Scalable Workload Management for Water Efficiency in Data Centers

Lanchao Liu*, Shaolei Ren[†] and Zhu Han*

*Department of Electrical and Computer Engineering, University of Houston, Houston, Texas 77004

[†]School of Computing and Information Sciences, Florida International University, Miami, Florida 33199

Abstract—The huge demand for data center computing nowadays has resulted in a significant amount of electricity consumption as well as environmental impacts. While current works mainly focus on the energy cost of data centers, the severity of water consumption problem in data centers is largely neglected. In this paper, we propose an optimization framework for the workload management of data centers, which takes the efficiency of water usage into account. The workload management is formulated as a revenue maximization problem. To solve the large-scale optimization problem with scalability, the alternating direction method of multipliers (ADMM) is utilized. The optimization problem is decomposed into independent subproblems, which can be solved in a parallel fashion on distributed computing units and coordinated through dual variables. We evaluate the performance of proposed algorithm by simulations, and numerical results validate the effectiveness of the proposed algorithm.

I. INTRODUCTION

The cloud computing has become a part of people's daily lives and has caused a large demand for data center computing. Nowadays, people are used to access various services such as search engine, email, GPS navigation, streaming video and social networks deployed on data centers. Meanwhile, small and medium enterprises seize the opportunity to utilize the data center computing ability as a flexible and economically efficient solution for service provisioning. It has great potential for cloud service providers to generate a huge amount of revenues without investing much capitals for building and maintaining their own infrastructures. Hence, it is of great incentives for service providers to expand both the number and scale of their data centers worldwide.

The energy consumption of data centers have drawn significant attentions from both industry and academia. Currently, significant research efforts have been made to reduce electricity usage and carbon footprint of data centers. However, the importance of water efficiency, which is a critical concern for future data centers' sustainability, has been undervalued. The AT&T's cooling towers in data center facilities consumed 1 billion gallons of water in 2012, which is approximate 30% of the entire company's water consumption [1]. When operational, the Utah Data Center [2] is expected to consume 1.7 million gallons of water a day in order to cool the people and facilities, which is equivalent to 10,000 households' daily usage. Moreover, the extended droughts have been a potential threats for the daily operation of data centers. On January 17,

2014, California [3] proclaimed a state-wide drought emergency urging a water usage cut of 20% for business and may take mandatory measures if drought condition continuously worsens. It is imperative to consider water consumption and water efficiency for workload management in data centers.

The workload management problem for data centers can be seen as a network utility maximization (NUM) problem [4], which describes a unifying framework for understanding and designing distributed control as well as resource allocation in communication networks. [5] presents a novel workload management mechanism which exploits the time variation in electricity price, the availability of renewable energy, and the efficiency cooling. A significant amount of energy cost can be reduced by the proposed method. [6] considers a temperature-aware workload management in geographically distributed datacenters for energy efficiency in cooling systems. [7] considers a jointly request mapping and response routing for cloud services, which can reduce the energy usage and enhance the quality of cloud services deployed on data centers.

The energy-efficient workload management approaches are insufficient to improve the water efficiency in data centers since the temporal and spatial diversities of data center water efficiency are neglected. The significance of water efficiency in data centers has been noticed in recent works [8]–[10]. The benefits and possibilities of improvements in power and water usage effectiveness, cost, and operation for data centers are discussed in [8]. [9] presents technologies, design strategies, and operational approaches to improve the energy efficiency and sustainability of data centers. [10] exploits the temporal and spatial diversities of water efficiency, and proposed a novel approach for geographical load balancing.

In this paper, we present a distributed workload management framework for data centers, which takes the impact of water consumption into account. In particular, we consider the direct water usage in data centers especially for the cooling infrastructure, which is a significant factor for water efficiency in data centers. The management tasks are formulated as a revenue maximization problem, in which the users' service requests are optimally located to provide qualified service at a reasonable cost. To solve the resultant large scale optimization problem with manageable complexity, the alternating direction method of multipliers (ADMM) is utilized. The optimization problem is decomposed into a set of independent subproblems, which can be solved in parallel on distributed nodes and coordinated through dual variables. We evaluate the performance

This work is partially supported by US NSF CNS-1443917, ECCS-1405121, CNS-1265268, and CNS- 0953377

of proposed algorithm by simulations, and the numerical tests validate the effectiveness of the proposed algorithm.

The remaining of this paper is organized as follows. Section II presents the problem formulation. The proposed algorithm is described in Section III, and the numerical tests are given in Section IV. Finally, Section V concludes the paper.

II. PROBLEM FORMULATION

We first present the background of water efficiency in data centers, and then describe the workload management problem.

A. Background

Nowadays chillers typically use the most energy in a data center's cooling infrastructure, so large data centers, including AT&T [1] and Google [11], employ evaporative cooling as an energy-efficient substitution. In particular, large cooling towers are used to vent the heat generated by the servers into the environment through water evaporation, and then the evaporated water is replenished with makeup water. Currently, around 40% [12] of large data centers utilize cooling towers for the temperature control, and majority of those data centers rely on fresh water for supplement.

Generally, water consumption in cooling towers consists of two parts: evaporation and blow down [9]. The first one happens when exhaust air leaving a cooling tower also carries away droplets of water. The second one is the dumping of the cooling tower water to sewage. Blow down is necessary for a cooling tower due to the water accumulating dissolved minerals as the water evaporates. The data center water efficiency can be described by the water usage effectiveness (WUE), which is defined by the ratio of total water consumption to the information technology (IT) energy usage as follows:

$$\text{WUE} = \frac{\text{Total Water Usage}}{\text{IT Equipment Energy}}. \quad (1)$$

The value of WUE depends on the water flow rate, outside humidity, temperature, supply/return water temperature difference as well as system configurations. Similarly, the data center power usage efficiency (PUE) is defined as:

$$\text{PUE} = \frac{\text{Total Power Usage}}{\text{IT Equipment Energy}}. \quad (2)$$

A record of PUE and WUE of Facebooks data center at Prineville, OR from Apr. 2013 to Mar. 2014 is shown in Fig 1.¹ Remark that the WUE of data centers diversifies temporally and spatially, and most of the time the PUE and WUE at the same location demonstrate little correlation to each other.

B. System Model

We consider a set $\mathcal{J}$ of geographically distributed data centers, indexed by $j = 1, \dots, J$. A set $\mathcal{I}$ of users are requesting services from data centers. In this work, a user $i = 1, \dots, N$ is defined as an aggregated group of customers

¹Though Facebook utilize the air economizer for cooling, water is still used for temperature and humidity control in data centers. For example, the WUE at a cooling tower-based data center is 2.52L/kWh (eBay [14]), and 0.42L/kWh for a "free" air cooling data center (Facebook [13]).

Fig. 1. PUE and WUE of Facebooks data center in Prineville, OR from Apr. 2013 to Mar. 2014 [13].

which have the same IP prefix. The service requests from customers are first handled by the portal server. After that, one data center at a specific location is assigned to process the request. The service request at user i is denoted by λ_i and we assume that user's workload is precisely predicted [15].

We use $a_{i,j}$ as user i 's application requests processed by data center j . The allocation decision variables of user i is denoted as $\mathbf{a}_i = (a_{i,1}, \dots, a_{i,J})^T$. In practices, a user can be authoritative DNS servers or HTTP ingress proxies. Additionally, in this work we assume an agent has the fine-grained control of the network traffic, which is a reasonable assumption in nowadays commercial products and techniques [16]–[18]. The capacity of data center j is C_j , and the total workload from all users is $\sum_{i \in \mathcal{I}} a_{i,j}$. Once data center j finished processing the service request, the response traffic will be routed back to the user. A propagation latency $L_{i,j}$ is observed over wired connection between user i and data center j , which can be obtained by active measurements.

The performance objective of user i is characterized by a utility function $F_i(\cdot)$, and the utility functions can be different among different users. In this work, $F_i(\cdot)$ is designed to be a non-decreasing, non-negative and concave function in $\mathbf{a}_i$. A simple example is $F_i(\mathbf{a}_i) = \sum_{j \in \mathcal{J}} a_{i,j}(L_{max} - L_{i,j})$, where L_{max} is the maximum tolerable latency for service, or can be a more general class of functions that represent the elasticity of the service request and/or determine the fairness of resource allocation [19].

We characterize the performance objectives of each data center j by function G_j , which accounts for the power and water consumptions and parameterized by β_j as follows:

$$G_j(\mathbf{a}_i) = \gamma_{j,1}(\mathbf{a}_i) + \beta_j \gamma_{j,2}(\mathbf{a}_i), \quad (3)$$

where $\beta_j > 0$ highlights the importance of the water usage in cost function, and can be different across locations. $\gamma_{j,1}(\cdot)$ and $\gamma_{j,2}(\cdot)$ account for the electricity cost and water consumption of data center j , respectively, which scale with the total workload assigned to data center j . Here, $\gamma_{j,1}(\cdot)$ and $\gamma_{j,2}(\cdot)$ are designed to be non-decreasing, non-negative and convex function in $\sum_{i \in \mathcal{I}} a_{i,j}$. For example, the electricity bills and

water cost at data center j are modeled as follows [20], [21]:

$$\gamma_{j,1}(\mathbf{a}_i) = Pr_j \times \varepsilon_j \times [P_{idle} + (P_{peak} - P_{idle}) \sum_{i \in \mathcal{I}} a_{i,j}], \quad (4)$$

$$\gamma_{j,1}(\mathbf{a}_i) = \varepsilon_j \times [P_{idle} + (P_{peak} - P_{idle}) \sum_{i \in \mathcal{I}} a_{i,j}], \quad (5)$$

where Pr_j is the spot electricity price at the location of data center j , and ε_j is the PUE at data center j . ε_j is the WUE at data center j . P_{peak} and P_{idle} are the server peak power and server idle power, respectively.

C. Workload Management Problem

The workload management problem can be formulated as an optimization problem to maximize the revenue as follows:

$$\underset{\{\mathbf{a}_i\}_{i=1}^N}{\text{minimize}} \quad \sum_{j \in \mathcal{J}} G_j(\mathbf{a}_i) - \alpha \sum_{i \in \mathcal{I}} F_i(\mathbf{a}_i) \quad (6)$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{i,j} = \lambda_i, \quad \forall i, \quad (7)$$

$$\sum_{i \in \mathcal{I}} a_{i,j} \leq C_j, \quad \forall j, \quad (8)$$

$$a_{i,j} \geq 0 \quad \forall i, j, \quad (9)$$

where (7) is the workload conservation constraint for each user, and (8) is the data center capacity constraint. The problem (6) is not readily solved in a distributed fashion due to coupling of $a_{i,j}$ across data centers in constraint (7), and the coupling of $a_{i,j}$ across users in constraint (8). In the next section, the ADMM is utilized to decouple those constraints and design a distributed approach for (6).

III. PROPOSED ALGORITHM

In this section, we first introduce the background of the ADMM. Then we present the proposed method for distributed workload management for data centers.

A. An Introduction to ADMM

The general form of ADMM [22] can be described as follows. Consider an optimization problem

$$\underset{\mathbf{x}, \mathbf{z}}{\text{minimize}} \quad f(\mathbf{x}) + g(\mathbf{z}) \quad (10)$$

$$\text{subject to} \quad \mathbf{Ax} + \mathbf{Bz} = \mathbf{c}, \quad (11)$$

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{z} \in \mathbb{R}^m$ and $\mathbf{c} \in \mathbb{R}^p$. Matrices $\mathbf{A} \in \mathbb{R}^{p \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times m}$. Functions f and g are convex, close and proper. The scaled augmented Lagrangian can be expressed as:

$$\mathcal{L}_\rho(\mathbf{x}, \mathbf{z}, \boldsymbol{\mu}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{Ax} + \mathbf{Bz} - \mathbf{c} + \boldsymbol{\mu}\|_2^2, \quad (12)$$

where $\rho > 0$ is the penalty parameter and $\boldsymbol{\mu}$ is the scaled dual variable. The $\mathbf{x}$ and $\mathbf{z}$ are updated in a Gauss-Seidel fashion using the scaled dual variable. At each iteration k , the update process can be expressed as:

$$\mathbf{x}^k = \arg \min_{\mathbf{x}} f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{Ax} + \mathbf{Bz}^k - \mathbf{c} + \boldsymbol{\mu}^k\|_2^2,$$

$$\mathbf{z}^{k+1} = \arg \min_{\mathbf{z}} g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{Ax}^{k+1} + \mathbf{Bz} - \mathbf{c} + \boldsymbol{\mu}^k\|_2^2.$$

Fig. 2. The proposed distributed optimization framework for workload management in data centers.

Finally, the scale dual variable is updated by:

$$\boldsymbol{\mu}^{k+1} = \boldsymbol{\mu}^k + \mathbf{Ax}^{k+1} + \mathbf{Bz}^{k+1} - \mathbf{c}.$$

B. Distributed Scalable Design

By introducing a set of auxiliary variables $\mathbf{b}_j = (b_{1,j}, \dots, b_{N,j})^\top$ for $j \in \mathcal{J}$ and enforcing $b_{i,j} = a_{i,j}$, the optimization problem (6) can be reformulated as:

$$\underset{\{\mathbf{a}_i\}_{i=1}^N}{\text{minimize}} \quad \sum_{j \in \mathcal{J}} G_j(\mathbf{b}_j) - \alpha \sum_{i \in \mathcal{I}} F_i(\mathbf{a}_i) \quad (13)$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{i,j} = \lambda_i, \quad \forall i, \quad (14)$$

$$\sum_{i \in \mathcal{I}} b_{i,j} \leq C_j, \quad \forall j, \quad (15)$$

$$a_{i,j} = b_{i,j}, \quad \forall i, j, \quad (16)$$

$$a_{i,j} \geq 0 \quad b_{i,j} \geq 0 \quad \forall i, j. \quad (17)$$

Now a distributed approach can be designed to solve the optimization problem (13). Specifically, the decision variables $\mathbf{a}_i$ and $\mathbf{b}_j$ are arranged into two groups, which correspond to user updates and data center updates, respectively. During the optimization, the variables of each group are optimized in a distributed and parallel fashion. In particular, each mobile service agent i solves $\mathbf{a}_i$ and each data center j obtains $\mathbf{b}_j$, and those two groups of decision variables are coordinated through auxiliary variables. The architecture of the proposed mechanism is illustrated in Fig. 2.

For compactness, we define sets $\mathcal{A}_i = \{\mathbf{a}_i | \sum_{j \in \mathcal{J}} a_{i,j} = \lambda_i, a_{i,j} \geq 0, \forall j \in \mathcal{J}\}$ and $\mathcal{B}_j = \{\mathbf{b}_j | \sum_{i \in \mathcal{I}} b_{i,j} \leq C_j, b_{i,j} \geq 0, \forall i \in \mathcal{I}\}$. Accordingly, $\mathcal{A} = \{\bigcup_{i=1}^N \mathcal{A}_i\}$ and $\mathcal{B} = \{\bigcup_{j=1}^J \mathcal{B}_j\}$. By applying the ADMM to solve the optimization problem (13), we first calculate the partial Lagrangian, which introduces the Lagrange multipliers only for constraint (16):

$$\begin{aligned} \mathcal{L}_\rho(\{\mathbf{a}_i\}_{i=1}^N, \{\mathbf{b}_j\}_{j=1}^J, \{\mu_{i,j}\}_{i=1,j=1}^{N,J}) \\ = \sum_{j \in \mathcal{J}} G_j(\mathbf{b}_j) - \alpha \sum_{i \in \mathcal{I}} F_i(\mathbf{a}_i) + \frac{\rho}{2} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \|a_{i,j} - b_{i,j} + \mu_{i,j}\|_2^2, \end{aligned} \quad (18)$$

Algorithm 1 Distributed workload management

Input: $\{G_j(\cdot)\}_{j=1}^J$, $\{F_i(\cdot)\}_{i=1}^N$, α , $\{\lambda_i\}_{i=1}^N$, $\{C_j\}_{j=1}^J$
Initialize: $\{\mathbf{a}_i\}_{i=1}^N$, $\{\mathbf{b}_j\}_{j=1}^J$, $\{\mu_{i,j}\}_{i=1,j=1}^{N,J}$;
while not converge **do**

$\{\mathbf{a}_i\}_{i=1}^N$ -update, distributively at each user computing node:
Parallel for i ;
minimize $\mathbf{a}_i \in \mathcal{A}_i \frac{\rho}{2} \sum_{j \in \mathcal{J}} \|a_{i,j} - b_{i,j}^k + \mu_{i,j}^k\|_2^2 - \alpha F_i(\mathbf{a}_i)$;

$\{\mathbf{b}_j\}_{j=1}^J$ -update, distributively at each data center computing node:
Parallel for j ;
minimize $\mathbf{b}_j \in \mathcal{B}_j G_j(\mathbf{b}_j) + \frac{\rho}{2} \sum_{i \in \mathcal{I}} \|a_{i,j}^{k+1} - b_{i,j}^{k+1} + \mu_{i,j}^k\|_2^2$;
 $\mu_{i,j}^{k+1} = \mu_{i,j}^k + a_{i,j}^{k+1} - b_{i,j}^{k+1}$;
Adjust penalty parameter ρ if necessary;
 $k = k + 1$;

end while
return $\{\mathbf{a}_i\}_{i=1}^N$, $\{\mathbf{b}_j\}_{j=1}^J$;
Output $\{\mathbf{a}_i\}_{i=1}^N$, $\{\mathbf{b}_j\}_{j=1}^J$.

where $\mu_{i,j}$ is the scaled Lagrange multiplier. The decision variables $\mathbf{a}_i$ and $\mathbf{b}_j$ are arranged into two groups and updated in an iterative fashion.

User updates: At iteration k , $\mathbf{a}_i$ is updated by

$$\underset{\{\mathbf{a}_i\}_{i=1}^N \in \mathcal{A}}{\text{minimize}} \quad \frac{\rho}{2} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \|a_{i,j} - b_{i,j}^k + \mu_{i,j}^k\|_2^2 - \alpha \sum_{i \in \mathcal{I}} F_i(\mathbf{a}_i). \quad (19)$$

In (19) the optimization is performed to maximize the utility at users under a square regularization term. The optimization problem (19) can be handled at independent computing unit corresponding to each user. Each computing unit i solves a stochastic optimization problem as follows:

$$\underset{\mathbf{a}_i \in \mathcal{A}_i}{\text{minimize}} \quad \frac{\rho}{2} \sum_{j \in \mathcal{J}} \|a_{i,j} - b_{i,j}^k + \mu_{i,j}^k\|_2^2 - \alpha F_i(\mathbf{a}_i). \quad (20)$$

Data center updates: The decision variables $\mathbf{b}_j$ at data centers are updated by:

$$\underset{\{\mathbf{b}_j\}_{j=1}^J \in \mathcal{B}}{\text{minimize}} \quad \sum_{j \in \mathcal{J}} G_j(\mathbf{b}_j) + \frac{\rho}{2} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \|a_{i,j}^{k+1} - b_{i,j}^{k+1} + \mu_{i,j}^k\|_2^2. \quad (21)$$

Problem (21) aims at minimizing the total cost among all data centers, which can be readily solved at each data center in parallel. Each data center j will determine $\mathbf{b}_j$ independently by solving the following optimization problem:

$$\underset{\mathbf{b}_j \in \mathcal{B}_j}{\text{minimize}} \quad G_j(\mathbf{b}_j) + \frac{\rho}{2} \sum_{i \in \mathcal{I}} \|a_{i,j}^{k+1} - b_{i,j}^{k+1} + \mu_{i,j}^k\|_2^2. \quad (22)$$

Finally, the dual variables are updated as:

$$\mu_{i,j}^{k+1} = \mu_{i,j}^k + a_{i,j}^{k+1} - b_{i,j}^{k+1}. \quad (23)$$

A description of proposed distributed workload management can be found in Algorithm 1. ADMM is a primal-dual algorithm where each computing unit solves its own subproblem,

Fig. 3. The convergence of the proposed algorithm at different number of users 100, 500 and 1000.

the variations to constraint (16) are systematically penalized at certain prices through the scaled dual variable to each individual subproblem. Remark that in the ADMM framework for distributed computing the dual variables, or price, are not uniformly set for all nodes, which will require costly synchronization. For the convex optimization problem, ADMM converges to the optimum geometrically [23].

IV. EVALUATION

The proposed algorithms are evaluated by numerical simulations from the perspective of computation performance and service management performance. The evaluation setup is introduced briefly, and then the numerical results are presented.

A. Evaluation Setup

We consider N users in our evaluation, $N \in \{100, 200, \dots, 1000\}$. There are 10 geographically distributed data centers to deliver services to users. The workload of each user is generated from a uniform distribution $\mathcal{U}(4000, 8000)$, with a mean of 6000 data unit. The capacity of each data center is generated in a similar fashion such that the total capacity of data centers is 1.6 times of mobile service agents total capacity. The latency $L_{i,j}$ is randomly generated from $\mathcal{U}(25, 300)$ for simplicity, with a unit of milliseconds.

To describe the service cost at data centers, the 2011 annual average day-ahead on peak prices at 10 different local markets as the electricity price for data centers [7] is used. The server peak power and server idle power are set to 200W and 100W, respectively. The power usage efficiently is 1.5. Due to the lack of access to direct WUE for data centers, we generate synthetic WUE base on the data from Facebook [13].

B. Convergence Performance

The convergence performance of the proposed algorithm is shown in Fig. 3. Here, the scaled relative error is used to demonstrate the results. Suppose r^k is the value of the objective function at the k^{th} iteration, and r^* is the optimal solution of the workload management problem, the scaled

Fig. 4. Comparisons of average revenue, cost and utility for the proposed algorithm and water-oblivious management approach.

relative error e is defined as

$$e = \left| \frac{r^t - r^*}{r^0 - r^*} \right|. \quad (24)$$

It is shown in Fig. 3 that the proposed algorithm converges for different number of users. When the number of mobile service agents is 100, the proposed algorithm takes a moderate number of iterations to converge. Furthermore, after very few iterations, the proposed algorithm yields close objective value to the optimal solution, which demonstrates the effectiveness and potential of proposed algorithm for solving the optimization problem approximately. Similar performance can be found when the numbers of mobile service agents are 500 and 1000. As the number of mobile service agents increases, the proposed algorithm converges with only small increment of iterations, which demonstrates the scalability of the proposed algorithm.

C. Workload Management Performance

In the following we show the effectiveness of the proposed algorithm on workload management. We compare the proposed algorithm with the water-oblivious workload management approach. In this case, we specify the number of users to 1000, and the performance comparison are discussed below.

The comparisons of average latency, revenue, cost and utility for two workload management approaches are shown in Fig. 4. It is shown in Fig. 4(a) that the proposed algorithm outperforms the water-oblivious management approach in average revenue. The proposed algorithm can largely reduce the average operational cost of data centers while only incurring a small loss in average utility. A detail analysis of the average cost is shown in Fig. 4(b). It is shown that the proposed algorithm can effectively reduce the water consumptions of data centers, while only leads to a negligible increase in the average power usage.

Next, we compare the average water usage performance of the proposed algorithm with the water-oblivious workload management approach by varying the number of users from 100 to 1000. The water usage percentage is calculated by dividing the water usage of proposed algorithm by that of the water-oblivious workload management. It is shown in Fig. 5 that when taking the water usage of data centers into consideration, the proposed workload management algorithm achieves a water usage reduction around 20% at different

Fig. 5. Comparisons of water usage for service management with/without consideration of water efficiency at different number of users.

number of users, compared with the water-oblivious workload management approach. The box plot, which depicts groups of numerical data through their quartiles, and the plot of the mean of the water usage reduction in Fig. 5 shows that such kind of improvement is relatively stable.

V. CONCLUSION

In this work, we have presented a distributed workload management mechanism for data centers, and taken the effect of data center water efficiency into account. We have formulated a service revenue maximization problem, in which the users' requests are optimally located to provide qualified service at a reasonable cost. We have used the alternating direction method of multipliers to solve the large scale stochastic optimization problem with manageable complexity. We have shown that, our algorithm can decompose the optimization problem into a set of independent subproblems, which can be solved in parallel on distributed nodes and coordinated through dual variables. Our numerical tests validate the effectiveness of the proposed algorithm.

REFERENCES

- [1] AT&T 2012 sustainability report: water, <http://www.att.com/gen/landing-pages?pid=24188>.
- [2] Water use hints at problems at Utah Data Center, <http://www.sltrib.com/>.
- [3] Office of governor in California. Governor brown declares drought state of emergency, 2014, <http://gov.ca.gov/news.php?id=18368>.
- [4] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: a mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255-312, January, 2007.
- [5] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwab, and C. Hyser, "Renewable and cooling aware workload management for sustainable data centers," in *proceedings of SIGMETRICS*, London, UK, June 2012.
- [6] H. Xu, C. Feng, and B. Li, "Temperature aware workload management in geo-distributed datacenters," in *proceedings of USENIX ICAC*, San Jose, CA, June 2013.
- [7] H. Xu and B. Li, "Joint request mapping and response routing for geo-distributed cloud services," in *proceedings of IEEE INFOCOM*, Turin, Italy, April 2013.
- [8] E. Frachtenberg, "Holistic datacenter design in the open compute project," *Computer*, vol.45, no. 7, pp. 83-85, July 2012.

- [9] D. Alger, *Grow a Greener Data Center*, Pearson, New Jersey, 2009.
- [10] S. Ren, "Optimizing Water Efficiency in Distributed Data Centers," *International Conference on Cloud and Green Computing*, Karlsruhe, Germany, September 2013.
- [11] Google Data Centers, <http://www.google.com/about/datacenters/>.
- [12] Uptime Institute. Data center industry survey, 2013. <http://uptimeinstitute.com/2013-survey-results>
- [13] Prineville Data Center, <http://www.facebook.com/PrinevilleDataCenter/>.
- [14] eBay, <http://tech.ebay.com/dashboard>
- [15] K. Papagiannaki, N. Taft, Z. Zhang, and C. Diot, "Long-term forecasting of Internet backbone traffic: observations and initial methods", in *proceedings of the IEEE INFOCOM*, San Francisco, CA, March 2003.
- [16] P. Wendell, J. W. Jiang, M. J. Freedman, and J. Rexford, "DONAR: decentralized server selection for cloud services," in *proceedings of the ACM SIGCOMM*, New Delhi, India, August 2010.
- [17] G. Wang, J. Wu, G. Zhou, and G. Li, "Collision-tolerant media access control for asynchronous users over frequency-selective channels," *IEEE Transactions on Wireless Communications*, vol. 12, no. 10, pp. 5162-5171, October 2013
- [18] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: enabling innovation in campus networks," in *proceedings of the ACM SIGCOMM*, Seattle, WA, August 2008.
- [19] M. Jeonghoon and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556-567, October, 2000.
- [20] 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 *proceedings of IEEE INFOCOM*, San Diego, CA, March 2010.
- [21] X. Fan, W. Weber, and L. A. Barroso, "Power provisioning for a warehouse-size computer," in *Proceedings of the ACM International Symposium on Computer Architecture*, San Diego, CA, June 2007.
- [22] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundation and Trends in Machine Learning*, vol. 3, no. 1, pp. 1-122, November 2010.
- [23] W. Deng and W. Yin, "On the global and linear convergence of the generalized alternating direction method of multipliers," *Rice CAAM technical report* 12-14, 2012.