2017. *The SunShot 2030 Goals*. <https://bit.ly/2AUqexf>.
- [46] C. Stewart and K. Shen. 2009. Some Joules are More Precious than Others: Managing Renewable Energy in the Datacenter. In *In HotPower*.
- [47] Time and Date 2019. *Climate and Weather Averages in Anchorage*. <https://bit.ly/2HlmswT>.
- [48] H. Zeng, C. S. Ellis, A. R. Lebeck, and A. Vahdat. 2002. ECOSystem: Managing Energy As a First Class Operating System Resource. *SIGPLAN Not.* 37, 10 (Oct. 2002), 123–132. <https://doi.org/10.1145/605432.605411>
- [49] Y. Zhang, Y. Wang, and X. Wang. 2011. GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy. In *Middleware (Lecture Notes in Computer Science)*, Fabio Kon and Anne-Marie Kermarrec (Eds.), Vol. 7049. Springer, 143–164. <http://dblp.uni-trier.de/db/conf/middleware/middleware2011.html#ZhangWW11>

Appendix A GREEDY ALGORITHM PSEUDOCODE

Algorithm 1 Greedy Algorithm Pseudocode

```

1: function GREENHEURISTIC( )
2:    $dcs \leftarrow [1, 2, \dots, m]$  ▷ datacenters
3:    $sortedpeers \leftarrow [p_1, p_2, \dots, p_m]$  ▷ sorted list of dc peer dcs
   in increasing order of dist
4:    $time \leftarrow [1, 2, \dots, n]$  ▷ time periods
5:    $r = \text{max radius of load movement}$ 
6:    $b \leftarrow [b_1, b_2, \dots, b_m]$  ▷ battery charge
7:   for  $i$  in  $time$  do
8:      $sload \leftarrow [l_{11}, l_{12}, \dots, l_{mn}]$  ▷ server load for time period
9:      $cload \leftarrow [c_{11}, c_{12}, \dots, c_{mn}]$  ▷ cooling load for time
   period
10:     $oac \leftarrow [o_{11}, o_{12}, \dots, o_{mn}]$  ▷ oac available y/n?
11:     $solarenergy \leftarrow [s_{11}, s_{12}, \dots, s_{mn}]$  ▷ local solar energy
12:     $surpluslist \leftarrow []$  ▷ to store dcs with surplus solar
   energy
13:     $deficitlist \leftarrow []$  ▷ to store dcs using brown energy
14:    for  $j$  in  $dcs$  do
15:      if  $o_{ij} = y$  then
16:         $c_{ij} \leftarrow 0$  ▷ if there is OAC, cooling load is zero
17:         $excessSolar_{ij} \leftarrow s_{ij} + b_j - (l_{ij} + c_{ij})$  ▷ determine
   excess solar
18:        if  $l_{ij} + c_{ij} > s_{ij}$  then
19:           $b_j \leftarrow b_j - (l_{ij} + c_{ij} - s_{ij})$  ▷ use battery if solar
   energy falls short
20:        if  $excessSolar_{ij} > 0$  then
21:           $surpluslist \leftarrow surpluslist \cup [j]$  ▷ add dc to
   surplus list
22:        else if  $excessSolar_{ij} < 0$  then
23:           $deficitlist \leftarrow deficitlist \cup [j]$  ▷ add dc to
   deficit list
24:        for  $j \in deficitlist$  do ▷ first iteration
25:          for  $p \in sortedpeers$  do
26:            if  $p \in surpluslist \wedge o_{ip} = y \wedge dist(j, p) \leq r$  then
27:              move load to  $p$  and adjust variable values
28:          for  $j \in deficitlist$  do ▷ second iteration
29:            for  $p \in sortedpeers$  do
30:              if  $p \in surpluslist \wedge o_{ip} = n \wedge dist(j, p) \leq r$  then
31:                move load to  $p$  and adjust variable values

```
