A Supply-Rail-Coupled eTextiles Transceiver for Body-Area Networks

Patrick P. Mercier, *Student Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—This paper presents a transceiver that communicates over electronic textiles as an alternative, energy-efficient communication medium for body-area network (BAN) applications. The proposed eTextiles network architecture consists of a two-wire conductive yarn medium, body-worn nodes, and a basestation used for data collection and medium-access control. Fabricated in 0.18 μm CMOS technology, the eTextiles transceiver employs supply-rail-coupled differential signaling to efficiently time-share the eTextiles medium between communication and remote charging activities. Remote charging achieves up to 96% power transfer efficiency when a basestation battery is used to charge remote ultra capacitors, which are used as the power supplies of body-worn nodes. Operating at 0.9 V and at 10 Mb/s, the receiver and transmitter front-ends together consume 3.2 pJ/bit over 1 m, which is at least 20X more efficient than conventional BAN receiver front-ends. The transceiver also contains an integrated digital baseband and medium access controller, which, together with the receiver front-end, consume 110 μW during continuous operation.**

***Index Terms*—BAN, body-area networks, conductive fabrics, electronic textiles, eTexiltes, PAN, personal area networks, power line communication, supply-rail coupled, transceiver.**

I. Introduction

ADVANCES in sensor technologies and integrated electronics are revolutionizing how humans access and receive health care. One current thrust in this area is to transition medical monitoring operations away from expensive hospital environments and move them towards the home. So-called *telemedicine* or *connected health* systems employ the use of body-worn sensor nodes for remote monitoring of vital signs or other physiological signals. Monitoring of this nature is not only useful for patients with chronic conditions, but can also be valuable to those at high risk of medical emergencies (e.g., a monitoring system could alert an ambulance at the onset of a heart attack), or those undergoing physical stress (e.g., monitoring dehydration in athletes or soldiers).

A common constraint shared by many sensor-based systems is the need for devices that are physically small and have long operational lifetimes between battery charges. This constraint is of particular importance to body-worn applications, where aesthetics and anatomy necessitate small form factors. Another common requirement shared by many sensor-based systems is their need for communication. Unfortunately, communication circuits often dominate the power budget of sensor nodes, thereby limiting the achievable node lifetime per unit battery weight. This is particularly evident when using wireless circuits for communication distances in excess of 10 m, where power amplifier (PA) power consumption scales directly with transmit distance and ends up dominating the entire system power budget [1].

From a system perspective, it makes sense to push complexity away from locations where weight and volume (and as a result, energy) are constrained, and move this complexity to locations where energy is more abundant. Thus, rather than having each sensor node communicate directly over a cellular or WiFi-type of network, it is beneficial to have body-worn sensor nodes communicate over a local body-area network (BAN). This limits communication distances to under a few meters, pushing the bulk of the transmit power requirements to an energy-rich local relay (e.g., a cellular phone or pager-type device).

Unfortunately, even when communication distances are on the order of one meter, communication circuit power still often dominates the power budget of sensor nodes. Due to short transmit distances (and therefore limited required radiated power), the power budget is not completely dominated by the transmitter's PA in this situation; instead, the receiver front-end constitutes a large fraction of system power [2].

As a result, it is still imperative to increase the energy efficiency of communication circuits in order to achieve the desired goal of small, long-lasting body-worn sensor nodes, even in the limit of short transmit distances. This paper addresses this requirement by exploring the use of electronic textiles (or eTextiles) as a communication medium for body-area networks. A transceiver is designed and fabricated to demonstrate the concept, resulting in communication circuit energy efficiencies that are at least 20X higher than previous BAN approaches.

II. One-Meter BAN Communication

This section describes various means of achieving 1-m body-area network communication, and examines the strengths and weaknesses of each approach.

A. Wireless Communication

From a user's perspective, the untethered mobility gained by wireless communications can, in certain cases, significantly improve freedom of patient movement. Unfortunately, the inherent broadcasting nature of wireless communication raises significant privacy, security, and reliability concerns. To overcome

Manuscript received July 22, 2010; revised January 05, 2011; accepted February 08, 2011. Date of current version May 25, 2011. This work was supported in part by the FCRP Focus Center for Circuit & System Solutions (C2S2), under contract 2003-CT-888.

The authors are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139 USA (e-mail: pmercier@mit.edu).

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

Digital Object Identifier 10.1109/JSSC.2011.2120690

these issues, the BAN must be designed to be robust to eavesdroppers (privacy), malicious attackers (security), and interference from other users (reliability). As an example to combat privacy concerns, an AES engine could be used to encrypt data before transmission, at the cost of increased circuit complexity and energy consumption [3].

The path loss faced by radiated electromagnetic (EM) waves around the human body can also be significant, measured to vary between 40 dB and 100 dB at 2.4 GHz, depending on body movement [4], [5]. As a result, a BAN system employing wireless communication must amplify signals, either at the transmit or receive side, to compensate for the dramatic path loss and its variation. Unfortunately for sensor node size considerations, RF amplification is typically power-expensive and inefficient, especially if good linearity and noise performance is required [6].

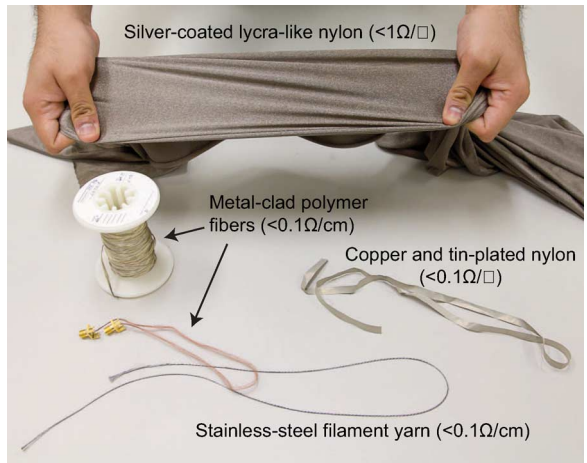

Fig. 1. Various conductive fabrics for use as eTextiles communication mediums.

B. Body-Coupled Communication

Another emerging BAN technique employs the human body as a communication medium. This scheme, termed body-coupled communication (BCC), uses near-field electrostatic coupling to transmit information between nodes that are on or near the human body [7]. Such a technique is a good solution to some of the broadcasting issues associated with traditional wireless techniques, since only a small amount of electromagnetic radiation is generated when carrier frequencies are sufficiently low.

It has been experimentally shown that the path loss of BCC systems operating at near optimal frequencies (30–120 MHz) is on the same order as 1-m wireless systems [8], [9], thereby requiring power-hungry transmit and/or receive amplification. Additionally, since the human body acts as a receiving antenna at these frequencies, the inherently single-ended BCC signaling scheme is particularly susceptible to common-mode interference (for example, from FM radio stations). This requires designs that transition away from low-complexity modulation schemes (such as digital NRZ) [10], and towards more complex interference-robust schemes (such as correlation-based schemes, or frequency-shift keying with channel-hopping) [2], [11], [12], thereby increasing power consumption.

C. Electronic Textiles

As an alternative to high path-loss broadcasting schemes, electronic textiles, or eTextiles, integrates electronics with clothing using conductive fabrics. Such fabrics are often much more flexible and/or stretchable than traditional wires, easing integration into clothing that can be worn during daily-use activities. Conductive fabrics have essentially zero path loss, are relatively cheap to manufacture, and can potentially survive many washing cycles [13]; Fig. 1 presents several common conductive fabrics. Electronic textiles have previously been used in several applications, including wearable motherboards [14], LED-based fashion [15], and biomedical sensing [16], [17]. It is also naturally possible to use eTextile wires purely for BAN purposes, for example as short-range inductive links or direct node-to-node wiring [18], [19].

Compared to wireless and BCC systems, which are inherently single-ended, a wired eTextiles link also provides an opportunity to use multiple parallel wires. As discussed in Sections IV and V, this permits the use of low-complexity two-wire differential modulation schemes that are robust to common-mode interference, requiring little increase in circuit power consumption. Additionally, such a two-wire interface may be used as a vehicle to transfer energy between storage elements at extremely high efficiencies.

By virtue of these advantages, this work employs an eTextiles network as a communication and power delivery medium for BAN applications. Unlike [18] and [19], which optimize the efficiency of communication and power transfer *between* clothing layers (using wireless cm-range inductive links), the goal of this work is to optimize the efficiency over a *single* clothing layer for much longer distances (e.g., 1 m).

III. System Architecture

A. Network-on-Shirt Architecture

The implemented eTextiles-based body-area network system is shown in Fig. 2 [20]. The eTextiles network-on-shirt consists of stainless-steel filament yarn sewn into an argyle pattern on a plain cotton T-shirt. The pattern is designed such that there are two conducting grids: an exterior grid (the visible diamond-patterned wires) and an interior grid (the visible X-patterned wires), corresponding to nodes v_- and v_+, respectively.[1] The use of two separate conductive channels permits the use of low-complexity differential signaling and high-efficiency remote charging.

The X-crossing locations are electrically connected using conductive epoxy to ensure conductivity throughout the grid. The resistance between X-crossings is measured to be 1.8 Ω. The maximum resistance between any two points on the shirt is measured to be 3.1 Ω.

Body-worn sensor nodes (henceforth referred to as "sensor nodes") attach to the grid using metallic button-snaps that are commonly found in fabric stores. One half of each button is electrically connected to the grid using conductive epoxy, while

[1]This is a very simple pattern; it is certainly feasible to design more complex patterns with multiple layers in future work.

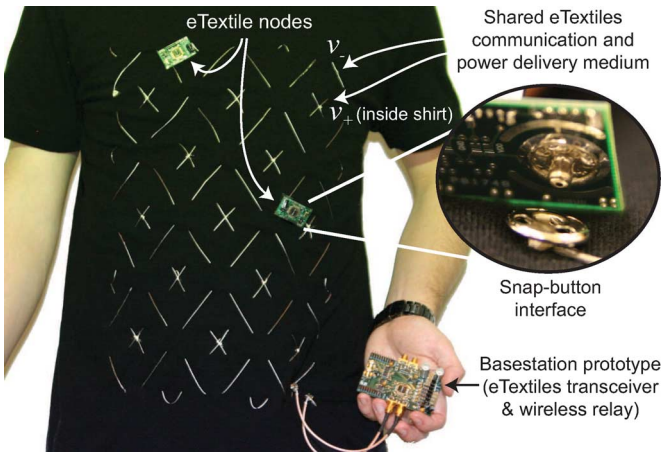

Fig. 2. Implemented eTextiles-based body-area network system using conductive yarn as a time-shared communication and remote charging medium.

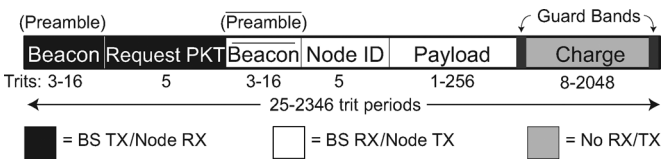

Fig. 3. Packet diagram indicating when sensor nodes and the basestation transmit, receive, and enter the remote charging phase.

the other half is soldered to a standard printed circuit board (PCB). As a result, sensor nodes can only reside in fixed locations. However, sensor nodes directly broadcast over the entire grid network, easing snap button placement restrictions.[2] Broadcasting also provides fault-tolerance on the network, since there are many redundant paths for signals to traverse.

A basestation node is electrically connected to the network-on-shirt, and is the primary destination for sensor node transmissions. The basestation controls medium access using a time-domain multiple-access (TDMA) scheme by ensuring all sensor nodes strictly adhere and synchronize to the packet structure shown in Fig. 3. Section VI-B discusses the medium access scheme in more detail.

B. Transceiver Architecture

Communication and remote charging is performed by an eTextiles transceiver chip, whose block diagram is shown in Fig. 4. Note that this same chip, albeit configured differently, is used in both sensor nodes and in the basestation.

The two main I/O nodes, v_+ and v_-, connect to two metallic snap buttons soldered to the PCB, which then connect to the two-wire eTextiles network. Signaling on the network is performed differentially and directly at baseband, using NRZ ternary digits (trits) as data symbols $\{+1, 0, -1\}$. As discussed in Section VI, trits are used in the preamble to improve synchronization performance, while binary payload data is mapped to trits using the following convention: $0 \rightarrow -1$ and $1 \rightarrow +1$ (ternary '0' is a "don't care" when decoding binary data).

On-chip, nodes v_+ and v_- feed directly into the receiver front-end, consisting of four time-interleaved acquisition (AQ) blocks. Each AQ block contains a sample and hold network, a 1.5 bit ADC, and a ternary encoder. The AQ outputs feed two separate receiver back-end (RX BE) blocks: a correlation module, used for packet synchronization, and a demodulation module, used for decoding ternary payloads into bits. An on-chip digital baseband controls the flow of trits and bits from the receiver back-end, and an on-chip medium access controller (MAC) ensures the chip is synchronized to the packet diagram enforced by the basestation node. Nodes v_+ and v_- are also the output of differential capacitively coupled transmitters. Since the receiver inputs are naturally high impedance when turned off and the transmitter outputs are capacitively coupled, no explicit transmit/receive switch is required.

At the end of the communication phase of a packet, nodes v_+ and v_- are effectively left floating. Remote charging is enabled at this point by turning on charging transistors M1 and M2, which connect nodes v_+ and v_- to the supply terminals of each chip on the network. Each sensor node chip is equipped with a large capacitor, C_{DD}, across its supply terminals. The network is synchronized such that the basestation, which has a battery across its supply terminals, turns on its charging transistors at the same time as the sensor nodes, thereby remotely charging C_{DD}.

The required size of C_{DD} is ultimately dictated by the length of time during any potential power outages (for example, the time required between programming the node off of the network, and when the node is placed and synchronized to the network). Related the capacitor size, the absolute amount of time dedicated to remote charging is set by the basestation during network configuration, and is dictated by the amount of time spent in the communication packet phases (i.e., the duty cycle between communication and remote charging), the power consumption of the chip during the active and charging states, the effective resistance between battery terminals and the C_{DD} terminals, the relative sizes of various C_{DD}s on the network (usually set for the worst case), and the acceptable amount of voltage droop during the active phase. For system demonstration purposes (i.e., Fig. 2), a 0.47 F ultra capacitor is used, allowing for many minutes of node operation during basestation power outages or off-network programming. Alternatively, a much smaller 100 μF capacitor can be used, providing several seconds of node operation during basestation power outages. A charging/communication duty cycle of 80% is typically used, resulting in a remote charging period of approximately 100 μs.

Since nodes v_+ and v_- are nominally centered at V_{DD} and GND, respectively,[3] transmission gates are not required to pass charge. Instead, PMOS-only and NMOS-only pass transistors are used for nodes v_+ and v_-, respectively. The pass transistors are sized to be large in order to minimize the effective charging resistance. Since the eTextiles network is again left floating after

[2]Future work could involve replacing snap buttons with cm-range wireless inductive links for the sensor node-to-eTextiles link, as is done in [19], to further ease sensor node location restrictions. Once data is on the eTextiles network, the proposed transceiver could be used as an efficient way to move information and energy across the 1-m BAN.

[3]Large resistors help keep v_+ and v_- at these potentials during network start up and during sleep periods. These resistors are sized such that they will not substantially affect the network potential during long strings of consecutive $+1$'s and -1's over the course of a 256-bit payload.

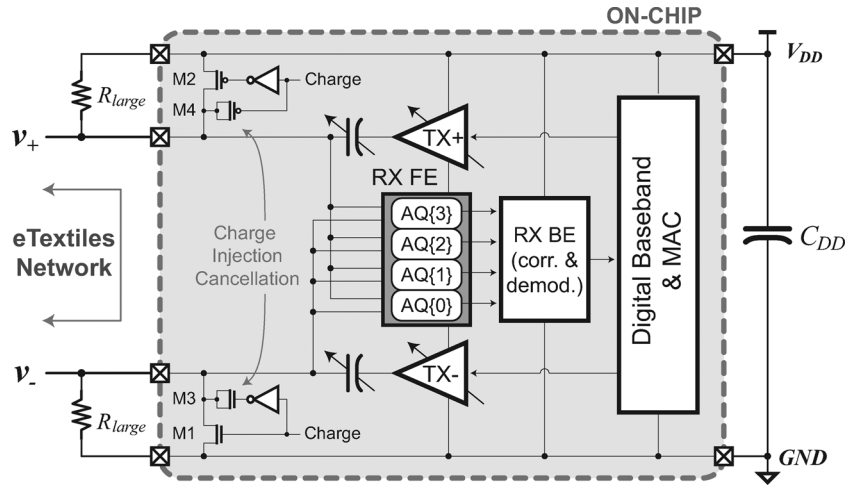

Fig. 4. eTextiles transceiver block diagram.

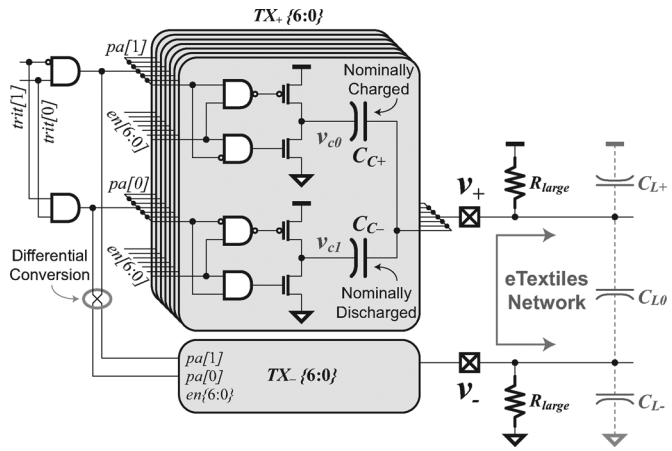

Fig. 5. Supply-rail-coupled differential ternary transmitter.

the remote charging phase, charge injection effects can potentially disrupt the floating voltage on the network. Equally-sized charged-injection-cancellation transistors, M3 and M4, are used to attenuate this effect.

The primary design goal of the eTextiles transceiver is to minimize the energy required to communicate over the BAN for personal sensing applications. Such applications typically require data rates on the order of 10 to 100 kb/s per sensor. Given multiple sensors per network, in addition to coding and duty-cycling overhead, an instantaneous data rate of 10 Mb/s is chosen for the transceiver.

IV. Transmitter Front-End

The transmitter front-end is shown in Fig. 5, and is comprised of entirely digital circuits. Two separate driver blocks, $TX_+\{6{:}0\}$ and $TX_-\{6{:}0\}$, are used to generate differential signals on the eTextiles network. Unlike wireless and BCC systems, differential signaling permits direct digital modulation schemes that have inherent resilience to common-mode interferers. Up-conversion of the 10 Mb/s baseband signal to a higher frequency is not performed, since this would require increased power consumption due to oscillator and mixers stages. Additionally, frequencies above 100 MHz have wavelengths approaching the physical size of the eTextiles network, potentially producing undesired EM radiation and multi-path effects.

To save CV^2 energy, low-swing digital transmissions are used. Specifically, capacitive-driving is employed, permitting the use of simple tri-state inverter drivers to couple signals on to the network [21], [22]. There are several advantages to this approach. For instance, a second voltage supply is not necessary to generate the low-swing waveforms; instead, the resulting output swings seen on the network are proportional to the capacitive divider ratios between the coupling capacitors, C_{C+} and C_{C-}, and the parasitic eTextiles network capacitance, C_{L+}, C_{L0}, and C_{L-}. For the design in Fig. 2, C_{L0} is measured to be 35 pF, with ± 5 pF variation when the user moves around. Since C_{L0} is bound to differ amongst different users, the voltage swing can be made programmable by simply varying the coupling capacitor sizes. In this implementation, seven parallel sets of binary-weighted coupling capacitors are driven by binary-weighted inverters in each differential transmitter, resulting in 7 bits of output swing programmability. The parallel drivers are individually enabled through an on-chip configuration register which sets the $en[6{:}0]$ signals. To overcome any capacitance variation induced by user motion, the output swing must be configured with sufficient margin in order to guarantee reliable communication.[4] Capacitive-driving also permits the use of smaller, more efficient drivers, since the effective load seen by each driver is not the entire eTextiles network, but is rather the load of the much smaller coupling capacitor; this allows for reduced parasitic CV^2 switching energy in the driver buffer circuits. When the tri-state inverters are disabled, their output impedances are very high, thereby ensuring that the loading of the network by unused drivers is dominated by the parasitic capacitance at nodes v_{c0} and v_{c1}, instead of the much larger coupling capacitance, C_{C+} and C_{C-}.

Another benefit of using capacitively coupled transmitters is that the network can more efficiently switch between remote charging and communication states by using *supply-rail-coupled* differential signaling. To illustrate, first recall that

[4]Real-time output swing calibration could be performed in future work by adding an ADC on the output node and feeding back swing information. This would likely contribute minimal parasitic capacitance relative to the existing network capacitance, and can be aggressively duty-cycled to minimize power.

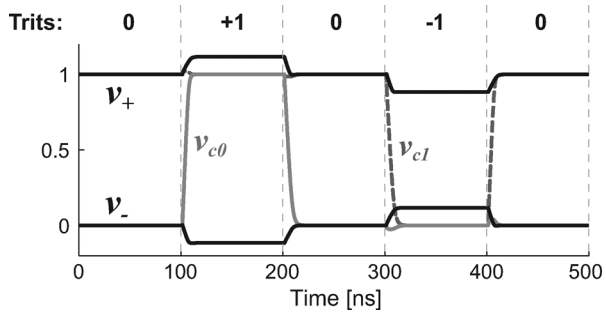

Fig. 6. Simulated supply-rail-coupled transmitter timing diagram for transmitter TX_+.

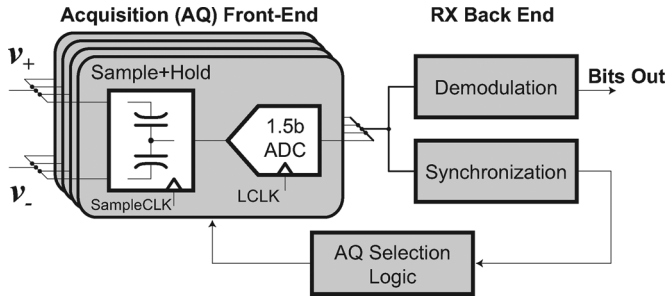

Fig. 7. Receiver block diagram.

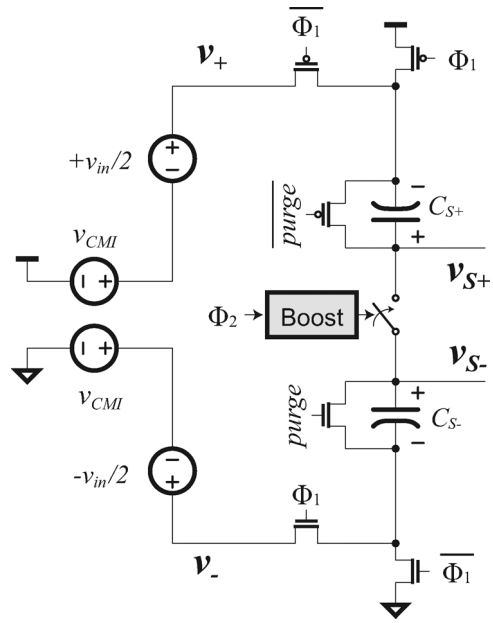

Fig. 8. Common-mode independent supply-rail-coupled sample-and-hold schematic.

the eTextiles medium is left floating after the remote charging phase. Specifically, nodes v_+ and v_- are left floating at $V_{\rm DD}$ and GND, respectively. Conventional differential signaling schemes are designed to operate over a pair wires centered at the same potential (nominally GND); disrupting the floating voltage on the eTextiles network by discharging v_+ to GND would require consuming unnecessary CV^2 energy. Since the transmitter outputs are capacitively coupled onto the network, nodes v_+ and v_- can instead be left floating at their prior DC potentials throughout a packet.

By using two coupling capacitors, C_{C+} and C_{C-} that are normally charged and discharged, respectively, the supply-rail-coupled technique can additionally permit bipolar signaling [23]. This technique can be used to generate ternary data, which aids in receiver synchronization.

The transmitter generates bipolar signals as follows. Ternary digits are generated in the digital back-end, and are sent to the transmitter in a 2-bit two's complement binary representation (*trit*[1:0]). The trits are applied to the input of the driving tri-states, ultimately changing the voltages at nodes v_{c0} and v_{c1}. As shown in Fig. 6, for an input ternary value of '0', or *trit*[1:0] = 00, node $v_{c0} = GND$ and $v_{c1} = V_{\rm DD}$ for transmitter TX_+. Since node v_+ is nominally at $V_{\rm DD}$, capacitors C_{C+} and C_{C-} are nominally charged and discharged, respectively.

In order to transmit a ternary '+1' (*trit*[1:0] = 01), v_{c0} is pulled up to $V_{\rm DD}$, thereby discharging capacitor C_{C+}, and creating a small positive voltage swing on the output node v_+. Node v_{c1} is left unchanged here. To return to a zero state, capacitor C_{C+} is recharged by de-asserting v_{c0}. In order to transmit a ternary '−1' (*trit*[1:0] = 11), v_{c1} is pulled down to GND. This charges capacitor C_{C-} (which is nominally discharged), creating a small negative voltage swing on the output node v_+. Node v_{c0} is left unchanged here. Again, to return to a zero state, capacitor C_{C-} is re-discharged by asserting v_{c1}.

A similar arrangement is made for TX_-, thereby completing the generation of supply-rail-coupled differential ternary signals.

V. Receiver Front-End

A block diagram of the receiver architecture is shown in Fig. 7. Since the path loss across the eTextiles medium is effectively zero, low-noise pre-amplification is not required. Supply-rail-coupled differential signals are instead directly sampled and digitized by four time-interleaved acquisition (AQ) blocks in the receiver front-end. Each AQ block contains dedicated sample-and-hold units and 1.5-bit ternary analog-to-digital converters (ADCs). Depending on the packet phase, output samples of the AQ blocks either flow into a synchronization module, used to select the optimal set AQ of front-end units, or into a demodulator, used to decode payload bits. No circuits in the receiver require static bias currents; as a result, only CV^2 switching and subthreshold leakage energies are consumed.

A. Supply-Rail-Coupled Sampling

A schematic of the AQ block sample-and-hold unit is shown in Fig. 8. The sampler consists of PMOS- and NMOS-only transistors as sampling switches for nodes v_+ and v_-, respectively, exploiting the fact that the differential supply-rail-coupled signals are biased at opposite supply rails (otherwise, transmission gates would be required). The top-plate sampling switch is implemented as a transmission gate since, as will be shown shortly, the capacitor top plates are nominally at half-supply. When operating at supply voltages below 1 V, it is necessary to boost the gate of the NMOS sampling switch device to a larger voltage in order to maintain a sufficiently low overall on-resistance; this

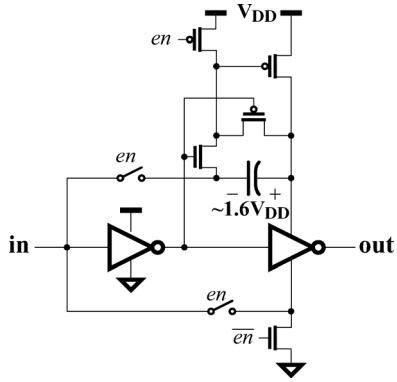

Fig. 9. Sampling switch boosting circuit, enabled when operating at supply voltages below 1 V.

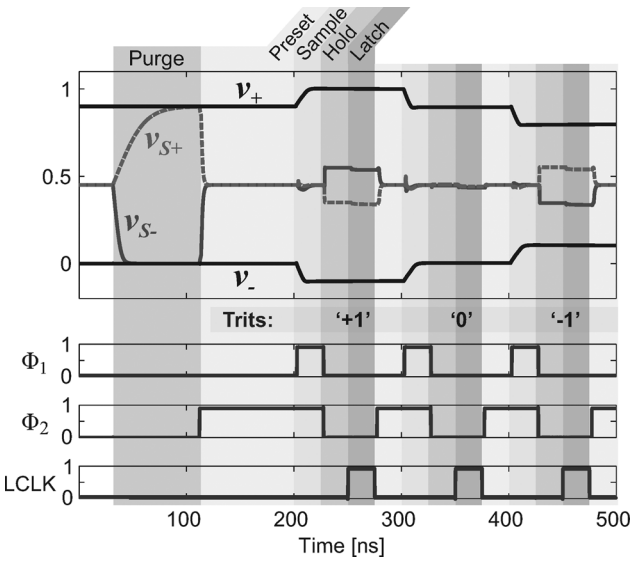

Fig. 10. Simulated supply-rail-coupled sampling transient waveforms.

is accomplished using the charge-pump circuit shown in Fig. 9, such that a second dedicated supply voltage is not required [24].

The left side of Fig. 8 shows a representation of the supply-rail-coupled input signals, v_+ and v_-:

$$v_+ = V_{\rm DD} + v_{\rm CMI} + v_{\rm in}/2$$
$$v_- = GND + v_{\rm CMI} - v_{\rm in}/2$$

where $v_{\rm CMI}$ is a common-mode interference signal that may have coupled onto the eTextiles network.

Simulated waveforms illustrating the sampling procedure are shown Fig. 10, using a supply voltage of $V_{\rm DD} = 0.9$ V, and, for this illustration, $v_{\rm CMI} = 0$.

At the beginning of packet reception, the sampling capacitors, C_{S+} and C_{S-}, are purged of their charge by shorting their respective terminals to a fixed supply rail. Purging must be performed semi-occasionally, as the passively-set capacitor top-plate voltage will eventually be disrupted by leakage currents. Experimental results show that over 5000 bit cycles are required before bit errors start to accrue because of this effect. For reasonably sized packets (e.g., in this case, payloads of 256 bits), leakage has minimal effect on the bit error rate (BER); as a result, it is sufficient to perform purging only at the beginning of every packet.

During normal sampling operation, the sample-and-hold circuit cycles through four stages: preset, sample, hold, and latch. During the preset phase, $\Phi_1 = 0$ and $\Phi_2 = 1$, thus connecting the bottom capacitor plates to opposite rails while connecting their top plates together. As a result, a potential $\alpha V_{\rm DD} = [C_{S+}/(C_{S+} + C_{S-})] \cdot V_{\rm DD}$ is developed at nodes v_{S+} and v_{S-}. If C_{S+} and C_{S-} are matched, $\alpha = 0.5$, leaving equal and opposite amounts of charge on each sampling capacitor:

$$Q_{S+} = -C_{S+}(\alpha V_{\rm DD})$$
$$Q_{S-} = C_{S-}(\alpha V_{\rm DD}).$$

During the sample phase, Φ_1 is asserted, connecting the bottom capacitor plates directly to the eTextiles network. Since the top capacitor plates remain tied together, and the differential supply-rail-coupled signals move in equal and opposite directions, differential-mode charge is accumulated on the capacitors:

$$Q_{S+} = -C_{S+}\left((1-\alpha)V_{\rm DD} + v_{\rm in}/2 + v_{\rm CMI}\right)$$
$$Q_{S-} = C_{S-}(\alpha V_{\rm DD} + v_{\rm in}/2 - v_{\rm CMI}).$$

Note that the top plate potentials, v_{S+} and v_{S-}, remain at $\alpha V_{\rm DD}$ during this phase.

The hold phase begins by de-asserting Φ_1 and Φ_2, connecting the bottom plates back to the supplies and opening the top plate switch. The top plate switch is designed to open before the bottom plate switches by passing Φ_1 through a four-stage inverter-based delay chain with non-minimum length devices. This timing is important for two main reasons. First, this avoids the scenario where the sampled charge would be disrupted by the supply rails (if the bottom plates were connected to the supplies before the top plate switch was opened). Second, since the bottom plate switches are sized such that the total impedance is symmetric around the top plate nodes, the channel charge accumulated in the top-plate switch ideally distributes evenly across both sampling capacitors, resulting in only common-mode charge-injection effects [25].

After the hold phase and into the latch phase, the charge on the sampling capacitors remains the same as in the sampling phase. However, after some settling time, the bottom capacitor plates are now referenced to the supplies, rather than the eTextiles network. This creates the desired differential voltages on the top plate nodes:

$$v_{S+} = \alpha V_{\rm DD} - v_{\rm in}/2 - v_{\rm CMI}$$
$$v_{S-} = \alpha V_{\rm DD} + v_{\rm in}/2 - v_{\rm CMI}.$$

The salient features of this approach are that the top plate voltages are automatically centered at DC potentials suitable for the ensuing comparator inputs, and, taken differentially, the common-mode interference terms and capacitor mismatch terms cancel out ($v_{S+} - v_{S-} = -v_{\rm in}$, regardless of the value of $V_{\rm CMI}$ or if $\alpha = 0.5$). Similar common-mode independent sampling approaches have been previously proposed in ADC design, albeit without supply-rail-coupled input biases [25].

Once the sampler outputs have settled, the results may be latched by the ensuing ternary ADC on the rising edge of LCLK.

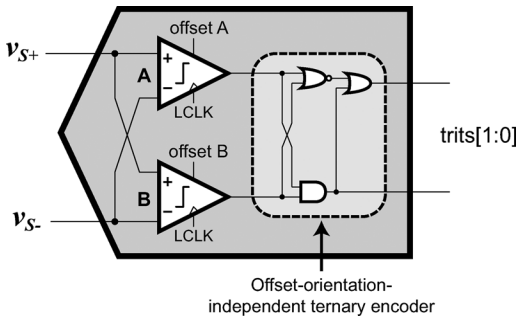

Fig. 11. 1.5-bit flash-like ternary ADC.

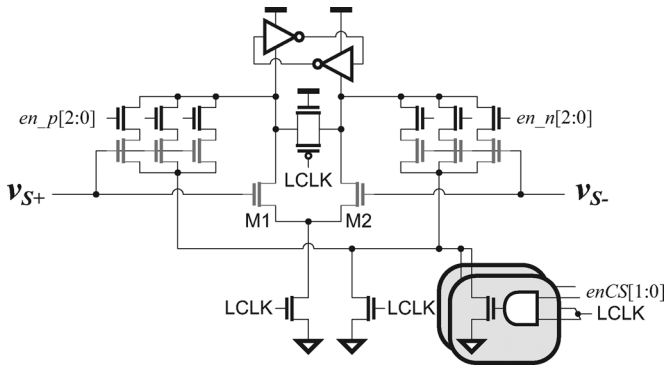

Fig. 12. Programmable-offset comparator (pre-charge transistors not shown).

Timing signals Φ_1, Φ_2, and LCLK are all generated off of the master system clock using a simple 2-bit counter consisting of dual-edge-triggered flip-flops [26].

B. Ternary ADC

The differential outputs of the sampler feed directly into a 1.5-bit ternary ADC. The ADC, shown in Fig. 11, consists of two regenerative clocked comparators followed by some simple encoding logic.

The comparator, shown in Fig. 12, is a Strong-ARM inspired sense amplifier with a primary differential input pair consisting of transistors M1 and M2 [27]. The primary devices are sized for a nominal 3σ input-referred offset voltage of 25 mV and have a common-mode input range of 0.35 to 0.9 V. Offset tuning is attained using weighted supplementary differential input devices, controlled by six enable signals, $en\_p[2:0]$ and $en\_n[2:0]$, similar to [28]. However, unlike [28], these supplementary input devices use a separate, tunable tail current source, configured with 2 bits of tuning signals $enCS[1:0]$. This expands the overall tuning range by varying the strength of the supplementary devices relative to the primary devices, for a total of eight tuning bits. Capacitive trimming could have been used to achieve the large desired tuning range [29], however at the expense of additional CV^2 switching energy. The tuning calibration routine is discussed in more detail in the Appendix.

The ADC is flash-like, in the sense that the input signal is applied in parallel to the comparators, which are intentionally configured to have equal and opposite non-zero offsets. A sampled signal whose differential amplitude falls below the absolute offset level (i.e., has low amplitude) is converted to ternary '0'.

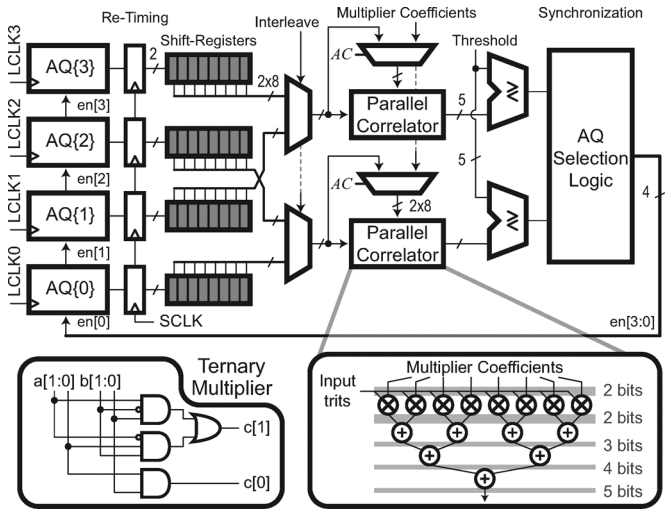

Fig. 13. Receiver back-end used for synchronization between the transmitter and receiver.

On the other hand, if the differential input voltage $|v_{S+} - v_{S-}|$ is greater than the comparator offset voltage, the output will be a ternary '+1' or '−1', depending on the sign of the input signal.

The conversion between the comparator outputs and the 2-bit binary-coded ternary signals is performed by an offset-orientation-independent ternary encoder. Consisting of a few simple logic gates, the encoder permits the comparator pair to swap roles. That is, if during calibration it is found that comparator A is better suited to be configured with a negative offset than comparator B,[5] then the two comparators may be configured as such without requiring any additional configuration of the encoder. The measured benefits of this type of comparator configuration redundancy are discussed in the Appendix.

Since the comparators are programmed to have equal and opposite offsets, any metastable states would only occur at transitions between ternary values[6] and would affect only a single comparator. The encoder is designed such that valid logic levels should propagate in case this unlikely scenario is encountered.

VI. Digital Baseband and Medium Access Controller

A. Synchronization

As in most communication systems, synchronization between the transmitter and receiver is essential for overall system performance. In this work, synchronization more precisely refers to three separate steps: 1) *detection* of valid incoming signals; 2) *phase-synchronization* to align clocks between the transmitter and receiver; and 3) *packet-synchronization* to align the demodulator with the beginning of the payload. The computations for these three steps are performed simultaneously in the receiver back-end, shown in Fig. 13, during the beacon (preamble) phase of the packet (Fig. 3).

As described in Section V-A and illustrated in Fig. 14, each AQ block samples for one half clock period every two cycles.

[5]In this case, comparators A and B are normally configured with positive and negative offsets, respectively.

[6]In other words, the differential input voltage would have to *exactly* equal one of the comparator input-referred offset voltages.

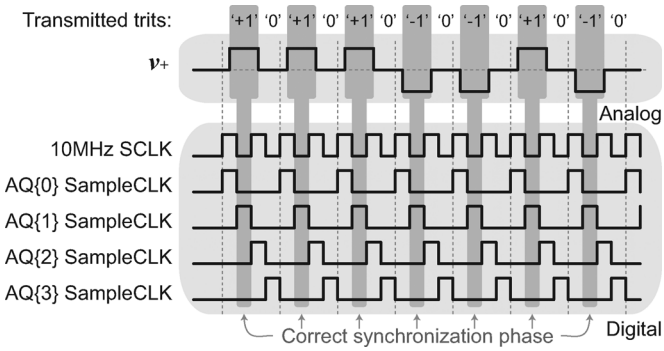

Fig. 14. An illustrative beacon transmission and the corresponding interleaved sampling clocks. Since each AQ unit samples for one half clock period every two cycles, two beacon codes are interleaved such that only a single AQ unit captures the appropriate data.

For data reception at the 10 MHz system clock (SCLK) frequency, two time-interleaved AQ blocks, spaced apart by one clock cycle, are required.

When only two time-interleaved AQ blocks are used, input signals are still only sampled for one half clock period every cycle. If a transmitter and receiver are phase-aligned, this situation is sufficient for payload demodulation. However, such alignment is not guaranteed (nor expected) at the beginning of every packet. Thus, phase-synchronization is required in order to accurately align transmitter and receiver clocks. This is accomplished by using an extra pair of AQ blocks, such that sampling covers the entire clock period.

Synchronization is accomplished by correlating input samples from all four AQ blocks with a pre-determined code; the AQ block that registers the highest correlation value determines the implied phase- and packet-synchronization, provided the correlation value exceeds a pre-programmed threshold (for detection purposes). The use of ternary signaling introduces some softness in the data and decision making process.

During the synchronization phase, the basestation's transmitter broadcasts a beacon code across the network. For example, the code used in Fig. 14 is a 7-trit sequence, S_1, that is interleaved with a second code, S_0, which in this case consists entirely of erasures (i.e., ternary '0's) for illustrative purposes,[7] to create the following 14-trit sequence:

$$S_b = [+1\ 0\ +1\ 0\ +1\ 0\ -1\ 0\ -1\ 0\ +1\ 0\ -1\ 0].$$

Since AQ units sample for one half clock period every two cycles, each AQ unit effectively only observes every second trit in the overall beacon code, S_b. To illustrate, AQ{1} in Fig. 14 correctly samples the desired 7-trit code S_1, while AQs{2–3} sample seven consecutive erasures (code S_0). In this illustration, AQ{0} ideally also has a chance of correctly sampling S_0. However, the period of data covered by its sampling window is beyond the specifications of the system; as a result, decoding errors may occur, resulting in a likely lower correlation value than AQ{1}. Due to finite transmitter rise times, this is rarely an issue in practice. However, extra precaution could be added into the system by using non-overlapping sampling clocks or adding very small guard-bands to each non-zero ternary transmission.

[7] S_0 should ideally be a code with the greatest Hamming distance to S_1; this would, however, unnecessarily complicate this illustrative example.

The actual correlations are performed as follows. Ternary symbols from the AQ blocks, which are clocked using different phases, are re-timed to SCLK, and then shifted into serial-to-parallel registers every time new symbols are available, as shown in Fig. 13. Once the registers are full, the parallel data is then interleaved into two parallel correlators. The correlators multiply the parallel input trits with up to eight parallel ternary multiplier coefficients (with the example code, S_b as used in Fig. 14, only seven coefficients are required). Since ternary multiplication has several "don't care" results, a custom multiplication structure is used, saving 2 bits in each ensuing adder stage over a traditional two's complement implementation. If one of the correlation results crosses a pre-programmed threshold, detection is achieved. The AQ selection logic block then feeds back control signals to the receiver front-end for phase synchronization purposes, indicating which pair of AQ blocks to gate for the remainder of the packet. Immediately after the threshold is exceeded, the receiver waits for a fixed number of clocks cycles until the payload begins, thereby achieving packet synchronization.

B. Medium Access

The basestation controls medium access using a TDMA scheme by acting as the master node on the star-configured eTextiles network. Sensor nodes are configured as slave devices, each pre-programmed with a unique 5-trit ID code.

Communication on the eTextiles network strictly follows the protocol outlined in the packet diagram (Fig. 3); sensor nodes do not change their internal wake-up schedule while connected to the network, unless directed to do so by the basestation (e.g., if a new node is added to the network and additional charging time is required). Synchronization is performed by each node once per packet. Since the signaling scheme is directly at baseband, active clock phase and frequency tracking is not required between nodes. For example, up to 50 ns of phase error can be tolerated between a transmitter and receiver before trit errors occur. When operating off of a 10 MHz crystal with 30 ppm accuracy, this occurs after 50,000 clock cycles (5 ms). The maximum packet length is 29.8 μs—well within these bounds.

The basestation begins each packet at the exact same time by broadcasting a beacon signal across the network, after which the receivers of all sensor nodes are synchronized with the basestation's transmitter. Immediately following the beacon, the basestation transmits a 5-trit request code, containing the ID code of a particular sensor node. The sensor nodes who do not receive their own ID codes at this point enter a sleep mode. The sensor node that does receive its ID code proceeds to retransmit an inverted beacon across the network, followed again by its own ID code. This is done to ensure the basestation's receiver is synchronized to the sensor node's transmitter, as well as to provide an additional means of false packet detection (for instance, if the wrong node ID is sent back to the basestation). The sensor node then transmits its payload to the basestation. Following the payload, all nodes on the network wake up from sleep using their watchdog timers and enter the remote charging phase, as described in Section III-B. Small guard bands are inserted at this

TABLE I
COMPARISON OF BODY-AREA NETWORK COMMUNICATION SYSTEMS

	[30]	[32]	[12]	[18]	[18]	**This work**
Type	Wireless UWB	Wireless MICS	BCC	Wireless eTextiles	Wired eTextiles	Wired eTextiles
Range	10m	2m	1.8m	0.012m	0.85m	1m
CMOS Technology	0.13μm	90nm	0.18μm	0.25μm	0.25μm	0.18μm
Supply Voltage	1.2V/0.95V	0.7V	1.0V/1.5V	2.5V/3.3V	2.5V/3.3V	0.9V
Data Rate	39.1Mbps	120kbps	10Mbps	10Mbps	10Mbps	10Mbps
RX FE energy	69pJ/bit	10,000pJ/bit	225pJ/bit	1.1pJ/bit	–	2pJ/bit
TX FE energy	–*	2900pJ/bit	90pJ/bit	2.9pJ/bit	–	0.7-to-18pJ/bit
RX and baseband power	4.2mW	–*	3.7mW	>0.17mW*	2.7mW	0.11mW

*transmitter, digital baseband, and/or network controller not included.

point to avoid charging contention on the network, and ease the synchronization procedure.

Sensor nodes can be dynamically added to the network at any time. Upon initial power up, each sensor node IC alternates between synchronization and sleep states until it synchronizes with the beacon. However, the basestation does not know the node has been added to the network, and as a result, will not request any packets from it. To add new nodes to the basestation queue, the basestation will periodically broadcast a special request code following the beacon, called the *welcome request* code. Any node on the network who has not yet been recognized will attempt to respond to the welcome request code by following the standard packet convention: broadcasting an inverted beacon, followed by its ID code, and an empty payload. To avoid contention on the network, sensor nodes in this state only respond with a certainly probability, similar to a p-persistance scheme. At this point, the basestation and sensor nodes handshake over the next two packets, formally adding the sensor node to the network.

In the absence of a master basestation node, the sensor nodes can instead be programmed to use a carrier-sense multiple-access (CSMA) MAC with collision avoidance. Network activity is sensed in the receiver back-end (Fig. 13) by asserting the AC signal, turning the correlators into autocorrelators.[8] If the auto-correlation result is sufficiently low (i.e., the network is quiet), the sensing node proceeds to transmit a beacon and payload, again using p-persistence to avoid contention.

VII. EXPERIMENTAL RESULTS

The eTextiles transceiver chip is fabricated in a 0.18 μm process, and occupies a core area of 0.83 mm^2; a die photo is shown in Fig. 15. Both digital and analog circuits are designed to operate down to voltages as low as 0.9 V while being clocked at 10 MHz. Table I summarizes the chip results while comparing the results to other 1-m BAN solutions.

Fig. 16 presents measured transient results across the network shown in Fig. 2, demonstrating both transceiver and network operation. The measurement directly probes nodes v_+ and v_- on the eTextiles network, as well as the demodulated data stream from the basestation's receiver back-end.

With 7 bits of binary weighted coupling capacitors, the transmitter achieves voltage swings ranging from 6 to 260 mV across

[8]This is another important reason for using ternary data; autocorrelation results are more robust on ternary data than on binary data.

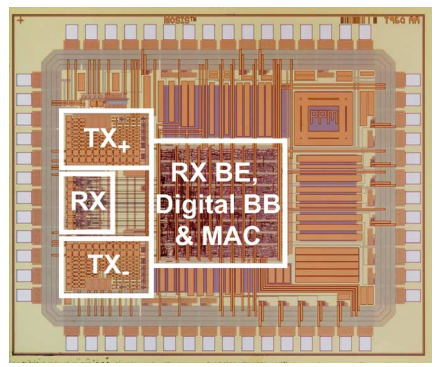

Fig. 15. Die photograph of the eTextiles transceiver.

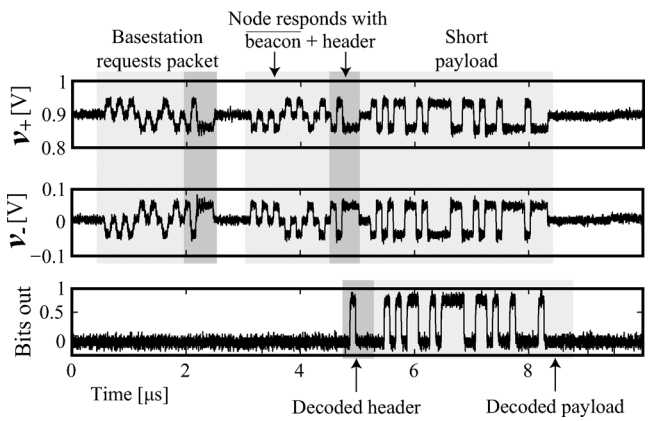

Fig. 16. Measured transient response across the eTextiles network, illustrating a typical packet request by the basestation.

a 1-m eTextiles network.[9] This results in a transmitter front-end energy consumption of 0.7 to 18 pJ per transmitter bit. The receiver front-end consumes 2 pJ/bit, regardless of the input signal amplitude. As shown in Table I, these results are up to 35X more efficient than an already very efficient wireless ultra-wideband (UWB) receiver [30], greater than 100X more efficient than comparable BCC systems [11], [12], and on the same order as a wireless eTextiles approach that operates over 1.2-cm distances [18]. Due to the energy-efficient front-end combined with low-voltage digital baseband operation, the chip consumes 110 μW during full receive-mode duty cycle (including receiver front-end, back-end, digital baseband, and I/O power), which is at least 20X lower than other fully integrated receiver systems

[9]A 1-m eTextiles link was used, as this facilitates reasonable comparison to other BAN approaches, and is also a good length for a generic size-medium t-shirt. Longer and/or shorter designs are certainly possible.

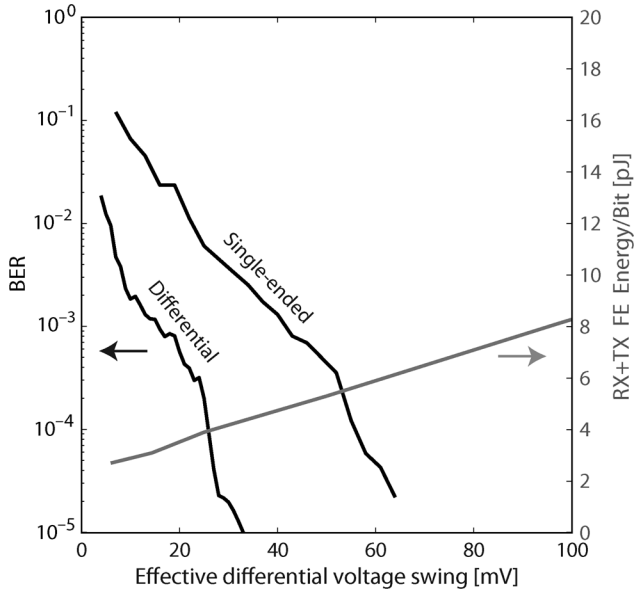

Fig. 17. Measured receiver BER curves and combined receiver and transmitter front-end energy consumption as a function of the effective differential voltage swing seen on the 1-m eTextiles network.

(Table I), including the wired portion of the eTexiles transceiver in [18]. Unlike wireless and BCC approaches which operate over inherent media, the proposed transceiver does, however, require the use of an eTextiles network which must be worn during operation.

To quantify the communication performance of the system, the receiver BER is characterized with respect to varying transmitter voltage swings for both differential and single-ended operation across a 1-m eTextiles network. Measurement results, taken in a noisy lab environment with significant RF (e.g., AM and FM radio, 802.11, cellular, etc.) and powerline interference, are shown in Fig. 17. Note that single-ended measurements were made with double the transmitted voltage swing to facilitate fair comparisons to the differential results. To achieve a BER of 10^{-3}, the differential case requires 15 mV, while the single-ended case requires 42 mV. This 2.8X difference illustrates the communication advantage of differential signaling in noisy environments. The right axis of Fig. 17 also shows the combined receiver and transmitter energy consumption in differential operation. At the receiver sensitivity of 15 mV, the transmitter consumes 1.2 pJ/bit, resulting in a total front-end energy consumption of 3.2 pJ/bit.

Although the sensitivity of the receiver could potentially be improved by adding linear amplification in front of the comparators, this would likely come with an overall system energy penalty. For instance, adding linear amplification would increase the energy consumption of the receiver front-end, increasing its sensitivity, thereby requiring lower transmitter voltage swings. However, since the receiver already constitutes a large fraction of the total front-end energy consumption, marginally decreasing the transmitter energy consumption while raising the increasingly dominant receiver energy consumption would not likely result in overall energy savings.

When operating at a charging/communication duty cycle of 80%, the remote charging scheme achieves 96% power transfer efficiency, measured over a 1-m distance as the ratio of power going into the sensor node storage capacitor over the basestation's output power.[10] At a 50% duty cycle, 91% power transfer efficiency is achieved. Existing wireless power transfer schemes achieve upwards of 55% power transfer efficiencies over much shorter distances [18].

VIII. Conclusion

To address the broadcasting and/or path loss issues of wireless and BCC systems, this paper proposes an eTextiles transceiver system as an inherently secure and private communication scheme. The proposed eTextiles transceiver uses supply-rail-coupled differential signaling to efficiently time-share the eTextiles medium between communication and remote-charging circuits. Fabricated in 0.18 μm CMOS technology and operating at 0.9 V, the resulting transmitter and receiver front-ends together consume 3.2 pJ/bit, which is greater than 20X more efficient than conventional 1-m BAN approaches. At 100% receive-mode duty cycle operating at a data rate of 10 Mb/s, the transceiver consumes 110 μW, including the power from an integrated digital baseband. Remote charging, where a basestation's battery is directly connected to a sensor node's power supply through the eTextiles network and remote charging transistors, achieves up to 96% power transfer efficiency, thereby enabling small sensor nodes which last for long time intervals between charges.

Appendix
Comparator Calibration

For reliable system operation, the absolute comparator offsets are typically set to be slightly less than the absolute voltage swing on the network. For example, given ± 15 mV voltage swings, comparators might be configured to have ± 10 mV offsets. In this implementation, offsets are configured manually at start-up.[11] To configure each comparator with a specific offset, it is first necessary to characterize the offset profile of each comparator across all of its tuning bits. This is accomplished using an offline characterization routine, where for each comparator setting, an external voltage ramp is applied to the input and the comparator trip points are recorded. Fig. 18(a) shows an example offset profile for two comparators located in the same AQ unit. In this case, the two comparators achieve an 8-bit offset tuning range of ± 60 mV at a supply voltage of 0.9 V.

Calibration is performed by sorting the offset profiles, and configuring each comparator with the tuning code that most closely results in the desired offset. In the unoptimized case, comparators A and B (Fig. 12) are always programmed to have positive and negative offsets, respectively. This results in the offset error distribution, measured as the difference between the best attainable offset and the desired offset, shown in Fig. 18(b).

In some instances, comparators A and B might be better suited to have negative and positive offsets, respectively. The

[10]This is similar to the measurement of a low drop-out regulator's (LDO's) efficiency.

[11]Future work could involve a heavily-duty-cycled feedback ADC to measure incoming voltage swings, permitting once-per-packet offset configuration based on a fixed voltage-swing/offset ratio. As discussed in Section IV, such an ADC could also be reused to dynamically configure the transmitter voltage swings.

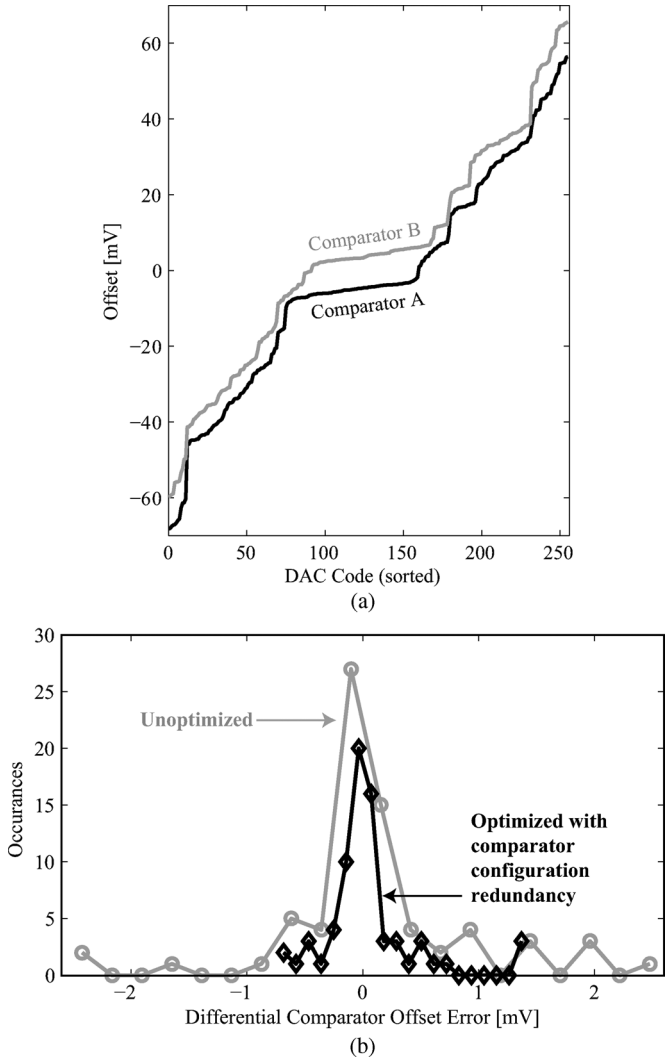

Fig. 18. Comparator calibration results. (a) Measured offset profile of two comparators in an AQ unit. (b) Offset error distribution of 72 different comparators, measured across nine different chips. The comparator configuration redundancy improves the standard deviation of offset errors by 2.1X in this case.

offset-independent ternary encoder enables the comparators to swap roles as such, with no extra configuration required. Thus, the calibration routine can leverage this inherent comparator redundancy in its optimization search: for instance, it can try to minimize the RMS offset error between the two comparator. Fig. 18(b) illustrates the optimized offset error distribution when the calibration routine minimizes the largest offset error. After calibration optimization, the standard deviation of offset errors is improved by 1.5 to 2.5X depending on the optimization criterion. Since two comparators are necessary for proper operation regardless of the optimization, no area overhead is incurred. Similar types of comparator redundancy schemes have been proposed for ADCs and SRAM applications [24], [31].

ACKNOWLEDGMENT

The authors thank Jose Bohorquez and Marcus Yip for design and CAD support.

REFERENCES

[1] A. P. Chandrakasan, R. Min, M. Bhardwaj, S. Cho, and A. Wang, "Power aware wireless microsensor systems," in *Proc. European Solid-State Circuits Conf. (ESSCIRC)*, 2002, pp. 47–54.

[2] N. Cho, J. Bae, and H.-J. Yoo, "A 10.8 mW body channel communication/MICS dual-band transceiver for a unified body sensor network controller," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3459–3468, Dec. 2009.

[3] M. Feldhofer, J. Wolkerstorfer, and V. Rijmen, "AES implementation on a grain of sand," in *IEE Proc. Inf. Security*, Oct. 2005, vol. 152, pp. 13–20.

[4] J. Ryckaert, P. De Doncker, R. Meys, A. de Le Hoye, and S. Donnay, "Channel model for wireless communication around human body," *Electron. Lett.*, vol. 40, no. 9, pp. 543–544, Apr. 2004.

[5] E. Reusens, W. Joseph, B. Latré, B. Braem, G. Vermeeren, E. Tanghe, L. Martens, I. Moerman, and C. Blondia, "Characterization of on-body communication channel and energy efficient topology design for wireless body area networks," *IEEE Trans. Inf. Technol. Biomed.*, vol. 13, no. 6, pp. 933–945, Sep. 2009.

[6] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[7] T. Zimmerman, "Personal area networks (PAN): Near-field intra-body communication," M.S. thesis, Massachusetts Inst. Technol., Cambridge, MA, 1995.

[8] N. Cho, J. Yoo, S. Song, J. Lee, S. Jeon, and H.-J. Yoo, "The human body characteristics as a signal transmission medium for intrabody communication," *IEEE Trans. Microw. Theory Tech.*, vol. 55, no. 5, pp. 1080–1086, May 2007.

[9] T. Schenk, N. Mazloum, L. Tan, and P. Rutten, "Experimental characterization of the body-coupled communications channel," in *Proc. IEEE Int. Symp. Wireless Commun. Syst.*, Oct. 2008, pp. 234–239.

[10] S. Song, N. Cho, and H.-J. Yoo, "A 0.2-mW 2-Mb/s digital transceiver based on wideband signaling for human body communications," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 2021–2033, Sep. 2007.

[11] A. Fazzi, S. Ouzounov, and J. van den Homberg, "A 2.75 mW wideband correlation-based transceiver for body-coupled communication," in *IEEE ISSCC Dig. Tech. Papers*, 2009, pp. 204–205.

[12] N. Cho, L. Yan, J. Bae, and H.-J. Yoo, "A 60 kb/s–10 Mb/s adaptive frequency hopping transceiver for interference-resilient body channel communication," *IEEE J. Solid-State Circuits*, vol. 44, no. 3, pp. 708–717, Mar. 2009.

[13] H. Kim, Y. Kim, B. Kim, and H.-J. Yoo, "A wearable fabric computer by planar-fashionable circuit board technique," in *Proc. IEEE Int. Workshop Wearable and Implantable Body Sensor Networks*, Jun. 2009, pp. 282–285.

[14] S. Park, K. Mackenzie, and S. Jayaraman, "The wearable motherboard: A framework for personalized mobile information processing (PMIP)," in *Proc. IEEE Design Automation Conf.*, 2002, pp. 170–174.

[15] L. Buechley and M. Eisenberg, "Fabric PCBs, electronic sequins, and socket buttons: Techniques for e-textile craft," *J. Personal and Ubiquitous Computing*, vol. 13, no. 2, pp. 133–150, Feb. 2009.

[16] E. Scilingo, A. Gemignani, R. Paradiso, N. Taccini, B. Ghelarducci, and D. De Rossi, "Performance evaluation of sensing fabrics for monitoring physiological and biomechanical variables," *IEEE Trans. Inf. Technol. Biomed.*, vol. 9, no. 3, pp. 345–352, Sep. 2005.

[17] H. Kim, Y. Kim, Y. Kwon, and H.-J. Yoo, "A 1.12 mW continuous healthcare monitor chip integrated on a planar-fashionable circuit board," in *IEEE ISSCC Dig. Tech. Papers*, 2008, pp. 150–151.

[18] J. Yoo, S. Lee, and H.-J. Yoo, "A 1.12 pJ/b inductive transceiver with a fault-tolerant network switch for multi-layer wearable body area network applications," *IEEE J. Solid-State Circuits*, vol. 44, no. 11, pp. 2999–3010, Nov. 2009.

[19] J. Yoo, L. Yan, S. Lee, Y. Kim, and H.-J. Yoo, "A 5.2 mW self-configured wearable body sensor network controller and a 12 μW 54.9% efficiency wirelessly powered sensor for continuous health monitoring system," *IEEE J. Solid-State Circuits*, vol. 45, no. 1, pp. 178–188, Jan. 2010.

[20] P. P. Mercier and A. P. Chandrakasan, "A 110 μW 10 Mb/s eTextiles transceiver for body area networks with remote batter power," in *IEEE ISSCC Dig. Tech. Papers*, 2010, pp. 496–497.

[21] R. Ho, T. Ono, R. Hopkins, A. Chow, J. Schauer, F. Liu, and R. Drost, "High speed and low energy capacitively driven on-chip wires," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 52–60, Jan. 2008.

[22] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "Low-power, high-speed transceivers for network-on-chip communication," *IEEE Trans. VLSI Syst.*, vol. 17, no. 1, pp. 12–21, Jan. 2009.

[23] P. P. Mercier, D. C. Daly, and A. P. Chandrakasan, "An energy-efficient all-digital UWB transmitter employing dual capacitively-coupled pulse-shaping drivers," *IEEE J. Solid-State Circuits*, vol. 44, no. 6, pp. 1679–1688, Jun. 2009.
[24] D. C. Daly and A. P. Chandrakasan, "A 6-bit, 0.2 V to 0.9 V highly digital flash ADC with comparator redundancy," *IEEE J. Solid-State Circuits*, vol. 44, no. 11, pp. 3030–3038, Nov. 2009.
[25] N. Verma and A. P. Chandrakasan, "An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes," *IEEE J. Solid-State Circuits*, vol. 442, no. 6, pp. 1196–1205, Jun. 2007.
[26] R. Llopis and M. Sachdev, "Low power, testable dual edge triggered flip-flops," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 1996, pp. 341–345.
[27] J. Montanaro, R. Witek, K. Anne, A. Black, E. Cooper, D. Dobberpuhl, P. Donahue, J. Eno, W. Hoeppner, and D. Kruckemyer *et al.*, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.
[28] W. Ellersick, C. Yang, M. Horowitz, and W. Dally, "GAD: A 12-GS/s CMOS 4-bit A/D converter for an equalized multi-level link," in *IEEE Symp. VLSI Circuits Dig.*, 1999, pp. 49–52.
[29] M. Lee, W. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1591–1599, Nov. 2000.
[30] N. Van Helleputte, M. Verhelst, W. Dehaene, and G. Gielen, "A reconfigurable, 130 nm CMOS 108 pJ/pulse, fully integrated IR-UWB receiver for communication and precise ranging," *IEEE J. Solid-State Circuits*, vol. 45, no. 1, p. 69, Jan. 2010.
[31] N. Verma and A. P. Chandrakasan, "A 256 kb 65 nm 8 T subthreshold SRAM employing sense-amplifier redundancy," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, p. 141, Jan. 2008.
[32] J. L. Bohorquez and A. P. Chandrakasan, "A 350 μW CMOS MSK transmitter and 400 μW OOK super-regenerative receiver for medical implant communications," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1248–1259, Apr. 2009.

Patrick P. Mercier (S'04) received the B.Sc. degree in electrical and computer engineering from the University of Alberta, Edmonton, AB, Canada, in 2006, and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2008, where he is currently pursuing his doctoral degree.

From May 2008 to August 2008, he worked in the Microprocessor Technology Lab at Intel Corporation, Hillsboro, OR, designing spatial encoding and low-swing circuits for multi-core on-chip communication networks. His research interests include the design of energy-efficient digital systems, RF circuits, and power converters for biomedical and implantable applications.

Mr. Mercier was a corecipient of the 2009 ISSCC Jack Kilby Award for Outstanding Student Paper at ISSCC 2010. He also received a Natural Sciences and Engineering Council of Canada (NSERC) Julie Payette fellowship in 2006, NSERC Postgraduate Scholarships in 2007 and 2009, and an Intel Ph.D. Fellowship in 2009.

Anantha P. Chandrakasan (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. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor 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 coeditor 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).

Dr. Chandrakasan was a corecipient 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, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence and the ISSCC Jack Kilby Award for Outstanding Student Paper (2007, 2008, 2009), and the 2009 Semiconductor Industry Association (SIA) University Researcher Award. 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, the Technology Directions Sub-committee Chair for ISSCC 2004–2009, and the Conference Chair for ISSCC 2010. He is the Conference Chair for ISSCC 2011. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on SSCS AdCom from 2000 to 2007 and he was the meetings committee chair from 2004 to 2007.