Geographical Load Balancing with Renewables

Zhenhua Liu, Minghong Lin,
Adam Wierman, Steven H. Low
CMS, California Institute of Technology, Pasadena
{zliu2,mhlin,adamw,slow}@caltech.edu

Lachlan L. H. Andrew
Faculty of ICT
Swinburne University of Technology, Australia
landrew@swin.edu.au

ABSTRACT

Given the significant energy consumption of data centers, improving their energy efficiency is an important social problem. However, energy efficiency is necessary but not sufficient for sustainability, which demands reduced usage of energy from fossil fuels. This paper investigates the feasibility of powering internet-scale systems using (nearly) entirely renewable energy. We perform a trace-based study to evaluate three issues related to achieving this goal: the impact of geographical load balancing, the role of storage, and the optimal mix of renewables. Our results highlight that geographical load balancing can significantly reduce the required capacity of renewable energy by using the energy more efficiently with "follow the renewables" routing. Further, our results show that small-scale storage can be useful, especially in combination with geographical load balancing, and that an optimal mix of renewables includes significantly more wind than photovoltaic solar.

1. INTRODUCTION

Energy consumption of data centers is of growing concern to both operators and society. Electricity for internet-scale systems costs millions of dollars per month [17] and, though ICT uses only a small percentage of electricity today, the growth of electricity in ICT exceeds nearly all sectors of the economy. For these reasons, and more, ICT must play its part in reducing our dependence on fossil fuels.

This can be achieved by using renewable energy to power data centers. Already, data centers are often powered by a "green" portfolio of energy [9, 12, 14]. However, most studies of powering data centers *entirely* with renewable energy have focused on powering individual data centers, e.g., [4, 3]. These have shown that it is challenging to power a data center using only local wind and solar energy without large-scale storage, due to the intermittency and unpredictability of these sources.

The goal of this paper is to illustrate that the geographical diversity of internet-scale services significantly improves the efficiency of the usage of renewable energy. This numerical study complements testbeds such [16]. Further, we illustrate that algorithmic solutions, such as geographical load balancing (GLB), can play a vital role in reducing the necessary capacity of renewable energy installed.

Specifically, we perform numerical experiments using real traffic workloads and real data about the availability of renewables combined with an analytic model for a geographically distributed system. Using this setup, we investigate issues related to the feasibility of powering an internet-scale system (nearly) completely with renewable energy.

Our study yields three key insights.

First, GLB significantly reduces the capacity of renewables needed to move toward a "green" system, since GLB allows the system to use "follow the renewables" routing, which reduces both the financial cost and the "brown" non-renewable energy usage. However, we also show the importance of using a fast control time-scale for GLB. If routing and capacity decisions are made only once an hour, then significantly more brown energy is consumed than if the adjustments are made every ten minutes. Unfortunately, adjustments at this faster time-scale may not be feasible due to server wear-and-tear and other concerns.

Second, we investigate the value of storage when using renewable energy. Often large-scale storage is viewed as essential for moving toward a completely renewable energy portfolio. However, our study shows that small-scale storage in combination with GLB is sufficient in moving to a portfolio of nearly completely renewable energy sources. This is particularly exciting since the UPSs in use at data centers today could be used to provide small-scale storage with few engineering changes.

Third, we find that wind is more valuable than solar for internet-scale systems, especially when GLB is used, because wind has little correlation across locations, and is available during both night and day. Thus, if one aggregates over many locations, there is much less variation in the availability [2]. The optimal mix seems to be dominated by wind, but include some solar to handle the peak workload around noon. For the traces we consider, the optimal portfolio is 80% wind and 20% solar. This ratio may depend on the quality of the local wind resources; moreover, workloads with higher (lower) diurnal peak-to-mean ratios may benefit from a higher (lower) solar component.

2. SETUP

Our numeric experiments combine analytic models with real traces for workload and renewable availability, to allow controlled experimentation but provide realistic findings. We now explain the setup, which extends that of [11].

2.1 The workload

Our workload model considers a set J of sources of requests, with one source at the center of each of the 48 continental states in the US.

We consider one hour time slots for a week. At the start of each slot, the routing and capacity of each data center are updated. We use a slot length of 1 hour so that servers are not turned on and off too frequently given the significant wear-and-tear costs of power-cycling. Let $L_j(t)$ denote the mean arrival rate from source j at time t.

To provide realistic estimates, we use real-world traces to define $L_j(t)$. The workload of each source is scaled proportionally to the number of internet users in the corresponding state. The workload is taken from a trace at Hewlett-Packard Labs [3] and is shifted in time to account for time zone of each state. The workload per internet user used in this paper is shown in Figure 1.
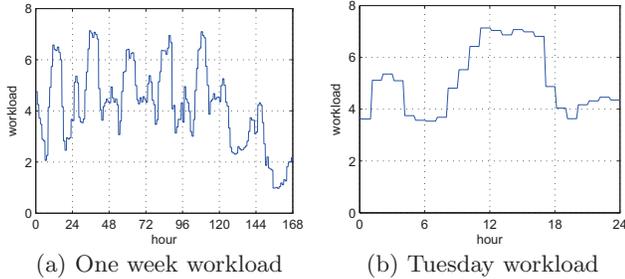
(a) One week workload   (b) Tuesday workload

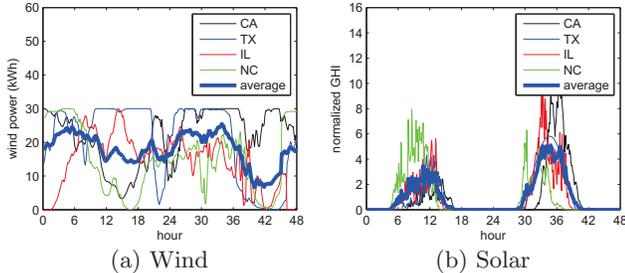**Figure 1: HP workload trace.**

(a) Wind   (b) Solar

Figure 2: Renewable generation for two days.

2.2 The availability of renewable energy

To capture the availability of wind and solar energy, we use traces of wind speed and Global Horizontal Irradiance (GHI) obtained from [6, 7] that have measurements every 10 minutes for a year. The traces of four states (CA, TX, IL, NC) are illustrated in Figure 2 (GHI is normalized to have average of 1). Note that as "solar", *we only consider photovoltaic generation and not solar thermal*, because of the significant infrastructure required for solar thermal plants. Since solar thermal plants typically incorporate a day's thermal storage [15], the results would be very different if solar thermal were considered.

These figures illustrate two important features of renewable energy: spatial variation and temporal variation. In particular, we see that wind energy does not exhibit a clearly predictable pattern throughout the day and that there is little correlation across the locations considered. In contrast, solar energy has a predictable peak during the day and is highly correlated across the locations.

In our investigation, we scale the "capacity" of wind and solar. When doing so, we scale the availability of wind and solar linearly, which is suitable when considering scaling the size of the wind farm or solar installation. Throughout, we measure the capacity of renewables as the ratio of the average renewable generation to the minimal energy required to serve the average workload. Thus, a capacity of $c = 2$ means that the average renewable generation is twice the minimal energy required to serve the average workload. In the following we set capacity to $c = 2$ by default, but vary it in Figures 5–7.

2.3 The internet-scale system

We model the internet-scale system as a set N of 10 data centers, placed at the centers of states known to have Google data centers [13], namely California, Washington, Oregon, Illinois, Georgia, Virginia, Texas, Florida, North Carolina, and South Carolina. Data center $i \in N$ contains M_i homogeneous servers, where M_i is twice the minimal number of servers required to serve the peak workload of i under a scheme which routes traffic to the nearest data center. Further, the renewable availability at each data center is defined
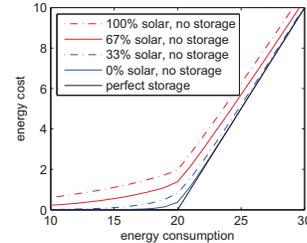

Figure 3: Energy cost as a function of total energy consumption, when the average renewable energy in a one-hour slot is 20.

by the trace from a nearby location; this was usually within the same state, but in five cases it was a different trace from a nearby state.

The two key control decisions of geographical load balancing are (i) determining $\lambda_{ij}(t)$, the amount of traffic routed from source j to data center i; and (ii) determining $m_i(t) \in \{0, \ldots, M_i\}$, the number of active servers at data center i. The objective is to choose $\lambda_{ij}(t)$ and $m_i(t)$ to minimize the "cost" of the system. The cost is the sum of the *energy cost* and the *delay cost* (lost revenue), below. We neglect the cost of altering the number of active servers $m_i(t)$; this can be incorporated using an approach similar to that of [10].

2.3.1 Delay cost

The delay cost captures the lost revenue incurred because of the delay experienced by the requests, where the delay includes both the network delay from source j to data center i, d_{ij}, and the queueing delay at i. We model d_{ij} to be the distance between source and data center, divided by the speed of $200 \, \text{km/ms}$ plus a constant ($5 \, \text{ms}$), resulting in delay ranging in $[5 \, \text{ms}, 56 \, \text{ms}]$. We model the queueing delays using parallel M/G/1/Processor Sharing queues with the total load $\lambda_i(t) = \sum_j \lambda_{ij}(t)$ divided equally among the $m_i(t)$ active servers, each having service rate $\mu_i = 0.1(\text{ms})^{-1}$. This parameter setting makes the average delay $20 \, \text{ms}$ when the utilization is 0.5, which is reasonable. Note that this model is not new, and is referred to as the GLB-LIN-Q model in [11].

2.3.2 Energy cost

To capture the effect of integrating renewable energy, we model the energy cost as the number of active servers *excluding* those that can be powered by renewables. Note that this assumes that data centers operate their own wind and solar generations and pay no marginal cost for renewable energy. Further, it ignores the installation and maintenance costs of renewable generation.

Quantitatively, if the renewable energy available at data center i at time t is $r_i(t)$, measured in terms of number of servers that can be powered, then the energy cost is

$$p_i(m_i(t) - r_i(t))^+ \tag{1}$$

The p_i for each data center is constant, and equals to the industrial electricity price of each state in May 2010 [8]. This contrasts with the total power $p_i m_i$ used typically, e.g., in [11]. Although data center could instead sell the renewable supply to the grid, this will have only a small effect on p_i in a future dominated by renewable energy; when the renewable supply is high at the data center location, the local spot price of power is also likely to be low.

2.3.3 Storage

The above formula for the energy cost is simplistic because it assumes the cost is zero if fewer servers are provisioned than the *average* renewable generation in a time slot. While this is true if one can do "perfect" smoothing of the

renewables using storage, it is not necessarily true in practice. Note that the fluctuations in available power are usually considerably less than the total power draw. For our traces, "perfect" smoothing can be achieved using the amount of storage that can support the whole data center for several minutes. This is feasible, since this amount of storage is provided by currently deployed Uninterruptible Power Supplies (UPSs).

Without storage, there will be some cost incurred due to the variability of renewable availability within a time slot. Here the actual energy cost becomes

$$\mathbb{E}_\tau \left[p_i \left(m_i(t) - r_i(t, \tau) \right)^+ \right] \tag{2}$$

where (t, τ) is the sub-slot in time-slot t and $\mathbb{E}_\tau(r_i(t, \tau)) = r_i(t)$. Figure 3 shows several curves for California. The lowest curve is (1), i.e., the price for a system with enough storage to smooth renewable generation over one interval, and the upper ones for (2), i.e. the case of no storage, which is derived directly from the renewable traces as described above. The figure illustrates the curves in the case of an average renewable generation of 20 kW with $p_i = 1$, $\tau = 10$ min, and includes different mixtures of solar and wind. Interestingly, the added cost for not having storage increases as the percentage of solar increases in the renewable portfolio. The figures are almost the same for different $\mathbb{E}_\tau(r_i(t, \tau))$.

2.3.4   Total cost

Combining the discussions above, we can now write the total cost for an internet-scale system. In particular, for the case where the data centers have small-scale storage, the cost optimization that the system seeks to solve at time t becomes:

$$\min_{\mathbf{m}(t), \boldsymbol{\lambda}(t)} \quad \beta \sum_{i \in N} p_i \left(m_i(t) - r_i(t) \right)^+ \tag{3a}$$

$$+ \sum_{j \in J} \sum_{i \in N} \lambda_{ij}(t) \left(\frac{1}{\mu_i - \lambda_i(t)/m_i(t)} + d_{ij} \right)$$

$$\text{s.t.} \quad \sum_{i \in N} \lambda_{ij}(t) = L_j(t), \qquad \forall j \in J \tag{3b}$$

$$\lambda_{ij}(t) \geq 0, \qquad \forall i \in N, j \in J \tag{3c}$$

$$0 \leq m_i(t) \leq M_i, \qquad \forall i \in N \tag{3d}$$

$$\lambda_i(t) \leq m_i(t)\mu_i \qquad \forall i \in N. \tag{3e}$$

Here β determines the relative importance of energy and delay to the system. In the simulations β is set to 1 by default, but we also vary $\beta \in [1, 10]$ in Figure 4 to show its impact. Measurements [1] show that a 500 ms increase in delay reduces revenue by 20%, or 0.04%/ms. In a system where energy expenditure is 10% of the revenue, $\beta = 1$ or 10 corresponds to the more conservative values of 0.5 or 0.05%/ms, respectively.

When data center i can cheaply serve data from many sources, the upper bound $m_i(t) \leq M_i$ in (3d) can become binding, then $\lambda_{ij}(t)$ is strictly between 0 and $L_j(t)$.

When data centers have no energy storage, only the energy cost component of the optimization changes from (1) to (2). However, we cannot write the optimization in closed form for that case because the energy cost curves are determined by the renewable energy traces, as shown in Figure 3. To formulate and solve the optimization in this case we use a piecewise-linear approximation of the curves in Figure 3.

2.3.5   Geographical load balancing

The above sections describe the cost optimization that the internet-scale system seeks to solve; however they do not describe how the system actually performs geographical load balancing to solve the optimization. In [11], decentralized algorithms are presented, which can be used to achieve the optimal cost in the setting described above. Thus, in this paper we present as "GLB" the routing and capacity provisioning decisions which solve the cost optimization problem. Note that the results do not rely on the algorithms in [11]; they simply consider the optimal allocation.

As a benchmark for comparison, we consider a system that does no geographical load balancing, but instead routes all requests to the nearest data center and optimally adjusts the number of active servers at each location. We call this system 'LOCAL' and use it to illustrate the benefits that come from using geographical load balancing.

3.   RESULTS

With the setup described in the previous section, we have performed a number of numerical experiments to evaluate the feasibility of moving toward internet-scale systems powered (nearly) entirely by renewable energy. We focus on three issues: (i) the impact of geographical load balancing, (ii) the role of storage, and (iii) the optimal mix of wind and solar.

3.1   The impact of geographical load balancing

Geographical load balancing is known to provide internet-scale system operators significant energy cost savings, at the expense of small increases in network delay due to the fact that requests can be routed to where energy is cheap or renewable generation is high. This behavior is illustrated in Figure 4, which shows the average cost and delay under GLB versus LOCAL as the cost (β) of brown energy relative to delay is increased. The novelty of Figure 4 is in (a), which shows the reduction in consumption of brown energy. This illustrates that the reduction in brown consumption is significantly larger even than the reduction in cost, and that it is significant even when energy is cheap (β is small).

Next, we consider Figure 5, which illustrates the differences between LOCAL and GLB as a function of the capacity of renewable energy. Interestingly, Figure 5 highlights that when there is little capacity of renewables, both GLB and LOCAL can take advantage of it, but that as the capacity of renewables increases GLB is much more efficient at using it. This is evident because of the significantly lower brown energy consumption of GLB that emerges at capacities ¿1.5. Figure 5 also illustrates that increased capacity provides significant reductions in both the total cost and the average delay under both GLB and LOCAL.

Finally, let us consider the effect that GLB has on capacity provisioning of renewable energy. Figure 6 illustrates the capacity of solar and wind necessary to achieve certain brown energy reduction targets. We see in (a) that, under LOCAL, the capacity of renewables necessary to hit an 85% reduction of brown energy is extreme. However, in (b) we see that the use of GLB can provide significant reductions due to its more efficient use of the renewable capacity. If energy is expensive ($\beta = 10$), then the required capacity is nearly viable.

The discussion above highlights that GLB can be extremely useful for the adoption of renewable energy into internet-scale systems. There are certainly challenges that remain for the design of GLB. One particularly important challenge is that of adjusting the routing and capacity decisions at a fast enough time scale to react to the availability of renewable energy. In particular, the plots we have described so far use a control time-scale of 1 hour. If this control time-scale is slower, or if information about the availability of renewable energy is stale, then the benefits provided by GLB degrade, as illustrated in Figure 6(c). This highlights the importance of reducing the energy and wear-and-tear costs of switching servers into and out of active mode, since it is the magnitude of these costs that most often limits the control time-scale.
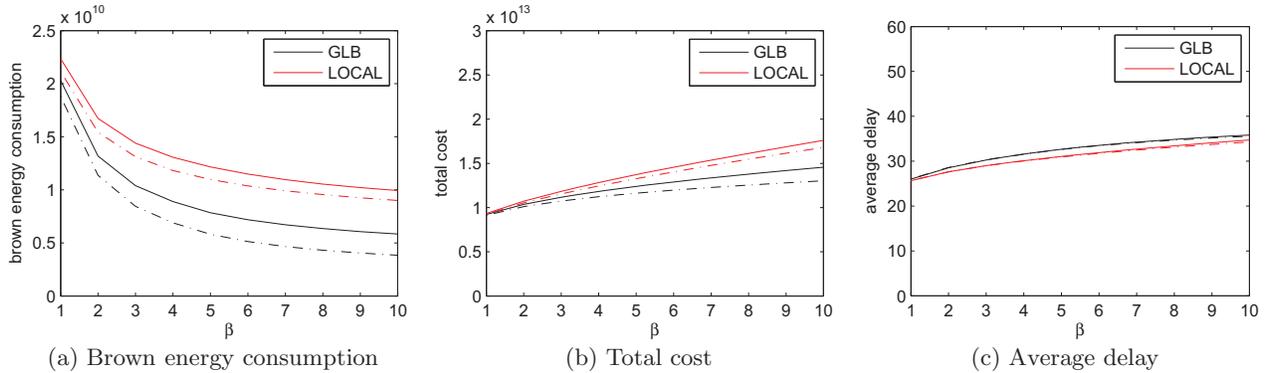
(a) Brown energy consumption   (b) Total cost   (c) Average delay

Figure 4: Comparison of GLB and LOCAL as a function of β. The renewable capacity is 2. The dashed line is the same setting except with storage.

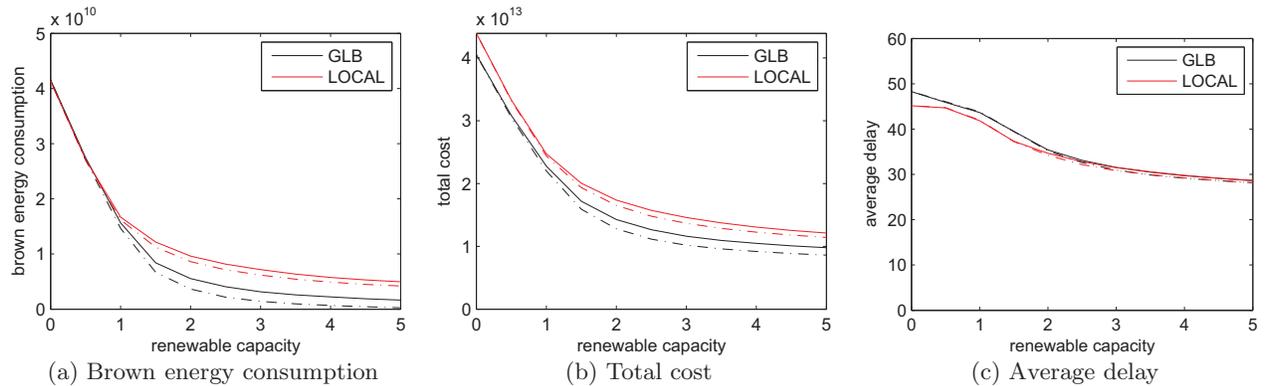(a) Brown energy consumption   (b) Total cost   (c) Average delay

Figure 5: Comparison of GLB and LOCAL with different renewable capacities and $\beta = 10$. The dashed line is the same setting except with storage.

3.2 The role of storage

In addition to GLB, another important tool to aid the incorporation of renewable energy into internet-scale systems is storage, e.g. UPSs. In this paper, we limit our consideration to small-scale storage, which is only able to power the data center over the time-scale of a few minutes. This is due to the opportunity provided by the UPSs already available in data centers today, and due to the fact that the cost of large scale storage is prohibitive at this point.

Throughout Figures 4, 5, 6(a), 6(b), and 7, the impact of storage can be seen through the difference between the corresponding dashed and solid lines. A few trends that are evident in these plots are the following: (i) storage becomes more valuable with either higher capacities of renewables, i.e., ¿ 1.5, or larger β; and (ii) storage plays a more significant role under GLB than under LOCAL. Both points are clearly illustrated in Figures 4 and 5, and Figure 5 also highlights that storage allows brown energy consumption to be almost completely eliminated when using GLB.

3.3 The optimal renewable portfolio

We now move to the question of what mix of solar and wind is most effective. A priori, it seems that solar may be the most effective, since the peak of solar availability is closely aligned with that of the data center workload. However, the fact that solar is not available during the night is a significant drawback. Also, once GLB is used, it becomes possible to aggregate wind availability across geographical locations. This provides significant benefits because wind availability

is not correlated across large geographical distances, and so when aggregated, the availability smoothes considerably, as illustrated in Figure 2(a). As a result, it seems that wind should be quite valuable to internet-scale systems.

Our results lend support to the discussion above. As illustrated in Figures 6 and 7, the optimal renewable portfolio for brown energy reduction is around 80% wind, and extra solar capacity provides little benefit beyond that point, in keeping with the findings of [5]. More specifically, Figure 7 shows the brown energy usage as a function of the fraction of energy coming from wind, for three values of total generating capacity, c. Keeping c fixed implicitly assumes that the cost per kWh of solar and wind installation are equal.

Interestingly, this optimal portfolio is robust to many factors including storage, renewable capacity, and even whether the system seeks to optimize brown energy consumption, total cost, or average delay. This last point is important, since it highlights that the system operators' goal is aligned with both the users' experience and society's interest. However, the optimal portfolio is affected by the workload, specifically, the peak-to-mean ratio. For large diurnal peak-to-mean ratios the optimal portfolio can be expected to use a higher percentage of solar. Another interesting direction is to include seasonal effect [5], e.g., during summer the solar supply is higher and wind supply is lower, and the cooling power demand is higher.

4. CONCLUSIONS

This paper performs a trace-based study to investigate the
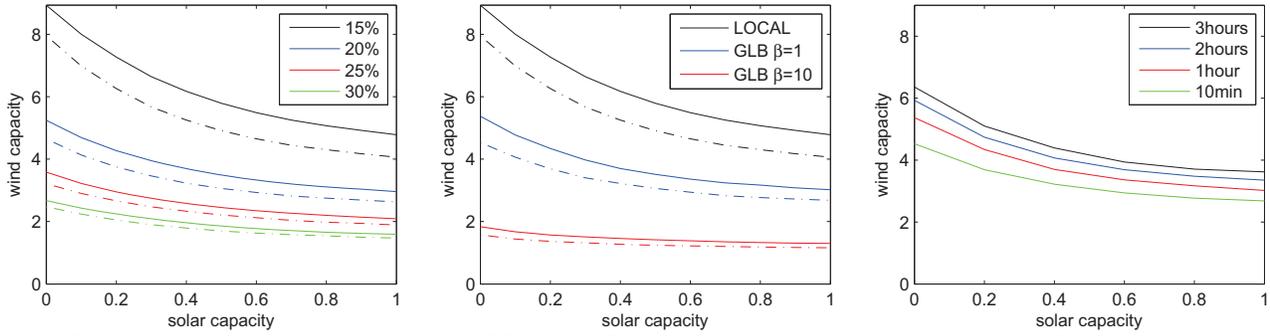
(a) LOCAL for different brown energy targets

(b) LOCAL and GLB for 15% brown energy target

(c) GLB with different control timescales and 15% brown energy target

Figure 6: Wind and solar capacity required to reduce brown energy usage to 15–30% of the baseline, as a function of solar capacity. The dashed line is the same setting except with storage.
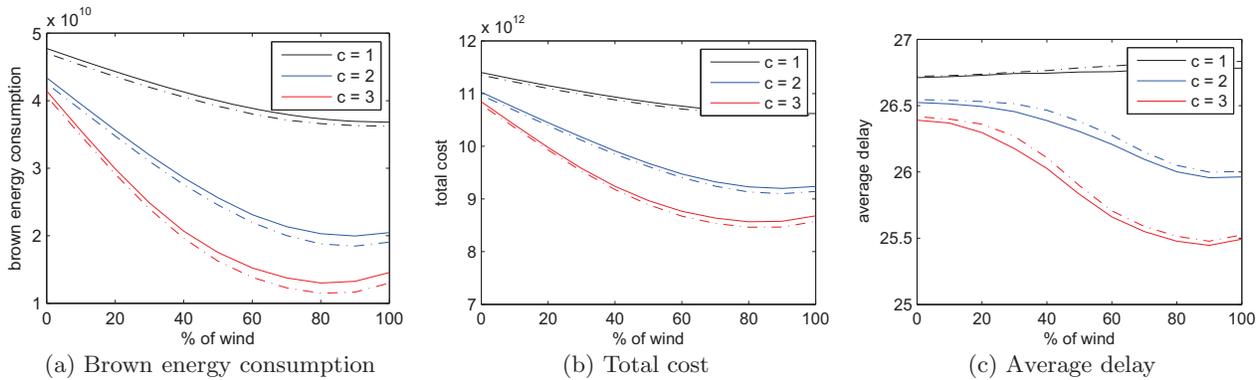
(a) Brown energy consumption

(b) Total cost

(c) Average delay

Figure 7: Impact of the mix of renewable energy used under GLB with $\beta = 1$ when the total renewable capacity is $c = 1, 2, 3$. The dashed line is the same setting except with storage.

feasibility of powering internet-scale systems using (nearly) entirely renewable energy. Our results highlight the vital role of geographical load balancing in achieving this goal.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] The cost of latency. *James Hamilton's Blog*, 2009.
[2] C. L. Archer and M. Z. Jacobson. Supplying baseload power and reducing transmission requirements by interconnecting wind farms. *J. Applied Meteorology and Climatology*, 46:1701–1717, Nov. 2007.
[3] D. Gmach, Y. Chen, A. Shah, J. Rolia, C. Bash, T. Christian, and R. Sharma. Profiling sustainability of data centers. In *Proc. ISSST*, 2010.
[4] D. Gmach, J. Rolia, C. Bash, Y. Chen, T. Christian, and A. Shah. Capacity planning and power management to exploit sustainable energy. In *Proc. of CNSM*, 2010.
[5] D. Heide, L. von Bremen, M. Greiner, C. Hoffmann, M. Speckmann, and S. Bofinger. Advances in solar thermal electricity technology. *Renewable Energy*, 35(11):2483–2489, 2010.
[6] http://rredc.nrel.gov. 2010.
[7] http://wind.nrel.gov. 2010.
[8] http://www.eia.doe.gov. 2010.
[9] M. LaMonica. Google data center to get boost from wind farm. *CNET News*, 21 April 2011.
[10] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. In *Proc. of INFOCOM*, 2011.
[11] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew. Greening geographic load balancing. In *Proc. ACM Sigmetrics*, 2011.
[12] C. Miller. Solar-powered data centers. *Datacenter Knowledge*, 13 July 2010.
[13] R. Miller. Google data center FAQ. *Datacenter Knowledge*, 27 March 2008.
[14] R. Miller. Facebook installs solar panels at new data center. *Datacenter Knowledge*, 16 April 2011.
[15] D. Mills. Advances in solar thermal electricity technology. *solar Energy*, 76:19–31, 2004.
[16] K. K. Nguyen, M. Cheriet, M. Lemay, B. St. Arnaud, V. Reijs, A. Mackarel, P. Minoves, A. Pastrama, and W. Van Heddeghem. Renewable energy provisioning for ICT services in a future internet. In *Future Internet Assembly. LNCS 6656*, pages 419–429. 2011.
[17] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. In *ACM SIGCOMM*, Aug. 2009.