Concentrating Renewable Energy in Grid-Tied Datacenters

Nan Deng, Christopher Stewart, and Jing Li
The Ohio State University

***Abstract*—Datacenters, the large server farms that host widely used Internet services, account for a larger fraction of worldwide carbon emissions each year. Increasingly, datacenters are reducing their emissions by using clean, renewable energy from rooftop solar panels to partially power their servers. While some customers value renewable-powered servers, many others are indifferent. We argue that renewable energy produced on site should be concentrated as much as possible on the servers used by green customers. This paper introduces a new metric, the *renewable-powered instance*, that measures the concentration of renewable energy in datacenters. We conducted a simulation-based study of renewable-energy datacenters, focusing on the grid tie—the device most commonly used to integrate renewable energy. We found that grid-tie placement has first-order effects on renewable-energy concentration.**

I. Introduction

Today's datacenters comprise thousands of networked servers that collectively use a lot of electricity [12]. In 2006, datacenters in the US used more electricity (61 billion kWh [25]) than 6 states. By 2020, datacenters worldwide are on track to emit 340 metric megatons of CO_2 annually, more than entire countries [17]. Small but motivated groups of concerned datacenter customers have demanded that datacenters reduce their carbon emissions, for example:

- Over 300,000 Facebook users have petitioned the social networking site to use renewable energy to power its datacenter [26] and
- a Rackspace Hosting survey found that 35% of their customers would pay a premium for services that feature renewable energy or carbon offsets [9].

In fact, consumer studies show that green computing, in general, is an under served niche market and that carbon awareness offers a "strategic competitive edge" [22].

Some datacenters reduce their carbon footprint by producing renewable energy on site via solar panels or wind turbines [1], [4], [11]— 7 such renewable-energy datacenters were announced in 2009 [11]. While the prospect of renewable-powered servers pleases green customers, other customers (today, the majority) may be indifferent. In the interest of maximizing the competitive advantage, we believe that renewable energy produced on site should power only the servers belonging to green customers. Datacenters that offer such differentiated service allow customers to "go green to be seen [15]", satisfying the altruistic motives of green customers. Further, such differentiated service allows datacenters to incrementally deploy on-site renewable energy, making investments proportional to the demand from green customers.

For this paper, we studied the device most commonly used to integrate on-site renewable energy: *the grid tie*. Grid ties transform electricity from a primary power source (e.g., rooftop solar panels) into electricity with the same frequency and voltage as a secondary power source (e.g., the electric grid). After passing through a grid tie, electricity produced on site can safely replace electricity drawn from the electric grid, reducing the net amount of grid energy used by servers downstream. The grid fills any unmet power needs when on-site sources are insufficient and it accepts excess energy when the on-site sources over produce. Thus, the concentration of renewable energy in grid-tied datacenters depends on 1) the amount of energy produced by the primary source and 2) the number of servers that are downstream to the grid tie. If every server is downstream, the secondary source will probably need to supply some power, diluting the concentration of renewable energy per server. If just one server connects to the grid tie, only that server is guaranteed to have high concentrations of renewable energy. *Our goal was to quantify the impact of grid-tie placement on the concentration of renewable energy.*

Alternative Approaches to Green Computing: Renewable-energy credits allow datacenters to offset their carbon footprint without producing electricity on site. In this model, datacenters use electricity from dirty-energy sources while paying others to produce an equal amount of clean energy somewhere else. In contrast, renewable-energy datacenters use clean energy directly, reducing both their net and *gross* carbon footprint. The two approaches are qualitatively different, both deserving systems level research.

II. Related Work

Prior work focused on renewable-energy datacenters that use transfer switches instead of grid ties. Transfer switches isolate electricity produced on site, keeping it separate from the electric grid. At any point in time, servers in a transfer-switched datacenter can receive power from only 1 source. Grid ties allow multiple sources to power a server. In a grid-tied datacenter, on-site sources can contribute to only a portion of total energy needs, eliminating the need for costly batteries. To the best of our knowledge, this is the first paper to study on grid-tied datacenters.

In prior work, Gmach et al. [13] observed that popular Internet services could dynamically change their energy needs to match the production of on-site renewable energy. In a similar approach, Li et al. [19] dynamically reconfigured multi-core processors in response to the production patterns

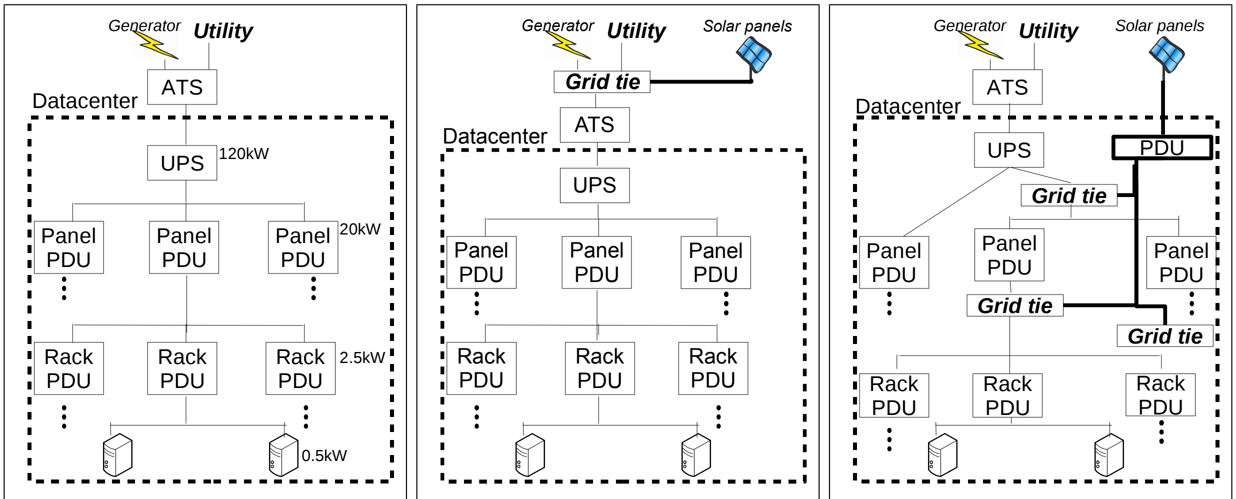

(a) The power delivery subsystem in our department today.

(b) A large grid tie integrated just above the UPS.

(c) Multiple grid ties integrated and controlled.

Fig. 1: Simplified views of power delivery and grid-tie integration.

of solar panels. These works sacrifice performance to avoid wasting renewable energy during low production periods. Gmach et al. [14] studied another alternative: purchasing on-site batteries. Their studied provided a systematic way to look for sweet spots in the cost/benefit tradeoff. More generally, this work uses a research methodology similar to the above prior work. We combine datacenter workload traces with energy production traces from renewable sources, using simulation to understand the effects system changes.

Stewart and Shen [24] provided an early study on renewable-energy datacenters, describing the core research challenges, architectural features, and systems infrastructure. They found that transfer switches, commonly used for off-grid buildings, were not ideal for datacenters, because they presented an undesirable tradeoff between safety from brownouts and renewable energy usage, arguing instead for a device (the grid tie) that could combine electricity from renewable sources and the grid on the same circuit. Sharma et al. [23] proposed a new energy abstraction to allow datacenter customers to program for intermittent power.

III. Grid-tied Power Delivery

The servers and network switches inside datacenters operate at very low voltages, using much less power than the entire facility. These devices are powered on parallel circuits—i.e., each device has its own receptacle for 120V AC power. These parallel circuits are combined using higher voltage PDUs (208V, 240V, or 480V) that eventually connect to uninterruptible power supplies (UPS) that serve the whole facility [5], [12], [24]. Grid ties can be integrated into any of the above devices. Each integration point places certain requirements on the grid tie during manufacturing, e.g., UPS integration normally requires that grid ties output 480V, 3-phase AC. Most grid-tie manufacturers offer at least 1 product line for each integration point [7], creating a lot of options for the placement of grid ties. Below, we formally describe the grid-tie placement problem and make a case for a new metric to study the concentration of renewable-energy in datacenters.

A. Grid-Tie Placement

Power delivery components (P) and independently powered servers (D) in a datacenter can be represented as vertices in a directed acyclic graph, $G = (P \cup D, E)$. Cords in the datacenter are edges (E). Grid-tie placements are represented by a function $gt[e]$ that is 1 only if a grid tie is integrated at edge e; the function is 0 otherwise. Let $e = (v_1 \rightarrow v_2)$, power from the grid tie will first flow downstream to descendants of v_2, i.e., the servers that draw power from v_2's circuit. The amount of renewable energy flowing into an integration point is represented as $r[e]$. The power draw of a vertex v and its descendants is represented by the function $f[v]$. In this model, grid energy used is: $g = f[v_0] - \Sigma_{e \in E} r[e] \cdot gt[e]$.

B. Renewable-Powered Instances

After grid ties transform electricity produced on-site, the concentration of renewable electrons versus dirty electrons can not be measured directly. It is physically impossible to tag renewable electrons. However, prior power engineering research [6] provides a reasonable principle: Electricity from multiple sources is distributed proportionally from sources to loads. Adapted to the datacenter environment, the percentage of renewable energy powering downstream servers equals the percentage of renewable energy used upstream at the grid tie's integration point. More precisely, the percentage of renewable energy powering a datacenter server d_i is:

$$c[d_i] = MAX_j\left(\frac{r[(v_x, v_j)] \cdot gt[v_x, v_j]}{f[v_j]}\right)$$
$$\text{where } d_i \in D \wedge \ d_i \in descendants(v_j)$$

We define a k renewable-powered instance (k-RPI) as a server that gets at least $k\%$ of its energy needs from renewable energy. Here, k reflects a minimum level of concentration supplied to green customers and the result is a metric that can be used to compare grid-tie placements.

$$\begin{aligned} k-rpi &= \{d_0, d_1, ...d_n\} \\ &\text{where } d_i \in D \wedge \\ \exists v_j, v_x: \; & d_i \in descendants(v_j) \wedge \\ & \frac{r[(v_x, v_j)] \cdot gt[v_x, v_j]}{f[v_j]} > k \end{aligned}$$

Figure 1(a) shows a portion of the power delivery subsystem for our department's two server rooms (a tier-1 datacenter). There is one UPS operating at 3-phase 480V with a power capacity of 120KW. The UPS delivers power to 7 PDU panels which ultimately deliver 120V-power to 93 rack-level PDUs. These rack-level PDUs power 332 servers, networking equipment, and (in 4 cases) other power distribution strips attached to monitors and tower-style servers. On average, there are 32 cores per rack. Figure 1(b) shows our subsystem with a grid tie placed between the transfer switch and the utility provider. A real datacenter with on-site renewable energy uses a very similar grid-tie integration strategy [18]. Figure 1(c) shows our subsystem with a smart PDU controlling multiple grid ties. In Section IV, we will see that this particular placement performs well.

IV. Study of Renewable-Powered Instances

In our work, we used a trace-driven simulation approach to explore the effect of different grid-tie placements under 2 datacenter power delivery subsystems. The inputs to our simulator were 1) a graph of the power delivery subsystem, 2) a trace of per-server energy needs, and 3) a trace of renewable energy production. Our default datacenter subsystem and trace comes from our department's datacenter which comprises more than 104 rack- and panel-level PDUs. It is possible to integrate a grid tie at each, meaning that even for our small datacenter there are thousands of ways to place multiple grid ties. We used solar and wind renewable-energy traces collected from different places in the U.S.

Our department's power delivery subsystem was simulated with 500 randomly selected grid-tie placements. Each placement had at most three grid ties with randomly chosen integration points. After the integration points were selected, we chose the grid tie from [7] that most closely matched the capacity of the integration point (e.g., 240V, 6KW PDU). Data from [7] is used to simulate the grid tie's energy needs (i.e., its inefficiency) as a function of its load. We collected the energy usage of rack-level PDUs in our department from March to June 2010 and looped this trace to get 1 year of virtual time entries for our simulator. For the trace of renewable energy production, we used a 1-year (2004) trace of wind energy production from Cheyenne, WY [21], the site of a well-known datacenter with on-site renewable energy [3]. The trace came from the National Renewable Energy Laboratory [20]. We linearly scaled the production trace to produce 20% of the energy used by the datacenter, reflecting an incremental approach to deploying on-site renewable energy.

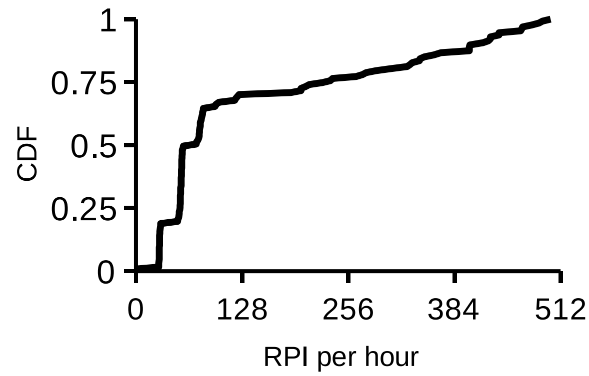

(a) CDF of the 100% 1-hour RPI per hour

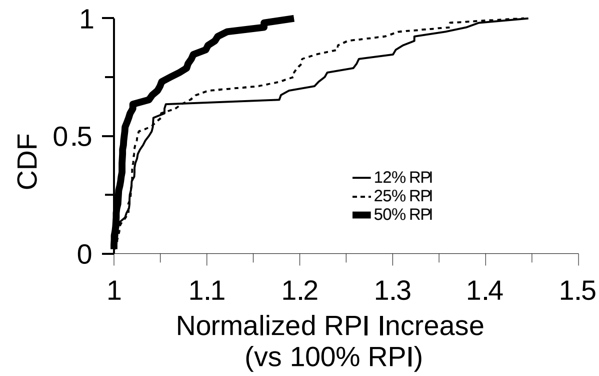

(b) RPI = X% renewable power for 1-hour

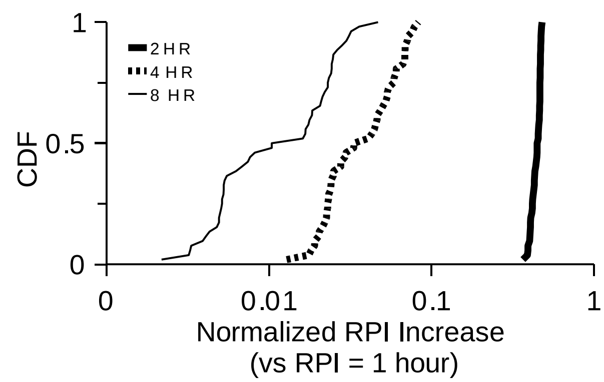

(c) RPI = 100% renewable power X-hours

Fig. 2: Performance of 500 randomly selected grid-tie placements under the WY production trace, our department's energy usage trace, and 20% renewable to energy ratio. An RPI indicates that 100% of a compute core's [2] energy for 1 hour came from an on-site renewable source. This our default test setup, unless otherwise mentioned the reader can assume that experiments in this section have these settings.

Figure 2(a) plots the distribution of 100% renewable-powered instances (100-RPI, or RPI for simplicity) across the tested placements. RPI production varies a lot; some placements allow datacenter managers to track 499 RPI per hour while other placements can't report any RPI. Figure 2(a) shows that grid-tie placement *can* affect a datacenter's ability

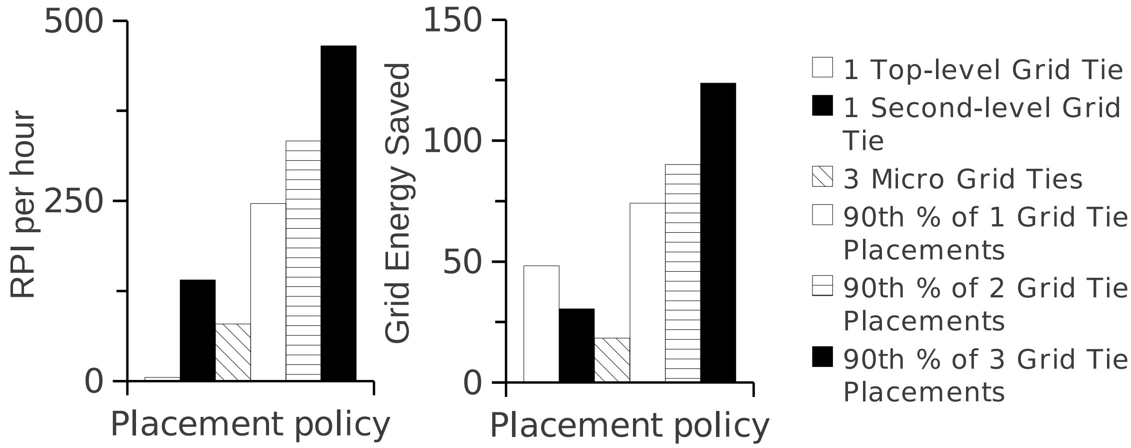

Fig. 3: Comparing commonly used grid-tie placements to the 90^{th} percentile of RPI producers.

to concentrate renewable energy. In the remainder of this subsection, we examine factors that could vary from one datacenter to another to see if they change this result.

An RPI in our default setup required a grid-tie placement to power a compute core with 100% renewable energy for an hour. Both of these parameters, the percentage of renewable energy and the duration of delivery, could vary from datacenter to datacenter. Figure 2(b) shows that decreasing the first parameter improved some placements much more than others. Looking further into the data, we found that the best placements were getting better. When 12% renewable power constituted an RPI, we observed that the 90^{th} percentile of grid-tie placements could produce 539 more RPI per hour than the 10^{th} percentile. This was 1.39 times larger than the difference between the 90^{th} and 10^{th} percentiles when 100% renewable power was required.

Figure 2(c) shows that increasing the duration required for an RPI can significantly reduce the total number of RPI (note the x-axis is log scale). As a result, the absolute difference between grid-tie placements is also reduced. When the required duration of an RPI is 8 hours, the 90^{th} percentile produces only 11 more RPI per hour than the 10^{th} percentile—but these placements only produce 12 RPI per hour total.

A. Comparison of Common Grid-Tie Placements

Next, we asked "how well do the grid-tie placements used in practice today perform?" We spoke with 2 datacenter managers who were able to share high-level information about their power delivery structure [16], [18]. Both of these datacenters integrated only one grid tie near the top of their power delivery hierarchy. One integrated the grid tie above their UPS and ATS systems (see Figure 1(b)), meaning that renewable power could support their entire facility. The other integrate the grid tie on the B-side of their datacenter. Here, the grid tie could only power part of the datacenter at a time. In contrast to these 1-grid-tie placement strategies, we also considered a new trend in practice: the use of multiple micro grid ties [10]. Because grid ties consume energy during the process of making primary and secondary sources compatible, and the consumption differs across grid ties with different capacity or under different load [7], placing several grid ties at multiple levels of power delivery subsystem may affect the grid energy used. To study the impact of multiple gird-tie placement, the idea is to integrate renewable energy into the system at easy to install 120V wall outlets rather than larger capacity UPS or PDU integration points.

We applied these grid-tie placements to our department's datacenter and compared the results to the 90^{th} percentile of our randomly selected placements. Note, Figure 1(c) shows the 90^{th} percentile placement for the placements with 3 grid ties. Figure 3 shows that the placements used in practice performed considerably worse. Below we highlight important results that further confirm the impact of grid-tie placement:

- The naive approach of placing grid ties near the utility provider can significantly decrease RPI production. Placing 1 grid tie wisely at the rack-level or panel-level increased RPI per hour by 1.76X.
- The integration of multiple grid ties improved RPI production by 1.89X while using 60% less grid energy. As shown in Section III, simply turning on grid ties consumes a lot of power. When renewable-energy sources produce low amounts of power, big large capacity grid ties can't be turned on—explaining the poor savings from the top-level grid tie. Adding multiple grid ties allows the power delivery subsystem to flexibly use such renewable energy during low production periods.
- Micro grid ties produced 17% of RPI produced by the 90^{th} percentile of 3 grid ties. This was because the micro grid ties were unable to bring in enough renewable power. A high-capacity grid tie ensures that power produced during windy surges flowed downstream.
- The best placements have significant overhead, almost 9% of the grid energy savings in some cases.

Figure 4(a) shows that the use of multiple grid ties increases the disparity between grid-tie placements. The difference between the 90^{th} percentile and the 10^{th} percentile of placements

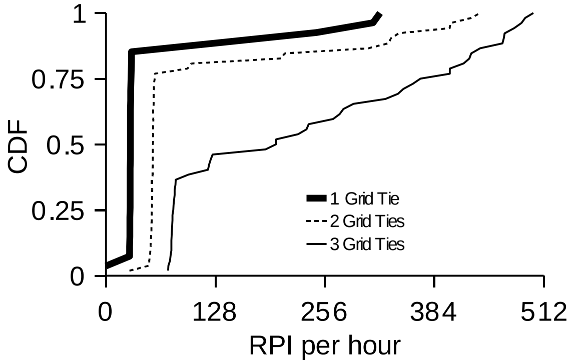

(a) Impact of multiple grid ties

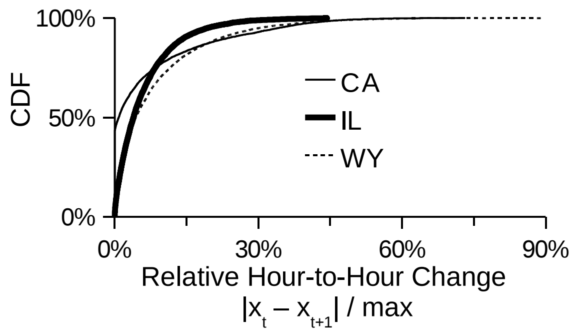

(b) The volatility of renewable-energy sources in Chicago,*IL*, Cheyenne,*WY*, and Rohnert Park,*CA*.

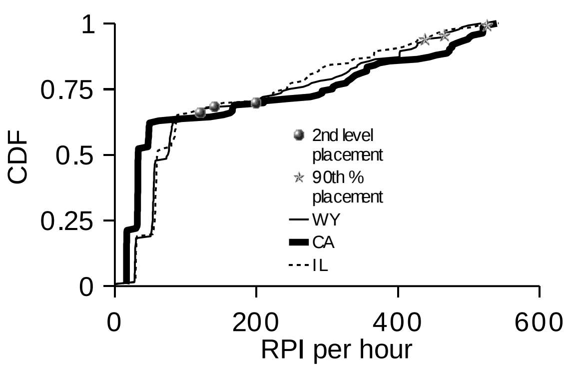

(c) Impact of Geographic Location

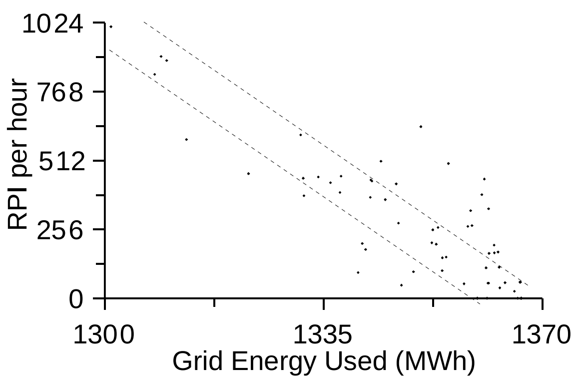

(d) Impact of Datacenter Size and Workload

Fig. 4: The impact of practical design parameters.

with 3 grid ties was 388 RPI per hour, more than 2.5 times larger than the difference for placements with 1 grid tie. Paired with Figure 1(b), these results highlight the importance of grid-tie placement in a new way. Yes, using multiple grid ties can lead to better performance, but it can also increase the opportunity cost from poor placements.

B. Renewable Energy Production Patterns

We examined renewable energy production in 3 areas used by today's renewable-energy datacenters: Woodstock, IL [18], Rohnert Park, CA [11] and Cheyenne, WY [3]. The CA trace reflects solar-energy production (taken from the Solar Advisor Model [8]), and the IL and WY traces reflect wind-energy production (taken from NREL [21]). Each trace captures hourly snapshots of the monitored energy source's performance (i.e., production divided by capacity) for 2004. The traces were scaled linearly to have the same aggregate energy output.

Not only do these areas reflect different real-world site locations, the production traces offer fundamentally different production patterns. It is well known that solar energy exhibits strong daily cycles and wind energy does not. The auto correlation function with a 24-hour delay (a measure of daily periodicity) was more than 3 times larger (0.88) for the CA trace than for the WY and IL traces (0.09 and 0.27). The WY and CA traces exhibited larger changes from hour-to-hour than the IL trace, shown in Figure 4(b). We observe that the largest hour-to-hour changes in CA and WY (72% and 89% respectively) are more than 60% larger than the largest change in the IL trace (44%).

Figure 4(c) shows that studied production patterns had only a small affect on the full distribution of grid-tie placements. The median solar placement produced only 1.05X more RPI than the median wind placements, between the wind sites, the distributions were essentially the same. These results suggest that datacenter managers in a wide range of areas can expect similar returns on investments in renewable energy. Figure 4(c) also explores the performance of two specific placements, the 2^{nd} level placement approach described earlier and the 90^{th} percentile placement in Cheyenne, WY. We were surprised to see the the performance of the individual placements varied a lot, producing 1.5X and 1.8X more RPI in the solar study. When we investigated this result, we found that these placements both place a grid tie at the largest panel-level PDU which happens to have energy consumption patterns that closely match the daily peaks of solar energy production. If our simulated solar panel system produced just 10% more energy, the placements would send a substantial amount of renewable energy upstream and would be less effective relative to their performance under the wind sites. This results shows that the performance of individual grid-tie placements can vary from location to location.

Next, we examined the impact of amount of renewable energy produced on site. We scaled our production traces to produce 5%, 20% (default) and 100% of the energy used by the datacenter's energy consumption traces. We found that the tail of the distribution became much heavier as we increased the amount of renewable energy produced on site.

The difference between the 90^{th} percentile and 10^{th} percentile was 1.56 times larger when the amount of on-site production was equal to the on-site consumption compared to the default case. Comparatively, it was only 6% smaller when the on-site production was 5% of aggregate energy needs.

C. Datacenter Size and Workload

We concluded our study by examining grid tie placement on a larger datacenter. We monitored the energy usage and mapped the power delivery subsystem for a datacenter serving our entire university (one of the largest university's in the nation). The workload at this datacenter differed from the mostly research workloads in our department, running more enterprise-oriented applications, like SAP and BlackBoard. Further, this datacenter uses virtualization and other energy efficiency techniques. We compared the hour-to-hour volatility of the energy usage between these datacenters and found that the average hour-to-hour change in the university-wide datacenter was 10 times more than in our department datacenter which we attribute to fluctuating workload and higher energy efficiency. The median hourly change of university-wide PDUs was 15% of its peak consumption.

Figure 4(d) plots the grid energy used against the RPI for the university-wide datacenter. This result provides further confirmation of our hypothesis that grid-tie placements can significantly affect the ability to concentrate renewable energy. We once again observe that the worst placements essentially eliminate the ability to track when servers are powered by renewable energy (i.e., RPI). In future work, we plan to study the economic value of the RPI metric.

V. Conclusion

In this paper, we presented a simulation-based study on the concentration of renewable energy in datacenters. Our study focused on the grid tie— the device most commonly used to integrate renewable energy into the datacenter's power delivery system. We proposed *renewable-powered instances* as a metric to understand the concentration of renewable energy. We found that proper grid-tie placement can increase the degree of concentration and reduce the use of grid energy. Our study found that such grid-tie placement mattered across a variety of parameters including the geographic location of datacenter, its size and the amount of renewable energy produced on site.

VI. Acknowledgements

We would like to thank undergraduate summer researchers Joel Poualeu and Nelson Burgos for their contributions to our datacenter traces. Daniel Gmach and Amip Shah provided feedback on early ideas. This work was funded by Christopher Stewart's start-up funds at The Ohio State University.

References

[1] Aiso.net: Web hosting as nature intended. http://aiso.net.

[2] Amazon elastic compute cloud. http://aws.amazon.com/ec2/pricing/.

[3] Green House Data: Greening the data center. http://www.greenhousedata.com/.

[4] Google solar panel project. http://www.google.com/corporate/solarpanels/home, June 2007.

[5] L. Barroso and U. Holzle. *The Datacenter as a Computer – An Introduction to to the Design of Warhouse-Scale Machines*. Morgan and Claypool Publishers, 2009.

[6] J. Bialek. Tracing the flow of electricity. *IEE Proceedings on Generation, Transmission, and Distribution*, 143(4), 1996.

[7] W. Bower, C. Whitaker, W. Erdman, M. Behnke, and M. Fitzgerald. Performance test protocol for evaluating inverters used in grid-connected photovoltaic systems. http://www.gosolarcalifornia.org/, 2004.

[8] C. Cameron, W. Boyson, and D. Riely. Comparison of pv system performance-model predictions with measured pv system performance. In *IEEE Photovoltaic Specialists Conference*, May 2008.

[9] J. Chattaway. Rackspace green survey. In *Rackspace Hosting White Paper*, June 2008.

[10] E. Energy. Enphase grid tie inverters. http://www.enphase.com/.

[11] Environmental Leader. Data centers power up savings with renewable energy. http://www.environmentalleader.com/2009/07/29/data-centers-power\-up-savings-with-renewable-energy/.

[12] X. Fan, W. Weber, and L. Barroso. Power provisioning for a warehouse-sized computer. In *Int'l Symp. on Computer Architecture*, June 2007.

[13] D. Gmach, Y. Chen, A. Shah, J. Rolia, C. Bash, T. Christian, and R. Sharma. Profiling sustainability of data centers. In *ISSST*, 2010.

[14] D. Gmach, J. Rolia, C. Bash, Y. Chen, T. Christian, A. Shah, R. Sharma, and Z. Wang. Capacity planning and power management to exploit sustainable energy. In *International Conference on Network and Services Management*, Oct. 2010.

[15] V. Griskevicius, J. M. Tybur, and B. Van den Bergh. Going green to be seen: status, reputation, and conspicuous conservation. *J Pers Soc Psychol*, 98(3):392–404, 2010.

[16] John Mattson, Emerson Inc. Personal communication. http://emerson-datacenter.com.

[17] J. Kaplan, W. Forrest, and N. Kindler. Revolutionizing data center energy efficiency. McKinsey and Company, 2008.

[18] Larry O'Connor, CEO of OWC.net. Personal communication. http://owc.net.

[19] C. Li, W. Zhang, C. Cho, and T. Li. Solarcore: Solar energy driven multi-core architecture power management. In *High-Performance Computing Architecture*, Feb. 2011.

[20] National Renewable Energy Laboratory. NREL: Western wind resources dataset. http://wind.nrel.gov/Web_nrel/, 2009.

[21] C. Potter, D. Lew, J. McCaa, S. Cheng, S. Eichelberger, and E. Grimit. Creating the dataset for the western wind and solar integration study (u.s.a.). In *International Workshop on Large Scale Integration of Wind Power and on Transmission Networks for Offshore Wind Farms*, May 2008.

[22] SAP. Sap insight: Product environmental footprint and ecolabeling. Aug. 2009.

[23] N. Sharma, S. Barker, D. Irwin, and P. Shenoy. Blink: managing server clusters on intermittent power. In *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems*, ASPLOS '11, pages 185–198, New York, NY, USA, 2011. ACM.

[24] C. Stewart and K. Shen. Some joules are more precious than others: Managing renewable energy in the datacenter. In *Workshop on Power Aware Computing and Systems(HotPower)*, Sept. 2009.

[25] US Environment Protection Agency. Report to congress on server and data center energy efficiency. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

[26] G. USA. Facebook status update: Renewable energy now. http://www.greenpeace.org/usa/news/facebook-update-renewable-ene.