

CPU Scheduling in Data Centers Using Asynchronous Finite-Time Distributed Coordination Mechanisms

Andreas Grammenos¹, Themistoklis Charalambous², *Senior Member, IEEE*, and Evangelia Kalyvianaki

Abstract—We propose an asynchronous iterative scheme that allows a set of interconnected nodes to distributively reach an agreement within a pre-specified bound in a finite number of steps. While this scheme could be adopted in a wide variety of applications, we discuss it within the context of task scheduling for data centers. In this context, the algorithm is guaranteed to *approximately* converge to the optimal scheduling plan, given the available resources, in a finite number of steps. Furthermore, by being asynchronous, the proposed scheme is able to take into account the uncertainty that can be introduced from straggler nodes or communication issues in the form of latency variability while still converging to the target objective. In addition, by using extensive empirical evaluation through simulations we show that the proposed method exhibits state-of-the-art performance.

Index Terms—CPU, distributed coordination, finite-time termination, optimization, ratio consensus, scheduling.

I. INTRODUCTION

CLOUD computing provides software and hardware resources on demand via the Internet and has become the predominant model for application deployment. The backbone of modern Cloud infrastructure consists of a network of data centers, each equipped with thousands of server machines, running diverse application workloads, supporting uncoordinated and heterogeneous users and their applications [1]. Data center resource management is the fundamental task of allocating resources (e.g., CPU, memory, network bandwidth, and disk space) to workloads such that their performance objectives are satisfied and the overall data center utilization is kept high [2]. Notably, even slight deviations from the desired objectives can have substantial detrimental effects with millions of dollars in revenue potentially lost [3]. Therefore, scheduling in data centers

is the most fundamental operation responsible for allocating resources to workloads while satisfying their performance requirements [4]. In doing so, scheduling aims to find the best placement of jobs within the available compute nodes that maximizes the overall utilization of resources and which can ultimately lead to a massive reduction in operational and capital costs.

More formally, scheduling can be viewed as an *optimization problem* in which workloads are allocated to server machines such that a performance goal is optimized while all constraints are satisfied [5], [6]. In this paper, we focus on *minimizing the sum of CPU utilization across servers*. In other words, the workload should be shared proportionally across servers based on their hardware, such that they all use the minimum percentage of their capacity and essentially the total workload at each server node is balanced and proportional to its available resources. The main reason for this formulation is to avoid overloading specific servers and so to efficiently serve workloads. Solving a scheduling optimization problem in such a large-scale system in a centralized fashion is challenging due to the size of the network and the dynamic nature of resource requirements of incoming and existing workloads. In general, centralized approaches in multi-component systems require the collection of measurements or other information to a central location (at possibly high communication and computational cost), the computation of quantities of interest at this central location, and then the dissemination of these quantities to (a subset of) the components. This approach, as it is the case in Clouds, is often inefficient. This is because centralised approaches focus the entire load towards a single node. This can not only be a point of failure, but also create congestion in the network causing often causing delays and spikes in response times [4], [7], [8]. Cooperative distributed coordination algorithms have therefore received tremendous attention, especially during the last two decades. Several diverse research communities (e.g., biology, physics, control, communication and computer science) have made several contributions that have resulted in many recent advances in so called consensus-based approaches (see, for example, [9]) and in distributed computation of functions of geographically dispersed data, also known as in-network computation (see, for example, [10] and references therein). Classical approaches in distributed coordination algorithms typically assume timely and reliable exchange of information between neighboring components of a given multi-component

Manuscript received 6 January 2022; revised 6 November 2022; accepted 1 January 2023. Date of publication 11 January 2023; date of current version 16 June 2023. Recommended for acceptance by Chi Harold Liu. (*Corresponding author: Themistoklis Charalambous.*)

Andreas Grammenos is with the Department of Computer Science and Technology, University of Cambridge, Cambridge CB2 1TN, U.K., and also with Alan Turing Institute, London NW1 2DB, U.K. (e-mail: ag926@cl.cam.ac.uk).

Themistoklis Charalambous is with the Department Electrical and Computer Engineering, School of Engineering, University of Cyprus, 1678 Nicosia, Cyprus, and also with the Department of Electrical Engineering and Automation, School of Electrical Engineering, Aalto University, 02150 Espoo, Finland (e-mail: themistoklis.charalambous@aalto.fi).

Evangelia Kalyvianaki is with the Department of Computer Science and Technology, University of Cambridge, Cambridge CB2 1TN, U.K. (e-mail: ek264@cl.cam.ac.uk).

Digital Object Identifier 10.1109/TNSE.2023.3236214

system. These assumptions are not necessarily valid in practical settings due to varying delays that might affect transmissions at different times, as well as possible changes in the underlying interconnection topology (e.g., due to unexpected cluster changes as nodes randomly fail and/or abnormal runtime behaviors due to software or configuration faults and resource contention) [1], [11]. In this work, we propose a distributed coordination protocol to overcome these limitations. To this end, we posit a novel scheme that takes in account these potential latency variations in the form of explicit delays in the communication links during planning, while still remaining asynchronous in its operation and we guarantee that it will converge in finite-time.

A. Contributions

For the context of this work, we formulate the CPU scheduling as a distributed optimization problem and solve it using distributed coordination mechanisms. More concretely, the contributions of the paper are as follows.

- First, using existing theory from optimization, we provide the closed form solution, which requires the knowledge of global parameters, such as, the total capacity of the network and the total incoming workload.
- Second, it is shown that the problem can be solved in a distributed fashion.
- When the updates of the nodes are synchronous, we adopt a mechanism which uses a well-known consensus algorithm (namely, ratio consensus) proposed in [12], with which an *approximate* solution is reached in a finite number of steps.
- When the updates of the nodes are asynchronous, we adopt a mechanism, of similar flavor to the one proposed in [13], in which finite-time average consensus is achieved in the presence of bounded time-varying delays. More specifically, our proposed algorithm allows the nodes to distributively compute the optimal value to within an error bound in a finite number of steps. The methodology builds upon (i) robustified ratio consensus [14], [15], a distributed iterative algorithm in which each node maintains two state variables where the ratio of the states converges asymptotically to a constant that is equal for all the nodes, and (ii) asynchronous max – consensus algorithm [16].
- Finally, numerical examples and evaluations show the efficacy of the proposed solutions.

The main benefit of our approach is that the global optimization problem is decomposed into local objectives and the problem is then solved in a distributed manner via our proposed distributed coordination mechanisms, which provide a way for the nodes to terminate iterations simultaneously, while ensuring at the same time that the worst-case error lies within the pre-specified bound. These properties make these mechanisms suitable for applications in which (repeated) optimization problems have to be solved fast and in a finite number of steps. Moreover, contrary to methods such as ADMM our scheme requires significantly less resources for its computation to reach similar objectives as can be seen from the results put forth in

recent studies [17], [18]. This property can be particularly useful as most scheduling operations assume minimal processing latency to reach a solution for the optimal placement of tasks. To the best of our knowledge, this is the first algorithm with finite-time termination guarantees that can handle delays and provide asynchronous consensus.

B. Related Work

1) *Data Center Scheduling*: Centralized data center schedulers such as [19], [20], [21], [22], [23], [24], as well as the well known centralized approaches of Google’s Borg and Kubernetes [25] and Microsoft’s Resource Central [2], are able to provide optimized scheduling decisions under specific constraints and goals. More recently, there have been some centralized schedulers that tackle utilisation optimization focusing on energy efficiency [26], [27], [28]. However, they require continuous transferring of resource information at the centralized scheduler which increases data center network traffic. Furthermore, centralized schedulers typically lack of large-scale scalability and they can be a single point of failure. In contrast, our distributed approach requires each node to send its estimated utilization to its out-neighbors only reducing therefore the total amount of information sent and uses the most up-to-date resource estimates for more accurate scheduling.

Popular decentralized schedulers such as [4], [7], [8], [29], [30], as well as the most recent primary autoscaler that Google uses on its internal cloud [31], aim to tackle data center scalability by allowing different scheduling decisions to occur in parallel by multiple schedulers. Such approaches span a wide spectrum of schedulers’ coordination—from schedulers operating independently from each other (e.g., [8]) to schedulers sharing some global resource information (e.g., [4], [29])—and they also differ in the way they detect and resolve conflicts in the allocation of shared resources. We remark that, while these solutions exhibit good empirical performance, lack formal guarantees and largely work by using heuristics [32]. Unfortunately, this can be problematic when volatile or unpredictable workloads are encountered [8]. Additionally, state sharing can be problematic in case of delays as such schedulers attempt to globally infer the state of the cluster and normally are not able to tolerate delays. This in turn can lead to suboptimal performance [33]. In contrast in our distributed approach all nodes/schedulers coordinate *asynchronously* by design to find optimal allocations at scheduling time without facing any conflicts.

Multi-resource allocation of tasks to data center nodes is known to be a APX-Hard [21]. Most scheduling approaches employ heuristics to solve the problem in reasonable timeframes [4], [21], [22], [34]. Fewer approaches tackle the problem using appropriate centralized solvers (e.g., IBM’s CPLEX in [22]) albeit for small problem sizes compared to today’s data center sizes of thousands of nodes. Such approaches highly depend on the compute and memory capacity of the centralized solver to handle hundreds of thousands of constraints typically present in such problem formulations.

Our approach is to formulate the problem of CPU task scheduling in data centers as a distributed optimization one to solve it using distributed coordination mechanisms. An approximate (not accurate) solution can be computed in a finite number of steps and is guaranteed to complete while exhibiting graceful scaling. These properties enable its application to data center sized scheduling problems containing even tens of thousands of participating nodes.

2) *Distributed Finite-Time Average Consensus*: A distributed system or network consists of a set of components (nodes) that can share information via connection links (edges), forming a directed interconnection topology (directed graph). In general, the objective of a consensus problem is to have all agents agree upon a certain (*a priori* unknown) quantity of interest that is typically a function of some values that the nodes initially possess. When the agents (asymptotically) reach an agreement to the same value, we say that the distributed system (asymptotically) reaches consensus. Such problems include network coordination problems involving self-organization, formation of patterns, parallel processing, and distributed optimization. The problem of convergence of discrete-time consensus algorithms was initially targeted by Tsitsikis et al. [35] and subsequently by many other researchers (see, for example, [36], [37], [38], [39], [40], [41]). Convergence of consensus algorithms can usually be established under relatively weak requirements. Common challenges include the handling of node failures, transmission delays on the transfer of data between agents, packet losses in wireless communication networks, and inaccurate sensor measurements. As a result, it is imperative to address agreement problems that consider networks of dynamical agents, possibly with directed information flow, under delays and/or changing topologies. One of the most well known consensus problems is the so-called *average consensus* problem in which agents aim to reach the average of their initial values (see, for example, [42], [43]). This work is based on *synchronous* and *asynchronous* finite-time average consensus algorithms. There have been several works on synchronous finite-time average consensus algorithms due to their use *i*) in resource-constrained applications (such as wireless sensor networks) since they save energy and computational resources, and *ii*) in applications in which the result of the consensus algorithm is used in real-time to perform other subsequent tasks (such as smart energy networks). Nevertheless, there have not been any works for the asynchronous case when consensus is achieved in a finite number of steps.

The model of asynchrony considered herein allows for heterogeneous, but bounded computation and communication delays, thus quantifying the degree of asynchrony by a bound on the time-delays. It is highlighted that the nodes are not required to know the bound for the execution of the algorithm. Finite-time average consensus in the presence of delays in directed graphs has been studied mainly by [44] for exact average consensus and more recently by [45] for approximate average consensus. Moreover, the bound provided has linear dependency to the maximum delay within the network multiplied with its diameter. This is a powerful result, as not only allows its applicability in traditional data centers where consensus can be achieved quickly but also in delay tolerant networks. This particular

category includes numerous types of networks with some notable examples being collaborative autonomous agents, mobile phones, IoT clusters, and others.

C. Organization

The remainder of the paper is organized as follows. In Section II, we give the necessary notation and describe the model of the system. In Section III, we provide the necessary background knowledge needed for the development of our results. In Section IV, we first provide the problem under consideration and then we modify it so that it is formulated as a distributed coordination. Next, in Sections V and VI we propose a synchronous and an asynchronous finite-time distributed algorithm, respectively, that solve the problem approximately. In Section VII, we demonstrate the efficacy of our proposed algorithms. In Section VIII, we provide a quantitative discussion of the contributions herein and discuss our findings. In Section IX we draw conclusions and discuss possible directions for future work.

II. NOTATION AND SYSTEM MODEL

A. Notational Conventions

The set of real (integer) numbers is denoted by \mathbb{R} (\mathbb{Z}) and the set of non-negative real (integer) numbers is denoted by \mathbb{R}_+ (\mathbb{Z}_+). Vectors are denoted by small letters whereas matrices are denoted by capital letters. A^T denotes the transpose of matrix A . The i^{th} component of a vector x is denoted by x_i . For $A \in \mathbb{R}^{n \times n}$, a_{ij} denotes the entry in row i and column j .

In multi-component systems with fixed communication links (edges), the exchange of information between components (nodes) can be conveniently captured by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ of order n ($n \geq 2$), where $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ is the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. An edge from node v_i to node v_j is denoted by $\varepsilon_{ji} = (v_j, v_i) \in \mathcal{E}$ and represents a communication link that allows node v_j to receive information from node v_i . A graph is said to be undirected if and only if $\varepsilon_{ji} \in \mathcal{E}$ implies $\varepsilon_{ij} \in \mathcal{E}$. A digraph is called connected if there exists a path from each vertex v_i of the graph to each vertex v_j ($v_j \neq v_i$). The diameter D of a graph is the longest shortest path between any two nodes in the network.

In *digraphs*, nodes that can transmit information to node v_j directly are said to be in-neighbors of node v_j and belong to the set $\mathcal{N}_j^- = \{v_i \in \mathcal{V} \mid \varepsilon_{ji} \in \mathcal{E}\}$. The cardinality of \mathcal{N}_j^- , is called the *in-degree* of v_j and is denoted by $\mathcal{D}_j^- = |\mathcal{N}_j^-|$. The nodes that receive information from node v_j belong to the set of out-neighbors of node v_j , denoted by $\mathcal{N}_j^+ = \{v_l \in \mathcal{V} \mid \varepsilon_{lj} \in \mathcal{E}\}$. The cardinality of \mathcal{N}_j^+ , is called the *out-degree* of v_j and is denoted by $\mathcal{D}_j^+ = |\mathcal{N}_j^+|$.

In the type of algorithms we will consider, we will associate a positive weight p_{ji} for each edge $\varepsilon_{ji} \in \mathcal{E} \cup \{(v_j, v_j) \mid v_j \in \mathcal{V}\}$. The nonnegative matrix $P = [p_{ji}] \in \mathbb{R}_+^{n \times n}$ (with p_{ji} as the entry at its j^{th} row, i^{th} column position) is a weighted adjacency matrix (also referred to as weight matrix) that has zero entries at locations that do not correspond to directed edges (or self-edges) in the graph. In other words, apart from the main diagonal,

the zero-nonzero structure of the adjacency matrix P matches exactly the given set of links in the graph.

B. System Model

In our setup, we assume a set \mathcal{V} of server compute nodes, denoted by $v_i \in \mathcal{V}$, which also operate as resource schedulers; this is a frequent occurrence in modern data-centers. All participating schedulers are interconnected with bidirectional communication links and, thus, the network topology forms a connected undirected graph.

A job is defined as a group of tasks and \mathcal{J} as the set of all jobs to be scheduled. Each job $b_j \in \mathcal{J}$, $j \in \{1, \dots, |\mathcal{J}|\}$ requires ρ_j cycles to be executed and their individual estimated cost is assumed to be known before the optimization starts. The time horizon of the optimization (denoted by T_h) is defined as the time period for which the optimization is considering the jobs to be running on the server nodes, before the next optimization decides the next allocation of resources. Hence, the CPU capacity of each node, considered during the optimization, is computed as

$$\pi_i^{\max} := c_i T_h, \quad (1)$$

where c_i is the sum of all clock rate frequencies of all processing cores of node v_i given in cycles/second. The CPU availability for node v_i at optimization step m (i.e., at time mT_h) is given by

$$\pi_i^{\text{avail}}[m] := \pi_i^{\max} - u_i[m], \quad (2)$$

where $u_i[m]$ is the number of unavailable/occupied cycles due to predicted or known utilization from already running tasks on the server over the time horizon T_h at step m .

Assumption 1: Since the time horizon T_h is a parameter chosen, we assume that the time horizon is chosen such that the total amount of resources demanded at a specific optimization step m , denoted by $\rho[m] := \sum_{b_j[m] \in \mathcal{J}[m]} \rho_j[m]$, is smaller than the total capacity of the network available, given by $\pi^{\text{avail}}[m] := \sum_{v_i \in \mathcal{V}} \pi_i^{\text{avail}}[m]$, i.e., $\rho[m] \leq \pi^{\text{avail}}[m]$.

This assumption indicates that there is no more demand than the available resources. In case this assumption is violated, the solution will be that all resources are being used and some workloads will not be scheduled, due to lack of resources. The workloads selected to be discarded are arbitrary and the purging does not adhere to any particular priority policy; the jobs are scheduled on a first-come, first-scheduled basis. A more sophisticated priority mechanism could be deployed whose task would be to allocate a subset of the requests only, based on some optimization problem (taking into account deadlines, etc). However, since the prioritization problem is out of the scope of this paper, this problem will be addressed separately.

In particular, in the context of this work, we focus on CPU utilization as it is one of the most important and precious resource for many workloads. We note that while communication costs are also important, as data center networking becomes faster, we believe that CPU remains the most important resource. Notably, forecasting resource demands can be challenging and costly [25], [29], without necessarily providing the expected

gains. Further, there are certain types of commonly encountered jobs (e.g., recurring batch processing workloads) that have known a-priority demands. In fact, such workloads are frequently encountered in enterprise environments as such current data center scheduler's (e.g., Google's Kubernetes [46]) operate on known resource demands.

III. PRELIMINARIES

A. Average Consensus

In a synchronous setting, each node v_j updates and sends its information to its out-neighbors (and also receives similar information from its in-neighbors) at discrete times $t(0), t(1), t(2), \dots$. We index nodes' information states and any other information at time $t(k)$ by k . We use $x_j[k] \in \mathbb{R}$ to denote the information state of node v_j at time t_k . In our setup, each node v_j updates and sends its information regarding its input workload ℓ_j (ℓ_j is the summation of workloads at node v_j), estimated needed utilization for other tasks u_j , and capacity π_j^{\max} to its out-neighbors.

At each step, node v_j updates its information state $x_j[k]$ by combining the available information received by its neighbors $x_i[k]$ ($v_i \in \mathcal{N}_j^-$) using a weighted linear combination, i.e.,

$$x_j[k+1] = p_{jj}[k]x_j[k] + \sum_{v_i \in \mathcal{N}_j^-} p_{ji}[k]x_i[k], \quad k \geq 0, \quad (3)$$

where $x_j[0] \in \mathbb{R}$ is the initial state of node v_j . The positive weights $p_{ji}[k]$ capture the weight of the information inflow from node v_i to node v_j at time k (note that unspecified weights in P correspond to pairs of nodes (v_j, v_i) that are not connected and are set (without loss of generality) to zero, i.e., $p_{ji}[k] = 0$, $\forall v_i \notin \mathcal{N}_j$). If we let $x[k] = (x_1[k] \ x_2[k] \ \dots \ x_n[k])^T$ and $P[k] = [p_{ji}[k]] \in \mathbb{R}_+^{n \times n}$, then (3) can be written in matrix form as

$$x[k+1] = P[k]x[k], \quad (4)$$

where $x[0] = (x_1[0] \ x_2[0] \ \dots \ x_n[0])^T \equiv x_0$. In this work, we consider a static network; as a result, the graph remains invariant. In this case, the weights can be chosen to be constant for all times k (i.e., $p_{ji}[k] = p_{ji} \ \forall k$), and equation (4) can be expressed as

$$x[k+1] = Px[k]. \quad (5)$$

We say that the nodes asymptotically reach average consensus if $\lim_{k \rightarrow \infty} x_j[k] = \frac{\sum_i x_i[0]}{n}$ for all $v_j \in \mathcal{V}$. The necessary and sufficient conditions for (5) to reach average consensus are the following [43]: (a) P has a simple eigenvalue $\lambda_i(P) = 1$ with left eigenvector $\mathbf{1}^T$ and right eigenvector $\mathbf{1}$, and (b) all other eigenvalues of P ($\lambda_j(P), j \neq i$) have magnitude less than 1 ($|\lambda_j(P)| < 1$). If $P \geq 0$ (as in our case), the necessary and sufficient condition is that P be a primitive doubly stochastic matrix.

B. Ratio Consensus

Dominguez-García and Hadjicostis in [47], propose an algorithm that solves the average consensus problem in a directed graph in which each node v_j distributively sets the weights on

its self-link and outgoing-links to be $p_{lj} = \frac{1}{1+\mathcal{D}_j^+}, \forall (v_l, v_j) \in \mathcal{E}$, so that the resulting weight matrix P is column stochastic, but not necessarily row stochastic. Average consensus is reached by using this weight matrix to run two iterations with appropriately chosen initial conditions. The algorithm is stated below for the specific choice of weights mentioned above (which assumes that each node knows its out-degree). Note, however, that the algorithm also works for any set of weights that adhere to the graph structure and form a primitive column stochastic weight matrix.

Lemma 1 ([47]): Consider a strongly connected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Let $y_j[k]$ and $z_j[k]$ (for all $v_j \in \mathcal{V}$ and $k = 0, 1, 2, \dots$) be the result of the iterations

$$x_j[k+1] = p_{jj}x_j[k] + \sum_{v_i \in \mathcal{N}_j^-} p_{ji}x_i[k], \quad (6a)$$

$$y_j[k+1] = p_{jj}y_j[k] + \sum_{v_i \in \mathcal{N}_j^-} p_{ji}y_i[k], \quad (6b)$$

where $p_{lj} = \frac{1}{1+\mathcal{D}_j^+}$ for $v_l \in \mathcal{N}_j^{+}$ (zeros otherwise), and the initial conditions are $x[0] = (x_0(1) \ x_0(2) \ \dots \ x_0(|\mathcal{V}|))^T \equiv x_0$ and $y[0] = \mathbf{1}$. Then, the solution to the average consensus problem can be asymptotically obtained as

$$\lim_{k \rightarrow \infty} \mu_j[k] = \frac{\sum_{v_\ell \in \mathcal{V}} x_0(\ell)}{|\mathcal{V}|}, \quad \forall v_j \in \mathcal{V},$$

where $\mu_j[k] = x_j[k]/y_j[k]$.

C. Synchronous max-consensus

It is desired that each node $v_j \in \mathcal{V}$ of a network reaches consensus on the maximum value of the initial states/measurements under the assumption that all the nodes have a single real-valued state that they update based on local received states. Each node should reach the value $x_{\max} = \max_{v_j \in \mathcal{V}} x_j[0]$. The max-consensus algorithm is a simple algorithm for computing the maximum value (of, e.g., initial measurements in a sensor network) in a distributed fashion [48]. When the updates are synchronous, in the absence of communication noise (as it is the case in this work), max-consensus can be done by having each node $v_j \in \mathcal{V}$ update the state value with the largest received value in every iteration; the update rule is as follows:

$$x_j[k+1] = \max_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} \{x_i[k]\}. \quad (7)$$

It has been shown (see, e.g., [16, Theorem 5.4]) that this algorithm converges to the maximum value among all nodes in a finite number of steps s , $s \leq D$. Similar results hold for the min-consensus algorithm.

D. Optimization Problem

In a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ of $N = |\mathcal{V}|$ nodes, each node is endowed with a scalar quadratic cost function $f_i : \mathbb{R}^N \mapsto \mathbb{R}$. Most cases consider a quadratic cost function of the form:

$$f_i(z) = \frac{1}{2} \alpha_i (z - \rho_i)^2, \quad (8)$$

where $\alpha_i > 0$ and $\rho_i \in \mathbb{R}$ (in our case it is the demand in node v_i and it is a non-negative number). Parameter z is a function of the workload and it will be explained shortly. The global function $f : \mathbb{R}^N \mapsto \mathbb{R}$ is the sum of the cost function (8) of each node v_i . The main goal of the nodes is to allocate the jobs in order to minimize the cost function in a distributed fashion, by communicating with their neighbors only. Each node is thus required to solve the following optimization problem:

$$z^* = \arg \min_{z \in \mathcal{Z}} \sum_{v_i \in \mathcal{V}} f_i(z), \quad (9)$$

where \mathcal{Z} is the set of feasible values of parameter z . The solution of (9) in closed form can be expressed as

$$z^* = \frac{\sum_{v_i \in \mathcal{V}} \alpha_i \rho_i}{\sum_{v_i \in \mathcal{V}} \alpha_i}. \quad (10)$$

Note that by setting $\alpha_i = 1$ for all $v_i \in \mathcal{V}$, the solution is the average consensus.

IV. PROBLEM FORMULATION

A. Problem Statement

In our case, we are interested in finding a solution in which the total workload at each server node is balanced. This translates to having all server nodes having the same percentage of capacity being utilized during the execution of the tasks, i.e.,

$$\begin{aligned} \frac{w_i^*[m] + u_i[m]}{\pi_i^{\max}} &= \frac{w_j^*[m] + u_j[m]}{\pi_j^{\max}} \\ &= \frac{\rho[m] + u_{\text{tot}}[m]}{\pi^{\max}} \quad \forall v_i, v_j \in \mathcal{V}, \end{aligned} \quad (11)$$

where $w_i^*[m]$ is the *optimal* workload to be added to server node v_i at optimization step m , $\pi^{\max} := \sum_{v_i \in \mathcal{V}} \pi_i^{\max}$ and $u_{\text{tot}}[m] = \sum_{v_i \in \mathcal{V}} u_i[m]$.

The aim of this work is to find the optimal solution at every optimization step m via a distributed coordination algorithm run for a finite number of steps.

B. Modification of the Optimization Problem

To achieve the requirement set in (11), we modify (8) accordingly. Let

$$z[m] := \frac{w_i[m] + u_i[m]}{\pi_i^{\max}}. \quad (12)$$

For simplicity of exposition, and since we consider a single optimization step, we drop the index m . Then, the cost function $f_i(z)$ in (8) is given by

$$f_i(z) = \frac{1}{2} \pi_i^{\max} \left(z - \frac{\rho_i + u_i}{\pi_i^{\max}} \right)^2, \quad (13)$$

and the solution to problem (9) according to (10) is

$$z^* = \frac{\sum_{v_i \in \mathcal{V}} \pi_i^{\max} \frac{\rho_i + u_i}{\pi_i^{\max}}}{\sum_{v_i \in \mathcal{V}} \pi_i^{\max}} = \frac{\rho + u_{\text{tot}}}{\pi^{\max}}. \quad (14)$$

In other words, the nodes find the proportion of workload that each of them should have. From that each node is able to deduce




Fig. 1. At every optimization step of horizon T_h , the resource allocation optimization requires $T_c[m]$ steps to converge. Note that $T_c[m]$ is much smaller in duration than the time horizon of the optimization.

the workload w_i^* to receive, i.e.,

$$w_i^* = \frac{\rho + u_{\text{tot}}}{\pi^{\text{max}}} \pi_i^{\text{max}} - u_i. \quad (15)$$

V. A SYNCHRONOUS DISTRIBUTED ALGORITHM

The solution that we are aiming for should satisfy the balance condition in (11). For each node to be able to compute the optimal workload w_i^* in (15), the total workload ρ , the total estimated utilization needed for other tasks u_{tot} , and the total capacity of the network π^{max} are needed. For solving the problem in a distributed fashion we assume the following:

Assumption 2: The graph is static and strongly connected.

This assumption is, in general, valid even for large data-centers, since their topology is not expected to change for prolonged periods of time and remains mostly static, since failures are rare. Also, fault diagnostic mechanisms can be used to detect such failures and restore the connectivity of the network. Moreover, changes in the network can be handled by our algorithm, provided that the number of outgoing links can be found at each node in a distributed fashion, either because the links are bidirectional or because specific recovery schemes are deployed; see, for example, [49], [50].

Under Assumption 2, running the ratio consensus algorithm (6a) with initial conditions $y_j[0] = \ell_j + u_j$ and $z_{ij}[0] = \pi^{\text{max}}$, we obtain

$$\lim_{k \rightarrow \infty} y[k] = \lim_{k \rightarrow \infty} P^k y[0] = c \mathbf{1}^T y[0] = c(\rho + u_{\text{tot}}),$$

$$\lim_{k \rightarrow \infty} z[k] = \lim_{k \rightarrow \infty} P^k z[0] = c \mathbf{1}^T z[0] = c \pi^{\text{max}},$$

where c is a vector (the left eigenvector of column matrix P). Therefore,

$$\lim_{k \rightarrow \infty} \mu_j[k] = \lim_{k \rightarrow \infty} \frac{x_j[k]}{y_j[k]} = \frac{c_j(\rho + u_{\text{tot}})}{c_j \pi^{\text{max}}} = \frac{\rho + u_{\text{tot}}}{\pi^{\text{max}}}.$$

A. Finite-Time Implementation

Since the optimization is repeated periodically, the consensus algorithm should stop way before the beginning of the next optimization cycle, since the resources should be allocated and have the tasks allocated (and process as many of them as possible) before the next bunch of tasks arrives; see Fig. 1. However, often it is impossible or undesirable to predetermine the number of steps needed to stop the iterations. Towards this end, we deploy an algorithm that allows the nodes to distributively stop iterations in a finite number of steps, tolerating some deviation from the




Fig. 2. A simple example of a network consisting of 3 nodes. In the timeline of each node, blue ticks indicate an iteration for node v_i and the arrows indicate the transmissions. The time in between transmissions is the processing delay. The time from the beginning of the transmission to the end (arrow) is the transmission delay.

exact optimal solution. Before we proceed with the finite time implementation, we make the following assumption:

Assumption 3: The diameter of the network D is known to all server nodes.

Under Assumption 3, Cady et al. in [12] proposed an algorithm which is based on the ratio-consensus protocol [47] and takes advantage of min- and max-consensus iterations to allow the nodes to determine the time step, k_0 , when their ratios, $\{\mu_j[k_0] | v_j \in \mathcal{V}\}$, are within ϵ of each other.

First, we present the synchronous case, in order to demonstrate the main idea before we present the asynchronous case. Towards this end, we adopt the algorithm proposed by Cady et al. in [12] *mutatis mutandis*. More specifically, the algorithm makes use of the following ideas:

- Each node v_j runs ratio consensus iteration, as described in Lemma 1; in our case, we use initial conditions $y_j[0] = \ell_j + u_j$ and $z_j[0] = \pi_j^{\text{max}}$.
- At the same time, each node maintains two auxiliary states, $m_j[k]$ and $M_j[k]$, which are updated using min- and max-consensus, respectively.
- Every D steps (where D is the diameter of the graph) each node checks whether $|M_j[k] - m_j[k]| < \epsilon$. If this is the case, then the ratios for all nodes are close to the asymptotic value and it stops iterating. Otherwise, $m_j[k]$ and $M_j[k]$ are reinitialized to $\mu_j[k]$.

The algorithm, adopted to our case, is described in Algorithm 1 for digraphs (which means it holds for undirected graphs as well, that we consider in this case).

Remark 1: The number of iterations needed for the distributed algorithm to terminate at optimization step m , $T_c[m]$, is a multiple of the diameter of the network. As it will be shown in the simulations, the distributed algorithm converges fast and it only needs a fraction of the optimization step of horizon T_h ; see an illustration in Fig. 1.

VI. AN ASYNCHRONOUS DISTRIBUTED ALGORITHM

Resource allocation in data centers gives rise to large-scale problems and networks, which naturally call for asynchronous solutions. Let $t(0) \in \mathbb{R}_+$ the time at which the iterations for

Algorithm 1: Distributed Finite-Time Termination for Ratio Consensus.

Input: A strongly connected digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node $v_j \in \mathcal{V}$ knows its out-degree \mathcal{N}_j^+ . Initial values are $y_j[0] = \ell_j + u_j$ and $z_j[0] = \pi_j^{\max}$, and tolerance ϵ .

set
 $M_j[0] = +\infty$, $m_j[0] = -\infty$, $\text{flag}_j[0] = 0$, $\mu_j = \frac{y_j[0]}{z_j[0]}$

set $p_{lj} = \frac{1}{1+d_j^{\text{out}}}$, $\forall v_l \in \mathcal{N}_j^+ \cup \{v_j\}$ (zero otherwise)

for $k \geq 0$ **do**
 while $\text{flag}_j[k] = 0$ **do**
 if $k \bmod D = 0$ and $k \neq 0$ **then**
 if $|M_j[k] - m_j[k]| < \epsilon$ **then**
 set $\text{flag}_j[k] = 1$
 end if
 set $M_j[k] = m_j[k] = \mu_j[k] = \frac{y_j[k]}{z_j[k]}$
 end if
 broadcast to all $v_l \in \mathcal{N}_j^+$:
 $p_{lj}y_j[k]$, $p_{lj}z_j[k]$, $M_j[k]$, $m_j[k]$
 receive from all $v_i \in \mathcal{N}_j^-$:
 $p_{ji}y_i[k]$, $p_{ji}z_i[k]$, $M_i[k]$, $m_i[k]$
 compute
 $y_j[k] \leftarrow \sum_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} p_{ji}y_i[k]$
 $z_j[k] \leftarrow \sum_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} p_{ji}z_i[k]$
 $M_j[k] \leftarrow \max_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} M_i[k]$
 $m_j[k] \leftarrow \min_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} m_i[k]$
 end while
end for

the optimization start. We assume that there is a set of times $\mathcal{T} = \{t(1), t(2), t(3), \dots\}$ at which one or more nodes transmit some value to their neighbors. A message that is received at time $t(k_1)$ and processed at time $t(k_2)$, $k_2 > k_1$, experiences a process delay of $t(k_1) - t(k_2)$ (or a time-index delay $k_2 - k_1$). In Fig. 2, we show through a simple example how the time steps evolve for each node in the network; with $t_j(k)$ we denote the time step at which iteration k takes place for node v_j .

Assumption 4: There exists an upper bound B on the time-index steps that is needed for a node to process the information received from another node.

A. Asynchronous max-consensus

When the updates are asynchronous, for any node $v_j \in \mathcal{V}$, the update rule is as follows [16]:

$$x_j[t_j(k+1)] = \max_{v_i \in \mathcal{N}_j^-[t_j(k+1)] \cup \{v_j\}} \{x_i[t_j(k) + \theta_{ij}(k)]\},$$

where $x_i[t_j(k) + \theta_{ij}(k)]$ are the states of the in-neighbors $\mathcal{N}_j^-[t_j(k+1)]$ available at the time of the update. Variable $\theta_{ij}(k) \in \mathbb{R}$, evaluated with respect to the update time $t_j(k)$, is used here to express asynchronous state updates occurring at the neighbors of node v_j , between two consecutive updates of the state of node v_j . It has been shown in [16, Lemma 5.1] that this algorithm converges to the maximum value among all nodes in a finite number of steps s , $s \leq BD$.

B. Asynchronous (Robustified) Ratio Consensus

An adaptation of the above approach to a protocol where each node updates its information state $x_j[k+1]$ by combining the available (possibly delayed) information received by its neighbors $x_i[s]$ ($s \in \mathbb{Z}$, $s \leq k$, $v_i \in \mathcal{N}_j^-$) using constant positive weights p_{ji} was developed in [15]. Integer $\bar{\tau}_{ji}[k] \geq 0$ is used to represent the delay of a message sent from node v_i to node v_j at time instant k . We require that $0 \leq \tau_{ji}[k] \leq \bar{\tau}_{ji} \leq \bar{\tau}$ for all $k \geq 0$ for some finite $\bar{\tau} = \max\{\bar{\tau}_{ji}\}$, $\bar{\tau} \in \mathbb{Z}_+$. We make the reasonable assumption that $\tau_{jj}[k] = 0$, $\forall v_j \in \mathcal{V}$, at all time instances k (i.e., the own value of a node is always available without delay). Each node updates its information state according to:

$$x_j[k+1] = p_{jj}x_j[k] + \sum_{v_i \in \mathcal{N}_j^-} \sum_{r=0}^{\bar{\tau}} p_{ji}x_i[k-r]I_{k-r,ji}[r],$$

for $k \geq 0$, where $x_j[0] \in \mathbb{R}$ is the initial state of node v_j ; $p_{ji} \forall \varepsilon_{ji} \in \mathcal{E}$ form $P = [p_{ji}]$ that adheres to the graph structure, and is primitive column stochastic; and

$$I_{k,ji}(\tau) = \begin{cases} 1, & \text{if } \tau_{ji}[k] = \tau, \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

Lemma 2: [15, Lemma 2] Consider a strongly connected digraph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Let $y_j[k]$ and $z_j[k]$ (for all $v_j \in \mathcal{V}$ and $k = 0, 1, 2, \dots$) be the result of the iterations

$$y_j[k+1] = p_{jj}y_j[k] + \sum_{v_i \in \mathcal{N}_j^-} \sum_{r=0}^{\bar{\tau}} y_{ji}[k-r]I_{k-r,ji}[r], \quad (17)$$

$$z_j[k+1] = p_{jj}z_j[k] + \sum_{v_i \in \mathcal{N}_j^-} \sum_{r=0}^{\bar{\tau}} z_{ji}[k-r]I_{k-r,ji}[r], \quad (18)$$

with $y[0] = (y_0(1) \ y_0(2) \ \dots \ y_0(|\mathcal{V}|))^T \equiv y_0$ and $z[0] = \mathbf{1}$; $I_{k,ji}$ is an indicator function that captures the bounded delay $\tau_{ji}[k]$ on link (v_j, v_i) at iteration k (as defined in (16), $\tau_{ji}[k] \leq \bar{\tau}$). Then, the solution to the average consensus problem can be asymptotically obtained as

$$\lim_{k \rightarrow \infty} \mu_j[k] = \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|}, \quad \forall v_j \in \mathcal{V},$$

where $\mu_j[k] = y_j[k]/z_j[k]$.

C. Finite-Time Asynchronous Ratio Consensus

As it is the case for the synchronous distributed algorithm (see § V), the consensus algorithm should terminate before the next optimization step and in a distributed fashion. In what follows, we propose a distributed termination protocol for the asynchronous case, based on the one used for the synchronous case. We believe, that this is the first termination algorithm that can handle delays and perform asynchronous consensus.

The proposed termination algorithm has the same principles as before [12]. However, in order to make the ideas put forth in [12] applicable into the asynchronous case we expand upon

Algorithm 2: Distributed Finite-Time Termination for Asynchronous Ratio Consensus.

Input: A strongly connected digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node $v_j \in \mathcal{V}$ knows its out-degree \mathcal{N}_j^+ . Initial values are $y_j[0] = \ell_j + u_j$ and $z_j[0] = \pi_j^{\max}$, and tolerance ϵ .

set

$$M_j[0] = +\infty, m_j[0] = -\infty, \text{flag}_j[0] = 0, \mu_j = \frac{y_j[0]}{z_j[0]}$$

set $p_{lj} = \frac{1}{1+d_j^{\text{out}}}, \forall v_l \in \mathcal{N}_j^+ \cup \{v_j\}$ (zero otherwise)

for $k \geq 0$ **do**

while $\text{flag}_j[k] = 0$ **do**

if $k \bmod (1 + \bar{\tau})D = 0$ and $k \neq 0$ **then**

if $|M_j[k] - m_j[k]| < \epsilon$ **then**

set $\text{flag}_j[k] = 1$

end if

set $M_j[k] = m_j[k] = \mu_j[k] = \frac{y_j[k]}{z_j[k]}$

end if

broadcast to all $v_l \in \mathcal{N}_j^+$:

$p_{lj}y_j[k], p_{lj}z_j[k], M_j[k], m_j[k]$

receive from all $v_i \in \mathcal{N}_j^-$:

$p_{ji}y_i[k], p_{ji}z_i[k], M_i[k], m_i[k]$

compute

$$y_j[k] \leftarrow p_{jj}y_j[k] + \sum_{v_i \in \mathcal{N}_j^-} \sum_{r=0}^{\bar{\tau}} y_{ji}[k-r] I_{k-r,ji}[r]$$

$$z_j[k] \leftarrow p_{jj}z_j[k] + \sum_{v_i \in \mathcal{N}_j^-} \sum_{r=0}^{\bar{\tau}} z_{ji}[k-r] I_{k-r,ji}[r]$$

$$M_j[k] \leftarrow \max_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} \{M_i[t_j(k) + \theta_{ij}(k)]\}$$

$$m_j[k] \leftarrow \max_{v_i \in \mathcal{N}_j^- \cup \{v_j\}} \{m_i[t_j(k) + \theta_{ij}(k)]\}$$

end while

end for

them using several innovations. More concretely, when compared to the synchronous case the aforementioned innovations are outlined below:

- The min and max – consensus algorithm converge in $(1 + \bar{\tau})D$ steps [16].
- Every $(1 + \bar{\tau})D$ steps each node checks whether $|M_j[k] - m_j[k]| < \epsilon$. If this is the case, then the ratios for all nodes are close to the asymptotic value and it stops iterating. Otherwise, $m_j[k]$ and $M_j[k]$ are reinitialized to $\mu_j[k]$.

The algorithm is described in Algorithm 2 or digraphs; note that this implies it also holds for *undirected* graphs as well, that we consider in this case.

Theorem 1: Algorithm 2 converges in finite time.

Proof: From Lemma 2, we know that $\lim_{k \rightarrow \infty} \mu_j[k] = (\sum_{v_\ell \in \mathcal{V}} y_0(\ell)) / \mathcal{V}$, for all $v_j \in \mathcal{V}$. Therefore, it follows that

$$\lim_{k \rightarrow \infty} \left| \max_{v_j \in \mathcal{V}} \mu_j[k] - \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|} \right| = 0, \quad (19)$$

which means that essentially $\lim_{k \rightarrow \infty} M[k] = \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|}$. Additionally, k_0 exists, such that for all $k \geq k_0$, we have

$$\left| \mu_j[k] - \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|} \right| < \epsilon, \quad \forall v_j \in \mathcal{V}. \quad (20)$$

Therefore, it follows that

$$\left| \max_{v_j \in \mathcal{V}} \mu_j[k] - \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|} \right| < \epsilon, \quad (21)$$

In turn, this implies that there exists k_0 , such that for all $k \geq k_0$,

$$\left| M[k] - \frac{\sum_{v_\ell \in \mathcal{V}} y_0(\ell)}{|\mathcal{V}|} \right| < \epsilon. \quad (22)$$

Similar arguments hold for $m[k]$. Since $\{M[r(1 + \bar{\tau})D]\}_{r \in \mathbb{N}}$ and $\{m[r(1 + \bar{\tau})D]\}_{r \in \mathbb{N}}$ are sub-sequences of sequences that converge (due to the fact that asynchronous max – consensus converges within $(1 + \bar{\tau})D$ steps), then they converge to the same limit. Therefore, there exists r_0 , such that for all $r \geq r_0$, $|M[r(1 + \bar{\tau})D] - m[r(1 + \bar{\tau})D]| < \epsilon$.

Remark 2: We stress that similar results were proposed in [13] for guaranteeing convergence to approximate average consensus in a finite number of steps, allowing for time-varying bounded delays in information transmission and reception between agents. Nevertheless, apart from the fact that our results are obtained for an optimization problem for CPU scheduling, there are some additional differences:

- we use the consensus algorithm in the concept of asynchronous operation, rather than synchronous operation with delays, despite the fact that the mathematical analysis relies on similar concepts;
- the window used for updating the min/max value of the agents is different (for us this is $(1 + \bar{\tau})D$ while for them is $(1 + \bar{\tau})D + \bar{\tau}$), and
- we show via simulation that the lemmas (and, hence, the proofs) in [13] are incorrect (see also the discussion in Section VIII).

VII. SIMULATIONS

To validate our scheme, we divide our evaluation into three separate segments. The first focuses on simulating the performance using a simple, easy to understand, network of five nodes. The second one presents a thorough quantitative evaluation using simulations for various randomly generated graphs and latencies. The last one, provides a large scale evaluation with network graphs and simulation parameters that would be applicable in large scale data centers having thousands of nodes. To our knowledge this is the first work that tackles the problem at this scale in this setting while also providing a thorough evaluation and theoretical guarantees. All experiments are computed on a workstation using an AMD 3970X CPU with 32 cores at 4.0 GHz, 256 GB 3600 MHz DDR4 RAM, and Matlab R2022b (build 9.13.0.2080170)¹.

A. Evaluation Using a Small Network

The digraph is comprised out of $|\mathcal{V}| = 5$ vertices and has a diameter equal to $D = 4$; for helping exposition the exact digraph is shown in Fig. 3.

¹To foster reproducibility both code and datasets used for our numerical evaluation are publicly available at: <https://github.com/andylamp/federated-capacity-consensus>.




Fig. 3. The strongly connected digraph network comprised out of five nodes which is used to evaluate the validity of our results though an indicative, small-scale example.




Fig. 4. A simple example of a network of five nodes as described in Fig. 3 when the links experience time-varying delays with maximum delay ($\bar{\tau}$) of 4. The figure shows the evolution of the converge ratios across all nodes along with the min-consensus (dashed blue) and the max-consensus (dashed red).

All nodes are set with equal capacities and the workload vector ρ is set to $\rho = [1, 2, 3, 4, 5]$ in all runs. Further, we set the convergence threshold for the absolute difference of the quantity $|M_j[k] - m_j[k]| < \epsilon$ to $\epsilon = 10^{-5}$. Throughout our experiments, at each interval the workload to be scheduled is generated for each node independently. Concretely, each randomly chooses a job cost from a uniform distribution bounded between an acceptable cost range which is provided upon initialisation. These values are then concatenated to generate a workload vector which has a value between that range for each of the nodes.

Then in order to study the impact of increased delay in the number of total iterations required, we evaluate our proposed algorithm when using $\bar{\tau} = [4, 9]$. We start by showing the results for $\bar{\tau} = 4$ in Fig. 4. In this figure, we observe that convergence happens after 120 iterations which is $4(1 + \bar{\tau})D$, meaning that in total four rounds are required. Following, we shift our attention to Fig. 5 in which we show the results of the same experiment when using a delay value of $\bar{\tau} = 9$. Concretely, we see that the increased delay has an impact on the total iterations required to converge increasing them by a factor of about ≈ 1.6 when compared to the previous experiment.

We see that both figures converge in multiples of $(1 + \bar{\tau})D$ which requires six rounds when having $\bar{\tau} = 5$ and four rounds




Fig. 5. Converge ratios, when using the network of five nodes as described in Fig. 3 when the links experience time-varying delays with maximum delay ($\bar{\tau}$) of 9. The min-consensus and max-consensus are depicted by the dashed blue and red lines respectively.

when using $\bar{\tau} = 10$. Notably, as delay grows the round size increases linearly assuming we operate on the same graph (hence the diameter D remains the same). Indeed, the round size for $\bar{\tau} = 5$ is 20 iterations whereas in the case of $\bar{\tau} = 10$ the round size is 40 iterations. Quantitatively speaking, we observe that as the round size increases the number of rounds required to converge decreases. We conjecture that this can be attributed to the fact that as the round size increases the information has an elongated iteration window to propagate throughout the graph which in turn helps to converge with fewer rounds. However, since the results are simulated centrally even if the aggregated simulation cost is large, the amortised cost (e.g. the actual computation that would be required per node) is practically very low - even in the presence of large delays.

Remark 3: Note that there are some nodes $v_j \in \mathcal{V}$ for which the state $\mu_j[k']$ is larger than the maximum $M(k)$, where $k' > k$ and $k \bmod D = 0$ (note that this constitutes a counterexample to Lemma IV.2 in [45]). Despite the fact that the ratio is not monotonically decreasing (due to the nonlinearity imposed by the ratio), the main properties that guarantee the convergence of this algorithm is that the ratio is guaranteed to converge and the max-consensus algorithm converges within $(1 + \bar{\tau})D$ steps.

B. Evaluation Using Varying Delays and Network Sizes

The previous example is indicative on how our scheme performs in a tangible, small-scale scenario. In this section, we evaluate the performance of our proposed algorithm across a broader range of parameters reflecting realistic deployments, as such our generated topologies attempt to replicate ones that would be in real data-centers. To that end, we create a test suite monitoring both convergence and actual simulation execution time for varying graph sizes and delays. check this again: Concretely, for a given amount of trials, graph size dictated by $|\mathcal{V}|$, and a range of delays upper bounds we create a random graph for different unique pairs $(|\mathcal{V}|, \bar{\tau})$. The values considered for graph sizes and delays upper bounds are $|\mathcal{V}| = [20, 50, 100, 200, 300, 600]$ and $\bar{\tau} = [1, 5, 10, 15, 20, 30]$, respectively, which result in the




Fig. 6. Total number of iterations required to converge for each unique $(|\mathcal{V}|, \bar{\tau})$ pair averaged across 10 trials. The x -axis shows the different delays $(\bar{\tau})$ while each line represents the number of nodes $(|\mathcal{V}|)$ that exist within each graph.

evaluation of 36 unique $(|\mathcal{V}|, \bar{\tau})$ pairs. More specifically, for each unique $(|\mathcal{V}|, \bar{\tau})$ pair we perform 10 trials and average the results for each pair. We also note, that throughout our experiments, More concretely, as long as we are able to generate a connected random graph, all trial instances converge within the maximum iteration limit set; this value is set to 4000 iterations across all runs. Additionally, while the randomly generated network topologies are guaranteed to be strongly connected and of relatively low diameter they are not necessarily assured to be fat-tree and/or spine-leaf compliant topologies. We begin by presenting the number of iterations required to converge, on average, across 10 runs for each $(|\mathcal{V}|, \bar{\tau})$ pair; results are shown in Fig. 6.

Fig. 6 indicates that smaller networks require more iterations than larger ones to converge, which are still multiples of $(1 + \bar{\tau})D$. At first glance this observation might seem as counter-intuitive, however, we conjecture that such behaviour is encountered because the round size for smaller networks is smaller thus the system has fewer iterations to reach a steady state within each round. Indeed, similarly to the delay, recall that each round length is dictated by $(1 + \bar{\tau})D$; thus, fixing the delay $\bar{\tau}$ and increasing the diameter D —as is the case when the graph network grows—results in linear inflation of the round size. Notably, even if the round size increases this does not mean that the execution time is less. In fact it is quite the opposite since the total simulation time is higher as the network size increases. However, the extrapolated actual cost per node is much less. This is because, the workload for each can be parallelized and is asynchronous.

Fig. 7 shows the average execution time required to converge and the average iteration execution time for the same experiments discussed previously. As we can see from Fig. 7, the execution time scales exponentially as both delay and graph size increase. More importantly, this graph shows in practice that larger graphs take more time to converge than smaller ones given the same delay even if the actual rounds to converge are less as graph size increases. This is because, as we noted previously, even if the iterations are fewer each iteration within a larger graph takes significantly more time to complete in practice. However,

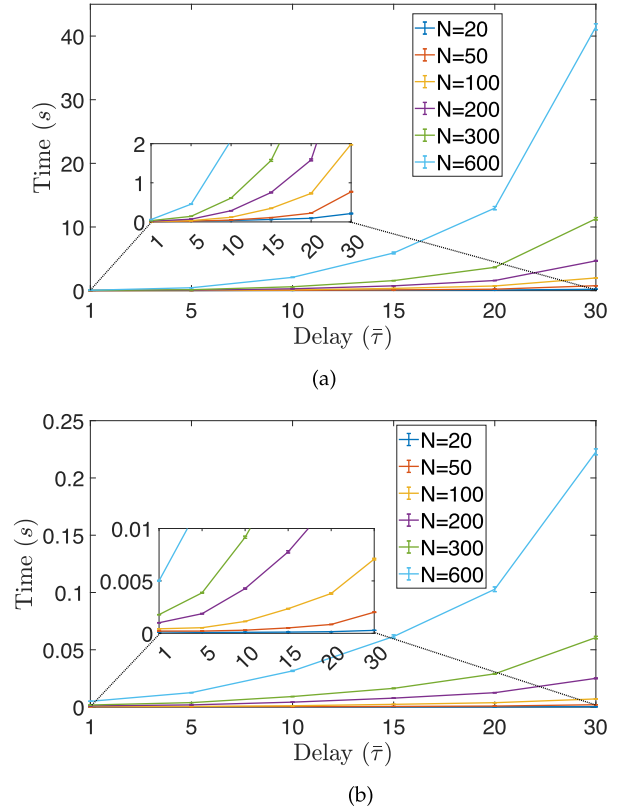


Fig. 7. Top (Fig. 7(a)), we present the total execution time required to converge for each unique $(|\mathcal{V}|, \bar{\tau})$ pair averaged across 10 trials. The x -axis shows the different delays upper bounds $(\bar{\tau})$ while each line represents the number of nodes $(|\mathcal{V}|)$ that exist within each graph. Bottom (Fig. 7(b)), we show the average iteration time for each of the different configurations. (a) Total simulation time. (b) Iteration time.

as a general trend we observe that regardless of the network size used in our experiments, if the delay remains below $\bar{\tau} = 10$, then it converges relatively quickly. Conversely, it seems that for delays greater than $\bar{\tau} = 15$ then the time to converge scales exponentially.

C. Data Center Scale Evaluation

Previous examples evaluate the performance of the algorithm in practical small-scale deployment. However, these experiments do not capture the scale of modern data centers which contain thousands of server machines. To that end, to evaluate the data center scalability of our scheme we perform experiments on thousands of nodes. We assume that in data centers most nodes are few hops away from each other, so we use graphs with a small diameter [51]. Further, we assume that the latency within data centers is near zero as shown before in order to satisfy the needs of modern workloads [52], [53]. To sum up, in order to provide a realistic data center scale representation, we create a simulation configuration that scales to thousands of nodes; considers graphs of a small diameter; and finally assumes low, even if variable, network delays upper bounds. Concretely, the values considered for the graph sizes and delays upper bounds are $|\mathcal{V}| = [20, 200, 500, 1000, 5000, 10000]$ and $\bar{\tau} = [1, 2, 3, 45]$ respectively; which result in the evaluation of




Fig. 8. Example run of a network comprised of 1000 nodes having a diameter equal to 2 and using a delay upper bound $\bar{\tau}$ of 1. The network converges to the optimal solution in very few iterations.




Fig. 9. Mean iterations to converge for different network sizes and delay values. Delay plays a larger role in smaller networks (< 200 nodes) whereas as network size increases the delay impact is lower.

30 unique $\langle |\mathcal{V}|, \bar{\tau} \rangle$ pairs. We note, however, that in order for modern data centers to maintain very low network communication delays, it is desirable to have just a couple of hops between nodes and, hence, we consider graphs with small diameter [51], [54]. As previously, for each unique $\langle |\mathcal{V}|, \bar{\tau} \rangle$ pair we perform 5 trials and average the results for each pair.

Fig. 8 illustrates the results of an example run of a network size of 1000 and a delay $\bar{\tau} = 1$. We can see that our scheme is able to converge to the optimal solution in very few iterations. This is attributed to the diameter of the graph which was equal to $D = 2$ and to low delays ($\bar{\tau} = 1$).

In the next data center scale experiment we vary the number of nodes from 20 to data center scale of 10000. We also vary the upper bound on the delay $\bar{\tau}$. Results are shown in Figs. 9 and 10. Fig. 9 shows the converge scaling with respect to the iterations required as the delays upper bound and network size grow. Fig. 10 shows the average total simulation time and per iteration time required per each network size and delays upper bound. Note, that the simulation indicates the *aggregated times* required to complete each round since for the context of this work we simulate our scheme centrally for all networks. In practice, in a real system, the actual execution cost per node would be much




Fig. 10. The above figures show the average time to converge (Fig. 10(a)) and the average iteration time (Fig. 10(b)) in the data centre scale experiments. We can see that as we increase the number of nodes, iterations take longer but overall we require less iterations to converge. This can be attributed to the low diameter of the graph, which allows more paths of communication between the nodes as their overall count in each network topology increases. (a) Total simulation time. (b) Iteration time.




Fig. 11. Converge statistics in the presence of both low (Fig. 11(a)) and higher upper ((Fig. 11(b))) bounds on delays ($\bar{\tau}$), equal to 1 and 5 respectively. As the network size grows the window which the first (*min*) and the last (*max*) node converges becomes zero. This indicates that as the network size grows we require fewer iterations to converge and all nodes will converge at the same iteration. Note, that as the delay scales delay results in an increase, on average, by a factor of $\approx 2x$ to the number of iterations required to converge. However, it is worth pointing out that the window to converge remains still very low, and sometimes zero as network size increases. (a) Delay $\bar{\tau} = 1$. (b) Delay $\bar{\tau} = 5$.

less since the workload would be executed asynchronously and concurrently.

The same trend can be seen in the converge statistics in Fig. 11(a) and (b). We define as the “min” the iteration in which the first node successfully converges and the “max” the iteration where the last node converges. Note, that mean is the “average” converge iteration for all nodes and the converge “window” is the difference between the “max” and “min”. As we can see from Fig. 11(a) and (b) the window size *decreases* as the network size grows. In the presence of low delays (Fig. 11(a)) the window is practically zero indicating that the “min” and “max” converge iteration coincides.

Practically speaking, this indicates that the converge variability is low in large networks and is expected to converge in few iterations. This means that tasks can be scheduled in a timely fashion and with optimal placement for the given set of jobs. This is highly important for any modern data center scheduler aiming to schedule thousands of jobs at-a-time on thousands of nodes in a timely fashion.

VIII. DISCUSSION

In this paper, we proposed a finite-time asynchronous algorithm for distributively computing a value which a network of nodes can use to make local control decisions. Contrary to

prior work, our approach is able to operate asynchronously and, as a consequence, also able to handle delays by construction. To our knowledge this is the first proposed algorithm able to provide finite-time guarantees in the combined delay tolerant and asynchronous setting.

The proposed scheme uses the industry standard CPU utilization model and is able to balance the workload allocation such that each node is allocated tasks proportional to its capabilities. Concretely, this model defines that the utilization of each CPU core is measured in the bounded range of $[0, 100]$ and indicates the utilization percentage for each individual core within a specified machine [55]. This effectively allows us to evenly distribute to load across all of the available network nodes loading to better overall cluster utilization. Practically speaking, it is standard practice in data centers to share the load across the available nodes. We emphasize that our algorithm algorithm is able to handle both regular workloads as well as bursty ones. This is achieved because our optimization algorithm works in buckets; where each bucket is filled with incoming jobs. At the time of scheduling each job within the bucket is attempted to be placed into a suitable node, while guaranteeing the balancing of the overall cluster load. The time of scheduling is fixed to be at regular intervals or if the bucket is filled. This behaviour is beneficial for a number of reasons: firstly, bursty workloads are able to be handled gracefully, and secondly, even if there are not enough jobs within the bucket they will still be scheduled in a timely manner. Note, that our experiments are designed to reflect practical data center deployments which implies that the network graphs considered will be of low diameter and have good connectivity.

Interestingly, as per Algorithm 2 and a corollary of Theorem 1 the convergence rate is only bounded by the network diameter and its maximum delay. More importantly, our particular setting implies that packet loss is assumed to be minimal in such deployments but not *delays*. The delays can be attributed to processing and communication delays. Experiencing processing delays is common in data centers and in the presence of over-provisioned or straggler nodes. Communication delays are mainly because of re-transmissions due to packet losses. However, packet losses are not so common and, for this reason, we do not consider them in this work. Nevertheless, in case one wishes to consider packet losses as well, this can be achieved by establishing probabilistic guarantees for convergence based on the packet loss distribution. However, that is beyond the context of this work and is left for future work.

We note that our scheme is asynchronous but in order to successfully operate it implies that the internal clocks of all nodes are paced similarly. This requirement is necessitated as each node needs to be able to recognize when the appropriate iterations have elapsed. As noted previously these checks happen every $(1 + \hat{\tau})D$ iterations. Consistent pacing of each node's clock ensures that the check for convergence at each node will happen at roughly the same time [56]. However, this does not imply that we actually need to synchronize each of the nodes' time-zones nor their actual clocks but, rather, their internal clocks must have similar pacing [57]. This is especially the case after the introduction of High Precision Event Timers in the




Fig. 12. Violation in a network consisting of 20 nodes with a diameter equal to $D = 5$ when using a delay of $\tau = 20$. Indicatively, circles indicate violations of the claim in [13, Lemma 3.2].

low-level firmware of most commodity computers and servers alike [58], [59]. Notably, this is common practice and present in most modern computers as the clock pacing specification is defined within the Advanced Configuration and Power Interface (ACPI) specifications [60], [61]. In fact, to address these issues in a standardized way hardware manufacturers created the Unified Extensible Firmware Interface (UEFI) forum [62], which is responsible for defining the characteristics and functionality regarding the most basic, low-level functions that each modern computer or server should support.

As aforementioned in Remark 2, a similar approach was proposed in [13] in the context of average consensus with bounded time-varying delays. Apart from the differences in the application and the fact that we consider asynchronous operation of the nodes, the approach is similar. However, for proving convergence of their proposed algorithm they claim a form of monotonicity of the maximum and minimum values of the states. Specifically, it is claimed [13, Lemma 3.2] that if the value held by an agent v_i at the present instant of time is strictly lesser (greater) than the maximum (minimum) over the current and delayed values over a horizon $\bar{\tau}$ of all the nodal states, then, the value of agent v_i continues to be strictly lesser (greater) than this maximum (minimum) for all future instants. Notably, we found several examples of networks for which that statement is not valid. Practical examples of networks that exhibit such violations are presented in Figures 12 and 13.

Concretely, in Fig. 12 we present a violation that happens in a network comprising of 20 nodes with a diameter $D = 5$ and a delay $\tau = 20$. Interestingly, as we can observe in Fig. 13 this violation is also observed when dealing with larger networks. In this particular example presented below the issue is manifested in a network of 50 nodes with a diameter of $D = 4$ and a delay of $\tau = 20$.

Throughout our experiments we observed this behaviour to be more frequent with medium sized networks that had delays greater than $\tau = 5$. On the other hand, the diameter seems to be not a major contributing factor; at least for the values considered in our experiments (e.g., D between 1 and 10).

Our solution is able to gracefully handle this situation and still converge into the optimal solution. The effectiveness of




Fig. 13. Another example of a violation using a larger network consisting of 50 nodes with a diameter equal to $D = 4$ when using a delay of $\tau = 10$. As in the previous figure, circles indicate violations of the claim in [13, Lemma 3.2].

our asynchronous finite-time algorithm was demonstrated on CPU resource allocation in data centers, which can result in better overall system utilization. However, one important aspect of such approaches, including our own, is the way they compare against more complex optimization problems. In particular against ones that do not have a closed form solution and require complex solvers to be approximated such as ADMM [18]. As formulated, our problem is able to tackle placement of jobs using the most commonly used CPU utilization model in practical deployments. Furthermore, due to its problem formulation the problem admits a closed-form solution. This enables our method to reach the optimization objective significantly faster when compared to more sophisticated solvers such as ADMM; especially as the network sizes scale [17]. Other approaches have been proposed as well for the same problem formulation [63], but the termination of the optimization cannot be synchronized and re-initiating the optimization with the new requests is not possible. More importantly, we note that our proposed method could also be exploited across multiple domains where asynchronous distributed coordination is desirable (e.g., distributed frequency regulation in microgrids, decentralized computation networks, and voltage control in distribution systems).

IX. CONCLUSION AND FUTURE DIRECTIONS

A. Conclusion

In this paper, we proposed a finite-time asynchronous algorithm for distributively computing a value which a network of nodes can use to make local control decisions. Contrary to previously-proposed algorithms, our approach works also asynchronously. We evaluated our proposed solution using networks of varying delays and diameters which reflected practical data center installations as per common deployment guidelines. The effectiveness of our asynchronous finite-time algorithm was evaluated against the CPU resource allocation in data centers. In turn, more efficient allocation of resources can lead to better overall system responsiveness and utilization.

B. Future Directions

Our work can be easily extended to more general convex optimization problems, using gradient-consensus methods, as

in [45], but our solution will allow for asynchronous operation and will be able to tolerate delays.

Part of ongoing research focuses on considering deadline constraints and cases for which the workloads exceed the available resources and exploit the heterogeneity of resource units available (e.g. CPU, GPU, and/or accelerators). In such instances a more sophisticated rejection policy can take place based on inferred resource demands, task priorities, or introduce partial scheduling plans based on either priorities or further, more complex, constraints.

REFERENCES

- [1] L. A. Barroso, U. Hölzle, and P. Ranganathan, "The datacenter as a computer: Designing warehouse-scale machines, 3rd Edition," *Synth. Lectures Comput. Archit.*, San Rafael, CA, USA: Morgan & Claypool Publishers., vol. 13, no. 3, 2018, pp. 1–189.
- [2] E. Cortez, A. Bonde, A. Muzio, M. Russinovich, M. Fontoura, and R. Bianchini, "Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms," in *Proc. 26th Symp. Operating Syst. Princ.*, 2017, pp. 153–167.
- [3] G. Amvrosiadis, J. W. Park, G. R. Ganger, G. A. Gibson, E. Baseman, and N. DeBardeleben, "On the diversity of cluster workloads and its impact on research results," in *Proc. USENIX Annu. Tech. Conf.*, 2018, pp. 533–546.
- [4] E. Boutin et al., "Apollo: Scalable and coordinated scheduling for cloud-scale computing," in *Proc. 11th USENIX Symp. Operating Syst. Des. Implementation*, 2014, pp. 285–300.
- [5] P. Moritz et al., "Ray: A distributed framework for emerging AI applications," in *Proc. 13th USENIX Symp. Operating Syst. Des. Implementation*, 2018, pp. 561–577.
- [6] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at Google with Borg," in *Proc. 10th Eur. Conf. Comput. Syst.*, 2015, pp. 1–17.
- [7] B. Hindman et al., "Mesos: A platform for fine-grained resource sharing in the data center," in *Proc. USENIX Symp. Netw. Syst. Des. Implementation*, 2011, pp. 295–308.
- [8] K. Ousterhout, P. Wendell, M. Zaharia, and I. Stoica, "Sparrow: Distributed, low latency scheduling," in *Proc. ACM Symp. Operating Syst. Princ.*, 2013, pp. 69–84.
- [9] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.
- [10] A. Giridhar and P. R. Kumar, "Toward a theory of in-network computation in wireless sensor networks," *IEEE Commun. Mag.*, vol. 44, no. 4, pp. 98–107, Apr. 2006.
- [11] P. A. Misra, M. F. Borge, I. Goiri, A. R. Lebeck, W. Zwaenepoel, and R. Bianchini, "Managing tail latency in datacenter-scale file systems under production constraints," in *Proc. EuroSys Conf.*, 2019, pp. 1–15.
- [12] S. T. Cady, A. D. Domínguez-García, and C. N. Hadjicostis, "Finite-time approximate consensus and its application to distributed frequency regulation in islanded AC microgrids," in *Proc. IEEE 48th Hawaii Int. Conf. Syst. Sci.*, 2015, pp. 2664–2670.
- [13] M. Prakash, S. Talukdar, S. Attree, V. Yadav, and M. V. Salapaka, "Distributed stopping criterion for consensus in the presence of delays," *IEEE Trans. Control Netw. Syst.*, vol. 7, no. 1, pp. 85–95, Mar. 2020.
- [14] C. N. Hadjicostis and T. Charalambous, "Asynchronous coordination of distributed energy resources for the provisioning of ancillary services," in *Proc. IEEE Annu. Allerton Conf. Commun., Control, Comput.*, 2011, pp. 1500–1507.
- [15] C. N. Hadjicostis and T. Charalambous, "Average consensus in the presence of delays in directed graph topologies," *IEEE Trans. Autom. Control*, vol. 59, no. 3, pp. 763–768, Mar. 2014.
- [16] S. Giannini, D. D. Paola, A. Pettiti, and A. Rizzo, "On the convergence of the max-consensus protocol with asynchronous updates," in *Proc. IEEE Conf. Decis. Control*, 2013, pp. 2605–2610.
- [17] T. H. Chang, W. C. Liao, M. Hong, and X. Wang, "Asynchronous distributed ADMM for large-scale optimization x2014;Part II: Linear convergence analysis and numerical performance," *IEEE Trans. Signal Process.*, vol. 64, no. 12, pp. 3131–3144, Jun. 2016.
- [18] W. Jiang, A. Grammenos, E. Kalyvianaki, and T. Charalambous, "An asynchronous approximate distributed alternating direction method of multipliers in digraphs," in *Proc. 60th IEEE Conf. Decis. Control*, 2021, pp. 3406–3413.

- [19] M. Isard, V. Prabhakaran, J. Currey, U. Wieder, K. Talwar, and A. Goldberg, "Quincy: Fair scheduling for distributed computing clusters," in *Proc. ACM SIGOPS 22nd Symp. Operating Syst. Princ.*, 2009, pp. 261–276.
- [20] M. Gog, M. Schwarzkopf, A. Gleave, R. N. M. Watson, and S. Hand, "Firmament: Fast, centralized cluster scheduling at scale," in *Proc. Operating Syst. Des. Implementation*, 2016, pp. 99–115.
- [21] H. Mao, M. Schwarzkopf, S. B. Venkatakrisnan, Z. Meng, and M. Alizadeh, "Learning scheduling algorithms for data processing clusters," in *Proc. ACM Special Int. Group Data Commun. (SIGCOMM Series)*, Beijing, China, 2019, pp. 270–288.
- [22] A. Tumanov, T. Zhu, J. W. Park, M. A. Kozuch, M. Harchol-Balter, and G. R. Ganger, "TetriSched: Global rescheduling with adaptive plan-ahead in dynamic heterogeneous clusters," in *Proc. 11th Eur. Conf. Comput. Syst.*, (EuroSys Series), 2016, pp. 1–16.
- [23] R. Grandl, G. Ananthanarayanan, S. Kandula, S. Rao, and A. Akella, "Multi-resource packing for cluster schedulers," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 455–466, 2014.
- [24] Y. Wang and X. Zuo, "An effective cloud workflow scheduling approach combining PSO and idle time slot-aware rules," *IEEE/CAA J. Automatica Sinica*, vol. 8, no. 5, pp. 1079–1094, May 2021.
- [25] B. Burns et al., "Borg, Omega, and Kubernetes: Lessons learned from three container-management systems over a decade," *Queue*, vol. 14, no. 1, pp. 70–93, 2016.
- [26] H. Yuan, J. Bi, M. Zhou, Q. Liu, and A. C. Ammari, "Biobjective task scheduling for distributed green data centers," *IEEE Trans. Autom. Sci. Eng.*, vol. 18, no. 2, pp. 731–742, Apr. 2021.
- [27] H. Yuan, M. Zhou, Q. Liu, and A. Abusorrah, "Fine-grained resource provisioning and task scheduling for heterogeneous applications in distributed green clouds," *IEEE/CAA J. Automatica Sinica*, vol. 7, no. 5, pp. 1380–1393, Sep. 2020.
- [28] H. Yuan, J. Bi, and M. Zhou, "Geography-aware task scheduling for profit maximization in distributed green data centers," *IEEE Trans. Cloud Comput.*, vol. 10, no. 3, pp. 1864–1874, Jul.–Sep. 2022.
- [29] M. Schwarzkopf, A. Konwinski, M. Abd-El-Malek, and J. Wilkes, "Omega: Flexible, scalable schedulers for large compute clusters," in *Proc. 8th ACM Eur. Conf. Comput. Syst.*, 2013, pp. 351–364.
- [30] V. K. Vavilapalli et al., "Apache hadoop YARN: Yet another resource negotiator," in *Proc. 4th Annu. Symp. Cloud Comput.*, 2013, pp. 1–16.
- [31] K. Rzadca et al., "Autopilot: Workload autoscaling at Google," in *Proc. Fifteenth Eur. Conf. Comput. Syst.*, 2020, pp. 1–16.
- [32] Y. Gan et al., "An open-source benchmark suite for microservices and their hardware-software implications for cloud & edge systems," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 3–18.
- [33] M. Jeon, S. Venkataraman, A. Phanishayee, J. Qian, W. Xiao, and F. Yang, "Analysis of large-scale multi-tenant GPU clusters for DNN training workloads," in *Proc. USENIX Annu. Tech. Conf.*, 2019, pp. 947–960.
- [34] R. Grandl, S. Kandula, S. Rao, A. Akella, and J. Kulkarni, "GRAPHENE: Packing and dependency-aware scheduling for data-parallel clusters," in *Proc. USENIX Symp. Operating Syst. Des. Implementation*, Savannah, GA, 2016, pp. 81–97.
- [35] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. 31, no. 9, pp. 803–812, Sep. 1986.
- [36] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.
- [37] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, May 2005.
- [38] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 169–182, Feb. 2005.
- [39] D. Angeli and P.-A. Bliman, "Stability of leaderless discrete-time multi-agent systems," *Math. Control, Signals, Syst.*, vol. 18, no. 4, pp. 293–322, 2006.
- [40] P.-A. Bliman and G. Ferrari-Trecate, "Average consensus problems in networks of agents with delayed communications," *Automatica*, vol. 44, no. 8, pp. 1985–1995, Aug. 2008.
- [41] K. Cai and H. Ishii, "Average consensus on general digraphs," in *Proc. IEEE Conf. Decis. Control Eur. Control Conf.*, 2011, pp. 1956–1961.
- [42] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [43] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, Sep. 2004.
- [44] T. Charalambous, Y. Yuan, T. Yang, W. Pan, C. N. Hadjicostis, and M. Johansson, "Distributed finite-time average consensus in digraphs in the presence of time delays," *IEEE Trans. Control Netw. Syst.*, vol. 2, no. 4, pp. 370–381, Dec. 2015.
- [45] V. Khatana, G. Saraswat, S. Patel, and M. V. Salapaka, "Gradient-consensus method for distributed optimization in directed multi-agent networks," in *Proc. IEEE Amer. Control Conf.*, 2020, pp. 4689–4694.
- [46] D. Bernstein, "Containers and cloud: From lxc to docker to kubernetes," *IEEE Cloud Comput.*, vol. 1, no. 3, pp. 81–84, Sep. 2014.
- [47] A. D. Domínguez-García and C. N. Hadjicostis, "Coordination and control of distributed energy resources for provision of ancillary services," in *Proc. IEEE First Int. Conf. Smart Grid Commun.*, 2010, pp. 537–542.
- [48] J. Cortés, "Distributed algorithms for reaching consensus on general functions," *Automatica*, vol. 44, no. 3, pp. 726–737, Mar. 2008.
- [49] C. N. Hadjicostis, A. D. Domínguez-García, and N. H. Vaidya, "Resilient average consensus in the presence of heterogeneous packet dropping links," in *Proc. IEEE Conf. Decis. Control*, 2012, pp. 106–111.
- [50] C. N. Hadjicostis, N. H. Vaidya, and A. D. Domínguez-García, "Robust distributed average consensus via exchange of running sums," *IEEE Trans. Autom. Control*, vol. 61, no. 6, pp. 1492–1507, Jun. 2016.
- [51] A. Singla, P. B. Godfrey, and A. Kolla, "High throughput data center topology design," in *Proc. 11th USENIX Symp. Networked Syst. Des. Implementation*, 2014, pp. 29–41.
- [52] C. Guo et al., "Pingmesh: A large-scale system for data center network latency measurement and analysis," in *Proc. ACM Conf. Special Int. Group Data Commun.*, 2015, pp. 139–152.
- [53] M. Alizadeh, A. Kabbani, T. Edsall, B. Prabhakar, A. Vahdat, and M. Yasuda, "Less is more: Trading a little bandwidth for ultra-low latency in the data center," in *Proc. 9th USENIX Symp. Netw. Syst. Des. Implementation*, 2012, pp. 253–266.
- [54] L. Popa, S. Ratnasamy, G. Iannaccone, A. Krishnamurthy, and I. Stoica, "A cost comparison of datacenter network architectures," in *Proc. 6th Int. Conf. Emerg. Netw. Experiments And Technol.*, 2010, pp. 1–12.
- [55] L. VMware, "Performance Troubleshooting CPU Ready Time," VMware, Oct. 2018. [Online]. Available: <https://learnvmware.online/2018/03/08/performance-troubleshooting-cpu-ready-time/>
- [56] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," in *Proc. Concurrency: Works Leslie Lamport*, 2019, pp. 179–196.
- [57] M. Nyström, M. Nicholes, and V. J. Zimmer, "UEFI networking and pre-os security," *Intel Technol. J.*, vol. 15, no. 1, pp. 80–100, 2011.
- [58] J. Ridoux, D. Veitch, and T. Broomhead, "The case for feed-forward clock synchronization," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, pp. 231–242, Feb. 2012.
- [59] P. Orosz and T. Skopko, "Performance, evaluation of a high precision software-based timestamping solution for network monitoring," *Int. J. Adv. Soft.*, vol. 4, no. 1, pp. 181–188, 2011.
- [60] U. Forum, "Advanced configuration and power interface (ACPI) specification—ACPI specification 6.4 documentation," 2021. [Online]. Available: https://uefi.org/htmlspecs/ACPI_Spec_6_4_html/
- [61] V. Zimmer, M. Rothman, and S. Marisetty, *Beyond BIOS: Developing With the Unified Extensible Firmware Interface*. Berlin, Germany: Walter de Gruyter GmbH & Co KG, 2017.
- [62] "Unified extensible firmware interface specifications," publication Title: Unified Extensible Firmware Interface Forum, 2021. [Online]. Available: <https://uefi.org/>
- [63] A. I. Rikos, A. Grammenos, E. Kalyvianaki, C. N. Hadjicostis, T. Charalambous, and K. H. Johansson, "Optimal CPU scheduling in data centers via a finite-time distributed quantized coordination mechanism," in *Proc. IEEE Conf. Decis. Control*, 2021, pp. 6276–6281.



Andreas Grammenos received the B.Sc. degree from the Technical University of Crete, Kounoupidiana, Greece, and the Ph.D. degree from the Department of Computer Science and Technology, University of Cambridge, Cambridge, U.K., where he was also affiliated with the Turing Institute. His research interests include large scale Data and streaming analytics with a focus on topics around distributed, federated, & decentralized systems. He was a Founding Member of COVID-19 Sounds, an audio based study which builds predictive models for COVID-19 detection using crowdsourced smartphone respiratory recordings.



Themistoklis Charalambous (Senior Member, IEEE) received the B.A. and M.Eng. degrees in electrical and information sciences from Cambridge University, Cambridge, U.K., and the Ph.D. degree from the Control Laboratory of the Engineering Department, Cambridge University. Following his Ph.D., he was a Postdoctoral Researcher with the Human Robotics Group, Imperial College London, London, U.K., the Department of Automatic Control of the School of Electrical Engineering, KTH the Royal Institute of Technology, Stockholm, Sweden, and with

the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden. In January 2017, he joined the Department of Electrical Engineering and Automation, School of Electrical Engineering, Aalto University, Espoo, Finland, as a tenure-track Assistant Professor. In September 2018, he was awarded the Academy of Finland Research Fellowship and in July 2020 he was appointed as a tenured Associate Professor. In September 2021, he joined the Department of Electrical and Computer Engineering, University of Cyprus as a tenure-track Assistant Professor and he remains associated with Aalto University as a Visiting Professor. His primary research targets the design and analysis of (wireless) networked control systems that are stable, scalable, and energy-efficient. The study of such systems involves the interaction between dynamical systems, their communication, and the integration of these concepts. As a result, his research is interdisciplinary combining theory and applications from control theory, communications, network, and distributed optimization.



Evangelia Kalyvianaki received the B.Sc. the M.Sc. degrees from the Computer Science Department, University of Crete, Rethymno, Greece, and the Ph.D. degree from the Computer Laboratory, Cambridge University, Cambridge, U.K. She is currently a Senior Lecturer (Assistant/associate professor) with the Department of Computer Science and Technology, University of Cambridge and a Turing Fellow. Before, she was a Lecturer with the Department of Computer Science, City University London, London, U.K., and a Postdoctoral Researcher with the Department of

Computing, Imperial College London, London. Her research interests include cloud computing, Big Data processing, autonomic computing, distributed systems and systems research in general. She is interested in the design and management of next generation, large-scale applications in the cloud.