Alison Burdett

Ultra-Low-Power Wireless Systems

Energy-Efficient Radios for the Internet of Things

An ultra-low-power wireless system is one which has to operate for an extended period of time with only a limited power source available, and is typically constrained to a limited size (if size is not a limitation then a larger battery could be used). Clearly the term "ultra-low-power wireless" covers a broad range of applications which may have different key drivers as illustrated in Figure 1. For a fitness device such as a heart rate monitor, the number one driver is often cost as these are basic consumer products. For a bio-implant such as a smart pacemaker, battery life is critical as battery replacement typically requires surgery. For a smart home system such as automatic climate control, cost and battery life are important but the system also needs to support a relatively large number of devices and the communication range should be large—throughout the whole building. Finally for a gaming application such as a wireless headset, high data rate and very low latency are key, so as not to ruin the high-speed gaming experience.

An ultra-low-power wireless system comprises a number of functional blocks in addition to the wireless transceiver; Figure 2 shows a block diagram of a typical Internet of Things (IoT) "node." The energy source will be a battery or energy harvesting unit; the power management unit generates the required supply voltages and handles power sequencing of the various blocks; the sensors and analog front end generate and condition the environmental data which is processed typically by a microcontroller. In an ultra-low-power wireless system, all the blocks shown in Figure 2 must be carefully specified and designed to minimize the total system power, however it is often the

Digital Object Identifier 10.1109/MSSC.2015.2417095
Date of publication: 25 June 2015

1943-0582/15©2015IEEE

Application / Importance	Fitness	Bio-Implant	Smart Home	Gaming
Cost	High	Low	High	Medium
Battery Life	Medium	High	High	Low
Data Rate	Medium	Low/Medium	Low	High
Range	Low	Low	High	Medium
Latency	Low	Low/Medium	Medium	High
# of Nodes	Low	Low	High	Low

FIGURE 1: Ultra-low-power wireless application examples.

wireless transceiver (which is the focus of this article) which consumes the highest power when active of all the blocks.

Achieving a very low average power for a wireless system typically makes extensive use of duty cycling. The aim is to reduce the radio "on" time to a short communication burst, and then between these active periods have the radio enter a sleep mode when power consumption is minimized. To minimize radio "on" time, an obvious step is to maximize the data rate. Figure 3 shows the typical sequence required when a radio wakes up from sleep to active. First, the crystal oscillator needs to settle so that there is a stable frequency reference for the radio—this can take several milliseconds. Once the reference is stable, the RF synthesizer is turned on and that itself needs to settle, and may require calibration. Finally the main receiver can wake up and listen for communication, then we have a turn-around period as we switch from receive to transmit (or vice versa), and after the data exchange is completed there may be some final processing of received data or commands before the radio can re-enter its sleep state. If the average data rate is fairly low, then Figure 3 shows firstly that the radio on time may be dominated by wake up and settling periods as much as actual communication. Second, if the average data rate is fairly low and the transceiver is heavily duty cycled, then the average power will actually be dominated by the sleep current.

Energy Budget

Since an ultra-low-power wireless system is size constrained, Figure 4 shows the energy available from batteries with relatively small volumes. The first three columns are fairly well-known battery types—a lithium coin cell, and manganese dioxide and zinc button cells—followed by a lithium thin film "printed" battery and a supercapacitor. By far the most commonly used type today for ultra-low-power systems is the lithium coin cell. Although Zinc oxide batteries have a much higher energy density, once activated they have high leakage and so will self-discharge in just a few weeks, whereas the Lithium coin cell lifetime can be 10 years or more.

When selecting a battery, it is important to understand the real battery discharge characteristics so that the optimal load current can be understood. Take for example a CR2032, whose battery capacity is typically quoted as 240 mAh. This battery capacity is the energy

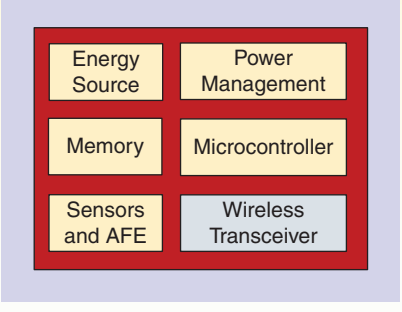

FIGURE 2: Block diagram of an ultra-low-power wireless system.

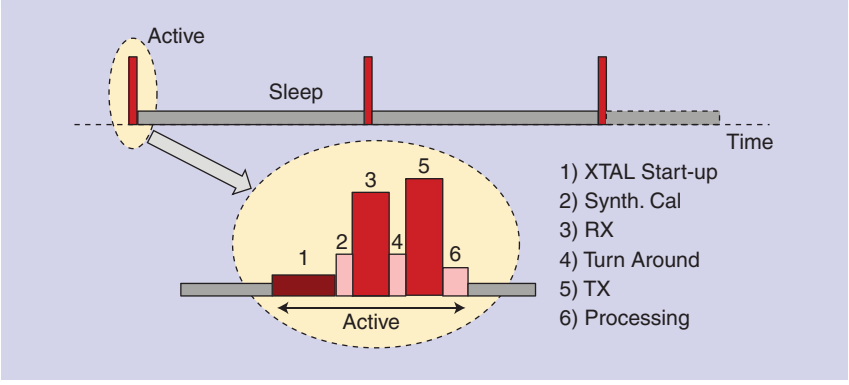

FIGURE 3: Duty cycling.

	CR2032 (Coin)	LR44 (Button)	PR44 (Button)	Thin Film	Super Capacitor
Technology	Li/MnO_2	MnO_2	ZnO_2	Li Thin Film	Supercap
Nominal Voltage	3.0 V	1.5 V	1.4 V	4.1 V	2.75 V
Capacity	240 mAh	153 mAh	620 mAh	2 mAh	0.6 mAh
Volume	1 cm^3	0.6 cm^3	0.5 cm^3	0.25 cm^3	0.5 cm^3
Peak Current	< 15 mA	< 10 mA	< 10 mA	< 5 mA	~30 A
Internal Resistance	10–40 Ω	1–5 Ω	2–8 Ω	~15 Ω	< 0.02 Ω
Self-Discharge	1%/yr (>10 Years)	2%/yr (5+ Years)	High (Few Weeks)	1%/yr (> 10 Years)	2 uA (Few Days)

FIGURE 4: Characteristics of small batteries.

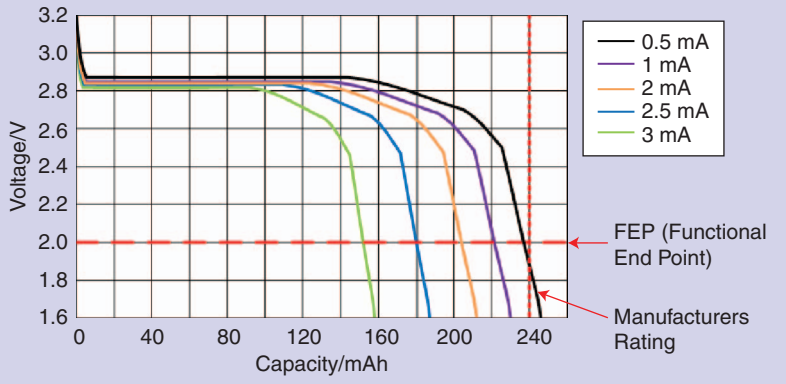

FIGURE 5: Constant current discharge curves for a CR2032.

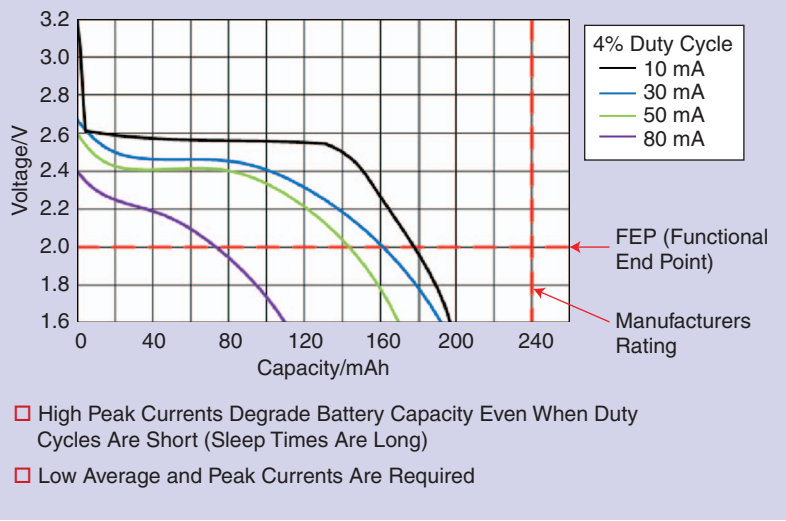

FIGURE 6: Pulsed current discharge curves for a CR2032.

capacity of the battery, when discharged to a defined functional endpoint, at a given constant discharge current. Figure 5 shows that for a CR2032 discharged down to a functional end point voltage of 2 V, the total charge delivered will be 240 mAh when the constant discharge current of 0.5 mA. However if the constant load current increases to 3 mA, the battery capacity is degraded to around 150 mAh. So clearly to extract the maximum charge from the battery the constant discharge current—typically the sleep current—should be minimized.

What about the effect of pulsed currents, since the goal is to operate in a duty-cycled mode? Figure 6 shows typical CR2032 discharge characteristics for pulsed currents at a 4% duty cycle; it can be seen that a peak pulsed current of about 10 mA will degrade the battery capacity significantly, and with 50 mA peak current the battery capacity has dropped to around half the 240 mAh rating. Clearly to maximize battery capacity, both average and peak currents should be minimized.

Operating Frequency versus Range

In order to understand the optimal carrier frequency for an ultra-low-power system, Figure 7 states Friis' equation for free space propagation, expressed in dB. This expression defines the ratio between the power detected at the receiver antenna to the power radiated from the transmitter antenna, and suggests that to double the transmit range (d) without increasing transmit power (Pt), then the RF frequency should be halved (i.e., the wavelength, λ will double). However it's important to note that this is only true if the antenna gains (Gt, Gr) remain constant. For a fixed antenna aperture—related to the antenna size—the gain will actually *decrease* as $(1/\lambda)^2$; since there are two antenna gains in Friis' equation, then if the antenna aperture is kept constant as the frequency is decreased, the range will actually decrease. Clearly, if the system is

size constrained, then the antenna dimension will be limited; assuming that antenna dimensions are constrained to the cm range, it can be shown that the low GHz region is a fairly "sweet spot" in terms of operating frequency [1].

As a practical example, assume transmit and receive antenna gains are –6 dB, operating frequency is 2.4 GHz, and line of sight range is 50 m. From Friis' equation, the difference between the transmitted power Pt and received power Pr is calculated as 86 dB, which is known as the link budget (since effectively it is the power lost over the wireless link between the two antennas). Assume first a receiver design that requires a received power Pr of at least –76 dBm for successful detection and demodulation (this value is 6 dB better than the minimum required by the Bluetooth low energy standard). Since the link budget is 86 dB, the transmitted power must be +10 dBm, which is 10 mW. Assuming the transmitter efficiency is significantly less than 100%, we can conclude that the peak power of our transmitter is going to be much greater than 10 mW—not really consistent with an ultra-low power wireless implementation.

So alternatively consider a receiver which is 10 dB more sensitive such that Pr = −86 dBm. Increasing sensitivity can usually be achieved without a significant overhead in power consumption, especially if high quality passive devices are available. The required transmitted power similarly reduces to 0 dBm, or 1 mW. Even with a modest transmit efficiency, peak transmitter power can now be below 10 mW, which is more in-line with the stated low peak power requirements.

In summary for a given range and link budget, the aim should be to maximize receiver sensitivity. This enables the transmitted power to be reduced, resulting in significant power savings during transmitter operation.

Friis' formula states how the operating frequency affects the achievable range, but in reality most ultra-low-power wireless systems have to operate within sections of the frequency spectrum allocated to short range devices. Devices operating within these bands have to follow certain rules defining maximum transmit power, bandwidth and channel spacing, and duty cycle. The aim of these rules is to ensure that as many users as possible can coexist within the same spectrum.

The actual spectrum allocated for short range wireless devices tends to vary with geography; as shown in Figure 8, most countries have allocation in the low hundreds of MHz, 300 or 400 MHz, and also allocation just below 1 GHz, around 800/900 MHz. 2.4 GHz is a worldwide band, and while this band is crowded, the simplicity of worldwide operation means that it is currently very popular for low power wireless implementation. There are also worldwide ISM bands at 5, 24 and 60 GHz—these bands are currently not targeted for ULP wireless because the high frequencies translate to high power consumption for the transceiver circuits, but it is likely that further developments in technology and techniques will change this. After all, even 10 years ago it was unthinkable that you could design a very low power wireless system operating at 2.4 GHz!

An ultra-low-power wireless system comprises a number of functional blocks in addition to the wireless transceiver.

Deep Asleep

A very low power wireless system which has a low duty cycle may spend the majority of the time asleep.

In this state, a major function which remains operational is the sleep timer, which counts elapsed time and produces an interrupt when

$$20 \log d = Gr(\text{dB}) + Gt(\text{dB}) + Pt(\text{dBm}) - Pr(\text{dBm}) + 20 \log \left(\frac{\lambda}{4\pi} \right)$$

FIGURE 7: Friis' transmission equation.

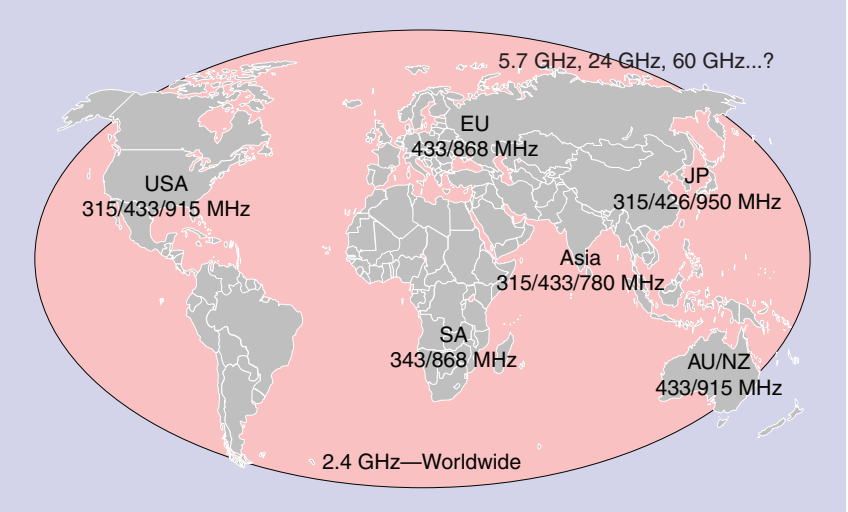

FIGURE 8: Allocated frequency bands for low power wireless.

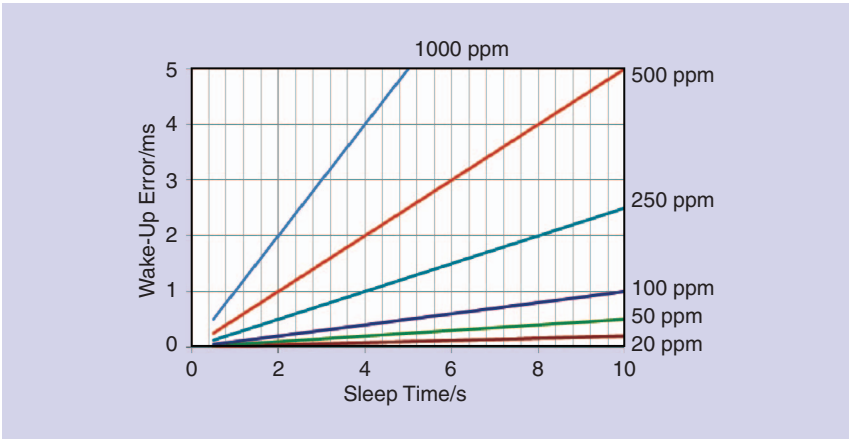

FIGURE 9: Sleep time error versus timer accuracy.

the rest of the system needs to wake up at the end of the sleep period. A sleep timer is simply a low-power oscillator and a counter, but the design of this low power oscillator is very important in ultra-low-power wireless systems; the important design parameter, in addition to nanopower consumption, is the accuracy of the oscillator reference.

Figure 9 illustrates sleep time (or counting period) against timing error, for different accuracy references ranging from 20 ppm to 1,000 ppm. The longer the sleep time, or the less accurate the reference, the larger will be the time error between when the device was supposed to wake up, and when it actually does wake up. Take the example of a sleep timer reference with a modest tolerance of around 500 ppm (allowing the use of a lower cost crystal reference). Over a modest 4-second sleep time, a 500 ppm reference error will result in a wake-up time error of ±2 ms. Either the device has woken up early, and has to waste power waiting around for the scheduled communication, or it has woken up late and may have missed the communication event.

If cost is not a major driver, then the solution is simply to use a very low tolerance, low frequency crystal reference, a typical watch crystal. It is possible to implement very low power crystal oscillator designs that consume just 10 s of nanowatts of power [2], [3]. In many applications the additional cost of an accurate low power crystal is not an option; instead, it may be preferable to implement a very low power on-chip sleep timer and calibrate this device against a more accurate frequency reference.

The major source of error in sleep timer oscillation frequency will be process variation; take the example in Figure 10 of a current starved ring oscillator whose frequency is determined by the switching delay per stage. This switching delay depends on the absolute value of the bias current as well as the transistor threshold voltages and the loading capacitance—all of which vary with process—and can lead to a large frequency offset. This frequency offset needs to be trimmed by initial calibration, i.e., by tuning the bias current until the desired frequency is obtained.

The sleep timer can be calibrated to correct for these frequency offsets due to process variation as shown in Figure 11. The calibration circuit has two counters, one which is clocked by the accurate reference oscillator and one clocked by the low power sleep oscillator. The counter values are set such that if the sleep oscillator is accurate, both counters

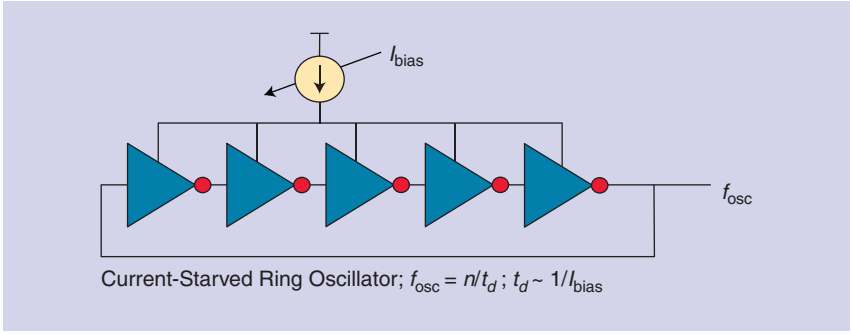

FIGURE 10: Current-starved ring oscillator.

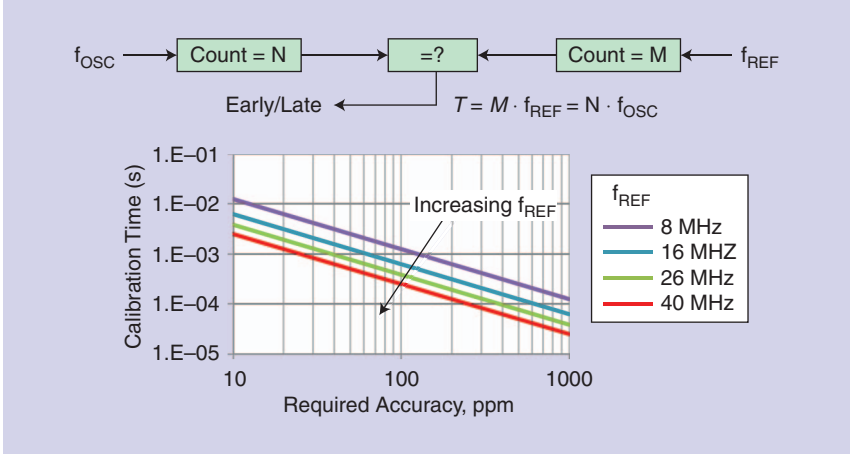

FIGURE 11: Sleep oscillator calibration.

should reach their target count at the same time. If the sleep counter terminates first then the sleep oscillator is running fast (early), while if the reference counter terminates first then the sleep oscillator is running slow (late). This early/late flag is used to trim the sleep oscillator, and then the whole process reruns until the sleep oscillator reaches its target value.

The sleep oscillator can be trimmed to within one period of the reference, over the total count time T. This gives a calibration accuracy which is the ratio of the reference oscillator period t_{REF} to the count period T. To increase the accuracy in ppm, it is necessary either to increase the reference frequency, or increase the calibration cycle time, or both. The resulting tuning accuracy is only achieved in relation to the inherent accuracy of the reference. So calibrating to within 50 ppm of a 50 ppm reference will give an overall sleep timer accuracy of ±100 ppm.

Each time a Calibration cycle is completed, the Sleep Timer frequency will be trimmed upwards or downwards until the target frequency is reached. How many bits of trimming are needed? This will depend on the initial frequency offset, and the final accuracy required. Assume we start with an initial tolerance of ±20%, which gives an error range of 400,000 ppm. The final target for accuracy is to calibrate the sleep oscillator to within 100 ppm of the reference. This requires a total of $\log 2(400{,}000/100) = 12$ bits of tuning range to vary the oscillator frequency—for example by trimming the bias current in the current-starved ring oscillator. The overall total calibration time will thus depend on the time T per calibration cycle, plus the number of tuning bits and the tuning algorithm selected.

Wireless System Architectures: The Receiver

One of the first widely adopted receiver architectures is the superheterodyne or dual-conversion approach, illustrated in Figure 12. In this implementation, a first variable local oscillator (LO) signal down-converts the wanted RF signal to a fixed first intermediate frequency (IF), and then a second down-conversion occurs using a fixed LO to a second IF. The dual conversion approach allows the first IF to be high, which relaxes the front-end image filter requirements, while the second IF is low, which simplifies channel selection. Although the superheterodyne approach enables both good sensitivity and selectivity, the penalty is a high power consumption since the architecture needs two separate synthesizers to generate the LO signals.

When selecting a battery, it is important to understand the real battery discharge characteristics so that the optimal load current can be understood.

Another popular receiver architecture is direct conversion, shown in Figure 13, where the LO frequency is set to equal the wanted RF channel, and so a single stage of frequency conversion shifts the wanted channel directly to 0 Hz baseband. Since there is only one stage of frequency conversion (mixing) and thus one frequency synthesizer, this seems attractive for low power implementation, and the first very low power CMOS transceivers used this direct conversion approach [4]. However there are also disadvantages associated with direct conversion. A major challenge is that direct frequency conversion requires quadrature LO

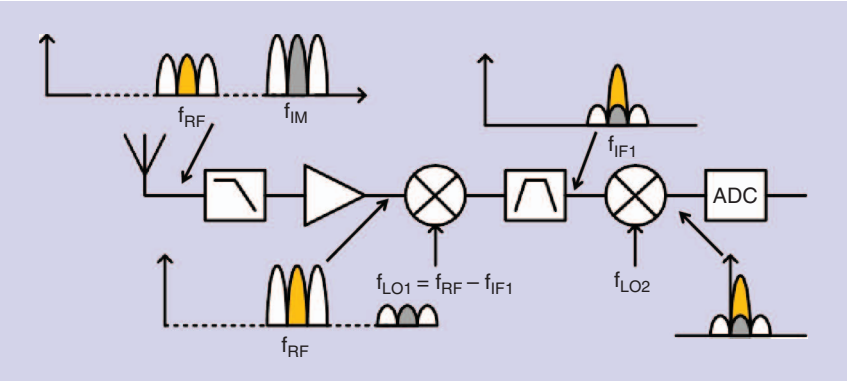

FIGURE 12: Superheterodyne receiver architecture.

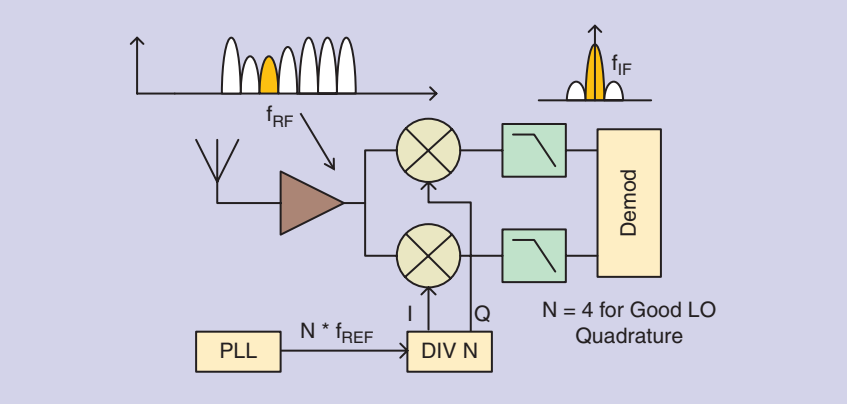

FIGURE 13: Direct conversion receiver architecture.

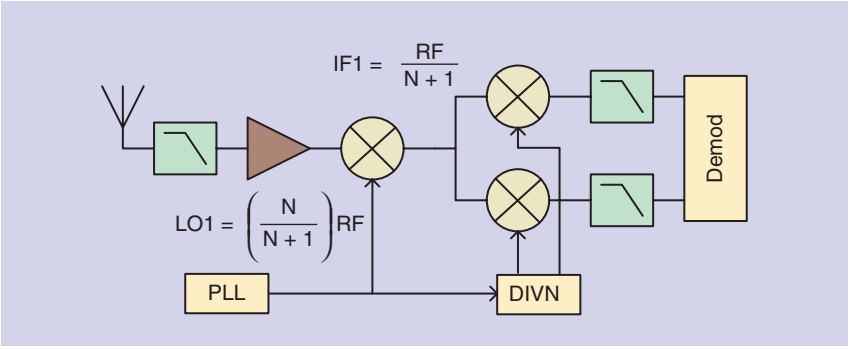

FIGURE 14: Sliding-IF receiver architecture.

signals, which are usually generated by running the frequency synthesizer at a higher frequency and then dividing down to get signals with a 90° phase shift. This requires a divide ratio of at least two and ideally four, which means the phase-locked loop (PLL) has to run at four times the RF frequency, with a corresponding high power penalty.

A receiver architecture suited for low power implementation is the "sliding IF" architecture, which inherits benefits from both the superheterodyne and direct conversion approaches. As shown in Figure 14, this is a dual conversion architecture, so again the image rejection is relaxed by selecting a suitably high first IF. The second mix brings the wanted signal down to a low or zero IF for channel selection and demodulation. This architecture differs from direct conversion or superheterodyne in the generation of the LO signals. The second LO is generated by an integer divide of the first LO, so there is no need for independent PLLs—essentially the second LO comes "for free." Furthermore the second frequency conversion is usually complex, but choosing the divide ratio N to be a multiple of two (or preferably four) allows the I and Q LO signals to be easily generated. For a given divide ratio N, the first LO will be $N/(N+1)$ times the RF frequency, or just below the wanted frequency—this is in contrast to direct conversion, where the PLL may have to generate two or four times the wanted RF in order to get good quadrature. The second LO is then $1/(N+1)$ times the RF, which thus shifts the wanted signal down to 0 Hz. As the wanted RF frequency varies, so does the IF, hence the name "sliding IF" since the position of the wanted signal at the first mixer output "slides around" as the RF channel varies. As a result of these benefits, the sliding IF architecture has recently become popular as an optimal receiver architecture for very low power wireless receiver implementation [5]–[9].

A receiver architecture that has recently emerged for ultra-low-power implementation is the super-regenerative receiver illustrated in Figure 15 (in practice this architecture was proposed in the early 1920s around the same time as the superheterodyne receiver, but the superheterodyne became dominant because of its superior selectivity and flexibility). At the heart of the super-regenerative receiver is an oscillator which is tuned to oscillate at the RF input frequency, and which is coupled to the RF signal received at the antenna. This "super-regenerative oscillator" (SRO) is also controlled by a second lower frequency quench oscillator which as its name suggests, periodically quenches or damps the SRO signal causing the oscillations to stop.

Achieving a very low average power for a wireless system typically make extensive use of duty cycling.

When the quench oscillator releases the SRO, the high frequency SRO oscillations begin to build up. If there is no RF input signal present, then the SRO has to start oscillating entirely on its own and the amplitude of oscillations will build up slowly and will not have reached a significant level before the quench signal damps the process again. However, if the SRO is released and an input RF signal at the SRO frequency is present, the RF signal will couple into the SRO and force oscillations to build up very quickly before the quench signal damps them again. This architecture therefore detects the presence or absence of an input RF signal, so is used for detection of an amplitude shift keyed

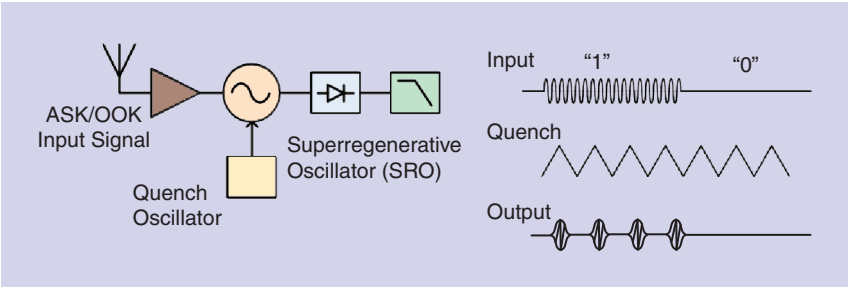

FIGURE 15: Super-regenerative receiver architecture.

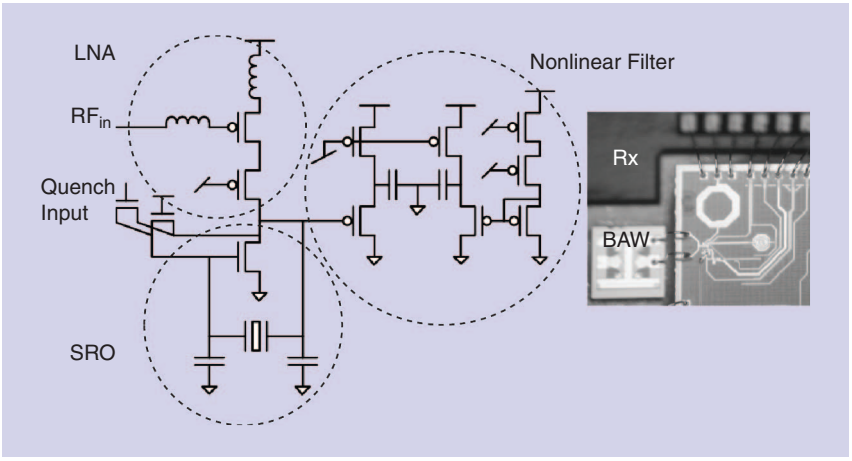

FIGURE 16: Ultra-low-power super-regenerative receiver.

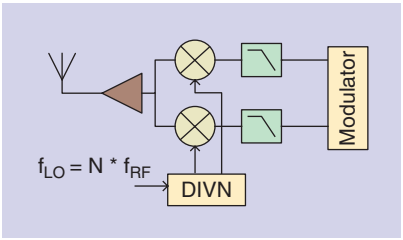

FIGURE 17: Direct conversion transmitter.

or ON-OFF amplitude modulation signal, as illustrated in Figure 16. The receiver architecture is simple, comprising just an amplifier and oscillator, followed by a baseband detector which is a simple envelope detector, and a number of ultra-low-power implementations have recently been proposed [10]–[15]. A recent practical example of this approach demonstrated entire receiver power consumption of less than 0.5 mW for a 5 kbps data rate [10].

Wireless System Architectures: The Transmitter

The most common approach to transmitter implementation is to generate a modulated signal at baseband, then up-convert to the RF frequency in a single stage as shown in Figure 17, i.e., direct up-conversion. A complex mix is needed to avoid generating a spurious signal at the image frequency; thus we need quadrature LO signals. Similar to the case of direct down-conversion, the simplest way of doing this is to implement a synthesizer, which generates an LO signal at two or four times the RF, and then divide down to get accurate quadrature.

A major issue for transmitters is that the power amplifier (PA) output power is high, and this modulated signal may couple into other parts of the circuit. If LO signals are generated from a voltage controlled oscillator (VCO) running at a harmonic of the RF frequency, this PA power coupling may cause frequency "pulling" of the VCO, severely degrading performance. VCO frequency pulling arises because the PA output is a modulated signal; that is, it has a varying instantaneous frequency, while the VCO is at a fixed frequency. The varying PA frequency couples into the VCO and pulls it away from its design value. However is it possible to avoid this problem by having the VCO frequency track the variable PA frequency? That way, any PA coupling back into the VCO would have no effect, and the VCO and PA could thus operate at harmonically related frequencies, or even at the same frequency.

This is the principle behind the direct loop modulator shown in Figure 18; the idea is that the phase modulation signal is "injected" directly into the frequency synthesizer, so that rather than producing a single frequency signal, the PLL directly generates a modulated signal at the RF frequency, which is then amplified by the PA for transmission. The entire transmitter is therefore just a synthesizer and PA, which

The implementation of an ultra-low-power wireless system requires an integrated design approach that considers many requirements.

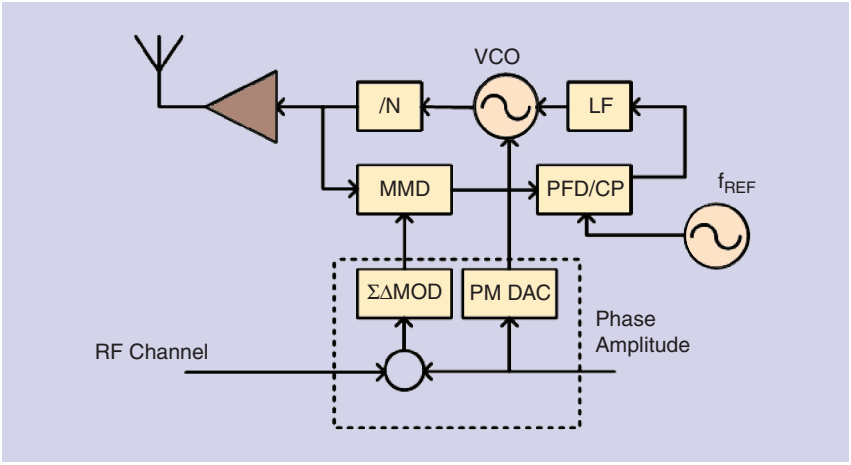

FIGURE 18: Direct modulation transmitter architecture.

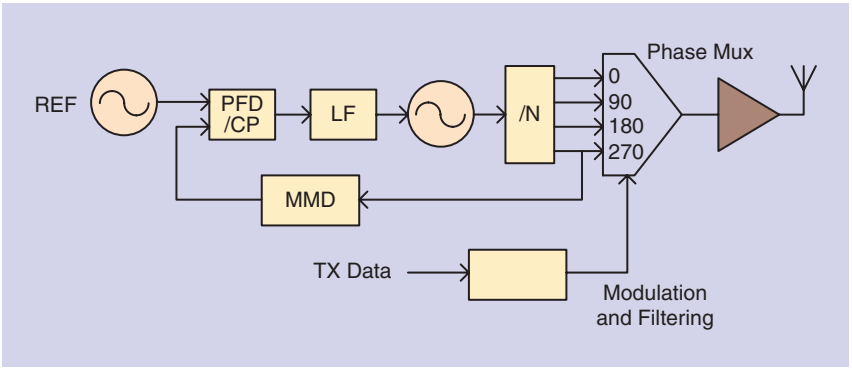

FIGURE 19: Direct phase modulation transmitter.

enables an efficient and low power implementation.

The starting point is the fractional-N sigma-delta synthesizer architecture. This frequency synthesizer architecture achieves a frequency step size, which is a fraction of the reference frequency by switching the divide ratio of the multimodulus divider between two different values; for example, if we divide by 2,000 for 15 cycles and then divide by 2,001 for one cycle, the average divide ratio is 2000 + 1/16. In practice if we switch between these two divide ratios at a relatively slow rate, then frequency "spurs" at this switching frequency are superimposed on the VCO output, which is undesirable if the aim is to generate a pure LO tone. The solution is to dither between the two divide ratios at a much higher frequency, controlled by a sigma-delta modulator. The higher frequency switching noise is then filtered by the loop filter, giving a much cleaner VCO output spectrum.

Since varying the divide ratio at a low frequency (within the bandwidth of the loop filter) will cause the VCO frequency to vary, i.e., to exhibit frequency modulation, this is exactly how the modulation signal can be "injected" directly onto the VCO. The sigma-delta modulator has a second phase modulation input as shown in Figure 18 that slowly controls the divide ratio such that the VCO frequency is not just a pure tone but varies according to the input modulation signal.

The limitation of this approach is that any phase modulation injected at the feedback divider will get filtered by the loop filter, thus the peak frequency deviation (and thus the data rate) becomes limited by the lop filter bandwidth. This limitation loop can be overcome by adding a second modulation injection point directly at the VCO control node as shown in Figure 18. Modulation signals injected at this point are subject to a high pass characteristic through the synthesizer loop; the net result is that injecting the modulation at two separate points in the loop allows frequency modulation of the VCO at data rates which exceed the bandwidth of the loop filter. The polar loop approach is attractive for low power wireless implementation because it simply requires the injection of baseband modulation data into the existing frequency synthesizer, and all that is additionally required is the PA. For this reason many recent low power wireless transmitters have been based on this approach [5]–[9].

Although the loop modulator is currently state of the art in terms of power consumption versus performance for commercial ultra-low power wireless systems, new approaches are being proposed to try and push the power/performance envelope even further. One such approach shown in Figure 19 is direct transmitter phase selection [16]–[19]. The idea is that rather than injecting the phase modulation into the synthesizer loop, a series of constant frequency VCO signals with defined phase offsets are generated. Modulation of the transmitted signal is then simplified to the control of a multiplexer which selects the appropriate signal phases in turn. For binary phase-shift keyed (PSK) modulation only two phases are required (at 180° offset), while transmission of GMSK requires four signals at 90° phase offset. Good quadrature signals can be produced by running the VCO at four times the wanted RF and then dividing by four, but clearly this is more power hungry. Furthermore higher-order phase modulation schemes require an increased number of phases. To overcome this problem, researchers have looked at methods to generate the phase offset signals without requiring frequency multiplication and division, for example the use of either polyphase filters or delay locked loops to produce the required multiple phases.

A hybrid approach shown in Figure 20 further reduces power consumption by generating the RF

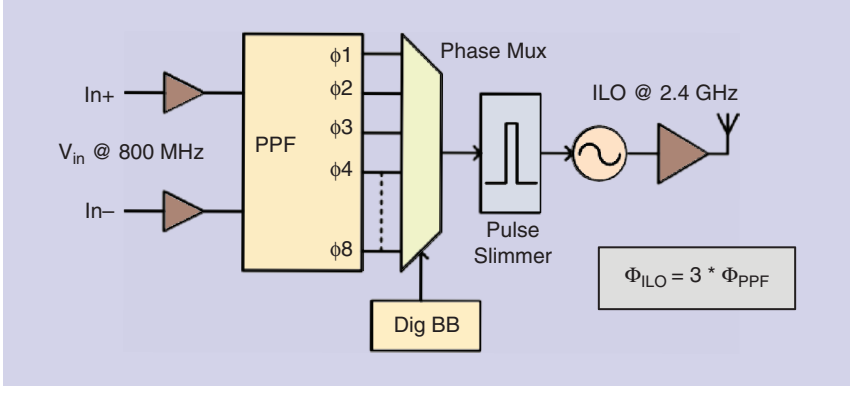

FIGURE 20: Subharmonic direct phase modulated transmitter with injection locking.

signal at a lower frequency and then employing an injection locked oscillator (ILO) [16]. In this example a polyphase filter (PPF) is used to generate the required signal phases but at a sub-harmonic frequency, in this case one third of the wanted RF. A multiplexer is employed to synthesize the modulated signal which is used to drive an injection locked oscillator tuned to the third harmonic, to generate an output signal at 2.4 GHz. A pulse slimmer is used to enhance the third harmonic component of the subharmonic signal, which increases the injection locking efficiency.

Circuit Architectures

With an appropriate system architecture selected, appropriate low power circuit implementations are required. "Good practice" for low power wireless circuit implementation should consider the following general goals:

- Reduce the maximum frequency at which the circuits need to operate, since circuits operating at lower frequencies can get away with smaller bias currents
- Circuits which do need to operate at RF should be as simple as possible; minimize the number of components
- Operate with very low supply voltages to reduce overall power consumption, provided your system can implement efficient dc-dc conversion
- Designing a circuit to get good performance across temperature, voltage, and process typically requires an element of "overdesign," which means that devices are larger and take more current then is really required at nominal, to ensure they still perform at the corner extremes. To avoid this, make use of digital calibration which ensures circuits are only biased to be "good enough" at the particular P, V, T operating point.
- Make use of high quality passives; lossy inductors and capacitors need to be compensated by much higher bias currents, and so if power consumption is a major driver you should ensure your process has good RF passives available.

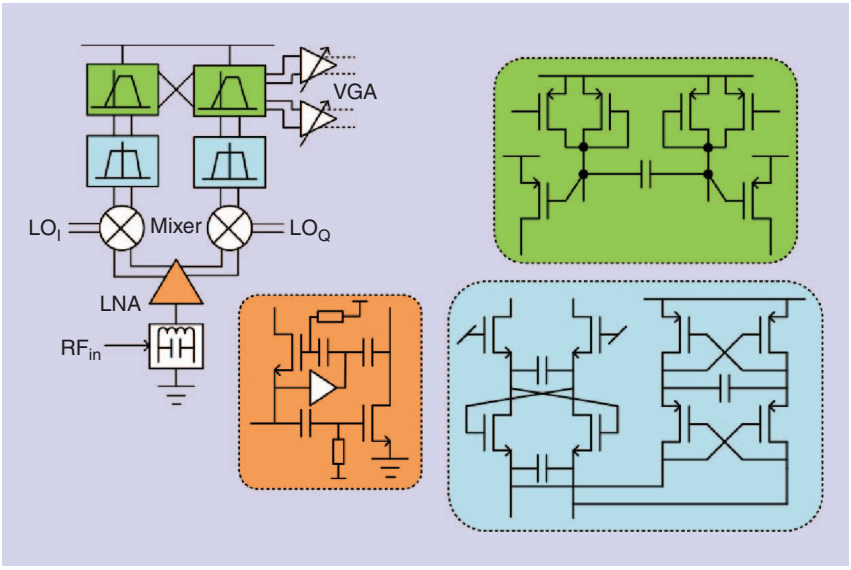

FIGURE 21: Current-reuse receiver architecture.

Current reuse is an interesting circuit technique for low power circuit implementation. In a typical wireless system, each of the various circuit blocks (LNA, mixers, filters, etc.) is designed independently and then connected to its neighbor circuits in a "cascade" arrangement. Current reuse reverses this approach and instead of cascading, the circuit blocks are effectively "cascoded" by stacking them on top of each other to share a single bias current [20]–[21]. Figure 21 shows a direct conversion receiver front-end where the LNA, mixer, and filter blocks are all stacked on top of each other and share a single bias current. A similar current-reuse approach has also been proposed for VCO design as shown in Figure 22 [22]–[24]. Both these techniques can lead to an overall reduction in total power consumption.

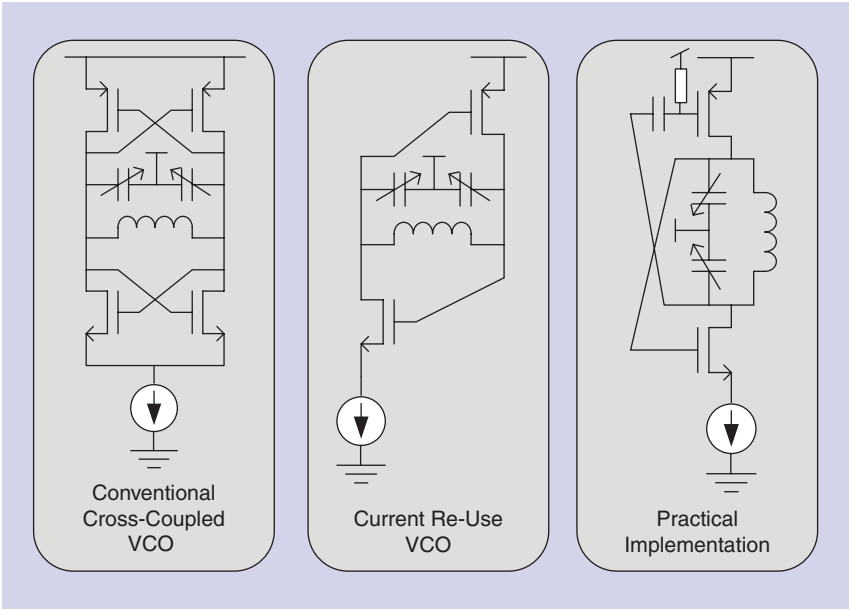

FIGURE 22: Current-reuse voltage-controlled oscillator.

Summary

To minimize operating power and achieve maximum battery lifetime, the implementation of an ultra-low-power wireless system requires an integrated design approach that considers many requirements including battery source, active and sleep mode energy requirements, system architectures, and circuit implementations. In practice, these issues lead to trade-offs, which may require numerous iterations to arrive at an optimal solution for the desired application. With the recent explosion of interest in ultra-low-power wireless systems for the Internet of Things and wearable devices, the current rate of innovation in the development of ultra-low-power wireless systems is sure to continue.

References

[1] D. C. Yates, A. S. Holmes, and A. J. Burdett, "Optimal transmission frequency for ultralow-power short-range radio links," *IEEE Trans. Circuits Syst. I*, vol. 51, no. 7, pp. 1405–1413, 2004.

[2] W. Thommen, "An improved low power crystal oscillator," in *Proc. 25th European Solid-State Circuits Conf. (ESSCIRC'99)*, 1999, pp. 146–149.

[3] D. Ruffieux, "A high-stability, ultra-low-power quartz differential oscillator circuit for demanding radio applications," in *Proc. 28th European Solid-State Circuits Conf., ESSCIRC 2002*, pp. 85–88.

[4] A.-S. Porret, T. Melly, E. A. Vittoz, and C. C. Enz, "Tradeoffs and design of an ultra low power UHF transceiver integrated in a standard digital CMOS process," in *Proc. 2000 Int. Symp. Low Power Electronics and Design (ISLPED'00)*, 2000, pp. 273–278.

[5] A. C. W. Wong, M. Dawkins, G. Devita, N. Kasparidis, A. Katsiamis, O. King, F. Lauria, J. Schiff, and A. J. Burdett, "A 1 V 5 mA multimode IEEE 802.15.6/Bluetooth low-energy WBAN transceiver for biotelemetry applications," *IEEE J. Solid-State Circuits*, vol. 8, no. 1, pp. 186–198, 2013.

[6] Y.-H. Liu, X. Huang, M. Vidojkovic, A. Ba, P. Harpe, G. Dolmans, and H. de Groot, "A 1.9nJ/b 2.4GHz multistandard (Bluetooth Low Energy/Zigbee/IEEE802.15.6) transceiver for personal/body-area networks," in *IEEE Int. Solid-State Circuits Conf. Dig. Technical Papers (ISSCC)*, 2013, pp. 446–447

[7] E. Le Roux, N. Scolari, B. Banerjee, C. Arm, P. Volet, D. Sigg, P. Heim, J.-F. Perotto, F. Kaess, N. Raemy, A. Vouilloz, D. Ruffieux, M. Contaldo, F. Giroud, D. Severac, M. Morgan, S. Gyger, C. Monneron, T.-C. Le, C. Henzelin, and V. Peiris, "A 1V RF SoC with an 863-to-928MHz 400kb/s radio and a 32b Dual-MAC DSP core for wireless sensor and body networks," in *IEEE Int. Solid-State Circuits Conf. Dig. Technical Papers (ISSCC)*, 2010, pp. 464–465.

[8] G. Devita, A. C. W. Wong, M. Dawkins, K. Glaros, U. Kiani, F. Lauria, V. Madaka, O. Omeni, J. Schiff, A. Vasudevan, L. Whitaker, S. Yu, and A. Burdett, "A 5mW multi-standard Bluetooth LE/IEEE 802.15.6 SoC for WBAN applications," in *Proc. 40th European Solid State Circuits Conf. (ESSCIRC)*, 2014, pp. 283–286.

[9] F. Pengg, D. Barras, N. Scolari, and A. Vouilloz, "A low power miniaturized 1.95 mm^2 fully integrated transceiver with fastPLL mode for IEEE 802.15.4/Bluetooth smart and proprietary 2.4GHz applications," in *Proc. IEEE RFIC Symp.*, 2013, pp. 71–74.

[10] B. Otis, Y. H. Chee, and J. Rabaey, "A 400 µW-RX, 1.6mW-TX super-regenerative transceiver for wireless sensor networks," in *IEEE Int. Solid-State Circuits Conf. Dig. Technical Papers (ISSCC)*, 2005, pp. 396–606.

[11] M. Vidojkovic, S. Rampu, K. Imamura, P. Harpe, G. Dolmans, and H. de Groot, "A 500uW 5Mbps ULP super-regenerative RF front-end," in *Proc. ESSCIRC,* 2010, pp. 462–465.

[12] F. X. Moncunill-Geniz, P. Pala-Schonwalder, C. Dehollaini, N. Joehl, and M. Declercq, "An 11-Mb/s 2.1-mW synchronous superregenerative receiver at 2.4 GHz," *IEEE Trans. Microwave Theory Tech.*, vol. 55, no. 6, pp. 1355–1362.

[13] C. Ma, C. Hu, J. Cheng, L. Xia, and P.Y. Chiang, "A near-threshold, 0.16 nJ/b OOK-transmitter with 0.18 nJ/b noise-cancelling super-regenerative receiver for the medical implant communications service," *IEEE Trans. Biomed. Circuits Syst.*, vol. 7, no. 6, pp. 841–850, 2013.

[14] K. Kim, S. Yun, S. Lee, and S. Nam, "Low-power CMOS super-regenerative receiver with a digitally self-quenching loop," *IEEE Microwave Wireless Components Lett.*, vol. 22, no. 9, pp. 486–488, 2012.

[15] H.-G. Park, J. Lee, J.-a Jang, J.-H. Jang, D.-S. Lee, H. Kim, S. J. Kim, S.-G. Lee, and K.-Y. Lee, "An ultra-low-power super regeneration oscillator-based transceiver with 177 mW leakage-compensated PLL and automatic quench waveform generator," *IEEE Trans. Microwave Theory Tech.*, vol. 61, no. 9, pp. 3381–3390, 2013.

[16] M. Rahman, M. Elbadry, and R. Harjani, "A 2.5 nJ/bit multiband (MBAN & ISM) transmitter for IEEE 802.15.6 based on a hybrid polyphase-MUX/ILO based modulator," in *Proc. IEEE Radio Frequency Integrated Circuits Symp.*, 2014, pp. 17–20.

[17] Y.-H. Liu, C.-L. Li, and T.-H. Lin, "A 200-pJ/b MUX-based RF transmitter for implantable multichannel neural recording," *IEEE Trans. Microwave Theory Tech.*, vol. 57, no. 10, pt. 2, pp. 2533–2541, 2009.

[18] S. Diao, Y. Zheng, Y. Gao, S.-J. Cheng, X. Yuan, M. Je, and C.-H. Heng, "A 50-Mb/s CMOS QPSK/O-QPSK transmitter employing injection locking for direct modulation," *IEEE Trans. Microwave Theory Tech.*, vol. 60, no. 1, pp. 120–130, 2012.

[19] S.-J. Cheng, Y. Gao, W.-D. Toh, Y. Zheng, M. Je, and C.-H. Heng, "A 110pJ/b multichannel FSK/GMSK/QPSK/p/4-DQPSK transmitter with phase-interpolated dual-injection DLL-based synthesizer employing hybrid FIR," in *IEEE Int. Solid-State Circuits Conf. Dig. Technical Papers (ISSCC)*, 2013, pp. 450–451.

[20] Z. Lin, P.-I. Mak, and R. P. Martins, "A 2.4 GHz ZigBee receiver exploiting an RF-to-BB-current-reuse Blixer + hybrid filter topology in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 6, pp. 1333–1344, 2013.

[21] T. Song, H.-S. Oh, E. Yoon, and S. Hong, "A low-power 2.4-GHz current-reused receiver front-end and frequency source for wireless sensor network," *IEEE J. Solid-State Circuits*, vol. 42, no. 5, pp. 1012–1022, 2007.

[22] J. Shi, J. Pandey, and B. P. Otis, "A 1.5 GHz high-Q receiver based on current reuse," in *Proc. 2011 IEEE Int. Symp. Circuits and Systems (ISCAS)*, 2011, pp. 1800–1803.

[23] W. Ying, P. Qin, J. Jin, and T. Mo, "A 1 mW 5 GHz current reuse CMOS VCO with low phase noise and balanced differential outputs," in *Proc. 13th Int. Symp. Integrated Circuits (ISIC)*, 2011, pp. 543–546.

[24] T. Siriburanon, W. Deng, K. Okada, and A. Matsuzawa, "A current-reuse Class-C LC-VCO with an adaptive bias scheme," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, 2013, pp. 35–38.

About the Author

Alison Burdett (Alison.Burdett@toumaz.com) has over 25 years of experience in semiconductor design, particularly in the field of ultra-low-power wireless communication. She joined Toumaz in 2001 as technical director, and is currently chief technology officer, responsible for delivering silicon and healthcare technology programmes within the company. Prior to joining Toumaz, Alison spent time both in industry as an integrated circuit designer, and also in academia (as senior lecturer in analogue IC design at Imperial College London). Dr. Burdett is a Chartered Engineer, a fellow of the Institute of Engineering and Technology, and a Senior Member of the IEEE. She was European Regional Chair (2013–2015) of the Technical Programme Committee for the IEEE International Solid-State Circuits Conference (ISSCC). She is a member of the U.K. National Microelectronics Institute Microelectronics Design Advisory Board and a visiting researcher at the Institute of Biomedical Engineering, Imperial College.

SSC