

Carbon-Aware Computing in a Network of Data Centers: A Hierarchical Game-Theoretic Approach

Enno Breukelman¹, Sophie Hall², Giuseppe Belgioioso², and Florian Dörfler²

Abstract—Over the past decade, the continuous surge in cloud computing demand has intensified data center workloads, leading to significant carbon emissions and driving the need for improving their efficiency and sustainability. This paper focuses on the optimal allocation problem of batch compute loads with temporal and spatial flexibility across a global network of data centers. We propose a bilevel game-theoretic solution approach that captures the inherent hierarchical relationship between supervisory control objectives, such as carbon reduction and peak shaving, and operational objectives, such as priority-aware scheduling. Numerical simulations with real carbon intensity data demonstrate that the proposed approach successfully reduces carbon emissions while simultaneously ensuring operational reliability and priority-aware scheduling.

I. INTRODUCTION

Between 2010 and 2020, the global compute load on data centers (DCs) has increased more than 9 times [1]. As a consequence, the number of hyper-scale DCs has doubled between 2015 and 2021, reaching 700 installed facilities worldwide [2]. For their operation, DCs require a significant supply of electricity from the grid, about 1-1.5% of the global electricity demand in 2022, corresponding to 230-340 TWh [3]. Unfortunately, the electricity they run on is still coming from predominantly carbon-intensive sources. Incorporating more renewable energy into the electricity mix is challenging due to supply variability, which depends on both time of day and geographical location. Remarkably, a large share of the global compute load is not time-sensitive, allowing a delayed execution, nor bound to a specific DC, allowing it to be executed at a different DC. Therefore, various compute-demand management and real-time routing mechanisms exploit this flexibility to mitigate the carbon impact and provide ancillary services to the power grid.

Most of the existing work focuses on real-time routing of temporally inflexible compute jobs [4], such as on-demand services and search engine requests. Furthermore, loads with either spatial or temporal flexibility [5], [6] are considered, but rarely the combination of both. Compute jobs are also mostly considered part of an aggregate compute load, which does not differentiate between individual jobs

[7]. Typical objectives considered by these allocation mechanisms include reducing monetary expenses, carbon taxes [8], [4], carbon emissions [9], and inducing peak shaving [5], [10]. Additional modeling features investigated in existing research include: (i) the interaction with electricity utilities and generators, [11], [5], (ii) more detailed power modeling of the DC facilities [12], and (iii) the geographical location and load migration among a network of DCs [13].

Within this large body of literature, the carbon-aware computing platform proposed in [14] stands out as it is currently implemented by Google to operate its inter-continental fleet of DCs. Google uses carbon intensity forecasts, and predicts future compute demand for their DCs. This information is used to generate so-called Virtual Capacity Curves (VCCs), which limit the hourly resource of their DCs in the fleet such that temporally flexible compute load is pushed to less carbon-intensive hours. Thus, by harnessing the compute load predictions, they can minimize carbon emissions.

In this paper, we design a novel day-ahead scheduling mechanism to allocate batch compute jobs over an interconnected fleet of DCs, inspired by Google's carbon-intelligent platform in [14]. The approach in [14] is scheduler-agnostic, meaning VCCs are computed independently of the actual compute job schedule. In contrast, we co-design the VCCs and compute job allocation using a hierarchical game-theoretic approach that distinguishes between separate individual compute jobs. Additionally, our mechanism considers not only temporal shifting but also spatial migration.

The contributions of this paper are threefold. Firstly, we formalize the optimal allocation problem of compute jobs with spatial and temporal flexibility, over a network of DCs as a single-leader multiple-follower Stackelberg game. At the lower level, the owners of compute jobs compete for the computational resources in the DC network to process their jobs as soon as possible. At the upper level, the DC operator generates the virtual capacities to induce a competitive allocation that reduces carbon emissions and induces peak shaving. Thirdly, we derive an efficient ad-hoc algorithm to solve the resulting large-scale Stackelberg game by adapting the method proposed in [15], which bears similarities to the approach of solving a Stackelberg game in [16]. Finally, we numerically validate the proposed allocation mechanism using real carbon intensity data for a selection of Google's DC locations.

Our numerical findings show that temporally shifting and spatially migrating flexible compute jobs significantly reduce carbon emissions. In contrast to scheduler-agnostic approaches, co-designing the VCCs can improve the homo-

¹Enno Breukelman is with the KTH Royal Institute of Technology, School of Electrical Engineering and Computer Science, Division of Decision and Control Systems, Malvinas väg 10, SE-100 44 Stockholm, Sweden. cebre@kth.se

²Sophie Hall, Giuseppe Belgioioso, and Florian Dörfler are with the Automatic Control Laboratory, Department of Electrical Engineering and Information Technology, ETH Zurich, Physikstrasse 3, 8092 Zürich, Switzerland. {gbelgioioso, shall, dorfler}@ethz.ch

This work is supported by the SNSF via NCCR Automation (Grant Number 180545).

generality of waiting times in the job schedule.

Notation: \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\mathbb{R}_{> 0}$ denote the set of real, nonnegative real, and positive real numbers, respectively. Given N scalars a_1, \dots, a_N , $\text{diag}(a_1, \dots, a_N)$ denotes the diagonal matrix with a_1, \dots, a_N on the main diagonal. Given N column vectors $x_1, \dots, x_N \in \mathbb{R}^n$, $x = \text{col}(x_1, \dots, x_N) = [x_1^\top, \dots, x_N^\top]^\top$ denotes their vertical concatenation. The Euclidean projection onto the set \mathcal{X} is denoted as $\mathbb{P}_{\mathcal{X}}[\cdot]$. The partial derivative of a function f with respect to its i -th argument is denoted as: $\nabla_i f(x_1, \dots, x_n) = \frac{\partial f(x_1, \dots, x_n)}{\partial x_i}$.

II. PROBLEM STATEMENT

We consider a set $\mathcal{I} := \{1, \dots, I\}$ of batch compute jobs to allocate across a fleet of DCs $\mathcal{D} := \{1, \dots, D\}$, over a planning horizon $\mathcal{T} := \{1, \dots, T\}$. Each batch job $i \in \mathcal{I}$ is uploaded by a customer, or team, at an initial location $d^i \in \mathcal{D}$, at which the data required for the computation is assumed to be physically stored. Each batch job is characterized by a predicted compute volume v^i and a priority parameter τ^i , describing its time urgency. The DCs are physically interconnected via a fiber network that allows compute jobs to be transferred across the fleet. This network is modeled via a weighted graph $\mathcal{G} := \mathcal{G}(\mathcal{D}, \mathcal{E})$, where the vertices are the DCs, and the edges are the fiber connections $e_{k,\ell} \in \mathcal{E}$ directly connecting neighboring DCs $k, \ell \in \mathcal{D}$. The weights of the edges encode the price of migrating the data, established by the internet service providers for using their routers.

A. Team's allocation game

Each team shall choose the allocation of their batch job $i \in \mathcal{I}$, represented by the vector $y^i = \text{col}(y_1^i, \dots, y_D^i)$, whose entries $y_{d,t}^i \in \mathbb{R}_{\geq 0}$ describe the share of the compute volume v^i allocated to DC $d \in \mathcal{D}$ at time $t \in \mathcal{T}$. In other words, y^i determines where and when to reserve DC capacities to execute compute job i . The allocation chosen by each team must satisfy various operational constraints. Intuitively, it must cover the total compute volume v^i , yielding

$$\sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} y_{d,t}^i = v^i, \quad \forall i \in \mathcal{I}. \quad (1)$$

A job $i \in \mathcal{I}$ may be transferred from its initial DC d^i , to any alternative DC $j \in \mathcal{D}$ via a predetermined shortest path $\omega_{i,j}$ over the fiber-optic network \mathcal{G} . This path is uniquely identified and computed in advance, to ensure minimal transit time and resource usage, and consists of a sequence of fiber connections $e \in \mathcal{E}$. In Fig. 1, we illustrate an example of the considered setup, where the path $\omega_{1,3} = (e_{1,2}, e_{2,3})$ uses the fiber connections $e_{1,2}$ and $e_{2,3}$.

Any migrated share of v^i appears in a migration variable $z^i = \text{col}(z_1^i, \dots, z_D^i)$, whose entries $z_{j,t}^i \in \mathbb{R}_{\geq 0}$ describe how much predicted load is migrated from d^i to DC $j \in \mathcal{D}$ at time step $t \in \mathcal{T}$, over the path $\omega_{i,j}$. The entry $z_{i,t}^i$ describes the queue of unprocessed compute volume at this DC. Since there is no migration before the first time step, there are $T-1$ migration steps. The share of a predicted compute volume v^i

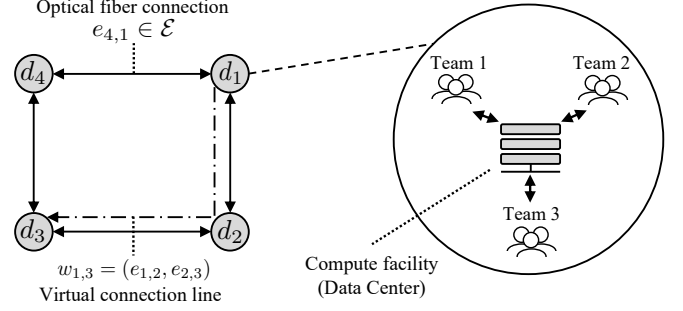


Fig. 1. DC network featuring 4 DC locations interconnected by physical connection lines (solid line). A job migration between DC 1 and DC 3 is realized via a path through the fiber network. Right: Teams are associated with a DC location, where they initially submit their compute jobs.

executed at a DC other than the initial d^i , i.e., $d \in \mathcal{D} \setminus \{d^i\}$, must be migrated in the time step before

$$y_{d,t+1}^i = z_{d,t}^i, \quad \forall i \in \mathcal{I}, \forall d \in \mathcal{D} \setminus \{d^i\}, \forall t \in \mathcal{T} \setminus \{T\}. \quad (2)$$

Any share of a compute job v^i , that has not yet been computed at the initial DC d^i or migrated to another DC $z_{j \neq i, \ell}^i$ is stored in the queue for the next time step $z_{i,t}^i$

$$v^i - \sum_{\ell=1}^t \left(y_{d^i, \ell}^i + \sum_{j=1, j \neq i}^D z_{j, \ell}^i \right) = z_{i,t}^i, \quad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \setminus \{T\}. \quad (3)$$

Finally, the total amount of compute job volume allocable at each DC $d \in \mathcal{D}$ is limited by

$$\sum_{i \in \mathcal{I}} y_{d,t}^i \leq x_{d,t}, \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}, \quad (4)$$

where $x_{d,t}$ is a virtual capacity, set by the DC operator. These virtual capacities, considered over the complete planning horizon and all DCs in the network, constitute the VCCs.

Let $\mathbf{y}^i = \text{col}(y^i, z^i)$ denote the stacked vector of the local decision variables (namely, allocation and migration) of team i . All local operational constraints (1)–(4) can be compactly represented via the set-valued mapping

$$\mathcal{Y}^i(x, \mathbf{y}^{-i}) := \{(y^i, z^i) \mid (1) - (4) \text{ hold}\}, \quad (5)$$

which depends on the VCCs $x = \text{col}(x_1, \dots, x_D)$ viewed as parameters, as well as the allocations of the other compute jobs $\mathbf{y}^{-i} = \text{col}(\mathbf{y}^1, \dots, \mathbf{y}^{i-1}, \mathbf{y}^{i+1}, \dots, \mathbf{y}^I)$.

The objective of each team is to choose an allocation that minimizes execution time, migration cost, and deviation from a pre-determined allocation profile \hat{y}^i , formulated as

$$J^i(\mathbf{y}^i) = \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} \left(\tau_t^i y_{d,t}^i + z_{d,t}^i \left(\sum_{e \in \omega_{i,d}} \sigma_e^i \right) \right) + \frac{1}{2} \epsilon \|\mathbf{y}^i - \hat{\mathbf{y}}^i\|^2. \quad (6)$$

The first term in (6) penalizes allocating compute jobs incrementally with the time steps and depends on the priority parameter $\tau_t^i = \tau^i l_t$, where $l_t = t/T$ is a time-dependent

weight that penalizes delayed allocations. The second term in (6) penalizes the migration of compute jobs, as transferring data over the network comes with a time penalty for each team. Therein, the total price of migration is the sum of all the prices at each fiber connection $\sigma_e^i = \tau^i \sigma_e$, multiplied by the priority parameter to further penalize migrations of urgent compute jobs. The quadratic terms penalize the deviation from a predefined allocation and migration $\hat{\mathbf{y}}^i = \text{col}(\hat{y}^i, \hat{z}^i)$.

Overall, each team solves the following optimization problem to find their optimal allocation

$$(\forall i \in \mathcal{I}) : \min_{\mathbf{y}^i} J^i(\mathbf{y}^i) \quad \text{s.t.} \quad \mathbf{y}^i \in \mathcal{Y}^i(x, \mathbf{y}^{-i}). \quad (7)$$

The collection of these inter-dependent optimization problems constitutes a generalized game [17], parametric in x , the VCCs. Note that the other teams' allocations \mathbf{y}^{-i} enter the local constraints in (7) due to the resource constraint (4), rendering the game a generalized game. Furthermore, with the teams' objectives in (6) being decoupled, this game structure corresponds to an exact potential game [18].

A meaningful solution concept for (7) is the Generalized Nash Equilibrium (GNE), which is a set of allocations $\bar{\mathbf{y}} = (\bar{\mathbf{y}}^1, \dots, \bar{\mathbf{y}}^I)$ that simultaneously solve the optimization problems in (7). Here, we focus on the special subclass of variational GNEs (v-GNEs) due to their computational tractability and economic fairness [17]. This subclass of equilibria corresponds to the solution set of the parameterised variational inequality $\text{VI}(F, \mathcal{Y}(x))$, namely, the problem of finding a vector $\bar{\mathbf{y}}$ such that

$$\langle F(\bar{\mathbf{y}}), (\mathbf{y} - \bar{\mathbf{y}}) \rangle \geq 0, \quad \forall \mathbf{y} \in \mathcal{Y}(x), \quad (8)$$

where, $F(\mathbf{y}) := \text{col}(\{\nabla J_i(\mathbf{y}^i)\}_{i \in \mathcal{I}})$ is the *pseudo-gradient* mapping of the game (7), and $\mathcal{Y}(x)$ collects the operational constraints (5) and depends explicitly on the VCCs x . We denote by $\mathbf{y}_*(\cdot)$ the parameter-to-solution mapping that, given the VCCs x , returns the set of solutions $\mathbf{y}_*(x)$ to the VI in (8), namely, a set of strategically-stable allocations. It can be shown that $\mathbf{y}_*(x)$ is single-valued for any x that makes (8) feasible, meaning that the allocation game has a unique solution for given VCCs, as the pseudo-gradient mapping F is strongly monotone.

B. Supervisory Objectives

The DC operator can influence the outcome of the allocation game (7) by manipulating the virtual capacities $x_{d,t} \in \mathbb{R}_{\geq 0}$, limiting the total load at each DC and in each time step. As for the allocations, also the VCCs must satisfy a series of operational constraints. Firstly, the total available capacity defined by the VCCs must accommodate the total compute demand, yielding

$$\sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} x_{d,t} \geq \sum_{i \in \mathcal{I}} v^i. \quad (9)$$

Additionally, the VCCs cannot exceed the physical computational capacity of the DCs, yielding

$$0 \leq x_{d,t} \leq x_{d,t}^{\max}, \quad \forall d \in \mathcal{D}, t \in \mathcal{T}, \quad (10)$$




Fig. 2. An example of sequential allocation of two compute jobs y^1 and y^2 on a data center $d \in \mathcal{D}$, where job 1 is of higher priority than job 2. The virtual capacity curve (red dashed line) limits the allocable load at each time slot. The maximum capacity of the DC $x_{d,t}^{\max}$ is obtained by subtracting the inflexible load from the maximum capacity of that DC x_d^{\max} (dashed black line).

where the maximum capacity is denoted by $x_{d,t}^{\max}$. This value is obtained by subtracting the inflexible demand, which we cannot shift or migrate but assume to be forecast perfectly, from the physical capacity x_d^{\max} . Fig. 2 displays an exemplary job schedule for two jobs y^1 and y^2 , with the VCC $x_{d,t}$ as their upper bound and physical constraint $x_{d,t}^{\max}$. The operational constraints on the DC operator's decision are compactly represented via

$$\mathcal{X} := \{x \mid (9), (10) \text{ hold}\}. \quad (11)$$

The DC operator utilizes its influence on the teams' allocations to reduce the carbon impact and peak usage of the DCs. This goal is modeled by the multi-objective function

$$\phi(x, \mathbf{y}) = \phi^{\text{carb}}(\mathbf{y}) + \phi^{\text{peak}}(\mathbf{y}) + \xi \phi^{\text{migr}}(\mathbf{y}). \quad (12)$$

The term ϕ^{carb} penalizes carbon emissions and is defined as

$$\phi^{\text{carb}}(\mathbf{y}) = \sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} \rho_{d,t}^{\text{carb}} \left(\sum_{i \in \mathcal{I}} y_{d,t}^i \right), \quad (13)$$

where $\rho_{d,t}^{\text{carb}}$ is the carbon intensity of the supplied power at DC d and time step t , modeling the carbon impact of utilizing computing power in a linear relationship, as it is done in [7], [19]. The second term ϕ^{peak} penalizes peak usage of DCs computational resources per time step and is defined as

$$\phi^{\text{peak}}(\mathbf{y}) = \sum_{d \in \mathcal{D}} \left(\sum_{t \in \mathcal{T}} \left(\sum_{i \in \mathcal{I}} y_{d,t}^i \right)^p \right)^{1/p}, \quad (14)$$

for some $p \in \mathbb{N}$ large enough to approximate the infinity norm. The third term penalizes the aggregate migration of compute jobs and is defined as

$$\phi^{\text{migr}}(\mathbf{y}) = \sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} z_{d,t}^i \left(\sum_{e \in \omega_{i,d}} \sigma_e^i \right). \quad (15)$$

This last term is necessary to reduce network traffic from the DC operator's perspective, as the DC operator is responsible for paying network routing fees to internet service providers. Finally, the parameter $\xi \geq 0$ in (12) regulates the impact of migration on the final allocation. In contrast to an exclusively

monetary objective, the multi-objective in (12) addresses sustainability aspects directly.

Overall, the DC operator aims to solve the following single-leader multiple-follower Stackelberg game

$$\min_{x, \mathbf{y}} \phi(x, \mathbf{y}) \quad (16a)$$

$$\text{s.t. } x \in \mathcal{X} \quad (16b)$$

$$\langle F(\mathbf{y}), (\mathbf{y}' - \mathbf{y}) \rangle \geq 0, \quad \forall \mathbf{y}' \in \mathcal{Y}(x). \quad (16c)$$

III. BILEVEL GAME SOLUTION APPROACH

We solve (16) using a customized version of a first-order algorithm named BIG Hype, recently proposed in [15].

A. Hyper-gradient based algorithm

The first-order algorithm in [15] is based on the idea of substituting the equilibrium constraints (16c) into the objective function (16a), by exploiting the solution mapping $\mathbf{y}_*(\cdot)$ and the fact that solutions to (8) are unique. The resulting non-convex non-smooth optimization problem reads as

$$\begin{aligned} \min_x \quad & \phi(x, \mathbf{y}_*(x)) =: \phi_e(x) \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned} \quad (17)$$

Then, a local solution to (17) is obtained by relying on projected ‘‘gradient’’ descent. Whenever $\mathbf{y}_*(x)$ is differentiable at x , one can obtain a gradient by applying the chain rule¹

$$\nabla \phi_e(x) = \nabla_1 \phi(x, \mathbf{y}_*(x)) + \mathbf{J}\mathbf{y}_*(x)^\top \nabla_2 \phi(x, \mathbf{y}_*(x)), \quad (18)$$

where $\mathbf{J}\mathbf{y}_*(x)$ is the Jacobian of the solution mapping $\mathbf{y}_*(\cdot)$ at x , commonly known as the sensitivity.

The proposed algorithm is summarized in Algorithm 1 and consists of three main steps. In step **1**, the DC operator uses the current estimates of the equilibrium and sensitivity to compute a hypergradient (18), and runs a projected hypergradient descent step. Based on the new VCCs x^{k+1} , in step **2**, the teams compute the resulting optimal allocation $\mathbf{y}_*(x^{k+1})$ and subsequently, in step **3**, its sensitivity $\mathbf{J}\mathbf{y}_*(x)$.

Algorithm 1: Customized BIG Hype

Initialize $x^0 \in \mathcal{X}$, $\mathbf{y}^0 \in \mathbb{R}_{\geq 0}^{n_y}$, $s^0 = \mathbf{0} \in \mathbb{R}^{n_y \times n_x}$,

$k = 1$, and $\{\alpha^k\}_{k \in \mathbb{N}}$

repeat until convergence

1. DC operator’s projected hypergradient step:

$$\begin{aligned} \nabla \phi_e^k &\leftarrow \nabla_1 \phi(x^k, \mathbf{y}^k) + (s^k)^\top \nabla_2 \phi(x^k, \mathbf{y}^k) \\ x^{k+1} &\leftarrow \mathbb{P}_{\mathcal{X}}[x^k - \alpha^k \nabla \phi_e^k] \end{aligned}$$

2. Equilibrium seeking step:

$$\mathbf{y}^{k+1} \leftarrow \mathbf{y}_*(x^{k+1})$$

3. Sensitivity computation step:

$$s^{k+1} \leftarrow \mathbf{J}\mathbf{y}_*(x^{k+1})$$

Under the considered problem setup, the convergence of Algorithm 1 to critical points of (17) follows by Theorem 2 in [15]. We omit the details here for the sake of brevity. In

¹If not differentiable, the standard Jacobians are replaced by elements of the conservative Jacobians. For detailed explanations and proofs, see [15].

the following subsections, we describe how to perform steps **2** and **3** of Algorithm 1 for the specific game setup in (7) in a computationally efficient manner.

B. Equilibrium Seeking Step

To efficiently find a solution of the allocation game (7), we consider the surrogate optimization problem

$$\min_{\mathbf{y}} \sum_{i \in \mathcal{I}} J^i(\mathbf{y}^i) \quad (19a)$$

$$\text{s.t. } \mathbf{y} \in \mathcal{Y}(x), \quad (19b)$$

whose minimizer corresponds to the v-GNE of (7). The individual objective functions (6) are decoupled, such that the sum of all teams’ objectives constitutes the potential function [18, p. 243], allowing us to rewrite (7) as (19). A formal proof can also be obtained by comparing the Karush-Kuhn-Tucker (KKT) conditions [20, Th. 4.8], but is omitted here due to space limitations. The sum in (19a) can be written as

$$J(\mathbf{y}) = \sum_{i \in \mathcal{I}} J^i(\mathbf{y}^i) = q^\top \mathbf{y} + \frac{1}{2} \epsilon (\mathbf{y} - \hat{\mathbf{y}})^\top (\mathbf{y} - \hat{\mathbf{y}}), \quad (20)$$

while the collection of constraints (19b) can be expressed as

$$\mathcal{Y}(x) = \{\mathbf{y} \in \mathbb{R}^{n_y} | A\mathbf{y} = b, G\mathbf{y} \leq h + Hx\}. \quad (21)$$

To improve the computational efficiency, we reduce the problem size by eliminating variables, as outlined in [21, Sec. 10.1.2]. We substitute $\mathbf{y} = F_T \tilde{\mathbf{y}} + \mathbf{y}^\dagger$ into $J(\mathbf{y})$, where the transformation matrix F_T consists of basis vectors of the nullspace of A and \mathbf{y}^\dagger is any solution to $A\mathbf{y}^\dagger = b$. By choosing $\tilde{\mathbf{y}} = \mathbf{y}^\dagger$, the objective reduces to $J(\tilde{\mathbf{y}}) = (F_T q)^\top \tilde{\mathbf{y}} + \frac{1}{2} \epsilon \tilde{\mathbf{y}}^\top F_T^\top F_T \tilde{\mathbf{y}}$. The new optimization variable $\tilde{\mathbf{y}}$ is reduced in dimension by the number of rows in A , i.e., the number of equality constraints. Finally, the optimization problem reads as follows

$$\begin{aligned} \min_{\tilde{\mathbf{y}}} \quad & \|F_T \tilde{\mathbf{y}} + q/\epsilon\|_2^2 \\ \text{s.t.} \quad & \tilde{G} \tilde{\mathbf{y}} \leq \tilde{h} + Hx, \end{aligned} \quad (22)$$

where $q, \tilde{G} = GF_T, \tilde{h} = h - G\mathbf{y}^\dagger, H$ depend on the problem data. Problem (22) is an inequality-constrained quadratic program (QP) with a convex objective, that can be efficiently solved using off-the-shelf solvers.

C. Sensitivity Computation

The equilibrium sensitivity $\mathbf{J}\mathbf{y}_*(x)$ corresponds to the sensitivity of the solutions to the QP (22) with respect to changes in x , which we compute using the approach in [22]. First, we find the total differentials of the KKT conditions of (22), more specifically of the stationarity condition and complementarity slackness at the optimal point

$$0 = \tilde{G}^\top d\boldsymbol{\lambda} + d\tilde{\mathbf{y}}, \quad (23)$$

$$0 = \text{diag}(d\boldsymbol{\lambda})(\tilde{G}\tilde{\mathbf{y}}_* - \tilde{h} - Hx) + \text{diag}(\boldsymbol{\lambda}_*)(\tilde{G}d\tilde{\mathbf{y}} - Hd\tilde{\mathbf{y}}), \quad (24)$$

with the Lagrange multipliers $\boldsymbol{\lambda}$ for the inequality constraints and the differentials $d\boldsymbol{\lambda}$, $d\tilde{\mathbf{y}}$ and $d\tilde{\mathbf{y}}$. The columns of the sensitivity matrix $\mathbf{J}\mathbf{y}_*(x)$ are the solutions of $d\tilde{\mathbf{y}}$ when

replacing dx by the columns of the identity matrix. Mathematically, for each single constraint, the optimal solution $\tilde{\mathbf{y}}_*$ and the differential $d\lambda$ do not influence $d\tilde{\mathbf{y}}$ and dx , because either $\tilde{G}\tilde{\mathbf{y}}_* - \tilde{h} - Hx = 0$ (active constraint), or $\lambda_* = 0$ (inactive constraint), or both are zero. Intuitively, since the DC operator's decisions only influence the lower level through the constraints, the sensitivity does not depend on the teams' objectives. Therefore, we only consider the rows of \tilde{G} and H which correspond to the active inequality constraints and solve the following system of linear equations

$$\tilde{G}_{k,\cdot} d\tilde{\mathbf{y}} - H_{k,\cdot} dx = 0, \quad \forall k \text{ with } \lambda_{*,k} > 0. \quad (25)$$

IV. SIMULATION RESULTS

A. Simulation Setup

We consider a network of 12 DCs distributed over 4 continents, whose locations we source from Google's fleet of DCs [23]. We use real carbon intensity data for those locations, obtained from Electricity Maps [24], which covers 24 hours with a sampling rate of 5 hours, starting on February 22nd, 2023, at 10 o'clock. The maximum computational capacities $x_{d,t}^{\max}$ of the DCs vary due to inflexible loads that fluctuate in a sinusoidal shape, similar to [25]. We set the parameter for the quadratic penalty in the team's objective to a small value $\epsilon = 2 \times 10^{-8}$, which means that teams virtually have no preferred allocation, but only care about the processing time and migration cost. The p -norm parameter for the peak shaving cost in (14) is set to $p = 6$, a value approximating the infinity norm while retaining numerical stability. To improve the conditioning of the DC operator's cost function (12), we include the uniform term $\frac{1}{2} \mathbf{1}_{n_x}^\top x$, that ensures that $\nabla_1 \phi(x, \mathbf{y}_*(x))$ is nonzero. Moreover, to ensure the feasibility of the parametrized allocation game (7), at each step of Algorithm 1, we impose the following: In the first time step, the VCCs at each DC shall not exceed the cumulative predicted load of all teams, that have their data uploaded at this DC, i.e., $x_{d,1} \leq \sum_{\{i \in \mathcal{I} \mid d=d^i\}} v^i$, for all $d \in \mathcal{D}$.

B. The Impact of Temporal Shifting & Spatial Migration

To analyze the impact of temporal shifting, we compare our allocation mechanism with a naive algorithm that allocates the compute jobs according to their priority, always using the full available capacity $x_{d,t}^{\max}$. We consider three load scenarios, each characterized by different compute job volumes: (a) large compute jobs that typically require more than one time step at a DC location; (b) multiple small compute jobs that can be processed during a single time step; and (c) a mixture of large and small compute jobs. As shown in Fig. 3, temporally shifting compute loads significantly reduces carbon emissions, depending on the load scenario.

To investigate the impact of spatial migration on carbon emissions, we modify the migration price ξ in the DC operator's objective function. A very large ξ corresponds to disabling the migration of compute jobs. As shown in Fig. 4, increasing ξ increases carbon emissions. This demonstrates the potential to save carbon emissions that comes with enabling spatial migration.




Fig. 3. Carbon emission savings due to time-shifting. Bilevel game vs. naive approach (full capacity utilization) in three scenarios.




Fig. 4. Carbon emissions, normalized by compute volume, vs migration price ξ , for three scenarios. The growth demonstrates a positive impact of spatial migration.

C. Co-design vs. Sequential Optimization

In this case study, we compare our bilevel approach for co-designing VCCs and allocations with a sequential optimization approach, similar to the scheduler-agnostic scheme used by Google [14]. The latter consists of two sequential steps: In the first step, optimal VCCs are computed by solely using a forecast of the compute load and carbon intensity. In the second step, the teams find the optimal allocation given the previously computed VCCs. In this sequential approach, the DC operator must optimize over the worst-case allocation scenarios, namely, the VCCs, rather than the actual allocations. The sequential approach does not take the migration of compute loads into account. Thus, the resulting cost function of the DC operator reads as

$$\hat{\phi}(x) = \sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} \rho_{d,t}^{\text{carb}} x_{d,t} + \sum_{d \in \mathcal{D}} \left(\sum_{t \in \mathcal{T}} (x_{d,t})^p \right)^{1/p}, \quad (26)$$

where the actual $\sum_{i \in \mathcal{I}} y_{d,t}^i$ have been replaced by the VCCs $x_{d,t}$. Overall, this sequential optimization scheme reads as

$$\text{Step 1: } \bar{x} = \arg \min_x \hat{\phi}(x), \quad (27)$$

$$\text{Step 2: } \mathbf{y} = \mathbf{y}_*(\bar{x}). \quad (28)$$

Further, we introduce the following two metrics to evaluate the performance of the allocations:

$$1) \text{ allocation fairness: } \psi \left(\left\{ \tau^{\text{time},i}(y^i)/v^i \right\}_{i \in \mathcal{I}} \right),$$

$$2) \text{ total waiting time: } \sum_{i \in \mathcal{I}} \tau^{\text{time},i}(y^i),$$

where $\psi(\cdot)$ is the empirical standard deviation and $\tau^{\text{time},i}(y^i)$ denotes the time cost, i.e., the first linear term in (6). The fairness criterion represents the heterogeneity of the waiting times among the teams. In contrast, the second criterion is the total wait time for all jobs, weighted by their time priority parameter τ^i .

In Fig. 5, we show the outcome of simulations with varying migration prices to showcase the differences between the bilevel game (16) and the sequential approach (28). Generally, fairness and waiting time are improved when co-designing VCCs and the compute job schedule. The difference becomes even more apparent with higher values of ξ , as this directly addresses the shortcoming of not featuring migration in the sequential approach. However, as demonstrated by the results in Fig. 6, the sequential

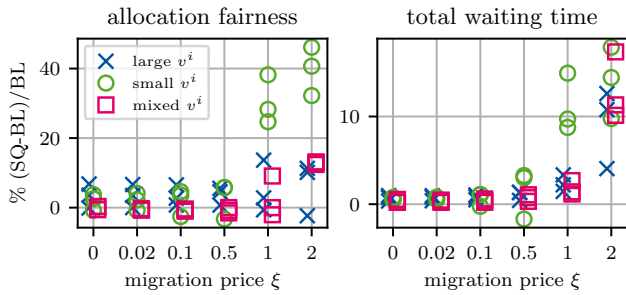


Fig. 5. Fairness and total waiting time in a direct comparison between our approach (BL) and sequential optimization (SQ). Sequential optimization scores higher (worse) values in both, and the difference grows with increasing migration price.

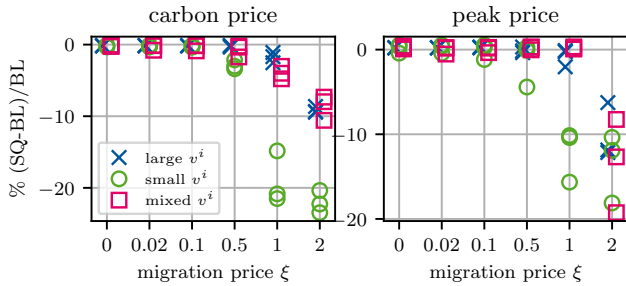


Fig. 6. Carbon emissions and peak price in a direct comparison between our approach (BL) and sequential optimization (SQ). SQ scores lower (better) values in both for high migration prices in all three scenarios.

optimization approach scores lower carbon emissions and better peak-shaving performance for large migration prices.

In summary, if we enable spatial shifting, the co-design of VCCs and allocations allows us to perform similarly to the sequential approach regarding decarbonization and peak-shaving. Yet, co-design can lead to reduced and fairer waiting times for the teams. This may incentivize users to participate in the proposed coordination mechanism.

V. CONCLUSIONS

We modeled the problem of co-designing the allocation of flexible batch compute jobs and virtual capacity curves of DCs as a bilevel game. This formulation promotes the supervisory control objectives of the DC operator and models each team as a player competing for the computational resources in the DCs. A local solution of the resulting game is found by deploying the recently developed algorithm [15], with some ad-hoc modifications that improve efficiency. Simulation results show that allowing for spatial migration and temporal shifting of compute jobs has a high potential to reduce carbon emissions in the operation of DCs. Moreover, compared with standard sequential optimization, our hierarchical approach can reduce the total waiting time of compute jobs and improve the fairness of their allocation.

REFERENCES

- [1] International Energy Agency, “Global trends in internet traffic, data centres workloads and data centre energy use, 2010-2020, iea, paris,” 2020. [Online]. Available: <https://t.ly/R5Gy>
- [2] Synergy Research Group, “Number of hyperscale data centers worldwide from 2015 to 2021,” 2021.

- [3] International Energy Agency, “Data centres and data transmission networks,” 2023. [Online]. Available: <https://t.ly/bmyWg>
- [4] A. Khosravi, L. L. H. Andrew, and R. Buyya, “Dynamic vm placement method for minimizing energy and carbon cost in geographically distributed cloud data centers,” *IEEE Transactions on Sustainable Computing*, vol. 2, no. 2, pp. 183–196, 2017.
- [5] 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,” *Performance Evaluation*, vol. 70, no. 10, pp. 770–791, 2013.
- [6] R. Rahmani, I. Moser, and A. L. Cricenti, “Inter-continental data centre power load balancing for renewable energy maximisation,” *Electronics*, vol. 11, no. 10, p. 1564, 2022.
- [7] M. Xu and R. Buyya, “Managing renewable energy and carbon footprint in multi-cloud computing environments,” *Journal of Parallel and Distributed Computing*, vol. 135, pp. 191–202, 2020.
- [8] N. Buchbinder, N. Jain, and I. Menache, “Online job-migration for reducing the electricity bill in the cloud,” in *Networking*, 2011, Conference Proceedings.
- [9] J. Shuja, A. Gani, S. Shamshirband, R. W. Ahmad, and K. Bilal, “Sustainable cloud data centers: A survey of enabling techniques and technologies,” *Renewable and Sustainable Energy Reviews*, vol. 62, pp. 195–214, 2016.
- [10] M. Dabbagh, B. Hamdaoui, A. Rayes, and M. Guizani, “Shaving data center power demand peaks through energy storage and workload shifting control,” *IEEE Transactions on Cloud Computing*, vol. 7, no. 4, pp. 1095–1108, 2019.
- [11] M. T. Takcı, T. Gözel, and M. H. Hocaöglü, “Quantitative evaluation of data centers’ participation in demand side management,” *IEEE Access*, vol. 9, pp. 14 883–14 896, 2021.
- [12] A. Radovanovic, B. Chen, S. Talukdar, B. Roy, A. Duarte, and M. Shahbazi, “Power modeling for effective datacenter planning and compute management,” *IEEE Transactions on Smart Grid*, vol. 13, no. 2, pp. 1611–1621, 2022.
- [13] J. Doyle, R. Shorten, and D. O. Mahony, “Stratus: Load balancing the cloud for carbon emissions control,” *IEEE Transactions on Cloud Computing*, vol. 1, no. 1, pp. 1–1, 2013.
- [14] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care, S. Talukdar, E. Mullen, K. Smith, M. Cottman, and W. Cirne, “Carbon-Aware Computing for Datacenters,” *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, Mar. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/9770383/>
- [15] P. D. Grontas, G. Belgioioso, C. Cenedese, M. Fochesato, J. Lygeros, and F. Dörfler, “Big hype: Best intervention in games via distributed hypergradient descent,” *arXiv preprint arXiv:2303.01101*, 2023.
- [16] M. Maljkovic, G. Nilsson, and N. Geroliminis, “On finding the leader’s strategy in quadratic aggregative stackelberg pricing games,” in *2023 European Control Conference (ECC)*, 2023, pp. 1–6.
- [17] G. Belgioioso, P. Yi, S. Grammatico, and L. Pavel, “Distributed generalized Nash equilibrium seeking: An operator-theoretic perspective,” *IEEE Control Systems*, vol. 42, no. 4, pp. 87–102, 2022.
- [18] F. Facchinei, V. Piccialli, and M. Scandrone, “Decomposition algorithms for generalized potential games,” *Computational Optimization and Applications*, vol. 50, no. 2, pp. 237–262, 2011.
- [19] N. H. Tran, D. H. Tran, S. Ren, Z. Han, E. N. Huh, and C. S. Hong, “How geo-distributed data centers do demand response: A game-theoretic approach,” *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 937–947, 2016.
- [20] F. Facchinei and C. Kanzow, “Generalized Nash equilibrium problems,” *Annals of Operations Research*, vol. 175, no. 1, pp. 177–211, 2010.
- [21] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [22] B. Amos and J. Z. Kolter, “Optnet: Differentiable optimization as a layer in neural networks,” in *International Conference on Machine Learning*. PMLR, 2021, Conference Proceedings, pp. 136–145.
- [23] Google, “Data center locations,” 2023. [Online]. Available: <https://www.google.com/about/datacenters/locations/>
- [24] Electricity Maps, “Organising the world’s electricity data,” 2023. [Online]. Available: <https://www.electricitymaps.com>
- [25] N. Hogade, S. Pasricha, and H. J. Siegel, “Energy and network aware workload management for geographically distributed data centers,” *IEEE Transactions on Sustainable Computing*, vol. 7, no. 2, pp. 400–413, 2022.