

2025

Distributed economic dispatch in power networks incorporating data center flexibility

A. Tsiligkaridis, P. Andrianesis, A.K. Coskun, M.C. Caramanis, I.C. Paschalidis. 2025.

"Distributed Economic Dispatch in Power Networks Incorporating Data Center Flexibility" IEEE Transactions on Sustainable Computing, pp.1-15. <https://doi.org/10.1109/tsusc.2025.3538817>

<https://hdl.handle.net/2144/50540>

"Downloaded from OpenBU. Boston University's institutional repository."

Distributed Economic Dispatch in Power Networks Incorporating Data Center Flexibility

Athanasios Tsiligkaridis, *Member, IEEE*, Panagiotis Andrianesis, *Member, IEEE*, Ayse K. Coskun, *Senior Member, IEEE*, Michael C. Caramanis, *Senior Member, IEEE*, and Ioannis Ch. Paschalidis, *Fellow, IEEE*

Abstract—We consider Data Centers (DCs) as flexible loads that can alter their power consumption to alleviate congestion in the electric power network. We model DCs using a queuing-theoretic view and we form a Quality of Service (QoS)-based cost function that signifies how well a DC can carry out its workload given an amount of active servers. We integrate DCs in a centralized economic dispatch problem that determines, apart from power generation, DC workload shifting and server utilization, while respecting transmission line constraints. We further present a tractable decentralized formulation obtained via Lagrangian decomposition, which we solve using a dual gradient ascent algorithm. Experimental results on a standard power network explore the system-wide benefits of DC flexibility in “coupled” data and power networks, emphasizing on the trade-offs between the DC location, QoS, and efficiency.

Index Terms—Data centers, flexibility, economic dispatch, Lagrangian decomposition.

1 INTRODUCTION

THE amount of data in the world has been increasing at an exponential rate [1] and is predicted to grow to more than 180ZB by 2025, as compared to 61.2ZB in 2020 [2]. This growth can be attributed to a myriad of factors, such as the digitization of many parts of the economy [3], the shift to remote work due to the Covid pandemic [4], the advent of internet-connected devices [5], technological advancements in developing countries, and society’s increased interest in artificial intelligence [6]. In the data-driven world, Data Centers (DCs) have emerged as important entities for handling, storing, stockpiling, and processing data / user-submitted (computationally demanding) jobs [7], and providing cloud computing services [8]. With the advancement of cloud computing, DCs have evolved from physical infrastructures to virtual entities that service workload by harnessing cloud-based resources [9]. The global DC market is predicted to grow at compound annual rates of over 2% and increase by \$615.96B from 2020 to 2025 [10]; DC market revenues of over \$69B are expected by 2024, in the US alone [11].

The proliferation of data render DCs more and more prevalent as heavy power consumers [12]. Globally, DCs consumed around 200TWh in 2015; in 2021, this amount

increased to the range of 220 to 320TWh, and it is expected to grow further [13]. However, DCs can also serve as flexible loads in the power network that can alleviate network congestion, which is becoming more frequent due to the increased penetration of intermittent renewable energy sources (e.g., solar, wind). DCs can offer the much needed flexibility [14], [15], by shifting workload over time and space, e.g., by delaying their computing jobs and transferring jobs between each other (provided they are governed by a common entity). The tech giant Google has recently implemented DC workload shifting in time — non-urgent workloads are delayed to times when renewable energy is present [16] — and has also proposed workload shifting over space [17] — where DCs can transfer jobs based on the availability of renewable energy.

The focus of this paper is on the interaction of DCs with power networks and the ability of DCs to serve as flexible loads through alteration of their workloads that can alleviate network congestion in an economic dispatch problem. Since DCs can transfer workload between each other, controlling both DC and generator outputs via some demand response mechanism that respects power generation and network constraints, while also accounting for DC Quality of Service (QoS) can achieve system-wide benefits. From the DC perspective, we aim to explore the benefits of their incorporation in an economic dispatch problem, and demonstrate that they can achieve significant cost savings for the entire system (data and power networks), enabling the IT sector to contribute to societal sustainability efforts.

The main contributions of this paper are:

- We model DCs leveraging results from queuing theory and form QoS-based cost functions, for which we provide convexity guarantees under certain conditions.
- We introduce a novel centralized economic dispatch problem with full network representation, which involves DCs as flexible loads with QoS considerations,

Research was partially supported by the DOE under grants DE-EE0009696 and DE-AR-0001282, by the NSF under grants CCF-2200052, DMS-1664644, and IIS-1914792, by the ONR under grant N00014-19-1-2571, by France 2030 under grants ANR-22-PETA-0004 and ANR-22-PETA-0009 in the frameworks of the AI-NRGY and FlexTASE projects, and by the Boston University Hariri Institute for Computing and Computational Science & Engineering.

A. Tsiligkaridis, A.K. Coskun, and I.C. Paschalidis are with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215, e-mail: {atsili, acoskun, yannis}@bu.edu.

P. Andrianesis is with the PERSEE Center, Mines Paris - PSL University, Sophia Antipolis 06904, France, and the Systems Engineering Division, Boston University, Boston, MA 02215, e-mail: panagiotis.andrianesis@minesparis.psl.eu.

M.C. Caramanis is with the Department of Mechanical Engineering and the Division of Systems Engineering, Boston University, Boston, MA 02215, e-mail: mcaraman@bu.edu.

and solves for both DC and generator outputs.

- We provide a decentralized version of the economic dispatch problem employing Lagrangian decomposition and a primal-dual algorithm, which can cater for both the power network constraints, and, most importantly, the DC workload shifting, in a distributed manner that scales for the “coupled” data and power networks.
- We present experimental results that provide useful insights on the system-wide benefits from leveraging DC flexibility in the economic dispatch problem, highlighting the trade-offs between the DC locations, their efficiencies, and QoS costs.

The remainder of this paper is organized as follows. Section 2 reviews related work, and Section 3 presents the proposed DC model. Section 4 introduces the centralized economic dispatch problem, whereas Section 5 provides a decentralized formulation and solution algorithm. Section 6 presents experimental results and discusses the key insights. Lastly, Section 7 concludes and provides future research directions.

2 RELATED WORK

Sustainable/green DCs [18] aim at reducing their carbon footprint by leveraging their workload scheduling flexibility to reduce their energy consumption and/or increase the usage of renewable energy [19]. Dou *et al.* [20] propose an on-line workload scheduling algorithm to trade-off electricity cost and performance of delay tolerant workloads. Hogade *et al.* [21] study the energy minimization problem for geo-distributed DCs and propose a set of workload management techniques that account for detailed DC cooling power, collocation interference, renewable energy, and electricity costs. Lin *et al.* [22] estimate cloud DC power consumption — key for energy-aware scheduling — employing an artificial neural network, to enhance adaptability to the changes and fluctuation of different workloads. Zhou *et al.* [23] optimize DCs with deferrable computation requests — represented by a queuing model — and controllable air conditioning to reduce total operation costs that consider time-varying electricity prices and server degradation. DC energy efficiency improvement is also explored via the hybridization of two effective approaches [24]: (i) combined cooling, heating, and power, and (ii) waste heat reuse systems. Aiming at reducing the carbon footprint of cloud DCs by mitigating brown energy usage and maximizing renewable energy usage, while ensuring the QoS of workloads, Xu *et al.* [25] propose a self-adaptive approach with a brownout-based algorithm for interactive workloads and a deferring algorithm for batch workloads. Assuming no grid connection, de Nardin *et al.* [26] explore a clean — renewable only — DC architecture and evaluate scheduling and power capping online heuristics that can handle fluctuating power profiles while accounting for job QoS. Considering virtualized networked DCs, Zhou *et al.* [27] propose an adaptive algorithm that minimizes energy consumption by computational resources as well as virtual machine configuration and communication and storage resources, while meeting agreed user QoS requirements. Exploring mini DC, da Silva *et al.* [28] propose an optimization approach for energy-efficient resource

allocation that exploits both virtual machines and energy migrations between green compute nodes. In order to avoid server over-provisioning, Ahmed *et al.* [29] introduce a reliability index that quantifies the probability of insufficient spare servers due to failures, while balancing between QoS and energy consumption.

Early DC demand response works suggest combining workload scheduling and local power generation to mitigate load peaks and reduce energy costs [30], and compare the potential of DCs with large-scale storage [31], presenting a prediction-based pricing design of DC demand response. Chen *et al.* [32] propose a pricing mechanism based on parameterized supply function bidding for the provision of emergency demand response from multi-tenant collocation DCs. Niu *et al.* [33] use coalitional game theory along with a fair payoff allocation to model the cooperation of DC aggregation that mitigates their demand response capacity. Real-time pricing for DC demand response and a game-theoretic model are employed by Bahrami *et al.* [34], along with a distributed algorithm that guarantees convergence to a stable outcome. Chen *et al.* [35] provide a comprehensive survey on several aspects of demand response in internet DCs. An optimal bidding strategy is proposed by Fu *et al.* [36], which combines DC demand response while taking advantage of DC waste heat, as well as various forms of distributed energy resources. Prior work in the area has explored the provision of regulation service reserve along with adaptive policies and QoS assurances [37], [38], and has focused on price-based models for demand response, proposing a framework where DCs interact with an aggregator via some price incentive [39], [40], but has not considered the joint optimization of DCs in a power network architecture. In this work, our attention shifts to the integration of DCs in an economic dispatch problem with full network representation.

There are a few works that consider the collocation of DCs with renewable generation and their interaction with the power network. Wang *et al.* [41] propose that the optimal placement of a new DC should consider — in addition to capital and operational costs — the cost of transmission system losses, as well as constraints for avoiding line overloading and voltage variations. Interestingly, it is shown that the collocation of DC and wind generation may not reduce overall costs; hence, the impact of location on the costs of both DCs and the grid should be jointly considered. Similar findings are obtained by Kim *et al.* [42]. Economic dispatch principles are used in a DC uninterrupted power supply control strategy [43] to leverage DC flexibility aiming at DC operation cost minimization. Recent DC research has also focused on space-time load shifting in a setting of geographically-distributed DCs. Zhang *et al.* [44] model the ability for shifting computing loads to capture flexibility using the concept of virtual links in electricity market clearing processes. Niu *et al.* [45] formulate a stochastic economic dispatch problem — solved with Benders decomposition — as a means of incorporating geo-distributed DCs to assist in renewable integration, through geographical transferring of workload via the internet, while considering a QoS requirement as a bound on the jobs’ average delay. DC flexibility is also harvested using machine learning models that predict energy consumption and server temperatures, integrated

TABLE 1: Main Notation

Symbol	Description
DC	Data Center
QoS	Quality of Service
\mathcal{I}	Set of I DCs, indexed by i, j
\mathcal{G}	Set of G generators, indexed by g
\mathcal{D}	Set of D inflexible loads, index by d
\mathcal{B}	Set of B buses, indexed by $b(\cdot)$
\mathcal{K}	Set of K lines, indexed by k
μ_{A_i}, σ_{A_i}	Arrival process statistics for DC i
μ_{H_i}, σ_{H_i}	Service process statistics for DC i
μ_{B_j}, σ_{B_j}	Per server statistics for departures from DC j
C_i^{QoS}	QoS-based cost function for DC i
$\rho_{1,i}, \rho_{2,i}$	QoS cost coefficients for DC i
θ_i	QoS-related parameter for DC i
β_i	Power consumption rate of active server at DC i
$N_{i \rightarrow j}$	Amount of servers at DC j used by DC i (vector \mathbf{n}_i)
N_j^{active}	Amount of active servers at DC j
$\bar{N}_j^{\text{active}}$	Active server capacity limit for DC i
C_i^{2w}, κ_i^{2w}	Cost/scalar used to prevent 2-way transfer of jobs
c_g	Cost of generator g
P_g	Power produced by generator g (vector \mathbf{p})
$\underline{P}_g, \bar{P}_g$	Min/max capacity limits of generator g
P_d	Power used by inflexible load d
$a_{b,k}$	Shift factor of bus b , line k
\bar{F}_k	Capacity limit of line k .
λ, ν, ζ	Dual variables

within a power system scheduling framework for potential demand shifting based on (system) carbon intensity [46]. The potential activation of DC temporal, spatial and integrated energy flexibility is approached via a centralized electricity–heat coordinated operation problem [47], which is decomposed into a bilevel model with an incentivized profit-sharing mechanism. Zhang *et al.* [48] present a mixed integer programming model for DC scheduling and temporal flexibility that results in reduced peak demand charges carbon emissions.

Arguably, the use of distributed algorithms [49] is indispensable for “coupled” data and power networks, whereas DC QoS modeling and integration in an economic dispatch problem is key to properly schedule and inherently price DC flexibility. To the best of our knowledge, prior work has not considered QoS-based cost function modeling and integration in an economic dispatch problem, which jointly optimizes DC flexibility and power network assets. This work leverages queuing theory to form a DC QoS cost function, and integrates this function in a distributed economic dispatch problem, that effectively discovers system-wide efficient allocations for both DCs and the power network assets.

3 DATA CENTER MODEL

In this section, we present preliminary material on stochastic processes (in Subsection 3.1) that is relevant for the DC QoS modeling (in Subsection 3.2), and the DC service process scaling (in Subsection 3.3). We further present our main analytical results under a normal distribution that is employed in this paper (in Subsection 3.4), which we extend to a Poisson distribution (in Subsection 3.5) to showcase the wider applicability of our queuing theory-based approach.

3.1 Preliminaries

The main notation used throughout the paper is listed in Table 1. Vectors are represented using bold font (i.e., \mathbf{v}), \mathbb{E} is used to denote the expectation operation, \mathbb{P} represents probability, and $1_{\{a>b\}}$ is the indicator function that takes the value 1 if the condition $a > b$ is satisfied or 0 otherwise.

Consider a discrete stochastic process $X = \{X_1, X_2, \dots\}$. The Moment Generating Function (MGF), $M_{X_m}(\theta)$, of some random variable X_m , with $m = 1, 2, \dots$, and some real parameter θ , is defined as follows:

$$M_{X_m}(\theta) = \mathbb{E}[e^{\theta X_m}]. \quad (1)$$

Let S_n denote the n -th partial sum of random variables X_m , i.e., $S_n = \sum_{m=1}^n X_m$, and define $\Lambda_n(\theta) = \frac{1}{n} \log \mathbb{E}[e^{\theta S_n}]$. Let $\Lambda_X(\theta)$ denote the limiting log-MGF of X , implying, with some abuse of notation, that the partial sum process S_n maps to a discrete stochastic process X , with:

$$\Lambda_X(\theta) \triangleq \lim_{n \rightarrow \infty} \Lambda_n(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}[e^{\theta \sum_{m=1}^n X_m}]. \quad (2)$$

We make the following assumption, though it can be relaxed to include stochastic processes with appropriate mixing conditions:

Assumption A. *The random variables X_m , $m = 1, 2, \dots$, are independent and identically distributed (iid), and Gaussian with parameters μ_X (mean) and σ_X (standard deviation).*

Lemma 1. *If Assumption A holds, the limiting log-MGF of X , $\Lambda_X(\theta)$, is given by:*

$$\Lambda_X(\theta) = \mu_X \theta + \frac{1}{2} \sigma_X^2 \theta^2. \quad (3)$$

Proof. From (2), the limit of the log-MGF is given by:

$$\Lambda_X(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[\prod_{m=1}^n e^{\theta X_m} \right].$$

Since X_m is iid, $\mathbb{E} \left[\prod_{m=1}^n e^{\theta X_m} \right] = \prod_{m=1}^n \mathbb{E}[e^{\theta X_m}] = \mathbb{E}[e^{\theta X}]^n$,

and, hence, the limit $\Lambda_X(\theta)$ becomes:

$$\begin{aligned} \Lambda_X(\theta) &= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}[e^{\theta X}]^n = \log \mathbb{E}[e^{\theta X}] \\ &= \log M_X(\theta). \end{aligned} \quad (4)$$

If X is Gaussian, the MGF in (1) is given by [50, Example 10.16]:

$$M_X(\theta) = e^{\left(\mu_X \theta + \frac{1}{2} \sigma_X^2 \theta^2\right)}. \quad (5)$$

By substituting (5) into (4), we get (3). \square

We next present the DC QoS modeling, which allows workload shifting from one DC to another.

3.2 Data Center QoS Modeling

We model each DC as a G/G/1 queuing system, where jobs arrive with inter-arrival times following a general (arbitrary) distribution and are serviced by a (different) general distribution. This model considers the DC as a pooled (single server) resource consisting of all active servers at the DC and maintains generality in terms of the considered probability distribution.

Let $\mathcal{I} \equiv \{1, \dots, I\}$ denote the set of DCs that are indexed by i, j , where I is the number of DCs. Unless otherwise mentioned, $i \in \mathcal{I}, j \in \mathcal{I}$. We prefer to use index i for the (from) DC that shifts the workload and index j for the (to) DC that receives the workload. We represent the workload shifting from DC i to DC j through a shared server quantity $N_{i \rightarrow j}$, which denotes the amount of servers belonging to (and are physically located at) DC j that are being used by DC i . We view the idea of workload shifting between DCs through the lens of server sharing, which occurs when $i \neq j$. We define \mathbf{n}_i to describe all local and foreign (i.e., shared) servers being used by DC i , i.e., $\mathbf{n}_i \triangleq (N_{i \rightarrow j}; \forall j)$. We also define N_j^{active} to describe the amount of active servers at DC j , which can be determined as $N_j^{\text{active}} = \sum_i N_{i \rightarrow j}$, i.e., the sum of all servers that are physically located at DC j and are being used by DC $i, \forall i$; said differently, all of the active servers physically located at DC j that are either used by itself (DC j) or by other DCs (DC $i \neq j$).

Let A_j denote a discrete stochastic process that describes the arrival process at DC j . Specifically, we have $\{A_{j,1}, A_{j,2}, \dots\}$, where $A_{j,t}$ is a random variable representing the number of jobs that arrive at DC j during time slot t . Let H_j denote a discrete stochastic process that describes the service process of DC j . Specifically, we have $\{H_{j,1}, H_{j,2}, \dots\}$, where $H_{j,t}$ is a random variable representing the (per DC) maximum number of departing jobs from DC j during time slot t .

For DC i , we define an individual QoS constraint using the aforementioned per DC arrival and service processes, A_i and H_i , respectively. Qualitatively, QoS represents the effectiveness of a DC in completing jobs. The QoS is modeled using the probability that the queue length L in the system exceeds or equals some value U . We want this probability to be as small as possible. By applying [51, Theorem 6.1.1], we can approximate this probability as:

$$\mathbb{P}[L \geq U] \sim e^{-\theta_i U}, \quad (6)$$

where $\theta_i > 0$ is a scalar that depends on the arrival and service distributions, defined as the positive root to the following equation:

$$\Lambda_{A_i}(\theta_i) + \Lambda_{H_i}(-\theta_i) = 0, \quad (7)$$

where the limiting log-MGFs $\Lambda_{A_i}(\cdot)$ and $\Lambda_{H_i}(\cdot)$ are given by (2) for processes A_i and H_i .

We then proceed to define a QoS cost function for DC i , using the QoS parameter θ_i , also providing the intuition for our selection. Naturally, we endeavor to make $\mathbb{P}[L \geq U]$ small. This requires a large θ_i . Hence, we use this QoS parameter to design the cost function for DC i such that it yields a small cost for a large θ_i , and vice versa. Therefore, a *convex, non-increasing* cost function for DC i , $C_i(\theta_i)$, seems as a natural fit, which can be interpreted as providing better

QoS as θ_i grows. Examples of valid cost functions include scaled exponential or negative logarithmic functions.¹ In this work, we define the QoS cost function for DC i as:

$$C_i(\theta_i(\mathbf{n}_i)) \triangleq \rho_{1,i} e^{-\rho_{2,i} \theta_i(\mathbf{n}_i)}, \quad (8)$$

where $\theta_i(\mathbf{n}_i)$ expresses the dependence of the QoS parameter on the workload shifting, and $\rho_{1,i}$ and $\rho_{2,i}$ are empirically defined DC cost coefficients — see e.g., our prior work [40] where a similar cost function structure was used.

The following proposition elaborates on the convexity of the QoS function.

Proposition 1. *If $\theta_i(\mathbf{n}_i)$ is concave in $N_{i \rightarrow j}$, the QoS cost function $C_i(\theta_i(\mathbf{n}_i))$ given by (8), is convex in $N_{i \rightarrow j}, \forall j$.*

Proof. We first show that function $C_i(\theta_i)$ is *convex* and *non-increasing* in θ_i . The first derivative is:

$$\frac{\partial C_i(\theta_i)}{\partial \theta_i} = -\rho_{1,i} \rho_{2,i} e^{-\rho_{2,i} \theta_i} \leq 0,$$

and hence C is a *non-increasing* function. The second derivative is:

$$\frac{\partial^2 C_i(\theta_i)}{\partial \theta_i^2} = \rho_{1,i} \rho_{2,i}^2 e^{-\rho_{2,i} \theta_i} \geq 0,$$

and hence C is a *convex* function. Since $C_i(\theta_j(\mathbf{n}_i))$ is a composition of two functions, the outer of which, $C_i(\theta_i)$, is *convex* and *non-increasing*, and the inner of which, $\theta_i(\mathbf{n}_i)$, is assumed to be *concave*, the composition, $C_i(\theta_i(\mathbf{n}_i))$, is a *convex* function [52]. \square

3.3 DC Service Process Scaling

So far, we considered the DC as a pooled resource (single server). However, the DC service process can be modeled using the per server statistics, scaled according to the amount of active servers at the DC. Assuming that the number of jobs serviced per server and per unit of time is constant, on average, the DC service process, H , can be determined by linearly scaling the *per server* statistics with the amount of active servers.

Let B_j denote a discrete stochastic process that describes the number of jobs that are serviced by each server of DC j . Specifically, we have $\{B_{j,1}, B_{j,2}, \dots\}$, where $B_{j,t}$ is a random variable representing the maximum number of departing jobs from each server of DC j during time slot t .

For the rest of this paper, and since we do not consider workload shifting over time in this work, we can focus on a specific time slot and de-clutter the notation by removing subscript t . Hence, we can also refer to A_j, H_j , and B_j as random variables. To avoid confusion, we reiterate that A_j and H_j are defined per DC j , whereas B_j is defined per server of DC j . Random variable H_j can thus be considered as a scaled random variable, with $H_j = N_j^{\text{active}} B_j$.

In what follows, we assume that the queuing system that describes the DC experiences Gaussian arrival and departure processes, noting, however, that other distributions

1. We acknowledge that the exact DC QoS cost structure and parametrization is a research topic in itself, which is outside the scope of this paper. Nevertheless, we note that our setup is general and allows employing any convex and non-increasing (in θ_i) cost function that can be empirically parameterized.

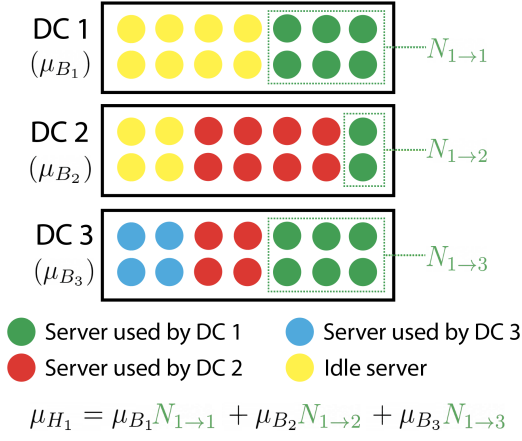


Fig. 1: Visualization of DC server partitioning and per DC service process statistics (mean departure parameter, DC 1).

can be chosen, since the G/G/1 queue model allows for a general distribution. Our choice allows us to derive the statistics per server and DC, which describe the arrival and departure/service processes. Specifically, let μ_{A_i} and σ_{A_i} denote the mean and standard deviation of random variable A_i for DC i . Also, let μ_{B_j} and σ_{B_j} denote the mean and standard deviation of random variable B_j , for each server of DC j . The service process statistics of the scaled random variable H_i , μ_{H_i} (mean) and $\sigma_{H_i}^2$ (variance), for DC i , are given by:

$$\mu_{H_i} = \sum_j \mu_{B_j} N_{i \rightarrow j}, \quad (9)$$

$$\sigma_{H_i}^2 = \sum_j \sigma_{B_j}^2 N_{i \rightarrow j}, \quad (10)$$

i.e., the sum (over j) of the per server statistics of DC j , μ_{B_j} and $\sigma_{B_j}^2$, scaled by the amount of servers at DC j that are used by DC i , $N_{i \rightarrow j}$.

In Fig. 1, we explain the concept with an example of 3 DCs. Each DC i has a certain per server mean departure parameter (μ_{B_i}) and a certain partition of idle (yellow) and active local servers (green, red, blue). The active server partition reflects the amount of servers that are used by itself as well or by other DCs ($N_{1 \rightarrow i}$, $N_{2 \rightarrow i}$, $N_{3 \rightarrow i}$). For illustration, we show how μ_{H_1} , the (per DC) mean number of jobs departing from DC 1, is formed — see (9). DC 1 uses some of its local servers, $N_{1 \rightarrow 1}$, with each server carrying out a mean amount of μ_{B_1} jobs, as well as (foreign) servers that are located at DC 2 and 3 ($N_{1 \rightarrow 2}$ and $N_{1 \rightarrow 3}$), with each server carrying out a mean amount of μ_{B_2} and μ_{B_3} jobs, respectively. Thus, the aggregate (per DC) service rate for DC 1, i.e., the mean amount of DC 1 serviced jobs, μ_{H_1} , is given by the sum of the servers used by DC 1 (see the green partitions in Fig. 1) weighted by the mean amount of jobs per server that refers to the physical location of the server. A similar depiction can be made for the aggregate (per DC) service rates of DC 2 (see red partitions) and 3 (see blue partition), as well as for the aggregate (per DC) variance, $\sigma_{H_i}^2$, for $i = 1, 2, 3$, shown in (10).

3.4 Data Center QoS Results under Assumption A

The QoS parameter θ_i is given by the following Proposition.

Proposition 2. *If Assumption A holds for discrete stochastic processes A , B , and H , the QoS parameter θ_i is a function of \mathbf{n}_i given by:*

$$\theta_i(\mathbf{n}_i) = 2 \left(\frac{\sum_j \mu_{B_j} N_{i \rightarrow j} - \mu_{A_i}}{\sum_j \sigma_{B_j}^2 N_{i \rightarrow j} + \sigma_{A_i}^2} \right). \quad (11)$$

Proof. If Assumption A holds for processes A , B , and H , using (3) of Lemma 1, (7) becomes:

$$(\mu_{A_i} - \mu_{H_i})\theta_i + \frac{1}{2}(\sigma_{A_i}^2 + \sigma_{H_i}^2)\theta_i^2 = 0. \quad (12)$$

From [51], we know that $\theta_i > 0$, hence (12) becomes:

$$\theta_i = 2 \left(\frac{\mu_{H_i} - \mu_{A_i}}{\sigma_{H_i}^2 + \sigma_{A_i}^2} \right). \quad (13)$$

Using (9) and (10), (13) yields (11). \square

The non-negativity of θ_i implies that DC i , which we model as a queuing system, has the ability to service its incoming jobs and does not experience an input buildup where incoming jobs cannot be carried out by the servers that DC i uses. By looking at the structure of θ_i in (11) and using (9), we see that $\mu_{H_i} \geq \mu_{A_i}$ must be satisfied to obtain a non-negative θ_i . The preceding inequality signifies that the mean amount of departing jobs from DC i must exceed or equal the mean amount of arriving jobs at DC i . If the preceding inequality is reversed, i.e., $\mu_{H_i} \leq \mu_{A_i}$, a queue buildup occurs and θ_i can become negative. For the rest of this paper, we assume that the non-negativity condition for $\theta_i(\mathbf{n}_i)$, $\theta_i(\mathbf{n}_i) > 0$, is always satisfied. To guarantee the non-negativity, a lower bound on the amount of servers that DC i uses, $\mathbf{n}_i = (N_{i \rightarrow j}; \forall j)$, could be explicitly imposed. However, we do not explicitly include this constraint in our forthcoming formulations, as it is implicitly satisfied because of our QoS cost function, which induces \mathbf{n}_i , $\forall i$, to be large so as to avoid high cost penalties for low QoS, thus discouraging small amounts of active servers. Larger amounts of active servers result in more jobs being serviced and, in turn, smaller QoS penalty costs.

The function $\theta_i(\mathbf{n}_i)$ in (11) is *concave* under the condition discussed in the following Lemma:

Lemma 2. *The QoS parameter $\theta_i(\mathbf{n}_i)$ is concave in $N_{i \rightarrow j}$ if the following condition holds:*

$$\theta_i(\mathbf{n}_i) \leq \frac{2\mu_{B_j}}{\sigma_{B_j}^2}. \quad (14)$$

Proof. For concavity, we check the second derivative of $\theta_i(\mathbf{n}_i)$ w.r.t. $N_{i \rightarrow j}$. For notational simplicity, we use index j' in the summations of parameter $\theta_i(\mathbf{n}_i)$ in (11), and drop the dependence on \mathbf{n}_i as follows:

$$\theta_i = 2 \frac{\sum_{j' \in \mathcal{I}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i}}{\sum_{j' \in \mathcal{I}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2}. \quad (15)$$

We further simplify by denoting the numerator and denominator in (15) as Y and Z , respectively, i.e., $\theta_i = 2\frac{Y}{Z}$, with

$$Y = \sum_{j' \in \mathcal{I}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i}, \quad (16)$$

$$Z = \sum_{j' \in \mathcal{I}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2, \quad (17)$$

and take the partial derivatives *w.r.t.* $N_{i \rightarrow j}$. Since

$$\frac{\partial Y}{\partial N_{i \rightarrow j}} = \mu_{B_j} \quad \text{and} \quad \frac{\partial Z}{\partial N_{i \rightarrow j}} = \sigma_{B_j}^2,$$

the first derivative of θ_i *w.r.t.* $N_{i \rightarrow j}$ is given by

$$\frac{\partial \theta_i}{\partial N_{i \rightarrow j}} = 2 \cdot \frac{\mu_{B_j} \cdot Z - \sigma_{B_j}^2 \cdot Y}{Z^2}. \quad (18)$$

For simplicity, we denote the numerator in (18) by W , i.e.,

$$W = \mu_{B_j} \cdot Z - \sigma_{B_j}^2 \cdot Y. \quad (19)$$

The second derivative of θ_i *w.r.t.* $N_{i \rightarrow j}$ is given by

$$\frac{\partial^2 \theta_i}{\partial N_{i \rightarrow j}^2} = 2 \frac{\frac{\partial W}{\partial N_{i \rightarrow j}} \cdot Z^2 - W \cdot 2 \cdot Z \cdot \sigma_{B_j}^2}{Z^4}. \quad (20)$$

Since the partial derivative of W *w.r.t.* $N_{i \rightarrow j}$ is $\frac{\partial W}{\partial N_{i \rightarrow j}} = 0$, (20) becomes

$$\frac{\partial^2 \theta_i}{\partial N_{i \rightarrow j}^2} = -4\sigma_{B_j}^2 \frac{W}{Z^3}. \quad (21)$$

Concavity of θ_i requires $\frac{\partial^2 \theta_i}{\partial N_{i \rightarrow j}^2} \leq 0$, and hence the condition $W \geq 0$, since $Z > 0$. By replacing W from (19), the condition for concavity of θ_i is $\mu_{B_j} \cdot Z - \sigma_{B_j}^2 \cdot Y \geq 0$, which can be rewritten as $\frac{Y}{Z} \leq \frac{\mu_{B_j}}{\sigma_{B_j}^2}$. Since $\theta_i = 2\frac{Y}{Z}$, we see that the preceding inequality implies a condition on θ_i (in addition to the condition $\theta_i \geq 0$); thus, θ_i is convex in $N_{i \rightarrow j}$ if:

$$\theta_i(\mathbf{n}_i) \leq \frac{2\mu_{B_j}}{\sigma_{B_j}^2}. \quad \square$$

Interestingly, the following Proposition elaborates on the aforementioned concavity condition:

Proposition 3. *The condition for concavity of $\theta_i(\mathbf{n}_i)$ in $N_{i \rightarrow j}$ shown in (14) does not depend on the value of $N_{i \rightarrow j}$.*

Proof. Given the expression for $\theta_i(\mathbf{n}_i)$ in (11) and the QoS parameter concavity result presented in (14), we want to show that the condition in (14) does not depend on $N_{i \rightarrow j}$ itself. Substitute (11) into (14) and get:

$$\frac{\sum_{j' \in \mathcal{I} \setminus \{j\}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i}}{\sum_{j' \in \mathcal{I} \setminus \{j\}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2} \leq \frac{\mu_{B_j}}{\sigma_{B_j}^2},$$

break the sums into index sets $j' \in \mathcal{I} \setminus \{j\}$ and $\{j\}$, and incorporate $N_{i \rightarrow j}$ into the *rhs* to obtain:

$$\frac{\left(\sum_{j' \in \mathcal{I} \setminus \{j\}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i} \right) + \mu_{B_j} N_{i \rightarrow j}}{\left(\sum_{j' \in \mathcal{I} \setminus \{j\}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2 \right) + \sigma_{B_j}^2 N_{i \rightarrow j}} \leq \frac{\mu_{B_j} N_{i \rightarrow j}}{\sigma_{B_j}^2 N_{i \rightarrow j}}.$$

The above inequality $\left[\frac{\alpha+\gamma}{\beta+\delta} \leq \frac{\gamma}{\delta} \right]$ is equivalent to $\left[\frac{\alpha}{\beta} \leq \frac{\gamma}{\delta} \right]$, and we get:

$$\frac{\left(\sum_{j' \in \mathcal{I} \setminus \{j\}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i} \right)}{\left(\sum_{j' \in \mathcal{I} \setminus \{j\}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2 \right)} \leq \frac{\mu_{B_j} N_{i \rightarrow j}}{\sigma_{B_j}^2 N_{i \rightarrow j}}.$$

By simplifying the *rhs*, we have:

$$\frac{\sum_{j' \in \mathcal{I} \setminus \{j\}} \mu_{B_{j'}} N_{i \rightarrow j'} - \mu_{A_i}}{\sum_{j' \in \mathcal{I} \setminus \{j\}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2} \leq \frac{\mu_{B_j}}{\sigma_{B_j}^2}.$$

The above condition for concavity of $\theta_i(\mathbf{n}_i)$ in $N_{i \rightarrow j}$, where $\mathbf{n}_i = (N_{i \rightarrow j}; \forall j \in \mathcal{I})$, does not depend on the value of $N_{i \rightarrow j}$ itself, but on all other $N_{i \rightarrow j'}$, where $j' \in \mathcal{I} \setminus \{j\}$. \square

The following Proposition elaborates on the convexity of the QoS cost function.

Proposition 4. *Under the condition of Lemma 2, the QoS cost function $C_i(\theta_i(\mathbf{n}_i))$ given by (8), where $\theta_i(\mathbf{n}_i)$ is given by Proposition 2, is convex in $N_{i \rightarrow j}$, $\forall j$.*

Proof. This result is obtained directly applying Proposition 1 and using Lemma 2 for the concavity of $\theta_i(\mathbf{n}_i)$ in $N_{i \rightarrow j}$. \square

Considering DCs of the same type (i.e., with same per server departure process statistics), the following Proposition proves that the QoS cost function is always convex.

Proposition 5. *If all DCs have the same per server departure statistics, $\mu_{B_i} = \mu_{B_j}$ and $\sigma_{B_i}^2 = \sigma_{B_j}^2 \forall i, j$, the QoS cost function $C_i(\theta_i(\mathbf{n}_i))$ given by (8), where $\theta_i(\mathbf{n}_i)$ is given by (11), is convex in $N_{i \rightarrow j}$, $\forall j$.*

Proof. If all DCs have the same per server departure statistics, $\mu_B = \mu_{B_i} = \mu_{B_j}$ and $\sigma_B^2 = \sigma_{B_i}^2 = \sigma_{B_j}^2 \forall i \in \mathcal{I}, j \in \mathcal{I}$, we see that the first derivative of $\theta_i(\mathbf{n}_i)$ becomes:

$$\frac{\partial \theta(\mathbf{n}_i)}{\partial N_{i \rightarrow j}} = 2 \frac{\mu_B \sigma_{A_i}^2 + \mu_{A_i} \sigma_B^2}{(\sigma_B^2 \sum_{j' \in \mathcal{I}} N_{i \rightarrow j'} + \sigma_{A_i}^2)^2}.$$

Next, we compute the second derivative of $\theta_i(\mathbf{n}_i)$ and we see that it satisfies the concavity condition for any positive values of μ_B and σ_B^2 :

$$\frac{\partial^2 \theta(\mathbf{n}_i)}{\partial N_{i \rightarrow j}^2} = \frac{-4\sigma_B^2 (\mu_B \sigma_{A_i}^2 + \mu_{A_i} \sigma_B^2)}{\left(\sigma_B^2 \sum_{j' \in \mathcal{I}} N_{i \rightarrow j'} + \sigma_{A_i}^2 \right)^3} \leq 0. \quad \square$$

3.5 DC QoS Results under the Poisson Distribution

As already noted, we presented the results for the normal distribution (Gaussian); however, other distributions can be similarly chosen. In this respect, and although it is beyond the scope of this paper, we extend our analysis and provide results for the Poisson distribution to showcase the applicability of our queuing theory-based approach.

Let ξ_{A_i} denote the arrival rate for DC i , and ξ_{B_j} denote the per server departure rate for DC j — where we use a different symbol (ξ) to avoid confusion with the Gaussian processes. Similarly to (9), the scaled parameter ξ_{H_i} that refers to the service process rate of DC i is given by

$$\xi_{H_i} = \sum_j \xi_{B_j} N_{i \rightarrow j}. \quad (22)$$

Naturally, the illustration in Fig. 1, replacing μ by ξ , applies straightforwardly to (22).

Proposition 6. *If the random variables A , B , and H , follow the Poisson distribution, the QoS parameter $\theta_i(\mathbf{n}_i)$ is given by:*

$$\theta_i(\mathbf{n}_i) = \xi_{A_i} - \sum_j \xi_{B_j} N_{i \rightarrow j}. \quad (23)$$

Proof. From [51, Eq. 2.62], we have

$$\Lambda_{A_i}(\theta_i) = \log\left(\frac{\xi_{A_i}}{\xi_{A_i} - \theta_i}\right), \quad \Lambda_{H_i}(\theta_i) = \log\left(\frac{\xi_{H_i}}{\xi_{H_i} - \theta_i}\right).$$

Applying (7), similarly to [51, Eq. 2.63], which refers to an M/M/1 queue, we get:

$$\frac{\xi_{A_i}}{\xi_{A_i} - \theta_i} \cdot \frac{\xi_{H_i}}{\xi_{H_i} + \theta_i} = 1,$$

which yields $\theta_i = 0$ or $\theta_i = \xi_{A_i} - \xi_{H_i}$. Taking the positive root and using (22), we get (23). \square

Since we will not further employ the Poisson distribution in this work, we leave to the interested reader the (arguably straightforward) proof of the QoS cost function convexity.

4 CENTRALIZED PROBLEM FORMULATION

In this section, we formulate a centralized economic dispatch optimization problem, with a full network representation, and enriched with DCs.

We consider a system with B buses (nodes), K transmission lines, G generators, and D inflexible loads. Let \mathcal{B} denote the set of buses, indexed by b , with $\mathcal{B} = \{1, \dots, B\}$ and \mathcal{K} denote the set of lines, with $\mathcal{K} = \{1, \dots, K\}$. Let $a_{b,k}$ denote the fraction of power injected in bus b that flows over line k , *a.k.a.* shift factor. Let $\mathcal{G} = \{1, \dots, G\}$ denote the set of generators, indexed by g , and $\mathcal{D} = \{1, \dots, D\}$ denote the set of inflexible loads, indexed by d . Let $b(i)$, $b(g)$, and $b(d)$, denote the bus at which DC i , generator g , and load d is located, respectively. In what follows, and unless otherwise mentioned, $b \in \mathcal{B}$, $k \in \mathcal{K}$, $g \in \mathcal{G}$, and $d \in \mathcal{D}$.

We aim to minimize the total system cost, which includes the total QoS costs (for all DCs), the total power generation costs, and a penalty cost related to DC job transfer:

$$\text{Total System Cost} = \sum_i C_i^{\text{QoS}}(\mathbf{n}_i) + \sum_g c_g P_g + C^{2w}, \quad (24)$$

where P_g is the output of generator g , with cost c_g , and $C_i^{\text{QoS}}(\mathbf{n}_i)$ is given by (8) and (11) as follows:

$$\begin{aligned} C_i^{\text{QoS}}(\mathbf{n}_i) &= C_i(\theta_i(\mathbf{n}_i)) = \rho_{1,i} e^{-\rho_{2,i} \theta_i(\mathbf{n}_i)} \\ &= \rho_{1,i} e^{-2\rho_{2,i} \left(\frac{\sum_j \mu_{B_j} N_{i \rightarrow j} - \mu_{A_i}}{\sum_j \sigma_{B_j}^2 N_{i \rightarrow j} + \sigma_{A_i}^2} \right)}. \end{aligned} \quad (25)$$

In order to prevent two-way transfer of jobs between DCs, we amend the QoS cost in the objective function (24) with a penalty cost:

$$C^{2w} = \kappa^{2w} \sum_{(i,j) \in \mathcal{I}^2 \cap i \neq j} N_{i \rightarrow j} N_{j \rightarrow i}, \quad (26)$$

where κ^{2w} is a scalar. Note that the minimization of (24) will drive C^{2w} to zero, with one of $N_{i \rightarrow j}$ or $N_{j \rightarrow i}$ going to zero.

The constraints of the centralized problem are listed below. Dual variables are shown in parentheses.

The power balance constraint takes the form:

$$\sum_i (\beta_i \sum_j N_{j \rightarrow i}) + \sum_d P_d - \sum_g P_g = 0, \rightarrow (\lambda) \quad (27)$$

where β_i is the power consumption rate of an active server at DC i , and P_d is the inflexible load d that must be serviced.

The transmission line constraints are:

$$\begin{aligned} -\bar{F}_k &\leq \sum_g a_{b(g),k} P_g - \sum_d a_{b(d),k} P_d \\ &- \sum_i a_{b(i),k} \beta_i \sum_j N_{j \rightarrow i} \leq \bar{F}_k, \forall k, \rightarrow (\nu_k^-, \nu_k^+) \end{aligned} \quad (28)$$

where \bar{F}_k is the maximum flow limit of line k .

Each DC has the following active server capacity limit:

$$\sum_j N_{j \rightarrow i} \leq \bar{N}_i^{\text{active}}, \quad \forall i, \rightarrow (\zeta_i) \quad (29)$$

where $\bar{N}_i^{\text{active}}$ is the maximum amount of active servers in DC i .

All generators are subject to bound constraints:

$$\underline{P}_g \leq P_g \leq \bar{P}_g, \quad \forall g, \quad (30)$$

where \underline{P}_g and \bar{P}_g represent minimum and maximum capacity limits for generator g .

Summarizing (24)–(30), the centralized economic dispatch problem, which solves for the DC output $N_{i \rightarrow j}$, $\forall i, j$, and the generator output P_g , $\forall g$, with \mathbf{n} referring to \mathbf{n}_i , $\forall i$, and \mathbf{p} referring to P_g , $\forall g$, is formulated as follows:

Centralized Economic Dispatch problem:

$$\begin{aligned} \min_{\mathbf{n}, \mathbf{p}} \quad & \sum_i \rho_{1,i} e^{-2\rho_{2,i} \left(\frac{\sum_j \mu_{B_j} N_{i \rightarrow j} - \mu_{A_i}}{\sum_j \sigma_{B_j}^2 N_{i \rightarrow j} + \sigma_{A_i}^2} \right)} + \sum_g c_g P_g \\ & + \kappa^{2w} \sum_{(i,j) \in \mathcal{I}^2 \cap i \neq j} N_{i \rightarrow j} N_{j \rightarrow i}, \\ \text{s.t.:} \quad & (27) - (30), \end{aligned} \quad (31)$$

with $N_{i \rightarrow j} \geq 0$, $\forall i, j$, and $P_g \geq 0$, $\forall g$.

Notably, the sum of the QoS cost functions in the objective (first term) is not separable for each DC due to the potential server sharing between DCs, whereas the sum of generation costs (second term) is separable for each generator. The two-way DC cost (third term) is not separable as it considers servers belonging to multiple DCs. The power balance constraint (27) and transmission constraints (28) couple DC and generator variables. The DC active server capacity limit (29) couples variables for each DC; the generator limits (30) are generator-specific constraints.

Using the dual variables of the economic dispatch problem, we determine the Locational Marginal Prices (LMPs), which represent the marginal cost of serving an additional unit of load at a given location in the power network. For bus b , the LMP is defined as:

$$\text{LMP}_b = \lambda - \sum_k a_{b,k} (\nu_k^+ - \nu_k^-). \quad (32)$$

When there is no network congestion, *i.e.*, all line flow constraints are not binding, their dual variables will be zero

and all LMPs will be equal across all buses. The equality of all LMPs signifies that power can flow unimpeded through the network and the cost of serving an additional unit of load is the same at every bus.

Next, we transform the centralized economic dispatch optimization problem into a decentralized problem that can be solved using Lagrangian decomposition and a primal-dual algorithm.

5 DECENTRALIZED PROBLEM FORMULATION AND PROPOSED SOLUTION APPROACH

In this section, we sketch the proposed solution approach using the Lagrangian and dual functions (in Subsection 5.1), we derive the closed form gradients (in Subsection 5.2), and we describe the primal-dual algorithm (in Subsection 5.3).

5.1 Sketch of the Proposed Solution Approach

Using (31), we form the Lagrangian dualizing the power balance constraint (27), the transmission line constraints (28), and the DC active server capacity limit (29), as follows:

$$\begin{aligned} \mathcal{L}(\mathbf{n}, \mathbf{p}, \lambda, \boldsymbol{\nu}, \boldsymbol{\zeta}) &= \sum_i C_i^{\text{QoS}}(\mathbf{n}_i) + \sum_g c_g P_g + C^{2w} \\ &+ \lambda \left[\sum_i \left(\beta_i \sum_j N_{j \rightarrow i} \right) + \sum_d P_d - \sum_g P_g \right] \\ &+ \sum_k \nu_k^- \left(- \sum_g a_{b(g),k} P_g + \sum_d a_{b(d),k} P_d \right. \\ &\quad \left. + \sum_i a_{b(i),k} \beta_i \sum_j N_{j \rightarrow i} - \bar{F}_k \right) \\ &+ \sum_k \nu_k^+ \left(\sum_g a_{b(g),k} P_g - \sum_d a_{b(d),k} P_d \right. \\ &\quad \left. - \sum_i a_{b(i),k} \beta_i \sum_j N_{j \rightarrow i} - \bar{F}_k \right) \\ &+ \sum_i \zeta_i \left(\sum_j N_{j \rightarrow i} - \bar{N}_i^{\text{active}} \right), \end{aligned} \quad (33)$$

where $\mathbf{n} = \{\mathbf{n}_i; \forall i\}$, $\mathbf{p} = \{P_g; \forall g\}$, $\boldsymbol{\nu} = \{\nu_k^-, \nu_k^+; \forall k\}$, $\boldsymbol{\zeta} = \{\zeta_i; \forall i\}$. The dual problem is given as follows:

$$\max_{\lambda \in \mathbb{R}, \boldsymbol{\nu}, \boldsymbol{\zeta} \geq \mathbf{0}} q(\lambda, \boldsymbol{\nu}, \boldsymbol{\zeta}), \quad (34)$$

where $q(\lambda, \boldsymbol{\nu}, \boldsymbol{\zeta})$ is the dual function given by:

$$\begin{aligned} q(\lambda, \boldsymbol{\nu}, \boldsymbol{\zeta}) &= \inf_{\mathbf{n} \geq \mathbf{0}, \mathbf{p} \in \mathcal{P}} \mathcal{L}(\mathbf{n}, \mathbf{p}, \lambda, \boldsymbol{\nu}, \boldsymbol{\zeta}) \\ &= \inf_{\mathbf{n} \geq \mathbf{0}} \left[\sum_i C_i^{\text{QoS}}(\mathbf{n}_i) + C^{2w} + \lambda \sum_i \beta_i \sum_j N_{j \rightarrow i} + \right. \\ &\quad \left. \sum_k (\nu_k^- - \nu_k^+) \sum_i a_{b(i),k} \beta_i \sum_j N_{j \rightarrow i} + \sum_i \zeta_i \sum_j N_{j \rightarrow i} \right] \\ &+ \inf_{\mathbf{p} \in \mathcal{P}} \left[\sum_g (c_g - \lambda) P_g + \sum_k (\nu_k^+ - \nu_k^-) \sum_{g \in \mathcal{G}} a_{b(g),k} P_g \right. \\ &\quad \left. + \lambda \sum_d P_d + \sum_k (\nu_k^- - \nu_k^+) \left(\sum_d a_{b(d),k} P_d - \bar{F}_k \right) \right], \end{aligned} \quad (35)$$

with \mathcal{P} representing the generator capacity limits, i.e., $\mathcal{P}_g = [\underline{P}_g, \bar{P}_g]$, $\forall g$. Note that the dual function (each infimum) is separable in the DC and generator variables.

To solve the decentralized problem, we propose a gradient algorithm that iteratively computes estimates of the primal and dual variables, which are approximate solutions to (31) and (34). The qualitative procedure is as follows:

- 1) Compute primal variable estimates using gradient descent steps that minimize the Lagrangian (33) for a given set of dual variables.
- 2) Compute dual variable estimates using gradient ascent that moves in the direction of maximizing the dual function (35) given the primal variables.
- 3) Repeat 1) and 2) until convergence.

Problem (34) can be loosely interpreted, similarly to [40], considering an aggregator that sends price signals via optimization of a social cost metric and a set of DCs that respond to these prices by collectively determining their server allocations. Specifically, the first inner infimum problem over the responses \mathbf{n} in (35) can be viewed as a local DC optimization for the server responses where generator output and multipliers (i.e., prices) are fixed. Similarly, the second inner infimum over the generator output \mathbf{p} in (35) can be viewed as a local generator optimization for the generator output where server responses and multipliers (i.e., prices) are fixed.

5.2 Gradient Closed Form Expressions

For the gradient descent steps in primal variables, \mathbf{n} and \mathbf{p} , we derive the following closed form expressions from (33).

For the gradient $\frac{\partial \mathcal{L}}{\partial N_{i \rightarrow j}}$, we have:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial N_{i \rightarrow j}} &= \frac{\partial C_i^{\text{QoS}}(\mathbf{n}_i)}{\partial N_{i \rightarrow j}} + \kappa^{2w} N_{j \rightarrow i} \mathbf{1}_{\{i \neq j\}} \\ &\quad + \underbrace{\beta_i \left[\lambda - \sum_{k \in \mathcal{K}} a_{b(i),k} (\nu_k^+ - \nu_k^-) \right]}_{\text{LMP}_{b(i)}} + \zeta_i \mathbf{1}_{\{i=j\}}, \end{aligned} \quad (36)$$

where

$$\frac{\partial C_i^{\text{QoS}}(\mathbf{n}_i)}{\partial N_{i \rightarrow j}} = \frac{\partial C_i^{\text{QoS}}(\mathbf{n}_i)}{\partial \theta_i(\mathbf{n}_i)} \cdot \frac{\partial \theta_i(\mathbf{n}_i)}{\partial N_{i \rightarrow j}}. \quad (37)$$

Using (25), we get:

$$\frac{\partial C_i^{\text{QoS}}(\mathbf{n}_i)}{\partial \theta_i(\mathbf{n}_i)} = -\rho_{2,i} C_i^{\text{QoS}}(\mathbf{n}_i). \quad (38)$$

whereas using (15)–(18), we get:

$$\frac{\partial \theta_i(\mathbf{n}_i)}{\partial N_{i \rightarrow j}} = \frac{\sigma_{B_j}^2}{\sum_{j' \in \mathcal{I}} \sigma_{B_{j'}}^2 N_{i \rightarrow j'} + \sigma_{A_i}^2} \left[\frac{2\mu_{B_j}}{\sigma_{B_j}^2} - \theta_i(\mathbf{n}_i) \right]. \quad (39)$$

For the gradient $\frac{\partial \mathcal{L}}{\partial P_g}$, we have:

$$\frac{\partial \mathcal{L}}{\partial P_g} = c_g - \underbrace{\left[\lambda - \sum_{k \in \mathcal{K}} a_{b(g),k} (\nu_k^+ - \nu_k^-) \right]}_{\text{LMP}_{b(g)}}. \quad (40)$$

For the gradient ascent steps in dual variables, λ , $\boldsymbol{\nu}$, and $\boldsymbol{\zeta}$, we derive the following closed form expressions from (35).

For the gradient $\frac{\partial q}{\partial \lambda}$, we have:

$$\frac{\partial q}{\partial \lambda} = \overbrace{\sum_{i \in \mathcal{I}} (\beta_i \sum_{j \in \mathcal{I}} N_{j \rightarrow i}) + \sum_{d \in \mathcal{D}} P_d - \sum_{g \in \mathcal{G}} P_g}^{\text{Power balance}}. \quad (41)$$

For the gradients $\frac{\partial q}{\partial \nu_k^-}$ and $\frac{\partial q}{\partial \nu_k^+}$, we have:

$$\frac{\partial q}{\partial \nu_k^-} = -F_k - \bar{F}_k, \quad \frac{\partial q}{\partial \nu_k^+} = F_k - \bar{F}_k, \quad (42)$$

where

$$F_k = \overbrace{\sum_{g \in \mathcal{G}} a_{b(g),k} P_g - \sum_{d \in \mathcal{D}} a_{b(d),k} P_d}^{\text{Flow along line } k} - \sum_{i \in \mathcal{I}} a_{b(i),k} \beta_i \sum_{j \in \mathcal{I}} N_{j \rightarrow i}. \quad (43)$$

For the gradient $\frac{\partial q}{\partial \zeta_i}$, we have:

$$\frac{\partial q}{\partial \zeta_i} = \overbrace{\sum_{j \in \mathcal{I}} N_{j \rightarrow i} - \bar{N}_i^{\text{active}}}^{\text{Active servers}}. \quad (44)$$

5.3 Primal-Dual Algorithm Description

Let α denote a stepsize, which can be selected individually for each variable update, using common stepsize rules. Let Φ denote the number of gradient descent iterations to obtain primal variable estimates, indexed by ϕ . Let τ denote the iteration counter of the primal-dual algorithm, with T_{\max} being the maximum number of iterations. Let E_{\min} be a threshold used as a termination condition. Let $w|_{\tau}$ denote the evaluation of some quantity w at iteration τ ; similarly $w|_{\phi,\tau}$ is evaluated at the ϕ -th (inner) iteration of iteration τ . Let also $[w]^+ = \max\{0, w\}$, and $[w]_{\mathcal{P}_g} = \min\{\max\{w, \underline{P}_g\}, \bar{P}_g\}$.

Algorithm 1 details the proposed primal-dual algorithm. The primal updates (45) — see lines 5 to 7 — are executed Φ times, using the gradient expressions in (36)–(40). We note that each computation within (inner) primal update iteration ϕ in both (45a) and (45b) can be performed in parallel. The dual updates (46) — see line 8 — are then executed once, given the primal variable estimates obtained earlier, and can be also executed in parallel. The algorithm terminates (see line 4) when the maximum number of iterations T_{\max} is reached or the maximum norm ϵ (computed in line 9 as the squared Euclidean norm of the dual variable differences in adjacent iterations) is below the threshold E_{\min} .

6 NUMERICAL EXPERIMENTS

In this section, we demonstrate our approach on a standard power network from the literature and quantify the system-wide benefits from the incorporation of DC flexibility in the economic dispatch problem. We compare two settings: No Sharing (NS), in which the DCs *cannot* share server resources, and With Sharing (WS), in which DCs *can* share server resources.

In what follows, we present results for a Base Case with identical DCs (in Subsection 6.1), we explore variants for different QoS and efficiency parameters (in Subsection 6.2), and we validate the performance of our approach for a higher number of DCs (in Subsection 6.3).

Algorithm 1 Proposed Primal-Dual Algorithm

1: **Input:** α (step size), Φ (primal iterations), T_{\max} (maximum number of primal-dual iterations), E_{\min} (threshold for maximum dual norm), $\mathbf{n}|_0, \mathbf{p}|_0, \lambda|_0, \nu|_0, \zeta|_0$ (initial estimates of primal and dual variables).

2: **Output:** $\mathbf{n}, \mathbf{p}, \lambda, \nu, \zeta$ (primal and dual variables).

3: Set $\tau = 0$.

4: **While** $\epsilon > E_{\min}$ **and** $\tau \leq T_{\max}$, **do:**

5: *Primal Updates:*

6: **For** $\phi = 1 \dots \Phi$, **do:** Using (36)–(40), derive

$$N_{i \rightarrow j}|_{\tau+1} = \left[N_{i \rightarrow j}|_{\tau} - \alpha \frac{\partial \mathcal{L}}{\partial N_{i \rightarrow j}} \Big|_{\phi, \tau} \right]^+, \quad \forall i, j, \quad (45a)$$

$$P_g|_{\tau+1} = \left[P_g|_{\tau} - \alpha \frac{\partial \mathcal{L}}{\partial P_g} \Big|_{\phi, \tau} \right]_{\mathcal{P}_g}, \quad \forall g. \quad (45b)$$

7: **end**

8: *Dual Updates:* Using (41)–(44), derive

$$\lambda|_{\tau+1} = \lambda|_{\tau} + \alpha \frac{\partial q}{\partial \lambda} \Big|_{\tau}, \quad (46a)$$

$$\nu_k^-|_{\tau+1} = \left[\nu_k^-|_{\tau} + \alpha \frac{\partial q}{\partial \nu_k^-} \Big|_{\tau} \right]^+, \quad \forall k, \quad (46b)$$

$$\nu_k^+|_{\tau+1} = \left[\nu_k^+|_{\tau} + \alpha \frac{\partial q}{\partial \nu_k^+} \Big|_{\tau} \right]^+, \quad \forall k, \quad (46c)$$

$$\zeta_i|_{\tau+1} = \left[\zeta_i|_{\tau} + \alpha \frac{\partial q}{\partial \zeta_i} \Big|_{\tau} \right]^+, \quad \forall i. \quad (46d)$$

9: $\epsilon = \max(\{\|\lambda|_{\tau+1} - \lambda|_{\tau}\|_2^2, \|\nu_k^-|_{\tau+1} - \nu_k^-|_{\tau}\|_2^2, \|\nu_k^+|_{\tau+1} - \nu_k^+|_{\tau}\|_2^2, \|\zeta_i|_{\tau+1} - \zeta_i|_{\tau}\|_2^2\})$, $\tau \leftarrow \tau + 1$.




Fig. 2: Depiction of the 5-bus PJM system.

6.1 Base Case

In this subsection, we describe the Base Case power network and the DC allocation and parameters (in Subsection 6.1.1), we present the computational setup (in Subsection 6.1.2), and we discuss the Base Case results (in Subsection 6.1.3).

6.1.1 Input Data

TABLE 2: Generator and network parameters, 5-bus system.

Generator	G1	G2	G3	G4	G5
Cost (\$/MWh)	14	15	30	40	10
Capacity (MW)	40	170	520	200	600
Shift Factors:					
	Bus				
Line	A	B	C	D	E
AB	0.1939	-0.4759	-0.3490	0	0.15954
AD	0.4376	0.2583	0.1895	0	0.3600
AE	0.3685	0.2176	0.1595	0	-0.5195
BC	0.1939	0.5241	-0.3490	0	0.1595
CD	0.1939	0.5241	0.6510	0	0.1595
DE	-0.3685	-0.2176	-0.1595	0	-0.4805

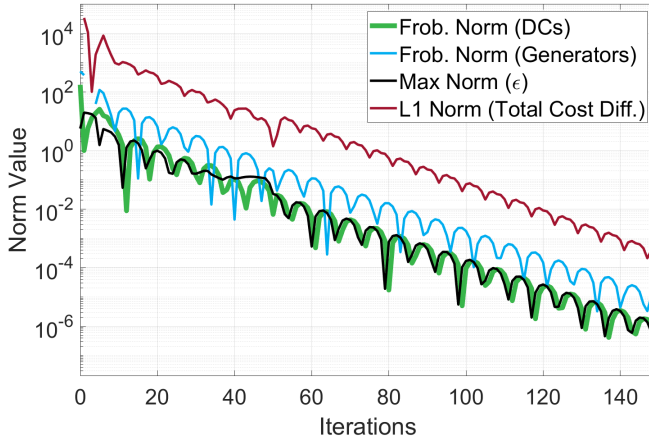


Fig. 3: Convergence illustration of Algorithm 1, Base Case, WS setting.

The Base Case employs the 5-bus PJM system displayed in Fig. 2 [53]. The system includes 5 buses (labeled as A, B, C, D, and E), 6 transmission lines (labeled as AB, BC, CD, DE, AE, and AD), 5 generators (labeled as G1, G2, G3, G4, and G5), and 3 inflexible loads (labeled as L1, L2, and L3) located at buses B, C, and D, with values 300, 300, and 400MW, respectively. Transmission lines DE and AB have capacity limits of 240 and 400MW, respectively. Generator costs and capacities, as well as shift factors, are shown in Table 2.

In the Base Case, we allocated 3 DCs, labeled as DC1, DC2, and DC3, at buses A, B, and C, respectively, with identical characteristics. Specifically, $\forall i$, DC i has QoS cost coefficients $\rho_{1,i} = 7,500$ and $\rho_{2,i} = 0.002$, arrival and departure quantities (per hour) $\mu_{A_i} = 100$, $\sigma_{A_i}^2 = 0.5$, $\mu_{B_i} = 10$, and $\sigma_{B_i}^2 = 0.02$, active server power consumption rate $\beta_i = 2$, and capacity limit $\bar{N}_i^{\text{active}} = 300$.

6.1.2 Computational Setup

We implemented Algorithm 1 in Matlab R2024a and ran it on an i9 processor at 2.4 GHz with 64 GB RAM. Algorithm 1 parameters were set as follows: $\alpha = 0.05$, $\Phi = 100$, $E_{\min} = 10^{-7}$, and $T_{\max} = 10,000$. Initial values of primal variables were set to zero. Initial values of dual variables were randomly selected between 0 and 1.

In Fig. 3, we demonstrate the execution of Algorithm 1, for the Base Case, WS setting. We present the value of the maximum norm ϵ (stopping criterion) per iteration whose threshold (E_{\min}) is reached in about 150 iterations. We also show the Frobenius norms (between adjacent iterations) for DC and generator outputs, which, upon convergence,

TABLE 3: LMP, generator and DC outputs; Base Case.

No Sharing (NS)						
Bus	A	B	C	D	E	
LMP (\$/MWh)	16.98	26.38	30	39.94	10	
Generator	G1	G2	G3	G4	G5	
Gen. (MWh)	40	170	492.77	0	543.75	
Tot. Gen. (MWh)	1,246.52				Tot. Gen. Cost (\$)	23,300
	Avg. Gen. Cost (\$/MWh)				18.69	
	DC1	DC2	DC3	Total		
Active Servers	48.60	38.61	36.05	123.26		
DC Load (MWh)	97.20	77.22	72.10	246.52		
DC Cost (\$)	2,627.5	3,050.5	3,194.7	8,872.7		
With Sharing (WS)						
Bus	A	B	C	D	E	
LMP (\$/MWh)	30	30	30	30	30	
Generator	G1	G2	G3	G4	G5	
Gen. (MWh)	40	170	406.3	0	600	
Tot. Gen. (MWh)	1,216.3				Tot. Gen. Cost (\$)	21,299
	Avg. Gen. Cost (\$/MWh)				17.51	
DC To / From	DC1	DC2	DC3	Total		
DC1	36.05	32.92	34.29	102.48		
DC2	0	3.13	0.16	4.14		
DC3	0	0	1.60	1.53		
Used Servers	36.05	36.05	36.05	108.15		
DC Load (MWh)	204.96	8.28	3.06	216.30		
DC Cost (\$)	3,194.7	3,194.7	3,194.7	9,584.1		

reach values in the order of 10^{-7} and 10^{-5} , respectively. Furthermore, we present the L1 norm (between adjacent iterations) for the total cost difference, which reaches values in the order of 10^{-3} . Computational time for this admittedly small example was in the order of milliseconds (msec). We further elaborate on the computational performance of our approach in Subsection 6.3.

6.1.3 Base Case Results (Identical DCs)

In Table 3, we present the Base Case outputs for LMPs, generators and DCs. When the DCs do not share server resources, i.e., in the NS setting, line DE reaches its limit of 240MW, which prevents the full dispatch of the cheapest generator G5 (note that G5, with $c_5 = \$10/\text{MWh}$ produces 543.75MWh instead of at full capacity, i.e., 600MWh). LMPs are different at each bus ranging from \$10 to \$39.94/MWh, hence, DCs effectively see a different price at each location, i.e., DC1, DC2, and DC3 see a price of \$16.98, \$26.38, and \$30/MWh, respectively. Since all DCs have identical characteristics, DC1 that sees the lowest price has more active servers (48.6), followed by DC2 (38.61) and then by DC3 (36.05).

In Fig. 4, we visualize the DC QoS cost function and scaled derivative (*w.r.t.* $N_{i \rightarrow j}$) for the Base Case, NS setting (DCs have the same cost function parameters, hence the single curve). Derivative scaling is performed with $\frac{1}{\beta_i}$ to account for the energy consumed (in MWh). Indeed, the magnitudes of the QoS cost scaled derivative at the optimal amount of servers match the LMP values at the buses where each respective DC is located at. This is expected since LMPs represent marginal costs of delivering electricity at a given location.²

When DCs can share resources, i.e., in the WS setting, we observe that line DE congestion is relieved and LMPs

2. For clarity, set the derivative in (36) equal to zero and discard the 2-way transfer terms (zero) and ζ_i (also zero since the active server capacity limit is not binding), to obtain the aforementioned equality.

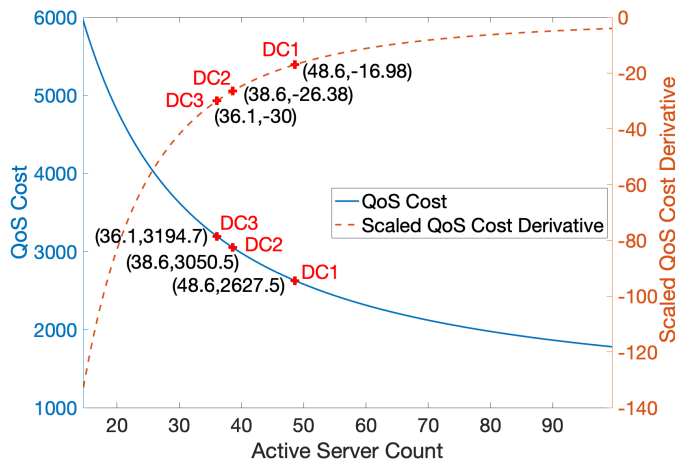


Fig. 4: DC QoS cost function and scaled derivative, Base Case, NS setting.

TABLE 4: LMP, generator and DC outputs; higher DC QoS cost coefficient ($\rho_1 = 15,000$).

		No Sharing (NS)					
Bus		A	B	C	D	E	
LMP (\$/MWh)		16.99	26.42	30.04	40	10	
Generator		G1	G2	G3	G4	G5	
Gen. (MWh)		40	170	520	30.41	591.30	
Tot. Gen. (MWh)		1,351.70	Tot. Gen. Cost (\$)			25,839	
		Avg. Gen. Cost (\$/MWh)					19.09
Active Servers		DC1	DC2	DC3	Total		
DC Load (MWh)		68.88	55.23	51.74	175.85		
DC Cost (\$)		137.76	110.46	103.48	351.70		
		4,278.7	4,857.8	5,054.2	14,190.7		
		With Sharing (WS)					
Bus		A	B	C	D	E	
LMP (\$/MWh)		30	30	30	30	30	
Generator		G1	G2	G3	G4	G5	
Gen. (MWh)		40	170	500.64	0	600	
Tot. Gen. (MWh)		1,310.64	Tot. Gen. Cost (\$)			24,129	
		Avg. Gen. Cost (\$/MWh)					18.41
DC To / From		DC1	DC2	DC3	Total		
DC1		51.76	0	42.25	118.61		
DC2		0.01	51.77	0.38	32.83		
DC3		0	0	9.14	3.87		
Used Servers		51.77	51.77	51.77	155.31		
DC Load (MWh)		237.22	65.66	7.74	310.64		
DC Cost (\$)		5,052.2	5,052.2	5,052.2	15,156.6		

are equal to \$30/MWh at all buses (G3 is the marginal unit). The cheapest generator G5 is now dispatched at full capacity producing 600MWh. In addition, the total generation in the WS setting is about 30MWh less than that of the NS setting, due to the smaller amount of active servers in the WS setting. Additionally, server sharing results in the average generation cost decreasing from \$18.69/MWh to \$17.51/MWh. Note that the total amount of active servers when DCs share resources is equal in each DC. Since all DCs see the same price and they have identical characteristics, this outcome is expected. The congestion relief in line DE is achieved by using more active servers in DC1, indicating the favorable location of DC1 under the specific system conditions. In turn, this allows for the cheapest generator G5 to be dispatched at full capacity. The larger DC cost (by \$711.4) is more than offset by the lower generation cost (by \$2,001), resulting in an overall cost decrease equal to \$1,289.6.

TABLE 5: DC outputs; one DC with efficiency $\beta = 1$.

		No Sharing (NS)			
Effic. DC1		DC1	DC2	DC3	Total
	Active servers	68.91	38.61	36.05	143.57
	DC Load (MWh)	68.91	77.22	72.1	218.23
Tot. DC Cost (\$)	8,384.1	Tot. Gen. Cost (\$)		22,850	
		With Sharing (WS)			
Effic. DC1	Used servers	DC1	DC2	DC3	Total
	Active servers	68.34	68.34	68.34	205.02
	DC Load (MWh)	205.02	0	0	205.02
	Tot. DC Cost (\$)	6,446.1	Tot. Gen. Cost (\$)		20,961
		Tot. Benefit of Sharing (\$)			3,827
		No Sharing (NS)			
Effic. DC2		DC1	DC2	DC3	Total
	Active servers	48.60	55.26	36.05	139.91
	DC Load (MWh)	97.2	55.26	72.1	224.56
Tot. DC Cost (\$)	8,250.2	Tot. Gen. Cost (\$)		22,751	
		With Sharing (WS)			
Effic. DC2	Used servers	DC1	DC2	DC3	Total
	Active servers	55.26	55.26	55.26	165.78
	DC Load (MWh)	0	165.78	0	165.78
	Tot. DC Cost (\$)	7,283.4	Tot. Gen. Cost (\$)		21,854
		Tot. Benefit of Sharing (\$)			1,863.8
		No Sharing (NS)			
Effic. DC3		DC1	DC2	DC3	Total
	Active servers	48.60	38.61	51.77	138.98
	DC Load (MWh)	97.2	77.22	51.77	226.2
Tot. DC Cost (\$)	8,204.1	Tot. Gen. Cost (\$)		22,721	
		With Sharing (WS)			
Effic. DC3	Used servers	DC1	DC2	DC3	Total
	Active servers	51.77	51.77	51.77	155.31
	DC Load (MWh)	0	0	155.31	155.31
	Tot. DC Cost (\$)	7,578.3	Tot. Gen. Cost (\$)		22,140
		Tot. Benefit of Sharing (\$)			1,206.8

6.2 Base Case Variants

In this subsection, we present and discuss results for Base Case variants with identical DCs but higher QoS cost (Subsection 6.2.1), one DC with higher efficiency (Subsection 6.2.2), and one DC (namely DC1) exhibiting both higher efficiency and QoS cost (Subsection 6.2.3).

6.2.1 Identical DCs with Higher QoS Cost

In Table 4, we present the results for a Base Case variant with a higher QoS cost coefficient for all DCs, $\rho_{1,i} = 15,000, \forall i$. Higher QoS cost coefficients result in more active servers (175.85 vs. 123.26, and 155.31 vs. 108.15, for the NS and WS settings, respectively) compared to the results in Table 3, in order to keep QoS costs low. More active servers result in higher DC loads (by about 105.18 and 94.34MWh, for the NS and WS settings, respectively). In the NS setting, line DE is congested and the most expensive generator G4 is dispatched; LMPs at buses A, B, and C are similar to the Base Case. In the WS setting, the total generation drops by about 40MWh, congestion is relieved (all LMPs become equal), and the total amount of active servers drops (compared to the NS setting). The larger DC cost (by \$965.9), in the WS setting, is more than offset by the lower generation cost (by \$1,721), thus resulting in an overall cost decrease equal to \$755.1, which is lower than that of the Base Case (due to the higher amount of servers).

6.2.2 One DC with Higher Efficiency

Next, we present variants that alter the Base Case DC efficiencies in Tables 5 and 6. Efficiency is defined based on

TABLE 6: DC outputs; one DC with efficiency, $\beta = 1.5$.

		No Sharing (NS)				
		DC1	DC2	DC3	Total	
Effic. DC1	Active servers	55.26	38.61	36.05	130.92	
	DC Load (MWh)	84.39	77.22	72.1	233.71	
	Tot. DC Cost (\$)	8,647.3	Tot. Gen. Cost (\$)		23,113	
			With Sharing (WS)			
			DC1	DC2	DC3	Total
	Used servers	45.56	45.56	45.56	136.68	
Active servers	136.68	0	0	136.68		
DC Load (MWh)	205.02	0	0	205.02		
Tot. DC Cost (\$)	8,211.9	Tot. Gen. Cost (\$)		20,961		
		Tot. Benefit of Sharing (\$)			2,587.4	
		No Sharing (NS)				
		DC1	DC2	DC3	Total	
Effic. DC2	Active servers	48.60	44.90	36.05	129.55	
	DC Load (MWh)	97.2	67.36	72.1	236.65	
	Tot. DC Cost (\$)	8,585	Tot. Gen. Cost (\$)		22,751	
			With Sharing (WS)			
			DC1	DC2	DC3	Total
	Used servers	44.06	44.06	44.06	132.18	
Active servers	94.71	37.47	0	132.18		
DC Load (MWh)	189.42	56.20	0	245.62		
Tot. DC Cost (\$)	8,391	Tot. Gen. Cost (\$)		22,178		
		Tot. Benefit of Sharing (\$)			767	
		No Sharing (NS)				
		DC1	DC2	DC3	Total	
Effic. DC3	Active servers	48.60	38.61	42.00	129.21	
	DC Load (MWh)	97.2	77.22	63.00	237.42	
	Tot. DC Cost (\$)	8,563.6	Tot. Gen. Cost (\$)		23,057	
			With Sharing (WS)			
			DC1	DC2	DC3	Total
	Used servers	42.00	42.00	42.00	126.00	
Active servers	102.51	0	23.49	126.00		
DC Load (MWh)	205.02	0	35.22	240.24		
Tot. DC Cost (\$)	8,656.8	Tot. Gen. Cost (\$)		22,017		
		Tot. Benefit of Sharing (\$)			946.8	

the value of the power consumption rate of an active server, β_i . Each Table presents DC outputs for variants where a single DC is set to be the most efficient; it is segmented into three sub-Tables each denoting in the first column the DC that becomes the most efficient (gets assigned the smallest β_i). Conceptually, a smaller β_i denotes more efficient servers at DC i , and, in turn, more active servers at DC i can be used to carry out workload.

Table 5 presents a variant with the most efficient DC having $\beta = 1$. As an example, when DC1 is the most efficient, it has more active servers than DC2 and DC3 in the NS setting, whereas all jobs are sent to DC1 in the WS setting — see the zero active server quantities. Comparing the WS to the NS setting, the total amount of active servers is increased, the total DC load is reduced, and both the total DC and generation costs are reduced. The same trends appear in variants where DC2 or DC3 are the most efficient. Notably, however, line DE remains congested under all variants, although there is some generator re-dispatching in favor of the cheapest generator G5. Furthermore, we note that the location of the most efficient DC does matter. In the NS setting, we observe more active servers when DC1 (at bus A) is the most efficient, followed by DC2 (at bus B), followed by DC3 (at bus C), with respective benefits of sharing (in the WS setting), \$3,827, \$1,863.8, and \$1,206.8.

Table 6 presents a variant with the most efficient DC having $\beta = 1.5$, i.e., it is less efficient compared to Table 5. Similarly to Table 5, more jobs are directed to the most efficient DC, however, the amount of active servers is re-

TABLE 7: LMP, generation and DC outputs; DC1 higher efficiency and QoS cost coefficient ($\beta_1 = 1$, $\rho_{1,1} = 37,500$).

		No Sharing (NS)					
Bus		A	B	C	D	E	
LMP (\$/MWh)		16.98	26.38	30	39.94	10	
			DC1	DC2	DC3	Total	
	Active Servers	151.40	38.61	36.05	226.06		
	DC Load (MWh)	151.40	77.22	72.1	300.72		
Tot. DC Cost (\$)		13,792.3			Tot. Gen. Cost (\$)		24,251
		With Sharing (WS)					
Bus		A	B	C	D	E	
LMP (\$/MWh)		30	30	30	30	30	
			DC1	DC2	DC3	Total	
	Used Servers	114.78	51.77	51.77	218.33		
	Active Servers	218.33	0	0	218.33		
DC Load (MWh)	218.33	0	0	218.33			
Tot. DC Cost (\$)		13,426.1			Tot. Gen. Cost (\$)		21,360
		Tot. Benefit of Sharing (\$)				3,257.2	

duced — though still higher than the Base Case (see Table 3). Comparing the WS to the NS setting, we still observe that all jobs are sent to DC1 when this is the most efficient, however, DC1 still receives most of the jobs even when DC2 or DC3 are most efficient. Apparently, the location of DC1 plays an important role in terms of reducing the overall cost, when the difference in efficiency is smaller, highlighting the trade-off between the DC location and efficiency. Sharing leads to a larger amount of active servers when either DC1 or DC2 are most efficient, and to smaller DC load only when DC1 is most efficient. However, the lower generation cost in all variants results in a positive monetary benefit (higher when DC1 is most efficient), although congestion is not relieved (but still re-dispatching reduces the generation cost).

6.2.3 One DC with Higher Efficiency and Higher QoS Cost

Lastly, we present in Table 7 a variant with DC1 having higher efficiency, $\beta_1 = 1$, whereas $\beta_2 = \beta_3 = 2$, but also a higher QoS cost coefficient $\rho_{1,1} = 37,500$, whereas $\rho_{1,2} = \rho_{1,3} = 7,500$. Essentially, this compares to Table 5 with DC1 most efficient, but with higher QoS cost, which in turn results in an increase of its active servers (in the NS setting, 151.40 vs. 68.91). In the WS setting, DC2 and DC3 send their entire workload to the most efficient DC1, and, surprisingly, due to the higher DC load at bus A (as a result of the higher QoS cost of DC1) congestion is relieved — see LMPs — illustrating the much involved impact on the network flows.

6.3 Extensions

In what follows, we present additional experiments that aim at validating the proposed algorithm for a higher number of DCs with different characteristics and locations.

In this respect, we increased the DCs at each location by a factor of 10, i.e., 30 DCs in total, and by a factor of 100, i.e., 300 DCs in total. To keep the costs scaled, we decreased the QoS cost function parameter ρ_1 and efficiency parameter β by the same factors. We further randomly selected parameters μ_A , μ_B , and ρ_1 by sampling within a 10% and 50% margin around the nominal values of the Base Case, and we reduced the threshold E_{\min} to 10^{-5} . For the 300 DCs, WS setting, we reduced the stepsize α to 0.005, and the number of primal update iterations Φ to 20. The results are shown in Table 8.

TABLE 8: Results for a Higher Number of DCs at 3 locations.

Total Number of DCs	30		300	
	10%	50%	10%	50%
No Sharing (NS)				
Iterations	5,189	5,349	8,605	8,794
Time per Iter. (msec)	0.4	0.5	19.4	20.2
Tot. Gen. Cost (\$)	23,338	23,220	23,318	23,176
Tot. DC Cost (\$)	9,015	10,094	8,848	9,138
With Sharing (WS)				
Iterations	5,152	5,734	10,000	10,000
Time per Iter. (msec)	25.2	28.4	551.6	520.6
Tot. Gen. Cost (\$)	21,302	21,020	21,340	21,225
Tot. DC Cost (\$)	8,759	6,226	8,771	6,046
Benefit of Sharing (\$)	2,292	6,069	2,055	5,043

Iterations (τ) have now increased in the thousands. For the NS setting, computational times per iteration τ are still minimal (around half a msec) for 30 DCs, whereas they increase to about 20 msec for 300 DCs. Note that the time consuming part of the algorithm per iteration τ refers to the primal updates, which are executed multiple times ($\Phi = 100$) compared to the dual updates that are executed only once. An increase from 30 to 300 DCs increases the primal variables in (45a) by a factor of 10, whereas the gradient expressions also involve the computations of larger sums over DCs (also by a factor of 10, which in turn affects the time per variable update). For the WS setting, computational times have now increased to about 28 msec for 30 DCs and to about 550 msec for 300 DCs. We note, in a parenthesis, that we further experimented with different random seeds and the variations we observed were negligible. Compared to the NS setting, the WS setting increases the number of DC primal variables by a factor that equals to the number of DCs, but executes 5 times less primal update iterations (because $\Phi = 20$ instead of 100). Note, however, that the variable updates are amenable to parallelization. Updating DC primal variables in parallel would reduce the computational time of each iteration ϕ by a factor equal to the square of the number of DCs. For example, the computational times for 300 DCs could be reduced by at least four orders of magnitude. Hence, running an algorithm in parallel for 10,000 iterations (which in this case reached a value of the maximum norm ϵ in the order of 10^{-4}) would converge in less than 1 sec.

In Table 9, we present results considering 10 and 100 DCs in all 5 bus locations, i.e., in total 50 and 500 DCs, respectively. The setup and parameters for 50 and 500 DCs are selected similarly to the aforementioned for 30 and 300 DCs, respectively. In addition, we reduced the limit of line AB from 400MW to 200MW to increase congestion (which now also appears in line AB). Although DC sharing does not eliminate congestion (only line DE congestion is relieved for the 50 DCs, WS setting), the benefit of sharing is roughly proportional to the number of DCs (compared to the results in Table 8). Computational times are in the same order of magnitude and somewhat increased due to the increase in the number of DCs (whereas the impact on performance from our experimentation with different random seeds was negligible). However, parallelization can still drastically reduce the computational time. For illustration purposes, we present the convergence of Algorithm 1 for 50 DCs, 50% margin, WS setting in Fig. 5. The results show that when the

TABLE 9: Results for a Higher Number of DCs at 5 locations.

Total Number of DCs	50		500	
	10%	50%	10%	50%
No Sharing (NS)				
Iterations	2,408	2,155	2,419	2,331
Time per Iter. (msec)	0.4	0.4	13.5	13.0
Tot. Gen. Cost (\$)	24,776	24,790	24,767	24,738
Tot. DC Cost (\$)	18,861	20,522	18,723	19,181
With Sharing (WS)				
Iterations	5,713	3,435	10,000	10,000
Time per Iter. (msec)	76.7	62.4	1165.1	1123.9
Tot. Gen. Cost (\$)	24,710	24,381	24,710	24,359
Tot. DC Cost (\$)	15,293	10,215	15,462	10,676
Benefit of Sharing (\$)	3,633	10,716	3,318	8,884

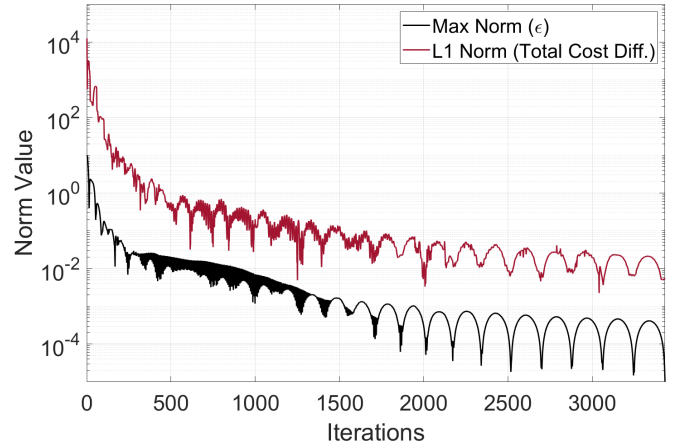


Fig. 5: Convergence illustration of Algorithm 1, 50 DCs, 50% margin, WS setting.

maximum norm ϵ reaches the 10^{-5} threshold, the L1 norm of the total cost difference is already in low values, in the order of 10^{-2} .

7 CONCLUSION

In this work, we proposed a DC model based on queuing theory and a convex QoS-based cost function. We presented a novel centralized economic dispatch formulation involving DCs as flexible loads, whose solution yields optimal generator and DC workload outputs, while respecting transmission constraints. We then showcased a distributed approach to solving the economic dispatch problem using a Lagrangian decomposition and a dual gradient ascent algorithm. Finally, we presented experimental results and demonstrated the system-wide benefits from incorporating DC flexibility in an economic dispatch problem.

Our results highlighted the trade-offs between DC location, QoS, and efficiency. Indeed, the location of a DC should be viewed in the context of “dynamic” power network conditions. Congestion depends on the location of loads and generation, and the intermittent nature of renewable energy creates “unsystematic” congestion and “dynamic” conditions, which may render one location more favorable at a certain time, and another location at another time. Hence, it is imperative to consider the power network constraints in an economic dispatch problem, and harvest the DC flexibility to shift workload over space given the

specific conditions at each time period, to ensure system-wide efficient and “greener” allocations. Arguably, our results provided useful insights on the “coupled” data and power networks.

Future work is directed to fewer assumptions about the DC models, such as no prior knowledge about the DC cost function model and/or coefficients, and learning methods employing network interaction to estimate DC cost function structure and parameters in the spirit of our prior work [40]. Shifting DC workload over time, thus considering the DC storage-like capabilities in the economic dispatch problem, is another interesting direction for further research, which will also benefit greatly from distributed approaches.

REFERENCES

- [1] A. Holst, “Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2025,” *Statista*.
- [2] C. Petrov, “25+ impressive big data statistics for 2022,” *Tech Jury*.
- [3] D. Reinsel, J. Gantz, and J. Rydning, “The digitization of the world - from edge to core,” Nov. 2018, iDC White Paper - #US44413318.
- [4] G. Levanon, “Remote work: The biggest legacy of Covid-19,” *Forbes*.
- [5] A. Jay, “Number of internet of things (iot) connected devices worldwide 2022/2023: Breakdowns, growth & predictions,” *Finances Online*.
- [6] “IDC forecasts improved growth for global AI market in 2021,” *International Data Corporation (IDC)*.
- [7] “Data center infrastructure design,” 2022, whitepaper from FS.
- [8] S. U. Khan and A. Y. Zomaya, *Handbook on Data Centers*. Springer, 2015.
- [9] A. Velimirovic, “What is a virtual data center (VDC)?” *Phoenix-NAP*, July 2021.
- [10] “Data center market by component and geography - forecast and analysis 2022-2026,” Feb. 2022, technavio; 120 page report.
- [11] S. Gray, “The data center industry is booming,” *Grayway*, Apr. 2020.
- [12] H. Rong, H. Zhang, S. Xiao, C. Li, and C. Hu, “Optimizing energy consumption for data centers,” *Renew. Sustain. Energy Rev.*, vol. 58, pp. 674–691, 2016.
- [13] G. Kamiya, “Data centres and data transmission networks,” *International Energy Agency*, Sep. 2022.
- [14] R. E. Brown *et al.*, “Report to congress on server and data center energy efficiency: Public law 109-431,” Aug. 2007.
- [15] G. Ghatikar, V. Ganti, N. Matson, and M. A. Piette, “Demand response opportunities and enabling technologies for data centers: Findings from field studies,” Aug. 2012.
- [16] D. Leprince-Ringuet, “Green tech: Google shifts data center workloads to follow the sun - and the wind,” *ZDNet*.
- [17] Y. Sverdlik, “Google to shift workloads between data centers to follow clean energy,” *Data Center Knowledge*.
- [18] Z. Cao, X. Zhou, X. Wu, Z. Zhu, T. Liu, J. Neng, and Y. Wen, “Data center sustainability: Revisits and outlooks,” *IEEE Transactions on Sustainable Computing*, vol. 9, no. 3, pp. 236–248, 2024.
- [19] W. E. Gribga, A. Blavette, and A.-C. Orgerie, “Renewable energy in data centers: The dilemma of electrical grid dependency and autonomy costs,” *IEEE Transactions on Sustainable Computing*, vol. 9, no. 3, pp. 315–328, 2024.
- [20] H. Dou, Y. Qi, W. Wei, and H. Song, “Carbon-aware electricity cost minimization for sustainable data centers,” *IEEE Trans. Sustain. Comput.*, vol. 2, no. 2, pp. 211–223, Apr. 2017.
- [21] N. Hogade, S. Pasricha, H. J. Siegel, A. A. Maciejewski, M. A. Oxley, and E. Jonardi, “Minimizing energy costs for geographically distributed heterogeneous data centers,” *IEEE Trans. Sustain. Comput.*, vol. 3, no. 4, pp. 318–331, Oct. 2018.
- [22] W. Lin, G. Wu, X. Wang, and K. Li, “An artificial neural network approach to power consumption model construction for servers in cloud data centers,” *IEEE Trans. Sustain. Comput.*, vol. 5, no. 3, pp. 329–340, July 2020.
- [23] B. Zou, H. Zhang, G. Chen, and Y. Song, “Optimal power scheduling of data centers with deferrable computation requests,” in *Proc. IEEE 5th Conf. Energy Internet and Energy Syst. Integration (EI2)*, Oct. 2021, pp. 721–726.
- [24] J. Wan, J. Zhou, and X. Gui, “Sustainability analysis of green data centers with chcp and waste heat reuse systems,” *IEEE Trans. Sustain. Comput.*, vol. 6, no. 1, pp. 155–167, Jan 2021.
- [25] M. Xu, A. N. Toosi, and R. Buyya, “A self-adaptive approach for managing applications and harnessing renewable energy for sustainable cloud computing,” *IEEE Trans. Sustain. Comput.*, vol. 6, no. 4, pp. 544–558, Oct. 2021.
- [26] I. F. de Nardin, P. Stolf, and S. Caux, “Evaluation of heuristics to manage a data center under power constraints,” in *Proc. IEEE 13th Intl. Green Sustain. Comput. Conf. (IGSC)*, Oct. 2022, pp. 1–8.
- [27] Z. Zhou, K. Li, J. Abawajy, M. Shojafar, M. Chowdhury, F. Li, and K. Li, “An adaptive energy-aware stochastic task execution algorithm in virtualized networked datacenters,” *IEEE Trans. Sustain. Comput.*, vol. 7, no. 2, pp. 371–385, Apr. 2022.
- [28] M. D. M. da Silva, A. Gamatié, G. Sassatelli, M. Poss, and M. Robert, “Optimization of data and energy migrations in mini data centers for carbon-neutral computing,” *IEEE Trans. Sustain. Comput.*, vol. 8, no. 1, pp. 68–81, Jan 2023.
- [29] K. M. U. Ahmed, M. H. J. Bollen, and M. Alvarez, “A stochastic approach to determine the optimal number of servers for reliable and energy efficient operation of data centers,” *IEEE Trans. Sustain. Comput.*, vol. 8, no. 2, pp. 153–164, Apr. 2023.
- [30] 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, 2013.
- [31] Z. Liu, I. Liu, S. Low, and A. Wierman, “Pricing data center demand response,” in *Proc. ACM Intl. Conf. Measurement Modeling Computer Syst. (SIGMETRICS)*, Austin, TX, USA, June 2014.
- [32] N. Chen, X. Ren, S. Ren, and A. Wierman, “Greening multi-tenant data center demand response,” *Perform. Eval.*, vol. 43, no. 2, 2015.
- [33] L. Niu and Y. Guo, “Enabling reliable data center demand response via aggregation,” in *7th Intl. Conf. Future Energy Syst. - ACM e-Energy '16*, Waterloo, Ontario, Canada, June 2016.
- [34] S. Bahrami, V. W. S. Wong, and J. Huang, “Data center demand response in deregulated electricity markets,” *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 2820–2832, May 2019.
- [35] M. Chen, C. Gao, M. Song, S. Chen, D. Li, and Q. Liu, “Internet data centers participating in demand response: A comprehensive review,” *Renew. Sustain. Energy Rev.*, vol. 117, p. 109466, 2020.
- [36] C. Fu, J. Wang, G. Li, M. Zhou, X. Wang, and Z. Liu, “Optimal bidding strategy for energy hub incorporating data center flexibility,” in *Proc. Power Syst. Green Energy Conf. (PSGEC)*, Aug 2021, pp. 139–144.
- [37] A. Pahlevan, M. Zapater, A. K. Coskun, and D. Atienza, “Ecogreen: Electricity cost optimization for green datacenters in emerging power markets,” *IEEE Trans. Sustain. Comput.*, vol. 6, no. 2, pp. 289–305, Apr. 2021.
- [38] Y. Zhang, D. C. Wilson, I. C. Paschalidis, and A. K. Coskun, “HPC data center participation in demand response: An adaptive policy with QoS assurance,” *IEEE Trans. Sustain. Comput.*, vol. 7, no. 1, pp. 157–171, Jan. 2022.
- [39] A. Tsiligkaridis, I. C. Paschalidis, and A. Coskun, “Data center demand response pricing using inverse optimization,” in *Proc. 10th Intl. Conf. Future Energy Syst. - ACM e-Energy '19*, Phoenix, AZ, USA, June 2019.
- [40] —, “Coordinated demand response by data centers using inverse optimization,” in *Proc. IEEE Intl. Conf. Comm., Control, Comput. Technologies Smart Grids (SmartGridComm 2020)*, Nov. 2020.
- [41] X. Wang, D. K. Chandrashekhara, M. Haque, I. Goiri, R. Bianchini, and T. Nguyen, “Grid-aware placement of datacenters and wind farms,” in *Proc. 6th Intl. Green Sustain. Comput. Conf. (IGSC)*, Dec 2015, pp. 1–8.
- [42] K. Kim, F. Yang, V. M. Zavala, and A. A. Chien, “Data centers as dispatchable loads to harness stranded power,” *IEEE Trans. Sustain. Energy*, vol. 8, no. 1, pp. 208–218, Jan. 2017.
- [43] H. Ma, D. W. Gao, B. Wang, and D. Liu, “Control strategy of UPS for data center based on economic dispatch,” in *Proc. IEEE/PES Transm. Distrib. Conf. Expo. (T&D)*, 2018, pp. 1–6.
- [44] W. Zhang, L. A. Roald, A. A. Chien, J. R. Birge, and V. M. Zavala, “Flexibility from networks of data centers: A market clearing formulation with virtual links,” *Electr. Power Syst. Res.*, vol. 189, 2020.
- [45] T. Niu, B. Hu, K. Xie, C. Pan, H. Jin, and C. Li, “Spatial coordination between data centers and power system considering uncertainties of both source and load sides,” *Int. J. Electr. Power Energy Syst.*, vol. 124, p. 106358, 2021.

- [46] M. S. Misaghian, G. Tardioli, A. G. Cabrera, I. Salerno, D. Flynn, and R. Kerrigan, "Assessment of carbon-aware flexibility measures from data centres using machine learning," *IEEE Trans. Ind. Appl.*, vol. 59, no. 1, pp. 70–80, Jan 2023.
- [47] Z. Wu, L. Chen, J. Wang, M. Zhou, G. Li, and Q. Xia, "Incentivizing the spatiotemporal flexibility of data centers toward power system coordination," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 3, pp. 1766–1778, May 2023.
- [48] W. Zhang, L. Roald, and V. Zavala, "Exploring the impacts of power grid signals on data center operations using a receding-horizon scheduling model," *IEEE Transactions on Sustainable Computing*, vol. 8, no. 2, pp. 245–256, 2023.
- [49] D. K. Molzahn *et al.*, "A survey of distributed optimization and control algorithms for electric power systems," *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 2941–2962, 2017.
- [50] C. M. Grinstead and L. Snell, *Introduction to Probability, 2nd edition*. American Mathematical Society, 2003.
- [51] I. C. Paschalidis, "Large deviations in high speed communication networks," Ph.D. dissertation, MIT, 1996.
- [52] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [53] F. Li and R. Bo, "Small test systems for power system economic studies," in *IEEE PES General Meeting*, Minneapolis, 2010.



Athanasios Tsiligkaridis received his B. Sc. (summa cum laude), M. Sc., and Ph.D. degrees in Electrical Engineering from Boston University in 2017, 2021, and 2023, respectively. He has had research internships at MITRE (2017) and Mitsubishi Electric Research Laboratories (2019, 2020, 2021). Dr. Tsiligkaridis is currently a Senior Research Scientist at Systems & Technology Research (STR) in Woburn, MA. His research interests include machine learning, artificial intelligence, optimization, algorithms, and

signal processing.



Panagiotis Andrianesis is a graduate of the Hellenic Army Academy, also holding a B.Sc. degree in economics (2004) from the National and Kapodistrian University of Athens, and a Diploma degree in electrical and computer engineering (2010) from the National Technical University of Athens, Greece. He received his M.Sc. degree in production management (2011) and his Ph.D. degree (2016) in the area of design and analysis of electricity market mechanisms from the University of Thessaly, Greece. He is currently an Associate Professor at the Center for Processes, Renewable Energy and Energy Systems (PERSEE), Mines Paris, PSL University, France. He has been a Research Associate Professor in the Division of Systems Engineering, at Boston University and a consultant and research associate with ECCO International Inc. His research interests include power system economics, electricity markets, operations research, optimization, and applied mathematics. Dr. Andrianesis is a Member of IEEE and INFORMS. He is the recipient of the 2010 IEEE APS Pre-Doctoral Research Award and the 2024 ERC Starting Grant.



Ayse K. Coskun received the PhD degree in computer science and engineering from UC San Diego. She is a full professor in Electrical and Computer Engineering Department, Boston University and the director of the Center for Information and Systems Engineering. Her research focuses broadly on design automation and computer systems, particularly focusing on energy efficiency, thermal challenges, and using analytics for intelligent system management. Coskun led multi-institutional projects, authored more than 120 technical papers, taught classes on computer systems and software, and delivered many invited talks and tutorials. Her research outcomes are widely recognized and culminated in several technical awards, including the IEEE CEDA Ernest Kuh Early Career Award and an IBM Faculty Award.



Michael Caramanis received the B.S. in Chemical Eng. from Stanford University, Palo Alto, CA, USA, in 1971 and the M.S. and Ph.D. degrees in Engineering from Harvard University, Cambridge, MA, USA, in 1972 and 1976, respectively. Since 1982 he is at Boston University where he is Professor of Systems and Mechanical Eng. He chaired the Greek Regulatory Authority for Energy and the International Energy Charter's Investment Group (2014-2008), was personally involved in power market implementations in England (1989-90) and Italy (2000-03), and his written work has influenced Power Market design in the U.S. and Europe. His current application domain focus is Marginal Costing and Dynamic Pricing on smart Power grids, grid topology control for congestion mitigation, and the extension of power markets to include distribution connected loads, generation, and resources. He is coauthor of "Spot Pricing of Electricity," Kluwer, 1987, and has more than 100 refereed publications. His disciplinary background is in Mathematical Economics, Optimization, and Stochastic Dynamic Decision Making.



Ioannis Ch. Paschalidis received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1991, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1993 and 1996, respectively. Since 1996, he has been with Boston University, Boston, MA, USA, where he is currently a Distinguished Professor of Engineering having appointments with the Department of Electrical and Computer Engineering, Division of Systems Engineering, Department of Biomedical Engineering, and Faculty of Computing and Data Sciences. He is the Director of the Rafik B. Hariri Institute for Computing and Computational Science & Engineering, Boston University. He has held visiting appointments with the MIT and Columbia University, New York, NY, USA. His current research interests include systems and control, networks, machine learning, optimization, computational biology, and medical informatics. Dr. Paschalidis was the recipient of the National Science Foundation CAREER Award in 2000, several Best Paper and Best Algorithmic Performance awards, and a 2014 IBM/IEEE Smarter Planet Challenge Award. He was an Invited Participant at the 2002 Frontiers of Engineering Symposium, organized by the U.S. National Academy of Engineering and the 2014 U.S. National Academies Keck Futures Initiative Conference. From 2013 to 2019, he was the Founding Editor-in-Chief for IEEE Transactions on Control of Network Systems.