High-Level Energy Estimation in the Sub-V_T Domain: Simulation and Measurement of a Cardiac Event Detector

Omer Can Akgun, *Member, IEEE*, Joachim Neves Rodrigues, *Senior Member, IEEE*, Yusuf Leblebici, *Fellow, IEEE*, and Viktor Öwall, *Member, IEEE*

Abstract—This paper presents a flow that is suitable to estimate energy dissipation of digital standard-cell based designs which are determined to operate in the subthreshold regime. The flow is applicable on gate-level netlists, where back-annotated toggle information is used to find the minimum energy operation point, corresponding maximum clock frequency, as well as the dissipated energy per clock cycle. The application of the model is demonstrated by exploring the energy efficiency of pipelining, retiming, and register balancing. Simulation results, which are obtained during a fraction of SPICE simulation time, are validated by measurements on a wavelet-based cardiac event detector that was fabricated in 65-nm low-leakage high-threshold technology. The mean of the absolute modeling error is calculated as 5.2%, with a standard deviation of 6.6% over the measurement points. The cardiac event detector dissipates 0.88 pJ/sample at a supply voltage of 320 mV.

Index Terms—Cardiac pacemaker, energy model, high-level energy estimation, QRS detection, R-wave, subthreshold.

I. INTRODUCTION

LOW ENERGY design is a very crucial design constraint for biomedical implants. Significant reduction in energy dissipation is achieved by lowering the supply voltage [1]. This is possible by relaxing the constraints of classical strong-inversion operation of metal–oxide semiconductor field-effect transistors (MOSFETs), and by accepting the notion that transistors operate well below threshold in the subthreshold (weak-inversion) regime.

In subthreshold (Sub-V_T) mode, the supply voltage may be scaled aggressively and, thus, power consumption is decreased by magnitudes. The subthreshold operation of static CMOS logic has been analyzed by using the EKV model in [2]. In this analysis, it is shown that static CMOS logic may operate with a supply voltage as low as 50 mV at ambient temperature.

There are several successful implementations of digital circuits operating in the subthreshold regime in the literature, such as a fast Fourier transform (FFT) processor that is operational down to 180 mV [3] and a subthreshold static random-access memory (SRAM) which operates with a supply voltage of 160 mV [4]. Circuits operating at these extreme low supply voltages work at much lower speeds, as an example, the FFT processor presented in [3] works with a maximum clock frequency of 10 kHz with a power supply of 350 mV. Their extremely low power consumption results in excellent power delay product (PDP) values, making such circuits very interesting candidates for ultra-low-power applications which do not have very high processing requirements; one of these applications is cardiac pacemakers.

In the subthreshold regime, the subthreshold current of the transistors is used for computation. The subthreshold current depends on the supply voltage exponentially, resulting in an exponential increase in the circuit delay and lower leakage energy dissipation for lower supply voltages. Due to an exponential dependence of the leakage energy and quadratic dependence of the switching energy in the subthreshold regime, subthreshold operation has an energy-minimum operating voltage (EMV).

Several energy estimation models for sub-V_T operation were published [2], [5], [6]. In [2], Vittoz investigated and proved the energy-minimum operation property of subthreshold logic. In the developed model, an expression for the EMV was not derived and the energy-minimum operating point was shown by numerically inverting the duty factor for minimum energy. In [5], the occurrence of EMV was shown but the equation for EMV was solved by curve fitting. In [6], Calhoun solved the subthreshold EMV analytically. In Calhoun's model, the average switched capacitance and average leakage current were specified as parameters and extracted from SPICE level simulation results. However, previous models are only applicable on synchronous designs, and furthermore, none of the mentioned models is feasible for high-level design exploration.

The extended and experimentally validated energy model presented in this paper provides accurate results on sub-V_T design characteristics, without requiring computation and time-intensive SPICE simulations [7]. High-level synthesis engines and power simulation tools typically use lookup tables from the leaf-cell characterization libraries supplied by the chip manufacturing companies. However, even with lower than nominal supply voltage recharacterized libraries (i.e., 0.3, 0.4, 0.5 V), a synthesis will only provide a suboptimal result, since the library

is recharacterized for one specific supply voltage. To find the EMV and frequency of a design and to apply further optimizations, a continuous model, which does not depend on precharacterized libraries and which is compatible with the standard design flow employed, is required. Due to its simplicity and nondependency on costly simulations, the presented model is usable at an early stage in the design flow and is applicable at the gate level and, thus, it is possible to characterize and optimize subthreshold energy efficiency by architectural optimizations. Moreover, only a single synthesis at the nominal library voltage is required for characterizing the design. With application of the model, the supply voltage is easily swept for subthreshold values in a continuous manner, and design characteristics, such as energy dissipation and operation speed, are easily obtained. The model is capable of modeling asynchronous and synchronous designs. In this paper, the energy model as well as its integration into a standard design flow is presented. In addition, the model is extended to include suboptimal clocking (i.e., working at a speed slower than the clock frequency at EMV). A sample case study for the model, pipelining exploration in the subthreshold regime, is presented with the application of the model. Furthermore, the model is validated with measurements of a cardiac event detector in a 65-nm process. Temperature measurements are performed and the validity of the model is confirmed for different temperature corners.

In Section II, the theory of the energy model is presented. The implementation flow of the energy model is presented in Section III. Section IV introduces the architecture of a cardiac event detector, which is used as a reference design to validate the energy model. Model application for energy analysis and a high-level design exploration example are presented in Section V and chip implementation details are given in Section VI. The energy model is validated by measurements in Section VII, and, finally, conclusions are presented in Section VIII.

II. THEORY

This section presents the theory behind the energy estimation flow, applicable to asynchronous and synchronous circuits. Furthermore, in order to simplify mathematical operations and derivation of equations, the following assumptions are made:

Assumption-1: The (asynchronous or synchronous) design operates with the highest possible throughput (i.e., new data are applied as soon as processing is finished (asynchronous)), or with the next clock (synchronous) that operates at critical path speed. Thus, the idle time of the hardware is minimized.

Assumption-2: Energy dissipation and processing delays of the circuit per computation are randomly distributed, guaranteed by processing a randomly distributed data set.

Assumption-1 guarantees that synchronous and asynchronous operations are compared fairly in terms of energy dissipation. Circuits run at the maximum speeds possible, and, hence, leakage energy dissipation is minimized, while working at the EMV. The claim in Assumption-2 is necessary to simplify the statistical properties of the energy model. Moreover, as long as energy dissipation and processing delay of a circuit are randomly distributed with a mean, the model may be applied to any circuit that operates with an arbitrary data set.

A. Energy Model Derivation

The energy model is developed by initial observations of an asynchronous design, and subsequently, the model is modified to enable numerical simulations of synchronous designs as well. In an asynchronous system, the operation of the system is dictated by the switching and delay properties as well as the external request and acknowledge signals. A conceptual illustration of an asynchronous operation is shown in Fig. 1(a). The design is observed for an arbitrary time frame T, where four sets of input data are processed. The time intervals where the circuit is purely leaking (waiting for handshake completion) are denoted with T_{Li}, and the time interval where the gates are both switching (processing data) and leaking are denoted as T_{Si}. It is assumed that as soon as the *RequestOut* signal is lowered, *RequestIn* goes high (Assumption-1) and new data are processed. Thus, the purely leaking time intervals (T_{Li}s) are fixed due to the asynchronous protocol overhead.

Fig. 1(b) shows the timing diagram of a synchronous design, which processes the same data as depicted in Fig. 1(a). The combinational processing times T_{Si} remain unchanged. However, the interval T_{Li}, where the synchronous timed design is in idle mode, is compared to its asynchronous counterpart. Thus, the synchronous implementation will have a lower throughput and longer idles times, while operating at the same supply voltage. These observations will be considered while developing the energy model.

1) Asynchronous Designs: To calculate the energy dissipation of an arbitrarily long time frame T, the operation cycles of a design need to be monitored. During this time frame, it is assumed that the number of switchings is N. In any static digital CMOS, the total energy dissipation is calculated as

$$E_T = E_{\text{dynamic}} + E_{\text{leakage}} + E_{\text{short-circuit}}, \qquad (1)$$

where E_{dynamic} is the dynamic energy due to the switching of the capacitances, E_{leakage} is the leakage energy during the time the gates are supplied by an energy source, and $E_{\text{short-circuit}}$ is the energy due to the short-circuit current flowing from the supply to the ground during switching. In this analysis, the contribution of the short circuit is neglected, as it is known to contribute only a small portion of the overall energy [2]. Dynamic energy dissipation in (1) during the ith time interval is specified as

$$E_{\text{dynamic}_i} = e_i C_{\text{tot}} V_{\text{DD}}^2 \qquad (2)$$

where e_i is a scaling parameter that defines the switching property of a design circuit for a specific input data transition, and C_{tot} is the maximum possible switched capacitance of the design. The switching energy scaling parameter e_i is in the range $[0, 1]$. Without loss of generality, e_i is expressed as a single value in a random process e (Assumption-2). Hence, it is possible to specify a mean μ_e and, thus, (2) for N computations may be rewritten as

$$E_{\text{dynamic}} = N\mu_e C_{\text{tot}} V_{\text{DD}}^2. \qquad (3)$$

In (2) and (3), the total capacitance C_{tot} may be normalized in terms of the total inverter capacitance using a capacitance
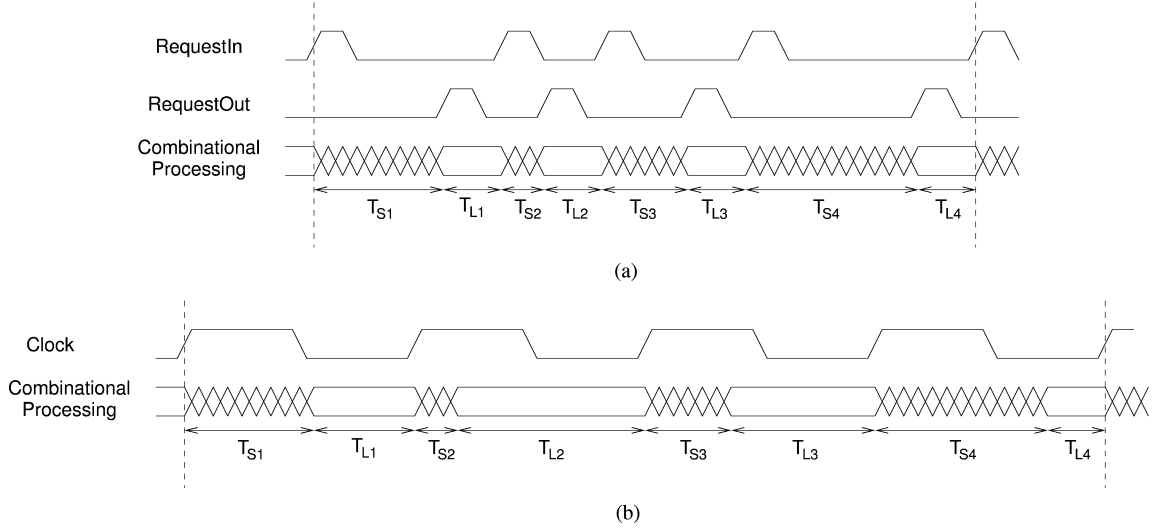
Fig. 1. Timing diagram showing the same logic block operating in (a) asynchronous and (b) synchronous modes.

scaling factor k_{cap} as $C_{\text{tot}} = k_{\text{cap}}C_{\text{inv}}$, where C_{inv} is the switched capacitance of an inverter. Furthermore, even while switching, the gates are leaking, and, consequently, leakage energy during the observation period T is defined as

$$E_{\text{leak}} = k_{\text{leak}}I_0V_{\text{DD}}T \qquad (4)$$

where k_{leak} is the average leakage scaling factor over all gates, and I_0 represents the average leakage current of a single inverter. Total average leakage current of a design is calculated as $k_{\text{leak}}I_0$ from (4). The average leakage parameter k_{leak} is obtained from synthesis results by computing the mean of the leakage current for all combinations of input vectors applied to the logic gate inferred in the gate-level netlist, and normalizing the result to the average leakage current of a single inverter. Insertion of (3) and (4) in (1), results in

$$E_T = N\mu_e k_{\text{cap}}C_{\text{inv}}V_{\text{DD}}^2 + k_{\text{leak}}I_0V_{\text{DD}}T \qquad (5)$$

which specifies the energy dissipation during a time interval T. The combinational processing time is specified as

$$T_S = \sum_{i=1}^{S} T_{s_i} \qquad (6)$$

and based on the switching/timing statistics of the design. T_{s_i} in Fig. 1(a) is calculated as

$$T_{s_i} = d_i k_{\text{crit}}T_{sw\_inv} \qquad (7)$$

where d_i in the range $[0, 1]$ is a scaling parameter that defines the delay properties during switching, k_{crit} defines the critical path delay per sample in terms of an inverter delay, and T_{sw_inv} is the delay of an inverter. Furthermore, by modeling d as a random process with the mean μ_d (Assumption-2), the total time spent during switching is computed as

$$T_S = N\mu_d k_{\text{crit}}T_{sw\_inv}. \qquad (8)$$

During the observation frame T, N switchings and N handshakes take place, thus T is expressed as

$$\begin{aligned} T &= T_S + T_L \\ &= N\mu_d k_{\text{crit}}T_{sw\_inv} + Nk_{oh}k_{crit}T_{sw\_inv} \end{aligned} \qquad (9)$$

where k_{oh} is a parameter that defines the overhead caused by the asynchronous communication in terms of critical path delay. k_{oh} depends on the chosen asynchronous communication protocol and its implementation specifics, and differs between different implementations. The delay of an inverter working in the subthreshold regime is given in [2] as

$$T_{sw\_inv} = \frac{C_{\text{inv}}V_{\text{DD}}}{I_0 e^{V_{\text{DD}}/(nU_t)}} \qquad (10)$$

where n is a process-dependent term called slope factor typically in the range of 1.3–1.5 in modern CMOS technologies, I_0 is the saturation off current, and U_t is the thermal voltage (26 mV at 300 K). The introduction of (10) into (9) gives

$$T = Nk_{\text{crit}}\frac{C_{\text{inv}}V_{\text{DD}}}{I_0 e^{V_{\text{DD}}/(nU_t)}}(\mu_d + k_{\text{oh}}) \qquad (11)$$

which is the total observation time. Finally, the insertion of (11) in (5), results in

$$\begin{aligned} E_{T-async} = NC_{\text{inv}}V_{\text{DD}}^2 [\mu_e k_{\text{cap}} \\ + k_{\text{crit}}k_{\text{leak}}(\mu_d + k_{oh})e^{-V_{\text{DD}}/(nU_t)}] \end{aligned} \qquad (12)$$

which is the total energy dissipation for N switchings. The average energy dissipation per clock cycle is calculated by setting $N = 1$ in (12). The optimal operating voltage for minimum energy operation is found by setting the derivative of (12) with respect to V_{DD} to zero. Thus, the EMV is given in (13) as

$$V_{\text{opt}-async} = 2nU_t - nU_tW_{-1}\left[-\frac{2e^2k_{\text{cap}}\mu_e}{k_{crit}k_{\text{leak}}(k_{oh}+\mu_d)}\right] \qquad (13)$$

where W_{-1} is the -1 branch of the LambertW function [8]. All of the k-parameters in (12) and (13) are extracted from synthesis results, and the μ-parameters are obtained from toggle

information generated by gate-level simulations with back-annotated timing information.

2) Synchronous Designs: A similar modeling approach is developed for synchronous designs, which operate as depicted in Fig. 1(b). The model is developed according to Assumption-1, which constraints that the clock period is equal to the critical path, i.e., $(T = k_{\mathrm{crit}}T_{sw\_inv})$; the energy per clock cycle is derived by modification of (11) and (12) as

$$E_T = C_{\mathrm{inv}}V_{\mathrm{DD}}^2 \times \left[\mu_e k_{\mathrm{cap}} + k_{\mathrm{crit}}k_{\mathrm{leak}}e^{-V_{\mathrm{DD}}/(nU_t)} \right]. \quad (14)$$

As in the asynchronous case, by taking the derivative of (14), we obtain the minimum energy operating point as

$$V_{\mathrm{opt-sync}} = 2nU_t - nU_tW_{-1}\left[-\frac{2e^2 k_{\mathrm{cap}}\mu_e}{k_{\mathrm{crit}}k_{\mathrm{leak}}} \right]. \quad (15)$$

So far, it was assumed that the design operates at the maximum frequency imposed by the operating voltage, hence operating with minimum leakage energy is possible at that voltage. Usually, this is not the case in real-world applications, where operation frequency is dictated by external constraints. Thus, (14) cannot be used to calculate the energy dissipation of a circuit in such a scenario. Therefore, a model not constraining the leakage time by the maximum operating frequency needs to be developed. For externally constrained systems, which work below the speed achievable at the energy-minimum operating point, (5) is modified to

$$E_T = \mu_e k_{\mathrm{cap}}C_{\mathrm{inv}}V_{\mathrm{DD}}^2 + k_{\mathrm{leak}}I_0 V_{\mathrm{DD}}T_{\mathrm{CLK}} \quad (16)$$

where T_{CLK} is the period of the clock.

B. Maximum Operating Frequency Calculation

From the introduced model, the maximum operating frequency of a circuit may be easily calculated. The operating frequency of a synchronous design is defined as

$$f_{\max} = \frac{1}{k_{\mathrm{crit}}T_{sw\_inv}} \quad (17)$$

and by introducing (10) into (17), the maximum operating frequency of a synchronous circuit is found as

$$f_{\max} = \frac{I_0 e^{V_{\mathrm{DD}}/(nU_t)}}{k_{\mathrm{crit}}C_{\mathrm{inv}}V_{\mathrm{DD}}}. \quad (18)$$

For an asynchronous design, the average throughput is a more important parameter, which can be obtained by the inclusion of the $\mu_d + k_{oh}$ factor as

$$f_{\mathrm{average}} = \frac{1}{\mu_d + k_{oh}}\frac{I_0 e^{V_{\mathrm{DD}}/(nU_t)}}{k_{\mathrm{crit}}C_{\mathrm{inv}}V_{\mathrm{DD}}}. \quad (19)$$

When (18) and (19) are compared, it can be seen that as long as $\mu_d + k_{\mathrm{oh}}$ is less than 1, the average throughput of an asynchronous implementation will be higher than that of a synchronous implementation.

III. Model Implementation Flow

The application flow of the energy model derived in Section II is shown in Fig. 2 and explained in this section with emphasis on the generation of the required parameters. The flow requires
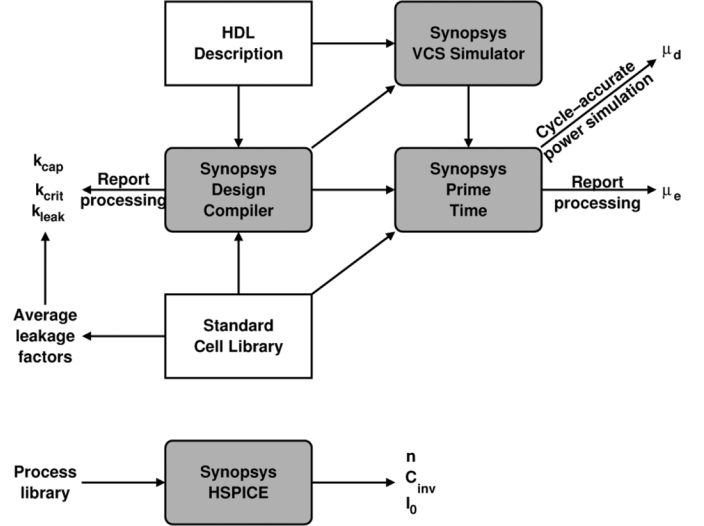

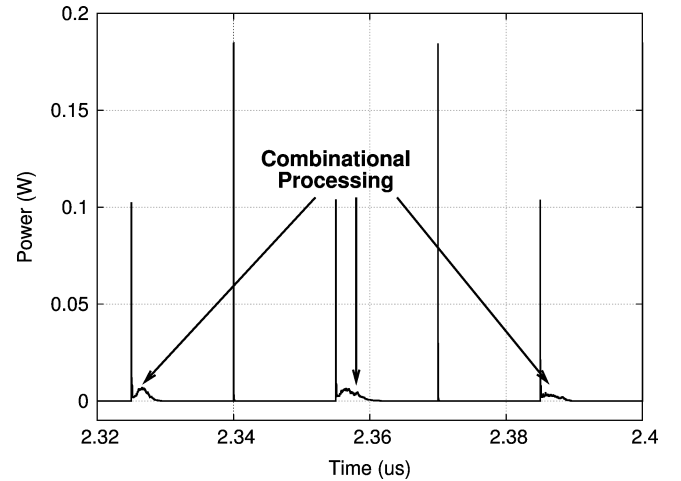Fig. 2. Model application flow emphasizing the tools used.

Fig. 3. Synopsys PrimeTime Power cycle-accurate power waveform to emphasize the power-consuming combinational operations. Combinational processing timing information is used for calculating the delay distribution of a mixed combinational-sequential circuit.

a so-called precharacterization of the process and the cell libraries, which only needs to be done once for the used technology. First, the cell libraries are characterized to compute the average leakage factor of each cell, normalized to a chosen inverter implementation. This is achieved by processing the .lib file supplied by the foundry (or generated during full custom layout design), with custom-developed scripts [9]. Second, the process library is characterized for the slope factor n, the inverter internal switched capacitance C_{int}, and saturation off current I_0. These parameters are not directly available in the .lib file. This is carried out by HSPICE simulations, where inverter characterization and single transistor testbenches are created.

Digital designs may be implemented by any hardware description language (HDL). Gate mapping is carried out by a synthesis tool (e.g., Synopsys Design Compiler). During synthesis, detailed reports of circuit properties are generated by the tool. For cycle accurate timing and power simulation, delay information of the design needs to be written in a standard delay
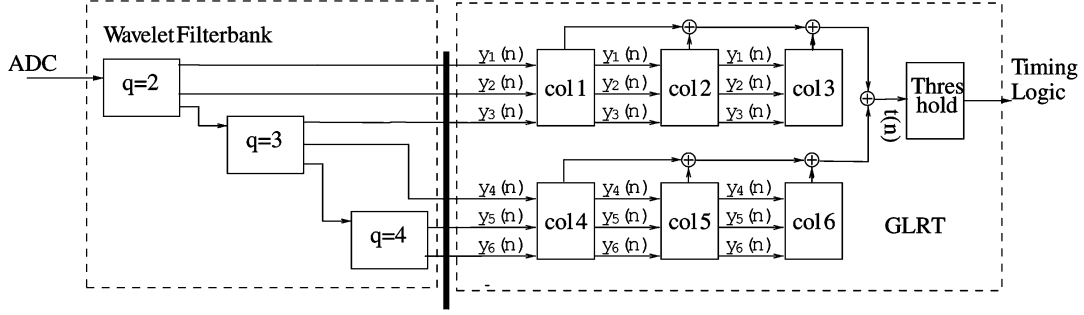
Fig. 4. Block diagram of the wavelet filterbank and GLRT. The manually placed pipeline stage, which is explained in Section V, is shown. The pipeline stage is emphasized with a vertical thick line.

format (`.sdf`) file. Thereafter, the synthesis reports are processed using custom-developed scripts in combination with average leakage factors obtained during precharacterization that are used to calculate k_{leak}. Finally, the values for k_{cap}, k_{crit}, and k_{leak} are obtained, see Fig. 2.

Input data-dependent toggle information of the design is obtained by gate-level netlist simulations using back-annotated timing information from the `sdf` file. HDL simulators, such as Synopsys VCS or Modelsim, may be used to create a value change dump (vcd) file. The gate-level netlist is subjected to Synopsys Prime Time to perform power simulation. Here, the vcd information is used in combination with the .sdf file to carry out cycle-accurate power simulation. By postprocessing the data, it is possible to extract the delay distribution and its average (μ_d), as well as the energy distribution and its mean (μ_e). A sample waveform from the Synopsys PrimeTime Power simulations is presented in Fig. 3. Power spikes just before combinational processing occurs are due to the positive edge of the clock signal and spikes that are not followed by combinational processing are due to the negative edge of the clock signal. Thus, the distributions and the average values of delay and energy characteristics of the circuit are obtained.

The proposed flow is applied during hardware optimization and implementation of a cardiac event detector for implantable pacemaker therapy.

IV. DIGITAL HARDWARE IMPLEMENTATION OF A CARDIAC EVENT DETECTOR

In current pacemaker families, cardiac event detection is composed of a bandpass filter followed by a programmable threshold level, implemented in analog circuitry [10]. Together with a low-power analog-to-digital converter (ADC) [11], digital implementation has the advantage of accommodating more extensive signal processing, such as features for morphology classification (e.g., in implantable cardioverter defibrillators (ICDs) and data compression for postanalysis [12], [13]). In this study, a digital cardiac event detector, which serves as a reference design for model validation, was implemented in 65-nm CMOS technology, targeting subthreshold operation.

This section briefly presents the theory and architecture of a three-scaled wavelet filter bank, that scales and conditions the signal for hypothesis testing in a generalized likelihood ratio test
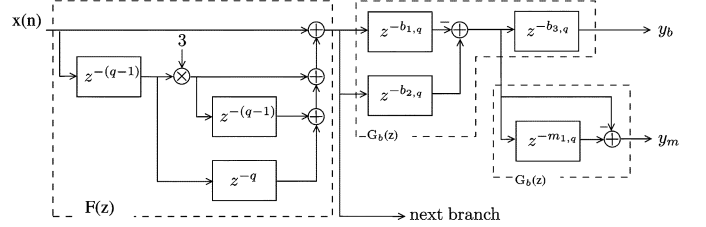

Fig. 5. Data-flow diagram of a wavelet branch using Mallat's algorithm ($q = 2$).

(GLRT), see Fig. 4. A more thorough description of the cardiac event detector may be found in [14].

A. Architecture

To achieve 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 [15]. The implemented wavelet filterbank consist of three branches $q = 2, 3, 4$, that scale and filter the signal $x(n)$, see Figs. 4 and 5. The first biphasic branch realizes a straightforward implementation as

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

and

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

Reusing $G_b(z)$ implements a monophasic filterbank using a single branch for one scale factor and realizes the output of the filterbank. 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. 5. The impulse responses of the filterbank are presented in Fig. 6. The wavelet-based structure offers high flexibility for various cardiac morphologies.

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

$$T(n) = \mathbf{x}^T(n)\mathbf{H}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{x}(n) \qquad (20)$$

where $\mathbf{H}$ holds the coefficients of the bi and monophasic filter functions. Since $\mathbf{x}^T(n)\mathbf{H} = \mathbf{H}^T\mathbf{x}(n)$, the remaining part of (20) to be implemented is the multiplication by $(\mathbf{H}^T\mathbf{H})^{-1}$, a
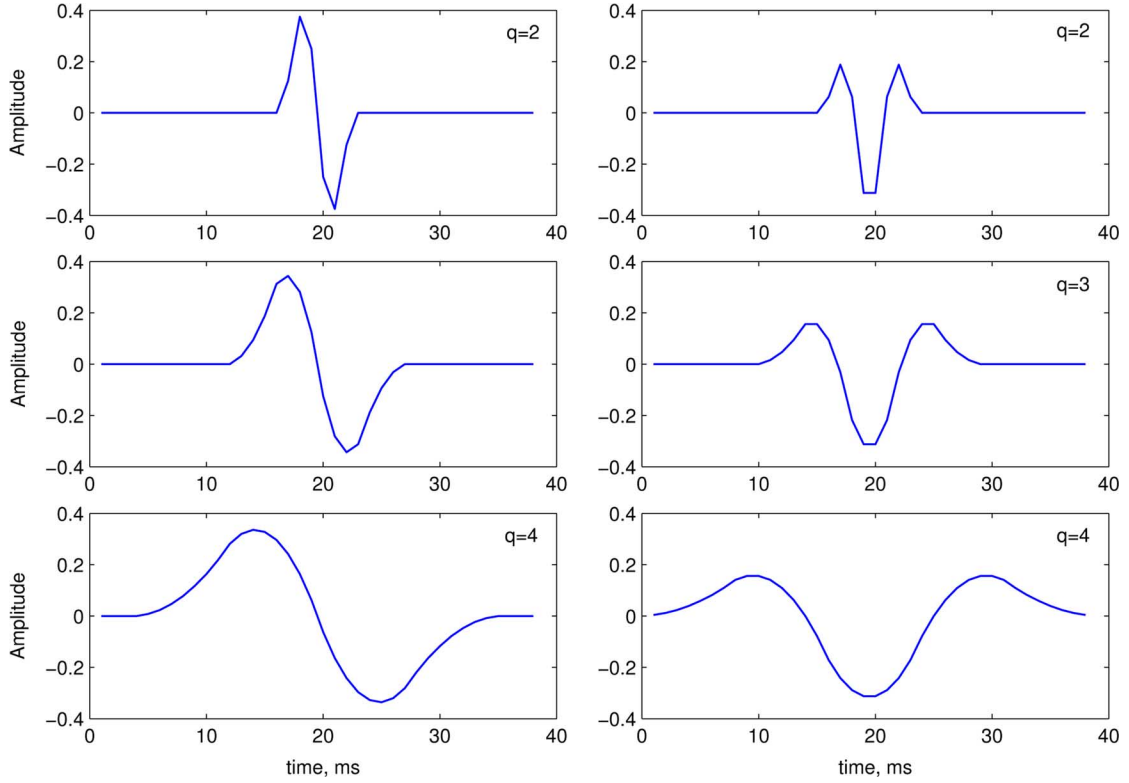
Fig. 6. Impulse responses of the wavelet filterbank. The biphasic impulse responses $y_{b,k}(n)$ for $q = 2, 3, 4$ are displayed in the left panel and the monophasic impulse responses $y_{m,k}(n)$ in the right panel.
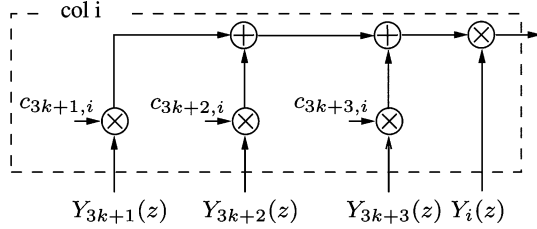

Fig. 7. Data-flow diagram of a direct-mapped block in the GLRT.

matrix which is symmetric and sparse with half of its elements equal to zero

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

$$= \begin{bmatrix} 4.3 & -2.8 & 0.7 & 0 & 0 & 0 \\ -2.8 & 4.5 & -1.8 & 0 & 0 & 0 \\ 0.7 & -1.8 & 1.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4.8 & -2.3 & 0.6 \\ 0 & 0 & 0 & -2.3 & 4.2 & -1.4 \\ 0 & 0 & 0 & 0.6 & -1.4 & 1.7 \end{bmatrix}. \quad (21)$$

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. 7, 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 wavelet scale and GLRT is mapped as illustrated in Figs. 5 and 7, 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. 7 represents the block referred to as *col i* in Fig. 4, which needs to be replicated six times to realize the multiplication with the columns of the matrix $(H^TH)^{-1}$ in (21). To simplify the implementation, the matrix coefficients $c_{i,i} \cdots c_{i,i+2}$ are replaced with rounded integer values, which do not degrade detection 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 GLRT.

V. Model Application and Process Variation Simulations

This section presents the characterization of the cardiac event detector in the sub-V_T domain with application of the energy model. In addition, the impact of pipelining on energy dissipation is investigated in order to illustrate a typical use-case for the energy model. Furthermore, simulations on process variations and reliability are performed and results are presented in order to find out the voltage interval where the model results will match those of the measurements.

A. Model Application

The model is applied to the cardiac event detector presented in Section IV. The characteristics of the design are summarized

TABLE I
DESIGN CHARACTERISTICS IN THE SUB-V$_T$ DOMAIN

Parameter	Value
k_{cap}	17820
k_{crit}	608
k_{leak}	13358
μ_{d}	0.47
μ_{e}	0.29

TABLE II
COMPOSITION OF THE CARDIAC EVENT DETECTOR IN TERMS OF
COMBINATIONAL LOGIC GATES AND REGISTERS

Logic element	kLeak	Area
Register	6987.74	5670.60
Combinational	6370.26	8578.44
Total	13358.00	14249.04

in Table I, and computed according to the flow presented in Section III. All of the model-based calculations are made using the data in Table I.

The signal that was supplied to the application-specific integrated-circuit (ASIC) input is shown in Fig. 8(a), and 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. 8(b) is the postprocessed/reconstructed signal at the cardiac event detector output. Running the switch-level simulation by using real-world data, μ_{d} and μ_{e} values are extracted as explained in the previous sections. Both the model-based analysis and measurements are made by using the same data set.

The gate composition of the event detector includes a substantial number of registers, see Table II. According to the synthesis results, leakage contribution of the combinational gates and registers is similar at 47.69% and 52.31%, respectively. Another use case for the application of the model is the exploration of pipelining on the energy dissipation of the cardiac event detector.

During the first step of our pipelining exploration, the register banks are implemented to divide the critical path of the circuit into shorter paths. Using this manual pipelining technique, a new architecture is generated and synthesized. Fig. 4 shows the pipelining applied to the circuit.

In the second phase of our pipelining exploration, pipeline optimization of the digital circuit synthesis engine is employed. For synthesis and optimization of the circuits in this paper, the synopsys design compiler (DC) was used. However, since the presented model employs textual synthesis reports as input data, any synthesis engine, which can provide the required data, may be used for the synthesis of circuits. For our pipelining exploration, the design is optimized by either register balancing (only timing optimization) or by timing and register area optimization. Following the synthesis of all different versions of the event detector, data extraction and simulations as explained in Section III are applied to all designs. The scaling parameters k and switching properties of the circuits are presented inTable III.

By using the design characteristics, the energy estimation model is applied to all of the architectures, and their energy profiles are generated. The results are presented in Fig. 9. By
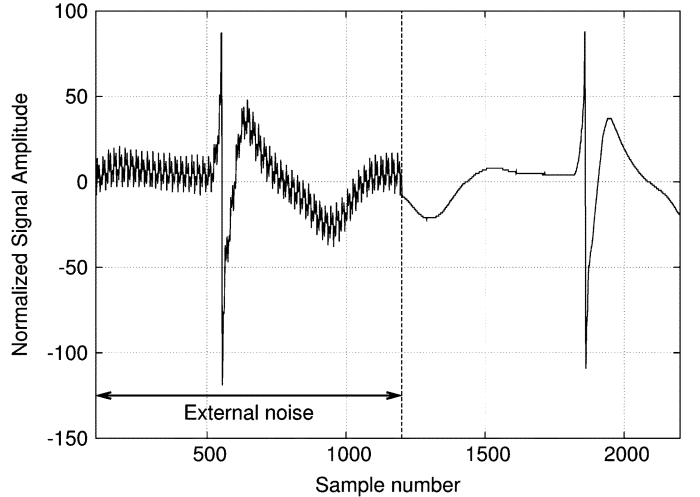
(a)
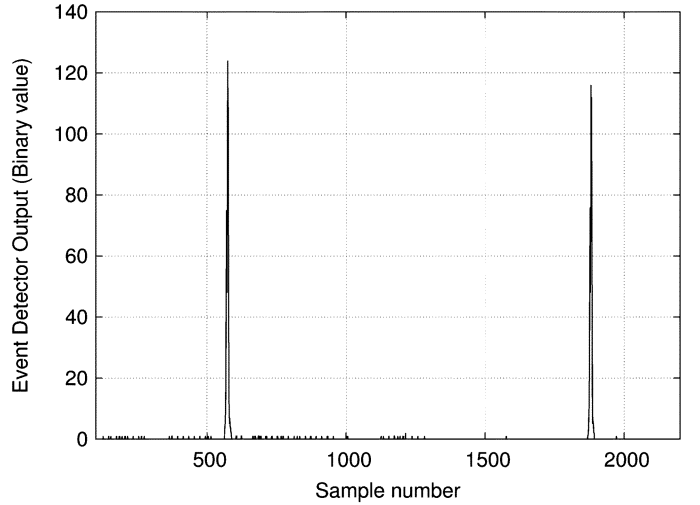
(b)

Fig. 8. (a) Electrocardiogram fed to the ASIC. (b) Measured ASIC output data.

TABLE III
COMPOSITION PROPERTIES OF MULTIPLE VERSIONS OF PIPELINED CIRCUITS

Circuit	kCap	kLeak	kCrit	μ_e
Manual Pipelined	18033.7	14150.2	382.0	0.29
Balanced Registers	18752.4	15356.9	229.4	0.27
Optimized Registers	17216.0	12460.	283.5	0.29
Non-pipelined	17820.0	13358.	608.0	0.29

analyzing the simulation results in the figure and the table, it is found that manual pipelining results in mediocre energy dissipation reduction. The lowest energy dissipation is achieved by the optimized registers circuit followed by balanced registers. The shortest path is achieved by the balanced registers implementation; however, the hardware overhead due to the extra registers makes this technique less efficient than the optimized registers implementation. This design exploration demonstrates how the proposed model is applied to evaluate different optimization strategies.

To view the effects of changes made to the cardiac event detector circuit in a compact way, we may visualize the circuits in terms of their normalized $k_{\text{Leak}}k_{\text{Crit}}$ and $\mu_e k_{\text{Cap}}$ parameters.
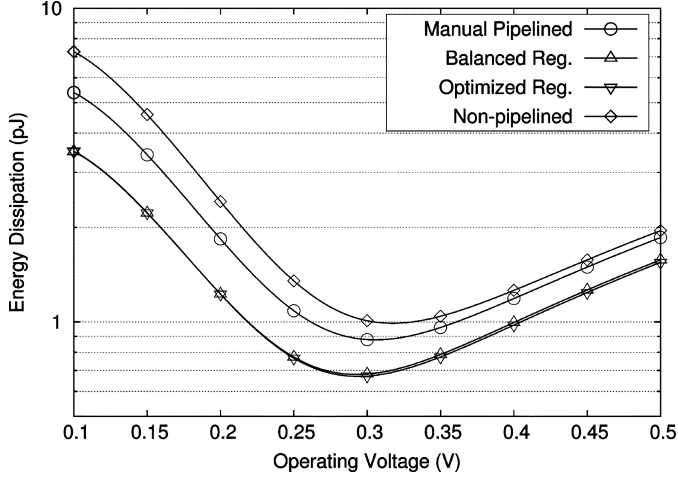
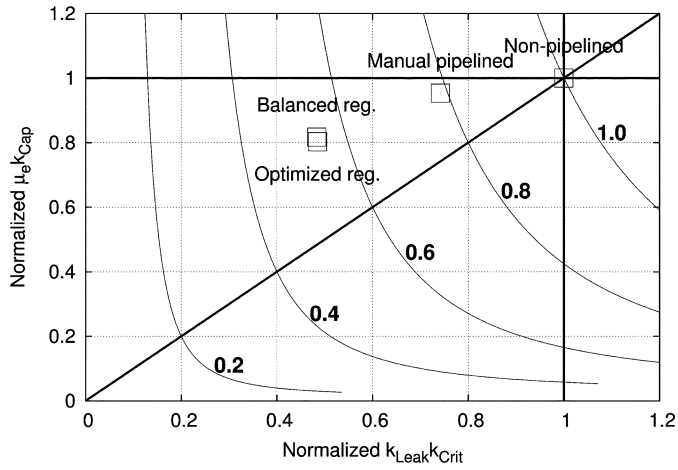Fig. 9. Energy curves for different pipelining levels.

Fig. 10. Pipeline comparison based on normalized k-factors with normalized equal energy contours overplotted.

These values are normalized to the basic architecture (i.e., non-pipelined circuit). Normalizing these values and plotting them as in Fig. 10 gives quick insight into the results of the optimizations. Lines $x = 1$, $y = 1$, and $y = x$ are added as references. The values for $\mu_e k_{\mathrm{Cap}}$ and $k_{\mathrm{Leak}} k_{\mathrm{Crit}}$ represent the switching energy and leakage energy components of the energy profile, respectively. Any value greater than 1 means that the modification resulted in overhead. For example, if the normalized $k_{\mathrm{Leak}} k_{\mathrm{Crit}}$ parameter of a design is greater than 1, this means for the same operating voltage, leakage energy of the modified design is greater. Furthermore, the $y = x$ line is used to deduce whether the modification resulted in lowering the energy-minimum voltage or not.

Due to the properties of the LambertW function in (15), as the ratio of $(\mu_e k_{\mathrm{Cap}})/(k_{\mathrm{Leak}} k_{\mathrm{Crit}})$ increases, the EMV shifts to a lower value due to higher relative switching energy. This case is shown in the plot by the region over the $y = x$ line. In Fig. 10, all of the pipelined designs are above the $y = x$ line; thus, they have lower EMV than the nonpipelined version, as already illustrated in Fig. 9. Furthermore, being on the upper region of the $y = x$ line signifies that the new architecture has a higher switching/leakage energy ratio when compared to the reference design.

Normalized equal energy contours are plotted in the figure to show the change of energy efficiency with the change in circuit parameters. From the figure, the effects of incremental changes to the reference design may be visualized.

In this section, we have presented an example of high-level design exploration using the proposed model. As can be seen, the presented application of the model allows a designer to conduct rapid high-level design exploration by using the synthesis results. Just by extracting the k-parameters and the switching properties of any design and with the application of the model, similar analysis for any high-level design exploration is possible.

B. Process Variation in 65-nm CMOS

Subthreshold circuits are more susceptible to random process variations than their above-threshold counterparts. In some cases, these random process variations cause functional failure. In modern nanometer technologies, random dopant fluctuation (RDF) and global process variations cause a shift from the nominal threshold voltage. Due to the fact that the subthreshold drain current depends on the threshold voltage exponentially, any change in the threshold voltage dominates other process variation effects in the subthreshold regime [16]. Functional failure of the static CMOS circuits due to random process variation (i.e., threshold voltage variation) may be investigated by using the static noise margin values derived from the voltage transfer characteristic (VTC) curves of digital gates. The aim of the conducted process variation analysis is to set a lower limit for our simulations and measurements.

To investigate functional failure of digital gates, the used process is simulated following the methodology described in [17] and [18], where the static noise margin (SNM) of the SRAM cells is calculated by using butterfly plots. Butterfly plots for two gates are formed by superimposing the VTC of one gate over the mirrored VTC of the other. Superimposed VTC curves of an inverter are presented in Fig. 11(a). A case where the diagonals of the drawn rectangles are nonexistent (i.e., at least one of the rectangles cannot be drawn), means that at least one of the logic levels cannot be sensed by the gate, resulting in functional failure at the simulated voltage. This formation simulates whether a logic level can be regenerated in case these two successive gates are in a logic path one after another. A sample testbench, which includes the static-noise sources, is shown in Fig. 11(b). To simulate the scenario that different gates are connected in back-to-back fashion, testbenches which extract the SNM data automatically are set up according to the methodology in [17]. The SNM square diagonals are calculated by rotating the voltage axes by 45°. In the rotated plane, the diagonal of the SNM square for any voltage is calculated by subtracting one VTC curve from the other. If the calculated value is negative, the SNM is negative and the logic value cannot be regenerated.

The SNM failure rates of the gates are extracted from 10k-point Monte Carlo simulations. The simulations are run for supply voltage values which vary between 0.1 V and 0.35 V with 25-mV steps. Simulated functional failure rates are presented in Fig. 12. Supply-voltage values, which realize operation with less than 0.001 failure rate, are taken as the minimum
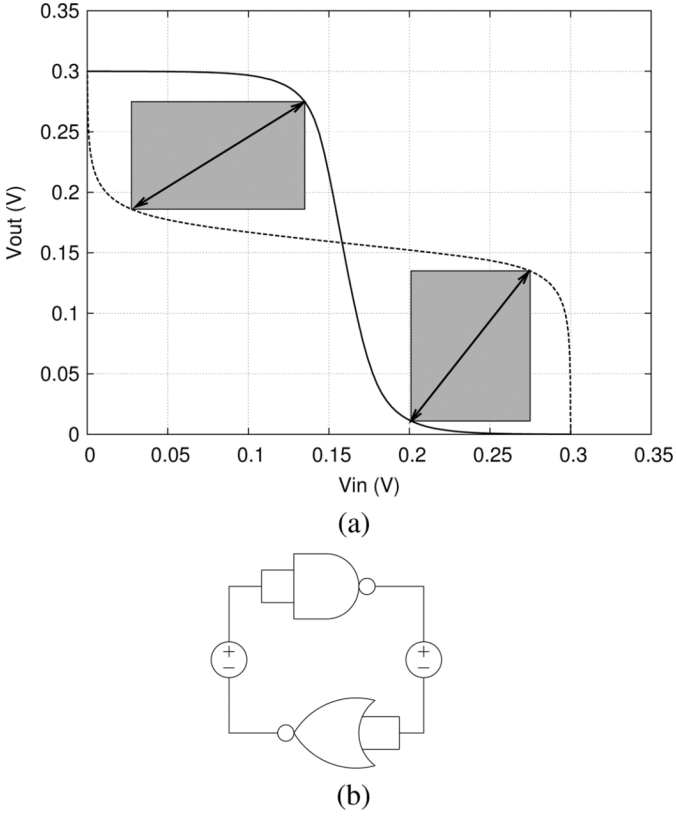
(a)

(b)

Fig. 11. (a) Butterfly SNM curves for a subthreshold inverter with SNM diagonals emphasized. (b) SNM testbench for reliability simulations.
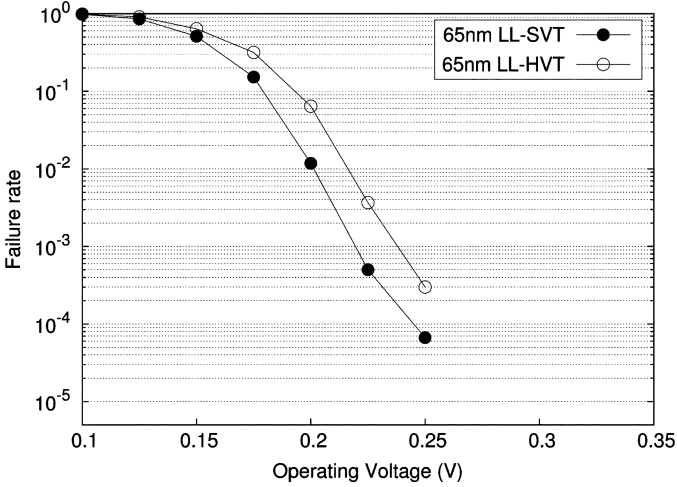

Fig. 12. Simulated reliability for 65-nm CMOS.

Fig. 13. Cardiac event detector and peripherals.
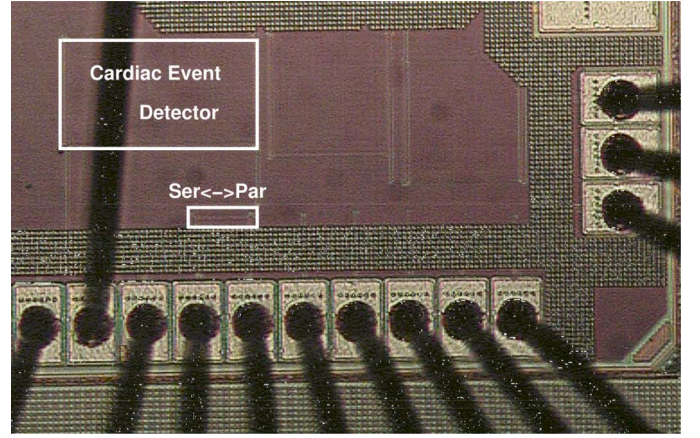
Fig. 14. Chip microphotograph. The area of the cardiac event detector is 19 425 μm^2.

reliable operating voltage (ROV) in this analysis. It is shown that standard-threshold technology (SVT) has lower failure rates at all supply voltages. SVT technology has a lower failure rate because this technology has a lower threshold voltage and, thus, the ratio of variation of the threshold voltage of the transistors to the mean threshold voltage value is lower, resulting in less functional variation. Consequently, it is concluded that the design will malfunction if the supply voltage is below ROV.
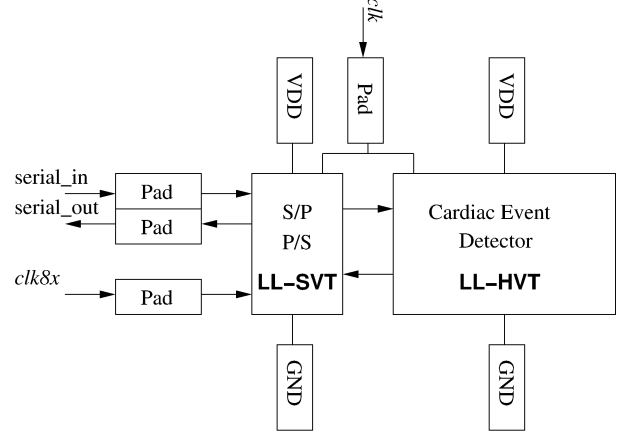
VI. TEST CHIP IMPLEMENTATION

This section presents the details of the cardiac event detector test chip implementation for validating the energy model. Both synthesis and silicon implementation details are presented. The design is part of a multi project tape-out, where several different implementations are accommodated on the same pad-limited die.

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 source, see Fig. 13. The sequential logic is triggered by *clk*. Furthermore, timing is not a design constraint and, therefore, the clock is routed as an ordinary signal. The cardiac event detector consist of 727 registers and 4200 NAND2 equivalent gates, occupying an area of 19 425 μm^2.

In order to reduce silicon area and overcome pad limitation, the data to/from the ASIC are serially supplied/sampled, see Fig. 13. This is achieved by a module that receives serial input data and converts the bits to 8-b 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 input and
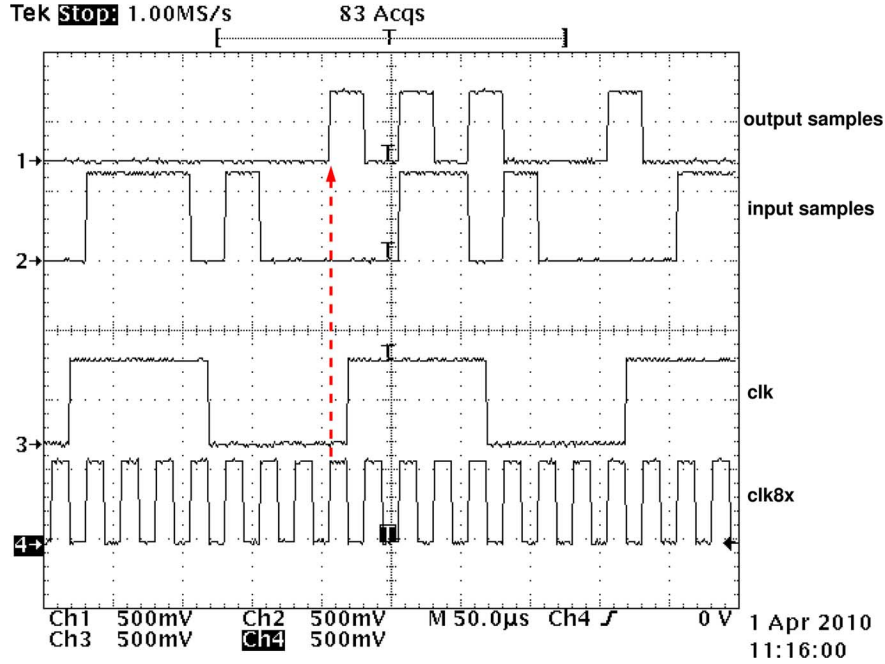
Fig. 15. Oscilloscope measurements of the ASIC input and output signals.

output data, the number of pads is reduced from 19 to 8. Moreover, the wordlength of the output is truncated to 8 b which simplifies clocking (i.e., the module is triggered by a clock that is eight times faster than the clock that triggers the ASIC). Furthermore, low-leakage standard-V_T (LL-SVT) cells in the S/P-P/S block are chosen for two reasons: 1) in order to drive the load of the pads and external measurement equipment and 2) to operate at a faster clock frequency at similar voltage levels. The serialization module is accommodated on an independent power domain, which allows accurate energy measurements on the cardiac event detector. Moreover, by setting the supply voltage of the serialization module to a higher voltage, sharper signals can be observed at the chip outputs. Although there are two separate power domains which are supplied by different voltages, during the measurements, the voltage difference between these two domains was always kept low (swept together) to guarantee that lower voltage domain can always cause a state change in the higher voltage domain. Furthermore, the gates at the input side of the boundary of two different voltage domains were chosen to have lower gate threshold voltages to guarantee a state change even with a lower voltage signal at the input. This was simulated and verified with SPICE and confirmed by our measurements. The chip micrograph is shown in Fig. 14, where the serialization module is labeled with Ser<->Par.

VII. MODEL VALIDATION

This section presents the measurement on the reference design (i.e., the direct-mapped cardiac event detector) that was fabricated in 65-nm LL-HVT technology. The measured data are used to evaluate the accuracy of the energy model for energy dissipation and operation speed at different temperatures. Before the manufacturing of the test chip, the energy model was compared to the SPICE simulation results of the ISCAS85

[19] benchmark circuits in the same process. The mean of absolute error was found to be 3.83% with a maximum error of 8.74%. For older technology (i.e., 180 nm), measurement results on other test circuits were presented in [20]. For all SPICE-level simulations, BSIM models supplied by the foundries were used. Good accuracy between the simulations and the measurements in 65 nm (this paper) and 180 nm [20] technologies were observed.

A. Measurement Setup

The measurements are carried out by sweeping 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 partially in-house developed energy acquisition equipment is applicable to slowly clocked ($<\ = 10$ kHz) subthreshold ASICs while measuring the ASICs in a sweeping fashion (i.e., frequency-supply voltage grid). However, this limitation in clock frequency is sufficient for the validation of the energy model.

The energy is measured on ASIC samples, which were available in an earlier tapeout [21], [22], where the ASIC cores could be triggered by a synchronous clock or in a self-timed manner depending on the implementation. The measured cores are identical; however, the ASIC in [21] has a sensing transistor in the supply rails. Speed degradation of 40% occurs due to this transistor and is taken into consideration in our analysis. This core is used to deliberately slow down the operation so that a larger measurement grid can be employed for speed measurements at varying temperatures. Energy dissipation per data sample is measured by sweeping V_{DD} from 220 to 350 mV, in steps of 10
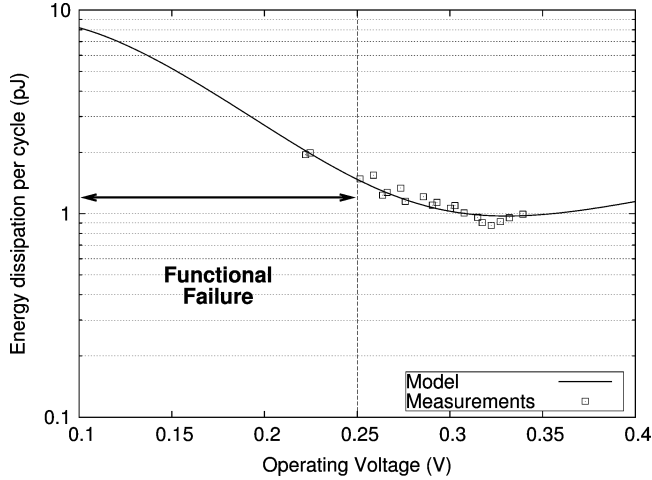
Fig. 16. Measured and simulated data of the reference design.

mV, while *clk* is increased from 1000 to 10000 Hz. The frequency step size is 1 kHz up to 10 kHz. The supplied clock signals as well as a sequence of input and output samples are presented in Fig. 15. It may be observed that the 8th rising clock edge of *clk8x* occurs before the rising edge of *clk* (dashed arrow). This guarantees that 8 b 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.

The supplied signals from the pattern generator have an amplitude of 550 mV, see Fig. 15. 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 postprocessing of the output data.

B. Sub-V_T Energy and Failure-Rate Measurements

The measured values as well as the data obtained by the energy model are plotted in Fig. 16. Simulation data are represented by the solid curve, and measured data are indicated by squares. It is shown that the measured data are in the near vicinity of the simulated data. The mean of the absolute modeling error is calculated as 5.2%, with a standard deviation of 6.6%.

At an EMV of 320 mV, the cardiac event detector dissipates as little as 0.88 pJ/per sample, operating at a clock speed of 20 kHz. The measurements in Fig. 16 are made in a point-by-point fashion (i.e., frequency-voltage pairs generated from the model are used for the setup, and corresponding energy dissipation is measured).

Furthermore, to compare the maximum operating speed of the implemented circuit to the model, the ASIC output samples are monitored for their correctness. This analysis is carried out for different die temperatures (i.e., 0°, 27°, and 37°). For these
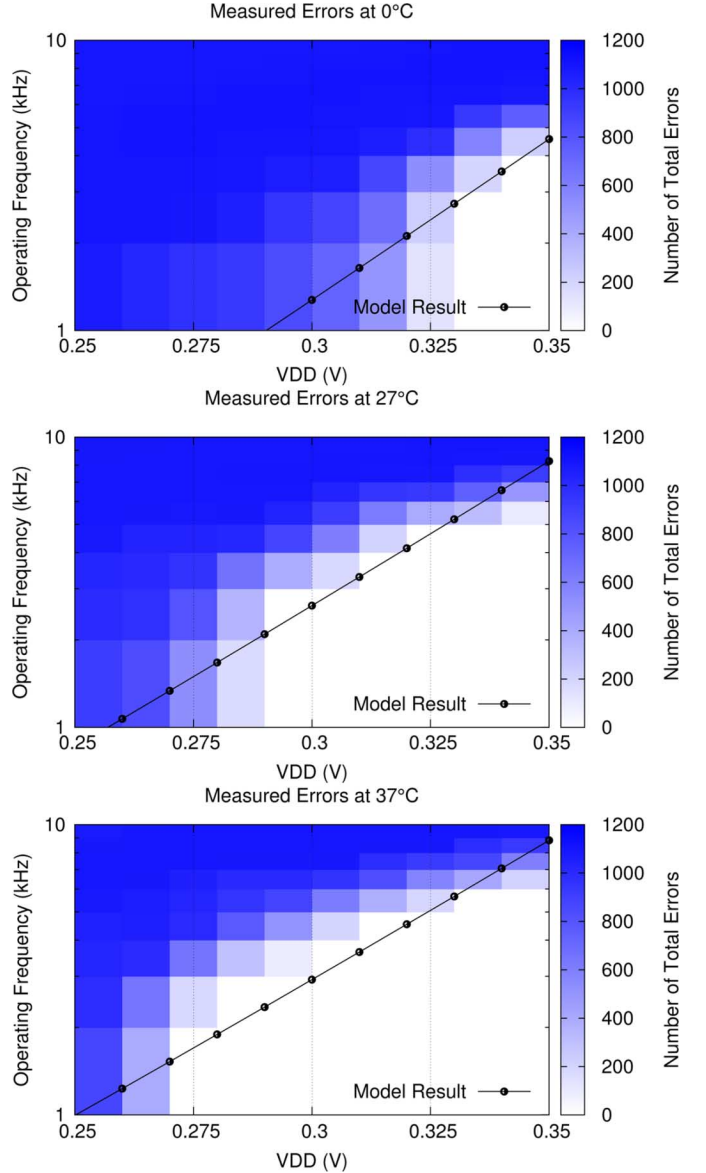

Fig. 17. Error measurements for different temperatures with the operation frequency calculation from the model plotted over.

measurements, V_{DD} is increased in steps of 10 mV from 250 to 350 mV, while sweeping f from 1 to 10 kHz.

The corresponding failure-rate measurement plots are presented in Fig. 17. Functional correctness of the ASIC is indicated by the white area, whereas the number of wrong samples is indicated by the shaded area. The highest measured failure rate is 1200, which represents the number of samples which were incorrectly computed. The simulated clock frequency from the model (18), which is constrained by the propagation delay of the critical path, is plotted as a solid line. To calculate the maximum operating speed for different operating temperatures, I_0 in (18) needs to be extracted for the required temperature values and U_t needs to be recalculated. The highest failure is measured when the ASIC operates at 0°, where the ROV is above 325 mV. With a higher temperature, ROV is reduced to 290 and 270 mV at 27°, and 37°, respectively. From this observation, it becomes evident that the failure rate decreases while increasing
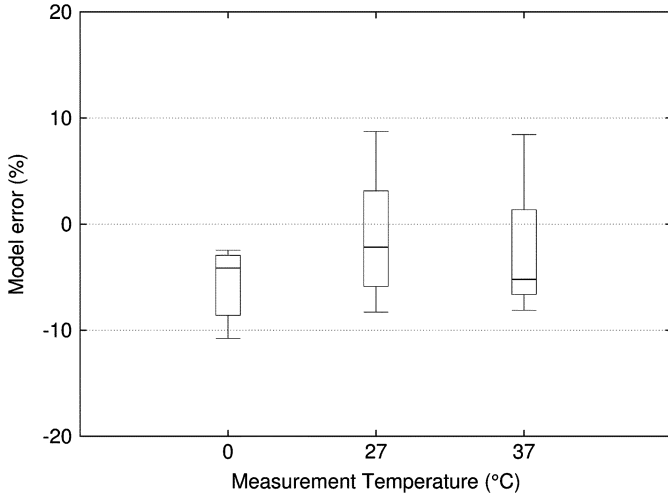
Fig. 18. Energy model accuracy for different operating temperatures. For each temperature error plot, energy values and measurements on the energy curve from (14) are used.

the die temperature. This is expected as the leakage current of a CMOS gate increases with increasing temperature [23], resulting in faster operation according to (18). Moreover, it is shown that simulated maximum frequency (imposed by V_{DD}) matches well with the frequency where the ASIC starts to malfunction. Thus, with these observations, it is shown that the energy model provides reliable simulation data for different ASIC operation temperatures.

As in the maximum operating frequency, energy dissipation of a circuit also varies with changing temperature according to (14). Similar to the maximum frequency calculation, the value of the thermal voltage U_t in (10) needs to be updated to retrieve data by simulations. The accuracy of the energy model for different operating temperatures is evaluated with respect to measurements, and the resulting error box-and-whisker plot is presented in Fig. 18, where the modeling error is lower than 10% at different operating temperatures.

VIII. CONCLUSION

A high-level energy estimation flow for ASICs that operate in the subthreshold domain is proposed. All parameters are directly obtained by applying a traditional synthesis and power estimation flow, and by postprocessing data that were generated during synthesis. The model is suitable for design and technology space exploration and is orders of magnitude faster than SPICE simulations. The model has been extended to include suboptimal clocking. A sample case study for the application of the model, pipelining exploration for energy efficiency improvement, is shown. An ASIC in 65-nm LL-HVT technology for the cardiac pacemaker application was fabricated. The accuracy of the model is validated by ASIC measurements under various operating conditions (i.e., V_{DD}, f_{clk}, and temperature). It is shown that the proposed model is able to accurately simulate the energy dissipation and maximum operating frequency of the implemented design.

REFERENCES

[1] A. Chandrakasan, S. Sheng, and R. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, no. 4, pp. 473–484, Apr. 1992.

[2] E. Vittoz, *Low-Power Electronics Design*. Boca Raton, FL: CRC, 2004, ch. 16.

[3] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.

[4] J. P. Kulkarni, K. Kim, and K. Roy, "A 160 mV robust schmitt trigger based subthreshold SRAM," *IEEE J. Solid-State Circuits*, vol. 42, no. 10, pp. 2303–2313, Oct. 2007.

[5] B. Zhai, D. Blaauw, D. Sylvester, and K. Flautner, "Theoretical and practical limits of dynamic voltage scaling," in *Proc. 41st Annu. Conf. Design Automation*, New York, 2004, pp. 868–873.

[6] B. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, Sep. 2005.

[7] O. C. Akgun and Y. Leblebici, "Energy efficiency comparison of asynchronous and synchronous circuits operating in the sub-threshold regime," *J. Low Power Electron.*, vol. 3, no. 3, pp. 320–336, 2008.

[8] R. Corless, G. Gonnet, D. Hare, D. Jeffrey, and D. Knuth, "On the LambertW function," *Adv. Comput. Math.*, vol. 5, no. 1, pp. 329–359, 1996.

[9] A. Robbins, *Effective Awk Programming*. Sebastopol, CA: O'Reilly & Associates, 2001.

[10] S. A. P. Haddad, W. A. Serdijn, S. A. P. Haddad, and W. A. Serdijn, "The evolution of pacemakers: An electronics perspective," in *Ultra Low-Power Biomedical Signal Processing*, ser. Analog Circuits and Signal Processing. New York: Springer, 2009, pp. 13–31.

[11] X. Zou, X. Xu, L. Yao, and Y. Lian, "A 1-V 450-nW fully integrated programmable biomedical sensor interface chip," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1067–1077, Apr. 2009.

[12] A. Auricchio, W. Hartung, C. Geller, and H. Klein, "Clinical relevance of stored electrograms for implantable cardioverter-defibrillator (ICD) troubleshooting and understanding of mechanisms for ventricular tachyarrhythmias," *Amer. J. Cardiol.*, vol. 78, pp. 33–41, Sep. 1996.

[13] B. Nowak, "Taking advantage of sophisticated pacemaker diagnosis," *Amer. J. Cardiol.*, vol. 83, pp. 172–179, Mar. 1999.

[14] J. Rodrigues, L. Olsson, T. Sörnmo, and V. Öwall, "Digital implementation of a wavelet-based event detector for cardiac pacemakers," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 52, no. 12, pp. 2686–2698, Dec. 2005.

[15] M. Åström, S. Olmos, and L. Sörnmo, "Wavelet-based event detection in implantable cardiac rhythm management devices," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 3, pp. 478–484, Mar. 2006.

[16] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. Int. Symp. Low Power Electron. Design*, New York, 2005, pp. 20–25.

[17] E. Seevinck, F. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol. SSC-22, no. 5, pp. 748–754, Oct. 1987.

[18] J. Kwong and A. Chandrakasan, "Variation-driven device sizing for minimum energy sub-threshold circuits," presented at the Int. Symp. Low Power Electronics and Design, New York, 2006.

[19] M. Hansen, H. Yalcin, J. Hayes, D. Syst, and I. Kokomo, "Unveiling the ISCAS-85 benchmarks: A case study in reverseengineering," *IEEE Design Test Comput.*, vol. 16, no. 3, pp. 72–80, Jul.–Sep. 1999.

[20] O. C. Akgun, "Energy efficiency enhancement of sub-threshold digital cmos—modeling, technology selection, and architectural exploration," Ph.D. dissertation, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, 2009.

[21] O. C. Akgun, J. Rodrigues, and J. Sparsø, "Minimum-energy subthreshold self-timed circuits: Design methodology and a case study," in *Proc. 16th IEEE Int. Symp. Asynchron. Circuits Syst.*, 2010, pp. 41–51.

[22] J. Rodrigues, O. C. Akgun, and V. Owall, "A <1 pJ sub-vt Cardiac Event Detector in 65 nm LL-HVT CMOS," in *Proc. 18th IEEE/IFIP Int. Conf. VLSI Syst.-on-Chip*, 2010, pp. 253–258.

[23] C. Enz, F. Krummenacher, and E. Vittoz, "An analytical mos transistor model valid in all regions of operation and dedicated to low-voltage and low-current applications," *Analog Integr. Circuits Signal Process.*, vol. 8, pp. 83–114, Jul. 1995.

Omer Can Akgun (S'01–M'11) received the M.Sc. degree in electrical engineering from the Ohio State University, Columbus, in 2004 and the Ph.D. degree in microelectronics from École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2009.

He was a Postdoctoral Researcher at Lund University, Lund, Sweden, until 2010. His main research interests are the design and design automation of ultra-low energy mixed-signal subthreshold circuits, asynchronous and energy harvesting systems, brain-circuit interfaces, and implantable self-powered biomedical systems.

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

Currently, he is an Assistant Professor in the Department of Electrical and Information Technology, Lund University. From 2005 to 2008, he was an ASIC Process Leader in the Digital ASIC Department, Ericcsson Mobile Platforms, Lund, and he rejoined his current department in 2008. He has contributed to 20 ASICs in industry and academia. His main research interest is the modeling and implementation of digital and mixed-mode microelectronics and architectures for high-performance ultra-low-power design which may be operated with an aggressively scaled supply voltage, with a focus on biomedical circuits and systems.

Dr. Rodrigues has been a TC member of the biomedical circuits and systems society since 2010.

Yusuf Leblebici (M'90–SM'98–F'09) received the B.Sc. and M.Sc. degrees in electrical engineering from Istanbul Technical University, Istanbul, Turkey, in 1984 and in 1986, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign (UIUC) in 1990.

Between 1991 and 2001, he was a faculty member at UIUC, at Istanbul Technical University, and at Worcester Polytechnic Institute (WPI), Worcester, MA. In 2000–2001, he was the Microelectronics Program Coordinator at Sabanci University. Since 2002, he has been a Chair Professor at the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, and Director of the Microelectronic Systems Laboratory. His research interests include the design of high-speed CMOS digital and mixed-signal integrated circuits, computer-aided design of very-large-scale integrated (VLSI) systems, intelligent sensor interfaces, modeling and simulation of semiconductor devices, and VLSI reliability analysis. He is the coauthor of four textbooks, namely, *Hot-Carrier Reliability of MOS VLSI Circuits* (Kluwer, 1993), *CMOS Digital Integrated Circuits: Analysis and Design* (McGraw Hill, 2002), *CMOS Multichannel Single-Chip Receivers for Multi-Gigabit Optical Data Communications* (Springer, 2007), and *Fundamentals of High Frequency CMOS Analog Integrated Circuits* (Cambridge University Press, 2009), as well as more than 200 articles published in various journals and conferences. He was an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II: EXPRESS BRIEFS and IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATED SYSTEMS.

Dr. Leblebici was also General Co-Chair of the 2006 European Solid-State Circuits Conference, and the 2006 European Solid State Device Research Conference (ESSCIRC/ESSDERC). He has been elected as Distinguished Lecturer of the IEEE Circuits and Systems Society for 2010–2011.

Viktor Öwall (M'90) received the M.Sc. and Ph.D. degrees in electrical engineering from Lund University, Lund, Sweden, in 1988 and 1994, respectively.

From 1995 to 1996, he was a Postdoctoral Researcher with the Electrical Engineering Department, University of California, Los Angeles, where he mainly worked in the field of multimedia simulations. Since 1996, he has been with the Department of Electrical and Information Technology, Lund University, which he is currently heading. His main research interest is in the field of digital hardware implementation, especially algorithms and architectures for wireless communication, image processing, and biomedical applications. His current research projects include combining theoretical research with hardware implementation aspects in the areas of pacemakers, baseband processing, video processing, and digital holography.

Dr. Öwall was an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II: ANALOG AND DIGITAL SIGNAL PROCESSING from 2000 to 2002 and of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I: REGULAR PAPERS from 2007 to 2009.