

Selective Edge Computing for Mobile Analytics

Apostolos Galanopoulos*, George Iosifidis†,
Theodoros Salonidis‡, Douglas J. Leith*

*School of Computer Science and Statistics, Trinity College Dublin

†Delft University of Technology, The Netherlands

‡IBM T. J. Watson Research Center

Abstract—An increasing number of mobile applications rely on Machine Learning (ML) routines for analyzing data. Executing such tasks at the user devices saves the energy spent on transmitting and processing large data volumes at distant cloud-deployed servers. However, due to memory and computing limitations, the devices often cannot support the required resource-intensive routines and fail to accurately execute such tasks. In this work, we address the problem of edge-assisted analytics in resource-constrained systems by proposing and evaluating a rigorous selective offloading framework. The devices execute their tasks locally and outsource them to cloudlet servers only when they predict a significant performance improvement. We consider the practical scenario where the offloading gains and resource costs are time-varying; and propose an online optimization algorithm that maximizes the service performance without requiring to know this information. Our approach relies on an approximate dual subgradient method combined with a primal-averaging scheme, and works under minimal assumptions about the system stochasticity. We fully implement the proposed algorithm in a wireless testbed and evaluate its performance using a state-of-the-art image recognition application, finding significant performance gains and cost savings.

Index Terms—Edge Computing, Data Analytics, Network Optimization, Resource Allocation, Subgradient method

I. INTRODUCTION

A. Background and Motivation

The recent demand for mobile machine learning (ML) analytic applications, such as image recognition, natural language translation and health monitoring, has been unprecedented [2]. These services collect data streams generated by hand-held or other Internet of Things (IoT) devices, and analyze them locally or at cloud servers. The challenge with such services is that they are both resource intensive and delay sensitive. On the one hand, the cloud offers powerful ML models and abundant compute resources, but requires data transfers which consume network bandwidth and device power, as well as induce significant delays, e.g., due to intermittent connectivity [3]. On the other hand, executing these services directly at the devices, as in [4], economizes network bandwidth but degrades their performance due to the devices' limited resources. For example, these nodes may have insufficient memory to support accurate deep-learning neural networks.

A promising approach to tackle this problem is to follow a middle-ground solution where the devices outsource their

tasks to nearby cloudlets [5]. These edge servers are typically deployed in locations close to cellular base stations or Wi-Fi access points, and hence are in proximity with the users. Therefore, they can increase the service performance by *augmenting* the devices' ML components with more accurate models, while offering tolerable communication and execution delay. Nevertheless, the success of such solutions requires intelligent decision algorithms for selecting which tasks from each device will be outsourced in order to maximize the aggregate accuracy. This is a new problem that raises intricate challenges for the network and the involved computing nodes.

Namely, the cloudlets, unlike the cloud, have limited computing capacity and hence cannot support the requests from all devices. If overloaded, they will eventually become unresponsive. At the same time, task execution often involves the transfer of large data volumes. This calls for prudent transmission decisions in order to avoid wasting the energy of devices and congesting the network when link bandwidth is also a bottleneck. Furthermore, unlike general computation offloading solutions [6], in ML analytics it is imperative to identify and outsource only the tasks which can significantly benefit from cloudlet execution. Otherwise, the system will spend resources only to gain marginal performance improvements. Finally, these decisions need to be made in a dynamic fashion accounting for the time-varying network conditions, user requests and cloudlet availability; while the statistical properties of these random parameters are unknown in practice.

Our goal is to design and evaluate an online decision framework that supports edge-augmented mobile analytic services. While prior works have studied the problem of offloading computation-intensive tasks and others proposed system architectures for mobile analytics, see. Sec. VII, we lack an analytical framework for maximizing the performance of such services under resource (un)availability, and time-varying network conditions. Our solution works under such practical limitations (which we measure experimentally) and is general enough to be applied to different architectures and services.

B. Methodology and Contributions

In detail, we consider a service where a cloudlet improves upon request the execution quality of data analytic tasks that are generated by small user devices. We use as an exemplary service the processing of image frames captured by nodes

Part of this work appeared in the Proceedings of IEEE ICC 2020 [1]. This publication has emanated from research conducted with the financial support of the European Commission through Grant No. 101017109 (DAEMON)

such as wireless IoT cameras or small robots, that need to be processed for classifying objects of interest. Each device has a low-precision classifier, while the cloudlet can possibly execute the task with higher precision. The devices classify the received objects upon arrival and decide whether to transmit them to the cloudlet for further processing. This decision requires an assessment of the potential performance gains, which are measured in terms of accuracy improvement. To this end, we propose the usage of a *predictor* that is installed at each device and leverages the local classification results.

In terms of resource constraints, we focus on power consumption, a bottleneck issue in small devices; and the computing capacity of the cloudlet which - unlike the cloud - is finite. The former couples the decisions of each device across time, while the latter ties the decisions of all devices sharing the cloudlet. We consider the practical case where resource availability is unknown and possibly time-varying and we observe their instantaneous values. We aim to design an algorithm that enables the coordination of devices and dictates the task outsourcing policy by carefully tuning the trade-off between maximizing the aggregate analytics accuracy and constraining resource consumption.

We formulate the system operation as an optimization program with unknown parameters appearing both in the objective (performance gains) and constraints (power and capacity), which are learned in an online fashion. This program is decomposed via Lagrange relaxation to device-specific problems and this enables its distributed solution through an *approximate* - due to the unknown parameters - dual ascent method. Leveraging the ϵ -(sub)gradient information that is produced in the dual space by each device, we calculate primal solutions which are applied in real time. Our approach is inspired by primal averaging schemes for *static* problems [7], [8], and yields a *tunable* optimality bound compared to the hypothetical benchmark policy that has access to an oracle. The designed algorithm is lightweight in terms of communication overheads and adapts to resource availability and user requests. Importantly, it offers deterministic performance bounds (i.e., for each sample path) and works under minimal assumptions for the stochastic perturbations of the resources and task requests. This is in contrast with extensively-used stochastic optimization toolboxes which presume i.i.d. or Markov-modulated perturbations and offer only average guarantees; see [9] and references therein.

Finally, the framework is extended for when the bottleneck is the wireless link capacity and for services that optimize jointly the accuracy and execution delay. Other scenario-specific amendments are also possible, e.g., considering multi-stage services, multiple cloudlets, or other related constraints such as the cloudlet energy budget. Given that this is a new problem, we investigate experimentally its properties in a wireless testbed; and assess our algorithm using real datasets [10], [11] and carefully selected benchmarks. Hence, the contributions of this work are the following:

- *Edge-Augmented Analytics*. We introduce the novel problem of augmenting the performance of mobile analytics using edge infrastructure (e.g., cloudlets), which is increasingly important for mobile computing services and IoT networks.

Our model can be tailored to different system architectures, types of analytic services, and resource constraints.

- *Decision Framework*. A task outsourcing policy is proposed that achieves near-optimal performance while being oblivious to the system’s statistics. We fully characterize the performance of the algorithm, i.e., its optimality gap, as a function of the system parameters, perturbations and the employed step rule. To the best of our knowledge, our algorithm is the first to offer deterministic performance bounds with discrete actions, under such general conditions; and this is a result of independent interest.
- *Implementation & Evaluation*. The solution is evaluated in a wireless testbed with a typical ML service and real datasets. We show that our algorithm can be implemented as a lightweight protocol, increasing task accuracy (up to 15%) and reducing the energy costs (down to 50%) compared to carefully selected greedy benchmarks.

Concluding, this work proposes a new problem, designs a novel optimization algorithm which is tailored to its needs, and uses a fully-fledged implementation in a wireless testbed in order to evaluate the proposal.

Organization. Sec. II introduces the model and problem, Sec. III presents the algorithm and Sec. IV analyzes its performance. We discuss practical extensions in Sec. V and Sec. VI presents the system implementation, a series of experiments and trace-driven simulations. We discuss related work in Sec. VII and conclude in Sec. VIII. The proofs of the various lemmas can be found in the Appendix, Sec. IX.

II. MODEL AND PROBLEM FORMULATION

We introduce the system model, the problem and the respective mathematical program. Table I summarizes the key notation we use throughout the paper. We use calligraphic capital letters for sets, bold typeface letters for vectors, and $\|\cdot\|$ denotes the Euclidean norm.

A. Task Model

Time is slotted and we index the slots. There is a set \mathcal{C} of C disjoint object classes and a set \mathcal{N} of N devices. Each device n may receive at slot t an object $s_{nt} \in \mathcal{S}$ for classification, where \mathcal{S} is the set of possible objects, e.g., images captured by its camera. In case a device n does not produce an image in slot t (no task), we set $s_{nt} = \emptyset$. Every device n is equipped with a *local* classifier J_n which outputs the inferred class $J_n(s_{nt}) \in \mathcal{C}$ of object s_{nt} and a normalized *accuracy* (or, confidence) value $d_n(s_{nt}) \in [0, 1]$ for that inference¹. There is also a classifier J_0 at the cloudlet which can classify any object $s_{nt} \in \mathcal{S}$ with confidence $d_0(s_{nt})$. The local classifiers may have different performance, e.g., due to possibly different ML components or training datasets, while the cloudlet classifier has the highest accuracy, i.e., $d_0(s_{nt}) \geq d_n(s_{nt}), \forall s_{nt} \in \mathcal{S}$. Parameter $\phi_{nt} \in [0, 1]$ denotes the accuracy improvement when the cloudlet classifier is used:

$$\phi_{nt}(s_{nt}) = d_0(s_{nt}) - d_n(s_{nt}), \quad \forall s_{nt} \in \mathcal{S}.$$

¹The classifier might output only the class with the highest confidence or a vector with the confidence for each class and allow the user to decide - typically selecting the more likely class. Our analysis works for both cases.

It is worth stressing that several services, e.g., see YOLO [12], provide in real-time feedback on the confidence about the accuracy of inferences, without requiring labeled data, which exhibit indeed strong correlation with the actual accuracy [13]. Finally, every device is also equipped with a predictor² that is trained with the outcomes of the local and cloudlet classifiers. This predictor can estimate the cloudlet's improvement $\tilde{\phi}_n(s_{nt})$ for each object s_{nt} , where this assessment might not be exact, i.e., $\tilde{\phi}_n(s_{nt}) \neq \phi_{nt}(s_{nt})$; and we denote with $\sigma_{nt}(s_{nt}) \in [0, 1]$ the normalized predictor confidence.

B. Wireless System

The devices access the cloudlet through high-capacity cellular or Wi-Fi links, see Fig. 1, that do not impose data transfer constraints (we relax this assumption in Sec. V). Each device n has an *average power budget* of B_n Watts that it can spend on transmitting the images to cloudlet.³ Average power consumption is a key limitation in such systems [15], because: the devices might have a small energy budget to spend; their small form-factor imposes power consumption limitations; there are protocol-induced transmission constraints; or users might impose constraints on the power cost of this service. Similarly, the cloudlet has an *average processing capacity* of H cycles/sec. This resource is shared by all devices and when the total load exceeds H the task delay increases fast, eventually rendering the system non-responsive.

When an image is transmitted in slot t from device n to cloudlet, it consumes o_{nt} Watts of the device's power budget. This quantity might change across slots due to channel conditions variations, shadowing effects, interference from other transmission, and so on; and follows a random process $\{o_{nt}\}_{t=1}^{\infty}$, where $o_{nt} \in \mathcal{O} = \{o_1, \dots, o_{|\mathcal{O}|}\}$ is drawn from a set of possible values. Also, each transmitted image s_{nt} requires a number of cloudlet processing cycles h_{nt} , which might vary with time, e.g., due to different image sizes, and possibly stems from a random process $\{h_{nt}\}_{t=1}^{\infty}$, with $h_{nt} \in \mathcal{H} = \{h_1, \dots, h_{|\mathcal{H}|}\}$. We also define $\mathbf{o}_t = (o_{nt}, n \in \mathcal{N})$ and $\mathbf{h}_t = (h_{nt} \leq, n \in \mathcal{N})$. Our model is general as the requests, power and computing costs per request can be arbitrarily time-varying and with unknown statistics.

The devices wish to involve the cloudlet only when they expect high classification precision gain with high confidence. When the cloudlet does not offer high-enough gains or, even worse, lower accuracy, the devices need to refrain from offloading their tasks. Otherwise, they risk consuming the cloudlet's capacity and their own power without significant benefits. Therefore, the outsourcing decision for each object s_{nt} is based on the *weighted improvement gain*⁴:

$$w_{nt}(s_{nt}) = \tilde{\phi}(s_{nt}) - v_n \sigma_{nt}(s_{nt}), \quad (1)$$

²This can be a model-based or model-free solution, e.g., a regressor or a neural-network; our analysis and framework work for any of these solutions. In the implementation we used a mixed-effects regressor [14].

³Local classifications can induce non-negligible energy costs to devices but these are not considered for B_n since every object undergoes local classification anyway.

⁴Whenever the cloudlet has lower expected accuracy from the device, then we set $w_{nt} = 0$, and decide not to offload.

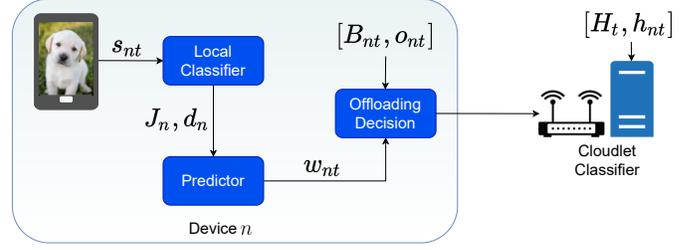


Fig. 1: System model including the local/cloudlet classifiers and predictors. Each device is constrained by its average power budget, and the cloudlet has a limited computation capacity.

where $v_n \geq 0$ a *risk aversion* parameter set by the system designer or each user. For example, assuming normal distribution for the outputs of the predictor we could set $v_n = 1$ and use a threshold rule of 1 standard deviation. The improvement gains follow an unknown random process $\{w_{nt}\}_{t=1}^{\infty}$, where⁵ $w_{nt} \in \mathcal{W} = \{w_1, \dots, w_{|\mathcal{W}|}\}$.

C. Problem Definition and Assumptions

Our goal is to maximize the long-term accuracy improvement gains for all devices while satisfying the average capacity constraints. Let us first define the set of possible system states

$$\mathcal{J} = \mathcal{O}^N \times \mathcal{H}^N \times \mathcal{W}^N,$$

and introduce parameter $\pi_t \in \mathcal{J}$ that indicates the system state at slot t . We assume the system operation can be described by the stationary probability distribution $\rho = (\rho^j, j = 1, \dots, M)$, where $M = |\mathcal{J}|$. We introduce variables $y_n^j \in [0, 1], \forall n \in \mathcal{N}, j \in \mathcal{J}$ that indicate the outsourcing probability of objects from each device n when the system is in state j . We also define the vector $\mathbf{y} = (y_n^j : n = 1, \dots, N, j = 1, \dots, M)$ and the set $\mathcal{Y} = [0, 1]^{NM}$. Henceforth we use superscript j to indicate the values of the random variables when $\pi_t = j$.

Putting the above together, our overall (i.e., long-term) optimization goal can be expressed with the following program:

$$\mathbb{P}_1 : \quad \underset{\mathbf{y} \in \mathcal{Y}}{\text{maximize}} \quad \sum_{j=1}^M \sum_{n=1}^N w_n^j y_n^j \rho^j \quad (2)$$

$$\text{s.t.} \quad \sum_{j=1}^M y_n^j o_n^j \rho^j \leq B_n, n \in \mathcal{N}, \quad (3)$$

$$\sum_{j=1}^M \sum_{n=1}^N y_n^j h_n^j \rho^j \leq H. \quad (4)$$

Constraints (3) impose the average power budget⁶ of each device and (4) bounds the cloudlet utilization. Additional constraints can be included if needed; and we can also replace the linear objective with any other convex function. For instance,

⁵We note that most systems use such quantized values for the prediction gains, and the number of possible values depends on the granularity.

⁶To capture the total power consumption we should add a term related to the computation energy cost at the LHS of (3). However, this term is independent of the decision variable y_n^j , since the local classifier is used either way and thus it is omitted.

Description	Parameter / Variable
Classification confidence of n (cloudlet)	d_n (d_0)
Actual (predicted) offloading improvement	ϕ_{nt} (Q_n)
Average power (computing) constraint	B_n (H)
Power (Computing) resource consumption of task s_{nt} at state j	σ_n^j (h_n^j)
Improvement gain for device n at slot t (quantized value at state j)	w_{nt} (w_n^j)
Probability of system being in state j	ρ^j
Outsourcing probability for task in state j	$y_n^j \in [0, 1]$

TABLE I: Key Parameters and Variables.

we might wish to enforce a fairness criterion by using α -fair functions [16] or an objective that maximizes accuracy while minimizing the total delay. We elaborate on these extensions in Sec. V. Finally, \mathbb{P}_1 can also account for time varying capacities as we can replace B_n with the time average term $\sum_j B_n^j \rho^j$, and similarly for parameter H ; it suffices to augment the state space \mathcal{J} accordingly.

An important comment is in place here. If one knew in advance the value of ρ , then we could solve \mathbb{P}_1 to obtain the optimal offloading solution \mathbf{y}^* . This solution can then be implemented as a randomized policy to maximize the performance of the service. Namely, in each slot t we observe the state π_t and decide to offload or not based on the respective element of \mathbf{y}^* . Nevertheless, in practice one does not have access to ρ and hence cannot devise that optimal static policy, i.e., cannot solve problem \mathbb{P}_1 . In line with the standard approach in stochastic optimization, cf. [9], we will use the unknown solution of \mathbb{P}_1 as the performance benchmark that our online algorithm aims to meet while being oblivious to the task statistics and the system parameters.

III. DECISION FRAMEWORK AND ONLINE ALGORITHM

Our solution approach is the following: we replace the unknown parameters in \mathbb{P}_1 with their running averages, which we calculate in runtime; and we solve the modified problem with approximate gradient ascent in the dual space and perform primal averaging. This gives us an online policy that can be implemented in real time. More importantly, we will prove that this policy ensures asymptotically optimal performance and constraint satisfaction; a desirable result that cannot be achieved by greedy or other heuristic approaches.

A. Problem Decomposition and Algorithm

To streamline presentation we define the functions:

$$f(\mathbf{y}) = -\sum_{n=1}^N \sum_{j=1}^M y_n^j w_n^j \rho^j, \quad g_n(\mathbf{y}) = \sum_{j=1}^M y_n^j \sigma_n^j \rho^j - B_n, \quad \forall n \in \mathcal{N},$$

$$g_{N+1}(\mathbf{y}) = \sum_{n=1}^N \sum_{j=1}^M y_n^j h_n^j \rho^j - H,$$

that appear in problem \mathbb{P}_1 , and we further collect all constraints in function $g(\mathbf{y}) : \mathbb{R}^{NM} \rightarrow \mathbb{R}^{N+1}$. Since we can only observe the current system state, we define the respective t -slot functions that aggregate this information up to t :

$$f_t(\mathbf{y}_t) = -\sum_{n=1}^N \sum_{j=1}^M y_{nt}^j w_n^j \rho_t^j, \quad g_{nt}(\mathbf{y}_t) = \sum_{j=1}^M y_{nt}^j \sigma_n^j \rho_t^j - B_n,$$

$$g_{N+1,t}(\mathbf{y}_t) = \sum_{n=1}^N \sum_{j=1}^M y_{nt}^j h_n^j \rho_t^j - H$$

where $\rho_t^j = 1/t \sum_{\tau=1}^t \mathbf{1}_{\{\pi_\tau=j\}}$ measures the relative occurrences of state j up to slot t and serves as a prediction for the respective ρ^j parameter. Our goal is to use the above *proxy* functions in order to find a dynamic policy $\{\mathbf{y}_t\}_{t=1}^T$ such that the realized performance $\sum_{t=1}^T f(\mathbf{y}_t)/T$ approaches $f(\mathbf{y}^*)$, and similarly the induced constraint violation $\sum_{t=1}^T g(\mathbf{y}_t)/T$ approaches $g(\mathbf{y}^*) \preceq 0$, for any value of time horizon T .

It is worth noticing that the t -slot functions can be expressed as perturbations of the actual unknown functions:

$$f_t(\mathbf{y}) = -\sum_{j=1}^M \sum_{n=1}^N y_n^j w_n^j \rho^j + \sum_{j=1}^M \sum_{n=1}^N y_n^j w_n^j (\rho^j - \rho_t^j)$$

$$\triangleq f(\mathbf{y}) + \epsilon_t(\mathbf{y}),$$

with $\epsilon_t(\mathbf{y}) = \sum_{j=1}^M \sum_{n=1}^N y_n^j w_n^j (\rho^j - \rho_t^j)$. Similarly, we write:

$$g_t(\mathbf{y}) = g(\mathbf{y}) + \delta_t(\mathbf{y}), \quad \delta_t(\mathbf{y}) = (\delta_{nt}(\mathbf{y}), n = 1, \dots, N+1),$$

$$\text{where: } \delta_{nt}(\mathbf{y}) = \sum_{j=1}^M y_n^j \sigma_n^j (\rho_t^j - \rho^j), \quad \forall n \in \mathcal{N},$$

$$\text{and } \delta_{N+1,t}(\mathbf{y}) = \sum_{j=1}^M \sum_{n=1}^N y_n^j h_n^j (\rho_t^j - \rho^j).$$

Next, we can define a new problem for each slot t :

$$\mathbb{P}_2(t) : \underset{\mathbf{y} \in \mathcal{Y}}{\text{maximize}} \quad f_t(\mathbf{y}) \quad \text{s.t.} \quad g_t(\mathbf{y}) \preceq 0. \quad (5)$$

And we will use $\{\mathbb{P}_2(t)\}_t$ to perform a dual ascent and obtain the $\{\mathbf{y}_t\}_t$ that applied in real time.

First, we dualize $\mathbb{P}_2(t)$ and introduce the Lagrangian:

$$L_t(\mathbf{y}, \boldsymbol{\lambda}) = f_t(\mathbf{y}) + \boldsymbol{\lambda}^\top g_t(\mathbf{y})$$

where $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_N, \mu)$ are the non-negative dual variables that will be acting as *shadow prices* for the $N+1$ time-average (or, budget) constraints. The dual function is:

$$V_t(\boldsymbol{\lambda}) = \min_{\mathbf{y} \in \mathcal{Y}} L_t(\mathbf{y}, \boldsymbol{\lambda}).$$

The basis of our approach is the application of a dual-ascent algorithm where the iterations are in sync with the system's time slots t . Specifically, in each iteration t we minimize the Lagrangian by executing:

$$y_n^{j,*} = \arg \min_{y_n^j \in [0,1]} y_n^j (-w_n^j + \lambda_{nt} \sigma_n^j + \mu_t h_n^j) \rho_t^j. \quad (6)$$

This yields the currently optimal offloading policy for each state $j \in \mathcal{J}$, based on which we derive an easy-to-implement offloading rule. That is, denoting with j_t the state at slot t , we write for the offloading decision of each device $n \in \mathcal{N}$:

$$y_n^{j_t} = \begin{cases} 1 & \text{if } \lambda_{nt} o_n^{j_t} + \mu_t h_n^{j_t} < w_n^{j_t} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Note that in practice, state j_t is not entirely known to each device $n \in \mathcal{N}$, but it rather refers to the partial system state regarding the device. This is possible since each device knows its own expected power consumption of the current slot by, e.g. estimating the channel state and also the expected cloudlet resource consumption through the image's file size. Eq. (7) dictates an offloading when the expected accuracy gain $w_n^{j_t}$ exceeds the weighted resource cost $\lambda_{nt} o_n^{j_t} + \mu_t h_n^{j_t}$ that is calculated based on the constraint violation which, in turn, is reflected on the shadow prices.

Then, we improve the current value of $V_t(\boldsymbol{\lambda})$ by updating the dual variables:

$$\lambda_{n,t+1} = \left[\lambda_{nt} + a_t \left(\sum_{j=1}^M o_n^j \rho_t^j y_n^j - B_n \right) \right]^+, \quad \forall n \in \mathcal{N}, \quad (8)$$

$$\mu_{t+1} = \left[\mu_t + a_t \left(\sum_{n=1}^N \sum_{j=1}^M h_n^j \rho_t^j y_n^j - H \right) \right]^+ \quad (9)$$

where $[u]^+ = \max\{0, u\}$ and a_t is the dual step.

The online task outsourcing algorithm, henceforth called *OnAlgo*, is based on eq. (7)-(9). The details are presented in Algorithm 1. When each device n receives an object s_{nt} in slot t , it uses its classifier to predict its class and the predictor to estimate the cloudlet's classification improvement (Steps 5-7). Then, the device uses its threshold decision rule (Step 9) that compares the expected benefits for state j_t with the outsourcing costs for the device and cloudlet. If the cloudlet is not expected to offer satisfactory gains (or, even worse, has lower accuracy), the devices refrain from offloading their tasks. The devices receive the updated state distribution from the cloudlet (Step 12), and update their local dual variable for the power constraint (Step 13). The cloudlet initially evaluates the current system state and sends it to the devices (step 15). Then, it classifies the received objects and updates its congestion variable (Step 17), which is sent to devices.

It is interesting to observe that OnAlgo is lightweight in its computation and communication requirements. The offloading decision are made simply by using an intuitive threshold rule that weights the expected performance gains with the expected costs where the latter are captured in a systematic way via the dual multipliers. And this rule can be employed by each device independently. Similarly, the updates of the dual variables involve summation of scalars and projection onto the non-negative orthant, i.e., keeping only the positive result or setting equal to zero otherwise.

IV. PERFORMANCE ANALYSIS

The gist of our approach is that, as time evolves, the sequence of problems $\{\mathbb{P}_2(t)\}_t$ that aggregate the statistical information up to slot t , approaches the original problem \mathbb{P}_1 .

Algorithm 1: OnAlgo

```

1: Initialization:  $t = 0, \lambda_0 = 0, \forall n, j$ 
2: while True do
3:   for each device  $n \in \mathcal{N}$  do
4:      $y_{nt}^j = 0, \forall j$ 
5:     Receive object  $s_{nt}$ 
6:     Classify objects and obtain  $J_n(s_{nt}), d_n(s_{nt}), \forall s_{nt}$ 
7:     Use classification results on predictor to obtain  $w_{nt}$ 
8:     Observe partial current state  $j_t$  and send it to cloudlet
9:     if  $\lambda_{nt} o_n^{j_t} + \mu_t h_n^{j_t} < w_n^{j_t}$  then
10:       $y_n^{j_t} \leftarrow 1$  % Send object to cloudlet
11:     end if
12:     Receive updated distribution  $\rho_t^j$  from the cloudlet
13:      $\lambda_{n,t+1} \leftarrow [\lambda_{nt} + \alpha_t (\sum_{j=1}^M o_n^j \rho_t^j y_n^j - B_n)]^+, \forall n \in \mathcal{N}$ 
14:   end for
15:   Cloudlet: Receive partial system states from devices, and send back  $\rho_t^j$ 
16:   Compute tasks received from all devices
17:    $\mu_{t+1} \leftarrow [\mu_t + \alpha_t (\sum_{n=1}^N \sum_{j=1}^M h_n^j \rho_t^j y_n^j - H)]^+$ 
18:   Send  $\mu_{t+1}$  to devices
19:    $t \leftarrow t + 1$ 
20: end while

```

We note that the following analysis is general as it holds for different functions $f(\mathbf{y})$ and $g(\mathbf{y})$ than the above, as long as they are convex. We first introduce formally the necessary assumptions and then present a set of technical Lemmas that lead to our main Theorem.

Assumption 1. *The constraint functions and objective functions of $\{\mathbb{P}_2(t)\}_t$ satisfy: $|f_t(\mathbf{y})| \leq \sigma_f, \|g_t(\mathbf{y})\| \leq \sigma_g, \forall t, \mathbf{y} \in \mathcal{Y}$ for some bounded positive parameters σ_f, σ_g .*

Assumption 2 (Slater Condition). *There exists a vector $\mathbf{y}_s \in \mathcal{Y}$ such that $g_t(\mathbf{y}_s) \prec 0, \forall t$.*

These conditions are mild as they, essentially, require that the performance gains and costs are bounded; and that there is an offloading strategy, \mathbf{y}_s , that satisfies strictly the constraints – e.g., transmitting a very small number of images would satisfy this requirement.

A. Complementary Slackness and Constraint Bounds

Lemma 1 (Complementary Slackness Lower Bound). *Under the dual update (8)-(9) it holds:*

$$-\sum_{t=1}^T \boldsymbol{\lambda}_t^\top g_t(\mathbf{y}_t) \leq \frac{\sigma_g^2}{2} \sum_{t=1}^T a_t + \frac{1}{2} \sum_{t=1}^T \|\boldsymbol{\lambda}_t\|^2 \left(\frac{1}{a_t} - \frac{1}{a_{t-1}} \right) - \frac{\|\boldsymbol{\lambda}_{T+1}\|^2}{2a_T} \quad (10)$$

The next result bounds the constraint violation of OnAlgo.

Lemma 2 (Bounded Constraint Violation). *Under the dual update (8)-(9) it holds:*

$$\left\| \frac{1}{T} \sum_{t=1}^T g(\mathbf{y}_t) \right\| \leq \frac{\|\boldsymbol{\lambda}_{T+1}\|}{T a_T} + \frac{1}{T} \sum_{t=1}^T \|\boldsymbol{\lambda}_t\| \left(\frac{1}{a_{t-1}} - \frac{1}{a_t} \right) + \frac{1}{T} \sum_{t=1}^T \|\delta_t(\mathbf{y}_t)\|. \quad (11)$$

B. Approximate Primal Averaging Bounds

The basic idea is that OnAlgo converges to an approximate saddle point. Approximate complementary slackness then allows us to bound the performance gap. We use the next lemma.

Lemma 3 (Approximate Saddle Point). *When $\{\mathbf{y}_t\}_t$ are selected using (7), the t -slot Lagrangian is bounded by:*

$$\frac{1}{T} \sum_{t=1}^T L_t(\mathbf{y}_t, \boldsymbol{\lambda}_t) - f(\mathbf{y}^*) \leq \frac{1}{T} \sum_{t=1}^T \left(\epsilon_t(\mathbf{z}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \right) \quad (12)$$

where $\mathbf{z}_t \in \arg \min_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{y}) + \boldsymbol{\lambda}_t^\top g(\mathbf{y})$.

We can now state and prove the main theorem.

Theorem 1 (Performance Bounds). *OnAlgo ensures:*

$$\begin{aligned} \text{(a)} : \quad & \frac{1}{T} \sum_{t=1}^T f(\mathbf{y}_t) - f(\mathbf{y}^*) \leq C_T + \frac{\sigma_g^2}{2T} \sum_{t=1}^T a_t \\ & - \frac{\|\boldsymbol{\lambda}_{T+1}\|^2}{2Ta_T} + \frac{1}{2T} \sum_{t=1}^T \|\boldsymbol{\lambda}_t\|^2 \left(\frac{1}{a_t} - \frac{1}{a_{t-1}} \right) \\ \text{(b)} : \quad & \left\| \frac{1}{T} \sum_{t=1}^T g(\mathbf{y}_t) \right\| \leq \frac{\|\boldsymbol{\lambda}_{T+1}\|}{Ta_T} + \frac{1}{T} \sum_{t=1}^T \|\delta_t(\mathbf{y}_t)\| \\ & + \frac{1}{T} \sum_{t=1}^T \|\boldsymbol{\lambda}_t\| \left(\frac{1}{a_{t-1}} - \frac{1}{a_t} \right) \end{aligned}$$

$$\text{where } C_T = \frac{1}{T} \sum_{t=1}^T \left(\epsilon_t(\mathbf{z}_t) - \epsilon_t(\mathbf{y}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \right),$$

$$\text{with } \mathbf{z}_t \in \arg \min_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{y}) + \boldsymbol{\lambda}_t^\top g(\mathbf{y}).$$

Proof: Replacing the definition of the t -slot Lagrangian, $L_t(\mathbf{y}_t, \boldsymbol{\lambda}_t) = f_t(\mathbf{y}_t) + \boldsymbol{\lambda}_t^\top g_t(\mathbf{y}_t)$, in Lemma 3 and subtracting $(1/T) \sum_{t=1}^T \boldsymbol{\lambda}_t^\top g_t(\mathbf{y}_t)$ from both sides we can write:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T f_t(\mathbf{y}_t) - f(\mathbf{y}^*) & \leq \frac{1}{T} \sum_{t=1}^T \left(\epsilon_t(\mathbf{z}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \right) \\ & - \frac{1}{T} \sum_{t=1}^T \boldsymbol{\lambda}_t^\top g_t(\mathbf{y}_t), \end{aligned}$$

and expanding $f_t(\mathbf{y}_t) = f(\mathbf{y}_t) + \epsilon_t(\mathbf{y}_t)$ and using Lemma 1, we eventually get:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T f(\mathbf{y}_t) - f(\mathbf{y}^*) & \leq \frac{1}{T} \sum_{t=1}^T \left(\epsilon_t(\mathbf{z}_t) - \epsilon_t(\mathbf{y}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \right) \\ & + \frac{\sigma_g^2}{2T} \sum_{t=1}^T a_t + \frac{1}{2T} \sum_{t=1}^T \|\boldsymbol{\lambda}_t\|^2 \left(\frac{1}{a_t} - \frac{1}{a_{t-1}} \right) - \frac{\|\boldsymbol{\lambda}_{T+1}\|^2}{2Ta_T}. \end{aligned}$$

The second claim of the Theorem follows from Lemma 2. ■

Theorem 1 characterizes the optimality gap and time-average constraint violation for any time horizon T . And we see that in the long run we are guaranteed to approach the (unknown) optimal solution and ensure that the constraints will be respected, i.e., we will not exceed the devices' power budget, nor the cloudlet's capacity. The steps can be constant,

e.g., $a_t = a$ as in [7] or diminishing, e.g., $a_t = a/t^\beta$, with $\beta \in (0, 1)$. Also, the theorem reveals how the error terms of the proxy functions affect the convergence; and it is valid even if one uses other types of estimators, e.g., employing Gaussian Processes to approximate the objective and constraints.

C. Convergence Analysis

The final step of our analysis is to study the convergence of the proposed algorithm. First, it is important to see that Theorem 1 provides a full characterization of the performance gap, and demonstrates how this depends on the time horizon, the system parameters (e.g., σ_g), the system perturbations (errors), and the update steps $\{a_t\}_t$. The convergence rate depends on all these factors. We start by proving that $\boldsymbol{\lambda}_t$ is bounded $\forall t$, which is a technical requirement for our analysis.

1) *Boundedness of Multipliers:* For $\boldsymbol{\lambda}_t$ to remain bounded we need $g_t(\mathbf{y}_t)$ to converge to 0 sufficiently quickly or to be negative sufficiently often. We start with the following result:

Lemma 4 (Bounded level set). *Under Assumptions 1-2 and defining $q := \min_t q_t$ with $q_t = \min_n \{-g_{nt}(\mathbf{y}_s)\} > 0$, it holds:*

$$\sum_{n=1}^{N+1} \lambda_n \leq (\sigma_f - v)/q, \quad \forall \boldsymbol{\lambda} \in \mathcal{Q}_v := \{\boldsymbol{\lambda} \succeq 0 \mid V_t(\boldsymbol{\lambda}) \geq v\}$$

Lemma 5 (Dual vector bound). *Under Assumptions 1-2, the dual update (8)-(9) ensures $\|\boldsymbol{\lambda}_t\|$ is uniformly bounded.*

Obtaining an upper bound for the norm $\|\boldsymbol{\lambda}_t\|$ ensures that the respective terms $\|\boldsymbol{\lambda}_T\|/T$ appearing on the bounds of Theorem 1, are guaranteed to diminish with time.

2) *Error Terms:* Finally, we characterize the aggregate error terms that are induced by the employed approximate dual method, and which affect the bounds of Theorem 1. We write:

$$\sum_{t=1}^T \epsilon_t(\mathbf{z}_t) - \epsilon_t(\mathbf{y}_t) = \sum_{t=1}^T \sum_{j=1}^M \sum_{n=1}^N w_n^j (z_{nt}^j - y_{nt}^j) (\rho_t^j - \rho^j) \quad (13)$$

and similarly the other error-related terms of the performance and constraint bounds:

$$\begin{aligned} \sum_{t=1}^T \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) & = \sum_{t=1}^T \sum_{j=1}^M \sum_{n=1}^N o_n^j \lambda_{nt} z_{nt}^j (\rho_t^j - \rho^j) \\ & + \sum_{t=1}^T \sum_{j=1}^M h_n^j \mu_t z_{nt}^j (\rho_t^j - \rho^j), \end{aligned}$$

$$\begin{aligned} \sum_{t=1}^T \delta_t(\mathbf{y}_t) & = \sum_{t=1}^T \sum_{j=1}^M \sum_{n=1}^N o_n^j y_{nt}^j (\rho_t^j - \rho^j) \\ & + \sum_{t=1}^T \sum_{j=1}^M h_n^j y_{nt}^j (\rho_t^j - \rho^j). \end{aligned}$$

Now, we can upper bound the above terms by their norms and observe that, since parameters $w_n^j, \forall n, j$ and the offloading variables are uniformly bounded, their overall behavior depends on terms $|\rho_t^j - \rho^j|, \forall j \in \mathcal{J}$. Hence, as long as the running average of the realizations for each state j converge

to the respective mean value, the errors gradually diminish to zero. The conditions that ensure this convergence range from the random variables $\{\mathbf{1}_{\{\pi_t=j\}}\}_{t,j}$ being i.i.d. where the Law of Large Numbers applies; to more general settings where they are independent and we can use Hoeffding's inequality [59, Theorem 1] to obtain:

$$\text{Prob}\left(|\rho_t^j - \rho^j| > \kappa\right) < \frac{1}{e^{2\kappa^2 t}}. \quad (14)$$

And one can further relax the assumptions regarding the system state statistics, e.g., to allow for a martingale-type weakly dependence across successive states, and employ the Azuma inequality for a similar bound. These conditions generalize the stricter requirements of i.i.d. statistics that other network optimization frameworks require [9].

Concluding, it is interesting to consider some special cases in order to shed light on the favorable convergence properties of our algorithm. Namely, for the case where we use the step $a_t = a/\sqrt{t}$, Theorem 1(a) shows that the average gap closes at a rate of $\mathcal{O}(1/\sqrt{T})$. To see this, first note that it holds:

$$\frac{\sigma_g^2}{2T} \sum_{t=1}^T a_t \leq \frac{\sigma_g^2}{2T} \sum_{t=1}^T \frac{a}{\sqrt{t}} \leq \frac{\sigma_g^2}{2T} 2a\sqrt{T} = \frac{2a\sigma_g^2}{\sqrt{T}} = \mathcal{O}(T^{-1/2})$$

We can also bound the next term:

$$\begin{aligned} \frac{1}{2T} \sum_{t=1}^T \|\lambda_t\|^2 \left(\frac{1}{a_t} - \frac{1}{a_{t-1}} \right) &= \frac{1}{2aT} \sum_{t=1}^T \|\lambda_t\|^2 \left(\sqrt{t} - \sqrt{t-1} \right) \\ &\stackrel{(a)}{\leq} \frac{\|\lambda_{max}\|^2}{2aT} \sum_{t=1}^T \left(\sqrt{t} - \sqrt{t-1} \right) = \frac{\|\lambda_{max}\|^2}{2a\sqrt{T}} = \mathcal{O}(T^{-1/2}) \end{aligned}$$

where (a) follows from Lemma 5 (dual vectors uniformly bounded). And similarly we can bound the last RHS term in Theorem 1(a) by $\mathcal{O}(T^{-1/2})$. Now it remains to bound C_T . Indeed, when the perturbations are i.i.d. the gap between the running average of the state probabilities (ρ_t^j) and their mean values (ρ^j) diminishes at the rate of $\mathcal{O}(T^{-1/2})$. Hence, using the fact that $\sum_{t=1}^T 1/\sqrt{t} \leq 2\sqrt{T}$, and that all variables are bounded in $[0, 1]$ the error term in (13) diminishes with rate $\mathcal{O}(1/\sqrt{T})$ as well. Finally, it is easy to see that with a similar argument we find that the constraint violation diminishes with rate $\mathcal{O}(1/\sqrt{T})$ in this case. Hence, overall the algorithm converges with that rate, both w.r.t. the optimality gap and the constraint violation.

V. MODEL AND ALGORITHM EXTENSIONS

We extend our framework by jointly optimizing prediction accuracy and total execution delay, since the latter can also be crucial for many edge services. Then, we explain how it can cope with massive demand scenarios where the wireless bandwidth becomes a bottleneck or the cloudlet's energy cost is significant; and finally we elaborate on alternative designs/usages of the predictor.

Joint Accuracy and Delay Optimization: We extend our model to capture both the accuracy gains and the impact of offloading decisions on delay. We do so by adding the total delay for processing the tasks of all users in the objective function and using a scaling parameter $\zeta \in [0, 1]$ to balance

between the two objectives. In detail, we can express the total delay as:

$$D_{tot}(\mathbf{y}) = \sum_{n=1}^N \sum_{j=1}^M (1 - y_n^j) D_n^{pr} + y_n^j (D_n^{pr} + D_0^{pr} + D_n^{tr}),$$

where D_n^{pr} , D_0^{pr} are the delays for processing images at device n or the cloudlet, respectively; and D_n^{tr} the delay for transmitting images to cloudlet. These quantities can vary with time, similarly to the other system parameters, because each image has different size or the wireless medium changes. The processing delays can be modeled with linear functions as we enforce the processing capacity constraints. That is, we can write $D_n^{pr} = k_n/H_{d,n}$, where k_n is the number of CPU cycles required for processing the images of device n , and $H_{d,n}$ is the processing speed of device n (cycles/sec). Similarly, we can define the processing delay at the cloudlet as $D_0^{pr} = k_0/H$ which may vary with time; we refer the reader also to [17] and references therein.

Regarding the transmission delay, this depends on the actual system architecture. For example, if different channels are employed for the users, we can express it as $D_n^{tr} = \ell_n/(r_n W)$, where ℓ_n is the size of each image, r_n the channel gain for user n , and W the link bandwidth. If there is a CSMA-type network where users need to share their links, we need to replace W with the actual airtime W_n that user n receives; and in the case we have a fair round-robin (vanilla version of CSMA) we can approximate this with $W_n = \sum_{j=1}^M y_n^j / \sum_{n=1}^N \sum_{j=1}^M y_n^j$. This model has been used extensively in Wi-Fi service allocation, see [18], and in mobile code offloading, e.g., in [17].

Following the analysis in Sec. II we can replace in \mathbb{P}_1 the new objective function $f(\mathbf{y}) = \sum_{j=1}^M \sum_{n=1}^N w_n^j y_n^j \rho^j - \zeta D_{tot}(\mathbf{y})$, and by following the same process obtain the offloading rule:

$$y_{nt}^j = \begin{cases} 1 & \text{if } \lambda_{nt} o_n^{j_t} + \mu_t h_n^{j_t} < w_n^j - \zeta (D_{nt}^{tr} + D_0^{pr}) \\ 0 & \text{otherwise} \end{cases}, \quad (15)$$

where we observe that the device execution delay is nullified since it is independent of the offloading decision, and the condition in line 10 of Algorithm 1 will be replaced by (15).

Wireless Bandwidth and Energy Cost Constraints: We have assumed the system operation is constrained by the devices' power budget and the computing capacity of the cloudlet. Indeed, most often these are the bottleneck resources [5], [6], [20]. However, in scenarios of massive demand the wireless link capacity might also be a bottleneck constraint. Our analysis can be readily extended for this case. If we denote with $\{W_t\}_{t=1}^\infty$ the link capacity process (uniformly bounded; well-defined mean value W) assuming a wireless link shared by all devices⁷, we can add to \mathbb{P}_1 the constraint:

$$\sum_{n=1}^N \sum_{j=1}^M y_n^j \rho^j \ell_n \leq W. \quad (16)$$

Eq. (16) can be handled as the computing constraint (4) and

⁷This can be either an OFDM-based cellular link or a coordinated access WiFi link; in the case we have a CSMA-type of mechanism, one needs to account for the additional bandwidth loss due to collisions, etc.

will only affect the convergence bounds. Similarly, we can include other constraints that couple the actions of all devices, such as the energy cost at the cloudlet which increases with the aggregated offloaded tasks from all devices and might depend on time-varying energy prices.

Alternative System Architectures: A different mechanism is possible, where the devices send objects to the cloudlet before using their own classifier. This approach can reduce the consumed energy, since it avoids low-accuracy local classifications. However, it requires a different type of a predictor, namely one that can estimate the expected accuracy gain using some basic features of the object (e.g., its file size), and without requiring input from the local classifier. In this case, modeling the power consumption of the devices would modify constraint (3) of \mathbb{P}_1 as:

$$\sum_{j=1}^M \left(y_n^j \rho^j o_n + (1 - y_n^j) \rho^j \nu_n \right) \leq B_n, \quad \forall n \in \mathcal{N},$$

where the second term indicates the power ν_n consumed by each device when only local classification is performed. OnAlgo can be extended to this case by changing the predictor.

Similarly, it is possible to have services that are executed in multiple stages, e.g., a video stream is compressed, then frames of interest are selected, and objects are identified on each frame. In this case, the devices might decide to outsource some of the tasks in the first stage, some others after the second stage, and so on. Again, our optimization algorithms can be extended to include these decisions, by defining a separate set of variables for each stage while accounting for the costs and properties (e.g., data volumes) in each case. In specific, (3) would be transformed to:

$$\sum_{j=1}^M \left(y_n^j \rho^j o_n + (1 - y_n^j) \rho^j \nu_n^{cl} \right) \leq B_n, \quad \forall n \in \mathcal{N},$$

where ν_n^{cl} is the classification computing cost, which is significantly smaller than ν_n . Observe that the computing load of stage 1, i.e. feature extraction, is not accounted for since it is again induced regardless of the offloading decision.

VI. IMPLEMENTATION AND EVALUATION

We have implemented the proposed architecture in a small wireless testbed using 4 Raspberry Pis, which we also leveraged to create synthetic traces so as to run additional larger-scale simulations. The evaluation of OnAlgo was conducted using images from two publicly-available datasets of images, and compared against benchmarks from previous works – to which we provide pointers. Our evaluation plan has four goals: (i) investigate the accuracy performance of well-known classifiers for different sizes of training datasets, hence revealing why *edge augmentation* is needed; (ii) Measure the energy and computing costs of image classification tasks; (iii) Perform a parameter-sensitivity analysis of OnAlgo; and (iv) Compare OnAlgo with several benchmark algorithms.

A. Experiments Setup

1) *Testbed and Measurements:* We used 4 Raspberry Pis (RPs) as end-nodes, and a cloudlet with specs as in [21],

see Fig. 2a. The RPs are placed in different distances from the cloudlet, and all plots are using data from at least 50 experiments. We measured energy using a Monsoon Power Monitor, and used Python libraries and TensorFlow for the classifiers. We have used *vanilla* versions of libraries and classifiers so as to facilitate observation of the results.⁸

We measured the average power consumption when a RP transmits with different rates, Fig. 2b. Then we fitted a linear regression model that estimates the consumed power (Watts) as a function of the rate r , $p(r) = -0.00037r^2 + 0.0214r + 0.1277$. This result is used by OnAlgo to estimate the energy cost for each image, given the data rate in each slot (which might differ for the RPs). Also, we measured the average computing cost (cycles/task) for the classification task for a convolutional neural network (CNN) in the RPs and cloudlet. Since the images have different sizes, we observed that the computation load varies, with a mean of 441 *Mcycles* and std. 90 *Mcycles* for the cloudlet (see Fig. 2c), and a mean of 3044 *Mcycles* and std. 173 *Mcycles* for RPs. Regarding the delays, we measured device and cloudlet average processing and transmission delays and found that $D_n^{pr} = 2.537$, $D_0^{pr} = 0.191$ and $D_n^{tr} = 0.157$ *ms*. This result suggests that local processing is about 10 times slower than offloading in our system. Hence, it is possible that the extra offloading delay experienced by the devices can be worth trading off for the enhanced accuracy of the cloudlet.

2) *Data Sets and Classifiers:* We focus on image classification, a widely employed analytic task, and use two well-known data sets: (i) **MNIST** which consists of 28×28 pixel handwritten digits, and includes 60K training and 10K test examples; (ii) **CIFAR-10** that consists of 50K training and 10K test examples of 32×32 color images of 10 classes. We used two very different classifiers: the normalized-distance weighted k -nearest neighbors (KNN) algorithm [22], and the more sophisticated Convolutional Neural Network (CNN), implemented with TensorFlow [23]. Both classifiers output a vector where each coordinate represents the probability that the object belongs to the respective class. These classifiers differ substantially in their performance and resource requirements, hence allowing us to build diverse experiment scenarios. Our goal is to evaluate both and determine which one is more suitable depending on other system parameters like the number of available training samples at each location.

The predictors are trained with labeled images and the outputs of the local ($d_n(s_{nt})$) and cloudlet ($d_0(s_{nt})$) classifiers. We implemented an ordinary least squares regressor and a model-free random forest that estimate ϕ_{nt} (dependent variables) based on the classifier outputs (independent variables). Recall that the dependent variables are calculated using (1). We have used training sets of different sizes and two different regressors: (i) a general model, where the prediction does not consider the locally inferred class as an independent variable; and (ii) a class specific model that is based on the output of the local classifier.

⁸For instance, the memory footprint of NNs can be made smaller [27], [28] but such actions possibly affect their performance. Our analysis is orthogonal to such interventions.

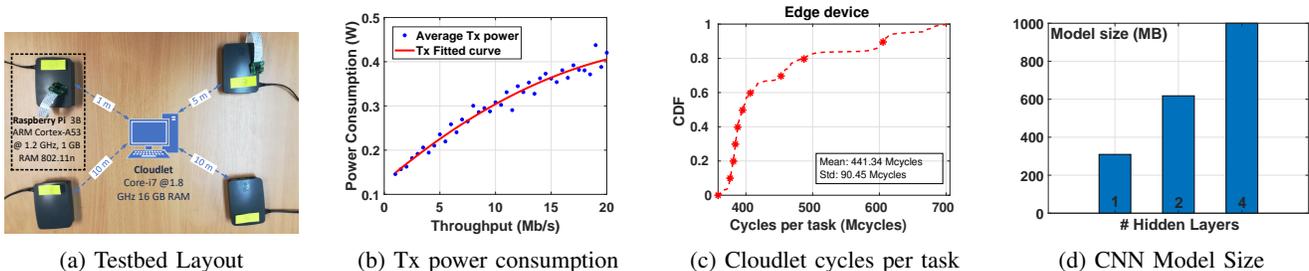


Fig. 2: (a): Testbed: 4 RPs and a cloudlet (laptop). (b): Transmit power consumption measurements and the fitted curve for the RPs. (c) CDF of computing cycles per task for the cloudlet. (d) Increasing the number of layers in CNN increases the model size (MB) up to 100%.

3) *Benchmark Algorithms*: We compare OnAlgo with three different algorithms:

- *Accuracy-Threshold Offloading (ATO)*, where a task is offloaded when the confidence of the local classifier is below a threshold, without considering the resource consumption. This is basically the non-distributed version of [24], where if the local result is not sufficiently reliable, further CNN layers in the edge or cloud are invoked.
- *Resource-Consumption Offloading (RCO)*, where a task is offloaded when there is enough energy, without considering the expected classification improvement. This is a type of sophisticated greedy algorithm that takes myopic decisions in each slot.
- *Online Code Offloading and Scheduling (OCOS)* [25], where the devices always try to exploit the cloudlet’s classifier, and the cloudlet tries to schedule as many tasks as possible in each slot, given its available resources. That is, the cloudlet follows a greed strategy.

B. Initial Measurements

1) *Limitations of Mobile Devices*: We used our testbed to verify these small resource-footprint devices require the assistance of a cloudlet. These findings are in line with previous studies, e.g., [24], [26]. The performance of a CNN model increases with the number of layers (as we will show next), but so does the model size, see Fig. 2d. We find that, even with 4 layers, a CNN trained for CIFAR has 1GB size and hence cannot be stored in the RPs (e.g., even more so in a smaller IoT node). Similar conclusions hold for the KNN classifier, the accuracy of which is directly linked to the number of labeled local data (KNN needs the training data available locally). Clearly, despite the efforts to reduce the size of ML models by using, for instance, compression [27] or residual learning [28], the increasing complexity of analytics and the small form-factor of devices will continue to raise the local versus cloudlet execution trade off.

2) *Classifier and Predictor Assessment*: Here we evaluate the different classifier and predictor designs towards building a more efficient system. This provides interesting insights on how these system elements work in practice. In detail, in Fig. 3a we see that the accuracy (defined as the ratio of correct predictions over the sum of all predictions) of the KNN classifier improves with the size K_n of labeled data when applied to MNIST. Figure 3b depicts the accuracy

improvement for CNN as more hidden layers are added. The performance increase is higher for the digits that are more difficult to recognize (e.g., 4 and 5), up to about 20%. Notice, that the performance of the CNN classifier is superior to KNN, when we use fewer layers, or samples respectively. In addition, we present the CNN performance on CIFAR, for 1, 2 and 4 hidden layers in Fig. 3c. CIFAR is more complex than MNIST due to the properties of its objects (colored images, etc.), and this results in lower accuracy. Overall, we see that the classifier performance depends on the algorithm (KNN, CNN, etc.), the settings (datasets, layers, etc.), and differ also for each object class. Hence, an algorithm is required that can adapt to all these parameters (as OnAlgo does). Since we have verified the superiority of CNN classifiers, we continue our evaluation using only these, instead of KNN.

Finally, we studied the training dataset impact on the predictor’s error, using both general and class-specific (i) linear regressors and (ii) random forests. In Fig. 4, we plot the prediction error of the accuracy improvement for both cases of general and class-specific predictors for CNN local device and cloudlet classifiers. We observe that the random forest is superior to the simpler linear regressor only when the number of samples is small. Moreover, random forests display an inconsistency when comparing general to class-based models as the number of training samples varies. The class specific regressor for 5K samples achieves the lowest average absolute error, thus it is used throughout the following experiments, while its error is rapidly decreasing from 35% for 100 points to 12.3% for 5K points on the CIFAR dataset.

C. Performance Evaluation

Next, we evaluate the performance of OnAlgo in terms of achieved accuracy, offloading frequency and resource consumption. First, we evaluate OnAlgo for different values of the power consumption constraint B_n . Then, we use a variable non-i.i.d. traffic load to compare its performance against the competitors, by considering these different criteria. The traffic load is an exponentially distributed sequence of task bursts, with a uniform duration of 5 – 10 seconds. This way we emulate the real-world scenario of sensor-activated cameras that generate images for short time periods.

1) *Resource Availability*: We evaluate OnAlgo, by using a 1-layer CNN for the RPs and a 4-layer CNN for the cloudlet. In Fig. 5 we show the average accuracy achieved by the four devices, as well as the fraction of requests offloaded to the

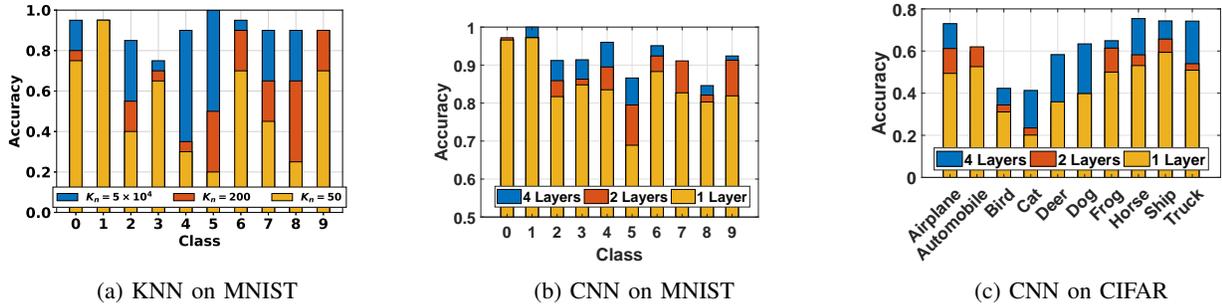


Fig. 3: Per class Accuracy of MNIST and CIFAR-10 for KNN and CNN classifiers of various labeled data sizes and hidden layers. For details of training and validation see Section V.A.

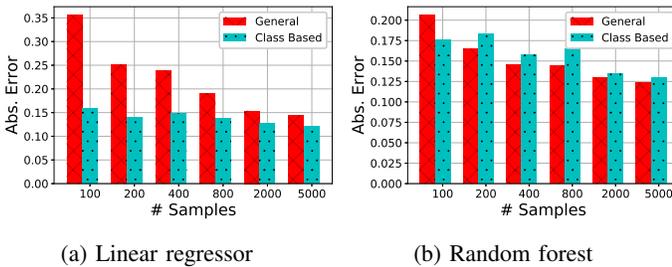


Fig. 4: Predictor assessment.

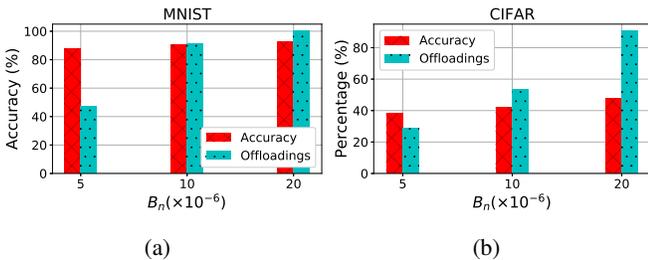


Fig. 5: Accuracy and offloading percentage of OnAlgo for various resource constraints, on MNIST and CIFAR-10.

cloudlet when we vary the devices’ power budget B_n , for MNIST and CIFAR. Evidently, as B_n increases there are more opportunities for exploiting the cloudlet and obtaining a better result than the local classifier. Furthermore, some interesting remarks can be made by comparing the two datasets. As shown in Fig. 3(b-c), MNIST is easier to classify and the gain of using a better classifier is not as important as on the CIFAR dataset. In particular, with MNIST the gains are about 6% in accuracy as the resources (and thus the offloaded tasks) increase. With CIFAR, on the other hand, the potential performance gain when using the cloudlet is higher; and as B_n increases, the accuracy gains are up to 15%. These two experiments demonstrate *the agility of our algorithm, which assesses the potential accuracy gains and shapes accordingly the offloading strategy, based on resource availability.*

2) *Comparison with Benchmarks:* Next, we compare OnAlgo to ATO, RCO and OCOS for a varying non-i.i.d. traffic load in Fig. 6 using the criterion of accuracy and power consumption. Ideally, we would like an algorithm to perform well

in both these dimensions. To ensure a realistic comparison, we set the rule for all algorithms that the cloudlet will not serve any task if the computing capacity constraint is violated; while for RCO the availability of energy is determined by computing the average consumption by each device during the experiment. We employ two scenarios to demonstrate the algorithms’ performance and energy costs under different data sets and resource availability states. The experiments below reveal the, rather expected, finding of an optimization-based dynamic algorithm outperforming the respective greedy (or, myopic) algorithms we use as benchmarks.

Scenario 1: *Low accuracy improvement; high resources.* In this case, we set⁹ $B_n = 0.02 \text{ mW}$, $H = 2 \text{ GHz}$ allowing the devices to offload many tasks, and the MNIST dataset (has small improvement between 1 layer and 4 layer CNNs). We depict the average accuracy achieved by the devices and the average power consumption versus the task load in bursts per minute in Fig. 6a and 6b respectively. We observe that OnAlgo shows a smaller slope in the decrease of accuracy, as the load increases than all the competitors. The performance of ATO quickly drops because the cloudlet’s resources are insufficient for high loads. RCO’s performance is good for the most part, but it quickly deteriorates for high task loads as the devices refrain from offloading due to the power constraints. It is interesting also to note that RCO even outperforms OnAlgo in terms of accuracy (by approx. 2%) but this happens at the expense of larger energy cost, namely it spends more than double the energy of OnAlgo. OCOS performs similarly to RCO since performance degradation is caused by cloudlet resource exhaustion. The problem with both algorithms is that they do not offload intelligently, based on both the improvement potential *and* the availability of resources. Hence, when considering both performance and energy cost criteria, and especially in the non-trivial higher load cases they are significantly outperformed by OnAlgo.

Scenario 2: *High accuracy improvement; low resources.* The settings for this scenario are $B_n = 0.01 \text{ mW}$, $H = 500 \text{ MHz}$ not allowing many offloadings and cloudlet classifications. We used the CIFAR dataset that demonstrates a substantial performance difference between local and cloudlet

⁹We have explicitly set a small power budget so as to highlight the impact of power constraints on the system performance; higher power budgets will still be a bottleneck for higher task request rates or images of larger size.

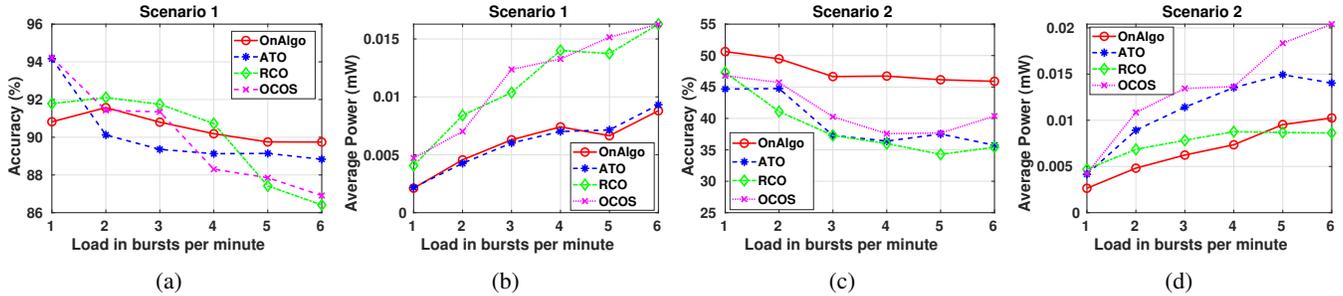


Fig. 6: Comparison of different offloading algorithms w.r.t. their accuracy and energy cost, under different task load conditions.

classifiers. We see from Fig. 6c that OnAlgo is up to 12% more accurate than ATO/RCO for high task load, and in any case significantly higher than in Scenario 1. OCOS performs slightly better than ATO/RCO, but at the cost of very high power consumption. Since the potential of improvement is higher in Scenario 2, ATO marginally outperforms RCO by spending up to 50% more power than RCO (see Fig. 6d). OnAlgo consumes about 50% less power than OCOS since the latter always tries to offloads tasks but does not leverage the cloudlet efficiently due to the lack of computing capacity.

Summing up the 2 scenarios above, we see that OnAlgo achieves a smooth performance across varying traffic loads, while its competitors struggle, especially as the load increases. Moreover, it achieves reasonable power consumption regardless of the resource availability as opposed to RCO in Scenario 1, ATO in Scenario 2, and OCOS in both scenarios. Even when in some cases OnAlgo is being outperformed by some competitor with respect to one criterion (e.g., by RCO w.r.t. accuracy in Scenario 1), this happens at the expense of losing at much larger rates w.r.t. the other criterion (power consumption).

3) *Trade-off Analysis*: Next we demonstrate the trade-offs between number of offloadings, accuracy and resource consumption between OnAlgo and the competitor algorithms using net graphs. Fig. 7a displays the performance of OnAlgo for low medium and high task load. Observe that as the load increases, OnAlgo rapidly increases resource consumption to maintain high accuracy. For instance, comparing low to high load, we see that performance drops only by about 7% as the computing and power consumption is increased by 75%. In Fig. 7b we compare the same metrics for high traffic load, and the different competitors. Observe that OnAlgo achieves the highest accuracy, while being (closely) second best (behind RCO) in terms of computing resource and power consumption. Moreover it achieves high accuracy despite offloading less frequently than OCOS, due to the intelligent way it makes the offloading decisions. In summary, OnAlgo achieves the *highest* accuracy between the competitors, and at the same time has a *moderate* resource consumption.

Next, in Fig. 8 we explore the accuracy-resource consumption-delay trade-off when problem (P3) is considered, i.e. total delay is jointly optimized with accuracy. Notice in Fig. 8a, that the increasing traffic load will not only result in lower accuracy (about 20%) and higher resource consumption, but also in significantly higher delay (up to

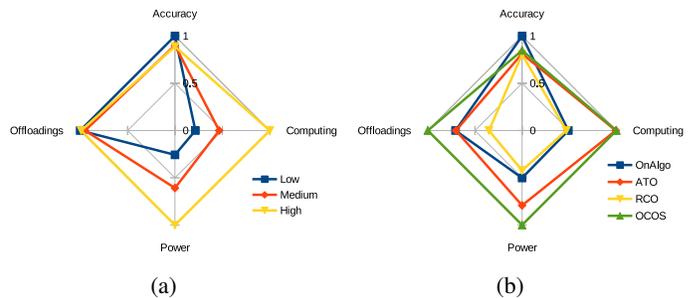


Fig. 7: Comparison of different key metrics (normalized): (a) OnAlgo for low, medium and high traffic load. (b) Algorithm comparison for high load in scenario 2.

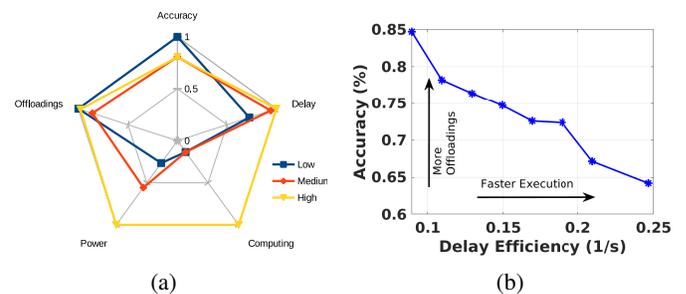


Fig. 8: (a) OnAlgo performance for problem (P3). (b) Pareto front between accuracy and delay efficiency obtained by tuning ζ .

25%). Hence, despite consuming extra resources in high load cases, OnAlgo still maintains high accuracy standards. Finally, Fig. 8b displays the Pareto front between accuracy and delay¹⁰. This shows the effect of parameter ζ (ranging from 0.1 to 0.3) on the resulted offloading policy and consequently on the performance of accuracy and delay. For instance, in order to double the delay efficiency (from 0.1 to 0.2), we would have to sacrifice roughly 10% accuracy, by offloading less frequently.

VII. RELATED WORK

Edge & Distributed Computing. Most solutions partition compute-intense mobile applications and offload them to cloud [29]. This approach cannot support applications with

¹⁰In fact delay is inverted (1/s) so that increasing the value towards either the x-axis or the y-axis yields better performance with respect to the relevant metric.

stringent requirements due to possible large delays in data transfers [30]. Cloudlets on the other hand, achieve lower delay by leveraging edge computing [5], [20] but have limited serving capacity. A different line of work proposes the distribution of tasks among collaborating nodes [53]–[55] using intuitive allocation metrics or static optimization models. Hence, there is need for an intelligent cloudlet offloading strategy and this idea lies at the core of our proposal which, unlike previous works: (i) considers the quality of outcome (accuracy) and resource costs of devices and cloudlet; and (ii) is adaptive and oblivious to statistics of system parameters and user requests.

Previous works in this area consider simple performance criteria, such as reducing the computation load and focus on the architecture design. For example, Misco [31] and CWC [32] implement frameworks for parallel task execution on mobile devices; and similarly MobiStreams [33], Swing [34] and [35], focus on collaborative data stream computations. These systems either do not optimize the offloading policy [33] or use heuristics that do not cater for task accuracy [34], [35]. Instead, OnAlgo ensures optimal performance, subject to resource availability, even when the latter is unknown.

System Designs for Mobile Analytics. The increasing importance of these services has motivated the design of wireless systems that can execute such tasks. For instance, [36]–[38] tailor deep neural networks for execution in mobile devices, and [39] focuses on how to maximize accuracy in edge-cloud deployments. These works however, focus either only on execution delay or accuracy. Glimpse [40] reduces delay in video tracking applications using an active cache of frames at the device; Cachier [41] uses edge servers as caches for image recognition requests so as to minimize latency; and Precog [42] prefetches trained classifiers on devices to accelerate image recognition. In a different approach, [43] selects in runtime the the DNN size, in order to balance accuracy and resource consumption. Similarly, [44] considers a richer set of decisions, including model selection; image compression; and frame rate, aiming to maximize the accuracy of frames. Finally, [45] minimizes execution time and energy cost for a single device, for known system parameters and task loads; while [51] optimizes again delay but through the orchestration of the edge resources. The plethora of such system proposals underlines the necessity for an *analytical framework* for task outsourcing that can optimize performance.

Optimization of Analytics. Prior analytical works in the context of computation offloading focus on different metrics, such as the number of served requests, e.g., see [46] and overview in [6], and hence are not applicable here. In our previous work [47], we proposed a *static* optimization framework for a peer-to-peer collaborative task execution scheme, which does not employ predictions of gains nor accounts for computation constraints. The authors of [48] employ a Lyapunov optimization approach to configure a video analytics application towards balancing the accuracy and energy costs, under i.i.d. requests and system dynamics. In [49], video quality and computing resources are selected to maximize the approximate analytics accuracy. FastVA [50] is a video analytics system that leverages neural processing units at the mobile devices and proposes a heuristic offloading

policy towards maximizing accuracy. Other works that cater for accuracy either rely on heuristics or static models and complete knowledge of system parameters [43], [44], [52].

Clearly, these assumptions are invalid for many practical systems where the expected accuracy improvements, power availability, wireless channels, and cloudlet resources not only vary with time, but often do not follow an i.i.d. process. This renders the application of max-weight type of policies [9] inefficient. Our approach is fundamentally different and leads to a more robust algorithm that converges as long as the perturbations are bounded (in each slot), and have well defined mean values (which can be unknown). Our methodology is inspired by dual averaging and primal recovery algorithms for static problems, see [7], [8], [19]. We have extended here this idea and succeeded in obtaining deterministic bounds and for a broad range of perturbations. It is also important that the employed algorithm is lightweight and amenable to distributed execution, hence can be implemented as a network protocol. This is in contrast with other optimization approaches, e.g. using Bayesian learning [13], [56], which require a centralized computation-demanding execution.

Improvement of ML Models. Clearly, despite the efforts to improve the execution of analytics at small devices, e.g., by compressing NN models [27] or using residual learning [28], the trade off between low-accuracy local and high-accuracy cloudlet execution is still important due to the increasing number and complexity of these tasks. This observation has spurred efforts for designing fast multi-tier deep neural networks [24]; for dynamic model selection [43], [57]; and for threshold-based task allocation to DNNs [26]; see also discussion in [44]. These works are orthogonal to our approach and can be readily incorporated in our framework.

VIII. CONCLUSIONS

This work introduced the idea of augmenting the execution of data analytics at mobile devices with more accurate libraries available at nearby cloudlets. The performance evaluation showcased particularly high gains, both in terms of accuracy and energy, when intelligent offloading is implemented. To that end, we designed a dynamic and distributed algorithm that decides outsourcing the tasks based on the expected performance improvement, and the available device and cloudlet resources. It was shown, theoretically and via experiments, that this *joint performance-cost* design outperforms other proposals that do not cater for the accuracy or resource availability; and it often improves both criteria. Our algorithm achieves near-optimal performance (with tunable gap) in a deterministic fashion, and under minimal assumptions about the system behavior; it suffices the perturbations to be bounded in each slot and have well-defined means. This makes it ideal for the problem at hand where the stochastic effects (e.g., expected accuracy gains) might not follow an i.i.d. or a Markov-modulated process. Finally, we note that the proposed online algorithm is novel as it extends prior primal-averaging techniques [7], [19], and can optimize the network operation under more general conditions than pertinent optimization frameworks, cf. [9]; hence it is of independent interest with potential applications beyond this edge computing.

REFERENCES

- [1] A. Galanopoulos, et al., “Improving IoT Analytics through Selective Edge Execution”, in *Proc. of IEEE ICC*, 2020.
- [2] E. Siow, T. Tiropanis, W. Hall, “Analytics for the Internet of Things: A Survey”, *ACM Comp. Surv.*, vol. 1, no. 1, 2018.
- [3] Cisco White Paper, “Cisco Global Cloud Index: Forecast and Methodology, 2016–2021”, Document ID:1513879861264127.
- [4] Microsoft HoloLens, [Online]: <https://www.microsoft.com/en-ie/hololens>
- [5] M. Satyayan, et al., “The Case for VM-Based Cloudlets in Mobile Computing”, *IEEE Per. Comp.*, vol. 8, no. 4, 2009.
- [6] Y. Mao, C. You, J. Zhang, K. Huang and K. B. Letaief, “A Survey on Mobile Edge Computing: The Communication Perspective”, in *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, 2017.
- [7] A. Nedic, and A. Ozdaglar, “Approximate Primal Solutions and Rate Analysis for Dual Subgradient Methods”, *SIAM Journal on Optimization*, vol. 19, no. 4, 2009.
- [8] K. Kiwiel, et al., “Lagrangian Relaxation via Ballstep Subgradient Methods”, *Math. Oper. Res.*, vol. 32, no. 3, 2007.
- [9] L. Georgiadis, M. J. Neely, and L. Tassiulas, “Resource Allocation and Cross-Layer Control in Wireless Networks”, *Foundations and Trends in Networking*, vol. , no. 1, pp. 1-144, 2006.
- [10] Y. Lecun, et al., “Gradient-based Learning Applied to Document Recognition”, *Proc. of the IEEE*, vol. 86, no. 11, pp. 2278 - 2324, 1998.
- [11] Alex Krizhevsky, “Learning Multiple Layers of Features from Tiny Images”, Tech Report, 2009.
- [12] J. Redmon, and A. Farhadi, “YOLOv3: An Incremental Improvement”, *arXiv:1804.02767*, 2018.
- [13] A. Galanopoulos, et al., “AutoML for Video Analytics with Edge Computing”, in *Proc. of IEEE INFOCOM*, 2021.
- [14] A. Gelman, and J. Hill, “Data Analysis Using Regression and Multilevel Models”, *Cambridge University Press*, ISBN:9780511790942, 2012.
- [15] M. Neely, “Energy Optimal Control for Time-Varying Wireless Networks”, *IEEE Trans. on Inform. Theory.*, vol. 52, no. 7, pp. 2915 - 2934, 2006.
- [16] T. Lan, D. Kao, M. Chiang, A. Sabharwal, “An Axiomatic Theory of Fairness in Network Resource Allocation”, in *Proc. of IEEE INFOCOM*, 2010.
- [17] X. Chen et al., “Efficient Multi-User Computation Offloading for Mobile-Edge Cloud Computing”, *IEEE/ACM Trans. on Networking*, vol. 24, no. 5, 2015.
- [18] L. Li, et al., “Proportional Fairness in Multi-rate Wireless LANs”, in *Proc. of IEEE INFOCOM*, 2008.
- [19] V. Valls, D. Leith, “A Convex Optimization Approach to Discrete Optimal Control”, *IEEE Trans. Aut. Cont.*, vol. 64, no. 1, pp. 35 - 50, 2018.
- [20] J. He, et al., “Multi-tier Fog Computing with Large-scale IoT Data Analytics”, *IEEE Internet of Things*, vol. 5, no. 2, pp. 677-686, 2018.
- [21] Z. Chen, et al., “An Empirical Study of Latency in an Emerging Class of Edge Computing Applications for Wearable Cognitive Assistance”, in *Proc. of IEEE/ACM SEC*, 2017.
- [22] S. Dudani, “The Distance-weighted k-Nearest-Neighbor Rule”, *IEEE Trans. on Systems, Man, & Cyb.*, vol. 6, no. 4, 1976.
- [23] M. Abadi, et al., “TensorFlow: A System for Large-scale Machine Learning”, in *Proc. of USENIX OSDI*, 2016.
- [24] S. Teerapittayanon, B. McDanel and H. T. Kung, “Distributed Deep Neural Networks Over the Cloud, the Edge and End Devices”, in *Proc. of IEEE ICDCS*, 2017.
- [25] B. Zhou et al., “An Online Algorithm for Task Offloading in Heterogeneous Mobile Clouds”, *ACM Trans. Internet Technol.*, vol. 18, no. 2, 2018.
- [26] C. Lo, Y. Y. Su, C. Y. Lee and S. C. Chang, “A Dynamic Deep Neural Network Design for Efficient Workload Allocation in Edge Computing”, in *Proc. of IEEE ICCD*, 2017.
- [27] V. Chandrasekhar et al., “Compression of Deep Neural Networks for Image Instance Retrieval”, *IEEE Data Comp.*, 2017.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition”, *arXiv:1512.03385*, Dec. 2015.
- [29] B. Chun, et al., “CloneCloud: Elastic Execution Between Mobile Device and Cloud”, in *Proc. of ACM EuroSys*, 2011.
- [30] K. Ha, P. Pillai, G. Lewis, S. Simanta, S. Clinch, N. Davies, and M. Satyanarayanan, “The Impact of Mobile Multimedia Applications on Data Center Consolidation”, in *Proc. of IEEE IC2E*, 2013.
- [31] A. Dou, et al., “Misco: A MapReduce Framework for Mobile Systems”, in *Proc. of ACM Petra*, 2010.
- [32] M. Arslan, et al., “Computing While Charging: Building a Distributed Computing Infrastructure Using Smartphones”, in *Proc. of ACM CoNEXT*, 2012.
- [33] H. Wang, et al., “MobiStreams: A Reliable Distributed Stream Processing System for Mobile Devices”, *IEEE IPDPS*, 2014.
- [34] S. Fan, T. Salonidis, and B. Lee, “Swing: Swarm Computing for Mobile Sensing”, in *Proc. of IEEE ICDCS*, 2018.
- [35] D. O’Keefe, T. Salonidis, P. Pietzuch, “Frontier: Resilient Edge Processing for the IoT”, *PVLDB*, vol. 11, no. 10, 2018.
- [36] X. Ran, H. Chen, Z. Liu, J. Chen, “Delivering Deep Learning to Mobile Devices via Offloading”, in *Proc. ACM VR/AR Network Workshop*, 2017.
- [37] N. Lane, et al., “DeepX: A Software Accelerator for Low-power Deep Learning Inference on Mobile Devices”, in *Proc. of IEEE/ACM IPSN*, 2016.
- [38] A. G. Howard et al., “MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications”, in *CoRR*, abs/1704.04861, 2017.
- [39] Y. Wang et al., “Enabling Edge-Cloud Video Analytics for Robotics Applications”, in *Proc. of IEEE INFOCOM*, 2021.
- [40] T. Chen, et al., “Glimpse: Continuous, Real-time Object Recognition on Mobile Devices”, in *Proc. of ACM SenSys*, 2015.
- [41] U. Drolia, et al., “Cachier: Edge-Caching for Recognition Applications”, in *Proc. of IEEE ICDCS*, 2017.
- [42] U. Drolia, et al., “Precog: Prefetching for Image Recognition Applications at the Edge”, in *Proc. of IEEE/ACM SEC*, 2017.
- [43] S. Han, et al., “MCDNN: An Approximation-Based Execution Framework for Deep Stream Processing Under Resource Constraints”, in *Proc. of ACM Mobisys*, 2016.
- [44] X. Ran, H. Chen, X. Zhu, Z. Liue and J. Chen, “DeepDecision: A Mobile Deep Learning Framework for Edge Video Analytics”, in *Proc. of IEEE INFOCOM*, 2018.
- [45] Y. Li, Y. Chen, T. Lan, G. Venkataramani, “MobiQoR: Pushing the Envelope of Mobile Edge Computing via Quality-of-Result Optimization”, in *Proc. of IEEE ICDCS*, 2017.
- [46] X. Chen, L. Jiao, W. Li, and X. Fu, “Efficient Multi-User Computation Offloading for Mobile-Edge Cloud Computing”, *IEEE/ACM Trans. Netw.*, vol. 24, no. 5, pp. 2795 - 2808, 2016.
- [47] A. Galanopoulos, T. Salonidis and G. Iosifidis, “Cooperative Edge Computing of Data Analytics for the Internet of Things”, in *IEEE Trans. on Cognitive Communications and Networking*, vol. 6, no. 4, pp. 1166-1179, 2020.
- [48] S. Zhang, et al., “Adaptive Configuration Selection and Bandwidth Allocation for Edge-Based Video Analytics”, *IEEE/ACM Trans. on Networking*, 2021, Early Access.
- [49] P. Yang, F. et al., “Edge Coordinated Query Configuration for Low-Latency and Accurate Video Analytics”, in *IEEE Trans. on Industrial Informatics*, vol. 16, no. 7, pp. 4855-4864, 2020.
- [50] T. Tan, and G. Cao, “Deep Learning Video Analytics Through Edge Processing and Neural Processing Units in Mobile Devices”, *IEEE Trans. on Mobile Computin*, 2021, Early Access.
- [51] W. Zhang, et al., “Hetero-Edge: Orchestration of Real-time Vision Applications on Heterogeneous Edge Clouds”, in *Proc. of IEEE INFOCOM*, 2019.
- [52] S. Yi, Z.g Hao, Q. Zhang, Q. Zhang, W. Shi, and Q. Li. Lavea, “Latency-aware Video Analytics on Edge Computing Platform”, in *Proc. of IEEE/ACM Symp. on Edge Computing*, 2017.
- [53] J. C. S. Dos Anjos, et al., “Data Processing Model to Perform Big Data Analytics in Hybrid Infrastructures”, *IEEE Access*, vol. 8, pp. 170281-170294, 2020.
- [54] A. M. Ghosh, and K. Grolinger, “Edge-Cloud Computing for Internet of Things Data Analytics: Embedding Intelligence in the Edge With Deep Learning”, *IEEE Trans. on Industrial Informatics*, vol. 17, no. 3, 2021.
- [55] H. Jin, L. Jia, and Z. Zhou, “Boosting Edge Intelligence With Collaborative Cross-Edge Analytics”, *IEEE Internet of Things Journal*, vol. 8, no. 4, 2021.
- [56] J. Romero-Ayala, A. Garcia-Saavedra, X. Costa-Perez, and G. Iosifidis, “EdgeBOL: Automating Energy Savings for Mobile Edge AI”, in *Proc. of ACM CoNEXT*, 2021.
- [57] L. Liu, and J. Deng, “Dynamic DNNs: Optimizing Accuracy-Efficiency Trade-offs by Selective Execution”, *arXiv:1701.00299*, Mar. 2018.
- [58] H. Uzawa, “Iterative Methods in Concave Programming”, in *Studies in Linear and Nonlinear Programming*, K. Arrow, L. Hurwicz, and H. Uzawa, eds., Stanford University Press, pp. 154-165, 1958.
- [59] W. Hoeffding, “Probability Inequalities for Sums of Bounded Random Variables”, *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13-30, 1963.

IX. APPENDIX

Proof of Lemma 1: For any $\boldsymbol{\theta} \in \mathbb{R}_+^{N+1}$ we can write:

$$\begin{aligned} \|\boldsymbol{\lambda}_{t+1} - \boldsymbol{\theta}\|^2 &= \|[\boldsymbol{\lambda}_t + \alpha_t g_t(\mathbf{y}_t)]^+ - \boldsymbol{\theta}\|^2 \leq \|\boldsymbol{\lambda}_t + \alpha_t g_t(\mathbf{y}_t) - \boldsymbol{\theta}\|^2 \\ &= \|\boldsymbol{\lambda}_t - \boldsymbol{\theta}\|^2 + \alpha_t^2 \|g_t(\mathbf{y}_t)\|^2 + 2\alpha_t (\boldsymbol{\lambda}_t - \boldsymbol{\theta})^\top g_t(\mathbf{y}_t), \end{aligned}$$

where we used the non-expansiveness property of the Euclidean projection. Rearranging:

$$\|\boldsymbol{\lambda}_{t+1} - \boldsymbol{\theta}\|^2 - \|\boldsymbol{\lambda}_t - \boldsymbol{\theta}\|^2 \leq \alpha_t^2 \|g_t(\mathbf{y}_t)\|^2 + 2\alpha_t (\boldsymbol{\lambda}_t - \boldsymbol{\theta})^\top g_t(\mathbf{y}_t).$$

Dividing with a_t , setting $\boldsymbol{\theta} = \mathbf{0}$, and applying the telescopic summation we obtain the final result.

Proof of Lemma 2: We have $\boldsymbol{\lambda}_{t+1} = [\boldsymbol{\lambda}_t + a_t g_t(\mathbf{y}_t)]^+ \succeq \boldsymbol{\lambda}_t + a_t g_t(\mathbf{y}_t)$, and dividing by a_t we get:

$$\frac{\boldsymbol{\lambda}_{t+1}}{a_t} - \frac{\boldsymbol{\lambda}_t}{a_t} \succeq g_t(\mathbf{y}_t).$$

Summing telescopically for the first T slots and setting $\boldsymbol{\lambda}_1 = \mathbf{0}$, we obtain:

$$\sum_{t=1}^T g_t(\mathbf{y}_t) \preceq \frac{\boldsymbol{\lambda}_{T+1}}{a_T} + \sum_{t=1}^T \boldsymbol{\lambda}_t \left(\frac{1}{a_{t-1}} - \frac{1}{a_t} \right).$$

Expanding $g_t(\mathbf{y}_t) = g(\mathbf{y}_t) + \delta_t(\mathbf{y}_t)$, dividing with T and taking the norms yields the result.

Proof of Lemma 3: Recall that we defined: $L_t(\mathbf{y}, \boldsymbol{\lambda}) =$

$$f_t(\mathbf{y}) + \boldsymbol{\lambda}^\top g_t(\mathbf{y}) = f(\mathbf{y}) + \boldsymbol{\lambda}^\top g(\mathbf{y}) + \epsilon_t(\mathbf{y}) + \boldsymbol{\lambda}^\top \delta_t(\mathbf{y}). \quad (17)$$

Next, we bound the t -slot dual function $V_t(\boldsymbol{\lambda}_t) = \min_{\mathbf{y} \in \mathcal{Y}} L_t(\mathbf{y}, \boldsymbol{\lambda}_t)$ in terms of the dual function of problem \mathbb{P} , $V(\boldsymbol{\lambda}_t) = \min_{\mathbf{y} \in \mathcal{Y}} f(\mathbf{y}) + \boldsymbol{\lambda}_t^\top g(\mathbf{y})$. Since $\mathbf{y}_t \in \arg \min_{\mathbf{y} \in \mathcal{Y}} L_t(\mathbf{y}, \boldsymbol{\lambda}_t)$, we have:

$$\begin{aligned} V_t(\boldsymbol{\lambda}_t) &= f(\mathbf{y}_t) + \boldsymbol{\lambda}_t^\top g(\mathbf{y}_t) + \epsilon_t(\mathbf{y}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{y}_t) \\ &\stackrel{(a)}{\leq} f(\mathbf{z}_t) + \boldsymbol{\lambda}_t^\top g(\mathbf{z}_t) + \epsilon_t(\mathbf{z}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \\ &= V(\boldsymbol{\lambda}_t) + \epsilon_t(\mathbf{z}_t) + \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \end{aligned}$$

where (a) follows from the minimality of \mathbf{y}_t . Hence:

$$\begin{aligned} f(\mathbf{y}^*) &= V(\boldsymbol{\lambda}^*) \stackrel{(a)}{\geq} \frac{1}{T} \sum_{t=1}^T V(\boldsymbol{\lambda}_t) \\ &\geq \frac{1}{T} \sum_{t=1}^T V_t(\boldsymbol{\lambda}_t) - \epsilon_t(\mathbf{z}_t) - \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \\ &\stackrel{(b)}{\geq} \frac{1}{T} \sum_{t=1}^T \left(L_t(\mathbf{y}_t, \boldsymbol{\lambda}_t) - \epsilon_t(\mathbf{z}_t) - \boldsymbol{\lambda}_t^\top \delta_t(\mathbf{z}_t) \right) \end{aligned}$$

where (a) follows from the maximality of $\boldsymbol{\lambda}^*$ and (b) due to our primal update.

Proof of Lemma 4: $\forall \boldsymbol{\lambda} \in \mathcal{Q}_v$ we have:

$$\begin{aligned} v \leq V_t(\boldsymbol{\lambda}) &= \min_{\mathbf{y} \in \mathcal{Y}} \{ f_t(\mathbf{y}) + \boldsymbol{\lambda}^\top g_t(\mathbf{y}) \} \leq f_t(\mathbf{y}_s) + \boldsymbol{\lambda}^\top g_t(\mathbf{y}_s) \\ &= f_t(\mathbf{y}_s) + \sum_{n=1}^{N+1} \lambda_n g_{nt}(\mathbf{y}_s) \end{aligned}$$

Hence it holds: $-\sum_{n=1}^{N+1} \lambda_n g_{nt}(\mathbf{y}_s) \leq f_t(\mathbf{y}_s) - v$. Since

$g_{nt}(\mathbf{y}_s) < 0$, and $\lambda_n \geq 0$ we get:

$$\min_n \{-g_{nt}(\mathbf{y}_s)\} \sum_{n=1}^{N+1} \lambda_n \leq f_t(\mathbf{y}_s) - v \Rightarrow \sum_{n=1}^{N+1} \lambda_n \leq \frac{f_t(\mathbf{y}_s) - v}{q_t}.$$

Using that $|f_t(\mathbf{y}_s)| \leq \sigma_f$, the definition of q and the fact that $q_t > 0$, we arrive at the result.

Proof of Lemma 5: We use an induction argument to show:

$$\begin{aligned} \|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\| &\leq \lambda_{max} := \frac{2\sigma_f + \|\boldsymbol{\lambda}\|\sigma_g}{q} + \frac{\sigma_g^2}{2q} + \frac{\epsilon}{q} + \|\boldsymbol{\lambda}\| \quad (18) \\ &\quad + \|\boldsymbol{\lambda}_1\| + a\sigma_g, \quad \forall \boldsymbol{\lambda} \succeq \mathbf{0}. \end{aligned}$$

Trivially, $\|\boldsymbol{\lambda}_1 - \boldsymbol{\lambda}\| \leq \|\boldsymbol{\lambda}\| + \|\boldsymbol{\lambda}_1\| \leq \lambda_{max}$, and assume (18) holds at t . We consider two cases.

Case (i): $V_t(\boldsymbol{\lambda}_t) < V_t(\boldsymbol{\lambda}) - \frac{a_t \sigma_g^2}{2}$. Then we can write: $2a_t(V_t(\boldsymbol{\lambda}_t) - V_t(\boldsymbol{\lambda})) < -a_t^2 \sigma_g^2 \Rightarrow -2a_t(V_t(\boldsymbol{\lambda}) - V_t(\boldsymbol{\lambda}_t)) + a_t^2 \sigma_g^2 < 0$. Hence, we have:

$$\begin{aligned} \|\boldsymbol{\lambda}_{t+1} - \boldsymbol{\lambda}\|^2 &\leq \|\boldsymbol{\lambda}_t + a_t g_t(\mathbf{y}_t) - \boldsymbol{\lambda}\|^2 \\ &\leq \|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\|^2 + 2a_t g_t(\mathbf{y}_t)^\top (\boldsymbol{\lambda}_t - \boldsymbol{\lambda}) + a_t^2 \sigma_g^2 \\ &\stackrel{(a)}{\leq} \|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\|^2 - 2a_t (V_t(\boldsymbol{\lambda}) - V_t(\boldsymbol{\lambda}_t)) + a_t^2 \sigma_g^2 \\ &\stackrel{(b)}{<} \|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\|^2, \end{aligned}$$

where (a) follows from the fact that $g_t(\mathbf{y}_t)$ is a subgradient of $V_t(\boldsymbol{\lambda}_t)$ and (b) from the assumptions of the case considered.

Hence, it holds $\|\boldsymbol{\lambda}_{t+1} - \boldsymbol{\lambda}\| \leq \lambda_{max}$.

Case (ii): $V_t(\boldsymbol{\lambda}_t) \geq V_t(\boldsymbol{\lambda}) - \frac{a_t \sigma_g^2}{2}$.

$$\begin{aligned} \|\boldsymbol{\lambda}_{t+1} - \boldsymbol{\lambda}\| &= \|[\boldsymbol{\lambda}_t + a_t g_t(\mathbf{y}_t)]^+ - \boldsymbol{\lambda}\| \leq \|\boldsymbol{\lambda}_t + a_t g_t(\mathbf{y}_t) - \boldsymbol{\lambda}\| \\ &\leq \|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\| + a_t \sigma_g \\ &\leq \|\boldsymbol{\lambda}_t\| + \|\boldsymbol{\lambda}\| + a_t \sigma_g \leq \sum_{n=1}^{N+1} \lambda_{nt} + \|\boldsymbol{\lambda}\| + a_t \sigma_g \\ &\leq \frac{\sigma_f - v}{q} + \|\boldsymbol{\lambda}\| + a_t \sigma_g \\ &\leq \frac{\sigma_f}{q} - \frac{1}{q} \left(\sigma_f + \|\boldsymbol{\lambda}\|\sigma_g - \frac{a_t \sigma_g^2}{2} \right) + \|\boldsymbol{\lambda}\| + a_t \sigma_g \\ &\leq -\frac{\|\boldsymbol{\lambda}\|\sigma_g}{q} + \frac{a_t \sigma_g^2}{2q} + \|\boldsymbol{\lambda}\| + a_t \sigma_g \\ &\leq -\frac{\|\boldsymbol{\lambda}\|\sigma_g}{q} + \frac{a_1 \sigma_g^2}{2q} + \|\boldsymbol{\lambda}\| + a_1 \sigma_g \triangleq \lambda_{max} \end{aligned}$$

where we used that, by Holders inequality, $\|\boldsymbol{\lambda}_t\| \leq \sum_{n=1}^{N+1} \lambda_{nt}$; and applied Lemma 4 with:

$$v = V_t(\boldsymbol{\lambda}) - \frac{a_t \sigma_g^2}{2} \leq \sigma_f + \|\boldsymbol{\lambda}\|\sigma_g - \frac{a_t \sigma_g^2}{2}$$

and used that $\alpha_1 \geq \alpha_t, \forall t$. It follows that $\|\boldsymbol{\lambda}_t - \boldsymbol{\lambda}\| \leq \lambda_{max}$ and so $\|\boldsymbol{\lambda}_t\| \leq \lambda_{max} + \|\boldsymbol{\lambda}\|$.



Apostolos Galanopoulos is currently with Genesys in Ireland. He received the Diploma and M.Sc. degrees from the Department of Electrical and Computer Engineering, University of Thessaly, Volos, Greece in 2014 and 2016, respectively. He has been a Researcher with the Industrial Systems Institute, Athena Research and Innovation Center from 2014 until 2016, and obtained his PhD from Trinity College Dublin, Republic of Ireland in 2021. His research topics of interest are wireless communications, edge computing networks, and application of optimization theory and machine learning in communication networks.



George Iosifidis is an Assistant Professor with Delft University of Technology. He received a Diploma in Electronics and Telecommunications Engineering from the Greek Air Force Academy (Athens, 2000), and a PhD degree from University of Thessaly in 2012. He was an Assistant Prof. with Trinity College Dublin from 2016 until 2020. His research interests lie in the broad area of network optimization and economics; more information can be found at www.FutureNetworksLab.net.



Theodoros Salonidis (S'98 - M'04 - SM'17) received the Diploma degree in Electronic and Computer Engineering from the Technical University of Crete, Greece, in 1997, and M.Sc./Ph.D. in Electrical Engineering from the University of Maryland, College Park, in 1999 and 2004, respectively. He was a Post-Doctoral Researcher at Rice University from 2004 to 2006, a Researcher at Intel Research, Cambridge, U.K., in 2006, and a Researcher at Thomson/Technicolor, Paris, France, from 2007 to 2012. He is a Research Staff Member at the IBM

Thomas J. Watson Research Center, Yorktown Heights, NY, USA. His current research interests are in the areas of automated AI, distributed analytics, machine learning, and performance analysis, design, and implementation of cloud, edge and mobile computing systems.



Doug Leith (SM'01) graduated from the University of Glasgow in 1986 and was awarded his PhD, also from the University of Glasgow, in 1989. Prof. Leith moved to the National University of Ireland, Maynooth in 2001 to establish the Hamilton Institute (www.hamilton.ie) of which he was founding Director from 2001-2014. Towards the end of 2014, Prof Leith moved to Trinity College Dublin to take up the Chair in Computer Systems in the School of Computer Science and Statistics. His current research interests include wireless networks, congestion control, optimization and data privacy.