

 Open access • Proceedings Article • DOI:10.1109/ACSSC.2016.7869524

## Exploiting system configurability towards dynamic accuracy-power trade-offs in sensor front-ends — [Source link](#)

O. Laura I. Galindez, Komail Badami, V. Rajesh Pamula, Steven Lauwereins ...+2 more authors

**Institutions:** Katholieke Universiteit Leuven

**Published on:** 01 Jan 2016 - Asilomar Conference on Signals, Systems and Computers

**Topics:** Data stream, Voice activity detection, Feature extraction and Signal reconstruction

Related papers:

- [Biomedical sensor design using analog compressed sensing](#)
- [Design and Analysis of a Hardware-Efficient Compressed Sensing Architecture for Data Compression in Wireless Sensors](#)
- [Hibernets: Energy-Efficient Sensor Networks Using Analog Signal Processing](#)
- [Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing](#)
- [Time-Encoding-Based Ultra-Low Power Features Extraction Circuit for Speech Recognition Tasks](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/exploiting-system-configurability-towards-dynamic-accuracy-4ophjcsk2y>

<b>Citation</b>	Laura I. Galindez O. , Komail Badami , V. Rajesh Pamula , Steven Lauwereins , Wannes Meert , Marian Verhelst , (2016), <b>Exploiting System Configurability towards Dynamic Accuracy-Power Trade-offs in Sensor Front-ends</b> 50th Asilomar Conference on Signals, Systems and Computers ,pages: 1027-1031.
<b>Archived version</b>	Author manuscript: the content is identical to the content of the published paper, but without the final typesetting by the publisher
<b>Published version</b>	<a href="http://dx.doi.org/10.1109/ACSSC.2016.7869524">http://dx.doi.org/10.1109/ACSSC.2016.7869524</a>
<b>Journal homepage</b>	<a href="http://www.asilomarsscconf.org/">http://www.asilomarsscconf.org/</a>
<b>Author contact</b>	Laura.Galindez@esat.kuleuven.be + 32 (0)16 323362

*(article begins on next page)*

# Exploiting System Configurability towards Dynamic Accuracy-Power Trade-offs in Sensor Front-ends

Laura I. Galindez O.\*, Komail Badami\*, V. Rajesh Pamula\*<sup>†</sup>, Steven Lauwereins\*, Wannes Meert<sup>‡</sup>, Marian Verhelst\*<sup>†</sup>

\*ESAT-MICAS, KU Leuven, Leuven, Belgium

<sup>†</sup>IMEC, Leuven, Belgium

<sup>‡</sup>CS-DTAI, KU Leuven, Leuven, Belgium

**Abstract**—Analog-to-information converters and Compressed Sampling (CS) sensor front-ends try to only extract the relevant, information-bearing elements of an incoming data stream. Information extraction and recognition tasks can run directly on the compressed data stream without needing full signal reconstruction. The accuracy of the extracted information or classification is strongly determined by the front-end settings and tolerated level of hardware impairments. Exploiting this, allows to dynamically tune accuracy for power consumption. This paper discusses this trade-off and introduces a theoretical framework to guide the selection of optimal hardware settings under given power or accuracy constraints. This is illustrated with two circuit realizations: 1) an analog-to-information converter for voice activity detection (VAD), and 2) a CS photoplethysmographic (PPG) heart rate (HR) extraction application.

## I. INTRODUCTION

Reactive sensing applications on portable devices, such as phones or smart-watches, need a variety of always-on integrated sensors. Existing sensor front-ends immediately translate the observed physical analog signals into the digital domain after only basic analog signal conditioning as shown in Fig.1(a). Such a system generates raw digital data at the rate ( $> 2 \cdot sig_{BW} \cdot 2^{nbits}$ ) bits/sec, where  $sig_{BW}$  is the raw signal bandwidth. With multiple sensors on every device, this analog to digital translation of raw sensory signals generates a huge amount of data to be stored, processed and transmitted. Continuous information extraction from this sensory data deluge reduces the battery life of the device, prohibiting thus its continuous always-on operation. Analog-to-information converters [1] or CS front-ends [2] overcome this by only extracting a relevant feature subset from the incoming sensory stream. In such systems, early data discrimination is ensured by extracting features/information as close to the sensor as possible, as highlighted in Fig.1(b). Such approach results in a rate of  $2 \cdot info_{BW} \cdot 2^{nbits}$  bits/sec, where  $info_{BW}$  is the information bandwidth which is normally much lower than  $sig_{BW}$  [1]. The power consumption of the sensing system, as well as the information content of the retained data stream, strongly depend on the computational precision of extracted features. Tuning this precision enables an interesting, dynamic, trade-off between system power consumption and information or classification accuracy.

This work strives to maximize the relevance of extracted information under a given operating context through a sensing




Fig. 1: (a) Nyquist sensor systems (b) Information sensing systems.

architecture that dynamically trades off feature extraction power against classification accuracy, as presented in Section II. This trade-off is then illustrated with a context aware VAD system in Section III and a CS based heart-rate extraction from PPG signals in Section IV.

## II. FEATURE TUNABLE SENSING SYSTEMS

Fig. 2 shows the proposed feature tunable sensing system. The features generated by the extraction unit are fed to a classifier or information extractor that performs the application specific information estimations, i.e. speech identification or HR estimation. Confidence on resulting estimations strongly depends on environmental conditions such as signal dynamics or signal interference. A feedback loop is thus required to dynamically tune the specifications of the feature extraction unit under given power or performance constraints. In the system highlighted by Fig. 2, a control logic block that dictates the feature subset's properties to the extraction unit is inserted in the loop. This block trades application performance for power savings to determine the feature requirements. The remainder of this section describes how to enable the dynamic tunability of the feature extraction unit. Secondly, it describes how this tunability allows dynamic power scaling.

### A. Feature extraction unit

The feature extraction unit in Fig. 2 reduces the generated data rate by decomposing the raw signal into a set of useful features. This unit is dynamically tunable in terms of the number of features sampled, their precision (Section III) and their sampling frequency (Section IV). As Sections III and IV detail, varying computational precision and selectively computing a context-aware feature subset allows to continuously scale the




Fig. 2: Self-tunable variable precision sensing systems.

systems' power consumption, achieving  $10\text{--}50\times$  savings over existing state-of-the-art implementations [3].

### B. System performance evaluation and dynamic power scaling

Aforementioned power saving opportunities present an interesting trade-off against extracted information accuracy. Among other aspects, task accuracy is determined by the environmental (i.e. ambient noise) and computational noise (i.e. data resolution) that features are subjected to. Our framework allows to dynamically assess this impact and tune the associated computational noise in favor of power scaling. The VAD application discussed in Section III exploits this flexibility by means of a Bayesian Network based framework that enables feature subset selection and feature precision tunability. In its turn, the Lomb-Scargle periodogram (LSP) based HR estimation application illustrated in Section IV exploits signal acquisition rate reduction opportunities provided by CS. Both applications will demonstrate how the sensor front-end's adaptiveness allows to dynamically tune the confidence in the extracted information with its power consumption in order to efficiently deal with the stochastic nature of ambient noise and external interference.

## III. FEATURE PRECISION TUNABILITY FOR VAD

Always-awake VAD systems must operate continuously to identify speech under changing ambient noise conditions, which results in a significant power budget restriction. A mixed-signal, highly tunable VAD front-end was introduced and proved in silicon in [4]. This system enables power-scalable feature extraction by 1) extracting information bearing features in the analog domain, 2) enabling individual feature (de)activation and 3) enabling individual feature-precision scaling. The VAD system's top-level block diagram is shown in Fig. 3(a). The analog feature-extractor described in Section III-A translates the input signal to a set of features which are discretized by a slow ADC for subsequent processing. In this work we exploit the tunability of the VAD's feature extractor by means of a Bayesian Network classifier and a feature precision selection block, detailed in Section III-B. The Pareto-optimal feature extraction power vs classification accuracy trade-off resulting from this scheme is then presented in Section III-C. An example of a confidence metric that can be used by the Control Logic block in Figure 2 to dynamically reconfigure the feature extraction's properties is also described in this Section.




Fig. 3: (a) Block diagram for VAD (b) Variable precision feature extraction for VAD.

### A. Variable-precision feature extraction

Analog features are computed by decomposing the acoustic signal into different frequency bands and estimating the power in each frequency band. Fig. 3(b) shows the feature extractor's schematic for one of the 16 bands. The amplified signal from a passive microphone is fed to a Band Pass Filter (BPF), whose output is rectified and averaged to obtain the analog feature [4]. The center frequency of the BPF increases exponentially from 75 Hz to 5 kHz from band 1 to band 16. Further, each band has variable amplification for precision-tunability. Each band is capable of extracting the feature set  $F_i = \{F_{i,8b}, F_{i,4b}, F_{i,2b}, F_{i,1b}\}$ , where the subscript  $nb$  denotes precision (expressed in number of effective bits) of the feature and  $i$  ranges from 1 to 16. Barring the overhead for low noise amplification for passive microphone, the feature extractor's power consumption scales: a) exponentially from feature 1 to 16, and b) proportionally with each additional bit of resolution [4], [5].

### B. Context and feature precision aware Bayesian Network classifier

To enable dynamically-scalable VAD, the tunable feature extraction unit is interfaced with a Bayesian Network classifier. This classifier assigns a 'voice'/'no voice' class label  $c$  to a class variable  $C$  that maximizes the posterior probability  $Pr(C|F_1, \dots, F_n)$  where  $\mathbf{U} = \{F_1, \dots, F_n\}$  are observed features. For the Naive Bayes (NB) classifier in Fig. 4(a), this computation is efficient due to the feature conditional independence assumption [6]. We will here extend the NB classifier with variable precision children of the feature nodes, as in [7], to build a dynamically tunable VAD classifier. As shown in Fig. 4(b), every feature is represented with a high precision (8b) node, as well as with lower precision nodes (4b, 2b, 1b), and a feature pruning node (0b). The resulting graph hence allows to infer the joint probability over a subset of observed features  $F_{i,nb}$  at various precisions expressed in number of bits ( $nb$ ), marginalized over all higher precision non-observed nodes:

$$Pr(C, F_{1,8b}, \dots, F_{16,0b}) = \prod_{i=1}^{16} \prod_{nb=0b}^{4b} Pr(F_{i,nb}|F_{i,nb+1}) \cdot Pr(F_{i,8b}|C) \cdot Pr(C), (1)$$




Fig. 4: (a) Naive Bayes Classifier (b) Precision scalable naive Bayes classifier.

where the probabilistic relations  $Pr(F_{i,nb}|F_{i,nb+1})$  are deterministic and the relations  $Pr(F_{i,8b}|C)$  are modeled as discretized Gaussian distributions. This model moreover allows to individually assess feature usefulness in accordance to the computational and ambient noise they bear. This is of particular relevance to applications such as VAD, where the features' relative information content is highly context-dependent (e.g. background noise). The wide power vs accuracy trade-off space resulting from the extended NB model can be traversed with a heuristic that iteratively prunes a Pareto-optimal feature precision set  $\mathbf{U}_{select}$ , as proposed in [7]. At each iteration, a greedy neighborhood search is performed over  $n$  pruning candidates, each with a feature precision reduction of one level ( $\mathbf{U}_{candidate_i} \leftarrow \mathbf{U}_{select} \setminus F_{i,nb_i} \vee \{F_{i,nb_i-1}\}$ ). The candidates' impact on classification accuracy is estimated from the correctly classified instances in a validation set and its impact on power consumption is estimated from the hardware's characteristics defined in Section III-A. The selected feature set maximizes the cost function

$$CF = \log \left( \frac{\Delta power}{\max(power)} \right) - \log(\Delta accuracy), \quad (2)$$

where the term  $\Delta$  is the predicted state difference between each candidate  $i$  and the currently selected set. Hence, the greedy neighborhood search in this heuristic enforces maximal power consumption reduction in return for accuracy loss.

### C. VAD system results

Fig. 5 presents the feature extraction power vs accuracy trade-off achieved by this framework for the VAD application subjected to the noise inside a train under three SNR environments. The results show potential power consumption savings in the feature extraction unit of more than three orders of magnitude for an estimated accuracy degradation of less than 5% in each setting. Note that the optimal settings are strongly dependent on the current environmental conditions. For example, assume that the application requires a performance between 80% and 82% accuracy. When environmental conditions are not favorable (+3dB SNR), 11 features must be extracted at the precision specified in Table I resulting in the consumption of  $2.16 \mu W$ . When the conditions improve to +9dB and +15 dB SNR, power savings of one and three




Fig. 5: Power vs accuracy trade-off for the train context under different SNR conditions.

TABLE I: Selected feature sets for train context

Feature number	Number of bits		
	3dB SNR	9dB SNR	15dB SNR
1 to 5	-	-	-
6	1 bit	-	-
7	4 bits	4 bits	-
8	8 bits	8 bits	4 bits
9 and 10	8 bits	-	-
11	1 bit	-	-
12	8 bits	1 bit	-
13	2 bits	-	-
14 and 15	8 bits	-	-
16	8 bits	1 bit	-
Power	$2.16 \mu W$	$0.30 \mu W$	$0.0004 \mu W$

orders of magnitude can be achieved by extracting the feature subsets specified in Table I.

The aforementioned trade-off analysis determines the local optimal operating points available to the feature extractor. At run-time, the Control Logic block described in Figure 2 must make decisions in function of environmental variations which result in different levels of prediction confidence. We define such confidence metric as a variation of the log odds ratio:

$$Confidence = 10 \cdot \log_{10} \left( \frac{Pr(C = c_{max}|F)}{Pr(C = c_{min}|F)} \right), \quad (3)$$

where  $Pr(C = c_{max}|F)$  and  $Pr(C = c_{min}|F)$  are the posterior probabilities for each of the two class labels. Fig. 6 shows that, although for some settings estimated accuracy (blue) stays within the same range, confidence (red) differs significantly, again indicating the importance of the feedback loop that dynamically controls the feature extraction unit.

## IV. PERFORMANCE SCALABILITY FOR CS PPG BASED HR ESTIMATION

A PPG is a volumetric measure of an organ commonly obtained by illuminating the skin and measuring changes in light absorption. The LED stimulation of such systems dominates the overall power consumption [8], leading to the desire for early data rate reduction. Therefore, PPG greatly benefits from CS techniques, which acquire the signal at a rate proportional to  $info_{BW}$  rather than the Nyquist bandwidth [2], as illustrated in Figure 1. To circumvent the intensive signal reconstruction associated to CS based acquisition systems, Section IV-B presents an algorithm capable of estimating the

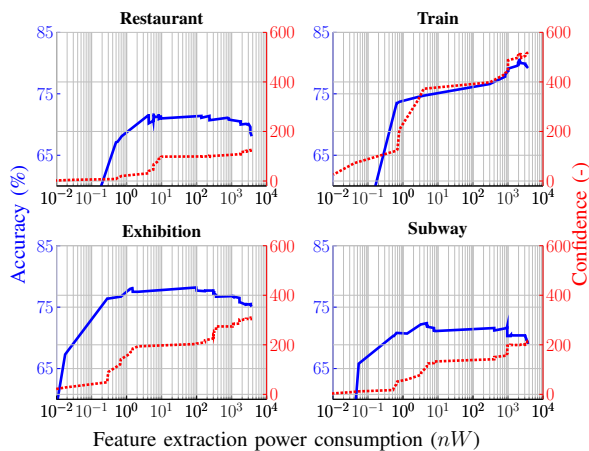


Fig. 6: Power vs accuracy and confidence trade-off for different environment settings at +3dB SNR.

average HR directly from CS PPG signal [9]. The confidence over such estimation is then be used to implementation the self-tunable system of Fig. 2. Finally, Section IV-C presents the trade-offs relevant to the CS based PPG acquisition system with integrated feature extraction.

#### A. Tunable CS in analog domain

For systems where power consumption is dominated by the sensor, like PPG acquisition, CS for reducing power consumption is possible only when implemented in the analog domain, as previously reported in the context of biomedical signal acquisition [10]. The signal's information rate, which determines the acquisition rate in CS, depends on its sparsity on a given basis  $\Psi$ . In practice, CS for PPG signals is equivalent to randomly sub-sampling the signal in time domain since PPG signals are sparse on frequency basis [9]. For an  $N$ -dimensional signal vector  $X$  that is  $K$ -sparse of a given basis  $\Psi$ , near exact information sampling is possible by acquiring a projected  $M$ -samples, provided the projection transform  $\Phi$  satisfies the restricted isometry property (RIP) and is maximally incoherent with the basis  $\Psi$ . The number of samples required for faithful information sampling is given by

$$M = \mathcal{O} \left( K \log \frac{N}{K} \right). \quad (4)$$

To quantify the data reduction, compression ratio,  $CR = N/M$ , is used. A configurable CR enables dynamic information extraction as in Figure 2.

#### B. Compressed domain information extraction for PPG read-out

When implemented in hardware, signal recovery techniques can consume up to 10 mWs of power [11], offsetting the savings obtained by compressively acquiring the signal. It is therefore desirable to extract the features of interest directly from CS data as proposed in [9] and shown in Fig. 7. This approach relies on least square spectral fitting technique




Fig. 7: Compressively sampled photoplethysmographic (PPG) readout with integrated feature extraction.

through LSP to estimate the power spectral density (PSD) of the randomly sub-sampled signal. The average HR can then be estimated from the frequency corresponding to the peak in the PSD. In practice, LSP is evaluated over a finite frequency range which is discretized. Associated with each frequency bin is a confidence or significance interval  $1 - \alpha$ . For the normalized periodogram,  $\alpha$  is given by  $\alpha = 1 - (1 - e^{-\eta})^k$ , where  $k$  is the number of bins and  $\eta$  the desired frequency value. A higher value of  $1 - \alpha$  for a bin indicates a higher likelihood of the true presence of the frequency bin in the signal, similarly to the confidence in (3). As described in III-C, tuning the information extraction block's properties (via configurable CR in this case) provides the means to trade off estimation confidence for system power consumption. The details of such a trade-off under different operational conditions are given next.

#### C. CS system results

The system's configuration determines the power consumption and the confidence on the HR estimation. While some of the configuration parameters are determined during the design phase (number of frequency bins for LSP, signal input frequency etc.), other parameters can be modified at runtime (CR). The number of bins to be computed for LSP determine the minimum resolvable HR and the complexity of the hardware. While reducing the number of bins is favorable for reduced latency and energy consumption, the accuracy and minimum resolvable HR are affected. For a typical frequency range of 0.5Hz – 3.5Hz, which encompasses a HR range of 30bpm – 210bpm, a 32 point LSP gives a resolution of 6 bpm, which is higher than the requirement of 5 bpm specified by ANSI-AAMI standards for HR monitors [12]. A 64 point LSP was therefore chosen for the implementation. To quantify the accuracy of HR detection with respect to number of bins, an input sinusoid, which serves as a proxy for PPG signal, with a frequency of 1.2 Hz (corresponding to 72 bpm) is randomly sub-sampled with varying CR (8x, 10x and 30x). An LSP is then performed to estimate the HR from the acquired signal with number of bins varying from 32 to 512. Fig. 8(a) shows that a higher number of bins does not lead to significant improvement in the accuracy of estimated HR. The performance of HR estimation for different CRs is simulated by randomly sub-sampling a sinusoid, whose frequency is swept from 0.5 Hz - 3.5 Hz, with CRs 8x, 10x and 30x. A 64 bin, fixed point LSP is then performed on the randomly sub-




Fig. 8: (a) Estimated HR with varying number of LSP bins at different CRs (b) Estimated HR obtained from CS sinusoidal signal whose frequency is swept from 0.5 Hz to 3.4 Hz through LSP.




Fig. 9: Significance Interval and LED drive's power consumption of the estimated HR when CR is varied from 2x to 40x.

sampled signal and HR is estimated from the resulting PSD. Fig. 8(b) shows the extracted HR for different CRs, which indicates that the maximum error in estimated HR is less than 7 bpm at 30x compression for a nominal HR of 120 bpm. This error is well within the accuracy requirements specified in [12].

In order to implement the dynamically controlled loop of Figure 2, the effect of CR on the confidence of the estimated HR must be closely analyzed. Fig. 9 shows how the significance of the estimated peak in the PSD depends on CR. The LED driver power consumption of the CS PPG acquisition system presented in [13] at different CRs is also shown. The significance on HR estimation is high (> 90%) for low CRs but it drops significantly at higher CRs. This translates to uncertainty over the true presence of the tone at a specific frequency in the signal and hence uncertainty over the HR estimation. It should however be noted that the power required to drive the LED scales linearly with CR, from 1200  $\mu W$  to 43  $\mu W$  for CR of 1x and 40x respectively. Under such scaling opportunities the system is capable of preventing significant losses in the confidence of HR estimation for a LED drive power savings of more than 30 $\times$ . This allows the self-tunable system (Fig. 2) to dynamically optimize the power consumption with respect to the required accuracy and confidence.

## V. CONCLUSION

The feature-tunable information extraction paradigm proposed in this work targets the operational limitations of power

constrained always-awake sensing systems, by dynamically computing only the relevant information bearing elements in the incoming data stream. In this system, a feedback loop selects the optimal set of feature/extracted information specifications in function of environmental conditions and application accuracy constraints. Two circuit realizations present significant energy savings. 1) Up to 3 orders of magnitude in the feature extraction unit of a VAD analog-to-information with precision-scalable NB classifier. 2) Up to 30x in the LED driver power consumption of a variable rate CS PPG readout, without significant accuracy losses.

## VI. ACKNOWLEDGEMENTS

This work has been supported by the EU ERC project RESENSE under grant agreement ERC-2016-STG-715037.

## REFERENCES

- [1] M. Verhelst and A. Bahai, "Where analog meets digital: Analog to information conversion and beyond," *IEEE Solid-State Circuits Magazine*, vol. 7, pp. 67–80, Summer 2015.
- [2] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21–30, 2008.
- [3] K. Badami, S. Lauwereins, W. Meert, and M. Verhelst, "24.2 context-aware hierarchical information-sensing in a 6  $\mu W$  90nm cmos voice activity detector," in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, pp. 1–3, Feb 2015.
- [4] K. M. H. Badami, S. Lauwereins, W. Meert, and M. Verhelst, "A 90 nm cmos, 6  $\mu W$  power-proportional acoustic sensing frontend for voice activity detection," *IEEE Journal of Solid-State Circuits*, vol. 51, pp. 291–302, Jan 2016.
- [5] K. Uyttenhove and M. S. J. Steyaert, "Speed-power-accuracy tradeoff in high-speed cmos adcs," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 49, pp. 280–287, Apr 2002.
- [6] N. Friedman, D. Geiger, and M. Goldszmidt, "Bayesian network classifiers," *Journal of Machine Learning*, vol. 29, no. 2, pp. 131–163, 1997.
- [7] L. Galindez, W. Meert, H. Bruyninckx, and M. Verhelst, "Extending naive bayes with precision-tunable feature variables for resource-efficient sensor fusion," in *2nd Workshop on Artificial Intelligence and Internet of Things co-located with ECAI 2016*, pp. 23–30.
- [8] J. G. Webster, *Design of pulse oximeters*. Institute of Physics Bristol PA, 1997.
- [9] V. R. Pamula, M. Verhelst, C. Van Hoof, and R. F. Yazicioglu, "A novel feature extraction algorithm for on the sensor node processing of compressive sampled photoplethysmography signals," in *SENSORS, 2015 IEEE*, pp. 1–4, IEEE, 2015.
- [10] H. Mamaghanian, N. Khaled, D. Atienza, and P. Vandergheynst, "Design and exploration of low-power analog to information conversion based on compressed sensing," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, no. 3, pp. 493–501, 2012.
- [11] F. Ren and D. Marković, "18.5 A configurable 12-to-237ks/s 12.8 mW sparse-approximation engine for mobile ExG data aggregation," in *2015 IEEE International Solid-State Circuits Conference-(ISSCC) Digest of Technical Papers*, pp. 1–3, IEEE, 2015.
- [12] ANSI/AAMI-EC13, "American national standards for cardiac monitors, heart rate meters and alarms," 2002.
- [13] P. V. Rajesh, J. M. Valero-Sarmiento, L. Yan, A. Bozkurt, C. Van Hoof, N. Van Helleputte, R. F. Yazicioglu, and M. Verhelst, "22.4 A 172 $\mu W$  compressive sampling photoplethysmographic readout with embedded direct heart-rate and variability extraction from compressively sampled data," in *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 386–387, IEEE, 2016.