A 290 mV Sub- $V_{\rm T}$ ASIC for Real-Time Atrial Fibrillation Detection

Oskar Andersson, Student Member, IEEE, Ki H. Chon, Senior Member, IEEE, Leif Sörnmo, Senior Member, IEEE, and Joachim Neves Rodrigues, Senior Member, IEEE

Abstract—A real-time detector for episodes of atrial fibrillation is fabricated as an application specific integrated circuit (ASIC). The basis for detection is a set of three parameters for characterizing the RR interval series, i.e., turning point ratio, root mean square of successive differences, and Shannon entropy. The developed hardware architecture targets ultra-low voltage operation, suitable for implantable loop recorders with ultra-low energy requirements. Algorithmic and architectural optimizations are performed to minimize area and energy dissipation, with a total area footprint reduction of 44%. The design is fabricated in 65-nm CMOS low-leakage high-threshold technology. Measurements with aggressively scaled supply voltage (V_{DD}) in the subthreshold (sub- V_T) region show energy savings of up to 41 X when operating at 1 kHz with a V_{DD} of 300 mV compared to a nominal V_{DD} of 1.2 V.

Index Terms—Atrial fibrillation, loop recorder, low-power, sub-threshold, ultra-low voltage.

I. INTRODUCTION

N ONINVASIVE methods for detection of atrial fibrillation (AF) have a long history which began several decades ago [1]–[8]. The vast majority of proposed detectors rely on information conveyed by the ventricular response, i.e., the RR interval series. The main reason for relying on ventricular information only is evidently that the atrial activity exhibits a much lower amplitude than does the ventricular activity, and therefore accurate atrial measurements are more difficult to produce. The analysis of atrial activity becomes particularly problematic in ambulatory recordings because of a noise level which generally is much higher than that of recordings made at rest.

The use of a handheld device or a smartphone for AF detection has recently received much attention since it facilitates home-based screening at low or no extra cost [9]–[12]. For example, the detector considered in the present study was previ-

Manuscript received March 17, 2014; revised June 08, 2014; accepted August 05, 2014. Date of publication October 16, 2014; date of current version May 22, 2015. This work was supported in part by Swedish Research Council (621-2011-4540), and Swedish VINNOVA Industrial Excellence Centre (SoS). K. Chon was supported by a grant from NIH R15HL121761. This paper was recommended by Associate Editor J. Georgiou.

O. Andersson and J. Neves Rodrigues are with the Department of Electrical and Information Technology, Lund University, Lund SE-221 00, Sweden (e-mail: oskar.andersson@eit.lth.se; joachim.rodrigues@eit.lth.se).

L. Sörnmo is with the Department of Biomedical Engineering and Center for Integrative Electrocardiology, Lund University, Lund SE-221 00, Sweden (e-mail: leif.sornmo@bme.lth.se).

K. H. Chon is with the Department of Biomedical Engineering, University of Connecticut, Storrs, CT 06269-3237 USA (e-mail: kchon@engr.uconn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TBCAS.2014.2354054

ously implemented in software operating on the processor of a smartphone, the detector input being the pulsatile signal measured on the human fingertip using the built-in camera and flash [9]. This type of device provides valuable intermittent information on AF occurrence which may be acquired on a regular basis during the day or when symptoms are experienced.

The length of an ambulatory ECG recording is usually limited to 24 hours, and consequently, only limited knowledge is provided on the progression of AF when considering that the time course of AF from silent to a sustained form may span several years [13]. Furthermore, concerning length, ambulatory recordings are not particularly effective for assessing the outcome of different therapies, e.g., cardioversion and antiarrhythmic drugs. With the use of implantable loop recorders, however, it is possible to continuously monitor the burden of AF during extended time periods, using an algorithm that performs AF detection.

While an AF detector, originally developed for use in a noninvasive environment, may be a candidate for implementation in an implantable loop recorder, the energy requirements associated with the detector are likely to be incompatible with available battery capacity. Therefore, the starting point for detector development depends on the operating environment: an "invasive" AF detector has a structure which, from a computational viewpoint, need to be less costly than that of a "noninvasive" detector. It has recently been pointed out that further improvement of AF detection in implantable loop recorders is needed since existing ones do not offer satisfactory performance with respect to sensitivity/specificity, or are limited by their computational complexity [14].

The implementation of an invasive AF detector not only involves high requirements on detection performance, but it is equally important to minimize energy dissipation in both idle and active mode. Idle energy is dominated by the leakage drawn by memories retaining data, and active energy is minimized by reducing complexity of computations. For AF detectors operating on the RR intervals series, and thus, have a low incoming data rate, computational complexity is less of a concern than the minimization of required memory.

Since a key requirement of this study is to minimize memory cost, the detector by Dash *et al.* [5] is chosen for implementation since it does not require storage of training data. The RR intervals are characterized by the turning point ratio, the root mean square of successive differences, as well as the Shannon entropy. Inspired by the good performance reported in [5], an investigation at the algorithmic and architectural levels was car-

1932-4545 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications standards/publications/rights/index.html for more information.

ried out with the goal to develop an appropriate design for hardware (H/W) implementation.

Detectors suitable for implementation in a chronic implantable monitor include the ones proposed by Sarkar *et al.* [15], see also [14] and [16]. These detectors explore either the series of differenced, successive RR intervals, or both the original and differenced RR interval series. Aspects on H/W implementation appear not to have been treated in any subsequent publication.

The RR interval series, used as input to the AF detector, can be delivered from an cardiac event detector, e.g., [17] or [18]. These two approaches differ in terms of flexibility, where the approach in [17] offers a flexible solution to process a complex detection algorithm at a cost of higher power consumption and area footprint. In contrast, the proposed design in [18] offers a highly energy efficient solution but lacks flexibility in terms of signal processing. Both implementations are suitable candidates to be used with the proposed AF detector, as they can be operated at ultra-low voltages (ULV).

Contribution: An AF detector architecture was optimized for ULV operation. Algorithmic and architectural improvements are performed, i.e., resource sharing of arithmetic units, and adoption of an advantageous preprocessing scheme that reduces the memory capacity requirement. An ASIC was fabricated and measurements confirm simulation results. To the best of the authors' knowledge, this paper is the first to detail aspects of H/W implementation and measurements from an AF detector.

This paper is organized as follows. Section II gives a brief summary of the AF detector structure, and its performance is compared to that of other detectors. Section III describes the hardware architecture of the system, including different optimization issues. Hardware performance in terms of area, speed, and energy dissipation is presented in Section IV, and followed by a discussion in Section V. Conclusions are found in Section VI.

II. METHODS

The AF detector developed by Dash *et al.* is here considered for ASIC implementation, and is therefore briefly described in the following; a detailed description is found in [5]. The detector processes successive, contiguous segments of the RR interval series, where the N RR intervals of a segment are denoted $r(0), \ldots, r(N-1)$.

Prior to AF detection, ectopic beats are eliminated by preprocessing r(n) using a simple interval-based criteria. An ectopic beat is detected when the ratio r(n)/r(n-1) belongs to the shortest 1% of all ratios and the interval ratio r(n+1)/r(n) belongs to the longest 1%. The two intervals adjacent to the ectopic beat are excluded from r(n). Similar percentile criteria are employed for excluding long RR intervals. The preprocessed RR intervals are denoted as $r_p(n)$.

In [5], the 25*th* percentile was also analyzed, however, computation of this percentile requires extra memory, and was therefore omitted in the final implementation of the present study; this omission did not lead to any degradation in detection performance.

A. Detector Structure

The detector structure embraces three parameters for characterizing a segment of the RR interval series, namely, the turning point ratio (TPR), the root mean square of successive differences (RMSSD), and the Shannon entropy (SE). The three parameters are subjected to threshold detection.

A turning point in the RR interval series is said to occur at $r_p(n)$ whenever

$$r_p(n-1) > r_p(n) < r_p(n+1)$$
 or
 $r_p(n-1) < r_p(n) > r_p(n+1)$ (1)

where $r_p(n)$ denotes the preprocessed RR intervals. The TPR is defined as

$$TPR = \frac{N_{TP}}{N}$$
(2)

where N_{TP} indicates the number of turning points in a segment. A nonparametric statistical test is employed in which the number of turning points is compared to the confidence interval of the expected number of turning points, serving as a means to quantify the degree of randomness of the RR interval series. The RMSSD is defined by

RMSSD =
$$\sqrt{\frac{1}{N-1} \sum_{n=0}^{N-2} (r_p(n+1) - r_p(n))^2}$$
. (3)

The Shannon entropy is defined by

$$SE = \frac{1}{\log_2 N_b} \sum_{k=1}^{N_b} \hat{p}(k) \log_2 \hat{p}(k)$$
(4)

where $\hat{p}(k)$ is a probability estimated from a RR interval histogram once outliers have been excluded; outliers are defined as the N_o longest and N_o shortest intervals. The histogram is constructed from the remaining RR intervals by sorting them into N_b equally spaced bins, with limits defined by the shortest and longest RR intervals. The probability is estimated by

$$\hat{p}(k) = \frac{N_k}{N - 2N_o} \tag{5}$$

where N_k denotes the number of RR intervals in the k^{th} bin. The same values as in [5] were used, i.e., $N_b = 16$, $N_o = 8$.

Atrial fibrillation is detected in a segment when the following three conditions are jointly fulfilled:

$$\eta_{\rm TPR_{high}} > {\rm TPR} > \eta_{\rm TPR_{low}},$$
 (6)

$$\frac{\text{LMSSD}}{-} > \eta_{\text{RMSSD}},\tag{7}$$

$$E > \eta_{SE}$$
 (8)

where $\eta_{\text{TPR}_{\text{low}}}, \eta_{\text{TPR}_{\text{high}}}, \eta_{\text{RMSSD}}$, and η_{SE} denote fixed thresholds, and \overline{r} denotes the mean RR interval length.

TABLE I DATABASES USED FOR DETECTOR EVALUATION, CONTAINING AF AND OTHER ARRHYTHMIAS [19]

Database	#Records	Duration	AF
MIT-BIH AF (Atrial Fibrillation)	25	10 h	Yes
MIT-BIH Arrhythmia (100-series)	25	30 min	No
MIT-BIH Arrhythmia (200-series)	23	30 min	Yes

TABLE II RECOMMENDED THRESHOLD VALUES FOR THE DETECTOR

Threshold	Value
N	128
$\eta_{\mathrm{TPR}_{\mathrm{low}}}$	0.46
$\eta_{\mathrm{TPR}_{\mathrm{high}}}$	0.84
η_{RMSSD}	0.15
η_{SE}	0.84

B. Update of Preprocessing Parameters

In the present study, preprocessing of r(n) is considered in two variants:

- one in which the percentile limits are both computed from and applied to the *current* segment (CurrSeg), and
- another computed from the *preceding* segment and applied to the current segment (PrecSeg).

The latter variant is proposed and investigated since it leads to substantial memory savings. In its original description, the entire RR interval series was preprocessed off-line prior to AF detection [5], whereas the present approach performs preprocessing in almost real-time, i.e., with a delay of one segment.

C. Detector Evaluation

The performance of the AF detector is studied with a bitaccurate fixed-point implementation, and validated using the same ECG databases as those in [5], see Table I. The performance measures sensitivity (Se) and specificity (Sp) are defined as

$$Se = \frac{N_{TP}}{N_{TP} + N_{FN}} \tag{9}$$

and

$$Sp = \frac{N_{TN}}{N_{FP} + N_{TN}} \tag{10}$$

where N_{TP} denotes the number of correctly detected AF episodes ("true positives"), N_{TN} the number of correctly detected non-AF episodes ("true negatives"), N_{FP} the number of falsely detected AF episodes ("false positives"), and N_{FN} the number of falsely detected non-AF episodes ("false negatives").

The four thresholds $\eta_{\text{TPR}_{\text{low}}}$, $\eta_{\text{TPR}_{\text{high}}}$, η_{RMSSD} , η_{SE} , and the segment length N were all varied over a grid identical to that used in [5], to determine those values which produced a specificity equal to or better than that in [5]. The grid search was performed for both CurrSeg- and PrecSeg-based percentiles.

TABLE III AF DETECTOR PERFORMANCE. SENSITIVITY IS NOT COMPUTED FOR THE 100-SERIES DUE TO THE ABSENCE OF AF EPISODES

Database	Configuration	Sensitivity	Specificity	
MIT-BIH AF	Present impl., PrecSeg	94.9 %	95.8 %	
	Present impl., CurrSeg	90.6 %	97.6 %	
	[5]	94.4 %	95.1 %	
	[2]	94.4 %	97.2 %	
MIT-BIH Arrhythmia (100-series)	Present impl., PrecSeg	N/A	97.0 %	
	Present impl., CurrSeg	N/A	99.4 %	
	[5]	N/A	96.2 %	
	[2]	N/A	99.9 %	
MIT-BIH Arrhythmia (200-series)	Present impl., PrecSeg	98.8%	79.0 %	
	Present impl., CurrSeg	91.0%	81.8 %	
	[5]	96.5 %	69.4 %	
	[2]	88.2 %	87.6%	



Fig. 1. Detector architecture in the context of AF classification. The AF detector uses the three statistical measures: TPR, RMSSD, and SE. The FIFO stores new RR intervals that arrive while the RMSSD and SE are operating. All three measures need to indicate positive for AF.

The resulting parameter values are referred to as "recommended values;" see Table II.

The sensitivity and specificity results are presented in Table III, showing that the present implementation offers a slight improvement over [5]. This improvement is observed for all three databases, and is due to a minor alteration in the preprocessing and the re-evaluation of threshold values. Both sensitivity and specificity increase when PrecSeg is used, whereas an increase in specificity and a decrease in sensitivity is observed for CurrSeg.

III. HARDWARE ARCHITECTURE

The algorithm is designed with a focus on low energy dissipation, and the architecture is optimized with knowledge learnt from detector evaluation. Design choices are taken to



Fig. 2. Conceptual architectural implementation that illustrates how the blocks of the preprocessing module interact while using the PrecSeg scheme. There exists one instance of each element which is reused over time, i.e., time-multiplexed. The ratios of the previous segment are used to compute p_1 and p_{99} . The first segment of RR intervals is discarded in the PrecSeg preprocessing scheme as p_1 and p_{99} are not initialized yet. Delay element is indicated with D.

achieve a small area footprint and reduced memory capacity. The memory is facilitated using a random-access memory (RAM) implemented with standard cells in order to allow for aggressive supply voltage scaling down to the subthreshold region [20].

The detector architecture is displayed in Fig. 1. Once available, RR intervals are preprocessed for ectopic beats. Thereafter, RR intervals are used for computing the detection parameter TPR, as well as stored in the RAM, for use in the other detection parameters, i.e., RMSSD and SE. If TPR indicates AF, the RMSSD computation is started, which in turn triggers SE, if AF is detected. RR intervals, which become available during RMSSD or SE processing, are stored in a first-in first-out (FIFO) memory and sent to the RAM afterwards. The three detection parameters rely on a time-multiplexed CORDIC divider, shared multiplier (SM) and a time-multiplexed log₂ implementation (LOG-2). The different thresholds are applied in order to determine whether an AF episode is present.

A. Preprocessing

As described in Section II, ectopic beats need to be removed from the RR interval series r(n). Furthermore, two preprocessing methods are proposed, CurrSeg and PrecSeg, that differ with respect to the segment from which percentiles used for filtering are computed. The architecture of the implementation using the PrecSeg scheme is shown in Fig. 2. The two ratios r(n)/r(n-1) and r(n+1)/r(n) are computed in the processing block 'Ratios' using the 'Divider' block. These ratios are compared with the 1% (p_1) and 99% (p_{99}) percentiles (of the previous segment) in the 'Ectopic Filter' block, to approve or discard r(n). In order to approve r(n) already when r(n+1)has arrived (i.e., perform on-the-fly processing), the values of p_1 and p_{99} are computed from the ratios of the previous segment. Therefore, the two largest and smallest ratios of a segment are stored and used to compute p_1 and p_{99} for the next segment. To clarify, the steps performed whenever a sample arrives are:

• r(n)/r(n-1) and r(n+1)/r(n) are computed.



Fig. 3. Architectural overview of Turning Point Ratio, which characterizes the randomness of the RR interval series. The shared multiplier is denoted SM, and a delay element with D.

- The two largest and smallest ratios are updated.
- r(n-1) is classified using p_1 and p_{99} of the previous segment.
- When 128 samples of r(n) have arrived, p_1 and p_{99} are updated.

Consequently, the ratio r(n)/r(n-1) is no longer needed (discarded) once r(n) is classified, and thus, no memory is needed to store these ratios.

In the second variant of the preprocessing scheme CurrSeg, the 'Ectopic Filter' compares the ratios of r(n) with the percentiles $(p_1 \text{ and } p_{99})$ of the current segment. This requires that the computed ratios of r(n) of an entire segment are stored in an additional RAM as the ratios are first used to compute the percentiles $(p_1 \text{ and } p_{99})$ and afterwards used in the 'Ectopic Filter' to filter r(n). The preprocessing scheme PrecSeg does not require any extra storage and is therefore chosen.

B. Turning Point Ratio

The TPR is computed by comparing $r_p(n)$, $r_p(n-1)$, and $r_p(n-2)$ to each other; see Fig. 3. The three RR intervals are stored in delay elements (*D* flip-flops). The accumulated value is divided with *N*, where *N* denotes the segment size, using a shared multiplier described below.

C. Root Mean Square of Successive Differences

The implementation of the RMSSD is costly in terms of area footprint due to the square root and division operations are expensive in H/W implementation. Therefore, to simplify hard-



Fig. 4. Architectural overview of Root Mean Square of Successive Differences, which flags for AF when a high variability is present.



Fig. 5. Shannon Entropy architecture, which flags positive when a low degree of periodicity is found.

ware implementation, the square root operation is eliminated by squaring both sides of (7), and thus, yielding

$$\frac{1}{\eta_{\text{RMSSD}}^2(N-1)} \sum_{n=0}^{N-2} (r_p(n+1) - r_p(n))^2 > \overline{r}^2.$$
(11)

The terms $(r_p(n+1)-r_p(n))^2$ and \overline{r} are both accumulated using partial sums and normalized by $\eta^2_{\text{RMSSD}}(N-1)$ and N, respectively. N indicates the segment size. Partial sums and normalization is utilized in order to restrict the required word length; see Fig. 4. The normalized values are then accumulated and \overline{r} is squared to form (11).

D. Shannon Entropy

The SE computation is accomplished by three different modules, illustrated in Fig. 5: removal of outliers, computation of the RR interval histogram, and computation of the sum in (4). The block 'Outlier Removal' identifies the N_o longest and N_o shortest RR intervals of the current segment, and uses the resulting values to determine the upper and lower limits of the histogram. The outliers are detected by iterating through the current segment, which is stored in the RAM. The RAM addresses of the outlier candidates are stored in a shift register of size $2(N_o+1) \times 7$ bits, thus a total of $2(N_o+1) = 18$ memory positions are used, which each require 7 bits. A number of $N_o + 1$ is required as the 9th longest and 9th shortest intervals are needed to determine the boundaries of the histogram.

The block 'Histogram Computation' constructs a histogram using the limits determined by the first module. The histogram requires a storage capacity of $N_b \times 8$ bits to account for the possible outcomes of the histogram computation. RR intervals are placed in their appropriate bin by iterating through the RAM and increasing the comparison in steps of the distance between bins. H/W optimization is performed in terms of memory minimization by reusing the storage (Outlier & Histogram Storage)



Fig. 6. Logarithm architecture, used within the Shannon Entropy module. Shift register is denoted as SR.

used by the first module, after modification into the biggest requirement for both, i.e., $2(N_o + 1) = 18$ positions of size 8 bits. Re-usage is possible as the boundaries of the histogram are defined from the excluded outliers. Therefore, all outliers will reside in the boundary bins, and are easily subtracted, i.e., excluded from SE computation. This is performed instead of tracking all outlier addresses which would need an additional storage of 112 bits. The block 'SE Computation' computes the SE by iterating over the bins and determining the estimated probabilities in (5), through inverse multiplication as the denominator is constant. These estimated probabilities are used to compute (4) using multiplications are performed in a shared multiplier. Division with $\log_2 N_b$ can be replaced by a logic shift since $N_b = 16$.

E. Logarithm Computation (LOG-2)

For the logarithm function a time-multiplexed algorithm is chosen to avoid the use of a look-up table (LUT), and thus, reduces memory cost [21]. This avoidance is permissible since the requirement on computational speed of the AF detector is low. The original algorithm operates on floating-point numbers, however, for a H/W implementation fixed-point representation is preferred. Thus, the incoming fixed-point number x is transformed into an exponent e and normalized mantissa m within the range [1 2). Multiplication and division with power of two is implemented as logic shift and truncation, respectively. This transformation is illustrated in the upper part of Fig. 6; the lower part account for the computation of $\log_2 m$, i.e., realization of (15) and (16) in the Appendix. The multiplication required for squaring is realized using a shared multiplier. The intermediate result, i.e., $\log_2 m$, is stored in a shift register and added to the exponent e to produce the result of the logarithm computation, i.e., $\log_2 x$, see (12).

F. Shared Multiplier

A total of 8 multiplications are needed for the computation of the TPR, RMSSD, SE, and LOG-2. In order to reduce area cost, a time-multiplexing scheme that facilitates a generic multiplier is chosen. A relaxed timing constraint makes resource sharing of a single multiplier feasible, see Fig. 7. The output data is



Fig. 7. Multiplexing network for the shared multiplier unit, used in all three measures.

rounded to the same word length as the input data, and overflow can be avoided since the input data is known.

G. Hardware Optimizations

The detector is optimized on both algorithmic and architectural levels to minimize area footprint and energy dissipation. This is achieved by minimizing memory cost, reducing internal word length, and adopting resource sharing and time-multiplexing for arithmetic operations. The total amount of processing is reduced by exploring the property that AF is only detected if all criteria are jointly fulfilled, i.e., SE is only computed if the RMSSD exceeds η_{RMSSD} , and the RMSSD only computed if the TPR is within the interval defined by $\eta_{\text{TPR}_{\text{low}}}$ and $\eta_{\text{TPR}_{\text{high}}}$.

1) Memory Minimization Technique: Table III shows that the best detection performance is obtained for PrecSeg, which therefore is adopted in the present implementation. By computing percentiles from the preceding segment, ratios are not stored but rather compared to the percentiles computed of the preceding segment and used to compute new percentiles. Avoiding storage of these ratios eliminates the need for storing an entire segment of ratios, i.e., $N \times W$ bits, where W denotes word length. This is the same size as the RAM needed to store the RR intervals of one segment, i.e., RAM storage is reduced by 50% by including only one of the RAMs.

2) Resource Sharing and Time-Multiplexing: With several multiplications required in the modules for TPR, RMSSD, SE, and LOG-2, a shared multiplier leads to reduced area footprint. Furthermore, the adoption of a time-multiplexed algorithm for the divisions required in the preprocessor will reduce the area footprint relative to straightforward division. As aforementioned, the logarithm function is implemented using an area efficient time multiplexing and does not rely on a LUT, which further reduces the area footprint. The speed penalty due to time-multiplexing and resource sharing does not degrade detector performance in terms of sensitivity and specificity due to the low rate of incoming samples, i.e., RR intervals.

3) Internal Word Length Reduction: The word length is selected by studying detection performance for different fractional bits in a bit-accurate fixed-point implementation; see Fig. 8. It is observed that the performance is largely constant for 12 or more fractional bits, whereas it deteriorates for fewer bits due to precision mismatch. Hence, the word length is chosen to be 12 fractional bits. The rational part of the word is internally from



Fig. 8. Word length evaluation (for MIT-BIH databases), with the chosen precision encircled, diminishing returns after 12 bits, where *Se* denotes sensitivity and *Sp* denotes specificity.

TABLE IV Area Cost After Gate Mapping (Synthesis) in 65-nm Low-Leakage High-Threshold CMOS Technology

Design	Configuration	Area [μ m ²]	NAND2 eq.
Shared Multiplier	PrecSeg	50,000	24,038
Dedicated Multipliers	CurrSeg	88,300	42,452

1 to 5 bits, and the data storage for an RR interval is 1 rational bit; see RAM in Fig. 1. Using N = 128, 12 instead of 16 fractional bits, and 1 rational bit, the RAM capacity reduces from 2176 to 1664 bits, i.e., a reduction of 23.5%.

IV. HARDWARE IMPLEMENTATION

The proposed architecture operates at a clock frequency of 1 kHz, and consequently, energy dissipation is dominated by leakage currents. Therefore, design considerations to reduce leakage are carried out while mapping the H/W architecture to silicon for manufacturing, e.g., clock-tree buffer sizing, gates with low fanout, and regular buffer sizing. The proposed architecture is implemented using 65-nm CMOS low-leakage high-threshold technology. The design constraints are on reduced leakage current cost and area footprint, while timing constraints are relaxed. The area cost reported after gate-level synthesis of using either PrecSeg or CurrSeg is presented in Table IV, together with the area gain of sharing a single multiplier. For comparison the NAND2 equivalents are stated as well. The H/W optimization results in a reduction of 43%, by combining the shared multiplier and PrecSeg.

A. Silicon Measurements

The chip is manufactured with separate supply voltages $(V_{\rm DD}s)$ for the RAM and the remaining AF detector (Core). Thereby, it is possible to investigate individual shares on the power consumption; the measurement results are presented in Table V. The total power consumption of the core is measured as $P_{\rm total(Core)}$, where leakage power $P_{\rm leak(Core)}$, at nominal $V_{\rm DD}$ (1.2 V), accounts for ~ 93%. Similarly, ~ 99% of the memory's total power consumption $P_{\rm total(Mem)}$ is leakage power. Consequently, to efficiently reduce the power consumption, the issue of leakage power needs to be seriously addressed. The most effective method to reduce both dynamic- and leakage

 TABLE V

 Measured Energy Dissipation With a Fixed Frequency of 1 kHz and A Scaled V_{DD} , Together With Gain Compared to Nominal V_{DD}

	C	ore	Memory		Core + Memory	
$V_{\rm DD}$ [V]	$P_{\text{leak}(\text{Core})}[nW]$	$P_{\rm total(Core)}[nW]$	$P_{\text{leak}(\text{Mem})}[nW]$	$P_{\rm total(Mem)}[nW]$	$P_{\text{total}(\text{Core}+\text{Mem})}[nW]$	Gain
1.2	216.7	233.4	99.4	100.5	333.9	N/A
0.7	24.9	30.0	10.6	11.4	41.4	8.1 X
0.4	7.8	9.4	3.2	3.5	12.9	25.7 X
0.3	3.9	5.9	2.1	2.2	8.1	41.1 X



Fig. 9. Chip microphotograph of the fabricated AF detector. Area footprint is 0.1 mm^2 (excluding pads).

power consumption is to lower $V_{\rm DD}$ to ultra-low voltages (ULV), and operate in the near-threshold voltage (NTV) region or the subthreshold (sub- $V_{\rm T}$) region, where $V_{\rm T}$ is the threshold voltage. The NTV region is referred to as the region in the vicinity of the $V_{\rm T}$ of the transistor, around 500–700 mV in this technology. The sub- $V_{\rm T}$ region is the operation domain below 500 mV. Operating in these regions has the drawback of a reduced maximum operational frequency, down to tens of MHz (NTV region) or hundreds of kHz (sub- $V_{\rm T}$ region).

Advantageously, the operational frequency of the AF detector is 1 kHz, and thus, this design is a suitable candidate for voltage scaling all the way down to sub- $V_{\rm T}$ operation. Design considerations result in a large slack on the critical path, and thus, process variations seen at aggressively scaled $V_{\rm DD}$ do not endanger timing. Furthermore, regular static random-access memories (SRAMs) are not operational at a scaled $V_{\rm DD}$. Therefore, the RAM is implemented using standard cells that are operational in the NTV to sub- $V_{\rm T}$ regions, as shown in [22], [23]. By lowering the $V_{\rm DD}$ down to the sub- $V_{\rm T}$ region, i.e., 300 mV, a gain of up to 41 × is seen; see Table V. Furthermore, the optimized preprocessing scheme PrecSeg, which reduces the memory capacity, results in power savings of $\sim 23 - 21\%$ if operated at nominal $V_{\rm DD}$ or 300 mV.

B. On-Chip Verification

Fig. 9 shows a microphotograph of the fabricated chip. Due to the small core area only a limited amount of pads were available for this design, (which was part of a multi-project die). Therefore, an on-chip peripheral circuit was incorporated, performing the asynchronous serial (single-bit) input and output communication, and converting parallel input for the AF detector. The on-chip peripheral circuit is operating at a $\Delta V_{\rm DD}$ of 140 mV higher than $V_{\rm DD}$ of the AF detector, acting as a level shifter and a buffer for the output signals. The chip is stimulated with the same test patterns used for detector evaluation, i.e., the RR intervals of the ECG databases in Table I. In order to observe the



Fig. 10. Post-processed measured chip output: the arrows indicate when the statistical measures are turned on/off. One output value is received per segment of 128 RR intervals.



Fig. 11. Oscilloscope screenshot illustrating input stimuli together with output. Peripheral serial communication is shown at the bottom.

functionality of the system, intermediate data signals (i.e., approved RR intervals, TPR, RMSSD and SE results) are sent over the serial interface. Post-processing of the intermediate data is shown in Fig. 10. Additionally, an oscilloscope screenshot together with the peripheral interface is shown in Fig. 11, where RR intervals are supplied as input while the circuit sends out approved RR intervals. The screenshot was taken when the AF detector operated at the lowest $V_{\rm DD}$ of 290 mV. The total silicon area, including the peripheral circuit, is 0.10 mm², seen in Fig. 9. The increase in area compared to after gate-level synthesis is due to a high congestion when routing signal wires.

C. System Perspective

To quantify the power consumption of the detector, the requirements of an implantable pulse generating pacemaker are considered [24]. The battery used in such a device has a typical rating of 2 Ah in energy capacity and delivers 2.5 V, which translates to a total energy capacity of 18 kJ [24]. With the assumption that this device is operated for a total of 10 years, an average energy budget of $57\mu J/s$ ($57\mu W$) is available. Assuming that efficient DC-DC converters are available (to reduce V_{DD} to 1.2 V or 0.3 V) the power consumption of the detector is calculated to 0.58% and 0.014% of the total power consumption for a $V_{\rm DD}$ of 1.2 V and 0.3 V, respectively. Thus, this represents a minor contribution compared to other more power consuming components of the system, such as analog sensing devices, A/D converters and radio transmitters. As a result, the AF detector can become a supplement to a health monitoring system, adding relatively small area and power overhead.

V. DISCUSSION

The present study provides detailed information on the implementation of an AF detector, and demonstrates that a detector with a rather complex structure can be considered for use in an implantable loop recorder, e.g., a complete system on an application specific integrated circuit (ASIC) or field programmable gate array (FPGA), depending on application requirements. The results suggest that the energy required for long-term operation of the detector, i.e., for several years, is well within what is provided by the battery of an implantable device.

The performance of the present block-based RR interval analysis was also compared to that of sliding window analysis (not presented). Since detection performance was found to be about the same for the two types of segmenting in RR interval analysis, block-based analysis (PrecSeg) was preferred as it is associated with fewer computations and thus lower power consumption.

From an algorithmic viewpoint, it is experimentally established that RMSSD is the more powerful parameter in determining whether an AF episode has occurred or not. From a low power H/W viewpoint, however, TPR is the least expensive in terms of power, and therefore TPR is computed first to reduce the amount of RMSSD computations.

VI. CONCLUSION

An ultra-low energy ASIC for real-time AF detection is presented. Certain alterations of the preprocessing are proposed that improve detection performance and reduce area cost by reducing memory footprint. Functional units are time-multiplexed to enable sharing of commonly-used resources, i.e., multiplication, division, and logarithm. The optimized detector is fabricated in 65-nm low-leakage high-threshold CMOS process. Measurements at aggressively scaled supply voltage down to ultra-low voltages, i.e., subthreshold (sub- V_T), show substantial energy savings compared to nominal supply voltage.

APPENDIX COMPUTATION OF THE LOGARITHM

Logarithms in H/W are often implemented using LUT to minimize computation time. However, due to the relaxed time requirements and memory limitations of the present implementation, another approach is adopted which avoids LUT [21]. This approach operates on a normalized mantissa m from a floating-point number in the range [1 2). The exponent e of the floating-point number is not needed to process according to the rule

$$\log_2 x = e + \log_2 m \tag{12}$$

where x is a real number. Since m is within [1 2), the result of the logarithm operations resides within [0 1), expressed as

$$\log_2 m = 0.b_1 b_2 \cdots b_i \cdots b_n \tag{13}$$

where b_i indicates a bit position in a word, and n is the word length. Rewriting (13), the corresponding anti-logarithm is

$$\log_2 m = b_1 2^{-1} + b_2 2^{-2} + \dots + b_i 2^{-i} + \dots + b_n 2^{-n}$$
$$m = 2^{b_1 2^{-1} + b_2 2^{-2} + \dots + b_i 2^{-i} + \dots + b_n 2^{-n}}.$$
(14)

To compute the value of b_1 , the mantissa m is squared, and henceforth, referred to as m_1 , which after restructuring (14) becomes

$$m_1^2 = 2^{b_1} 2^{b_2} \cdots 2^{b_i} \cdots 2^{b_n}. \tag{15}$$

From (15) it is obvious that $m_1^2 \ge 2$ if and only if $b_1 = 1$. To compute b_2 removal of b_1 is necessary, by dividing m_1^2 by 2, in case $b_1 = 1$, otherwise no division is required. Hence

$$m_2 = \begin{cases} \frac{m_1^2}{2} & b_1 = 1\\ m_1^2 & b_1 = 0 \end{cases}$$
(16)

after which

$$m_2 = 2^{b_2 2^{-1} + b_3 2^{-2} + \dots + b_i 2^{-i+1} + \dots + b_n 2^{-n+1}}.$$
 (17)

The pattern of (17) is identical to (14), except that it starts with b_2 , thus processing will continue until the desired precision is achieved.

ACKNOWLEDGMENT

The authors would like to thank STMicroelectronics for chip manufacturing, and they are grateful to D. Rehman and M. López Picazo for their contributions to the study at an early stage.

REFERENCES

- G. Moody and R. Mark, "A new method for detecting atrial fibrillation using R-R intervals," in *Computers in Cardiology 1983, Volume* 10. Aachen, Germany: IEEE Computer Soc. Press, 1983, vol. 10, pp. 227–230.
- [2] K. Tateno and L. Glass, "Automatic detection of atrial fibrillation using the coefficient of variation and density histograms of RR and Δ RR intervals," *Med. Biol. Eng. Comput.*, no. 39, pp. 664–671, 2001.

- [3] D. Duverney, J.-M. Gaspoz, V. Pichot, F. Roche, R. Brion, and A. A. J.-C. Barthélémy, "High accuracy of automatic detection of atrial fibrillation using wavelet transform of heart rate intervals," *Pacing Clin. Electrophysiol.*, vol. 25, pp. 457–462, 2002.
- [4] A. Bollmann, D. Husser, L. Mainardi, F. Lombardi, P. Langley, A. Murray, J. J. Rieta, J. Millet, S. B. Olsson, M. Stridh, and L. Sörnmo, "Analysis of surface electrocardiograms in atrial fibrillation: Techniques, research, and clinical applications," *Europace*, vol. 8, pp. 911–926, 2006.
- [5] S. Dash, K. Chon, S. Lu, and E. Raeder, "Automatic real time detection of atrial fibrillation," *Ann. Biomed. Eng.*, vol. 37, pp. 1701–1709, 2009.
- [6] J. Park, S. Lee, and M. Jeon, "Atrial fibrillation detection by heart rate variability in Poincaré plot," *Biomed. Eng. Online*, vol. 8, p. 38, 2009.
- [7] K. J. Jang, G. Balakrishnan, Z. Syed, and N. Verma, "Scalable customization of atrial fibrillation detection in cardiac monitoring devices: Increasing detection accuracy through personalized monitoring in large patient populations," in *Proc. Annu. Int. Conf. IEEE Engineering in Medicine and Biology Soc.*, 2011, vol. 33, pp. 2184–2187.
- [8] R. B. Shouldice, C. Heneghan, and P. de Chazal, "Automatic detection of paroxysmal atrial fibrillation," in *Atrial Fibrillation-Basic Research and Clinical Applications*, J. Choi, Ed. Rijeka, Croatia: Intech, 2012, ch. 7, pp. 125–146.
- [9] J. Lee, B. Reyes, D. McManus, O. Mathias, and K. Chon, "Atrial fibrillation detection using an iPhone 4S," *IEEE Trans. Biomed. Eng.*, vol. 60, pp. 203–206, Jan. 2013.
- [10] C. G. Scully, J. Lee, J. Meyer, A. M. Gorbach, D. Granquist-Fraser, Y. Mendelson, and K. H. Chon, "Physiological parameter monitoring from optical recordings with a mobile phone," *IEEE Trans. Biomed. Eng.*, vol. 59, pp. 303–306, 2012.
- [11] C.-W. Tseng, G.-H. Lin, C.-H. Chang, H.-Y. Chan, C.-L. Tsai, Y.-D. Lin, and K.-P. Lin, "Automatic detection of atrial fibrillation based on handheld ECG device," in *Proc. 5th Eur. Conf. Int. Federation for Medical and Biological Engineering*, 2012, vol. 37, pp. 506–509.
- [12] M. Stridh and M. Rosenqvist, "Automatic screening of atrial fibrillation in thumb-ECG recordings," in *Proc. Computing in Cardiology Conf.*, Sep. 2012, pp. 193–196.
- [13] "Guidelines for the management of atrial fibrillation," *Europace*, vol. 12, pp. 1360–1420, 2010, The Task Force for the Management of Atrial Fibrillation of the European Society of Cardiology (ESC).
- [14] J. Lian, L. Wang, and D. Muessig, "A simple method to detect atrial fibrillation using RR intervals," *Amer. J. Cardiol.*, vol. 107, pp. 1494–1497, 2011.
- [15] S. Sarkar, D. Ritscher, and R. Mehra, "A detector for a chronic implantable atrial tachyarrhythmia monitor," *IEEE Trans. Biomed. Eng.*, vol. 55, pp. 1219–1224, 2008.
- [16] R. Mehra, J. Gillberg, P. Ziegler, and S. Sarkar, "Algorithms for atrial tachyarrhythmia detection for long-term monitoring with implantable devices," in *Understanding Atrial Fibrillation: The Signal Processing Contribution*, L. Mainardi, L. Sörnmo, and S. Cerutti, Eds. San Rafael, CA, USA: Morgan Claypool, 2008, ch. 8, pp. 175–214.
- [17] J. Hulzink, M. Konijnenburg, M. Ashouei, A. Breeschoten, T. Berset, J. Huisken, J. Stuyt, H. de Groot, F. Barat, J. David, and J. Van Ginderdeuren, "An ultra low energy biomedical signal processing system operating at near-threshold," *IEEE Trans. Biomed. Circuits Syst.*, vol. 5, no. 6, pp. 546–554, Dec. 2011.
- [18] J. Rodrigues, O. Akgun, and V. Öwall, "A < 1 pJ sub-V_T cardiac event detector in 65 nm LL-HVT CMOS," in *Proc. 18th IEEE/IFIP* VLSI System on Chip Conf., 2010, pp. 253–258.
- [19] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, pp. E215–220, 2000.
- [20] M. Alioto, "Ultra-low power VLSI circuit design demystified and explained: A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 59, pp. 3–29, Jan. 2012.
- [21] Y. F. Ho, Fast Logarithm Converter for Fixed-Point Numbers without Look-Up Table, 2004 [Online]. Available: http://www.winnyefanho. net/research/Log2.pdf
- [22] P. Meinerzhagen, S. Sherazi, A. Burg, and J. Rodrigues, "Benchmarking of standard-cell based memories in the sub-V_T domain in 65-nm CMOS technology," *IEEE J. Emerg. Sel. Topic Circuits Syst.*, vol. 1, pp. 173–182, Jun. 2011.
- [23] P. Meinerzhagen, O. Andersson, B. Mohammadi, Y. Sherazi, A. Burg, and J. Rodrigues, "A 500 fW/bit 14 fJ/bit-access 4 kb standard-cell based sub-V_T memory in 65 nm CMOS," in *Proc. Eur. Solid-State Circuits Conf.*, 2012, pp. 321–324.

[24] V. S. Mallela, V. Ilankumaran, and N. S. Rao, "Trends in cardiac pacemaker batteries," *Indian Pacing Electrophysiol. J.*, vol. 4, pp. 201–212, 2004.



Oskar Andersson (S'11) received the M.Sc. degree in computer science and engineering from Lund University, Lund, Sweden, in 2010.

Currently, he is working toward the Ph. D. degree in the digital ASIC research group in the Electrical and Information Technology Department, Lund University. His research interests include power optimization of ultra-low voltage circuits, power-efficient circuits techniques, and biomedical circuits for implantable devices.



Ki H. Chon (SM'08) received the B.S. degree in electrical engineering from the University of Connecticut, Storrs, CT, USA, the M.S. degree in biomedical engineering from the University of Iowa, Iowa City, IA, USA, and the M.S. degree in electrical engineering and the Ph.D. degree in biomedical engineering from the University of Southern California, Los Angeles, CA, USA.

He spent three years as an NIH Postdoctoral Fellow at the Harvard-MIT Division of Health Science and Technology, one year as a Research

Assistant Professor in the Department of Molecular Pharmacology, Physiology, and Biotechnology at Brown University, Providence, RI, USA, and four years as an Assistant Professor and Associate Professor in the Department of Electrical Engineering at the City College of the City University of New York, NY, USA. He was Professor in the Department of Biomedical Engineering at SUNY Stony Brook, Stony Brook, NY, USA. Most recently, he was a Professor and Chair of Biomedical Engineering at Worcester Polytechnic Institute, Worcester, MA, USA, for four years. Currently, he is a Professor and Founding Head of Biomedical Engineering at the University of Connecticut. His current research interests include medical instrumentation, biomedical signal processing, biomedical instrumentation, mobile health diagnostics, wearable sensors and identification, and modeling of physiological systems.

Dr. Chon was an Associate Editor of the IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING from 2007–2013. He has chaired many international conferences, including the role of Program Co-Chair for the IEEE Engineering in Medicine and Biology Society (EMBS) Conference in New York City in 2006.



Leif Sörnmo (S'80–M'85–SM'02) received the M.Sc. and Ph.D. degrees in electrical engineering from Lund University, Lund, Sweden, in 1978 and 1984, respectively.

From 1983 to 1995, he was a Research Fellow in the Department of Clinical Physiology, Lund University, where he was involved with research on ECG signal processing. Since 1990, he has been with the Signal Processing Group, Department of Biomedical Engineering, Lund University, where he is currently a Professor of biomedical signal processing and re-

sponsible for the BME program. He is the author of *Bioelectrical Signal Processing in Cardiac and Neurological Applications* (New York, NY, USA: Elsevier, 2005). His research interests include statistical signal processing, modeling of biomedical signals, methods for analysis of atrial fibrillation, multimodal signal processing in hemodialysis, and power-efficient signal processing in implantable devices.

Dr. Sörnmo is an elected Fellow of International Academy of Medical and Biological Engineering. He is an Associate Editor of IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, *Journal of Electrocardiology*, and *Medical and Biological Engineering & Computing*. He is on the editorial board of the *Journal of Biomedical Engineering*. He was an Associate Editor of *Computers in Biomedical Research* (1997–2000).



Joachim Neves Rodrigues (S'00–M'05–SM'11) received the Ph.D. degree in electroscience from Lund University, Lund, Sweden, in 2005.

Currently, he is an Associate Professor in the Department of Electrical and Information Technology, Lund University. From 2005 to 2008, he acted as ASIC process lead in the Digital ASIC Department at Ericsson Mobile Platforms, Lund, Sweden. He rejoined his current department in 2008, and is currently the Program Director for System-on-Chip. His main research interests are modeling and implementation of digital and mixed-mode microelectronics, architectures for high performance ultra-low voltage designs, with a focus on biomedical circuits and systems.

Dr. Rodrigues is a technical committee member of the Biomedical Circuits and Systems Society since 2010, and Vice-Chair of the Swedish SSC chapter.