A Batteryless 19 μW MICS/ISM-Band Energy Harvesting Body Sensor Node SoC for ExG Applications

Yanqing Zhang, Fan Zhang, Yousef Shakhsheer, Jason D. Silver, Alicia Klinefelter, Manohar Nagaraju, *Student Member, IEEE*, James Boley, Jagdish Pandey, *Student Member, IEEE*, Aatmesh Shrivastava, Eric J. Carlson, Austin Wood, Benton H. Calhoun, *Senior Member, IEEE*, and Brian P. Otis, *Senior Member, IEEE*

***Abstract*—This paper presents an ultra-low power batteryless energy harvesting body sensor node (BSN) SoC fabricated in a commercial 130 nm CMOS technology capable of acquiring, processing, and transmitting electrocardiogram (ECG), electromyogram (EMG), and electroencephalogram (EEG) data. This SoC utilizes recent advances in energy harvesting, dynamic power management, low voltage boost circuits, bio-signal front-ends, subthreshold processing, and RF transmitter circuit topologies. The SoC is designed so the integration and interaction of circuit blocks accomplish an integrated, flexible, and reconfigurable wireless BSN SoC capable of autonomous power management and operation from harvested power, thus prolonging the node lifetime indefinitely. The chip performs ECG heart rate extraction and atrial fibrillation detection while only consuming 19 μW, running solely on harvested energy. This chip is the first wireless BSN powered solely from a thermoelectric harvester and/or RF power and has lower power, lower minimum supply voltage (30 mV), and more complete system integration than previously reported wireless BSN SoCs.**

***Index Terms*—Body-area sensors, energy harvesting, subthreshold, ultra-low power.**

I. Introduction

BODY SENSOR NODES (BSNs) promise to provide significant benefits to the healthcare domain by enabling continuous monitoring, actuation, and logging of patient bio-signal data, which can help medical personnel to diagnose, prevent, and respond to various illnesses such as diabetes, asthma, and heart attacks [1]. Though they show great potential, BSNs have many design challenges that impede their widespread adoption including node operating lifetime, small form factor for wearability, and affordable cost. One of the most critical issues is node lifetime. In many applications, such as long-term monitoring of chronic illnesses, limited battery lifetimes severely undermine the deployment of BSNs, since the required node operating lifetime is effectively indefinite. Supplying the node with sufficient energy over a long lifetime poses a challenge. A large battery increases the form factor of the node, making the node unwearable or uncomfortable, while a small battery requires frequent changing and reduces wearer compliance. Energy harvesting from ambient energy sources, such as thermal gradients or mechanical vibrations, potentially provides indefinite lifetime. Examples such as [2], [3] have shown that commercial thermopiles can be applied to BSNs. To eliminate battery changing, nodes can operate solely from energy harvesting instead of using a battery, although this introduces new challenges. The full system must consume less energy than the amount harvested, high power components such as the transmitter must be heavily duty-cycled, and the node must cope with time varying harvested energy profiles [4].

To ensure sustained operation of the node using harvested energy, on-node processing to reduce the amount of data transmitted, power management, and ultra-low power (ULP) circuits are key. Since commercial-off-the-shelf (COTS) based BSNs (e.g. Zigbee-type radios) can consume 10's of mA during operation and consume additional board area, it is infeasible to use COTS components to build a batteryless system. Instead, we pursue an integrated ULP SoC approach. Recent advances in ultra-low power chip design techniques, with many targeting wireless sensor networks, have enabled a push toward long lifetime BSN devices performing complex applications. For example, [5] presents an ECG acquisition and processing SoC with a 3-channel analog front-end (AFE) and flexible, generic DSP components. Clock-gating and duty-cycling are used to reduce power, but without voltage reduction techniques (V_{DD} = 1.2 V), the minimum power is 31.1 μW. The SoC does not include a transmitter or voltage regulators. An SoC for EEG seizure detection integrates a COTS radio and does not include voltage regulation or power management [6]. The system in [7] integrates several chips with solar cells and a battery to accomplish near-perpetual operation for measuring intraocular pressure. The system consumes 3.3 fW/bit at 400 mV when taking one measurement every hour. These works show that energy harvesting could provide a viable power supply for ULP BSN circuits. However, integration of a complete wireless, flexible, easily deployable BSN node on an SoC that supports closed-loop power management and energy harvesting has yet to be demonstrated.

Manuscript received May 01, 2012; revised June 11, 2012; accepted August 28, 2012. Date of current version December 31, 2012. This paper was approved by Guest Editor Uming Ko. This work was supported in part by the National Science Foundation (NSF) under #1035771.

Y. Zhang, Y. Shakhsheer, A. Klinefelter, J. Boley, A. Shrivastava, and B. Calhoun are with the University of Virginia, Charlottesville, VA 22904 USA (e-mail: yanqing@virginia.edu).

F. Zhang, J. Silver, M. Nagaraju, J. Pandey, E. Carlson, A. Wood, and B. Otis are with the University of Washington, Seattle, WA 98195 USA.

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

Digital Object Identifier 10.1109/JSSC.2012.2221217

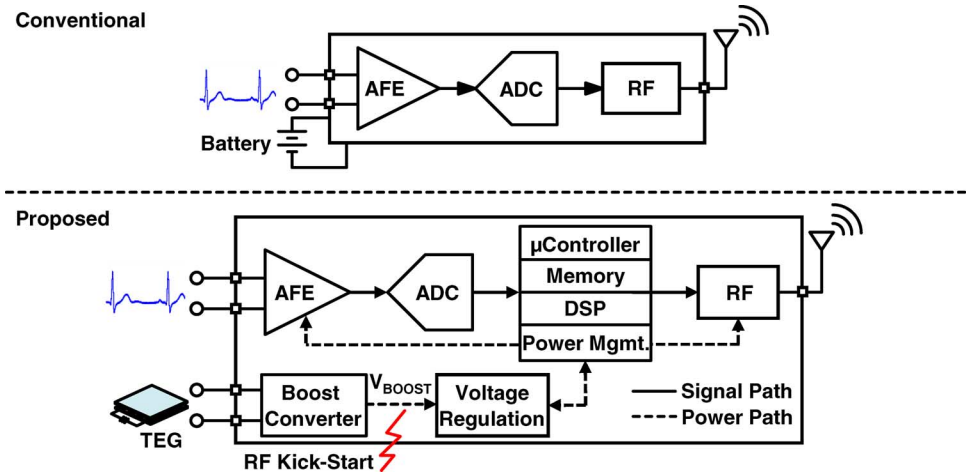

Fig. 1. High-level diagram of conventional BSN solution (top) and the proposed (bottom) solution with energy harvesting, integrated power management, and ultra-low power flexible DSP architecture [8].

We utilize recent advances in energy harvesting, low voltage boost circuits, dynamic power management, subthreshold processing, bio-signal front-ends, and low power RF transmitters to realize an integrated reconfigurable wireless BSN SoC for ECG, EMG, and EEG applications with autonomous power management for completely battery-free operation [8]. This SoC can run indefinitely from energy harvested from body heat while worn, and potentially decreases cost by having high integration and targeting a wide range of bio-electric sensing applications.

We introduce the system architecture in Section II, focusing on the integration and features of circuit blocks that enable batteryless operation. Section III explains the energy harvesting mechanism and related circuits. Section IV presents a closed-loop power management scheme that maximizes node lifetime. Section V details the circuits that comprise a flexible, re-programmable data signal path for ExG (ECG, EMG, and EEG). Finally, Section VI reports measured chip results showing full system operation. Section VII concludes our paper.

II. System Overview

Conventional wireless sensors use batteries (Fig. 1), limiting node lifetime and reducing user compliance due to the requirement for charging or replacing batteries. These sensors often require high data transmission rates, which quickly drain battery energy. More sophisticated approaches that include on-node processing and duty-cycling of the power-hungry radio are available [9], but the lifetime of such COTS nodes is usually limited to a few days, and custom BSN chips have not yet integrated radios, processing, and power management with energy harvesting. In contrast, we propose a wireless BSN chip powered by energy harvested from human body heat using a thermoelectric generator (TEG). This, in conjunction with ULP circuits, intelligent duty cycling of power-hungry blocks (e.g. the transmitter), and a programmable power management system allows for indefinite operation of the chip. To demonstrate, we present a chip targeting ExG applications.

To achieve flexible data acquisition and processing while operating the node solely from harvested energy, we propose a system architecture, illustrated in Fig. 2, which comprises four subsystems. First, the energy harvesting/supply regulation section boosts a harvested supply input as low as 30 mV up to a regulated 1.35 V using an off chip storage capacitor. It provides five regulated voltage supplies to the rest of the chip, and generates a bandgap reference. Second, the four-channel AFE subsystem provides bio-signal acquisition with programmable gain and sampling rate, amplifying bio-signals as low as a few μV's while consuming $< 4\mu\mathrm{W/channel}$. A variable gain amplifier (VGA) maximizes the signal at the input to an 8-bit successive-approximation (SAR) analog to digital converter (ADC), reducing the ADC resolution requirement. Third, the acquired data is sent to a subthreshold digital processing subsystem that also performs mode control and power management (including power/clock-gating of blocks and dynamic voltage scaling (DVS)) based on the available energy on the storage capacitor. The digital section includes a custom digital power management (DPM) processor, general purpose microprocessor (MCU), programmable FIR, 1.5 kB instruction SRAM/ROM, 4 kB data memory FIFO, and dedicated accelerators for ECG heart rate (R-R) extraction, atrial fibrillation (AFib) detection, and EEG band energy calculation. The DPM is responsible for power management, node control, data flow management, and overseeing all processing on-node. Finally, a sub-mW 400/433 MHz MICS/ISM band frequency-multiplying transmitter (TX) performs BFSK transmission up to 200 kbps. The TX has low instantaneous power consumption to avoid the need of large filtering capacitors on the supplies and is intelligently duty-cycled to achieve low average power consumption.

III. Energy Harvesting/Supply Regulation Subsystem

The energy harvesting subsystem is designed to harvest energy from RF, thermoelectric, or solar power sources and to provide regulated voltages to the rest of the chip.

A. Harvesting Energy From a Thermoelectric Generator (TEG)

Thermoelectric generators (TEGs) intended to be used in environments with low temperature differentials (i.e., the human body in a room temperature setting) are typically constructed of

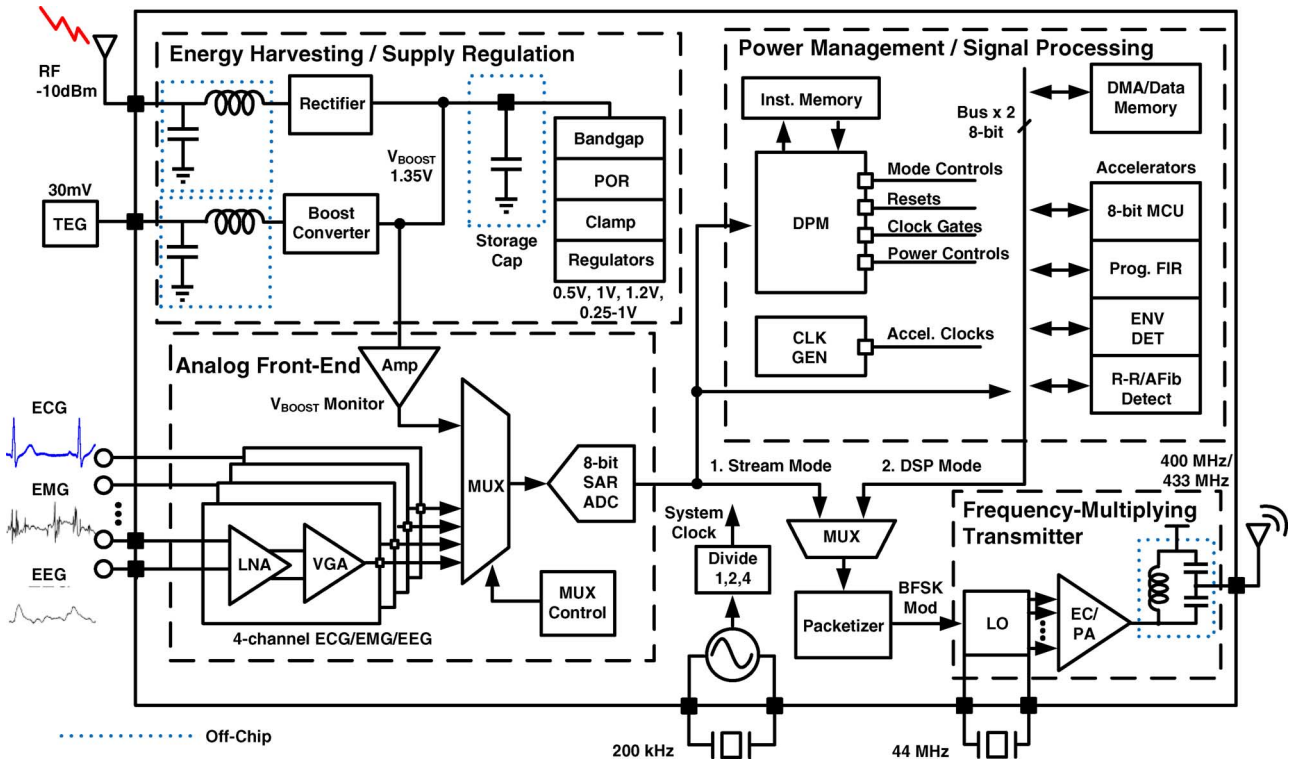

Fig. 2. System block diagram for the proposed chip comprising the energy harvesting/supply regulation, analog front-end (AFE), subthreshold digital signal processing, and transmitter subsystems [8].

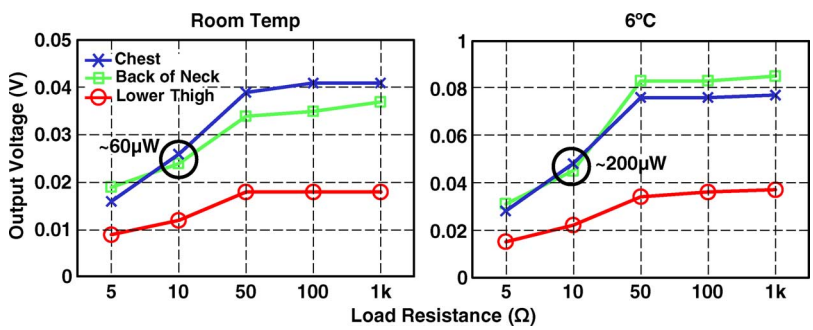

Fig. 3. Measured results from an on-body thermal-harvesting experiment with 4×4 cm^2 COTS TEG.

a large number of thermopiles in parallel, and few in series (to limit their output resistance). This arrangement places a bound on the maximum voltage achievable from a given temperature difference for a given size. The Seebeck coefficient of a conventional thermocouple (bismuth telluride) is ± 0.2 mV/°C. Assuming a temperature gradient of 1°C, a 1×1 cm^2 TEG will generate much less than 1 V. In addition, the surrounding air presents a large thermal resistance that dramatically reduces the effective temperature gradient across the thermopiles, further limiting the voltage available at the TEG output [3]. To quantify how much power can be harvested, we placed a COTS TEG of 4×4 cm^2 [10] on different parts of the human body. Fig. 3 shows that this TEG can harvest approximately 60 μW at room temperature and 200 μW at 6°C. However, the voltage available at the TEG output is only tens of mV, requiring a high conversion ratio boost converter to generate a usable supply voltage.

B. Boost Converter and RF Kick-Start

This work employs the boost converter architecture proposed in [11] to utilize the low voltages available from a body-worn TEG. Fig. 4 shows the diagram of the energy harvesting subsystem. Due to its increased efficiency at high conversion ratios and minimum usable input voltage (30 mV in this work), the converter is well suited for harvesting the TEG energy input. By modulating the two power switches in the boost converter, energy is first stored in the inductor and then transferred onto the storage capacitor. The precise timing of turning on/off the switches is critical for high conversion efficiency. The switching frequency is adjusted in a feedback loop to synchronize the PMOS turn-off with the inductor current zero-crossing, preventing current flowing backward through the PMOS switch. The measured efficiency is 38% on our chip when converting

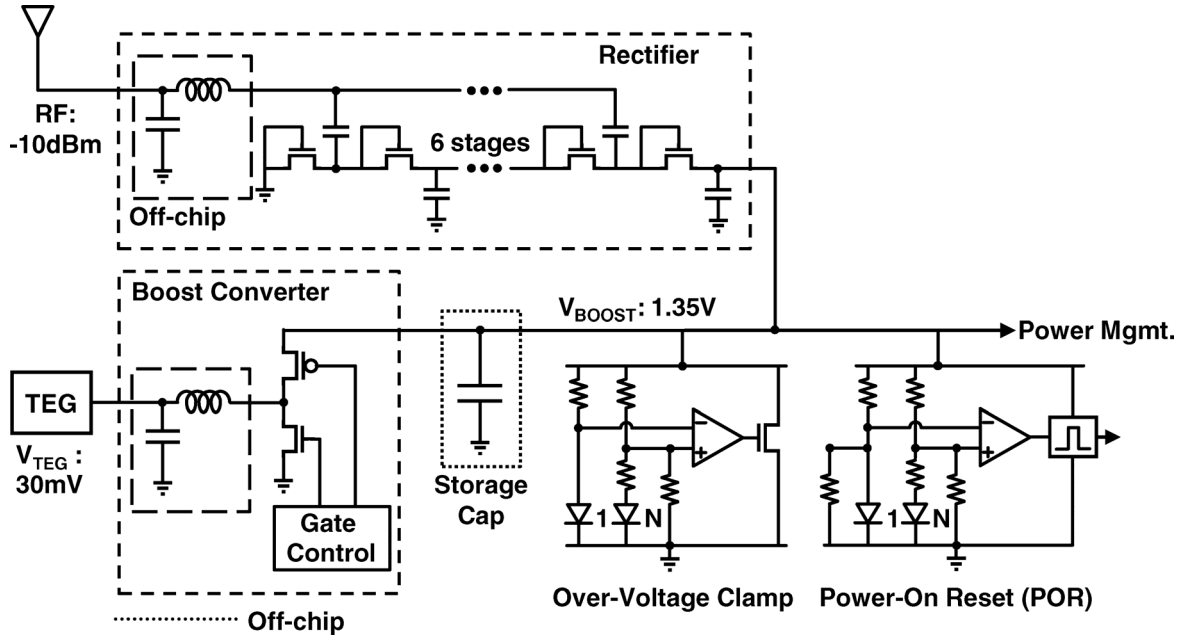

Fig. 4. Schematic of the energy harvesting section.

from 30 mV to 1.35 V. For low power levels at high supply voltages, the dynamic power dissipation caused by the switches and control logic in the designed converter becomes a significant fraction of the overall power consumption, limiting the ability to achieve high efficiency.

While the boost converter supports a 30 mV power source, the internal oscillator and control logic need 600 mV for startup. This requires a one-time pre-charge of the V_{BOOST} node. Previously reported start-up mechanisms use batteries and mechanical switches [11], [12], requiring bulky off-chip components. Instead, we use wireless RF power for the kick-start. Incident RF power around -10 dBm is provided wirelessly for 1–2 seconds and rectified through an RF rectifier front-end consisting of a 6-stage charge pump (Fig. 4). We chose 6 stages to optimize for sensitivity instead of efficiency. In the case of kick-starting the boost converter, we charge a (storage) capacitor instead of powering a relatively constant voltage supply. Fewer stages exhibit higher efficiency, shorten the charging time, but require a larger RF input. As an alternative, self-synchronous rectifiers could be more efficient if the forward conduction loss is less than threshold voltage of the diode-connected devices [13]. However, they require fully-differential RF inputs. A shunt regulator clamps V_{BOOST} to 1.35 V to prevent over-voltage during the RF powering. A low-voltage bandgap-based power-on-reset (POR) resets the chip after V_{BOOST} reaches 1 V. Hysteresis in the POR trigger levels allows POR resetting only when V_{BOOST} drops below a critical voltage where the chip fails to function correctly. V_{KILL} is determined by the minimum V_{BOOST} voltage required to generate correct reference voltages and sustain conversion. Fig. 5 shows a measurement of the RF kick-start. After the voltage at the TEG output settles, a short RF burst wirelessly charges the storage capacitor at V_{BOOST}. Shortly after the voltage reaches 600 mV, the boost converter turns on and charges the capacitor to a regulated voltage. The node can then

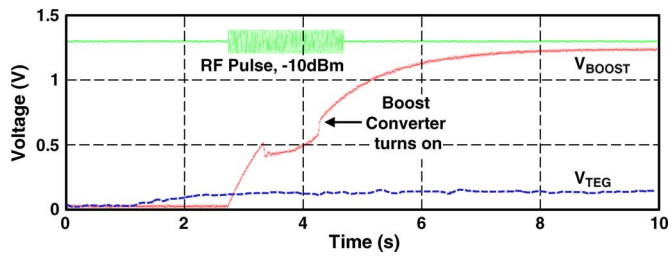

Fig. 5. Measured startup sequence for the chip. An RF pulse kick-starts V_{BOOST}, allowing the boost converter to turn on. The chip is able to sustain a usable voltage from harvested energy (V_{TEG}) from then on [8].

continue to run indefinitely from the TEG, unless V_{BOOST} drops below V_{KILL} due to a prolonged period of consumption exceeding the harvested energy.

C. Chip Resuscitation

In the case that V_{BOOST} decreases below V_{KILL}, the chip will blackout and will automatically shut off. The chip can be 'revived' with the same startup sequence of RF powering and harvesting from the TEG. A 90 instruction (132 byte total) subthreshold ROM can re-boot the chip and execute a default AFib detection algorithm. We chose to store an AFib detection algorithm as an indicative example of our chip performing long term illness monitoring, though objectively it is possible to miss the rare event when the chip is in shut-off mode. The chip is able to recognize its program state through a latch structure that stores this information.

D. Supply Regulation

The low voltage TEG input is boosted and subsequently regulated. Fig. 6 details the supply regulation circuits. All biases are generated on-chip. On-chip supply regulation is provided by four sub-μW linear regulators: 1.2 V (AFE), 0.5 V (DSP), 1.0 V (TX local oscillator (LO)), and 0.5 V (TX power amplifier

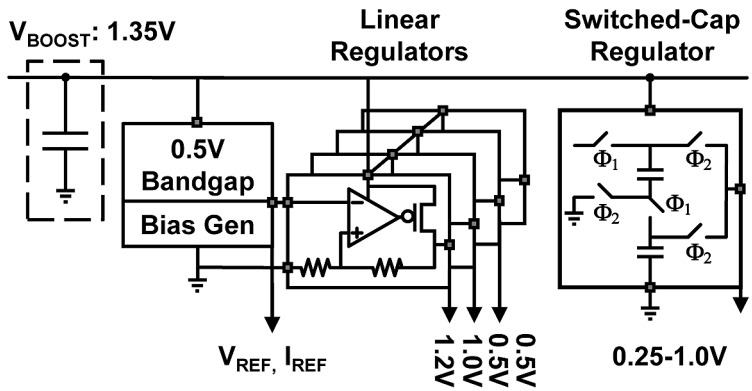

Voltage Domain	Blocks Powered
1.2V	Pads, AFE
1.0V	TX LO
0.5V	TX PA
0.5V	DPM, MEM, Accels
SC Reg. (0.25-1.0V)	Accels for DVS

Fig. 6. Top level diagram of the supply regulation subsystem.

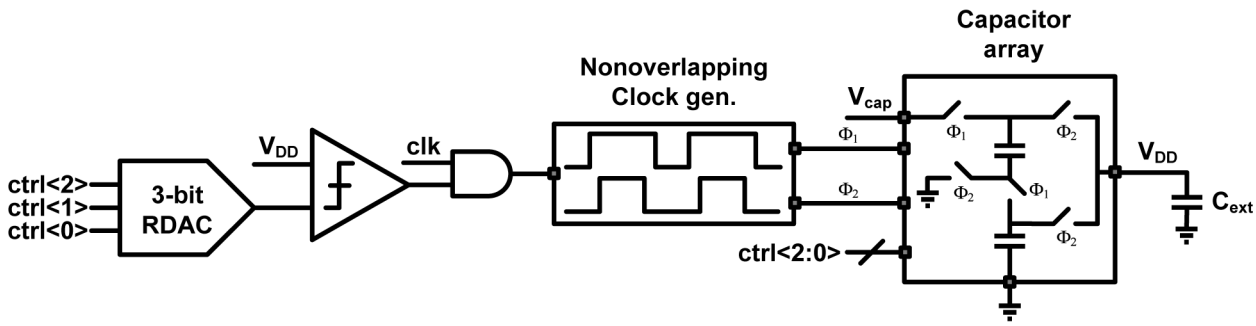

Fig. 7. Diagram of the switched-capacitor variable DC-DC converter, which is similar to [14].

(PA)). A programmable switched-capacitor DC-DC converter provides an output from 0.25 V to 1 V in 50 mV steps (Fig. 7). A 3-bit resistor DAC (RDAC) generates a reference for the desired output level based on a control word from the DPM. The arrangement of the capacitors in the array varies according to the desired output range, in a manner similar to [14]. An external capacitor $(\mathrm{C_{ext}})$ ensures that the voltage ripple due to the switching operation is minimized. The subthreshold DSP accelerators can either be connected to the 0.5 V supply or variable supply (switched-capacitor DC-DC converter) through PMOS headers, enabling DVS for additional power savings.

IV. Closed-Loop Power Management

To prevent node blackout, the SoC must be aware of the available energy on the storage capacitor and adjust its mode of operation accordingly. If $\mathrm{V_{BOOST}}$ is decreasing, the SoC must switch modes and consume less power and energy. When harvested energy is abundant again, the chip should recover itself to a mode of full operation. The always-on DPM (digital power manager) is a custom-ISA (Instruction Set Architecture) chip controller that implements the closed-loop power management scheme. The DPM is also responsible for node control, data flow management, and overseeing all processing on the node, as described in Section V.

The DPM monitors $\mathrm{V_{BOOST}}$ through the ADC, where it is sampled after first being scaled to the ADC's full input range (Fig. 8). To do so, $\mathrm{V_{BOOST}}$ is halved by a resistive divider, buffered to reduce the output impedance, and compared with a reference voltage. The difference is amplified (gain of 4) through a difference amplifier. Both amplifiers in Fig. 8 are implemented as two-stage differential-input, single-ended-output OTAs. The DPM issues an instruction that selects an input channel of a 5-input mux (the other 4 channels belong to the 4 channel AFE as explained in Section V) that enables the ADC to digitize the scaled $\mathrm{V_{BOOST}}(\mathrm{V_{CAP\_DIG}}[7:0])$ and send it to the DPM.

TABLE I
Summary of DPM 'Stoplight' Policy

	Inst. MEM	AFE	Data MEM	Accel. Blocks	Transmit
Red	On	Off	Off	Off	Off
Yellow	On	On	On	On	Duty Cycle
Green	On	On	On	On	On

Based on the $\mathrm{V_{CAP\_DIG}}$ value, the DPM selects the node operating mode in a 'stoplight' fashion based on programmable digital threshold values. The DPM compares $\mathrm{V_{CAP\_DIG}}$ to two 8-bit threshold values (green threshold, yellow threshold) to set the DPM operating mode (green, yellow, or red). The DPM is capable of jumping from the current mode to any mode, regardless of its previous state (i.e. green to red mode). Each operating mode limits the subset of blocks that can be powered-on (Table I), capping the maximum power consumption for the mode. Normal operation is green mode (e.g. $\mathrm{V_{BOOST}} > 1.3$ V), which allows all blocks to be on. In yellow mode (e.g. $1.1\ \mathrm{V} < \mathrm{V_{BOOST}} < 1.3$ V), the DPM duty cycles the transmitter based on available energy. In red (e.g. $\mathrm{V_{BOOST}} < 1.1$ V),

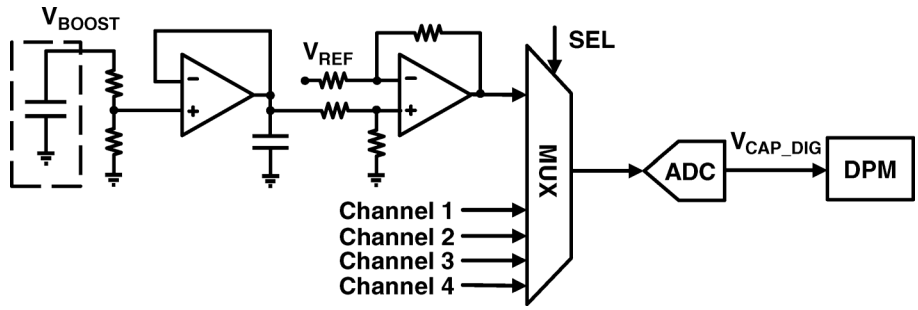

Fig. 8. Diagram of the V_{BOOST} monitoring path.

```
while (1) {
 run_EMG(raw_data); //run EMG activity algorithm

 if (EVENT_DETECTED) { //rare event detected?
  check_Vcap; //yes. Transmit immediately.
  if (green)   //but check to see we can TX first
   do_green(); //TX sequence
  else
   do_yellow(); //change to yellow mode
 }
 else
  ; //no, keep running algorithm
}

do_green() { //so long we are still in green, keep TXing
 transmit(stored_data); //TX a packet of data
 check_Vcap; //make sure we can continue
 if (green)
    do_green(); //keep TXing
 else
    do_yellow(); //change to yellow mode
 return;
}

do_yellow() {
 sleep_mode(); //turn everything not needed off
 sample_only(raw_data); //all we do is sample
 check_Vcap; //check to see if we have recuperated
 if (green){//we are green again! Do the recuperate sequence
    setDVS("high_speed"); //set DVS to highest speed
    run_EMG(stored_data); //quickly run algorithm on stored data
    setDVS("normal"); //ok done! resume normal operation
    return; //back to normal! this program will start again!
 }
 else
    do_yellow(); //stay in yellow mode
}
```

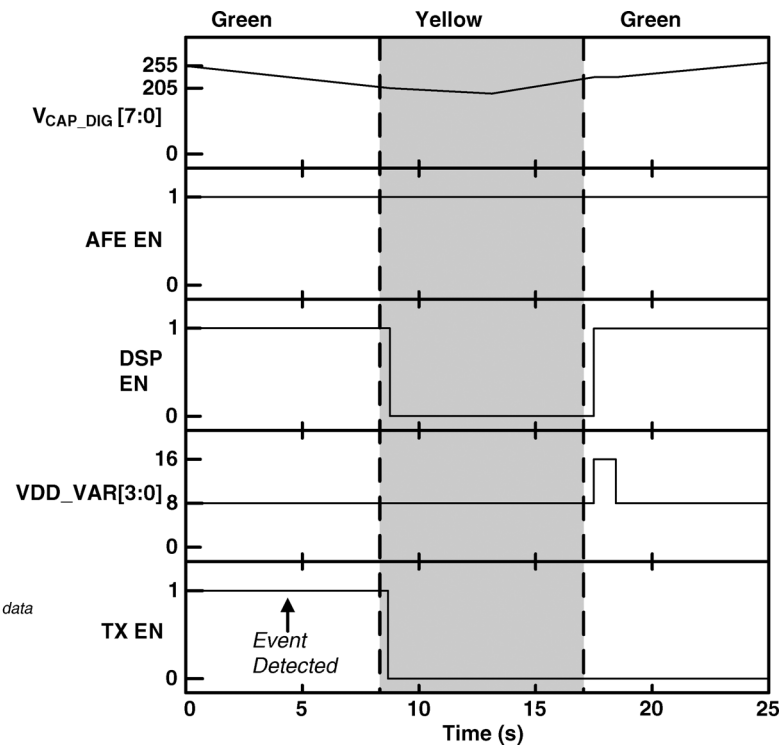

Fig. 9. An example that makes use of the 'stoplight' scheme's implementation flexibility. Here, the green threshold is set to $V_{CAP\_DIG} = 205$. A subroutine for yellow mode ensures timely processing of stored data when the chip recuperates back to green mode.

the transmitter, accelerator blocks, and the AFE are clock- and power-gated. Even if the node program requires one of those blocks to be turned on and functional, the 'stoplight' power management will override this request. The DPM transitions immediately from mode to mode when the 8-bit digital V_{CAP_DIG} value updates without requiring an additional instruction. By sampling V_{BOOST} and reacting immediately, the DPM closes the loop for a sustainable power management scheme.

This power management scheme provides a flexible platform. The programmer can set the mode threshold values and the timing of instructions to sample V_{BOOST} for possible mode transition. Thus, the programmer can bypass power management restrictions if transmitting or processing the data is deemed more important than keeping the node alive (e.g. the node has detected a rare EMG muscle movement event and must transmit the real-time data immediately using the power-hungry wireless transmitter, regardless of the capacitor energy). When the 'stoplight' scheme is used in parallel with DPM branch instructions, the programmer can also control how the node reacts during a mode change. Fig. 9 displays an example scenario that uses the DPM 'stoplight' flexibility by sending EMG data when an event is detected during green mode until the transition to yellow (at $V_{BOOST} = 0.96$ V in this example). In yellow, a subroutine samples and stores the raw input data and shuts off other blocks. When V_{BOOST} rises above the green threshold, another subroutine post-processes the buffered raw data at a higher frequency using DVS before recuperating the chip back to green mode.

Fig. 10 shows measured waveforms of the closed-loop DPM 'stoplight' scheme cycling through the three modes as the boost converter input is swept from 250 mV to 20 mV and back.

V. Flexible Bio-Signal Datapath

A. DPM as a Flexible Signal Path Controller

In addition to the DPM's power management responsibility, the DPM manages the data signal path. The DPM executes instructions from a 1.5 kB instruction memory and provides a

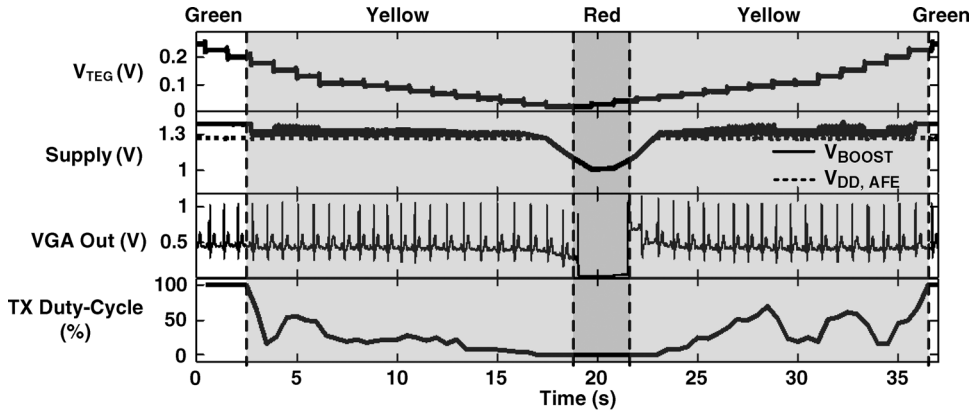

Modes	AFE	DSP Accels	Data Mem.	Inst. Mem.	Transmit	V_{THR}
						1.35V
Green	✓	✓	✓	✓	✓	
						1.3V
Yellow	✓	✓	✓	✓	×/✓	
						1.1V
Red	×	×	×	✓	×	
						0.9V

Fig. 10. Measured closed-loop power management 'stoplight' scheme. $\mathrm{V_{BOOST}}$ is sampled at a rate of 256 Hz. The node automatically updates its 'stoplight' status according to the $\mathrm{V_{BOOST}}$ value [8].

TABLE II
SUMMARY OF DPM ENERGY SAVINGS WHEN COMPARED TO OUR GENERAL PURPOSE PROCESSOR MCU IMPLEMENTATION

DPM Operation	DPM Energy	MCU Equivalent
NOP	0.7pJ	1.46pJ
Control Signals	2.8pJ	2.92pJ
Branch Commands	2.9pJ	4.38pJ

lower energy alternative to using generic MCUs for controlling the node (Table II).

Table III shows the DPM's custom instruction set architecture (ISA), which is designed to facilitate node management. The DPM controls the data memory, input channels of the AFE and ADC, sampling rate, transmission rate, clock frequency creation and distribution, bus management for flexible and timing-defined data flow, time delays, clock-gating and DVS voltage of the digital blocks, as summarized in Fig. 11.

B. Analog Front End (AFE)

The chip has four independently selectable bio-signal input channels, each with a fully-differential chopper-stabilized low-noise amplifier (LNA) and variable-gain amplifier (VGA) [15]. We chose 20 kHz as the chopper clock frequency to be significantly higher than the flicker-noise corner of the OTA for effective flicker-noise reduction. Input chopper switches are placed before the input capacitors. Compared to a topology where the switches come after the input capacitors, our arrangement reduces the amplification of any OTA offsets that might saturate its output. Any mismatch in the input capacitors results in common-mode to differential-mode gain. Since this amplifier is effectively AC-coupled, an off-chip capacitor and resistor are used to block any DC offset voltage at the electrode interface. One drawback of the topology is that the input impedance is relatively low. Periodic steady-state simulations reveal that the input impedance is a few $\mathrm{M\Omega s}$, sufficient for our ECG, EMG, and EEG applications. A programmable Gm-C filter reduces the switching ripple to below the noise floor. Along with the VGA, our amplifiers provide 7-step digitally-programmable gain (40–78 dB) from DC to 320 Hz at 3 $\mu\mathrm{W/channel}$. A 5-input mux allows the sub-μW 8-bit SAR ADC to sample any of the four channels as well as the $\mathrm{V_{BOOST}}$ node for monitoring stored energy.

C. Subthreshold Digital Signal Processing Subsystem

Fig. 12 shows the subthreshold DSP subsystem. To achieve ultra-low power consumption, we implement ASIC versions of the heart rate extractor (R-R), atrial fibrillation detector (AFib), and energy band extractor/envelope detector (ENV DET). An 8-bit RISC ISA general purpose processor (GPP) MCU executes generic computations, and a re-programmable FIR performs digital filtering. A digital packetizer streams serial data to the transmitter. Two memory arrays store the program (Instruction Memory, IMEM) and bio-signal data (Data Memory, DMEM). A DMA achieves easy FIFO control and low memory latency for the DMEM. Two 8-bit switch-box busses, controlled by the DPM, configure the connections of all the processing accelerators, MCU, DMA, and packetizer. Each input/output bus port has a 4-bit address. Having two busses eases data steering and simplifies the control instructions. To support the 'stoplight' scheme described in Section IV, each processing element has a clock-gate and two PMOS headers [16], one connected to 0.5 V, the other to the variable voltage

TABLE III
SUMMARY OF THE DPM ISA

Code	Description	Code	Description
NOP	No operation	SAVEPC	Stores current PC values
STALL	NOPs until new ADC sample available	RESTOREPC	Restores PC value from the last saved PC value
TIMER	NOPs for a set number of cycles	SETVAR	Sets voltage of variable voltage
EN	Enable/disables, sets voltage of, clock gates, and resets block specified in instruction.	BUS1/BUS2	Controls connections to the respective bus
DMA	Sets DMA control to read or write	SETCLK	Sets clock to accelerators and DMA
ADCCHAN	Selects ADC channel input	CJMP	Conditional jump based on amount of energy, amount of data stored, detection of AFib, or if the node has been programmed
CTRL	Sets flag for intended destination of instruction	JMP	Absolute jump to literal value

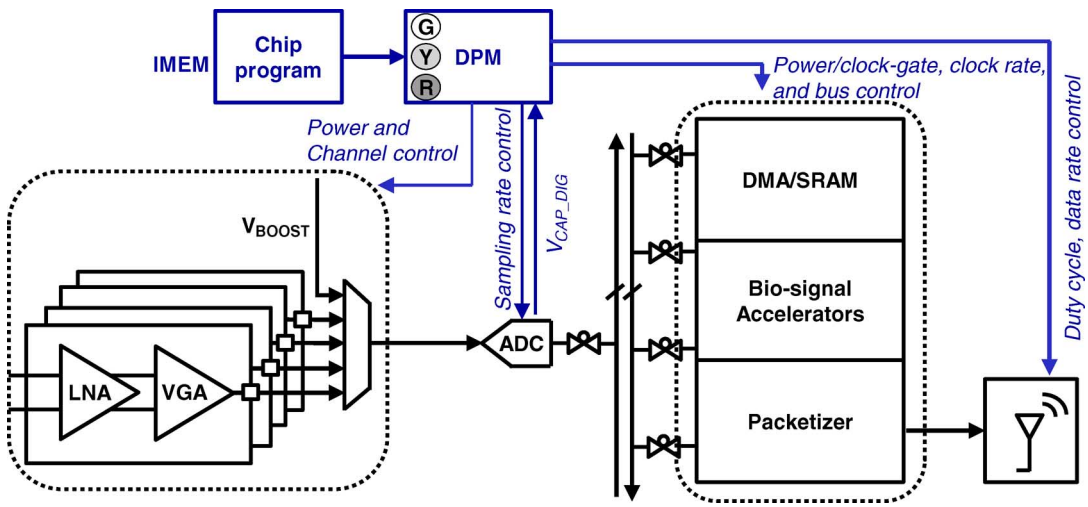

Fig. 11. Summary of DPM functions, which include power management ('stoplight' scheme) and node management across whole signal path.

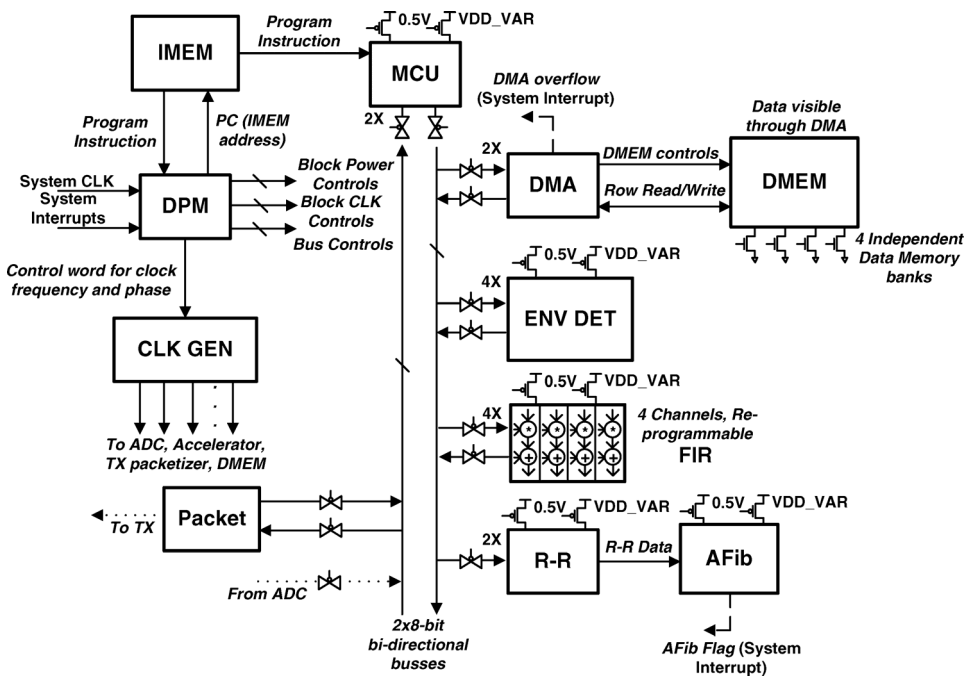

Fig. 12. Block diagram for the proposed subthresholddata processingsubsystem.

for DVS. The clock generator block (CLK GEN) distributes a programmable clock signal (frequency and phase) to each of the processing units. The chip can process data flexibly with the MCU, use the highly efficient hardware accelerators, or

cascade accelerators with MCU processing. It also can stream data on the transmitter, store and burst data, or do event based transmission. The main components work as follows:

1) *MCU*: The 8-bit GPP MCU is a subthreshold RISC based on the PIC series [17]. The MCU is designed to run arbitrary programs and functions down to 0.26 V, 1.2 kHz. Fig. 13 shows the energy-delay (E-D) curve for the MCU. The MCU consumes 0.7 nW to 1.4 μW measured power (0.26–0.55 V) and 1.5 pJ/op at the default 0.5 V, 200 kHz setting. The MCU shares the IMEM with the DPM. Fig. 14 summarizes the organization of memory, MCU, and DPM, and their capabilities. A multiplexer steers each instruction to either the MCU or DPM (INST_STEER) based on a special code word. When the MCU is executing instructions, the DPM automatically goes into a low power sleep mode. When the DPM is executing instructions, the MCU is either turned off or clock gated to save state. In this way, we retain the energy efficiency of the DPM as a chip controller and the generic flexibility of the MCU without requiring extra instruction memory space. The MCU's instructions are programmed at the same time as the DPM during the chip's pre-deployment.
2) *Instruction and Data Memories*: Because the chip operates in subthreshold in an N-strong technology, the SRAMs use an 8 T bitcell and the zero leakage read-buffer from [18]. To eliminate half-select instability during a write, both reads and writes apply to full rows of memory. The DMEM is split into four 1 kB banks that can be individually power gated by NMOS footers being overdriven to 1.2 V when active to ensure low levels of ground bounce. Measured results show reliable operation down to 0.3 V at 200 kHz with IMEM read energy of 12.1 pJ per read at 0.5 V and leakage energy per cycle at 200 kHz of 6.6 pJ.
3) *DMA*: The DMA is an efficient subthreshold accelerator to interface between the DMEM and the rest of the SoC. It is programmed by one instruction of the DPM and effectively treats the DMEM as a FIFO to support efficient streaming. A clock multiplexor synchronizes the DMA clock rate to the component it interfaces to by choosing between several clock rates. A memory controller uses separate DMEM banks during green and yellow modes for easier data management. To solve the half-select stability issue during writes, we use a row buffer and only write a row when all words are ready. When the difference between the write pointer address and read pointer address is greater than or equal to 4 bytes, a DMA_flag is raised, which signifies to the DPM that there is a full packet of data. This simple and efficient mechanism of interrupting for transmission limits overflows.
4) *Programmable FIR Filter*: A four-channel, programmable, max-30 tap, and synthesizable filter was designed to enable operation in the subthreshold regime down to 300 mV (measured). The FIR filter is capable of filtering all four-channels of the AFE. The programmable options include coefficient selection, number of taps, and number of filters. When power is critical and data fidelity can be compromised, a half-taps mode allows for a 15-tap filter. The direct-form implementation of an FIR requires as many adders and multipliers as there are taps, costing area and leakage. Due to the small sampling rate for ExG signals, each result can instead be computed serially over multiple faster clock cycles using only one multiplier and one adder. This architecture results in a 30x reduction in area per channel and a measured 1.1 pJ per tap at 350 mV. The architecture saves valuable chip area and reduces leakage current. For further power reduction, each individual channel can be clock-gated. A measured energy-delay curve is given in Fig. 13.

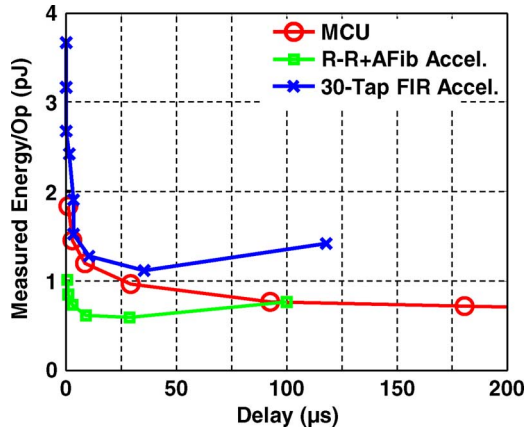

Fig. 13. Measured energy-delay curves for MCU, RR+AFib accelerator, and 30-Tap, 1-Channel FIR.

5) *Envelope Detector*: For EEG signals, knowing the signal power within a specific frequency band is useful for determining neural activity in the α, β, γ, and low-γ frequency bands [19]. The ENV DET circuit computes the average signal power within a specified frequency band. This block receives data directly from the FIR filter and has four input channels corresponding to the channel outputs of the filter. Equation (1) computes the average signal power p_x of input signal x,

$$p_x = \frac{1}{N}\sum_{n=0}^{N-1}|x[n]|^2 \tag{1}$$

where N is the summing window size. To reduce the computation complexity, N is set to powers of two $(2^2 - 2^7)$, which allows for division to be implemented as a simple right shifting of the data. Further, x values are rounded to the nearest power of four and the square results come from a lookup table. The rounding reduces the number of bits required during data transformation as the lower two bits are always 0. The ENV DET consumes 3.5 nW (measured) at 0.5 V and 200 kHz.
6) *R-R Extraction*: The heart-rate extractor accelerator is a simple version of the popular Pan-Tomkins algorithm [20]. This R-R algorithm calculates the heart rate by means of time windowing and thresholding, after an initial 4 second time frame where the R-R accelerator gains a baseline DC value for the heart waveform. The time stamp given to two consecutive peaks is the difference in the number of samples between them. For this reason, we can achieve a desired accuracy by changing the sampling rate and using

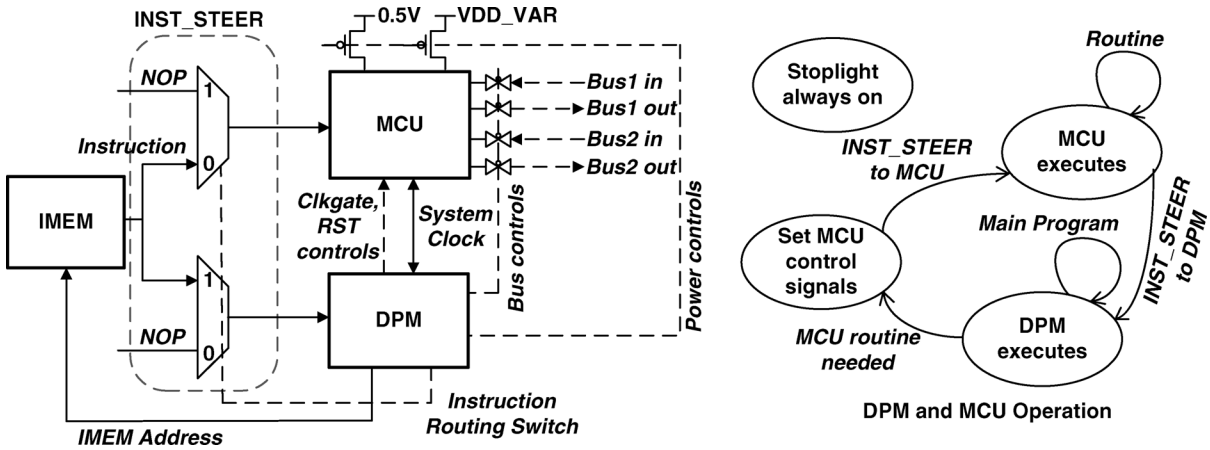

Fig. 14. Integration of DPM and GPP MCU. DPM and MCU share the same memory, and instructions are steered to the right processor. The MCU can be power-gated when idle. In this way, both efficient control of the node and generic processing are implemented.

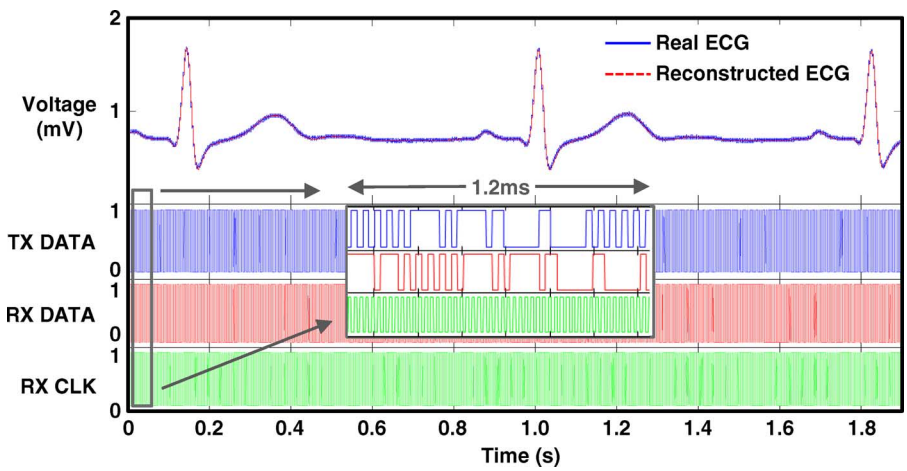

Fig. 15. Measured system experiment showing correct data acquisition and streaming from the transmitter. Total power of 397 μW prevents long use of this mode from harvested power [8].

the R-R accelerator in a DVS fashion to accommodate the faster or slower processing rate needed. Once an R-R time has been calculated, a pulse is output, which signifies to the AFib accelerator a new R-R sample is ready.

7) *AFib Accelerator*: The atrial fibrillation detector is an ASIC accelerator that detects the arrhythmia using an implementation of the clinically validated algorithm described in [21]. It receives its inputs from the R-R accelerator and outputs an $AFib\_flag$ signal to the DPM signifying the detection of atrial fibrillation. The algorithm uses only 12 R-R intervals [21]. Many variables in the algorithm, such as the margin of error, are programmable. The algorithm uses a pattern recognition scheme that quantifies the entropy in these 12 R-R intervals. If the entropy is more than a programmable threshold, then an AFib event is reported.

D. Low Power RF Transmitter

To allow operation from harvested power, the peak current consumption of the chip must be minimized. We utilize a frequency-multiplying transmitter architecture to reduce the synthesizer power by operating the LO at 1/9 the carrier frequency [22]. We use equally spaced edges generated from the cascaded ring oscillators to drive the edge-combiner (EC) embedded PA to perform the frequency multiplication.

The use of frequency multiplication allows harmonic injection-locking from the crystal oscillator. Instead of using a PLL, injection-locking a low-frequency ring oscillator to an on-chip crystal reference eliminates the longer settling times, therefore allowing aggressive duty-cycling of the transmitter to further save power. Directly injection-locking the multi-phase ring oscillator using a single-phase reference introduces significant mismatch. We used cascaded multi-phase injection-locking to correct the phase and amplitude mismatches.

On-chip BFSK modulation is accomplished by pulling the quartz reference clock. By modulating the load capacitor, we can pull the crystal frequency by 200 ppm. After 9x multiplication, the resulting frequency deviation is approximately 100 kHz, achieving $>$100 kbps datarate. The TX circuit consumes 160 μW when transmitting at its maximum data rate of 200 kbps [22]. In bio-signal raw data mode, the transmitter operates at a 100% duty-cycle, while the R-R extraction mode, explained in Section VI, reduces the duty cycle and average transmitter power consumption to 0.013% and 190 nW respectively. The packetizer contains a programmable packet header and CRC to allow compatibility with commercial receivers.

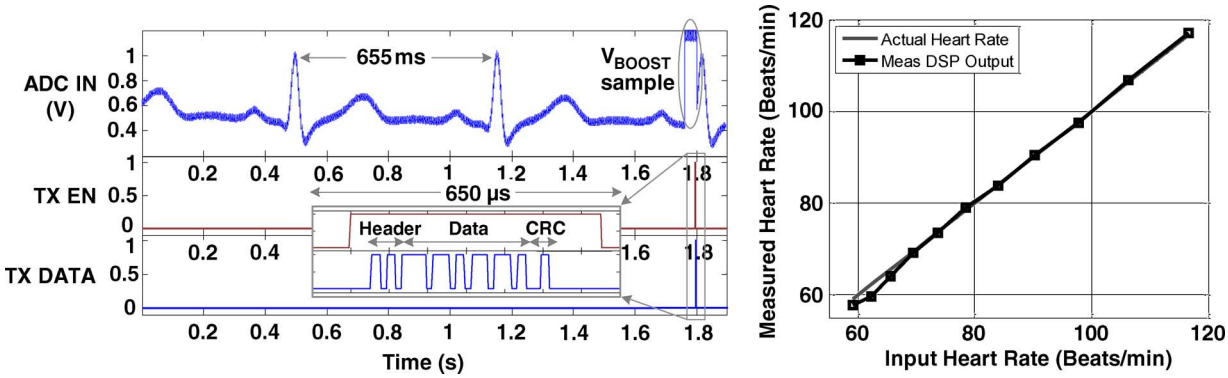

Fig. 16. Measured system results for an R-R extraction algorithm. Measured system results are for acquiring ECG, extracting R-R intervals, and sending RF updates for R-R every 5 s. Total power in this mode is 19 μW drawn from V_{BOOST}, powered exclusively by a 30 mV harvestor input. Measured accuracy results for R-R also shown [8].

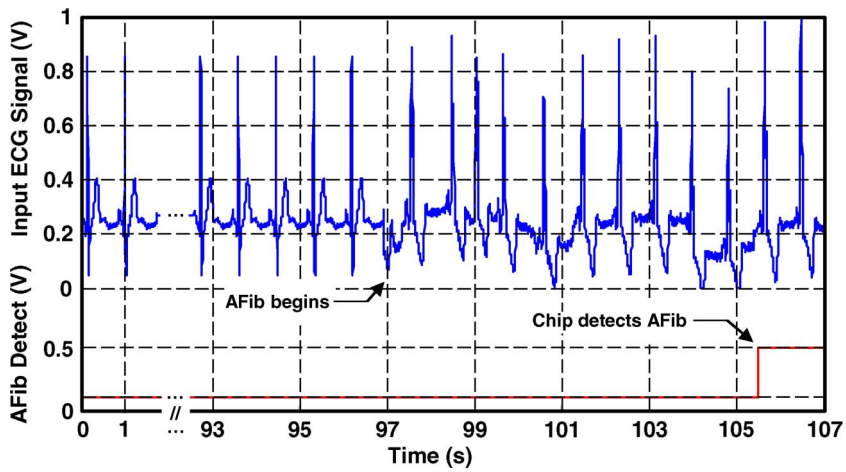

Fig. 17. Measured system AFib demo experiment using R-R extractor and AFib accelerator. Normal and atrial fibrillation heart waveforms from MIT-BIH database [23]. Last 8 beats of raw ECG are stored in DMEM and streamed over TX if AFib is detected. Total chip power in this mode is 19 μW from V_{BOOST}, powered exclusively by a 30mV harvestor input.

TABLE IV
MEASURED PERFORMANCE SUMMARY [8]

Energy Harvesting		**Supply Regulation**		**AFE (1 CH, ADC)**		**DSP**		**TX**	
V_{in}	30mV	**V_{unreg}**	>1.25V	**Current**	4µA	**Op. range**	0.3-1.2V	**Current**	280µA
Kick-Start	RF@ -10dBm	**$I_{quiescent}$**	3µA	**Supply**	1.2V	**MCU E/op @ 0.5V**	1.5pJ	**Supply**	1V (LO) 0.5V (PA)
V_{out}	1.35V	**$V_{analog,digital}$**	0.5V, 1V	**Gain**	40-78dB	**IMEM E/rd/Inst.**	1.0pJ	**Data-rate**	200kbps
Efficiency	38%	**V_{PA}** **V_{DVS}**	1.2V 0.3-0.5V	**$V_{ni,rms}$**	<2µV_{rms}	**FIR FOM***	0.27	**E/b**	0.8nJ/b
System Power				**Bandwidth**	0-320 Hz	**AFib E/sample @ 0.5V**	6pJ	**Output Power**	-18.5dBm
Current	14µA(R-R) (0.013% TX duty cycle)		294µA (Stream)	**CMRR**	>70dB	**ENV DET E/sample**	0.53pJ	**Band**	400 MHz MICS 433 MHz ISM
Supply	1.35V					**Sub-V_T DVS: 0.3V to 0.6V**	2 kHz-1.7 MHz	**Modulation**	BFSK
FIR FOM*: Power (nW) / frequency (MHz) / # of taps / input bit length / coefficient bit length									

VI. SYSTEM MEASUREMENTS

An ECG experiment was performed on a healthy human subject. In our in-vivo experiments, we used self-adhesive surface electrodes (Kendal Meditrace 535) to acquire ECG signals on the order of a few mVs. One electrode is attached to the chest, while the second (reference) electrode is attached to the abdomen. The two electrode leads are twisted before interfacing with the SoC to minimize interference. First, the chip was set to ECG raw data mode (consuming 397 μW from the 1.35 V V_{BOOST} node) (Fig. 15). Our chip was paired to an unmodified TI CC1101 receiver and a wireless link was successfully established in the 433 MHz ISM band. The reconstructed ECG (dashed, Fig. 15) closely matched the actual ECG. Also shown is the transmitted data, received data and clock waveforms. The zoomed-in section shows one 44 b packet of data, including 9 b header, 32 b data, and 3 b CRC.

TABLE V
PERFORMANCE COMPARISON WITH STATE-OF-THE-ART BSN NODES

	This Work	[5]	[24]	[6]	[25]	[7]
Sensors	ECG, EMG, EEG	ECG	Neural, ECG, EMG, EEG	EEG	ECG, TIV	Temp, Pressure
Supply Voltage	30mV, -10dBm	1.2V	1V	1V	1.2V	0.4V/0.5V
E Harvesting	Thermal, RF	×	×	×	×	Solar
Supply Reg.	✓	×	×	×	×	✓
AFE	4-channel	3-channel	1-channel	18-channel	4-channel	N/A
Power Mgmt.	DPM, Clock gating, Power gating	Clock gating	×	×	×	Power gating
GPP MCU	1.5pJ/Instr. @ 200 kHz (8-bit RISC ISA)	×	×	×	×	28.9pJ/Instr. @ 73 kHz (32-bit CORTEX-M3)
Accelerators	Programmable FIR, AFib, MCU, ENV DET, DMA, Packetizer	ASIC DSP (4x SIMD), FIR, Encryption, DMA	×	ASIC DSP	FIR, Packetizer, Compression	×
Memory	5.5kB (0.3V-0.7V)	42kB (1.2V)	×	×	20kB (1.2V)	5kB (0.4V)
DVS	✓	×	×	×	×	×
Digital Power	2.1μW	~12 μW	N/A	2.1μW	500 μW	2.1μW (MCU)
TX (datarate)	200 kb/s	×	100 kb/S	×	1Mbps (on-body link)	×
TX P_{DC} (100% on)	160 μW	×	400 μW	×	2.8 mW	×
TX P_{OUT}	-18.5dBm	×	-16dBm	×	-6dBm	×
TX band	402/433 MHz	×	402/433 MHz	×	20-40 MHz	×
Total Chip Power	19 μW	31.1μW	500 μW	77.1μW	2.4 mW	7.7 μW
Note on Total Power (includes):	1-channel AFE, 8-bit ADC, DSP (R-R extraction), and TX duty-cycled at 0.013%	AFE, 12-bit ADC, DSP (heart beat detection)	1-channel AFE, 8-bit ADC, and streaming TX with 100% duty-cycle	18-channel AFE, 12-bit ADC, and DSP (EEG feature extraction)	4-channel AFE, 10-bit ADC, DSP (data compression, FIR), SRAM, TX at 5% duty cycle	Data acquisition, DSP (DFT), storage in SRAM
Technology	130nm	180nm	130nm	180nm	180nm	180nm

Next, the chip ran an R-R interval extraction algorithm on the MCU and transmitted measured heart-rate every 5 s operating from a 30 mV harvested input (Fig. 16). Every 5 s, V_{BOOST} is sampled to check for sufficient available energy, in which case the crystal oscillator is enabled for 20 ms before the TX transmission, which takes 650 μs including turn-on time and transmission of a 24-bit packet. The heart-rate extractor algorithm measures the R-R interval with a time resolution of (1/128)s (Fig. 16). The extremely low duty-cycle of TX and crystal oscillator dramatically reduces the power consumption of the TX and crystal oscillator to negligible amounts.

In AFib detection mode, the R-R and AFib accelerators enable the TX and transmit the last 8 beats of raw ECG (buffered in the DMEM) only when a rare AFib event occurs. Measurement results for the AFib demo are presented in Fig. 17. A pre-recorded set of AFib data from MIT-BIH database is used for this demo [23]. Detection occurs 12 R-R intervals after the inception of an AFib event. A pattern recognition algorithm determines if an AFib has occurred [21]. The total chip power in both the R-R and AFib modes is 19 μW, and the chip is powered exclusively from a 30 mV harvested input.

Fig. 18 presents a current breakdown of the R-R extraction demo. The current is nearly evenly distributed among different components, and selective transmission significantly reduces the average power consumption of the transmitter. Fig. 19 shows the micrograph of the 2.5 mm $\times$ 3.3 mm batteryless BSN SoC (130 nm CMOS), and Table IV gives a performance summary.

VII. CONCLUSIONS

Table V shows a performance comparison table with recent BSN SoCs. This work is the first wireless bio-signal processing chip enabling battery-free operation. The chip can be powered from an input as small as 30 mV, enabling thermal energy harvesting. For on-chip power management we augment power and clock-gating by using the DPM, a custom chip controller, to in-

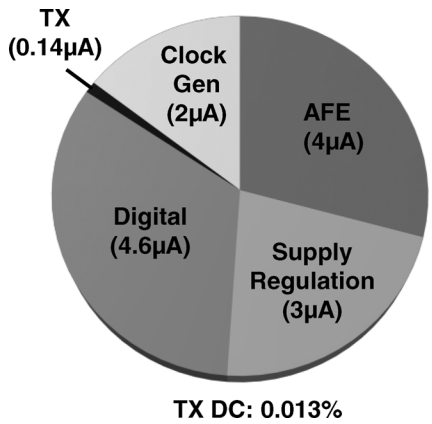

Fig. 18. Current breakdown for R-R extraction experiment. The current distribution is roughly evenly distributed amongst main contributors, and the originally power-hungry transmitter consumption is now nearly mitigated.

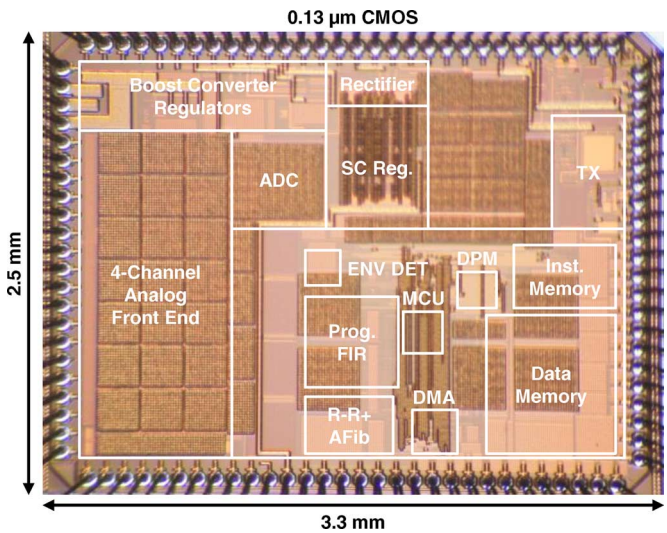

Fig. 19. Annotated chip die photo.

telligently handle energy consumption based on the available energy. The closed-loop power management 'stoplight' scheme enables potentially indefinite operation while the node is worn. Our 5.5 kB of on-chip memory remains operational down to a subthreshold voltage of 0.3 V and is compatible with the flexible subthreshold datapath. In the heart-rate extraction mode where the transmitter is duty-cycled, the entire chip, including V_{DD} regulation, only consumes 19 μW. To the best of the authors' knowledge, this system has lower power, lower minimum input supply voltage, and more complete system integration than other reported wireless BSN SoCs to date.

REFERENCES

[1] G. Z. Yang, *Body Sensor Networks*. London, U.K.: Springer-Verlag, 2006.

[2] B. Gyselinckx, C. Van Hoof, J. Ryckaert, R. F. Yazicioglu, P. Fiorini, and V. Leonov, "Human++: Autonomous wireless sensors for body area networks," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2005, pp. 13–19.

[3] V. Leonov, T. Torfs, P. Fiorini, and C. Van Hoof, "Thermoelectric converters of human warmth for self-powered wireless sensor nodes," *IEEE Sensors J.*, vol. 7, no. 5, pp. 650–657, May 2007.

[4] B. H. Calhoun, S. Khanna, Y. Zhang, J. Ryan, and B. Otis, "System design principles combining subthreshold circuits and architectures with energy scavenging mechanisms," in *Proc. IEEE Int. Symp. Circuits and Systems*, 2010, pp. 269–272.

[5] H. Kim, R. F. Yazicioglu, S. Kim, N. Van Helleputte, A. Artes, M. Konijnenburg, J. Huisken, J. Penders, and C. Van Hoof, "A configurable and low-power mixed signal soc for portable ECG monitoring applications," in *2011 Symp. VLSI Circuits Dig.*, June 2011, pp. 142–143.

[6] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. P. Chandrakasan, "A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 804–816, Apr. 2010.

[7] G. Chen, M. Fojtik, K. Daeyeon, D. Fick, J. Park, M. Seok, M.-T. Chen, Z. Foo, D. Sylvester, and D. Blaauw, "Millimeter-scale nearly perpetual sensor system with stacked battery and solar cells," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2010, pp. 288–289.

[8] F. Zhang, Y. Zhang, J. Silver, Y. Shakhsheer, M. Nagaraju, A. Klinefelter, J. Pandey, J. Boley, E. Carlson, A. Shrivastava, B. Otis, and B. Calhoun, "A batteryless 19 μW MICS/ISM-band energy harvesting body area sensor node SoC," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2012, pp. 298–300.

[9] A. T. Barth, M. A. Hanson, H. C. Powell, and J. Lach, "TEMPO 3.1: A body area sensor network platform for continuous movement assessment," in *Proc. 6th Int. Workshop on Wearable and Implantable Body Sensor Networks*, 2009, pp. 71–76.

[10] Laird Technologies, "Thermoelectric module," ZT Series ZT4, 12, F1, 3030 datasheet, 2009 [Online]. Available: http://www.lairdtech.com/

[11] E. J. Carlson, K. Strunz, and B. P. Otis, "A 20 mV input boost converter with efficient digital control for thermoelectric energy harvesting," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, Apr. 2010.

[12] Y. K. Ramadass and A. P. Chandrakasan, "A batteryless thermoelectric energy harvesting interface circuit with 35 mV startup voltage," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 333–341, Jan. 2011.

[13] S. O'Driscoll, A. Poon, and T. H. Meng, "A mm-sized implantable power receiver with adaptive link compensation," in *IEEE ISSCC Dig.*, 2009, pp. 294–295, 295a.

[14] Y. Ramadass and A. P. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Proc. IEEE Power Electronics Specialists Conf.*, 2007, pp. 2353–2359.

[15] F. Zhang, A. Mishra, A. Richardson, and B. Otis, "A low-power EEG/EEG processing IC with integrated multiband energy extractor," *IEEE Trans. Circuits Syst.*, vol. 58, no. 9, pp. 2069–2082, Sep. 2011.

[16] L. Di, M. Putic, J. Lach, and B. H. Calhoun, "Power switch characterization for fine-grained dynamic voltage scaling," in *Proc. Int. Conf. Computer Design*, Aug. 2008, pp. 605–611.

[17] Microchip, "8-bit PIC Microcontrollers," PIC16C5X datasheet, 2002 [Online]. Available: http://www.microchip.com/

[18] N. Verma and A. P. Chandrakasan, "A 256 kb 65 nm 8T subthreshold SRAM employing sense-amplifier redundancy," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 141–149, Jan. 2008.

[19] F. Zhang, A. Mishra, A. G. Richardson, S. Zanos, and B. Otis, "A Low-Power Multi-Band EEG/EEG Interface IC," in *Proc. IEEE Custom Integrated Circuits Conf.*, Sep. 2010.

[20] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, Mar. 1985.

[21] D. Lake and J. R. Moorman, "Accurate estimation of entropy in very short physiological time series: The problem of atrial fibrillation detection in implanted ventricular devices," *Am. J. Physiol. Heart and Circulatory Physiol.*, Jan 2011.

[22] J. Pandey and B. Otis, "A sub-100 μW MICS/ISM band transmitter based on injection-locking and frequency multiplication," *IEEE J. Solid-State Circuits*, vol. 46, no. 5, pp. 1049–1058, May 2011.

[23] Date accessed: May, 2010. [Online]. Available: http://ecg.mit.edu/

[24] S. Rai, J. Holleman, J. N. Pandey, F. Zhang, and B. Otis, "A 500 μW neural tag with 2 μVrms AFE and frequency-multiplying MICS/ISM FSK transmitter," in *IEEE ISSCC Dig.*, 2009, pp. 212–213, 213a.

[25] Y. Long, J. Bae, S. Lee, T. Roh, K. Song, and H.-J. Yoo, "A 3.9 mW 25-electrode reconfigured sensor for wearable cardiac monitoring system," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 353–364, Jan. 2011.

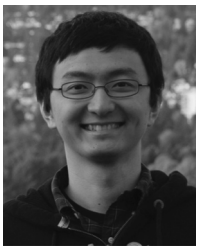

Yanqing Zhang is currently in his fourth year working towards the Ph.D. degree in electrical engineering, after joining the RLP-VLSI group at the University of Virginia, Charlottesville, in 2009.

He was a circuit research scientist intern with nVidia in 2012. His research interests include low power SoCs and SoC architecture, synthesis and timing closure for subthreshold circuits, standard cell library design for robustness, and low power, robust block level design methods.

Fan Zhang received the B.S. degree in electrical engineering and computer sciences from the University of California at Berkeley in 2007, and the M.S. degree in electrical engineering from the University of Washington, Seattle, in 2009. She is currently working toward the Ph.D. degree at University of Washington.

From 2006 to 2007, she interned at Agilent Technologies, Santa Clara, CA. In 2009, she interned at Intel Corporations, Hillsboro, OR, and worked on modeling of sigma-delta phaselocked loops. She has worked on low-noise analog front-end and system-level designs for biomedical sensing applications. Her current research interests include low-power/low-voltage RF circuit design for biomedical applications.

Yousef Shakhsheer received the B.S. degree in computer and electrical engineering from the University of Virginia, Charlottesville, in 2008. He is currently pursuing the Ph.D. degree at the University of Virginia.

His research interests include energy constrained chip architecture, body sensor nodes architecture, energy-efficient circuits, and energy harvesting-specific power management.

Jason D. Silver received the B.S. and M.S. degrees in electrical engineering from Arizona State University, Tempe, in 2006 and 2009, respectively. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Washington, Seattle.

In 2008 and 2009, he held internships with ON Semiconductor in Phoenix, AZ. In 2012, he held an internship at Broadcom Corporation. His research interests include mixed signal integrated circuit design with emphasis on low power telemetry and signal processing for sensor interfaces.

Alicia Klinefelter received the B.S. degree in computer and electrical engineering from Miami University, Oxford, OH, in 2010. She is currently pursuing the Ph.D. degree at the University of Virginia, Charlottesville.

Her research interests include low-power circuit design for wireless sensor nodes, novel architectures for sub-threshold DSPs, and the modeling of low-power system architectures for SoC design.

Manohar Nagaraju (S'09) received the B.S. degree from R.V. College of Engineering, Bangalore, India, and the M.S. degree from the University of Utah, Salt Lake City. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Washington, Seattle.

He is currently an RFIC intern in the Mobile Wireless Group (MWG) at Intel, Hillsboro, OR, where he is working on integrated RF receiver frontends. His research interests include analog and RF integrated circuits for low-power wireless sensors.

James Boley received the B.S. degree in electrical engineering from the University of Virginia, Charlottesville, in 2010. He is a third year student working towards the Ph.D. degree in electrical engineering at the University of Virginia.

He interned at ARM in the summers of 2011 and 2012. His research interests include subthreshold circuit design, low power SRAM design, and SRAM design automation and optimization.

Jagdish Pandey (S'07) received the B.Tech. degree from the Indian Institute of Technology, Madras, in 2003, and the M.Sc. degree from the Indian Institute of Science, Bangalore, in 2007. He received the Ph.D. degree from the University of Washington, Seattle, in 2011.

He was an RFIC design intern at Broadcom Corp. in 2010. He is currently with Qualcomm Inc., San Diego. His research interests are in the area of energy efficient communication circuits.

Aatmesh Shrivastava received the B.Tech. degree in electronics and communication engineering from Birla Institute of Technology, India, in 2006. He worked as a Senior Design Engineer at Texas Instruments from 2006 to 2010. He joined the RLP-VLSI group at the University of Virginia, Charlottesville, in June 2010 where he is working towards the Ph.D. degree.

His research interests include low-power circuit design, clock and energy harvesting circuits for ultra-low power systems like wireless/body sensor networks.

Eric J. Carlson received the B.S. and M.S. degrees from the University of Washington, Seattle, in 2006 and 2008, respectively. He was awarded the Grainger Foundation fellowship in 2006, and the Robert Rushmer fellowship in 2008.

He joined National Semiconductor in 2009 (now Texas Instruments), where he currently works in the RF Power group. His technical interests include DC-DC converters, power management for low power wireless circuits, and energy harvesting.

Austin Wood received the Bachelor degree in electrical engineering from the University of Washington, Seattle, with a focus on embedded computing systems.

He currently works at a small engineering firm outside Seattle that designs control systems.

Benton H. Calhoun (S'02–M'05–SM'12) received the B.S. degree in electrical engineering from the University of Virginia, Charlottesville, in 2000. He received the M.S. degree and Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, in 2002 and 2006, respectively.

In January 2006, he joined the faculty at the University of Virginia in the Electrical and Computer Engineering Department, where he is now an Associate Professor. He is a Campus Director and Technical Thrust Leader in the NSF Nanosystems Engineering Research Center for Advanced Self-Powered Systems of Integrated Sensors and Technologies (ASSIST). His research interests include body area sensor nodes (BSN), wireless sensing, low power digital circuit design, sub-threshold digital circuits, SRAM design for end-of-the-roadmap silicon, variation tolerant circuit design methodologies, and low energy electronics for medical applications. Dr. Calhoun is a co-author of "Sub-threshold Design for Ultra Low-Power Systems" (Springer, 2006), author of "Design Principles for Digital CMOS Integrated Circuit Design" (NTS Press, 2012), and has over 90 peer reviewed publications.

Brian P. Otis (S'96–M'05–SM'10) received the B.S. degree in electrical engineering from the University of Washington, Seattle, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley.

He is currently an Associate Professor at the University of Washington. His research interests are ultra-low power RFIC design, FBAR-based clocks, and bioelectrical interface circuits and systems. He has held positions at Intel, Agilent Technologies, and Google Inc.

Dr. Otis has served as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II and is a member of the ISSCC Technical Program Committee. He received the U. C. Berkeley Seven Rosen Funds award for innovation in 2003, was co-recipient of the 2002 ISSCC Jack Raper Award for an Outstanding Technology Directions Paper, received the National Science Foundation CAREER award in 2009, and was awarded the University of Washington College of Engineering Junior Faculty Innovator Award in 2011.