

Design and Evaluation of a Person-Centric Heart Monitoring System over Fog Computing Infrastructure

Extended Abstract

Orestis Akrivopoulos
Spark Works ITC Ltd., United Kingdom
akribopo@sparkworks.net

Athanasios Antoniou
Spark Works ITC Ltd., United Kingdom
a.antoniou79@sparkworks.net

Dimitrios Amaxilatis
University of Patras, Greece
amaxilat@ceid.upatras.gr

Ioannis Chatzigiannakis
University of Rome “La Sapienza”, Italy
ichatz@dis.uniroma1.it

ABSTRACT

Heart disease and stroke are becoming the leading cause of death worldwide. Electrocardiography monitoring devices (ECG) are the only tool that helps physicians diagnose cardiac abnormalities. Although the design of ECGs has followed closely the electronics miniaturization evolution over the years, existing wearable ECG have limited accuracy and rely on external resources to analyze the signal and evaluate heart activity. In this paper, we work towards empowering the wearable device with processing capabilities to locally analyze the signal and identify abnormal behavior. The ability to differentiate between normal and abnormal heart activity significantly reduces (a) the need to store the signals, (b) the data transmitted to the cloud and (c) the overall power consumption. Based on this concept, the HEART platform is presented that combines wearable embedded devices, mobile edge devices, and cloud services to provide on-the-spot, reliable, accurate and instant monitoring of the heart. The performance of the system is evaluated concerning the accuracy of detecting abnormal events and the power consumption of the wearable device. Results indicate that a very high percentage of success can be achieved in terms of event detection ratio and the device being operative up to a several days without the need for a recharge.

CCS CONCEPTS

• **Human-centered computing** → **Mobile devices**; • **Computing methodologies** → Machine learning;

KEYWORDS

IoT, Healthcare, mHealth, Wearable, Prototype, Evaluation

ACM Reference Format:

Orestis Akrivopoulos, Dimitrios Amaxilatis, Athanasios Antoniou, and Ioannis Chatzigiannakis. 2017. Design and Evaluation of a Person-Centric Heart Monitoring System over Fog Computing Infrastructure: Extended Abstract.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HumanSys'17, November 5, 2017, Delft, Netherlands

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5480-6/17/11...\$15.00

<https://doi.org/10.1145/3144730.3144736>

In *Proceedings of The First International Workshop on Human-centered Sensing, Networking, and Systems (HumanSys'17)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3144730.3144736>

1 INTRODUCTION

Heart disease and stroke cause the death of about 17 million people annually – approximately 33% of all deaths worldwide (one every 60 seconds) [6]. The only tool available to physicians and cardiologists to evaluate irregular heart rate, heart rhythm, and diagnose cardiac abnormalities are Electrocardiography monitoring devices (ECG). Such high-accuracy devices that are capable of recording the heart's patterns are available only in hospitals and are bulky for usage in outside place, e.g., in home environments. As many heart abnormalities take place at random intervals, assessments based on recordings at given the time of day are inadequate in spotting heart problems.

The design of electrocardiogram devices has followed closely the electronics miniaturization evolution over the years [13]. Wearable monitoring devices that are able to provide accurate and precise measurements are gaining support from physicians and healthcare units. Yet, currently available portable ECG devices rely on 1-3 leads which have *limited accuracy and reliability*, as electromagnetic noise may be either of environmental (e.g., mobile phones, electrical wires, appliances) or internal origin (e.g., produced by muscular contractions) and remains a frequent source of abnormally appearing ECGs, encountering for repeat test or false positive test [9]. This problem is addressed by introducing more than one recording points (three, six, twelve ECG leads) in order to acquire more information from more than one heart regions. The additional recording points increase the power consumption of the device and produce a high-definition signal that requires bigger memory capacity to be stored [2]. This creates a trade-off between accuracy and longevity of the wearable device.

Another critical aspect of existing wearable solutions is their *lack of computational power to locally process the ECG recordings and detect abnormal behaviour*. The recorded signal is transmitted to a nearby gateway device that relays the recording over to the cloud where advanced analysis algorithms are executed [2, 10]. Depending on the number of leads used, even a trace of a short period can require large amounts of data to be transmitted over the wireless network, affecting the battery performance of the wearable device and possibly also that of the gateway (e.g., in case of a smartphone).

Given the nature of the heart activity, the large majority of the trace periods correspond to normal heart activity. In most cases, large amounts of data transmitted over the network infrastructure and processed at the cloud level are simply dropped, leading to a significant waste not only of the battery resources of the wearable device but also of network and cloud resources. It is evident that minimal possible latency, network bandwidth preservation, and efficient data storage resource utilization are elements of paramount importance for any IoT-related application [4, 18].

The paper examines the design approach in which the wearable devices become capable of executing advanced alerting algorithms, so that only the segments that correspond to an abnormal behavior are required to be stored and transmitted to the cloud infrastructure. Based on this approach, the wearable device would conserve battery power and minimize memory requirements [1]. The concept of combining the resource-bound last-mile sensors of any IoT-related application with computational capabilities is receiving increasing attention from researchers and practitioners. The so-called Fog computing approach extends the cloud computing paradigm by migrating data processing closer to the production site, accelerates system responsiveness to events along with its overall awareness, by eliminating the data round-trip to the cloud. Offloading large datasets to the core network is no longer a necessity, consequently leading to improved resource utilization and quality of experience (QoE) [19].

In this paper, the HEART platform is presented that combines wearable embedded devices, mobile edge devices, and cloud services to provide on-the-spot, reliable, accurate, and instant monitoring of the heart. Initially, a medical specialist remotely trains the wearable device to learn to identify critical events related to the particular patient. Following this training session, the wearable device becomes capable of interpreting a large number of heart abnormalities with very high accuracy, independently of the cloud and edge services. The platform establishes a connection between the patient, the caregiver, and the physician allowing them to continuously fine-tune the system's functionality. The patient can receive customized feedback without having their caregiver continuously present, and by doing so learning to better manage their own condition. In this way, the independence of the patients is significantly increased. At the same time, continuous communication between patients and their caregivers/physicians is enabled in emergency situations or in cases when the patient requires their full support.

The rest of the paper is structured as follows. In Sec. 2 selected relevant literature are presented. Sec. 3 discusses the person-centric heart monitoring use case scenario that is used to identify key requirements and performance indicators of the system. Sec. 4 presents the key aspects of the design of the HEART system. Sec. 5 provides the results of the evaluation of the performance of the system in a real-world environment in terms of accuracy to detect abnormal events and power consumption. Finally Section 6 concludes and presents future work directions.

2 RECENT AND RELATED PREVIOUS WORK

The design of efficient health monitoring systems has been a topic of active research over the last few years, mostly reinvigorated by the numerous advances in communication protocols and access

technologies. Most researchers tried to merge wireless sensor networks with smart gateways [3, 11], and also use smartphones as gateways for developing personal health monitoring systems [2, 20]. Specialized diagnosis techniques are utilized for specific types of health monitoring applications [8] or platforms using dedicated IoT communication protocols [21], however, all aforementioned solutions somehow fail to harness the full extent of capabilities fog computing offers towards providing a holistic solution that will be applicable to a larger number of use cases, involving several actors.

Existing algorithms for computerized detection of abnormal heart activity based on the interpretation of ECG recordings are known to have very low accuracy [17]. For this reason, diagnosis is performed by expert medical specialists and physicians. Very recently, [15] investigated the use of machine learning techniques in order to reduce the errors in the computerized electrocardiogram interpretation of cardiac rhythm. An application of K-Nearest Neighbor (KNN) algorithm as a classifier for detection of QRS-complex in ECG produced very high detection rates. More recently [14] apply a 34-layer convolutional neural network which maps a sequence of ECG samples to a sequence of rhythm classes. They compare the performance of the deep learning network with certified cardiologists in detecting a wide range of heart arrhythmias from electrocardiograms recorded with a wearable monitor. Interestingly their system exceeds the average cardiologist performance in both recalls (sensitivity) and precision (positive predictive value).

In 2016 a number of asymmetric multicore embedded processors have been introduced that provide a dedicated core for hardware-assisted pattern matching, such as the NXP MPC8572E module [16] and the Intel® Curie™ module [5]. These embedded platforms allow the execution of machine learning techniques with significantly reduced power consumption requirements. To the best of our knowledge, our work is the first attempt to exploit the novel computational capabilities introduced by these modules in order to allow machine learning based data analysis to be relocated to the ECG device. Given the limited capabilities of the hardware-based implementation of the K-Nearest Neighbor (KNN) algorithm for supervised learning, we develop a system that achieves very high detection rates that in some cases match the accuracy of the algorithm implementation for generic purpose hardware processors.

3 A PERSON-CENTRIC USE-CASE SCENARIO

We study a heart monitoring scenario that targets end-users that might be either patients that have recuperated from an incident and need to be monitored periodically, or other individuals that need to check and be aware of their medical condition on a regular basis as part of a preventive health monitoring and early diagnosis of potentially alarming medical conditions. Such individuals could be athletes or people that have the record of hereditary conditions in their family history etc. The users carry the wearable device and periodically consult with their physician to acquire feedback on how to improve their health and avoid certain activities.

Initially, the wearable ECG goes through a *learning phase* in order to collect an adequate number of ECG recordings based on which we will train the pattern matching engine to match the needs of the particular patient. This is a necessary first step as the evaluation of an ECG depends on anthropometric data (body height, body weight)

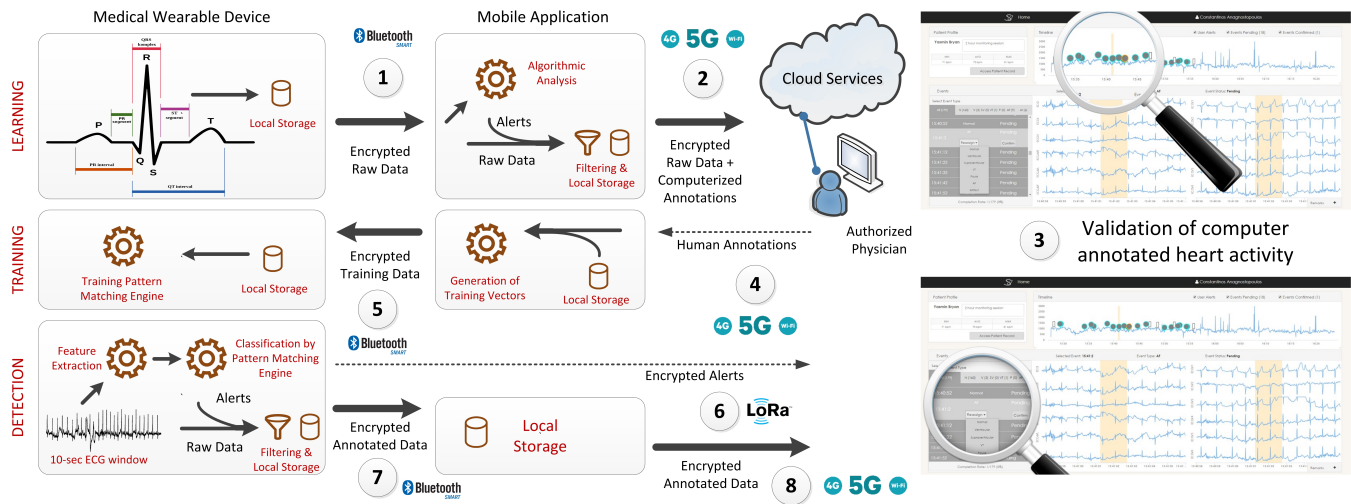


Figure 1: Information exchange across different layers during the three phases of Learning, Training and Detection

on age and sex of a patient. Moreover, we allow the physician to choose the specific abnormal events that need to be monitored. During this phase, the wearable device continuously stores the ECG recordings and periodically forwards them to the nearby edge device (see Fig. 1, exchange 1). We exploit the processing capabilities of the edge device (i.e., the smartphone) to analyze the received signal and forward it to the cloud services along with computerized annotations (see Fig. 1, exchange 2). The authorized physician may view reports, search traces and examine ECG alerts aggregated on the patient's health records remotely after they are synced with the cloud platform. The physician goes through the annotated recordings and validates or rejects the computerized interpretations depending on her expert assessment.

When an adequate number of normal and abnormal sessions have been identified by the authorized physician, the system is now ready to enter the *training phase*. The human-curated annotations are forwarded to the edge device (see Fig. 1, exchange 4) where the corresponding sessions are extracted from the local storage, are analyzed to extract a carefully selected set of features. The features of each session along with the annotations constitute the training vectors for the pattern matching engine of the wearable device. The wearable device receives the training set (see Fig. 1, exchange 5) and uses it to train the pattern matching engine. When the training completes, the wearable device is ready to enter the *detection phase*.

During the *detection phase* the signals collected from the ECG leads are analyzed and the features are extracted using the local processor. The resulting vector is passed to the pattern matching engine for classification. In case an abnormal event is detected, the wearable device is in a position to immediately notify the authorized physician via a short message exchange including only the alert type (see Fig. 1, exchange 6). As soon as the patient visits the physician at her infirmary, or while he goes through an emergency submission to the hospital, the complete ECG recordings are relayed to the nearby edge device (see Fig. 1, exchange 7). The ECG recordings are finally uploaded on the cloud platform (see Fig. 1, exchange

8) where they become available to the authorized physician for examination and assessment.

The above-described cycle of learning, training, and detecting is repeated periodically to re-evaluate the operation of the pattern matching of the wearable device and fine-tune its performance.

A wearable device that is used for diagnosing and monitoring heart diseases needs to be capable of collecting high-quality ECG traces. Existing diagnostic algorithms typically require a sampling rate of 500Hz and a resolution of 16 bit with $1\mu V$ per lead of the measured value. The input values can, however, be acceptable (after being properly adjusted), if the input sampling rate is at least 320Hz and the real resolution is better than $3\mu V$ per digit.

During the *learning phase*, the wearable device is continuously storing the ECG traces on the internal memory. During this off-line monitoring period (also referred to HOLTERR mode), the device is expected to store a recorded session of 24 hours. For a 12-channel recording produced by sampling 8 leads at 500Hz and a resolution of 16 bit, 24 hours would require about 3.2Gbit memory space (using a standard EDF+ format). Dropping the sampling rate to 320Hz would reduce the requirements to 2.6Gbit memory capacity.

During the *detection phase* the limiting factors of the operation of the wearable devices are the standby time achieved. The standby time is directly correlated with (a) the battery performance (while transmitting data and conducting computations), and (b) the memory limitations (while storing the collected signals).

4 PROTOTYPE DESIGN

The fundamental components of the HEART Platform prototype can be categorized into three different groups, mostly due to their physical location and functionality namely (i) the Wearable Device, (ii) the Mobile Application (executed on the edge device, e.g., the user's smartphone) and (iii) the cloud platform. In this work, we focus on the Fog Computing elements introduced mainly on the Wearable Device.

The wearable device software is comprised of several components that (a) interface with the ADC in order to acquire samples

from the ECG leads and the battery analogue input, in predefined intervals as specified by the sampling rate configuration setting, (b) handle the persistence of the configuration settings, file system parameters and structures, as well as the encryption, storage and retrieval of ECG recorded sessions and ECG metadata, (c) provide the routines for handling each of the valid commands that are sent remotely from the edge device, (d) carry out communication with the edge devices and the cloud platform based on BLE, (e) analyse the recorded signal and (f) produce alerts that are transmitted directly to the cloud platform based on LORA. Data stored within the wearable device and message exchanges across the different layers are always performed in a secure and encrypted manner making sure that the confidentiality of the user data is always respected. However, due to the limited space, we will not provide detailed information on these aspects. In this work, we focus on the analysis of the recorded signal and the person-centering training of the analysis module to identify with high-accuracy and low-energy consumption abnormal events.

The wearable device software is based on the Intel® Curie™ module developed using the Zephyr [7], a small real-time operating system for connected, resource-constrained devices supporting multiple architectures and released under the Apache License 2.0. We also relied on the Intel® Curie™ ODK¹ that includes the software, tools and documentation for developers to build boards.

The main concept for our ECG analysis algorithm is that the ECG recordings are analyzed, and the values for a set of characteristics like the heart rate, time intervals for the PQRST waveform, and amplitude are mapped to a predefined set of categories. Our ECG analysis module works in collaboration with the Pattern Matching Engine (PME) provided by the Intel® Curie™ module. The module segments the ECG traces collected from the sensors and identifies certain points of interest within the particular segment (i.e., extract features) that are used to detect possible abnormal events. The PME provides two methods for the classification of events based on the KNN (K-Nearest Neighbor) algorithm and the RBF (Radial Basis Function) algorithm. The ECG analysis module is comprised of two submodules:

Feature Extraction Module – The first submodule is the feature extraction module, where the ECG sampled data are processed and a set of descriptive features is extracted. The submodule determines the onset and offset of the *P* wave, onset and offset of *RST* and offset of the *T* wave. Recognition of onset and offset points are most important for all ECG measurements. All these intervals and wave durations are diagnostically relevant. Further on, the wave recognition points serve also as references for all amplitude measurements.

Training & Classification Module – The second submodule is the training and classification module, where the PME is first trained with the extracted features from a portion of the ECG recording and after this training stage is complete, the extracted features from the rest of the ECG recording is submitted for classification.

During the training phase, when the training vectors are transmitted to the wearable device, the mappings are submitted to the

Pattern Matching Engine to be stored as neurons. The PME offers 128 neurons which can contain up to 128 8-bit feature values (dimensions). The Intel® Curie™ module provides hardware implementations of the RBF and the KNN algorithms. We used the RBF algorithm for the purposes of the training phase since this allows for assigning more than one neurons per output category, whereas the KNN algorithm would only assign one neuron.

We implemented a window of 10 seconds based on which we extract certain feature (see Sec. 5 for more details) in order to construct the vector that will be classified based on the categories defined by the physician. Assuming (3) predefined categories: *normal*, *abnormal* and *noise*, a segment is assigned to the normal category if all beats within it are annotated as normal. An abnormal segment is one that contains at least one abnormal (non-normal) beat and no noise annotations. Finally, a noise segment is one that contains at least one noise annotation.

5 PROTOTYPE EVALUATION

The goal of the experimentation is to provide evidence that the wearable prototype will be capable of detecting, analyzing and interpreting a wide number of ECG data types and traces on board. The study examines the components developed on the wearable device and measures the performance of the components as individual/stand-alone parts and also in combination with the other parts of the system. The evaluation environment is setup in such a way in order to better reflect the usage scenario described in Sec. 3.

An Arduino/Genuino 101 development board² is used in order to conduct the first set of experiments based on the Intel® Curie™. The Genuino 101 offers an on-board 6 channel analog-to-digital (A/D) converter, which has a 10-bit resolution. The Arduino/Genuino 101 development board is constantly connected to the power supply. A USB connector is used for programming the microcontroller. The Intel® Curie™ ODK Z tree firmware³ is used that targets the Arduino 101 as host core.

The evaluation uses input data the recordings publicly available from the MIT-BIH arrhythmia database [12] that contains 48 half-hour excerpts of two-channel ambulatory ECG recordings, obtained from 47 subjects. Twenty-three recordings were chosen at random from a set of 4000 24-hour ambulatory ECG recordings collected from a mixed population of inpatients (about 60%) and outpatients (about 40%); the remaining 25 recordings were selected from the same set to include less common but clinically significant arrhythmias that would not be well-represented in a small random sample. Remark that there are 21 ECG recordings where there are no abnormal vectors were detected or with only one. It was decided to exclude these files as the pattern matching algorithm would perform poorly since there were not enough or any patterns for abnormal segments to match with.

The recordings were digitized at 360 samples per second per channel with 11-bit resolution over a 10mV range. The samples include two leads – the modified limb lead-II and one of the modified lead V_1 , V_2 , V_4 or V_5 with a resolution of 5 *lV/bit*. Two or more cardiologists independently annotated each record; disagreements were resolved to obtain the computer-readable reference annotations for

¹<https://software.intel.com/en-us/node/674972>

²<https://store.arduino.cc/genuino-101>

³<https://github.com/01org/CODK-Z>

each beat (approximately 110,000 annotations in all) included in the database. The size of these recordings, at about 1.9MBs, is ideal for our experiments, since the Arduino 101 has a persistent storage NAND Flash chip with a total capacity of 2MBs.

A lightweight feature extraction module is implemented that detects features such as the R locations, R-peak amplitude, R-R intervals, projected heartbeat, S locations and R-S amplitude. A subset of the extracted features of this algorithm is aggregated over the period of the corresponding segment providing the average values for R-peak amplitude, R-R interval, projected heartbeat, and R-S amplitude. Additionally, the count of R-peaks within a segment is also considered as a feature for the ECG segment. In total nine features are extracted for each segment and stored in the feature vector: the count of R-peak, average R-peak amplitude, average R-R interval, average projected heartbeat, last five R-peak amplitude.

In this study, three classification categories are used: normal, abnormal and noise. The initial 65% of each ECG recording is used for the training purposes: it is split it into segments and run the feature extraction module for each sample. The training segments are submitted to the Pattern Matching Engine of the Intel[®] Curie[™] module to be stored as neurons using the RBF algorithm. Then, the remaining 35% of each ECG recording is used to evaluate the accuracy of the classification. Again, it is split into segments and executed the feature extraction module, however now, the resulting feature vector is submitted to the (trained) Pattern Matching Engine and the output is the decision for the classification to one of the predefined categories. Since each trace includes 30 minute ECG recordings which are parsed in segments of ten seconds each, there is a total of 180 segments to classify per file. For these experiments, we have used 65% of these segments, i.e., 117 segments, for the training phase.

Tab. 1 presents the results achieved. The first column refers to the ECG recording used as input. The second and third columns contain the number of actual total beats, and the number of beats that were missed by the feature extraction algorithm. The following three columns contain the allocation of the segments into their actual categories based on the annotation files provided by the MIT-BIH database and the processing of those in the first step of our methodology. The "Neur" column shows the number of PME neurons trained during the training phase in RBF mode. Finally, the TN column shows the percentage of classified normal segments, and the TP column the classified abnormal or noise segments. These results are with the KNN classification algorithm. Arguably, the values for the TP percentages are the most significant, since they indicate the success rate in correctly classifying segments as abnormal. We observe that in a large number of the ECG recordings that percentages for TP are quite high (80% or higher) with only two files that achieve a lower TP percentage.

The experimental study indicates that there is a strong correlation between the feature extraction and the pattern matching modules. The extracted features should be specifically chosen so as to help provide an indication of an abnormal event in the ECG data and make a decision about whether an ECG segment is normal or not. It is also preferable that the set of these features should be as extended as possible since that means a better description and representation of the corresponding ECG segment. Some features

Table 1: Results achieved by basic feature extraction

#	Beats	Miss	Norm	Abn	Noise	Neur	TN	TP
100	2273	5	149	31	0	26	61.1	80.6
102	2187	615	5	175	0	14	100	100
104	2229	405	1	165	14	63	100	100
105	2572	591	102	34	44	70	66.7	91
106	2027	579	82	98	0	114	100	80.6
108	1774	331	151	19	10	99	96.7	69
114	1879	420	136	39	5	116	90.4	90.9
116	2412	193	93	83	4	56	60.2	96.6
119	1987	819	21	159	0	90	95.2	93.7
200	2601	295	11	157	12	109	100	95.9
201	2000	464	63	115	2	112	84.1	90.6
202	2136	237	145	33	2	57	78.6	91.4
203	2980	1267	16	133	31	117	62.5	96.3
208	2955	332	1	166	13	97	0	100
209	3005	68	71	104	5	79	76.1	83.5
210	2650	584	69	106	5	115	89.9	86.5
212	2748	248	47	129	4	46	85.1	68.4
213	3251	264	56	124	0	58	64.3	78.2
215	3363	682	72	106	2	64	77.8	95.4
219	2287	459	125	55	0	116	91.2	76.4
220	2048	20	137	43	0	35	75.9	90.7
221	2427	324	17	162	1	115	17.6	96.9
222	2483	1324	95	79	6	106	87.4	81.2
223	2605	704	63	116	1	90	79.4	82.1
228	2053	661	30	129	21	117	86.7	98
231	1573	132	47	133	0	55	97.9	95.5

of the ECG signal are easy to detect, namely the R peaks locations and the S location (typically the local lowest point). However, existing algorithms are not 100% accurate in the detection of these features, due to noise or the ECG abnormality conditions themselves introducing artifacts and atypical behavior in the ECG signal. Increasing the accuracy of these algorithms requires implementing extra filters for noise and amplifying the ECG signal, as well as allocating large buffers for keeping a history window of ECG data and optionally executing the algorithm in multiple passes. These increase the complexity of the algorithm and its requirements in program memory resources (RAM) as well as offline storage of the data (for the case of multiple pass algorithms). Additionally, multi-pass algorithms are not suitable for the paradigm of real-time detection of events in an ECG data stream.

During the tests performed, the power consumption of the wearable device was measured in relation to the Pattern Matching Engine. These data allow to identify the maximum energy requirements of the wearable device, those cases where the battery will be drained faster, as well as the best configuration for the operation of the device with the lowest energy consumption.

Learning Phase – Initially the wearable device is recording the sensor data received from the ADC. The data are stored in chunks of 2048 bytes after being encrypted. In this case, the PME core is idle and the BLE radio is deactivated. The power consumption is at 42.61 mA. When the device has collected

enough data (e.g., after 24/48 hours) the data are streamed over the BLE connection to a specific edge device. Since the data forwarded to the BLE radio are already encrypted, no further processing is required. The PME core remains reactivated however the BLE radio is activated. The power consumption is at 40.51 mA.

Training Phase – The training data set has already received a training set through the BLE radio and is stored locally. The data are passed to the PME core that is in training mode. The power consumption is at 48.26 mA.

Detection Phase – Sensor data are continuously received from the ADC and are processed for feature extraction. We measure the power consumption of the feature extraction module (by deactivating the classification methods of the PME core) at 49.25 mA. In the sequel, the generated vectors are classified using the KNN method provided by the PME core. We now measure the power consumption of the PME core (by deactivating the feature extraction module) at 53.87 mA.

Based on the above analysis we understand that the power requirements of the wearable device can be accommodated by an off-the-shelf reasonably sized battery that will sustain its continuous operation for 24 hours or 48 hours, a typical duration existing HOLTER products.

6 CONCLUSIONS AND FUTURE WORK

This paper has presented some of the basic issues a system architect must consider when designing, implementing and deploying a person-centric end-to-end healthcare application which includes IoT nodes and cloud computing backend services, that leverages the benefits of the Fog computing approach. Being an intermediate layer between end-user devices and the remote cloud datacenters, Fog alleviates a series of issues in the areas of bandwidth consumption reduction, latency decrease, and seamless operation. As shown by the evaluation of a functional end-to-end application prototype designed by the latest trends of IoT and Fog computing, presented in Section 5, analysis of high-quality recordings based on the computational resources of the wearable devices, without the support of a nearby edge/mobile device, and in particular using only the PME core of the Intel® Curie™ module, can potentially lead to high levels of accuracy. Our findings indicate that it is possible to totally avoid storing and transmitting sensor readings with no significant value to domain experts. We are confident that our approach can help address the actual problems derived from the fundamentals of IoT applications and in particular address the problems of colossal data transfer between the network nodes.

There is a lot of future work required to study this new computing paradigm and apply it in different application domains. For the case studied here, we wish to look into alternative implementations of the feature extraction module in order to further increase accuracy. There are a number of parameters and subroutines that can be tweaked to change the feature extraction algorithm behavior, as well as the pattern matching engine behavior. These include the selection and number of the extracted features, the duration of the ECG segments where the feature extraction (and the PME classification) is applied, the selection of (more than three) categories for the PME, the process of lowering the resolution of the extracted feature

values to 8-bit values for the training vectors, the percentage of the recording that is used for training the PME module. Additionally, although current bibliography mostly focuses on QRS detection (and hence, heart-rate calculation) there are also other approaches that we wish to investigate and implement on a constrained resources device. Finally, we also wish to evaluate our device using different data sets in order to identify specific hard problem instances.

REFERENCES

- [1] I. Chatzigiannakis, A. Kinalis, and S. Nikolettseas. 2005. Power conservation schemes for energy efficient data propagation in heterogeneous wireless sensor networks. *Proceedings - Simulation Symposium* (2005), 60–71. <https://doi.org/10.1109/ANSS.2005.37>
- [2] Ioannis Chatzigiannakis, Emil Stoyanov Valchinov, Athanasios Antoniou, Athanasios P. Kalogeras, Christos E. Alexakos, and Panagiotis Konstantinopoulos. 2015. Advanced observation and telemetry heart system utilizing wearable ECG device and a Cloud platform. In *2015 IEEE Symposium on Computers and Communication, ISCC 2015, Larnaca, Cyprus, July 6-9, 2015*. IEEE Computer Society, 25–30. <https://doi.org/10.1109/ISCC.2015.7405449>
- [3] Y. Chen, W. Shen, H. Huo, and Y. Xu. 2010. A Smart Gateway for Health Care System Using Wireless Sensor Network. In *SENSORCOMM 2010*. 545–550. <https://doi.org/10.1109/SENSORCOMM.2010.88>
- [4] Cisco Systems Inc. 2017. Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are. http://www.cisco.com/c/dam/en_us/solutions/trends/iot/docs/computing-overview.pdf. (2017).
- [5] Intel Corporation. 2017. Intel® Curie™ Module. <https://www.intel.com/content/dam/support/us/en/documents/boardsandkits/curie/intel-curie-module-datasheet.pdf>. (2017).
- [6] European Heart Network. 2017. European Cardiovascular Disease Statistics. <http://www.ehnheart.org/component/downloads/downloads/2452>. (2017). [Online].
- [7] Linux Foundation. 2017. Zephyr Project. <https://www.zephyrproject.org/>. (2017).
- [8] T. N. Gia et al. 2015. Fog Computing in Healthcare Internet of Things: A Case Study on ECG Feature Extraction. In *2015 IEEE CIT/IUCC/DASC/PICOM*. 356–363. <https://doi.org/10.1109/CIT/IUCC/DASC/PICOM.2015.51>
- [9] Periyasamy M. Mariappan, Dhanasekaran R. Raghavan, Shady H.E. Abdel Aleem, and Ahmed F. Zobaa. 2016. Effects of electromagnetic interference on the functional usage of medical equipment by 2G/3G/4G cellular phones: A review. *Journal of Advanced Research* 7, 5 (2016), 727 – 738.
- [10] Fen Miao, Yayu Cheng, Yi He, Qingyun He, and Ye Li. 2015. A Wearable Context-Aware ECG Monitoring System Integrated with Built-in Kinematic Sensors of the Smartphone. *Sensors* 15, 5 (2015), 11465–11484.
- [11] Subasish Mohapatra and K Smruti Rekha. 2012. Sensor-cloud: a hybrid framework for remote patient monitoring. *Intl. Journal of Computer Applications* 55, 2 (2012).
- [12] George B Moody and Roger G Mark. 1992. MIT-BIH Arrhythmia Database. (1992).
- [13] E. Nemati, M. J. Deen, and T. Mondal. 2012. A wireless wearable ECG sensor for long-term applications. *Communications Magazine, IEEE* 50, 1 (Jan. 2012), 36–43. <https://doi.org/10.1109/mcom.2012.6122530>
- [14] Pranav Rajpurkar, Awni Y. Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y. Ng. 2017. Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks. *CoRR* abs/1707.01836 (2017).
- [15] Indu Saini, Dilbag Singh, and Arun Khosla. 2013. QRS detection using K-Nearest Neighbor algorithm (KNN) and evaluation on standard ECG databases. *Journal of Advanced Research* 4, 4 (2013), 331 – 344. <https://doi.org/10.1016/j.jare.2012.05.007>
- [16] NXP Semiconductors. 2016. MPC8572E PowerQUICC III Integrated Processor. <http://www.nxp.com/docs/en/data-sheet/MPC8572EEC.pdf>. (2016).
- [17] Atman P. Shah and Stanley A. Rubin. 2007. Errors in the computerized electrocardiogram interpretation of cardiac rhythm. *Journal of Electrocardiology* 40, 5 (2007), 385 – 390. <https://doi.org/10.1016/j.jelectrocard.2007.03.008>
- [18] E. Theodoridis, G. Mylonas, and I. Chatzigiannakis. 2013. Developing an IoT Smart City framework. *IISA 2013 - 4th International Conference on Information, Intelligence, Systems and Applications* (2013), 180–185. <https://doi.org/10.1109/IISA.2013.6623710>
- [19] Christos Tselios and George Tsolis. 2016. On QoE-awareness through Virtualized Probes in 5G Networks. In *Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), 2016 IEEE 21st International Workshop on*. 1–5.
- [20] Sungwon Yang and M. Gerla. 2011. Personal gateway in mobile health monitoring. In *2011 IEEE PERCOM Workshops*. 636–641. <https://doi.org/10.1109/PERCOMW.2011.5766966>
- [21] Ding Yi, Fan Binwen, Kong Xiaoming, and Ma Qianqian. 2016. Design and implementation of mobile health monitoring system based on MQTT protocol. In *2016 IEEE Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*. 1679–1682. <https://doi.org/10.1109/IMCEC.2016.7867503>