

Agent Coordination via Contextual Regression (AgentCONCUR) for Data Center Flexibility

Vladimir Dvorkin, *Member, IEEE*

Abstract—A network of spatially distributed data centers can provide operational flexibility to power systems by shifting computing tasks among electrically remote locations. However, harnessing this flexibility in real-time through the standard optimization techniques is challenged by the need for sensitive operational datasets and substantial computational resources. To alleviate the data and computational requirements, this paper introduces a coordination mechanism based on contextual regression. This mechanism, abbreviated as AgentCONCUR, associates cost-optimal task shifts with public and trusted contextual data (e.g., real-time prices) and uses regression on this data as a coordination policy. Notably, regression-based coordination does not learn the optimal coordination actions from a labeled dataset. Instead, it exploits the optimization structure of the coordination problem to ensure feasible and cost-effective actions. A NYISO-based study reveals large coordination gains and the optimal features for the successful regression-based coordination.

Index Terms—Contextual learning, data centers, feature selection, regression, sustainable computing, system coordination

NOMENCLATURE

The main symbols used in this paper are stated below. Additional symbols are defined in the paper where needed.

A. Dimensions

- b Number of power system buses
- l Number of transmission lines
- m Number of data center users
- n Number of data centers forming a network
- k Number of virtual links connecting data centers
- q Length of a training dataset
- t Length of the day-ahead planning horizon

B. Optimization Parameters

- A Incidence matrix of the network of data centers
- c Vector of the 1st-order cost coefficients
- C Matrix of the 2nd-order cost coefficients
- c^{on} Vector of start-up cost coefficients
- d Vector of the electric power loads
- F Matrix of the power transfer distribution factors
- \bar{f} Vector of the maximum power transmission capacity
- G Matrix of distances between users and data centers
- \bar{p} Vector of the maximum generation capacities
- \underline{p} Vector of the minimum generation capacities

The author is with the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, MI, USA. e-mail: dvorkin@umich.edu. The project is supported in part by the Marie Skłodowska-Curie Actions COFUND Program, Grant Agreement No. 101034297 – project Learning ORDER.

- $\bar{p}^{\uparrow/\downarrow}$ Vectors of maximum ramp up/down capacity
- \bar{p}_{up} Vectors of maximum ramp up capacity at start up
- \bar{p}_{dw} Vectors of maximum ramp down capacity at shot down
- s Vector of the linear load shedding cost coefficients
- x Vector of contextual features
- α Latency loss parameters (% of the nominal latency)
- Γ Matrix converting computing loads to electric loads
- δ Vector of user computing demands
- ε Regularization parameter for feature selection
- ϱ Regularization parameter for task allocation
- ω Vector of renewable power injections

C. Optimization Variables

- e^{up} Vector of auxiliary variables to model start-up
- ℓ Vector of electric power load shedding
- p Vector of generation dispatch decisions
- u Vector of unit commitment decisions
- W Matrix allocating user demand among data centers
- β_0 Vector of regression intercept coefficient
- β° Matrix of regression coefficients associate w/ feature \circ
- β Vector stacking β_0 and vectorized matrices β°
- $\tilde{\delta}$ Vector of time-shifted user computing demands
- ϑ Vector of nominal computing loads of data centers
- $\tilde{\vartheta}$ Vector of time-shifted computing loads of data centers
- φ Vector of computing task shifts in a data center network

D. Notation

Lower- and upper-case letters denote vectors and matrices, respectively. For some matrix A , a_{ij} denotes its element at position (i, j) . Symbol \top stands for transposition, and \hat{x} denotes the optimal value of x . Vectors $\mathbf{0}$ and $\mathbf{1}$ are of zeros and ones, respectively. Operator $\langle \cdot \rangle_F$ is the Frobenius inner product, and $\|\cdot\|_p$ denotes the p -norm.

I. INTRODUCTION

COORDINATED operations of bulk power systems and coupled infrastructures allow for leveraging their complementarity and offsetting operational and economic inefficiencies, thus leading to enhanced performance. Coordination schemes have been proposed to synchronize grid operations with power distribution [1], natural gas [2], water [3], and district heating [4] systems, and more recently, a large coordination potential has emerged from the networks of spatially distributed data centers (NetDC) [5]. The unique coordination advantage of such networks is in *spatial flexibility*, which distributed data centers provide by shifting computing tasks among electrically remote locations. This flexibility resource will be important for future power grids, as electricity demand

of data centers is rapidly growing, and is expected to reach 35 GW by 2030 in the U.S. alone [6]. Even at present, the coordination potential is significant: training a single GPT-3 language model – the core of the popular ChatGPT chatbot – consumes as much as 1.3 GWh [7]. Allocating such energy-intensive tasks in the NetDC is thus likely to predetermine the dispatch cost and emission intensity in adjacent grids.

The growing environmental footprint of computing has encouraged large internet companies to optimize NetDC operations in a carbon-aware manner. Using online emission monitoring tools, such as WATTTIME.ORG and ELECTRICITYMAPS.COM, they smartly time and allocate computing tasks in regions with the least emission footprint [8], [9]. However, the sole reliance on limited emission data is the form of *grid-agnostic* coordination, which respects NetDC constraints yet ignores those of power grids. For *grid-aware* coordination, the literature offers three coordination mechanisms: demand response [5], enhanced market design [10]–[12], and co-optimization of grid and NetDC operations [13]. For the former, the work in [14] develops a comprehensive data center management protocol which accounts for a broader range of signals from the power grid, e.g., including not only carbon intensity but also electricity prices and loads; a closely related work also includes [15]–[18]. While demand response protocols follow power grid signals, the benefits from this type of coordination are still limited, as power grid and NetDC operations are still optimized separately. The second set of mechanisms integrate the flexibility of data centers within market-clearing algorithms. Although they feature robust market properties guaranteeing theoretical efficiency of system operations [19], [20], it remains challenging to fully represent complex data center objectives (e.g., quality of service) and constraints (e.g., latency) via single utility function. The third set of mechanisms propose to co-optimize the grid and NetDC operations within operational planning routines [13], with the full representation of operational constraints, akin to coordination models for conventional energy infrastructures [1]–[4]. While this approach suits the day-ahead planning, compensating for any deviation from the planning parameters in real time by resolving this optimization is challenging due to large data and computational requirements.

This paper develops a new, regression-based mechanism for grid-aware coordination of power systems and NetDC, termed as AgentCONCUR. The grid-awareness is achieved by embedding the decision-making of NetDC within the OPF optimization of a system operator, and jointly optimizing the objectives of the two systems within one optimization problem. However, unlike optimization-based coordination, AgentCONCUR shifts all computational burden to the planning stage, where a linear regression is trained to act on many operational scenarios. The coordination in the real time is then carried out by a trained regression, which solely acts on available contextual grid information, such as electricity prices. This way, AgentCONCUR resembles industry practices in [8] and [9] by relying on limited grid data, while also leveraging the optimization structure of power system's and NetDC's problems, thereby ensuring feasibility and economic efficiency of the Power–NetDC coordination. More specifi-

cally, this paper provides the following technical contributions:

- 1) To estimate economic potential of the Power-NetDC coordination under grid and communication constraints, the paper develops a bilevel co-optimization where power system decision-making is constrained by that of NetDC. Similar to grid-aware models in [12]–[20], this model takes the power system perspective, but it represents the NetDC via an embedded optimization problem with customer-oriented objectives and constraints, i.e., minimizing communication latency. This co-optimization provides the ideal coordination solution, suitable for the day-ahead operational planning of the two systems.
- 2) To enable fast and data-light implementation of such coordination in real-time, the paper devises a contextual regression policy that approximates the ideal coordination. The policy feasibility and cost-consistency is ensured by the new training optimization which inherits the bilevel optimization structure. Using sufficiently many operational scenarios in training allows for robust and cost-consistent performance across testing scenarios. Furthermore, the proposed training allows for the optimal coordination feature selection, such that the coordination can be made possible at different data requirements.

To demonstrate the cost-saving potential of the Power–NetDC coordination, we present a case study based on the system of the New York Independent System Operator (ISO). First, we demonstrate the performance of the day-ahead planning, where unit commitment and dispatch decisions are optimized by leveraging the space-time flexibility of NetDC. We then present the results of AgentCONCUR for peak-hour coordination by means of spatial computing task redistribution. The results reveal practical trade-offs between the amount of contextual information (features) and the efficiency of coordination. Interestingly, the case study found that the economic coordination between the two systems can be achieved by relying on a single contextual feature, i.e., the electric power load of the largest electricity demand center – New York City. However, there are more substantial economic gains that come with more data requirements.

The AgentCONCUR mechanism contributes to the area of optimization with contextual information [21], also referred to as smart predict-then-optimize [22] or decision-focused learning [23]. While most of the prior work in power systems has focused on decision-optimal *data* predictions, e.g., demand [21] or wind power [24], AgentCONCUR enables contextual predictions of the optimal coordination *decisions* instead. The mechanism also relates to learning-based grid coordination approaches, such as the one developed for transmission-distribution system coordination in [25]. However, the distinct features of the AgentCONCUR mechanism include its small computational requirements at the time of coordination, ability to preserve operational feasibility for coordinating agents, and the optimal coordination feature selection.

In the remainder, Section II details decision-making of power grid and NetDC operators, and then introduces the bilevel optimization problem for Power-NetDC coordination. Section III introduces the contextual regression approach for

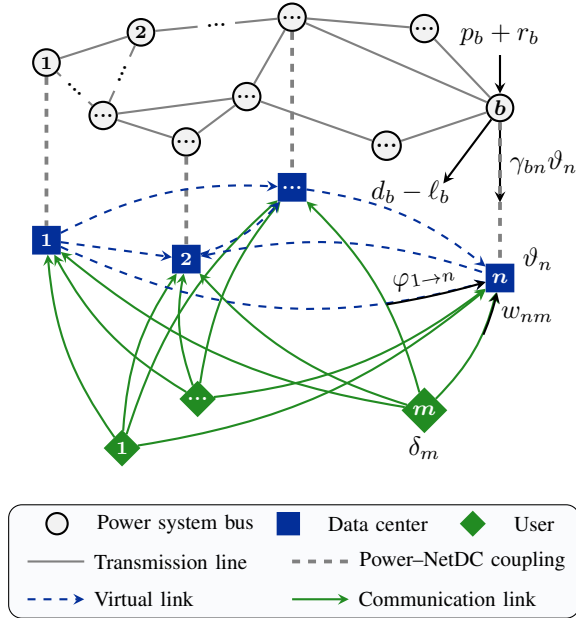


Fig. 1. Interfaces and notation of the power system, network of data centers (NetDC), and communication network between data centers and users.

AgentCONCUR. Section IV applies AgentCONCUR to New York ISO system and Section V concludes.

II. OPTIMIZING POWER-NETDC COORDINATION

We consider the power-NetDC coordination problem, in which agents interface as illustrated in Fig. 1. The NetDC operator selects the spatial allocation of computing tasks among data centers, receiving these tasks from a spatially distributed population of users. The main criterion for allocation is the minimization of network *latency* – a time delay between sending, executing, and returning the result of a computational task for every user. The resulting task allocation subsequently shapes electricity demand, which is then incorporated into the optimal power flow (OPF) problem for power system dispatch. Thus, these two problems can be simultaneously solved in a coordinated fashion to achieve minimal dispatch costs.

The coordination is performed by means of spatial shifts of computing tasks using *virtual links* connecting data centers into a network [12]. These shifts must be coordinated to satisfy both power system and NetDC objectives and constraints. To enable such coordination, we formulate the following bilevel optimization, where the power system operator acts as a leader, whose decision space is constrained by the NetDC operator, acting as a follower:

$$\begin{aligned}
 & \underset{p, \varphi}{\text{minimize}} && c_{\text{pwr}}(p) && \triangleright \text{Dispatch cost} \\
 & \text{subject to} && p \in \mathcal{P}_{\text{pwr}}(\vartheta) && \triangleright \text{Grid feasibility} \\
 & && \underset{\vartheta}{\text{minimize}} && c_{\text{net-dc}}(\vartheta) && \triangleright \text{Latency loss} \\
 & && \text{subject to} && \vartheta \in \mathcal{W}_{\text{net-dc}}(\varphi) && \triangleright \text{NetDC feasibility}
 \end{aligned}$$

Here, the task shift φ is the *coordination variable*, which coordinates the decision-making of the two operators. The lower-level problem takes request φ as input, and finds the

new task allocation ϑ among data centers that minimizes the latency loss. The new allocation must also be feasible for the set $\mathcal{W}_{\text{net-dc}}$ of NetDC operational constraints, parameterized by φ . The optimized allocation then enters the set \mathcal{P}_{pwr} of power system constraints, and the system operators compute the new generation dispatch p which minimizes the dispatch cost. The optimal solution φ^* thus achieves cost-optimal and feasible for the two systems coordination. Importantly, this hierarchical optimization structure separates the decision-making of the two operators, and coordinates them via a shared variable φ .

The rest of this section details decision-making of the two systems, and then presents the coordination problem in detail.

A. Power System Optimization

The power system planning is based on the unit commitment (UC) problem, which computes the least-cost commitment $u_t \in \{0, 1\}^b$ and dispatch $p_t \in [p, \bar{p}] \in \mathbb{R}_+^b$ decisions, that satisfies electricity net demand – load $d_t \in \mathbb{R}_+^b$ subtracted by non-dispatchable renewable generation $\omega_t \in \mathbb{R}_+^b$. The dispatch cost is modeled using a quadratic function with the first- and second-order coefficients $c \in \mathbb{R}_+^b$ and $C \in \mathbb{S}_+^b$, respectively. The power flows are computed using the matrix of power transfer distribution factors $F \in \mathbb{R}^{b \times l}$, and must respect the line capacity $\bar{f} \in \mathbb{R}_+^l$. In rare cases, when generation lacks to satisfy all loads, we model load shedding $\ell_t \in \mathbb{R}^b$ with the most expensive cost $s \gg c$. The UC problem takes the form:

$$\begin{aligned}
 & \underset{p_t, \ell_t, u_t, c_t^{\text{up}}}{\text{minimize}} && \sum_{t=1}^{\tau} (c^\top p_t + p_t^\top C p_t + s^\top \ell_t + \mathbb{1}^\top c_t^{\text{up}}) && (1a) \\
 & \text{subject to} && \mathbb{1}^\top (p_t + \omega_t + \ell_t - d_t - \Gamma \vartheta_t) = 0, && (1b) \\
 & && |F(p_t + \omega_t + \ell_t - d_t - \Gamma \vartheta_t)| \leq \bar{f}, && (1c) \\
 & && c_t^{\text{up}} \geq 0, \quad c_t^{\text{up}} \geq c_t^{\text{on}} \cdot (u_t - u_{t-1}), && (1d) \\
 & && \underline{p} \cdot u_t \leq p_t \leq \bar{p} \cdot u_t, && (1e) \\
 & && p_t - p_{t-1} \leq \bar{p}^\uparrow \cdot u_{t-1} \\
 & && \quad + \bar{p} \cdot (1 - u_t) + \bar{p}_{\text{up}} \cdot (u_t - u_{t-1}), && (1f) \\
 & && p_{t-1} - p_t \leq \bar{p}^\downarrow \cdot u_t \\
 & && \quad + \bar{p} \cdot (1 - u_{t-1}) + \bar{p}_{\text{dw}} \cdot (u_{t-1} - u_t), && (1g) \\
 & && 0 \leq \ell_t \leq d_t, \quad \forall t = 1, \dots, \tau, && (1h)
 \end{aligned}$$

which minimizes the total unit commitment and dispatch cost (1a) across τ hours subject to power balance condition (1b) and grid limits. Specifically, constraint (1c) ensures that the power flows remain within the maximum transmission capacity through the scheduling horizon. The two inequalities in (1d) model the cost of generation unit commitment using an auxiliary variable c_t^{up} . Generation constraints (1e)–(1g) use the binary logic to ensure that the generation units, when committed, produce within their production and ramping limits. Here, \bar{p}^\uparrow and \bar{p}^\downarrow are the maximum inter-hour ramping capacities, which can be different from ramping limits \bar{p}_{up} and \bar{p}_{dw} at a startup or shutdown of generating units, respectively.

Modelling Power-NetDC coordination, we distinguish between conventional loads d and power consumption by data centers $\Gamma \vartheta_t$ in constraints (1b) and (1c), where auxiliary matrix $\Gamma \in \mathbb{R}^{b \times n}$ converts computing loads $\vartheta_t \in \mathbb{R}^n$ of

n data centers into electrical loads. Although restrictive, this linear conversion model is consistent with power consumption models under different utilization regimes of data centers [26].

B. NetDC Optimization

The NetDC operator allocates computing tasks of m users among n data centers. For some computing demand $\delta_t \in \mathbb{R}^m$ at hour t , allocation $W_t \in \mathbb{R}^{n \times m}$ is optimized to satisfy conservation conditions $\vartheta_{ti} = \sum_{j=1}^m w_{tij}$ and $\delta_{tj} = \sum_{i=1}^n w_{tij}$, enforced for each data center i and user j , respectively. The goal is to minimize the latency, which is proportional to geodesic distance $G \in \mathbb{R}^{n \times m}$ between users and data centers [27]. The proxy function for aggregated latency is defined as

$$\mathcal{L} : \mathbb{R}^{n \times m} \mapsto \mathbb{R}, \quad \mathcal{L}(W) = \sum_{i=1}^n \sum_{j=1}^m g_{ij} w_{ij}. \quad (2)$$

The latency-optimal task allocation problem for each hour t then becomes:

$$\underset{W_t, \vartheta_t \geq 0}{\text{minimize}} \quad \mathcal{L}(W_t) + \frac{\rho}{2} \|W_t\|_2^2 \quad (3a)$$

$$\text{subject to} \quad W_t^\top \mathbf{1} = \delta_t, \quad (3b)$$

$$W_t \mathbf{1} = \vartheta_t, \quad (3c)$$

which minimizes latency subject to task conservation conditions (3b) and (3c). The objective function (3a) additionally includes a quadratic term that evenly allocates tasks among equally remote data centers, using a small parameter $\rho > 0$.

Although the resulting data center loading ϑ_t is latency-optimal, it is ignorant of the processes in the power system and may shape an expensive electricity demand allocation $\Gamma \vartheta_t$. In this scenario, the power system may leverage the space-time load shifting flexibility of data centers to reduce the cost of grid operations. To model NetDC flexibility, consider a vector of load shifts $\varphi \in \mathbb{R}^k$ along $k = \left(\frac{n(n-1)}{2}\right)\tau + n(\tau-1)$ virtual links, connecting n data centers in space and across τ time steps. Without loss of generality, we assume that tasks can be shifted to any other data center within a given time step or postponed to the adjacent time step at the same data center. To model the data center connectivity, consider a space-time incidence matrix $A \in \mathbb{R}^{n\tau \times k}$ of the NetDC graph, with entries

$$a_{ij} = \begin{cases} +1, & \text{if } i = n \\ -1, & \text{if } i = n' \end{cases} \quad \forall j = (n, n') \in 1, \dots, k, \quad (4)$$

so that the change in data center loading is computed as

$$A\varphi = \begin{bmatrix} \vartheta_1 \\ \dots \\ \vartheta_\tau \end{bmatrix} - \begin{bmatrix} \dot{\vartheta}_1 \\ \dots \\ \dot{\vartheta}_\tau \end{bmatrix}, \quad (5)$$

where $\tilde{\vartheta}_t$ is the vector of new computing loads at a particular hour t . Equation (5) preserves the conservation of computing loads of data centers across space and time. Advancing or postponing computing tasks of individual users results in a new task schedule $\tilde{\delta}_1, \dots, \tilde{\delta}_\tau \in \mathbb{R}^m$, and requires enforcing the following conservation condition:

$$\sum_{t=1}^{\tau} (\tilde{\delta}_t - \delta_t) = 0. \quad (6)$$

Aswering to the load shift request φ , the NetDC operator solves the following computing task re-allocation problem:

$$\underset{\tilde{\vartheta}_t, \tilde{\delta}_t, W_t \geq 0}{\text{minimize}} \quad \sum_{t=1}^{\tau} \left(\frac{1}{2} \|\mathcal{L}(W_t - \dot{W}_t)\|_2^2 + \frac{\rho}{2} \|\tilde{\delta}_t - \delta_t\|_2^2 \right) \quad (7a)$$

$$\text{subject to} \quad \mathcal{L}(W_t - \dot{W}_t) \leq \alpha \mathcal{L}(\dot{W}_t), \quad (7b)$$

$$W_t^\top \mathbf{1} = \tilde{\delta}_t, \quad W_t \mathbf{1} = \tilde{\vartheta}_t, \quad \forall t = 1, \dots, \tau, \quad (7c)$$

$$\text{equations (5) and (6)}, \quad (7d)$$

where the objective function (7a) minimizes the overall latency loss as the result of altering the latency-optimal allocation \dot{W}_t , plus the regularization term which penalizes the deviation of tasks from the original schedule using a small parameter $\rho > 0$. In case the time re-allocation of tasks incurs costs to users, ρ can also be used to assign those costs. While the objective steers the solution in a latency-aware fashion, inequality (7b) imposes the hard constraint on the latency loss, such that if the task shift request φ_t increases the latency beyond the $\alpha\%$ of the original latency, the request is considered infeasible. Inequality constraints in (7c) simply reinstate the conservation requirements for the new task schedule and data center loading, respectively.

C. Bilevel Optimization for Power and NetDC Coordination

Since the vector of task shifts affects the UC costs in (1) and the latency optimality loss in (7) simultaneously, φ is considered a coordination variable between power system and NetDC operators. We propose to optimize the coordination variable using the following bilevel optimization problem:

$$\begin{aligned} & \underset{p_t, \ell_t, u_t, c_t^{\text{up}}, \varphi}{\text{minimize}} \quad \sum_{t=1}^{\tau} (c^\top p_t + p_t^\top C p_t + s^\top \ell_t + \mathbf{1}^\top c_t^{\text{up}}) \\ & \text{subject to} \quad \mathbf{1}^\top (p_t + \omega_t + \ell_t - d_t - \Gamma \tilde{\vartheta}_t) = 0, \\ & \quad |F(p_t + \omega_t + \ell_t - d_t - \Gamma \tilde{\vartheta}_t)| \leq \bar{f}, \\ & \quad c_t^{\text{up}} \geq 0, \quad c_t^{\text{up}} \geq c_t^{\text{on}} \cdot (u_t - u_{t-1}), \\ & \quad \underline{p} \cdot u_t \leq p_t \leq \bar{p} \cdot u_t, \\ & \quad p_t - p_{t-1} \leq \bar{p}^\uparrow \cdot u_{t-1} \\ & \quad \quad + \bar{p} \cdot (1 - u_t) + \bar{p}_{\text{up}} \cdot (u_t - u_{t-1}), \\ & \quad p_{t-1} - p_t \leq \bar{p}^\downarrow, \\ & \quad \quad + \bar{p} \cdot (1 - u_{t-1}) + \bar{p}_{\text{dw}} \cdot (u_{t-1} - u_t), \\ & \quad 0 \leq \ell_t \leq d_t, \quad \forall t = 1, \dots, \tau, \end{aligned} \quad (\text{BL.U})$$

where the data center loading $\tilde{\vartheta}_1, \dots, \tilde{\vartheta}_\tau$ comes from:

$$\begin{aligned} & \underset{\tilde{\vartheta}_t, \tilde{\delta}_t, W_t \geq 0}{\text{minimize}} \quad \sum_{t=1}^{\tau} \left(\frac{1}{2} \|\mathcal{L}(W_t - \dot{W}_t)\|_2^2 + \frac{\rho}{2} \|\tilde{\delta}_t - \delta_t\|_2^2 \right) \\ & \text{subject to} \quad \mathcal{L}(W_t - \dot{W}_t) \leq \alpha \mathcal{L}(\dot{W}_t), \\ & \quad W_t^\top \mathbf{1} = \tilde{\delta}_t, \quad W_t \mathbf{1} = \tilde{\vartheta}_t, \\ & \quad \forall t = 1, \dots, \tau, \quad \text{equations (5) and (6)}. \end{aligned} \quad (\text{BL.L})$$

In this bilevel problem, the task shift request φ is the variable in (BL.U), that parameterizes the lower-level problem (BL.L), which, in turn, responds with the new data center loading. This way, the bilevel optimization identifies the cost-optimal and feasible for the two system tasks shifts. A common solution strategy for this problem is to replace (BL.L) with

its Karush–Kuhn–Tucker (KKT) conditions [28], yielding a mixed-integer problem reformulation (see Appendix A).

While solving problem (BL) requires collecting the sensitive parameters of the two systems into one optimization problem, the privacy risks for individual NetDC users are limited, as the NetDC operator submits aggregated computing demands without disclosing any individual computational request.

The bilevel problem (BL) can be solved within the day-ahead planning routines to coordinate the space-time task allocation that favors the UC costs while respecting the latency constraints of NetDC users. However, in case of deviations from the planning parameters (e.g., due to computing or electricity demand forecast errors), re-solving problem (BL) in real-time is challenging because of large data requirements, the possible lack of real-time coordination interfaces between the power grid and NetDC operators, and the computational burden of bilevel programming that may fail to provide the solution within narrow time frames. Motivated by these challenges, the next section introduces the AgentCONCUR framework for efficient coordination in real-time.

III. AGENT COORDINATION VIA CONTEXTUAL REGRESSION (AGENTCONCUR) FOR REAL TIME

To bypass the real-time coordination challenges, we adopt a *contextual* regression approach, which consists of two stages. At the first, offline stage, a regression policy is trained to learn the cost-optimal tasks shifts using contextual, easy-to-access features. They may include partial yet strongly correlated with grid conditions features, such as aggregated demand, electricity prices, network flows, and renewable power generation¹. After being trained on many representative scenarios of grid operations, e.g., renewable power generation scenarios, the regression policy *instantly* maps the unseen realization of features into task shifts at the second, operational stage.

We present the AgentCONCUR framework for peak-hour coordination, which is modeled as a single period problem, in which data centers only use the spatial flexibility, considering that computational tasks may arrive on demand and can not be delayed (e.g., online streaming), yet could be dispatched elsewhere at some latency loss. In this case, we can drop index t in presentation, and ignore the second term in (7a) and temporal task conservation constraint (6). Moreover, in real time, the commitment of generation units is known and the UC problem modeled in (BL.U) reduces to a DC-OPF problem [30]. Let x be a vector of contextual features, and let $\phi(x)$ be the affine regression policy of the form:

$$\phi(x) \triangleq \beta_0 + \beta_1 x \in \mathbb{R}^k,$$

where $\beta = (\beta_0, \beta_1)$ are regression parameters subject to optimization. Once they are optimized, for some feature vector \hat{x} , the coordination in real-time proceeds as follows:

$$\varphi = \begin{cases} \phi(\hat{x}), & \text{if feasible for NetDC and OPF} \\ \mathbf{0}, & \text{otherwise.} \end{cases} \quad (8)$$

¹Importantly, such contextual information is available online from many power system operators, including the NYISO [29] studied in Sec. IV.

That is, implement the regression solution if the task shifts are feasible for NetDC operations and also produce an OPF-feasible electricity load profile. Otherwise, proceed with a typically more expensive yet feasible non-coordinated solution.

Remark. The coordination rule in (8) is based on DC power flow equations. In practice, the shifts under policy $\phi(\hat{x})$ must also be verified to respect the AC power flow equations, to ensure voltage, apparent, active and reactive power feasibility. This can be achieved by solving a load flow problem, which projects the task shifts, and resulting bus injections, onto the AC power flow equations.

In the remainder, we first present the base regression training, used as a benchmark approach to optimizing policy $\phi(x)$. Then, we present the proposed training optimization problem at the core of AgentCONCUR.

A. Base Regression

The base approach to optimize policy $\phi(x)$ is two-fold:

1. Collect a training dataset $\{(x_i, \varphi_i)\}_{i=1}^q$ of q records, where each record i includes contextual features x_i and the optimal solution φ_i to the single-period problem (BL), specific to a particular record i .
2. Training the linear regression by solving

$$\underset{\|\beta\|_1 \leq \varepsilon}{\text{minimize}} \quad \frac{1}{q} \sum_{i=1}^q \|\beta_0 + \beta_1 x_i - \varphi_i\|_2^2 \quad (9)$$

which minimizes the regularized mean squared error over q historical records. Here, we chose L_1 -regularization, know as *Lasso* [31], which encourages sparsity of β up to selected regularization parameter $\varepsilon \in \mathbb{R}_+$. For any given value ε , optimization (9) selects optimal coordination features and minimizes the prediction error simultaneously.

While being a *data-only* approximation of the bilevel problem solution, this approach suffers from at least two drawbacks that may prevent its practical implementation. First, although it minimizes a prediction error, it may result in large decision errors in terms of OPF costs, e.g., when under- and over-predictions of task shifts have asymmetric cost implications. This may result in a large regret, i.e., the average distance between the OPF costs induced by trained policy ϕ and the OPF costs of the bilevel problem (BL). Second, optimization (9) is myopic to the upper- and lower-level feasible regions, thus risking violating operational limits of both power system and NetDC. These two observations motivate us to internalize the cost and feasibility criteria into model training.

B. Cost- and Feasibility-Aware Regression

Optimizing policy $\phi(x)$ for AgentCONCUR, we leverage the optimization structure of bilevel model (BL) to guarantee the least-cost and feasible regression-based coordination across available historical records. The proposed policy optimization for real-time coordination is based on the DC-OPF re-dispatch problem and takes the following form:

$$\underset{\mathcal{P} \cup \mathcal{D} \cup \varphi \cup \beta}{\text{minimize}} \quad \frac{1}{q} \sum_{i=1}^q \left((\hat{p}_i + r_i)^\top C (\hat{p}_i + r_i) \right)$$

$$+ c^\top (\dot{p}_i + r_i) + s^\top \ell_i \quad (10a)$$

$$\text{subject to } \varphi_i = \beta_0 + \beta_1 x_i, \|\beta\|_1 \leq \varepsilon, \quad (10b)$$

$$\mathbb{1}^\top (\dot{p}_i + r_i + \omega_i + \ell_i - d_i - \Gamma \tilde{\vartheta}_i) = 0, \quad (10c)$$

$$|F(\dot{p}_i + r_i + \omega_i + \ell_i - d_i - \Gamma \tilde{\vartheta}_i)| \leq \bar{f}, \quad (10d)$$

$$\underline{p} \leq \dot{p}_i + r_i \leq \bar{p}, |r_i| \leq \bar{r}, \mathbf{0} \leq \ell_i \leq d_i, \quad (10e)$$

$$\text{KKT conditions of single-period (BL.L),} \quad (10f)$$

$$\forall i = 1, \dots, q,$$

where set $\mathcal{P} = \{(r_i, \ell_i, W_i, \tilde{\vartheta}_i)\}_{i=1}^n$ includes DC-OPF and NetDC primal optimization variables, and set \mathcal{D} includes the dual variables of the single-period lower-level problem (BL.L). Relative to the base regression training in (9), this bilevel problem requires a larger training dataset, consisting of power system, NetDC, and contextual parameters, each specific to a particular scenario i , and minimizes the sample average OPF cost, subject to policy constraints in (10b), the set of upper-level OPF constraints in (10c)–(10e) and a set of KKT conditions of the lower-level problems (BL.L).

In problem (10), \dot{p}_i is a vector of fixed generator dispatch from the day-ahead planning optimization, and $r_i \in \mathbb{R}^b$ is the real-time re-dispatch of generation units, which is bounded by the re-dispatch limit $\bar{r} \in \mathbb{R}^b$. There are two scenarios in which the coordination policy assists with real-time re-dispatch. First, when the real-time conditions deviate from the planning parameters, e.g., due to stochastic renewable generation and loads, the trained coordination policy offsets the cost incurred by re-dispatching conventional generation. Second, if the day-ahead dispatch \dot{p}_i has been optimized in an uncoordinated fashion, the trained regression may still deliver coordination benefits in real time by re-dispatching conventional generation within the limit \bar{r} . Hence, the flexibility of NetDC is limited by \bar{r} , which provides important trade-offs in policy optimization. A larger \bar{r} allows for more flexibility and cost savings. However, larger \bar{r} also withholds more reserve capacity from generation units, thus exposing grid operations to more risks. It should thus be considered as a tuning parameter of Power-NetDC coordination.

Constraint (10b) couples many instances of bilevel coordination via regression policy and its role is two-fold: it structures task shifts and selects the optimal coordination features using the L_1 -regularization. This type of regularization also prevents overfitting and ensures robustness of the trained coordination policy in the sense of [32]. Furthermore, this regularization bounds the optimal solution, which is necessary when $\varphi_i = \beta_0 + \beta_1 x_i$ is a rank-deficient system of equations, i.e., having more features than virtual links. Without this regularization, a typical optimization solver returns the status “dual infeasible” or “unbounded”.

Similar to the base regression, the task shifts are restricted to the affine policy of contextual information. However, problem (10) also anticipates how the affine restriction affects the average OPF costs. Indeed, the choice of parameters β affects the task shift requests $\varphi_1, \dots, \varphi_q$, which then alter electricity load of data centers $\vartheta_1, \dots, \vartheta_q$ as they are coupled through the KKT conditions of the lower-level problem (BL.L). Thus, the optimal solution of problem (10) returns regression parameters

that are cost-optimal on average under the affine restriction.

Moreover, by solving problem (10), we also guarantee the feasibility of power system and NetDC operations across historical records. Indeed, $\beta = \mathbf{0}$ is always a feasible choice, which corresponds to the latency-optimal, non-coordinated solution from problem (3). Hence, in the worst-case, problem (10) chooses a non-coordinated solution to ensure feasibility for both systems. We can also reason about the feasibility of $\phi(x)$ for unseen operational scenarios in a probabilistic sense. Indeed, the theory of sample-based approximation of stochastic programs suggests that feasibility on unseen, out-of-sample scenarios improves as the sample size q increases [33], [34]. In the numerical experiments, we investigate the relationship between the sample size and the out-of-sample feasibility of the optimized coordination policy $\phi(x)$.

The training optimization (10) is solved at the operational planning stage using the similar mixed-integer reformation from Appendix A. Although the problem is NP-hard, modern optimization solvers, e.g., *Gurobi* [35], make the optimization more practical than its worst-case complexity would imply. Then, at the real-time stage, the trained regression model instantly maps contextual features into computing task shifts.

IV. NYISO CASE STUDIES

A. Data and Settings

We use an 11-zone aggregation of the New York ISO power system depicted in Fig. 2, sourcing data from [36]. This zonal layout corresponds to the granularity of the contextual data from the ISO website [29], which is used to train coordination policies. The power system includes approximately 30 GW of electricity demand, supplied by approximately 42 GW of conventional generation (oil, gas, and hydro) and by 1.7 GW of renewable generation (wind and solar). We install $n = 5$ data centers in the West, Central, North, NYC, and MillWd zones, serving customers in all $m = 11$ aggregation zones, as further shown in Fig. 2. For the day-ahead coordination, we consider that NetDC offers space-time flexibility during the evening hours from 4pm to 8pm. Computing loads can thus be shifted using $k = \binom{n(n-1)}{2} \tau + n(\tau - 1) = 70$ virtual links. For the peak-hour coordination by means of AgentCONCUR, we have $k = n(n - 1)/2 = 10$ virtual links. The task shifts outside the NY state area are not permitted. The computing demand δ_i is assumed to be proportional to the maximum peak load d_i in the i^{th} area, and will be scaled to achieve different NetDC penetration levels in the range from 5% to 30% of the peak system load. The operational data spans the period from January 1, 2018, to June 30, 2019, and includes 546 records, 24 hours each. All data, codes and default settings needed to replicate the results are available at

<https://github.com/wdvorkin/AgentCONCUR>

B. Day-Ahead Power-NetDC Coordination

The average cost-saving potential of the day-ahead coordination is illustrated in Fig. 3, for the 20% NetDC penetration level. Here, the left plane illustrates a \$137,812 cost-saving potential of the temporal NetDC flexibility, when latency loss

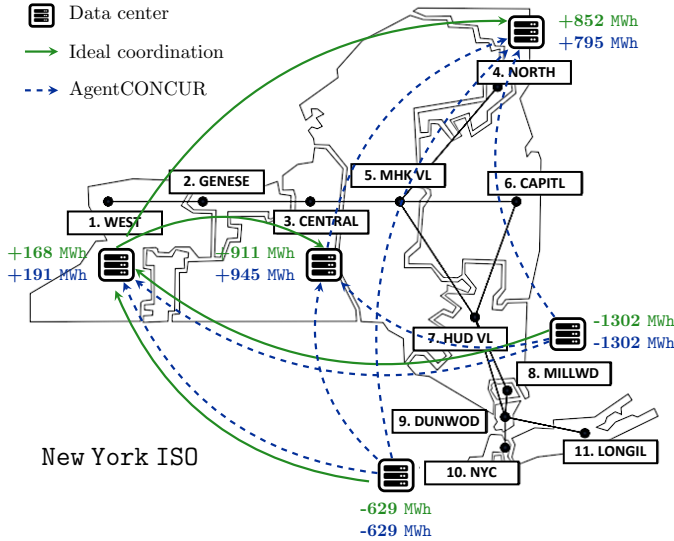


Fig. 2. 11-zone New York ISO system with 5 data centers. The arrows show active virtual links for peak-hour, real-time coordination under different coordination solutions for the 20% NetDC penetration level and 100% maximum latency loss. The change of NetDC electricity demand is given as the average across the test dataset.

requirement is set to zero, $\alpha \rightarrow 0$, thus restricting a spatial task re-allocation. The cost-saving potential then triples when invoking the spatial flexibility, by allowing the latency loss of at most 25% (middle plane). While with such a small latency loss the users are likely to observe no difference in the quality of service, it brings a substantial economic benefit to the power grid. Further increasing the allowable latency loss results in additional savings, yet the effect is diminishing with growing α , as shown in the right plane in Fig. 3. This is because the additional flexibility in NetDC would result in a steep net load profile, which would require more units to commit, hence increasing the overall cost of operation. Indeed, relaxing UC constraints, i.e., $u \in \{1,0\} \Rightarrow 0 \leq u \leq 1$, increases the cost-saving potential from \$445,059 to \$486,785 for the latency loss of $\alpha = 100\%$. This illustrates the importance of accounting for unit commitment constraints in harnessing the flexibility of data centers at the day-ahead stage.

Figure 4 illustrates how the baseline electricity consumption of data centers changes with coordination. Relative to the baseline, the coordinated solution shifts computing tasks from the South, with the largest electricity demand center in NYC, to the less-loaded data centers in the Central, Western, and Northern parts of the state. At the same time, tasks are shifted in time such that the data center loading in NYC and MillWd is minimal during the peak demand hour at 6pm.

If the NetDC operator directly participates in the day-ahead market, there are implicit financial incentives to provide the grid with NetDC flexibility. Table I reports the average cost for NetDC electricity consumption in a non-coordinated case and three coordination cases, revealing the opportunity to reduce the cost of electricity from 6% and 34.5%, depending on the maximum allowable latency optimality loss.

For each representative 24-hour horizon, we solve a bilevel optimization problem (BL) with 2,693 continuous and 936

TABLE I
AVERAGE CHARGE FOR NETDC ELECTRICITY DEMAND DURING THE PERIOD FROM 4 TO 20 PM, IN \$1,000. THE COORDINATED SOLUTION IS PROVIDED FOR VARYING LATENCY OPTIMALITY LOSS REQUIREMENT.

non-coordinated solution	coordinated solution		
	$\alpha = 0\%$	$\alpha = 25\%$	$\alpha = 100\%$
2,804	2,636 (-6.0%)	2,328 (-17.0%)	2,117 (-34.5%)




Fig. 3. UC costs throughout the 24-hour planning horizon for three latency loss bounds. The thick lines depict the average cost and the shaded area is the confidence band. The colored area is the period of Power-NetDC coordination. The cost resulting from coordination is in green, and the non-coordinated solution is in gray. The area between the curves is the cost-saving potential.

binary variables, and 5,611 constraints. The latency loss $\alpha = 100\%$ required longer solution time, taking 19.9s on average, with a standard deviation of 10.7s, on a standard laptop.

C. Power-NetDC Coordination via AgentCONCUR

In Sec. IV-C through Sec. IV-F, we consider the case of real-time peak-hour coordination by means of AgentCONCUR. We consider that the day-ahead schedule \hat{p} is optimized for the nominal latency solution \hat{v} , and the parameter \bar{r} is set to generation capacity. To train the policy, we use a dataset, where each record contains the following contextual features, which are readily available on the New York ISO website [29]:

- ▶ Zonal real-time electricity demand (d);
- ▶ Zonal electricity prices (λ);
- ▶ Total renewable generation, then disaggregated by zones using data on existing renewable installations (r);
- ▶ Power flows between aggregation zones (f).




Fig. 4. Average electricity consumption of five data centers from 4 pm to 8 pm. The horizontal lines depict the constant, baseline consumption without any coordination. The bars are hourly consumption under coordination.

Each record includes 45 contextual features, so that the coordination policy based on these features takes the form:

$$\phi \triangleq \beta_0 + \beta_1^d \begin{bmatrix} d_1 \\ \vdots \\ d_{11} \end{bmatrix} + \beta_1^\lambda \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_{11} \end{bmatrix} + \beta_1^r \begin{bmatrix} r_1 \\ \vdots \\ r_{11} \end{bmatrix} + \beta_1^f \begin{bmatrix} f_1 \\ \vdots \\ f_{12} \end{bmatrix}$$

To optimize and test the policy, we randomly select $q = 250$ records for training and reserve the remaining 296 records for testing, unless stated otherwise. The performance of the trained coordination policies is discussed using the unseen, test set. The remaining settings include default regularization parameters $\varrho = 10^{-5}$ and $\varepsilon = 10$.

The dispatch costs are compared in four cases:

- ▶ *No coordination*: NetDC electricity demand obeys the latency-optimal solution from problem (3);
- ▶ *Ideal coordination*: NetDC demand obeys the ideal coordination solution obtained from bilevel problem (BL);
- ▶ *Base regression*: NetDC demand is shifted according to the base regression policy optimized in (9);
- ▶ *AgentCONCUR*: NetDC demand is shifted according to the regression policy optimized in (10).

Our results reveal that the New York ISO system benefits from coordinating spatial tasks shifts in amount of ≈ 1.9 GWh from the densely populated South towards the Central, Northern, and Western parts of the state, as shown in Fig. 2. Noticeably, the ideal coordination consistently uses the same 4 out of 10 virtual links, while the AgentCONCUR coordination policy enjoys more active links. This difference is due to less flexible, affine policy structure, which results in more used links to ensure feasibility across the entire training dataset simultaneously, as opposed to per-scenario feasibility satisfaction provided by the ideal coordination.

Figure 5 illustrates the discrepancies in dispatch costs in all four cases. As the penetration of NetDC increases, the non-coordinated solution demonstrates rapid, quadratic growth of dispatch costs in the New York ISO dominated by conventional generation. On the other hand, the ideal coordination demonstrates a rather linear growth (e.g., see the bottom plot) of dispatch costs thanks to the cost-aware allocation of computing tasks. However, the extent of cost reduction significantly depends on the maximum allowable latency loss α , specified by the NetDC operator. For a small loss of 25%, users are likely to observe no difference in the quality of service. However, this enables savings of up to 24.5% of dispatch costs in the ideal coordination case, depending on the penetration level. The cost-saving potential increases to 49.0% and 56.7% in the case of double and tripled latency loss, respectively, when users experience more noticeable delays during peak-hour operations of the power system.

This cost-saving potential is exploited by both base regression and AgentCONCUR coordination policies. However, the base regression policy, which ignores power system and NetDC operational constraints, often results in substantively higher dispatch costs, which tend to stay closer to the non-coordinated solution than to the ideal one. On the other hand, the AgentCONCUR policy, which is aware of constraints of both systems, efficiently approximates the ideal solution,




Fig. 5. Average NYISO dispatch cost across the testing dataset under different coordination models for the varying NetDC penetration level and maximum allowable latency loss. The area between the dashed lines defines the cost-saving potential for regression-based coordination.

i.e., staying relatively close to the ideal solution in many cases depicted in Fig. 5. However, it tends to show a larger approximation gap with the allowable latency loss and NetDC penetration increase.

D. Feasibility of Regression-Based Coordination

The approximation gaps reported in Fig. 5 are due to infeasible task shifts, i.e., the shifts that violate power system constraints, NetDC constraints, or both. Whenever the task shift is infeasible in real-time, the two operators resort to a more expensive yet feasible non-coordinated solution. However, the feasibility of regression-based coordination improves with a larger size of the training dataset, as illustrated in Fig. 6. The AgentCONCUR policy dominates the base one and achieves zero violations of power system constraints (e.g., no load shedding) with sample size $q \geq 150$. Moreover, for $q \geq 150$, it keeps infeasibility of NetDC operations below 7%. The dominance of AgentCONCUR is *consistent*, which is important when the size of the training dataset with representative records is limited. We also observed similar results across other NetDC penetration and latency parameters.

E. CPU Times

While increasing the size of the training dataset improves feasibility, it also increases the computational burden of problem (10), as further shown in Fig. 7. For a reasonable choice of $q = 150$, the CPU times are approximately 8 hours. However,

NetDC penetration levels revealed 24.5–56.7% cost-saving potential, most of which has shown to be delivered by regression policies at different data-intensity preferences.

There are some notable limitations that motivate several directions for future work. For example, while the optimization-based coordination remunerates data center flexibility via duality theory [20], as such, the duality-based solution is unavailable under regression policies. It is thus relevant to study the integration of regression policies into real-time electricity markets. Furthermore, although the proposed mechanism does not require any private data exchange at the time of coordination, it still needs sensitive data from the power system and NetDC for training at the planning stage. One potential solution to remove this practical bottleneck is the use of differential private algorithms for data obfuscation [37], yet it will require additional modifications to the training procedure to eliminate the effect of noise.

ACKNOWLEDGEMENTS

This work was in part supported by the Marie Skłodowska-Curie Actions COFUND Postdoctoral Program, Grant Agreement No. 101034297 – project Learning ORDER.

APPENDIX

A. Mixed-Integer Reformulation of the Bilevel Problem

The Karush–Kuhn–Tucker conditions of the lower-level problem (BL.L) are derived from the following Lagrangian:

$$\begin{aligned} \max_{\mu} \min_{W, \vartheta, \tilde{\delta}} & \sum_{t=1}^{\tau} \frac{1}{2} \|\mathcal{L}(W_t) - \mathcal{L}(\dot{W}_t)\|_2^2 + \sum_{t=1}^{\tau} \frac{\varrho}{2} \|\tilde{\delta}_t - \delta_t\|_2^2 \\ & - \sum_{t=1}^{\tau} \mu_{\delta t}^{\top} (W_t^{\top} \mathbb{1} - \tilde{\delta}_t) - \sum_{t=1}^{\tau} \mu_{\vartheta t}^{\top} (W_t \mathbb{1} - \tilde{\vartheta}_t) \\ & - \sum_{t=1}^{\tau} \langle \mu_{wt}, W_t \rangle_{\text{F}} - \sum_{t=1}^{\tau} \mu_{\alpha t} (-\mathcal{L}(W_t - \dot{W}_t) + \alpha \mathcal{L}(\dot{W}_t)) \\ & - \mu_{\varphi}^{\top} \left(A\varphi - \begin{bmatrix} \tilde{\vartheta}_1 \\ \dots \\ \tilde{\vartheta}_{\tau} \end{bmatrix} + \begin{bmatrix} \dot{\vartheta}_1 \\ \dots \\ \dot{\vartheta}_{\tau} \end{bmatrix} \right) - \mu_{\tau} \left(\sum_{t=1}^{\tau} (\tilde{\delta}_t - \delta_t) \right) \end{aligned}$$

The stationarity conditions are as follows:

$$\tilde{\vartheta}_t: \mu_{\vartheta t} + \mu_{\varphi} = \mathbf{0}, \quad (11a)$$

$$\tilde{\delta}_t: \varrho(\tilde{\delta}_t - \delta_t) + \mu_{\delta t} - \mu_{\tau} = \mathbf{0}, \quad (11b)$$

$$\begin{aligned} w_{tij}: g_{ij}(\mathcal{L}(W_t - \dot{W}_t)) - \mu_{\delta t j} - \mu_{\vartheta t i} - \mu_{wtij} + \mu_{\alpha t} g_{ij} &= 0, \\ \forall t = 1, \dots, \tau, i = 1, \dots, n, j = 1, \dots, m. \end{aligned} \quad (11c)$$

The dual feasibility conditions are:

$$\mu_{wtij} \geq 0, \quad \forall i = 1, \dots, n, j = 1, \dots, m, \quad (11d)$$

$$\mu_{\alpha t} \geq 0, \quad \forall t = 1, \dots, \tau. \quad (11e)$$

The non-convex complementarity slackness conditions are addressed using an equivalent SOS1-based reformulation [38]:

$$\{\mu_{wtij}, w_{tij}\} \in \text{SOS1}, \quad \forall i = 1, \dots, n, j = 1, \dots, m, \quad (11f)$$

$$\{\mu_{\alpha t}, \mathcal{L}(W_t - \dot{W}_t) - \alpha \mathcal{L}(\dot{W}_t)\} \in \text{SOS1}, \quad \forall t = 1, \dots, \tau. \quad (11g)$$


where notation $\{x, y\} \in \text{SOS1}$ means that at most one variable may be nonzero. The equivalent reformulation of problem (BL) is then obtained when the lower-level problem (BL.L) is replaced with constraints (7b)–(7d) and (11).

REFERENCES

- [1] A. Papavasiliou and I. Mezghani, “Coordination schemes for the integration of transmission and distribution system operations,” in *Power Systems Computation Conference*, 2018, pp. 1–7.
- [2] A. Ratha *et al.*, “Affine policies for flexibility provision by natural gas networks to power systems,” *Electr. Power Syst. Res.*, vol. 189, 2020.
- [3] A. Stuhlmacher and J. L. Mathieu, “Water distribution networks as flexible loads: A chance-constrained programming approach,” *Electr. Power Syst. Res.*, vol. 188, p. 106570, 2020.
- [4] L. Mitridati and J. A. Taylor, “Power systems flexibility from district heating networks,” in *Power Systems Computation Conference*, 2018.
- [5] A. Wierman *et al.*, “Opportunities and challenges for data center demand response,” in *Internat. Green Computing Conference*, 2014, pp. 1–10.
- [6] McKinsey and Company. (2023) Investing in the rising data center economy. [Online]. Available: <http://tinyurl.com/mppbbh4h>
- [7] D. Patterson *et al.*, “Carbon emissions and large neural network training,” *arXiv preprint arXiv:2104.10350*, 2021.
- [8] A. James and D. Schien, “A low carbon kubernetes scheduler,” in *6th International Conference on ICT for Sustainability*, 2019.
- [9] A. Radovanović *et al.*, “Carbon-aware computing for datacenters,” *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1270–1280, 2022.
- [10] J. Ruddy and M. O’Malley, “Global shifting of data centers demand,” in *IEEE PES Innovative Smart Grid Technologies, Europe*, 2014, pp. 1–5.
- [11] G. Fridgen *et al.*, “Shifting load through space—The economics of spatial demand side management using distributed data centers,” *Energy Policy*, vol. 109, pp. 400–413, 2017.
- [12] W. Zhang *et al.*, “Flexibility from networks of data centers: A market clearing formulation with virtual links,” *Electr. Power Syst. Res.*, vol. 189, p. 106723, 2020.
- [13] K. Kim *et al.*, “Data centers as dispatchable loads to harness stranded power,” *IEEE Trans. Sustain. Energy*, vol. 8, no. 1, pp. 208–218, 2016.
- [14] 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 Trans. Sustain. Comput.*, vol. 8, no. 2, pp. 245–256, 2022.
- [15] D. Paul, W.-D. Zhong, and S. K. Bose, “Demand response in data centers through energy-efficient scheduling and simple incentivization,” *IEEE Systems Journal*, vol. 11, no. 2, pp. 613–624, 2015.
- [16] Q. Fang, J. Wang, H. Zhu, and Q. Gong, “Using model predictive control in data centers for dynamic server provisioning,” *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 9635–9642, 2014.
- [17] L. Parolini *et al.*, “A cyber-physical systems approach to data center modeling and control for energy efficiency,” *Proc. of IEEE*, vol. 100, no. 1, pp. 254–268, 2011.
- [18] H. Wang and Z. Ye, “Renewable energy-aware demand response for distributed data centers in smart grid,” in *2016 IEEE green energy and systems conference (IGSEC)*. IEEE, 2016, pp. 1–8.
- [19] L. Werner, A. Wierman, and S. H. Low, “Pricing flexibility of shiftable demand in electricity markets,” in *Proceedings of the 12th ACM International Conference on Future Energy Systems*, 2021, pp. 1–14.
- [20] W. Zhang and V. M. Zavala, “Remunerating space-time, load-shifting flexibility from data centers in electricity markets,” *Appl. Energy*, vol. 326, p. 119930, 2022.
- [21] M. Muñoz, S. Pineda, and J. Morales, “A bilevel framework for decision-making under uncertainty with contextual information,” *Omega*, vol. 108, p. 102575, 2022.
- [22] A. N. Elmachtoub and P. Grigas, “Smart “predict, then optimize,”” *Management Science*, vol. 68, no. 1, pp. 9–26, 2022.
- [23] J. Mandi *et al.*, “Decision-focused learning: Foundations, state of the art, benchmark and future opportunities,” *arXiv preprint arXiv:2307.13565*, 2023.
- [24] V. Dvorkin and F. Fioretto, “Price-aware deep learning for electricity markets,” *arXiv preprint arXiv:2308.01436*, 2023.
- [25] J. M. Morales, S. Pineda, and Y. Dvorkin, “Learning the price response of active distribution networks for TSO-DSO coordination,” *IEEE Trans. Power Syst.*, vol. 37, no. 4, pp. 2858–2868, 2021.
- [26] A. Radovanovic *et al.*, “Power modeling for effective datacenter planning and compute management,” *IEEE Trans. Smart Grid*, vol. 13, no. 2, pp. 1611–1621, 2021.
- [27] I. Narayanan, A. Kansal, and A. Sivasubramaniam, “Right-sizing geodistributed data centers for availability and latency,” in *37th International Conference on Distributed Computing Systems*, 2017, pp. 230–240.

- [28] D. Pozo, E. Sauma, and J. Contreras, “Basic theoretical foundations and insights on bilevel models and their applications to power systems,” *Ann. Oper. Res.*, vol. 254, pp. 303–334, 2017.
- [29] NYISO. (2023) Energy market & operational data. [Online]. Available: <https://www.nyiso.com/energy-market-operational-data>
- [30] S. Chatzivasileiadis, “Lecture notes on optimal power flow (OPF),” *arXiv preprint arXiv:1811.00943*, 2018.
- [31] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. R. Stat. Soc., B: Stat.*, vol. 58, no. 1, pp. 267–288, 1996.
- [32] D. Bertsimas and M. S. Copenhaver, “Characterization of the equivalence of robustification and regularization in linear and matrix regression,” *Eur. J. Oper. Res.*, vol. 270, no. 3, pp. 931–942, 2018.
- [33] A. Nemirovski and A. Shapiro, “Convex approximations of chance constrained programs,” *SIAM Journal on Optimization*, vol. 17, no. 4, pp. 969–996, 2007.
- [34] K. Margellos, P. Goulart, and J. Lygeros, “On the road between robust optimization and the scenario approach for chance constrained optimization problems,” *IEEE Trans. Automat. Contr.*, vol. 59, no. 8, pp. 2258–2263, 2014.
- [35] Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2023. [Online]. Available: <https://www.gurobi.com>
- [36] H. A. U. Khan, J. Kim, and Y. Dvorkin, “Risk-informed participation in T&D markets,” *Electr. Power Syst. Res.*, vol. 202, p. 107624, 2022.
- [37] V. Dvorkin and A. Botterud, “Differentially private algorithms for synthetic power system datasets,” *IEEE Control Systems Letters*, 2023.
- [38] S. Siddiqui and S. A. Gabriel, “An SOS1-based approach for solving MPECs with a natural gas market application,” *Netw. Spat. Econ.*, vol. 13, pp. 205–227, 2013.



Vladimir Dvorkin Jr. (S'18, M'21) is an assistant professor in the Electrical Engineering and Computer Science Department at the University of Michigan, Ann Arbor. He leads the ptiML research group, which develops optimization and machine learning algorithms to support the transition of power grids to digital and carbon-free systems. Before joining Michigan, he was a postdoctoral fellow at the Massachusetts Institute of Technology affiliated with the Energy Initiative and Laboratory for Information & Decision Systems. He earned a PhD degree in electrical engineering at the Technical University of Denmark in 2021, and he was visiting the School of Industrial and Systems Engineering of Georgia Tech in 2019. His work is recognized by several awards, including the Marie Skłodowska-Curie Actions & Iberdrola postdoctoral fellowship and the IEEE Transactions on Power Systems Best Paper Award.