Optimal Energy Procurement for Geo-distributed Data Centers in Multi-timescale Electricity Markets

Tan N. Le[1,2], Jie Liang[1], Zhenhua Liu[1],
Ramesh K. Sitaraman[3,4], Jayakrishnan Nair[5], Bong Jun Choi[1,2]
Stony Brook Univ.[1], SUNY Korea[2], Univ. of Massachusetts at Amherst[3], Akamai Tech.[4], IIT Bombay[5]

ABSTRACT

Heavy power consumers, such as cloud providers and data center operators, can significantly benefit from multi-timescale electricity markets by purchasing some of the needed electricity ahead of time at cheaper rates. However, the energy procurement strategy for data centers in multi-timescale markets becomes a challenging problem when real world dynamics, such as the spatial diversity of data centers and the uncertainty of renewable energy, IT workload, and electricity price, are taken into account. In this paper, we develop energy procurement algorithms for geo-distributed data centers that utilize multi-timescale markets to minimize the electricity procurement cost. We propose two algorithms. The first algorithm provides provably optimal cost minimization while the other achieves near-optimal cost at a much lower computational cost. We empirically evaluate our energy procurement algorithms using real-world traces of renewable energy, electricity prices, and the workload demand. Our empirical evaluations show that our proposed energy procurement algorithms save up to 44% of the total cost compared to traditional algorithms that do not use multi-timescale electricity markets or geographical load balancing.

1. INTRODUCTION

Data centers are becoming the largest and the fastest growing consumers of electricity in the United States. It is reported that US data centers consumed 91 billion kilowatt-hours (kWh) in 2013, which is more than twice of the electricity consumed by households in New York City (see [15]). In the same report, the electricity consumption of the data centers is estimated to reach 140 billion kWh in 2020 due to the explosion of demand for cloud computing and other Internet-scale services.

Multi-timescale electricity markets have been proposed to improve the efficiency of electricity markets [4]. Multi-timescale electricity markets encompass both *forward* (futures) and *spot* (real-time) markets. While energy is procured at the time of consumption in a spot market, forward markets allow customers to buy electricity a day ahead or even several months ahead of when it is consumed. Forward electricity markets reduce the risk for both the supplier and consumer by reducing the quantity of energy being traded in the real-time spot markets [2]. Furthermore, purchasing electricity ahead of time can facilitate the expansion of renewable energy sources [6].

Utilizing multi-timescale markets has great potential for electricity cost savings of cloud providers who operate one or more data centers. There has been much recent work that exploits the variation of real-time electricity prices in the temporal and spatial dimensions to reduce the total electricity cost [12, 13]. Other papers exploit temporal variation in the real-time energy price and use energy storage to reduce the electricity costs [7, 17], i.e., the storage device is charged during the times when the electricity price is low and discharged when the price is high. However, while these works focus on traditional real-time markets, our paper studies the potential of using multi-timescale markets in the context of a cloud provider.

In particular, using forward markets to lower the electricity cost for a cloud provider is challenging for multiple reasons. The optimal amount of electricity that a cloud provider should purchase in advance for a particular location depends on the workload, the availability of onsite renewables, and the real-time electricity price at that location. The main challenge is that the future workloads, renewables, and real-time electricity prices are not perfectly predictable and are subject to significant forecasting errors. Note that if the cloud provider is too conservative and buys too little from the forward market, any shortfall in electricity would need to be covered by purchasing it from the more expensive[1] real-time market. Likewise, if the cloud provider is too aggressive and buys too much from the forward market, any excess in electricity will go to waste. In addition, the ability of a cloud provider to move the load from one data center to another, possibly incurring a performance penalty that we characterize as the "delay cost", adds an additional level of complexity that needs to be optimized. In this work, we provide an optimization framework for tackling the aforementioned challenges. There are a few recent papers that consider forward markets; these papers deal with the financial risk of a single data center arising from the uncertainties in electricity prices and workload [14, 16]. Geographical load balancing systems with both day-head market and real-time market have been studied in a recent publication [5]. However, the proposed solution is somewhat restrictive to particular distributions to facilitate stochastic optimization

[1]In some cases, the prices in the forward markets might be (on average) higher than real-time prices. If so, instead of saving electricity expenditure, the cloud provider can participate in forward markets to reduce cost variations. Our model can be extended to handle either case.

GreenMetrics '17 IL, USA

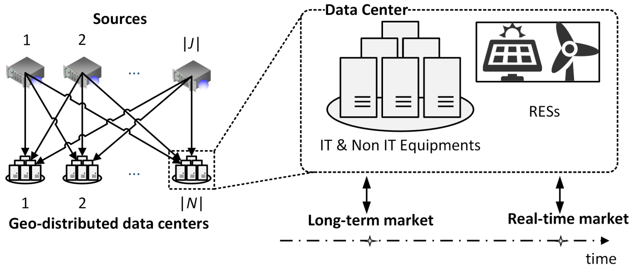

Figure 1: Geo-distributed data centers in long-term and real-time markets.

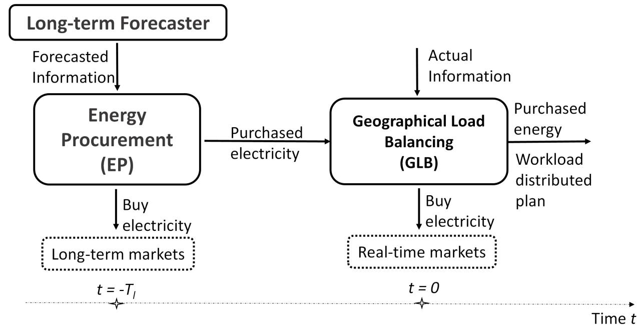

Figure 2: Energy Procurement System (EPS) Architecture for geo-distributed data centers.

and does not provide any optimality guarantee.

Our contributions are two-fold.

(1) Optimal algorithm development. We develop two algorithms for a cloud provider with geo-distributed data centers to buy electricity in multi-timescale markets: one algorithm provides optimality guarantees, while the other is simpler in that it uses limited predictions but achieves near-optimal performance. To develop the energy procurement system, we first model the problem of procuring electricity for geo-distributed data centers in multi-timescale markets in Section 2. The system model is general and applicable to any global cloud provider with access to multi-timescale electricity markets. We focus on two-timescale markets that consist of one long-term market and one real-time market. However, our model and algorithms can be extended to handle multiple markets at various timescales. We present the characteristics of the objective functions and the optimal solution in Section 3, which forms the theoretical basis for our algorithm design. The two algorithms that we design, prediction based algorithm (PA) and stochastic gradient based algorithm (SGA), are described in Section 4.

(2) Empirical evaluation. We carry out a detailed empirical evaluation of our proposed energy procurement systems using real world traces. In Section 5, we demonstrate that SGA can converge to the optimal solution in a small number of iterations. Moreover, we show that PA, our heuristic algorithm, surprisingly achieves a near-optimal solution. This is partially because the real-time optimization takes into consideration the trade-off between energy cost and delay cost, and is able to compensate for some prediction errors by redirecting workloads. The proposed energy procurement systems are compared with other comparable energy procurement strategies to highlight their benefits.

2. MODEL

2.1 System model

Two-timescale markets. A service provider operating geo-distributed data centers can purchase electricity in two markets – a long-term market and a real-time market. The electricity consumed at time $t = 0$ must be procured from the real-time market at t=0 and/or from the long term market *ahead of time* at $t = -T_l$.

Geo-distributed data centers. We consider a set N of geo-distributed data centers serving workload demands from a set J of sources as illustrated in Figure 1. The workload demand from each source is split between the $|N|$ data centers. Here, a source can represent the aggregate demand from a group of local users, such as users of a particular city, ISP, or geographical region. Each data center has access to renewable energy sources. Further, each data center participates in a (local) long-term electricity market and a (local) real-time electricity market. In other words, each data center i can buy electricity ahead of time in its long-term market, and can also buy additional electricity in its real-time market if necessary.

Energy procurement system (EPS). Our proposed energy procurement system for geo-distributed data centers is depicted in Figure 2. There are three main components, namely, the long-term forecaster, the energy procurement (EP) in long-term markets and the geographical load balancing (GLB). The long-term forecaster provides the forecasted information for the energy procurement. The forecasted information includes the predicted values and the prediction error distributions of IT workload, renewable energy generations, and electricity prices. The EP component procures electricity for each data center in the corresponding long-term markets (at time $t = -T_l$) based on the electricity prices in the long-term markets and *forecasts* of real-time prices, workload, and renewable generation. The GLB component (at time $t = 0$) distributes (routes) the realized workload from sources to data centers, provisions the required computing capacity at each data center, and procures additional electricity as needed in the real-time markets.

Data center. Let M_i denote the number of servers in data center i. The number of active servers at real-time (time $t = 0$) is denoted by m_i, which is a control parameter. In practice, there can be more than a hundred thousand servers in a single data center. Thus, we treat m_i as a continuous number satisfying $0 \leq m_i \leq M_i$.

At time $t = 0$, the power consumption of data center i is denoted by d_i^r. In general, the power consumption of data center i is dependent on the number of active servers m_i and the workload arrival λ_i. For simplicity, we assume that $d_i^r = m_i$, which implies that the power consumption is proportional to the number of active servers, and is independent of the workload λ_i.[2]

Workload. Workload demand from source j in real-time ($t = 0$) is denoted as L_j^r. We assume that the exact realization of the random vector $\boldsymbol{L}^r = (L_j^r, j \in J)$ is known to the cloud provider at time $t = 0$, and is an input to GLB. Let λ_{ij} denote the distributed workload arrival from

[2]The proportionality constant relating the number of active servers and the power consumption is taken to be 1 without loss of generality. Also, our analysis can be easily generalised to the case $d_i^r = O_i(m_i, \lambda_i)$, where the function O_i is continuously differentiable, convex, and non-decreasing in each coordinate.

source j to data center i at time $t = 0$ (set by GLB). Thus, $L^r_j = \sum_{i\in N} \lambda_{ij} \quad (j \in J)$, and $\lambda_i = \sum_{j\in J} \lambda_{ij} \quad (i \in N)$.

Renewable energy. Data centers can utilize their integrated renewable energy sources. Let w^r_i denote the renewable energy generation at data center i in real-time ($t = 0$). We assume that the exact realization of the random vector $\boldsymbol{w}^r = (w^r_i, i \in N)$ is known at time $t = 0$, and is an input to GLB.

Electricity price. For each data center, the cloud provider can purchase electricity at time $t = -T_l$ in the local long-term market and then purchase any additional electricity needed in the local real-time market at time $t = 0$. For data center i, let p^l_i denote the long-term price for 1 unit of electricity, and p^r_i denote the real-time price for 1 unit of electricity. We assume that $\boldsymbol{p}^l = (p^l_i, i \in N)$ is fixed (or equivalently, is known at the time of the long-term procurement), and $\boldsymbol{p}^r = (p^r_i, i \in N)$ is a random vector whose exact value is known at time $t = 0$ and is an input to GLB.

2.2 Cost model

The total cost of operating geo-distributed data centers is composed of a delay cost and an energy cost.

Delay cost. The delay cost represents the monetary cost incurred due to the delay experienced by the sources. We model the delay cost $h_{ij}(m_i, \lambda_i)$ of routing and processing each unit of workload from source j to data center i as follows.

$$h_{ij}(m_i, \lambda_i) = \beta\left(\frac{1}{\mu_i - \lambda_i/m_i} + \pi_{ij}\right) \qquad (\lambda_i < m_i\mu_i) \quad (1)$$

Here, the parameter β weighs the delay relative to the energy cost. The first term above captures queuing delay at delay center i, which is based on the well-known mean delay formula for the M/GI/1 processor sharing queue; μ_i is the service rate of servers in data center i. The second term captures the network delay from source j to data center i. Note that for stability, we need that $\lambda_i < m_i\mu_i$. Our delay cost model assumes a linear relationship between delay and its associated monetary cost, as is suggested in [3].

Energy cost. Let q^l_i and q^r_i respectively denote the amount of electricity purchased in the long-term market and the real-time market by data center i. Here, we require that sufficient electricity is procured to process the workload routed to each data center as

$$q^r_i + w^r_i + q^l_i \geq d^r_i = m_i \quad (i \in N).$$

The electricity bills of data center i in the long-term market and the real-time market are respectively computed as

$$R^l_i(q^l_i) = p^l_i q^l_i \qquad i \in N,$$
$$R^r_i(q^r_i) = p^r_i q^r_i \qquad i \in N.$$

2.3 Formulation of optimal energy procurement in multi-timescale markets

Recall that the total cost of operating geo-distributed data centers in our two-timescale market setting is the sum of the energy cost and the delay cost, given by

$$F = \sum_{i\in N} R^l_i(q^l_i) + \sum_{i\in N} R^r_i(q^r_i) + \sum_{i\in N, j\in J} \lambda_{ij} h_{ij}(m_i, \lambda_i).$$

We seek to minimize $\mathbb{E}[F]$ subject to the aforementioned constraints. Note that this optimization is performed on two timescales, with different sets of information available at each. The EP component optimizes the long-term procurements $\boldsymbol{q}^l = (q^l_i, i \in N)$ given only distributional information of the real-time workload $\boldsymbol{L}^r$, the real-time renewable generation $\boldsymbol{w}^r$, and the real-time electricity prices $\boldsymbol{p}^r$. The GLB component optimizes the workload routing $\boldsymbol{\lambda} = (\lambda_{ij}, i \in N, j \in J)$, the number of active servers $\boldsymbol{m} = (m_i, i \in N)$ at the data centers, and the real-time procurements $\boldsymbol{q}^r = (q^r_i, i \in N)$ given the prior long-term procurements $\boldsymbol{q}^l$, and the exact realization of $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$. Below, we first formalize the real-time optimization, followed by the long-term optimization.

Geographical load balancing in real-time markets. Note that in real-time, GLB optimizes the real-time procurements $\boldsymbol{q}^r$, the numbers of active servers $\boldsymbol{m}$, and the workload routing $\boldsymbol{\lambda}$, given the long-term procurements $\boldsymbol{q}^l$ and the realization of the random vector $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$. The total cost as seen by GLB is

$$F^r(\boldsymbol{q}^r, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{p}^r) := \sum_{i\in N} R^r_i(q^r_i) + \sum_{i\in N, j\in J} \lambda_{ij} h_{ij}(m_i, \lambda_i).$$

Thus, the real-time optimization is defined as follows.

$$\textbf{GLB-RT:} \quad \min_{\boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{q}^r} F^r(\boldsymbol{q}^r, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{p}^r)$$

$$\text{s.t.} \quad \lambda_{ij} \geq 0 \quad \forall i \in N, j \in J \qquad (2a)$$

$$\sum_{i\in N} \lambda_{ij} = L^r_j \quad \forall j \in J \qquad (2b)$$

$$\lambda_i \leq m_i\mu_i, \quad \forall i \in N \qquad (2c)$$

$$0 \leq m_i \leq M_i \quad \forall i \in N \qquad (2d)$$

$$q^r_i \geq 0, \quad \forall i \in N \qquad (2e)$$

$$m_i - q^r_i - w^r_i \leq q^l_i \quad \forall i \in N. \qquad (2f)$$

Since $p^r_i \geq 0$, it easily follows that any solution of the above optimization problem satisfies $q^r_i = [m_i - w^r_i - q^l_i]_+$, where $[x]_+ := \min\{0, x\}$. Thus, the real-time objective can be re-written as

$$\begin{aligned}\tilde{F}^r(\boldsymbol{q}^l, \boldsymbol{m}, \boldsymbol{\lambda}, \boldsymbol{p}^r, \boldsymbol{w}^r) = &\textstyle\sum_{i\in N} p^r_i [m_i - w^r_i - q^l_i]_+ \\ &+ \textstyle\sum_{i\in N, j\in J} \lambda_{ij} h_{ij}(m_i, \lambda_i).\end{aligned} \qquad (3)$$

GLB-RT problem is a convex optimization problem and hence can be solved efficiently using standard methods [11].

Energy procurement in long-term markets. At time $t = -T_l$, the cloud provider purchases electricity $\boldsymbol{q}^l$ in long-term markets that will be used at real-time. Note that optimization of the long-term procurements has to be performed based only on distributional information for the random vector $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$, and subject to the real-time optimization that will be subsequently performed based on the realization of the random vector $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$.

Let us denote the optimal value of the optimization GLB-RT by $F^{*r}(\boldsymbol{q}^l, \boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$. The long-term objective is thus defined as

$$F^l(\boldsymbol{q}^l) := \sum_{i\in N} R^l_i(q^l_i) + \mathbb{E}\left[F^{*r}(\boldsymbol{q}^l, \boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)\right].$$

Note that the above expectation is with respect to the random vector $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$. The long-term optimization problem is then given by:

$$\textbf{EP-LT:} \quad \min F^l(\boldsymbol{q}^l)$$
$$\text{s.t } \boldsymbol{q}^l \in \mathbb{R}^N_+.$$

The above optimization is more challenging than GLB-RT. In Section 3, we prove that EP-LT is a convex optimization and characterize the gradient of the objective function. These results are then used to arrive at a provably optimal stochastic gradient algorithm in Section 4.

3. CHARACTERIZING THE OPTIMA

In this section, we collect useful properties of the optimizations EP-LT and GLB-RT. These are important for understanding the behavior of the energy procurement system, and also for proving convergence of the stochastic gradient algorithm for EP-LT in Section 4.

Our first result is that EP-LT is indeed a convex optimization, which suggests that EP-LT is a tractable optimization.

THEOREM 1. *$F^l(\boldsymbol{q}^l)$ is convex over $\boldsymbol{q}^l \in \mathbb{R}_+^N$.*

We provide the proof of Theorem 1 in Appendix A.1 of the full version [9]. Next, we characterize the gradient of the EP-LT objective function as follows.

THEOREM 2. *The gradient of $F^l(\cdot)$ is characterised as follows.*

$$\begin{aligned}\frac{\partial F^l(\boldsymbol{q}^l)}{\partial q_i^l} &= p_i^l + \mathbb{E}\left[\frac{\partial F^{*r}(\boldsymbol{q}^l, \boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)}{\partial q_i^l}\right] \\ &= p_i^l - \mathbb{E}\left[\varrho_i(\boldsymbol{q}^l, \boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)\right],\end{aligned}$$

where $\varrho_i(\boldsymbol{q}^l, \boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$ is the unique Lagrange multiplier of GLB-RT corresponding to the constraint (2f).

Note that the first equality in the theorem statement asserts that the order of an expectation and a partial derivative can be interchanged. The second equality relates the partial derivative of F^{*r} with respect to q_i^l to a certain Lagrange multiplier of GLB-RT. We provide the proof of Theorem 2 in Appendix A.2 of the full version [9].

We note that Theorem 2 does not enable us to compute the gradient of the $F^l(\cdot)$ exactly. Indeed, the expectation the Lagrange multiplier ϱ_i with respect to $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$ would in general be analytically intractable. However, Theorem 2 does enable a noisy estimation of the gradient of the $F^l(\cdot)$ via Monte Carlo simulation as follows. Suppose we simulate a finite number, say $\mathbb{S}$, of samples from the distribution of $(\boldsymbol{p}^r, \boldsymbol{L}^r, \boldsymbol{w}^r)$. In practice, we can obtain these samples by using real-world traces . For each sample, the Lagrange multipliers $(\varrho_i, i \in N)$ can be computed efficiently by solving GLB-RT. By averaging the $\mathbb{S}$ instances of $(\varrho_i, i \in N)$ thus obtained, we get an unbiased estimate of the gradient of $F^l(\cdot)$. This, in turn, enables us to solve EP-LT using a stochastic gradient descent method; details follow in Section 4.

As there are two timescales in optimization, it is critical to investigate how EP-LT affects the operation of geographical load balancing in real-time. We start by answering the following question: how does the long-term procurement $\boldsymbol{q}^l$ impact the number of active servers m_i in data center i?. Formally, we have the following intuitive result:

LEMMA 3. *At any data center i, an optimal solution always utilizes the long term energy procurement q_i^l and renewable generation w_i^r as much as possible. It is simply represented by*

$$\begin{cases} m_i \geq w_i^r + q_i^l & \text{if } w_i^r + q_i^l < M_i \\ m_i = M_i & \text{if } w_i^r + q_i^l \geq M_i. \end{cases} \tag{5}$$

PROOF. Appendix A.3 of the full version [9]. □

The above lemma states that a data center i uses up the reserved electricity, including free renewable energy and prepurchased electricity, because doing so reduces the queueing delay.

4. ALGORITHM DESIGN

The energy procurement system needs algorithms for both energy procurement in long-term (EP-LT) and geographical load balancing in real-time (GLB-RT). GLB-RT is a convex optimization problem that can be solved efficiently in real-time by standard techniques [11]. Thus, we focus on designing algorithms for energy procurement in the long-term markets. Note that even though EP-LT is a convex optimization (see Theorem 1), neither the objective function nor its gradient admit a closed-form representation, which presents significant challenges.

4.1 Prediction based Algorithm (PA)

Prediction based algorithm (PA) relies on the mean values of renewable generation, workload, and electricity price. The predicted values $\hat{L}_j^r$, $\hat{w}_i^r$, and $\hat{p}_i^r$ are the estimates of the mean values of renewable generation, workload, and electricity price.

PA computes the long-term procurement $\boldsymbol{q}^l$ by solving EP-LT and GLB-RT at the same time, with the random variables w_i^r, L_j^r, and p_i^r replaced by their predicted values. Formally, this is done by solving the following deterministic convex optimization problem.

$$\begin{aligned}\text{LT-PA: } \min_{\mathbf{m}, \boldsymbol{\lambda}, \mathbf{q}^l} & \sum_{i=1}^N p_i^l q_i^l + \sum_{i=1}^N \hat{p}_i^r [m_i - \hat{w}_i^r - q_i^l]_+ \\ & + \beta \sum_i \sum_j h_{ij}(m_i, \lambda_{ij}) \\ \text{s.t } & (2a), (2c)\text{–}(2e) \\ & \sum_{i \in N} \lambda_{ij} = \hat{L}_j^r \quad \forall j \in J \\ & q_i^l \geq 0 \ \forall i \in N\end{aligned}$$

The objective function of LT-PA is similar to that of the EP-LT without the expectation operation. The constraints over $\boldsymbol{m}$, $\boldsymbol{\lambda}$, and $\boldsymbol{q}^l$ of LT-PA are identical to those of GLB-RT and EP-LT. LT-PA is a convex optimization problem and can be solved efficiently by standard techniques.

4.2 Stochastic Gradient-based Algorithm (SGA)

Although PA can offer a quick heurictic decision, it is desirable to have an algorithm that optimally procures electricity in long-term markets. To this end, we exploit the gradient characterization of the long-term objective (see Theorem 2) to design a stochastic gradient descent algorithm. The algorithm, namely, SGA, is summarized in Algorithm 1. The main idea of the algorithm is to compute a noisy estimate of the gradient of the long-term objective by averaging the gradient of the (random) total cost over a finite number of sample paths. This noisy gradient is used to perform a stochastic gradient descent. Stochastic approximation theory can then be used to prove convergence to the set of optimal solutions, as long as the step-size sequence is appropriately diminishing [8].

Algorithm 1 Stochastic Gradient based Algorithm (SGA).

Input: Obtain $\boldsymbol{p}^l$ from the $|N|$ long-term electricity markets.
Prepare $\mathbb{S}$ samples of $(\boldsymbol{w}^r, \boldsymbol{L}^r, \boldsymbol{p}^r)$ based on prediction error distributions.
Output: $q_i^l\ \forall i \in N$
Initialize: $q_i^l = 0, \forall i \in N$.
Step: $\tau = 1$.
while 1 **do**
for all k such that $1 \leq k \leq \mathbb{S}$ **do**
Solve: GLB-RT for kth sample of $(\boldsymbol{w}^r, \boldsymbol{L}^r, \boldsymbol{p}^r)$ with long-term procurement $\boldsymbol{q}^l$
Obtain: The Lagrange multipliers $\varrho_i^{(k)}$ corresponding to constraint (2f), $\forall i \in N$
end for
Compute: $\hat{\varrho}_i = \frac{1}{\mathbb{S}}\sum_{k=1}^{\mathbb{S}} \varrho_i^{(k)}, \forall i \in N$
Update: $q_i^l = [q_i^l - \eta_\tau(p_i^l - \hat{\varrho}_i)]_{[0,M_i]}$ for $\forall i \in N$. $[z]_{[0,M_i]}$ indicates the projection of z onto the set $[0, M_i]$.
Increase: $\tau = \tau + 1$.
end while

We prove that SGA converges to the set of optimal solutions of EP-LT under the following standard assumption on the step-size sequence.

Assumption 1. $\sum_{\tau=1}^{\infty}(\eta_\tau) = \infty$ *and* $\sum_{\tau=1}^{\infty}(\eta_\tau)^2 < \infty$.

The convergence of SGA is asserted by the following theorem.

Theorem 4. *Under Assumption 1, almost surely, the iterates $\boldsymbol{q}^l$ generated by SGA converge to the set of optimal solutions of EP-LT as $\tau \to \infty$.*

We give the proof of Theorem 4 in Appendix B of [9].

Note that SGA requires samples from the joint distribution of $(\boldsymbol{w}^r, \boldsymbol{L}^r, \boldsymbol{p}^r)$. This means that SGA can be solved in an entirely data-driven manner, without needing to actually model the distributions of workload, renewable generation, and electricity price, or the complex inter-dependencies between these quantities. This makes it particularly suitable in today's 'big-data' era. The bottleneck of SGA is the computation of the noisy gradient estimate, which involves solving $\mathbb{S}$ instances of GLB-RT. Moreover, the diminishing step-size sequence implies that SGA requires a large number of iterations to compute a near-optimal solution. However, it is important to note that since this algorithm is only used for long-term procurement, its computation time would not be a bottleneck in practice.

5. NUMERICAL RESULTS

Experimental Setup. There are 10 logical data centers which are located in 10 different states. We assume that there are one million servers distributed across the ten logical data centers. The peak power consumption for each server is 300W. We consider 40 sources, corresponding to 40 states of the US; the corresponding workload data is obtained from Akamai Technologies. The average workload is 30% of the total capacity of the data centers. The network delays are estimated to be proportional to the distance between sources and data centers. The importance of delay is estimated according to the fact that 100 ms latency costs 1% of Amazon in sales [10].

The electricity prices in real-time markets are the industrial electricity prices of each state in May 2010 [1]. We set the long-term prices such that the ratio $\frac{\mathbb{E}[p_i^r]}{p_i^l} = 2.5$. To simulate the uncertainties, the error distributions from the full version [9] are used to generate the samples of renewable energy generation (wind generation), workload, and electricity price. The mean absolute errors (MAE) of prediction errors for wind generation, electricity price, and workload demand are 65%, 40%, and 35%, respectively. The penetration of renewable penetration is 50%.

Table 1: Baseline algorithms.

	EP-LT	fixed LT	GLB-RT	Nearest
nLTnGLB	no	no	no	yes
fLTnGLB	no	50%	no	yes
nLT	no	no	yes	no
fLT	no	50%	yes	no

Cost savings. We highlight the benefit of our proposed system by comparing with the baselines in Table 1. Nearest is the naive load balancing scheme that forwards the workload to the closest data centers. Fixed LT just pre-purchases electricity for a certain percentage of workload in long-term markets. In addition to the baseline algorithms, we compare our algorithms to *Oracle Algorithm (OA)*. OA is an unrealizable algorithm that is given to the absolute performance limit by assuming assumes all realizations of renewable energy, workload, and electricity prices are fully known apriori. Similarly to PA, the problem of long-term procurement can then be solved efficiently. The cost of OA is measured by averaging its output over many realizations.

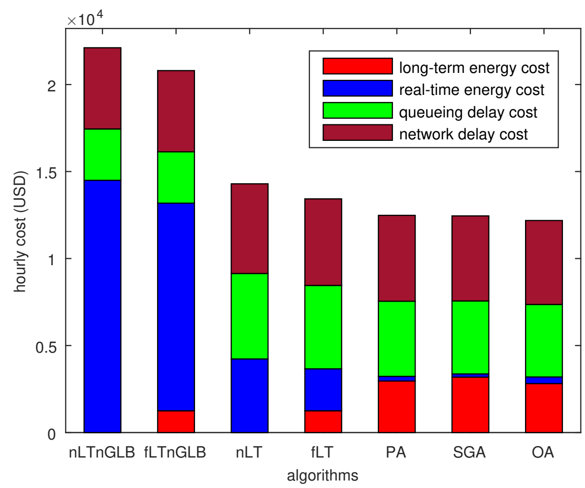

Figure 3: The proposed algorithms PA and SGA are very close to the lower bound, OA and outperform the traditional methods up to 44%.

Figure 3 compares the cost performance among our proposed algorithms and the traditional algorithms. PA and SGA save up to 44% compared to other simpler algorithms, and are comparable to the oracle algorithm (OA), the impractical lower bound. It also shows the significant benefits for cloud providers to participate long-term markets. Surprisingly, the performance of PA is very close to that of the SGA. The small performance gap between PA, SGA and OA is due to the trade-off between energy cost and delay cost, and the compensation of GLB-RT at real-time markets. The details can be found in the full version [9].

Why do our proposed algorithms perform so well? The intuition behind the small performance gaps between PA, SGA and OA is the compensation of GLB-RT at real-time markets. In particular, GLB-RT can utilize the available renewable energy and cheap electricity to partially compensates for performance gap caused by the prediction errors in long-term. More interestingly, PA and SGA are noticeably aggressive in long-term markets as in Figure 3. In addition, PA and SGA are even more aggressive than OA. In fact, Lemma 3 allows PA, SGA, and OA to purchase a lot of electricity in long-term markets, because the over-provisioned energy can be used up to reduce queuing delay in real-time. Thus, there is the trade-off between the energy costs and delay costs that helps our proposed methods become close to OA.

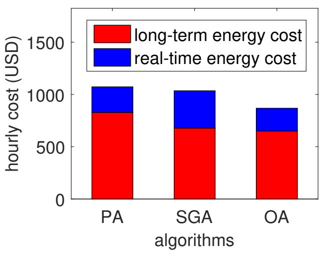

(a) $\beta = 0$

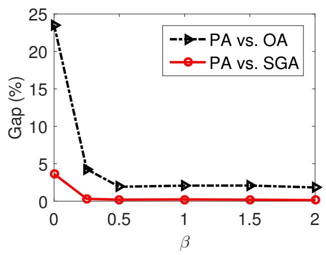

(b) Vary β

Figure 4: The impact of delay on the proposed algorithms.

How does the trade-off between energy costs and delay costs benefit our proposed algorithms? To answer this question, we vary the constant factor β that weighs the delay costs relative to energy costs. When $\beta = 0$, i.e., the delay costs are ignored, the cost breakdown are shown in Figure 4a. The performance gap between PA and OA is 24% that is much larger than the 2% gap in Figure 3 ($\beta = 1$). In this setting, SGA outperforms PA by 4%. We observe that PA is more aggressive compared to SGA in long-term procurement. Figure 4b shows the performance gaps of PA versus OA and PA versus SGA with varying β. In this figure, the x-axis shows a scaled β, where a value of 1 corresponds to the default value. We note that the performance gaps are significant when β is small (< 0.25). However, the gaps are very small when β is relatively large (≥ 0.5).

6. CONCLUDING REMARKS

In this paper, we present a systematic study of optimal energy procurement systems for geo-distributed data centers that utilize multi-timescale electricity markets. The contributions of this paper are three-fold: (i) designing algorithms for long-term electricity procurement in multi-timescale markets; (ii) analyzing long-term prediction errors using real-world traces; and (iii) empirically evaluating the benefits of our proposed procurement systems. In particular, we proposed two algorithms, PA and SGA, both of which save up to 44% of the energy procurement cost compared to traditional algorithms that do not use long-term markets or geographical load balancing. While SGA provably converges to an optimal solution, PA surprisingly achieves a cost that is nearly optimal with much less computing effort.

7. ACKNOWLEDGEMENTS

This research is supported by NSF grants CNS-1464388, CNS-1617698, and CNS-1413998. This research was partially funded by the MSIP, Korea, under the “ICT Consilience Creative Program” (IITP-2015-R0346-15-1007) supervised by the IITP and by the NRF, Korea, under the “Basic Science Research Program” (NRF-2015R1C1A1A0 1053788).

8. REFERENCES

[1] `http://www.eia.gov/`.

[2] L. M. Ausubel and P. Cramton. Using forward markets to improve electricity market design. *Utilities Policy*, 18(4):195–200, 2010.

[3] H. Beheshti and A. Croll. Performance impact: How web speed affects online business KPIs. In *Velocity Online Conference*. O'Reilly, 2009.

[4] P. Cramton. Colombia's forward energy market. 2007.

[5] M. Ghamkhari, A. Wierman, and H. Mohsenian-Rad. Energy portfolio optimization of data centers. *IEEE Transactions on Smart Grid*, 2016.

[6] Google. Renewable energy. `http://www.google.com/about/datacenters/renewable/index.html`, 2015. [Online; accessed 25-April-2015].

[7] Y. Guo, Z. Ding, Y. Fang, and D. Wu. Cutting down electricity cost in internet data centers by using energy storage. In *Proc. IEEE GLOBECOM*, pages 1–5, 2011.

[8] H. J. Kushner and G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2003.

[9] T. N. Le, J. Liang, Z. Liu, R. K. Sitaraman, J. Nair, and B. J. Choi. Optimal energy procurement for geo-distributed data centers in multi-timescale markets. `https://goo.gl/7OUB9h`, March 2017.

[10] J. Liddle. Amazon found every 100ms of latency cost them 1% in sales. *The GigaSpaces*, 27, 2008.

[11] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. Andrew. Greening geographical load balancing. In *Proc. ACM SIGMETRICS*, pages 233–244, 2011.

[12] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. In *Proc. ACM SIGCOMM*, volume 39, pages 123–134, 2009.

[13] 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*, pages 1–9, 2010.

[14] L. Rao, X. Liu, L. Xie, and Z. Pang. Hedging against uncertainty: A tale of internet data center operations under smart grid environment. *Smart Grid, IEEE Transactions on*, 2(3):555–563, 2011.

[15] J. Whitney and P. Delforge. Data center efficiency assessment. *Issue paper on NRDC (The Natural Resource Defense Council)*, 2014.

[16] L. Yu, T. Jiang, Y. Cao, S. Yang, and Z. Wang. Risk management in internet data center operations under smart grid environment. In *Proc. IEEE SmartGridComm*, pages 384–388, 2012.

[17] W. Zheng, K. Ma, and X. Wang. Exploiting thermal energy storage to reduce data center capital and operating expenses. In *Proc. IEEE HPCA*, pages 132–141, 2014.