IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received September 5, 2018, accepted October 9, 2018, date of publication October 23, 2018, date of current version December 7, 2018.

Digital Object Identifier 10.1109/ACCESS.2018.2877679

ECG Signal Reconstruction on the IoT-Gateway and Efficacy of Compressive Sensing Under Real-Time Constraints

MOHAMMED AL DISI[1], HAMZA DJELOUAT[1], CHRISTOS KOTRONIS[2], ELENA POLITIS[2], ABBES AMIRA[1], (Senior Member, IEEE), FAYCAL BENSAALI[3], (Senior Member, IEEE), GEORGE DIMITRAKOPOULOS[2], AND GUILLAUME ALINIER[4,5]

[1]Department of Computer Science and Engineering, College of Engineering, Qatar University, Doha 2713, Qatar
[2]Department of Informatics and Telematics, Harokopio University of Athens, 17676 Athens, Greece
[3]Department of Electrical Engineering, College of Engineering, Qatar University, Doha 2713, Qatar
[4]Ambulance Service, Hamad Medical Corporation, Doha 3050, Qatar
[5]School of Health and Social Work, University of Hertfordshire, Hatfield AL109AB, U.K.

Corresponding author: Abbes Amira (abbes.amira@qu.edu.qa)

This work was supported by the National Priorities Research Program (NPRP) from the Qatar National Research Fund (a member of Qatar Foundation) under Grant 9-114-2-055. The statements made herein are solely the responsibility of the authors. The publication of this article was funded by the Qatar National Library.

ABSTRACT Remote health monitoring is becoming indispensable, though, Internet of Things (IoTs)-based solutions have many implementation challenges, including energy consumption at the sensing node, and delay and instability due to cloud computing. Compressive sensing (CS) has been explored as a method to extend the battery lifetime of medical wearable devices. However, it is usually associated with computational complexity at the decoding end, increasing the latency of the system. Meanwhile, mobile processors are becoming computationally stronger and more efficient. Heterogeneous multicore platforms (HMPs) offer a local processing solution that can alleviate the limitations of remote signal processing. This paper demonstrates the real-time performance of compressed ECG reconstruction on ARM's big.LITTLE HMP and the advantages they provide as the primary processing unit of the IoT architecture. It also investigates the efficacy of CS in minimizing power consumption of a wearable device under real-time and hardware constraints. Results show that both the orthogonal matching pursuit and subspace pursuit reconstruction algorithms can be executed on the platform in real time and yield optimum performance on a single A15 core at minimum frequency. The CS extends the battery life of wearable medical devices up to 15.4% considering ECGs suitable for wellness applications and up to 6.6% for clinical grade ECGs. Energy consumption at the gateway is largely due to an active internet connection; hence, processing the signals locally both mitigates system's latency and improves gateway's battery life. Many remote health solutions can benefit from an architecture centered around the use of HMPs, a step toward better remote health monitoring systems.

INDEX TERMS Connected health, compressed sensing, energy efficiency, heterogeneous multicore platforms, internet of things, mobile real-time health monitoring, multicore processing, remote monitoring, wearable sensors.

I. INTRODUCTION

An internet of things (IoT) based remote health monitoring system is becoming less of a luxury and more of a normal commodity. Due to the growing older adult population and unhealthy lifestyle trends, there is a recent rise in chronic diseases, escalating the demand for continuous clinical supervision and consequently amplifying healthcare management costs [1]. Adopting a connected health monitoring solution offers the ability to keep an eye on patients and provide on-demand or automated assistance in case of emergencies—while keeping costs in check [2]. Coupled with seamless assimilation, it allows the patient to live normally without regular and frequent hospital visits, improving their quality of life.

Unfortunately, a real-time mobile health monitoring system has many design challenges such as energy consumption of wearable medical devices [3]. Attempts to extend their battery life lead to an interest in compressive sensing (CS) within the context of physiological signals. CS is a signal acquisition and compression scheme capable of digitizing signals using fewer measurements than that required by the Shannon-Nyquist theorem [4]. Considering electrocardiogram (ECG) signals, CS has been shown to be advantageous over state-of-the-art wavelet-based compression schemes due to its encoder simplicity [5]–[7]. However at the decoder end, it requires a signal recovery process that is orders of magnitude more complex than traditional compression [8]. Thus, a CS implementation addresses energy consumption at one end but accumulates it onto another, in addition to exacerbating the latecy associated with cloud computing [9]. Even though CS enhances privacy and security of transmissions by bundling in free-of-charge encryption with compression [10]–[13], it does not address other challenges associated with cloud-based systems such as instability and high energy consumption at the gateway due to the continuous stream of data and reliance on internet connectivity, and big data management [9], [14].

Meanwhile, computing capacity of mobile processors is on an exponential rise and is simultaneously becoming more energy efficient by utilizing processing heterogeneity. Heterogeneous multicore platforms (HMPs) combine slow power-efficient and fast power-hungry processors and dynamically allocates tasks between them to optimize for either computational resources or energy efficiency. Hence, a remote health monitoring architecture based on using HMPs as both the system gateway and main processing unit could effectively address many limitations. Real-time signal processing and assessment of the patient's health is performed locally, emergency services are automatically alerted at the detection of adverse events, and information about the patient health can be uploaded at fixed intervals or on-demand.

In the development of our real-time health monitoring system, we implement and analyze implementations of different data processing frameworks on HMPs. This paper investigates (a) real-time CS ECG signal reconstruction on the IoT gateway and (b) whether CS is still capable of reducing energy consumption on wearable sensors under real-time and hardware constraints. It also compares the gateway from the conventional viewpoint of a simple router against the utilization as the primary processing unit of the IoT architecture.

The reminder of the paper is organized as follows. Related work and the main contributions of the paper are presented in section II. Section III gives a brief background about CS and relevant reconstruction algorithms. An overview of our gateway-centric connected health system is given in section IV. Sections V and VI discuss the employed performance metrics and obtained results respectively. Finally, the paper is concluded in section VII.

II. RELATED WORK

IoT-based remote monitoring systems consist of miniaturized and networked sensors, continuously measuring and transmitting signals of interest to a nearby gateway through low-power wireless communications [15]. The gateway—typically a mobile phone—routes the data to a remote IoT platform where it is analyzed and then visualized by the end user. [15].

As the technology for collecting, analyzing and transmitting data in the IoT continues to grow and evolve, more IoT-driven healthcare applications, services and systems emerge. In [16], the need of an integration of IoT technologies (e.g. wearable devices) and e-Health solutions was addressed with focus on an integrated system for the continuous monitoring of students at risk of high blood pressure as well as quick treatment and consultation from medical experts at a distance. Riazul Islam *et al.* analyzed a variety of medical IoT applications, such as remote health monitoring, fitness programs, and elderly care [17]. The authors in [18] experimented with portable devices and different communications protocols and models for the creation of e-Health applications, like CardioNet. Moreover, a tested real-time monitoring platform that uses IoT-gateways and medical devices showed the effectiveness real-time e-Health services and introduced its need of edge computing [19]. In a similar context, the exploitation of fog computing in healthcare IoT systems is proposed in [20], where an implementation of a Smart e-Health Gateway (UT-GATE) suitable for the deployment in clinical environments was demonstrated.

Mobile devices have been recently levered for more than data routing. In the simplest form, they are used to provide feedback to patients by displaying remotely extracted information through web-based applications [21]. The authors in [22] extended their system's gateway with a radio frequency identification (RFID) reader to detect tagged patients and extracted eight ECG features locally, reducing server-side computational costs. Also, Gradl *et al.* [23] and Oresko *et al.* [24] demonstrated real-time ECG QRS complex detection, feature extraction, and heart beat classification on an Android and Windows-based mobile phones respectively.

The CS scheme has been thoroughly studied for ECG signals. Dixon *et al.* experimented with their compressibility with varying sensing matrix designs and reconstruction algorithms, demonstrating ECG dimensionality reduction up to 16 times [25]. In [26], Zigei *et al.* established a clinically acceptable tolerance for distortion in ECG signals. With that tolerance in mind, the authors in [27] and [6], [5], and [28] reported 2.5, 3.44, and 5 compression factors respectively. Additionally, there is also continuous development of CS reconstruction algorithms to reduce error-rates or address special applications in remote monitoring. Zhang *et al.* recently proposed a reconstruction algorithm that capitalizes on the signals' structure and its intra-block correlation, showing its accuracy advantages over popular greedy algorithms, and its use for fetal ECG monitoring [29]–[31]. Although, in the previous studies, signal recovery from compressed vectors is done offline and on workstation computers using time-complex and accurate algorithms. This allows for better compression performance but prohibits real-time applications.

There are three studies closely related to this work. First, Mamaghanian *et al*. demonstrated the superiority of CS over wavelet-based compression on the Shimmer™ wearable device [5]. They also report that CS extends the sensor battery life by 9.7% compared to no-compression alternatives. However, CS reconstruction is performed offline by the SPGL1 solver [32], [33]. Furthermore, the analysis uses MIT ECG records, fed into the sensors through a wire, not signals collected from the wearable device itself which are typically more noisy and subsequently less compressible. Second, the authors in [34] showcased real-time reconstruction of ECG signals on an iPhone using a modified version of the iterative shrinkage-thresholding algorithm (ISTA), but do not consider trade-offs between execution time, power consumption, allocated computational resources, and signal dimensions. Those limitations were addressed in [35] and [36] where a real-time, single-thread implementation of the orthogonal matching pursuit (OMP) and focal underdetermined system solver (FOCUSS) is thoroughly studied on ARM's big.LITTLE™ HMP. Additionally, real-time energy-efficient reconstruction has been demonstrated on an ARM's Cortex-M4F microcontroller [37]. In [35]–[37], compression is carried out using a custom-designed rakeness-based sensing matrix. This approach exploits the uneven distribution of information in sparse signals to further reduce the number of samples needed for faithful acquisition. Consequently, the computational complexity of reconstruction is lowered in comparison to traditional sensing matrices [37].

In light of the aforementioned works, this article focuses on offering a practical viewpoint of CS under real-time and hardware constraints. It also signifies the advantages of local processing on the gateway and how it might be able to enhance connected health. The main contributions of this paper are as follows:

- **At the level of the signal:** ECG compression performance is assessed using signals obtained from the wearable device.
- **At the level of the wearable device:** a more recent Shimmer™ wearable device is used. Signal acquisition, CS real-time constraints, signal window sizes, reconstruction accuracy, and transmission of non-compressed delayed real-time signals are factored in the energy analysis. Finally, an improved compression implementation that reduces processing time is presented.
- **At the level of the gateway:** the subspace pursuit (SP) reconstruction algorithm is implemented and compared to OMP including their multithreaded performance and energy consumption trade-offs. By simulating a real-world remote monitoring scenario, we demonstrate the energy saving advantages and feasibility of a gateway-processing for connected heath architectures.

III. COMPRESSIVE SENSING BACKGROUND

The CS framework can capture and recover sparse signals using fewer measurements than that required by the Nyquist rate [25]. A signal $\boldsymbol{x} \in \mathbb{R}^n$ is sparse if it can be presented as:

$$\boldsymbol{x} = \sum_{i=1}^{n} \boldsymbol{\alpha}_i \boldsymbol{\psi}_i = \boldsymbol{\psi}\boldsymbol{\alpha} \tag{1}$$

where $\boldsymbol{\alpha} \in \mathbb{R}^n$ is a column vector with only few non-zero elements, and $\boldsymbol{\psi} \in \mathbb{R}^{n\times n}$ is a matrix containing a domain's coefficients (Fourier, wavelet, etc.). ECG signals have sparse representations in multiple domains and hence can be compressively acquired using CS.

In practice, CS can be done directly on analog signals as an alternative to conventional data acquisition systems or applied to signals sampled at the Nyquist rate for dimensionality reduction (i.e. digital CS). This study belongs to the latter where the compressed signal is expressed as follows:

$$\boldsymbol{y} = \boldsymbol{\varphi}\boldsymbol{x} = \boldsymbol{\varphi}\boldsymbol{\psi}\boldsymbol{\alpha} = \boldsymbol{A}\boldsymbol{\alpha} \tag{2}$$

where $\boldsymbol{y} \in \mathbb{R}^m$ is the compressed signal, $\boldsymbol{\varphi} \in \mathbb{R}^{m\times n}$ is a sensing matrix with $m << n$, and $\boldsymbol{A} = \boldsymbol{\varphi}\boldsymbol{\psi} \in \mathbb{R}^{m\times n}$ is a matrix linking the sparse representation of $\boldsymbol{x}$ with the compressed signal $\boldsymbol{y}$.

To reconstruct signals faithfully, the sensing matrix must meet the restricted isometry property (RIP) and be incoherent with the sparsfying basis [38]. Those requirements make their design non-trivial, but fortunately, it has been shown that a sensing matrix populated with independent identically distributed (i.i.d) values such as those obtained for a Gaussian or a Bernoulli distribution satisfy the two with overwhelming probability [38].

Recovering the sparse signal $\boldsymbol{\alpha}$ from the compressed signal $\boldsymbol{y}$ requires solving an under-determined system which has infinite solutions [4]. But since $\boldsymbol{\alpha}$ is known to be sparse, a solution can be obtained by searching for the sparsest version of $\boldsymbol{\alpha}$ that satisfies (2). This can be achieved with relaxed convex optimization such as [39]:

$$\hat{\boldsymbol{\alpha}} = \arg\min \ \|\boldsymbol{\alpha}\|_1 \quad \text{subject to } \|\boldsymbol{A}\boldsymbol{\alpha} - y\|_2 < \epsilon \tag{3}$$

where $\hat{\boldsymbol{\alpha}} \in \mathbb{R}^n$ is the recovered version of $\boldsymbol{\alpha}$, $\|\cdot\|_1 = \sum_{i=1}^{n} |\cdot_i|$ denotes the ℓ_1-norm, and ϵ relaxes the optimization problem. Finally, the signal can be reconstructed by computing $\hat{\boldsymbol{x}} = \boldsymbol{\psi}\hat{\boldsymbol{\alpha}}$.

Linear programing approaches can solve (3) in near polynomial time but that is still too slow for a real-time implementation, especially considering dimensionally large signals [4]. Greedy algorithms are a family of solvers that are typically fast and less complex making them preferred for a real-time systems. The caveat is however, that they provide suboptimal results. Two greedy algorithms are considered in this study: OMP and SP [40]

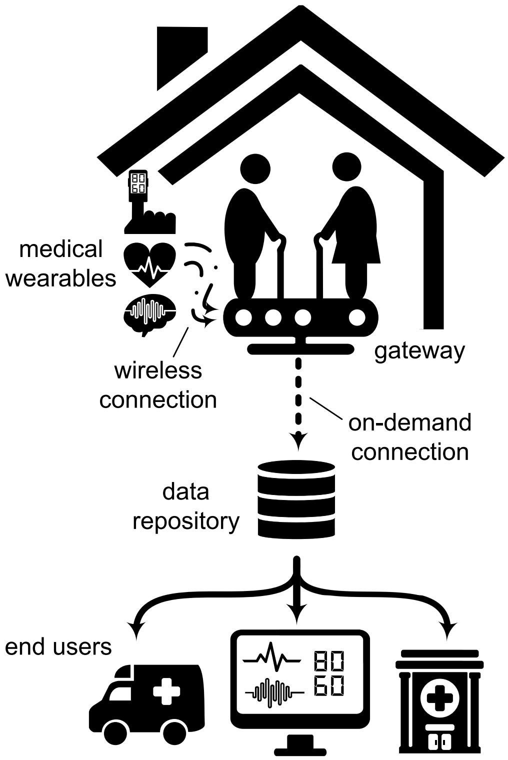

FIGURE 1. Overview of REMS.

The OMP mechanism selects one column of the sensing matrix that highly correlates with a residual at each iteration [41]. Selected columns are added to an active set, and removed from $\boldsymbol{\varphi}$ so they are not reselected [4]. Afterwards, a solution is calculated from the active set using a least square operation and is subsequently subtracted from the residual. The algorithm stops when residual ℓ_2-norm meets a user-defined criterion. SP selects multiple columns per iteration and does not remove them from $\boldsymbol{\varphi}$ [42]. The active set is updated at each iteration. Hence, suboptimal columns added in earlier iterations (that are proven to be not well correlated with the residual) can be removed later [42].

IV. SYSTEM OVERVIEW

The proposed remote elderly monitoring system (REMS) consists of three components as in conventional IoT architectures and is illustrated in Fig. 1. Wearable medical sensors, equipped with state-of-the-art power saving techniques (e.g. CS), transmit physical and physiological data to a nearby gateway. The gateway is the main component of this system. It is a HMP that is responsible for most of the processing and data treatment which includes CS signal reconstruction—as in the scope of this paper—and other data analysis algorithms such as heart beat detection and classification, fall detection, biometric identification, among others. As for the "remote" part of this architecture, the gateway stores day to day activities and the health status of the patient and transmits summarized reports in fixed intervals or by demand of the end user. This allows for trends of the various physical and/or physiological monitored parameters to be analyzed over time. Additionally, through its classification, it notifies emergency services and caregivers of any unfortunate events and then transmits physiological data in real-time for early diagnosis, and can host physician to patient video or audio calls for early and remote assistance. As mentioned earlier, this approach eliminates system latency and addresses stability concerns associated with cloud-computing and internet connectivity dependence. It also facilitates large scale deployment of connected health. When new subjects are introduced to the system, they are equipped with their own processing unit, and hence, there is no additional load onto the cloud-platform, which reduces the infrastructure costs of the system. Finally, HMPs provide a plethora of computational resource configurations, enabling customizability and optimizations on a case to case basis.

In this implementation, the wearable device is the Shimmer3™ ECG/EMG unit [43], [44]. Hardware-wise, it operates on a Texas Instruments MSP430F5437A—a 16-bit microcontroller running at 24 MHz with 16 kB of RAM and 256 kB of flash. The unit communicates through an RN-42 Bluetooth module, is powered by a 450 mAh battery, and weighs 31 grams. As for software, the manufacturer provides the firmware of the device written in C [45] which can be compiled using Texas Instruments Code Composer Studio (v7.4) [46], and written it into the Shimmer using its dock. REMS uses three ECG electrodes to acquire left arm (LA) - right arm (RA) lead signals sampled in 16-bits (0.074 mV resolution) and with a gain of four, at 360 samples per second.

The gateway is a HardKarnel's Odroid XU4 board [47] featuring ARM's big.LITTLE™ heterogeneous multicore solution (Fig. 2). Running it is Samsung's Exynos 5422 octa-core processor which houses a cluster of four Cortex-A15 cores (big) and a cluster of four Cortex-A7 cores (LITTLE). This architecture can be found in modern mobile phones (e.g. Samsung's Galaxy S5), and hence considered a good candidate for simulating real-world mobile connected health scenarios. There are no built-in wireless options and consequently USB Bluetooth and WiFi modules were used for connectivity. It operates on Ubuntu 16.04 (Linux 3.10 + armv7)—algorithms were implemented in C++ using the `Armadillo` (v8.3) library [48] for linear algebra and complied using `gcc` (v5.4).

V. PERFORMANCE METRICS

One aspect for comparison between OMP and SP is the compression ratio (CR = n/m) achieved for a particular reconstruction accuracy. Reconstruction accuracy is computed using the average reconstruction signal to noise ratio (ARSNR) defined as [37]:

$$\text{ARSNR} = \boldsymbol{E}_{\boldsymbol{x}}[(\frac{\|\boldsymbol{x}\|_2}{\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2})_{dB}] \tag{4}$$

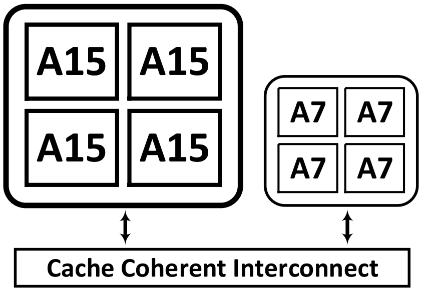

FIGURE 2. ARM's big.LITTLE heterogeneous multicore solution.

where $\hat{\boldsymbol{x}}$ is the reconstructed version of the original signal $\boldsymbol{x}$—with its dc component removed—and $\boldsymbol{E_x}$ indicates the averaging over all considered instances of $\boldsymbol{x}$.

Three ECG signal durations are considered: a half-second ($n = 180$), one-second ($n = 360$), and 1.42-seconds ($n = 512$)—capped due to limitations of the Shimmer. In REMS, transmission of real-time ECG windows to health-care services is limited to few cases, and hence, so is the need of medical grade signals. Furthermore, signal analysis algorithms (e.g. heart rate detection and beat classification) can accommodate higher levels of signal distortion while remaining effective. For those reasons, two levels of reconstruction accuracies are defined: high quality (HQ) and low quality (LQ). They were both classified by visual inspection—with Zigei *et al.* [26] clinical grade ECG distortion in mind—such that:

- **HQ:** ECG signals are reconstructed with minimal error. Suitable for accurate clinical diagnosis.
- **LQ:** ECG features can be clearly identified. Maintains precision for wellness applications.

Figure 3 depicts a Shimmer recorded ECG and its two operative points, which are defined as follows:

- **HQ:** Probability[RSNR $\geq$ 11.5] $\leq$ 97%
- **LQ:** Probability[RSNR $\geq$ 7.5] $\leq$ 97%

On the platform, the execution time of reconstruction is assessed by averaging an algorithm's execution time for 50 consecutive windows using Armadillo's built-in timer. Reconstruction is performed using the two core types (A15 big cores and A7 LITTLE cores), processor frequency ranges between $0.8 - 1.4$ GHz, and $1 - 4$ cores. For a fair comparison between the three ECG dimensions, the average reconstruction time is normalized over the real-time window (i.e. time gap between two consecutive windows) which equals the window duration. Reconstruction power consumption is measured by logging the readings of a digital wattmeter in series with the gateway's power source. After selecting a reconstruction algorithm, and setting the processor type, frequency, and number of cores, the idle power draw is calculated and subsequently subtracted from the active (i.e. reconstruction) power draw. It has been observed that window sizes do not affect the algorithms power consumption, so instead, we considered energy consumption per window, which is computed as average power consumption $\times$ reconstruction time.

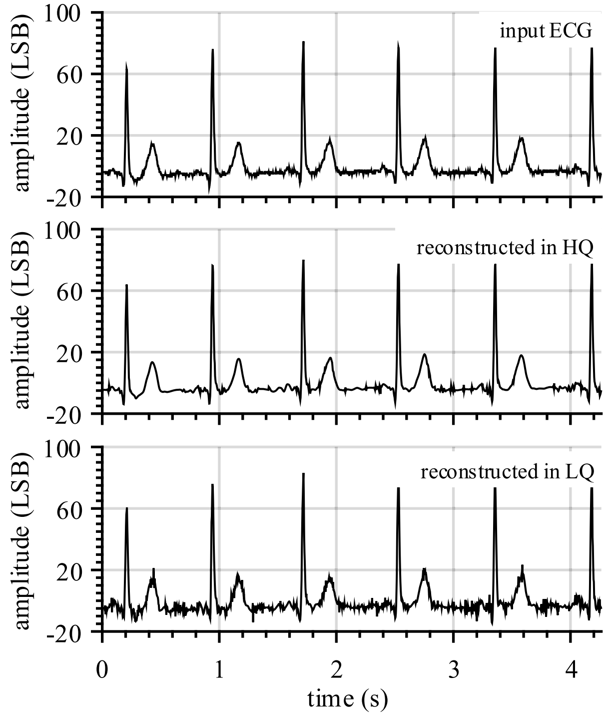

FIGURE 3. A Shimmer recorded ECG segment and its HQ and LQ reconstruction points computed with window size $n = 512$ and OMP.

TABLE 1. Highest CR that achieves the LQ and HQ targets with probability of 97% for SP and OMP at ECG window sizes (n).

	OMP			SP		
	$n = 180$	$n = 360$	$n = 512$	$n = 180$	$n = 360$	$n = 512$
HQ	1.04	1.48	1.56	**1.25**	**1.64**	**1.71**
LQ	**2.25**	**2.90**	**3.12**	1.91	2.60	2.60

On the Shimmer, the compression processing time is calculated from the difference between two timestamps, recorded before and after processing, sent in packets' header. Each packet also contains the battery voltage of the sensor which is subsequently converted into battery percentage at the receiver. After running the sensor for 10 minutes under particular settings (e.g. CS or no CS) while continuously recording the battery percentage, the data is linearly fitted, obtaining the percentage drop per packet. The experiment is repeated three times for each test case, and the average percentage drop per packet is scaled by the number of packets per hour and then multiplied by the battery capacity to estimate the Shimmer's power consumption in mAh.

VI. RESULTS

A. RECONSTRUCTION PERFORMANCE

A MATLAB simulation is performed to compare between SP and OMP by investigating the highest achievable CR with respect to the predefined qualities of service and ECG

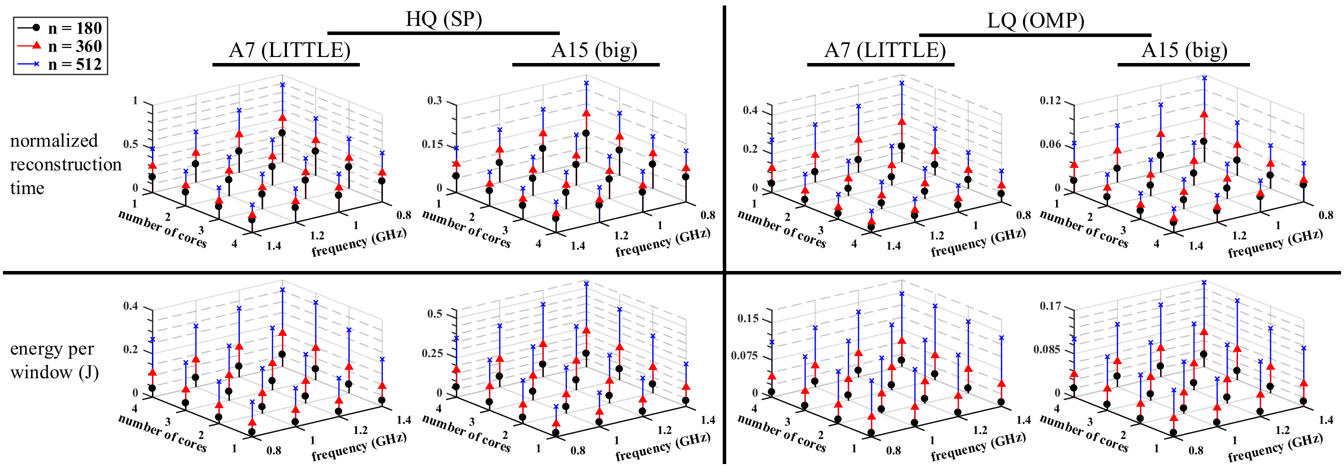

FIGURE 4. Per window normalized reconstruction time (top) and energy consumption (bottom) of HQ and LQ ECG reconstruction using SP (left) and OMP (right) with variable signal lengths, cores type, number, and frequency.

window sizes. The procedure involves computing ARSNR with a confidence interval of 97% from reconstructed ECG windows segmented from a 45-minute Shimmer ECG recording. Table 1 reports the obtained results and the data clearly indicate the superiority of SP for the HQ target while OMP remains more suited for the LQ operative point. Hence, for the remainder of this paper, HQ and LQ refers to ECG signals—compressed at their corresponding highest CR—reconstructed with SP and OMP respectively.

Larger ECG windows improve compressibility, by scaling the number of samples from $n = 180$ to $n = 512$, CR increases by nearly 28% with the majority gain ($\approx$ 23%) occurring at the first step (180 $\rightarrow$ 360). It is worth mentioning that, in comparison to the literature, the CRs are low, which we believe is caused by (i) the higher amounts of non-idealities in the 3-electrode Shimmer recording and (ii) the less accurate but real-time friendly greedy algorithms.

B. RECONSTRUCTION ON THE PLATFORM

Execution time of the algorithms, and the consequent energy required per window depends on many variables: the original signal length (n), the compressed signal length (m), and the type, number, and frequency of the cores assigned to the process. The performance of the platform and the effects of all the aforementioned variables are comprehensively summarized in Fig. 4. It should be noted that the time and energy plots have reversed x and y axes for easier viewing and normalized reconstruction time is plotted at the top row.

At first glance, it is clear that there is an exponential growth against the use of larger windows, and hence, potential benefits of higher compression on the sensing node are met with higher energy requirements at the gateway. Big (A15) cores expedite processing by an average of 1.7 to 4 times without affecting energy consumption significantly, which can be easily observed by the plot peaks. At a fixed minimum frequency and $n = 512$, HQ reconstruction is 3.25 times faster using a single A15 and only consumes 19% more energy than its single LITTLE core counterpart. In fact, for LQ, it speeds up the process by 3.52 and consumes 12% less energy as the speed gained is significant enough to compensate for the increased average power consumption of the big core.

OMP has many parallelizable instructions and benefits greatly from utilizing more cores. On the other hand, SP has a greater number of sequential instructors. This indicates that using higher CPU frequencies rivals or is better than higher core number. At a fixed frequency, on average, and for HQ and LQ respectively, going from 1×A7 to (2 $\rightarrow$ 3 $\rightarrow$ 4) provides (24% $\rightarrow$ 10% $\rightarrow$ 5%) and (60% $\rightarrow$ 25% $\rightarrow$ 9%) faster reconstruction while consuming (32% $\rightarrow$ 22% $\rightarrow$ 11%) and (4% $\rightarrow$ 7% $\rightarrow$ 8%) more energy. In a similar manner, increasing the core frequency from 0.8 to (1.0 $\rightarrow$ 1.2 $\rightarrow$ 1.4) GHz speeds up processing by (23% $\rightarrow$ 18% $\rightarrow$ 16%) and (20% $\rightarrow$ 15% $\rightarrow$ 12%) with (10% $\rightarrow$ 12% $\rightarrow$ 20%) and (4% $\rightarrow$ 8% $\rightarrow$ 9%) growth in energy requirements. Subsequent increments to the core frequency or their number returns lower and lower improvements to reconstruction time, deteriorating more steeply for the number of cores. Interestingly, for dimensionally large signals, the energy consumption of OMP is almost constant against the number of LITTLE cores and, in fact, decreases at the 2-core mark. However, a single A15 core at 0.8 GHz is still a better option for OMP, as it consumes equivalent or less energy of a LITTLE core but reconstruct signals much faster.

Maximum energy efficiency on that gateway can be achieved by using shorter ECG segments and a single A7 and A15 for HQ and LQ reconstruction respectively. However, this results in reduced signal compressibility, and theoretically, increased power consumption of the sensing node. Moreover, a 1× A15 core can halve or quarter HQ execution time, especially for larger windows, while consuming slightly more energy than its 1× A7 counterpart. Hence, optimal real-time and energy reconstruction performance for the gateway can be achieved using a single A15 core at minimum frequency.

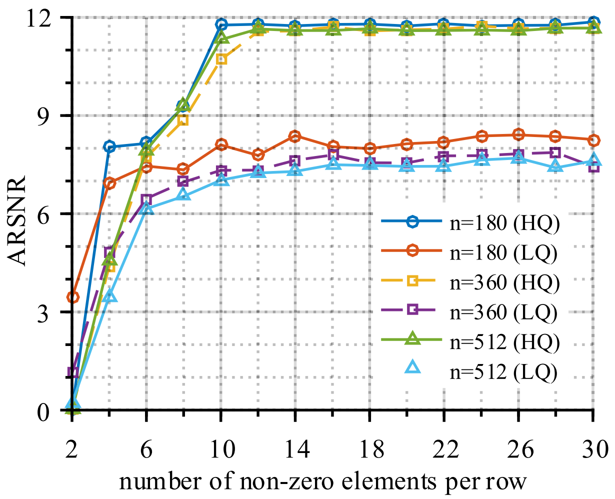

FIGURE 5. ARSNR verses the number of non-zero elements per row of the sensing matrix.

C. COMPRESSION ON THE SHIMMER

A digital CS implementation on the sensing node is carried out using the following steps. First, ECG signals must be buffered to their desired window size and then a matrix multiplication is performed on it to reduce its dimensions prior to transmission. If a sensing matrix with binary elements is used, the multiplications turn into additions, further simplifying CS deployment. Since CS inherently requires a delayed real-time (DRT) implementation, this section compares CS to both real-time and non-compressed DRT transmission modes.

The main challenge is the storage of the sensing matrix and computing compression in real-time. A generic binary sensing matrix generated from a Gaussian distribution has approximately ($nm/2$) non-zero elements, which is impossible to store nor it allows for real-time compression. Thankfully, the authors in [5] experimentally showed—using MIT arrhythmia database—that a sparse binary sensing matrix designed by populating a fixed number of randomly distributed 1s in each column can provide compression performance equivalent to conventional sensing matrices. In a system with two compression points (HQ and LQ), a sparse sensing matrix where non-zero elements are distributed per row is more memory efficient. Since rows are independent from each other, the Shimmer only needs to store the larger HQ matrix and can utilize it for the LQ point by evaluating less compressed samples (i.e. use less rows from the stored matrix). The appropriate number of non-zero elements per row is determined experimentally by computing the reconstruction accuracy against the number of non-zero elements in a row (Fig. 5). Clearly, the output ARSNR saturates between 10 to 12 non-zero elements for the HQ operative point and 16 to 18 for LQ. Since one matrix is used for both operation points, a single matrix with 18-ones in each row—generated independently using MATLAB's `sprandn()` function—is used for compression for the remainder of this implementation.

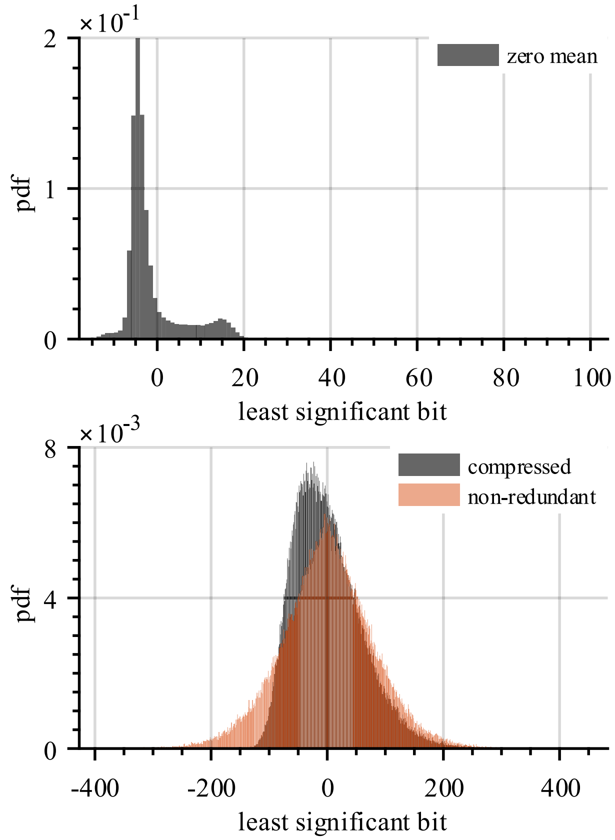

FIGURE 6. Probability density function ECG sample values with their mean removed (top) and their corresponding compressed packets compared to [5]'s redundancy removal method (bottom) computed from 256000 samples.

The Shimmer stores ECG samples in 16-bit integers, however, it was observed that much of the information contained in those variables pertain to the dc-level of the signal and not its morphology. And since CS requires a buffering, the dc-component can be removed to reduce the number of bits needed to represent the ECG. Figure 6 concludes that 8-bit are sufficient to represent zero-mean DRT ECGs and 10-bits are required for the corresponding compressed samples. This reduction of bit-range was reported on in [5], however, it was attributed to interpacket redundancy compressed packets and addressed by subtracting a packet from the one sent before it. But as Fig. 6 shows, redundancy removal does not further decrease the bit-range. The cause of the lower bit-range is dc-removal, meaning that this property can be extended to DRT ECGs and can be addressed in a simpler manner. Digital CS is primarily concerned with reducing the number of bytes prior to transmission. Since compressed samples are 25% larger than their zero-mean DRT counterpart, our HQ and LQ CRs are reduced by the same amount to $\{1.00, 1.39, 1.46\}$ and $\{2.25, 2.65, 2.87\}$ respectively.

Compression is typically computed in bursts after the ECG buffer has been filled. Using this approach, our most computationally complex case ($n = 512$ HQ) executes in under 18 ms which is under real-time constraints but is higher than the sampling period. Consequently, unless a secondary

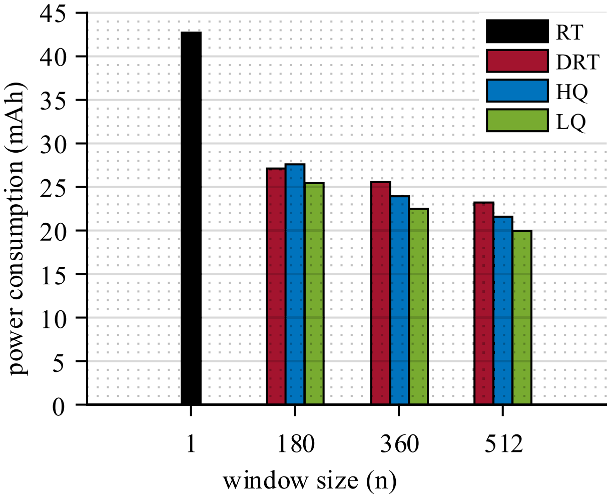

FIGURE 7. Shimmer power consumption for real-time (RT), DRT and CS transmission modes and varying window sizes.

TABLE 2. Shimmer battery lifetime estimation for real-time (RT), DRT and CS transmission modes and varying window sizes.

	Battery Lifetime (h)			
	$n=1$	$n=180$	$n=360$	$n=512$
RT	10.5	—	—	—
DRT	—	16.6	17.8	19.4
HQ	—	16.3	18.8	20.8
LQ	—	17.7	20.0	22.5

buffer and an interrupt implementation is used, some samples between two consecutive windows will be lost, which could be problematic if they were in the QRS region. We optimized CS execution time by taking advantage of the sampling period to perform a small number of additions for the compressed samples. The sensing matrix is stored as the locations of ones per column and once an ECG sample is measured it is immediately added to the compressed samples where it is present. When the buffer is ready, the compressed vector is computed in 1.7 ms, which is the time associated with m-subtractions for mean removal and $5m/4$ shift operations for bit reduction. Finally, the output of CS was validated by test packets containing both the original and compressed versions and computing the compressed output in MATLAB, where identical results were observed.

The Shimmer's power consumption in real-time, DRT, and CS modes are shown in Fig. 7. Most of the power saving benefits can be obtained by a simple DRT implementation as the figure shows a significant drop in comparison to real-time transmission. Furthermore, using a larger window is observed to reduce power consumption by a constant of $\approx$ 1.5 mAh, perhaps due to increased idle vs active ratio of the Bluetooth connection. Relative to real-time transmission, DRT improves the device's power consumption by (36, 40, 45) percent for 180, 360, 512 ECG windows respectively. Obviously, HQ "compression" at $n = 180$ drains more of the battery since it transmits an equivalent number of bytes but adds computational costs. Although, the small increase of 1.8% attests to CS encoder simplicity. The battery lifetime of the sensing node in the different modes is estimated in Table 2. HQ CS provides at most 1.4 hours extension to the sensor battery life (6.6%); it could be argued that this small amount is not worth the added ECG distortion and the reconstruction costs at the gateway. LQ, on the other hand, can add up to 3.1 hours of battery life extension (15.4%) and makes the device 1.5 hours short of continuous 24 hours operation. Maximum energy efficiency of the wearable device can be achieved using $n = 512$ ECGs with HQ and LQ compression.

D. SYSTEM EVALUATION

In real-world remote health monitoring systems, the gateway would be responsible for much more than signal reconstruction. It establishes Bluetooth and WiFi/cellular network connections, routes data from the sensor to the cloud, and in REMS—continuously performs analysis on the signal. This section compares the use of HMPs as a simple router and as a main processing unit (i.e. in REMS) by simulating real-world configuration and measuring and averaging power consumption at the gateway. Estimation of its battery lifetime is based on 10.78 Wh battery such as in [36]. In REMS, the transmission of clinical grade ECGs is limited, and hence, the battery life is estimated assuming a 25% time allocation for remote streaming of HQ data and 75% for LQ with local analytics.

Bluetooth and WiFi reception and routing of data is done through a Python script and Dropbox [49] to simulate the remote server. The script is assigned to $1\times$ A7 at 0.8 GHz to minimize its energy footprint. In REMS scenarios, three additional processes are implemented in C++ which are filtering, QRS detection using Pan and Tompkins [50], and K-nearest neighbour classifications (KNN). Note that the accuracy of those algorithms is not within the scope of this work, they are simply used to artificially load the CPU and more accurately simulate real-world performance. Four configurations are considered: conventional router, and three REMS scenarios that minimize or optimize the energy consumption at the wearable device and the gateway. The configurations, and their objectives and descriptions are summarized in Table 3.

The power consumption and the corresponding battery life for each configuration are depicted in Fig. 8 and Table 4 respectively. By comparing CFG1 with any DRT configuration, it is concluded that most energy consumption at the gateway is due to WiFi utilization and not CS reconstruction or signal processing. It must be noted that even in LQ cases, the gateway is still connected to WiFi but idling, meaning that it could still receive notifications from the server or quickly switch to medical grade ECG transmission in case of accidents. Also, the observed computational time for the three aforementioned additional signal processing is 15ms on average. Hence, more operations can be implemented on the gateway without worry about real-time performance. If the signal is both processed and transmitted, energy consumption

TABLE 3. Configurations, and their objective and description.

Configuration	Objective	Description
CFG1	**Conventional**: gateway as a router; used as a benchmark.	Real-time reception of data from the Shimmer, buffered in 1-seconds chucks and uploaded to cloud using a 1×A7 at 0.8 GHz.
CFG2	**REMS**: Minimize energy consumption at the wearable device	Clinical grade ECGs in HQ at $n = 512$, reconstructed and analyzed using 1×A15 at 0.8 GHz and transmitted to cloud every 3 seconds. LQ ECGs follow the same setup but without cloud transmission.
CFG3	**REMS**: Optimize energy consumption considering the wearable device and gateway.	Clinical grade ECGs in DRT at $n = 512$, analyzed using 1×A7 at 0.8 GHz and transmitted to cloud every 3 seconds. LQ ECGs are at $n = 360$, reconstructed and analyzed using 1×A15 at 0.8 GHz
CFG4	**REMS**: Minimize energy consumption at the gateway	Clinical grade ECGs in DRT at $n = 180$, analyzed using 1×A7 at 0.8 GHz and transmitted to cloud every 3 seconds. LQ ECGs are at $n = 180$, reconstructed and analyzed using 1×A15 at 0.8 GHz

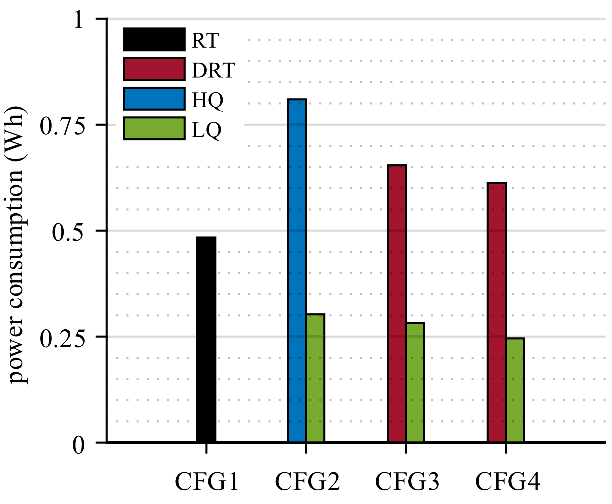

FIGURE 8. Gateway power consumption under the different configurations.

TABLE 4. Battery life of the wearable device and the gateway in different configurations.

	Battery Lifetime (h)			
	CFG1	CFG2	CFG3	CFG4
Wearable device	10.5	22.1	19.9	17.5
Gateway	22.3	25.1	28.7	31.9

at the gateway rises in comparison with the conventional approach (CFG1), but by limiting the transmission period (i.e. only transmitting when abnormalities are detected or at the request of the end user), the overall battery life of the gateway is extended. CFG2 provides the best outcomes for a connected health system because both the wearable device and the gateway have approximately 24 hours of operation.

VII. CONCLUSIONS

Traditional IoT-based remote health monitoring systems generally suffer from latency and reliability issues associated with cloud computing. which can be addressed by the mobile and efficient processing power of HMPs. This work presented REMS, a mobile real-time remote health monitoring architecture centered around the utilization of ARM's big.LITTLE^{TM} heterogenous multicore platform for signal processing, focusing on the efficacy of ECG CS under real-time and hardware constraints. Compressively sensed signal recovery can be performed on the gateway in real-time while minimizing energy consumption of the gateway in comparison to the traditional gateway-as-router point of view. With a single A15 (big) core, reconstruction occurred in under 30% CPU utilization and consumed an equivalent amount of energy to the energy-efficient LITTLE cores. On the wearable device, CS reduced the power consumption of the wearable device by up to 15.4% when considering ECGs suitable for wellness applications. However, performance was generally suboptimal for clinical grade signals. Power consumption at the gateway was largely due to the continuous wireless transmission of data, hence, a system that relies on the gateway as a signal processing unit benefits from an ample extension in battery life. Due to its capacity for real-time and energy efficient computation and ever-growing potential of adoption, HMPs can be considered a suitable technology for connected health, leading the way towards a more reliable remote health monitoring system.

REFERENCES

[1] X.-F. Teng, Y.-T. Zhang, C. C. Y. Poon, and P. Bonato, "Wearable medical systems for P-health," *IEEE Rev. Biomed. Eng.*, vol. 1, no. 1, pp. 62–74, Jan. 2008.

[2] E. Spanò, S. Di Pascoli, and G. Iannaccone, "Low-power wearable ECG monitoring system for multiple-patient remote monitoring," *IEEE Sensors J.*, vol. 16, no. 13, pp. 5452–5462, Jul. 2016.

[3] A. Milenković, C. Otto, and E. Jovanov, "Wireless sensor networks for personal health monitoring: Issues and an implementation," *Comput. Commun.*, vol. 29, nos. 13–14, pp. 2521–2533, Aug. 2006.
[4] M. Rani, S. Dhok, and R. Deshmukh, "A systematic review of compressive sensing: Concepts, implementations and applications," *IEEE Access*, vol. 6, pp. 4875–4894, 2018.
[5] H. Mamaghanian, N. Khaled, D. Atienza, and P. Vandergheynst, "Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 9, pp. 2456–2466, Sep. 2011.
[6] D. H. Chae, Y. F. Alem, S. Durrani, and R. A. Kennedy, "Performance study of compressive sampling for ECG signal compression in noisy and varying sparsity acquisition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2013, pp. 1306–1309.
[7] C. Karakus, A. C. Gurbuz, and B. Tavli, "Analysis of energy efficiency of compressive sensing in wireless sensor networks," *IEEE Sensors J.*, vol. 13, no. 5, pp. 1999–2008, May 2013.
[8] D. Craven, B. McGinley, L. Kilmartin, M. Glavin, and E. Jones, "Compressed sensing for bioelectric signals: A review," *IEEE J. Biomed. Health Inform.*, vol. 19, no. 2, pp. 529–540, Mar. 2015.
[9] M. M. Baig, H. Gholamhosseini, and M. J. Connolly, "A comprehensive survey of wearable and wireless ECG monitoring systems for older adults," *Med. Biol. Eng. Comput.*, vol. 51, no. 5, pp. 485–495, 2013.
[10] T. Bianchi, V. Bioglio, and E. Magli, "On the security of random linear measurements," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2014, pp. 3992–3996.
[11] Y. Zhang, L. Y. Zhang, J. Zhou, L. Liu, F. Chen, and X. He, "A review of compressive sensing in information security field," *IEEE Access*, vol. 4, pp. 2507–2519, 2016.
[12] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "On known-plaintext attacks to a compressed sensing-based encryption: A quantitative analysis," *IEEE Trans. Inf. Forensics Security*, vol. 10, no. 10, pp. 2182–2195, Oct. 2015.
[13] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Low-complexity multiclass encryption by compressed sensing," *IEEE Trans. Signal Process.*, vol. 63, no. 9, pp. 2183–2195, May 2015.
[14] F. Samie, V. Tsoutsouras, L. Bauer, S. Xydis, D. Soudris, and J. Henkel, "Computation offloading and resource allocation for low-power IoT edge devices," in *Proc. IEEE 3rd World Forum Internet Things (WF-IoT)*, Dec. 2016, pp. 7–12.
[15] D. Minoli, *Building the Internet of Things With IPv6 and MIPv6: The Evolving World of M2M Communications*. Hoboken, NJ, USA: Wiley, 2013.
[16] T. Takpor and A. A. Atayero, "Integrating Internet of Things and EHealth solutions for students' healthcare," in *Proc. World Congr. Eng.*, London, U.K., vol. 1, 2015, pp. 265–268.
[17] S. M. Riazul Islam, D. Kwak, M. Humaun Kabir, M. Hossain, and K.-S. Kwak, "The Internet of Things for health care: A comprehensive survey," *IEEE Access*, vol. 3, pp. 678–708, Jun. 2015.
[18] G. Sebestyen, A. Hangan, S. Oniga, and Z. Gál, "eHealth solutions in the context of Internet of Things." in *Proc. IEEE Int. Conf. Automat., Qual. Test., Robot.*, May 2014, pp. 1–6.
[19] H. Moustafa, E. M. Schooler, G. Shen, and S. Kamath, "Remote monitoring and medical devices control in eHealth," in *Proc. 12th Int. Conf. Wireless Mobile Comput., Netw. Commun. (WiMob)*, Oct. 2016, pp. 1–8.
[20] A. M. Rahmani *et al.*, "Exploiting smart e-health gateways at the edge of healthcare Internet-of-Things: A fog computing approach," *Future Gener. Comput. Syst.*, vol. 78, pp. 641–658, Jan. 2018.
[21] A. Bansal *et al.*, "Remote health monitoring system for detecting cardiac disorders," *IET Syst. Biol.*, vol. 9, no. 6, pp. 309–314, Dec. 2015.
[22] D.-H. Shih, H.-S. Chiang, B. Lin, and S.-B. Lin, "An embedded mobile ECG reasoning system for elderly patients," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 3, pp. 854–865, May 2010.
[23] S. Gradl, P. Kugler, C. Lohmüller, and B. Eskofier, "Real-time ECG monitoring and arrhythmia detection using Android-based mobile devices," in *Proc. IEEE Annu. Int. Conf. Eng. Med. Biol. Soc.*, Aug./Sep. 2012, pp. 2452–2455.
[24] J. J. Oresko *et al.*, "A wearable smartphone-based platform for real-time cardiovascular disease detection via electrocardiogram processing," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 3, pp. 734–740, May 2010.
[25] A. M. R. Dixon, E. G. Allstot, D. Gangopadhyay, and D. J. Allstot, "Compressed sensing system considerations for ECG and EMG wireless biosensors," *IEEE Trans. Biomed. Circuits Syst.*, vol. 6, no. 2, pp. 156–166, Apr. 2012.
[26] Y. Zigel, A. Cohen, and A. Katz, "The weighted diagnostic distortion (WDD) measure for ECG signal compression," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 11, pp. 1422–1430, Nov. 2000.
[27] A. J. Casson and E. Rodriguez-Villegas, "Signal agnostic compressive sensing for Body Area Networks: Comparison of signal reconstructions," in *Proc. IEEE Annu. Int. Conf. Eng. Med. Biol. Soc.*, Aug./Sep. 2012, pp. 4497–4500.
[28] F. Ansari-Ram and S. Hosseini-Khayat, "ECG signal compression using compressed sensing with nonuniform binary matrices," in *Proc. 16th CSI Int. Symp. Artif. Intell. Signal Process. (AISP)*, May 2012, pp. 305–309.
[29] Z. Zhang and B. D. Rao, "Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation," *IEEE Trans. Signal Process.*, vol. 61, no. 8, pp. 2009–2015, Apr. 2013.
[30] Z. Zhang, T.-P. Jung, S. Makeig, and B. D. Rao, "Compressed sensing for energy-efficient wireless telemonitoring of noninvasive fetal ECG via block sparse Bayesian learning," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 2, pp. 300–309, Feb. 2013.
[31] Z. Zhang, S. Wei, D. Wei, L. Li, F. Liu, and C. Liu, "Comparison of four recovery algorithms used in compressed sensing for ECG signal processing," in *Proc. Comput. Cardiol. Conf. (CinC)*, Sep. 2016, pp. 401–404.
[32] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 890–912, Jan. 2009.
[33] E. van den Berg and M. P. Friedlander. (Jun. 2007). *SPGL1: A Solver for Large-Scale Sparse Reconstruction*. [Online]. Available: http://www.cs.ubc.ca/labs/scl/spgl1
[34] K. Kanoun, H. Mamaghanian, N. Khaled, and D. Atienza, "A real-time compressed sensing-based personal electrocardiogram monitoring system," in *Proc. Design, Automat. Test Eur.*, Mar. 2011, pp. 1–6.
[35] D. Bortolotti, A. Bartolini, M. Mangia, R. Rovatti, G. Setti, and L. Benini, "Energy-aware bio-signal compressed sensing reconstruction: FOCUSS on the WBSN-gateway," in *Proc. IEEE 9th Int. Symp. Embedded Multicore/Many-Core Syst.-Chip (MCSoC)*, Sep. 2015, pp. 120–126.
[36] D. Bortolotti, M. Mangia, A. Bartolini, R. Rovatti, G. Setti, and L. Benini, "Energy-aware bio-signal compressed sensing reconstruction on the WBSN-gateway," *IEEE Trans. Emerg. Topics Comput.*, vol. 6, no. 3, pp. 370–381, Jul./Sep. 2018.
[37] F. Pareschi *et al.*, "Energy analysis of decoders for rakeness-based compressed sensing of ECG signals," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 6, pp. 1278–1289, Dec. 2017.
[38] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008.
[39] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.
[40] H. Baali, R. Akmeliawati, M. J. E. Salami, A. Khorshidtalab, and E. Lim, "ECG parametric modeling based on signal dependent orthogonal transform," *IEEE Signal Process. Lett.*, vol. 21, no. 10, pp. 1293–1297, Oct. 2014.
[41] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
[42] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230–2249, May 2009.
[43] *Wearable ECG Sensor | Wireless ECG Sensor | Electrocardiogram*. Accessed: May 2018. [Online]. Available: http://www.shimmersensing.com/products/shimmer3-ecg-sensor
[44] A. Burns *et al.*, "SHIMMER—A wireless sensor platform for noninvasive biomedical research," *IEEE Sensors J.*, vol. 10, no. 9, pp. 1527–1534, Sep. 2010.
[45] *Shimmer3 Applications for Code Composer Studio*. Accessed: May 2018. [Online]. Available: https://github.com/ShimmerResearch/shimmer3
[46] *Code Composer Studio (CCS) Integrated Development Environment (IDE)*. Accessed: May 2018. [Online]. Available: http://www.ti.com/tool/CCSTUDIO
[47] *ODROID | Hardkernel*. Accessed: May 2018. [Online]. Available: http://www.hardkernel.com
[48] C. Sanderson and R. Curtin, "Armadillo: a template-based C++ library for linear algebra," *J. Open Source Softw.*, vol. 1, no. 2, p. 26, Jun. 2016.
[49] *Dropbox*. Accessed: May 2018. [Online]. Available: https://www.dropbox.com/
[50] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, Mar. 1985.

MOHAMMED AL DISI received the B.S. degree in electrical engineering from Qatar University. He is currently a Research Assistant with Qatar University working on heterogeneous multicore platforms for connected health applications. His research interests include biomedical signal processing and connected health implementations. He was a recipient of the Best Poster Presentation and Oral Presentation Awards at the Extracorporeal Life Support Organization South West Asia Conference 2017 for the works Using Thermochromic Ink for Medical Simulation and Design and implementation of a modular ECMO simulator.

HAMZA DJELOUAT received the B.Eng. degree in telecommunication from the Institute of Electrical and Electronic Engineering, Boumerdes, Algeria, in 2013, and the M.Sc. degree in electronic systems majoring telecommunication from the Ecole Militaire Polytechnique, Algiers, Algeria, in 2015. He is currently a Research Assistant with Qatar University. His research interests are mainly in digital signal processing and the applications of compressive sensing in wireless networks and connected health.

CHRISTOS KOTRONIS received the B.Sc. degree in informatics and telematics from Harokopio University in 2014 and the M.Sc. degree in Web engineering from the Department of Informatics and Telematics, of Harokopio University in 2016, where he is currently pursuing the Ph.D. degree under the supervision of Prof. M. Nikolaidou. His Ph.D. research—systems-of-systems engineering: an integrated model-driven methodology for complex systems—proposes model-driven approaches, utilizing SysML as a modeling language to automatically generate performance and cost models. Case studies on transportation, e-health, and sensor networks, are explored, taking into account the dynamic behavior of such systems-of-systems under real-time conditions. He has been a Teaching Assistant with the Department of Informatics and Telematics, Harokopio University of Athens, since 2016. He received the three-year (2017–2019) Engineering and Technology Sciences Scholarship from the General Secretariat for Research and Technology (GSRT) and the Hellenic Foundation for Research and Innovation (HFRI) for conducting his Ph.D. research.

ELENA POLITIS received the B.Sc. degree in physics from the National and Kapodistrian University of Athens and the M.Sc. degree (Hons.) in informatics and telematic services from Harokopio University, where she is currently pursuing the Ph.D. degree with the Department of Informatics and Telematics. Her dissertation is dedicated to the Internet of Underwater Things. She speaks fluent Greek and English and is very good in the German language. Her research interests revolve around the Internet of Things, related applications and services, and optimization algorithms for autonomous underwater vehicles.

ABBES AMIRA (S'99–M'01–SM'07) received the Ph.D. degree in the area of computer engineering from Queen's University Belfast, Belfast, U.K., in 2001.

He took many academic and consultancy positions, including his recent positions as an Associate Dean for Research and Graduate Studies at Qatar University and also a Professor of visual communications with the University of the West of Scotland, U.K. He published over 300 papers in top journals and conferences and secured significant funding from government and industry. His research interests include reconfigurable computing, signal processing, IoT, connected health and security. He is an IET fellow, an HEA fellow, a Senior Member of the ACM.

FAYCAL BENSAALI (S'03–M'06–SM'15) received the Dipl.-Ing. (M.Eng.) degree in electronics from the University of Constantine and the Ph.D. degree in electronic and computer engineering from Queen's University Belfast, U.K. He is currently an Associate Professor in electrical engineering with Qatar University. He took other academic positions at the University of Hertfordshire, U.K., and Queen's University Belfast. His research interests are mainly in embedded systems and high performance computing, custom computing using FPGAs, image and vision systems, and connected health.

Dr. Bensaali serves as a Guest Editor of the IEEE IoT journal and the General Chair, the Workshop Chair, and a TPC member of a number of international conferences and workshops. He has authored/co-authored over 100 scientific papers in international journals and conference proceedings. He is an HEA Associate Member.

GEORGE DIMITRAKOPOULOS received the bachelor's degree in electrical and computer engineering from the National Technical University of Athens in 2002 and the Ph.D. degree from the University of Piraeus in 2007. He has been an Assistant Professor with the Department of Informatics and Telematics, Harokopio University of Athens, since 2010, a Senior Consultant and Project Manager with Intrasoft International S.A. since 2015, and a Coordinator of funded projects at Infineon Technologies A.G. since 2018. He is the author of about 120 publications in international journals and conferences, whereas he has been involved in over 20 internationally funded research and development projects. His research interests include the design and development of safety and emergency management algorithms for smart devices, and the design and development of strategies for the optimization of wireless networks based on cognitive networking principles.

GUILLAUME ALINIER received the M.Phys. (Hons.) degree and the Ph.D. degree. He received the PGCert. He is a National Teaching Fellow in 2006, Senior Fellow of the Higher Education Academy, a member of the Institute of Physics and Engineering in Medicine, and a member of the Institute of Physics.

He is currently the Director of research with Hamad Medical Corporation Ambulance Service, Doha, Qatar. He has been a Visiting Fellow at Northumbria University, Newcastle upon Tyne, U.K., since 2009, and a Professor of simulation in healthcare education with the University of Hertfordshire, Hatfield, U.K., since 2011. He is a long standing senior member of several healthcare simulation societies and has contributed to over 350 simulation-related peer-reviewed publications, book chapters, and conference presentations and posters. He has received several awards from his employers and the Higher Education Academy, U.K., for his work in relation to learning and teaching. He has editorial responsibilities with the journals of *Advances in Simulation*, *Simulation in Healthcare, Paramedic Practice, International Paramedic Practice, BMJ Simulation and Technology Enhanced Learning*, and the *Mediterranean Journal of Emergency Medicine*. He is currently involved in several innovative projects across the Hamad Medical Corporation and also in collaboration with Qatar University, and serves on various international committees in relation to standards, certification, and accreditation.