A

Rodrigues, Joachim; Akgun, OmerCan; Öwall, Viktor

Link to publication

A $< 1\,\mathrm{pJ}$ sub-V_T Cardiac Event Detector in 65 nm LL-HVT CMOS

Joachim Neves Rodrigues *Member*, IEEE, Omer Can Akgun *Member*, IEEE, and Viktor Öwall, *Member*, IEEE

Electrical and Information Technology, Lund University
22100 Lund, Sweden
Email: {joachim.rodrigues, omercan.akgun, viktor.owall}@eit.lth.se

Abstract— **This paper presents the hardware implementation of a wavelet based event detector for cardiac pacemakers. A high level energy estimation flow was applied to evaluate energy efficiency of standard-cell based designs, over several CMOS technology generations, from 180 to 65 nm, operated in the sub-threshold domain. The simulation results indicate a 65 nm low-leakage high-threshold (LL-HVT) CMOS technology as the favourable choice. Accordingly, the design was fabricated in 65 nm LL-HVT CMOS. Measurements validate the simulation results and prove that the circuit is fully functional down to a supply voltage of 250 mV. At the energy minimum voltage of 320 mV the circuit dissipates 0.88 pJ per sample at a clock rate of 20 kHz.**

I. INTRODUCTION

The application of implantable biomedical appliances has progressed tremendously during the last decades due to advances in CMOS technology scaling. The functionality of cardiac pacemakers has evolved from the steady-rate pacing in 1958, to programmable rate-responsive operation [1]. Traditionally, sensing, amplification and filtering of cardiac activity in the μV signal range is performed in the analog domain, before the signal is digitized [1], [2]. However, pacemaker functionality may be enhanced by performing signal processing in the digital domain, with the advantage of deploying more advanced algorithms.

The application of digital CMOS for cardiac event detection in favour of analog circuitry has previously been discarded because of the constraint on energy dissipation [1]. Technology scaling reduces dynamic power consumption due to smaller capacitive parasitics. However, disadvantageously leakage current has emerged as a major design constraint. Thus, leakage dissipation is seen as the major barrier, if targeting smaller technology nodes. However, if leakage is aggressively addressed, the overall energy dissipation may be competitive to analog circuitry.

Power consumption in digital CMOS is decreased significantly by aggressive supply voltage scaling. Several successful implementations of digital circuits operating in the sub-threshold regime are available in the literature [3]–[6]. Circuits operating at these extreme low supply voltages work at much lower speeds, i. e., the FFT processor presented in [6] has a maximum clock frequency of 10 kHz with a power supply of 350 mV. In the weak inversion regime, sub-threshold current of the transistors is used for computation. The sub-threshold

leakage current depends on the supply voltage exponentially, resulting in exponential increase in the circuit delay and lower leakage energy dissipation for lower supply voltages. Consequently, the extreme low power consumption results in excellent power delay product, making such circuits very interesting candidates for ultra-low power applications which do not have very high processing requirements.

The proposed architecture of a 3-scaled wavelet filterbank feeds a generalized likelihood ratio test (GLRT), and qualifies for pacemaker applications by reliable detection performance in moderate to high noise levels [7]. The architecture is implemented with 65 nm low leakage-high threshold (LL-HVT) CMOS technology. Energy efficiency is evaluated over several CMOS technology generations by deploying a SPICE-accurate energy model on the gate-level netlists [8], and confirmed by measurements.

In Sec. II the theory of the cardiac event detector is presented. The energy model used for evaluation is briefly presented together with CMOS process comparison in Sec. III. ASIC implementation is presented in Sec. IV, the simulation results are confirmed by measurements in Sec. V, and conclusions are presented in Sec. VI.

II. DIGITAL HARDWARE IMPLEMENTATION OF A CARDIAC EVENT DETECTOR

This section briefly presents the theory and architecture of a 3-scaled wavelet filterbank, that scales and conditions the signal for hypothesis testing in the GLRT, see Fig. 1. A more thorough description of the cardiac event detector may be found in [7].

To achieve a power-efficient hardware mapping, short filters with integer values are chosen, i.e., first order difference, and the impulse response was chosen as a third order binomial function. A more detailed description of the wavelet filterbank and the GLRT is found in [9]. The implemented wavelet filterbank consist of three branches, $q = 2, 3, 4$, that scale and filter the signal $x(n)$ from the analog-to-digital converter, see Fig. 1 and 2. The first biphasic branch realizes a straightforward implementation as

$$F(z) = 1 + 3z^{-(q-1)} + 3z^{-(2q-2)} + z^{-(2q-1)}$$
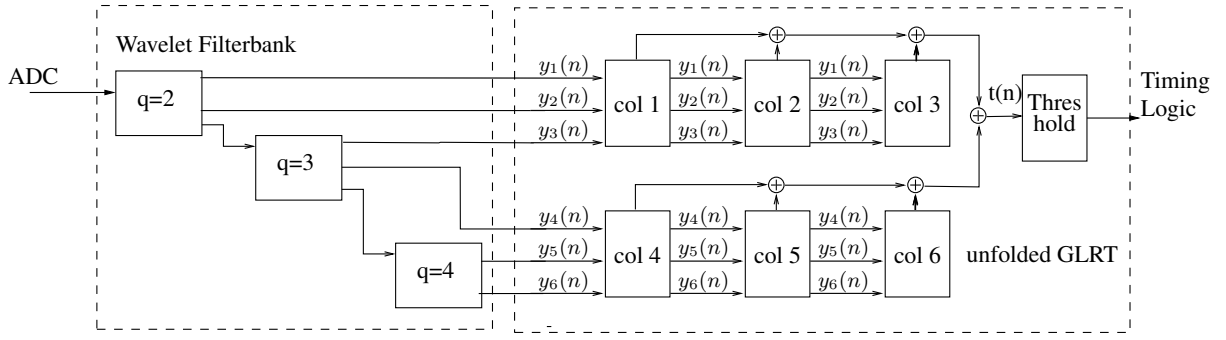
and

$$G_b(z) = -1 + z^{-q}.$$

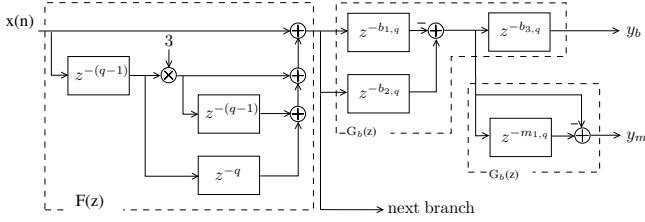Fig. 1: Direct mapped architecture of the wavelet filterbank and GLRT.

Fig. 2: Data flow diagram of the first wavelet filterbank branch using Mallat's algorithm, ($q = 2$).

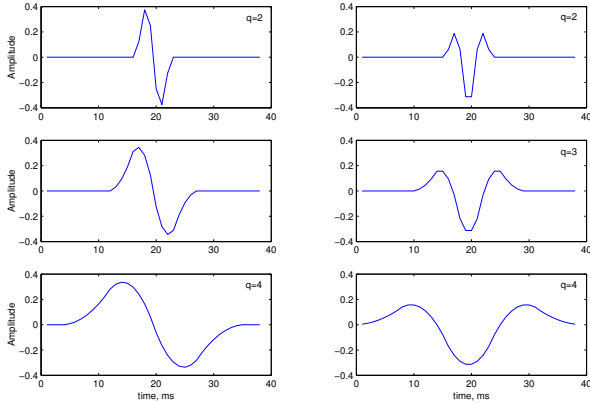Fig. 4: Data flow diagram of a direct-mapped block in the GLRT.

Fig. 3: Impulse responses of the wavelet filterbank. The biphasic impulse responses $y_{b,k}(n)$ for $q = 2, 3, 4$ are displayed in the left panel and the monophasic impulse responses $y_{m,k}(n)$ in the right panel.

Reusing $G_b(z)$ implements a monophasic filterbank using a single branch for one scale factor and realizes the output of the filterbank. However, in order to center the functions to the longest propagation delay in the third branch, it is necessary to introduce additional delays in $G_b(z)$, see Fig. 2. The impulse responses of the filterbank are presented in Fig. 3. It may be observed that the wavelet-based structure offers a high flexibility for various cardiac morphologies.

The decision signal $T(n)$ is computed by the GLRT as

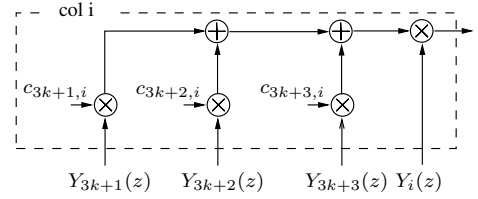$$T(n) = \mathbf{x}^T(n)\mathbf{H}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{x}(n), \qquad (1)$$

where $\mathbf{H}$ holds the coefficients of the bi- and mono-phasic filter functions. Since $\mathbf{x}^T(n)\mathbf{H} = \mathbf{H}^T\mathbf{x}(n)$, the remaining part of (1) to be implemented is the multiplication by $(\mathbf{H}^T\mathbf{H})^{-1}$, a matrix which is symmetric and sparse with half of its elements equal to zero,

$$(\mathbf{H}^T\mathbf{H})^{-1} =$$

$$\begin{bmatrix} 4.3 & -2.8 & 0.7 & 0 & 0 & 0 \\ -2.8 & 4.5 & -1.8 & 0 & 0 & 0 \\ 0.7 & -1.8 & 1.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4.8 & -2.3 & 0.6 \\ 0 & 0 & 0 & -2.3 & 4.2 & -1.4 \\ 0 & 0 & 0 & 0.6 & -1.4 & 1.7 \end{bmatrix}. \qquad (2)$$

The multiplication of $\mathbf{y}(n)$ with the first column of $(\mathbf{H}^T\mathbf{H})^{-1}$ and the first element of $\mathbf{H}^T\mathbf{x}(n)$ is carried out as depicted in Fig. 4, where $c_{3k+j,i}$ are elements of $(\mathbf{H}^T\mathbf{H})^{-1}$ and $y_{3k+j}(n)$ the output of the filterbank, with $k = 0, 1$ and $j = 1, 2, 3$.

The architecture of a single wavelet scale and an element of the GLRT is mapped as illustrated in Fig. 2 and 4, respectively. Three elements of the wavelet scale are cascaded to realize the scaling factors $q = [2, 3, 4]$ of the wavelet filterbank. The schematic in Fig. 4 represents the block referred to as *col i* in Fig. 1, which needs to be replicated six times to realize the multiplication with the columns of the matrix $(H^TH)^{-1}$ in (2). To simplify the implementation the matrix coefficients $c_{i,i} \cdots c_{i,i+2}$ are replaced with rounded integer values, which does not degrade performance. Thus, the multiplications are realized by *shift-add* instructions. Hence, the hardware realization of the GLRT requires six generic multipliers and 17 adders. Furthermore, the architecture is optimized by register

minimization, numerical strength reduction, and internal word-length optimization, which, in turn, results in narrower adders and multipliers in the following GLRT.

III. SUB-V_T ENERGY MODEL AND PROCESS COMPARISON

This section summarizes a high-level energy estimation flow [8]. The flow is used to model standard-cell based designs in the sub-V_T domain, and in this study, used to evaluate energy efficiency over several CMOS technology generations.

A. Energy Model

The model employed in this study uses parameters derived from high level simulations. The proposed model delivers SPICE-accurate data, but requires only a fraction of SPICE simulation runtime to obtain the internal energy dissipation of a single inverter equivalent capacitance value, which is not directly available in the synthesis library.

The total energy dissipation of static digital CMOS circuits is specified by the following well-known equation

$$E_{total} = \underbrace{\alpha C_{tot} V_{DD}^2}_{E_{dyn}} + \underbrace{I_{leak} V_{DD} T_{clk}}_{E_{leak}} + \underbrace{I_{peak} t_{sc} V_{DD}}_{E_{sc}}, \quad (3)$$

where E_{dyn} and E_{leak} are the average switching and leakage energy dissipated during a clock cycle T_{clk}, respectively. The contribution of short circuit energy (E_{sc}) in the sub-V_T regime is neglected, as it is known to contribute only a small share of the overall energy dissipation [10]. In (3), E_{dyn} during one clock period is specified by the switching activity factor α, and the maximum possible switched capacitance of the circuit C_{tot}. The total capacitance C_{tot} is normalized in terms of total inverter capacitance using a capacitance scaling factor k_{cap} as $C_{tot} = k_{cap} C_{inv}$, where C_{inv} is the switched capacitance of an inverter. To calculate k_{cap}, the total capacitance obtained by the synthesis is normalized by the gate capacitance value of an inverter from the synthesis library. The leakage energy dissipation during a clock period T_{clk} is defined as

$$E_{leak} = k_{leak} I_0 V_{DD} T_{clk}, \quad (4)$$

where k_{leak} and I_0 are the average leakage scaling factor of all gates in a design and the average leakage current of a single inverter, respectively. The value for k_{leak} is obtained from the synthesis results by summing the individual average leakage currents of the gates, where the average leakage current is the mean of the leakage current for all the combinations of input vectors applied to the logic gate, and normalizing the result to the average leakage current of a single inverter.

The critical path that constraints the maximum clock frequency is specified as

$$T_{clk} = k_{crit} T_{sw\_inv}, \quad (5)$$

where k_{crit} is a coefficient that defines the critical path delay of the circuit in terms of the inverter delay T_{sw_inv}. The parameter k_{crit} is calculated by dividing the critical path from the synthesis results by the average delay of the inverter while operating at nominal supply.

TABLE I: Supply voltage, operating frequency and the leakage current of the cardiac event detector at EMV.

Process	VDD (V)	f_{max} (kHz)	Leakage (nA)
180nm	0.34	49.3	142.1
130nm LL	0.32	13.7	24.2
90nm LL-HVT	0.30	29.2	44.4
65nm LL-HVT	0.33	20	15.5

The delay of an inverter operating in the sub-V_T regime is given in [10] as

$$T_{sw\_inv} = \frac{C_{inv} V_{DD}}{I_0 e^{V_{DD}/(nU_t)}}. \quad (6)$$

By introducing (6) into (5), the clock period is specified as

$$T_{clk} = k_{crit} \frac{C_{inv} V_{DD}}{I_0 e^{V_{DD}/(nU_t)}}, \quad (7)$$

and by combining (3), (4) and (7), the final total energy dissipation while working at the maximum clock frequency at a given supply voltage is specified as

$$E_T = C_{inv} V_{DD}^2 \left[\alpha k_{cap} + k_{crit} k_{leak} e^{-V_{DD}/(nU_t)} \right]. \quad (8)$$

The energy minimum operating voltage (EMV) is found by taking the derivative of (8) with respect to V_{DD}.

So far, it was assumed that the design operates at the maximum frequency that is imposed by the operating voltage, and hence operating with minimum leakage energy possible at that voltage. Usually, this is not the case in real world applications, where an operation frequency is dictated by external constraints. Thus, (8) cannot be used to calculate the energy dissipation of a circuit. A model that is not constraining the leakage time by the maximum operating frequency needs to be developed. For externally speed constrained systems which work below the speed that is achievable at the energy-minimum operating point, (8) is modified to

$$E_T = \mu_e k_{cap} C_{inv} V_{DD}^2 + k_{leak} I_0 V_{DD} T_{CLK}, \quad (9)$$

where μ_e is the average switching activity over N samples, and T_{CLK} is the period of the clock.

B. Process Comparison

Different foundry supplied standard cell libraries over different CMOS technology generations were analyzed for their energy efficiency, i.e., 180, 130, 90 and 65 nm. Low-leakage (LL) and high-V_T (HVT) options were used whenever available, and no special techniques for leakage reduction such as power gating or reverse body-bias were considered.

In sub-V_T circuits, EMV depends on the circuit and process properties. The EMV imposes a corresponding clock frequency of the circuit which is not practical, if the required clock frequency is dictated by external design constraints. Operating at any supply voltage value other than the energy-minimum voltage results in energy dissipation overhead when compared to the energy-minimum operation. In Table I, EMV, operating frequency and the leakage current of the cardiac event detector implemented in different technologies
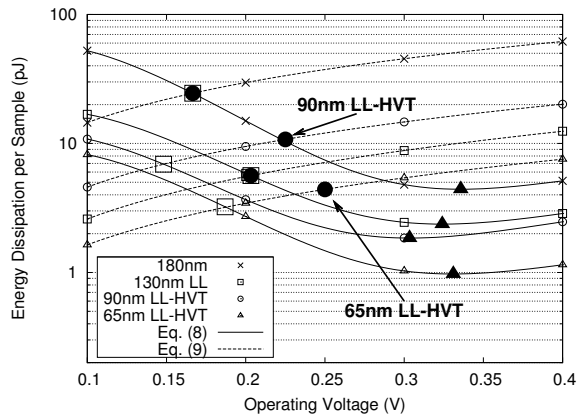
Fig. 5: Energy dissipation change for varying supply voltages across technology nodes, f=1 kHz.

Fig. 6: Cardiac event detector and peripherals.

are shown. It may be seen that at EMV, all the operating frequencies are much higher than target low-end frequency of the cardiac event detector, i.e., 1 kHz. Furthermore, it is found that despite that 90 nm process offers LL-HVT option, the leakage current is higher than in 130 nm LL and 65 nm LL-HVT processes. Hence the investigated 90 nm LL process is unattractive for low-speed, low-energy applications.

The investigated CMOS technology generations of 130 nm and below, come with LL process options. The drawback of LL and HVT technologies is the speed penalty. However, biomedical applications usually have a low requirement on processing speed, which in turn makes LL and HVT technologies a favourable process choice.

The circuit structure in the energy model is captured by k_{cap}, k_{crit}, k_{leak}. The activity factor α in (3) depends on the data processed by the circuit, and varies insignificantly across technology. These small variations are due to minor energy characteristic differences imposed by varying properties of leaf cells in different technologies. Moreover, since all k values are normalized to a basic inverter in a library, EMV varies insignificantly across technologies if the circuit structure is not changed. The sub-threshold slope factor n is the main parameter that changes the optimum supply voltage across technologies. In the technologies examined it was found that n varied between 1.38 and 1.49, which resulted in changes in EMV from 0.3 V to 0.34 V (Table I). If the circuit needs to operate at a lower speed than imposed by EMV, the supply voltage needs to be set to the maximum of either [11]:

- supply voltage which the circuit can operate at the required frequency, or
- supply voltage that satisfies a defined failure rate in ultra-low voltage operation mode.

The energy curves for the examined processes are shown in Fig. 5. The solid curves characterize energy dissipation if the circuit is operated at critical path speed, imposed by the scaled supply voltage, according to (8). The dashed curves represent energy dissipation if the circuit is operated at a lower speed as constraint by the supply voltage, i.e., the circuit is operated
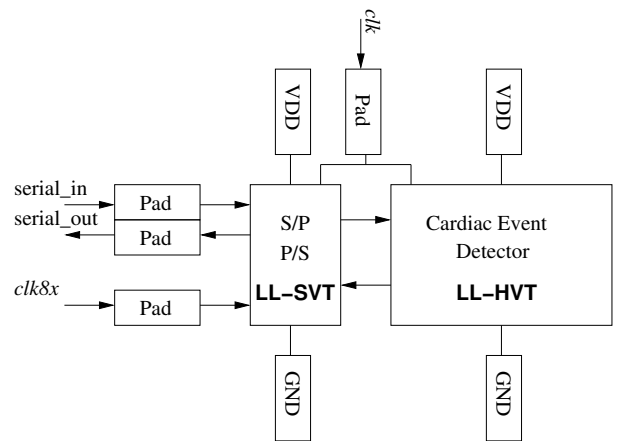
with an idle time on the critical path. On the curves, EMVs are marked with solid triangles. The target clock frequency of the proposed detector, i.e., 1 kHz, is lower than the maximum operating frequency at EMV. Thus, operating the cardiac event detector at the optimum voltage is not beneficial due to the excessive leakage energy as leakage energy scales with the operating voltage for a fixed frequency. In such a case, it is recommended that the circuit is operated at a supply voltage that is adequate to both the clock speed and the reliability requirements [12]. The supply voltages that satisfy a 1 kHz clock are indicated by squares in Fig. 5. Furthermore, due to ultra-low voltage operation, the static noise margin (SNM) degrades sharply with lower supply voltages, imposing a higher reliable operating voltage [12]. The operating points for 1 kHz operation when reliability is taken into consideration are marked with solid dots. For the 180 and 130 nm processes, the supply voltages that corresponds to a 1 kHz clock are higher than the operating voltages that satisfy the reliability condition. Contrarily, for 90 and 65 nm processes the supply voltages that satisfy reliability is higher than 1 kHz operation voltage. The system clock frequency is fixed at 1 kHz, however, these higher supply voltages need to be used to satisfy the reliability criterion. This non-optimal operating points results in excessive leakage energy, in 90 and 65 nm processes, and, thus, the operating points move off-curve. Hence, the expression in (8) is not applicable anymore and (9) needs to be applied.

As shown in Fig. 5, energy dissipation is reduced by 21.8% if migrating from 130 to 65 nm, for 1 kHz, and even more beneficial, i.e., 59%, if migrating from 90 to 65 nm. The analysis shows, if smaller process nodes are well-tuned for low-leakage operation, migrating to smaller technologies is beneficial as the energy curves in Fig. 5 shift according to C_{inv} in (8) which scales with the technology feature size.

IV. CHIP IMPLEMENTATION

This section presents how the cardiac event detector and peripheral hardware is mapped on an ASIC. The design is part of a multi project tapeout, where several different
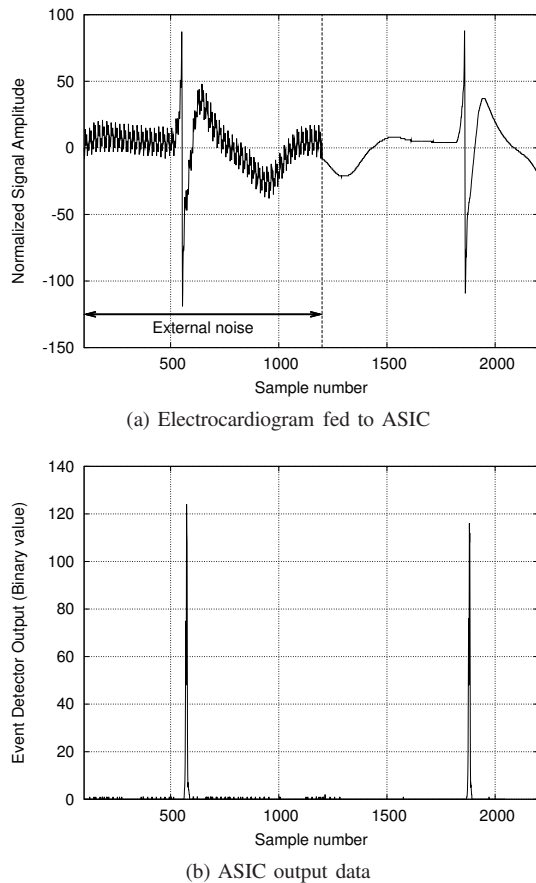
(a) Electrocardiogram fed to ASIC

(b) ASIC output data

Fig. 7: Measured ASIC input and output signals.

Fig. 8: Oscilloscope measurements of the ASIC in- and output signals.

implementations are accommodated on the same pad-limited die.

A. Cardiac Event Detector

The cardiac event detector was implemented with a 65 nm LL-HVT standard cell library, using constraints for minimum area and leakage. The gates are supplied by an independent power domain, where the power pads are isolated from any other power sources, see Fig. 6. The sequential logic is triggered by *clk*. Furthermore, timing is not a design constraint, and therefore, the clock is routed as an ordinary signal.

B. Parallelization and Serialization Module

To reduce silicon area, both the data to/from the ASIC are serially supplied/sampled, see Fig. 6. This is achieved by a module that receives serial input data and converts the bits to 8-bit words (S/P), and concurrently, the output of the ASIC is serialized (P/S). Two clocks, *clk* and *clk8x*, are connected to the module. By serializing the data, the number of pads is reduced from 19 to 8. The wordlength of the output is truncated to 8 bits which simplifies clocking, i.e., the module is triggered by a clock that is eight times faster than the the clock that triggers the ASIC. Furthermore, low-leakage standard-V_T (LL-SVT) cells are used in order to be able to drive the load of the pads and external measurement equipment. Thus, no
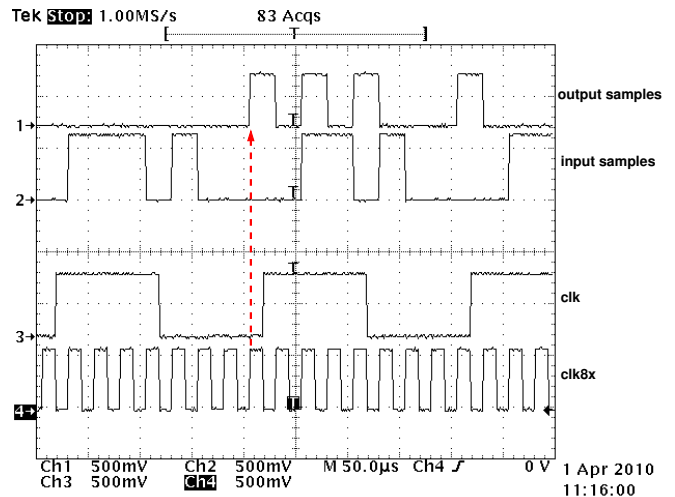
level shifters were required. The logic is accommodated on an independent power domain, which allows accurate energy measurements.

V. ASIC MEASUREMENTS AND SIMULATION COMPARISON

This section presents the measurement on the direct-mapped cardiac event detector, fabricated in 65 nm LL-HVT technology.

1) Measurement Setup: The measurements are carried out by sweeping both the supply voltage and clock frequency of the cardiac event detector. The former is supplied by a programmable voltage source, and the latter is generated by a XILINX Spartan-3 FPGA. The output of the circuit is monitored by a logic analyzer. Furthermore, the current drawn by the ASIC is measured with an integrator IC, which is accommodated on a custom made printed circuit board (PCB) that supplies the ASIC core. The signal subjected to the ASIC core is presented in Fig. 7a, which is a typical electrogram that is distorted with noise for 1200 samples. Using a sequence where the input signal is partially distorted is supposed to represent an average use case. The signal in Fig. 7b is the post-processed/re-constructed signal at the cardiac event detector output.

2) Sub-V_T Energy Measurements: Energy dissipation per sample is measured by sweeping V_{DD} from 250 to 350 mV, in steps of 10 mV, while *clk* is increased from 1000 to 20000 Hz. The step size is 1 kHz up to 10 kHz, and 10 kHz up to the limit. The supplied clock signals, as well as a sequence of input and output samples are presented in Fig. 8. It may be observed that the 8[th] rising clock edge of *clk8x* occurs before the rising edge of *clk* (dashed arrow). This guarantees that 8 bits are stored in the registers before the 1 kHz clock submits the input sample to the design. The output samples are bitwise fed with *clk8x* by the serialization module, as indicated by the dashed arrow.
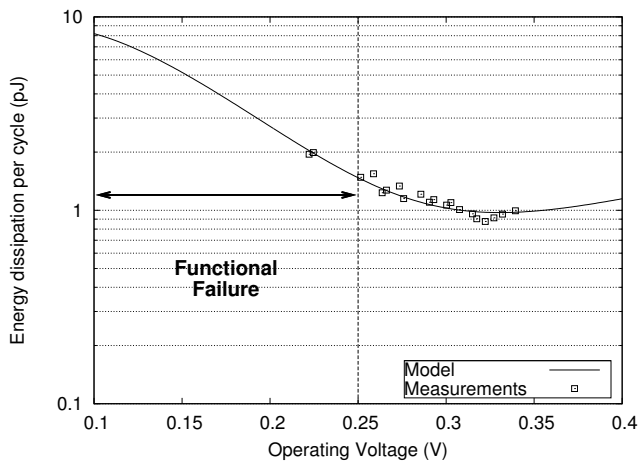
Fig. 9: Simulated and measured energy dissipation of the cardiac even detector.

Fig. 10: Chip microphotograph. The area of the cardiac event detector is $19425\mu m^2$.

The area of the design is $19425\mu m^2$, and poses an competitive alternative to state-of-the art analog event detectors.

The supplied signals from the pattern generator have an amplitude of 550 mV, see Fig. 8. These low voltage levels are obtained by external level shifters between the pattern generator pods and the PCB. The amplitude of the output samples is kept at the same level, as the independent power domain of the serialization module is connected to a 550 mV voltage source. Consequently, the captured samples have clear and sharp pulses. The output samples are captured and saved by a logic analyzer, and afterwards, correctness of the signal is approved by post-processing of the output data.

The measured energy values as well as the data obtained by the energy model are plotted in Fig. 9. Simulation data is represented by the solid curve, and measured data is indicated by squares. It is shown that the measured data is in the near vicinity of the simulated data. The mean of the absolute modelling error is calculated as 5.2 %, with a standard deviation of 6.6 %. A better matching with the simulation data may be achieved by a smaller step size of *clk*. Any supply voltage below 250 mV results in a high failure rate and is referred to as *functional failure*.

At an EMV of 320 mV the cardiac event detector dissipates as little as 0.88 pJ/per clock-cycle (sample), operating at a clock speed of 20 kHz. If the circuit is operated at a supply voltage that is adequate to a 1 kHz clock, the energy per sample is 4.4 pJ at 250 mV supply voltage. A chip photograph is presented in Fig. 10.

VI. CONCLUSIONS

This paper presents a cardiac event detector that was fabricated in 65 nm LL-HVT CMOS technology. The target technology was chosen based on simulation results of a high-level energy model that is able to capture design characteristics of standard-cell based design in the sub-V_T domain. Measurements prove that the ASIC is functional down to a supply voltage of 250 mV. The energy minimum voltage is 320 mV where the circuits dissipates 0.88 pJ per clock cycle.
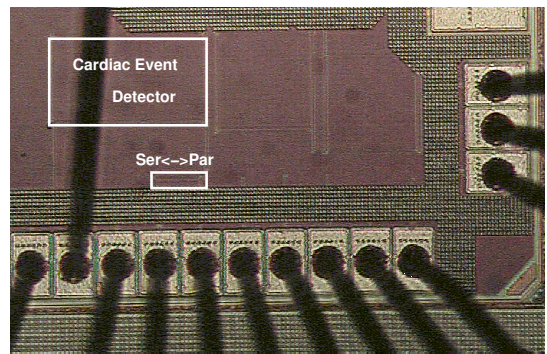
REFERENCES

[1] S. Haddad, R. Houben, and W. Serdijin, "The evolution of pacemakers," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 25, pp. 38–48, May-June 2006.
[2] L. Wong, S. Hossain, A. Ta, J. Edvinsson, D. Rivas, and H. Naas, "A very low-power CMOS mixed-signal IC for implantable pacemaker applications," *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 2446–2456, Dec. 2004.
[3] B. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, 2005.
[4] S. C. Jocke, J. F. Bolus, S. N. Wooters, T. N. Blalock, and B. H. Calhoun, "A 2.6 μw sub-threshold mixed-signal ECG SoC," in *ISLPED '09: Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design*, (New York, NY, USA), pp. 117–118, ACM, 2009.
[5] J. P. Kulkarni, K. Kim, and K. Roy, "A 160 mv robust schmitt trigger based subthreshold SRAM," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 2303–2313, Oct. 2007.
[6] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, 2005.
[7] J. Rodrigues, L. Olsson, T. Sörnmo, and V. Öwall, "Digital implementation of a wavelet-based event detector for cardiac pacemakers," *IEEE Transactions on Circuits and Systems I: Regular Papers,*, vol. 52, pp. 2686–2698, Dec. 2005.
[8] O. C. Akgun and Y. Leblebici, "Energy Efficiency Comparison of Asynchronous and Synchronous Circuits Operating in the Sub-Threshold Regime," *J. Low Power Electronics*, vol. 3, no. 3, pp. 320–336, 2008.
[9] M. Åström, S. Olmos, and L. Sörnmo, "Wavelet-based event detection in implantable cardiac rhythm management devices," *IEEE Trans. Biomed. Eng.*, vol. 53, March 2006.
[10] E. Vittoz, *Low-Power Electronics Design*, ch. 16. CRC Press LLC, 2004.
[11] D. Bol, R. Ambroise, D. Flandre, and J. Legat, "Analysis and minimization of practical energy in 45nm subthreshold logic circuits," in *IEEE International Conference on Computer Design, 2008. ICCD 2008*, pp. 294–300, 2008.
[12] J. N. Rodrigues, O. C. Akgun, P. Acharya, A. de la Calle, Y. Leblebici, and V. Öwall, "Energy dissipation reduction of a cardiac event detector in the sub-V_T domain by architectural folding.," in *PATMOS*, vol. 5953 of *Lecture Notes in Computer Science*, pp. 347–356, Springer, 2009.