

LEAF + AIO: Edge-Assisted Energy-Aware Object Detection for Mobile Augmented Reality

Haoxin Wang[✉], Member, IEEE, BaekGyu Kim[✉], Member, IEEE,
Jiang Xie[✉], Fellow, IEEE, and Zhu Han[✉], Fellow, IEEE

Abstract—Today very few deep learning-based mobile augmented reality (MAR) applications are applied in mobile devices because they are significantly energy-guzzling. In this paper, we design an edge-based energy-aware MAR system that enables MAR devices to dynamically change their configurations, such as CPU frequency, computation model size, and image offloading frequency based on user preferences, camera sampling rates, and available radio resources. Our proposed dynamic MAR configuration adaptations can minimize the per frame energy consumption of multiple MAR clients without degrading their preferred MAR performance metrics, such as latency and detection accuracy. To thoroughly analyze the interactions among MAR configurations, user preferences, camera sampling rate, and energy consumption, we propose, to the best of our knowledge, the first comprehensive analytical energy model for MAR devices. Based on the proposed analytical model, we design a LEAF optimization algorithm to guide the MAR configuration adaptation and server radio resource allocation. An image offloading frequency orchestrator, coordinating with the LEAF, is developed to adaptively regulate the edge-based object detection invocations and to further improve the energy efficiency of MAR devices. Extensive evaluations are conducted to validate the performance of the proposed analytical model and algorithms.

Index Terms—Augmented reality, mobile edge computing, object detection

1 INTRODUCTION

WITH the advancement in *Deep Learning* in the past few years, we are able to create intelligent machine learning models to accurately detect and classify complex objects in the physical world. This advancement has the potential to make *Mobile Augmented Reality* (MAR) applications highly intelligent and widely adaptable in various scenarios, such as tourism, education, and entertainment. Thus, implementing MAR applications on popular mobile architectures is a new trend in modern technologies.

However, only a few MAR applications are implemented in mobile devices and are developed based on deep learning frameworks because (i) performing deep learning algorithms on mobile devices is significantly energy-guzzling; (ii) deep learning algorithms are computation-intensive, and executing locally in resource limited mobile devices may not provide acceptable performance for MAR clients [1]. To solve these issues, a promising approach is to transfer MAR

input image/video frames to an edge server that is sufficiently powerful to execute the deep learning algorithms.

Motivations. Although compared to running a deep learning algorithm locally on a mobile device, edge-based approaches may extend the device's battery life to certain extents, it is still considerably energy consuming due to conducting multiple pre-processes on the mobile device, such as camera sampling, screen rendering, image conversion, and data transmission [2]. For instance, based on the measurement from our developed MAR testbed, a 3000 mAh smartphone battery is exhausted within approximately 2.3 hours for executing our developed MAR application which continuously transmits the latest camera sampled image frames to an edge server for object detection. Therefore, the energy efficiency of MAR devices becomes a bottleneck, which impedes MAR clients to obtain better MAR performance. For example, decreasing the energy consumption of an MAR device is always at the cost of reducing the object detection accuracy. Therefore, improving the energy efficiency of MAR devices and balancing the tradeoffs between energy efficiency and other MAR performance metrics are crucial to edge-based MAR systems.

Challenges. An accurate analytical energy model is significantly important for understanding how energy is consumed in an MAR device and for guiding the design of energy-aware MAR systems. However, to the best of our knowledge, there is no existing energy model developed for MAR devices or applications. Developing a comprehensive MAR energy model that is sufficiently general to handle any MAR architecture and application is very challenging. This is because (i) interactions between MAR configuration parameters (e.g., client's CPU frequency and computation model size) and MAR device's energy consumption are

- Haoxin Wang and BaekGyu Kim are with the Toyota Motor North America (TMNA) R&D InfoTech Labs, Mountain View, CA 94043 USA. E-mail: hhwang50@uncc.edu, baekgyu.kim@toyota.com.
- Jiang Xie is with the University of North Carolina at Charlotte, Charlotte, NC 28223 USA. E-mail: linda.xie@uncc.edu.
- Zhu Han is with the Department of Electrical, Computer Engineering, University of Houston, Houston, TX 77004 USA, and also with the Department of Computer Science, Engineering, Kyung Hee University, Seoul 446-701, South Korea. E-mail: hanzhu22@gmail.com.

Manuscript received 29 Sept. 2021; revised 17 Mar. 2022; accepted 24 May 2022. Date of publication 3 June 2022; date of current version 31 Aug. 2023.

This work was supported in part by the U.S. National Science Foundation (NSF) under Grants 1718666, 1910667, 1910891, 2025284, 2107216, and 2128368, and funds from Toyota Motor North America.

(Corresponding author: Jiang Xie.)

Digital Object Identifier no. 10.1109/TMC.2022.3179943

complex and lack analytic understandings; (ii) interactions between these configurations and the device's energy consumption may also vary with different mobile architectures.

In addition, designing an energy-aware solution for mobile devices in edge-based MAR systems is also challenging, even after we obtain an analytical energy model. This is because: (i) complicated pre-processes on MAR devices increase the complexity of the problem. Compared to conventional computation offloading systems, besides data transmission, there are also a variety of pre-processing tasks (e.g., camera sampling, screen rendering, and image conversion) necessarily to be performed on MAR devices, which are also energy consuming. For example, over 60% of the energy is consumed by camera sampling and screen rendering, based on observations from our developed testbed. Therefore, we have to take into account the energy efficiency of these pre-processing tasks while designing an energy-aware approach for MAR clients. (ii) Considering the user preference constraint of individual MAR clients also increases the complexity of the problem. For example, maintaining a high object detection accuracy for a client who prefers a precise MAR while decreasing its energy consumption is very challenging. As stated previously, reducing the energy consumption of the MAR device without degrading other performance metrics is no easy task. (iii) In practical scenarios, an edge server is shared by multiple MAR clients. Individual client's energy efficiency is also coupled with the radio resource allocation at the edge server. Such a coupling makes it computationally hard to optimally allocate radio resources and improve each client's energy efficiency.

Our Contributions. In this paper¹, we study these research challenges and design a user preference based energy-aware edge-based MAR system. The novel contributions of this paper are summarized as follows:

- 1) We design and implement an edge-based object detection for MAR systems to analyze the interactions between MAR configurations and the client's energy consumption. Based on our experimental study, we summarize several insights which can potentially guide the design of energy-aware object detection.
- 2) We propose, to the best of our knowledge, the first comprehensive energy model which identifies (i) the tradeoffs among the energy consumption, service latency, and detection accuracy, and (ii) the interactions among MAR configuration parameters (i.e., CPU frequency and computation model size), user preferences, camera sampling rate, network bandwidth, and per frame energy consumption for a multi-user edge-based MAR system.
- 3) We propose an energy-efficient optimization algorithm, LEAF, which guides MAR configuration adaptations and radio resource allocations at the edge server, and minimizes the per frame energy consumption while satisfying variant clients' user preferences.
- 4) We develop and implement an image offloading frequency orchestrator that coordinates with the LEAF

algorithm to avoid energy-consuming continuous repeated executions of object detections and further improve the energy efficiency of MAR devices.

2 RELATED WORK

Energy Modeling. Energy modeling has been widely used for investigating the factors that influence the energy consumption of mobile devices. [3] and [4] propose energy models of WiFi and LTE data transmission with respect to the network performance metrics, such as data and retransmission rates, respectively. [5], [6], [7], [8], [9] propose multiple power consumption models to estimate the energy consumption of mobile CPUs. Tail energy caused by different components, such as disk, Wi-Fi, 3 G, and GPS in smartphones has been investigated in [9], [10]. However, none of them can be directly applied to estimate the energy consumed by MAR applications. This is because MAR applications introduce a variety of (i) energy consuming components (e.g., camera sampling and image conversion) that are not considered in the previous models and (ii) configuration variables (e.g., computation model size and camera sample rate) that also significantly influence the energy consumption of mobile devices.

Computation Offloading. Most existing research on computation offloading focuses on how to make offloading decisions. [11] and [12] coordinate the scheduling of offloading requests for multiple applications to further reduce the wireless energy cost caused by the long tail problem. [13] proposes an energy-efficient offloading approach for multicore-based mobile devices. [14] discusses the energy efficiency of computation offloading for mobile clients in cloud computing. However, these solutions cannot be applied to improving the energy efficiency of mobile devices in MAR offloading cases. This is because (i) a variety of pre-processing tasks in MAR executions, such as camera sampling, screen rendering, and image conversion, are not taken into account and (ii) besides the latency constraint that is considered in most existing computation offloading approaches, object detection accuracy is also a key performance metric, which must be considered while designing an MAR offloading solution. In addition, although some existing work proposes to study the tradeoffs between the MAR service latency and detection accuracy [2], [15], [16], [17], [18], [19], [20], [21], [22], none of them considered (i) the energy consumption of the MAR device and (ii) the whole processing pipeline of MAR (i.e., starting from the camera sampling to obtaining detection results).

CPU Frequency Scaling. Our work is also related to CPU frequency scaling. For modern mobile devices, such as smartphones, CPU frequency and the voltage provided to the CPU can be adjusted at run-time, which is called Dynamic Voltage and Frequency Scaling (DVFS). Prior work [11], [23], [24], [25] proposes various DVFS strategies to reduce the mobile device energy consumption under various applications, such as video streaming [11] and delay-tolerant applications [24]. However, to the best of our knowledge, there has been little effort factoring in the energy efficiency of MAR applications in the context of mobile device DVFS.

1. This work is an extension of our previous conference paper: 10.1109/INFOCOM41043.2020.9155517

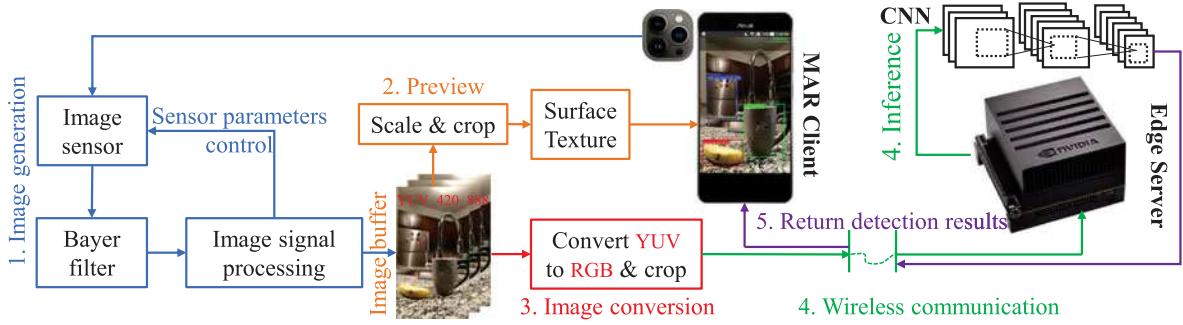


Fig. 1. The processing pipeline of the edge-based MAR system developed in this paper [26], [27].

3 EXPERIMENTAL RESULTS ON FACTORS AFFECTING MAR CLIENT ENERGY EFFICIENCY

In this section, we describe our preliminary experiments to evaluate the impact of various factors on the energy efficiency of an MAR client, service latency, and detection accuracy in an edge-based MAR system. Specifically, these experimental results provide (i) observations on interactions between energy consumption and MAR configuration parameters, such as MAR client's CPU frequency, computation model size, camera sampling rate, and user preference, (ii) bases of modeling the energy consumption of an MAR client, and (iii) insights on designing an energy-efficient optimization algorithm.

3.1 Testbed Setup

Our testbed consists of three major components: MAR client, edge server, and power monitor. Note that this paper focuses on the MAR application in which an MAR client captures physical environmental information through the camera and sends the information to an edge server for object detection. The detailed processing pipeline is shown in Fig. 1.

Processing Pipeline² *Image generation (phase 1):* The input to this phase is continuous light signal and the output is an image frame. In this phase, the image sensor first senses the intensity of light and converts it into an electronic signal. A Bayer filter is responsible for determining the color information. Then, an image signal processor (ISP) takes the raw data from the image sensor and converts it into a high-quality image frame. The ISP performs a series of image signal processing operations to deliver a high-quality image, such as noise reduction, color correction, and edge enhancement. In addition, the ISP conducts automated selection of key camera control values according to the environment (e.g., auto-focus (AF), auto-exposure (AE), and auto-white-balance (AWB)). The whole image generation pipeline in our implemented application is constructed based on `android.hardware.camera2` which is a package that provides an interface to individual camera devices connected to an Android device. `CaptureRequest` is a class in `android.hardware.camera2` that constructs the configurations for the capture hardware (sensor, lens, and flash), the processing pipeline, and the control algorithms. Therefore, in our implemented

application, we use `CaptureRequest` to set up image generation configurations. For example, `CaptureRequest.CONTROL_AE_MODE_OFF` disables AE and `CaptureRequest.CONTROL_AE_TARGET_FPS_RANGE` sets the camera FPS (i.e., the number of frames that the camera samples per second).

Preview (phase 2): The input to this phase is a latest generated image frame with `YUV_420_888` format³ (i.e., the output of Phase 1) and the output is a camera preview rendered on a smartphone's screen with a pre-defined preview resolution. In this phase, the latest generated image frame is first resized to the desired preview resolution and then buffered in a **SurfaceTexture** which is a class capturing frames from an image stream (e.g., camera preview or video decode) as an OpenGL ES texture. Finally, the camera preview frame in **SurfaceTexture** is copied and sent to a dedicated drawing surface, **SurfaceView**, and rendered on the screen. In our implemented application, the preview resolution is set via method `SurfaceTexture.setDefaultBufferSize()`.

Image conversion (phase 3): The input to this phase is a latest generated image frame with `YUV_420_888` format (i.e., the output of Phase 1) and the output is a cropped RGB image frame. In this phase, in order to further process camera captured images (i.e., object detection), an `ImageReader` class is implemented to acquire the latest generated image frame, where `ImageReader.OnImageAvailableListener` provides a callback interface for being notified that a new generated image frame is available and method `ImageReader.acquireLatestImage()` acquires the latest image frame from the `ImageReader`'s queue while dropping an older image. Additionally, the desired size and format of acquired image frames are configured once an `ImageReader` is created. In our implemented application, the desired size and the preview resolution are the same and the image format in `ImageReader` is set to `YUV_420_888`. Furthermore, an image converter is implemented to convert the `YUV_420_888` image to an `RGB` image, because the input to a CNN-based object detection model must be an `RGB` image. Finally, the converted `RGB` image is cropped to the size of the CNN model for object detections.

Wireless communication & inference (phase 4): The input to this phase is a converted and cropped image frame (i.e., the output of Phase 3) and the output is an object detection

2. GitHub: <https://github.com/WINSAC/Mobile-AR-in-Edge-Computing-Client>

3. For `android.hardware.camera2`, `YUV_420_888` format is recommended for YUV output [28].

result. In our implemented application, the object detection result contains one or multiple bounding boxes with labels that identify the locations and classifications of the objects in an image frame. Each bounding box consists of 5 predictions: (x, y, w, h) and a confidence score [29]. The (x, y) coordinates represent the center of the box relative to the bounds of the grid cell. The (h, w) coordinates represent the height and width of the bounding box relative to (x, y) . The confidence score reflects how confident the CNN-based object detection model is on the box containing an object and also how accurate it thinks the box is what it predicts. Our implemented application transmits the converted and cropped image frame to the edge server through a wireless TCP socket connection in real time. To avoid having the server process stale frames, the application always sends the latest generated frame to the server and waits to receive the detection result before sending the next frame for processing.

Detection result rendering (phase 5): The input to this phase is the object detection result of an image frame (i.e., the output of Phase 4) and the output is a view with overlaid augmented objects (specifically, overlaid bounding boxes and labels in this paper) on top of the physical objects (e.g., a cup).

Edge Server. The edge server is developed to process received image frames and to send the detection results back to the MAR client. We implement an edge server on an Nvidia Jetson AGX Xavier, which connects to a WiFi access point (AP) through a 1 Gbps Ethernet cable. The transmission latency between the server and AP can be ignored. Two major modules are implemented on the edge server⁴: (i) the *communication handler* which establishes a TCP socket connection with the MAR device and (ii) the *analytics handler* which performs object detection for the MAR client. In this paper, the analytics handler is designed based on a custom framework called Darknet [30] with GPU acceleration and runs YOLOv3 [29], a large Convolutional Neural Networks (CNN) model. The YOLOv3 model used in our experiments is trained on COCO dataset [31] and can detect 80 classes.

MAR Client. We implement an MAR client on a rooted Android smartphone, Nexus 6, which is equipped with Qualcomm Snapdragon 805 SoC (System-on-Chip). The CPU frequency ranges from 0.3 GHz to 2.649 GHz.

Power Monitor. The power monitor is responsible for measuring the power consumption of the MAR client. We use Monsoon Power Monitor [32], which can sample at 5,000 Hz, to provide power supply for the MAR device. The power measurements are taken with the screen on, with the Bluetooth/LTE radios disabled, and with minimal background application activity, ensuring that the smartphone's *base power* is low and does not vary unpredictably over time. The base power is defined as the power consumed when the smartphone is connected to the AP without any data transmission activity. The detailed energy measurement methodology is presented in our previous paper [27].

Key Performance Metrics. We define three performance metrics to evaluate the MAR system:

- *Per frame energy consumption:* The per frame energy consumption is the total amount of energy consumed

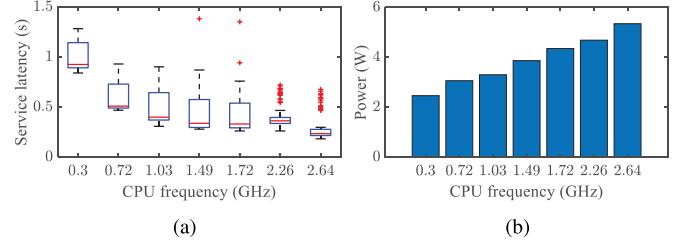


Fig. 2. CPU frequency versus service latency and power (computation model size: 320^2 pixels).

in an MAR client by successfully performing the object detection on one image frame. It includes the energy consumed by camera sampling (i.e., image generation), screen rendering (i.e., preview), image conversion, communication, and operating system.

- *Service latency:* The service latency is the total time needed to derive the detection result on one image frame. It includes the latency of image conversion, transmission, and inference.
- *Accuracy:* The mean average precision (mAP) is a commonly used performance metric to evaluate the detection accuracy of a visual object detection algorithm [33], where a greater accuracy is indicated by a higher mAP.

3.2 The Impact of CPU Frequency on Power Consumption and Service Latency

In this experiment, we seek to investigate how the CPU frequency impacts the power consumption of the MAR device and the service latency. We set the test device to the *Userspace Governor* and change its CPU frequency manually by writing files in the `/sys/devices/system/cpu/[cpu#]/cpufreq` virtual file system with root privilege. The camera FPS is set to 15 and the computation model size is 320^2 pixels. The results are shown in Fig. 2. The lower the CPU frequency, the longer service latency the MAR client derives and the less power it consumes. However, the reduction of the service latency and the increase of the power consumption is disproportional. For example, as compared to 1.03 GHz, 1.72 GHz reduces about 2% service latency but increases about 15% power consumption. As compared to 0.3 GHz, 0.72 GHz reduces about 60% service latency, but only increases about 20% power consumption.

Insight. This result advocates adapting the client's CPU frequency for the service latency reduction by trading as little increase of the power consumption as possible.

3.3 The Impact of Computation Model Size on Energy Consumption and Service Latency

In this experiment, we implement object detection based on the YOLOv3 framework with six different CNN model sizes. The test device works on the default CPU governor, *Interactive* and its camera FPS is set to 15. Increasing the model size always results in a gain of mAP. However, the gain on mAP becomes smaller as the increase of the model sizes [16]. In addition, the per frame energy consumption and the service latency boost 85% and 130%, respectively, when the model size increases from 128^2 to 608^2 pixels, as shown in Figs. 3a and 3b.

4. GitHub: <https://github.com/WINSAC/Mobile-AR-in-Edge-Computing-Server>

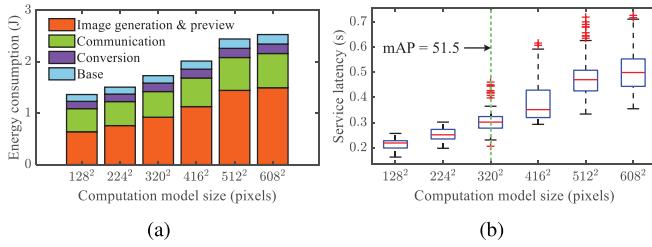


Fig. 3. Computation model size versus energy consumption and service latency.

Insight. This result inspires us to trade mAP for the per frame energy consumption and service latency reduction when the model size is large.

3.4 The Impact of Camera FPS on Power Consumption

In this experiment, we vary the MAR client's camera FPS to explore how it impacts the device's power consumption, where the camera FPS is defined as the number of frames that the camera samples per second. The computation model size is 320² pixels and the default CPU frequency is 1.49 GHz. Fig. 4a shows that a large camera FPS leads to a high power consumption. However, as shown in Fig. 1, not every camera captured image frame is sent to the edge server for detection. Because of the need (i) to avoid the processing of stale frames and (ii) to decrease the transmission energy consumption, only the latest camera sampled image frame is transmitted to the server. This may result in the MAR client expending significant reactive power for sampling non-detectable image frames. In Fig. 4b, we quantify the sampling efficiency with the variation of the camera FPS. As we expected, a large camera FPS leads to a lower sampling efficiency (e.g., less than 2% of the power is consumed for sampling the detectable image frames when the camera FPS is set to 30). However, in most MAR applications, users usually request a high camera FPS for a smoother preview experience, which is critical for tracking targets in physical environments. Interestingly, increasing CPU frequency can reduce the reactive power for sampling, as shown in Fig. 4b.

Insight. This result demonstrates that when a high camera FPS is requested, increasing CPU frequency can promote the sampling efficiency but may also boost the power consumption. Therefore, finding a CPU frequency that can balance this tradeoff is critical.

3.5 User Preference

An MAR client may have variant preferences in different implementation cases, including:

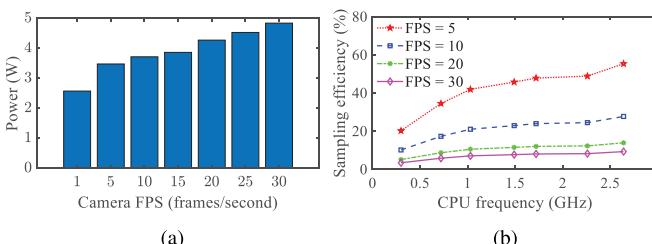


Fig. 4. Camera FPS versus power and sampling efficiency (computation model size: 320² pixels).

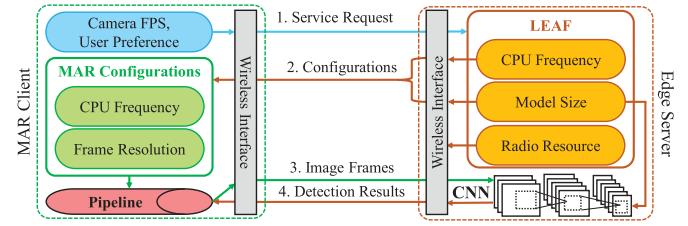


Fig. 5. Overview of the proposed edge-based MAR system.

- *Latency-preferred.* The MAR application of cognitive assistance [34], where a wearable device helps visually impaired people navigate on a street, may require a low service latency but can tolerate a relatively high number of false positives (i.e., false alarms are fine but missing any potential threats on the street is costly).
- *Accuracy-preferred.* An MAR application for recommending products in shopping malls or supermarkets may tolerate a long latency but requires a high detection accuracy and preview smoothness.
- *Preview-preferred.* The MAR drawing assistant application [35], where a user is instructed to trace virtual drawings from the phone, may tolerate a long latency (i.e., only needs to periodically detect the position of the paper where the user is drawing on) but requires a smooth preview to track the lines that the user is drawing.

Insight. This observation infers that the user preference's diversity may significantly affect the tradeoffs presented above. For instance, for the accuracy-preferred case, trading detection accuracy for the per frame energy consumption or service latency reduction works against the requirement of the user.

4 PROPOSED SYSTEM ARCHITECTURE

Based on the above insights, we propose an edge-based MAR system for object detection to reduce the per frame energy consumption of MAR clients by dynamically selecting the optimal combination of MAR configurations. To derive the optimal MAR configurations, we propose an optimization algorithm (LEAF) that supports low-energy, accurate, and fast MAR applications.

Fig. 5 shows the overview of our proposed system. In the first step, MAR clients send their service requests and selected camera FPS and user preferences to an edge server. In the second step, according to the received camera FPS and user preferences, the edge server determines the optimal CPU frequency, computation model size, and allocated radio resource for each MAR client using our proposed LEAF algorithm. The determined CPU frequency and computation model size are then sent back to corresponding MAR clients as MAR configuration messages. In the third step, MAR clients set their CPU frequency to the optimal value and resize their latest camera sampled image frames based on the received optimal computation model size. After the CPU frequency adaptation and image frame resizing, MAR clients transmit their image frames to the edge server for object detection. In the final step, the edge server returns detection results to corresponding MAR clients. The LEAF will be executed when (i) a new MAR client joins the system or (ii) the user changes the camera FPS or user preferences. The LEAF algorithm will be detailed in the next section.

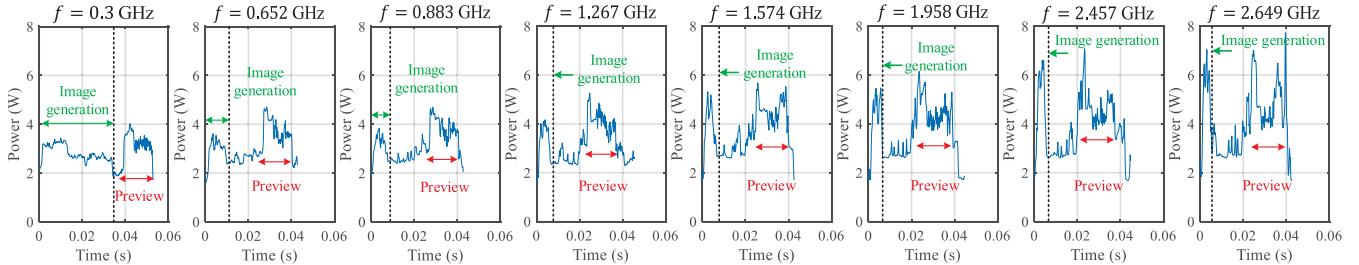


Fig. 6. The impact of CPU frequency on the power consumption of image generation and preview.

system; (ii) an MAR client leaves the system; or (iii) an MAR client re-sends the service request with a new user preference.

However, designing such a system is challenging. From the presented insights in the previous section, the interactions among the MAR system configuration variables, user preference, camera FPS, and the per frame energy consumption are complicated. (i) Some configuration variables improve one performance metric but impair another one. For example, a lower computation model size reduces the service latency but decreases the detection accuracy. (ii) Some configuration variables may affect the same metric in multiple ways. For example, selecting a higher CPU frequency can decrease the per frame energy consumption by increasing the sampling efficiency, but it increases the CPU power, which conversely increases the per frame energy consumption. Unfortunately, there is no analytical model for characterizing these interactions in the MAR system and it is not possible to design a prominent optimization algorithm without thoroughly analyzing these interactions.

5 PROPOSED ANALYTICAL MODEL AND PROBLEM FORMULATION

In this section, we thoroughly investigate the complicated interactions among the MAR configuration parameters, user preference, camera FPS, and the key performance metrics presented in Section 3. We first propose a comprehensive analytical model to theoretically dissect the per frame energy consumption and service latency. The proposed model is general enough to handle any MAR device and application. Then, using the proposed model, we further model multiple fine-grained interactions, whose theoretical properties are complex and hard to understand, via a data-driven methodology. Finally, based on the above proposed models, we formulate the MAR reconfiguration as an optimization problem.

5.1 Analytics-Based Modeling Methodology

We consider an edge-based MAR system with K MAR clients and one edge server, where clients are connected to the edge server via a single-hop wireless network. Denote \mathcal{K} as the set of MAR clients. The per frame service latency of the k th MAR client can be defined as

$$L^k = L_{cv}^k + L_{tr}^k + L_{inf}^k, \quad (1)$$

where L_{cv}^k is the image conversion latency caused by converting a buffered camera captured image frame from YUV to RGB; L_{tr}^k is the transmission latency incurred by sending

the converted RGB image frame from the k th client to its connected edge server; and L_{inf}^k is the inference latency of the object detection on the server. According to the MAR pipeline depicted in Fig. 1, the per frame energy consumption of the k th MAR client can be defined as

$$E^k = E_{img}^k + E_{cv}^k + E_{com}^k + E_{bs}^k, \quad (2)$$

where E_{img}^k is the image generation and preview energy consumption incurred by image sampling, processing, and preview rendering; E_{cv}^k is the image conversion energy consumption; E_{com}^k is the wireless communication energy consumption, which includes four phases: *promotion*, *data transmission*, *tail*, and *idle*; and E_{bs}^k is the MAR device base energy consumption.

The Model of Image Generation and Preview. Image generation is the process that an MAR client transfers its camera sensed continuous light signal to a displayable image frame. Preview is the process of rendering the latest generated image frame on the client's screen. As these two processes are executed in parallel with the main thread, their execution delays are not counted in the per frame service latency.

As depicted in Fig. 3a, the energy consumption of image generation and preview is the largest portion of the per frame energy consumption. To understand how energy is consumed in image generation and preview and what configuration variables impact it, we conduct a set of experiments. We find that *the power consumption of image generation and preview highly depends on the CPU frequency*. Fig. 6 shows the power consumption of image generation and preview under different CPU frequencies, where the camera FPS is set to 15. A higher CPU frequency results in a higher average power consumption. In addition, the image generation delay is also closely related to the CPU frequency, where a higher CPU frequency always leads to a shorter delay. However, the delay of rendering a preview is only related to the GPU frequency, which is out of the scope of this paper. Thus, we consider the preview delay as a fixed value with any CPU frequencies. We model the energy consumption of the k th MAR client's image generation and preview within a service latency as

$$E_{img}^k = \left(\int_0^{t_{gt}^k(f_k)} P_{gt}^k(f_k) dt + \int_0^{t_{prv}^k} P_{prv}^k(f_k) dt \right) \cdot fps_k \cdot L^k, \quad (3)$$

where P_{gt}^k , P_{prv}^k , t_{gt}^k , t_{prv}^k are the power consumption of image generation, preview, the delay of image generation, and preview, respectively. f_k is the CPU frequency. fps_k is the camera FPS. P_{gt}^k , P_{prv}^k , and t_{gt}^k are functions of f_k .

The Model of Image Conversion. Image conversion is processed through the MAR client's CPU; and hence, the conversion latency and power consumption highly depend on the CPU frequency. We define L_{cv}^k and E_{cv}^k a function of f_k . Therefore, the major source of the power consumption of the image conversion is the CPU computation. The power consumption of mobile CPUs can be divided into two components, $P_{cv}^k = P_{leak} + P_{dynamic}^k$ [7], where P_{leak} is independent and $P_{dynamic}^k$ is dependent upon the CPU frequency. (i) P_{leak} is the power originating from leakage effects and is in essence not useful for the CPU's purpose. In this paper, we consider P_{leak} a constant value ϵ . (ii) $P_{dynamic}^k$ is the power consumed by the logic gate switching at f_k and is proportional to $V_k^2 f_k$, where V_k is the supply voltage for the CPU. Due to the DVFS for the power saving purpose, e.g., a higher f_k will be supplied by a larger V_k , each f_k matches with a specific V_k , where $V_k \propto (\alpha_1 f_k + \alpha_2)$. α_1 and α_2 are two positive coefficients. Thus, the energy consumption of converting a single image frame of the k th MAR client can be modeled as

$$E_{cv}^k = P_{cv}^k L_{cv}^k = (\alpha_1^2 f_k^3 + 2\alpha_1 \alpha_2 f_k^2 + \alpha_2 f_k + \epsilon) \cdot L_{cv}^k(f_k). \quad (4)$$

The Model of Wireless Communication and Inference. Intuitively, the wireless communication latency is related to the data size of the transmitted image frame (determined by the frame resolution) and wireless data rate. As the data size of detection results is usually small, we do not consider the latency caused by returning the detection results [16]. In this paper, we use s_k^2 (pixels) to represent the computation model size of the k th MAR client. The client must send image frames whose resolutions are not smaller than s_k^2 to the edge server to obtain the corresponding detection accuracy. Thus, the most efficient way is to transmit the image frame with the resolution of s_k^2 to the server. Denote σ as the number of bits required to represent the information carried by one pixel. The data size of an image frame is calculated as σs_k^2 bits. Let B_k be the wireless bandwidth derived by the k th MAR client. We model the transmission latency of the k th client as

$$L_{tr}^k = \frac{\sigma s_k^2}{R_k}, \quad (5)$$

where R_k is the average wireless data rate of the k th client, which is a function of B_k .

In addition to the computation model size and wireless bandwidth, the transmission latency is also determined by the MAR client's CPU frequency. This is because the image transmission uses TCP as the transport layer protocol, and TCP utilizes substantial CPU capacity to handle congestion avoidance, buffer, and retransmission requests. For example, when the CPU frequency is low, the remaining CPU capacity may not be adequate to process the TCP task; and thus, the TCP throughput is decreased. Therefore, R_k is also a function of f_k , i.e., $R_k(B_k, f_k)$. In this paper, $R_k(B_k, f_k)$ is defined as

$$R_k(B_k, f_k) = r_k^{max}(B_k) \cdot r_k^*(f_k), \quad (6)$$

where $r_k^{max}(B_k)$ is the network throughput, which is not affected by the variation of the MAR client's CPU frequency,

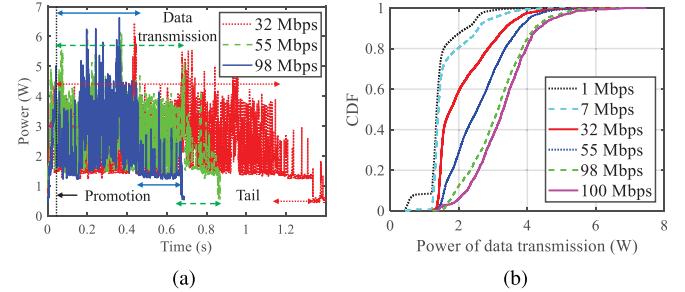


Fig. 7. MAR client's wireless interface power consumption.

and is only determined by the bandwidth (more comprehensive model of this part can be found in [3], which is out of the scope of this paper); $r_k^*(f_k)$ represents the impact of the CPU frequency on the TCP throughput.

In WiFi networks, when transmitting a single image frame, the MAR client's wireless interface experiences four phases: promotion, data transmission, tail, and idle. When an image transmission request comes, the wireless interface enters the promotion phase. Then, it enters the data transmission phase to send the image frame to the edge server. After completing the transmission, the wireless interface is forced to stay in the tail phase for a fixed duration and waits for other data transmission requests and the detection results. If the MAR client does not receive the detection result in the tail phase, it enters the idle phase and waits for the feedback from its associated edge server. Fig. 7 depicts the measured power consumption of the MAR client that transmits a $3,840 \times 2,160$ pixel image with different throughput. We find that the average power consumption of the data transmission phase increases as the throughput grows. However, the average power consumption and the duration of promotion and tail phases are almost constant. Therefore, we model the energy consumption of the k th MAR client in the duration that starts from the promotion phase to obtaining the object detection result as

$$E_{com}^k = P_{tr}^k(R_k(B_k, f_k))L_{tr}^k + P_{idle}^k t_{idle}^k + P_{prot}^k t_{prot}^k + P_{tail}^k t_{tail}^k, \quad (7)$$

where P_{tr}^k , P_{idle}^k , P_{prot}^k , and P_{tail}^k are the average power consumption of the data transmission, idle, promotion, and tail phases, respectively; t_{idle}^k , t_{prot}^k , and t_{tail}^k are the durations of the idle, promotion, and tail phases, respectively

$$P_{idle}^k t_{idle}^k = \begin{cases} 0, & L_{inf}^k(s_k^2) \leq t_{tail}^k, \\ P_{bs}^k \cdot (L_{inf}^k(s_k^2) - t_{tail}^k), & L_{inf}^k(s_k^2) > t_{tail}^k, \end{cases} \quad (8)$$

where P_{bs}^k is the MAR device's base power consumption; $L_{inf}^k(s_k^2)$ is the inference latency on the edge server, which is determined by the computation model size [16]. Note that our proposed wireless communication model can also be used in other wireless networks (e.g., LTE).

The Model of Base Energy. In this paper, the base energy consumption is defined as the energy consumed by the MAR clients' CPU without any workloads, except running its operating system, and the energy consumed by the screen without any rendering. Because the screen's brightness is not a critical factor that affects the object detection

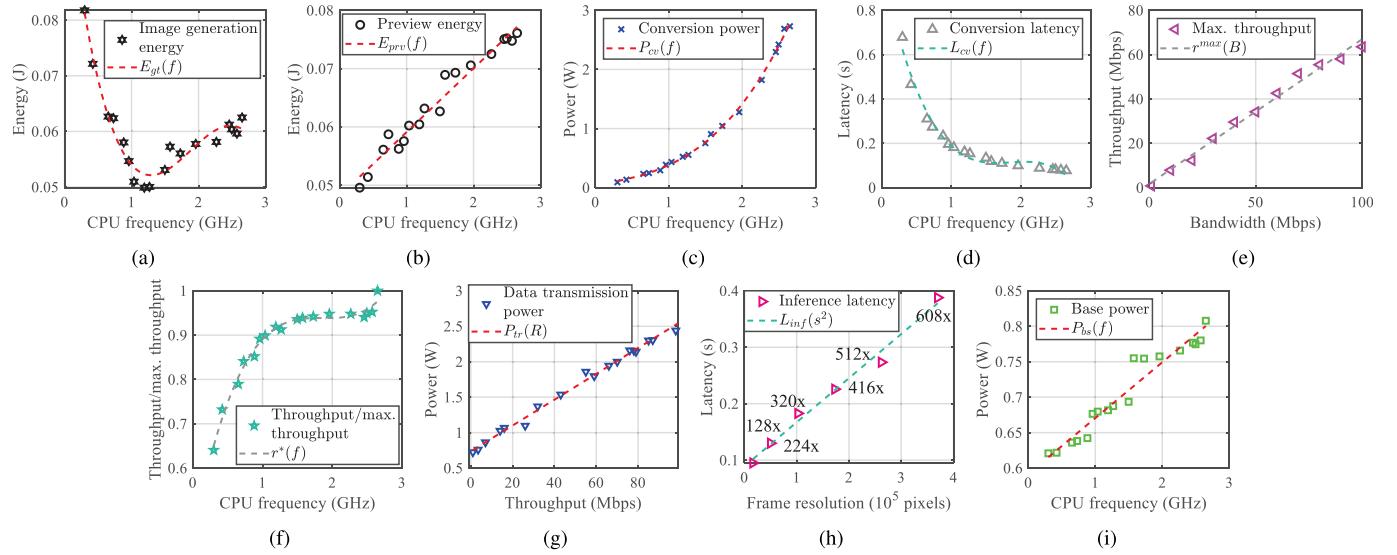


Fig. 8. The proposed data-driven analytic models for MAR devices, where each function, presented correspondingly in Table I, is trained offline via empirical measurements and regression analyses.

performance, it is considered as a constant value in our proposed power model. Thus, the base power consumption is only a function of the CPU frequency. We model the base energy consumption of the k th MAR client within a service latency as

$$E_{bs}^k = \begin{cases} P_{bs}^k(f_k) \cdot L^k, & L_{inf}^k(s_k^2) \leq t_{tail}, \\ P_{bs}^k(f_k) \cdot (L^k - L_{inf}^k(s_k^2) + t_{tail}), & L_{inf}^k(s_k^2) > t_{tail}. \end{cases} \quad (9)$$

5.2 Regression-Based Modeling Methodology

As shown in Subsection 5.1, some interactions or functions in our proposed analytical models still cannot be expressed clearly in an analytic form. This is because of (i) the lack of analytic understandings of some interactions and (ii) specific coefficients/functions that may vary with different MAR device models. For example, in (4), the specific coefficients in $P_{cv}^k(f_k)$ are unknown due to the lack of theoretical knowledge and vary with different MAR device models.

Therefore, we propose a data-driven methodology to address the above challenge, where those interactions with inadequate analytic understandings can be modeled and trained offline via empirical measurements and regression analyses. Note that regression-based modeling methodology is one of the most widely used approaches in developing mobile CPU's property models (e.g., CPU power and temperature variation modeling) and has shown to be effective in estimating CPU properties [5], [6], [11]. We use our testbed to collect measurements. The test MAR device is selected to work at 18 different CPU frequencies ranging from 0.3 to 2.649 GHz. In addition, in order to obtain fine-grained regression models and eliminate the interference among different workloads on the device power consumption, we develop three Android applications; each is applied with a specific function of the MAR client, which includes image generation and preview, image conversion, and image transmission applications. The developed regression models are shown in Fig. 8 and Table 1. Note that to obtain a statistical confidence in the experimental results, each

data point in Fig. 8 is derived by generating, transmitting, and detecting 1,000 image frames and calculating the average values. The root mean square error (RMSE) is applied for calculating the average model-prediction error in the units of the variable of interest [36].

5.3 Problem Formulation

Based on the above proposed models, we formulate the MAR reconfiguration as a multi-objective optimization problem [37]. We aim to *minimize the per frame energy consumption of multiple MAR clients in the system while satisfying the user preference* (as stated in Section 3.5) of each. We introduce two positive weight parameters λ_1^k and λ_2^k to characterize the user preference of the k th MAR client, where λ_1^k and λ_2^k can be specified by the client. We adopt the weighted sum method [38] to express the multi-object optimization problem as

$$\begin{aligned} \mathcal{P}_0 : \min_{\{f_k, s_k, B_k, \forall k \in \mathcal{K}\}} \quad & Q = \sum_{k \in \mathcal{K}} (E^k + \lambda_1^k L^k - \lambda_2^k A_k) \\ \text{s.t.} \quad & C_1 : \sum_{k \in \mathcal{K}} B_k \leq B_{max}; \\ & C_2 : L^k \leq L_{max}^k, \forall k \in \mathcal{K}; \\ & C_3 : F_{min} \leq f_k \leq F_{max}, \forall k \in \mathcal{K}; \\ & C_4 : s_k \in \{s_{min}, \dots, s_{max}\}, \forall k \in \mathcal{K}; \end{aligned} \quad (10)$$

TABLE 1
The Proposed Regression-Based Models

| Proposed models | RMSE | |
|-----------------|--|-------|
| $E_{gt}(f)$ | $-0.01071f^3 + 0.06055f^2 - 0.1028f + 0.107$ | 0.002 |
| $E_{prv}(f)$ | $0.01094f + 0.04816$ | 0.002 |
| $P_{cv}(f)$ | $0.1124f^3 + 0.01f^2 + 0.2175f + 0.04295$ | 0.041 |
| $L_{cv}(f)$ | $-0.145f^3 + 0.8f^2 - 1.467f + 0.996$ | 0.025 |
| $r^{max}(B)$ | $0.677B$ | 2.403 |
| $r^*(f)$ | $0.07651f^3 - 0.4264f^2 + 0.7916f + 0.4489$ | 0.013 |
| $P_{tr}(R)$ | $0.01821R + 0.7368$ | 0.052 |
| $L_{inf}(s^2)$ | $0.07816s^2 + 0.08892$ | 0.838 |
| $P_{bs}(f)$ | $0.07873f + 0.5918$ | 0.015 |

where A_k is an object detection accuracy function in terms of the k th MAR client selected computation model size s_k^2 (e.g., $A(s_k^2) = 1 - 1.578e^{-6.5 \times 10^{-3}s_k}$ [16]); L_{max}^k is the maximum tolerable service latency of the k th client; B_{max} is the maximum wireless bandwidth that an edge server can provide for its associated MAR clients. In practical scenarios, an edge server may simultaneously offer multiple different services for its associated users, e.g., video streaming, voice analysis, and content caching. Hence, the edge server may reallocate its bandwidth resource based on the user distribution. In this paper, we assume that B_{max} varies with time randomly. The constraint C_1 represents that MAR clients' derived bandwidth cannot exceed the total bandwidth allocated for the MAR service on the edge server; the constraint C_2 guarantees that the service latency of MAR clients are no larger than their maximum tolerable latency; the constraints C_3 and C_4 are the constraints of the MAR device's CPU frequency and computation model size configurations, where s_k is a discrete variable and its values depend on the available computation models in the MAR system.

6 PROPOSED LEAF OPTIMIZATION ALGORITHM

As shown in the previous section, problem \mathcal{P}_0 is a mixed-integer non-linear programming problem (MINLP) which is difficult to solve [39]. In order to solve this problem, we propose the LEAF algorithm based on the block coordinate descent (BCD) method [40].

To solve problem \mathcal{P}_0 , we relax the discrete variable s_k into continuous variable \hat{s}_k . The problem is relaxed as

$$\begin{aligned} \mathcal{P}_1 : \min_{\{f_k, \hat{s}_k, B_k, \forall k \in \mathcal{K}\}} \quad & Q = \sum_{k \in \mathcal{K}} (E^k + \lambda_1^k L^k - \lambda_2^k A_k) \\ \text{s.t.} \quad & C_1, C_2, C_3 \\ & \hat{C}_4 : s_{min} \leq \hat{s}_k \leq s_{max}, \forall k \in \mathcal{K}. \end{aligned} \quad (11)$$

According to the BCD method, we propose the LEAF algorithm which solves Problem \mathcal{P}_1 by sequentially fixing two of three variables and updating the remaining one. We iterate the process until the value of each variable converges.

$\nabla y(x)$ is denoted as the partial derivative of function y corresponding to variable x . Denote $\text{Proj}_{\mathcal{X}}(x)$ as the Euclidean projection of x onto \mathcal{X} ; $\text{Proj}_{\mathcal{X}}(x) \triangleq \arg \min_{v \in \mathcal{X}} \|x - v\|^2$.

The procedure of our proposed solution is summarized as:

- Given \hat{s}_k and B_k , we can derive a new f_k according to

$$f_k^{(j+1)} = \text{Proj}_{\mathcal{X}_f} \left(f_k^{(j)} - \gamma_k \nabla Q_k \left(f_k^{(j)} \right) \right), \forall k \in \mathcal{K}, \quad (12)$$

where $\gamma_k > 0$ is a constant step size and \mathcal{X}_f is the bounded domain constrained by C_3 . Based on the BCD method, we repeat (12) until the derived f_k is converged and then update f_k .

- Given f_k and B_k , we can derive a new \hat{s}_k according to

$$\hat{s}_k^{(\hat{j}+1)} = \text{Proj}_{\mathcal{X}_{\hat{s}}} \left(\hat{s}_k^{(j)} - \eta_k \nabla Q_k \left(\hat{s}_k^{(j)} \right) \right), \forall k \in \mathcal{K}, \quad (13)$$

where $\eta_k > 0$ is a constant step size and $\mathcal{X}_{\hat{s}}$ is the bounded domain constrained by \hat{C}_4 . Based on the

BCD method, we repeat (13) until the derived \hat{s}_k is converged and then update \hat{s}_k .

- Given f_k and \hat{s}_k , the problem is simplified to

$$\begin{aligned} \min_{\{B_k, \forall k \in \mathcal{K}\}} \quad & Q = \sum_{k \in \mathcal{K}} (E^k + \lambda_1^k L^k - \lambda_2^k A_k) \\ \text{s.t.} \quad & C_1 : \sum_{k \in \mathcal{K}} B_k \leq B_{max}; \\ & C_2 : L^k \leq L_{max}^k, \forall k \in \mathcal{K}; \end{aligned} \quad (14)$$

where constraints C_3 and \hat{C}_4 are irrelevant to this problem.

The Lagrangian dual decomposition method is utilized to solve the above problem, where the Lagrangian function is

$$\begin{aligned} \mathcal{L}(B_k, \mu, \beta) = \sum_{k \in \mathcal{K}} (E^k + \lambda_1^k L^k - \lambda_2^k A_k) + \mu \left(\sum_{k \in \mathcal{K}} B_k - B_{max} \right) \\ + \sum_{k \in \mathcal{K}} \beta_k (L^k - L_{max}^k), \end{aligned} \quad (15)$$

where μ and β are the Lagrange multipliers, (i.e., β is a Lagrange multiplier vector), corresponding to constraints C_1 and C_2 , respectively. The Lagrangian dual problem can therefore be expressed as

$$\begin{aligned} \max_{\{\mu, \beta\}} \quad & g(\mu, \beta) = \min_{\{B_k, \forall k \in \mathcal{K}\}} \mathcal{L}(B_k, \mu, \beta) \\ \text{s.t.} \quad & \mu \geq 0, \beta \geq 0. \end{aligned} \quad (16)$$

Here, $g(\mu, \beta)$ is concave with respect to B_k .

Lemma 1. *The problem \mathcal{P}_1 is convex with respect to B_k .*

Proof. For any feasible $B_i, B_j, \forall i, j \in \mathcal{K}$, we have

$$\frac{\partial^2 Q}{\partial B_i \partial B_j} = \begin{cases} 0, & i \neq j, \\ \Psi_i \cdot \frac{\partial^2 (1/r_{ij}^{max})}{\partial B_i \partial B_j}, & i = j, \end{cases} \quad (17)$$

where $\Psi_i = \frac{[fps_i(E_{gt}(f_i) + E_{prv}(f_i)) + P_{tr}^i(0) + P_{bs}(f_i) + \lambda_1^i] \sigma s_i^2}{r_i^*(f_i)}$ which is positive, and $\frac{\partial^2 (1/r_{ij}^{max})}{\partial B_i \partial B_j} = \frac{2}{0.677 B_i^3} > 0$. Thus, the Hessian matrix $\mathbf{H} = (\frac{\partial^2 Q}{\partial B_i \partial B_j})_{K \times K}$ is symmetric and positive definite. Constraint C_1 is linear and C_2 is convex with respect to B_k . Constraints C_3 and C_4 are irrelevant to B_k . Therefore, \mathcal{P}_1 is strictly convex with respect to B_k . \square

Therefore, based on the Karush-Kuhn-Tucker (KKT) condition [41], the sufficient and necessary condition of the optimal allocated bandwidth for the k th MU can be expressed as

$$B_k^* = \sqrt{\frac{\Phi(f_k, s_k, \beta_k)}{0.677 \mu}}, \quad (18)$$

$$\text{where } \Phi_k = \frac{[fps_i(E_{gt}(f_i) + E_{prv}(f_i)) + P_{tr}^i(0) + P_{bs}(f_i) + \lambda_1^i + \beta_k] \sigma s_k^2}{r_i^*(f_i)}.$$

Next, the sub-gradient method [41] is used to solve the dual problem. Based on the sub-gradient method, the dual variables of the k th MAR clients in the $(j+1)$ th iteration are

$$\begin{cases} \mu_k^{(j+1)} = \max \{0, [\mu^{(j)} + \vartheta_k^\mu \nabla g(\mu^{(j)})]\}, \forall k \in \mathcal{K}; \\ \beta_k^{(j+1)} = \max \{0, [\beta_k^{(j)} + \vartheta_k^\beta \nabla g(\beta_k^{(j)})]\}, \forall k \in \mathcal{K}; \end{cases} \quad (19)$$

where $\vartheta_k^\mu > 0$ and $\vartheta_k^\beta > 0$ are the constant step sizes.

Based on the above mathematical analysis, we propose an MAR optimization algorithm, LEAF, which can dynamically determines the CPU frequency of multiple MAR devices, selects the computation model sizes, and allocates the wireless bandwidth resources. The pseudo code of the proposed LEAF MAR algorithm is presented in Algorithm 6. First, the LEAF is initialized with the lowest CPU frequency, the smallest computation model size, and evenly allocated bandwidth resources among MAR devices. We then iteratively update f_k , \hat{s}_k , and B_k until the LEAF converges (i.e., line 7-8 in Algorithm 6). In addition, \hat{s}_k is a relaxed value of the computation model size. Thus, it may not match any pre-installed computation model in a real system. In this case, the LEAF selects the computation model size s_k that is the closest to the relaxed one \hat{s}_k (i.e., line 10 in Algorithm 6). Since the LEAF MAR algorithm is developed based on the BCD method and follows the convergence results in [40], we claim that the LEAF converges to a local optimal solution.

Algorithm 1. The LEAF MAR Algorithm

```

Input:  $\lambda_1^k, \lambda_2^k, L_{max}^k, B_{max}, fps_k$ , and  $\tau, \forall k \in \mathcal{K}$ .
Output:  $f_k, s_k$ , and  $B_k, \forall k \in \mathcal{K}$ .
1:  $B_k \leftarrow B_{max}/|\mathcal{K}|$ ,  $\hat{s}_k \leftarrow s_{min}, \forall k \in \mathcal{K}$ ,  $i \leftarrow 1$ ;
2: while True do
3:    $f_k \leftarrow$  solving  $\mathcal{P}_1$  with fixed  $\hat{s}_k$  and  $B_k$ ;
4:    $\hat{s}_k \leftarrow$  solving  $\mathcal{P}_1$  with fixed  $f_k$  and  $B_k$ ;
5:    $B_k \leftarrow$  solving  $\mathcal{P}_1$  with fixed  $f_k$  and  $\hat{s}_k$ ;
6:    $Q_i \leftarrow \sum_{k \in \mathcal{K}} (E^k + \lambda_1^k L^k - \lambda_2^k A_k)$ 
7:   if  $|(Q_i - Q_{i-1})/Q_i| \leq \tau$  then
8:     break;
9:    $i \leftarrow i + 1$ ;
10:   $s_k = \arg \min_{s \in \{s_{min} \dots s_{max}\}} |s - \hat{s}_k|, \forall k \in \mathcal{K}$ ;
11: return  $f_k, s_k$ , and  $B_k, \forall k \in \mathcal{K}$ .

```

7 IMAGE OFFLOADING FREQUENCY ORCHESTRATOR

In this section, an offloading frequency orchestrator with local object tracking is proposed to further reduce the energy consumption and latency of MAR devices by leveraging the model we developed for significant scene change estimation, based on our proposed LEAF algorithm.

7.1 Edge-Based Object Detection Versus Local Object Tracking

As presented in Section 6, our proposed LEAF is able to guide MAR configuration adaptations and radio resource allocations at the edge server to improve the energy efficiency of executing continuous image offloading and object detection. However, continuous repeated executions of offloading camera image frames to the edge server for object detection are unnecessary. This is because, although the positions of detected objects may slightly change in continuous camera captured frames due to the camera movement or detected object motions, the probability of significant changes to the scene or a new object appearing is low within a very short period. For example, as shown in Fig. 9, three image frames are extracted in a video stream. From Frame 1 to Frame 10, only the position of the detected dog in the scene changes.



Fig. 9. Continuous repeated executions of object detection are unnecessary.

Thus, sending every captured frame to the edge server for detecting objects (i.e., locating and recognizing objects in a frame) is extremely inefficient and will cause unnecessary energy expenditure even with our proposed LEAF.

To reduce the execution of continuous image offloading and object detection, one naive approach is to implement a local lightweight object tracker on the MAR device and invoke the tracker for updating the locations of the objects of interest that are achieved by performing a successful object detection, as is done in several prior works [1], [15], [20]. However, three essential questions are brought up here:

RQ 1. How much energy can the local object tracker save for an MAR device compared to performing edge-assisted object detection? It is intuitive that local lightweight object tracker consumes less battery than local CNN-based object detector due to the nature of CNNs, which contains tens to hundreds of computation-intensive layers. But how does it compare to the edge-based detectors, where the MAR device's on-board resource is not consumed by running CNNs?

RQ 2. How does the MAR device's hardware capacity (e.g., CPU frequency) impact the tracking performance and overhead (e.g., tracking delay and energy consumption)? It is critical to have the knowledge that whether the object tracker can help to improve the energy efficiency of MAR devices within the full or only a partial range of CPU frequencies.

RQ 3. How does the MAR device determine the frequency of the image offloading and object detection? The frequency of executing edge-based object detector is the most essential and challenging parameter of the MAR system. If the edge-based object detector (i.e., image offloading) is executed as often as possible, the MAR device may achieve a high object detection and tracking accuracy but a high energy expenditure. However, if the edge-based object detector is executed with a low frequency, for instance, executing an object detection only once at the beginning of tracking, the MAR device may achieve a high energy efficiency but unacceptable tracking accuracy (e.g., in our experiment, we observe that the tracking accuracy decreases or even the tracker loses objects of interest as the time interval between the current frame and reference frame performed object detection increases).

To the best of our knowledge, these questions lack pertinent investigations and sophisticated solutions in both academia and industry, such as ARCore [42] and ARKit [43]. To explore these questions, we implement a real-time lightweight object tracker on a Nexus 6 using JavaCV libraries⁵ [45]. The implemented lightweight object tracker in this paper is developed based on Kernelized Correlation Filter (KCF) [46] which is a tracking framework that utilizes

⁵ JavaCV is a collection of wrappers in Java for commonly used libraries in the field of computer vision such as OpenCV [44].

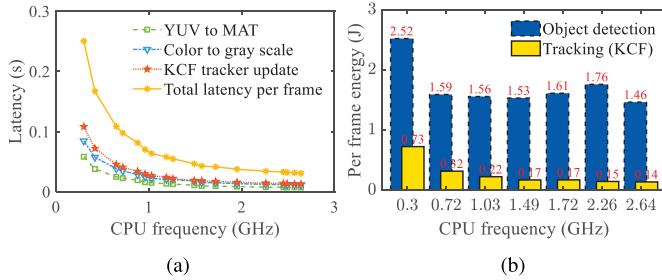


Fig. 10. CPU frequency versus latency and per frame energy consumption of local object tracking.

properties of circulant matrix to enhance the processing speed. KCF tracker has achieved impressive tracking delay and accuracy on Visual Tracker Benchmarks [47]. The latency of performing an object tracking on a single video frame contains (i) the latency of converting a camera captured raw YUV video frame produced by `IMAGEREADER` to an `MAT`⁶ object, (ii) the latency of converting a frame from color to gray scale (i.e., `COLOR_BGR2GRAY`), and (iii) the latency of executing the KCF object tracker.

Fig. 10 illustrates the average object tracking latency and per frame energy consumption when the MAR device runs on different CPU frequencies. In Fig. 10a, we observe that the average total latency of performing an object tracking is significantly reduced compared to the latency of edge-based object detection, as presented in Section 3. For instance, when the CPU frequency of Nexus 6 is 1.728 GHz, the average latency of object detection and tracking are 500 ms and 40 ms, respectively (RQ 1). Fig. 10b compares the per frame energy consumption of the edge-based object detection and the local object tracking, where we find that the lightweight local object tracker, KCF, can help to improve the energy efficiency of the Nexus 6 within the full range of CPU frequencies (RQ 2). The per frame energy consumption is decreased by over 80% comparing to the object detection ($s^2 = 320^2$) when the device's CPU frequency is not less than 1.032 GHz. *Therefore, implementing a lightweight local object tracker will not only help MAR devices to further mitigate the quick battery depletion, but also drop the latency substantially.*

The above experimental result and discussion advocate adding a local lightweight object tracker in our developed edge-based MAR system, as depicted in Fig. 5, for further improving the energy efficiency of MAR devices and reducing the latency. However, given the discussion on RQ 3, we argue that naively implementing an object tracker in such a system is inadequate, where an image offloading frequency orchestrator that balances the trade-off between MAR device's energy efficiency and tracking accuracy is essential and nonignorable. We design and implement such an orchestrator, which coordinates with the proposed LEAF to adaptively and intelligently adjust the image offloading frequency (i.e., execution of edge-based object detection) based on real-time scene change estimations.

7.2 Image Offloading Frequency Orchestrator

Fig. 11 provides an overview of how our proposed image offloading frequency orchestrator coordinates with the LEAF illustrated in Fig. 5. The proposed orchestrator is

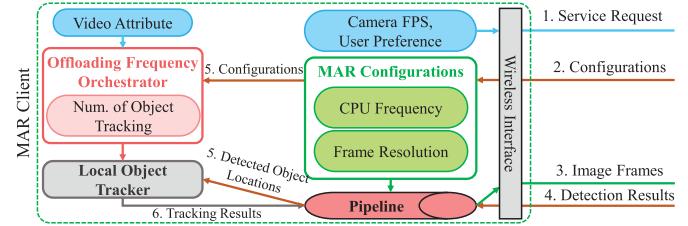


Fig. 11. Overview of the MAR client deployed with our proposed image offloading frequency orchestrator and how it coordinates with the rest of the edge-based MAR system with LEAF.

implemented in MAR devices. The MAR device invokes the orchestrator after it has successfully received the object detection results from the edge server. The inputs of the orchestrator are the optimal MAR configurations (i.e., CPU frequency and frame resolution) obtained from the proposed LEAF. The output is the estimated number of next successive image frames that will perform local object tracking, denoted by ρ . For instance, if the output of the orchestrator is 11, the next 11 continuous frames will not be eligible for offloading and will be transited to the local object tracker to perform tracking.

However, designing such an image offloading frequency orchestrator is challenging. Prior work [20] set a single threshold to determine whether the current image frame should be offloaded to the edge. However, (i) the value of the threshold is significantly experience-driven, which is unrealistic to handle all environment conditions with one single threshold; (ii) it lacks exploration of what is the optimal offloading solution for the MAR device in a time period. To tackle these, our orchestrator is designed based on two principles: (i) the detection/tracking decision will be made via a context-aware optimization algorithm, which is developed based on our proposed analytical model and LEAF; (ii) in order to achieve real-time results, considering the restricted computation capability of MAR devices, the designed algorithm should be as lightweight as possible.

To fulfill the first principle, it is necessary to predict how ρ will impact the object tracking accuracy within various scenarios, as tracking is not always accurate with respect to changes in object locations. The larger ρ the orchestrator provides, the less similarity between the current tracked image frame and the frame executed object detection, which raises the probability of tracking accuracy degradation. In addition, the attribute of the scenario (e.g., objects of interest being blurred) also heavily impacts the similarity among continuous tracked frames. To assess the tracking accuracy in terms of ρ , we choose to measure the similarity between two images using the peak signal-to-noise ratio (PSNR). PSNR provide a measure of the similarity between two images, I_1 and I_2 , and is given by $PSNR(I_1, I_2) = 20\log_{10}(\frac{MAX_I}{\sqrt{MSE(I_1, I_2)}})$,

where MAX_I is the maximum possible pixel value of the image. $MSE(I_1, I_2)$ is the mean squared error of two images and is calculated by $MSE(I_1, I_2) = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N (I_1^{ij} - I_2^{ij})^2$, where $M \times N$ is the image size; and i, j are the pixel locations within the images. Furthermore, the average intersection over union (IOU) is used for estimating the tracking accuracy. The IOU of object o in frame I is $IOU_o^I = \frac{R_o^G \cap R_o^P}{R_o^G \cup R_o^P}$, where R_o^G is the groundtruth region of object o , and R_o^P is the

⁶ MAT is a basic image container used in OpenCV.
Authorized licensed use limited to: University of North Carolina at Charlotte. Downloaded on June 20, 2024 at 19:13:39 UTC from IEEE Xplore. Restrictions apply.

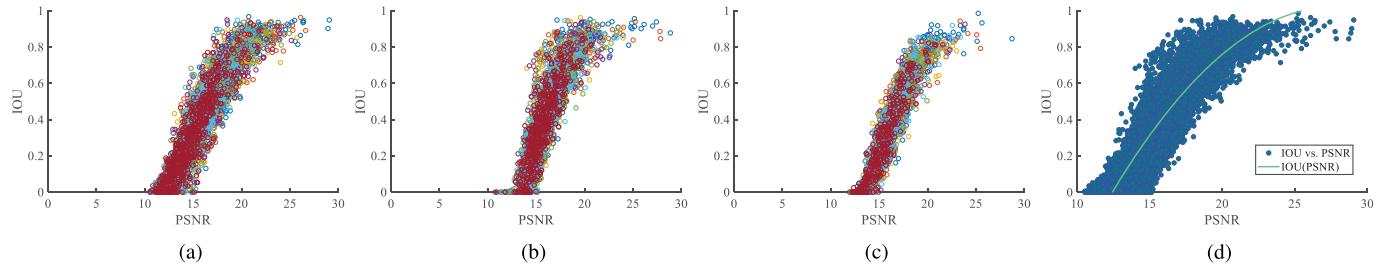


Fig. 12. Study on how the PSNR impacts the IOU degradation in videos with motion blurred attribute.

predicted region of object o . Therefore, to develop the context-aware optimization algorithm, we need to explore the interactions among ρ , $PSNR$, and IOU .

To explore such interactions, we leverage an open dataset [48] which contains 100 videos with nine different scene attributes, such as illumination variation (i.e., the illumination in the target region is significantly changed) and motion blur (i.e., the target region is blurred due to the motion of the target or camera). We measure how IOU varies in terms of $PSNR$ using videos with the same scene attribute, where we gradually increase the frame interval of two measured images, such as Frame 1 to Frame 3, Frame 1 to Frame 4, etc. Figs. 12a to 12c depict the measurement of three videos with motion blur. We observe that different videos that have the same major attribute obtain similar shape of $IOU(PSNR)$ (due to the page limit, we only show the results of three videos with motion blur). If we integrate all the samples into one figure, as illustrated in Fig. 12, we can achieve a regression-based model which describes an object tracking accuracy function in terms of $PSNR$ for a specific scenario (e.g., for motion blur, $IOU(PSNR) = -0.004335PSNR^2 + 0.2411PSNR - 2.328$).

Given the above discussion and analysis, we formulate the image offloading frequency decision as an optimization problem \mathcal{P}_2 , which aims to achieve an optimal ρ to balance the MAR device's energy consumption and tracking accuracy loss. E_{obj} and E_{trk} are estimated per frame energy consumption of edge-based object detection and local object tracking based on current MAR system configurations determined by LEAF, respectively. Two positive weight parameters θ_1 and θ_2 are introduced to characterize the offloading preference of an MAR client. For example, given a larger θ_1 and a smaller θ_2 , the decision made by the offloading frequency orchestrator will be more aggressive on saving MAR device's battery life, and vice versa. Current scene attribute is denoted by ϖ

$$\begin{aligned} \mathcal{P}_2 : \min_{\{\rho\}} \quad & J = \theta_1 \frac{E_{obj} + E_{trk}\rho}{1 + \rho} - \theta_2 IOU_{trk}^{\varpi}(PSNR(\rho)) \\ \text{s.t.} \quad & \rho \in \{0, 1, 2, \dots\}. \end{aligned} \quad (20)$$

Based on \mathcal{P}_2 , we develop an adaptive image offloading (AIO) algorithm implemented in the offloading frequency orchestrator, where the pseudo code of it is presented in Algorithm 7.2. The proposed AIO will be triggered after the MAR device executes a successful object detection and receives the corresponding detection results from the edge server. First, the AIO calculates the $PSNR$ of the current image frame and v_i (i.e., transient gradient of $PSNR$) and updates them in sets \mathcal{P} and \mathcal{V} , respectively (i.e., line 2-3 in Algorithm 7.2). We then

estimate the current scene change rate denoted by \bar{v} . To avoid a transient outlier that impacts the precision of estimation, \bar{v} is achieved by calculating the weighted mean of elements of \mathcal{V} that are the latest updated in a short range of time (e.g., in the past 2 seconds). In this paper, we use exponential function to calculate the weights, where the element that is updated later will be allocated with a larger weight. The scene attribute estimation is out of the scope of this paper, where we assume that the AIO knows the current $IOU_{trk}^{\varpi}()$ when solving \mathcal{P}_2 . In addition, to fulfill the second principle (i.e., reducing the workload on MAR devices), the per frame energy consumption of object detection E_{obj} is estimated by the edge server via our proposed analytical model, presented in Section 5, based on the LEAF guided configurations, while the per frame energy consumption of object tracking E_{trk} is estimated locally via a preset table. Thus, the AIO can efficiently achieve the value of E_{trk} through the table in terms of its current CPU frequency. Finally, the AIO outputs an optimal ρ and the MAR device will keep performing local object tracking until ρ decreases to 0.

Algorithm 2. The AIO Algorithm

Input: ρ , E_{obj} , E_{trk} , θ_1 , θ_2 , \mathcal{P} , and \mathcal{V} .
Output: ρ , \mathcal{P} , and \mathcal{V} .

- 1: **if** Object detection = True **and** $\rho = 0$ **then**
- 2: $\mathcal{P} \leftarrow \mathcal{P} \cup \{PSNR(I_{i-1}, I_i)\}$;
- 3: $\mathcal{V} \leftarrow \mathcal{V} \cup \left\{ v_i = \frac{PSNR(I_{i-1}, I_i) - PSNR(I_{i-2}, I_{i-1})}{2} \right\}$;
- 4: $\bar{v} \leftarrow \frac{\sum_{i=n}^{i=n} v_i w_i}{\sum_{i=n}^{i=n} w_i}$, $v_i \in \mathcal{V}$;
- 5: $IOU_{trk}^{\varpi}() \leftarrow IOU_{trk}^{\varpi}()$; ▷Scene attribute estimation
- 6: $\rho \leftarrow$ solving \mathcal{P}_2 with E_{obj} , E_{trk} , θ_1 , θ_2 , \bar{v} , and $IOU_{trk}^{\varpi}()$;
- 7: **if** $\rho = 0$ **then**
- 8: Object detection \leftarrow False;
- 9: **else if** $\rho \neq 0$ **then**
- 10: $\mathcal{P} \leftarrow \mathcal{P} \cup \{PSNR(I_{i-1}, I_i)\}$;
- 11: $\mathcal{V} \leftarrow \mathcal{V} \cup \left\{ v_i = \frac{PSNR(I_{i-1}, I_i) - PSNR(I_{i-2}, I_{i-1})}{2} \right\}$;
- 12: $\rho \leftarrow \rho - 1$;
- 13: **if** $\rho = 0$ **then**
- 14: Object detection \leftarrow False;
- 15: **return** ρ , \mathcal{P} , and \mathcal{V} .

TABLE 2
Power and Duration of Promotion and Tail Phases

| P_{pro} (W) | t_{pro} (s) | P_{tail} (W) | t_{tail} (s) |
|-----------------|-------------------|-----------------|-----------------|
| 1.97 ± 0.08 | 0.034 ± 0.004 | 1.61 ± 0.15 | 0.21 ± 0.02 |

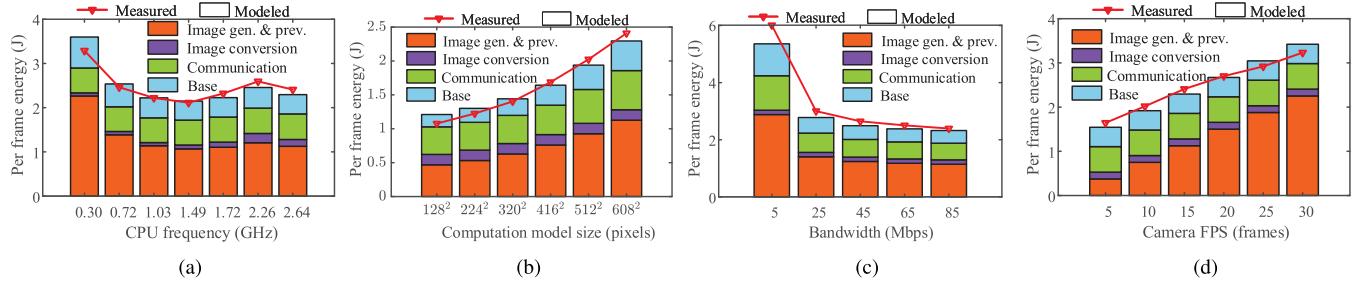


Fig. 13. Measured data versus estimated data from our proposed analytical model.

8 PERFORMANCE EVALUATION OF THE PROPOSED ANALYTICAL MODEL AND MAR SYSTEM

In this section, we evaluate both the proposed MAR analytical energy model as well as the proposed LEAF and AIO algorithms. We first validate our analytical model by comparing the estimated energy consumption with the real energy measurement (obtained from our developed testbed described in Section 3). The Mean Absolute Percentage Error (MAPE) is used for quantifying the estimation error. Then, we evaluate the per frame energy consumption, service latency, and detection accuracy of the proposed LEAF and AIO algorithms under variant bandwidth and user preferences through data-driven simulations.

8.1 Analytical Model Validation

The measured power and duration of promotion and tail phases in WiFi are shown in Table 2 (note that LTE has different values [49]). As shown in Fig. 13, we validate the proposed analytical model with respect to MAR client's CPU frequency, computation model size, allocated bandwidth, and camera FPS. Each measured data is the average of the per frame energy consumption of 1,000 image frames. The calculated MAPE of these four cases are $6.1\% \pm 3.4\%$, $7.6\% \pm 4.9\%$, $6.9\% \pm 3.9\%$, and $3.7\% \pm 2.6\%$, respectively. Therefore, our proposed energy model can estimate the MAR per frame energy consumption very well.

8.2 Performance Evaluation of LEAF

We simulate an edge-based MAR system with an edge server and multiple MAR clients. Each MAR client may select a different camera FPS, which is obtained randomly in the range of $[1, 30]$ frames.⁷ The default user preference is $\lambda_1 = 0.3$ and $\lambda_2 = 1.8$. We compare our proposed LEAF algorithm with two other algorithms summarized as follows:

- *FACT + Interactive*: It uses the FACT algorithm [16] to select the computation model size, which is optimized for the tradeoff between the service latency and the detection accuracy. As FACT does not consider the MAR client's CPU frequency scaling and radio resource allocation at the edge server, we use *Interactive* to conduct CPU frequency scaling and the radio resource is allocated evenly. Note that FACT

does not consider the energy efficiency of MAR clients either.

- *Energy-optimized only solution*: It selects the optimal CPU frequency, computation model size, and bandwidth allocation by minimizing the per frame energy consumption of MAR clients in the system without considering user preferences, which is named as MINE.

Optimality. We first validate the optimality of our proposed LEAF algorithm. As shown in Fig. 14, LEAF always obtains the minimal Q compared to the other two algorithms under variant maximum available bandwidth and user preference.

Comparison Under Variant Max. Bandwidth. We then evaluate the impact of the maximum available bandwidth on the performance of the proposed LEAF. As presented in

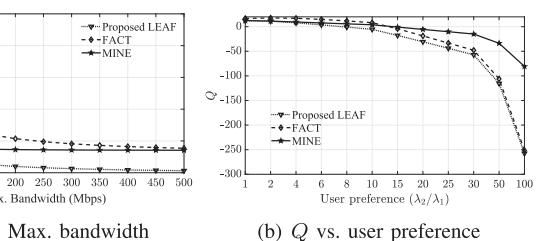


Fig. 14. Optimality.

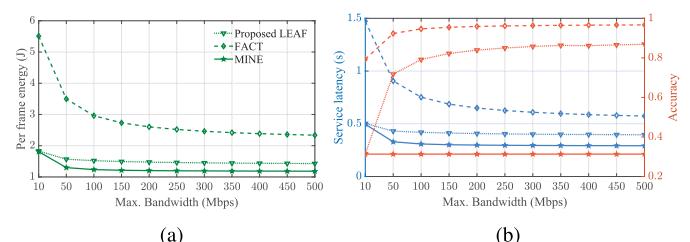


Fig. 15. System performance versus Max. bandwidth.

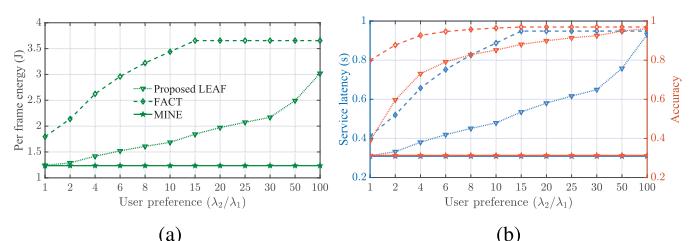


Fig. 16. System performance versus user preference.

7. Ten MAR devices are implemented in our simulation, where their camera FPS are 9, 30, 16, 23, 14, 17, 13, 2, 19, 5.

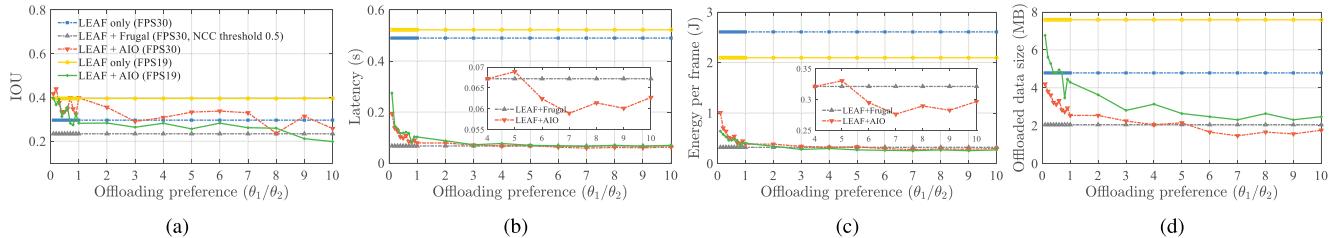


Fig. 17. System performance versus offloading preference.

Section 5.3, in practical environments, the maximum bandwidth at an edge server for serving its associated MAR clients may vary with the user distribution. For each MAR client, the value of the allocated bandwidth directly impacts not only the service latency and the per frame energy consumption but also the detection accuracy. The evaluation results are depicted in Fig. 15. (i) Compared to FACT, the proposed LEAF decreases up to 40% per frame energy consumption and 35% service latency with less than 9% loss of object detection accuracy when the max. bandwidth is 300 Mbps. The performance gap between LEAF and FACT is due to the gain derived through optimizing the clients' CPU frequency and the server radio resource allocation. (ii) Compared to MINE, the proposed LEAF significantly improves the detection accuracy at the cost of a slightly increase of the service latency and per frame energy. The performance gap between LEAF and MINE reflects the gain derived through considering the user preference.

Comparison Under Variant User Preferences. Finally, we evaluate the impact of the user preference on the performance of the proposed LEAF by varying the value of λ_2/λ_1 , as shown in Fig. 16. User preference impacts the tradeoffs among the per frame energy consumption, service latency, and detection accuracy. When λ_2/λ_1 grows, the MAR client emphasizes on the detection accuracy by trading the service latency and per frame energy. Since MINE does not consider the user preference, the variation of λ_2/λ_1 does not change its performance. (i) Compared to FACT, the proposed LEAF reduces over 20% per frame energy consumption while maintaining the same detection accuracy ($\lambda_2/\lambda_1 = 100$). (ii) Compared to MINE, the proposed LEAF is able to enhance over 50% accuracy while ensuring similar per frame energy and service latency ($\lambda_2/\lambda_1 = 2$). Fig. 16 also shows that, as compared to FACT, the proposed LEAF offers more fine-grained and diverse user preference options for MAR clients.

8.3 Performance Evaluation of LEAF + AIO

We implement our proposed image offloading orchestrator with the AIO in a device-to-device testbed that consists of an MAR device and an edge server. The MAR device works with the configurations achieved from the edge server and optimized by the proposed LEAF. In the experiment, we choose two MAR clients with camera FPS 30 and 19 for the evaluation. To make the experiment repeatable, we leverage the video frames from the open dataset [48] as the source of data ingestion. In addition, we compare our proposed AIO algorithm integrated with the LEAF with two other baselines summarized as follows:

- *LEAF + Frugal* [20]: It uses a preset normalized cross-correlation (NCC) threshold to trigger object detection

invocations. The value of NCC threshold is set to 0.5 which is experience-driven.

- *LEAF Only*: The MAR device offloads its camera captured image frames as many as possible and no local object tracker is deployed.

Comparison Under Variant Offloading Preferences. We evaluate the impact of the offloading preference on the performance of the proposed AIO by varying the value of θ_1/θ_2 , as illustrated in Fig. 17. Offloading preference influences the tradeoffs between the perception accuracy and energy efficiency of MAR devices. When θ_1/θ_2 increases, the image offloading orchestrator emphasizes on the energy efficiency by trading the perception accuracy. As Frugal sets the same trigger value for all scenarios and LEAF only does not consider the adaptive image offloading decision, the variation of θ_1/θ_2 does not change their performance. (i) Compared to Frugal, our proposed AIO improves the average IOU by 43% while decreasing the average service latency and per frame energy consumption by 12.3% and 13.9%, respectively ($\theta_1/\theta_2 = 7$). (ii) Compared to the LEAF only, our integration system not only significantly drops the latency and offloaded data size but also further improves IOU and energy efficiency of MAR devices.

9 CONCLUSION

In this paper, we proposed a user preference based energy-aware edge-based MAR system for object detection that can reduce the per frame energy consumption of MAR clients without compromising their user preferences by dynamically selecting the optimal combination of MAR configurations and radio resource allocations according to user preferences, camera FPS, and available radio resources at the edge server. To the best of our knowledge, we built the first analytical energy model for thoroughly investigating the interactions among MAR configuration parameters, user preferences, camera sampling rate, and per frame energy consumption in edge-based MAR systems. Based on the proposed analytical model, we proposed the LEAF optimization algorithm to guide the optimal MAR configurations and resource allocations. The performance of the proposed analytical model is validated against real energy measurements from our testbed and the LEAF algorithm is evaluated through extensive data-driven simulations. Additionally, we studied and implemented object tracking to further improve the energy efficiency of our proposed edge-based MAR system.

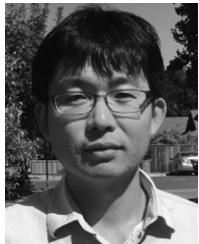
REFERENCES

- [1] L. N. Huynh, Y. Lee, and R. K. Balan, "DeepMon: Mobile GPU-based deep learning framework for continuous vision applications," in *Proc. ACM 15th Annu. Int. Conf. Mobile Syst., Appl. Serv.*, 2017, pp. 82–95.

- [2] H. Wang, B. Kim, J. Xie, and Z. Han, "How is energy consumed in smartphone deep learning apps? executing locally vs. remotely," in *Proc. IEEE Glob. Commun. Conf.*, 2019, pp. 1–6.
- [3] Y. Xiao *et al.*, "Modeling energy consumption of data transmission over Wi-Fi," *IEEE Trans. Mobile Comput.*, vol. 13, no. 8, pp. 1760–1773, Aug. 2013.
- [4] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4 G LTE networks," in *Proc. ACM 10th Int. Conf. Mobile Syst., Appl., Serv.*, 2012, pp. 225–238.
- [5] A. Shye, B. Scholbrock, and G. Memik, "Into the wild: Studying real user activity patterns to guide power optimizations for mobile architectures," in *Proc. IEEE/ACM Int. Symp. Microarchitecture*, 2009, pp. 168–178.
- [6] M. J. Walker *et al.*, "Accurate and stable run-time power modeling for mobile and embedded CPUs," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 36, no. 1, pp. 106–119, Jan. 2017.
- [7] K. DeVogeeler, G. Memmi, P. Jouvelot, and F. Coelho, "Modeling the temperature bias of power consumption for nanometer-scale CPUs in application processors," in *Proc. IEEE Int. Conf. Embedded Comput. Syst., Architect., Model., Simul.*, 2014, pp. 172–180.
- [8] F. Xu, Y. Liu, Q. Li, and Y. Zhang, "V-edge: Fast self-constructive power modeling of smartphones based on battery voltage dynamics," in *Proc. USENIX Symp. Netw. Syst. Des. Implementation*, 2013, pp. 43–55.
- [9] A. Pathak, Y. C. Hu, M. Zhang, P. Bahl, and Y.-M. Wang, "Fine-grained power modeling for smartphones using system call tracing," in *Proc. 6th Conf. Comput. Syst.*, 2011, pp. 153–168.
- [10] A. Pathak, Y. C. Hu, and M. Zhang, "Where is the energy spent inside my app? fine grained energy accounting on smartphones with eprof," in *Proc. 7th ACM Eur. Conf. Comput. Syst.*, 2012, pp. 29–42.
- [11] W. Hu and G. Cao, "Energy-aware CPU frequency scaling for mobile video streaming," in *Proc. IEEE 37th Int. Conf. Distrib. Comput. Syst.*, 2017, pp. 2314–2321.
- [12] W. Hu and G. Cao, "Energy optimization through traffic aggregation in wireless networks," in *Proc. IEEE Conf. Comput. Commun.*, 2014, pp. 916–924.
- [13] Y. Geng, Y. Yang, and G. Cao, "Energy-efficient computation offloading for multicore-based mobile devices," in *Proc. IEEE Conf. Comput. Commun.*, 2018, pp. 46–54.
- [14] A. P. Miettinen and J. K. Nurminen, "Energy efficiency of mobile clients in cloud computing," *HotCloud*, vol. 10, no. 4, 2010, Art. no. 19.
- [15] X. Ran, H. Chen, X. Zhu, Z. Liu, and J. Chen, "DeepDecision: A mobile deep learning framework for edge video analytics," in *Proc. IEEE Conf. Comput. Commun.*, 2018, pp. 1421–1429.
- [16] Q. Liu, S. Huang, J. Opadere, and T. Han, "An edge network orchestrator for mobile augmented reality," in *Proc. IEEE Conf. Comput. Commun.*, 2018, pp. 756–764.
- [17] J. Hanhriova, T. Kämäriäinen, S. Seppälä, M. Siekkinen, V. Hirvialo, and A. Ylä-Jääski, "Latency and throughput characterization of convolutional neural networks for mobile computer vision," in *Proc. 9th ACM Multimedia Syst. Conf.*, 2018, pp. 204–215.
- [18] L. Liu, H. Li, and M. Gruteser, "Edge assisted real-time object detection for mobile augmented reality," in *Proc. 25th ACM Annu. Int. Conf. Mobile Comput. Netw.*, 2019, pp. 1–16.
- [19] H. Wang, J. Xie, and T. Han, "A smart service rebuilding scheme across cloudlets via mobile AR frame feature mapping," in *Proc. IEEE Int. Conf. Commun.*, 2018, pp. 1–6.
- [20] K. Apicharttrisorn, X. Ran, J. Chen, S. V. Krishnamurthy, and A. K. Roy-Chowdhury, "Frugal following: Power thrifty object detection and tracking for mobile augmented reality," in *Proc. 17th ACM Conf. Embedded Netw. Sensor Syst.*, 2019, pp. 96–109.
- [21] H. Wang, B. Kim, J. Xie, and Z. Han, "E-auto: A communication scheme for connected vehicles with edge-assisted autonomous driving," in *Proc. IEEE Int. Conf. Commun.*, 2019, pp. 1–6.
- [22] A. Mallik and J. Xie, "H.264 video encoding-based edge-assisted mobile AR systems: Network and energy issues," in *Proc. IEEE Int. Conf. Commun.*, 2022, pp. 1–6.
- [23] J.-J. Chen, C.-Y. Yang, T.-W. Kuo, and C.-S. Shih, "Energy-efficient real-time task scheduling in multiprocessor DVFS systems," in *Proc. IEEE Asia South Pacific Des. Automat. Conf.*, 2007, pp. 342–349.
- [24] J. Kwak, O. Choi, S. Chong, and P. Mohapatra, "Dynamic speed scaling for energy minimization in delay-tolerant smartphone applications," in *Proc. IEEE Conf. Comput. Commun.*, 2014, pp. 2292–2300.
- [25] W. Y. Lee, "Energy-saving DVFS scheduling of multiple periodic real-time tasks on multi-core processors," in *Proc. 13th IEEE/ACM Int. Symp. Distrib. Simul. Real Time Appl.*, 2009, pp. 216–223.
- [26] H. Wang and J. Xie, "User preference based energy-aware mobile AR system with edge computing," in *Proc. Conf. Comput. Commun.*, 2020, pp. 1379–1388.
- [27] H. Wang, B. Kim, J. Xie, and Z. Han, "Energy drain of the object detection processing pipeline for mobile devices: Analysis and implications," *IEEE Trans. Green Commun. Netw.*, vol. 5, no. 1, pp. 41–60, Mar. 2020.
- [28] Android. ImageFormat, Oct. 2020. Accessed: Oct. 2020. [Online]. Available: <https://developer.android.com/reference/android/graphics/ImageFormat>
- [29] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," 2018, *arXiv:1804.02767*.
- [30] J. Redmon, Darknet: Open source neural networks in c, 2013–2016. [Online]. Available: <http://pjreddie.com/darknet/>
- [31] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [32] Monsoon power monitor, [Online]. Available: <https://www.msoon.com/>
- [33] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.
- [34] K. Ha, Z. Chen, W. Hu, W. Richter, P. Pillai, and M. Satyanarayanan, "Towards wearable cognitive assistance," in *Proc. 12th ACM Annu. Int. Conf. Mobile Systems, Appl., Serv.*, 2014, pp. 68–81.
- [35] SketchAR, Oct. 2020. [Online]. Available: <https://itunes.apple.com/us/app/sketchar-drawing-using-augmented-reality/id1221482822?l=ru&mt=8>
- [36] C. J. Willmott and K. Matsuura, "Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance," *Climate Res.*, vol. 30, no. 1, pp. 79–82, 2005.
- [37] K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*. Hoboken, NJ, USA: Wiley, 2001.
- [38] R. T. Marler and J. S. Arora, "The weighted sum method for multi-objective optimization: New insights," *Struct. Multidisciplinary Optim.*, vol. 41, no. 6, pp. 853–862, 2010.
- [39] P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan, "Mixed-integer nonlinear optimization," *Acta Numerica*, vol. 22, pp. 1–131, 2013.
- [40] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear gauss–seidel method under convex constraints," *Operations Res. Lett.*, vol. 26, no. 3, pp. 127–136, 2000.
- [41] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [42] ARCore, Oct. 2020. [Online]. Available: <https://developers.google.com/ar/develop>
- [43] ARKit, Oct. 2020. [Online]. Available: <https://developer.apple.com/augmented-reality/>
- [44] G. Bradski, "The OpenCV library," *Dr Dobb's J. Softw. Tools*, vol. 120, pp. 122–125, 2000.
- [45] JavaCV, Sep. 2020. Accessed: Sep. 10, 2020. [Online]. Available: <https://github.com/bytedeco/javacv>
- [46] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 3, pp. 583–596, Mar. 2015.
- [47] Y. Wu, J. Lim, and M.-H. Yang, "Online object tracking: A benchmark," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 2411–2418.
- [48] Y. Wu, J. Lim, and M.-H. Yang, "Online object tracking: A benchmark," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 2411–2418.
- [49] W. Hu and G. Cao, "Energy-aware video streaming on smartphones," in *Proc. IEEE Conf. Comput. Commun.*, 2015, pp. 1185–1193.



Haoxin Wang (Member, IEEE) received the B.S. degree in control science and engineering from the Harbin Institute of Technology, China, in 2015, and the PhD degree in electrical and computer engineering from The University of North Carolina at Charlotte, in 2020. He is currently a research scientist with Toyota Motor North America, InfoTech Labs, where he leads the "Edge Computing" project. His research interests include edge computing for connected and autonomous vehicles, applied machine learning for intelligent systems, and energy-efficient mobile computing systems.



BaekGyu Kim (Member, IEEE) received the BS and MS degrees in electrical engineering and computer science from Kyungpook National University, South Korea, and the PhD degree in computer science from the University of Pennsylvania. He is currently a principal researcher with Toyota Motor North America, InfoTech Labs. His research interests include software platform technologies for connected cars and model based software development for high-assurance systems.



Zhu Han (Fellow, IEEE) received the BS degree in electronic engineering from Tsinghua University, in 1997, and the MS and PhD degrees in electrical and computer engineering from the University of Maryland, College Park, in 1999 and 2003, respectively. From 2000 to 2002, he was an R&D engineer with JDSU, Germantown, Maryland. From 2003 to 2006, he was a research associate with the University of Maryland. From 2006 to 2008, he was an assistant professor with Boise State University, Idaho. Currently, he is a John and Rebecca Moores

professor with the Electrical and Computer Engineering Department, University of Houston, Texas. His research interests include wireless resource allocation and management, wireless communications and networking, game theory, Big Data analysis, security, and smart grid. He received an NSF Career Award, in 2010, the Fred W. Ellersick Prize of the *IEEE Communication Society*, in 2011, the EURASIP Best Paper Award for the *Journal on Advances in Signal Processing*, in 2015, IEEE Leonard G. Abraham Prize in the field of Communications Systems (best paper award in IEEE JSAC), in 2016, and several best paper awards in IEEE conferences. He was an IEEE Communications Society distinguished lecturer from 2015-2018, AAAS fellow since 2019, and ACM distinguished member since 2019. He is a 1% highly cited researcher since 2017 according to Web of Science. He is also the winner of the 2021 IEEE Kiyo Tomiyasu award, for outstanding early to mid-career contributions to technologies holding the promise of innovative applications, with the following citation: "for contributions to game theory and distributed management of autonomous communication networks."



Jiang Xie (Fellow, IEEE) received the BE degree from Tsinghua University, Beijing, China, the MPhil degree from the Hong Kong University of Science and Technology, and the MS and PhD degrees from the Georgia Institute of Technology, all in electrical and computer engineering. She joined the Department of Electrical and Computer Engineering, University of North Carolina, Charlotte (UNC Charlotte) as an assistant professor, in Aug. 2004, where she is currently a full professor. Her current research interests include resource and mobility

management in wireless networks, mobile computing, Internet of Things, and cloud/edge computing. She is on the editorial boards of the *IEEE Transactions on Wireless Communications*, *IEEE Transactions on Sustainable Computing*, and *Journal of Network and Computer Applications* (Elsevier). She received the US National Science Foundation (NSF) Faculty Early Career Development (CAREER) Award, in 2010, a Best Paper Award from IEEE Global Communications Conference (Globecom 2017), a Best Paper Award from IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2010), and a Graduate Teaching Excellence Award from the College of Engineering at UNC Charlotte, in 2007. She is a senior member of the ACM.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csl.