Robust Workload and Energy Management for Sustainable Data Centers

Tianyi Chen, Yu Zhang, *Student Member, IEEE*, Xin Wang, *Senior Member, IEEE*, and Georgios B. Giannakis, *Fellow, IEEE*

***Abstract*—A large number of geo-distributed data centers begin to surge in the era of data deluge and information explosion. To meet the growing demand in massive data processing, the infrastructure of future data centers must be energy-efficient and sustainable. Facing this challenge, a systematic framework is put forth in this paper to integrate renewable energy sources (RES), distributed storage units, cooling facilities, as well as dynamic pricing into the workload and energy management tasks of a data center network. To cope with RES uncertainty, the resource allocation task is formulated as a robust optimization problem minimizing the worst-case net cost. Compared with existing stochastic optimization methods, the proposed approach entails a deterministic uncertainty set where generated RES reside, thus can be readily obtained in practice. It is further shown that the problem can be cast as a convex program, and then solved in a distributed fashion using the dual decomposition method. By exploiting the spatio-temporal diversity of local temperature, workload demand, energy prices, and renewable availability, the proposed approach outperforms existing alternatives, as corroborated by extensive numerical tests performed using real data.**

***Index Terms*—Cloud computing, data centers, energy storage, geographical load balancing, renewable energy, robust optimization, smart grids.**

NOMENCLATURE

A. Abbreviations

RES	Renewable energy sources.
DC	Data center.
MN	Mapping node.
QoS	Quality-of-service.
IT	Information technology.
DW	Delay-tolerant workload.
CG	Conventional generator.
RG	Renewable generator.
RWEM	Robust workload and energy management.

Manuscript received July 31, 2015; revised November 14, 2015; accepted December 17, 2015. Date of publication February 3, 2016; date of current version March 15, 2016. This work was supported in part by NSF Grant 1509040, Grant 1509005, Grant 1423316, Grant 1442686, and Grant 1202135.

T. Chen, Y. Zhang, and G. B. Giannakis are with the Department of Electrical and Computer Engineering, Digital Technology Center, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: chen3827@umn.edu; zhan1220@umn.edu; georgios@umn.edu).

X. Wang is with the Department of Computer and Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL 33431 USA (e-mail: xwang11@fau.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSAC.2016.2525618

B. Indices, numbers, and sets

$T, t, \mathcal{T}, \mathcal{T}_{i,s}$	Number, index, set and sub-set of time slots.
$I, i, \mathcal{I}$	Number, index, and set of DCs.
$J, j, \mathcal{J}$	Number, index, and set of MNs.
Q, q	Number and index of delay-tolerant workloads.
$\mathcal{Q}, \mathcal{Q}_j$	Sets of all delay-tolerant workloads and those collected by MN j.
k, ℓ	Iteration indexes of the (dual) subgradient ascent.
$\mathcal{T}_q$	Set of active slots of qth delay-tolerant job.
$\mathcal{X}$	Feasible set of primal variables.
$\mathcal{E}_i$	Uncertainty set of renewable generation in DC i.

C. Constants

A_j^t	Demand of interactive workloads at MN j per slot t.
B_q	Demand of qth delay-tolerant job.
L_{ji}^t	Bandwidth from MN j to DC i per slot t.
$S_q,\ E_q$	Starting time and deadline for qth delay-tolerant job.
ϱ	Fraction of server peak consumption when in idle state.
γ, κ_i^t	Parameters of chilled-water and outside-air cooling.
$T_{i,RA}^t, T_{i,OA}^t$	IT rack and outside air temperatures at DC i per time t.
$\overline{P}_{i,a}^t$	Cooling capacity using outside-air in DC i per slot t.
$P_{i,TH}^t$	Cooling source selection threshold in DC i per slot t.
$C_i^0, \underline{C}_i, \overline{C}_i$	Initial, minimum and maximum energy levels of battery.
$\overline{P}_{\text{i,ch}}, \overline{P}_{\text{i,dis}}$	Maximum charging and discharging rates of battery at DC i.
$R_i^{\text{up}}, R_i^{\text{dw}}$	Ramping-up/down limits of CG at DC i.
$\overline{P}_{i,g}$	Generation capacity of CG at DC i.
$E_i^t, \mathbf{e}_i$	Renewable generation at DC i per slot t, and vector collecting $[E_i^1, \ldots, E_i^T]^\top$.
$\overline{E}_i^t, \underline{E}_i^t$	Upper and lower bounds of renewable generation at DC i per slot t.
$\overline{E}_{\mathcal{T}_{i,s}}, \underline{E}_{\mathcal{T}_{i,s}}$	Upper and lower bounds of total renewable generation at DC i over $\mathcal{T}_{i,s}$.
$\overline{E}_{\mathcal{T}_{i,s}}^{\text{avg}}$	Historical sample average of total RES over $\mathcal{T}_{i,s}$ at DC i.
$\Delta_s^{\text{up}}, \Delta_s^{\text{low}}$	Levels of robustness for upper and lower bounds in RES uncertainty set.

α_i^t, β_i^t Buying and selling energy prices in DC i at slot t.
φ^t, ϕ^t Auxiliary variables of energy prices.

D. Variables

a_{ji}^t Interactive workloads distributed from MN j to DC i.
$\tilde{b}_{i,q}^t$, $b_{i,q}^t$ Amounts of qth delay-tolerant job routed to, and being processed in DC i per slot t.
d_i^t Total IT workload demand in DC i at slot t.
m_i^t Number of active servers in DC i at slot t.
$P_{i,w}^t$, $P_{i,a}^t$ IT power consumption allocated for chiller and outside-air cooling.
$P_{\mathrm{i,dis}}^t$,$P_{\mathrm{i,ch}}^t$ (Dis)charging power from battery at DC i per slot t.
$P_{\mathrm{i,in}}^t$, $P_{\mathrm{i,out}}^t$ Total energy supply and consumption in DC i per slot t.
C_i^t Energy level in battery i at the beginning of slot t.
$P_{i,g}^t$ Output generation of CG in DC i per slot t.
R_i^t Auxiliary variable for market power trading.
$\mathbf{x}$, $\overline{\mathbf{x}}$ Matrix collecting all primal variables and its running average.
$\lambda_{i,q}^t$, v_i^t, π_i^t Lagrange multipliers.
$\boldsymbol{\varpi}$ Vectors collecting all Lagrange multipliers.

I. Introduction

In THE NEW era of big data analytics, cloud computing, and the Internet of Things, data centers are proliferating globally to provide important Internet services such as instant messaging, video distribution, and data backup. For the purposes of reliability and quality-of-service (QoS), a cloud service provider usually has multiple data centers (DCs) geographically distributed in different areas. For instance, Google currently operates seven DCs in the US, and fourteen all over the world [1]. Along with the ever-increasing demand of Internet applications, energy-consuming DCs incur a huge amount of electricity bills. According to [2], DCs among the US consumed about 91 billion kWh electricity in 2013, which is almost twice the amount of power needed by all the households in New York City. Such a consumption is projected to reach 140 billion kWh by 2020.

In order to reduce the electricity cost, considerable efforts from both industry and academia have been made over the last decade [3]. One typical direction is to reduce the energy usage in DCs, through e.g., dynamic thermal management, server speed scaling, and dynamic resizing [4]–[6]. Another line of research aims at information technology (IT) workload management by exploiting the spatio-temporal diversity of energy prices [7], [8]. Specifically, [7] proposed a stochastic optimization based approach to workload routing and server management for geo-distributed DCs, while [8] investigated the means of reducing a single DC's peak loads by shifting the batch workloads. To deal with the uncertainties of IT workloads and electricity prices, [9] advocated hedging techniques in the financial management to smooth the associated dynamics.

The carbon dioxide emission along with the large energy consumption will also become a major concern. For instance, Google's DCs in 2010 approximately emitted carbon comparable to that of 280,000 cars [10]. Hence, not only achieving energy efficiency is of essence, but also the sustainability of DCs should be targeted [11], [12]. Exploiting RES is clearly a key for sustainable operation of DCs [13]. The main challenge of renewable-integrated energy management is to account for its random and nondispatchable nature, which motivates the use of energy storage units [14]. Energy supply side management with distributed storage units was considered for a homogeneous DC [15], and geo-distributed DCs [16]. Taking advantage of RES, a two-timescale Lyapunov optimization technique was developed to control the energy supply in both ahead-of-time and real-time settings [14]. However, spatial diversity of workloads, prices, and renewables was not utilized therein. Furthermore, in [14]–[16] the server model was simplified while the cooling power consumption was ignored. Yet, a substantial amount of DC energy actually goes to a variety of cooling facilities [17], [18].

Existing approaches of dealing with RES uncertainty include the Lyapunov optimization based on the stochastic approximation technique [7], [14]–[16], and the scenario-based stochastic optimization approaches using RES samples from historical data or a given distribution [19], [20]. To guarantee the convergence and optimality, these methods typically require the condition of independent and identically distributed (i.i.d.) RES samples, which is generally unrealistic [21], [22]. For a large time scale, the distribution of the RES generation changes with the seasons. For a short time scale, hourly RES generation can be highly correlated. How to properly handle the RES uncertainty for DC daily operations is still a challenging problem.

In this paper, we consider the optimal workload and energy management for a cloud network consisting of multiple geo-distributed mapping nodes (MNs) and DCs. For each DC, practical models of power supply, cooling, and IT systems are considered. In particular, the power supply system comprises multiple energy sources; the cooling system combines two different cooling sources with time-varying cooling coefficients[1]; and the IT operating system is able to adaptively activate servers and schedule workloads with guaranteed hard deadlines. Distinct from existing works, a *deterministic* uncertainty set of the *unknown* renewable generation, as well as a two-way energy trading mechanism is introduced to account for the stochastic and nondispatchable nature of RES. The proposed uncertainty set of the RES generation only requires easy-to-obtain first-order statistics, and allows the possible sample correlation. Control parameters are further designed to trade off robustness for conservatism of the robust optimization formulation.

Built on practical models, the resource allocation task is formulated as a robust optimization problem, which minimizes the

[1]Cooling coefficient is the power consumed for cooling divided by the IT demand, thus representing the cooling efficiency.

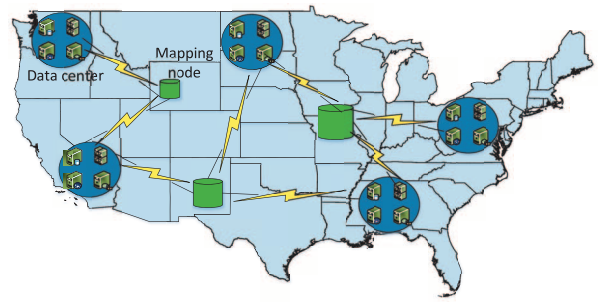

(a) A cloud network of geo-distributed DCs.

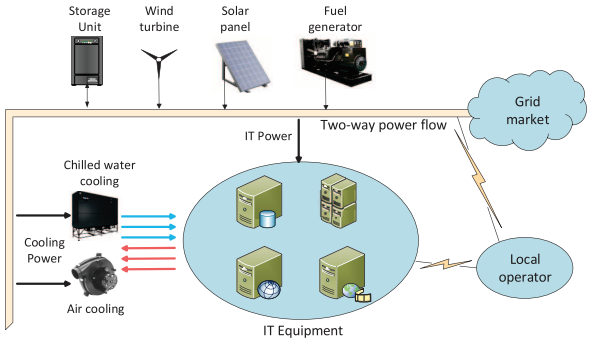

(b) A smart-grid powered sustainable DC.

Fig. 1. A sustainable cloud network design.

system's worst-case net cost subject to DCs' operational constraints. Leveraging the problem structure, we show that it can be cast as a convex program. Capitalizing on the dual decomposition approach, an efficient distributed solver is developed. It is shown that the proposed algorithm is guaranteed to obtain the desired robust workload and energy-management strategy, and could also facilitate distributed implementations among the MNs and DCs. Finally, extensive numerical results with real data corroborate the merits of the proposed framework and approaches.

The rest of the paper is organized as follows. The system models are described in Section II. The proposed robust management scheme is developed in Section III. Numerical results are provided in Section IV, followed by conclusions in Section V.

Notation. Boldface lower (upper) case letters represent column vectors (matrices); $(\cdot)'$ stands for vector (matrix) transpose. $\mathbb{R}$ represents the set of real numbers; and $[a]^+ := \max\{a, 0\}$.

II. System Models

Consider a network with geographically distributed MNs $\mathcal{J} := \{1, 2, \ldots, J\}$ and DCs $\mathcal{I} := \{1, 2, \ldots, I\}$; see Fig. 1(a). MNs first collect data requests from nearby areas, and then distribute them to different DCs. As shown in Fig. 1(b), each DC has three subsystems: an IT subsystem, a cooling (heat dissipation) subsystem, and a power supply subsystem.

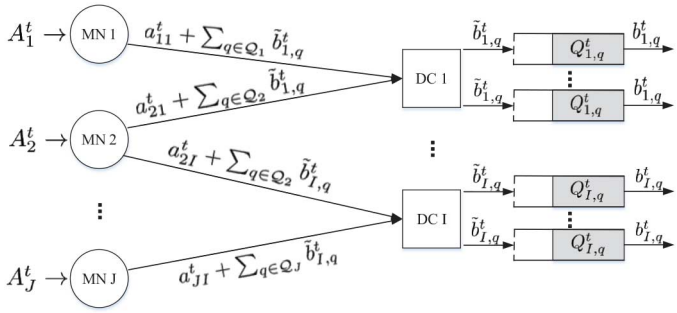

Fig. 2. A workload distribution diagram.

A. Network and Workload Models

In general, DC workloads are either delay-sensitive (interactive) or delay-tolerant [7]. The interactive workloads entailing real-time user requests must be attended to immediately; e.g., instant messages and voice services. In contrast, delay-tolerant workloads such as system updates and data backups are deferrable within a given time interval. This flexibility of delay-tolerant loads enables opportunistic workload management that can be adaptive to the time-varying energy prices and renewables.

Consider a discrete-time scheduling horizon[2] $\mathcal{T} := \{1, \ldots, T\}$. For interactive workloads, let A_j^t denote the arrival rate of service requests at MN j, and a_{ji}^t the workload directed from node j to DC i over slot t. For delay-tolerant workloads, let $\mathcal{Q}_j$ denote the jobs collected by node j, and $\mathcal{Q} := \bigcup_{j=1}^J \mathcal{Q}_j$ with $\mathcal{Q}_i \bigcap \mathcal{Q}_j = \emptyset, \forall i \neq j$, representing the set of all delay-tolerant jobs. The qth delay-tolerant job can be specified by its total demand B_q, and active interval $\mathcal{T}_q := \{S_q, \ldots, E_q\}$, with S_q and E_q denoting the start- and end-time slots. Let $\tilde{b}_{i,q}^t$ denote the amount of qth delay-tolerant job routed from[3] its MN to DC i, and $b_{i,q}^t$ the amount processed by DC i over slot t, respectively; and L_{ji}^t denote the link bandwidth from node j to DC i at time t. As shown in Fig. 2, these quantities must satisfy the following constraints:

$$\sum_{i=1}^{I} a_{ji}^t = A_j^t, \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{1}$$

$$\sum_{t=S_q}^{E_q} \sum_{i=1}^{I} \tilde{b}_{i,q}^t = B_q, \forall q \in \mathcal{Q} \tag{2}$$

$$a_{ji}^t + \sum_{q \in \mathcal{Q}_j} \tilde{b}_{i,q}^t \leq L_{ji}^t, \forall i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T} \tag{3}$$

where (1) ensures that interactive workloads are dispatched once arrived; (2) requires routing each delay-tolerant job before its deadline; and (3) captures the bandwidth limitation of data transfer. Clearly, $L_{ji}^t = 0$ if load transfer from node j to DC i is prohibited; e.g., when MN j is not physically linked with DC i.

[2]For convenience, the slot duration is normalized to unity; thus, the terms "energy" and "power" will be interchangeably used throughout the paper.

[3]Workload distribution delay is ignored here.

In each DC, interactive workloads are processed immediately, while delay-tolerant workloads are deferrable. The unserved portion of delay-tolerant workloads are buffered in separate queues obeying following dynamic recursions

$$Q_{i,q}^{t+1} = Q_{i,q}^{t} - b_{i,q}^{t} + \tilde{b}_{i,q}^{t}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}_q, q \in \mathcal{Q} \tag{4}$$

where $Q_{i,q}^{t}$ is the queue length of qth delay-tolerant job in DC i at the beginning of slot t. For the deadline completion requirements, queue length $Q_{i,q}^{t}$ must satisfy

$$Q_{i,q}^{t} \geq 0, \quad Q_{i,q}^{S_q} = Q_{i,q}^{E_q+1} = 0. \tag{5}$$

The total IT demand of DC i in slot t, is thus given by

$$d_i^t = \sum_{j=1}^{J} a_{ji}^t + \sum_{q \in \mathcal{Q}} b_{i,q}^t, \quad \forall t \in \mathcal{T}. \tag{6}$$

The aforementioned IT system mainly deals with the DC network workload balancing tasks by exploiting the heterogeneous server and bandwidth resources in the cloud. The control variables therein are $\{d_i^t, a_{ji}^t, b_{i,q}^t, \tilde{b}_{i,q}^t\}$, under constraints (1)–(6). In addition, the IT system also closely connects with the underlying power infrastructure through a power supply and demand relationship, which is instructive to detail in the next two subsections.

B. Power Demand Model

The power consumption of a DC generally comes from various sources, but mainly from the running servers and cooling systems [2].

Suppose that each DC i has a set of $\overline{M}_i$ homogeneous servers, so the number of active servers m_i^t at time t should be in the range $\underline{M}_i \leq m_i^t \leq \overline{M}_i$, where $\underline{M}_i$ stands for the minimum number of servers required for providing basic services. The consumption of each server can be generally modeled as a function of its running speed [23]

$$P_{i,s}(s_i^t) = \overline{P}_{i,s}\left(\varrho (s_i^t)^{\upsilon} + 1 - \varrho\right), \quad \varrho \in [0, 1]$$

where $\overline{P}_{i,s}$ denotes the peak power consumption of a server in DC i; $s_i^t \in [0, 1]$ is its actual speed (a.k.a. CPU usage); and parameter υ is typically around 2 for state-of-the-art servers [5]. Clearly, the fraction of peak consumption $1 - \varrho$ denotes the power consumed in the idle state (i.e., $s_i^t = 0$). When the server is in its highest speed $s_i^t = 1$, the actual consumption is $\overline{P}_{i,s}$.

Due to the convexity of $P_{i,s}(s_i^t)$, it readily follows that given a total IT demand d_i^t, uniform allocation of the workloads to each server is most energy efficient [7]. Accordingly, each server is running at a speed $d_i^t/(m_i^t D_i)$ with D_i denoting the server capacity under the required service level agreement (SLA), and the total power consumption in DC i becomes

$$P_{i,IT}(d_i^t, m_i^t) = \frac{\varrho \overline{P}_{i,s} {d_i^t}^2}{m_i^t D_i^2} + (1 - \varrho)\overline{P}_{i,s} m_i^t. \tag{7}$$

Since the number of servers is very large, m_i^t can be relaxed to be a positive real number for simplicity [6].

Along with the increasing density of IT equipment in DCs, a considerable amount of electricity is consumed by the cooling system that generally operates in two modes [4], [24]: outside-air and chilled-water cooling.

The energy usage of *outside-air cooling* is mainly the power consumed by blowers, which can be approximated as a cubic function of the blower speed [25]. As the blower speed under tight control is proportional to the IT consumption $P_{i,IT}$, the outside-air cooling power consumption can be modeled as a convex function of $P_{i,IT}$, namely

$$\mathcal{F}_{i,a}^t(P_{i,IT}) = \kappa_i^t (P_{i,IT})^3, \quad 0 \leq P_{i,IT} \leq \overline{P}_{i,a}^t \tag{8}$$

where $\kappa_i^t > 0$ depends on the temperature difference between the (hot) exhausting air temperature $T_{i,RA}^t$ from the IT racks and the outside air temperature $T_{i,OA}^t$ around DC i at time slot t. The capacity of outside-air cooling in (8) can be modeled as $\overline{P}_{i,a}^t = C(T_{i,RA}^t - T_{i,OA}^t)$, with $C > 0$ proportional to the maximal outside air flow rate.

The *chilled-water cooling* model here is established on the actual measurement of a practical chiller for which the power consumption can be approximated as [26]

$$\mathcal{F}_{i,w}(P_{i,IT}) = \gamma P_{i,IT} \tag{9}$$

where $P_{i,IT}$ is again the IT power consumption in (7) and $\gamma > 0$ is a constant depending on the specific chiller characteristics.

Due to different cooling efficiencies and capacities of the two approaches, for a given $P_{i,IT}$, there is an optimal allocation between outside-air cooling and chiller cooling. Let $P_{i,w}^t$ and $P_{i,a}^t$ denote the amounts of IT power consumption allocated for chiller and outside-air cooling, respectively. With (8) and (9), the optimal cooling power consumption is [24]

$$\begin{aligned}\mathcal{F}_i^t(P_{i,IT}) &= \min_{\substack{0 \leq P_{i,a}^t \leq \overline{P}_{i,a}^t \\ P_{i,a}^t + P_{i,w}^t = P_{i,IT}}} \mathcal{F}_{i,w}(P_{i,w}^t) + \mathcal{F}_{i,a}^t(P_{i,a}^t) \\ &= \min_{0 \leq P_{i,a}^t \leq \overline{P}_{i,a}^t} \gamma [P_{i,IT} - P_{i,a}^t]^+ + \kappa_i^t (P_{i,a}^t)^3. \end{aligned} \tag{10}$$

With a temperature-dependent threshold $P_{i,TH}^t := \min\{\overline{P}_{i,a}^t, \sqrt{\gamma/3\kappa_i^t}\}$, (10) admits a closed-form solution

$$\mathcal{F}_i^t(P_{i,IT}) = \begin{cases} \kappa_i^t \left(P_{i,IT}\right)^3, & P_{i,IT} \leq P_{i,TH}^t \\ \kappa_i^t \left(P_{i,TH}^t\right)^3 + \gamma \left(P_{i,IT} - P_{i,TH}^t\right), & \text{otherwise.} \end{cases}$$

For notational convenience, let $P_i^t(d_i^t, m_i^t)$ include the server and cooling consumptions in DC i per slot t as $P_i^t(d_i^t, m_i^t) := \mathcal{F}_i^t(d_i^t, m_i^t) + P_{i,IT}(d_i^t, m_i^t)$, and consider the following lemma.

Lemma 1: Function $P_i^t(d_i^t, m_i^t)$ is jointly convex in $\{d_i^t, m_i^t\}$.

Proof: See Appendix A. ■

C. Power Supply Model

A rapidly increasing use of microgrids characterizes the transformative change from our aging power grid to a smart grid over the last decade [27]. While the traditional geographical

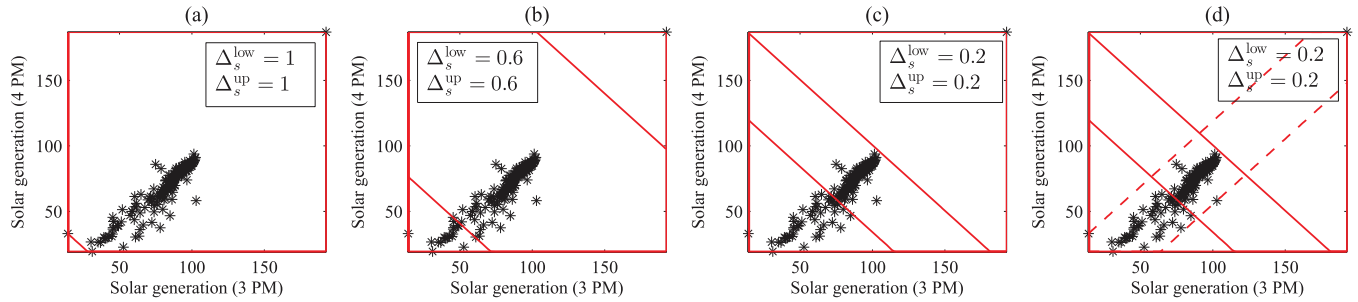

Fig. 3. A diagram of uncertainty sets of the CAISO solar generations. Figs. 3(a)–3(c) are for $\mathcal{E}_i$, and Fig. 3(d) is for $\check{\mathcal{E}}_i$. Points ($*$) denote the generation samples from historical data. Red lines represent the boundaries of the polyhedral uncertainty sets (13) and (14).

workload balancing operates separately from the local power balancing, recently advocated dynamic pricing and demand response programs motivate the interactions between them. In this context, we consider each DC to be supplied by a RES-integrated microgrid consisting of a conventional generator (CG) (e.g., a fuel generator), an on-site renewable generator (RG) (e.g., a wind or solar generator), and an energy storage unit (e.g., a battery).

Let $P_{i,g}^t$ denote the energy output of the CG in DC i per slot t, which is upper bounded by $\overline{P}_{i,g}$; that is,

$$0 \leq P_{i,g}^t \leq \overline{P}_{i,g}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}. \tag{11}$$

The change of CG energy output in two consecutive slots is bounded by the following so-termed ramping constraints:

$$P_{i,g}^t - P_{i,g}^{t-1} \leq R_i^{\text{up}},\ P_{i,g}^{t-1} - P_{i,g}^t \leq R_i^{\text{dw}},\ \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{12}$$

where R_i^{up} and R_i^{dw} are the ramping-up and ramping-down limits of CG at DC i.

Consider now the RES vector $\mathbf{e}_i := [E_i^1, \ldots, E_i^T]'$ generated at DC i across all slots. Due to the unpredictable and intermittent nature of RES, $\mathbf{e}_i$ is unknown a priori. In general, uncertain quantities can be modeled by postulating either an underlying probability distribution or an uncertainty region. Probability distributions (possibly mixed discrete/continuous) of the RES generation are seldom available in practice. Although (non-)parametric approaches can be used to learn these distributions, the processes can be very complicated due to the spatio-temporal correlations incurred by various meteorological factors [28]. On the contrary, the proposed method of postulating an uncertainty region provides the decision maker with ranges instead of point forecasts, which is essentially a distribution-free deterministic set and *robust* to prediction errors.

The actual RES generation $\mathbf{e}_i$ is assumed to lie in an uncertainty set $\mathcal{E}_i$, which can be obtained via forecasting or inference using historical measurements. In particular, the following polyhedral uncertainty set is considered (see also [29], [30])

$$\mathcal{E}_i := \Bigg\{ \mathbf{e}_i | \underline{E}_i^t \leq E_i^t \leq \overline{E}_i^t,\ \mathcal{T} = \bigcup_{s=1}^{S} \mathcal{T}_{i,s},\ \Delta_s^{\text{low}} \underline{E}_{\mathcal{T}_{i,s}} + (1 - \Delta_s^{\text{low}}) E_{\mathcal{T}_{i,s}}^{\text{avg}} \leq \sum_{t \in \mathcal{T}_{i,s}} E_i^t \leq (1 - \Delta_s^{\text{up}}) E_{\mathcal{T}_{i,s}}^{\text{avg}} + \Delta_s^{\text{up}} \overline{E}_{\mathcal{T}_{i,s}} \Bigg\} \tag{13}$$

where $\underline{E}_i^t$ ($\overline{E}_i^t$) denotes the lower (upper) bound on the actual E_i^t; $\mathcal{T}$ is partitioned into consecutive but non-overlapping sub-horizons $\mathcal{T}_{i,s}$, $s = 1, \ldots, S$; $E_{\mathcal{T}_{i,s}}^{\text{avg}}$ is the sample average of total renewables $\sum_{t \in \mathcal{T}_{i,s}} E_i^t$; the total renewables over $\mathcal{T}_{i,s}$ are bounded by $\underline{E}_{\mathcal{T}_{i,s}}$ and $\overline{E}_{\mathcal{T}_{i,s}}$; and the parameter $\Delta_s^{\text{low}}(\Delta_s^{\text{up}}) \in [0, 1]$ represents the *level of robustness* for the lower (upper) bound of the sth sub-horizon. Note that all the aforementioned statistics can be directly obtained using real RES generations from public sources.

Clearly, parameters Δ_s^{low}, Δ_s^{up} trade off robustness for conservatism of the resultant solutions. Based on the CAISO solar generations during Mar. 1 – Oct. 30, 2012 [31], examples are given in Fig. 3 for $|\mathcal{T}_{i,s}| = 2$; e.g., each sub-horizon consists of two time slots. When the robustness levels Δ_s^{low} and Δ_s^{up} are high, the uncertainty set is large, which includes most of historical samples thus usually leading to conservative solutions when the degree of uncertainty is high; if Δ_s^{low} and Δ_s^{up} are low, some of samples are excluded from the uncertainty set, which reduces the robustness of the resultant solutions.

Here it is also instructive to point out that by further capturing the maximum variation of RES over two consecutive slots $\overline{\partial E}_i^t$, a more accurate polyhedral set can be written as (see Fig. 3(d))

$$\check{\mathcal{E}}_i := \Bigg\{ \mathbf{e}_i | \underline{E}_i^t \leq E_i^t \leq \overline{E}_i^t,\ |E_i^t - E_i^{t-1}| \leq \overline{\partial E}_i^t,\ \mathcal{T} = \bigcup_{s=1}^{S} \mathcal{T}_{i,s},\ \Delta_s^{\text{low}} \underline{E}_{\mathcal{T}_{i,s}} + (1 - \Delta_s^{\text{low}}) E_{\mathcal{T}_{i,s}}^{\text{avg}} \leq \sum_{t \in \mathcal{T}_{i,s}} E_i^t \leq (1 - \Delta_s^{\text{up}}) E_{\mathcal{T}_{i,s}}^{\text{avg}} + \Delta_s^{\text{up}} \overline{E}_{\mathcal{T}_{i,s}} \Bigg\}. \tag{14}$$

This modification allows a decision maker to reduce the uncertainty region, without missing many potential samples. It makes sense intuitively since in a short time scale (e.g., 15 min, or, an hour), the RES are highly correlated over successive slots. Although the aforementioned practical models only capture RES uncertainty across the scheduling horizons per DC, our proposed approach could be easily extended to include joint spatio-temporal uncertainty models.

To mitigate the variability of RES, energy storage devices are recently considered so as to store the surplus renewables for later shortage [32]. We consider a storage unit with finite capacity $\overline{C}_i$, and let C_i^0 and C_i^t denote the initial energy level

of the storage unit in DC i at the beginning of time slot t. Since storage devices become unreliable with high depth-of-discharge[4] (DoD), a nonzero minimum level $\underline{C}_i$ can avoid high DoD. Such a level could also support the DC operation in the event of a grid outage. Let $P^t_{\mathrm{i,ch}}$ and $P^t_{\mathrm{i,dis}}$ denote the amounts of power charging and discharging the storage unit (battery) in DC i at slot t. Due to AC/DC power conversion during the (dis-)charging process, the power conversion losses need to be accounted for by the (dis-)charging efficiency $\delta \in (0, 1]$. In addition, dissipation losses due to battery energy leakages are captured by the efficiency coefficient $\eta \in (0, 1]$, which renders a decreasing energy level even if there is no (dis-)charging operation. In short, the energy storage unit can be compactly described as

$$C_i^{t+1} = \eta C_i^t + \delta P^t_{\mathrm{i,ch}} - \frac{P^t_{\mathrm{i,dis}}}{\delta}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{15}$$

$$\underline{C}_i \le C_i^t \le \overline{C}_i, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{16}$$

$$0 \le P^t_{\mathrm{i,dis}} \le \overline{P}_{\mathrm{i,dis}};\ 0 \le P^t_{\mathrm{i,ch}} \le \overline{P}_{\mathrm{i,ch}}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{17}$$

where the bounds $\overline{P}_{\mathrm{i,dis}}$ and $\overline{P}_{\mathrm{i,ch}}$ on (dis-)charging amounts are dictated by physical limits.

Let $P^t_{\mathrm{i,out}}$ denote the total energy consumption of DC i per slot t including the IT operating consumption, cooling power consumption, and battery charged power; that is, [cf. Lemma 1]

$$P^t_{\mathrm{i,out}} = P_i^t + P^t_{\mathrm{i,ch}}. \tag{18}$$

Likewise, the total energy supply $P^t_{\mathrm{i,in}}$ in DC i per slot t is given by

$$P^t_{\mathrm{i,in}} = P^t_{i,g} + E_i^t + P^t_{\mathrm{i,dis}}. \tag{19}$$

Under constraints (11)–(19), the power supply optimization variables are CG and battery power amounts $\{P^t_{i,g}, P^t_{\mathrm{i,dis}}, P^t_{\mathrm{i,ch}}, C_i^t\}$.

D. Cost-Revenue Model

In addition to the internal energy resources (namely, CG, RG, storage unit), DCs can resort to the main grid market in an on-demand manner. With a two-way energy trading facility, each DC can buy energy from external energy markets in the case of a deficit ($P^t_{\mathrm{i,out}} > P^t_{\mathrm{i,in}}$), or, sell energy to the markets in the case of a surplus ($P^t_{\mathrm{i,out}} < P^t_{\mathrm{i,in}}$). Clearly, the shortage energy that needs to be purchased by the DC is $[P^t_{\mathrm{i,out}} - P^t_{\mathrm{i,in}}]^+$; while the surplus energy that can be sold is $[P^t_{\mathrm{i,in}} - P^t_{\mathrm{i,out}}]^+$. Note that both the shortage and surplus energies are non-negative, and at most one of them is positive at any slot t.

Suppose that the energy can be purchased from the wholesale electricity market around DC i in period t at price α_i^t, while the energy is sold at price β_i^t. Notwithstanding, we shall always set $\alpha_i^t \ge \beta_i^t$ to avoid less relevant buy-and-sell activities of the DC for profit. For DC i, the *worst-case transaction cost* for the whole scheduling horizon is defined as

$$G_i(\{P^t_{\mathrm{i,out}}\}, \{P^t_{\mathrm{i,in}}\}) := \max_{\mathbf{e}_i \in \mathcal{E}_i} \sum_{t=1}^{T} \alpha_i^t \left[P^t_{\mathrm{i,out}} - P^t_{\mathrm{i,in}}\right]^+ - \beta_i^t \left[P^t_{\mathrm{i,in}} - P^t_{\mathrm{i,out}}\right]^+ \tag{20}$$

which is the point-wise maximum over any realization of the random RES generation in the uncertainty set.

In addition, let function $G_{C_i}(P^t_{i,g})$ denote the cost of CG at DC i in slot t, which is convex piecewise linear or smooth quadratic [33]. The revenue considered here comes from processing delay-tolerant workloads. Specifically, for the qth job, the revenue earned per slot t can be generally modeled as a concave function $U_q^t(b^t_{i,q})$, which reflects the diminishing marginal sensitivity of end users to the increasing gains.

III. Robust Workload and Energy Management

Based on the practical models in Section II, we pursue in this section a robust workload and energy management approach for the considered DC network. Over the scheduling horizon $\mathcal{T}$, the system operator per MN performs an (e.g. hour-) ahead-of-time schedule to optimize workloads routing $\{a^t_{ji}, \tilde{b}^t_{i,q}\}$, while the system operator in each DC optimizes servers and workloads scheduling $\{m_i^t, b^t_{i,q}\}$, CG generation $\{P^t_{i,g}\}$, and battery (dis-)charging energy $\{P^t_{\mathrm{i,ch}}, P^t_{\mathrm{i,dis}}\}$. The optimal management strategy minimizes the worst-case net cost $\tilde{\Psi}$, which includes the worst-case transaction cost, the CG cost and the revenue of delay-tolerant workloads, subject to DC operating and power supply constraints. Note that the worst-case net cost here is the maximum net cost for any realization of the random RES generation in the uncertainty set.

With $\tilde{\mathbf{x}}$ collecting all the optimization variables $\{a^t_{ji}, b^t_{i,q}, \tilde{b}^t_{i,q}, d_i^t, m_i^t, P^t_{i,g}, P^t_{\mathrm{i,ch}}, P^t_{\mathrm{i,dis}}, C_i^t\}$, we wish to solve

$$\tilde{\Psi}^* := \min_{\tilde{\mathbf{x}}} \sum_{i=1}^{I} G_i(\{P^t_{\mathrm{i,out}}\}, \{P^t_{\mathrm{i,in}}\}) + \sum_{t=1}^{T}\sum_{i=1}^{I} \left(G_{C_i}(P^t_{i,g}) - \sum_{q \in \mathcal{Q}} U_q^t(b^t_{i,q}) \right) \tag{21a}$$

subject to:

$$C_i^{t+1} = \eta C_i^t + \delta P^t_{\mathrm{i,ch}} - \frac{P^t_{\mathrm{i,dis}}}{\delta}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21b}$$

$$\underline{C}_i \le C_i^t \le \overline{C}_i, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21c}$$

$$0 \le P^t_{\mathrm{i,dis}} \le \overline{P}_{\mathrm{i,dis}}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21d}$$

$$0 \le P^t_{\mathrm{i,ch}} \le \overline{P}_{\mathrm{i,ch}}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21e}$$

$$0 \le P^t_{i,g} \le \overline{P}_{i,g}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21f}$$

$$P^t_{i,g} - P^{t-1}_{i,g} \le R_i^{\mathrm{up}}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21g}$$

$$P^{t-1}_{i,g} - P^t_{i,g} \le R_i^{\mathrm{dw}}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21h}$$

$$\underline{M}_i \le m_i^t \le \overline{M}_i, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21i}$$

$$0 \le d_i^t \le m_i^t D_i, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21j}$$

[4]DoD is the percentage of maximum charge removed during a discharge cycle.

$$a_{ji}^t + \sum_{q\in\mathcal{Q}_j} \tilde{b}_{i,q}^t \le L_{ji}^t, \quad \forall j \in \mathcal{J}, i \in \mathcal{I}, t \in \mathcal{T} \tag{21k}$$

$$\sum_{i=1}^{I} a_{ji}^t = A_j^t, \quad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{21l}$$

$$\sum_{t=S_q}^{E_q} \sum_{i=1}^{I} \tilde{b}_{i,q}^t = B_q, \quad \forall q \in \mathcal{Q} \tag{21m}$$

$$Q_{i,q}^{t+1} = Q_{i,q}^t - b_{i,q}^t + \tilde{b}_{i,q}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, q \in \mathcal{Q} \tag{21n}$$

$$P_{\text{i,out}}^t = P_i^t(d_i^t, m_i^t) + P_{\text{i,ch}}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21o}$$

$$P_{\text{i,in}}^t = P_{i,g}^t + E_i^t + P_{\text{i,dis}}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21p}$$

$$Q_{i,q}^{S_q} = Q_{i,q}^{E_q+1} = 0, Q_{i,q}^t \ge 0, \quad \forall i \in \mathcal{I}, q \in \mathcal{Q}, t \in \mathcal{T}_q \tag{21q}$$

$$d_i^t = \sum_{j=1}^{J} a_{ji}^t + \sum_{q\in\mathcal{Q}} b_{i,q}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{21r}$$

$$0 \le b_{i,q}^t \le B_q, \quad \forall t \in \mathcal{T}_q;\ b_{i,q}^t = 0,\ \forall i \in \mathcal{I}, q \in \mathcal{Q}, t \notin \mathcal{T}_q \tag{21s}$$

$$0 \le \tilde{b}_{i,q}^t \le B_q, \quad \forall t \in \mathcal{T}_q;\ \tilde{b}_{i,q}^t = 0,\ \forall i \in \mathcal{I}, q \in \mathcal{Q}, t \notin \mathcal{T}_q \tag{21t}$$

$$0 \le a_{ji}^t \le A_j^t, \quad \forall j \in \mathcal{J}, i \in \mathcal{I}, t \in \mathcal{T}. \tag{21u}$$

It is worth mentioning that thanks to the worst-case transaction cost $G_i(\{P_{\text{i,out}}^t\}, \{P_{\text{i,in}}^t\})$, the objective of (21) has an implicit min-max form, and the RES induced randomness can be eliminated; thus, (21) contains only deterministic variables. However, the objective of (21) is to minimize a point-wise maximum function, which is generally not differentiable when the maximum is attained by more than one solution. In addition, since P_i^t is a nonlinear function with respect to $\{d_i^t, m_i^t\}$, then (21o) are nonlinear equality constraints representing a nonconvex feasible set [34, Chap. 4]. Thus, problem (21) is nonsmooth and nonconvex, which is hard to be handled by existing solvers. To turn (21) into a tractable form, a reformulation relying on epigraph-based relaxation is pursued next.

A. Convex Reformulation

Define $\psi_i^t := (\alpha_i^t - \beta_i^t)/2$, $\phi_i^t := (\alpha_i^t + \beta_i^t)/2$, and $R_i^t = P_i^t + P_{\text{i,ch}}^t - P_{\text{i,dis}}^t - P_{i,g}^t$; and then rewrite (20) as

$$G_i(\{R_i^t\}) = \max_{\mathbf{e}_i\in\mathcal{E}_i} \sum_{t=1}^{T} \left(\psi_i^t |R_i^t - E_i^t| + \phi_i^t \left(R_i^t - E_i^t\right)\right). \tag{22}$$

In order to convexify the (21) and facilitate a distributed implementation, we define $\mathbf{x}$ collecting all the optimization variables $\{a_{ji}^t, b_{i,q}^t, \tilde{b}_{i,q}^t, m_i^t, d_i^t, R_i^t, P_{i,g}^t, P_{\text{i,ch}}^t, P_{\text{i,dis}}^t, C_i^t\}$, and rewrite (21) as (RWEM):

$$\Psi^* := \min_{\mathbf{x}} \sum_{t=1}^{T}\sum_{i=1}^{I} \left(G_{C_i}(P_{i,g}^t) - \sum_{q\in\mathcal{Q}} U_q^t(b_{i,q}^t)\right) + \sum_{i=1}^{I} G_i(\{R_i^t\}) \tag{23a}$$

subject to:

$$(21b) - (21m), (21s) - (21u) \tag{23b}$$

$$R_i^t \ge P_i^t(d_i^t, m_i^t) + P_{\text{i,ch}}^t - P_{\text{i,dis}}^t - P_{i,g}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{23c}$$

$$\sum_{\tau=S_q}^{E_q} \tilde{b}_{i,q}^\tau = \sum_{\tau=S_q}^{E_q} b_{i,q}^\tau,\ \sum_{\tau=S_q}^{t} \tilde{b}_{i,q}^\tau \ge \sum_{\tau=S_q}^{t} b_{i,q}^\tau, \quad \forall i, q \in \mathcal{Q}, t \in [S_q, E_q - 1] \tag{23d}$$

$$d_i^t = \sum_{j=1}^{J} a_{ji}^t + \sum_{q\in\mathcal{Q}} b_{i,q}^t, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}. \tag{23e}$$

Convexity of the worst-case net cost Ψ is established in the following proposition.

Proposition 1: If $\alpha_i^t \ge \beta_i^t$ holds for all $i \in \mathcal{I}$ and $t \in \mathcal{T}$, then RWEM problem (23) is convex and strong duality holds.

Proof: See Appendix B. ■

With $\tilde{\mathbf{x}}^*$ and $\mathbf{x}^*$ denoting the optimal solutions for (21) and (23), we arrive at the following claim.

Proposition 2: Problem (21) is equivalent to problem (23) in the sense that $\Psi^* = \tilde{\Psi}^*$, and $\mathbf{x}^* = \tilde{\mathbf{x}}^*$.

Proof: See Appendix C. ■

B. Lagrange Dual Approach

Notice that constraints (23c)–(23e) couple variables across MNs, DCs, workloads, and the RES, so a system operator over the entire network is essential to collect all the information and solve the problem in a centralized way, which may not be feasible in an Internet-scale network [35]. However, since (23) is a convex problem [cf. Proposition 1], a Lagrange dual approach can be developed to efficiently find its optimal dual solution with zero duality gap in a decentralized manner [36]. Let $\{\pi_i^t\}$, $\{\lambda_{i,q}^t\}$ and $\{\nu_i^t\}$ denote the Lagrange multipliers associated with the constraints (23c)–(23e). For notational convenience, let $\lambda_{i,q}^t = 0, \forall i, q \in \mathcal{Q}, t \notin \mathcal{T}_q$. And with $\boldsymbol{\varpi}$ collecting all the Lagrange multipliers, the partial Lagrangian function of (23) is

$$\begin{aligned}\mathcal{L}(\mathbf{x}, \boldsymbol{\varpi}) :=& \sum_{i=1}^{I}\left[G_i(\{R_i^t\}) + \sum_{t=1}^{T} G_{C_i}(P_{i,g}^t) - \sum_{q\in\mathcal{Q}} U_q^t(b_{i,q}^t)\right] \\ &+ \sum_{i=1}^{I}\sum_{t=1}^{T} \nu_i^t \left(d_i^t - \sum_{j=1}^{J} a_{ji}^t - \sum_{q\in\mathcal{Q}} b_{i,q}^t\right) \\ &+ \sum_{i=1}^{I}\sum_{q\in\mathcal{Q}}\sum_{t=1}^{T} \lambda_{i,q}^t \left(\sum_{\tau=S_q}^{t} b_{i,q}^\tau - \sum_{\tau=S_q}^{t} \tilde{b}_{i,q}^\tau\right) \\ &+ \sum_{i=1}^{I}\sum_{t=1}^{T} \pi_i^t \left(P_i^t + P_{\text{i,ch}}^t - P_{\text{i,dis}}^t - P_{i,g}^t - R_i^t\right).\end{aligned}$$

If $\mathcal{X}$ denotes the set of $\mathbf{x}$ satisfying constraints (21b)–(21m), and (21s)–(21u), the Lagrange dual function is given by

$$\mathcal{D}(\boldsymbol{\varpi}) := \min_{\mathbf{x}\in\mathcal{X}} \mathcal{L}(\mathbf{x}, \boldsymbol{\varpi}) \tag{24}$$

and the dual problem of (23) is

$$\begin{aligned}\max\ & \mathcal{D}(\{\pi_i^t\},\{\lambda_{i,q}^t\},\{\nu_i^t\})\\ \text{s. t.}\ & \pi_i^t\geq 0,\nu_i^t\in\mathbb{R},\forall i\in\mathcal{I},t\in\mathcal{T}\\ & \lambda_{i,q}^t\geq 0,\forall i\in\mathcal{I},q\in\mathcal{Q},t\in[S_q,E_q-1]\\ & \lambda_{i,q}^t\in\mathbb{R},\forall i\in\mathcal{I},q\in\mathcal{Q},t=E_q.\end{aligned}\tag{25}$$

For the dual problem (25), standard subgradient iterations can be employed to obtain the optimal $\boldsymbol{\varpi}^*$, namely

$$\pi_i^t(k+1)=[\pi_i^t(k)+\mu g_{\pi_i^t}(k)]^+,\ \ \forall i\in\mathcal{I},t\in\mathcal{T}\tag{26a}$$

$$\lambda_{i,q}^t(k+1)=[\lambda_{i,q}^t(k)+\mu g_{\lambda_{i,q}^t}(k)]^+,\ \ \forall i,q,t\in[S_q,E_q-1]$$

$$\lambda_{i,q}^t(k+1)=\lambda_{i,q}^t(k)+\mu g_{\lambda_{i,q}^t}(k),\ \ \forall i,q,t=E_q\tag{26b}$$

$$\nu_i^t(k+1)=\nu_i^t(k)+\mu g_{\nu_i^t}(k),\ \ \forall i\in\mathcal{I},t\in\mathcal{T}\tag{26c}$$

where k is the iteration index, and $\mu>0$ is a constant stepsize, and $\{g_{\pi_i^t}(k),g_{\lambda_{i,q}^t}(k),g_{\nu_i^t}(k)\}$ are the subgradients of (24) with respect to the Lagrange multipliers. Specifically, we have

$$g_{\pi_i^t}(k)=P_i^t(k)+P_{\mathrm{i,ch}}^t(k)-P_{\mathrm{i,dis}}^t(k)-P_{i,g}^t(k)-R_i^t(k)\tag{27a}$$

$$g_{\lambda_{i,q}^t}(k)=\sum_{\tau=S_q}^{t}b_{i,q}^{\tau}(k)-\sum_{\tau=S_q}^{t}\tilde{b}_{i,q}^{\tau}(k)\tag{27b}$$

$$g_{\nu_i^t}(k)=d_i^t(k)-\sum_{j=1}^{J}a_{ji}^t(k)-\sum_{q\in\mathcal{Q}}b_{i,q}^t(k)\tag{27c}$$

where primal variables $\mathbf{x}(k)$ can be obtained as

$$\begin{aligned}&\left\{a_{ji}^t(k),\tilde{b}_{i,q}^t(k)\right\}_{i\in\mathcal{I},t\in\mathcal{T},q\in\mathcal{Q}_j}\\ &\in\arg\min_{\{a_{ji}^t,\tilde{b}_{i,q}^t\}}\sum_{t=1}^{T}\sum_{i=1}^{I}\left[-a_{ji}^t\nu_i^t(k)-\sum_{q\in\mathcal{Q}_j}\tilde{b}_{i,q}^t\sum_{\tau=t}^{T}\lambda_{i,q}^{\tau}(k)\right]\\ &\qquad\text{s. t. (21k)}-\text{(21m)},\text{(21t)}-\text{(21u)}\end{aligned}\tag{28}$$

$$\begin{aligned}&\{b_{i,q}^t(k)\}_{t\in\mathcal{T},q\in\mathcal{Q}}\\ &\in\arg\min_{\{b_{i,q}^t\}}\sum_{t=1}^{T}\left[b_{i,q}^t\left(\sum_{\tau=t}^{T}\lambda_{i,q}^{\tau}(k)-\nu_i^t(k)\right)-U_q^t(b_{i,q}^t)\right]\\ &\qquad\text{s. t. (21s)}\end{aligned}\tag{29}$$

and (30), shown at the bottom of this page.

The subproblems (28)–(29) are linear programs (LPs) over $\{a_{ji}^t,\tilde{b}_{i,q}^t,b_{i,q}^t\}_{t=1}^T$; hence, they can be optimally solved using available efficient LP solvers. Due to the convexity of $G_i(\{R_i^t\})$, the subproblems (30) are convex per DC i. However, since $G_i(\{R_i^t\})$ is non-differentiable due to the absolute value operator and the maximization over $\mathbf{e}_i\in\mathcal{E}_i$, (30) still challenges existing solvers. To address this, consider splitting (30) into two subproblems, one being (31), shown at the bottom of this page, and the second solving

$$\{R_i^t(k)\}_{t=1}^T\in\arg\min_{\{\underline{R}_i\leq R_i^t\leq\overline{R}_i\}}G_i(\{R_i^t\})-\sum_{t=1}^{T}\pi_i^t(k)R_i^t.\tag{32}$$

Note that because of the exact relaxation [cf. Proposition 2], $\underline{R}_i$ and $\overline{R}_i$ are lower and upper bounds of the right hand side of (23c). Depending on the function $G_{C_i}(P_{i,g}^t)$, subproblem (31) is either an LP or a quadratic program. Hence, the optimal solution can be obtained by existing solvers. And for nonsmooth subproblems (32), a standard subgradient iteration can be employed to obtain the optimal solution as

$$R_i^t(\ell+1)=R_i^t(\ell)-\mu(\ell)g_{R_i^t}(\ell),\quad\forall t\in\mathcal{T}\tag{33}$$

where ℓ denotes iteration index, and $\{\mu(\ell)\}$ is a non-summable but square-summable stepsize sequence; while the partial subgradient of $G_i(\{R_i^t\})$ with respect to R_i^t is obtained as

$$\begin{aligned}g_{R_i^t}(\ell)&:=\frac{\partial\left(G_i(\{R_i^t\})-\sum_{t=1}^{T}\pi_i^t(k)R_i^t\right)}{\partial R_i^t}\\ &=\begin{cases}\alpha_i^t-\pi_i^t(k), & \text{if } R_i^t(\ell)\geq E_i^{t*}(\ell)\\ \beta_i^t-\pi_i^t(k), & \text{if } R_i^t(\ell)<E_i^{t*}(\ell)\end{cases}\end{aligned}\tag{34}$$

where $\mathbf{e}_i^*(\ell):=[E_i^{1*}(\ell),\ldots,E_i^{T*}(\ell)]'$ for the given $\{R_i^t(\ell)\}$ is found using

$$\mathbf{e}_i^*\in\arg\max_{\mathbf{e}_i\in\mathcal{E}_i}\sum_{t=1}^{T}\left(\psi_i^t|R_i^t(\ell)-E_i^t|+\phi_i^t(R_i^t(\ell)-E_i^t)\right).\tag{35}$$

It can be seen that the objective function in (35) is convex in $\mathbf{e}_i$ under the condition $\alpha_i^t\geq\beta_i^t,\ \forall t\in\mathcal{T}$. However, computing where the maximum of a convex function is attained can be *NP-hard*, in general. Fortunately, the globally optimal solution is attainable at the extreme points of $\mathcal{E}_i$ for convex maximization [37, Sec. 2.4]. Leveraging the polyhedral structure of $\mathcal{E}_i$,

$$\begin{aligned}&\{R_i^t(k),m_i^t(k),P_{i,\mathrm{ch}}^t(k),P_{i,\mathrm{dis}}^t(k),C_i^t(k),P_{i,g}^t(k),d_i^t(k)\}_{t=1}^T\\ &\in\arg\min_{\left\{m_i^t,P_{i,\mathrm{ch}}^t,P_{i,\mathrm{dis}}^t,R_i^t,P_{i,g}^t,d_i^t\right\}}G_i(\{R_i^t\})+\sum_{t=1}^{T}\left[\nu_i^t(k)d_i^t+G_{C_i}(P_{i,g}^t)+\pi_i^t(k)(P_i^t+P_{i,\mathrm{ch}}^t-P_{i,\mathrm{dis}}^t-P_{i,g}^t-R_i^t)\right]\\ &\qquad\text{s.t.}\quad\text{(21b)}-\text{(21j)}\end{aligned}\tag{30}$$

$$\begin{aligned}&\{m_i^t(k),P_{i,\mathrm{ch}}^t(k),P_{i,\mathrm{dis}}^t(k),C_i^t(k),P_{i,g}^t(k),d_i^t(k)\}_{t=1}^T\\ &\in\arg\min_{\{P_{i,\mathrm{ch}}^t,P_{i,\mathrm{dis}}^t,m_i^t,P_{i,g}^t,d_i^t\}}\sum_{t=1}^{T}\left[\pi_i^t(k)(P_i^t+P_{i,\mathrm{ch}}^t-P_{i,\mathrm{dis}}^t-P_{i,g}^t)+\nu_i^t(k)d_i^t+G_{C_i}(P_{i,g}^t)\right]\\ &\qquad\text{s.t.}\quad\text{(21b)}-\text{(21j)}\end{aligned}\tag{31}$$

Algorithm 1. Subgradient iteration for solving (32)

1: **Initialize:** Generate all the vertices of the polyhedral uncertainty set $\mathcal{E}_i$; choose a proper $\{R_i^t(0)\}$ and stepsize sequence $\mu(\ell)$
2: **repeat** $\ell = 0, 1, 2 \ldots$
3: Evaluate all the vertices in $\mathcal{E}_i$ and find $\mathbf{e}_i^*(\ell)$ in (35)
4: Calculate subgradients via (34)
5: Update $\{R_i^t(\ell)\}$ via (33)
6: **until** Convergence

Algorithm 2. Distributed workload and energy management

1: **Initialize:** Choose a proper $\boldsymbol{\varpi}(0)$ and stepsize μ
2: **repeat** $k = 0, 1, 2 \ldots$
3: Each DC obtains $\{b_{i,q}^t(k), R_i^t(k), m_i^t(k), P_{\mathrm{i,ch}}^t(k), P_{\mathrm{i,dis}}^t(k), C_i^t(k), P_{i,g}^t(k), d_i^t(k)\}$ by solving (29) and (31)–(32) separately
4: Each MN solves (28) and sends $\{a_{ji}^t(k), \tilde{b}_{i,q}^t(k)\}$ to each DC
5: DCs update $\boldsymbol{\varpi}(k)$ via (26) and send them to MNs
6: Run averages to recover primal variables via (36)
7: **until** Convergence

we utilize an efficient vertex enumerating algorithm to evaluate the objective in (35), and obtain $\mathbf{e}_i^*$ directly; see Algorithm 1. Although the number of vertices may increase exponentially with the number of variables and constraints, all the vertices of $\mathcal{E}_i$ need be generated only once before running Algorithm 1, which means that Algorithm 1 is computational affordable. In fact, our simulations in Section IV will corroborate that the vertex generating procedure can be completed within several seconds.

C. Optimality and Distributed Implementation

For the subgradient iterations (33), if a diminishing stepsize satisfying (i) $\sum_{\ell=0}^{\infty} \mu(\ell) = \infty$, and (ii) $\sum_{\ell=0}^{\infty} \mu(\ell)^2 < \infty$ is adopted, the sequence (33) converges as $\ell \to \infty$ to the optimal $\{R_i^t(k)^*\}$ [37]. As a constant stepsize μ is used in (26), the subgradient iterations will converge to a neighborhood of the optimal solution $\boldsymbol{\varpi}^*$ [37]. The size of the neighborhood is proportional to the stepsize μ. Since the objective (23a) is not strictly convex, running averages of the primal sequence $\{\mathbf{x}(k)\}$ can be used to recover the optimal primal solutions, which are given by

$$\overline{\mathbf{x}}(k) = \frac{1}{k}\mathbf{x}(k-1) + \frac{k-1}{k}\overline{\mathbf{x}}(k-1), \quad \forall k \tag{36}$$

where $\overline{\mathbf{x}}(k)$ is the average of all primal solutions up to iteration $k-1$. Since set $\mathcal{X}$ is convex and $\mathbf{x}(i) \in \mathcal{X}, 0 \le i \le k-1$, it turns out that $\overline{\mathbf{x}}(k)$ is a feasible point in $\mathcal{X}$. In addition, it can be shown that $\overline{\mathbf{x}}(k)$ is also asymptotically feasible for primal problem (23) [38].

It is also worth noting that RWEM can afford a distributed implementation, where optimization tasks are distributed among MNs and individual DCs; see Algorithm 2. In RWEM, dual variable updates (26) are all implemented at each DC locally. Subproblem (28) is solved by each MN operator, while each DC operator solves subproblems (29) and (31)–(32). To make these distributed implementations possible, a bidirectional message passing between MNs and DCs is necessary. At every iteration, workload routing variables $\{a_{ji}^t(k), \tilde{b}_{i,q}^t(k)\}$ are sent from each node to each DC, while the dual variables $\{\lambda_{i,q}^t\}$ and $\{\nu_i^t\}$ are fed back to each MN in turn to solve (28). Note that instead of real-time power and workload schedules, only Lagrange multipliers are sent back to MNs, which could be further leveraged to enhance privacy-preserving operations.

In addition, the worst-case complexity of solving a general convex program is on the order $\mathcal{O}\left(\max\{N_c, N_v\}^4 \sqrt{N_v} \log(1/\epsilon)\right)$, where N_v, N_c are the total number of variables and constraints, and $\epsilon > 0$ is the given accuracy [39]. Hence, solving subproblems (28)–(29), (31)–(32) in a distributed fashion incurs a markedly lower complexity than directly tackling (23) in a centralized fashion. Faster implementations are possible if we further decentralize (31), and let operators of conventional generation, storage units, delay-tolerant workloads, solve subproblems separately.

TABLE I
POWER SUPPLY PARAMETERS. THE UNITS ARE KW

DC i	$\overline{P}_{i,g}$	$\overline{C}_i$	$\underline{C}_i$	$\overline{P}_{i,\mathrm{ch}}$	$\overline{P}_{i,\mathrm{dis}}$	η	δ	R_i^{up}	R_i^{dw}
1	30	5	40	20	20	0.95	0.95	24	24
2	30	5	40	20	20	0.95	0.95	24	24
3	30	5	40	20	20	0.95	0.95	24	24
4	30	5	40	20	20	0.95	0.95	24	24

IV. NUMERICAL EVALUATION

In this section, results of simulated tests are presented to demonstrate the merits of the proposed approach.

A. Experiment Setup

The Matlab-based modeling package CVX 2.1 [40] is used to solve the optimization problems involved. The DC network includes 4 DCs and 4 MNs uniformly located in the eastern, central, mountain and western parts of the US. Each DC is connected to a microgrid, of which the power supply parameters are listed in Table I. A polyhedral uncertainty set (13) with a single sub-horizon (no partition, $\mathcal{T}_{i,s} = \mathcal{T}$) is considered for the RES. The upper/lower limits $\{\overline{E}_i^t, \underline{E}_i^t, \overline{E}_{\mathcal{T}_{i,s}}, \underline{E}_{\mathcal{T}_{i,s}}\}$ and the average RES over scheduling horizon $E_{\mathcal{T}_{i,s}}^{\mathrm{avg}}$ were rescaled from the CAISO solar generations during Mar. 1-Oct. 30, 2012; see [31] for detailed description. The upper and lower levels of robustness are set by default to $\Delta_s^{\mathrm{up}} = \Delta_s^{\mathrm{low}} = 1$. Table II lists the energy purchase prices α_i^t, which are obtained by scaling the hourly electricity prices of the New York City [41]. The selling price is set to $\beta_i^t = \xi\alpha_i^t$ with $\xi = 0.6$, while the CG cost is considered as $G_{c_i}(P_{i,g}^t) = \omega_i P_{i,g}^t$ with $\omega_i = (1/T)\sum_{t=1}^{T} \alpha_i^t$.

The total number of servers $\overline{M}_i$ is set to 80, with common $\varrho = 0.4$, $\overline{P}_{i,s} = 500$ W and $\underline{M}_i = 5$ for all DC i. The cooling parameters are set to $\gamma = 0.2$, $\kappa_i^t = 2 \times 10^{-9}$, $\overline{P}_{i,a}^t = 30\,\mathrm{kW}, \forall i \in \mathcal{I}, t \in \mathcal{T}$. For simplicity, the unit of workloads in

TABLE II
ENERGY PURCHASE PRICES. THE UNITS ARE $/KWH

Slot	1	2	3	4	5	6	7	8	9	10	11	12
α_1^t	0.53	0.50	0.50	0.53	0.65	0.90	0.82	0.73	0.68	0.59	0.51	0.45
α_2^t	0.49	0.46	0.43	0.43	0.45	0.56	0.77	0.70	0.63	0.58	0.51	0.44
α_3^t	0.43	0.42	0.39	0.36	0.37	0.37	0.46	0.65	0.58	0.53	0.48	0.43
α_4^t	0.69	0.68	0.65	0.61	0.58	0.57	0.60	0.75	1.02	0.94	0.84	0.78

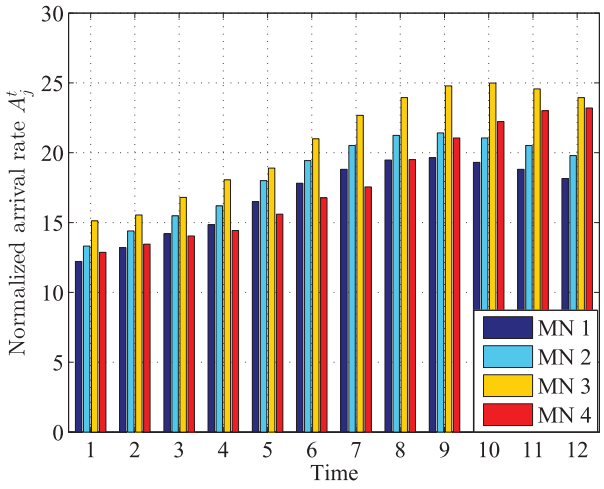

Fig. 4. Real-time arrival rate of interactive workloads.

TABLE III
DELAY-TOLERANT WORKLOADS PARAMETERS. THE UNITS FOR B_j ARE NORMALIZED, AND THE UNITS FOR $\bar{u}_q, \underline{u}_q$ ARE $/UNIT

	DW 1	DW 2	DW 3	DW 4	DW 5	DW 6	DW 7	DW 8
Node	1	1	2	2	3	3	4	4
B_j	75	100	87.5	112.5	100	75	75	150
$[S_q, E_q]$	2-9	1-4	4-10	4-6	1-12AM	2-6	2-7	1-5
$\bar{u}_q, \underline{u}_q$	0.95,0.6	1.1,0.8	0.67,0.61	0.68,0.58	0.6,0.05	0.5,0.5	0.4,0.4	0.3,0.3

this setting is normalized by the server is per-slot capacity D_i, with common D_i for all DCs. The interactive workload arrival rates at MNs are depicted in Fig. 4, rescaled from the real traffic of Wikipedia [42]. Eight different delay-tolerant workloads (DWs) are specified in Table III. For job q, the revenue function is considered $U_q^t(b_{i,q}^t) = -0.01u_q^t(b_{i,q}^t)^2 + u_q^t b_{i,q}^t$ $/unit, where u_q^t is linearly decreasing from to $\bar{u}_q$ across its active interval $[S_q, E_q]$. The bandwidth limits $L_j i^t$ are assumed to be time-invariant, given by

$$\mathbf{L}^t = \begin{bmatrix} 90 & 60 & 40 & 40 \\ 40 & 80 & 40 & 50 \\ 50 & 40 & 100 & 30 \\ 40 & 40 & 50 & 90 \end{bmatrix}, \quad \forall t \in \mathcal{T}.$$

Note that the homogenous settings of data centers are considered here in order to exemplify the impact of other factors (e.g., prices and RES).

The time horizon spans $T = 12$ hours, corresponding to the interval 1PM–12AM in Eastern Time Zone. Here we use the Eastern Time Zone for time-keeping, and the real data have been shifted to show the effect of time zone differences. As a result, the peaks of workload demands, RES and prices are different in the four areas, which provide an opportunity for spatio-temporal workload and energy management. Finally,

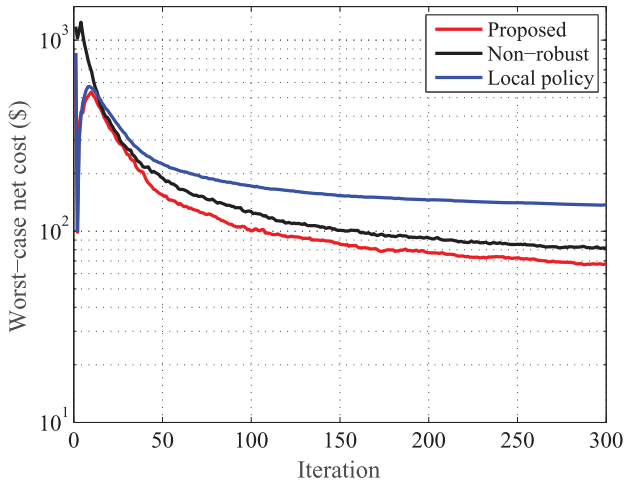

Fig. 5. Comparison of worst-case net costs.

two benchmarks are compared in this setting: A robust *local policy* allocating all workloads from each MN to its nearest DC, and a *non-robust* geographical load balancing policy, which predicts the RES at each slot via its sample mean from the historical dataset. Interestingly, our proposed approach can be reduced to the local policy if the bandwidth matrix $\mathbf{L}^t$ only has positive diagonal entries, and to the non-robust policy if we set $\Delta_s^{\text{up}} = 0$, $\Delta_s^{\text{low}} = 0$, and $|\mathcal{T}_{i,s}| = 1$ in (13).

B. Numerical Results

Fig. 5 depicts the evolution of the worst-case net cost for the proposed algorithm, as well as the two alternatives. Within 300 iterations, the proposed algorithm converges to a worst-case net cost 19% lower than that of the non-robust approach, and 51% lower than that of the local policy. Recall that the non-robust approach is sensitive to the RES prediction error, while the local policy can not perform geographical load balancing. In contrast, the proposed RWEM takes advantage of both factors, and purchases less amount of expensive energy from the spot market that results in a smaller worst-case net cost.

To better illustrate this point, sensitivity analysis to the level of robustness Δ_s^{low} is first studied in Fig. 6. As expected, the proposed RWEM outperforms the non-robust approach in all cases. Meanwhile, the worst-case net costs of both robust and non-robust scheme grow up as the Δ_s^{low} increases. This makes sense intuitively because a larger Δ_s^{low} implies a bigger uncertainty set [cf. (13)], which will eventually increase the worst-case net cost. Hence, the selection of Δ_s^{low} is critical for various scenarios. While a large Δ_s^{low} guarantees robustness of the resultant solution, a small one can moderately reduce its conservatism. We waive the analysis of Δ_s^{up} here, because it plays a less important role in cost minimization (21), due to the monotonicity of the objective (21a) with respect to E_i^t.

The optimal workload schedules d_i^t of the proposed RWEM and the local policy are compared in Figs. 7 and 8, respectively. The worst-case renewable generations $\{E_i^{t*}\}$ are shown to illustrate the principle of geographical workload distribution. Compared with the local policy, the proposed RWEM can

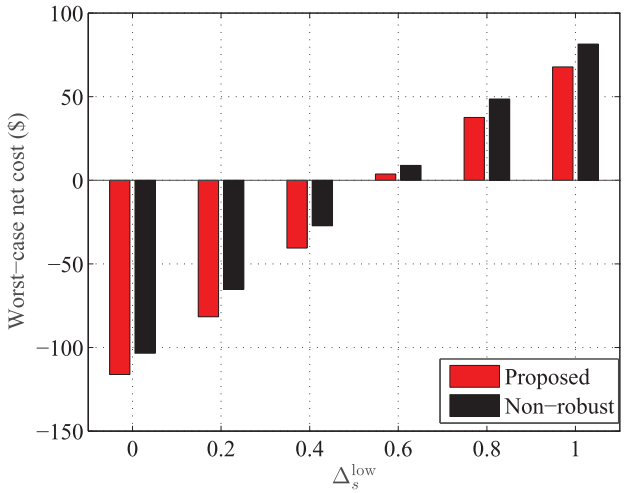

Fig. 6. Worst-case net costs versus the level of robustness Δ_s^{low}.

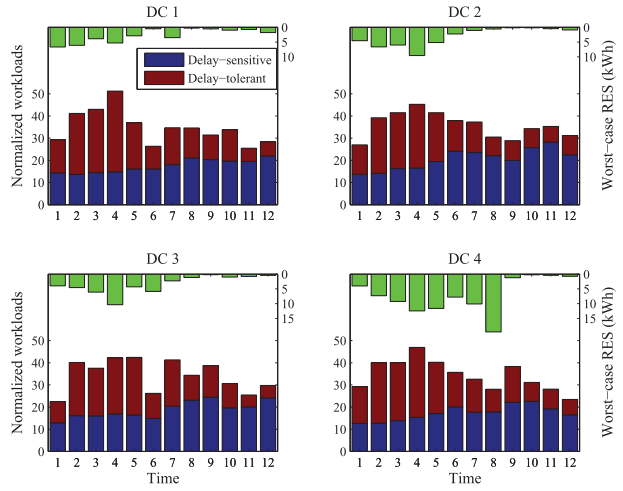

Fig. 7. Optimal workload schedule d_i^t of the proposed algorithm.

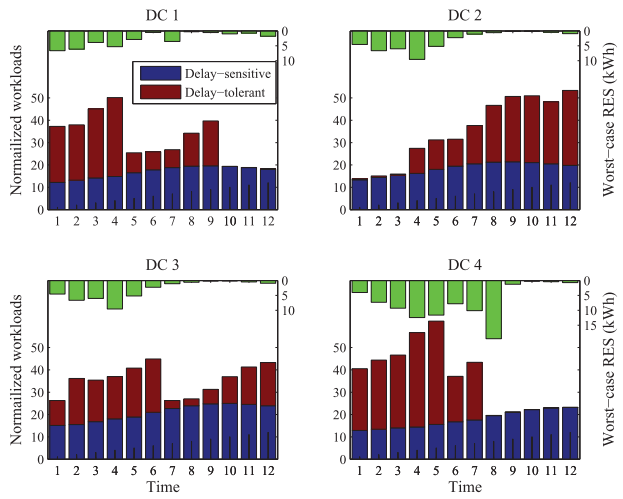

Fig. 8. Optimal workload schedule d_i^t of the local policy.

intelligently route workloads to a remote DC where the system IT demand is lower, RES availability is higher, or, the local energy price is more affordable. To see this, both interactive and flexible workloads are uniformly routed to each DC in

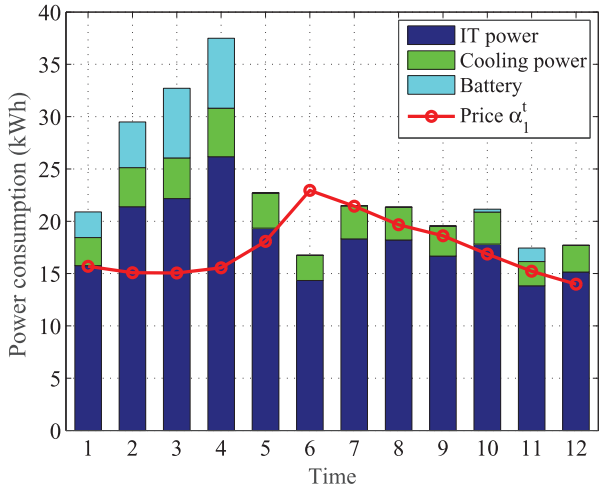

Fig. 9. Optimal power consumption schedule in DC 1.

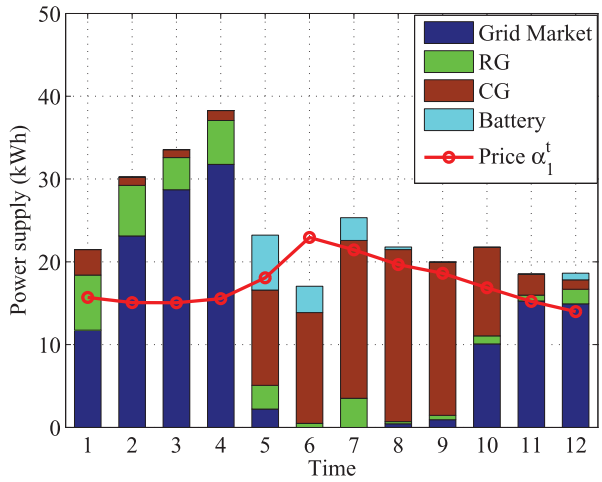

Fig. 10. Optimal power supply schedule in DC 1.

Fig. 7, thus the entire DC network can process more workloads when the RES generations are ample (1PM-6PM). An interesting observation is that even if DC 4 in the Western US enjoys a relatively higher RES, due to its high energy purchase price α_i^t [cf. Table II], the MNs are more likely to route workloads to the areas having lower prices, when there is no renewable surplus in DC 4.

In contrast to Fig. 7, the workload schedules in Fig. 8 are more isolated and thus inefficient. Without coordinating all DCs, the local policy cannot "smooth" the IT demand with the additional degree of freedom in space. Specifically, when the system demand is low in DC 1 as well as DC 3 (10-12 PM), no flexible workloads can be scheduled. Likewise, the RES is not fully utilized in DC 2 (1-3PM). This high fluctuation of system workloads will also cause switching on/off servers frequently, thus incurring an implicitly higher wear-and-tear cost.

The optimal power consumption and supply schedules of DC 1 are depicted in Figs. 9 and 10, respectively. The scaled fluctuation of energy purchase price α_1^t is also plotted to gain intuition on the optimal power schedules. Clearly, less power is consumed when α_1^t is higher (6PM). Using combined cooling sources, the average cooling coefficient of the proposed

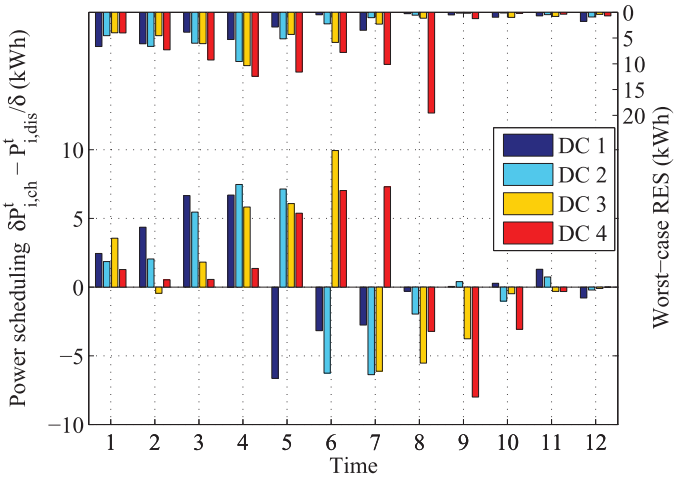

Fig. 11. Optimal battery (dis-)charging schedule in all DCs.

algorithm is around 0.17, which is more efficient than the simple chilled-water cooling with a constant coefficient $\gamma = 0.2$. Furthermore, with the goal of mitigating the high variability of RES, batteries are encouraged to charge when the worst-case renewable generations are high and the energy prices are low (1PM-4PM). Thus, batteries can be discharged when less renewables are available at night (e.g., 5-7PM in Fig. 10. Likewise, Fig. 11 shows that all the batteries exhibit a similar trend in response to RES and price fluctuations. From the power supply perspective, the lower purchase price α_1^t encourages purchasing more energy from the external grid market, while the peak of α_1^t results in a higher power usage from the CG.

V. Conclusions

Robust ahead-of-time workload and energy management for green DCs was considered in this paper. Practical models of IT, cooling, and power supply subsystems were first introduced. Taking into account the spatio-temporal variations of workloads, renewables and electricity market prices, a resource allocation problem was formulated to minimize the system net cost including the network operational cost and the worst-case energy transaction cost. Relying on the strong duality of the convex reformulation, a Lagrange dual based distributed solver was developed to yield the optimal solution. Extensive numerical tests further corroborated the effectiveness and merits of the proposed scheme. Without specific assumption on the underlying distribution of the renewable generation process, we offer a new perspective in dealing with the uncertainties involved in operating sustainable DCs. The novel robust optimization framework of the present paper opens up some interesting research directions, such as integrating additional uncertainties from the interactive workloads and real-time energy prices.

Appendix

A. Proof of Lemma 1

Function $P_{i,IT}(d_i^t, m_i^t)$ is jointly convex in $\{d_i^t, m_i^t\}$ because of its quadratic-over-linear form [cf. (7)]. Hence, the composite function $\mathcal{F}_i^t(P_{i,IT}(d_i^t, m_i^t))$ is jointly convex in $\{d_i^t, m_i^t\}$ since $\mathcal{F}_i^t(P_{i,IT}^t)$ is convex and nondecreasing [34, Sec. 3.2]; and, so is $P_i^t(d_i^t, m_i^t)$.

B. Proof of Proposition 1

Since the absolute value function is convex, and the operations of nonnegative weighted summation and pointwise maximum preserve convexity, it is easy to see that $G_i(\{R_i^t\})$ is convex in $\{R_i^t\}$. In addition, $G_{C_i}(P_{i,g}^t)$ is convex in $P_{i,g}^t$ and $U_q^t(b_{i,q}^t)$ is concave in $b_{i,q}^t$; hence, the objective function (23a) is jointly convex in $\{R_i^t, P_{i,g}^t, P_{\mathrm{i,ch}}^t, P_{\mathrm{i,dis}}^t\}$. The constraints except (23c) are linear, while (23c) defines a feasible set which is actually the epigraph of a convex function [cf. Lemma 1]. Since the epigraph of a convex function is a convex set [34, Chap. 3.1], it follows that (23) is a convex problem, and strong duality holds.

C. Proof of Proposition 2

Compared to (21), (21n)–(21q) are replaced by (23d), and (21o)–(21p) are substituted by (23c) to convexify the problem, and facilitate distributed implementation in (23). By summing (21n) from S_q to $t \in [S_q, \ldots, E_q - 1]$, we find $Q_{i,q}^t = Q_{i,q}^{S_q} - \sum_{\tau=S_q}^t b_{i,q}^\tau + \sum_{\tau=S_q}^t \tilde{b}_{i,q}^\tau$. Then, due to (21q), we have that $\sum_{\tau=S_q}^t \tilde{b}_{i,q}^\tau \geq \sum_{\tau=S_q}^t b_{i,q}^\tau$. Likewise, we obtain $\sum_{\tau=S_q}^{E_q} \tilde{b}_{i,q}^\tau = \sum_{\tau=S_q}^{E_q} b_{i,q}^\tau$, which establishes the equivalence of (21n)–(21q) with (23d). In addition, since the objective (23a) is monotonically increasing with R_i^t, it is easy to see that (23c) is always binding at the optimal solution $\mathbf{x}^*$, which implies that the optimal solution $\mathbf{x}^*$ is also an optimal solution (21); and thus, $\Psi^* = \tilde{\Psi}^*$.

References

[1] Google Inc. (2015). *Google Data Center Locations* [Online]. Available: http://www.google.com/about/datacenters/

[2] J. Whitney and P. Delforge. (2014, Aug.). *Data Center Efficiency Assessment*, Issue Paper, Aug. 2014 [Online]. Available: http://www.nrdc.org/energy/data-center-efficiency-assessment.asp

[3] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for Internet-scale systems," in *Proc. ACM SIGCOMM*, Barcelona, Spain, Aug. 2009, vol. 39, pp. 123–134.

[4] Active Power. (2007). *Data Center Thermal Runaway: A Review of Cooling Challenges in High Density Mission Critical Environments*, White Paper [Online]. Available: www.edsenerji.com.tr/dokuman_indir/16/

[5] A. Wierman, L. L. Andrew, and A. Tang, "Power-aware speed scaling in processor sharing systems," in *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, Apr. 2009, pp. 2007–2015.

[6] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1378–1391, Oct. 2013.

[7] Y. Yao, L. Huang, A. Sharma, L. Golubchik, and M. Neely, "Data centers power reduction: A two time scale approach for delay tolerant workloads," in *Proc. IEEE INFOCOM*, Orlando, FL, USA, Mar. 2012, pp. 1431–1439.

[8] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *Perform. Eval.*, vol. 70, no. 10, pp. 770–791, Oct. 2013.

[9] L. Rao, X. Liu, L. Xie, and Z. Pang, "Hedging against uncertainty: A tale of Internet data center operations under smart grid environment," *IEEE Trans. Smart Grid*, vol. 2, no. 3, pp. 555–563, Sep. 2011.

[10] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's not easy being green," in *Proc. ACM SIGCOMM*, Helsinki, Finland, Aug. 2012, vol. 42, pp. 211–222.
[11] Apple Inc. (2015). *Apple Environmental Responsibility* [Online]. Available: http://www.apple.com/environment/
[12] HP Development Company. (2015). *HP Labs Debuts Architecture for a Net-Zero Energy Data Center* [Online]. Available: http://www.hpl.hp.com/news/
[13] A. Rahman, X. Liu, and F. Kong, "A survey on geographic load balancing based data center power management in the smart grid environment," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 1, pp. 214–233, Feb. 2014.
[14] W. Deng, F. Liu, H. Jin, C. Wu, and X. Liu, "Multigreen: Cost-minimizing multi-source datacenter power supply with online control," in *Proc. ACM Int. Conf. Future Energy Syst.*, Berkeley, CA, USA, May 2013, pp. 149–160.
[15] R. Urgaonkar, B. Urgaonkar, M. Neely, and A. Sivasubramaniam, "Optimal power cost management using stored energy in data centers," in *Proc. ACM SIGMETRICS*, San Jose, CA, USA, Jun. 2011, pp. 221–232.
[16] Y. Guo and Y. Fang, "Electricity cost saving strategy in data centers by using energy storage," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 6, pp. 1149–1160, Jun. 2013.
[17] E. Samadiani, Y. Joshi, and F. Mistree, "The thermal design of a next generation data center: A conceptual exposition," in *Proc. IEEE Int. Conf. Thermal Issues Emerging Technol. Theory Appl.*, Cairo, Egypt, Jan. 2007, pp. 93–102.
[18] T. Chen, X. Wang, and G. B. Giannakis, "Cooling-aware energy and workload management in data centers via stochastic optimization," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 2, pp. 402–415, Mar. 2016.
[19] E. Yao, P. Samadi, V. Wong, and R. Schober, "Residential demand side management under high penetration of rooftop photovoltaic units," *IEEE Trans. Smart Grid*, 2016, to be published, doi: 10.1109/TSG.2015.2472523.
[20] A. Papavasiliou and S. S. Oren, "Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network," *Oper. Res.*, vol. 61, no. 3, pp. 578–592, May 2013.
[21] A. Shapiro, D. Dentcheva, and A. Ruszczynski, *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia, PA, USA: SIAM, 2009.
[22] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synth. Lect. Commun. Netw.*, vol. 3, no. 1, pp. 1–211, 2010.
[23] D. Xu and X. Liu, "Geographic trough filling for Internet datacenters," in *Proc. IEEE INFOCOM*, Orlando, FL, USA, Mar. 2012, pp. 2881–2885.
[24] Z. Liu *et al.*, "Renewable and cooling aware workload management for sustainable data centers," in *Proc. ACM SIGMETRICS*, London, U.K., Jun. 2012, vol. 40, pp. 175–186.
[25] R. Zhou, Z. Wang, A. McReynolds, C. E. Bash, T. W. Christian, and R. Shih, "Optimization and control of cooling microgrids for data centers," in *Proc. IEEE Conf. Thermal Thermomech. Phenom. Electron. Syst.*, San Diego, CA, USA, May 2012, pp. 338–343.
[26] C. Patel, R. Sharma, C. Bash, and A. Beitelmal, "Energy flow in the information technology stack," in *Proc. Int. Mech. Eng. Congr. Expo. (IMECE)*, Chicago, IL, USA, Nov. 2006, pp. 233–241.
[27] R. H. Lasseter and P. Paigi, "Microgrid: A conceptual solution," in *Proc. IEEE Power Electron. Spec. Conf.*, Aachen, Germany, Jun. 2004, pp. 4285–4290.
[28] Y. Zhang, S.-J. Kim, and G. B. Giannakis, "Short-term wind power forecasting using nonnegative sparse coding," in *Proc. IEEE Conf. Inf. Sci. Syst.*, Baltimore, MD, USA, Mar. 2015, pp. 1–5.
[29] D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng, "Adaptive robust optimization for the security constrained unit commitment problem," *IEEE Trans. Power Syst.*, vol. 28, no. 1, pp. 52–63, Feb. 2013.
[30] Y. Zhang, N. Gatsis, and G. B. Giannakis, "Robust energy management for microgrids with high-penetration renewables," *IEEE Trans. Sustain. Energy*, vol. 4, no. 4, pp. 944–953, Oct. 2013.
[31] (2013, Jun.). *California ISO Report for Wind/Solar Resources 2012* [Online]. Available: http://www.caiso.com/Documents/2012Report-Wind-SolarResourcesPostedJun5_2013.htm
[32] P. Harsha and M. Dahleh, "Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy," *IEEE Trans. Power Syst.*, vol. 30, no. 3, pp. 1164–1181, May 2015.
[33] A. J. Wood and B. F. Wollenberg, *Power Generation, Operation, and Control*, 2nd ed. Hoboken, NJ, USA: Wiley, 1996.
[34] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
[35] P. Wendell, J. W. Jiang, M. J. Freedman, and J. Rexford, "Donar: Decentralized server selection for cloud services," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 231–242, 2011.
[36] D. P. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1439–1451, Aug. 2006.
[37] D. P. Bertsekas, *Convex Optimization Theory*. Belmont, MA, USA: Athena Scientific, 2009.
[38] A. Nedic and A. Ozdaglar, "Approximate primal solutions and rate analysis for dual subgradient methods," *SIAM J. Optim.*, vol. 19, no. 4, pp. 1757–1780, 2009.
[39] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM Rev.*, vol. 38, no. 1, pp. 49–95, Mar. 1996.
[40] M. Grant, S. Boyd, and Y. Ye. (2012, Sep.). *CVX: Matlab Software for Disciplined Convex Programming*, version 2.1 [Online]. Available: http://cvxr.com/cvx
[41] National grid. (2015, Jan.). *Hourly Electric Supply Charges in New York* [Online]. Available: https://www.nationalgridus.com/
[42] H. Xu and B. Li, "Joint request mapping and response routing for geo-distributed cloud services," in *Proc. IEEE INFOCOM*, Turin, Italy, Apr. 2013, pp. 854–862.

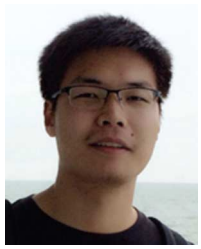

Tianyi Chen received the B.Eng. degree (with highest Hons.) in communication science and engineering from Fudan University, Shanghai, China, in 2014. He is currently pursing the Ph.D. degree in ECE at the University of Minnesota, Minneapolis, MN, USA. Since August 2014, he has been with SPiNCOM. His research interests include network optimization with applications to green communications, and sustainable data centers. He was the recipient of the Student Travel Grant from the IEEE Communications Society in 2013, the National Scholarship from China in 2013, and the UMN ECE Department Fellowship in 2014.

Yu Zhang (S'11) received the B.Eng. and M.Sc. degrees (both with highest Hons.) from Wuhan University of Technology, Wuhan, China, and from Shanghai Jiao Tong University, Shanghai, China, in 2006 and 2010, respectively, and the Ph.D. degree from the University of Minnesota, Minneapolis, MN, USA, in 2015, all in electrical engineering. He is currently a Postdoctoral Scholar with the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. His research interests include smart power grids, optimization theory, machine learning, and wireless communications. He was the recipient of the Huawei Scholarship and the Infineon Scholarship from the Shanghai Jiao Tong University (2009), the ECE Department Fellowship from the University of Minnesota (2010), and the Student Travel Awards from the SIAM and the IEEE Signal Processing Society (2014).

Xin Wang (SM'09) received the B.Sc. and M.Sc. degrees from Fudan University, Shanghai, China, in 1997 and 2000, respectively, and the Ph.D. degree from Auburn University, Auburn, AL, USA, in 2004, all in electrical engineering. From September 2004 to August 2006, he was a Postdoctoral Research Associate with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA. In August 2006, he joined the Department of Computer and Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL, USA, as an Assistant Professor and, in August 2010, was promoted and became an Associate Professor. His research interests include stochastic network optimization, energy-efficient communications, cross-layer design, and signal processing for communications. He served as an Associate Editor for the IEEE SIGNAL PROCESSING LETTERS. He currently serves as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING and as an Editor for the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY.

Georgios B. Giannakis (F'97) received the Diploma degree in electrical engineering from the National Technical University of Athens, Athens, Greece, 1981, and the M.Sc. degree in electrical engineering, the M.Sc. degree in mathematics, and the Ph.D. degree in electrical Engineering all from the University of Southern California (USC), Los Angeles, CA, USA, in 1983, 1986, and 1986, respectively. Since 1999, he has been a Professor with the University of Minnesota, Minneapolis, MN, USA, where he now holds an ADC Chair in Wireless Telecommunications with the ECE Department, and serves as the Director of the Digital Technology Center. His research interests include communications, networking and statistical signal processing—subjects on which he has published more than 385 journal papers, 650 conference papers, 22 book chapters, 2 edited books, and 2 research monographs (h-index 115). Current research focuses on learning from Big Data, wireless cognitive radios, and network science with applications to social, brain, and power networks with renewables. He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society. He is the (co-)inventor of 25 patents issued, and the (co-)recipient of eight Best Paper Awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He was also the recipient of the Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (2015).