Impact of workload and renewable prediction on the value of geographical workload management.*

Zahra Abbasi, Madhurima Pore, and Sandeep Gupta

Impact Laboratory,
School of Computing, Informatics and Decision Systems Engineering,
Arizona State University,
Firstname.Lastname@asu.edu

Abstract. There has been increasing demand for energy sustainable and low-cost operation in cloud computing. This paper proposes dynamic Geographical Load Balancing and energy buffering management (GLB) to achieve these goals which (i) shifts workload (particularly peak workload demand) toward Data centers that offer low utility rate or green energy at a time, and (ii) banks excess green and low-cost energy to shift peak workload demand away from high utility rate. Such a scheme needs to be aware of the workload intensity and the available renewable power of the cloud in future (over a relatively long prediction window such as a day). Existing solutions mainly focus on developing algorithms and demonstrating the cost efficiency of GLB, disregarding the prediction accuracy of the workload and the renewable power. However, erroneous information decreases the efficiency of GLB. This paper studies the performance of the online GLB solution when using time-series based prediction techniques (e.g., ARIMA) for the workload and the renewable power (i.e., solar and wind). The results of the simulation study using realistic traces highlight that GLB with and without prediction error is effective in reducing average energy cost and increasing sustainability of data centers. Further, GLB is shown to be significantly effective in shaving peak power draw from the grid (e.g., reducing peak power upto 100%), however the erroneous information due to the prediction error adversely affects its performance. Furthermore, the simulation study indicates that the optimal mix of the renewable power (i.e., wind and solar) to be leveraged by GLB, is achieved when data centers are powered from both the solar and the wind power.

Keywords: Cloud computing, Data Centers, Workload Prediction, Renewable power, electricity cost, Energy Storage, Energy Management

1 Introduction

Data center power consumption is raising concerns to both operators and society due to its huge electricity cost, scalability and detrimental impact on the environment. Particularly, there has been increasing push toward using renewable power in data centers from environmental activist [1, 2]. Large-scale Internet service providers, such as

*This work has been partly funded by NSF CRI grant #0855527, CNS grant #0834797, CNS grant #1218505 and Intel Corp.

Google, Microsoft and Yahoo! and other modern data centers have begun to partially power their data centers using on-site and other offline forms of the renewable energy resources [3, 4].

Recent works propose that dynamic Global/Geographical Workload Balancing (GLB) can potentially be a significant aid in maximizing renewable energy utilization and reducing energy cost without need for large scale Energy Storage Devices (ESDs) [5–12]. The idea is to leverage spatio-temporal variation of the workload, the renewable power and the electricity power to match the demand with low-cost and green power supply. However, the potential benefits of such a scheme have been a point of debate in the community [10, 11, 13]. On one hand, it is clear that, due to spatio-temporal variation of power supply and demand across data centers, GLB creates many energy management possibilities to lower electricity price, lower energy consumption, and efficiently manage the renewables [5–12]. On the other hand, there are also significant costs for its implementation in practice. These costs come in terms of the engineering challenges in implementing and designing efficient and automated algorithms. Fortunately, *cloud computing*, facilitating a dynamic, demand-driven allocation of computation, allows workload distribution across data centers. Additionally, the algorithmic challenges, in terms of time-efficiency, and making online decisions on workload distributions and energy buffering without requiring the knowledge of the future workload has been studied in literature [8, 12, 14]. Despite this body of the work, the question of characterizing the potential benefits of GLB has still not been properly addressed. Particularly, the proposed online algorithms, are based on simplification assumptions on the data center workload type [12], energy supply models and energy management objectives [8, 12, 14]. Also, it is shown that the efficiency of some of online solutions compared to the optimal offline solution depends on the prediction window length [13]. However, the existing solutions mainly rely on the predictability of data and ignore the possible impact of the prediction error on the performance of the algorithm. The predictability of input data is partially true, since some of the information, i.e., workload, electricity price, solar energy are shown to have nice cyclic behavior, and are thusly predictable. However, wind energy does not exhibit cyclic behavior and is thusly hard to predict. Also, generally, the prediction accuracy decreases with increasing the prediction window length. This raises concern that prediction error might be a significant downgrading factor on the efficiency of the global workload management.

Further, recent literature propose to utilize ESDs to shave the peak power draw from the grid, as it significantly contributes in the total electricity price of data centers [15, 16]. The idea is to use ESDs to smooth out the power draw from the grid. In this paper, we propose a comprehensive solution to integrate GLB with energy buffering management in order to shave peak power draw from the grid. Spatio-temporal variation of workload allows GLB to further smoothen the power draw from the grid. This necessitates a online workload management scheme with forecasting knowledge over a relatively long time in future [17].

Furthermore, recent works demonstrate that the wind energy is more valuable than the solar energy for Internet-scale systems, to be leveraged by global workload management [10, 18]. The reason is that the wind energy has little correlation across locations, and is available during both night and day. This suggests that the optimal renewable en-

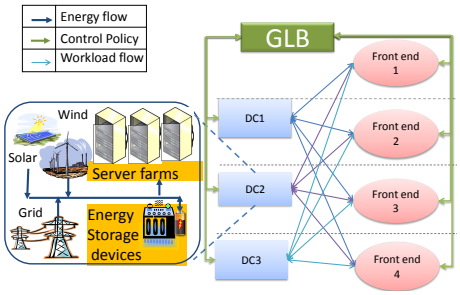

Fig. 1. GLB system model

ergy mix of the wind and the solar leveraged by GLB is dominated by the wind power. However, such a result is only true if we assume that the wind and the solar energy are perfectly predictable. We study such an effect in this paper.

In summary our main contribution is to study a workload management that coordinately manage the workload, energy buffering, and peak power draw of a cloud though an analytical model and illustrate its efficiency in presence of erroneous information of the workload, and the available renewable energy due to their prediction error. Our trace-based simulation study show that:

- GLB with and without prediction error yields higher renewable energy utilization (up to 20% for GLB with zero prediction error and up to 10% for GLB with prediction error) and consequently lower energy cost (up to 60% for GLB with zero prediction error and up to 50% for GLB with prediction error) compared to a conventional performance oriented load balancing scheme.
- The optimal renewable energy mix of the solar and the wind energy for GLB: (i) to increase the renewable utilization is achieved when both the wind and the solar energy contribute in the total renewable energy, (ii) to increase the GLB cost saving by reducing the impact of prediction error is achieved when there is more solar energy than wind energy.
- GLB when using forecast data increases the peak power draw from AC up to two times compared to when it accesses accurate data. These results are *pessimistic* since the values are compared to a baseline load balancing scheme which (i) manages local ESDs within data centers to shave peak power draw, and (ii) relies on the perfect knowledge of the workload and the renewables.

2 GLB system model and formal definition

Our simulation experiments combine analytical models with real traces for workload and renewable availability, to allow controlled experimentation but provide realistic findings. We now explain the system model and formulation of online GLB which extends the model presented in [6] to include energy buffering and to account the electricity cost per peak power draw from the grid.

GLB for interactive jobs can be generally modeled as a network flow optimization model on a bipartite graph (see Fig. 1). End users' requests arrive from $|A|$ geographically distributed front-ends (i.e., the sources) where $A = \{a_1 \dots a_j \dots a_{|A|}\}$ denotes the set of front-ends. The geographical front-ends may be network prefixes, or even geographic groupings (states and cities).

The workload must be distributed among the $|S|$ available data centers in the cloud (i.e., sink), where $S = \{s_i\}, i = 1, \dots, |S|$ denotes the set of available data centers in the cloud, and each s_i represents the total number of available servers in data center i. Also data centers may be provided with an energy storage of limited size, B^{size}, to smoothen renewable as much as possible.

There are many possible energy optimizations that can be developed considering various combination of factors such as workload split of data centers, power state of servers, migration of user state data, energy buffering level and performance level of applications. For simplicity, we only focus on the workload split, a two power state for servers (active and off), renewable harvesting and energy buffering. The goal is to perform workload consolidation over minimal number of servers in the most cost-efficient data center at a time. Extra servers are assumed to be turned off. GLB performs the optimization in a time-stepped system where the time is discretized into intervals, slots, and over a window of time intervals, denoted by T. We consider one hour slots for a window of a day. At the start we present mathematical modeling of GLB consisting energy consumption, cost, performance, workload, ESDs, and renewable energy models.

Table 1. Symbols and Definitions

Symbol	Definition
t	slot index
i	index of data centers
j	index of areas
S	set of data centers
A	set of front ends (areas)
T	Prediction window length
τ	length of slots (in second)
$p_{i,t}^{ESD}$	charging (for a positive value) and discharging (for a negative value)
$p^{max,discharge}$	maximum discharging rate
$p^{max,charge}$	maximum charging rate
p^{AC}	power draw from grid
$p^{total}i,t$	total power consumption
$r_{i,t}^{total}$	average available green power
$r_{i,t}$	renewable harvesting

Symbol	Definition
γ_i	cost per charging/discharging
B_i^{size}	energy storage capacity
ρ_i	energy loss coefficient
$\alpha_{i,t}$	cost per average power draw from the grid
β_i	cost per peak power draw from the grid
η	carbon footprint cap
f_i	function of DC's power consumption
$y_{i,t}$	number of active servers
μ_i	service rate
$\lambda_{j,t}$	workload arrival rate
d_i'	data center delay
$d_{i,j,t}''$	delay between areas and data centers
d^{ref}	total reference delay
d'^{ref}	service delay

Performance modeling We assume that delay experienced by a user, denoted by d should not exceed a reference delay, denoted by d^{ref}. Total delay, d, consists of the service delay d', i.e., data center delay, and the network delay d'', i.e., the delay between the front-end and the data center. As a result, the delay can be written as $d = d' + d''$. We model data center as a M/M/n queuing system and use its result to model service delay. We assumed a time-varying delay between every front-end j and data center i to model the network delay.

Workload modeling We model workload through its statistical parameters, i.e., average arrival rate and service time over every slot. Let $\lambda_{j,t}$ denote the mean arrival rate from front-end j at time t. To provide realistic estimates, we use real-world traces to denote $\lambda_{j,t}$. We consider four front-ends, corresponding to four time-zones in USA. We generate the workload of each front-end using NASA workload Internet data center trace (July and August, 1995) [19], such that it is shifted in time to account for time zone of each front-end, and scaled proportionally to the number of Internet users in the corresponding area (see Fig 2).

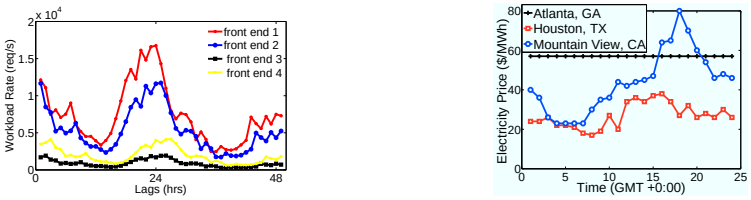

(a) (b)

Fig. 2. (a)NASA workload [19] scaled and shifted according to time zone for each front end, and (b) Hourly electricity price data for three locations on May 2nd, 2009 (data source [5]).

Energy consumption modeling We assume the total power consumption of the data center can be obtained by multiplying the total number of active servers (denoted by y) and power consumption value for each server (denoted by p). To keep the optimization framework linear, we set p as the power consumption of servers when they are utilized at their peak utilization.

ESD modeling We assume ESDs are associated with physical limitations on their size, B^{size}, measured in Joule, maximum charging discharging rate denoted by $p^{max,discharge}$, and $p^{max,charge}$, and energy efficiency (due to conversion) denoted by ρ. To account for the ESD cost we account the limits of the discharging cycles and state the capital cost of an ESD in terms of every discharge cycle. We denote such a cost as γ which incurs every discharging cycle. To model energy storage, we denote the energy storage level at time t by B_t with initial value B_0 and the charge/discharge at time t by p_i^{ESD}, where positive or negative values mean charge or discharge, respectively.

Renewable Energy Modeling We assume wind and solar energy as sources of renewable energy located on-site in a data center. To capture the availability of wind and solar energy, we use traces of [20]. We use wind speed and the rated power to calculate the wind power, and Global Horizontal Irradiance (GHI) and the ambient temperature to calculate the solar power using models described in [9]. The traces of three states for three days are illustrated in Fig 3. We scale the traces to study the efficiency of GLB under various configuration of the available renewable power.

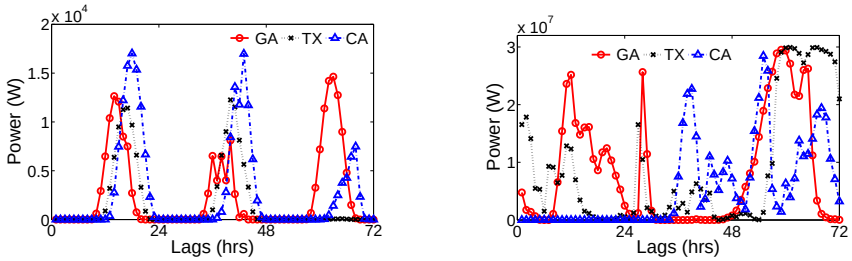

Fig. 3. Data traces showing variation of renewable energy generated at the three data center locations [20]: **(a)** Trace for solar energy, and **(b)** Trace for Wind energy.

Minimize $\sum_{i=1}^{|S|}\left(\sum_{t=1}^{T}\left(p_{i,t}^{AC}\tau\alpha_{i,t}+b_{i,t}\gamma_i\right)+\max_{1\le t\le T}(p_{i,t}^{AC}-p_0)^{+}\beta_i'\right),$

subject to:

[ESD const.] ,$\forall\, t,i:$ $r_{i,t}+p_{i,t}^{AC}=p_{i,t}^{total}+p_{i,t}^{ESD}$,
$B_{i,t+1}=\min(\eta_1(B_{i,t}+\rho_1 p_{i,t}^{ESD}\tau),B_i^{max})$,
$-b_{1,t}p_i^{max,discharge}\le p_{1,t}^{ESD}$,
$0\le b_{i,t}\le 1, r_{i,t}\le r_{i,t}^{total}$,
$p_{i,t}^{ESD}\le p_i^{max,discharge}$,
$-p_{i,t}^{ESD}\le p_1^{max,discharge}$.

[Capacity const.] ,$\forall\, t,i:$ $p_{i,t}^{total}=p_i y_{i,t},\ y_{i,t}\le s_i$.

[Service const.] ,$\forall\, t,j,:$ $\sum_{i=1}^{|S|}\lambda_{i,j,t}=\lambda_{j,t}$.

[Queuing const.] ,$\forall\, t,i:$ $y_{i,t}\mu_t\ge\sum_{j=1}^{|A|}\lambda_{i,j,t}$.

[Delay const.] ,$\forall\, t,i:$ $d'^{ref}\le\frac{1}{\mu}+\frac{1}{y_i\mu-\lambda_t}, d_{i,j,t}=d_i'^{\text{ref}}+d_{i,j,t}'',\ (d^{\text{ref}}-d_{i,j,t})\lambda_{i,j,t}\ge 0.$

Fig. 4. Linear Programming (LP) formulation of GLB problem.

GLB cost model We consider an electricity pricing model to account for both the electricity cost per average energy consumption, $\alpha_{i,t}$ (see Fig. 2(b)), and β_i per excess peak power draw from stipulated power (denoted by p_0) over a month. The later is used to penalize the peak power draw from the utility. Note that we consider a prediction window of length T for GLB problem that is less than a month. However, to use GLB to smooth the peak power, we incur a fraction of β, i.e., β'=Tβ/number of slots in a month per excess peak power during T to penalize peak power draw. A solution to this problem would specify, at each slot, how many servers in each data center should be assigned to the workload (i.e., $y_{i,t}$), what portion of each front-end's traffic should be assigned to which data center (i.e., $\lambda_{i,j,t}$), and how much is average power draw from AC, (i.e., $p_{i,t}^{AC}$), renewable (i.e., $r_{i,t}$), and energy storage (i.e., $p_{i,t}^{ESD}$, and $B_{i,t}$). We approximate $y_{i,t}$ as a real and solve this problem as a linear programming model. Cost minimization is subject to the following constraints:

- *ESD constraint* which assert the power demand and supply balance, energy level of ESDs over time which is affected by its charging/discharging (p^{ESD}), its energy efficiency, ρ, and its self-discharge ratio, η, and battery cost that is incurred per each discharging cycle using a linear approximation equation.
- *Service constraint* which asserts that all workload should be assigned to the data centers.
- *Capacity constraint* i.e., the number of assigned active servers in a data center should not exceed the available servers (denoted by s_i) in that data center. Further, each data center should supply power required for all of its active servers.
- *Queuing stability constraint* which asserts the M/M/n stability condition.
- *Delay constraint* i.e., the traffic of end users should be split among data centers whose network and service delay is less than the users' delay requirement.

3 GLB solution using prediction

In our solution, we evaluate the efficiency of the schemes by forecasting workload and renewable using time-series prediction methods (i.e., SARIMA and moving average) and the well known Rolling Horizon Control (RHC) technique. Consider a window of length T, RHC obtains GLB solutions at time t by solving the cost optimization over the window $(t,t+T)$, given the GLB solution at time $t-1$. The following sections discuss the workload and the renewable energy prediction (i.e., predicting $\lambda_{i,t}$, and $r_{i,t}^{total}$ over the window $t+T$) and their results. Note that the electricity price, i.e. α is usually known a day ahead [11], further electricity price, workload, and solar energy have daily variations. For that we choose the prediction window of a day (24 slots), and assume the actual electricity price is given for the entire window.

3.1 Workload prediction

We analyze NASA workload using different time series based prediction techniques and observe that Seasonal ARIMA (Auto Regressive Integrated Moving Average) captures the seasonal behavior of workload (see Fig. 2) with reasonable accuracy (compared to non-seasonal schemes). We use the workload data for the month of July of NASA trace for learning the SARIMA model, and use the trace of August in our numerical analysis. For every time slot, using workload information known up to that time slot, we forecast for the next 24 hours using the prediction model. For time slot of 1 hour we observer prediction error of 25.8% that increases up to 42.6% for higher lags as shown in Fig. 5(a).

3.2 Renewable power prediction

We use sample traces from the solar and wind energy sites in Georgia (GA), Texas (TX) and California (CA) for the months of February and March [20].

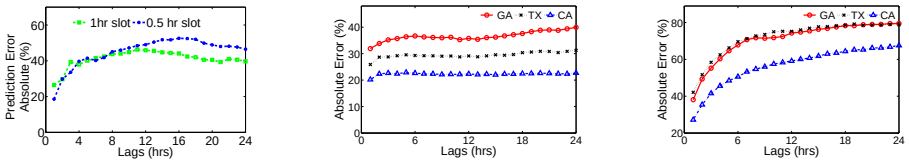

(a) (b) (c)

Fig. 5. Prediction error plots for different lags as obtained for 24 hours of forecast window **(a)** Workload, **(b)** Solar power, and **(c)** Wind power.

3.3 Solar Energy Prediction

The variation in solar power shows a daily seasonal behavior as seen in Fig. 3(a). This pattern is captured by SARIMA. For learning SARIMA model we use one month trace of February. The model is fitted by changing the parameters to minimize the error observed in the residual plots and the ACF and PACF (Auto Correlation and Partial Auto Correlation) plots. Predicted solar energy for one week is shown in Fig. 3. The error for different lags is almost constant starting with 23% for TX, and up to 25% for GA. The error obtained for different lags for prediction is shown in Fig. 5(b)

3.4 Wind Energy Prediction

The wind power shows a highly varying nature which depends on different factors, such as temperature, pressure, wind directions(see Fig. 3(b)). Hence prediction of wind using seasonal time series prediction is not so effective. We use the moving average technique to predict the wind energy. We observe that the average error for a given forecast window increases rapidly with increasing the prediction window length, starting with 30% for window of length one (hour) and increasing up to 80% for a window of length 24 (a day). Comparing Figs. 5(b) and (c), demonstrating the prediction error of solar and wind energy for different lags, we observe that the wind energy incurs much higher prediction error for lags greater than one than that of the solar energy.

4 Simulation study

4.1 Data center setup

We simulate a cloud consisting of three data centers as depicted in Fig. 1. The electricity pricing and physical characteristics, such as server power profiles, battery of data centers are set according to realistic data. To this end, we assume data centers are located at the following three locations: Atlanta, GA; Houston, TX; and Mountain View, CA, namely DC1, DC2 and DC3, respectively. These locations correspond to the location of three major Google data centers. We used the historical electricity prices for the above locations [5] (see Fig. 2(b)) as example in our setup. We perform experiments for both homogeneous and heterogeneous setting of data centers in terms of their servers' power efficiency as depicted in Table. 2. We assume each server can handle upto 2000 requests and that there is no network delay. The data sheet of Fload Lead Acidic (FLA)

batteries used in data centers is used for simulation study (see Table 3). We also assume each data center have 500 servers, and the workload intensity in our experiment is such that 600 active servers are needed on average. Our simulation environment is developed using MATLAB 2009. We use GNU Linear Programming Kit (GLPK) to solve GLB (i.e., optimization problem in Fig. 4).

Table 2. Data center characteristics

DC	Elec. price model	peak power(W)
DC1	Mountain View, CA.	395
DC2	Houston, TX	300
DC3	Atlanta, GA	450

Table 3. Specification of ESD parameters

Parameters	**FLA**
Capacity (KW)	115
Cost per discharge ($)	0.65
Cycle life of one cell (cycles)	1200
Discharge rate (W)	5387.5
Discharge-to-charge ratio	10
Efficiency (%)	80
Number of cells	53

4.2 Experiments performed

We design **Performance-oriented Load Balancing (PLB)** as a baseline scheme, which distributes workload equally across all data centers. For fair comparison, we assume PLB utilizes ESDs and determines number of active servers to optimize cost within each data center independently (i.e., solving problem in Fig. 4 for each data center). To evaluate the impact of prediction error on the efficiency of GLB we compare **GLB when using Perfect Prediction (GLB+PP)** with GLB when using our prediction technique (GLB+P). Note that GLB+PP accesses actual data over every prediction window of T. *Similarly, to characterize the maximum impact of the prediction error we assume that PLB accesses the actual data over every prediction window to perform energy buffering management in each individual data center*, in reality PLB also needs to predict the data, though. We evaluate the above schemes in terms of energy and ESD cost (i.e. the two first items in the objective function of Eq. 4), excess peak power draw from the grid, total energy consumption and renewable utilization with various cloud configuration (e.g., ESD size, and available renewable energy). to calculates peak power draw over a month, the stipulated power is set to 80KW, ($p_0 = 80kW$) a value that is 60% more than average power consumption of our simulated data center.

Available renewable energy We set the original renewable (wind+solar) traces such that the solar and the wind power contribute equally in the total renewable energy of data centers. Further, the renewable energy contributes 10% of the cloud total energy consumption. Then we scale the renewable traces from 1-16 as depicted in Fig. 6 to evaluate the efficiency of GLB with increasing availability of the renewable energy. We also set the ESD capacity to 12MJ for each data center, a value that is equivalent to the energy consumption of data centers for one minute when operating at their peak. Finally, we use a heterogeneous setting of data centers (see Table 2).

Fig. 6(a) shows that both GLB+PP and GLB+P outperform PLB in terms of energy and ESD cost. However, prediction error negatively affects the GLB cost saving, caus-

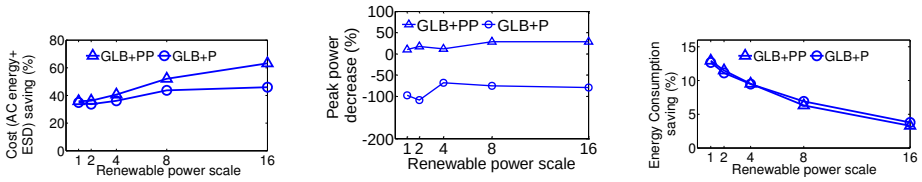

(a) (b) (c)

Fig. 6. Efficiency of GLB compared to PLB versus renewable power increase; **(a)** energy and ESD cost saving, **(b)** Peak power draw decrease, and **(c)** Total energy consumption decrease.

ing a decrease in cost saving from 60% for GLB+PP, down to 45% for GLB+P. Further, while the cost saving of GLB+PP increases with increase in the available renewable energy, this is not the case for GLB+P.

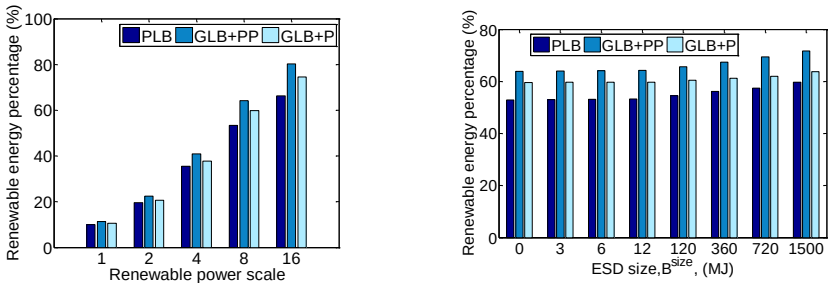

(a) (b)

Fig. 7. (a) Renewable energy percentage of total energy versus increase in available renewable energy, and **(b)** Renewable energy percentage of total energy versus increase in ESD size.

The results of Fig. 6(b) imply that the effect of prediction error is huge for peak power draw, as GLB+P increases the peak power draw from the grid up to 2 times compared to PLB that accesses a perfect prediction scheme. In reality, both PLB and GLB need to predict data. GLB+PP, however, significantly decreases the peak power compared to PLB. Particularly, it decreases the peak power (compared to PLB) from 10% up to 40% with increasing renewable energy.

Fig. 6(c), shows that both GLB+PP and GLB+P decrease the total energy consumption of the cloud compared to PLB, that is due to shifting workload to DC2 which has lesser peak power per server (see Table. 2). The figure shows that GLB+P has higher energy consumption saving than that of GLB+PP, this is because GLB+PP does not highly utilize DC2 to decrease its peak power. The energy saving decreases with increasing the available renewable, since GLB schemes shift the workload towards data centers that have more available renewable energy instead of shifting the workload to data centers that have more power efficient servers.

Finally, Fig. 7 (a) shows that both GLB+PP and GLB+P increase the utilization of the renewable energy compared to PLB.

ESD capacity In this experiment we scale the renewable power by the factor of eight, use heterogeneous setting of data centers (see Table 2), and vary ESD capacity from 0 up to 1500MJ, where 1500MJ is equivalent to the energy consumption of data centers for two hours when operating at their peak. Fig. 8(a) shows that both GLB schemes, i.e., with and without prediction error, induce lesser cost than PLB for zero size ESD capacity. The cost saving of GLB+P decreases with increasing ESD size, since PLB can leverage ESD to utilize low-utility rate electricity. This implies that by decreasing the opportunities for GLB to leverage spatio-temporal variation of renewable and low-cost electricity compared to workload management within data centers (in this case due to deploying large scale ESD in each data center), the impact of prediction error on downgrading the performance of GLB becomes higher.

Similar to Fig. 6(b), Fig. 8(b) shows that GLB+P increases the peak power up to 2 times compared to PLB. This is because PLB by using accurate data can shave the peak power using local ESD. However, the same figure shows that GLB+PP can decrease the peak power up to almost two times depending on the ESD size, compared to PLB. This result shows the value of GLB in shaving peak power draw from the grid when it uses an accurate prediction technique. Finally, Fig. 7 (b) shows that both GLB+PP and GLB+P increase the utilization of the renewable energy compared to PLB. However, the renewable utilization of GLB+PP increases with increasing ESD size, that is not the case for GLB+P.

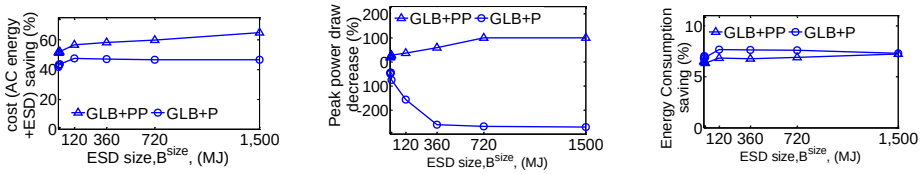

(a) (b) (c)

Fig. 8. Efficiency of GLB compared to PLB versus ESD capacity increase; **(a)** energy and ESD cost saving, **(b)** Peak power decrease, and **(c)** Total energy consumption decrease.

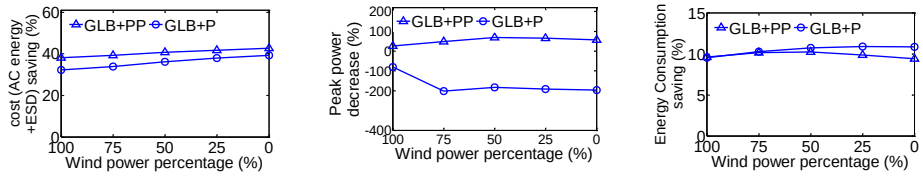

(a) (b) (c)

Fig. 9. Efficiency of GLB compared to PLB versus various mix of wind and solar; **(a)** Energy and ESD cost saving, **(b)** Peak power decrease, and **(c)** Total energy consumption decrease.

Wind and solar energy mix In this experiment, we fix the renewable scale to the factor of 8, ESD capacity to 360MJ (an energy storage for each data center to work for 30 minutes with server working at peak utilization), and use a heterogeneous setting of data centers. We vary the contribution of the wind power from 100% down to 0%.

The results shown in Fig. 9 indicate that the cost saving and energy saving of GLB schemes increase with increasing solar energy contribution. This means that PLB can not highly utilize excess solar energy with a relative high ESD size, since the solar energy is only available during the days. However GLB schemes, can leverage the spatio-temporal variation of the available renewable energy and efficiently utilize it. Interestingly, the cost saving of GLB+P increases with increasing the solar energy contribution, that is partly due to less prediction error of the solar energy compared to the wind energy (compare Fig. 5(b) and (c)). Finally, Fig. 10 (a) shows that the highest utilization of renewable for all schemes is achieved when there is a mix of both the wind and the solar (70-50% wind power), since with this mix, renewable power is almost available at all the time.

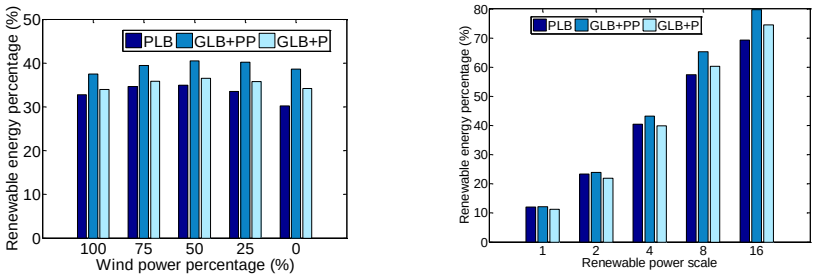

(a) (b)

Fig. 10. (a) Renewable energy percentage of total energy consumption versus various mix of wind and solar energy, and **(b)** Renewable energy percentage of total energy for homogeneous DCs versus increase in available renewable.

Homogeneous data centers We repeat the first experiment for the homogeneous setting of data centers where all data center consume 300W at peak per each server. Fig. 11 shows that for the homogeneous setting of data centers, GLB schemes can significantly save cost due to leveraging electricity price variation and available renewable across the cloud. Peak power draw plot (i.e., Fig.11(b)) has the same trend as that of (i.e., Fig.6(b)). Fig.11(c) shows that GLB schemes slightly (around 0.0005%) increase the energy consumption compared to PLB, that is because the GLB schemes increase ESD utilization and consequently its energy loss. Comparing Fig. 10(b), with Fig. 7(a) indicates that the renewable utilization trend of GLB+PP and PLB is the same. However, GLB+P does not always incur higher renewable utilization than that of PLB, confirming the previous results that when there is lower opportunity to leverage power and cost efficiency (in this case homogeneous setup of data centers), the impact of prediction error in decreasing the performance of GLB worsens.

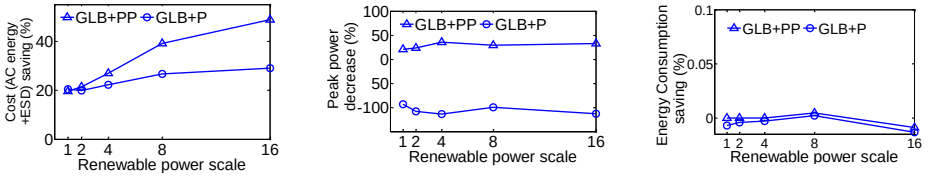

(a) (b) (c)

Fig. 11. Efficiency of GLB compared to PLB for homogeneous setting of data centers **(a)** energy and ESD cost saving, **(b)** Peak power draw decrease, and **(c)** Total energy consumption decrease.

4.3 Discussion

In summary, the above results show that, with the presence of erroneous information, the GLB scheme is still more effective in maximizing the utilization of renewable energy compared to the conventional performance oriented load balancing scheme (i.e., PLB). Particularly, it leverages spatio-temporal variation of power efficiency, power cost, and renewable in reducing energy cost, and increasing the renewable utilization without requiring a large-scale energy storage. However, pessimistically GLB with prediction error (GLB+P) increases the peak power drawn from AC compared to PLB with zero prediction error. Since peak power cost may be significant in some data centers, these results necessitate the need for developing: (i) online algorithms that can shave the peak power without having knowledge about the future, and (ii) prediction techniques to predict workload and renewable energy with very low prediction error.

The above results, however, heavily depend on the prediction accuracy. Although we do not claim superior performance of SARIMA over the other prediction techniques, it is a fairly stable and frequently used in the literature, particularly, for Internet workload prediction [7, 18, 21]. Further, we observe a prediction error that is very similar to what is reported by the previous works (for lag one prediction) [18, 22]. The prediction of renewable energy particularly for small window length and relatively short time intervals (e.g., one hour) can be improved using weather forecast data, however such approaches tend to be inaccurate at medium time scales (e.g., three hours to one week) [23].

5 Related Work

Similar to our results, related works highlight that global workload management can significantly reduce the electricity bill [5, 6, 8, 11, 12], and can potentially reduce the carbon footprint of data centers without requiring large-scale ESDs [7, 9, 10, 24, 25]. However, the focus of the above work are either demonstrating the efficiency of such a scheme using numerical study, or developing efficient algorithms. Particularly, Qureshi et al perform numerical analysis to show the efficiency of GLB on reducing the electricity price [11]. Le et al. demonstrate the carbon footprint reduction of the cloud using global workload management [7]. Akoush et al. and Stewart et al. show that workload management across data centers can increase the utilization of renewable [24, 25].

The algorithmic issues of GLB are studied in some recent literatures [5, 6, 8]. These literature assume a proactive and a time-stepped system for the scheme, where the workload intensity is assumed to be predicted ahead of time. However, the impact of the prediction error is not sufficiently studied. As such [18] proposes a cooling and renewable

aware workload management within data centers. The authors use a regression based predictor and k-nearest classification to predict workload and solar power, respectively. Similar to our result, they report an average error of 20% for both workload and solar power (for lag one). The authors show that the prediction error of data has negligible effect on the optimality of the solution. However, we show that when using a large prediction window, the impact of prediction error on shaving peak power is huge.

GLB when integrated with energy buffering management [14], workload migration overhead [6, 12], or server switching cost [13] can only be optimally solved using an offline algorithm. Accordingly, some online algorithms with guaranteed competitive ratio are proposed [12–14]. Particularly, Lin et al., derive a competitive ratio (i.e., online optimal solution over the offline optimal solution) and prove that the competitive ratio decreases with increasing the prediction window length. However, our results show that the prediction error increases with increasing the prediction window length (see Fig. 5). [14] proposes an online algorithm to utilize ESDs in order to leverage temporal variation of energy cost in data centers. However, more recent work propose to shave peak power draw using ESDs, since it is a big contributor of total electricity cost in data centers [15, 16]. Shaving peak power draw from grid requires an online workload management algorithm which uses forecast data for a relatively long window length [17]. Our results show that GLB is very effective in shaving peak power draw when it uses accurate data. However, erroneous information adversely affects its performance.

6 Conclusion

This paper proposes global workload management (GLB) integrated with energy buffering to increase renewable utilization, reduce average energy and peak power cost across data centers. The paper studies online GLB solution using SARIMA technique to predict the workload and the renewable energy over a window (i.e., 24 hours). The trace-based simulation study, shows that GLB with and without prediction error (i.e., GLB+P and GLB+PP) outperform over the performance oriented-load balancing, PLB, in increasing renewable energy utilization and reducing energy cost. The results are pessimistic since PLB is assumed to access accurate data as opposed to GLB+P. Results also highlight that GLB needs an accurate predictor to enable shaving peak power draw from the grid. Such prediction accuracy is not achieved using SARIMA based prediction, necessitating future work to design accurate prediction techniques.

References

1. H. Clancy, “Facebook becomes renewable energy activist website,” http://www.zdnet.com/blog/green/facebook-becomes-renewable-energy-activist/19656, cited Jan 2013.
2. greenpeace.org, “Greenpeace activists project messages on apple headquarters from supporters asking for cleaner cloud,” cited Jan 2013.
3. G. Demasi, “More renewable energy for our data centers,” cited Jan 2013. [Online]. Available: http://googleblog.blogspot.com.ar/2012/09/more-renewable-energy-for-our-data.html
4. R. Mcmillan, “Apple vows to build 100% renewable energy data center,” http://www.wired.com/wiredenterprise/2012/04/applerenewable/, cited Jan 2013.

5. L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing electricity cost: optimization of distributed Internet data centers in a multi-electricity-market environment," in *Proc. IEEE INFOCOM*, 2010, pp. 1–9.
6. Z. Abbasi, T. Mukherjee, G. Varsamopoulos, and S. K. S. Gupta, "DAHM: A green and dynamic web application hosting manager across geographically distributed data centers," *ACM Journal on emerging technology (JETC)*, 2012.
7. K. Le, O. Bilgir, R. Bianchini, M. Martonosi, and T. D. Nguyen, "Managing the cost, energy consumption, and carbon footprint of internet services," in *Proc. of ACM SIGMETRICS*, 2010, pp. 357–358.
8. Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Greening geographical load balancing," in *Proc. ACM SIGMETRICS, San Jose,USA*, june.
9. Y. Zhang, Y. Wang, and X. Wang, "Greenware: greening cloud-scale data centers to maximize the use of renewable energy," *Middleware 2011*, pp. 143–164, 2011.
10. Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Geographical load balancing with renewables," *ACM SIGMETRICS Perf. Eval. Rev.*, vol. 39, no. 3, pp. 62–66, 2011.
11. A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for Internet-scale systems," in *Proc. ACM SIGCOMM*, 2009, pp. 123–134.
12. N. Buchbinder, N. Jain, and I. Menache, "Online job-migration for reducing the electricity bill in the cloud," *NETWORKING 2011*, vol. 6640, pp. 172–185, 2011.
13. M. Lin, Z. Liu, A. Wierman, and L. L. H. Andrew, "Online algorithms for geographical load balancing," in *Proc. of IEEE IGCC*. IEEE, june 2012.
14. R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramanian, "Optimal power cost management using stored energy in data centers," in *Proc. of ACM SIGMETRICS*, 2011, pp. 221–232.
15. D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?" in *Proc. of the 12th ACM SIGMETRICS*, 2012, pp. 187–198.
16. V. Kontorinis, L. E. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, D. M. Tullsen, and T. S. Rosing, "Managing distributed UPS energy for effective power capping in data centers," in *Proc. of the 39th IEEE ISCA*, 2012, pp. 488–499.
17. A. Bar-Noy, M. P. Johnson, and O. Liu, "Peak shaving through resource buffering," in *Approximation and Online Algorithms*. Springer, 2009, pp. 147–159.
18. Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, "Renewable and cooling aware workload management for sustainable data centers," in *Proc. of the 12th ACM SIGMETRICS*, 2012, pp. 175–186.
19. [Online]. Available: http://ita.ee.lbl.gov/html/traces.html.
20. NREL, "Measurement and instrumentation data center MIDC," http://www.nrel.gov/midc/.
21. G. Jung, M. A. Hiltunen, K. R. Joshi, R. D. Schlichting, and C. Pu, "Mistral: Dynamically managing power, performance, and adaptation cost in cloud infrastructures," in *Proc. of 30th IEEE International Distributed Computing Systems (ICDCS)*.
22. Í. Goiri, R. Beauchea, K. Le, T. D. Nguyen, M. E. Haque, J. Guitart, J. Torres, and R. Bianchini, "GreenSlot: scheduling energy consumption in green datacenters," in *Proc. of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011, pp. 20:1–20:11.
23. N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy, "Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems," in *Proc. of 7th Annual IEEE conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010*, pp. 1–9.
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*, 2009.
25. S. Akoush, R. Sohan, A. Rice, A. W. Moore, and A. Hopper, "Free lunch: exploiting renewable energy for computing," in *Proceedings of HotOS*, 2011.