

## REVISION LETTER (SEPTEMBER 2020)

ISJ-RE-20-10082

---

Dear Dr. Amir Aghdam,

Thank you for the opportunity to revise and improve our manuscript “EAPS: Edge-Assisted Predictive Sleep Scheduling for 802.11 IoT Stations”. We appreciate the careful review and constructive suggestions. All the comments provided by reviewers have been addressed with corresponding changes made directly to the manuscript where appropriate. Appended to this letter is our point-by-point response to the comments raised by the reviewers. We have used blue color to show the new and revised materials.

### —ASSOCIATE EDITOR—

Recommendation:

*REQUIRED Comments to the Author: I have got comments from three reviewers, and their comments are rather mixed. Specifically, Reviewer 1 criticizes that the paper does not compare with the power saving mechanisms in the latest 802.11ba standard, which is known as extremely low power Wi-Fi for massive IoTs, but rather compare with 802.11ax (WiFi 6 that is for high speed access). I agree with this comment. Reviewer 1 also questioned about the infrastructure based IoT deployment scenario, but I think this scenario itself is acceptable, e.g., the NB-IoT based on cellular networks has seen huge success. Reviewer 1 also suggested to add performance comparisons with some other relevant mechanisms. This would be very helpful for improving the quality of the paper.*

*Reviewer 2 thinks that the novelty aspect of the approach is a bit trivial, but he/she appreciates the authors' contribution of using machine learning (ML) based predictive mechanisms for energy efficiency optimization, and suggested another advanced ML approach to use.*

*Reviewer 3 is generally happy with the paper's quality.*

*Based on my own reading, I feel that this paper is well aligned with the journal's scope by providing a system-level study based on a test-bed the authors built. Both Reviewer 1 and Reviewer 2 have provided high-quality comments, which I think the authors should seriously taken into consideration if they decide to undergo a major revision.*

We have made every attempt to fully address your suggestions and the reviewer's comments. Please find our replies and a summary of the improvements as follows.

### —REVIEWER 1—

Recommendation: Reject

**Comment 1.1.** *Paper ignores latest developments of IoT in the context of 802.11. Specifically, it doesn't take into account IEEE 802.11ba to be released this year that includes several mechanisms for power efficiency and duty cycle support including WUR (wake-up-radio).*

Dear Reviewer: Thank you for your comments. We have made every attempt to fully address your comments in the revised manuscript.

**Response 1.1.** We have modified the paper to address this comment. Before presenting the modifications, we would like to emphasize that the method proposed in this paper can be applied to a variety of 802.11 standards. Specifically, EAPS can be used to enhance 802.11n, 802.11ac, 802.11ax, and also can be integrated with and benefit from 802.11ba.

In order to address this comment we have made the following changes:

- We modified the Introduction section (please see page 1, line 54) to review the existing power saving mechanisms of 802.11 (please see page 1, line 80). Specifically, in addition to PSM and Adaptive PSM, we also discuss TWT, which is introduced in 802.11ah and available in 802.11ax. We also discuss 802.11ba. We emphasize that, although the power saving methods of 802.11 allow stations to duty cycle their radio operation, they do not consider the effect of buffering, interference, channel access delay, and traffic category on downlink delivery delay.
- To discuss how wake-up radio technologies can be used to enhance EAPS, we have added Section VI and a discussion on wake-up radio. Please see page 11, line 17: *Wake-up Radio (WUR) mechanisms such as 802.11ba [1], [2] can be used to enhance EAPS in several ways. For example, as soon as a station finishes sending its uplink packet(s), the primary radio*

can switch into sleep mode, and the station will receive the schedule message via its low-power WUR. The primary radio will then wake up at the scheduled time to receive the downlink packet. To further reduce the energy caused by prediction inaccuracy, the WUR can be scheduled to wake up at  $\delta_c - 2\sigma$  and wait for a command to wake up the primary radio. Another example, once the downlink packet arrives on the wired interface of the AP, the AP uses EAPS to compute the delivery delay of the packet ( $\delta_c$ ). Assuming that the wake-up delay of the primary radio is  $\beta$  [1], the AP sends the wake-up packet at  $\delta_c - \beta$  to make sure the station's primary radio will be awake on time for downlink delivery.

**Comment 1.2.** *802.11ax is by no means the best 802.11 standard for IoT applications (i.e. OFDM based)*

**Response 1.2.** Thank you for your comment. To show that the proposed method can be used to enhance a variety of power saving methods, we have made the following changes to the manuscript:

- We modified the Introduction section (page 2, line 5) to better clarify the various energy saving options available to stations. In the revised text we highlight that we do not associate the proposed method with a particular standard. For example, both APSD (available in 802.11n and 802.11ac) and TWT (available in 802.11ax and 802.11ah) can be used to implement EAPS.
- Also, in the Conclusion section (please see page 12, line 20): *EAPS can be considered as an enhancement of Automatic Power Save Delivery (APSD) and Target Wake-up Time (TWT) (introduced in 802.11ax). For example, with the next generation of IoT stations that support TWT, they can set up their wake up time based on the sleep schedule computed by the Access Point (AP). By protecting IoT stations against the effect of concurrent traffic and interference, EAPS is a particularly useful method in scenarios where both regular and IoT stations exist. For example, EAPS can reduce the energy cost of households and reduce the impact of IoT on global ICT footprint.*
- A few other changes to ensure writing consistency can be found in page 1, line 4.

**Comment 1.3.** *Due to coverage restrictions most IoT scenarios, even under 802.11, rely on multi-hop ad-hoc mesh topologies. The infrastructure deployment presented in this manuscript is not fully representative of IoT deployments.*

**Response 1.3.** Thank you for your comment. To address this comments, we have added a discussion to Section VI, page 11, line 33, as follows:

*As discussed in Section III-C, the primary types of networks considered in this paper are smart home environments where an AP is connected to an Internet modem, and campus and business deployments where APs communicate via an Ethernet infrastructure. EAPS can also be used in mesh deployments. In this case, the backbone communication between APs (mesh nodes) introduces a wireless-to-wireless switching delay. This delay primarily depends on the bandwidth difference between the backbone (AP-AP) and access (AP-station) links. For example, assume a 160 MHz channel (in the 5 GHz band) is used to form the backbone, while each AP operates on a 20 MHz or 40 MHz channel (in the 2.4 GHz band). This configuration is prevalent because most of the existing IoT stations operate in the 2.4 GHz band, and WiFi mesh systems are usually tri-band and dedicate a channel (in the 5 GHz band) to their backbone. With this configuration, the delay caused by the backbone would be negligible and a method similar to that mentioned in Section III-C can be used to measure the delay from each AP to the server. If backbone links suffer from congestion and significant interference, EAPS can be used to predict packet switching delay over the backbone. As an alternative, more efficient strategy, EAPS could run on a central machine and allow the stations to receive their downlink packet from the AP offering the lowest delay. We leave these enhancements as future works.*

**Comment 1.4.** *How does the mechanism behave in the context of other well known 802.11 improvements like channel bonding and packet bursting?*

**Response 1.4.** Thank you for your comment. The primary WiFi standard that we used in our testbed is 802.11n. We chose this because we noted that to ensure maximum compatibility with existing WiFi deployments, most of the IoT devices in the market use the 802.11n standard. For example, Ring cameras, Nest cameras, Schlage locks, and Lixar lights all use the 802.11n standard. We also noted that these devices do not leverage channel bonding and always operate on a 20 MHz channel.

In order to address this comment, and to account for current and future devices that will rely on channel bonding, we have tried to include the channel bandwidth used by each station as a feature in our delay prediction models. However, after digging into WiFi drivers and analyzing their codes, we found out that collecting this information requires accessing NIC registers and driver modification, which is mainly due to the dynamic nature of channel and rate adjustment algorithms used in the NICs. We plan to pursue this suggestion as one of our future research directions.

In terms of packet bursting, we do include features that allow our models to take it into account. Specifically, since packet bursting in the MAC layer is defined by the Transmission Opportunity (TXOP) of each Access Category (AC), we include the size of driver queues in our delay prediction models. To address this comment, we have made the following changes:

- In page 5, line 109: *Considering the size of these queues allows the prediction models to take into account the effect of packet aggregation and packet bursting. Specifically, for each Access Category (AC), packet aggregation is applied if a station's queue includes more than one packet. Also, depending on the AC, a burst of packets may be sent without having to contend for the channel on a per-packet basis.*

- In page 6, line 51: We clarify that packet bursting after receiving a NULL packet from the station is possible:  
*The AP responds by sending one or multiple downlink packets starting at  $t_s$ . As per the 802.11e amendment, stations can start a Transmission Opportunity (TXOP) for downlink delivery. For example, if the traffic belongs to the voice Access Category, the AP uses a 1.504 ms slot (as long as packets exist) for downlink delivery.*

We also highlight that we consider packet bursting in:

- The definition of a transaction. Please see page 3, line 86.
- Formulating our traffic characterization models. Please see page 6, line 76 and page 7, line 10.
- Discussing the LSTM model and why it is more suitable to capture the effect of packet bursting. Please see page 9, line 40.
- Discussing the importance of using incoming traffic rate over the wired interface to better capture the burstiness of incoming traffic. Please see page 5, line 39.

**Comment 1.5.** *How does the mechanism adapt to very low end 8-bit embedded devices?*

**Response 1.5.** Thank you for your comment. We would like to point out that the proposed method performs packet scheduling fully on access point. Therefore, no extra overhead is imposed on the stations. To clarify this point, we have revised page 2, line 76, as follows: *EAPS runs at network edge to ensure quick sleep schedule computation, which implies that all the necessary computations to calculate sleep schedules are performed by AP to avoid any additional processing overhead on resource-constrained IoT stations. Therefore, regardless of the IoT station's processor type, EAPS can be employed by any 802.11-based IoT devices.*

We also point out that we are not aware of any 8-bit platform using the 802.11 standard. Specifically, we noticed that the existing platforms are mostly based on 32-bit ARM processors (such as Cortex-M and Cortex-R series). Nevertheless, since EAPS is fully implemented on access points, it does not enforce any sort of limitation on IoT device's processor capabilities.

**Comment 1.6.** *How does the mechanism compare to that of other well known IoT phy/link layer technologies like 802.15.4 or BLE ?*

**Response 1.6.** Thank you for your comment. We have revised the paper to address this comment. Compared to 802.11, both 802.15.4 and BLE are used in scenarios where the properties of data flows are known, or when the data exchange size is small and sporadic. For example, 802.15.4 is mostly used to collect sensory data such as temperature and humidity. In contrast, the high data rate of 802.11 makes it a suitable technology for wireless connectivity of devices that require higher data rate, such as surveillance cameras. Also, since 802.11 access points are available everywhere, users usually prefer to use WiFi compatible devices, instead of purchasing new hubs to connect 802.15.4 and BLE devices to their home network. However, achieving both timeliness and energy efficiency is more challenging with the 802.11 standard because: these access points are used for both regular user connectivity and IoT connectivity, the traffic pattern is not known a priori (e.g., a camera streams when motion is detected), the amount of data exchanged with 802.11 access points is higher, and the type and size of traffic exchanged with the access points varies.

The improvements made to the manuscript are as follows:

- In page 1, line 27: *First, compared to 802.15.4 and BLE, the 802.11 standard offers higher data rates, and compared to cellular communication, 802.11 operates in unlicensed bands. In addition to reducing costs, these features are particularly useful in domains such as video surveillance, industrial control, and medical monitoring, where high bandwidth is necessary.*
- In page 1, line 44: *Nevertheless, achieving both energy efficiency and timeliness using the 802.11 standards is more challenging compared to 802.15.4 and BLE. Specifically, 802.15.4 and BLE are used in scenarios where the properties of data flows are known (e.g., WirelessHART), or when the data exchange size is small and sporadic (e.g., Nest BLE temperature sensors). In contrast, for example in a smart home scenario, an increasing number of devices share a wireless infrastructure to dynamically exchange various types of flows such as voice, video, and background.*

## —REVIEWER 2—

Recommendation: Major Revision

*Comments: Overall I think this is a well written paper. Although the novelty aspect of the technical approach is bit trivial, the ML approach has been thoroughly evaluated and also it took into account measuring all the necessary metrics, such as size of the dataset and value of time steps ( $k$ ).*

*Although, the results shown in the paper for LSTM outperform, in most aspects, other ML approaches they used (RFR, GBR, etc.), it still has the limitations of slower convergence time and the need for larger data set as compared to other Neural Network algorithms such as ANN (or MLPs). For this, I would suggest considering MLPs (Multi-Layer Perceptrons) since it converges much faster and does not require large dataset as with LSTM and it can also work well with time-series problems.*

Dear Reviewer: Thank you for your comments. We have made every attempt to fully address your comments in the revised manuscript.

To address this comment, and as suggested by the reviewers, we added MLP (Multi-Layer Perceptron) to the paper. The main changes applied to the paper are as follows:

- Figures 5, 6, 7, 8, and 10 are updated to include the performance comparison of MLP along with RFR, ETR, GBR, HBR, and LSTM. We have updated the figures with markers and modified their formats to ensure the presentation of results is clear. This allowed us to update and fit all the graphs in the revised manuscript without violating the 12 page limit enforced by the journal office.
- In page 8, line 39 we provide an overview of MLP: “MLP is one such variant of feed-forward neural networks that does not allow cycles, thereby resulting in the data progressing in a single direction over the network. Considering the one-direction data movement, one of the biggest drawbacks of using such a network is its lack of memory, i.e., it treats each instance of the input time-series independently.”
- In page 8, line 78 we discuss the convergence time of the prediction models. We show that, due to the higher complexity of neural networks, MLP requires 20000 and 35000 data points for ND and HD scenarios, respectively. Whereas LSTM requires 30000 and 50000 data points for its performance to converge under the ND and HD scenarios, respectively. Also, in page 9, line 28 we add that in scenarios where it is not possible to collect large datasets for training, either ETR or MLP can also be used.
- In page 9, line 4 we summarize the average MAE of all algorithms: “Averaged over all AC, the MAE (in millisecond) of algorithms in the ND scenario are: RFR: 1.43, ETR: 1.26, GBR: 1.49, HBR: 1.28, MLP: 1.24, and LSTM: 1.16. For the HD scenarios the MAE values are: RFR: 2.49, ETR: 2.17, GBR: 2.69, HBR: 2.27, MLP: 2.12, and LSTM: 2.01. On average for the ND and HD scenarios, the MAE of LSTM is 14% lower for all ACs, compared to the average MAE of all the other machine learning algorithms. Based on these results, in Section V we use LSTM for the empirical performance evaluation of EAPS on a testbed.

——REVIEWER 3——

Recommendation: Minor Revision

Comments: *This is a paper on using their EAPS makes the energy consumption of IoT stations is as low as PSM, whereas the delay of packet delivery is close to the case where the station is always awake. In the manuscript, Linux-based access point is relatively new.*

*In some places, the author did not make clear explanation or made mistakes. The details are as follows:*

**Comment 3.1.** *1. “the concept can easily be extended to other types of qdisc scheduling strategies as well” in page 5.*

Dear Reviewer: Thank you for your comments. We have made every attempt to fully address your comments in the revised manuscript.

**Response 3.1.** We have revised the sentence as follows (please see page 5, line 62.): *It is worth mentioning that, although our implementation utilizes the PRIO qdisc (the default policy used in several Linux distributions), the concept can easily be extended to other types of qdiscs, such as Hierarchical Token Bucket (HTB).*

**Comment 3.2.** *2. “Note that the station simply uses a timer to measure  $t_2 - t_1$ ” in page 5.*

**Response 3.2.** Thank you for your comments. We have revised the text as follows to clarify this sentence (please see page 6, line 41): *Note that the AP shares the relative wake-up time with the station. Since the AP has computed wake up schedule at time  $t_1$ , the station needs to measure  $t_2 - t_1$  and subtract it from the shared value. The station can simply use a timer to measure  $t_2 - t_1$ .*

**Comment 3.3.** *3. In my opinion, “ $g(bi)$ ” should be replaced by “ $g(bi-1)$ ” in formula (4) on page 6.*

**Response 3.3.** Thank you for pointing us to this error. We have fixed the problem in the revised version. Please see page 7, line 33.

**Comment 3.4.** *Finally, the structure of the article needs to be reconsidered. Related Work was put at the end of the article, which is very strange. My suggestion is to put this part behind the Introduction part, and make the previous literature into a table, which is more clear.*

**Response 3.4.** Thank you for your comments. To address this comment, we moved and placed the Related Work section after the Introduction. Due to the space limitation set by the journal office (12 pages), we were not able to add a table. However, we have revised the Related Work section to ensure the relevant works are clearly summarized.



# EAPS: Edge-Assisted Predictive Sleep Scheduling for 802.11 IoT Stations

Jaykumar Sheth\*, Cyrus Miremadi\*, Amir Dezfouli†, and Behnam Dezfouli\*

\*Internet of Things Research Lab, Department of Computer Science and Engineering, Santa Clara University, USA  
{jsheth, cmiremadi, bdezfouli}@scu.edu

†Commonwealth Scientific and Industrial Research Organisation (CSIRO), Sydney, Australia  
{amir.dezfouli@data61.csiro.au}

**Abstract**—The broad deployment of 802.11 (a.k.a., WiFi) access points and the significant energy-efficiency improvement of these wireless transceivers has resulted in increasing interest in building 802.11-based IoT systems. Unfortunately, the energy efficiency mechanisms of 802.11 fall short when used in IoT applications, especially because they do not take into account the delays caused by various factors such as buffering, interference, and round-trip delay. In this paper, we present *edge-assisted predictive sleep scheduling* (EAPS) to adjust the sleep duration of stations while they expect downlink packets. We first implement a Linux-based access point that enables us to collect parameters affecting communication latency. Using this access point, we build a testbed that, in addition to offering traffic pattern customization, replicates the characteristics of real-world environments. We then use multiple machine learning algorithms to predict downlink packet delivery. Our empirical evaluations confirm that with EAPS the energy consumption of IoT stations is as low as PSM, whereas the delay of packet delivery is close to the case where the station is always awake.

**Index Terms**—Energy Efficiency, Delay, Wireless Communication, Machine Learning, Edge Computing.

## I. INTRODUCTION

Nowadays, in addition to regular user devices such as phones, laptops, and tablets, many IoT devices such as security cameras, smart locks, and medical devices rely on the 802.11 standard. Several reasons support the importance of this standard for IoT applications: First, compared to 802.15.4 and BLE, the 802.11 standard offers higher data rates, and compared to cellular communication, 802.11 operates in unlicensed bands. In addition to reducing costs, these features are particularly useful in domains such as video surveillance, industrial control, and medical monitoring, where high bandwidth is necessary. Second, 802.11 base stations—known as APs—are broadly deployed in various environments and provide a ready-to-use infrastructure for IoT connectivity. Third, the power consumption of 802.11 stations has been considerably reduced during the past decade by introducing various power-save mechanisms and developing enhanced low-power RF transceiver technologies such as power gating and clock gating [3]–[5]. Accordingly, 802.11-based IoT devices are prevalent in the market, such as *Ring* security camera and doorbell [6], *Nest* security camera [7], *Schlage* locks [8], and *LIFX* light bulbs [9], to mention a few. Nevertheless, achieving both energy efficiency and timeliness using the 802.11 standards is more challenging compared to 802.15.4 and BLE. Specifically, 802.15.4 and BLE are used

in scenarios where the properties of data flows are known (e.g., WirelessHART), or when the data exchange size is small and sporadic (e.g., Nest BLE temperature sensors). In contrast, for example, in a smart home scenario, an increasing number of devices share a wireless infrastructure to dynamically exchange various types of flows such as voice, video, and background.

The 802.11 standard offers multiple power-saving mechanisms to support energy-constrained stations. Power Save Mode (PSM) enables the stations to wake up periodically and check if the AP has any buffered packet(s) for them. The AP periodically broadcasts beacon packets at a certain interval, called Beacon Interval (BI), to inform the stations about their buffered packets. Stations send PS-Poll packet to the AP to request downlink delivery. PSM significantly increases communication delay because stations can only receive downlink packets after each beacon instance. The delay problem further exacerbates with the concurrent transmission of PS-Poll packets and the accumulation of downlink delivery after each beacon instance [10], [11]. To reduce communication delay with AP, Adaptive PSM (APSM) requires a station waiting for downlink packets to stay awake for a *tail time* duration (e.g., 10 ms [3]) after each packet exchange with the AP [12]. The tail time may impose idle listening and further waste of energy, especially if the delay between uplink and downlink delivery is longer than tail time. Another enhancement of PSM is APSD, which is available in 802.11n, ac, and ax. APSD allows stations to request downlink packet delivery by sending NULL uplink packets to the AP [13]. A new power-saving mechanism introduced in 802.11ax is TWT, which was primarily introduced in 802.11ah for low-power IoT communication. Using TWT, pairwise agreements between AP and stations can be established, and stations are allowed to skip receiving beacon packets. To further reduce the overhead of periodic wake-up, the new 802.11ba standard specifies the addition of a low-power WUR [1], [2]. The primary radio only wakes up when the station receives a wake-up command over the WUR or when the station needs to perform uplink transmission. Although these power-saving mechanisms allow stations to reduce their energy consumption significantly, they do not consider the effect of communication delays caused by buffering, interference, channel access method, and traffic category.

Many IoT applications require the transmission of uplink reports by station and reception of commands from a server.

For example, consider a sample medical application where an IoT device reports an event and expects to receive actuation commands in return. Another example is a security camera that transfers an image whenever motion is detected and waits for a command to stream video if a particular object is detected.

After the transmission of uplink packet(s), the IoT station has five<sup>1</sup> options before receiving downlink packet(s):

- (i) use Continuously Active Mode (CAM),
- (ii) use PSM and return to sleep mode and wake up during the next beacon period,
- (iii) use APSM and stay in awake mode for a short time duration,
- (iv) use APSD or TWT and periodically check if the downlink packet has arrived,
- (v) use APSD or TWT and wake up when the downlink packet is about to be delivered.

Case (i) minimizes delay but does not offer any power efficiency. Case (ii) causes long end-to-end delays [14], [15] because the station has to wait until the next beacon instance, even if the actual downlink delivery delay is less than the time remaining until the next beacon. Besides, the delay considerably increases when the station and server need to complete multiple rounds of packet exchange to make a decision<sup>2</sup>. Case (iii) is effective if the delivery delay is short; otherwise, this case results in power waste in tail time. Case (iv) results in periodic wakeups and unnecessarily increases channel access contention. Therefore, none of these cases are suitable for applications where both delay and energy efficiency are the essential performance metrics. Applying case (v) requires an accurate estimation of the delay between uplink and downlink packets, which is composed of the following delay components: First, the uplink packet received over the wireless interface must be sent over the wired interface. The second component is the interval between the instance the packet leaves the AP until a response is received from the server. Third, once the reply is received, the packet must compete with other downlink packets and be delivered to the station in awake mode. In this paper, we show that computing the third delay component is particularly challenging because it depends on various factors, including channel utilization, the intensity of uplink and downlink communication, access category of packets, and AP's buffer status. This is also verified by the recent studies that show the delay experienced at AP is more than 60% of the total communication delay between a station and server [17], [18]. Besides, the buffering mechanisms employed in Linux's queuing discipline (qdisc) and wireless driver further complicate the modeling and prediction of these delay components [18]–[21]. For example, Intel's IWL and Qualcomm's ath9k and ath10 drivers perform packet scheduling; however, these algorithms are heuristically designed by vendors—further complicating delay estimations.

<sup>1</sup>We did not include WUR here because it has not been released yet, and we are not sure if it will be integrated with a wide variety of IoT devices. However, we will discuss in Section VI how using a WUR can be used to enhance the methods proposed in their paper.

<sup>2</sup>Not only for IoT applications, studies show that every 10 ms increase in network access results in a 1000 ms increase in page load time [16].

In this paper, we propose a novel mechanism called *edge-assisted predictive sleep scheduling* (EAPS) to reduce the idle listening time and energy consumption of stations when waiting for downlink packets. At a high level, EAPS works as follows: Once an uplink packet is received from an IoT station, the delivery delay is computed using machine learning techniques. The estimated delay is then conveyed to the station using a high-priority data-plane queue. The station then switches into sleep mode and wakes up at the scheduled time to retrieve downlink packet.

This paper introduces the following contributions: *First*, we present the implementation of a Linux-based AP with new and modified kernel and user-space modules to keep track of the system operation in terms of parameters such as incoming and outgoing transmissions, buffer status, and channel utilization. *Second*, we introduce multiple traffic characterization methods. Then, using the modified AP, we build a configurable testbed that allows us to generate various traffic patterns similar to real-world deployments. *Third*, the information collected across the 802.11 stack is fed into a user-space module, which estimates the delay components. We focus on wired-to-wireless switching delays and their prediction using various machine learning algorithms under different traffic scenarios. EAPS runs at network edge to ensure quick sleep schedule computation, which implies that all the necessary computations to calculate sleep schedules are performed by AP to avoid any additional processing overhead on resource-constrained IoT stations. Therefore, regardless of the IoT station's processor type, EAPS can be employed by any 802.11-based IoT devices. *Fourth*, we perform an empirical evaluation of the impact of delay prediction on both energy efficiency and timeliness in scenarios where IoT stations communicate with cloud and edge servers. In terms of delay, EAPS outperforms PSM by 45% in the cloud scenario and by 84% in the edge scenario. The energy consumption of EAPS is 26% lower in the cloud scenario and 6% in the edge scenario, compared to PSM. In the edge scenario, the delay of EAPS is close to that of APSM, while its energy efficiency is improved by 37%. In the cloud scenario, EAPS improves delay and energy efficiency by 41% and 46%, respectively, compared to APSM.

The rest of this paper is organized as follows. We overview the related works in Section II. We present delay components and implementation details of the AP in Section III. We present the edge-assisted sleep scheduling mechanism in Section IV. Section V presents empirical performance evaluations. Section VII concludes the paper.

## II. RELATED WORK

Peck et al. [22] propose PSM with adaptive wake-up (PSM-AW), which includes a metric called *PSM penalty* to enable the stations to establish their desired energy-delay trade-off. The authors define server delay as the total delay between sending a request to a server and receiving a reply. Based on Round-Trip Time (RTT) variations, the sleep duration is dynamically adjusted to satisfy the desired trade-off. Also, the size of the history window of server variations is dynamically tuned

based on the range of server delay variations. Compared to our work, PSM-AW [22] only considers AP-server RTT, thereby ignoring the variable, long impact of downlink wireless communication delay. Besides, since RTT sampling and averaging are performed by stations, it adds additional load on resource-constrained stations. Furthermore, delay estimation is directly affected by station-server communication, and estimation accuracy drops as the interval between transactions increases. In contrast, our work does not impose overhead on stations, and once a model is trained, it does not rely on ongoing communication to compute sleep schedules. Jang et al. [12] proposed an adaptive tail time adjustment mechanism by relying on inter-packet arrival delays. A moving average scheme is used to predict inter-packet arrival delay when a burst of packets arrives at a station. The station stays in the awake mode if the next packet arrival time is before the tail time expiry. If packet delivery is after the tail time expiry, the station might extend its tail time based on the outcome of an energy-delay trade-off model. In contrast to our work, neither [22] nor [12] considers the impact of buffering and channel access delay as variable, essential components of downlink delivery delay. Furthermore, the effectiveness of these approaches highly depends on the burst length and variability of end-to-end delays. Specifically, the moving average scheme employed in [12] would not be effective in IoT scenarios where most of the bursts are short-lived. Sui et al. [23] propose WiFiSeer, a centralized decision-making system to help stations choose the AP providing the shortest delay. WiFiSeer works in two phases: During the learning phase, a set of parameters (such as RSSI, RTT) are pulled every minute from all APs using SNMP. Then a random forest model is trained to generate a two-class learning model for classifying APs into high latency and low latency. A user agent installed on smartphones communicates with the controller and associates the station with the recommended AP. WiFiSeer is vertical to our solution to further reduce station-AP delay.

Jang et al. [24] study the overhead of radio switching and show that stations can achieve significant energy saving during inter-frame delays while the AP is communicating with other stations. The proposed AP-driven approach, called Snooze, utilizes the global knowledge of the AP to schedule sleep and awake duration of each station based on inter-packet delays and traffic load of the station. To distribute the schedule, the AP needs to exchange control messages with stations. Compared to Snooze, our work considers the sensing-actuation pattern of IoT applications and reduces the idle listening time between sensing and actuation. In addition, our work takes into account the impact of interference by measuring channel utilization perceived by the AP. Also, Snooze does not benefit from APSD. Sheth and Dezfouli [25] propose *Wiotap*, an AP-based, layer-3 scheduling mechanism for IoT stations that employ APSM. This mechanism uses an Earliest Deadline First (EDF) scheduling strategy to maximize the chance of packet delivery before tail time expiry. Rozner et al. [14] propose NAPman, which prioritizes the delivery of PSM traffic as long as other stations are not affected. Tozlu et al. [5] show that increasing AP load has a higher impact on packet loss and RTT compared to out-of-band interference. Pei et




Fig. 1. The end-to-end delay components between a station and a server. The prediction of  $\delta_c$  is particularly challenging because it is affected by several factors such as traffic rate, channel utilization, and buffering mechanisms employed by Linux's *qdisc* and wireless NIC's driver.

al. [17] demonstrate that 802.11 link delay comprises more than 60% of station-server RTT. They also show that more than 50% (10%) of packets experience more than 20 ms (100 ms) latency by station-AP links. For TCP traffic, the authors proposed an approach to measure the delay of wired latency as well as uplink and downlink channel access delay. Using the Kendall correlation, they also show that channel utilization, RSSI, and retry rate are the top three factors affecting station-AP delay. Then, a decision tree model is used to tune the parameters of APs and reduce the overall delay experienced by stations. This is in contrast to our work, which offers per-station, fine-granularity sleep scheduling. Primarily designed for VoIP traffic, Liu et al. [26] propose a mechanism to reduce contention among stations waking up using APSD to retrieve packets from the AP. After receiving a burst of voice data, the station measures the tolerable deadline of incoming packets and informs the AP about its wake up time before switching to sleep mode. The wake-up instance is approved only if the AP will be idle when the station will wake up.

### III. DELAY ANALYSIS AND AP DEVELOPMENT

In this section, we first analyze the delay components between uplink and downlink packet delivery, and then present our proposed AP architecture and implementation. The proposed AP enables us to collect features that are necessary for delay prediction and schedule dissemination to IoT stations.

#### A. Delay Components

As Figure 1 shows, at time  $t_1$  the station grasps the channel and transmits its uplink packet. This uplink packet may represent a single uplink packet sent by the station or the last packet of a burst of uplink packets. After this, the station waits to receive downlink packet(s) from the AP. We refer to the process of uplink and downlink packet exchange as a *transaction*. In event-driven applications, the downlink packet is usually a command message issued by a server in response to the message sent by the station. As discussed in Section I, multiple power saving strategies can be used by the station to save power while waiting for the downlink packet. The goal of this paper is to inform the station about the delivery time of downlink packet. Therefore, we enable the station to switch to



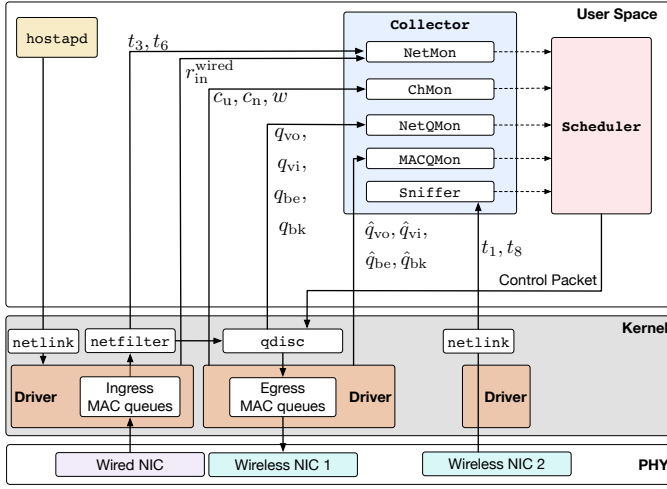


Fig. 2. The AP architecture developed and used in this paper. The Collector module communicates with various kernel and user-space components to collect a set of features required for delay prediction. The Scheduler estimates the sleep duration and conveys it to the station. This figure primarily focuses on the wired-to-wireless interfaces path to compute  $\delta_e$ . Some of the modules required to collect other delay components ( $\delta_a$  and  $\delta_b$ ) are not included in this figure.

sleep mode and wake up when the downlink delivery is about to happen.

To reduce the waiting time for downlink packet delivery, the station transitions into sleep mode after the reception of a control packet at  $t_2$  and wakes up at  $t_7$  to request and receive the downlink packet. The sleep duration is conveyed to the station by the AP through a control packet sent at  $t_2$ . Therefore, we need to estimate the delay between  $t_1$  to  $t_7$ . To this end, we first modify a Linux-based AP by adding new user-space and kernel-space modules.

### B. AP Development

Unfortunately, the current AP architectures do not provide the necessary tools to collect and apply predictive scheduling [27]. In this section, we present an AP architecture that allows us to collect the features necessary for predictive scheduling. Figure 2 presents the modules we developed on a Linux-based AP. The user-space components of the AP are Collector and Scheduler. The Collector is responsible for collecting all the features required to predict delay. This information is then shared with and used by the Scheduler to train a model, estimate the sleep duration, and dispatch the schedule. The information collected by the Collector is stored in the physical memory (using mmap) to reduce the delay of accessing this data through files. The Collector module includes the following modules: The Sniffer module utilizes the libpcap library to capture the timestamp of packets as soon as they are sent or received by the wireless NIC. The NetMon module records packet exchange instances over the wired interface as well as incoming data rate over this interface. The NetQMon and MACQMon are responsible for keeping track of the utilization of qdisc and MAC layer queues, respectively. The ChUMon module captures channel utilization.

To perform the standard AP functionalities, we use hostapd [28], which is a user-space daemon that handles beacon transmission, authentication, and association of stations. The underlying hardware includes an Atheros AR9462 wireless NIC, ath9k driver, a Core i3 processor, and 8 GB of RAM. The AP operates in 802.11n mode, uses two antennas, and supports up to 144 Mbps.

We explain the implementation detail and operation of the AP in the next three sections.

### C. Communication Delay Between AP and Server

Once the AP receives an uplink packet, it is stored in the qdisc of wired interface and then is sent over the wired interface. The qdisc is the scheduling mechanism employed by the kernel to schedule the transmission of packets while switching them between two interfaces. This buffering delay, denoted as  $\delta_a = t_3 - t_1$ , depends on the difference between the rate of incoming wireless uplink packets (destined to the wired interface) and the rate of transmitting these packets over the wired interface. The primary types of networks considered in this paper are smart home environments where an AP is connected to an Internet modem, and campus and business deployments where APs communicate via an Ethernet infrastructure. In such networks, the speed of the AP's wired interface is fixed and usually higher than the wireless interface. This are reasonable assumptions because: First, in enterprise environments, APs are connected to switches via Ethernet links supporting at least 1 Gbps. This may also be true in a residential environment where the AP is connected to a local processing server through Ethernet [29]. For residential environments, also, cable modems and fiber-to-the-home (FTTH) provide data rates higher than wireless. Second, the uplink speed between a home modem and an Internet provider is fixed. For example, DOCSIS employs a combination of TDMA and CDMA for deterministic channel access.

Based on these observations, we compute  $\delta_a$  by knowing the number of packets currently in the qdisc of wired interface (not shown in Figure 2). We have modified the qdisc module to communicate the number of packets in this buffer with NetMon. For each packet  $p_i$  in qdisc, the Scheduler computes  $\nu(p_i) = 8 \times (s(p_i) + h_{\text{mac}} + h_{\text{phy}}) / l_{\text{out}}^{\text{wired}}$ , where  $\nu(p_i)$  is the time required to transmit  $p_i$ ,  $s(p_i)$  is the packet size (bytes),  $h_{\text{mac}}$  is the MAC header size (bytes),  $h_{\text{phy}}$  is the physical header size (bytes), and  $l_{\text{out}}^{\text{wired}}$  is the transmission bit rate supported by the wired link. The switching delay is therefore computed as  $\delta_a = \sum_{p_i \in \text{qdisc}} \nu(p_i)$ .

The delivery delay between the AP and the server, i.e.,  $t_4 - t_3$  and  $t_6 - t_5$ , depend on various factors and primarily on the number of hops between these two nodes. Based on this number, we consider *edge* and *cloud* computing scenarios. Edge computing is employed in latency-sensitive and mission-critical applications to minimize the latency and overhead of communication over the wired network [30]. In the cloud computing scenario, the server is located at least a few hops away from the AP. For both cases, to measure this delay, we use a moving average, which is the standard approach used by various protocols such as TCP to estimate RTT [31], [32].



To this end, we modify the `netfilter` [33] kernel module to communicate with the `NetMon` module and timestamp  $t_3$  as the instance the packet is sent to the wired NIC, and  $t_6$  as the instance the packet has arrived in the AP.

#### D. AP to Station Delivery Delay

An incoming packet from the wired interface first passes through ingress driver queues. Subsequently, the packet is processed by the `netfilter` module. The packet is then queued in the qdisc. Finally, the packets are queued at the Enhanced Distributed Channel Access (EDCA) queues inside the wireless NIC's driver. These packets are served according to the channel contention parameters specified by the 802.11e standard. Each driver queue contends (individually) for channel access before packet transmission.

Here we mainly focus on the delay between the arrival of a downlink packet through the AP's wired interface and its transmission through the wireless interface. This delay is denoted as  $\delta_c = t_7 - t_6$ . It is particularly challenging to model and predict this delay because it is affected by several factors such as queuing strategy and queue utilization at the qdisc and MAC layer, channel utilization, number of stations, and link quality. However, in addition to the high complexity of buffering mechanisms implemented by wireless drivers such as `ath9k` and `ath10k` [19], [20], the actual operation of non-open source drivers is not known, which makes it impossible to develop a mathematical model of buffering delay. Therefore, we follow a data-driven approach to predict  $\delta_c$ . The predicted value is denoted as  $\delta'_c$ . The `Collector` module time stamps the switching delay between the wired and wireless interfaces. The time of packet arrival from the downlink transmission is determined by the `Sniffer` module, which in turn informs the `Collector`. During  $t_6$  to  $t_8$ , the `Collector` also collects statistics regarding the status of queues and channel condition. The collected parameters are explained in the following sections.

1) *Input traffic rate through wired interface:* The incoming traffic through wired interface, denoted as  $r_{in}^{wired}$  (bytes/second), impacts the current and future utilization of wireless interface's qdisc and driver queues. Hence, the `NetMon` module communicates with wired interface driver to collect incoming traffic rate.

2) *qdisc queues:* Every network interface is assigned a qdisc, which is `pfifo_fast` by default [18]. This mechanism contains three bands, and dequeuing from a band only happens when its upper bands are empty. The `PRIO` qdisc is a classful-configurable alternative of `pfifo_fast` and enables us to configure the number of bands. To enqueue the packets of each AC in its own queue, we implement four queues in this layer. These queues are denoted by  $\mathbf{Q} = \{q_{vo}, q_{vi}, q_{be}, q_{bk}\}$ . We have modified the `PRIO` kernel module to communicate with the `NetQMon` module to collect the number of packets in each qdisc band.

With `PRIO` qdisc, the queuing delay experienced by a packet enqueued in the lowest priority queue not only depends on the current utilization level of that queue, but also on the number of packets in the higher priority queues. In addition

to the four queues mentioned above, we have also included an additional queue—called *control queue*—that is assigned the highest priority level. We will utilize this queue in order to implement a higher priority data plane used to send the *control packet* that conveys sleep schedules to stations. We will explain this packet later. It is worth mentioning that, although our implementation utilizes the `PRIO` qdisc (the default policy used in several Linux distributions), the concept can easily be extended to other types of qdiscs, such as `Hierarchical Token Bucket` (HTB).

3) *Wireless channel condition:* Both interference and channel utilization are the main channel condition parameters that affect packet transmission delay. However, the duration and intensity of these parameters depend on various factors, such as the number of contending stations and APs, burst size, TXOP, and the transmission power of nearby stations and APs. Therefore, accounting for the effect of channel condition through measuring the factors (mentioned above) would be very challenging. Instead, we collect three parameters to capture the effect of interference and channel utilization on the delay of packet transmission. The first parameter is channel utilization ( $c_u$ ), which refers to the amount of time the AP or its associated stations are transmitting. The second parameter is the number of MAC layer retransmissions ( $w$ ) performed by the AP to deliver packets to stations. The third parameter is the channel's noise level ( $c_n$ ), which reflects the activity of nearby wireless devices (such as other APs and stations, ZigBee, and Bluetooth devices).

Most 802.11 drivers maintain counters that represent channel utilization rate. For example, the rate of updating `ch_time_busy` reflects channel utilization during a sampling interval. The `ChUMon` module is responsible to extract these counters from the driver. We realized that the interval of obtaining channel utilization impacts measurement accuracy. We obtained the peak accuracy, in terms of Kendall's correlation coefficient, when the frequency of polling  $c_u$  is 10 ms. Additionally, the granularity of the measurements also decreases as we increase the frequency of polling channel utilization. This is because the counters provided by the NIC are reported in milliseconds. For example, if the sampling frequency is 10 ms, the granularity of  $c_u$  obtained in percentage is 10%.

4) *Driver's transmission queues:* Using Enhanced Distributed Coordination Function (EDCF), packets arriving at the MAC layer are categorized and inserted into one of the four queues assigned to each station inside the driver. The categorization relies on the IP header's Type of Service (ToS) field. These queues are denoted by  $\hat{\mathbf{Q}} = \{\hat{q}_{vo}, \hat{q}_{vi}, \hat{q}_{be}, \hat{q}_{bk}\}$ . Each queue behaves like a virtual station that contends for channel access independently. In case of internal collision between two or more queues, the higher priority queue is granted the transmission opportunity. The status of these queues are monitored by the `MACQMon` module through communicating with the driver. Considering the size of these queues allows the prediction models to take into account the effect of packet aggregation and packet bursting. Specifically, for each AC, packet aggregation is applied if a station's queue includes more than one packet. Also, depending on the AC, a burst of packets may be sent without having to contend for the channel on a

per-packet basis.

5) *Summary of the features collected:* The Scheduler interacts with Collector to gather the features necessary for delay prediction. In summary, the developed AP enables us to collect the following features periodically:

$$\mathbf{C}_u, \mathbf{C}_n, \mathbf{R}_{in}^{wired}, \mathbf{W}, \mathbf{Q}_{vo}, \mathbf{Q}_{vi}, \mathbf{Q}_{be}, \mathbf{Q}_{bk}, \hat{\mathbf{Q}}_{vo}, \hat{\mathbf{Q}}_{vi}, \hat{\mathbf{Q}}_{be}, \hat{\mathbf{Q}}_{bk} \quad (1)$$

where  $\mathbf{C}_u$ ,  $\mathbf{C}_n$ ,  $\mathbf{R}$ ,  $\mathbf{Q}$ ,  $\hat{\mathbf{Q}}$ , and  $\mathbf{W}$  represent the list of channel utilization values, list of channel noise values, list of incoming traffic rate values over wired interface, list of MAC layer downlink retransmission values, list of the utilization values of qdisc queues (for each AC), and list of the utilization values of driver queues (for each AC), respectively. Each of these lists includes the values that have been periodically collected. For example, assuming that each list contains  $k$  values, list  $\mathbf{C}_u$  is represented as follows:

$$\mathbf{C}_u = [c_u(t' - (k - 1) \times \Delta), c_u(t' - (k - 2) \times \Delta), \dots, c_u(t' - \Delta), c_u(t')] \quad (2)$$

where  $c_u(t')$  is the last sampled value based on the periodical sampling technique, and  $\Delta$  refers to sampling interval. In our implementation, we collect samples every 10 ms ( $\Delta = 10$  ms), and each prediction uses the past 5 sampled values ( $k = 6$ ). We did not use a shorter sampling interval because of the significant increase in processor utilization ( $> 30\%$ ). Implementing a more efficient AP architecture is left as a future work.

In addition to the features collected periodically, we add two features that are computed once for each prediction. First, since each AC is treated differently at the driver, we include the AC of the transaction as a feature. Second, we use  $\delta_a + \delta_b$  because the predicted delay ( $\delta'_c$ ) depends on the interval between the uplink packet and the arrival of downlink packet over the wired interface. For example, if the server delay is expected to be 30 ms, the prediction for  $\delta_c$  should be made for a packet that would arrive at the AP in 30 ms.

#### E. Schedule Announcement

Once a predicted delay value is computed, the Scheduler creates a UDP control packet to send the value to the station. This packet includes  $\delta_a$ ,  $\delta_b$ , and  $\delta_c$ , as well as the standard deviation of prediction error, where each value is encoded as one byte. The value of each byte reflects duration in milliseconds. This data packet is sent using a dedicated queue with highest priority. Once this control packet reaches the station at  $t_2$ , the station immediately transitions into sleep state for  $\delta_a + \delta_b + \delta_c - (t_2 - t_1)$ . Note that the AP shares the relative wake-up time with the station. Since the AP has computed wake up schedule at time  $t_1$ , the station needs to measure  $t_2 - t_1$  and subtract it from the shared value. The station can simply use a timer to measure  $t_2 - t_1$ .

At the end of the sleep interval, the station wakes up and informs the AP about its transition into the awake mode. This is achieved by relying on APSD, which is supported by the state-of-the-art wireless NICs. To this end, at  $t_7$  (in Figure




Fig. 3. Traffic characterisation.

1), the station wakes up and sends a NULL packet to the AP, conveying that the station is ready for receiving a packet. The AP responds by sending one or multiple downlink packets starting at  $t_8$ . As per the 802.11e amendment, stations can start a Transmit Opportunity (TXOP) for downlink delivery. For example, if the traffic belongs to the voice AC, the AP uses a 1.504 ms slot (as long as packets exist) for downlink delivery.

#### IV. PREDICTIVE SCHEDULING

In this section, we first present our testbed setup for realistic traffic generation necessary for training and evaluation. Utilizing this dataset, we discuss various stages of the statistical learning and modeling process. Finally, we compare the performance of several machine learning approaches for delay prediction.

##### A. Traffic Generation

As explained in the previous sections, AP modification is necessary to collect the features required for predictive scheduling. Also, we need to introduce controlled changes in the traffic pattern of the environment to study the impact of these changes on prediction accuracy. Therefore, it is required to have a testbed that represents the traffic pattern of real-world environments as well as controllability over traffic generation parameters. To achieve this, we systematically characterize and compare the scenarios generated in our testbed with those collected in real-world environments.

A *burst*, denoted as  $b_i$ , is defined as a train of packets in either UL or DL direction with inter-arrival time less than a threshold value  $\theta$  [34]. Resembling 802.11 traffic, Figure 3 illustrates a series of bursts. The duration (in seconds) of a burst  $b_i$  is denoted by  $d(b_i)$ . And  $g(b_i)$  refers to the gap (in seconds) between two consecutive bursts  $b_i$  and  $b_{i+1}$ .

In order to generate traffic that represents various levels of *network dynamics* in real-world environments, we have developed a testbed that includes two categories of stations: (i) stations such as laptops, phones, and IoT devices, and (ii) four Raspberry Pi boards to control traffic generation pattern. Each RPi runs four threads, where each thread can be involved in a downlink, uplink, or bidirectional flow. This enables us to introduce up to 16 additional controlled flows into the network. The implementation of traffic control capability is composed of a set of Python scripts that use the iperf tool under the hood. A central controller is in charge of setting the parameters of traffic flows. Among the flow parameters, we can modify the access category, transport layer protocol, packet size, bit rate, burst size, and inter-burst interval. Also, we note that sharing flow characteristics by consecutive flows is more likely. To represent this behavior, after each burst, the controller either

repeats the process of traffic selection or chooses the same parameters for the next burst based on a *variability parameter* denoted by  $\nu$ . Specifically, a higher value of  $\nu$  results in a higher dynamicity. Hence, we use  $\nu = 0.9$  for generating flows that resembles *high dynamicity* (HD) traffic and  $\nu = 0.1$  for generating less diverse traffic referred to as *normal dynamicity* (ND). Also, for the voice and video ACs, UDP is preferred because it is the dominant transport protocol for real-time traffic.

As demonstrated in [34], capturing network dynamics can be achieved by focusing on characterizing burstiness. Additionally, Xiao et al. [35] characterized a flow as *regularly* bursty when the standard deviation of the inter-burst intervals ( $g(\cdot)$  seconds) and burst sizes ( $s(\cdot)$  bytes) are relatively smaller. Otherwise, the flow is characterized as *randomly* bursty. However, based on Xiao's metrics for burstiness, traffic with fewer bursts per unit time can still have a high standard deviation of  $s(\cdot)$  and  $g(\cdot)$ . Hence, we consider burst frequency (i.e., number of bursts per second) and burst size for calculating traffic burstiness. Additionally, due to the difference in the scale of those two parameters, we normalize the burst rate (per second) in the range  $[0, 1]$ . We define traffic *burstiness*, denoted by  $\mathcal{B}$ , as follows:

$$\mathcal{B} = \left(1 - \frac{1}{\mathcal{M}}\right) \times \left(\frac{\sum_{i=1}^N s(b_i)}{N}\right) \quad (3)$$

where  $N$  is the number of bursts in the dataset,  $s(b_i)$  is the size of burst  $b_i$  (in bytes), and  $\mathcal{M}$  is average number of bursts per second.

In addition to traffic burstiness, we define another metric that represents traffic *dynamicity* based on various aspects including burst size, burst duration, inter-burst interval, and the access category of the packets in each burst. This metric, which we refer to as *dynamicity* and denoted by  $\mathcal{D}$ , is defined as follows:

$$\mathcal{D} = \frac{1}{N} \sum_{i=2}^N \frac{|d(b_i) - d(b_{i-1})|}{g(b_{i-1})} + \frac{1}{N} \sum_{i=2}^N \frac{|s(b_i) - s(b_{i-1})|}{g(b_{i-1})} + \frac{1}{N} \sum_{i=2}^N \frac{|p(b_i) - p(b_{i-1})|}{g(b_{i-1})} + \frac{1}{N} \sum_{i=2}^N \frac{z(b_i)}{g(b_{i-1})} \quad (4)$$

where,

$$z(b_i) = \frac{|p_{vo}(b_i) - p_{vo}(b_{i-1})|}{p_{vo}(b_{i-1})} + \frac{|p_{vi}(b_i) - p_{vi}(b_{i-1})|}{p_{vi}(b_{i-1})} + \frac{|p_{be}(b_i) - p_{be}(b_{i-1})|}{p_{be}(b_{i-1})} + \frac{|p_{bk}(b_i) - p_{bk}(b_{i-1})|}{p_{bk}(b_{i-1})} \quad (5)$$

Here,  $p_x(b_i)$  is number of packets belonging to an AC  $x$  in a burst  $b_i$ . Parameter  $z(b_i)$  reflects the change in the number of packets belonging to each access category in each burst compared to that in the previous burst.

## B. Data Collection

Using the metrics mentioned above, we compare our testbed generated datasets against those collected in multiple real-world environments. Figure 4 presents the results. In general,

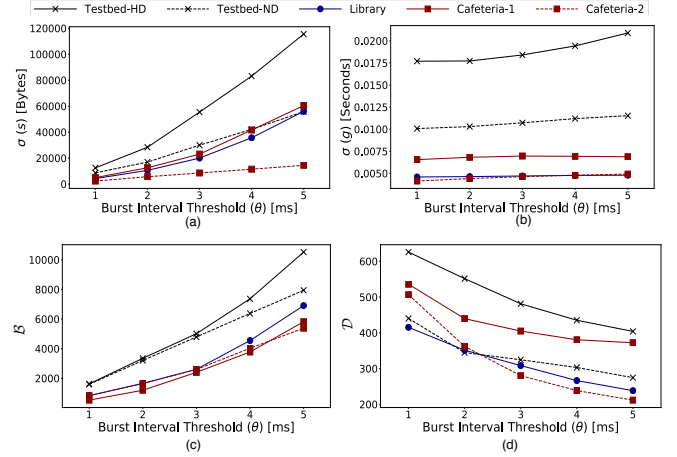


Fig. 4. (a) Standard deviation of burst size, (b) standard deviation of burst interval, (c) burstiness ( $\mathcal{B}$ ), and (d) dynamicity ( $\mathcal{D}$ ) of traffic generated by our testbed compared to traffic captured in real-world environments. ND and HD refer to normal and high dynamicity, respectively.

we can observe that our ND scenario resembles real traffic. The HD scenario offers higher network dynamics, which is essential to study the robustness of predictive scheduling.

When generating data in our testbed, the type of each transaction is selected from the voice, video, background, and best-effort ACs with equal probability. The inter-transaction delays are uniformly distributed between 1 ms and 500 ms. In addition to the features discussed in Section III, we also collect  $\delta_a$ ,  $\delta_b$ , and  $\delta_c$  values per transaction. We split each dataset, such that 70% of it is used for training and the remaining 30% is used for validation. We use independent datasets consisting of 10,000 data points for evaluating the performance and robustness of each modelling approach (test datasets) in the ND and HD scenarios.

## C. Data Pre-processing

We focus on delay prediction for  $\delta_c < 100$  ms, for two reasons: First, considering edge computing scenarios, observing RTTs more than 100 ms is very unlikely. Second, almost all commercial APs implement 102.4 ms as their beaconing period. Therefore, all stations wake up every 102.4 ms to synchronize with AP beacons and check if the AP has any buffered packets.

Although we have ascertained that there is no unnecessary user-space applications running on the AP, running kernel tasks, as well as context switching of existing processes, can lead to a higher operating system response time. Hence, we have noticed that the latency of data collection from the software modules discussed in Section III is impacted occasionally and results in longer polling intervals. We filter the datasets to remove such *stale* feature values that do not reflect the latest status of the channel or the queues.

The feature set varies in terms of ranges and units. For example,  $c_u$  varies from 0 to 100%, whereas  $c_n$  varies from  $-95$  dBm to  $-66$  dBm. Since this would result in disproportional treatment of the features by the machine-learning algorithms, we scale each of the features into the range  $\pm 1$ . Furthermore,



the dataset contained more samples (transactions) whose actual delay is within the range  $[1, 50]$  ms, compared to the samples whose actual delay was between  $[50, 100]$  ms. We under-sample the majority bins to prevent the machine-learning algorithm from generalizing the results for the packets whose actual delay is higher.

#### D. Regression models

Given the continuous nature of the target variable, we identify the problem of predicting  $\delta_c$  as a regression problem. The methods that we considered are Random Forest Regressor (RFR), Gradient Boosting Regressor (GBR), Extra Trees Regressor (ETR) and Histogram-Based Gradient Boosting Regressor (HBR), which are widely-used ensemble learning methods for regression. We also considered (deep) neural networks, which have been shown to be effective in different areas of machine learning such as prediction in time-series data and image processing.

In ensemble learning, final prediction can either be calculated by average of the predictions of the model trained on random subsets of data (bagging) or calculated via sequentially training the model using prediction success on the previous sample of the dataset (boosting). RFR is an example of bagging approach and functions by constructing several decision trees during training and makes predictions based on the outputs of the individual trees. RFR is a popular algorithm and runs efficiently on large and high-dimensional datasets.

GBR is an example of boosting approach. Each tree outputs a prediction value at different splits that can be added together, allowing subsequent models to modify error in predictions. HBR is a variant of GBR. Since it is a histogram-based estimator, HBR can reduce the number of splitting points by binning input samples, and therefore improves the performance when dealing with large datasets. ETR creates decision stumps at variable tree depths. The features and splits are selected randomly, and are less computationally expensive than other tree-based algorithms.

Neural Networks (NN) have been studied extensively in the past decade for their efficiency in learning complex data features for making predictions. Multilayer perceptrons (MLP) is one such variant of feed-forward neural networks that does not allow cycles, thereby resulting in the data progressing in a single direction over the network. Considering the one-direction data movement, one of the biggest drawbacks of using such a network is its lack of memory, i.e., it treats each instance of the input time-series independently. Recurrent Neural Networks (RNN) are a class of neural networks in which the predictions are based on the current and also past inputs, and therefore they are suitable for making predictions about  $\delta_c$  using historical network features. A specific variant of RNN is Long Short-Term Memory (LSTM) [36], which is able to track long-term dependencies of output predictions on input history. We compare the performance of traditional machine learning approaches, such as RFR, ETR, GBR, and HBR with neural networks, such as MLP and LSTM for predicting  $\delta_c$ .

For the ensemble learning methods, we use scikit-learn library [37] and tune the hyper parameters using grid search

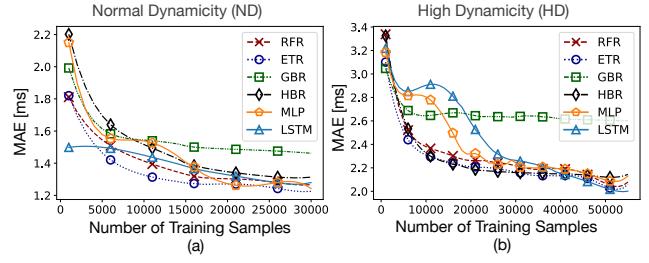


Fig. 5. MAE of machine learning algorithms versus the number of samples (transactions) in training dataset for (a) Normal Dynamicity (ND), and (b) High Dynamicity (HD) scenarios. Results are averaged over all ACs. ETR converges the fastest, and LSTM requires up to 3x more data points compared to ETR.

and a validation dataset to obtain the highest performance on the training data avoiding over-fitting. For the MLP and LSTM model, we use Tensorflow and Keras library [38]. Also, we utilize early stopping mechanism (on the validation dataset) to prevent over-fitting. The optimal MLP model contains five dense layers, each consisting of 32, 20, 16, 10, 8 neurons, respectively, and *ReLu* activation function. The optimal LSTM model contains one LSTM layer followed by three dense layers, each consisting of 20 neurons and *ReLu* activation function. Both the MLP and LSTM models contains an output layer and were trained using Adam optimiser [39] with learning rate of 0.01.

#### E. Model Evaluation

For all evaluations we used the test dataset. We first investigate the effect of training data size on model's performance. Figure 5 illustrates the Mean Absolute Error (MAE) of  $\delta_c$  as a function of the size of training data. For better visibility, we used Savitzky–Golay filter to reduce the noise and added markers at regular intervals in Figure 5. We observed that the performance of ETR converges at the fastest rate, utilizing 15000 and 20000 data points for training under the ND and HD traffic, respectively. Whereas, due to the higher complexity of neural networks, MLP requires 20000 data points in the ND scenario and 35000 data points in the HD scenarios. LSTM requires 30000 data points in the ND and 50000 data points in the HD scenario—showing slower convergence compared to MLP. Based on these results, for the rest of the evaluations presented in this paper, we use the number of data points that are required for each algorithm to demonstrate performance convergence.

Next, we quantify the effect of feature-history (represented by  $k$  in Eq. 2) on the overall performance of all algorithms. Figure 6 shows the results. MAE decreases significantly in both HD and ND scenarios for all the algorithms when we include the recent feature history as an input for training the model. This decrease continues up to 30 ms, beyond which it does not result in performance enhancement. Historical data of features helps the model to anticipate trends of features and accurately predict  $\delta_c$  that would be incurred by the downlink packet shortly. Therefore, in addition to the most recent record, we include the three preceding feature values (corresponding






Fig. 6. MAE of machine learning algorithms with respect to feature history in (a) Normal Dynamicity (ND), and (b) High Dynamicity (HD) scenarios. Results are averaged over all ACs.




Fig. 7. MAE of machine learning algorithms versus transaction's AC in (a) Normal Dynamicity (ND), and (b) High Dynamicity (HD) scenarios. MAE of VO and VI packets is lower than BK and BE packets.

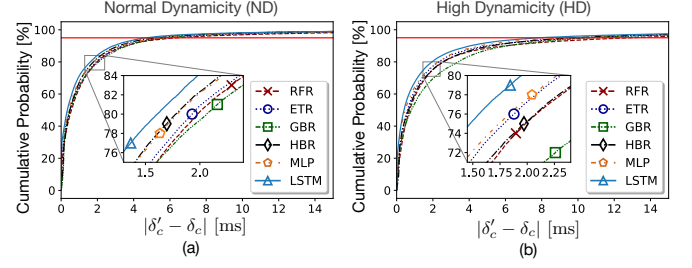


Fig. 8. ECDF of prediction errors ( $|\delta'_c - \delta_c|$ ) while utilizing various machine learning algorithms in (a) Normal Dynamicity (ND), and (b) High Dynamicity (HD) scenarios. All machine-learning algorithms are able to predict  $\delta'_c$  for 95% of the packets with an error of  $\pm 4.3$  ms in case of ND scenario, and  $\pm 8.2$  ms for the HD scenario. We have used markers in the inset graph for better visibility.




Fig. 9. MAE with respect to LSTM history in ND and HD scenarios.




Fig. 10. Processing time required to predict delay for each transaction.

to a total of the preceding 40 ms of feature values) to train the models.

Figure 7 compares the performances of the machine-learning algorithms versus the AC of transactions. Averaged over all ACs, the MAE (in millisecond) of algorithms in the ND scenario are: RFR: 1.43, ETR: 1.26, GBR: 1.49, HBR: 1.28, MLP: 1.24, and LSTM: 1.16. For the HD scenarios the MAE values are: RFR: 2.49, ETR: 2.17, GBR: 2.69, HBR: 2.27, MLP: 2.12, and LSTM: 2.01. On average for the ND and HD scenarios, the MAE of LSTM is 14% lower for all ACs, compared to the average MAE of all the other machine learning algorithms. Based on these results, in Section V we use LSTM for the empirical performance evaluation of EAPS on a testbed.

Figure 7 also shows that the MAE of VO and VI packets is lower than BK and BE packets. The reason is that the packets of these ACs are prioritized over higher ACs at the qdisc layer (using PRIO qdisc) as well as the driver's queues (using EDCA). This results in lower delays incurred by the downlink packets and lower unpredictability caused by the transmission of packets in higher priority queues.

Figure 8 presents the Empirical Cumulative Distribution Function (ECDF) of the deviation of  $\delta'_c$  from  $\delta_c$ . The 95th percentile of error ( $\delta'_c - \delta_c$ ) for all machine learning algorithms is less than  $\pm 4.3$  ms in case of ND scenario and  $\pm 8.2$  ms for the HD scenario. However, note that LSTM requires about 3x more training data for its performance to converge, compared to other algorithms (cf. Figure 5). Hence, in scenarios where it is not possible to collect large datasets for training, either ETR or MLP can also be used.

In LSTM networks, neuron activations are dependent on the previous network states. The network retains a memory

equal to the number of lookbacks (a.k.a., timesteps)—allowing the flow of information from all the previous timesteps [40]. Lookback is defined as the number of timesteps for which the LSTM is unfolded for back-propagation. Simply put, the number of timesteps reflects the number of previous transactions in the temporal domain that aids in predicting the delay of the current transactions by providing contextual information. This is particularly beneficial when multiple transactions occur during a single burst or bursts with similar characteristics. To this end, we estimate the effect of historical data of the previous transactions by including feature values of the past transactions and the current transaction. Figure 9 shows the MAE of LSTM versus the number of lookbacks. We observe that the MAE of the LSTM model decreases up to five lookbacks. This means, on average, five transactions occur during similar network conditions.

Figure 10 shows the prediction processing overhead of all the machine learning algorithms on a 2.4 GHz Core-i3 processor. Each marker shows the median of time taken to predict each data point in the test dataset, and the error bars present lower and upper quartiles. We observed that the HBR algorithm takes the least amount of time (24  $\mu$ s median and 0.046  $\mu$ s standard deviation) for prediction, whereas LSTM requires the most (48  $\mu$ s median and 3  $\mu$ s standard deviation). However, the time taken to predict the delay in case of LSTM is still considerably shorter than a packet transmission time. For example, with a 1400 bytes packet sent over a 54 Mbps link, the ratio of prediction duration to transmission duration is  $48\mu\text{s}/207\mu\text{s}$ .

## V. EMPIRICAL EVALUATION

In this section, we present an empirical evaluation of EAPS versus the power saving mechanisms of 802.11.

### A. Testbed

Our testbed includes four IoT stations, four Raspberry Pi boards, regular stations (smartphones and laptops), an AP, and a server. We simply refer to the IoT stations as *station*. Each station is a Cypress CYW43907 [3], which is a low-power 802.11n SoC designed for IoT applications. To represent a real-world scenario affected by variable background interference, the testbed is located in a residential environment surrounded by several APs belonging to different households. Also, the four Raspberry Pi boards are used to control network dynamicity and variability in  $\delta_c$ .

To represent the request-response behavior of IoT traffic, for each uplink UDP packet sent, the server responds by sending a downlink packet back to the station<sup>3</sup>. The exchange of an uplink packet and receiving its response is referred to as a *transaction*. In all the figures of this section, each marker shows the median of 1000 transactions and the error bars present lower and upper quartiles. We use the EMPIOT tool [41] to measure the energy and delay of each transaction. This tool samples voltage and current at approximately 500,000 samples per second. These samples are then averaged and streamed at 1 Ksps. The current and voltage resolution of this platform are 100  $\mu$ A and 4 mV, respectively.

We use two scenarios to evaluate the performance of EAPS with respect to varying AP-server delays (i.e.,  $\delta_b$  in Figure 1): *edge*, and *cloud* computing. In the former, the server is directly connected to the AP, and in the latter, we use an Amazon AWS server in Oregon. Note that in both cases the sleep schedules are computed at the edge and by the AP the station is associated with.

### B. Baselines and EAPS Variations

The baseline mechanisms used are PSM, APSM, and CAM. Using PSM, after an uplink packet, the station goes back into sleep mode and wakes up at each beacon instance to check for downlink packet delivery. With APSM, instead of going back into sleep right after uplink, the station stays in the awake mode for 10 ms. With CAM, the station always stays in awake mode. *Note that for CAM, to avoid the impact of inter-transaction time on results, we only measure the delay and energy consumption of transactions (only between the uplink and downlink packets).*

We consider three variants of EAPS based on prediction error. To justify these variations, we first present the distribution of prediction errors in Figure 11 for voice and background ACs. Based on distribution for each AC, the station can either choose to wake up at (i)  $\delta' - 2\sigma$ , (ii)  $\delta' + 2\sigma$ , or (iii)  $\delta'$ . We call these cases EAPS with Early wake-up (EAPS-E), EAPS with Late wake-up (EAPS-L), and EAPS with Mid wake-up (EAPS-M). Intuitively, EAPS-E reduces delay with




Fig. 11. Distribution of prediction error ( $\delta'_c - \delta_c$ ) for (a) voice, and (b) background ACs. Prediction error of voice AC is lower than that of background AC. Depending on the application's energy-delay trade-off, the station may wake up before, on, or after the predicted time.

a higher energy consumption, EAPS-L reduces energy with a longer delay, and EAPS-M establishes a trade-off between energy and delay. Please note that EAPS-E is only applicable if  $\delta' - 2\sigma > 0$ .

### C. Results

Figures 12 and 13 illustrate the average energy consumption and duration of transactions when the station is communicating with edge and cloud computing platforms under ND and HD conditions, respectively.

In the cloud computing scenario, CAM and EAPS incur an average round trip delay of 35 ms and 42 ms, respectively, while EAPS consumes 63% less energy. This is because EAPS conserves energy expenditure by switching to sleep mode and waking up right before the packet is ready for transmission at AP. In contrast, CAM needs to stay in awake mode until the response is received. Reduction in energy consumption of EAPS compared to CAM reduces to 30% in edge environment due to the shorter duration spent in awake mode to receive the downlink packet.

With PSM, the station immediately transitions to sleep mode after transmitting each uplink packet. While this results in less energy consumption compared to CAM, transactions suffer about 55 ms higher delay on average because the earliest opportunity for downlink packet delivery is after the next beacon instance. The transaction duration of EAPS is 62% lower compared to PSM on average across all the ACs. With APSM, the station remains in idle (wait) state for 10 ms after transmission or reception. This is beneficial only in specific scenarios. For example, in the edge scenario, the station receives its downlink packet within the tail time (similar to CAM). However, when the round trip delay is more than 10 ms, the station will have to wake up again to retrieve the downlink packet after the next beacon announcement, thereby resulting in both long delay and waste of more energy compared to PSM. On average, for both edge and cloud scenarios, the energy consumption of APSM is 30% higher compared to PSM. In contrast, the energy consumption of EAPS is 20% and 43% lower than PSM and APSM, respectively. Also, the transaction duration of PSM, APSM, and EAPS are 77 ms, 10 ms, and 12 ms in edge computing scenario, and 77 ms, 72 ms, and 42 ms in cloud computing scenario.

With EAPS, the station has the freedom to choose between EAPS-E, EAPS-M, or EAPS-L, according to application re-

<sup>3</sup>Note that the case where multiple uplink and downlink packets are exchanged is simply supported as explained in Section III.




Fig. 12. Performance comparison of EAPS with 802.11 power saving mechanisms in *ND* conditions for all ACs. (a) and (b) show the average *per-transaction* energy and duration in cloud scenario, respectively. (c) and (d) show the average *per-transaction* energy and duration in edge scenario, respectively.




Fig. 13. Performance comparison of EAPS with 802.11 power saving mechanisms in *HD* conditions for all ACs. (a) and (b) show the average *per-transaction* energy and duration in cloud scenario, respectively. (c) and (d) show the average *per-transaction* energy and duration in edge scenario, respectively.

quirements. We observed that with EAPS-E, the station suffers from slightly higher energy consumption because it wakes up early, waits for the packet to be received from the AP, and then transitions into sleep mode. In the case of EAPS-L, since the station wakes up  $2 \times \sigma$  after the predicted delay, the downlink packet will be buffered at the AP. However, the station can immediately transition to sleep mode once the packet is received. Thus, energy consumption of EAPS-L is 14% lesser compared to EAPS-E, whereas, the transaction duration of EAPS-L is 18% higher than EAPS-E. EAPS-M balances the trade off between energy consumption and transaction duration.

## VI. DISCUSSION

In this section we discuss the relevance and potential impact of the proposed solution with regards to other wireless communication standards.

**Wake-up radio.** WUR mechanisms such as 802.11ba [1], [2] can be used to enhance EAPS in several ways. For example, as soon as a station finishes sending its uplink packet(s), the primary radio can switch into sleep mode, and the station will receive the schedule message via its low-power WUR. The primary radio will then wake up at the scheduled time to receive the downlink packet. To further reduce the idle energy caused by prediction inaccuracy, the WUR can be scheduled to wake up at  $\delta'_c - 2\sigma$  and wait for a command to wake up the primary radio. Another example, once the downlink packet arrives on the wired interface of the AP, the AP uses EAPS to compute the packet delivery delay ( $\delta_c$ ). Assuming that the wake-up delay of the primary radio is  $\beta$  [1], the AP sends the wake-up packet at  $\delta'_c - \beta$  to make sure

the station's primary radio will be awake on time for downlink delivery.

**Mesh networks.** As discussed in Section III-C, the primary types of networks used in this paper are smart home environments where an AP is connected to an Internet modem, and campus and business deployments where APs communicate via an Ethernet infrastructure. EAPS can also be used in mesh deployments. In this case, the backbone communication between APs (mesh nodes) introduces a wireless-to-wireless switching delay. This delay primarily depends on the bandwidth difference between the backbone (AP-AP) and access (AP-station) links. For example, assume a 160 MHz channel (in the 5 GHz band) is used to form the backbone, while each AP operates on a 20 MHz or 40 MHz channel (in the 2.4 GHz band). This configuration is prevalent because most of the existing IoT stations operate in the 2.4 GHz band, and WiFi mesh systems are usually tri-band and dedicate a channel (in the 5 GHz band) to their backbone. With this configuration, the delay caused by the backbone would be negligible and a method similar to that mentioned in Section III-C can be used to measure the delay from each AP to the server. If backbone links suffer from congestion and significant interference, EAPS can be used to predict packet switching delay over the backbone. As an alternative, more efficient strategy, EAPS could run on a central machine and allow the stations to receive their downlink packet from the AP offering the lowest delay. We leave these enhancements as future works.

**Computation offloading.** In our current implementation, we perform both training and modeling on the AP. However, if the AP is not powerful enough to train the model, then



the training could be offloaded to a cloud or fog computing platform. In any case, edge computing is essential to perform scheduling immediately and convey the sleep schedule to the station.

## VII. CONCLUSION

In this paper, we presented the design and implementation of a predictive scheduling mechanism, named EAPS, to enable IoT stations to transition to sleep mode and wake up when the downlink packet is expected to be delivered. The proposed solution benefits from edge computing, meaning the computations of sleep scheduling are performed at the network edge and by the AP. We presented an AP architecture capable of collecting queues status, channel condition, and packet transmission and reception instances. Once the AP receives an uplink packet, a machine learning model is used to compute the sleep delay, and the station is informed about its schedule using a high-priority data plane. Using empirical evaluations, we confirm the significant enhancement of EAPS in terms of energy efficiency and transaction delay.

EAPS can be considered as an enhancement of APSD and TWT (introduced in 802.11ax). For example, with the next generation of IoT stations that support TWT, they can set up their wake up time based on the sleep schedule computed by the AP. By protecting IoT stations against the effect of concurrent traffic and interference, EAPS is a particularly useful method in scenarios where both regular and IoT stations exist. For example, EAPS can reduce the energy cost of households and reduce the impact of IoT on global ICT footprint.

## REFERENCES

- [1] D.-J. Deng, M. Gan, Y.-C. Guo, J. Yu, Y.-P. Lin, S.-Y. Lien, and K.-C. Chen, "Ieee 802.11 ba: Low-power wake-up radio for green iot," *IEEE Communications Magazine*, vol. 57, no. 7, pp. 106–112, 2019.
- [2] D. Bankov, E. Khorov, A. Lyakhov, and E. Stepanova, "Ieee 802.11 ba—extremely low power wi-fi for massive internet of things—challenges, open issues, performance evaluation," in *Proceedings of BlackSeaCom*. IEEE, 2019, pp. 1–5.
- [3] Cypress Semiconductor. CYW43907: IEEE 802.11 a/b/g/n SoC with an Embedded Applications Processor. [Online]. Available: <http://www.cypress.com/file/298236/download>
- [4] AVNET Inc. BCM4343W: 802.11b/g/n WLAN, Bluetooth and BLE SoC Module. [Online]. Available: [https://products.avnet.com/opasdata/d120001/medias/docus/138/AES-BCM4343W-M1-G\\_data\\_sheet\\_v2\\_3.pdf](https://products.avnet.com/opasdata/d120001/medias/docus/138/AES-BCM4343W-M1-G_data_sheet_v2_3.pdf)
- [5] S. Tozlu, M. Senel, W. Mao, and A. Keshavarzian, "Wi-Fi enabled sensors for internet of things: A practical approach," *IEEE Communications Magazine*, vol. 50, no. 6, pp. 134–143, 2012.
- [6] Ring. [Online]. Available: <https://ring.com>
- [7] Nest. [Online]. Available: <https://nest.com>
- [8] Schlage. [Online]. Available: <https://www.schlage.com>
- [9] Lifx. [Online]. Available: <https://www.lifx.com>
- [10] Y. He, R. Yuan, X. Ma, and J. Li, "Analysis of the impact of background traffic on the performance of 802.11 power saving mechanism," *IEEE Communications Letters*, vol. 13, no. 3, pp. 164–166, 2009.
- [11] J. Manweiler and R. Roy Choudhury, "Avoiding the rush hours: Wi-Fi energy management via traffic isolation," in *Proceedings of MobiSys*, 2011, pp. 253–266.
- [12] S. Y. Jang, B. Shin, and D. Lee, "An adaptive tail time adjustment scheme based on inter-packet arrival time for IEEE 802.11 WLAN," in *Proceedings of ICC*. IEEE, 2016, pp. 1–6.
- [13] A. Vinas, V. Bernardo, M. Pascoal Curado, and T. Braun, "Performance analysis and comparison between legacy-PSM and U-APSD," *Proceedings of CRC*, pp. 1–12, 2013.
- [14] E. Rozner, V. Navda, R. Ramjee, and S. Rayanchu, "NAPman: Network-assisted power management for WiFi devices," in *Proceedings of MobiSys*. ACM, 2010, pp. 91–106.
- [15] A. J. Pyles, X. Qi, G. Zhou, M. Keally, and X. Liu, "SAPSM: Smart adaptive 802.11 PSM for smartphones," in *Proceedings of UbiComp*. ACM, 2012, pp. 11–20.
- [16] S. Sundaresan, N. Magharei, N. Feamster, R. Teixeira, and S. Crawford, "Web performance bottlenecks in broadband access networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, no. 1, pp. 383–384, 2013.
- [17] C. Pei, Y. Zhao, G. Chen, R. Tang, Y. Meng, M. Ma, K. Ling, and D. Pei, "WiFi can be the weakest link of round trip network latency in the wild," in *Proceedings of INFOCOM*. IEEE, 2016, pp. 1–9.
- [18] T. Høiland-Jørgensen, P. Hurtig, and A. Brunstrom, "The good, the bad and the WiFi: Modern AQMs in a residential setting," *Computer Networks*, vol. 89, pp. 90–106, 2015.
- [19] T. Høiland-Jørgensen, M. Kazior, D. Täht, P. Hurtig, and A. Brunstrom, "Ending the anomaly: Achieving low latency and airtime fairness in WiFi," in *Proceedings of USENIX*, 2017, pp. 139–151.
- [20] A. Showail, K. Jamshaid, and B. Shihada, "Buffer sizing in wireless networks: challenges, solutions, and opportunities," *IEEE Communications Magazine*, vol. 54, no. 4, pp. 130–137, 2016.
- [21] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proceedings of INFOCOM*, vol. 2. IEEE, 2003, pp. 836–843.
- [22] B. Peck and D. Qiao, "A practical PSM scheme for varying server delay," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 1, pp. 303–314, 2015.
- [23] K. Sui, M. Zhou, D. Liu, M. Ma, D. Pei, Y. Zhao, Z. Li, and T. Moscibroda, "Characterizing and improving WiFi latency in large-scale operational networks," in *Proceedings of MobiSys*. ACM, 2016, pp. 347–360.
- [24] K.-Y. Jang, S. Hao, A. Sheth, and R. Govindan, "Snooze: Energy management in 802.11n WLANs," in *Proceedings of Co-NEXT*. ACM, 2011, p. 12.
- [25] J. Sheth and B. Dezfouli, "Enhancing the energy-efficiency and timeliness of IoT communication in WiFi networks," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 9085–9097, 2019.
- [26] L. Liu, X. Cao, Y. Cheng, and Z. Niu, "Energy-efficient sleep scheduling for delay-constrained applications over WLANs," *IEEE Transactions on Vehicular Technology*, vol. 63, no. 5, pp. 2048–2058, 2014.
- [27] F. Wilhelmi, S. Barrachina-Munoz, B. Bellalta, C. Cano, A. Jonsson, and V. Ram, "A flexible machine-learning-aware architecture for future WLANs," *IEEE Communications Magazine*, vol. 58, no. 3, pp. 25–31, 2020.
- [28] J. Malinen. hostapd: IEEE 802.11 AP, IEEE 802.1X/WPA/WPA2/EAP/RADIUS Authenticator. [Online]. Available: <https://w1.fi/hostapd/>
- [29] Y.-H. Wei, Q. Leng, S. Han, A. K. Mok, W. Zhang, and M. Tomizuka, "RT-WiFi: Real-time high-speed communication protocol for wireless cyber-physical control applications," in *Proceedings of RTSS*. IEEE, 2013, pp. 140–149.
- [30] C. Powell, C. Desiniotis, and B. Dezfouli, "The fog development kit: A development platform for SDN-based edge-fog systems," *IEEE Internet of Things Journal*, 2019.
- [31] V. Jacobson, "Congestion avoidance and control," in *SIGCOMM computer communication review*, vol. 18, no. 4. ACM, 1988, pp. 314–329.
- [32] R. Ludwig and K. Sklower, "The Eifel retransmission timer," *SIGCOMM Computer Communication Review*, vol. 30, no. 3, pp. 17–27, 2000.
- [33] R. Russell and H. Welte. Linux netfilter hacking howto. [Online]. Available: <http://www.netfilter.org/documentation/HOWTO/netfilter-hacking-HOWTO>
- [34] K.-c. Lan and J. Heidemann, "A measurement study of correlations of internet flow characteristics," *Computer Networks*, vol. 50, no. 1, pp. 46–62, 2006.
- [35] Y. Xiao, Y. Cui, P. Savolainen, M. Siekkinen, A. Wang, L. Yang, A. Ylä-Jääski, and S. Tarkoma, "Modeling energy consumption of data transmission over Wi-Fi," *IEEE Transactions on Mobile Computing*, vol. 13, no. 8, pp. 1760–1773, 2013.
- [36] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *The Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.
- [38] F. Chollet *et al.*, "Keras, github," *GitHub repository*, <https://github.com/fchollet/keras>, 2015.



- 1 [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization,"  
2 *arXiv preprint arXiv:1412.6980*, 2014.
- 3 [40] K. M. Koudjonou and M. Rout, "A stateless deep learning framework to  
4 predict net asset value," *Neural Computing and Applications*, pp. 1–19,  
5 2019.
- 6 [41] B. Dezfouli, I. Amirtharaj, and C.-C. Li, "EMPIOT: An energy mea-  
7 surement platform for wireless IoT devices," *Journal of Network and*  
8 *Computer Applications*, vol. 121, pp. 135–148, 2018.