A 350 μW CMOS MSK Transmitter and 400 μW OOK Super-Regenerative Receiver for Medical Implant Communications

Jose L. Bohorquez, *Student Member, IEEE*, Anantha P. Chandrakasan, *Fellow, IEEE*, and Joel L. Dawson, *Member, IEEE*

***Abstract*—Recent advances in the medical field are spurring the need for ultra-low power transceivers for wireless communication with medical implants. To deal with the growing demand for medical telemetry, the FCC commissioned the Medical Implant Communications Services (MICS) standard in 1999 in the 402–405 MHz band. This paper presents a 350 μW FSK/MSK direct modulation transmitter and a 400 μW OOK super-regenerative receiver (SRR) specifically optimized for medical implant communications. The transceiver is implemented in 90 nm CMOS and digitally tunes 24 MHz in frequency steps smaller than 2 kHz. The transmitter meets MICS mask specifications with data rates up to 120 kb/s consuming only 2.9 nJ/bit; the receiver has a sensitivity better than −99 dBm with a data rate of 40 kb/s or −93 dBm with a data rate of 120 kb/s consuming 3.3 nJ/bit. A frequency correction loop incorporating the base-station is prototyped to eliminate the need for a frequency synthesizer in the implant while still achieving frequency stability of less than 3 ppm.**

***Index Terms*—Capacitor predistortion, digitally-controlled oscillator, direct-modulation transmitter, frequency-control loop, frequency-shift keying, low power, medical implants, MICS, on-off keying, super-regenerative receiver.**

I. Introduction

UNTIL recently, few implantable medical devices existed and fewer still provided the capability for wireless transmission of information. Most devices capable of data transmission did so through inductive coupling, which requires physical contact with the base-station and only allows for low data rates [1]. In 1999, the FCC created the Medical Implant Communications Service (MICS) band in the range of 402–405 MHz specifically for medical telemetry [2]. The MICS band plan allows for RF communication between a medical implant and a base-station that is up to two meters away, giving patients the ability to move around freely while being monitored. A challenge that arises with medical implants, however, is that batteries must last over a decade without the possibility of recharging. As a result, medical implants and their respective transceivers must be optimized to consume as little energy as possible while achieving acceptable levels of performance.

There are a few critical observations that motivate this work. The first is that the human body is an excellent temperature regulator, and the second is that the MICS standard features relaxed output power and frequency stability specifications. Together, these observations point to simplified transceiver architectures that consume less power than those commonly employed for other applications. Furthermore, while it is critical that the implant consume minimal power in order to preserve battery life, the corresponding base-station is free to consume much more power. This observation motivates shifting complexity in the wireless link from the implant to the base-station.

With these observations in mind, we propose a simple, low-power topology where a digitally-controlled oscillator (DCO) is directly modulated using frequency-shift keying (FSK) [3]. Instead of using a PA to drive the loop antenna, the DCO incorporates it as its inductive element, radiating energy that would otherwise be lost as thermal heat [4]. To concurrently achieve an acceptably wide tuning range and fine frequency resolution, a sub-ranged capacitor array is used to divide 20 bits of frequency tuning into coarse, medium, and fine tuning capacitor banks. The capacitor banks are predistorted to achieve linear digital-to-frequency conversion and 14 effective bits of frequency resolution.

A challenge that arises from having the antenna attached to the DCO is that popular receiver architectures such as the superheterodyne or homodyne topologies cannot be used. Instead, we propose a super-regenerative architecture to demodulate on-off keying (OOK), achieving excellent sensitivity and good selectivity while consuming less than 400 μW. By optimizing the system holistically, we achieve data transmission consuming 2.9 nJ/bit and reception consuming 3.3 nJ/bit while meeting the MICS 300 kHz channel bandwidth requirements.

II. Architecture Overview

Fig. 1 shows the prototype transceiver comprising a simple digital baseband implemented in an FPGA, a direct modulation FSK transmitter, and a super-regenerative receiver. The transmitter and receiver are time-division multiplexed and share an external loop antenna implemented on the prototype PCB. The low radiation power requirements of MICS are exploited in the transmitter by eliminating the PA and incorporating the antenna

Manuscript received August 24, 2008; revised November 03, 2008. Current version published March 25, 2009. This work was supported by the Focus Center for Circuit and System Solutions (C2S2), one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program. Chip fabrication was provided by IBM through the Trusted Access Program Office.

The authors are with the Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: joselb@mit.edu; anantha@mtl.mit.edu; jldawson@mtl.mit.edu).

Digital Object Identifier 10.1109/JSSC.2009.2014728

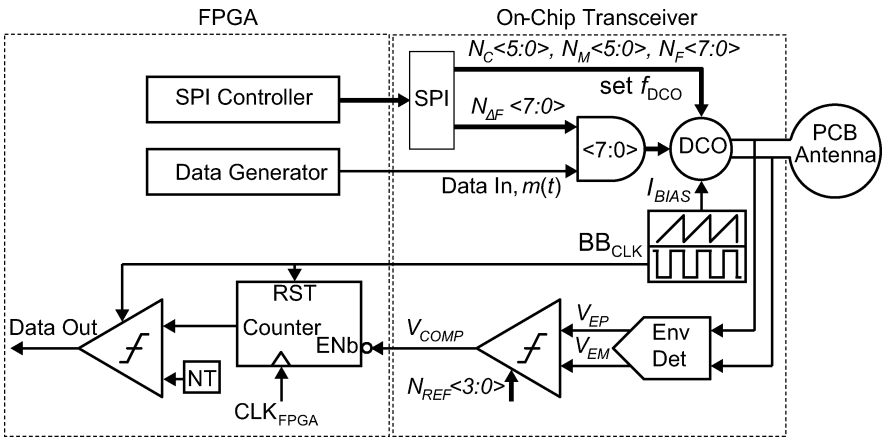

Fig. 1. Transceiver block diagram.

into the DCO.[1] On the receive side, the same DCO is used as a synchronous OOK super-regenerative receiver. As explained in Section IV, the SRR provides a tremendous amount of gain that results in relaxed input-referred noise specifications for subsequent stages. To exploit this benefit, the envelope detector and programmable comparator in Fig. 1 were biased in subthreshold and optimized for ultra-low power consumption.

The following subsections describe the unique features of medical implant communications that were exploited to achieve ultra-low power operation. Section III describes the theory and design of the direct modulation FSK transmitter. Section IV explains the key elements of super-regeneration theory and describes the receiver implementation in detail. Section V describe the circuit implementations of the transceiver. Section VI introduces a prototype frequency correction loop used to calibrate the DCO's frequency without the need for a frequency synthesizer in the implant. Finally, Sections VII and VIII show measurement results for the transceiver and conclude the paper.

A. Exploiting the Unique Features of Medical Implants

A fundamental difference between medical implant transceivers and more conventional radios is the environment in which they operate. Cellular phones, for example, must maintain a frequency stability of less than 1 ppm in the temperature range of -40: $+85\,^{\circ}$C, while the MICS standard only requires 100 ppm of stability in the range of 25–45 °C. Furthermore, while MICS calls for proper implant operation over a 20 °C range, the temperature of the human body rarely changes so drastically, and its moderate shifts occur very slowly. This quality of temperature regulation in the human body can be exploited to reduce the complexity, power consumption, and size of medical implants. For example, by choosing architectures that do not require stringent frequency accuracy in the implant, the classical frequency synthesizer can be replaced by a frequency-control loop that does not require a crystal oscillator in the implant. In Section VI, such a system is described.

Another quality of medical implants that can be exploited is the low bandwidth of most biomedical signals. Since low sample rates can be used to digitize these signals, the transmitter can be off most of the time as data is accumulated. The transmitter is then turned on briefly, sending data packets in short bursts. To maximize the battery life, the transmitter should be designed to minimize the energy consumed per transmitted bit. This criteria creates a tradeoff between complex architectures that are more spectrally efficient but consume more power, and simpler topologies that consume less power but are less spectrally efficient [5]. Direct modulation FSK transmitters have the advantage of consuming very little power due to their simplicity, while achieving good spectral efficiency when modulated using minimum-shift keying (MSK). Other modulation schemes (such as 256-QAM), while more spectrally efficient, require complex topologies that consume more power and may result in lower energy efficiency.

B. Link Analysis and Antenna Considerations

For MICS transceivers, antenna gain (the product of efficiency and directivity) inside the human body is considerably lower than in free-space for two main reasons. First, the size of the antenna is typically much smaller than the wavelength of the signal (75 cm), and second, human tissue is conductive causing large amounts of loss. Fortunately, MICS data links are only intended to function within a 2 meter radius limiting the amount of path loss. Measurements on MICS signal propagation reported in [6] show the worst case path loss in a hospital room to be 47 dB including fading margin. Using these results along with the MICS requirement limiting radiation to -16 dBm EIRP and reasonable assumptions about the base-station lead to requirements for the medical implant.

Assuming the base-station transmits at the maximum -16 dBm EIRP, means that at least -63 dBm of power will be available at any location in the hospital room. This places a restriction on the receiver's sensitivity and antenna gain. For example, if the sensitivity is -92 dBm, the antenna gain (including path loss in the body) must be at least -29 dBi ($P_{\mathrm{sens,imp}} \leq P_{\mathrm{tx,bs}} + G_{\mathrm{path}} + G_{\mathrm{ant}}$, in dB). Using the same antenna to transmit from the implant and assuming the sensitivity of the base-station receiver is -105 dBm (sensible for a 300 kHz bandwidth channel), the power delivered to the antenna by the implanted transmitter must be at least -29 dBm or 1.3 μW ($P_{\mathrm{tx,imp}} \geq P_{\mathrm{sens,bs}} - G_{\mathrm{path}} - G_{\mathrm{ant}}$).

[1]The maximum transmission power is 25 μW effective isotropic radiated power (EIRP), defined as the product of radiated power and the antenna gain.

A *small* loop antenna (relative to its signal's wavelength) can be modeled as the series combination of an inductor and a resistor. The resistance is composed of two elements: radiation resistance and loss [7]. The radiation resistance is desirable since it models the conversion of energy into electromagnetic waves. The loss, expectedly, is not desirable since it reduces the antenna's efficiency:

$$\epsilon_{\rm rad} = \frac{R_{\rm rad}}{R_{\rm rad} + R_{\rm loss}}. \tag{1}$$

The poor efficiency of small loop antennas is manifested in their low radiation resistance and results in their impedance exhibiting a high quality factor (Q). To achieve high power transfer, the inductive element of the antenna must be resonated out by a capacitor. The high Q, however, poses challenges since the frequency response of the resonator is narrow and even small impedance variations can result in significant mismatch.

This last challenge can be overcome, and actually exploited. Since the power needed to drive the antenna in an MICS transmitter is only a small portion of the power budget, a smaller emphasis can be placed on PA efficiency. Using the antenna as the inductive element of an LC-oscillator eliminates the need for an explicit PA since the resulting power oscillator drives the antenna. Further benefits include inherent impedance matching, low-power consumption, and low-noise design since the small loop antenna has a high Q.

A drawback of using the antenna as part of the oscillator is that changes in the antenna's environment lead to frequency pulling. To compensate for this effect, the bandwidth of the frequency-control loop (FCL) must be higher than the maximum rate of change of the implant's environment. Studies on human gait show that most human motion has frequency content of less than 10 Hz [8]. Therefore, as long as the frequency pulling does not exceed the lock-in range of the FCL and the FCL bandwidth is greater than 10 Hz, the effects of frequency pulling can be properly mitigated. To verify that frequency pulling does not pose a serious threat to the functionality of the system, we tested our chip with the simple antenna shown in Fig. 2(a). It uses a 1.6 mm FR4 substrate and superstrate with a single patch on the bottom side. HFSS simulations show this structure to exhibit a Q of 115 in free space, an inductance of 23 nH, and a radiation efficiency of 0.4%. We performed a few simple experiments such as waving a hand near the antenna and the FCL maintained correct center frequency with ease.

Perhaps a more serious threat to the functionality of the system is that the high conductance of human tissue makes a loop antenna lossy, lowering its Q and efficiency [9]. HFSS simulations show that high Q and acceptable efficiency can be maintained by using a substrate and superstrate with metal patches above and below as shown in Fig. 2(b). This antenna has a diameter of 2.3 cm, uses a 4 mm substrate and superstrate (with a relative permittivity of 3.5 and a dissipation factor of 0.0027). HFSS simulations of the structure immersed in a model of human tissue (relative permittivity of 42.8, conductivity of 0.65 S/m) show a worst-case Q of 130 and a radiation efficiency of 0.1%. Note that this Q is slightly better than the Q of the antenna used in the prototype. It follows that the difficulties of transmitting through human tissue will not compromise the performance of this system.

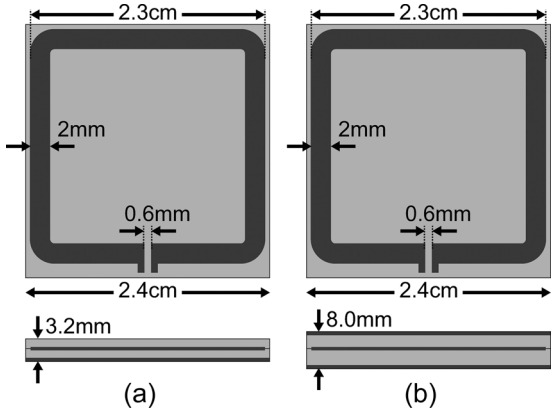

Fig. 2. (a) Loop antenna with FR4 substrate and superstrate, and copper patch below the substrate. (b) Loop antenna with metal patches above and below the substrate and superstrate.

III. FSK/MSK Transmitter

Connecting the antenna directly to an oscillator precludes the use of conventional up-conversion transmitters. Fortunately, frequency and phase modulation techniques can be used by modulating the oscillator directly. The result is a simple transmitter architecture that is spectrally efficient and consumes very little power.

A. FSK/MSK Theory

FSK is a constant envelope form of digital modulation and can be easily implemented by directly modulating the instantaneous frequency of an oscillator:

$$f_i(t) = f_c + \Delta F \times m(t) \tag{2}$$

where f_c is the carrier frequency, ΔF is the frequency deviation constant, and $m(t) \in [-1, 1]$ is the digital modulating signal [10]. For a bit rate R, setting $\Delta F = 0.25R$ results in MSK; the most spectrally efficient form of FSK that still produces orthogonal signaling and can be demodulated coherently. This relaxes SNR requirements in the base station receiver reducing output power requirements on the implanted transmitter. Theoretically, MSK can be used at a bit rate of 200 kb/s without any pre-filtering and meet the MICS spectral mask since its first nulls occur at $f_c \pm 0.75R = f_c \pm 150\,\text{kHz}$ and subsequent peaks are at least 25 dB below the main lobe.

B. Transmitter Implementation

Fig. 3(a) shows a simple FSK transmitter comprising a DCO that incorporates a small loop antenna as its inductive element, and is modulated directly with digital data. Fig. 3(b) shows the simplified schematic of the DCO with an equivalent parallel resistance of

$$R_p \approx Q\omega_0 L \tag{3}$$

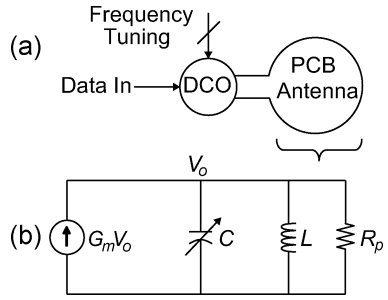

Fig. 3. (a) Direct modulation FSK transmitter and (b) simplified circuit model.

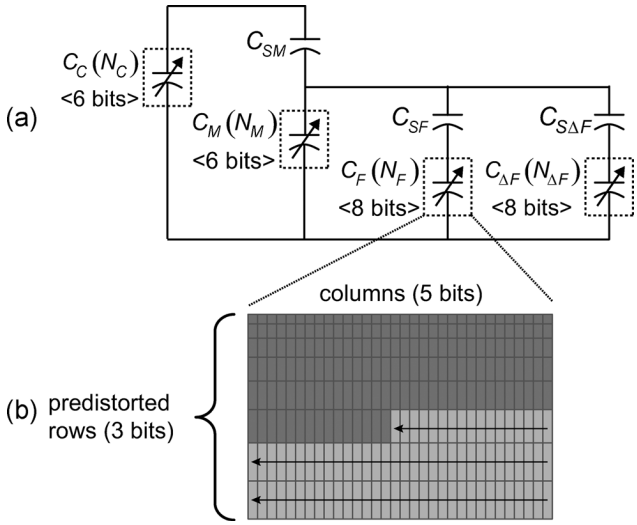

Fig. 4. (a) Sub-ranging capacitor array and (b) piece-wise linear predistortion.

where ω_0 is the tank's resonant frequency. Frequency tuning and modulation are done using switched capacitors that have much lower temperature coefficients than varactors and allow for a fully digital implementation. A challenge that arises is that many bits are required to tune a wide frequency range while achieving small frequency steps. Furthermore, classical implementations of capacitor arrays would require impractically small capacitors. For example, to tune 24 MHz (391–415 MHz) in 2 kHz steps requires 14 bits of resolution. Assuming an inductance value of 24 nH, the capacitance must tune from 6.128 pF to 6.904 pF with a minimum capacitor step size of 0.047 fF. Since it is impractical to implement such small capacitors in CMOS, a sub-ranging capacitor array was implemented to achieve very small effective capacitor step sizes while using practical capacitor values [11].

C. Capacitor Array With Predistortion

The digitally tunable capacitor in Fig. 3(b) is implemented using the four capacitor banks shown in Fig. 4(a). The carrier frequency (f_c) of the DCO is tuned with coarse, medium, and fine-tuning capacitor banks C_C, C_M, and C_F while the frequency deviation constant is set using $C_{\Delta F}$. Each capacitor bank is thermometer coded and predistorted as in [12] to improve linearity in digital-to-frequency conversion and guarantee monotonicity for each bank. This is critical since the DCO will be placed in a FCL and non-monotonicity could result in instability.

Capacitor bank C_C provides six bits of coarse frequency tuning and is directly connected across the inductor. As a result, any change in its capacitance results in an equal change in the total resonator capacitance. A small capacitor C_{SM} is connected in series with capacitor bank C_M to provide six bits of medium frequency tuning. Since C_{SM} is much smaller than C_M, the effective change in the resonator's capacitance is much smaller than changes made to the capacitor bank C_M. This allows switch capacitors on the order of 10 fF to be used while achieving incremental capacitance changes across the inductor on the order of 0.5 fF. Similarly, C_{SF} is used to further reduce the effective capacitance changes across the inductor when the fine tuning capacitor bank C_F is tuned. The result is incremental resonator capacitance changes on the order of 30 aF achieved using capacitors on the order of 10 fF.

Since the DCO frequency is proportional to $1/\sqrt{C}$, even a linearly tuned capacitor array would result in nonlinear digital-to-frequency conversion. Using series capacitors to reduce the effective capacitance change across the inductor has the undesirable effect of adding further nonlinearity. Fortunately, capacitor bank predistortion can be used to mitigate nonlinearity. Choosing proper predistortion was done heuristically through simulation in [12], but the design cycle can be shortened by solving for the optimum digital-to-capacitance curves mathematically. For example, to find the proper values of C_F, the other capacitor banks are set to their midrange values and the relationship between $C_F(N_F)$ and $f_{\mathrm{DCO}}(N_F)$ is found.

To achieve a perfectly linear digital-to-frequency relationship between N_F and f_{DCO}, the 256 desired frequencies are used to find the respective values of C_F. Implementing 256 unique incremental values of capacitance for C_F, however, would be impractical. Instead, a very good piece-wise linear approximation can be implemented using a thermometer-coded capacitor bank like the one shown in Fig. 4(b) where the capacitor bank is divided into 32 identical columns (5 bits) composed of 8 progressively larger capacitors (3 bits). Fig. 5 shows simulations of the frequency step sizes versus digital word N_F. Without predistortion (i.e., using linear C_F), the frequency steps are small for low values of N_F and much larger for high values of N_F. This means that extra bits would be required to achieve a desired frequency resolution while covering the desired frequency range. If a perfectly predistorted capacitor bank were used, the frequency steps would all be equal, but the implementation would be impractical. As a compromise, a piece-wise linear (PWL) approximation, which is much easier to implement, is used to bound the frequency step sizes well within one least significant bit.

By building a mathematical model in MATLAB, the values for all four capacitor banks could be easily recalculated to include the effects of parasitics, resulting in fewer iterations of lengthy SPICE simulations.

D. Setting f_C and ΔF

The carrier frequency and frequency deviation constant are set in two steps. First, $N_{\Delta F}$ is set to zero and the DCO frequency is calibrated to $f_{\mathrm{DCO}} = f_c - \Delta F$ using N_C, N_M, and N_F. Then the DCO frequency is set to $f_{\mathrm{DCO}} = f_c + \Delta F$ using $N_{\Delta F}$ and the value is stored in a register. To perform FSK modulation, each bit of $N_{\Delta F}$ is applied to one input of an *and* gate,

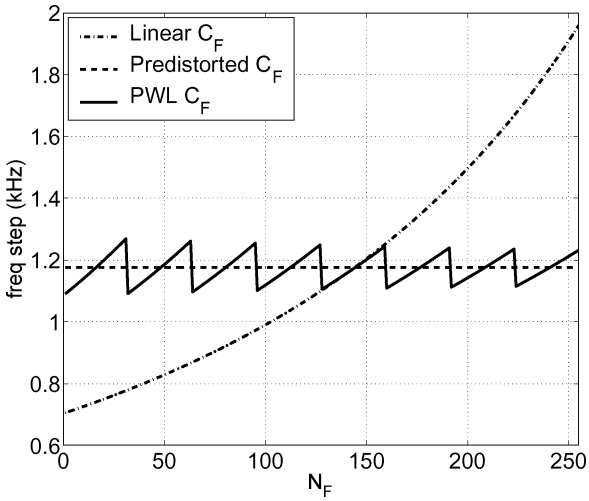

Fig. 5. Frequency steps versus N_F using linear, predistorted, and piece-wise linear (PWL) predistorted capacitor banks.

and unipolar nonreturn-to-zero modulation data $m_u(t) \in [0,1]$ is applied to the other input as shown in Fig. 1. The resulting instantaneous DCO frequency is

$$f_{\rm DCO}(t) = f_C - \Delta F + 2\Delta F \times m_u(t) \in [f_C \pm \Delta F] \quad (4)$$

which is equivalent to (2).

IV. Super-Regenerative Receiver

Using the loop antenna as the inductive element in the DCO has the benefits of reducing system complexity and power consumption for the transmitter. A challenge that arises, however, is that classical homodyne or superheterodyne receiver architectures cannot be used. In this section, the super-regenerative receiver (SRR) is introduced as an excellent low-power alternative that allows a direct connection between the antenna and the oscillator.

A. Super-Regeneration Theory

Super-regeneration was first introduced by Edwin Armstrong in 1922 and has had sporadic popularity in low power systems since then gaining renewed interest in recent years [13]–[19]. The SRR in its simplest form comprises a resonator, time-varying positive feedback, and an input current. Fig. 6(a) shows a simplified SRR circuit with (b) its corresponding feedback loop model. The combination of the Laplace variable s and time variable t is acceptable only because the rate at which $G_m(t)$ is varied is much slower than ω_0, the resonant frequency of Z_{RLC} (i.e., the system is quasi-static). The current $i_a(t)$ is induced by the antenna and serves as the input to the system. Using this model, the time-varying transfer function of the system can be written as

$$Z_{TV}(s,t) = \frac{V_o(s,t)}{I_a(s)} = \frac{Z_0\omega_0 s}{s^2 + 2\zeta(t)\omega_0 s + \omega_0^2} \quad (5)$$

where $Z_0 = \sqrt{L/C}$ is the characteristic impedance of the resonant tank, $\zeta_0 = 1/2Q_0$ is the quiescent damping factor, and

$$\zeta(t) = \zeta_0\,(1 - G_m(t)R) = -\beta t \quad (6)$$

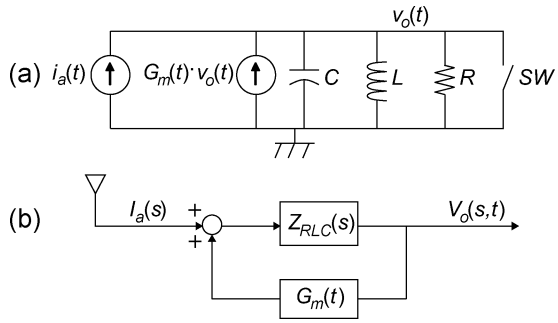

Fig. 6. (a) SRR circuit model, and (b) SRR feedback loop model.

is the damping function. This transfer function shows that the SRR is a second order system with time-varying poles. By changing $\zeta(t)$, the poles are periodically shifted from the left-hand-side of the complex plane to the right-hand-side, forcing the system to become temporarily unstable. The resulting output is a growing oscillation whose amplitude depends on the characteristics of the input current $i_a(t)$.

A thorough and general solution to the differential equation describing this system can be found in [16]. For the specific case of the ramp damping function described by (6), and a sinusoidal input

$$i_a(t) = I_a \sin(\omega_a t + \phi_a) \quad (7)$$

the output voltage during each cycle is

$$v_o(t) = I_a Z_0 \sqrt{\frac{\pi}{2}}\frac{\omega_a}{\Omega_s} \cdot e^{-\frac{(\omega_0-\omega_a)^2}{2\Omega_s^2}} \cdot e^{\frac{t^2}{2\sigma_s^2}} \sin(\omega_0 t + \phi_a). \quad (8)$$

The output voltage is a quickly growing oscillation that is proportional to the amplitude of the input signal I_a. The two exponential terms describe the filtering and time-dependent gain of the SRR. They are functions of Ω_s and σ_s, the SRR frequency and time constants, defined as

$$\Omega_s = \frac{1}{\sigma_s} = \sqrt{\omega_0\beta}. \quad (9)$$

The first exponential in (8) describes the Gaussian filtering quality of SRRs using a ramp damping function. The second exponential in (8) describes the tremendous time-dependent gain provided by SRRs.

Equation (8) is found using the assumption that the ramp damping function extends from $-\infty$ to ∞. In reality, the ramp only extends from some time t_a to t_b and the process starts all over again resulting in a saw-tooth damping function. As long as $t_a \ll -\sigma_s$ and $t_b \gg \sigma_s$, however, (8) is accurate because system is only sensitive to input currents during a narrow time window centered about the instant when $\zeta(t) = 0$. The rate at which this cycle is repeated is often referred to as the *quench frequency* because the output signal is usually quenched at the end of each cycle so that it does not affect the following cycle. The implementation presented here includes a CMOS switch placed across the resonator used to quench all oscillations at the end of every cycle [SW in Fig. 6(a)].

B. Receiver System Design

To maximize the bit rate of the receiver, the SRR is used to synchronously receive OOK data as in [17]. To ensure proper operation, the quench oscillator and the incoming data must be synchronized so that the instant when $\zeta(t) = 0$ occurs near the center of each bit period. As described in Section VI, the base-station has the task of synchronizing its baseband clock to the implant's baseband oscillator to reduce complexity in the implant.

As shown in Fig. 1, the SRR is followed by a fully differential envelope detector, a programmable comparator, a digital counter, and a digital comparator. At the start of each bit period (or quench cycle), the digital counter is reset and immediately begins to count. Receiving a *one* is equivalent to receiving a sinusoidal input with an oscillation frequency that is equal to the SRR's resonant frequency (i.e., $\omega_a = \omega_0$). Passing the output of the SRR through an envelope detector results in the envelope signal

$$v_e(t) = I_a Z_0 \sqrt{\frac{\pi}{2}} \frac{\omega_0}{\Omega_s} \cdot e^{\frac{t^2}{2\sigma_s^2}} \tag{10}$$

where I_a is the peak amplitude of the input current. The output of the envelope detector is connected to a comparator that has a programmable offset $V_{\rm OS}$. When the amplitude of the envelope becomes larger than $V_{\rm OS}$, its output changes states and disables the counter causing it to hold its value. This final value is a measure of the SRR's startup time (the time required for the SRR's envelope to reach $V_{\rm OS}$). For the receiver to have a low bit-error rate (BER), the amplitude of the current I_a must be significantly larger than the RMS value of equivalent input referred noise sources. This means that when a *one* is received, the SRR's startup time is faster and the counter's final count is lower. When a *zero* is received, the startup time is slower resulting in a higher count. In this manner, the counter functions as a time-integrating analog-to-digital converter that measures the startup time of the SRR which is a function of the input signal's amplitude. At the end of the bit period, the final count is compared with a digital threshold number NT, and a decision is made as to whether a *one* or a *zero* was received. NT is set by averaging the counter output over a 32-bit preamble with an equal number of *ones* and *zeros*.

Note that the counter was implemented in an FPGA and is clocked by an external oscillator, but could easily be integrated and clocked using an amplified version of the SRR's output. In such an implementation, the counter would start counting once the SRR's output envelope reached a threshold and the final count would be higher when a *one* was received. The performance of the receiver would be the same since the important parameter is the difference between the counter's output when a *one* is received versus a *zero*. Also, the power consumption of the digital blocks would be minimal due to their simplicity.

When a *zero* is received, the output of the SRR is also a growing oscillation similar to (8) except its amplitude is a random variable with an RMS value that depends on the noise density of active and passive noise sources in the system along with their effective noise bandwidths. The sensitivity of the receiver, therefore, is optimized by reducing the amplitude

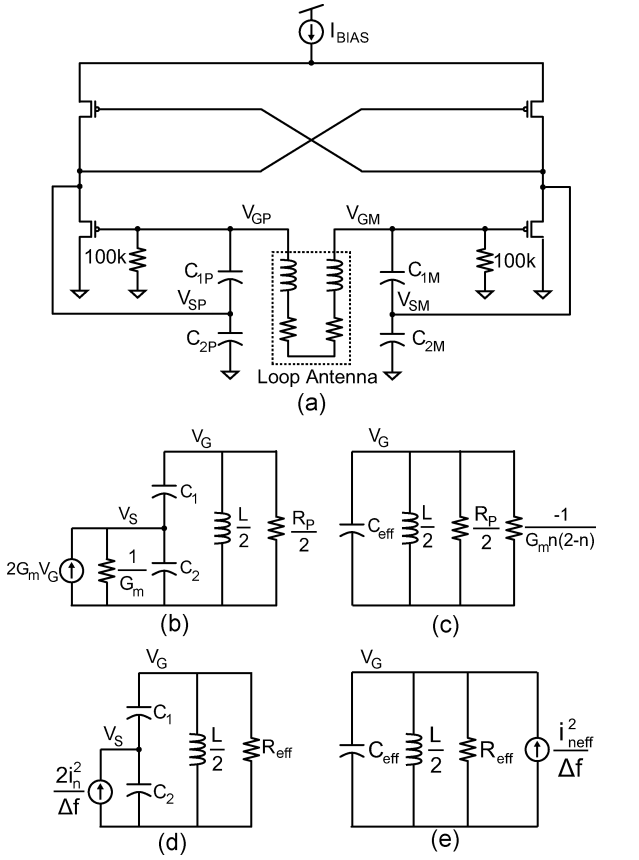

Fig. 7. (a) Differential Colpitts oscillator, (b) half circuit model, (c) equivalent half circuit, (d) noise model, and (e) equivalent noise model.

of the noise sources and the bandwidth of the SRR which is proportional to Ω_s. Since ω_0 is determined by the standard used (i.e., $\omega_0 \approx 2\pi \times 400$ MHz for MICS), the only design variable that can be altered to reduce the noise bandwidth is β (the slope of the damping function). The tradeoff, however, is that reducing Ω_s is equivalent to increasing σ_s which necessitates longer quench periods since the condition $|t_{a,b}| \ll \sigma_s$ must be kept. This results in a tradeoff between sensitivity and bit rate.

A concern that arises from having the SRR directly connected to the antenna is that the receiver actually radiates power. In architectures such as super-heterodyne receivers this is undesirable since the radiated signals can interfere with other receivers. For an SRR, however, the radiated power occupies its own channel and does not interfere with other receivers. Furthermore, the radiated power is well below the 25 μW allowed for transmission and does not violate spectral mask requirements.

V. Circuit Implementation

The full transceiver is shown in Fig. 1, where a portion of the baseband section was implemented in an FPGA for maximum flexibility. The receiver comprises five major blocks: a DCO/SRR, a quench/baseband oscillator, a fully differential envelope detector, a comparator with programmable offset, and a digital counter. Exploiting the high gain provided by the SRR allows all analog blocks to be biased in subthreshold for low-power operation.

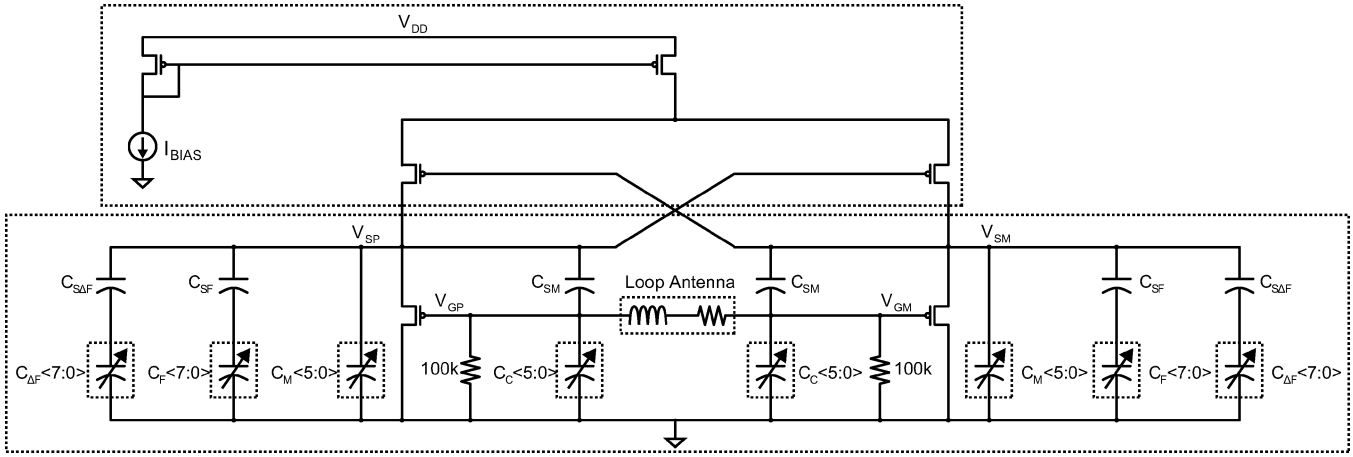

Fig. 8. Differential Colpitts digitally-controlled oscillator with predistorted sub-ranging capacitor banks and loop antenna.

A. DCO Implementation

Fig. 7(a) shows a simplified schematic of the differential Colpitts oscillator with switching current source similar to [20]. The loop antenna is shown as two inductors and resistors to facilitate understanding. Fig. 7(b) shows the equivalent small-signal model of the half-circuit. The transconductance is twice G_m because two pMOS transistors provide positive feedback in each half-circuit. Fig. 7(c) shows a simplified version of the model where the current source and $1/G_m$ resistor are placed in parallel with the antenna using capacitive impedance transformation [21]. The minimum value of G_m that will result in oscillation is

$$G_{m0} = \frac{2}{R_p n(2-n)} \tag{11}$$

where n is the impedance transformation ratio

$$n = \frac{C_1}{C_1 + C_2}. \tag{12}$$

As shown in (6), the damping function $\zeta(t)$ is controlled by varying $G_m(t)$, and $\zeta(t) = 0$ occurs at the instant when $G_m(t) = G_{m0}$. Using a Colpitts topology requires a larger G_{m0} than the more common cross-coupled transistor topologies, but offers a noise benefit. Fig. 7(d) shows the equivalent combined noise source from the two pMOS transistors in each half-circuit, and Fig. 7(e) shows the effective noise source in parallel with the resonator having a value of

$$\frac{i^2_{\text{neff}}}{\Delta f} = \frac{2n^2 i^2_n}{\Delta f}. \tag{13}$$

Since $n < 1$, the impedance transformation reduces the impact of transistor noise sources. The effective noise can be derived as

$$\frac{i^2_{\text{neff}}}{\Delta f} = \frac{4kT}{R_p}\left(\frac{n}{\kappa(2-n)}\right) \tag{14}$$

when the transistors are biased in subthreshold, where κ is a device parameter roughly equal to 0.7. By making n small, the effective noise from the transistors can be made much smaller than the noise from the passive components, improving the sensitivity of the receiver.

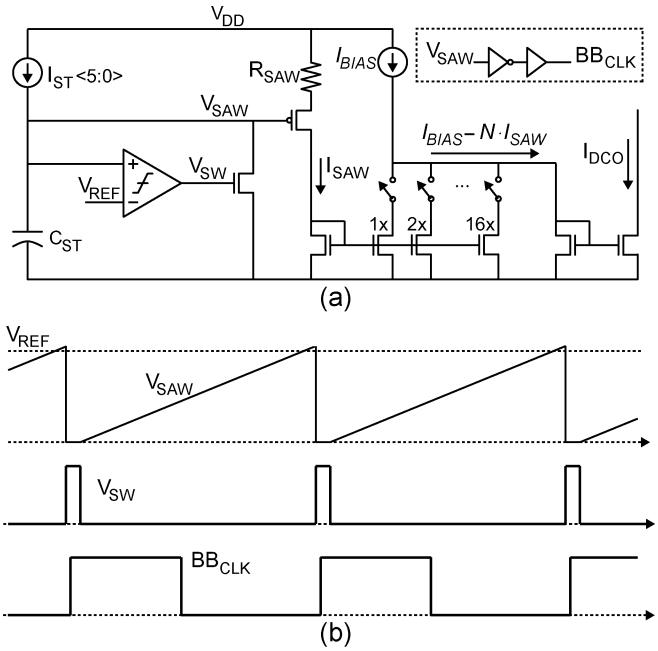

Fig. 9. (a) Sawtooth oscillator, SRR bias generator, and baseband clock. (b) Output signals.

The differential Colpitts oscillator also has some power transfer advantages for the transmitter. By using a tapped resonator, the peak voltage across the antenna is not restricted to the power supply. This means that a lower supply voltage can be used while still delivering sufficient power to the antenna. As shown in [20], the cross-coupled transistors perform current switching once the signal amplitude is large, resulting in twice as much efficiency delivering power to the antenna.

The full DCO with predistorted capacitor banks is shown in Fig. 8. By dividing the total capacitance into coarse, medium, and fine tuning capacitor banks, more than 14 bits of frequency resolution are achieved. Capacitor bank predistortion leads to improved digital-to-frequency conversion, and thermometer coding guarantees monotonicity for each bank. The differential Colpitts topology results in improved sensitivity for the receiver and higher power transfer capabilities for the transmitter.

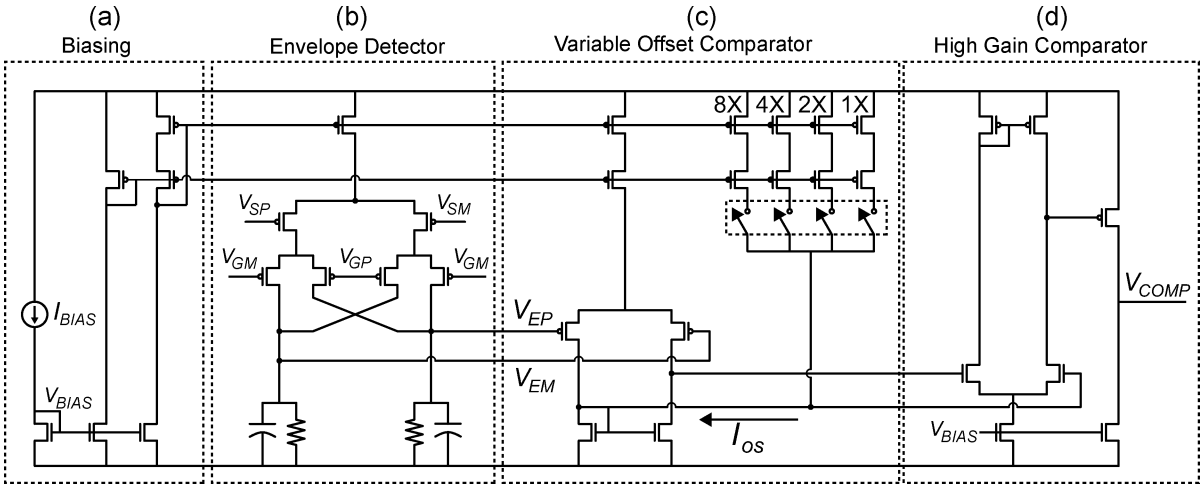

Fig. 10. Schematic for differential envelope detector and comparator with programmable input offset voltage. (a) Biasing; (b) envelope detector; (c) variable offset comparator; (d) high gain comparator.

B. Programmable Ramp Quench Oscillator

Fig. 9(a) shows a simplified schematic of the sawtooth oscillator used to generate the quench signal for the SRR and the baseband clock. The sawtooth voltage waveform $V_{\rm SAW}$ results from capacitor $C_{\rm ST}$ integrating the current $I_{\rm ST}$ over time. Once $V_{\rm SAW}$ exceeds $V_{\rm REF}$, a comparator turns on an nMOS switch that discharges the capacitor and forces $V_{\rm SAW}$ back to zero. This in turn causes the comparator to turn off the nMOS switch and restart the cycle. The comparator includes some delay stages that cause $V_{\rm SAW}$ to stay low and $V_{\rm SW}$ to stay high for a short period of time as shown in Fig. 9(b). The signal $V_{\rm SW}$ is used to drive the switch that briefly shorts the resonant tank of the DCO as shown in Fig. 6(a). $V_{\rm SAW}$ is passed through a pair of inverters to create a square wave with a duty cycle of roughly 50% that is used to clock the baseband circuits. The frequency of oscillation is

$$f_{\rm BB} \approx \frac{I_{\rm ST}}{V_{\rm REF} C_{\rm ST}} \tag{15}$$

and is tuned digitally using the programmable current $I_{\rm ST}$. The voltage waveform $V_{\rm SAW}$ is converted to a current $I_{\rm SAW}$ using a degenerated common-source pMOS transconductor. $I_{\rm SAW}$ is then used as the reference for a multiplying current-mode digital-to-analog converter (IDAC) which allows for digital tuning of the quench signal's slope. The output of the IDAC is subtracted from the bias current $I_{\rm BIAS}$ and mirrored to create the bias current of the SRR. In transmit mode, the IDAC is disabled, and the DCO bias current is $I_{\rm BIAS}$.

C. Envelope Detector

The task of the envelope detector is to create a one-to-one mapping between the SRR's output amplitude and the envelope detector's output voltage. This can be done by squaring the output of the SRR (using a Gilbert cell) and taking the average (using a low-pass filter). The Gilbert cell shown in Fig. 10(b) is biased in subthreshold at a very low bias current (10 μA), yet its input-referred noise has a negligible effect on the receiver's performance due to the high gain provided by the SRR. It is DC-coupled to the DCO/SRR in Fig. 8 and uses 60 kΩ load resistors resulting in a common-mode output voltage of 300 mV. A pair of 1 pF capacitors are used to filter the $2\omega_0$ term.

D. Programmable Comparator

The programmable comparator following the envelope detector in Fig. 1 has a digitally configurable offset that sets the differential input voltage threshold at which the comparator switches states. The offset is set by an IDAC and allows a fully differential connection between the envelope detector and the comparator, reducing the effects of common-mode noise. The programmable comparator comprises a low gain, variable offset comparator [Fig. 10(c)] and a two-stage op-amp used as a high gain comparator [Fig. 10(d)].

The input offset is created by feeding current into the low-impedance arm of the current mirror in (c). Since all of the circuits in Fig. 10 are biased in subthreshold, the transconductance of the input differential pair is

$$I_{\rm E} = I_{\rm EP} - I_{\rm EM} = I_{\rm DD} \tanh\left(\frac{\kappa V_{\rm E}}{2\phi_t}\right). \tag{16}$$

$I_{\rm DD}$ is the tail current of the differential pair and is set to $40 \times I_{\rm BIAS}$, and κ is a process parameter approximately equal to 0.66 for this technology. Adding the current

$$I_{\rm OS} = N_{\rm OS} \times I_{\rm BIAS} \tag{17}$$

to the low impedance arm of the current mirror, creates the input voltage offset

$$V_{\rm OS} = V_{\rm E} = \frac{2\phi_t}{\kappa} \tanh^{-1}\left(\frac{1}{40} N_{\rm OS}\right). \tag{18}$$

Since the maximum value of $N_{\rm OS}$ is 15, $V_{\rm OS}$ can be approximated at room temperature as

$$V_{\rm OS} \approx \frac{2\phi_t}{\kappa} \times \frac{N_{\rm OS}}{40} \approx 2.0\ {\rm mV} \times N_{\rm OS} \tag{19}$$

using the approximation $\tanh^{-1}(x) \approx x$. $V_{\rm OS}$ is proportional to absolute temperature and independent of $I_{\rm BIAS}$ as long as

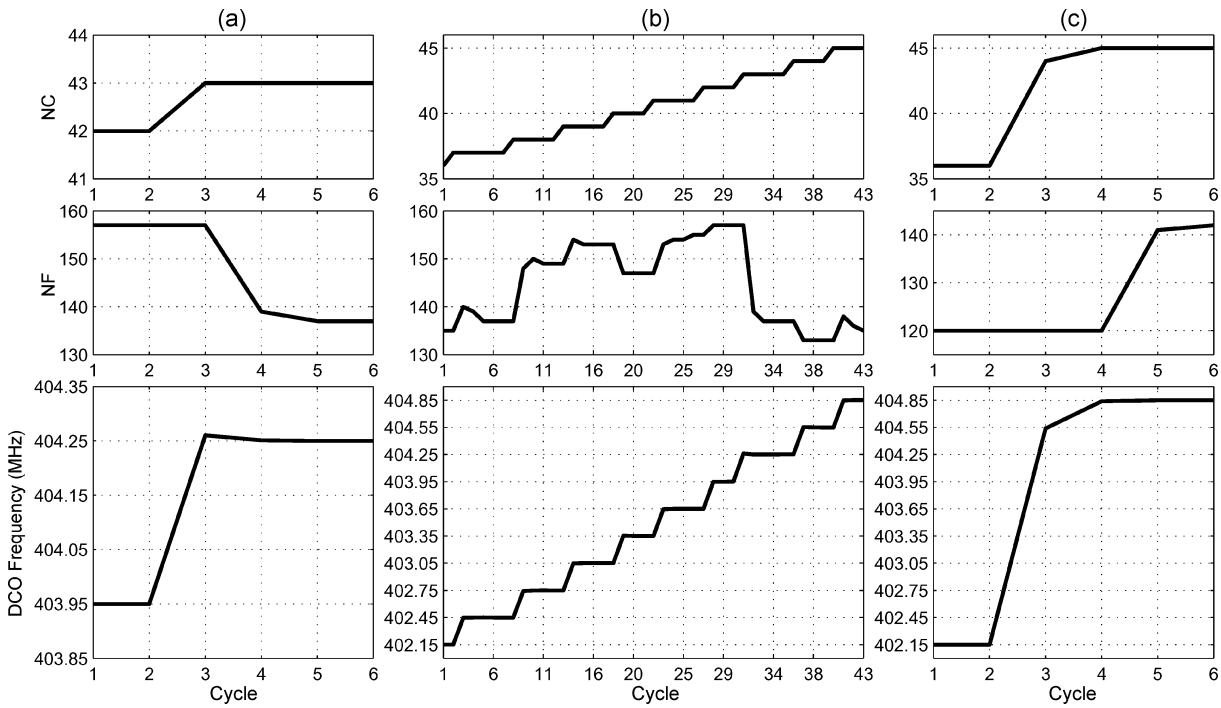

Fig. 11. Measurements of the frequency correction loop including the transition of the DCO frequency, coarse tuning coefficient, and fine tuning coefficient as the channel is changed from (a) 7 to 8, (b) 1 to 10 incrementally, and (c) 1 to 10 in a single command.

subthreshold operation is maintained. The comparator was designed to consume only 10 μW.

VI. Prototype Frequency Correction Loop

Modern transceivers typically use frequency synthesizers to phase-lock a voltage-controlled oscillator (VCO) to a very stable reference such as a crystal oscillator. For ultra-low power medical implants, however, a frequency synthesizer may not be necessary for two reasons: first, the frequency stability requirement for MICS ($\pm$100 ppm) is far more lax than typical systems and second, the human body provides excellent temperature regulation reducing both the amount of overall frequency drift and its rate of change. Instead of using a phase-lock loop, a frequency correction loop is proposed that pushes complexity from the implant to the base-station. In the proposed system, the base-station monitors the DCO's oscillation frequency and sends information to the implant to correct frequency errors. A crude implementation of such a system was prototyped using a spectrum analyzer, a PC, and an FPGA. The spectrum analyzer is used to measure the frequency error, and the information is sent to the PC. A simple MATLAB script determines what the new DCO settings (N values) should be and sends the information directly to the implant device using an FPGA. Fig. 11 shows the values of N_C and N_F for three scenarios (N_M did not change).

The first example shown in Fig. 11(a) uses the frequency calibration loop to jump from channel 7 to channel 8. The implant was initially set to 403.95 MHz and a command was sent to change the frequency to 404.25 MHz. First, the MATLAB script uses estimates of the coarse frequency step size to determine the correct value for N_C and the FPGA sends this information to the transceiver through a SPI bus. The spectrum analyzer determines the new DCO frequency and sends it to the PC. Since the new frequency error is small, a new value of N_F is sent to the transceiver and the cycle repeats until the frequency error is below a threshold of 1 kHz. After only five cycles, the final frequency is 404.2499 MHz (within 100 Hz of the desired frequency). Fig. 11(b) shows the progression of NC, NF, and the DCO frequency as the frequency control loop is used to change from one channel to the next. Finally, Fig. 11(c) shows the results from changing from channel 1 to channel 10 directly.

In all three of these examples, the frequency correction loop was used to correct large frequency errors and was able to do so within very few steps because of the fairly linear digital-to-frequency relationship of the DCO. Furthermore, because the digital-to-frequency relationship of the DCO is monotonic, convergence was always achieved over many measurements. Typical frequency corrections would actually be much smaller than the examples shown (on the order of 1–50 kHz), and even with this simple implementation, such corrections would require three cycles or less.

Calibrating the baseband oscillator is done similarly so that the SRR can synchronously demodulate OOK data. To do so, the implant transmits a short preamble prior to receiving data and the base-station uses it to lock to the implant. If the baseband clock frequency drifts beyond a limit, the transmitter sends a calibration command to correct it. Shifting the task of synchronization to the base-station results in power savings in the implant.

As explained in Section II-B, some frequency drift can result from human motion and correcting it requires that calibration be repeated every 100 ms. At a bit rate of 120 kb/s and assuming

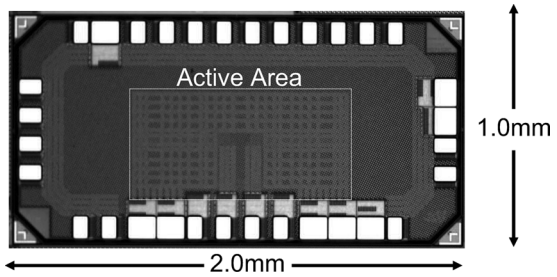

Fig. 12. Die photograph (1.0 × 0.5 mm active area).

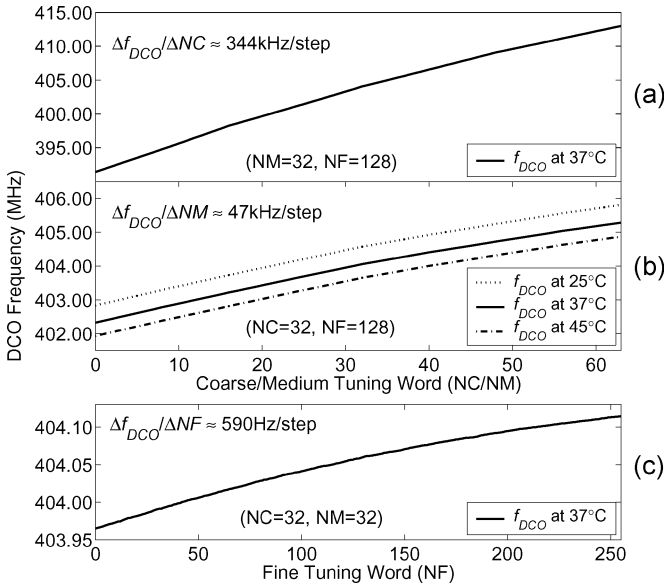

Fig. 13. DCO frequency for (a) coarse, (b) medium, and (c) fine tuning. Thermometer coding and predistortion of each capacitor bank lead to monotonic, linear tuning for each range.

128 bits are needed for each calibration cycle, the energy overhead is just over 1%. As a result, by taking advantage of the lax frequency stability requirements of MICS, the temperature regulation of the human body, and the robustness of the SRR to minor frequency mismatches, the frequency synthesizer can be completely replaced by a simple frequency correction loop that greatly reduces the power consumption of the implanted device.

VII. Measurement Results

The transceiver was fabricated in 90 nm CMOS and consumes a total of 0.5 mm^2 of active area as shown in Fig. 12. Fig. 13 shows the coarse, medium, and fine tuning digital-to-frequency curves. The downward concavity of the tuning curves is due to a slight overcompensation in the predistortion resulting from unaccounted parasitics. The DCO tunes from 391–415 MHz, and has an average frequency resolution of 590 Hz with a maximum fine-tuning step size of 1.4 kHz achieving 14 effective bits of resolution. Fig. 13(b) shows that C_M tunes the DCO 2.96 MHz which is well above the maximum coarse frequency step (450 kHz) providing ample overlap. The C_M tuning curve is plotted for the maximum and minimum temperature requirements and shows an average frequency shift of 923 kHz over the entire temperature range, corresponding to a temperature coefficient of 46.2 kHz/°C

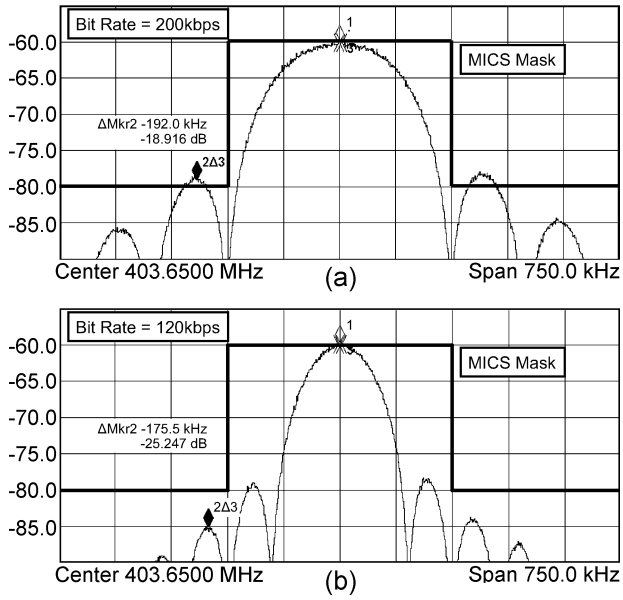

Fig. 14. Measured spectral mask of transmitter with a data rate of (a) 200 kb/s and (b) 120 kb/s. The signal is taken through an antenna 10 cm from the transmitter.

(115 ppm/°C). Since the frequency stability requirements of MICS call for frequency stability of ±100 ppm, the transceiver can tolerate temperature drifts of up to 0.87 °C without calibration. Correcting for such variations should be simple with a calibration rate of 10 calibrations per second. However, if the implant is in sleep mode for a long period of time, it may drift beyond the 100 ppm requirement. This may lead to a brief violation of the spectral mask and frequency accuracy requirements, but there is a very small possibility it would interfere with other implants due to the highly duty cycled nature of their operation.

A. Transmitter

Fig. 14(a) shows the output spectral mask of the transmitter measured through an antenna placed 10 cm away from the device. The first side lobe is roughly 6 dB higher than the theoretical value and violates the MICS spectral mask by 1.1 dB at a bit rate of 200 kb/s. This could easily be corrected with a very simple and low-power digital filter connected to the modulation capacitor bank $C_{\Delta F}$. At a bit rate of 120 kb/s and without pre-filtering, the MICS spectral mask requirement is met with 5 dB of headroom as shown in Fig. 14(b). Fig. 15 shows the power received at an antenna placed 20 cm from the transmitter and the DCO frequency as a function of bias current ($V_{\rm DD} = 700$ mV). As the DCO bias current is swept from 450 μA to 600 μA, the power at the receiving antenna varies from −65 dBm to −48 dBm with minimal variation in the DCO frequency. This ability to trade off power consumption for output power, means that as little as 315 μW of power can be consumed by the DCO when the path loss between the antennas is small. If the path loss is high, the base-station can send a command to the implant to increase the output power up to 17 dB without considerable frequency drift. Furthermore, the ability to control both the frequency and amplitude of the DCO opens the possibility of employing more spectrally efficient modulation techniques

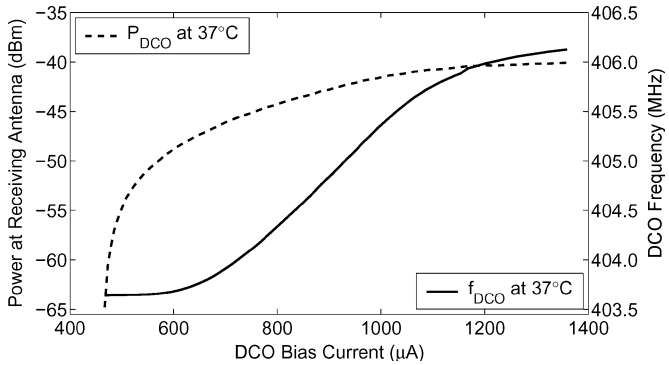

Fig. 15. Power received at an antenna placed 20 cm from the transmitter and DCO frequency versus bias current.

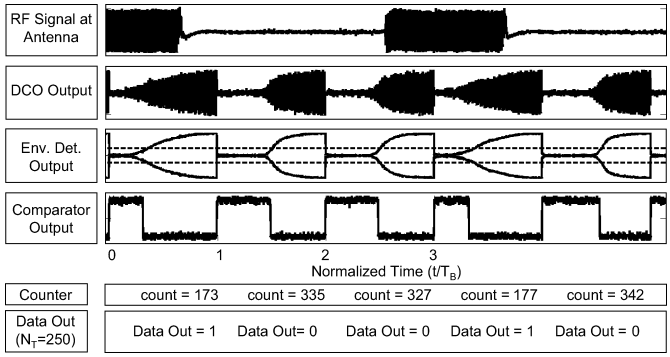

Fig. 16. Receiver chain time domain signals.

such as polar modulation. Including all circuits, the transmitter consumes less than 350 μW under normal operation.

B. Receiver

Fig. 16 shows measured time-domain signals for the receiver chain. As expected, when a *one* is received, the startup time of the DCO/SRR is faster. When the envelope crosses the threshold voltage of the comparator, the counter is disabled and the final count is compared to the threshold NT (250 in this case). If the final count is smaller than this threshold, the decision is made that a *one* was received and vice versa.

Since the transceiver was not designed to interface with a 50 Ω source or load, measuring sensitivity was challenging. To do so, a matching network was connected to each side of the resonator. At resonance, where the receiver works, the input impedance is real, so the matching network is meant to step up the impedance of a signal generator from 50 Ω to roughly 2 kΩ. Each matching network comprised a coupling capacitor, shunt capacitor, series inductor, and shunt capacitor connected to an SMA connector. The network had a negligible effect on the resonant frequency, but loaded the DCO requiring more bias current. A signal source was connected to a power splitter with two outputs that were each connected to one of the matching networks. The required DCO bias current increased from 450 μA to 650 μA due to the additional loading.

For bit rates of 40 kb/s and 120 kb/s, the measured sensitivity was -99 dBm and -93 dBm respectively (BER $= 0.1\%$), showing how bit rate can be traded off for sensitivity as discussed in Section IV. Without the matching networks, the power consumption in both cases was under 400 μW for the entire system resulting in an excellent energy-per-received-bit figure-of-merit 3.3 nJ/bit for a bit rate of 120 kb/s.

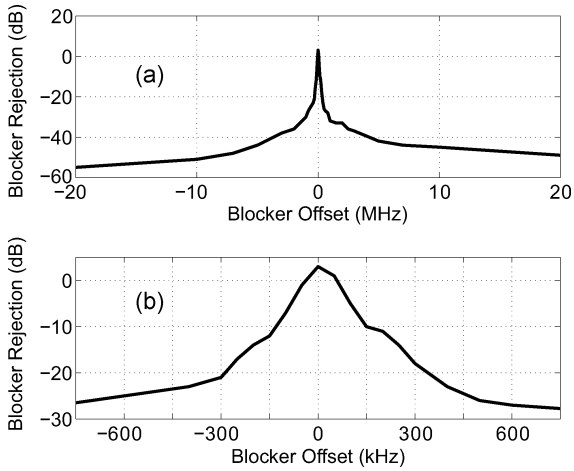

Fig. 17. (a) Far off and (b) close in CW blocker rejection. Input signal level set to sensitivity $+6$ dB $= -93$ dBm with a carrier frequency of 403.35 MHz and a bit rate of 40 kb/s.

TABLE I
SUMMARY OF MEASURED TRANSCEIVER PERFORMANCE

RF Frequency Range	391–415MHz
Frequency Resolution	$\leq 1.4kHz$
Rx Modulation	Sync. OOK
Rx Sensitivity @ 40kbps	–99dBm
Rx Sensitivity @ 120kbps	–93dBm
Rx Power Consumption	400μW
Rx Energy per Bit @ 40kbps	3.3nJ/bit
Rx Energy per Bit @ 120kbps	10nJ/bit
Tx Modulation	MSK
Tx Bit Rate	120kbps
Tx Power Consumption	350μW
Tx Energy per Bit	2.9nJ/bit

Fig. 17 shows the receiver's immunity to far-off and close-in blockers. For this measurement, the bit rate was set to 40 kb/s, the input signal to the SRR was set to sensitivity $+6$ dB (-93 dBm), and a the power of a CW blocker was swept for each frequency until the BER was degraded to 0.1%. The receiver shows excellent selectivity, with the ability to reject far off blockers by almost 60 dB and close in blockers by at least 10 dB. For example, at the center frequency of the adjacent channel (300 kHz offset from the desired channel), the receiver achieves more than 20 dB of rejection and at least 27 dB of rejection for the second adjacent channel (600 kHz away). Table I summarizes the transceiver performance.

VIII. CONCLUSION

A transceiver optimized for the MICS standard was presented that exploits unique features of medical implants to simplify its architecture and reduce power consumption. The central block in the transceiver is a digitally-controlled oscillator that uses

the loop antenna as its inductive element and a predistorted and sub-ranged capacitor array to improve linearity in digital-to-frequency conversion. The DCO has a differential Colpitts topology that improves power transfer to the antenna and reduces the effects of active device noise sources. It is used as a super-regenerative receiver with excellent sensitivity and selectivity performance for minimal power consumption. The transmitter uses direct MSK modulation to meet the MICS spectral mask requirements with relatively high data rates for such a simple topology. The transmitter and receiver achieve an energy-per-received-bit of 2.9 nJ/bit and 3.3 nJ/bit respectively while meeting the 300 kHz channel bandwidth requirements of MICS. Finally, a frequency correction loop was presented that eliminates the need for a frequency synthesizer in the implant, further reducing power consumption, device size, and system cost.

REFERENCES

[1] H. S. Savci, A. Sula, Z. Wang, N. Dogan, and E. Arvas, "MICS transceivers: Regulatory standards and applications [medical implant communications service]," in *Proc. SoutheastCon*, Apr. 2005, pp. 179–182.

[2] MICS Band Plan Federal Commun. Comm., Part 95, FCC Rules and Regulations, Jan. 2003.

[3] J. L. Bohorquez, J. L. Dawson, and A. P. Chandrakasan, "A 350 μW CMOS MSK transmitter and 400 μW OOK super-regenerative receiver for medical implant communications," in *Symp. VLSI Circuits Dig.*, Honolulu, HI, Jun. 2008, pp. 32–33.

[4] V. Karam, P. H. R. Popplewell, A. Shamim, J. Rogers, and C. Plett, "A 6.3 GHz BFSK transmitter with on-chip antenna for self-powered medical sensor applications," in *Proc. IEEE Radio Frequency Integrated Circuits Symp.*, Honolulu, HI, Jun. 2007, pp. 101–104.

[5] A. Wang, S. Cho, C. Sodini, and A. P. Chandrakasan, "Energy efficient modulation and MAC for asymmetric RF microsensor systems," in *Proc. ISLPED*, Huntington Beach, CA, Aug. 2001, pp. 106–111.

[6] A. H. Johansson, "Performance of a radio link between a base station and a medical implant utilizing the MICS standard," in *Proc. IEMBS*, Sep. 2004, pp. 2113–2116.

[7] D. B. Miron, *Small Antenna Design*, 1st ed. Oxford, UK: Elsevier Science and Technology, 2006.

[8] C. Angeloni, P. O. Riley, and D. E. Krebs, "Frequency content of whole body gait kinematic data," *IEEE Trans. Rehabil. Eng.*, vol. 2, no. 1, pp. 40–46, Mar. 1994.

[9] J. Kim and Y. Rahmat-Samii, "Implantable antennas inside a human body: Simulations, designs, and characterizations," *IEEE Trans. Microw. Theory Tech.*, vol. 52, no. 8, pp. 1934–1943, Aug. 2004.

[10] L. W. Couch, II, *Digital and Analog Communications Systems*, 6th ed. Upper Saddle River, NJ: Prentice-Hall, 2000.

[11] N. M. Pletcher and J. M. Rabaey, "A 100 μW, 1.9 GHz oscillator with fully digital frequency tuning," in *Proc. ESSCIRC*, Grenoble, France, Sep. 2005, pp. 387–390.

[12] J. Lin, "A low-phase-noise 0.004-ppm/step DCXO with guaranteed monotonicity in the 90-nm CMOS process," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2726–2734, Dec. 2005.

[13] E. H. Armstrong, "Some recent developments of regenerative circuits," *Proc. Inst. Radio Eng.*, vol. 10, pp. 244–260, Aug. 1922.

[14] J. Y. Chen, M. P. Flynn, and J. P. Hayes, "A fully integrated auto-calibrated super-regenerative receiver in 0.13-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1976–1985, Sep. 2007.

[15] P. Favre, N. Joehl, A. Vouilloz, P. Deval, C. Dehollain, and M. J. Declercq, "A 2-V 600-μA 1-GHz BiCMOS super-regenerative receiver for ISM applications," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2186–2196, Dec. 1998.

[16] F. X. Moncunill-Geniz, P. Palá-Schönwälder, and O. Mas-Casals, "A generic approach to the theory of superregenerative reception," *IEEE Trans. Circuits Syst. I*, vol. 52, no. 1, pp. 54–70, Jan. 2005.

[17] F. X. Moncunill-Geniz, P. Palá-Schönwälder, C. Dehollain, N. Joehl, and M. Declercq, "An 11-Mb/s 2.1 mW synchronous superregenerative receiver at 2.4 GHz," *IEEE Trans. Microw. Theory Tech.*, vol. 55, no. 6, pp. 1355–1362, Jun. 2007.

[18] B. Otis, Y. H. Chee, and J. Rabaey, "A 400 μW-RX, 1.6 mW-TX super-regenerative transceiver for wireless sensor networks," in *IEEE ISSCC Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 396–397, 606.

[19] A. Vouilloz, M. Declercq, and C. Dehollain, "A low-power CMOS super-regenerative receiver at 1 GHz," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, pp. 440–451, Mar. 2001.

[20] R. Aparicio and A. Hajimiri, "A noise-shifting differential Colpitts VCO," *IEEE J. Solid-State Circuits*, vol. 37, no. 12, pp. 1728–1736, Dec. 2002.

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

Jose L. Bohorquez (S'04) received the B.S. and M.S. degrees in electrical engineering from the University of Florida in 2002 and 2004, respectively. He then worked at the startup company BitWave Semiconductor designing analog and RF blocks for a reconfigurable transceiver. In 2006, he began doctoral studies at the Massachusetts Institute of Technology, where his research has focused on ultra-low-power systems for medical implants.

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 also 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 co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005).

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 2007 ISSCC Jack Kilby Award for Outstanding Student Paper. 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 Subcommittee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Subcommittee Chair for ISSCC 2004–2008. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on the SSCS AdCom from 2000 to 2007, and he was the meetings committee chair from 2004 to 2007. He is the Technology Directions Chair for ISSCC 2009.

Joel L. Dawson (M'97) received the S.B. degree in electrical engineering from MIT in 1996, and the MEng. degree from MIT in electrical engineering and computer science in 1997. He went on to pursue further graduate studies at Stanford University, where he received the Ph.D. degree in electrical engineering for his work on power amplifier linearization techniques.

He is currently an Assistant Professor in the Department of Electrical Engineering and Computer Science at MIT. Before joining the faculty at MIT in 2004, he spent one year at Aspendos Communications, a startup company that he cofounded. He continues to be active in the industry as both a technical and legal consultant.

Prof. Dawson received the National Science Foundation CAREER Award in 2008.