

Efficient Solar Provisioning for Net-Zero Internet-Scale Distributed Networks

Vani Gupta[†], Prashant Shenoy[†], Ramesh K. Sitaraman^{†‡}

[†]University of Massachusetts, Amherst [‡]Akamai Technologies

{vani, shenoy, ramesh}@cs.umass.edu

Abstract—Internet-scale Distributed Networks (IDNs) are global distributed systems consisting of a network of hundreds of thousands of servers located in hundreds of data centers around the world. A Content Delivery Network (CDN) is an example of an IDN that delivers content globally through a large network of servers. IDNs consume large amounts of energy and their energy requirements are projected to increase significantly in the future. With carbon emissions from data centers increasing every year, energy solutions for data centers powered by renewables are critical for the sustainability of data centers and for the environment. In this paper, we study the benefits of using solar energy to green IDNs. We study the impact of leveraging global data center locations with high solar potential on the number of solar panels needed to power a CDN. We develop optimal algorithms and heuristics to help minimize the number of solar panels provisioned across the CDN, while making it net-zero. We empirically evaluate our algorithms using extensive load traces from Akamai’s global CDN and solar data from PVWatts. Overall, with unrestricted load movement, we can reduce the number of solar panels by 36%, 68%, and 82% for net-zero year, month and week respectively. Our results show that our algorithms can significantly reduce the number of solar panels we need to power our CDN, thereby making sustainability of Internet-scale distributed networks more achievable.

I. INTRODUCTION

Modern Internet services are delivered using Internet-scale Distributed Networks (IDNs) that rely on a global network of hundreds of thousands of servers located all over the world. IDNs include cloud and Internet services such as content delivery networks (CDNs) that deliver web content, applications, and streaming media to clients via hundreds of thousands of servers located in thousands of data center locations throughout the world. Such large Internet-scale systems consume large amounts of energy and consequently incur large energy bills. A research study in 2010 [25] estimated that Google’s energy consumption was more than 1,120 GWh and their annual electricity bill exceeded \$67M. The environmental impact of these massive IDNs is also significant and is increasing every year. Carbon emissions from data centers are growing at the rate of 15% each year [13].

Given the large energy requirements of IDNs, data centers that are powered using renewables are gaining traction in the industry and in the research community. In just six years, Apple’s renewable energy usage to power its corporate facilities, stores, and data centers worldwide increased from 16% in 2010, to 96% in 2016 [1]. Apple has committed to powering all its facilities world-wide with 100% renewable energy. This year, Google announced that it will achieve the milestone of

purchasing 100% renewable energy to match consumption for its global operations, including its data centers and offices [2].

There has been a lot of research on making data centers greener by reducing energy consumption or using energy generated from renewable sources. Prior work includes energy reduction using server shutdown or low-power states during off-peak times [16] [19] [5] [26]. There is also work on job scheduling based on predicted solar [11] [10] and load balancing to encourage use of renewable energy [18] [17] [9]. Separately there has been a study [4] for selecting sites for and provisioning green data centers using a follow-the-renewables approach. Greening IDNs is now a necessity for sustainable growth of IDNs, for reducing environmental impact, and for lowering energy costs for companies. When a system produces enough green energy to off-set its brown energy use over a time period, it is said to be ‘net-zero’ over that time period. While prior work has addressed greening individual data centers, the problem of greening a large distributed network of data centers, such as an IDN, using renewables, has not received much research attention. To address this issue, we consider the following research question: how can we efficiently provision solar arrays across a global IDN with hundreds of locations to make it net-zero? We address this important question in our study by designing and implementing optimal algorithms and heuristics for provisioning solar panels, and evaluating those algorithms on a real world IDN trace and a year’s worth of PVWatts solar data.

There is a significant difference in solar output between different locations on the globe, due to differences in latitude, longitude, and weather. IDNs are deployed as a global network of data centers and the services they provide are replicated across those data centers. Proximity of servers to users is the main reason for global deployment. Therefore, data centers cannot be deployed only in locations where there is high solar output throughout the year, they also need to be close to users. However, replication of services in an IDN allows load movement between locations to leverage high solar outputs, and we use this feature extensively to optimize solar panel provisioning and energy usage

Contributions: Our contributions are listed below:

- *Determining Solar Potential for Global IDNs:* We conduct a comprehensive study to analyze the net-zero solar potential for existing global IDNs with data centers located in hundreds of locations throughout the world. In order to reduce the number of panels provisioned, we

leverage global locations that have high solar output. To determine the number of panels needed to be net-zero, we move load in an off-line fashion and ensure that the data center energy demand is matched by solar energy supply for the duration of the net-zero time period.

- *Algorithm Design:* We design optimal and heuristic algorithms to minimize the number of panels we need to serve the IDN's load by taking advantage of higher levels of solar across various regions on the globe. Firstly, we design our algorithms such that they can be generalized to different net-zero time periods, including net-zero week, month, and year. Secondly, we also design the LP to study the impact of restricting load movement within a certain radius when determining the number of panels to be provisioned. Finally, our LP can also be easily configured to restrict locations where solar panels are installed.
- *Extensive Trace-based Evaluation:* We evaluate our algorithms on an extensive load trace from one of the world's largest CDN. The month-long trace consists of five-minute information from 100,592 servers in 724 global data center locations from Akamai's CDN. We see significant reduction in the number of panels and also find that our heuristic algorithms perform well compared to the optimal. Overall, with unrestricted radius of load movement, we can reduce the number of solar panels by 36% for net-zero year. For net-zero month by about 67% to 68%, and for net-zero week by about 71% and 74% for heuristics, and by 82% for the optimal.

In the rest of the paper, we present background and define the problem in Section II and III respectively. We present our optimal and heuristic algorithms in Section IV and experimental methodology in Section V. In Section VI we describe our empirical results. We discuss related work in Section VII and finally present our conclusions in Section VIII.

II. BACKGROUND

Internet-scale Distributed Networks: Internet-scale distributed networks (IDNs) are large global networks that are comprised of data centers in hundreds of locations across the world. Content delivery networks (CDNs) are examples of IDNs and are used to deliver content, streaming audio, video, applications etc. on the web. Figure 1 shows some data center locations part of the Akamai CDN. Commercial CDNs use two levels of load-balancing in their systems: *local* and *global*. When a user requests content, the *global* load-balancer assigns the request to a server cluster located 'close-by' to minimize loss and latency [20]. The *local* load-balancer then maps the request to a specific server in the cluster. In order to assign users to nearby data centers and minimize loss and latency, CDNs replicate their services to have redundancy in the choice of data centers. This replication is also useful if load from one data center is assigned to another data center for other reasons, e.g., to leverage a local feature like high solar output.

Energy Consumption Model for an IDN: The primary source of energy consumption in IDNs are the numerous servers deployed in all the various data centers (*server energy*).




Fig. 1: Plot showing the diversity in geographic locations that can make up a global IDN

The energy consumed by a server is largely dependent on the amount of load it is serving, so we can model the energy consumed by a server as a function of its load (we use normalized load λ , $0 \leq \lambda \leq 1$, which is the actual load as a fraction of its capacity). However, although energy consumed is largely dependent on the load, servers are not energy proportional and still consume some energy, roughly 60%, when they are idle. For our work, we use the standard linear model of server power consumption [3] that defines power consumed by a server as $P_{idle} + (P_{peak} - P_{idle})\lambda$, where λ is the normalized load on the server, P_{idle} is the power consumed by server that is idle, and P_{peak} is the power consumed by the server that has peak load. We assume that we can move load between servers to consolidate load, and shut down idle servers, so as to use the minimum number of servers needed to serve the load [15]. Our assumption is that such consolidation of load is done at each of the data center locations for each time period. Using the linear model and consolidating load between servers, we then determine the power consumed by the data center. The power consumed (in watts) by the data center in each 5-minute time interval is then multiplied by the number of seconds ($5 \cdot 60$) to get the energy consumed by the data center in each time interval (in joules).

In addition to *server energy*, we also need energy to cool them (*cooling energy*). Heat dissipated by servers is a function of the energy they consume. The more heat they dissipate, the more energy is needed to cool them. So cooling energy is proportional to server energy. A recent study of data center energy consumption [21] showed that servers and cooling consumed 56% and 30% of the total energy respectively. Thus, most of the energy consumed by a data center is spent to power and cool servers. We assume a PUE of about 1.8 [28] and scale up server energy consumption to account for cooling energy.

Net-zero Systems: 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 [8]. That is, on-site renewable energy produced is at least the energy consumed over the period. Extending this basic definition, we define a 'net-zero' IDN for different time periods as below:

- A *Net-zero Year IDN* is net-zero over a year.
- A *Net-zero Month IDN* is net-zero for every month in a year.
- A *Net-zero Week IDN* is net-zero for every week in a year.

By definition, net-zero week is the most stringent requirement, followed by net-zero month, and finally net-zero year. So a net-zero week IDN will also be net-zero month and net-zero

year. A net-zero month IDN will also be net-zero year.

Solar Panels and Factors Affecting Solar Output: A solar panel is an electrical device that converts sunlight into electricity using the photovoltaic effect [22]. Several factors affect solar output, and we list some of them below:

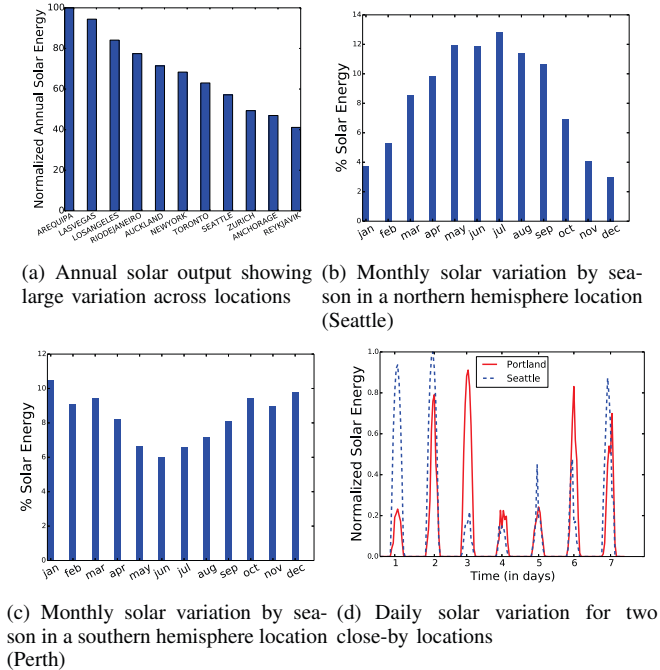


Fig. 2: Figure showing rich solar variations across the globe

- **Location:** As shown in Figure 2(a), there is a large variation in annual solar output based on location. Locations like Las Vegas have higher annual solar output, while locations like Anchorage have lower solar output.
- **Season:** Solar output also varies by season. Summer months tend to have higher levels of solar output than winter months. This can be seen in Figure 2(b) which shows the monthly solar output for Seattle, WA.
- **Hemisphere:** Figures 2(b) and 2(c) show monthly solar output for Seattle (northern hemisphere) and Perth (southern hemisphere). As we see, the trend in the levels of solar output is reversed for these locations, given the northern and southern hemispheres experience opposite seasons.
- **Daily variations:** Solar output goes to zero when the sun is not shining and so the hour of the day affects solar output. As Figure 2(d) shows, for each location there are hours when solar output is zero and then it rises steadily, peaks and then falls again once the sun sets.
- **Other factors:** All other factors like location, season, and time of day remaining constant, solar output can still change based on several factors that may include weather, cloud cover, pollution etc. As Figure 2(d) shows, within the same location, season, hemisphere and time of day, we see large variations between solar output from one day of the week to the next. Locations that are close-by e.g. Portland, OR and Seattle, WA also show large variations

in solar output, as we can see in the same figure.

Given the above analysis, we conclude that solar output is highly variable across time and space and is affected by several diverse factors, and their interplay. While this intermittency and variability is a challenge, it is also an opportunity in the context of a global IDN. Given an IDN has replicated services, we can leverage high levels of solar by moving load to locations that have high solar output. In our study, we leverage these complex variations to reduce the number of panels provisioned.

III. PROBLEM STATEMENT

Data centers consume energy to maintain, run, and cool servers and other equipment. For a net-zero data center, energy supply needs to be matched by the demand by using energy generated from renewables, like solar. There is a large variation in solar output across global locations, with certain locations being excellent for solar generation. In this paper, we address the problem of efficient solar panel provisioning for global net-zero IDNs. To provision panels efficiently, we move load to locations with high solar output. While defining the problem, we make three simplifying assumptions. First, we assume that it is possible to install as many solar panels as we need in any location. Second, we assume that it is possible to deploy as many servers as we need at any location. Finally, our approach assumes that the total energy produced over the net-zero period is greater than or equal to the total energy demand. That is, we assume that excess energy production can be net-metered to the grid, or stored locally, and consumed at a later time.

Our provisioning problem can be broadly stated as follows: *Using load movement, how can we efficiently provision solar panels for a global IDN so it is net-zero over a given time period such as a year, a month, or a week.* Importantly, being net-zero over a certain period implies total energy used in that period across the IDN equal total solar output for the IDN. It does not require instant power usage to be fully met by solar production at that instant.

Specifically, we study the following research questions:

- **How to achieve net-zero IDNs with and without performance constraints:** Firstly, in order to determine our full solar potential, we analyze how much reduction we can see in the number of panels if we allow load to be assigned to any location on the globe, without worrying about performance. This scenario yields best case results for reduction in the number of panels, and becomes a point of comparison for results under more constrained scenarios. Secondly, in order to reduce latency, it is important to serve load from locations close to users. Therefore, we place constraints on the radius within which we must operate while moving load to a data center with higher solar output. Although we do not explicitly add a relocation penalty, we do address latency by constraining load movement within certain radii.
- **How do results change if panels are assigned to top k locations only:** Solar panel installation may only be possible in large data centers in major metro areas. We

| dcid/time | 1 | 2... | n |
|------------|----------|-------------|----------|
| 1 | l_{11} | $l_{12}...$ | l_{1n} |
| 2 | l_{21} | $l_{22}...$ | l_{2n} |
| .. | .. | .. | .. |
| m | l_{m1} | $l_{m2}...$ | l_{mn} |
| total load | l_1 | $l_2...$ | l_n |

TABLE I: Demand Matrix

| dcid/time | 1 | 2... | n |
|-----------|----------|-------------|----------|
| 1 | s_{11} | $s_{12}...$ | s_{1n} |
| 2 | s_{21} | $s_{22}...$ | s_{2n} |
| .. | .. | .. | .. |
| m | s_{m1} | $s_{m2}...$ | s_{mn} |

TABLE II: Supply Matrix

analyze the number of solar panels needed to make IDNs net-zero if the panels could only be deployed in large data centers in major cities. For an IDN, data centers sizes vary by location and population. Generally speaking, areas with larger population tend to have larger data centers so servers can be proximal to users in order to reduce latency. We use the number of servers as an indicator for the size of the data center. In this scenario, we sort our data centers by the number of servers they have and consider only the top k locations as candidates for installing panels, ensuring panels are installed at bigger data centers. We study how the number of panels provisioned change as we vary k.

IV. ALGORITHMS FOR SOLAR PROVISIONING IN IDNS

We begin with an LP formulation to solve the solar provisioning problem under performance constraints. While the LP formulation is optimal, it is also computationally intensive. Therefore, to reduce run-time complexity in the no performance constraint scenario, we also define two heuristic algorithms that run faster while yielding comparable results.

Before defining the algorithm, we first define the main inputs to the algorithm.

- *Demand and Supply Matrices:* We set up the inputs as two matrices: one a demand matrix (shown in Table I) and the other a supply matrix (shown in Table II). The values in the demand matrix ' l_{ij} ' represent the energy used by the data center at the corresponding time. The values in the supply matrix ' s_{ij} ' represent the solar energy available at that location per panel. Both matrices have 'm' rows corresponding to data centers and 'n' columns corresponding to time periods.
- *Neighbors of a data center 'i':* We define $N_{i\delta}$ to be the set of all data centers within a radius of δ kms from i .

A. Optimal LP Formulation with Performance Constraints

We formulate an LP, which we refer to as LP_{perf} , to determine the minimum number of panels we can provision to meet demand, given the solar energy available at various locations. In addition to the inputs defined above, for each

data center i , we define variables l_{ijt} to be the load moved from datacenter i to datacenter j , at time t , $\forall j \in N_{i\delta}$. We also define p_i as the number of panels installed at data center i . Given this setup, we define the LP as shown below. We minimize the total number of solar panels provisioned in the objective function as below:

$$\text{Min: } \sum_{i=1}^m p_i \quad (1)$$

We then define constraints as below: Incoming load should be less than or equal to the solar supply:

$$\text{s.t.: } \sum_{i \in N_{j\delta}} l_{ijt} \leq s_{jt} * p_j, \quad \forall j, t \quad (2)$$

Total outgoing load, including load moved from the data center to itself should be equal to the starting load:

$$\sum_{j \in N_{i\delta}} l_{ijt} = l_{it}, \quad \forall i, t \quad (3)$$

In addition to these constraints, we also have non-negative constraints for each of the variables defined:

$$l_{ijt} \geq 0, \quad \forall i \in N_{j\delta}, \forall j, t \quad (4)$$

$$l_{it} \geq 0, \quad \forall i, t \quad (5)$$

$$s_{jt} \geq 0, \quad \forall j, t \quad (6)$$

$$p_j \geq 0, \quad \forall j \quad (7)$$

$$\delta \geq 0 \quad (8)$$

Without performance constraints, i.e. for an unlimited radius of load movement, where load can potentially be assigned from a data center to any other data center, the number of l_{ijt} variables and associated constraints, becomes very large. For net-zero month, for m data centers, the number of l_{ijt} type variables is $12m^2$. Therefore, we devise an alternative LP formulation for the no performance scenario.

B. Optimal LP Formulation With No Performance Constraints

In this section, we do not consider performance constraints in our formulation. In addition to the inputs and variables defined earlier, we define in_{it} as the load moved into datacenter i at time t . Similarly out_{it} is the load moved out of datacenter i at time t . Given this setup, we define the LP, referred to as LP_{nperf} , as shown below. We minimize the total number of solar panels installed:

$$\text{Min: } \sum_{i=1}^m p_i \quad (9)$$

Subject to four types of constraints: Total incoming load minus the outgoing load should be zero:

$$\text{s.t.: } \sum_{i=1}^m in_{it} - \sum_{i=1}^m out_{it} = 0, \quad \forall i, t \quad (10)$$

Incoming load should be less than or equal to the solar supply:

$$l_{it} + in_{it} - out_{it} \leq s_{it} * p_i, \quad \forall i, t \quad (11)$$

Total outgoing load should be less than or equal to the sum of the starting load and any incoming load:

$$l_{it} + in_{it} - out_{it} \geq 0, \quad \forall i, t \quad (12)$$

In addition to these constraints, we also have non-negative constraints for each of the variables defined:

$$l_{it} \geq 0, \quad \forall i, t \quad (13)$$

$$s_{it} \geq 0, \quad \forall i, t \quad (14)$$

$$p_i \geq 0, \quad \forall i \quad (15)$$

$$in_{it} \geq 0, \quad \forall i, t \quad (16)$$

$$out_{it} \geq 0, \quad \forall i, t \quad (17)$$

The LP_{nperf} does away with the l_{ijt} type variables that represent load moved from data center i to data center j at time t . Instead it only models load moved in and out of each data center at time t using variables in_{it} and out_{it} . So for the net-zero month case for m data centers for LP_{nperf} , where we had to contend with $12m^2$ variables, we now only have to work with $12m$ variables for the LP_{nperf} formulation.

Theorem 1. LP_{nperf} for unlimited δ and LP_{nperf} are equivalent.

Proof: We prove equivalence by showing that a solution for the first formulation with unlimited δ is also a solution for the second and vice versa. We omit the detailed proof due to space constraints, and refer the reader to the extended version of this paper [12]. ■

C. Greedy Heuristic Algorithms

Optimal LP formulations may be computationally very expensive for large system sizes, so we also develop efficient heuristics. We describe below our heuristic algorithms that are inspired at a high level by greedy approximation algorithms to the set-covering problem [6]. We loosely consider the load to be served as the set to be covered. The different amounts of load we can serve using solar energy generated at various locations are like the subsets that can cover the original set.

1) *Max Solar Per Panel Heuristic (MSP):* We now define a greedy heuristic algorithm that performs comparably and runs faster than the LP. For example, for the no performance net-zero week case, the heuristics took a few seconds to complete, while the LP_{nperf} took over 2.25 hours, LP_{perf} for $r=200$ kms ran for over 10.5 hours, and LP_{perf} for $r=700$ kms did not complete even in 28 days. We use the same demand and supply matrices that we defined earlier in Section IV. In order to minimize the number of panels, we note that we need to assign as much load as we can to a location that has the highest solar output. This would help us to cover the maximum load with the minimum number of panels for a given time slot. We greedily pick the maximum solar per panel location across time and space (i.e. across all time periods and all data centers), and assign the entire load for the time period to that location. Using the solar per panel value and the load, we determine the number of panels to place at that location. Once we install panels at a location, these panels can then be used to serve

demand for other time periods as well, so we accordingly adjust the demand values to reflect the extra supply for all other time periods. We continue to place panels in this way until we satisfy the entire demand in all time periods. For pseudocode, please see Algorithm 1.

Algorithm 1 MSP Heuristic Pseudocode

```

1: function SPHEURISTIC()
2:    $time \leftarrow [t_1, t_2, \dots, t_n]$  ▷ time periods
3:    $spp \leftarrow [s_{11}, s_{12}, \dots, s_{mn}]$  ▷ solar output
4:    $load \leftarrow [l_1, l_2, \dots, l_n]$  ▷ load for time period
5:    $provpanels \leftarrow []$  ▷ provisioned panels
6:   for  $t_y$  in  $time$  do
7:      $s_{xy} \leftarrow \min s_{ij}$  s.t.  $s_{ij} \in spp$  ▷ pick min solar
8:      $p_{xy} \leftarrow l_y / s_{xy}$  ▷ assign panels
9:      $provpanels \leftarrow provpanels \cup [p_{xy}]$  ▷ add to provisioned
panels
10:    for  $i$  in  $[1, 2, \dots, m]$  do
11:       $l_i \leftarrow |l_i - l_y|$  ▷ adjust other loads
12:       $load \leftarrow load - l_y$  ▷ delete assigned load
return  $provpanels$ 

```

2) *Min Number of Panels Heuristic (MNP):* We now describe our second heuristic algorithm. The basic structure of this algorithm is the same as the MSP algorithm, except we now use a different heuristic to make a decision on where to place panels. We first determine the number of panels for each location for each time period, by dividing the load for the time period for all locations, by the solar per panel for the corresponding time and location. This gives us the ‘Number of Panels Matrix’ shown in Table III. We then pick the lowest number of panels value and install those many panels at the corresponding location and time period. Like before, once any panels are installed at a location, they are also available for other time slots. So we accordingly adjust the demand to reflect the extra supply. We recompute the number of panels matrix for the adjusted loads, and start over. We do this exercise until all the demand is met. Pseudocode for this algorithm is detailed in Algorithm 2.

Algorithm 2 MNP Heuristic Pseudocode

```

function NPHEURISTIC()
   $time \leftarrow [t_1, t_2, \dots, t_n]$  ▷ time periods
   $spp \leftarrow [s_{11}, s_{12}, \dots, s_{mn}]$  ▷ solar output
   $load \leftarrow [l_1, l_2, \dots, l_n]$  ▷ load for time period
   $origpanels \leftarrow []$  ▷ original panels
   $provpanels \leftarrow []$  ▷ provisioned panels
  for  $s_{ij}$  in  $spp$  do
     $op_{ij} \leftarrow l_i / s_{ij}$  ▷ determine original num panels
     $origpanels \leftarrow origpanels \cup [op_{ij}]$  ▷ add to original panels
  for  $t_y$  in  $time$  do
     $op_{xy} \leftarrow \min op_{ij}$  s.t.  $op_{ij} \in origpanels$  ▷ pick min panels
     $p_{xy} \leftarrow l_y / s_{xy}$  ▷ assign panels
     $provpanels \leftarrow provpanels \cup [p_{xy}]$  ▷ add to provisioned panels
    for  $i$  in  $[1, 2, \dots, m]$  do
       $l_i \leftarrow |l_i - l_y|$  ▷ adjust other loads
       $load \leftarrow load - l_y$  ▷ delete assigned load
       $time \leftarrow time - t_y$  ▷ delete time column
       $origpanels \leftarrow []$  ▷ reset original panels
    for  $s_{ij}$  in  $spp$  do
       $op_{ij} \leftarrow l_i / s_{ij}$  ▷ determine num panels
       $origpanels \leftarrow origpanels \cup [op_{ij}]$  ▷ add to panels
return  $provpanels$ 

```

| dcid/time | 1 | 2... | n |
|-----------|----------|-------------|----------|
| 1 | p_{11} | $p_{12}...$ | p_{1n} |
| 2 | p_{21} | $p_{22}...$ | p_{2n} |
| .. | .. | .. | .. |
| m | p_{m1} | $p_{m2}...$ | m_n |

TABLE III: Number of Panels Matrix

| 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 IV: Parameters for PVWatts Data

V. EXPERIMENTAL METHODOLOGY

We conduct experiments on an extensive Akamai load trace spanning a month. The trace consists of load information from 100,592 servers in 724 global data center locations as shown in Figure 1. The dataset includes load, requests, and bytes served by each server every five minutes over a month-long trace. Further, it has detailed information about every data center, including the number of deployed servers, total server capacity, and the location of the data center including its latitude, longitude, city, state, and country.

For solar energy data, we use the PVWatts [23] hourly data of AC energy generation from solar radiation for a year. Assuming the power rating for solar panels is between 200 watts and 350 watts [7], we take an average value of 275 watts as the power rating per panel. Therefore, we use a system capacity of 0.275 kW for PVWatts in order to get the output generated by a single panel. For simplification, for all other required solar parameters, we use the values listed under 'Default Values' on page 3 in the PVWatts version 5 manual [24]. The required parameters used for downloading PVWatts data are detailed in Table IV.

The trace has 5-minute readings, whereas solar data is hourly. So we make an assumption that solar data does not change much during the hour and use the hour's reading for each of the five minutes that fall within that hour. Also, we have solar data for a year, but load trace data for a month. We assume that the load pattern for the CDN repeats monthly for the year. For each five minute interval, we convert load and solar output power (in watts) into energy units (joules).

For baseline comparison, we use the number of panels we need to serve the IDN load without any load movement. E.g. for net-zero week, for each week, we divide the week's load for the data center by the corresponding week's sum of solar per panel values. We then determine the number of original panels as the maximum of all the weekly number of panels.

VI. EMPIRICAL RESULTS

The following paragraphs describe our results for scenarios where we restrict load movement within certain fixed radii (*im-*

posing performance constraints), where we allow unrestricted load movement (*without any performance constraints*), and when panel provisioning is restricted to top k locations.

A. With and Without Performance Constraints


The goal of this experiment was to study the impact of load movement within a radius on the reduction in the number of panels when compared to the two extremes of unrestricted load movement and no load movement at all.

1) *Number of Panels*: Our results are shown in Figures 3(a) and 3(b). We list our observations below:


- *Heuristics perform comparably*: From Figure 3(a), we see that for an unlimited radius of load movement, the heuristic algorithms perform comparably with the LP.
- *Load movement helps reduce the number of required panels dramatically*: Figure 3(a) shows that with an unlimited radius, for net-zero year, number of panels decrease by 36% for all algorithms. For net-zero month, the MNP decreases panels by 66.94%, the MSP by 67.03%, and the LP by 68%. For net-zero week, the MNP decreases panels by 71.4%, the MSP by 73.7%, the LP by 82%. Figure 3(b) shows that even with performance constraints, we can see a significant reduction in the number of panels. For $r=500$ kms, we see a reduction of about 9.7%, 27%, and 53% for net-zero year, month, and week respectively.
- *Number of panels is inversely proportional to size of net-zero time window*: Figure 3(b) shows that the number of panels provisioned is highest for net-zero week, followed by net-zero month, followed by net-zero year. This is intuitive given we are averaging over a larger time period for net-zero year as compared to net-zero month. For net-zero year, we must match demand with supply over the entire year. For net-zero month we must match demand with supply for *each* month, however low our supply maybe and however high our demand may be for various months. Therefore, we must satisfy our net-zero condition for the 'worst' month in our list. Similarly for net-zero week. Therefore, the number of panels increase as we move from net-zero year, to net-zero month, to net-zero week. In Figures 3(b) and 3(c), missing bars for net-zero week correspond to input parameters for which LP_{perf} was computationally too intensive.

2) *Number of Locations*: We study the number of locations where solar panels are provisioned by radii for different net-zero time periods. Our results are show in Figures 3(c) and 3(d). Our main observations are:

- *Number of locations decrease with increase in radius*: Figure 3(c) shows that with an increase in the radius of load movement, the number of locations where solar panels are allocated decreases for each net-zero time period. This is because as the radius increases, load converges to locations that are globally high for solar output. Figure 4 shows locations that are chosen for various values of max radius for net-zero month. With load movement, we see the locations shrink and converge to the hubs for solar generation.



(a) For an unlimited radius, we see a dramatic reduction in the number of panels. Heuristics perform nearly as well as the optimal LP
 (b) Number of panels decrease as the radius of load movement increases. Number of panels is inversely proportional to the net-zero time period



(c) Number of locations where panels are allocated decreases with increase in radius, because load converges to global locations with high solar output
 (d) Distribution of panels for a radius of 500kms for different net-zero time periods. We observe that for all the net-zero time periods, panels are fairly evenly distributed with a few high peaks lying between about 12% and 16%

Fig. 3: Results with and without performance constraints

- *Number of locations is inversely proportional to the size of net-zero time period:* Figure 3(c) also shows that the number of locations are the maximum for net-zero week followed by net-zero month, and then net-zero year. Net-zero month and net-zero year are closer in the number of locations they pick. Figure 3(d) shows the distribution of panels for a radius of 500kms for different net-zero time periods. We observe that for all the net-zero time periods, panels are fairly evenly distributed with a few high peaks lying between about 12% and 16%.

With unconstrained load movement we observe that, for all the algorithms, load is moved to locations that are high in solar output. However, with these choices we find that load can end up in remote locations where solar panel installation may not be feasible. To address this, we restrict solar panel installation to locations having large data centers in major cities.

B. Restricting Panels to Top K Data Centers

We use the number of servers in a data center as the proxy for data centers that are large and are located in non-remote places with large populations. We restrict panel provisioning to top k locations sorted by the number of servers. Figure 5 shows the change in the number of panels across different values of k. We list our observations below:

- *Number of panels provisioned varies inversely with k:* First of all, we see that the number of panels provisioned increases when we restrict ourselves to fewer locations. This is intuitive considering that we are operating with more constraints, and therefore are not able to extract




Fig. 5: When restricting panel provisioning to top k data centers, the number of panels provisioned are highest for net-zero week, followed by net-zero month, and finally net-zero year

as much reduction from solar output as we could in an unconstrained setting. The larger the k, the more locations are in play for extracting solar savings.

- *Number of panels vary inversely with the size of the net-zero time window:* Once again, we see that the number of panels provisioned is the most for net-zero week, followed by net-zero month, and finally net-zero year. This trend is preserved across different values of k. See Section VI-A1 for a detailed explanation.
- *k=500 balances both panel optimization and top-k requirements:* For k=500, the number of panels are very close to the no performance panel provisioning scenario. Therefore, restricting panels to the top 500 data centers is a good middle ground for installing a near-optimal number of panels at non-remote locations.

VII. RELATED WORK

Recently, there has been a lot of research on reducing or greening energy consumption in data centers, including solutions focused on shutting down servers or clusters during off-peak periods and/or using low-power consumption states instead of powering them off in order to prevent wear and tear [16] [19] [5] [26]. Though these solutions provide significant savings in energy, they do not deal directly with the use of green energy to power data centers. Separately, a lot of work has been done in the area of renewables for data centers. Previous work has also looked into providing a solution for selecting sites for and provisioning green data centers using a follow-the-renewables approach [4]. However, their work focuses on setting up a data centers from scratch, where as our work explores how to use renewables to green an existing IDN. Work has also been done on job scheduling within a data center based on predicted solar and brown energy prices [11] [10]. Previous work has also modeled the potential of using renewable energy for data centers located in colder locations [27], while [14] proposes a solution for data center expansion using modular solar panels and distributed battery systems to have near-zero environmental impact. While this work takes advantage of renewables to reduce energy consumption, it does not deal with efficient provisioning of solar panels for an existing global IDN. Moving load across data centers to




Fig. 4: Locations where the LP with performance constraints places solar panels for net-zero month

increase the use of renewables has been studied before. Studies have also been done [18] [17] on how and to what extent geographical load balancing can encourage use of renewable energy and reduce the use of brown energy. Their distributed algorithm offers significant savings in cost (including energy cost and delay cost). Work has also been done on green solutions that control user traffic and direct it to different data center locations based on changes in workload and carbon footprint [9]. However, all these works [18] [17] [9] do not focus on efficient provisioning of solar panels.

VIII. CONCLUSIONS

We studied the problem of efficient solar panel provisioning for net-zero IDNs. Using our heuristic and optimal algorithms, we are able to significantly reduce the number of solar panels needed for creating net-zero IDNs. Overall, with unrestricted radius of load movement, for net-zero year, we can reduce the number of solar panels by 36%, between 67% to 68% for net-zero month, and for net-zero week about 71% and 74% using heuristics, and 82% using the optimal. We also show that if we restrict the radius of load movement, we can achieve a significant reduction in the number of panels for all net-zero time periods. For instance, for $r=500$ kms, we see a reduction of about 9.7%, 27%, and 53% for net-zero year, month, and week respectively. We saw that restricting to the top 500 data centers is a good middle ground for achieving near-optimal number of panels installed at major data center locations. In conclusion, we demonstrated that by leveraging locations with high solar output, we can significantly reduce the number of panels needed to serve load and achieve net-zero status for global IDNs. For this paper, we assumed a constant PUE value across different locations; however, cooler environments may have lower PUE values. As future work, we plan to study if we can combine renewables like solar and open air cooling to develop even more efficient solutions for greening IDNs.

Acknowledgments: This research is supported in part by NSF grants CNS-1422245, CNS-1413998 and by the Massachusetts Department of Energy Resources.

REFERENCES

- [1] 'Apple Data Centers'. https://images.apple.com/environment/pdf/Apple_Environmental_Responsibility_Report_2017.pdf.
- [2] 'Google Data Centers'. <https://www.google.com/about/datacenters/renewable/>.
- [3] L. Barroso and U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 2007.
- [4] Josep Lluís Berral, Iñigo Goiri, Thu D. Nguyen, Ricard Gavaldà, Jordi Torres, and Ricardo Bianchini. Building green cloud services at low cost. In *Proc. ICDCS 2014, Madrid, Spain, June 30 - July 3, 2014*, pages 449–460, 2014.
- [5] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, and Natarajan Gautam. Managing server energy and operational costs in hosting centers. In *Proc. ACM SIGMETRICS '05*, pages 303–314, New York, NY, USA. ACM.
- [6] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, third edition, 2009.
- [7] *Solar Panel Power Rating*. <https://www.energysage.com/solar/buyers-guide/solar-panel-efficiency/>.
- [8] M. Arlitt et al. Towards the design and operation of net-zero energy data centers. In *ITherm*, 2012.
- [9] Peter Xiang Gao, Andrew R. Curtis, Bernard Wong, and Srinivasan Keshav. It's not easy being green. *SIGCOMM Comput. Commun. Rev.*, 42(4):211–222, August 2012.
- [10] Iñigo Goiri, William Katsak, Kien Le, Thu D. Nguyen, and Ricardo Bianchini. Parasol and greenswitch: Managing datacenters powered by renewable energy. *SIGARCH Comput. Archit. News*, 41(1):51–64, March 2013.
- [11] Iñigo Goiri, Kien Le, Md. E. Haque, Ryan Beauchea, Thu D. Nguyen, Jordi Guitart, Jordi Torres, and Ricardo Bianchini. Greenslot: Scheduling energy consumption in green datacenters. In *Proc. SC '11*, pages 20:1–20:11, New York, NY, USA. ACM.
- [12] Vani Gupta, Prashant Shenoy, and Ramesh Sitaraman. Efficient solar provisioning for net-zero internet-scale distributed networks. Technical report, University of Massachusetts, Amherst, 2017.
- [13] J. Koomey. Growth in data center electricity use 2005 to 2010. 2010.
- [14] Chao Li, Yang Hu, Juncheng Gu, Yuan Jingling, and Tao Li. Oasis: Scaling out datacenter sustainably and economically. In *IEEE TPDS, vol. , no. , pp. 1, 5555*, 2016.
- [15] M. Lin, A. Wierman, L. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. In *Proc. IEEE INFOCOM*, 2011.
- [16] Minghong Lin, Adam Wierman, Lachlan L. H. Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Trans. Netw.*, 21(5):1378–1391, October 2013.
- [17] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H. Low, and Lachlan L. H. Andrew. Geographical load balancing with renewables. *SIGMETRICS Performance Evaluation Review*, 39(3):62–66, 2011.
- [18] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H. Low, and Lachlan L.H. Andrew. Greening geographical load balancing. In *Proc. ACM SIGMETRICS '11*, pages 233–244, New York, NY, USA. ACM.
- [19] Vimal Mathew, Ramesh K. Sitaraman, and Prashant J. Shenoy. Energy-efficient content delivery networks using cluster shutdown. In *Proc. IGCC 2013, Arlington, VA, USA*, pages 1–10, 2013.
- [20] E. Nygren, R.K. Sitaraman, and J. Sun. The Akamai Network: A platform for high-performance Internet applications. *ACM SIGOPS Operating Systems Review*, 44(3):2–19, 2010.
- [21] S. Pelley, D. Meisner, T. Wenisch, and J. VanGilder. Understanding and abstracting total data center power. In *Workshop on Energy-Efficient Design*, 2009.
- [22] *PV Education Solar Cell*. <http://www.pveducation.org/pvc/drom/solar-cell-structure>.
- [23] *PVWatts*. <https://developer.nrel.gov/docs/solar/pvwatts-v5/>.
- [24] *PVWatts V5 API*. <http://www.nrel.gov/docs/fy14osti/62641.pdf>.
- [25] A. Qureshi. Power-demand routing in massive geo-distributed systems. 2010.
- [26] Navin Sharma, Sean Barker, David Irwin, and Prashant Shenoy. Blink: Managing server clusters on intermittent power. *SIGARCH Comput. Archit. News*, 39(1):185–198, March 2011.
- [27] Enida Sheme, Neki Frashri, Simon Holmbacka, Sbastien Lafond, and Drazen Luzanin. Datacenters powered by renewable energy: A case study for 60 degrees latitude north. In *Proc. SoftCOM*, page 15, 2016.
- [28] *The Average PUE is 1.8: Uptime Institute*. <http://bit.ly/1eEqN9d>.