



ENOrMOUS: ENergy Optimization for MObile platform using User needS

Ismat Chaib Draa, Smail Niar, Emmanuelle Grislin-Le Strugeon, Morteza
Biglari-Abhari, Jamel Tayeb

► To cite this version:

Ismat Chaib Draa, Smail Niar, Emmanuelle Grislin-Le Strugeon, Morteza Biglari-Abhari, Jamel Tayeb. ENOrMOUS: ENergy Optimization for MObile platform using User needS. Journal of Systems Architecture, 2019, 97, pp.320-334. 10.1016/j.sysarc.2018.10.004 . hal-01879172

HAL Id: hal-01879172

<https://hal.science/hal-01879172>

Submitted on 22 Sep 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ENOrMOUS: ENergy Optimization for MObile platform using User needS

Ismat Chaib Draa, Smail Niar, Emmanuelle Grislin-Le Strugeon

Univ. Valenciennes, CNRS, UMR 8201 - LAMIH, F-59313 Valenciennes, France

Morteza Biglari-Abhari

University of Auckland, New Zealand

Jamel Tayeb

Intel Corporation, Portland, USA

Abstract

Optimizing energy consumption in modern mobile handled devices plays a crucial role as lowering the power consumption impacts battery life and system reliability. Recent mobile platforms have an increasing number of sensors and processing components. Added to the popularity of power-hungry applications, battery life in mobile devices is an important issue. However, the utilization pattern of large amount of data from the various sensors can be beneficial to detect the changing device context, the user needs and the running application requirements in terms of resources. When these information are used properly, an efficient control of power knobs can be implemented to reduce the energy consumption. This paper presents a framework for *ENergy Optimization for MObile platform using User needS* (ENOrMOUS). This framework is able to identify user contexts and to understand user habits, preferences and needs to improve the operating system power scheme. Machine Learning (ML) algorithms have been used to obtain an efficient trade-off between power consumption reduction opportunities and user satisfaction requirements. ENOrMOUS is a generic solution that manages the power knobs. When applied to the CPU frequency, the sound level, the screen brightness and the Wi-Fi, ENOrMOUS can lower the power consumption by up to 35% compared the out-of-the-box operating system power manager schemes with a negligible overhead.

Keywords: Mobile Systems, Mobile Power Consumption, Neural networks, Run-Time Analysis, Data Mining Algorithms

1. Introduction

Mobile and communicating devices are drastically changing our professional and personal activities. The number of smartphones users in western Europe has been predicted to increase from 240.3 millions in 2016 to 279.6 millions in 2019 ¹. The capabilities and hardware complexity of these handled devices are in constant improvement. They include a large number of cores,

¹www.statista.com/statistics/494554/smartphone-users-in-western-europe

a powerful GPU, large caches and a significant number of embedded sensors. For instance, the Samsung Galaxy S7 launched in 2016 contains 10 sensors, representing 6 additional sensors in comparison to the Samsung Galaxy S marketed in 2010. The number of cores has also increased from 1 to 8 cores. In addition, applications running on current and future mobile devices are very power-hungry and more complex. As a result, the user and applications needs in terms of computing, communication and storage can deplete the device's battery in a few hours and impact user satisfaction. Nowadays, battery life has become one of the biggest obstacles for mobile device advancements. For these reasons new power consumption management systems are needed to extend battery life without impacting the processing power and user's experience.

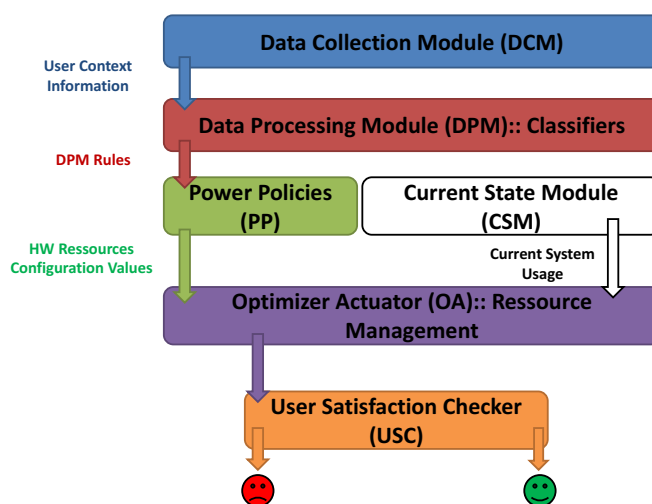
Most of the existing energy saving techniques take into account neither the user individual profiles nor the changing application needs and user requirements. In this paper, our work is based on the user context information and habits. In the literature, many definitions of the context are given [1]. In [2], the context is defined by any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves. Adapted from [2], we define the user context as a set of information that characterize the situation in which the mobile device is used. The information may be divided in 3 subsets depending on their subject: the user, the time and space environment, and the device. This includes the information that may impact the configuration of a hardware resource such as ambient luminosity for the screen brightness, the applications needs in terms of CPU and so on. ENOrMOUS exploits the large set of embedded sensors to collect, store and process the user context information and application requirements at run-time. The collected data are then exploited to generate power policies. The purpose here, is to improve the default power management policies of the operating system and to increase the battery life. In the worst case scenario, ENOrMOUS offers a similar energy management to the operating system. In this paper, we assume that a mobile system can be exploited by many users and to differentiate between them, we identify each user by his/her specific session. Our contribution can be summarized as follows:

1. We exploit the rich sensor hubs and the operating system's application programming interfaces (APIs) to collect a large set of data relative to the user, the launched applications, the consumed resources and so on to represent the user context. This data collection phase is performed by a software module consisted of several probes. Each software probe corresponds to a specific kind of information such as environmental information, system information, etc. These probes are implemented as background process. All the collected information are stored in unified XML way to be processed in the next phase.
2. In the second phase, we propose a new user context classification based on the collected information of the previous phase. The collected information will be processed through machine learning algorithms to extract usage pattern and regularities, identify the user context and predict the user's associated actions. The employed machine learning techniques allow us to use the opportunities to decrease the energy consumed by unused resources in some cases. The classification tends to find a trade-off between the user satisfaction, the running application requirements and power consumption reduction opportunities.
3. In this last phase, the designed classification techniques are used to implement new and efficient power policies to control the power knobs and to reduce the whole system power consumption. Techniques such as Dynamic Frequency Scaling (DFS), Wi-Fi management, screen brightness and sound level adjustments, can thus be efficiently implemented.

The aim is to offer the users, with different profiles and needs, a customized management ap-

proach depending on their habits, requirements, lifestyle and job. In ENOrMOUS, the power savings are achieved at run-time and are transparent to the user. The goal here, is to minimize the user involvement.

2. ENOrMOUS functional architecture



Details of these software modules are as follows:

- **User Probe (UP)**: this probe is linked to the launched applications and the user's preferences in terms of resources used during a certain period of time. The **UP** is executed at run-time and is transparent to the user. The probe indicates if the application is running in the background or foreground and its average running time in a given time period, [as shown in Section 2](#). We assume that users have different behaviors in weekdays and weekends, and also in different periods of a given day. This probe retrieves also the user mobility and permits us to know when the user is walking, running or in a stationary state. We use a two-week period to collect the required information through UP. The user can change this information collection period.

- Environment Probe (EP): this probe gathers information about the user's environment. It indicates: ambient luminosity, ambient noise, device stand, date and time and user position. These information are captured using available embedded sensors such as *Ambient Noise Sensor (ANS)*, *Ambient Light Sensor (ALS)*, *Accelerometer*, *GPS* and so on.
- System Probe (SP): this probe indicates the resources consumed by the system. The collected information are: average utilization of the CPU in percentage, the allocated CPU frequency by the OS, sound level, screen luminosity, Wi-Fi status, battery level and GPS state. This probe can also give information about the user's habits and needs in terms of system configuration preferences. For example, the CPU average utilization varies from a user to another depending on their needs.

The user context data for using each application is stored separately and later all contexts information are used to obtain the resource usage patterns as shown in Figure 2.

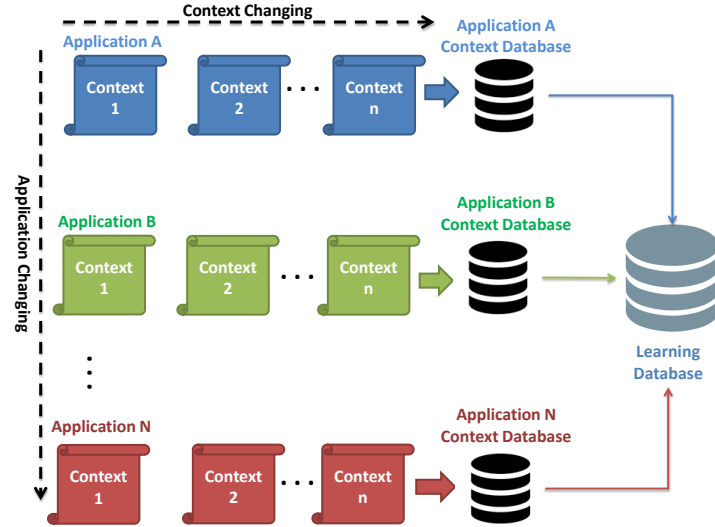


Figure 2: Context changes depending on the foreground application

DCM is the first executed module in ENOrMOUS and corresponds to the data collection phase. These data are captured and stored during two consecutive weeks in order to create a knowledge base that will be processed by the next module. DCM runs as background process and is executed when the mobile device is switched on. All these information are captured by interfacing the aforementioned probes with the embedded sensors and APIs offered by the operating system. The collected information are then stored in a data base in unified XML to be processed by the Data Processing module (DPM). Figure 3 shows the DCM structure.

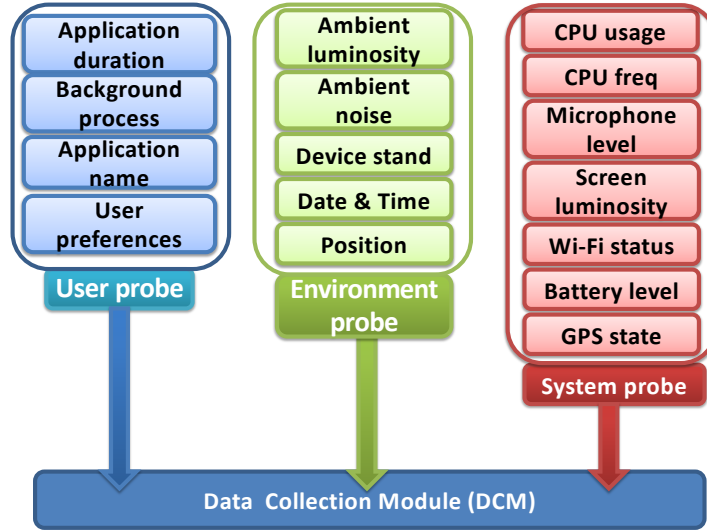


Figure 3: DCM collecting user context

As shown in the experimental section, the overheads due to DCM utilization is low. The **DCM** is launched only once, except when the user is not satisfied by the generated power policies. In this case, the data collection is redone to enrich the knowledge base as it will be detailed in the next sections.

- **Data Processing Module (DPM)**: this module is at the heart of our framework. It uses a large number of ad-hoc classifiers based on neural networks and data mining algorithms to process the stored DCM information and generate Power Policies (PP).
- **Power Policies (PP)**: the DPM's outputs represent a set of rules for each user context. These rules are applied to several power knobs to control them. For example, for managing screen brightness, a rule is the range of the most suitable screen brightness levels for the specified context. If the user is in a dark environment and the foreground application does not require high brightness level, the rule could be [0%-25%]. It means that the maximum value of screen brightness configuration in this context is 25%.
- **Current State Module (CSM)**: this module retrieves the current usage of the device. It permits to not impact negatively the user experience before applying the power policies by configuring some used hardware components. The retrieved information by the CSM are the background application and the consumed resources like the Wi-Fi connectivity, the screen brightness intensity, the sound volume, the mobility sensors and so on. Taking into account the current device usage will also increase the accuracy of the generated power policies.
- **Optimizer Actuator (OA)**: it performs resource management and subsystems configuration. The aim of ENOrMOUS is to propose a generic solution that will manage all the resources. In this paper, we focus on CPU frequency scaling, sound and brightness management and Wi-Fi interface configuration. These resources represent the most power-hungry

hardware components. The Optimizer Actuator takes into account the current power policy and the device's current state before applying the new power policy

- **User Satisfaction Checker (USC):** after adjusting the resources and configuring the hardware, the user's satisfaction is checked by analyzing his/her behavior as shown in Figure 4.

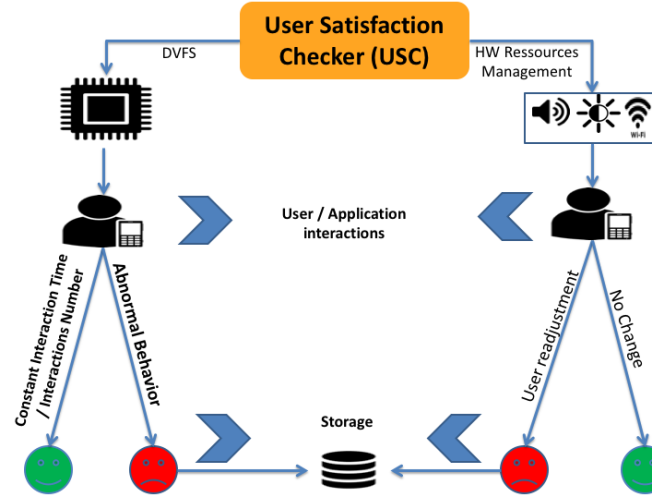


Figure 4: User Satisfaction Checker

The user satisfaction is checked as follows:

- If the user disables the resources management, ENOrMOUS is notified to modify the power policy that is responsible for the user's dissatisfaction. For example, when the CPU frequency is scaled, the user satisfaction checker verifies the acceptance of the proposed frequency by the user. If a new frequency is set by the user, this new value is stored and will be taken into account for the next run in the same context and for the same foreground application.
- In order to check the user satisfaction in terms of CPU configurations, for each specified context, hooks have been implemented as *ActionListners*. Specified contexts are defined by the executed foreground applications. For each foreground application, the evolution of the user's context is shown in figure 4. Thus, we have the different CPU frequencies allocated by ENOrMOUS. The number of contexts varies from one user to another depending on their behavior. We have on average 30 different contexts for the same application, knowing that we have on average 40 applications per user, the number of these contexts reaches 1200. The contexts are stored in an XML file that does not exceed 4 MB per user. The different stored contexts are enriched and redundancies are removed in order to reduce memory storage overheads. The storage of these different contexts is temporary, once the hardware configuration satisfies the user, All the context data is deleted. This operation reduces the impact on the memory consumption.

The ActionListeners identify interaction's speed and manner between the user and his mobile system. When the user launches an application during a given time, the number of times the user touches the screen and refreshes the application is collected and stored. The implementation of ActionListeners is simple and intuitive in the current version of ENOrMOUS. These ActionListeners have been implemented as Input libraries (ILs). The IL starts when a new foreground application is executed. If the user context has already been checked for user satisfaction, the ActionListener is not started and vice versa. Then, an average is calculated and is compared to the current number of interactions. If a difference is detected between these two values, the CPU frequency is increased by the OA and the same process is repeated until no gap between these values is detected. A difference threshold of 20% is used for the number of touches to detect user dissatisfaction. A decrease or an increase of 20% is synonymous with a too low frequency for the user. These 20 % was determined after an explicit questionnaire with users who mentioned that their dissatisfaction can be expressed by incessant tactile support or by an expectation. These 20 % can vary from one user to another, for this reason, as a perspective we will determine this value by using classification algorithms depending on the user interaction manner. This new frequency will be allocated to for the next similar context, thus correcting classifier outputs.

In this paper, we take into account the user's reactions. Indeed, the user's reaction to an optimization is quite indicative of her/his satisfaction. This process allows us to improve the classifiers accuracy and gauge user needs deeply.

3. User Context Classification for Power Reduction

The user context is continuously changing at run-time. The changes concern several parameters such as the foreground and background applications running, the available Wi-Fi connectivity, the mobility of the user, the brightness or noise of the environment, the CPU workload and so on. The first consequence of these variations is that the user behavior and applications needs in terms of resources are strongly affected. Indeed, depending on the context, the user and applications may require more or less resources than those allocated by the operating system. For example, when user is in a dark or dim environment and depending on his/her preferences in terms of luminosity, we can reduce the screen brightness. This reduction is not taken into account by the OS and will permit us to have a gain in terms of power consumption. Secondly, different users may react differently to the same context changes. In a given context, two distinct users launching the same application may need different levels of resources configuration. This behavioral difference stems from many factors such as job, lifestyle, activity, eyesight and so on. In order to take advantage of these characteristics and differences, a device/user context classification has been designed. As mentioned before, the classifiers are implemented in the Data Processing Module (DPM). The classification purpose is to determine which is the device current context of use to be able to generate adaptive power policies for each hardware component. The predictive model issued from the ENOrMOUS classifiers must lead to provide to the user only the hardware resources that are needed for the his/her current context. This resource calibration has to satisfy two major requirements:

- to provide all the possible power reduction opportunities. In the worst case, the proposed power scheme is similar to the OS power consumption;
- to produce power policies which will not impact negatively the user's satisfaction.

The output of the classification is thus used to select the appropriate power policy for each of the controlled resources according to the current context, as detailed in the following.

3.1. ENOrMOUS Resources Classifiers

The aim of the classification is to create a model able to predict the right target values for four resources: CPU, sound volume, screen brightness and Wi-Fi, according to the user context. As mentioned before, the context is defined by information related to the user, the environment and the device:

$$context = (User, Environment, Device)$$

with :

- *User* data include two values about the user's preferences:

$$User = (soundPref, lightPref)$$

From the user probes (see Section 2), the sound level and the brightness level are regularly recorded during one full day. Based on this information, their average level is computed and then assumed to be preferred by the user;

- *Environment* data include ambient noise and lux levels as sensed by the device:

$$Environment = (ambientNoise, ambientLux)$$

- *Device* data include the CPU current frequency, the ratio of the used CPU frequency to the maximum CPU frequency, the number of running applications, the application categories, the download and upload Wi-Fi rates, and the mobile use:

$$Device = (CPUfreq, freqRatio, \#Appli, appliCategories, WiFiRates, mobility)$$

These information constitute the input of the classification. We consider the four resources as independent targets for the classification. For this reason, the context identification is achieved by 4 classifiers, one ad-hoc classifier for each resource as shown in Figure 5. For each classifier, we have selected the input context data that are relevant given the output target resource.

3.1.1. CPU Classification

The CPU classifier uses data that cover application needs and user preferences in terms of computation:

- *CPUfreq*: this input is trivial and represents the CPU frequency allowed by the operating system. In our experiments, we consider four different values, 800 MHz, 1250 MHz, 1750 MHz and 2200 MHz, which are the CPU frequencies allowed by the mobile devices we used for the experiments.

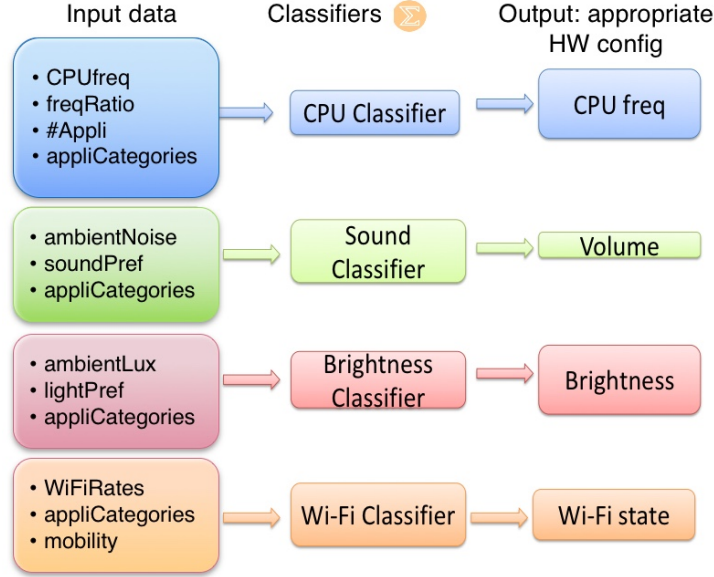


Figure 5: Resources *ad-hoc* classifiers

- *freqRatio*: this parameter indicates the ratio of the allowed OS CPU frequency relatively to the maximal frequency. It helps us to get an idea about the user preferences, when this ratio is high, it means that the user needs high computational resources and vice versa.
- *#Appli*: this input includes the number of running applications in background and foreground. It assists in finding the correlation between the number of running applications and the appropriate CPU frequency.
- *appliCategories*: it represents the category of the foreground application. Indeed, we made a specific classification of the applications into three distinct categories according to their requirements, as will be explained in the next section 3.1.5. The application category is used here to determine its requirements in terms of computation.

The output of the CPU classifier is an appropriate CPU frequency among four classes. These classes are defined by the four frequency values we chose:

- **L**: context requiring a low CPU frequency (under 800 MHz).
- **M**: context requiring a medium CPU frequency (between 800 MHz and 1.25 GHz).
- **H**: context requiring high computing resources (between 1.25 and 1.75 GHz).
- **VH**: context requiring very high computing resources (over 1.75GHz).

3.1.2. Sound Classification

Sound context is built according to three selected inputs:

- *ambientNoise*: this value indicates the noise level of the environment and is captured by the available embedded sensor on the platform. When the ambient noise is too high, the sound level must be increased and vice versa.
- *soundPref*: the sound level that is assumed to be preferred by the user (see Section 3.1). This data can be representative of the user hearing.
- *appliCategories*: the category of the foreground application is used here to determine its requirements in terms of sound level. For example the sound required by *Spotify* differs from the sound level needed by *Word*.

In the absence of a research background related to this question, we made the arbitrary choice of a four level scale for the sound volume, what constitutes a first attempt that should be refined. The output of this classifier is an appropriate sound level among four sound level classes:

- **VL** [0%-25%] : when the user needs a very low sound level.
- **L** [25%-50%]: when the need for the sound level is low.
- **M** [50%-75%]: when the need for the sound level is medium.
- **H** [75%-100%]: when the need for the sound level is high.

3.1.3. Brightness Classifier

The brightness classifier uses data that cover the ambient luminosity, the application needs and the user's preferences: Luminosity context is built according to three selected parameters from the probes.

- *ambientLux*: this input is common to the brightness management methods. For example in the Samsung Galaxy S7, the brightness adjustment is only based on the ambient luminosity [3]. This value indicates the luminosity of the environment and is captured by the available embedded sensor of the platform. This parameter value is very significant when the ambient luminosity is too high or when we need a high screen brightness level, and vice versa. The parameter values are available in the environment probe.
- *lightPref*: the screen brightness level that is assumed to be preferred by the user (see Section 3.1). These data may give information about the user's eyesight.
- *appliCategories*: the category of the foreground application is used here to determine its requirements in terms of screen brightness. For example, *Facebook* brightness needs differs from *Spotify* needs.

As for the sound, we used an arbitrary four level scale for the brightness. The output of this classifier is an appropriate brightness level among these four classes:

- **VL** [0%-25%] : when the user needs a very low screen brightness.
- **L** [25%-50%]: when the need for the screen brightness is low.
- **M** [50%-75%]: when the need for the screen brightness is medium
- **H** [75%-100%]: when the need for the screen brightness is high.

3.1.4. Wi-Fi Classification

For this classifier, we select 4 input data which are:

- *WiFiRates*: the rates of downloaded and uploaded data in kb/s.
- *appliCategories*: the category of the foreground application is used here to determine its requirements in terms of Wi-Fi.
- *mobility*: mobile use of the device, provided by an API that permits to know when the user is stationary, walking or running. It is calculated by using the accelerometer data combined with the compass. It will help us define whether the user is close to a Wi-Fi hotspot or not.

The output of this class is an appropriate Wi-Fi state among these three states:

- **1**: when the user needs to be connected.
- **1/0**: the Wi-Fi interface is on but disconnected.
- **0**: the Wi-Fi interface is disabled.

All these input data values are captured at the same time and stored into a database. The time-period of a collection is adjustable depending on the foreground application change.

3.1.5. Application categories

Each of the four classifiers includes in its input dataset the application category regarding the targeted resource. This parameter has been introduced to improve the classifiers accuracy. The application categories have been subjectively evaluated based on their functionality and our own usage, according to the CPU, Wi-Fi, watching and listening needs. This is illustrated in Table 1.

Table 1: Resource requirement categories for a selection of applications

Applications	CPU	Wi-Fi	Watching	Listening
Youtube	M	1	H	H
Word	L	0	H	L
Skype	L	1	M	H
Facebook	H	1	H	L
Messenger	M	1	H	L
Spotify	L	1/0	L	H
Adobe Reader	L	0	H	L
2048	M	0	H	L
Chrome	M	1	H	M
VLC	M	0	H	H
OneNote	L	0	H	L
Sudoku	L	0	H	L

- For CPU, we use three categories, L for low need, M for medium need and H for high need in terms of computation.

- For Wi-Fi, we also use three categories, 1 for applications that require internet connectivity and 0 for applications that do not. We have also 1/0 for the applications that can run in offline mode like *Spotify* or *Google Maps*.
- For brightness and sound needs, we have respectively evaluated screen watching and listening parameters with three values L for low need, M for medium and H for high needs.

The application categories are used by the four classifiers previously described. These classifiers are at first trained to build a predictive model using supervised machine learning, that means we use a set of instances for which we have the desired outputs (targets). The resulting model is then used to determine the proper power adjustments according to the input data that identify new situations.

In the following, we present how we create the predictive models.

3.2. Creation of the predictive models

The aim is to create four models dedicated to predict the power resource requirements according to the context. These four models correspond to the four classifiers presented previously. This is a classification problem since we know the target classes of the models. Among the numerous classification methods [4], we chose Artificial Neural Networks (ANN) because of our type of data and the uncertain aspect of user behavior [5]. Indeed, neural networks can be used to detect patterns and trends that are too complex to be noticed using statistics techniques. For example, in our case, the uncertain aspect of user behavior is linked mainly to his/her interaction manner following his context. These data may strongly impact the energy consumption of the mobile system and consequently the results of our proposed solution. In addition, we have various data that cover several axes to describe the user context as shown in section 2.

A neural network is made of input nodes, output nodes and hidden nodes that link the input to the output nodes. The hidden nodes are organized in one or more layers, what forms the architecture of the network.

- The inputs units receive information to be processed. In our case, it is the collected user context.
- The outputs units receive the results of the processing. In our case, they are four context classes for each CPU, sound and brightness. For The Wi-Fi, we have three outputs units corresponding to each state.
- The hidden units connect the input units to the output units. Their activity is determined by the activities of the input units and the weights on the connections between the input and the output units.

Many studies discussed the use of several ANN architectures and training algorithms for classification. In ENOrMOUS, since it is based on a set of known-correct outputs for each context, the backpropagation (BP) models are appropriate. The BP algorithm adapts the network according to the input-output associations that are expected.

The method used to create this training dataset is described in the following subsection.

3.2.1. Training dataset

The set of instances for which we have the desired outputs constitute the training dataset, which is based on the user's actions. Examples of appropriate brightness and sound level are inserted in the training dataset through the application of the change blindness mechanism [6]. Indeed, we exploit the change blindness concept by decreasing screen brightness and sound level gradually until it is no longer acceptable by the user. When the screen brightness and the sound volume levels are gradually lowered and the user does not change these values manually, we conclude that the user is satisfied with these new screen brightness and sound volume values. Then, we store these appropriate brightness/sound values as targets for the specified context. However, when our automatic adjustment is not tolerated and is modified manually by the user, the target value for the screen brightness and the sound volume are the last values that have not been modified manually.

A similar method is used to detect relevant CPU frequency: we decrease the frequency until a manual readjustment by the user is detected. Then, we store the last frequency that satisfied the user. In order to define the Wi-Fi targets, for each input sample we try to determine a corresponding connectivity configuration by experimentation. The user feedback is taken into account to correct the classification.

This process represents the only phase where the user changes resources configuration manually. This is necessary for the construction of the learning dataset. This is done just once before the start of the classification. All remaining phases are automatic and do not require any user action. The differences in the users' needs and preferences explain why the classifiers are trained specifically for each user.

Figure 6 shows the target classifiers selecting process

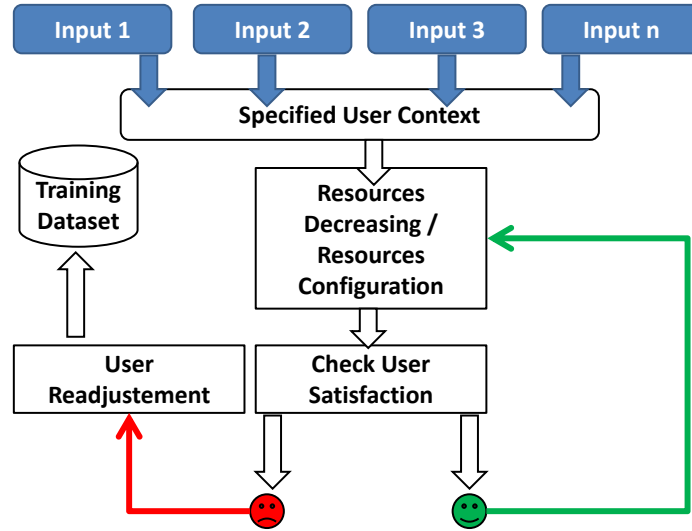


Figure 6: Target classifiers selecting process.

3.2.2. Network architecture

The back-propagation learning algorithm (BPLA) is one of the most studied and used algorithms for neural networks learning. It is a widely used method for ANN learning in many applications [7]. We used the sigmoid function to calculate the output of the neurons. The shape or architecture of the network is another parameter that must be decided upon to design an ANN. The architecture depends on the way the neuron layers are connected to each other, the number of hidden layers and the number of neurons per layer. Unfortunately, there is no generic method to determine the optimum values for these parameters.

About the connections between the layers, two architectures were tested: a simple Feed-forward BP architecture and a Cascade BP architecture [5]. Both of them are among the most popular multilayer ANN. Their network architectures differ by the number of connections between the neurons: in the feed-forward BP architecture, each hidden layer (n) node is connected to the previous layer (n-1) nodes, whereas in the cascade BP architecture, each hidden layer (n) node is connected to every previous layer (input layer-0, 1, ..., n-1). We simulate and compare the results obtained by each type of architecture in order to choose the right model and the suitable configuration for each classifier and each user. The most accurate architecture will be chosen to be embedded on the device.

We perform tests with ANNs containing two intermediate layers and with variations from 1 to 10 neurons in each hidden layer. Weights and biases were randomly initialized. We select 200 inputs samples for each resource and each user to test and choose which of the feed-forward or the cascade BP architecture were the most appropriate. Our networks were trained during 100 training cycles (or epochs). These 200 samples were divided 80% for training, and 20 % for validation. The evaluation of our ANNs has been done by experimentations in order to select different architectures (hidden layers size) and comparing obtained results.

We noticed that:

- In our case, changing the intermediate layers and neurons number did not affect the accuracy result. On the contrary, increasing the ANN size was counter-productive and causes latency with power overheads.
- We also notice that changing the number of input samples affect the accuracy of the ANN. By selecting less than 200 samples, the accuracy was decreased.
- Increasing the number of input samples can lead to a problem of over-learning.
- In the literature, there is no method to determine the ideal number of samples. The most recommended method is experimentation and comparison of different results. For ENorMOUS the number of selected samples is sufficient to demonstrate the feasibility of our solution.

Table 2 presents the neural networks architecture comparison for our four classifiers. This table represents results for one user in order to show how the ANN architecture choice is done. The table presents the mean square error, the number of iterations (Epochs) performed for the validation performance to reach a minimum and the regression results.

- The mean squared error (MSE) for each ANN with the number of epochs is shown. The MSE indicates the accuracy of our algorithm, which should be at an acceptable level.
- Regression results show the relationship between the outputs of the network and the targets during the two phases: Training and validation. If the training were perfect, the network

outputs and the targets would be exactly equal, but the relationship is rarely perfect in practice. The regression plot will confirm our choice concerning the right architecture and algorithm.

Table 2: Neural network accuracy for one user

Resource	ANN	MSE	Epochs	Training	Validation
CPU	Feed FBP	0.21	4	$0.49*Tar + 0.073$	$0.51*Tar + 0.17$
	Cascade FBP	0.062	12	$0.71*Tar + 0.077$	$0.69*Tar + 0.012$
Wi-Fi	Feed FBP	0.02	14	$0.84*Tar + 0.058$	$0.84*Tar + 0.044$
	Cascade FBP	0.11	10	$0.61*Tar + 0.11$	$0.56*Tar + 0.15$
Sound	Feed FBP	0.049	16	$0.92*Tar + 0.002$	$0.87*Tar + 0.039$
	Cascade FBP	0.044	3	$0.69*Tar + 0.085$	$0.75*Tar + 0.07$
Brightness	Feed FBP	4 e-08	19	$0.89*Tar + 0.006$	$0.89*Tar + 0.096$
	Cascade FBP	0.073	3	$0.7*Tar + 0.11$	$0.7*Tar + 0.13$

Training and validation represent the linear equation which shows the relationship between the targets and the obtained outputs for our two neural networks and for one user. The MSE and the regression results analysis provide the following information for the specific user:

- Cascade FBP provides better results than the simple Feed FBP for CPU and sound classifiers.
- Simple Feed FBP provides better results compared to the Cascade FBP for Brightness and Wi-Fi classifiers.
- The accuracy of our different neural networks by referring to the mean square error indicates that the number of samples is sufficient. Increasing this number may eventually improve accuracy, but this requires further data collection.

The ANN architecture selection and accuracy can also be different for each user because of the interaction manner, resources requirements, running application and so on. This difference indicates the need for separate experiments to determine the most suitable architecture for each user.

4. Experimental Results

This section presents the experimental results of our proposed solution. The purpose of these experiences is to validate our framework's architecture, evaluate the obtained energy consumption reduction and the cost of our solution. The presented evaluation is considered as preliminary results in an academic context. The section is divided in four parts. We first present the tools used in the implementation phase and in the experiments. In the second part, we present the obtained experimental results in terms of power consumption. The third part presents additional results regarding the power consumption results obtained with ENOrMOUS. We finally present the cost of our solution regarding the resource usage.

4.1. Tools and Experimental Environments

This section presents the experimental environment, the tools and their usage.

4.1.1. Measure and control tool

The Intel Energy Checker SDKit (IecSDK) [8] has been used to implement our solution. The SDK has been designed to measure and optimize applications energy efficiency. Two components of the SDK are leveraged in this work: the main driver (ESRV - Energy Server) and the Modeler. The Modeler provides the services required to implement data collection process and energy saving heuristics. Several data collection extension modules, a.k.a. Inputs Libraries (ILs), as well as Actuators Libraries (ALs) have been developed. The Modeler is composed of three components: the Front-End (FE), the Input Bus (IB), and the Back-End (BE).

- The **Front-End (FE)** collects the data through the **Inputs Libraries (ILs)**: CPU Utilization, display brightness, battery level, front end applications, etc. If necessary, new data can be collected by developing new ILs.
- Once collected by the ILs, data are made visible on the **Input Bus**. Any module connected to the bus has direct access to the metrics. The IB is the main interface between the FE and the BE.
- The **Back-end (BE)** provides core services e.g. a logger or a power-to-inputs automatic correlation, a watchdog as well as an interrupts and communications manager. The BE can be expanded via **Actuators Libraries (ALs)**. ALs are designed to perform specific actions such as dynamic OS configuration and dynamic platform configuration. Usually ALs are used to implement various optimization heuristics that are driven in real-time by the inputs provided by the FE.

Figure 7 depicts the architecture we used to implement our approach.

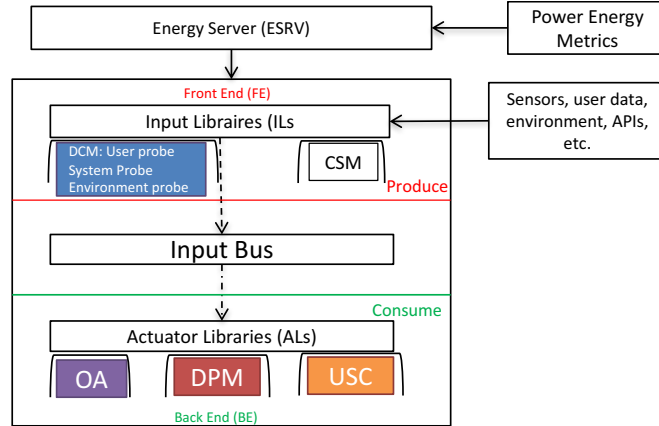


Figure 7: ENOrMOUS architecture based on IecSDK

ENOrMOUS has been implemented as:

- For the data collection module (DCM), we have three ILs corresponding to our three probes.

- For the data processing module (DPM), we have four ALs corresponding to each resources classifier.
- The current state module (CSM) has been implemented as an IL.
- Finally, we have four actuator libraries, one actuator library for each resource.

The developed Inputs Libraries and Actuator libraries are generic and cross platform. They can run under Linux/Android and iOS. They can a priori be embedded on any mobile devices thanks to their low cost in terms of power consumption, memory and CPU usage.

4.1.2. Development of the neural networks

For the neural networks setup for preliminary experiments as shown in section 3.2.1, Matlab neural Networks toolbox functions are used to simulate and compare our networks, adjusting neurons weights to have the most appropriate configuration. After the simulation on Matlab, we have implemented the neural network functions as Actuator Library using a C Wrapper.

4.1.3. Mobile device and OS

The experiments have been carried out on an Ultrabook running Windows 8.1 with a 2.50 GHz Intel dual-core i7-u3667U processor and 4GB of RAM. [The ultrabook is an Intel reference design 2 in 1 which can be used as tablet or ultrabook.](#) Power measurement has been done using the Yokogawa WT210 power analyzer. Table 3 below presents our mobile device hardware features.

Table 3: Intel 2in1 Ultrabook Features

Analyzed Platform	2in1 Intel Ultrabook
Average Battery Life (Hours)	8 Hours
Maximal Power Consumption (W)	23 W
Minimal Powe Consumption (W)	11.5 W
Average Power Consumption (W)	16 W
Number Of Cores	2
Number Of Threads	4
Processor Base Frequency	2.00 GHz
Max Turbo Frequency	3.20 GHz
Cache	4 MB SmartCache
RAM	4 GB

[The ENOrMOUS principle is usable in almost all mobile systems but will require some adjustments and modifications depending on the platform.](#) The main reason why the Ultrabook has been chosen to demonstrate the feasibility of our approach is because of the simplicity to connect it to our measurement device, the Yokogawa WT210 [9]. Other works like [10] propose a time energy model (TEM), which is a regression model to estimate the application energy consumption on real mobile devices. For more accuracy, in our work, we are using the Intel Energy Checker SDK and the Yokogawa WT 210 to measure the energy consumption of the whole process. This mobile device also contains a port for GSM Cards as well as a touch screen. In addition to these hardware features, with the Windows store, we have access to many metro style applications [11] such as *Facebook*, *Viber*, *Shazam*, *Instagram* and so on. These applications

will be used as an entire application and will not be used thru a web browser. In addition, these applications are widely used both on smartphones and tablets. These characteristics makes our solution useful for other mobiles devices.

4.1.4. Users population in the experiments

For our experiments mentioned in section 4.2, we selected six real users with the several aforementioned probes and we collected the required information over two weeks. These six users correspond to six Master students with different profiles and habits. We also simulate six synthetic users to give more details about our power management solution. In total we have a population of 12 users.

4.2. ENOrMOUS power management results

In this section, we measure the gain obtained in power consumption with the use of our solution. To demonstrate our solution efficiency, we made tests in different scenarios with different real user contexts. Recall that the classification output is an appropriate frequency for the CPU, a range of values for the sound and the brightness and one state among three for the Wi-Fi.

The optimizer actuator decreases screen brightness and sound level gradually by *delta* units every *lambda* seconds to the smallest value that satisfies the user. *Delta* and *lambda* have been fixed experimentally to 3% and 4 seconds respectively to not impact the user satisfaction and to converge as quickly as possible to the optimal value.

For the CPU frequency, the Optimizer Actuator decreases the frequency through the Windows API. Finally for the Wi-Fi, the interface is directly configured to the appropriate state (on, disconnected or off).

Table 4 presents a users contexts snapshot. This table gives an instantaneous user context information for our six users, we have four sets of information, we recall that each set corresponds to one of the four neural networks. We also give the foreground application name in addition to its needs in terms of the four resources. These information represent a sample for one specified context, when the foreground application will change, the context will be different and obviously the results presented in Table 5 will also change. For each [resource](#), Table 5 gives the corresponding classification results.

Table 4: Users Context Classification Snapshot

Context data	User 1	User 2	User 3	User 4	User 5	User6
CPU Freq	1750	1250	1250	1750	800	800
# Apps	5	3	2	4	1	2
App CPU CAT	M	L	L	L	L	L
Ambient Noise	35	57	76	15	42	0
Sound user pref	50	60	80	40	53	90
App Sound Need	L	M	H	L	M	H
Ambient Luminosity	25	58	85	35	65	47
Bright user pref	48	47	70	41	52	65
App Bright Need	M	M	M	H	M	H
Download (KB)	17	27	1.7	4.2	7.7	78
Upload (KB)	18	0.0	4.3	3.2	19.2	2.4
Wi-Fi App Cat	OFF	ON	OFF	OFF	ON	ON
Mobility	Low	High	Medium	High	Low	Low
Application	2048	Facebook	VLC	Foxit Reader	OneNote	Skype

Table 5: Classification Results

Classes	User 1	User 2	User 3	User 4	User 5	User 6
CPU freq (MHz)	1250	800	800	1250	800	800
Sound level (%)	[0;25]	[25;50]	[50;75]	[0;25]	[25;50]	[75;100]
Brightness level (%)	[0;25]	[25;50]	[50;75]	[25;50]	[25;50]	[50;75]
Wi-Fi state	Discon	Discon	OFF	OFF	Discon	Connec

We notice from the two tables above that:

- The CPU frequency allowed by ENOrMOUS is always lower or equal to the CPU frequency set by the OS.
- The Wi-Fi classification is highly correlated with the foreground application connectivity need and the user mobility. For example user 2 is running *Facebook*, however the classification's result is Wi-Fi disconnected which is due to the user mobility.
- We can have a similar classification results for a different user contexts for example, user 4 and user 5 have the same classification results for sound [25%-50%] but their context information are different. The same information is noticed for the brightness classification regarding user 2 and user 5.

ENOrMOUS may determine a different configuration to the one we have in Table 4. If the user mobility changes for example, the Wi-Fi configuration will be different than what is presented in Table 5. Recall that the outputs of the classifiers indicate the power policies. The Optimizer Actuator retrieves these outputs and consult the Current State Module. Then, it applies the corresponding power policy on the appropriate power knob. Figures 8, 10 and 9 represent respectively the results from the OS vs. the ENOrMOUS configurations for CPU, sound and brightness.




Figure 8: OS DFS configuration VS ENOrMOUS DFS configuration




Figure 9: OS brightness Vs ENOrMOUS management




Figure 10: OS sound Vs ENOrMOUS management

Table 6 represents the ENOrMOUS Wi-Fi management in comparison with the default OS Wi-Fi management for the specified user context 4. These results correspond to the classification results of the specified user context in the tables above.

Table 6: OS Wi-Fi Vs ENOrMOUS Wi-Fi management

Users	OS Wi-Fi Management	ENOrMOUS Wi-Fi Management
User 1	Connected	Disconnected
User 2	Connected	Connected
User 3	OFF	OFF
User 4	Connected	OFF
User 5	Connected	Disconnected
User 6	Connected	Connected

Figure 11 shows the whole system power consumption gains obtained with ENOrMOUS for the specified user context 4. The measurement tests have been done 10 times, then the average

has been calculated.




Figure 11: OS vs. ENOrMOUS power consumption (gain in %)

By analyzing the results, we notice:

- The total power reduction obtained by managing the CPU is relative to the running applications requirement, the user behaviors and the interaction manner as shown previously. Decreasing the frequency by 450 - 500 MHz corresponds to an average gain of 7% of the whole system power consumption.
- The power consumption reduction gains by decreasing the sound level is very small and corresponds to 4% of the whole power consumption for 100 % units of the volume.
- The gain in terms of power by managing the screen luminosity is linearly correlated with the screen brightness level. We made tests with no running applications in order to have an accurate measurement. 3% of screen brightness level represents approximately 0.155W for the Ultrabook power consumption.
- The power consumption gain by switching OFF the Wi-Fi interface is 1.6 W, which corresponds to 7% of the total power consumption. By disconnecting the Wi-Fi interface the gain is estimated at 3%.

4.3. Power reduction with application sequences

In this section, we carried out additional tests with six synthetic usage scenarios. For these tests, we set up user utilization scenarios illustrated in Figure 12. For each synthetic user, we present widely used applications and a specified context. As mentioned previously, these usage scenarios are synthetic but based on the six real users' data presented in Table 4: the synthetic users have been simulated by combining the collected information from our 6 master students and their answers to some questions about their satisfaction. Then, for each user, we simulate the specified scenario by executing the corresponding applications. We simulate each application for 10 minutes, then an average of the power consumption has been calculated.

We have chosen to use synthetic users to save time and propose a proof of concept concerning ENOrMOUS. We had also limitations concerning the availability of the testing platform and the wattmeter. The purpose of these results is to show how ENOrMOUS will act during a normal usage of the system and when the user is running applications sequences. The information presented in Figure 12 have been obtained as follow:

- First, we choose arbitrarily six days and six time intervals for each day, these information represent the temporal data.
- For each day, the applications sequences have been obtained by interviewing our six real users about their common application sequences. These applications represent widely used application. The sequence application may be composed by one or two applications, the first sequence is represented in blue and the second sequence in green. We also give the background application which are represented in gray.
- The CPU frequency set by the OS, the screen brightness and the volume levels and the Wi-Fi state are captured dynamically via the Current State Module (CSM) for each sequence.
- The user preferences in terms sound volume and screen brightness are those of our 6 real users. Recall that the user's preferences may vary depending on the foreground application and the temporal data.
- The environment information such ambient light, the ambient noise and the time intervals are also captured by the CSM during the experiments.

Figure 12 represents the user context for the 6 synthetic users.

Tuesday 18H – 19H35

Facebook

Instagram

Youtube

Spotify

Spotify

Instagram

CPU freq

Brightness

Vol

Mobility

Noise

Lux

Wi-Fi

1750

82 %

50 %

High

Quite

Dark

ON

1250

81 %

67 %

High

Quite

Bright

ON

Monday 8h30 – 11H

Bank App

Word

Photoshop

Skype

Power point

CPU Freq

Brightne

Vol

Mobility

Noise

Lux

Wi-Fi

1250

69 %

95 %

OFF

Noisy

Bright

OFF

1750

72 %

87 %

OFF

Quite

Bright

ON

Thursday 12h35 – 12H55

2048

Amazon

Messenger

Music

Excel

CPU Freq

Bright

Vol

Mobility

Noise

Lux

Wi-Fi

1250

67%

82%

OFF

Quite

Bright

ON

1250

53 %

52 %

High

Noisy

Indoor

OFF

Sunday 15 H – 16 H

PDF Reader

2048

VLC

Sudoku

2048

CPU freq

Bright

Vol

Mobility

Noise

Lux

Wi-Fi

1750

88%

84%

OFF

Noisy

Dark

ON

1750

88%

84%

ON

Noisy

Indoor

ON

Saturday 1H – 2H30

Tinder

Uber

Facebook

Instagram

Messenger

Viber

CPU Freq

Bright

Vol

Mobility

Noise

Lux

Wi-Fi

1250

64%

61%

Medium

Quite

Bright

ON

1250

62%

57%

Low

Noisy

Indoor

OFF

Friday 7H – 7H30

Photoshop

Facebook

Sudoku

2048

Need For Speed

Music

OS Freq

Bright

Vol

Mobility

Noise

Lux

Wi-Fi

1750

63%

79%

High

Noisy

indoor

ON

1750

63%

82%

Medium

Quite

Dark

ON

User 1

User 2

User 3

User 4

User 5

User 6

Figure 12: Utilization scenarios for the six synthetic users

The classification results for the users 1,2 and 3 are shown in Table 7.

Table 7: ENOrMOUS Classification results for users 1,2 and 3

Users	Sequences	CPU freq	Sound	Bright	Wi-Fi state
User 1	1st	1250	66 %	52 %	Discon
	2nd	/	/	/	OFF
User 2	1st	/	65 %	60 %	/
	2nd	/	/	/	OFF
User 3	1st	/	52 %	53 %	/
	2nd	/	/	43%	/

Figure 13 and 14 represent respectively the power consumption for the 6 users.




Figure 13: Power consumption for users 1,2 and 3

We notice for:

1. User 1: for the first sequence with *Facebook* and *Instagram*, the CPU frequency has been decreased from 1750 MHz to 1200 MHz. The screen brightness and the volume have been decreased respectively to 66% and 52% of the maximum. Finally, the Wi-Fi interface is disconnected. When the user launches *Youtube* and *Spotify*, for this second sequence, ENOrMOUS manages only the Wi-Fi by disabling the interface. The rest of the hardware components are not changed because of the context classification. By analyzing the results shown in Figure 13, we notice that the difference between the OS and ENOrMOUS is minor after the 65th minute which corresponds to Spotify launch. We conclude that we have more opportunities to save power when the user is launching several applications and requiring high hardware resources.

2. User 2: the effect of ENOrMOUS on the power consumption is relatively small in this case. However, the power consumption with the deployment of ENOrMOUS is always below the power consumption of the OS. For the first sequence with the *Bank App*, *Word*, and *Skype* in background, ENOrMOUS decreases the sound volume to 65% and the screen brightness to 60%. When the user launches *Photoshop* and *PowerPoint* in background, ENOrMOUS manages only the Wi-Fi interface, even though the user is not in motion, ENOrMOUS disables the Wi-Fi because the application does not require any connectivity.
 3. User 3: this user is the only one whose power consumption does not exceed 17 W. When user 3 is executing *2048*, *Amazon* and the *Music Player* in background, the CPU frequency and Wi-Fi interface state are not changed. ENOrMOUS reduces the volume and the screen brightness respectively to 52% and 53%. For the second sequence, ENOrMOUS decreases again the screen brightness to 43%. The volume has not been modified by ENOrMOUS because no application requires it.
- The classification results for the users 4,5 and 6 are shown in Table 8.

Table 8: ENOrMOUS Classification results for users 4, 5 and 6

Users	Sequences	CPU freq	Sound	Bright	Wi-Fi state
User 4	1st	800	/	38 %	OFF
	2nd	/	/	/	OFF
User 5	1st	/	/	/	Discon
	2nd	/	/	43 %	/
User 6	1st	/	/	/	Discon
	2nd	800	38 %	37 %	——




Figure 14: Power consumption for users 4,5 and 6

1. User 4: this user is the one with the most power reduction opportunities. For the first sequence, the CPU frequency has been decreased to 800 MHz, the Wi-Fi has been disabled and the screen brightness has been decreased to 38%. For the second sequence, ENOrMOUS keeps the same configurations in addition to sound volume decreasing, the reason is the running of the media application VLC.
2. User 5: the power reduction obtained with ENOrMOUS is relatively small. It is due mainly to the user context and especially the used application needs that do not offer power reduction opportunities. For the first sequence, the Wi-Fi is disconnected without being disabled because of the user mobility. For the second sequence, the screen brightness is decreased by 9%.
3. User 6: For this user, ENOrMOUS is very efficient. For the first sequence, ENOrMOUS disconnects the system without configuring the other hardware components. For the second sequence, the Wi-Fi interface is disabled, the sound volume has been decreased to 38% and the screen brightness to 37%. Finally, the CPU frequency has been decreased to its minimal value (800 MHz).

From these experiments, we can make the following remarks:

- The power reduction opportunities vary from a user to another depending on many factors such as the foreground application needs and the user mobility. All the information that form the context are likely to impact the power consumption.
- Users that require higher hardware resources offer more energy reduction opportunities.
- In all the scenarios, ENOrMOUS was able to decrease at least one resource among the four targeted resources.

Figure 15 shows the obtained results for our 6 users.




Figure 15: Energy saving results comparison

By analyzing the figure, we notice that the most important energy reduction is obtained with user 4. ENOrMOUS reduces the energy consumption by 25.9% for the first sequence and 35.10% for the second one. The energy reduction for the whole scenario reaches 32%. This user is the one who consumes the most energy and paradoxically, it represents the user for whom NORMOUS has been the most efficient. This information confirms that the users that require higher hardware resources offer more energy reduction opportunities.

For user 6, the hardware configuration used by ENOrMOUS offers a gain of 17.68% for the whole scenario. For this user, the obtained energy saving with the first sequence is very small and is 4.75%. However, for the second sequence, ENOrMOUS reduces the energy consumption by 29.1%. These reductions confirm our main hypothesis that the user context impacts the energy consumption. Indeed, for the same user, we have different energy consumption depending on the context.

For user 1, the energy reduction is about 13.29% for the whole scenario. For users 2, 3 and 5, the energy consumption has been reduced respectively by 6%, 4.60% and 6.28%. With these three users, the opportunities of energy reduction are rare because of the executed applications and the required resources.

4.4. ENOrMOUS Overhead Costs

Performing energy saving actions comes with a cost. In this section, we discuss the whole cost of ENOrMOUS in terms of resources consumption, power overheads and latency. We remind that we have three main phases: data collection, data processing and power policies application. We will present the cost of the ENOrMOUS components and modules.

The DCM is the first module used in ENOrMOUS. As mentioned previously, the data collection lasts for two weeks continuously for each user. The probes are implemented as Input Libraries and the collected data is stored in XML files. Figure 16 shows the DCM input libraries power consumption.




Figure 16: DCM power consumption cost over time

We notice that the deployment of the DCM causes a power overhead of 10 Watts during 3 seconds. These three seconds correspond to the probes interfacing with the different sensors. After the interfacing phase, the power consumption is stabilized. This power overhead occurs only when the framework is deployed and is common to all the users. The size of the collected data is relatively small and does not impact the user satisfaction. This size varies according to each user behavior and utilization time. Over the two weeks, the volume of collected data varies

between 22 MB and 34 MB. As mentioned previously the data collection is done once and is then redone depending on the power policies accuracy results. When the user is not satisfied by the proposed HW configurations, we readjust dynamically the CPU frequency, the luminosity level, the volume level and the Wi-Fi interface depending on his feedbacks. In addition, the user feedbacks will be considered as the desired output of each specified context, improving implicitly their satisfaction.

To measure the data processing cost, we measured ANNs training costs. ANNs learning represents the most greedy phase in terms of resources. The training is carried out for the first time at the end of the two weeks of data collection. Then, a new training phase is launched every time the dataset is modified when the platform battery is charging. The overheads are measured in terms of CPU usage, memory usage and power consumption. Table 9 shows the classifiers average training cost for each user with 200 samples. For every user, we calculated the training duration, the CPU utilization, the RAM and the power consumption overheads. We notice that no training took more than 10 seconds, the overheads in terms of CPU utilization, RAM consumption and power occur only during this period of time. Even if these values seem high, it does not seem to have affected the QoS because it occurred only once.

Table 9: Classifiers training cost measurements for our six real users

User	Avg learning time (Sec)	Avg CPU utilization overhead (%)	AVG RAM overhead (MB)	AVG power overhead (W)
User 1	4.2 sec	45 %	120 MB	11 W
User 2	8 sec	39 %	109 MB	10.8 W
User 3	7 sec	43 %	112 MB	11.3 W
User 4	3.7 sec	37 %	104 MB	9.9 W
User 5	6.2 sec	44 %	117 MB	11.7 W
User 6	9.7 sec	49 %	131 MB	10.4 W

These features allow us to embed the whole power optimizer in the frond-end, i.e. the mobile device. However, when the size of the user context data becomes large, this operation will not be possible. Indeed, when these data increases, the classification phase will last longer and will consume more hardware resources. For this reason, we are currently exploring the utilization of cloud computing for ANNs learning to find a trade-off between power consumption reduction, latency and overheads.

5. Related Works

Recently extensive works and projects have been developed for energy consumption optimization in mobile devices. However, very few of them take into account the dynamic aspect of user/application needs in the control of different resources. In this section, we summarize some of the existing works in energy consumption optimization for mobile systems at the application level. We mainly focus on the approaches based on user behavior and experience. Finally, we highlight the main differences with our approach.

One of the first works in this area is [12]. The authors demonstrated the benefit to study real user activities to characterize power consumption and to control the development of power optimization. Their experiments on an HTC ARM based mobile phone show important differences

between users behaviors. They also demonstrated that CPU and screen are the most demanding components in terms of energy. For the screen, the total utilization time is dominated by a small number of long intervals, with a duration of about 100 s.

Many works show the importance of considering the user's activities and behaviors to optimize power consumption in mobile systems. In [13], the authors proposed an approach which takes into account the user experience to apply different power optimization techniques. They developed a new *cpufreq governor*. Their dynamic clock scaling approach provides a mechanism to change the clock speed of the CPUs at run-time. Their proposed *cpufreq governor* analyzes the user perceived response time of the applications at runtime. Then this information is used to control the CPU frequency. The CPU energy consumption is reduced by up to 65.5% over the Android's default on-demand [14] *cpufreq governor*. The authors also exploited the characteristics of the interactions between the user and the system to minimize energy consumption. They used the elapsed time between two consecutive interactions to decrease the screen brightness during this interval in order to have a gain in terms of the whole system energy consumption. Authors in [15] and [16] proposed user activities and context information-based techniques and several management policies for each hardware component have been developed. In their approach the CPU frequency was adjusted dynamically depending on the workload. They also proposed to reduce background process life time depending on the obtained patterns. An energy consumption reduction of up to 20% in comparison with commercial solutions like Juice Defender [17] has been obtained. However their solution is not completely automatic and requires modifications in the running application source code.

Some of the existing works use machine learning techniques to classify the running applications or user activities. Targeting the Wi-Fi consumption, the approach proposed in [18] makes a selection among the applications to give priority to those with the highest network interactivity level. The applications are classified as high or low priority according to network traffic data with the help of a SVM (Support Vector Machine) classifier. On this basis, only the traffic from high priority applications is allowed in order to save energy. In [16], the authors proposed a classification of user activities in terms of screen brightness needs and they correlated these data with the current context information. Then, they used machine learning techniques to predict the required luminosity. CAPEd [19] improves the average satisfaction by 23.5% compared to the default scheme. In [20], the authors presented the power monitor which is a client-server architecture developed to collect usage logs from Android powered devices. Based on the utilization patterns, power saving profiles are generated and personalized to match the needs of each device in the system. The experimental results show that the power monitor can increase the battery life by almost 90%. However this solution has some privacy issues which are due to the exploitation of usage pattern generation. The survey section in [21] provides a useful list of studies that take into account human activities to save the energy of embedded and wearable sensing systems. Most of the listed studies use machine learning techniques, a large panel of different techniques such as Bayesian classifier, support vector machine and so on. However very few of the studies, allow to predict future users requirements to allow energy consumption minimization.

Approaches like [22], [23] and [24] focused on offloading computing to extend the battery life of the mobile devices. In [22], the authors proposed FMOCO, a framework which takes into account the offloading requests submitted by users during a specific period of time and proposed the best possible offloading strategy based on each individual user's computing power, bandwidth, as well as available cloud resources. In [23], the authors proposed a scheduling algorithm that adapts offloading decisions in fine granularity in dynamic wireless network conditions and verify its effectiveness through trace-driven simulations. In this paper, we focus on the neural

networks implementation in the device. However, we are currently working on a method to offload the classification process in a remote server to optimize the cost of our solution.

In [25] we proposed a software component based on application sequences prediction and offline application classification to improve software power management scheme. In [26], we propose to use a software agent whose goal is to save the energy of mobile devices with the lowest effect on QoS.

Compared to previous works, ENOrMOUS presents the following main new features:

- The user context is defined and classified by taking into account several sources of information. This classification leads to an efficient control of the resources such as CPU frequency, sound level, screen brightness and Wi-Fi.
- In our approach, one of the most important parameter is user satisfaction. By taking this parameter into account, ENOrMOUS is able to propose a customized power management solution. This approach has been rarely used in current commercial solutions.
- For screen brightness, ENOrMOUS differs from [27] and [28] in terms of objectives. Our work aims to improve OS policies independently from the screen hardware characteristics. In this sense, our approach is more applicable in real systems and will be more efficient in existing platforms.
- Finally, ENOrMOUS is modular due to the utilization of the Intel Input Libraries (ILs) and Actuator Libraries (ALs) [8]. This modularity along with the utilization of the Intel Energy Checker SDK, makes our solution flexible. Enhancing the optimizer by new data from a new sensor or adding a new actuator for managing a new hardware component is relatively easy by the use of these ILs and ALs.

6. Conclusion

In this paper, a new technique for energy consumption reduction in mobile systems has been proposed. Our approach is based on machine learning techniques. We also used the Intel Energy Server tool which allows the definition of new Input Libraries (ILs) for collecting data and Actuators Libraries (AL) for controlling hardware components such as screen brightness and CPU frequency. Compared with existing methods, where the users needs and behaviors are rarely taken into account, in our approach, we not only consider these elements but, we also take into account information from context and environment. In this paper we focus on CPU frequency, sound level, screen brightness and Wi-Fi classification. However, the proposed solution is generic and could be applied for any hardware component such the GPS, GPU utilization and so on. In comparison to the advanced energy management provided by Windows OS, for some situations the gain offered by our approach reaches 35% depending on the user's needs and habits.

This work will be extended at several levels:

- Considering more possible user patterns by adding information such as geographical position, battery level usage curve. We are also extending the set of tested applications. and more applications.
- Exploring the correlation between the user satisfaction, the training cost and the accuracy of power policies application.

- Adding a classifier for cellular data, which could take as input the application used, the applications in background, the mobility of the user and the use of Wi-Fi. The output of this classifier can determine the desired signal quality 2 G, 3 G, 4 G or a total deactivation of mobile connection.
- As long term perspective, we plan to realize a similar tool for low power mobile systems, such those based on ARM processors. Finally, some applications have varying HW resource needs at different runtime phases. For this reason, it could be interesting not to consider the needs as constant for the whole application execution time but also the variable needs at different execution phases of the application.

Acknowledgment

The Authors would like to thank Intel Corporation and especially the Intel Research Council for the support given to the project and the tools.

References

- [1] M. Baldauf, S. Dustdar, F. Rosenberg, A survey on context-aware systems, *International Journal of Ad Hoc and Ubiquitous Computing* 2 (4) (2007) 263–277.
- [2] A. K. Dey, Understanding and using context, in: *Personal Ubiquitous Computing*, Vol. 5, Springer-Verlag, London, UK, UK, 2001, pp. 4–7.
- [3] A. Pit, Ambient luminosity for screen brightness (2015).
URL <http://www.androidpit.fr/comment-controller-luminosite-ecran-optimiser-batterie>
- [4] S. Kotsiantis, Supervised machine learning: A review of classification techniques, in: *Informatica* (Ljubljana), Vol. 31, 2007.
- [5] D. Badde, A. Gupta, V. K. Patki, Cascade and feed forward back propagation artificial neural network models for prediction of compressive strength of ready mix concrete, *IOSR Journal of Mechanical and Civil Engineering* 3 (1).
- [6] D. J. Simons, D. T. Levin, Change blindness, in: *Trends in cognitive sciences*, Vol. 1, Elsevier, 1997, pp. 261–267.
- [7] Z.-G. Che, T.-A. Chiang, Z.-H. Che, et al., Feed-forward neural networks training: a comparison between genetic algorithm and back-propagation learning algorithm, *International Journal of Innovative Computing, Information and Control* 7 (10) (2011) 5839–5850.
- [8] I. Corp, Intel energy checker software development kit userguide (2011).
URL <https://software.intel.com/en-us/articles/intel-energy-checker-sdk>
- [9] N. Hirofumi, N. Naoya, T. Katsuya, Wt210/wt230 digital power meters, *Yokogawa Technical Report* 35.
- [10] L.-T. Duan, B. Guo, Y. Shen, Y. Wang, W.-L. Zhang, Energy analysis and prediction for applications on smart-phones, Vol. 59, Elsevier, 2013, pp. 1375–1382.
- [11] A. Asthana, R. Asthana, ios 5, android 4.0 and windows 8—a review, in: *IEEE Code of Ethics*, 2012.
- [12] A. Shye, B. Scholbrock, G. Memik, P. A. Dinda, Characterizing and modeling user activity on smartphones: Summary, in: *Proceed. of the ACM SIGMETRICS Int. Conf. on Measurement and Modeling of Computer Systems, SIGMETRICS '10*, ACM, New York, NY, USA, 2010, pp. 375–376.
- [13] W. Song, N. Sung, B.-G. Chun, J. Kim, Reducing energy consumption of smartphones using user-perceived response time analysis, in: *Proceed. of the 15th Workshop on Mobile Computing Systems and Applications, HotMobile '14*, ACM, New York, NY, USA, 2014, pp. 20:1–20:6.
- [14] V. Pallipadi, A. Starikovskiy, The ondemand governor, in: *Proceedings of the Linux Symposium*, Vol. 2, 2006, pp. 215–230.
- [15] S. Datta, C. Bonnet, N. Nikaein, Power monitor v2: Novel power saving android application, in: *Consumer Electronics (ISCE), 2013 IEEE 17th International Symposium on*, 2013, pp. 253–254. doi:10.1109/ISCE.2013.6570213.
- [16] M. Schuchhardt, S. Jha, R. Ayoub, M. Kishinevsky, G. Memik, Caped: Context-aware personalized display brightness for mobile devices, in: *Proceedings of the 2014 International Conference on Compilers, Architecture and Synthesis for Embedded Systems, CASES '14*, ACM, USA, 2014, pp. 19:1–19:10.
- [17] Laterdroid, Juice defender (2013).
URL <http://www.juicedefender.com/>

- [18] A. Pyles, X. Qi, G. Zhou, M. Keally, X. Liu, Sapsm: Smart adaptive 802.11 psm for smartphones, in: Proceed. of the 12th int. conf. on ubiquitous computing (UbiComp'12), 2012, pp. 11–20.
- [19] M. Schuchhardt, S. Jha, R. Ayoub, M. Kishinevsky, G. Memik, Caped: Context-aware personalized display brightness for mobile devices, in: Proceedings of the 2014 International Conference on Compilers, Architecture and Synthesis for Embedded Systems, ACM, 2014, p. 19.
- [20] S. Datta, C. Bonnet, N. Nikaein, Personalized power saving profiles generation analyzing smart device usage patterns, in: Wireless and Mobile Networking Conference (WMNC), 2014 7th IFIP, 2014, pp. 1–8. doi:10.1109/WMNC.2014.6878858.
- [21] D. Gordon, J. Czerny, M. Beigl, Activity recognition for creatures of habit, *Personal and Ubiquitous Computing* 18 (1) (2014) 205–221.
- [22] W. Liu, W. Gong, W. Du, C. Zou, Computation offloading strategy for multi user mobile data streaming applications, in: Advanced Communication Technology (ICACT), 2017 19th International Conference on, IEEE, 2017, pp. 111–120.
- [23] L. Zhang, D. Fu, J. Liu, E. C.-H. Ngai, W. Zhu, On energy-efficient offloading in mobile cloud for real-time video applications, *IEEE Transactions on Circuits and Systems for Video Technology* 27 (1) (2017) 170–181.
- [24] A. Ahmed, A. A. Hanan, K. Omprakash, M. Usman, O. Syed, Mobile cloud computing energy-aware task offloading (mcc: Eto), in: Communication and Computing Systems: Proceedings of the International Conference on Communication and Computing Systems (ICCCS 2016), Gurgaon, India, 9-11 September, 2016, CRC Press, 2017, p. 359.
- [25] I. C. Draa, S. Niar, J. Tayeb, E. Grislin, M. Desertot, Sensing user context and habits for run-time energy optimization, *EURASIP Journal on Embedded Systems* 2017 (1) (2017) 4.
- [26] I. C. Draa, E. Grislin-Le Strugeon, S. Niar, An energy-aware learning agent for power management in mobile devices, in: International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, Springer, 2017, pp. 242–245.
- [27] D. Shin, Y. Kim, N. Chang, M. Pedram, Dynamic voltage scaling of oled displays, in: Design Automation Conference (DAC), IEEE, 2011, pp. 53–58.
- [28] X. Chen, Y. Chen, Z. Ma, F. C. Fernandes, How is energy consumed in smartphone display applications?, in: Proceedings of the 14th Workshop on Mobile Computing Systems and Applications, ACM, 2013, p. 3.