Minimum Energy Tracking Loop With Embedded DC–DC Converter Enabling Ultra-Low-Voltage Operation Down to 250 mV in 65 nm CMOS

Yogesh K. Ramadass and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—Minimizing the energy consumption of battery-powered systems is a key focus in integrated circuit design. This paper presents an energy minimization loop, with on-chip energy sensor circuitry, that can dynamically track the minimum energy operating voltage of arbitrary digital circuits with changing workload and operating conditions. An embedded DC–DC converter which enables this minimum energy operation is designed to deliver load voltages between 0.25 V to 0.7 V. The minimum energy tracking loop along with the DC–DC converter and test circuitry were fabricated in a 65 nm CMOS process. The area overhead of the control loop is only 0.05 mm^2. Measured energy savings of the order of 50%–100% are obtained on tracking the minimum energy point (MEP) as it varies with workload and temperature. The DC–DC converter delivers load voltages as low as 250 mV and achieved an efficiency $>$80% while delivering load powers of the order of 1 μW and higher from a 1.2 V supply.**

***Index Terms*—DC-DC conversion, energy sensor circuitry, minimum energy point, minimum energy tracking loop, subthreshold logic, ultra-low-voltage operation.**

I. Introduction

THE energy efficiency of digital circuits continues to be a major factor in determining the size and weight of battery-operated electronics. Emerging applications like wireless microsensor networks [1], [2] and implantable medical electronics [3] are severely energy constrained. For applications like implantable medical devices that are battery operated, though the required speed of operation is low, the battery is expected to last through the lifetime of the device, without the possibility of a recharge. On the other hand, a key requirement in the design of sensor systems is constraining the power dissipation of the system below 10 μW [4] which will allow operation strictly using scavenged energy. So, irrespective of the mode of power delivery, there is a severe constraint on the energy consumed per desired operation of these devices. This forces us to look into system-level energy management techniques to further these applications. Subthreshold operation [5], [6] is a technique wherein the circuits are operated at a $V_{\rm DD}$ below the threshold voltage of its devices. While subthreshold operation promises an order of magnitude reduction in power dissipation over above-threshold approaches, it comes at the cost of circuit speed. Since the performance requirements are quite relaxed in many of these energy-constrained applications, subthreshold operation is a feasible approach to minimize energy.

Operating at the minimum energy operating voltage of digital circuits has been proposed as a solution for energy-critical applications [5]. The minimum energy point (MEP) is defined as the operating voltage at which the total energy consumed per desired operation of a digital circuit is minimized. Switching energy of digital circuits reduces quadratically as $V_{\rm DD}$ is decreased below $V_{\rm T}$, while the leakage energy increases exponentially. These opposing trends result in the minimum energy point. The MEP is not a fixed voltage for a given circuit, and can vary widely depending on its workload and environmental conditions (e.g., temperature). By tracking the MEP as it varies, energy savings of 50%–100% has been demonstrated [7] and even greater savings can be achieved in circuits dominated by leakage.

In this paper, a 65 nm CMOS circuit that can dynamically track the MEP of a digital circuit with varying operating conditions is presented. The tracking loop employs on-chip energy sensor circuitry that calculates the actual energy consumed per desired operation of the load circuit. Embedded within the tracking loop is a switching DC–DC converter that can efficiently deliver supply voltages down to 250 mV at ultra-low load power levels of 1 μW, thereby enabling minimum energy operation.

II. Minimum Energy Point

The total energy consumed per operation by a digital circuit can be split into two components: an active energy component and a leakage energy component. The active energy consumed per operation assuming rail to rail swing is given by

$$E_{\rm ACT} = C_{\rm eff} V_{\rm DD}^2 \tag{1}$$

where $C_{\rm eff}$ is the average effective switched capacitance per desired operation of the digital circuit. The active energy consumed per operation scales down quadratically with $V_{\rm DD}$, as can be seen by the active energy curve in Fig. 1 which shows the energy/operation breakdown for a 7-tap finite impulse response (FIR) filter in 65 nm CMOS.

The leakage energy per operation is due to the subthreshold leakage power of the digital circuit, which integrates over the time period of an operation, and is given by

$$E_{\rm LEAK} = (I_{\rm leak} V_{\rm DD}) T_{\rm op} = K W_{\rm eff} C_g L_{\rm DP} V_{\rm DD}^2 e^{\frac{-V_{\rm DD}}{nV_{\rm th}}} \tag{2}$$

Manuscript received April 27, 2007; revised November 5, 2007. This work was supported by the Defense Advanced Research Projects Agency (DARPA). Chip fabrication was provided by Texas Instruments.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: ryogesh@mtl.mit.edu).

Digital Object Identifier 10.1109/JSSC.2007.914720

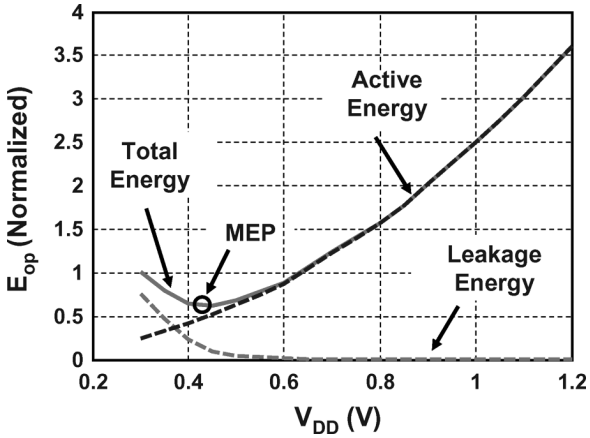

Fig. 1. Active, leakage, and total energy per operation curves showing the minimum energy point (0.42 V) for a 7-tap FIR filter implemented in 65 nm CMOS.

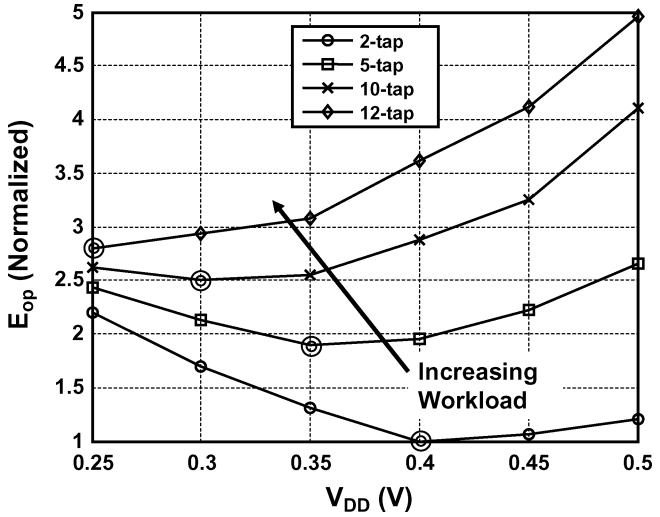

Fig. 2. Simulated energy/operation curves with change in number of taps of a 12-tap FIR filter at 27 °C.

where $T_{\rm op}$ is the time required to complete an operation, K is a delay-fitting parameter, $W_{\rm eff}$ is the average effective width of the leaking transistors, C_g is the output capacitance of a characteristic inverter, $L_{\rm DP}$ is the depth of the critical path in characteristic inverter delays, n is the subthreshold slope factor, and $V_{\rm th} = kT/q$. A more detailed description of the above-mentioned terms and how the equation is derived is given in [8]. The leakage energy component is negligible at higher voltages, but increases exponentially as $V_{\rm DD}$ is decreased close to subthreshold voltages, as can be seen from the leakage energy curve in Fig. 1. This is because of the exponential increase in the time to complete an operation $(T_{\rm op})$ at subthreshold voltages. The opposing trends of the active and leakage energy components with $V_{\rm DD}$ give rise to a minimum in the total energy consumed per operation. The minimum energy operating voltage usually falls in the subthreshold regime of operation and for the 7-tap FIR filter is at 420 mV. Adding equations (1) and (2), the total energy consumed per desired operation is given by

$$E_{\rm TOT} = V_{\rm DD}^2 \left[C_{\rm eff} + KW_{\rm eff} C_g L_{\rm DP} e^{\frac{-V_{\rm DD}}{nV_{\rm th}}} \right]. \qquad (3)$$

$E_{\rm TOT}$ is a nonlinear function of $V_{\rm DD}$ and can be minimized with respect to $V_{\rm DD}$ by differentiating (3) and equating it to 0. An analytical solution for the minimum energy operating voltage has been derived in [8]. The analytical solution shows that the minimum energy operating voltage for a particular circuit is not a fixed voltage. The MEP depends on temperature, $C_{\rm eff}$, $W_{\rm eff}$, and $L_{\rm DP}$. Any relative increase in the active energy component of the circuit $(C_{\rm eff})$ due to an increase in the workload or activity of the circuit decreases the minimum energy operating voltage. On the other hand, a relative increase of the leakage energy component due to an increase in temperature or the duration of leakage over an operation $(L_{\rm DP})$ pushes the minimum energy operating voltage to go up. This makes the circuit go faster, thereby not allowing it to leak for a longer time. Fig. 2 shows the simulated energy/operation curves with change in workload for a 12-tap FIR filter. The workload of the circuit is varied by turning ON or OFF the taps of the filter. The leakage energy/operation remains fixed as the number of taps are varied. It can be seen that the minimum energy operating voltage decreases as the number of taps and thereby the workload is increased. When the FIR filter operates with all 12 of its taps active, the MEP is at 250 mV. On decreasing the number of taps to two, the MEP increases to 400 mV. If the FIR filter working with two taps active continues to operate at the MEP corresponding to its 12-tap scenario (250 mV), the energy/operation curves show that 2.2X more energy is consumed, or in other words, a 120% savings in energy consumed can be achieved on tracking the MEP as it moves. In general, energy savings of the order of 50%–150% are possible by tracking the MEP as it changes with operating conditions and this motivates the design of a circuit that can dynamically track the MEP.

III. Minimum Energy Tracking Loop

Fig. 3 shows the architecture of the minimum energy tracking loop. The objective of this loop is to track the minimum energy operating voltage of the load circuit (FIR filter). The load circuit is powered from an off-chip voltage source (1.2 V) through a DC–DC converter and is clocked by a critical path replica ring oscillator which automatically scales the clock frequency of the FIR filter with change in load voltage. The energy sensor circuitry calculates on-chip the energy consumed per operation of the load circuit at a particular operating voltage. It then passes the estimate of the energy/operation $(E_{\rm op})$ to the energy minimizing algorithm, which uses the $E_{\rm op}$ to suitably adjust the reference voltage to the DC–DC converter. The DC–DC converter then tries to get $V_{\rm DD}$ close to the new reference voltage and the cycle repeats until the minimum energy point is achieved. The only off-chip components of this entire loop are the filter passives of the DC–DC converter.

A. Energy Sensing Technique

The key element in the minimum energy tracking loop is the energy sensor circuit which computes the $E_{\rm op}$ of the load circuit at a given reference voltage. Methods to measure $E_{\rm op}$, by sensing the current flow through the DC–DC converter's inductor, dissipate a significant amount of overhead power. The

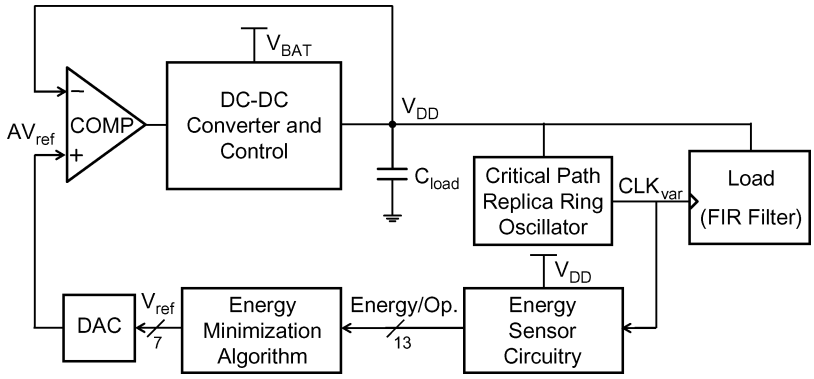

Fig. 3. Block diagram of the minimum energy tracking loop and embedded DC–DC converter.

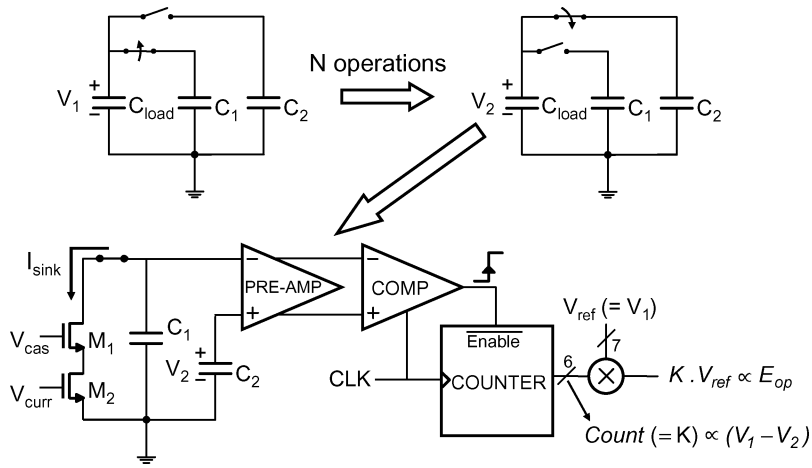

Fig. 4. Circuitry to compute energy/operation ($E_{\rm op}$) at a given operating voltage.

approach is more complicated at subthreshold voltages because the current levels are very low. Further, an estimate of the energy consumed per operation is what is required and not just the current which only gives an idea of the load power. Our proposed method to estimate $E_{\rm op}$ does not require any high-gain amplifiers or analog circuit blocks. The DC–DC converter while operating in steady state keeps the output voltage close to the reference voltage set. Just before the energy sense cycle begins, the DC–DC converter is disabled. The energy sense cycle consists of N operations of the digital circuit where the value N can be 32 or 64. Assuming that the voltage across the storage capacitor of the DC–DC converter, $C_{\rm load}$, falls from the reference voltage V_1 to V_2 in the course of N operations of the digital circuit, $E_{\rm op}$ at the voltage V_1 is

$$E_{\rm op} = \frac{C_{\rm load}\left(V_1^2 - V_2^2\right)}{2N}. \tag{4}$$

To measure $E_{\rm op}$ accurately, V_2 should be close in value (within 20 mV) to V_1. Measuring $E_{\rm op}$ by digitizing V_1 and V_2 using conventional ADCs would require at least 11 bits of precision in the ADC. This could prove costly in terms of power consumed. Our proposed energy-efficient approach to obtain $E_{\rm op}$ is to observe that, by design, V_1 is very close to V_2. Thus, the following simplification can be applied within an acceptable error.

$$\begin{aligned} E_{\rm op} &= \frac{C_{\rm load}(V_1+V_2)(V_1-V_2)}{2N} \\ &\approx \frac{C_{\rm load}V_1(V_1-V_2)}{N} \qquad (5) \\ E_{\rm op} &\propto V_1(V_1-V_2) \qquad (6) \end{aligned}$$

From (6), it can be seen that the energy consumed per operation is directly proportional to the product of V_1 and $V_1 - V_2$. Since the digital representation of V_1, which is the reference voltage to the DC–DC converter is already known, only the digital value for $V_1 - V_2$ is required to estimate $E_{\rm op}$.

Fig. 4 shows the voltage difference measuring circuitry. Before starting an N operation energy sense cycle, the voltage V_1 across $C_{\rm load}$ is sampled on a small 5 pF capacitor C_1. The DC–DC converter is then disabled. The digital circuit runs for N operations using the energy stored in $C_{\rm load}$, thereby bringing down the voltage across $C_{\rm load}$ to V_2 ($< V_1$). At the end of the N operation energy sense cycle, the value V_2 is sampled across another 5 pF capacitor C_2. Subsequently, the DC–DC converter is enabled and normal operation of the digital circuit continues. At this point, a current sink of high output impedance

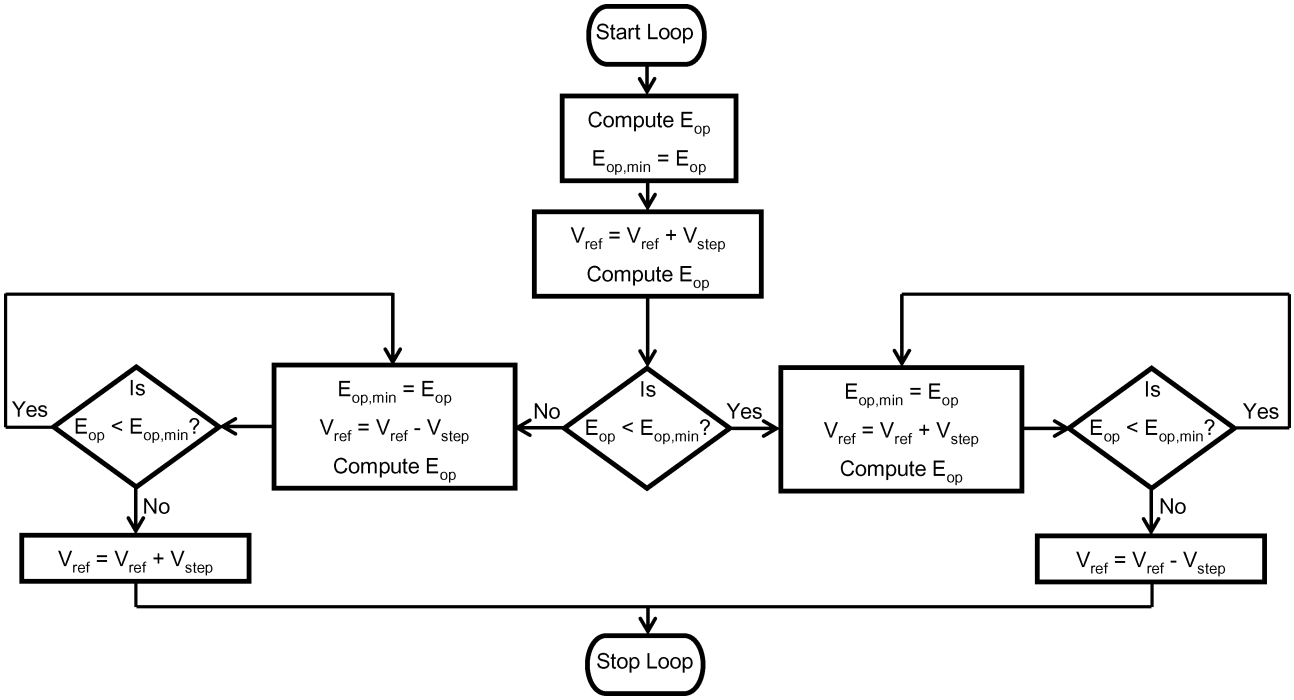

Fig. 5. Flowchart showing the operation of the minimum energy tracking algorithm.

(M_1, M_2) connected across C_1 turns ON, and the voltage across C_1 is allowed to droop. At the same time a counter is enabled. The counter is run by a fixed frequency clock. As the voltage across C_1 droops, a preamplifier-comparator system compares the voltage across C_1 with V_2. The same clock that is used in the counter is also used to clock the dynamic comparator. The fixed frequency clock, together with the constant current sink that drains C_1, quantizes voltage into time steps, as in an integrating ADC [9]. The number of fixed frequency clock cycles required for C_1 to droop down to V_2 is directly proportional to $V_1 - V_2$. This can be seen from the following equation:

$$C_1(V_1 - V_2) = I_{\text{sink}} \times T_{\text{droop}} \approx I_{\text{sink}}(K \times T_{\text{clk}}) \quad (7)$$

$$(V_1 - V_2) \propto K \quad (8)$$

where T_{droop} is the time taken for the voltage across C_1 to droop down to V_2, K is the number of clock cycles registered in the counter, and T_{clk} is the time period of the fixed frequency clock. The quantity K is then digitally multiplied with V_1 which is the reference voltage V_{ref} to the DC–DC converter. The product of these two quantities gives an estimate of the energy consumed per operation by the digital circuit at voltage V_1. This estimate is passed on to the energy minimization algorithm block. The estimate obtained is a normalized representation of the absolute value of the energy consumed per operation.

B. Energy Minimization Algorithm

Once the estimate of the energy per operation is obtained, the minimum energy tracking algorithm uses this to suitably adjust the reference voltage to the DC–DC converter. The minimum energy tracking algorithm is a slope-tracking algorithm which makes use of the single minimum, concave nature of the E_{op} versus V_{DD} curve. This can be seen from the E_{op} curves in Fig. 2. The algorithm described by the flowchart in Fig. 5 starts by setting the reference voltage V_{ref} to some initial value. The energy per operation at this voltage is computed and stored in a minimum energy register $(E_{\text{op,min}})$. The tracking loop then automatically increments V_{ref} by one voltage step. Once V_{DD} settles at this newly incremented voltage, E_{op} is computed again and is compared with the value stored in the minimum energy register. At this point, if the newly computed E_{op} is found to be smaller, the loop then just keeps incrementing V_{ref} at fixed voltage steps, while at the same time updating $E_{\text{op,min}}$ until the minimum is achieved. The other possibility is that the newly computed energy per operation is higher than that stored in the minimum energy register. In this case, the loop changes direction and begins to decrement V_{ref}. The loop keeps decrementing V_{ref} until the E_{op} calculated is higher than $E_{\text{op,min}}$ at which time the loop increments V_{ref} by one voltage step to get to the MEP and shuts down.

The voltage step used by the tracking algorithm is usually set to 50 mV. A large voltage step leads to coarse tracking of the MEP, with the possibility of missing the MEP. On the other hand, keeping the voltage step too small might lead to the loop settling at the non-minimum voltage due to errors involved in computing E_{op} which will be described in Section V. The E_{op} versus V_{DD} curve is shallow near the MEP and hence a 50 mV step leads to a very close approximation of the actual minimum energy consumed per operation. The MEP tracking loop can be enabled by a system controller as needed depending on the application, or periodically by a timer to track temperature variations.

IV. Low-Power DC–DC Converter

This section talks about the design of the DC–DC converter that enables minimum energy operation. Since the minimum en-

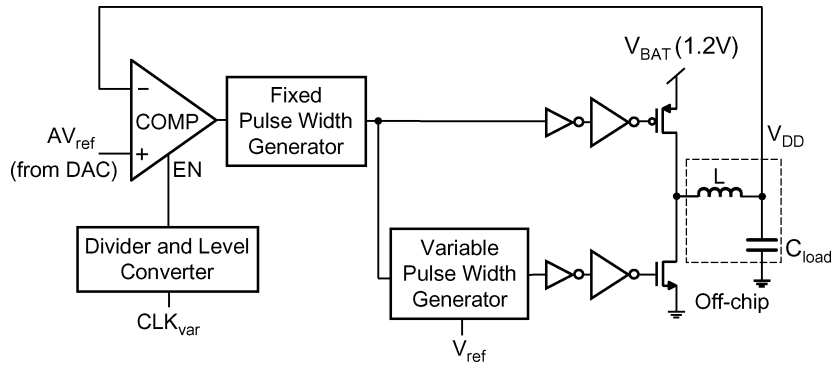

Fig. 6. Architecture of the DC–DC converter employing a pulse frequency modulation (PFM) mode of control.

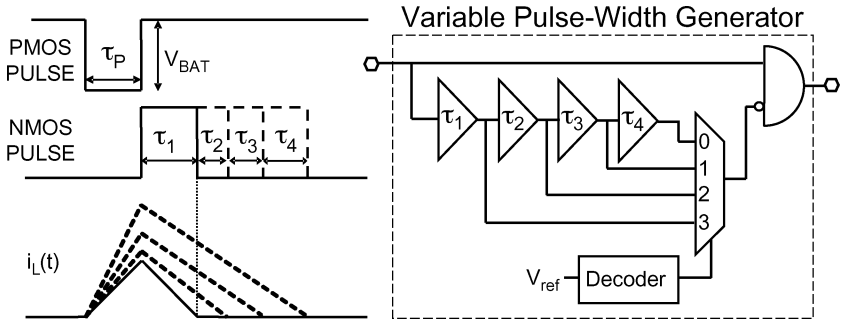

Fig. 7. NMOS pulse width determining circuitry using open loop control. The time delays are chosen to turn the nMOS power transistor OFF as the inductor current approaches zero.

ergy operating voltage usually falls in the subthreshold regime of operation, the DC–DC converter is designed to deliver load voltages from 250 mV to around 700 mV. The power consumed by digital circuits at these subthreshold voltages is exponentially smaller and hence the DC–DC converter needs to deliver efficiently load power levels of the order of microwatts. This demands extremely simple control circuitry design with minimal overhead power to get good efficiency. The DC–DC converter shown in Fig. 6 is a synchronous rectifier buck converter with off-chip filter components. It employs a pulse frequency modulation (PFM) [10] mode of control in order to get good efficiency at the ultra-low load power levels that the converter needs to deliver. The reference voltage to the converter is set digitally by the minimum energy tracking loop and is converted to an analog value by an on-chip DAC before it is fed to the comparator. The comparator compares $V_{\rm DD}$ with this reference voltage and when it is found to be smaller, generates a pulse of fixed width to turn the pMOS power transistor ON and ramp up the inductor current. A variable pulse width generator to achieve zero-current switching is used for the nMOS power transistor. The comparator is clocked by a divided and level converted version of the system clock which feeds the load FIR filter.

A. Approximate Zero-Current Switching

The ultra-low load power levels demand extremely simple control circuitry to achieve good efficiency. This precludes the usage of high gain amplifiers to detect zero-crossing and thereby do zero-current switching [11]. In order to keep the control circuitry simple and consume little overhead power, an all-digital open loop control as shown in Fig. 7 is used to achieve zero-current switching. The variable pulse width generator block which accomplishes this functions as follows. When the comparator senses that $V_{\rm DD}$ has fallen below the reference voltage, a pMOS ON pulse of fixed pulse width $\tau_{\rm P}$ is generated. This ramps up the inductor current from zero. Once the pMOS is turned OFF, the nMOS power transistor is turned ON after a fixed delay. This ramps down the inductor current. Ideally, in the discontinuous conduction mode (DCM) used in this implementation, the nMOS has to be turned OFF just when the inductor current reaches zero. The amount of time it takes for the inductor current to reach zero is dependent on the reference voltage set and in steady state, the ratio of the nMOS to pMOS ON-times is given by

$$\frac{\tau_N}{\tau_P} = \frac{V_{\rm BAT} - V_{\rm DD}}{V_{\rm DD}} \tag{9}$$

where $\tau_{\rm N}$ and $\tau_{\rm P}$ are the nMOS and pMOS ON-times and $V_{\rm BAT}$ is the battery voltage. Thus, by fixing $\tau_{\rm P}$, the values of $\tau_{\rm N}$ for specific load voltages can be predetermined. The variable pulse width generator block then suitably multiplexes these predetermined delays depending on the reference voltage set to achieve approximate zero-current switching. Increasing the number of these delay elements and the complexity of the multiplexer block gives a better approximation to zero-current switching. Since only the ratios of the nMOS and pMOS ON-time pulse widths need to match, this scheme is independent of absolute delay values and any tolerance in the inductor value.

B. Efficiency of the DC–DC Converter

The efficiency of the DC–DC converter operating in the PFM mode can be given by a simple equation as follows:

$$\eta = \frac{E_{\text{load}}}{E_{\text{load}} + E_{\text{cond}} + E_{\text{sw}} + E_{\text{para}} + E_{\text{control}}}. \tag{10}$$

Here, E_{load} is the energy delivered by the converter to the load every switching cycle. E_{cond} is the amount of energy lost every switching cycle due to conduction losses which occur primarily due to current flow through power transistors with finite on-state resistances. The gate-switching losses labeled as E_{sw} are due to the energy lost turning ON or OFF the gates of the power transistors every switching cycle. The energy (E_{para}) lost due to non-zero voltage switching of the intermediate parasitic capacitance at the drains of the power transistors dominates the timing error related losses [11]. The parasitic capacitance also includes the package and board parasitics and is close to 2 pF for the implemented converter. While zero-voltage switching was not actively tackled in this implementation of the DC–DC converter, E_{para} was reduced by introducing a finite delay between the pMOS and nMOS ON pulses. Also, the energy delivered to the load per cycle, E_{load} was increased by properly choosing the pMOS ON time τ_{P} and the inductor value. This helps in reducing the contribution of the parasitic power losses compared to E_{load}, thereby increasing the efficiency.

E_{control} is the energy lost every switching cycle due to switching and leakage related losses in the control circuitry. This loss mechanism affects significantly the low load efficiency of the DC–DC converter. PFM mode control reduces the switching frequency of the converter as load power drops. While the other loss mechanisms remain constant as the switching frequency is varied, the control circuitry leakage loss integrates over the time period of a switching cycle and hence increases in value as the load power and thereby the switching frequency decreases. This leads to a drop in efficiency as the load power decreases. In order to minimize the control circuitry losses, the DC–DC converter uses a simple all-digital PFM mode control which does not consume any static power. The approximate zero-current switching block removes the need for any high gain amplifiers. The clock for the reference voltage comparator is derived from the critical path replica ring oscillator which feeds the digital circuit. This allows the comparator clock to scale automatically with V_{DD} and hence the load power, eliminating unnecessary comparisons.

V. Error in Calculating Energy/Operation

The approximations introduced in the E_{op} formulation (5) and the non-idealities of the comparator introduce errors in the computed energy per operation. The major sources of error are analyzed below and an estimate of the energy difference that can be resolved is provided.

A. Error due to Voltage Approximation

The actual value of energy lost in the storage capacitor C_{load} when the voltage across it falls from V_1 to V_2 is proportional to $V_1^2 - V_2^2$. The error introduced in approximating V_2 to be equal to V_1 is given by

$$\begin{aligned}\delta E\text{op}_{\text{app}} &= C_{\text{load}} V_1 (V_1 - V_2) - \frac{1}{2} C_{\text{load}} \left(V_1^2 - V_2^2\right) \\ &= \frac{1}{2} C_{\text{load}} (V_1 - V_2)^2. \end{aligned} \tag{11}$$

The relative error introduced in the calculation of the energy is given by

$$\frac{\delta E\text{op}_{\text{app}}}{E\text{op}} = \frac{\frac{1}{2} C_{\text{load}} (V_1 - V_2)^2}{\frac{1}{2} C_{\text{load}} \left(V_1^2 - V_2^2\right)} = \frac{V_1 - V_2}{V_1 + V_2}. \tag{12}$$

The relative error is worst when V_1 is small, i.e., when E_{op} is calculated at the lowest allowed operating voltage. Since the lowest voltage the circuit is allowed to operate at is 270 mV and assuming V_1 falls from 270 to 260 mV in the course of N operations, the relative error introduced due to voltage approximation is equal to 1.9% (10/530). To get an estimate of the minimum energy difference that can be resolved, it is assumed that the actual E_{op} is the same at both 270 and 320 mV. A 10 mV drop at 270 mV has the same energy consumption as a 9 mV drop at 320 mV. The computed E_{op} at 320 mV would overestimate the actual energy consumed per operation by 1.4% (9/631). Thus, 0.5% of the energy consumed at 270 or 320 mV is the minimum that can be resolved correctly. This value is small enough to not considerably affect the accuracy of the minimum energy tracking circuitry.

B. Error due to Comparator Offset

Another major source of error in the E_{op} calculation is the finite offset error in the dynamic comparator that is used to compare voltages across the small capacitors C_1 and C_2. The devices that make up the comparator are sized up in area in order to reduce the effect of variations that decrease proportional to the increase in area of a MOSFET [12]. Further, the pre-amplifier that goes before the comparator helps in reducing the effect of the offset in the comparator. Even after taking these measures, the offset in the comparator is of the order of 1 mV. Assuming the offset is ΔV, which may be positive or negative, the error introduced in computing E_{op} is

$$\begin{aligned}\delta E\text{op}_{\text{off}} &= C_{\text{load}} V_1 (V_1 - V_2) - C_{\text{load}} V_1 (V_1 - V_2 - \Delta V) \\ &= C_{\text{load}} V_1 \Delta V. \end{aligned} \tag{13}$$

The relative error introduced in the calculation of E_{op} due to offset error in the comparator is given by

$$\frac{\delta E\text{op}_{\text{off}}}{E\text{op}} = \frac{C_{\text{load}} V_1 \Delta V}{C_{\text{load}} V_1 (V_1 - V_2)} = \frac{\Delta V}{V_1 - V_2}. \tag{14}$$

An offset of 1 mV leads to a 10% relative error in the calculation of the energy consumed per operation when a typical value of 10 mV for $V_1 - V_2$ is plugged in. The comparator offset is a static error and hence the same type of error occurs when E_{op} at a different voltage is computed. Assuming that the E_{op} is the same at operating voltages V_1 and V_3, there should be no difference in the computed E_{op} if the comparator was ideal. Due to the comparator offset of ΔV, the obtained energy difference

is $C_{\text{load}}(V_3 - V_1)\Delta V$. The relative error in energy difference is given by

$$\frac{\delta E\text{op}, \text{diff}_{\text{off}}}{E\text{op}} = \frac{(V_3 - V_1)\Delta V}{V_1(V_1 - V_2)}. \quad (15)$$

$(V_3 - V_1)$ is usually of the order of 50 mV, ΔV is 1 mV, V_1 can go as low as 270 mV, and $(V_1 - V_2)$ is around 10 mV. This gives a relative error of 1.85%. This sets a lower limit on how finely the minimum energy operating voltage can be tracked. Between voltages where the difference in energy consumption is less than 1.85% of each other, the loop cannot find the actual minimum. This is the reason why very small voltage steps can be detrimental. Since the error value of 1.85% is very low, not much is lost by operating at the non-optimal point in terms of total energy expended.

C. Effect of DC–DC Converter Efficiency Change With Load Power on the MEP

The MEP obtained by using the minimum energy tracking algorithm described in Section III does not take into account the efficiency of the DC–DC converter. This would not lead to any error if the efficiency plot of the DC–DC converter stays constant with load power. However, in reality the DC–DC converter efficiency drops with load power due to the leakage loss in the control circuitry. Since the load power itself drops with V_{DD}, it essentially leads the efficiency of the DC–DC converter to drop with V_{DD}. This introduces errors in calculating the actual MEP taking into account the efficiency loss of the DC–DC converter. Consider that the E_{op} is the same at 370 mV and at 420 mV. The power consumed by the FIR filter at 420 mV is 10 μW. From Fig. 11, it can be seen that the efficiency of the DC–DC converter at 10 μW is 84%. Thus, the actual energy extracted from the battery for every operation performed at 420 mV is 1.19 (1/.84) times the E_{op} computed by the energy sense circuitry. Assume that in the worst case, the power consumed by the FIR filter at 370 mV falls down to 1 μW. The efficiency of the DC–DC converter at 370 mV would then be 81% (Fig. 11), which means that the actual energy extracted from the battery is 1.23 (1/.81) times the E_{op} computed at 370 mV. Hence, in this worst case scenario, the minimum energy tracking algorithm would choose the wrong MEP as long as the energy difference between the two voltages in question is lesser than 4% (1.23–1.19). This value is much smaller than the potential energy savings that can be obtained by tracking the minimum energy operating voltage using the procedure described in Section III.

VI. Measurement Results

A 7-tap FIR filter shown in Fig. 8 was designed to act as the test load for the minimum energy tracking loop. The filter is capable of operation down to 250 mV. The workload of the filter is varied by changing the coefficient of the individual taps. Leakage energy per operation remains fixed as the number of taps is varied. Therefore, increasing the number of taps increases the ratio of active to leakage energy per operation and vice-versa. This feature is used to test the operation of the minimum energy tracking loop.

The test chip containing the minimum energy tracking loop, the DC–DC converter and the FIR filter was fabricated in Texas Instruments' 65 nm CMOS process. Fig. 9 shows a die photo of the implemented chip. The die area occupied is 1.05×1.12 mm^2 of which the active circuit area is only 0.23 mm^2. The minimum energy tracking circuitry which comprises the energy sense capacitors C_1 and C_2 and the energy minimizing block (EMB) occupies just 0.05 mm^2 in area.

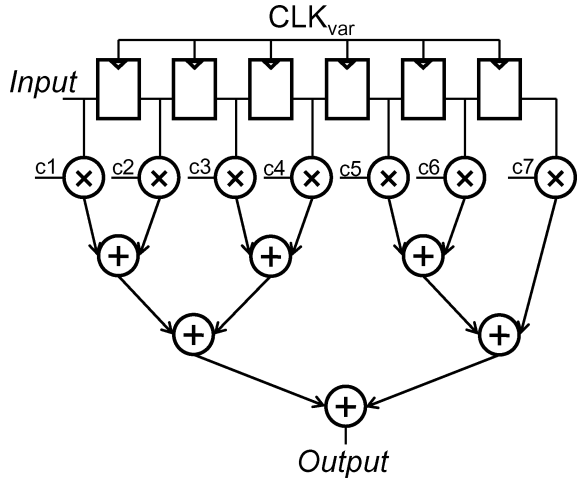

Fig. 8. Block diagram of a 7-tap FIR filter. c1, c2,. . ., c7 are the tap coefficients which can be set to zero if the tap is not used. The registers can be clock gated independent of each other.

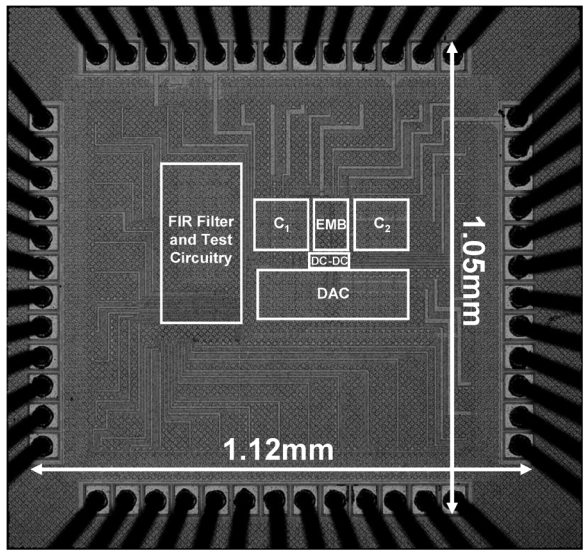

Fig. 9. Die photo of the test chip in 65 nm CMOS. EMB is the energy minimizing block which comprises the energy sensor circuitry and the energy minimization algorithm.

Fig. 10(a) shows a measured waveform of the minimum energy tracking loop in operation. For this measurement, only one of the taps of the FIR filter was activated. The loop begins by setting V_{DD} to 420 mV and the energy/operation is calculated at this voltage. In accordance to the minimum energy tracking algorithm (Fig. 5), the loop increments V_{DD}, by one voltage step, to 470 mV. A 50 mV voltage step is used for the measurement. Finding the E_{op} at 470 mV to be higher, the loop changes direction and starts reducing V_{DD}. Initially, V_{DD} is reduced to 370 mV, and the E_{op} is calculated again at this new voltage. The

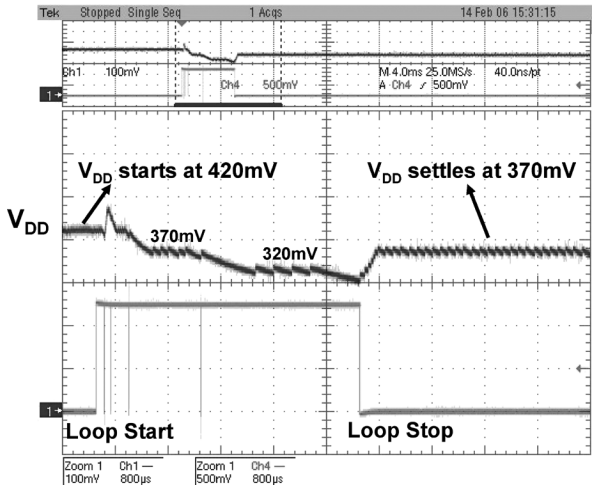

(a)

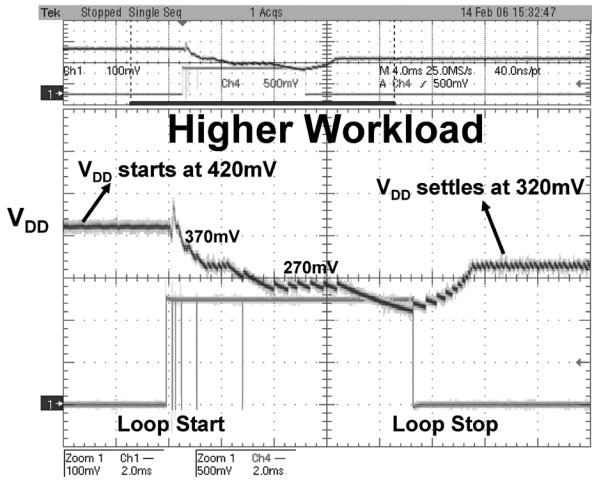

(b)

Fig. 10. (a) Measured waveform showing the minimum energy tracking loop in operation. $V_{\rm DD}$ starts at 420 mV and is then increased to 470 mV. The loop then changes direction and reduces $V_{\rm DD}$ to 370 mV and 320 mV before settling at the MEP of 370 mV. (b) Measured waveform showing the minimum energy tracking loop in operation for a higher workload scenario. Here, the loop settles at the MEP of 320 mV.

$E_{\rm op}$ calculated here is lower than that at 420 mV. So, the loop updates the minimum energy register and continues reducing $V_{\rm DD}$ down to 320 mV. It calculates $E_{\rm op}$ again at 320 mV where it finds that the newly computed $E_{\rm op}$ is higher. This makes the loop realize that the MEP is at 370 mV. The loop then sets $V_{\rm ref}$ to 370 mV, reverses direction again, and shuts down.

For the measurement in Fig. 10(b), the number of active taps of the FIR filter was increased to seven. This increase in workload pushes the minimum energy operating voltage down to 320 mV. Here again, the tracking loop performs operations similar to the scenario explained above. In this scenario, however, the $E_{\rm op}$ that the loop calculates at 320 mV is found to be smaller than the one at 370 mV. So, the loop decrements $V_{\rm ref}$ further to 270 mV where it finds that the $E_{\rm op}$ is higher. The loop then realizes that the MEP is at 320 mV, sets $V_{\rm ref}$ to 320 mV and shuts down until it is called into action again.

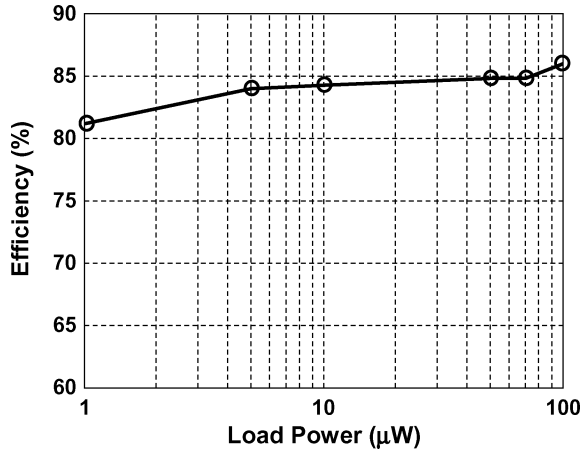

Fig. 11. Measured efficiency plot of the DC-DC converter while delivering a load voltage of 0.5 V.

The measured efficiency plot of the DC–DC converter implemented as a part of the minimum energy tracking loop is shown in Fig. 11. The DC–DC converter achieves an efficiency of around 86% at a load power level of 100 μW when delivering 0.5 V as the load voltage. At this power level, the efficiency is primarily determined by the conduction and switching loss mechanisms. As the load power levels decrease, the leakage power loss in the control circuitry begins to gain dominance and this brings the efficiency further down. The DC–DC converter is still able to achieve greater than 80% efficiency at 1 μW load power, due to the extremely simple control circuit design that consumes very little overhead power.

Fig. 12(a) shows the measured energy savings that can be obtained by tracking the minimum energy operating voltage with change in workload. It can be seen that as the workload of the 7-tap FIR filter is reduced by reducing the number of taps, the minimum energy operating voltage increases. The doughnut-shaped marks on the curves show the voltage at which the minimum energy tracking loop settles in measurements. Consider the curves corresponding to the 7-tap and 7-tap$+1$ μA scenarios. The 7-tap curve corresponds to all seven taps of the FIR filter being ON. A very small amount of leakage current of 1 μA is added in the scenario depicted by the 7-tap$+1$ μA curve. This scenario is common in cases where another circuit with very low activity is enabled. This additional leakage pushes the minimum energy operating voltage from 320 mV to around 420 mV. If the FIR filter continues to operate at the MEP corresponding to the 7-tap curve, it would consume 2.1X times more energy than the minimum necessary, or in other words, a 110% energy savings can be obtained by tracking the minimum energy operating voltage. Fig. 12(b) shows that the minimum energy operating voltage increases as ambient temperature is increased from 0 °C to 85 °C. This is as expected since an increase in temperature produces a relative increase in the leakage energy per operation. Here again, it can be seen from the curve corresponding to 85 °C, that a savings of 50% in energy can be achieved by tracking the minimum energy operating voltage. The energy savings obtained are highly circuit dependent and can be much larger in modern digital ICs which dissipate a significant portion of power in leakage.

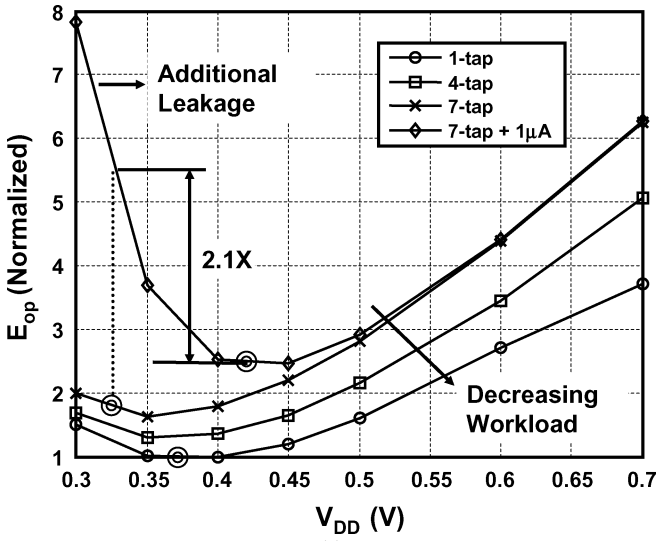

(a)

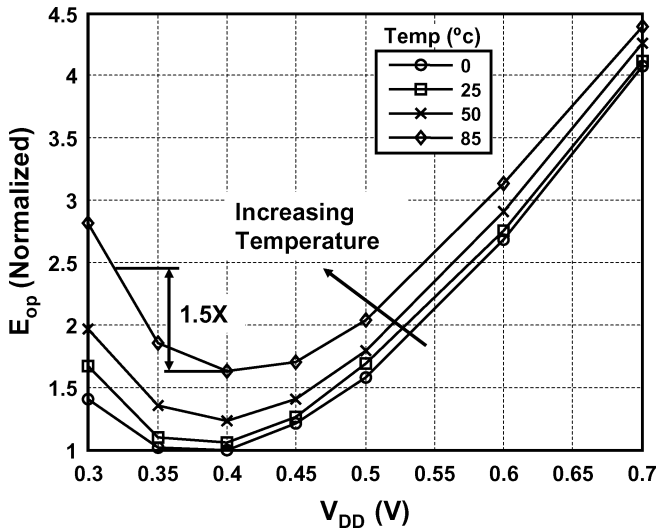

(b)

Fig. 12. (a) Measured $E_{\rm op}$ curves with change in workload for a 7-tap FIR filter. The curve with $\diamondsuit$ marks has an intentional 1 μA leakage current added to the maximum workload scenario (curve with x marks). The doughnut marks denote the measured voltage at which the minimum energy loop settles. (b) Measured $E_{\rm op}$ curves with change in ambient temperature.

The proposed minimum energy tracking loop is *nonintrusive*, thereby allowing the load circuit to operate without being shut down. At the same time, it calculates the energy per operation of the actual circuit and not of any replica. This eliminates the problems of designing a replica circuit that can track the energy behavior of a load circuit over varying operating conditions. The tracking methodology is independent of the size and type of digital circuit being driven and the topology of the DC–DC converter. The energy overhead associated with obtaining the MEP is equivalent to the energy consumed by 50 operations of the FIR filter at its MEP in the minimum workload scenario (1 tap). This combined with the small area overhead of 0.05 mm^2 makes it feasible to have multiple such minimum energy tracking loops for each distinct voltage domain in a complex digital system.

VII. Conclusion

In this paper we have presented a control loop to track the minimum energy operating voltage of arbitrary digital circuits. The tracking loop makes use of on-chip energy sensor circuitry and a slope-tracking algorithm to arrive at the minimum energy operating voltage. The nonintrusive tracking methodology presented is independent of the type of digital circuit being tracked. An embedded DC–DC converter which enables this minimum energy operation was designed to deliver load voltages between 0.25 V to 0.7 V. The DC–DC converter uses extremely simple all digital PFM mode control to deliver ultra-low load power levels. Measured energy savings of the order of 50%–100% were demonstrated on tracking the MEP as it varied with workload and temperature. The DC–DC converter was able to achieve an efficiency $>80\%$ while delivering load powers of the order of 1 μW and higher from a 1.2 V supply. The energy overhead associated with obtaining the MEP is equivalent to the energy consumed by 50 operations of a 1-tap FIR filter at its MEP. This combined with the small area overhead of 0.05 mm^2 makes it feasible to use multiple such tracking loops to monitor distinct voltage domains in a complex digital system.

References

[1] A. Sinha and A. Chandrakasan, "Dynamic power management in wireless sensor networks," *IEEE Des. Test Comput.*, vol. 18, no. 2, pp. 62–74, Mar. 2001.

[2] B. Zhai, L. Nazhandali, J. Olson, A. Reeves, M. Minuth, R. Helfand, S. Pant, D. Blaauw, and T. Austin, "A 2.6 pJ/Inst subthreshold sensor processor for optimal energy efficiency," in *Symp. VLSI Circuits Tech. Dig.*, 2006, pp. 192–193.

[3] O. Soykan, "Power sources for implantable medical devices," *Business Briefing: Medical Device Manufacturing and Technology 2002*, pp. 76–79, Jun. 2002.

[4] B. H. Calhoun, D. C. Daly, N. Verma, D. Finchelstein, D. Wentzloff, A. Wang, S.-H. Cho, and A. P. Chandrakasan, "Design considerations for ultra-low energy wireless microsensor nodes," *IEEE Trans. Comput.*, vol. 54, no. 6, pp. 727–740, Jun. 2005.

[5] 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.

[6] C. Kim, H. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Dec. 2003.

[7] Y. K. Ramadass and A. P. Chandrakasan, "Minimum energy tracking loop with embedded DC-DC converter delivering voltages down to 250 mV in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2007, pp. 64–65.

[8] B. H. Calhoun and A. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. ISLPED*, 2004, pp. 90–95.

[9] G. Bonfini *et al.*, "An ultralow-power switched opamp-based 10-B integrated ADC for implantable biomedical applications," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 1, pp. 174–177, Jan. 2004.

[10] J. Xiao, A. Peterchev, J. Zhang, and S. Sanders, "A 4 μA-quiescent-current dual-mode buck converter IC for cellular phone applications," in *IEEE ISSCC Dig. Tech. Papers*, 2004, pp. 280–281.

[11] A. J. Stratakos, "High-efficiency low-voltage DC-DC conversion for portable applications," Ph.D. dissertation, Univ. of California, Berkeley, 1998.

[12] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.

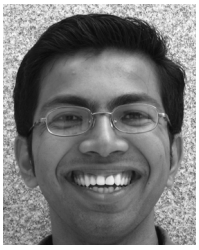

Yogesh K. Ramadass received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, in 2004, and the S.M. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, in 2006, where he is currently pursuing the doctoral degree.

From May 2007 to August 2007, he worked in the Wireless Analog Technology Center at Texas Instruments, Dallas, TX, designing power converters. His research interests include low-power circuit design, DC–DC converters, and energy-harvesting/processing circuits.

Mr. Ramadass was awarded the President of India Gold Medal in 2004, the Beatrice Winner Award for Editorial Excellence at ISSCC 2007, and the 7th International Low Power Design Contest Award at ISLPED 2007.

Anantha P. Chandrakasan (M'95–SM'01–F'04) received the B.S, M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. He is the Director of the MIT Microsystems Technology Laboratories. He is a co-author of *Low Power Digital CMOS Design* (Kluwer Academic Publishers, 1995), *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition), and *Sub-threshold Design for Ultra-Low Power Systems* (Springer 2006). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005). His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies.

Dr. Chandrakasan has been a co-recipient of several awards, including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 DAC Design Contest Award, the 2004 DAC/ISSCC Student Design Contest Award, and the ISSCC 2007 Beatrice Winner Award for Editorial Excellence. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Sub-committee Chair for ISSCC 2004–2007. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on the SSCS AdCom from 2000 to 2007 and he was the meetings committee chair from 2004 to 2007. He is the Technology Directions Chair for ISSCC 2008.