

Improving IoT Analytics through Selective Edge Execution

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

*School of Computer Science and Statistics, Trinity College Dublin

†Department of Electronic and Electrical Engineering, University College London

‡IBM T. J. Watson Research Center, New York

Abstract—A large number of emerging IoT applications rely on machine learning routines for analyzing data. Executing such tasks at the user devices improves response time and economizes network resources. However, due to power and computing limitations, the devices often cannot support such resource-intensive routines and fail to accurately execute the analytics. In this work, we propose to improve the performance of analytics by leveraging edge infrastructure. We devise an algorithm that enables the IoT devices to execute their routines locally; and then outsource them to cloudlet servers, only if they predict they will gain a significant performance improvement. It uses an approximate dual subgradient method, making minimal assumptions about the statistical properties of the system’s parameters. Our analysis demonstrates that our proposed algorithm can intelligently leverage the cloudlet, adapting to the service requirements.

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

I. INTRODUCTION

The recent demand for machine learning (ML) applications, such as image recognition, natural language translation, and health monitoring, has been unprecedented [1]. These services collect data streams generated by small devices, and analyze them locally or at distant cloud servers. There is growing consensus that such applications will be ubiquitous in Internet of Things (IoT) systems [2]. The challenge, however, with such services is that they are often resource intensive. On the one hand, the cloud offers powerful ML models and abundant compute resources but requires data transfers which consume network bandwidth and might induce significant delays [3]. On the other hand, executing these services at the devices economizes bandwidth but degrades their performance due to the devices’ limited resources, e.g. memory or energy.

A promising approach to tackle this problem is to allow the devices to outsource individual ML tasks to edge infrastructure such as cloudlets [4]. This can increase their execution accuracy since the cloudlet’s ML components are typically more complex, and hence offer improved results. Nevertheless, the success of such solutions presumes intelligent outsourcing algorithms. The cloudlets, unlike the cloud, have limited computing capacity and cannot support all requests. At the same time, task execution requires the transfer of large data volumes (e.g., video streams). This calls for prudent transmission decisions in order to avoid wasting device energy and bandwidth. Furthermore, unlike prior computation offloading

solutions [5], it is crucial to only outsource the tasks that can significantly benefit from cloudlet execution. Finally, these decisions need to be made in a dynamic fashion, accounting for the time-varying conditions of the system.

Our goal is to design an online framework that addresses the above issues and makes intelligent outsourcing decisions. We consider a system where a cloudlet improves the execution of analytics running on devices such as wireless IoT cameras. We focus on classification, but our study can be readily extended to other tasks. We assume that each device has a “low-precision” classifier while the cloudlet can execute the task with higher precision. The devices classify the received objects upon arrival, and decide whether to transmit them to the cloudlet or not, to get a better classification result. Making this decision requires an assessment of the potential performance gains, which are measured in terms of accuracy improvements. To this end, we propose the usage of a *predictor* at each device that leverages the local classification results.

We consider the practical case where the resources’ availability is unknown and time-varying, but their instantaneous values are observable. We design a distributed adaptive algorithm that decides the task outsourcing policy towards maximizing the long-term performance of analytics. To achieve this, we formulate the system’s operation as an optimization problem, which is decomposed via Lagrange relaxation to a set of device-specific problems. This enables its distributed solution through an *approximate* – due to the unknown parameters – dual ascent method, that can be applied in real time. Our method is inspired by primal averaging schemes for *static* problems, e.g., see [6], and achieves a *bounded and tunable* optimality gap using a novel approximate iteration technique. The contribution of this work can be summarized as follows:

- *Edge Analytics.* We study the novel problem of intelligently improving data analytics tasks using edge infrastructure, which is increasingly important for the IoT.
- *Decision Framework.* We propose an *online* task outsourcing algorithm that achieves near-optimal performance under very general conditions (unknown, non i.i.d. statistics). This is a novel analytical result of independent value.
- *Implementation & Evaluation.* The solution is evaluated in a wireless testbed using a real ML application, several classifiers and datasets. We find that our algorithm increases

the accuracy (up to 32%) and reduces the energy (down to 60%) compared to carefully selected benchmark policies.

Organization. Sec. II introduces the model and the problem. Sec. III presents the algorithm and Sec. IV the system implementation, experiments and trace-driven simulations. We discuss related work in Sec. V and conclude in Sec. VI. Although the paper is completely self-sufficient, the interested reader will find more results from the implementation of our system, as well as a more detailed version of the proof of our main analytical contribution in [7].

II. MODEL AND PROBLEM FORMULATION

Classifiers. There is a set \mathcal{C} of C disjoint object classes and a set \mathcal{N} of N edge devices. We assume a time-slotted operation where each device n receives at slot t a group of objects (or tasks) \mathcal{S}_{nt} to be classified, e.g., frames captured by its camera. We define $\mathcal{S}_n \supseteq \mathcal{S}_{nt}, \forall t$ as the set of objects that can arrive at n , and $\mathcal{S} = \cup \{\mathcal{S}_n\}_n$. Each device n is equipped with a *local classifier* $J_n: \mathcal{S}_n \rightarrow (\mathcal{C}_n, d_n(s_{nt}))$, which outputs the inferred class of an object s_{nt} and a normalized confidence value $d_n(s_{nt}) \in [0, 1]$ for that inference¹. The cloudlet has a classifier $J_0: \mathcal{S} \rightarrow (\mathcal{C}_0, d_0(s_{nt}))$ that can classify any object, and offers higher accuracy from all devices, i.e., $d_0(s_{nt}) \geq d_n(s_{nt}), \forall n \in \mathcal{N}$.

Let $\phi_{nt} \in [0, 1]$ denote the accuracy improvement when the cloudlet classifier is used:

$$\phi_{nt}(s_{nt}) = d_0(s_{nt}) - d_n(s_{nt}), \quad \forall n \in \mathcal{N}, s_{nt} \in \mathcal{S}_{nt}. \quad (1)$$

Every device is also equipped with a predictor² Q_n that is trained with the outcomes of the local and cloudlet classifiers. This predictor can estimate the accuracy improvement offered by the cloudlet for each object $s_{nt} \in \mathcal{S}_{nt}$:

$$Q_n: (J_n(s_{nt})) \rightarrow (\hat{\phi}_{nt}, \sigma_{nt}), \quad (2)$$

and, in general, this assessment might be inexact, $\hat{\phi}_{nt}(s_{nt}) \neq \phi_{nt}(s_{nt})$, and $\sigma_{nt} \in [0, 1]$ is the respective confidence value.

Wireless System. The devices access the cloudlet through high capacity cellular or Wi-Fi links. Each device n has an *average power budget* of B_n Watts. Power is a key limitation here because the devices might have a small energy budget due to protocol-induced transmission constraints, or due to user aversion for energy spending. The cloudlet has an *average processing capacity* of H cycles/sec which is shared by the devices, and when the total load exceeds H , the task delay increases and eventually renders the system non-responsive.

We consider the realistic scenario where the parameters of devices and the cloudlet change over time in an unknown fashion. Namely, they are created by random processes $\{B_{nt}\}_{t=1}^{\infty}$ and $\{H_t\}_{t=1}^{\infty}$, and our decision framework has access only to their instantaneous values in each slot. Unlike previous optimization frameworks [9] that assume i.i.d., or Markov modulated processes; here we only ask that these perturbations

¹The classifier might output only the class with the highest confidence, or a vector with the confidence for each class; our analysis holds for both cases.

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, see [8].

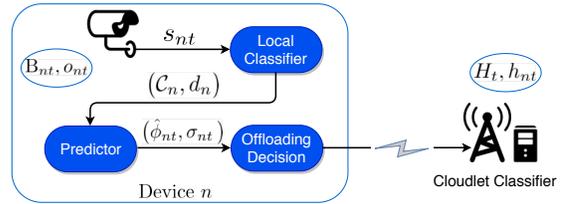


Fig. 1: Schematic of the basic notation and procedure followed by the system's devices.

are bounded in each slot, i.e. $H_t \leq H_{max}, B_{nt} \leq B_{max}, \forall t$ and their averages converge to some finite values which we do need to know, i.e., $\lim_{t \rightarrow \infty} \sum_{\tau=1}^t B_{nt}/t = B_n, \forall n$, and similarly for $\{H_t\}_{t=1}^{\infty}$. We also define $\mathbf{B}_t = (B_{nt}, n \in \mathcal{N})$.

When an object (say, image) is transmitted in slot t from device n to the cloudlet, it consumes³ part of the device's power budget B_n . We assume that this cost, denoted o_{nt} , follows a random process $\{o_{nt}\}_{t=1}^{\infty}$ that is uniformly upper-bounded and has well-defined mean values.⁴ Also, each transmitted object requires a number of processing cycles in the cloudlet which might also vary with time, e.g., due to the different type of the objects, and we assume it follows the random process $\{h_{nt}\}_{t=1}^{\infty}$, with $\lim_{t \rightarrow \infty} \sum_{\tau=1}^t h_{nt}/t = h_n$. We define $\mathbf{o}_t = (o_{nt} \leq o_{max}, n \in \mathcal{N})$, and $\mathbf{h}_t = (h_{nt} \leq h_{max}, n \in \mathcal{N})$. Our model is very general as the (i) requests, (ii) power and computing cost per request, and (iii) resource availability, can be arbitrarily time-varying, and with unknown statistics.

Problem Formulation. The IoT devices wish to involve the cloudlet only when they confidently expect high classification precision gains. Otherwise, they will consume the cloudlet's capacity and their own power without significant performance benefits. Therefore, we make the outsourcing decision for each object s_{nt} based on the *weighted improvement gain*:

$$w_{nt}(s_{nt}) = \hat{\phi}_{nt} - \rho_n \sigma_{nt}, \quad \forall n, t, \quad (3)$$

where $\rho_n \geq 0$ is a *risk aversion* parameter set by the system designer or each user. For example, assuming normal distribution for ϕ_{nt} , we could set $\rho_n = 1$ and use a threshold rule of 1 standard deviation. We use hereafter these modified parameters $w_{nt}, \forall n$, and partition the interval of their values $[-w_0, w_0]$ (w_0 being the maximum) into subintervals $\mathcal{I}^j, j = 1, \dots, M$ such that $\cup_{j=1}^M \mathcal{I}^j = [-w_0, w_0], \forall i \neq j$; with w_n^j being the center point of \mathcal{I}^j . This quantization facilitates the implementation of our algorithm in a real system, and is without loss of generality since we can use very short intervals. Finally, let λ_{nt}^j denote the number of objects with expected gain w_n^j that device n has created in slot t . These arrivals are generated by an unknown process $\{\lambda_{nt}^j\}_{t=1}^{\infty}$, with $\lim_{T \rightarrow \infty} 1/T \sum_{t=1}^T \lambda_{nt}^j = \lambda_n^j, \forall n, j$.

Our aim is to maximize the aggregate long-term analytics performance gains, for all objects and IoT devices. This can be formulated as a mathematical program. We define variables $y_n^j \in [0, 1], \forall n, j$ which indicate the long term ratio of objects

³Power budgets are also affected by the local classifier computations which are made for every object and thus do not affect the offloading decisions.

⁴This cost can reflect, e.g., the impact of time-varying channel conditions.

with expected gain of w_n^j that are sent to the cloudlet (with $y_n^j = 1$, when all objects of n in \mathcal{I}^j are sent), and formulate the convex problem:

$$\mathbb{P}_1 : \quad \underset{y_n^j \in [0,1]}{\text{maximize}} \quad \sum_{j=1}^M \sum_{n=1}^N w_n^j \lambda_n^j y_n^j \triangleq f(\mathbf{y}) \quad (4a)$$

$$s.t. \quad \sum_{j=1}^M y_n^j \lambda_n^j o_n \leq B_n, \quad n \in \mathcal{N}, \quad (4b)$$

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

where $\mathbf{y} = (y_n^j : \forall n, j)$. Eq. (4b) constrains the average power budget of each device and (4c) bounds the cloudlet utilization. Clearly, based on the specifics of each system we can add more constraints, e.g., for the average wireless link capacity in case bandwidth is also a bottleneck resource. Such extensions are straightforward as they do not change the properties of the problem, nor affect our analysis below.

The solution of \mathbb{P}_1 is a policy \mathbf{y}^* that maximizes the aggregate (hence also average) analytics performance in the system. Such policies can be randomized, with y_n^{j*} denoting the probability of sending each object of n in interval \mathcal{I}^j to the cloudlet (at each slot). However, in reality, the system parameters not only change with time, but are generated by processes that might not be i.i.d. and have unknown statistics (mean values, etc.). This means that in practice we cannot find \mathbf{y}^* . In the next section we present an online policy that is oblivious to the statistics of $\{\lambda_t\}$, $\{o_t\}$, $\{h_t\}$, $\{H_t\}$, $\{B_t\}$ but achieves indeed the same performance with \mathbf{y}^* .

III. ONLINE OFFLOADING ALGORITHM

Our solution approach is simple and, we believe, elegant. We replace the unknown parameters H , λ_n^j , B_n , o_n and h_n , $\forall n, j$ in \mathbb{P}_1 with their running averages (which we calculate as the system operates), solve the modified problem with gradient ascent in the dual space, and perform primal averaging. This gives us an online policy that applies in real time the solution $\mathbf{y}_t, \forall t$, while using only information made available by slot t .

A. Problem Decomposition & Algorithm Design

Let us first define the running-average function:

$$\begin{aligned} \bar{f}_t(\mathbf{y}) &\triangleq \sum_{j=1}^M \sum_{n=1}^N w_n^j y_n^j \bar{\lambda}_{nt}^j \\ &= \sum_{j=1}^M \sum_{n=1}^N w_n^j \lambda_n^j y_n^j + \sum_{j=1}^M \sum_{n=1}^N w_n^j y_n^j (\lambda_n^j - \bar{\lambda}_{nt}^j) \\ &= f(\mathbf{y}) + \mathbf{y}^\top \boldsymbol{\epsilon}_t, \end{aligned}$$

where $\bar{\lambda}_{nt}^j = \sum_{\tau=1}^t \lambda_{n\tau}^j / t$ is the running average of λ_n^j , and $\boldsymbol{\epsilon}_t = (w_n^j (\lambda_n^j - \bar{\lambda}_{nt}^j), \forall n, j) \in \mathbb{R}^{NM}$ is the vector of component-wise errors between $\bar{f}_t(\mathbf{y})$ and $f(\mathbf{y})$. Also, we denote $g(\mathbf{y}) \in \mathbb{R}^{N+1}$ the constraint vector of (4b)-(4c), and define

$$\bar{g}_t(\mathbf{y}) = g(\mathbf{y}) + \delta_t(\mathbf{y}), \quad (5)$$

with $\delta_t(\mathbf{y}) = (\delta_{nt}(\mathbf{y}), n = 1, \dots, N+1)$ and

$$\begin{aligned} \delta_{nt}(\mathbf{y}) &= B_n - \bar{B}_{nt} + \sum_{j=1}^M y_n^j (\bar{o}_{nt} \bar{\lambda}_{nt}^j - o_n \lambda_n^j), \quad n = 1, \dots, N, \\ \delta_{N+1,t}(\mathbf{y}) &= H - \bar{H}_t + \sum_{j=1}^M \sum_{n=1}^N y_n^j (\bar{h}_{nt} \bar{\lambda}_{nt}^j - h_n \lambda_n^j). \end{aligned}$$

$\bar{B}_{nt} = \sum_{\tau=1}^t B_{n\tau} / t$ is the running average of process $\{B_{nt}\}_{t=1}^\infty$, and similarly we define \bar{H}_t , \bar{o}_{nt} , and \bar{h}_{nt} . Note that $\bar{f}_t(\mathbf{y})$, $\bar{g}_t(\mathbf{y})$ can be calculated at each slot, while $f(\mathbf{y})$ and $g(\mathbf{y})$ are unknown. We can now define a new problem:

$$\mathbb{P}_2(t) : \quad \max_{\mathbf{y} \in [0,1]^{NM}} \bar{f}_t(\mathbf{y}) \quad s.t. \quad \bar{g}_t(\mathbf{y}) \leq 0$$

We will use the instances $\{\mathbb{P}_2(t)\}_t$ to perform a dual ascent method and obtain a sequence of decisions $\{\mathbf{y}\}_t$ that will be applied in real time and achieve performance that converges asymptotically to the (unknown) solution of \mathbb{P}_1 .

We first dualize $\mathbb{P}_2(t)$ and introduce the Lagrangian⁵:

$$\begin{aligned} L(\mathbf{y}, \boldsymbol{\mu}) &\triangleq \bar{f}_t(\mathbf{y}) + \boldsymbol{\mu}^\top \bar{g}_t(\mathbf{y}) = \sum_{j=1}^M \sum_{n=1}^N w_n^j y_n^j \bar{\lambda}_{nt}^j + \\ &\sum_{n=1}^N \mu_n \left(\sum_{j=1}^M y_n^j \bar{\lambda}_{nt}^j \bar{o}_{nt} - \bar{B}_{nt} \right) + \xi \left(\sum_{j=1}^M \sum_{n=1}^N y_n^j \bar{\lambda}_{nt}^j \bar{h}_{nt} - \bar{H}_t \right) \end{aligned}$$

where $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_N, \xi)$ are the non-negative dual variables for $\bar{g}_t(\mathbf{y}) \leq 0$. The dual function is:

$$V(\boldsymbol{\mu}) = \arg \min_{\mathbf{0} \leq \mathbf{y} \leq 1} L(\mathbf{y}, \boldsymbol{\mu}), \quad (6)$$

and the dual problem amounts to maximizing $V(\boldsymbol{\mu})$.

We apply a dual ascent algorithm where the iterations are in sync with the system's time slots t . Observe that $V(\boldsymbol{\mu})$ does not depend on \bar{B}_{nt} or \bar{H}_t , it is separable with respect to the primal variables, and independent of $\bar{\lambda}_{nt}^j$. Hence, in each iteration t we can minimize L by:

$$(y_n^j)^* \in \arg \min_{y_n^j \in [0,1]} y_n^j (-w_n^j + \mu_{nt} \bar{o}_{nt} + \xi_t \bar{h}_{nt}), \quad \forall n, j. \quad (7)$$

This yields the following easy-to-implement threshold rule:

$$y_{nt}^j = \begin{cases} 1 & \text{if } \lambda_{nt}^j > 0 \quad \text{and} \quad \mu_{nt} \bar{o}_{nt} + \xi_t \bar{h}_{nt} < w_n^j \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

which is a deterministic decision that offloads (or not) all requests of each device (at each t). Then we improve the current value of $V_t(\boldsymbol{\mu})$ by updating the dual variables:

$$\mu_{n,t+1} = \left[\mu_{nt} + \alpha \left(\sum_{j=1}^M \bar{o}_{nt} \bar{\lambda}_{nt}^j y_{nt}^j - \bar{B}_{nt} \right) \right]^+, \quad \forall n, \quad (9)$$

$$\xi_{t+1} = \left[\xi_t + \alpha \left(\sum_{n=1}^N \sum_{j=1}^M \bar{h}_{nt} \bar{\lambda}_{nt}^j y_{nt}^j - \bar{H}_t \right) \right]^+, \quad (10)$$

⁵For our system implementation, this relaxation means we install queues for the data transmission (at the devices) and image processing (at the cloudlet).

where $\alpha > 0$ is the update step size, and return to (7).

Algorithm 1: OnAlgo

```

1: Initialization:  $t = 0, \xi_0 = 0, \mu_0 = \mathbf{0}, \mathbf{y} = \mathbf{0}$ 
2: while True do
3:   for each device  $n \in \mathcal{N}$  do
4:     Receive objects  $\mathcal{S}_{nt} = \{s_{nt}\}$ ;
5:      $\hat{\phi}_{nt}, \sigma_{nt} \leftarrow Q_n(J_n(s_{nt})), \forall s_{nt} \in \mathcal{S}_{nt}$ 
6:     Calculate  $w_{nt}$  through (3);
7:     Observe  $o_{nt}, h_{nt}, B_{nt}$  and calculate  $\bar{o}_{nt}, \bar{h}_{nt}, \bar{B}_{nt}$ ;
8:     for  $j = 1, \dots, M$  do
9:       Observe  $\lambda_{nt}^j$  and calculate average  $\bar{\lambda}_{nt}^j$  and  $w_n^j$ ;
10:      Decide  $y_n^j$  by using (8);
11:    end for
12:    Update  $\mu_{n,t+1}$  using (9);
13:    Send averages  $\bar{\lambda}_{nt}^j, \forall j$ , to cloudlet;
14:  end for
15:  Cloudlet:
16:  Compute tasks and receive  $\bar{\lambda}_{nt}^j, \forall n$ ;
17:  Observe  $H_t$  and calculate  $\bar{H}_t$ ;
18:  Update  $\xi_{t+1}$  using (10), and send it to devices;
19:   $t \leftarrow t + 1$ ;
20: end while

```

The detailed steps that implement our online policy are as follows (with reference to *OnAlgo*, Algorithm 1). Each device n receives a group of objects \mathcal{S}_{nt} in slot t and uses its classifier to predict their classes, and the predictor to estimate the expected offloading gains (Steps 4-6). The device updates its statistics (step 7) and compares the expected benefits for each interval j with the outsourcing costs for the device and the cloudlet (Step 10). The devices update their local dual variable for the power constraint violation (Step 12). The cloudlet classifies the received objects (Step 16) and updates its parameter estimates (Step 17) and its congestion (Step 18), which is sent to the devices.

B. Performance Analysis

The gist of our approach is that, as time evolves, the sequence of problems $\{\mathbb{P}_2(t)\}_t$ approaches our initial problem \mathbb{P}_1 . This is true under the following mild assumption.

Assumption 1. *The perturbations of the system parameters are independent to each other, uniformly bounded, and their averages converge, e.g., $\lim_{t \rightarrow \infty} \bar{B}_{nt} = B_n$.*

Under this assumption it is easy to see that it holds:

$$\lim_{t \rightarrow \infty} \delta_t(\mathbf{y}) = 0, \quad \lim_{t \rightarrow \infty} \mathbf{y}^\top \epsilon_t = 0, \quad \forall \mathbf{y}.$$

Furthermore, note that due to boundedness of the parameters and $y_n^j \in [0, 1], \forall n, j$ we have that:

$$\|g(\mathbf{y})\|_2 \leq \sigma_g, \quad \|\delta_t(\mathbf{y})\|_2 \leq \sigma_{\delta_t}, \quad \forall t, \quad (11)$$

and using Minkowski's inequality, we get the bound:

$$\|\bar{g}_t(\mathbf{y})\| = \|g(\mathbf{y}) + \delta_t(\mathbf{y})\|_2 \leq \sigma_g + \sigma_{\delta_t}. \quad (12)$$

It is also easy to see that $\lim_{t \rightarrow \infty} \sigma_{\delta_t} = 0$. The following Theorem is our main analytical result.

Theorem 1. *Under Assumption 1, OnAlgo ensures the following optimality and feasibility gaps:*

$$(i) \lim_{t \rightarrow \infty} f(\bar{\mathbf{y}}_t) \leq f^* + \frac{a\sigma_g^2}{2}, \quad (ii) \lim_{t \rightarrow \infty} g(\bar{\mathbf{y}}_t) \leq 0,$$

where $\bar{\mathbf{y}}_t = \frac{1}{t} \sum_{i=1}^t \mathbf{y}_i$.

Proof. We drop bold typeface notation here, and use subscript $i = 1, \dots, t$ to denote the i -th slot. We first bound the distance of μ_{t+1} from vector $\theta \in \mathbf{R}^{N+1}$, i.e., $\|\mu_{t+1} - \theta\|_2^2 =$

$$\begin{aligned} & \|[\mu_t + a(g(y_t) + \delta_t(y_t))]^+ - \theta\|_2^2 \leq \\ & \|\mu_t - \theta\|_2^2 + a^2 \|g(y_t)\|_2^2 + a^2 \|\delta_t(y_t)\|_2^2 + 2a^2 \delta_t(y_t)^\top g(y_t) + \\ & 2a(\mu_t - \theta)^\top (g(y_t) + \delta_t(y_t)). \end{aligned} \quad (13)$$

(i) **Optimality Gap.** From the dual problem we can write:

$$\begin{aligned} V(\mu^*) & \geq \frac{1}{t} \sum_{i=1}^t V(\mu_i) \geq \frac{1}{t} \sum_{i=1}^t L(y_i, \mu_i) \\ & = \frac{1}{t} \sum_{i=1}^t \left(f(y_i) + y_i^\top \epsilon_i + \mu_i^\top (g(y_i) + \delta_i(y_i)) \right) \\ & \geq f(\bar{y}_t) + \frac{1}{t} \sum_{i=1}^t y_i^\top \epsilon_i + \frac{1}{t} \sum_{i=1}^t \left(\mu_i^\top (g(y_i) + \delta_i(y_i)) \right), \end{aligned} \quad (14)$$

where the last inequality follows from Jensen's inequality. Now, let $\theta = 0$ in (13). Using (11) and the Cauchy-Swartz inequality, and by summing over all t we obtain :

$$\begin{aligned} \|\mu_{t+1}\|_2^2 & \leq \|\mu_1\|_2^2 + a^2 t \sigma_g^2 + a^2 \sum_{i=1}^t \sigma_{\delta_i}^2 + \\ & 2a^2 \sigma_g \sum_{i=1}^t \sigma_{\delta_i} + 2a \sum_{i=1}^t \mu_i^\top (g(y_i) + \delta_i(y_i)). \end{aligned}$$

Dropping the non-negative term $\|\mu_{t+1}\|_2^2$, dividing by $2at$, setting $\mu_1 = 0$, and rearranging terms, yields:

$$-\frac{1}{t} \sum_{i=1}^t \mu_i^\top (g(y_i) + \delta_i(y_i)) \leq \frac{a\sigma_g^2}{2} + \frac{a}{2t} \sum_{i=1}^t \sigma_{\delta_i}^2 + \frac{a\sigma_g}{t} \sum_{i=1}^t \sigma_{\delta_i}.$$

Using the fact that $V(\mu^*) = f^*$, and combining the above with (14), we obtain:

$$f(\bar{y}_t) - f^* \leq -\frac{1}{t} \sum_{i=1}^t y_i^\top \epsilon_i + \frac{a\sigma_g^2}{2} + \frac{a}{2t} \sum_{i=1}^t \sigma_{\delta_i}^2 + \frac{a\sigma_g}{t} \sum_{i=1}^t \sigma_{\delta_i}.$$

All sums have diminishing terms and divided by t , hence converge to 0. Thus, we obtained the first part of the theorem.

(ii) **Constraint Violation.** If we apply recursively the dual variable update rule, we obtain:

$$\mu_{t+1} = \left[\mu_t + a(g(y_t) + \delta_t(y_t)) \right]^+ \succeq \mu_1 + a \sum_{i=1}^t (g(y_i) + \delta_i(y_i)).$$

Setting $\mu_1 = 0$, dividing by at , and using Jensen’s inequality for $g(\cdot)$, we get:

$$g(\bar{y}_t) + \frac{1}{t} \sum_{i=1}^t \delta_i(y_i) \preceq \frac{\mu_{t+1}}{at}. \quad (15)$$

The second term of the LHS converges to zero as $t \rightarrow \infty$. Our claim holds if the same is true for the RHS. Indeed, this is the case assuming the existence of a Slater vector, and the boundedness of the set of dual variables (see [6], [7]). \square

The theorem shows that OnAlgo asymptotically achieves zero feasibility gap (no constraint violation), and a fixed optimality gap that can be made arbitrarily small by tuning the step size.

IV. IMPLEMENTATION AND EVALUATION

A. Experimentation Setup and Initial Measurements

1) *Testbed and Measurements*: We used 4 Raspberry Pis (RPs) as end-nodes, placed in different distances from a laptop (cloudlet). We used a Monsoon monitor for the energy measurements, and Python libraries and TensorFlow for the classifiers.⁶ We first measured the average power consumption when RPs transmit data to the cloudlet with different rates, and then fitted a linear regression model that estimates the consumed power as a function of r . This model is used by OnAlgo to estimate the energy cost for each transmitted image, given the data rate in each slot (which might differ for the RPs). Also, we measured the average computing costs (h_n, h_0 cycles/task) of the classification tasks, to be used in simulations. For more details on the setup, see [7].

2) *Data Sets and Classifiers*: We use two well-known datasets: (i) MNIST [11] which consists of 28×28 pixel handwritten digits, and includes 60K training and 10K test examples; (ii) CIFAR-10 [12] with 50K training and 10K test examples of 32×32 color images of 10 classes. We used two classifiers, the normalized-distance weighted k -nearest neighbors (KNN) [13], and the more sophisticated Convolutional Neural Network (CNN) implemented with TensorFlow [14]. They output a vector with the probabilities that the object belongs to each class. These classifiers have different performance and resource needs, hence allow us to build diverse experiments. The predictors are trained with labeled images and the outputs of the local (f_n) and cloudlet (f_0) classifiers. These are the independent variables in our regression model that estimates ϕ_{nt} (dependent variables). Recall that the latter are calculated using (1), where we additionally use that $w_{nt} = d_0(s_{nt})$ if device n has given a wrong classification and $w_{nt} = -d_0(s_{nt})$ if the cloudlet is mistaken.

3) *Benchmarks*: We compare OnAlgo with two algorithms. The *Accuracy-Threshold Offloading* (ATO) algorithm, where a task is offloaded when the confidence of the local classifier is below a threshold, without considering the resource consumption. And the *Resource-Consumption Offloading* (RCO) algorithm, where a task is offloaded when there is enough energy, without considering the expected classification improvement.

⁶We used *vanilla* versions of the classifiers to facilitate observation of the results. The memory footprint of NNs can be made smaller [10] but this might affect their performance. Our analysis is orthogonal to such interventions.

4) *Limitations of Mobile Devices*: We used our testbed to verify that these small resource-footprint devices require the assistance of a cloudlet. Our findings are in line with previous studies, e.g., [15]. The performance of a CNN model increases with the number of layers. We find that, even with 4 layers, a CNN trained for CIFAR has 1GB size and hence cannot be stored in the RPs (see Fig. 2a). Similar conclusions hold for the KNN classifier that needs to locally store all training samples. Clearly, despite the successful efforts to reduce the size of ML models by, e.g. using compression [10]; the increasingly complex analytics and the small form-factor of devices will continue to raise the local versus cloudlet execution trade off.

5) *Classifier Assessment*: In Fig. 2b we see that the accuracy (ratio of successful over total predictions) of the KNN classifier improves with the size K_n of labeled data. Figure 2c presents the accuracy gains for CNN as more hidden layers are added. The gains are higher (up to 20%) for the digits that are difficult to recognize, e.g., 4 and 5. Fig. 2d shows the CNN performance on CIFAR, which is lower as this is a more complex dataset (colored images, etc.). Overall, we see that the classifier performance depends on the algorithm (KNN, CNN), the settings (datasets, layers), and the objects.

B. Performance Evaluation

1) *Resource Availability Impact*: Fig. 3 shows the average accuracy and fraction of requests offloaded to the cloudlet with OnAlgo when we vary their power budget. As B_n increases there are more opportunities to use the cloudlet (4-layer CNN) and obtain more accurate classifications than the local classifier (1-layer CNN). Furthermore, Fig. 2(c-d) show that MNIST is easier to classify and the gains of using a better classifier are smaller than with CIFAR. Hence, as B_n increases in Fig. 3 the ratio of offloaded tasks increases at a faster pace with CIFAR than with MNIST.

2) *Comparison with Benchmarks*: We compare OnAlgo to ATO and RCO. No-offloading (NO) serves as a baseline for these algorithms in Fig. 4. 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. For RCO, the availability of energy is determined by computing the running average consumption at each device during the experiment. We employ two testbed scenarios, and a simulation with larger number of devices.

Scenario 1: Low accuracy improvement; high resources. We set⁷ $B_n = 0.01mW$ and $H = 2GHz$ allowing the devices to offload many tasks, and the cloudlet to serve most of them; and used MNIST (has small improvement). We demonstrate the average accuracy and power consumption in Fig. 4a, where we see that OnAlgo outperforms both ATO and RCO by 5%. Regarding power consumption, ATO achieves the best result since it gets high enough confidence on its local classifier (rarely offloads). RCO however, offloads almost every task as it has enough resources and does not refrain even when

⁷We 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.

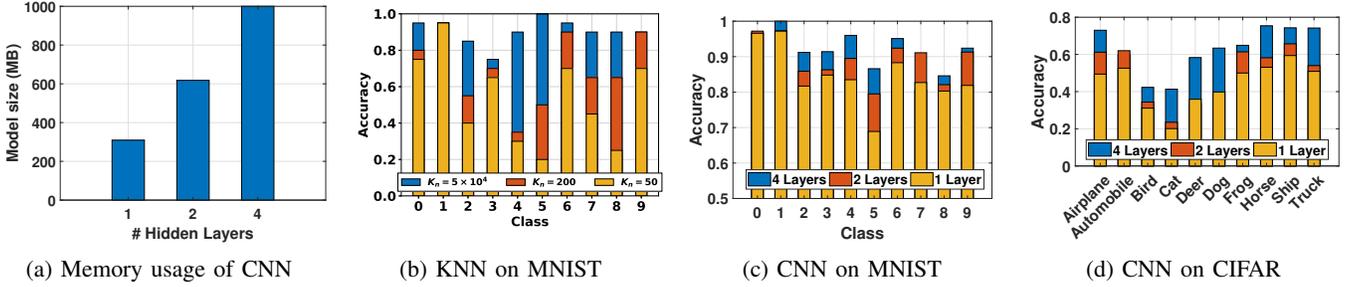


Fig. 2: CNN memory usage vs number of layers, and accuracy of MNIST and CIFAR-10 for KNN and CNN classifiers.

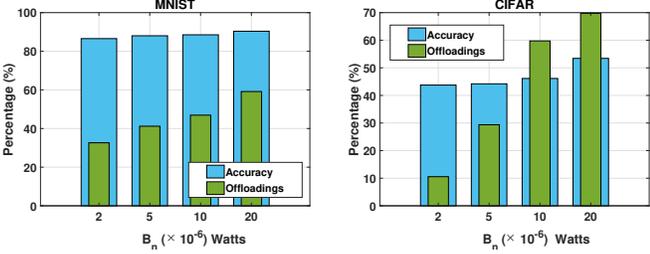


Fig. 3: Average accuracy and outsourcing of OnAlgo for different power budgets, on MNIST (left) and CIFAR (right).

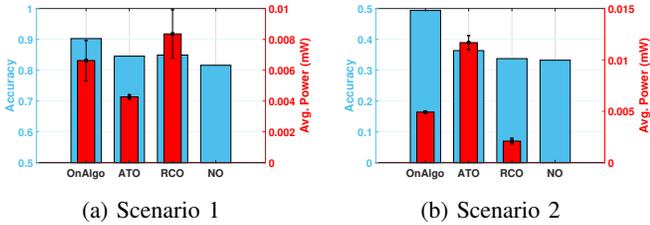


Fig. 4: Performance comparison of the offloading algorithms.

improvement is low. The reason it achieves lower accuracy than onAlgo is that it does not offload intelligently, and gets denied when the computing constraint is violated.

Scenario 2: High accuracy improvement; low resources. We set $B_n = 0.005mW$ and $H = 200MHz$ not allowing many offloadings and cloudlet classifications. We used the CIFAR dataset which has a large performance difference between local and cloudlet classifiers. We see from Fig. 4b that OnAlgo achieves 28%-32% higher accuracy than both competing algorithms. RCO is constrained to very few offloadings due to the limited power budget, while ATO is resource-oblivious and offloads tasks regardless of the cloudlet's capacity. This results in many denied offloadings that reduce ATO's accuracy and unnecessarily increase the power consumption. OnAlgo consumes 60% less power than ATO as it frequently offloads its low-confidence tasks.

Scenario 3: Large number of users. Finally, we simulated the algorithms for a large number of users while using the experimentally measured parameters. We observe in Fig. 5a that the accuracy gradually drops (for all algorithms) since now a smaller percentage of the tasks can be served by the cloudlet. OnAlgo constantly outperforms both ATO and RCO

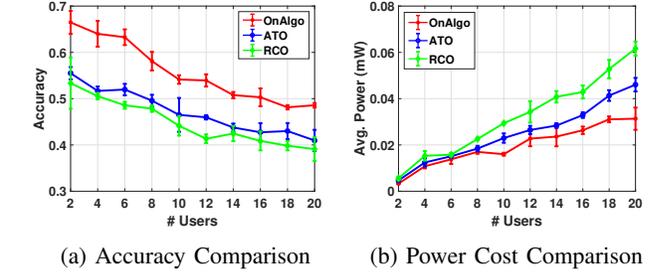


Fig. 5: Simulation results for increasing number of users on the CIFAR dataset. $B_n = 0.02 mW$, $H = 2 GHz$.

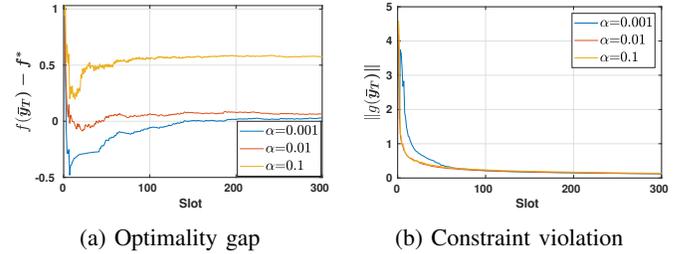


Fig. 6: The convergence properties of OnAlgo when $M = 6$, $N = 5$. Note that constraints are eventually satisfied, with some of them in a strict fashion (hence the norm is not zero).

by about 10%–25% since it adapts to the available resources. This is more evident in Fig. 5b that shows the fast-increasing energy cost of the two benchmark algorithms, as they either offload tasks that do not improve the performance, or offload tasks while the cloudlet is already congested (these tasks are dropped and energy is wasted). Power consumption of OnAlgo is up to 50% less than that of RCO.

3) Convergence of OnAlgo: Fig. 6 presents the convergence of OnAlgo for different step sizes α . Based on the system parameters the bound given by Theorem 1 is approximately 0.01, 0.2 and 1 for the three α values of Fig. 6. These are satisfied by the solution of OnAlgo in less than 300 iterations as observed in Fig. 6a. The convergence is faster for larger α , which however is achieved at the cost of smaller convergence accuracy. The constraint violation bound is also respected as shown in Fig. 6b with the constraints being violated more often for small α in the beginning, but improving as T increases.

V. RELATED WORK

Edge & Distributed Computing. Most solutions partition compute-intensive mobile applications and offload them to the cloud [16]; a solution that is unfit to enable low-latency applications. Cloudlets on the other hand, achieve lower delay [4] but have limited serving capacity, hence there is a need for an intelligent offloading strategy that we propose here. Previous works consider simple performance criteria, such as reducing computation loads [17], or power consumption [18] and focus on the architecture design. Also, Mobistreams [19] and Swing [20] focus on collaborative data stream computations. The above systems either do not optimize the offloading policy, or use heuristics that do not cater for task accuracy.

Mobile and IoT Analytics. The importance of analytics has motivated the design of wireless systems that can execute such tasks. For instance, [21], [22] tailor deep neural networks for execution in mobile devices, while [23] and [24] minimize the execution time for known system parameters and task loads. Finally, [25]–[27] leverage the edge architecture to effectively execute analytics for IoT devices. The plethora of such system proposals, underlines the necessity for our *online decision framework* that provides optimal execution of analytics.

Optimization of Analytics. Prior works in computation offloading focus on different metrics such as number of served requests, [5], [28], and hence are not applicable here. In our previous work [29], we proposed a *static* collaborative optimization framework, which does not employ predictions nor accounts for computation constraints. Other works, e.g. [22] either rely on heuristics or assume static systems and known requests. Clearly, these assumptions are invalid for many practical cases where system parameters not only vary with time, but often do not follow i.i.d. processes. This renders the application of max-weight type of policies [9] inefficient. Our approach is fundamentally different and leads to an online robust algorithm and is inspired by dual averaging and primal recovery algorithms for static problems, see [6].

Improvement of ML Models. Clearly, despite the efforts to improve the execution of analytics at small devices, e.g., by residual learning or compression [10], the trade off between local low-accuracy and cloudlet high-accuracy execution is still important due to the increasing number and complexity of these tasks. This observation has spurred efforts for designing fast multi-tier (cloud to edge) deep neural networks [15] and for dynamic model selection [30], among others. These works are orthogonal to our approach and can be directly incorporated in our framework.

VI. CONCLUSIONS

We propose the idea of improving the execution of data analytics at IoT devices with more robust instances running at cloudlets. The key feature of our proposal is a dynamic and distributed algorithm that makes the outsourcing decisions based on the expected performance improvement, and the available resources at the devices and cloudlet. The proposed algorithm achieves near-optimal performance in a deterministic fashion, and under minimal assumptions about the system

behavior. This makes it ideal for the problem at hand where, the stochastic effects (e.g., expected accuracy gains) have unknown mean values and possibly non-i.i.d. behavior.

REFERENCES

- [1] E. Siow, T. Tiropanis, and W. Hall, “Analytics for the internet of things: A survey,” *ACM Comput. Surv.*, vol. 51, no. 4, pp. 74:1–74:36, 2018.
- [2] C. Jiang *et al.*, “Machine learning paradigms for next-generation wireless networks,” *IEEE Wireless Comm.*, vol. 24, no. 2, pp. 98–105, 2017.
- [3] Cisco White Paper, “Cisco global cloud index: Forecast and methodology, document id:1513879861264127,” 2018.
- [4] M. Satyanarayanan *et al.*, “The case for vm-based cloudlets in mobile computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14–23, 2009.
- [5] Y. Mao *et al.*, “A survey on mobile edge computing: The communication perspective,” *IEEE Comm. Surv. Tut.*, vol. 19, no. 4, pp. 2322–2358, 2017.
- [6] A. Nedić and A. Ozdaglar, “Approximate primal solutions and rate analysis for dual subgradient methods,” *SIAM J. on Optimization*, vol. 19, no. 4, pp. 1757–1780, 2009.
- [7] A. Galanopoulos *et al.*, “Improving iot analytics through selective edge execution: Appendix,” 2019, <https://1drv.ms/b/s!Ao15IEO8XUP1iQIjflw0YeaUCa83?e=91W474>.
- [8] A. Gelman and J. Hill, “*Data Analysis Using Regression and Multi-level/Hierarchical Models*”. Cambridge University Press, 2007.
- [9] L. Georgiadis, M. J. Neely, and L. Tassiulas, “Resource allocation and cross-layer control in wireless networks,” *Found. Trends Netw.*, vol. 1, no. 1, pp. 1–144, 2006.
- [10] V. Chandrasekhar *et al.*, “Compression of deep neural networks for image instance retrieval,” in *Proc. of DCC*, 2017.
- [11] Y. Lecun *et al.*, “Gradient-based learning applied to document recognition,” *Proc. of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [12] A. Krizhevsky, “Learning multiple layers of features from tiny images,” Tech. Rep., 2009.
- [13] S. A. Dudani, “The distance-weighted k-nearest-neighbor rule,” *IEEE Trans. on Sys., Man, and Cybern.*, vol. 6, no. 4, pp. 325–327, 1976.
- [14] M. Abadi *et al.*, “Tensorflow: A system for large-scale machine learning,” in *Proc. of USENIX OSDI*, 2016.
- [15] S. Teerapittayanon *et al.*, “Distributed deep neural networks over the cloud, the edge and end devices,” in *Proc. of IEEE ICDCS*, 2017.
- [16] B.-G. Chun *et al.*, “Clonecloud: Elastic execution between mobile device and cloud,” in *Proc. of EuroSys*, 2011.
- [17] A. Dou *et al.*, “Misco: A mapreduce framework for mobile systems,” in *Proc. of PETRA*, 2010.
- [18] X. Lyu *et al.*, “Selective offloading in mobile edge computing for the green internet of things,” *IEEE Network*, vol. 32, no. 1, pp. 54–60, 2018.
- [19] H. Wang and L. Peh, “Mobistreams: A reliable distributed stream processing system for mobile devices,” in *Proc. of IEEE IPDPS*, 2014.
- [20] S. Fan, T. Salonidis, and B. Lee, “Swing: Swarm computing for mobile sensing,” in *Proc. of IEEE ICDCS*, 2018.
- [21] X. Ran *et al.*, “Delivering deep learning to mobile devices via offloading,” in *Proc. of VR/AR Network Workshop*, 2017.
- [22] X. Ran *et al.*, “Deepdecision: A mobile deep learning framework for edge video analytics,” in *Proc. of IEEE INFOCOM*, 2018.
- [23] Y. Li *et al.*, “Mobiqor: Pushing the envelope of mobile edge computing via quality-of-result optimization,” in *Proc. of IEEE ICDCS*, 2017.
- [24] W. Zhang *et al.*, “Hetero-edge: Orchestration of real-time vision applications on heterogeneous edge clouds,” in *Proc. of IEEE INFOCOM*, 2019.
- [25] G. Li *et al.*, “Data analytics for fog computing by distributed online learning with asynchronous update,” in *Proc. of IEEE ICC*, 2019.
- [26] S. K. Sharma and X. Wang, “Live data analytics with collaborative edge and cloud processing in wireless iot networks,” *IEEE Access*, vol. 5, pp. 4621–4635, 2017.
- [27] J. He *et al.*, “Multitier fog computing with large-scale iot data analytics for smart cities,” *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 677–686, 2018.
- [28] X. Chen *et al.*, “Efficient multi-user computation offloading for mobile-edge cloud computing,” *IEEE/ACM Trans. on Networking*, vol. 24, no. 5, pp. 2795–2808, 2016.
- [29] A. Galanopoulos, G. Iosifidis, and T. Salonidis, “Optimizing data analytics in energy constrained iot networks,” in *Proc. of WiOpt*, 2018.
- [30] L. Liu and J. Deng, “Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution,” *arXiv:1701.00299*, 2017.