

Aves: A Decision Engine for Energy-efficient Stream Analytics across Low-power Devices

Roshan Bharath Das, Marc X. Makkes, Alexandru Uta, Lin Wang, Henri Bal

Vrije Universiteit Amsterdam

{r.bharathdas, m.x.makkes, a.uta, lin.wang, h.e.bal}@vu.nl

Abstract—Today’s low-power devices, such as smartphones and wearables, form a very heterogeneous ecosystem. Applications in such a system typically follow a reactive pattern based on stream analytics, i.e., sensing, processing, and actuating. Despite the simplicity of this pattern, deciding where to place the processing tasks of an application to achieve energy efficiency is non-trivial in a heterogeneous system since application components are distributed across multiple devices.

In this paper, we present Aves – a decision-making engine based on a holistic energy-prediction model, with which the processing tasks of applications can be placed automatically in an energy-efficient manner without programmer/user intervention. We validate the effectiveness of the model and reveal several counter-intuitive placement decisions. Our decision engine’s improvements are typically 10-30%, with up to a factor 14 in the most extreme cases. We also show that Aves gives an accurate decision in comparison with real energy measurements for two sensor-based applications.

I. INTRODUCTION

With the introduction of Internet-of-Things and 5G, smart sensors will become ubiquitous and will generate massive amounts of data. These sensor-based Big Data are of increasing importance in many fields, such as smart homes, cities, or farming [1], [2]. Such sensors are building blocks of an extremely heterogeneous ecosystem: different devices have different processing power, battery capacity, networking capabilities, and programming environments. Therefore, it is essential to provide a programming environment with which energy-efficient sensor-based applications can be conveniently built.

To reduce application development complexity and aid processing task offloading decisions, we advocate a framework that helps programmers by providing a set of unified APIs that are easy to use, and allow for automatic processing task placement decisions based on device capabilities and application processing complexity.

One crucial deciding factor for processing-task placement is the maximum time the application can run with the remaining battery capacities of the involved devices. Nowadays, for battery-powered devices, such as smartphones or wearables, it is crucial to optimize battery usage. A typical solution to prolong battery life is to offload complex processing or actuation tasks to cloudlet or cloud platforms [3]. While in such one-to-one offloading scenarios optimizing the application lifetime can be simply mapped to maximizing the energy efficiency of single devices, this is not sufficient when we are facing a

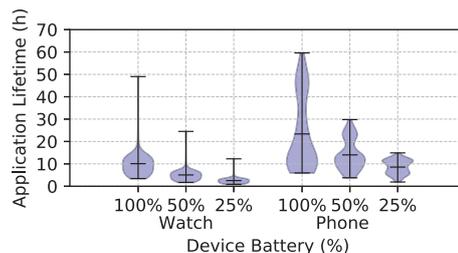


Fig. 1: Variation of application lifetime for different types of processing task for a device with a given initial battery level. This variation is further exacerbated by the available amount of battery and the type of device (smartphone or smartwatch).

distributed heterogeneous environment, where multiple devices are present with different properties.

Deciding where to place processing-tasks for maximizing the lifetime of an application is challenging, especially with the introduction of both local and remote sensing and actuation in a distributed environment. One reason is that the offloading decisions are typically difficult when multiple heterogeneous devices are involved, as depicted in Figure 1. The measurements are from a synthetic benchmark set. The difficulty is due to the variation in the application lifetime for different combinations of parameters such as the complexity of processing, the frequency of sensor data generation, the processing window size, and the battery charge of the devices involved. In addition, the fact that data needs to be moved along with the processing-task incurs communication overhead that also consumes considerable energy, which further complicates the problem. The location of sensing- and actuation-tasks plays a significant role in the offloading decision making. All these render manual decision making impractical.

In this paper, we introduce Aves, a decision engine built on top of our (extended) Cowbird framework [4] for energy-efficient stream analytics across low-power devices and clouds. Aves performs an automatic processing-task placement, based on efficiently exploring the large space of offloading policies and employing a holistic energy model to estimate how long a given application is able to run under a given policy. The energy model uses energy consumption measurements from a collection of synthetic benchmarks that are run once (‘offline’) for the given platform (smartphone, smartwatch). When an application is launched, Aves automatically decides where to do the processing, based on the application-specific computation complexity, the window size, the sensor frequency, and

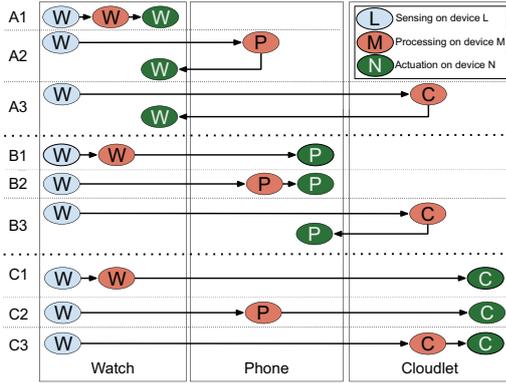


Fig. 2: Various stream analytics patterns in applications that are built across multiple heterogeneous devices with sensing-task on a smartwatch.

the current battery levels of the phone and watch. Here, Aves uses estimation among the pre-measured synthetic values, in combination with an energy model.

Our main contributions are summarized as follows:

- We propose and incorporate an automatic decision-making engine to the framework, which can place processing-tasks for applications that are based on stream analytics in an energy-efficient manner using a holistic energy model (Section III).
- We validate our energy model and decision-making engine with real-world scenarios, confirming the effectiveness of the decision engine (Section IV).

II. BACKGROUND AND MOTIVATION

In this section, we describe various sensor-based applications that can be built using our extended Cowbird framework [4]. Then, we discuss stream analytics patterns and the possibility of processing-task offloading. Further, various metrics for offloading decisions are described.

A. Applications

Context-aware applications can be built using various sensors from devices such as smartphones, smartwatches, or other IoT devices. The data collected from the sensor is processed, and some action is performed based on the result. There can be various applications that involve local and remote sensing and actuation. We provide five concrete scenarios that use sensors from a smartwatch and perform actuation remotely on a smartphone, or a cloud or cloudlet. In this paper, we will use cloudlet to denote either a remote cloud or a local cloudlet.

These scenarios require sensing on the smartwatch and remote actuation on the phone:

Elderly care app: In this application an elderly person's average heart rate needs to be continuously monitored to identify any abnormal behaviors and the caretaker should be notified when help is necessary. The heart rate sensing occurs on the elderly person's smartwatch and the actuation (notification) occurs remotely on the care taker's smartphone.

Coaching app: This application in the context of rowing gathers sensor data from the rower's watch and the coach is notified about the rower's statistics. In such a scenario, the movement (accelerometer) sensing occurs on the rower's smartwatch and the statistics update occurs remotely on the coach's smartphone.

These scenarios require sensing on the smartwatch and remote actuation on the cloudlet:

Environmental monitoring app: This application can gather a median of light sensor data along with GPS from multiple user's watches to spot dark areas in a city for safety purposes.

Noise detection app: It can gather sound sensor data from multiple watches in an area and perform Fast Fourier Transform (FFT) to identify noise pollution (such as airplane, explosion, etc).

Sleep pattern app: It concerns identifying a person's sleep pattern and co-relating it with some other data sets such as disease-related data that require processing over large data sets in the cloud.

Next, we discuss various stream analytics patterns in these applications and the possibility of offloading processing-tasks for some scenarios.

B. Stream Analytics Patterns

Figure 2 illustrates various cases in stream analytics for applications that are built across the smartwatch, the smartphone, and the cloudlet, as already covered in the above scenarios. When sensing is fixed on the smartwatch, we have 9 different cases. Similarly, we can build other cases with sensing on the smartphone or on the cloudlet.

For a sensing-processing-actuating scenario LMN , L represents sensing on the device L , M represents processing on the device M and N represents actuation on the device N . This notation will be used throughout the paper. Note that the devices L , M and N can be a watch, phone or cloudlet. Therefore, a scenario WPC implies that the sensing occurs on the watch (W), processing occurs on the phone (P) and actuation occurs on the cloudlet (C). When the processing-task could be on either phone, watch or cloudlet, we keep the notation 'X'. For example, for the scenario WXC, the processing occurs in a device X which could be either a watch, phone or cloudlet.

From figure 2, we observe three main patterns that include sensing and actuation tasks. Case A1 represents the first pattern (WWW) where both sensing and actuation occur locally. Case B1 (WWP) and C1 (WWC) are when sensing occurs locally on the smartwatch and actuation occurs remotely on the smartphone and the cloudlet, respectively. For these cases, we can offload the processing-task. For example, for A1, when we offload the processing-task to the smartphone, it becomes case A2 (WPW). Similarly, case A3 (WCW) is for offloading the processing-task to the cloudlet. We note that for a given scenario, there exist multiple options to offload the processing-task. Hence, it is important to know the most energy-efficient way to do the processing given a scenario at hand.

C. Metrics for Offloading Decisions

On offloading the processing-task, we can improve various metrics such as latency, throughput, and energy. The latency can be improved by performing simple processing locally on the device and offloading complex processing to remote devices that are nearby and have better capabilities in terms of computing resources. The throughput is another metric for offloading the processing-task. In case of stream analytics, when the processing frequency increases, the number of processing operations that need to be done also increases and it may reach the maximum sustainable throughput especially when data from multiple sensors needs to be processed. Energy is another important metric especially with battery-powered devices such as smartwatches and smartphones, it becomes important to offload processing to save energy. Especially with remote actuation, there is a possibility to do processing at the sensing device (WWP), at the actuation device (WPP) or a different device (WCP). In particular, we focus on a new metric related to energy efficiency, i.e., **maximizing the application lifetime**, by deciding where to place the processing-task given a scenario at hand. Next, we see how Aves can be used for this purpose.

III. AVES DECISION ENGINE

As discussed, the main goal for processing-task placement in Aves is to maximize the application lifetime. To this end, we built an energy model to enable the decision-making process, which involves finding the impact of various stream analytics parameters such as frequency, window size, and operation complexity on the energy measurement. For a given scenario, the energy is measured for a possible combination of parameters, and then curve fitting is applied on the training set. The trained model is used to predict the energy measurement for a given set of parameters. Then, the decision on where to do the processing is made based on the predicted energy measurement and current battery level for each device. The decision engine registers the relevant expression accordingly.

A. Extended Cowbird Framework

To implement Aves, we extend our Cowbird [4] framework. The Cowbird framework already runs on resource-constrained devices, and provides support for distributed sensing, offloads computation from phone to cloud, performs stream analytics on the cloudlet and makes use of an easy to use domain specific language called Swan-song [5]. However, the Cowbird middleware does not provide interaction between heterogeneous resource-constrained devices (phone and watch), and there is no support for actuation (both local and remote), which are vital functionalities to build applications that follow the sensing-processing-actuating model.

To this end, we implement sensors, actuators, middleware, and API for the watch on top of the Cowbird framework. In particular, we introduce distributed actuation by enabling interaction such as register or unregister expressions and send or receive sensor data between the phone, the watch, and the cloudlet. We also extend Swan-song [5] to enable

distributed actuation. Using this addition, we enable the easy building of context-aware applications that involve local and remote sensing, processing, and actuation on a combination of heterogeneous devices such as watches, phones, and cloudlets.

B. Applicability

The application developer registers an expression using our API on the phone. From the expression, the location and the type of sensors and actuators are inferred along with the type of operation and the window size. Aves has already classified the operations based on their complexity and the sensors based on their frequency, and it can also identify any additional configuration (sensor delay) set by the developer. An expression to measure the average heart rate over 10 seconds from the watch and to perform a vibration actuation on the phone can be written by the developer as:

```
watch@heartrate: value {MEAN,10s}  
THEN  
phone@vibrator: vibrate ? duration =500ms
```

where *watch* represents the location of sensing and *phone* represents the location of actuation, *heartrate : value* implies that it is a heart rate sensor, *MEAN* represents aggregation operation and *10s* represents the time window, *vibrate : vibrate* represents that it is a vibrate actuation and *duration = 500ms* represents the duration of vibration which is 500 milliseconds in this case.

After inferring the expression, Aves checks all possible options for placing the operation and detects the best possible scenario for the given expression. The best choice is identified based on an energy model and a decision engine. Using the energy model, the electric current for a given scenario with given parameters is estimated. The decision engine further uses the estimated electric current for all possible scenarios to identify the best possible choice with maximum application lifetime based on the present battery level of the devices that are involved. In the above scenario, the sensing occurs on the watch and the actuation occurs on the phone. Assuming that there is a nearby cloudlet, there is a choice to do the processing (*MEAN, 10s*) on the watch, on the phone or the cloudlet. The decision engine chooses the scenario that can maximize the application lifetime to do the processing based on the battery levels of the watch and of the phone and registers the relevant expression. In this case, we assume that the battery level of the cloudlet can be ignored as it is connected to the power supply. The decision on where to do the processing is complex as it changes based on multiple factors such as battery level, sensing frequency, operational complexity, window size and the location of sensing and actuation. Hence, Aves in combination with Cowbird, helps the developers to ease the application development process.

C. Evaluation Setup

We chose electric current instead of power to model energy for two reasons. First, the change in the voltage for both smartphones and smartwatches are minimal. Therefore, power

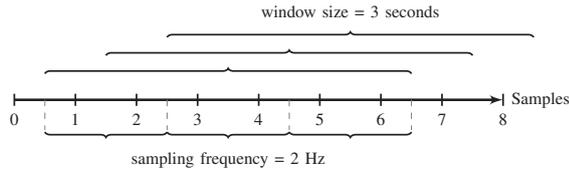


Fig. 3: Event processing scenario.

(voltage \times current) has an impact only on the electric current. Second, to measure the battery life for performing decision making, electric current is used instead of power (as shown in equation 4).

The electric current can be measured based on a hardware-based power monitoring tool such as Monsoon [6] or software-based such as Trepp profiler [7]. Trepp is more widely available, but the measurement from Monsoon is more accurate. For our experiments, we use Monsoon. It was attached to the watch and the phone to measure the electric current usage. We measure the average electric current over 5 seconds (30,000 data points) for various scenarios. The voltage was set to 4 V for all the experiments. As a watch, we use the Moto 360 2nd gen with a battery capacity of 375 mAh, and for phone, we use Nexus 5x with battery capacity 2700 mAh as shown in Table I. Both the watch and the phone communicate with each other using Bluetooth and with the cloudlet using WiFi communication.

D. Parameters

Figure 3 shows an event processing scenario that operates on a window size of 3 seconds and the sensor data collection rate of 2 Hz. On every new event, the operation is performed over the window size of 3 seconds. Since the frequency is 2 Hz, the window size will contain six samples. For the MEAN operation with complexity $\mathcal{O}(N)$ the average over six samples is calculated. Also, other operations with different complexities can be used depending on the scenario. Hence, the number of operations (ϕ) depends on the frequency (f), the window size (w) and the complexity of the processing (\mathcal{O}) and it can be shown as:

$$\phi = \mathcal{O}(g(n)) \quad (1)$$

where $n = f \cdot w$ and $g(n)$ is a function that represents the type of processing such as MEAN, MEDIAN that can have different time complexities.

E. Synthetic Workload

In general, it is not doable to do measurements and modeling for each new application. Our solution is to measure (once, offline) a broad spectrum of synthetic cases to build the energy model. Applications with a given set of parameters (complexity, frequency, window size) are estimated based on the model and the best option (e.g., WWP) that can maximize the application lifetime is chosen. In our previous work, we have already used an online profiler-based approach [8] to

improve the energy-efficiency. However, software profiling will take some time in the beginning to measure the electric current. With the offline model-based approach we can make fast decisions in the beginning and in combination with our earlier work we can also validate or change the decisions in the run-time.

Figure 4 shows the battery life of two scenarios (WXP and WXC) for various possible combinations of parameters. WXP indicates sensing on the watch and actuation on the phone and WXC indicates sensing on the watch and actuation on the cloudlet. The battery life in the figure represents the total time the application can run in hours when both the phone and the watch are at the battery level of 100%. On the x-axis, the label $10-10-N$ indicates a combination of sensor frequency, window size, and operational complexity, respectively. From the figure, in case of WXP and for $10-10-N$, there are three possibilities to perform processing: WWP (locally on the watch), WPP (remotely on the phone) and WCP (remotely on the cloudlet). WCP is the least preferred choice for that combination of parameters. We also note that for a different combination of parameters the best choice also changes.

The figures show five applications that were discussed in Section II-A: coaching, elderly care, sleep pattern, noise detection, and environmental monitoring. Such applications use different combinations of sensors, actuators, and parameters.

F. Energy Model

The energy model is built based on the two basic blocks: the computation on the device and the communication between the devices. The computation is represented in terms of the number of operations (ϕ) associated with a given analytics scenario and the communication is represented in terms of frequency of remote actuation (f). Here, we show the impact of computation and communication on the electric current for both the watch and the phone.

Figure 5 shows the average electric current on varying frequency and window size for both the watch and the phone. The figure is used to identify the relationship of various parameters (independent variables) to electric current (dependent variable). The average electric current is measured separately for each device that is involved in a scenario. To show the impact of frequency, out of all possible scenarios, we show the interesting scenarios such as local sensing, processing and actuation (WWW for watch, PPP for phone), local sensing and processing with remote actuation (WWP and WWC for watch, PPW and PPC for phone), remote sensing with local processing and actuation (PWW for watch, WPP for phone) and remote sensing and processing with local actuation (PCW for watch, WCP for phone). For measuring one parameter (e.g., frequency), we keep other parameters (window size and complexity) constant. To show the impact of window size we chose the scenario that involves only computation, i.e., local sensing, processing and actuation (WWW for watch, PPP for phone). We note that for various scenarios, the average electric current is different. The variation implies that offloading the processing-task will have an impact on the average electric

TABLE I: Device configuration.

Resource type	Device	Operating system	CPU	Memory	Network Communication	Battery Capacity (mAh)
Watch	Moto 360 2nd gen	Wear OS 2.0	1.2 GHz quad-core CPU (only one core enabled)	512 MB	Bluetooth 4.0 WiFi 802.11 b/g	375
Phone	LG Nexus 5X	Android 8.1.0	Hexa-core (4x1.4 GHz Cortex-A53 & 2x1.8 GHz Cortex-A57)	2 GB	Bluetooth 4.2 WiFi 802.11 a/b/g/n/ac	2700
Cloudlet	MacBook Pro	macOS Sierra 10.12.6	8 core (2.8 GHz Intel Core i7)	16 GB	WiFi 802.11a/b/g/n	N/A (connected to power supply)

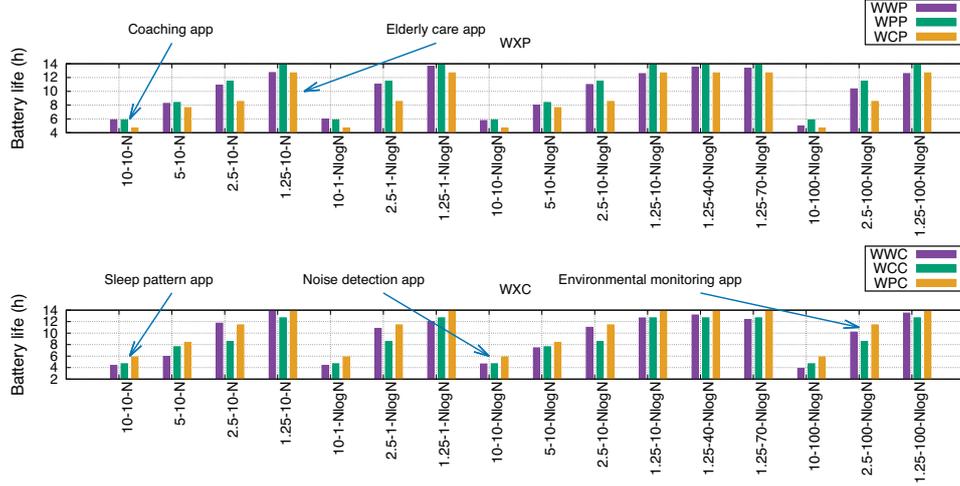


Fig. 4: Battery life comparison of three choices for two cases (WXP and WXC) using various synthetic workloads.

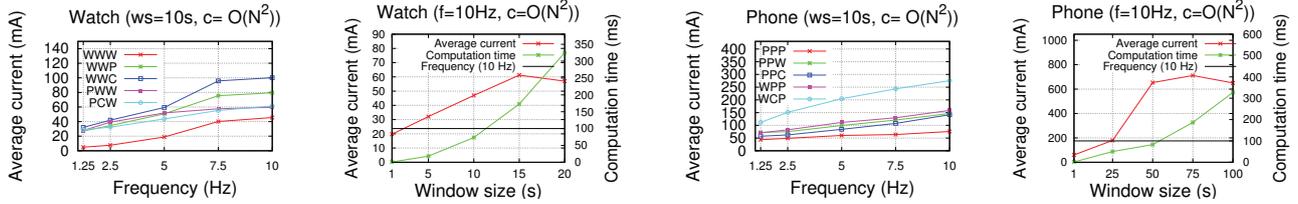


Fig. 5: The impact of frequency and window size on the electric current measurement for multiple scenarios on both watch and phone. The symbol f , ws and c represents frequency, window size and complexity respectively.

current. For a given scenario, we see a linear increase in the electric current used when both the frequency and the window size increase.

Based on the observations, we see a linear correction between various parameters and the electric current. In contrast, a non-linear model would have made it very difficult to interpret and use. Since we have a linear correlation, we can use a simple model using multiple linear regression analysis to model the electric current. For choosing the variables, we note that the electric current is mainly impacted by the sensing frequency, the processing frequency, the amount of processing and the actuation frequency. For expressions where the frequency for sensing, processing, and actuation are the same, we chose the frequency (f) as one variable. The other variable is the number of operations (ϕ) which indicates the amount of processing, and it is chosen based on the frequency, window size, and the complexity as seen in equation (1).

It is important to note that various hardware sensors and actuators will have varying energy usage. However, across multiple offloading possibilities, energy consumption will re-

main the same. For example, for a scenario WXP that uses accelerometer sensor on the watch and vibrate actuator on the phone, the operation offloading possibilities (WWP, WPP, and WCP) will consume the same amount of energy for accelerometer sensor and vibrate actuator. Hence, the energy consumed by the sensor will not have an impact on the energy model. Therefore, we measure the electric current using a test sensor and actuator for the synthetic cases. For making the model simple, we assume that the network transmission error is minimal. The model for measuring the electric current (δ) can be written as:

$$\delta = \beta_0 + \beta_1 f + \beta_2 \phi \quad (2)$$

For expressions where the frequency of sensing (f_1) and actuation (f_2) are different, the model will then be written as:

$$\delta = \beta_0 + \beta_1 f_1 + \beta_2 f_2 + \beta_3 \phi \quad (3)$$

From the model, the regression coefficients (β_n) are estimated based on several samples for each scenario. The

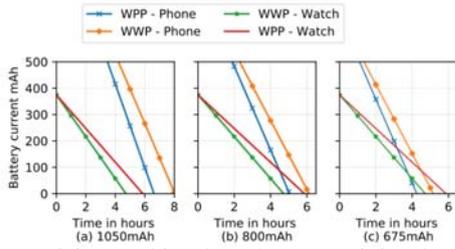


Fig. 6: Decision making based on remaining battery life.

estimation is done using curve fitting based on the least square method. It aims to minimize the difference between observed and predicted values.

G. Decision Making

While our decision engine decides where we should perform computation, we also have to take into account the remaining battery charge, as it can influence the decision where to process the sensor data. Figure 6 shows the influence of the remaining battery charge on the discharge rate of both watch and phone. Here, we fix the battery level of the watch and vary the initial battery charge of the phone. In addition, the figure shows two scenarios, where sensing and actuation are fixed. The first scenario shows two lines, i.e., sensing on the watch in combination with processing and actuation on the phone (WPP). The second scenario shows the lines for both sensing and processing on the watch and only actuation on the phone (WWP). For the watch, the average electric current on running the program while processing the data locally is 79.45 mA, not processing but offloading to the phone the electric current of the watch is 63.98 mA. In case of the phone, the electric current for processing on the watch is 130.69 mA (for using Bluetooth communication), and for processing locally on the phone, it is 158.61 mA. The overall goal of the decision engine is to run the application longest using these two devices.

Figure 6(a) shows the discharge rate of WPP and WWP for the phone exceeds the watch. The maximum lifetime is determined by the highest discharge rate of the devices involved for a given scenario, which in the case of WPP is WPP for the watch at 5.9 hours, and not WPP for the phone. Figure 6(b) shows the discharge rate of WPP for the phone exceeds WPP for the watch. However, we still achieve slightly longer lifetime using processing-task on the phone (WPP). On the other hand, Figure 6(c) shows that keeping the processing-task locally in the watch (WWP) will yield the longest application lifetime. Trivially, and not shown in Figure 6, if the lifetime of both WWP and WPP for the phone is lower than WWP and WPP for the watch, the application lifetime will be at most the lowest discharge rate of the phone, i.e., offloading processing-task to the watch (WWP). Hence, to span the lifetime of a distributed application, we have to take the remaining battery charge into account.

Now, we can generalize the decision-making process to more than two devices (namely phone and watch). A single scenario L_i is computed as follows:

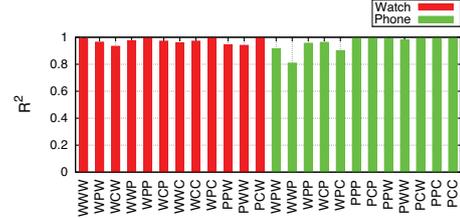


Fig. 7: R-squared value for various scenarios.

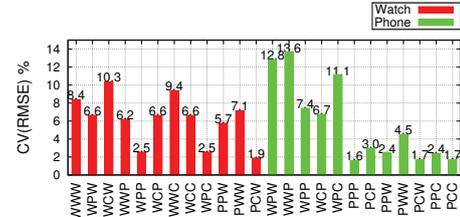


Fig. 8: CV(RMSE) for various scenarios.

$$L_i = \min \left(\frac{P_{B_1} \cdot C_{B_1}}{\delta_{B_{1i}}}, \frac{P_{B_2} \cdot C_{B_2}}{\delta_{B_{2i}}}, \dots, \frac{P_{B_m} \cdot C_{B_m}}{\delta_{B_{mi}}} \right), \quad (4)$$

where $B_{1...m} \in \{Phone, Watch, \dots\}$ (battery-powered devices), C_{B_m} is the battery capacity for the device B_m , P_{B_m} represents the present battery percentage for the device B_m , and $\delta_{B_{m_i}}$ (gathered from the equations 2 or 3) is the electric current for a scenario i for the device B_m . Equation 4 gives the result of the device with the shortest battery life for a given scenario. By applying the following equation:

$$\alpha = \max(L_1, L_2, \dots, L_n). \quad (5)$$

We will find the scenario α with the longest application lifetime out of the possible n scenarios. Note that the devices in the equation 4 are battery-powered. It is possible that for some devices, battery consumption is not an issue. For example, a laptop that is connected to the power supply may not necessarily be concerned about battery life. In such cases, we ignore the measurement of battery life for those devices in equation 4.

IV. EVALUATION

In this section, we give a brief summary of the various evaluations performed. We describe how the validation of the energy model is performed. Then, we show the impact of the battery percentage on the offloading decision and compare it with various offloading strategies. Finally, we compare the decision made by Aves with a hardware power monitor for two real-world applications.

A. Summary

We describe various evaluations performed using the decision engine and address the following questions: 1) What kind of model do we need? Is a linear model a good fit? 2) How accurately can the model predict the electric current measurements? 3) What is the impact of battery life and battery percentage of various devices on the offloading decision? 4)

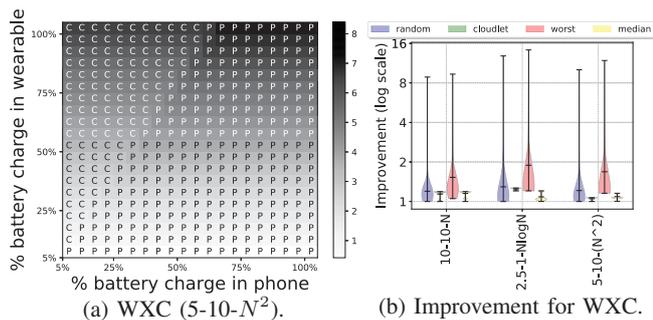


Fig. 9: Impact of battery percentage on the offloading decision.

How much is the maximum improvement of the application lifetime for a given scenario? Does offloading to the cloudlet save energy for any scenario? 5) How accurately can Aves estimate the best option for real-world applications?

B. Validation of the Model

We use 23 samples per scenario out of which 80% of the samples are used for training the model and 20% for prediction. We measure two aspects: coefficient of determination (R^2) [9] and the coefficient of variation of the root mean square error ($CV(RMSE)$). The R^2 value indicates the goodness of fit and is measured between 0 and 1. Figure 7 shows the R^2 value for various scenarios. We note that the R^2 value for 23 out of 24 scenarios is greater than 0.90 out of which the value for 18 scenarios is greater than 0.95. Only 1 scenario (WWP) has R^2 value of 0.80. The result implies that the linear regression model is a good fit.

Next, we measure the variability of errors between the observed values and the predicted values for each scenario using $CV(RMSE)$. It is measured in percentage, and a lower value indicates a higher accuracy for the model. Figure 8 shows the $CV(RMSE)$ for various scenarios. In the figure, we note that the $CV(RMSE)$ percentage is less than 10% for 20 scenarios out of which the value is less than 5% for 10 scenarios. For only 4 scenarios, the value is between 10% and 15% with the highest value for WWP (13.6%). The result indicates that the model is highly accurate.

C. Offloading Decision

Figure 9 shows the impact of battery percentage on the offloading decision. We choose the scenario: WXC. WXC does sensing on the watch and actuation on the cloudlet. Figure 9a shows the heat map of the chosen device to offload processing for a given battery percentage of the phone and the watch. The letter 'P','C' indicates that the processing on the phone, and the cloudlet respectively is the best decision. The intensity represents the application lifetime in hours. The parameter $5-10-N^2$ indicates a frequency of 5 Hz, a window size of 10 seconds and a complexity of $\mathcal{O}(N^2)$.

For scenario WXC, when the battery percentage of watch and phone is more than 75% and less than 50% respectively, it is better to do processing on the cloudlet. On the other hand, when the battery percentage for the watch is below 50% and

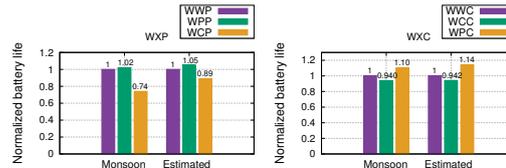


Fig. 10: Normalized battery life comparison of Monsoon hardware power monitor with the decision-engine's estimation for two real-world applications. The battery percentage for both phone and watch are at 100%.

the phone is above 50%, it is better to do processing on the phone. Hence, it is better to offload the processing-task to the phone or the cloudlet than processing locally on the watch. Note that this result is specific to a given scenario with given window size, frequency, and complexity.

Figure 9b shows the improvement (in number of times) for three different parameters. We compare the decision engine's choice of task placement with the random, the cloudlet-based, the worst, and the median-based choices. In the plot, the x-axis represents parameters with different frequency, window size, and complexity, the y-axis represents the improvement for different combinations (from 5% to 100%) of battery percentage of both the phone and the watch, the width shows the probability density of the data at different values and the horizontal stripes show the minimum, the mean and the maximum value. We observe a maximum improvement of 14.25 times. With the parameter $2.5-1-N\log N$, we notice an average improvement of 1.33, 1.31, 1.91 and 1.11 times. Hence, Aves always gives good decisions, whereas random or manual choices may suffer from much more variation.

D. Real Application Comparisons

As described in Section III-B, Aves makes automatic placement decisions for applications without doing new measurements (instead, it uses estimation based on pre-measured synthetic benchmarks). Here, we evaluate how well this automatic decision-making works. We study two applications, and compare the decisions made by Aves (through estimation) against the best decision (determined by measuring). The first application is an elderly care app that is used to measure the average heart rate over 10 seconds of an elderly person and displays it on the phone of the caretaker. This application follows a WXP scenario with parameters $1-10-N$. Since the benchmark set only contained measurements for the closest set of parameter $1.25-10-N$, the measurements for $1-10-N$ had to be estimated using equation 2.

The second application does simple environmental monitoring using $\mathcal{O}(N)$ analysis. The app measures the average light sensor data over 30 seconds from multiple smartwatches and sends it to the cloudlet for further analysis. This application follows a WXC scenario with parameter $3-30-N$. In this case, the benchmark set only contained measurements for the closest set of parameter $2.5-10-N$. Hence, the measurement had to be estimated using equation 1 and 2 for each scenario.

To evaluate our automatic decision engine, we have also actually measured the electric current for all combinations of the applications, using the Monsoon hardware power monitor. It is then compared against the estimated electric current for each scenario. Figure 10 shows the results of the comparison. The battery life measurement is normalized against the local processing (WXP and WWC) to remove the additional current usage caused by sensors and actuators. Removing the overhead will not have an impact on the decision-making process, as described in Section III-F.

In all cases, the decision engine and the measurement based on the Monsoon power monitor chose the same, i.e., WPP for WXP, and WPC for WXC. In the two cases, the maximum application lifetime is attained by using the remote device (phone). Hence, Aves helps the developers to build energy-efficient sensor-based applications.

V. RELATED WORK

Computation offloading has been intensively studied [10]. Its goal is typically reducing the execution time of smartphone-based applications or improving the energy efficiency of mobile devices. In most cases, computation from a phone is offloaded to a local server or a remote cloud over the Internet. While the benefit of faster processing can be generally expected, the impact on the energy efficiency of the smartphone seems less trivial because transferring the data needed for computation can also be very energy-consuming due to intensive communication.

A large body of work is focused on exploring energy-efficient computation offloading, and an overview can be found in [11]. Many projects [12], [13] assume a simple linear energy consumption model, based on which they explore the min-energy computation offloading scheme. Guo et al. also follow the linear energy consumption model and consider a more complex scenario where constraints on task dependencies and completion time deadlines are enforced [14]. Lyu provides a distributed integration architecture of the cloud, edge, and IoT devices, and propose a lightweight framework for energy-efficient selective computation offloading [15]. Offloading decision making for optimized energy efficiency has also been explored for specific applications such as real-time video analytics [16] and augmented reality [17]. Pablo et al. [18] provide a context-aware framework that take into account the energy efficiency in order to choose the best way to run specific tasks in a smartphone.

In contrast, our proposal provides a general way to measure the energy consumption of low power devices and to quickly build a realistic energy consumption model using the measurement results. We are among the first to consider the energy consumption of the whole sensing-processing-actuating cycle of sensor-based applications instead of looking at only the processing part when making computation offloading decisions.

VI. CONCLUSION

We designed Aves, a decision-making engine based on an energy model that automatically places the processing-task

in an energy-efficient manner, given a scenario at hand. The energy-model is built based on a set of synthetic workloads that are measured offline once.

We validate the model and show that our linear regression based model is a good fit and is accurate for most scenarios. We observe that the battery percentage of a device has an impact on offloading decisions. Aves significantly improves the application lifetime compared with different placement strategies. We show that the best choice estimated by Aves is the same as the one measured using a hardware power monitor for two real-world applications.

REFERENCES

- [1] A. Berger, A. Bischof, S. Totzauer, M. Storz, K. Lefevre, and A. Kurze, "Sensing home: Participatory exploration of smart sensors in the home," in *Social Internet of Things*. Springer, 2019, pp. 123–142.
- [2] Y. Sun, H. Song, A. J. Jara, and R. Bie, "Internet of things and big data analytics for smart and connected communities," *IEEE access*, vol. 4, pp. 766–773, 2016.
- [3] G. A. Lewis, "Software architecture strategies for cyber-foraging systems," Ph.D. dissertation, Vrije Universiteit Amsterdam, 2016.
- [4] R. B. Das, N. V. Bozdog, and H. Bal, "Cowbird: A flexible cloud-based framework for combining smartphone sensors and iot," in *Mobile Cloud Computing, Services, and Engineering (MobileCloud), 2017 5th IEEE International Conference on*. IEEE, 2017, pp. 1–8.
- [5] N. Palmer, R. Kemp, T. Kielmann, and H. Bal, "Swan-song: A flexible context expression language for smartphones," in *Proceedings of the Third International Workshop on Sensing Applications on Mobile Phones*. ACM, 2012, p. 12.
- [6] M. Solutions, "Monsoon power monitor," 2019. [Online]. Available: <https://www.msoon.com/>
- [7] Q. Technologies, "Trepn power profiler," 2019. [Online]. Available: <https://developer.qualcomm.com/software/trepn-power-profiler>
- [8] R. B. Das, N. V. Bozdog, M. X. Makkes, and H. Bal, "Kea: A computation offloading system for smartphone sensor data," in *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. IEEE, 2017, pp. 9–16.
- [9] G. R. Ruiz and C. F. Bandera, "Validation of calibrated energy models: Common errors," *Energies*, vol. 10, no. 10, p. 1587, 2017.
- [10] K. Kumar, J. Liu, Y. Lu, and B. K. Bhargava, "A survey of computation offloading for mobile systems," *MONET*, vol. 18, no. 1, pp. 129–140, 2013.
- [11] K. Kumar and Y. Lu, "Cloud computing for mobile users: Can offloading computation save energy?" *IEEE Computer*, vol. 43, no. 4, pp. 51–56, 2010.
- [12] W. Zhang, Y. Wen, and D. O. Wu, "Energy-efficient scheduling policy for collaborative execution in mobile cloud computing," in *INFOCOM*, 2013, pp. 190–194.
- [13] K. Zhang, Y. Mao, S. Leng, Q. Zhao, L. Li, X. Peng, L. Pan, S. Maharjan, and Y. Zhang, "Energy-efficient offloading for mobile edge computing in 5g heterogeneous networks," *IEEE Access*, vol. 4, pp. 5896–5907, 2016.
- [14] S. Guo, B. Xiao, Y. Yang, and Y. Yang, "Energy-efficient dynamic offloading and resource scheduling in mobile cloud computing," in *INFOCOM*, 2016, pp. 1–9.
- [15] X. Lyu, H. Tian, L. Jiang, A. V. Vinel, S. Maharjan, S. Gjessing, and Y. Zhang, "Selective offloading in mobile edge computing for the green internet of things," *IEEE Network*, vol. 32, no. 1, pp. 54–60, 2018.
- [16] L. Zhang, D. Fu, J. Liu, E. C. Ngai, and W. Zhu, "On energy-efficient offloading in mobile cloud for real-time video applications," *IEEE Trans. Circuits Syst. Video Techn.*, vol. 27, no. 1, pp. 170–181, 2017.
- [17] K. Chen, T. Li, H. Kim, D. E. Culler, and R. H. Katz, "MARVEL: enabling mobile augmented reality with low energy and low latency," in *SenSys*, 2018, pp. 292–304.
- [18] P. Graubner, C. Thelen, M. Körber, A. Sterz, G. Salvaneschi, M. Mezini, B. Seeger, and B. Freisleben, "Multimodal complex event processing on mobile devices," in *Proceedings of the 12th ACM International Conference on Distributed and Event-based Systems*. ACM, 2018, pp. 112–123.