Competitive Online Algorithms for Geographical Load Balancing in Data Centers with Energy Storage

Chi-Kin Chau
Masdar Institute of Science and Technology
ckchau@masdar.ac.ae

Lin Yang
The Chinese University of Hong Kong
yl015@ie.cuhk.edu.hk

ABSTRACT

Geographical load balancing takes advantage of the regional differences in dynamic electricity rates by shifting computing tasks among geographically distributed data centers. Since energy storage is becoming an integral part of data centers, one can maximize the benefit of the temporal and spatial fluctuations of electricity rates by combining geographical load balancing and energy storage management. Previously, the problem of integrated geographical load balancing with energy storage has been studied based on Lyapunov stochastic optimization approach, which relies on asymptotic analysis by averaging over infinite time horizon and arbitrarily large energy storage. In this paper, we present a *competitive online algorithmic* approach, which can be applied to finite time horizon and small-to-medium energy storage with a worst-case guarantee from the offline optimal solutions. By simulations on real-world data, it is observed that our competitive online algorithms can significantly outperform Lyapunov optimization algorithm.

CCS Concepts

•**Computer-Communication Networks** → *Network Architecture and Design;*

Keywords

Data Centers, Online Algorithms, Competitive Analysis

1. INTRODUCTION

Nowadays, data centers consume a substantial amount of electricity and generate ever increasing operation costs. The advent of dynamic electricity markets provides a novel way to alleviate the electricity cost of data centers. One vital option is the concept of *geographical load balancing* [10], by which computing tasks are forwarded among geographically distributed data centers to take advantage of the regional differences in dynamic electricity rates.

In addition to the spatial fluctuations of electricity rates, one can also harness the temporal fluctuations of electricity rates by employing energy storage management [6] at data centers to store electricity at a low electricity rate and discharge from energy storage at a high electricity rate. This naturally gives rise to an optimization problem combining geographical load balancing and energy storage management. Noteworthily, the decisions of energy storage management are carried out in an online manner, based on the currently revealed information without the knowledge of future electricity rates and computing task workload.

In the extant literature, the decision-making problems with uncertain future inputs are tackled by three common approaches. First, one can rely on the predictions of future inputs (e.g., electricity rates and computing task workload). This approach crucially relies on accurate prediction models or specifically trained classifiers for the particular environments, and is difficult to be adopted to new environments with noisy or limited historical data for calibration.

Second, one can utilize stochastic optimization, which relies on probability models to handle uncertainty or noisy data. The solutions are usually obtained in the sense of probabilistic expectation, which may deviate considerably from a particular sample outcome. In particular, a Lyapunov optimization approach has been proposed [11], by which a control policy is developed to asymptotically converge to the optimal policy, when inputs are assumed to be i.i.d. or stationary Markovian, and the storage size is large. Recently, [12] applies Lyapunov optimization approach to integrated geographical load balancing problem energy storage. However, in practice the inputs (workload and electricity rates) may be non-stationary, and the size of energy storage may be small or medium. Lyapunov optimization also relies on averaging over infinite time horizon, which may not be close to the optimal when applied to a finite time horizon.

As a departure from the aforementioned approaches, this paper pursues an *online competitive algorithmic* approach, which has been employed in a wide range of online decision-making problems [3,4,9], without relying on the information of future inputs. This approach can cope with arbitrary (stochastic or adversarial) future inputs, with a finite or infinite time horizon, and can provide a worst-case optimality assurance in terms of competitive ratio (by benchmarking with the offline optimal decisions based on full future inputs) without asymptotic assumptions. In this paper, we present competitive online algorithms for tackling integrated geographical load balancing problem with energy storage, which can be applied to finite time horizon and small-to-medium energy storage with a worst-case guarantee from the offline

E2DC, June 21-24 2016, Waterloo, ON, Canada

ISBN 978-1-4503-4421-0/16/06…$15.00

DOI: http://dx.doi.org/10.1145/2940679.2940680

optimal solutions. By simulations on real-world data, it is observed that our competitive online algorithms can significantly outperform Lyapunov optimization algorithm.

2. PROBLEM FORMULATION

In this paper, we consider a set of n data centers, denoted by $\mathcal{N}$. Each data center $j \in \mathcal{N}$ has local power grid and energy storage. The workload for computing tasks arrives at a centralized forwarder. Each data center makes online decisions for its energy management system, workload processing and forwarding operations to minimize the total cost. See Table 1 for a table of key notations.

2.1 Data Center System Model

Given certain workload of computing tasks, each data center $j \in \mathcal{N}$ orchestrates different energy sources (e.g., power grid, energy storage) to minimize the total cost, subject to satisfying the workload and operational constraints. The system model of such a scenario is depicted in Fig. 1 (a), which has been widely used in the literature [11]. A discrete-time model is considered, such that each time slot matches the timescale at which the energy management and workload forwarding decisions are updated (e.g., every minute). Without loss of generality, it is assumed that there are totally T time slots, and each has a unit length, where the inputs within one time slot are sufficiently quasi-static.

The system model of a single data center is consisted of the following components:

- **Workload:** Arbitrary arrivals of workload of computing tasks are considered, denoted by $a(t)$. We do not assume any specific stochastic model of $a(t)$. We normalize the unit of workload by the equivalent unit of required electricity for the processing corresponding computing tasks. The workload should be satisfied by the energy acquired from the grid or energy storage.

- **Power Grid:** The system can acquire electricity from the grid for unsatisfied workload in an on-demand manner. Let the market rate at time t of the grid at the j-th data center be $p^j(t)$, where $m^j \leq p^j(t) \leq M^j$. Denote the ratio between maximum and minimum rates by $\varphi^j \triangleq \frac{M^j}{m^j}$. We do not assume any specific stochastic model on $p^j(t)$. Denote the acquired energy for satisfying the workload directly by $v^j_{\rm a}(t)$ and the acquired energy to charge the energy storage by $v^j_{\rm b}(t)$. M^j and m^j can be estimated in advance, for example, based on historical data. Note that the proposed algorithm still applies, even when M^j and m^j are not known a-priori.

- **Energy Storage:** The energy storage can reduce the total electricity cost by exploiting rate fluctuations. The energy storage has a capacity B^j. The level of energy storage at time t is given by:

$$x^j(t+1) = x^j(t) + \eta_{\rm c} v^j_{\rm b}(t) - \eta_{\rm d} d^j(t) \tag{1}$$

  where $d^j(t)$ is the energy discharged from the energy storage, $v^j_{\rm b}(t)$ is the energy charged to the energy storage from the grid, respectively. Constants $\eta_{\rm c} \leq 1$ and $\eta_{\rm d} \geq 1$ are charging and discharging efficiencies, respectively. To capture the limitations in the charging and discharging rates, it is required that $d^j(t) \leq \mu_{\rm d}$ and $v^j_{\rm b}(t) \leq \mu_{\rm c}$. The level of energy storage is required to satisfy the boundary conditions, $x^j(0) = 0$ and $x^j(t) = B^j$.

Note that energy storage systems may bear other tear-and-wear and long-term maintenance costs. This paper considers the short-term operation of energy storage systems, such that the electricity cost considerably outweighs the maintenance costs.

Notation	Definition
$a(t)$	Incoming workload arriving at centralized forwarder at time t
$a^j(t)$	Workload processed at data center j at time t
$p^j(t)$	Market rate of electricity per unit of processed workload at data center j at time t (where $m^j \leq p^j(t) \leq M^j$)
q^j	Transmission cost per unit of workload from centralized forwarder to data center j
B^j	Capacity of energy storage at data center j
$x^j(t)$	Current level of energy storage at data center j at time t
$\eta_{\rm c}(\leq 1)$	Charging efficiency of energy storage
$\eta_{\rm d}(\geq 1)$	Discharging efficiency of energy storage
$\mu_{\rm c}$	Charging rate constraint of energy storage
$\mu_{\rm d}$	Discharging rate constraint of energy storage
$d^j(t)$	Energy discharged from energy storage at data center j at time t
$v^j(t)$	Energy acquired from the grid at data center j at time t
$v^j_{\rm a}(t)$	Energy acquired from the grid for satisfying workload at data center j at time t
$v^j_{\rm b}(t)$	Energy acquired from the grid to charge energy storage at data center j at time t

Table 1: Key notations.

2.2 Data Center Forwarding Model

When there are multiple data centers, we consider a centralized router that decides to which appropriate data center is the workload forwarded. See an illustration in Fig. 1 (b).

The cost minimization problem of geographical load balancing with energy storage management is formulated in `GLB-ES`.

$$(\texttt{GLB-ES}) \ \min \sum_{j\in\mathcal{N}} \sum_{t=1}^{T} p(t)\big(v^j_{\rm a}(t) + v^j_{\rm b}(t)\big) + \sum_{j\in\mathcal{N}} q^j a^j(t) \tag{2}$$

$$\text{s.t. } x^j(t+1) - x^j(t)$$

$$= \eta_{\rm c} v^j_{\rm b}(t) - \eta_{\rm d} d^j(t), \text{ for all } j \in \mathcal{N} \tag{3}$$

$$d^j(t) + v^j_{\rm a}(t) = a^j(t), \text{ for all } j \in \mathcal{N} \tag{4}$$

$$0 \leq x^j(t) \leq B^j, \text{ for all } j \in \mathcal{N} \tag{5}$$

$$v^j_{\rm b}(t) \leq \mu_{\rm c}, \text{ for all } j \in \mathcal{N} \tag{6}$$

$$d^j(t) \leq \mu_{\rm d}, \text{ for all } j \in \mathcal{N} \tag{7}$$

$$x^j(0) = 0, \ x^j(T) = B^j, \text{ for all } j \in \mathcal{N} \tag{8}$$

$$\sum_{j\in\mathcal{N}} a^j(t) = a(t) \tag{9}$$

$$\text{var. } a^j(t), x^j(t), d^j(t), v^j_{\rm a}(t), v^j_{\rm b}(t) \geq 0, \text{ for all } j \in \mathcal{N}$$

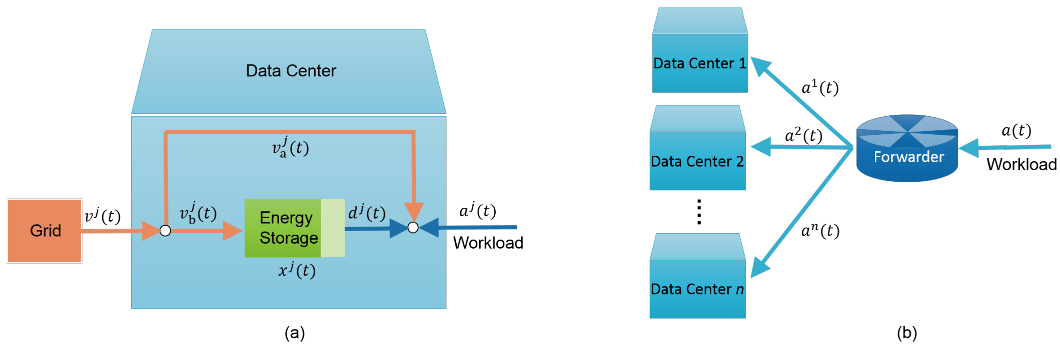

Figure 1: (a) A depiction of system model of data center with energy storage. (b) A depiction of forwarding model among data centers.

Note that our model is similar to that in [12], which also imposes QoS constraints on forwarding among data centers. Our model can easily incorporate QoS constraints. But for clarity, we focus on energy storage management decisions.

2.3 Online Algorithms

Let the inputs of the problem (i.e., the sequence of workload, market rates) be $\sigma = \big(a(t), (p^j(t))_{j\in\mathcal{N}}\big)_{t=1}^{T}$. The problem `GLB-ES` can be solved by linear programming, when all inputs σ are given in advance.

However, σ is revealed gradually over time in reality, which requires decisions to be made without future inputs. An algorithm is called *online*, if the decision at the current time only depends on the instantaneous inputs before or at the current time $t_{\rm now}$, namely, $\big(a(t), (p^j(t))_{j\in\mathcal{N}}\big)_{t\le t_{\rm now}}$.

Given input σ, let $\texttt{Cost}(\texttt{Algo}[\sigma])$ be the cost of algorithm `Algo`, and $\texttt{Opt}(\sigma)$ be the cost of an offline optimal solution (that may rely on an oracle to obtain all future inputs). In competitive analysis for online algorithms [3], *competitive ratio* is a common performance metric, defined as the *worst-case* ratio between the cost of the online algorithm `Algo` and that of an offline optimal solution, namely,

$$\texttt{CR}(\texttt{Algo}) \triangleq \max_{\sigma} \frac{\texttt{Cost}(\texttt{Algo}[\sigma])}{\texttt{Opt}(\sigma)} \tag{10}$$

`Algo` is called c-competitive, if $\texttt{CR}(\texttt{Algo}) = c$. This paper provides competitive online algorithms for solving `GLB-ES` with a worst-case guarantee.

2.4 k-Min Search

The problem `GLB-ES` is closely related to a classical online algorithmic problem called k-min search. In 1-min search problem, a trader is searching for the minimum rate of some asset. At each time slot t, the trader is presented a rate $p(t)$ and must decide whether or not to accept this rate. Once the trader decides to accept the rate $p(t)$, the search ends and the trader's cost is $p(t)$. If the trader does not accept any rate for the first $T-1$ time slots, he needs to accept any rate at time slot T.

According to [7], the online algorithm that accepts the first rate below threshold $\sqrt{Mm}$ has a competitive ratio $\sqrt{\varphi}$. Furthermore, any deterministic online algorithm attains a competitive ratio at least $\Omega(\sqrt{\varphi})$.

In more general k-min search problem [8], a trader is searching for the k-th minimum rate of some asset, when given a sequence of rates in an online fashion.

3. COMPETITIVE ONLINE ALGORITHMS

We denote the maximum rate, minimum rate, and its ratio for data center j by M^j, m^j, φ^j, respectively. For convenience, we assume that the data centers are ordered by $(j_1, j_2, ..., j_N)$, such that $\varphi^{j_1} \le \varphi^{j_2} \le ... \le \varphi^{j_N}$. We present online algorithms to solve `GLB-ES`.

3.1 Basic Online Algorithm

We first present a basic online algorithm based on 1-min search. Online algorithm $\texttt{Algo}_{\texttt{th}}$ proceeds as follows: (1) forward workload to any data center that has cheap nonzero energy storage level, (2) consume the energy from each energy storage, and (3) satisfy the unsatisfied workload by the acquired power from the data center that has the minimum market rate ($\min_{j\in\mathcal{N}} p^j(t)$). At each data center, the energy storage is charged from the grid, if the local market rate is below the respective threshold (i.e., $p^j(t) \le \theta$). Let $\underline{M} \triangleq \min_{j\in\mathcal{N}}\{q^j + M^j\}$ and $\underline{m} \triangleq \min_{j\in\mathcal{N}}\{q^j + m^j\}$.

Algorithm 1 $\texttt{Algo}_{\texttt{th}}\big[t, a(t), \big(\theta, B^j, p^j(t)\big)_{j\in\mathcal{N}}\big]$

Sort $\mathcal{N}$, such that $q^1 \le q^2 \le ... \le q^n$
for $j = 1$ **to** n **do**
  $x^j(t) \leftarrow x^j(t-1)$
  ▷ *Satisfy workload from energy storage*
  $d^j(t) \leftarrow \min\{a(t), \mu_{\rm d}, \frac{x^j(t)}{\eta_{\rm d}}\}$
  $a^j(t) \leftarrow d^j(t)$
  $a(t) \leftarrow a(t) - d^j(t)$
  $x^j(t) \leftarrow x^j(t) - \eta_{\rm d} d^j(t)$
  ▷ *Store energy from grid*
  if $p^j(t) \le \theta$ **then**
    $v_{\rm b}^j(t) \leftarrow \min\big\{[\frac{B^j - x^j(t)}{\eta_{\rm c}}]^+, \mu_{\rm c}\big\}$
    $x^j(t) \leftarrow x^j(t) + \eta_{\rm c} v_{\rm b}^j(t)$
  end if
end for
▷ *Satisfy workload from the cheapest grid*
$h = \arg\min_{j\in\mathcal{N}}\{p^j(t) + q^j\}$
$v_{\rm a}^h(t) \leftarrow a(t)$
$a^h(t) \leftarrow a^h(t) + v_{\rm a}^h(t)$
return $\big(a^j(t), x^j(t), d^j(t), v_{\rm a}^j(t), v_{\rm b}^j(t)\big)_{j\in\mathcal{N}}$

THEOREM 1. *If we set the threshold in* $\texttt{Algo}_{\texttt{th}}$ *by*

$$\theta = \frac{\sqrt{8\underline{M}\underline{m} + \underline{M}^2} - \underline{M}}{2} \cdot \frac{\eta_{\rm c}}{\eta_{\rm d}} \tag{11}$$

Then, the competitive ratio of $\texttt{Algo}_{\texttt{th}}$ *is bounded by*

$$\texttt{CR}(\texttt{Algo}_{\texttt{th}}) \leq \frac{\sqrt{8\underline{M}\underline{m} + \underline{M}^2} + \underline{M}}{2\underline{m}} \tag{12}$$

PROOF. (Sketch) Workload $(a(t))_{t=1}^{T}$ is called *one-shot*, if there is a unique time slot $t_{\mathrm{nz}} \in [1, T]$ such that

$$a(t) = \begin{cases} 0 & \text{if } t \neq t_{\mathrm{nz}} \\ \bar{a} & \text{if } t = t_{\mathrm{nz}} \end{cases} \tag{13}$$

where $\bar{a}$ is the peak of $(a(t))_{t=1}^{T}$. We define a function called *one-shot decomposition*, which can decompose any workload into a collection of one-shot workload:

$$\texttt{1sDecompose}\Big[(a(t))_{t=1}^{T}\Big] = (t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i})_{i=1}^{m} \tag{14}$$

where m is the number of decomposed one-shot workload, t_{nz}^{i} is the non-zero workload time slot and $\bar{a}^{i}$ is the peak of the i-th one-shot workload, and t_{s}^{i} $(\leq t_{\mathrm{nz}}^{i})$ is the minimum starting time slot for the i-th one-shot workload.

$\texttt{1sDecompose}$ is constructed as follows. Let $B \triangleq \sum_{j\in\mathcal{N}} B^{j}$. We define the accumulative workload curve by $\texttt{Acc}[a(t)] \triangleq \sum_{t'=0}^{t} a(t')$, and $\texttt{Acc}[a(t)] + \frac{B}{\eta_{\mathrm{d}}}$ is the upward shift by $\frac{B}{\eta_{\mathrm{d}}}$. The one-shot workload is constructed by the rectanglizing the region sandwiched between $\texttt{Acc}[a(t)] + \frac{B}{\eta_{\mathrm{d}}}$ and $\texttt{Acc}[a(t)]$. Each one-shot workload $(t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i})$ corresponds to a rectangle of $(t_{\mathrm{nz}}^{i} - t_{\mathrm{s}}^{i}) \times \bar{a}^{i}$, which is maximally inscribed in the sandwiched region. See Fig. 2 for an illustration.

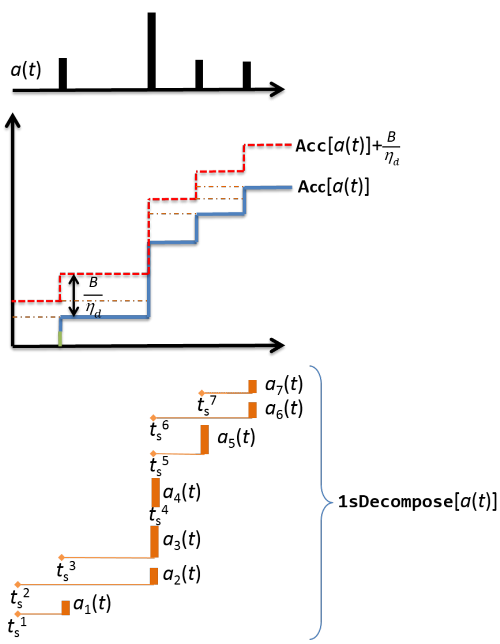

Figure 2: A illustration of $\texttt{1sDecompose}$**.**

Note that $\texttt{1sDecompose}$ has the following properties:

1. $a(t)$ can be reconstructed by the one-shot workload:

$$a(t) = \sum_{i:t_{\mathrm{nz}}^{i}=t} \bar{a}^{i} \quad \text{for all } t \tag{15}$$

2. There is a non-decreasing order on the starting time slots and non-zero workload time slots:

$$t_{\mathrm{s}}^{i} \leq t_{\mathrm{s}}^{i+1} \quad \text{and} \quad t_{\mathrm{nz}}^{i} \leq t_{\mathrm{nz}}^{i+1} \quad \text{for all } i \tag{16}$$

3. Let $\mathcal{D}$ be the set of one-shot workload that have non-zero duration, namely, $\mathcal{D} \triangleq \{i \mid t_{\mathrm{s}}^{i} < t_{\mathrm{nz}}^{i}\}$. Let $\mathcal{D}(i)$ be the set of one-shot workload in $\mathcal{D}$ other than i, such that the peak workload also lie within $[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}]$, namely, $\mathcal{D}(i) \triangleq \{j \in \mathcal{D}\backslash\{i\} \mid t_{\mathrm{s}}^{i} \leq t_{\mathrm{nz}}^{j} \leq t_{\mathrm{nz}}^{i}\}$. If $i \in \mathcal{D}$, then

$$\textstyle\sum_{j:\in\mathcal{D}(i)} \bar{a}^{j} + \bar{a}^{i} \leq \frac{B}{\eta_{\mathrm{d}}} \tag{17}$$

Eqn. (17) ensures that satisfying the other one-shot workload in $[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}]$ using the energy storage still leave sufficient capacity in the energy storage for the i-th one-shot workload. We set each t_{s}^{i} as minimum as possible subject to Eqn. (17). When $t_{\mathrm{s}}^{i} = t_{\mathrm{nz}}^{i}$, then the one-shot workload needs to acquire energy from the grid.

The basic idea is that each one-shot workload can be tackled separately as 1-min search, and hence, we obtain the competitive ratio of $\texttt{Algo}_{\texttt{th}}$ as that of 1-min search. Denote the offline optimal solution by $\texttt{Opt}$.

In the following, we first consider unconstrained charging and discharging rates, where $\mu_{\mathrm{c}}, \mu_{\mathrm{d}} \geq B$. With respect to each one-shot workload $(t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i})$, there are two cases:

Case 1: Market rate $p^{j}(t) > \theta$ for all $t \in [t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}]$ and all $j \in \mathcal{N}$. Then $\texttt{Algo}_{\texttt{th}}$ will not store energy from the grid. $\texttt{Algo}_{\texttt{th}}$ needs to acquire energy from the grid for an amount of $\bar{a}^{i}$ at a market rate at most $\underline{M}$ at time slot t_{nz}^{i}. $\texttt{Opt}$ needs to store energy from the grid for at least an amount of $\bar{a}^{i}\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}$ at a market rate at least θ within $[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i})$. Hence,

$$\texttt{Cost}\big(\texttt{Algo}_{\texttt{th}}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i}]\big) \leq \bar{a}^{i}\underline{M}$$
$$\texttt{Cost}\big(\texttt{Opt}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i}]\big) \geq \bar{a}^{i}\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}$$

Case 2: Market rate $p^{j}(t) \leq \theta$ for some $t \in [t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}]$ and some $j \in \mathcal{N}$. Let $\bar{a}^{i,j}$ be the amount of workload can be satisfied from energy storage at the j-th data center. If $\bar{a}^{i,j} > 0$, then $\texttt{Algo}_{\texttt{th}}$ stores energy from the grid for an amount of $\bar{a}^{i,j}\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}$ at a market rate at most θ. If $\bar{a}^{i,j} = 0$, then $\texttt{Algo}_{\texttt{th}}$ needs to acquire energy from the grid at a market rate at most $\underline{M}$. Note that this case is due to the fact that the energy storage has been used previously to acquire energy at a market rate at most θ. On average, the per unit cost is at most $\frac{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}+\underline{M}}{2}$. But $\texttt{Opt}$ needs to acquire energy from the grid for an amount of $\bar{a}^{i,j}$ at a market rate at least $\underline{m}$. Hence,

$$\texttt{Cost}\big(\texttt{Algo}_{\texttt{th}}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i,j}]\big) \leq \bar{a}^{i,j}\frac{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}+\underline{M}}{2}$$
$$\texttt{Cost}\big(\texttt{Opt}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i,j}]\big) \geq \bar{a}^{i,j}\underline{m}$$

Hence, the competitive ratio is obtained by

$$\texttt{CR}(\texttt{Algo}_{\texttt{th}}) = \max_{\sigma}\frac{\sum_{i,j}\texttt{Cost}\big(\texttt{Algo}_{\texttt{th}}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i,j}]\big)}{\sum_{i,j}\texttt{Cost}\big(\texttt{Opt}[t_{\mathrm{s}}^{i}, t_{\mathrm{nz}}^{i}, \bar{a}^{i,j}]\big)} \leq \max\Big\{\frac{\underline{M}}{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}}, \frac{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}+\underline{M}}{2\underline{m}}\Big\}$$

To minimize the competitive ratio, we set $\frac{\underline{M}}{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}} = \frac{\theta\frac{\eta_{\mathrm{d}}}{\eta_{\mathrm{c}}}+\underline{M}}{2\underline{m}}$. The positive root is $\theta = \frac{\sqrt{8\underline{M}\underline{m}+\underline{M}^2}-\underline{M}}{2} \cdot \frac{\eta_{\mathrm{c}}}{\eta_{\mathrm{d}}}$. Hence, the competitive ratio of $\texttt{Algo}_{\texttt{th}}$ is $\texttt{CR}(\texttt{Algo}_{\texttt{th}}) \leq \frac{\sqrt{8\underline{M}\underline{m}+\underline{M}^2}+\underline{M}}{2\underline{m}}$.

Finally, we consider constrained charging and discharging rates, where $\mu_{\mathrm{c}} < B$ or $\mu_{\mathrm{d}} < B$. It is straightforward to show that the competitive ratio with constrained charging and discharging rates is not higher than the one with unconstrained charging and discharging rates. □

3.2 Improved Online Algorithm

We next present an improved online algorithm based on k-min search [8]. In a more general setting, we divide the capacity of energy storage by k units. In algorithm $\texttt{Algo}^k_{\texttt{th}}$, at each time the ℓ-th unit of energy storage is charged from the grid at the j-th data center.

Algorithm 2 $\texttt{Algo}^k_{\texttt{th}}\big[t, a(t), \big((\theta_\ell)^k_{\ell=1}, B^j, p^j(t)\big)_{j\in\mathcal{N}}\big]$

Sort $\mathcal{N}$, such that $q^1 \le q^2 \le ... \le q^n$
  ▷ *ℓ^j is the number of charged units in energy storage*
$\ell^j \leftarrow 0$ for all $j \in \mathcal{N}$
for $j = 1$ to n do
  $x^j(t) \leftarrow x^j(t-1)$
  ▷ *Satisfy workload from energy storage*
  $d^j(t) \leftarrow \min\{a(t), \mu_{\rm d}, \frac{x^j(t)}{\eta_{\rm d}}\}$
  $a^j(t) \leftarrow d^j(t)$
  $a(t) \leftarrow a(t) - d^j(t)$
  $x^j(t) \leftarrow x^j(t) - \eta_{\rm d} d^j(t)$
  $\ell^j \leftarrow \min\{k\frac{x^j(t)}{B^j} + 1, k\}$
  ▷ *Store energy from grid*
  if $p^j(t) \le \theta_{\ell^j}$ then
    $v^j_{\rm b}(t) \leftarrow \min\big\{[\frac{B^j\cdot\ell^j/k - x^j(t)}{\eta_{\rm c}}]^+, \mu_{\rm c}\big\}$
    $x^j(t) \leftarrow x^j(t) + \eta_{\rm c} v^j_{\rm b}(t)$
    $\ell^j \leftarrow \min\{\ell^j + 1, k\}$
  end if
end for
▷ *Satisfy workload from the cheapest grid*
$h = \arg\min_{j\in\mathcal{N}}\{p^j(t) + q^j\}$
$v^h_{\rm a}(t) \leftarrow a(t)$
$a^h(t) \leftarrow a^h(t) + v^h_{\rm a}(t)$
return $\big(a^j(t), x^j(t), d^j(t), v^j_{\rm a}(t), v^j_{\rm b}(t)\big)_{j\in\mathcal{N}}$

THEOREM 2. *If we set the threshold in $\texttt{Algo}^k_{\texttt{th}}$ by*

$$\theta_\ell = \underline{M}\Big(1 - \frac{2k(1+\frac{1}{2ks^*})^\ell(1-s^*)}{1+2ks^*}\Big)\cdot\frac{\eta_{\rm c}}{\eta_{\rm d}} \qquad (18)$$

where s^* is a fixed-point solution to the following equation:

$$\frac{2\underline{m}s^* - \underline{M}}{\underline{M}} = 1 - 2\Big((1+\frac{1}{2ks^*})^k - 1\Big)(s^*-1) \qquad (19)$$

Then, the competitive ratio of $\texttt{Algo}^k_{\texttt{th}}$ is bounded by

$$\texttt{CR}(\texttt{Algo}^k_{\texttt{th}}) \le s^* \qquad (20)$$

PROOF. (Sketch) Most of the proof is similar to Theorem 1. But, with respect to each one-shot workload $(t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^i)$, there are $k+1$ cases:

Case 1: Market rate $p^j(t) > \theta_1$ for all $t \in [t^i_{\rm s}, t^i_{\rm nz}]$ and all $j \in \mathcal{N}$. Hence,

$$\texttt{Cost}\big(\texttt{Algo}^k_{\texttt{th}}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^i]\big) \le \bar{a}^i\underline{M}$$
$$\texttt{Cost}\big(\texttt{Opt}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^i]\big) \ge \bar{a}^i\theta_1\tfrac{\eta_{\rm d}}{\eta_{\rm c}}$$

Case 2: Market rate $\theta_2 \le p^j(t) \le \theta_1$ for some $t \in [t^i_{\rm s}, t^i_{\rm nz}]$ and some $j \in \mathcal{N}$. Let $\bar{a}^{i,j}$ be the amount of workload can be satisfied from energy storage at the j-th data center. At the j-th data center,

$$\texttt{Cost}\big(\texttt{Algo}^k_{\texttt{th}}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^{i,j}]\big) \le (2k-1)\tfrac{\bar{a}^{i,j}}{2k}\underline{M} + \tfrac{\bar{a}^{i,j}}{2k}\theta_1\tfrac{\eta_{\rm d}}{\eta_{\rm c}}$$
$$\texttt{Cost}\big(\texttt{Opt}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^{i,j}]\big) \ge \bar{a}^{i,j}\theta_2\tfrac{\eta_{\rm d}}{\eta_{\rm c}}$$

$$\vdots$$

Case k+1: Market rate $p^j(t) \le \theta_k$ for some $t \in [t^i_{\rm s}, t^i_{\rm nz}]$ and some $j \in \mathcal{N}$. Let $\bar{a}^{i,j}$ be the amount of workload can be satisfied from energy storage at the j-th data center. At the j-th data center,

$$\texttt{Cost}\big(\texttt{Algo}^k_{\texttt{th}}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^{i,j}]\big) \le \tfrac{\bar{a}^{i,j}\underline{M}}{2} + \textstyle\sum^k_{\ell=1}\tfrac{\bar{a}^{i,j}}{2k}\theta_\ell\tfrac{\eta_{\rm d}}{\eta_{\rm c}}$$
$$\texttt{Cost}\big(\texttt{Opt}[t^i_{\rm s}, t^i_{\rm nz}, \bar{a}^{i,j}]\big) \ge \bar{a}^{i,j}\underline{m}$$

Similar to the proof of Theorem 1, the competitive ratio is obtained by

$$\texttt{CR}(\texttt{Algo}^k_{\texttt{th}}) \le \max\Big\{\tfrac{k\underline{M}}{k\theta_1\frac{\eta_{\rm d}}{\eta_{\rm c}}}, \tfrac{(2k-1)\underline{M}+\theta_1\frac{\eta_{\rm d}}{\eta_{\rm c}}}{2k\theta_2\frac{\eta_{\rm d}}{\eta_{\rm c}}}, ..., \tfrac{k\underline{M}+\sum^k_{\ell=1}\theta_\ell\frac{\eta_{\rm d}}{\eta_{\rm c}}}{2k\underline{m}}\Big\}$$

To minimize the competitive ratio, we set

$$\tfrac{k\underline{M}}{k\theta_1\frac{\eta_{\rm d}}{\eta_{\rm c}}} = \tfrac{(2k-1)\underline{M}+\theta_1\frac{\eta_{\rm d}}{\eta_{\rm c}}}{2k\theta_2\frac{\eta_{\rm d}}{\eta_{\rm c}}} = \ldots = \tfrac{k\underline{M}+\sum^k_{\ell=1}\theta_\ell\frac{\eta_{\rm d}}{\eta_{\rm c}}}{2k\underline{m}} \qquad (21)$$

One can solve θ_ℓ from Eqn. (21) as

$$\theta_\ell = \underline{M}\Big(1 - \frac{2k(1+\frac{1}{2ks^*})^\ell(1-s^*)}{1+2ks^*}\Big)\cdot\frac{\eta_{\rm c}}{\eta_{\rm d}} \qquad (22)$$

where $s^* \triangleq \frac{\underline{M}}{\theta_1\frac{\eta_{\rm d}}{\eta_{\rm c}}}$. Namely, the competitive ratio of $\texttt{Algo}^k_{\texttt{th}}$ is $\texttt{CR}(\texttt{Algo}_{\texttt{th}}) \le s^*$.

By substituting Eqn. (22) into $s^* = \frac{k\underline{M}+\sum^k_{\ell=1}\theta_\ell\frac{\eta_{\rm d}}{\eta_{\rm c}}}{2k\underline{m}}$, one can obtain

$$\frac{2\underline{m}s^* - \underline{M}}{\underline{M}} = 1 - 2\Big((1+\frac{1}{2ks^*})^k - 1\Big)(s^*-1) \qquad (23)$$

□

The complete proofs can be found in the full paper [5].

4. EMPIRICAL EVALUATION

The performance of the proposed algorithms is evaluated based on the simulation studies using real-world traces. We consider the case of four data centers. We assume four data centers of Wikipedia are located in four different places whose spot electricity market rate are not necessarily equal. Every data center is equipped with a energy storage system which can shift the electricity peak during the peak periods of the spot price.

4.1 Parameters and Settings of Simulations

- *Workload*: The real-word workload data is based on the access traces published by Wikipedia [2]. One line of those traces corresponds to one web access, including: 1) a monotonically increasing counter (useful for sorting the trace in chronological order), 2) the timestamp of the request in Unix notation with millisecond precision, and 3) the requested URL etc. We assume that there are four Wikipedia data centers located in four different locations, namely, Chicago, New York, Palo Alto and Houston. According to the requested URL, we identify which data center is visited by the user. By counting the number of accesses every 10 minutes in a data center, we can approximately calculate the workload of a data center according to the model in [12]. The workload is measured by the amount of electricity in every five minutes with unit of KW. Note that we ignore the QoS model in [12], so the workload only depends on the frequency of web accessing.

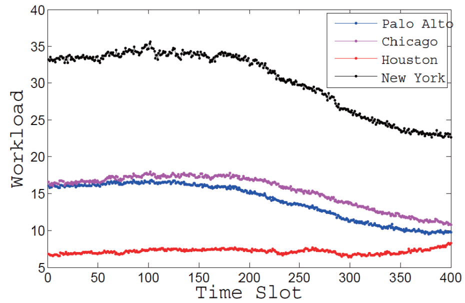

Figure 3: The workload of the four data centers in Chicago, New York, Palo Alto and Houston.

We parsed the trace files which record the website access requests during September, 2007, and select 400 slots whose length is 10 minutes to implement the online algorithm and Lyapunov optimization algorithm.

- *Electricity Rates*: In our simulation, the real-time price and workload data are both based on real-word raw data [1]. We use the 5-min locational marginal rates in four different locations: Chicago, New York, Palo Alto and Houston. The electricity price is unit of \$/MWh.

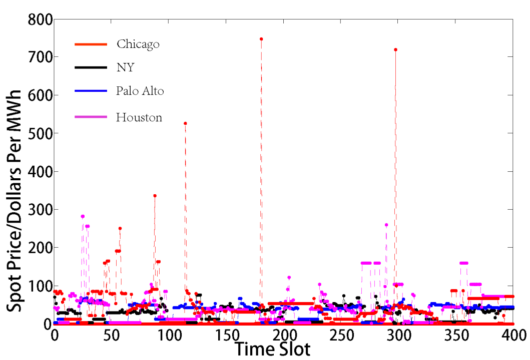

Figure 4: The spot prices of four locations, namely, Chicago, New York, Palo Alto and Houston.

- *Energy Storage*: For simplicity, It is assumed that $\eta_d = \eta_c = 1$. To compare with Lyapunov optimization, set the charging and discharging rate constraints μ_c and μ_d to be 10KWh per slot.

4.2 Performance Comparison

In the first evaluation, we set the energy storage capacity as 300 KWh, and the maximum charge and discharge rate are both 10 KWh. The initial state of batteries are set to be 150 KWh. In the improved online algorithm we partition the total capacity of batteries into forty equal parts, which means the parameter k in $\texttt{Algo}^k_{\texttt{th}}$ takes value 40. The performance result is shown in Figure 3. From this figure we can conclude that the basic online algorithm and improved online algorithm performs better than the Lyapunov optimization algorithm with around 40% and 50% improvement in cost saving, respectively.

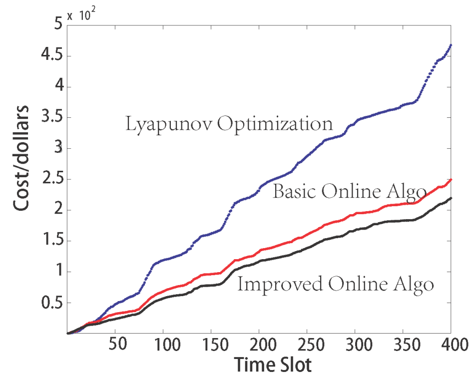

Figure 5: Performance comparison of Lyapunov optimization and proposed online algorithms.

Other than that, we also record the fluctuation of energy storage reserves under the Lyapunov optimization and our proposed online algorithm, respectively. See Figure 4 and 5. We analyze the behavior of the two algorithms and observe that the fluctuation scope under Lyapunov optimization is much smaller than that of the proposed online algorithm, which may lead to underutilization of the energy storage resource. Actually the fluctuation scope under Lyapunov optimization algorithm accounts for only about 15-25% of the total capacity, while ours can almost reach 100%. Actually, the Lyapunov optimization guarantees the fluctuation in a certain scope when it tries to improve the performance. In some case, it may be conservative and thus fails to gain further profit. This can explain why our proposed online algorithms outperform the Lyapunov optimization algorithm.

5. REFERENCES

[1] Energy regulatory commission. http://www.www.ferc.gov.

[2] Wikipedia access traces. http://www.wikibench.eu.

[3] A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.

[4] C.-K. Chau, M. Khonji, and M. Aftab. Online algorithms for information aggregation from distributed and correlated sources. *to appear in IEEE/ACM Trans. Networking*, 2016. http://arxiv.org/abs/1601.03147.

[5] C.-K. Chau and L. Yang. Competitive online geographical load balancing for data centers with energy storage and renewable energy. Technical report, Masdar Institute, 2016.

[6] C.-K. Chau, G. Zhang, and M. Chen. Cost minimizing online algorithms for energy storage management with worst-case guarantee. *to appear in IEEE Trans. Smart Grid*, 2016. http://arxiv.org/abs/1511.07559.

[7] R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. Optimal search and one-way trading online algorithms. *Algorithmica*, 2001.

[8] J. Lorenz, K. Panagiotou, and A. Steger. Optimal algorithms for k-search with application in option pricing. *Algorithmica*, 2009.

[9] L. Lu, J. Tu, C.-K. Chau, M. Chen, and X. Lin. Online energy generation scheduling for microgrids with intermittent energy sources and co-generation. In *ACM SIGMETRICS*, 2013.

[10] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the electric bill for internet-scale systems. In *ACM SIGCOMM*, 2009.

[11] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam. Optimal power cost management using stored energy in data centers. In *ACM SIGMETRICS*, 2011.

[12] G. Y. and F. Y. Electricity cost saving strategy in data centers by using energy storage. *IEEE Trans. Parallel and Distributed Systems*, 2013.