

# Frequency Offset Compensation for Crystal-free 802.15.4 Communication

Ankur M. Mehta, Kristofer S. J. Pister  
Department of Electrical Engineering and Computer Sciences  
UC Berkeley, Berkeley, CA, USA  
{mehtank,pister}@eecs.berkeley.edu

## *Abstract—*

The IEEE 802.15.4 standard is widely used in low power wireless sensor networks. Providing 250kbps raw data bandwidth, its physical layer (PHY) mandates tight constraints on the RF carrier frequency, calling for an absolute accuracy of  $\pm 40$  ppm. This specification necessitates the use of a crystal timing reference, thus requiring design of a multi-component PCB to implement any 802.15.4 compliant solution. However, in this work, a proposed frequency compensation algorithm can be added to the 802.15.4 PHY layer to relax the requirements on a timing reference to up to  $\pm 1000$  ppm relative frequency error between a transmitter and receiver, allowing for the elimination of the crystal and paving the way for a single-chip integrated 802.15.4 wireless node. A wide bandwidth channel-select filter allows for offset in the received signal carrier frequency, while an additional demodulator output estimates that frequency offset. An adaptive feedback loop can then adjust the receiver clock rate to center the received signal in the channel, following which a narrowband filter can be applied to restore noise performance. Such a system has been simulated, and the results presented in this work demonstrate the feasibility of standards-compliant wireless communication using inaccurate timing references.

## I. INTRODUCTION

Wireless sensors networks (WSNs) are becoming increasingly widespread, though they are yet far from being ubiquitous. One factor in this is that wireless sensor nodes are not cheap, requiring a full hardware solution to be effective: current radios all require a quartz crystal to provide an accurate frequency reference, increasing costs due to component and PCB requirements.

There has been some research that has produced crystal-free RF communication (e.g. [1]), however, such communication is typically broadband and not robust to interference. The IEEE 802.15.4 standard for Low-Rate Wireless Personal Area Networks [2] is a widely used protocol for WSN communication, and is meant to be scalable, low power, and robust. This work demonstrates the feasibility of crystal-free 802.15.4 compliant communication, the first step in moving towards an integrated single-chip wireless sensor mote.

Section II will describe the relevant aspects of the 802.15.4 standard, as well as give a brief background on frequency references in wireless network hardware. Section III will describe a first pass at addressing crystal-free communication, backed up by simulation results. A refinement will follow in Section IV outlining how to measure and compensate for frequency offsets between the transmitter and receiver.

## II. BACKGROUND

### A. 802.15.4 specifications

The most common wireless sensor network (WSN) protocol is the IEEE 802.15.4 standard [2], containing both medium access control (MAC) and physical (PHY) layer specifications. In particular, in the 2.4 GHz ISM band, the standard specifies 16 orthogonal channels spaced at 5 MHz centers. Data is encoded using 32 chip pseudo-noise sequences to represent 4 data bits, and with a raw chip rate of 2 Mbps achieves a data rate of 250 kbps. The chip sequence is modulated onto the RF carrier using offset quadrature phase-shift keying (OQPSK) with half-sine pulse shaping (HSS), giving a signal bandwidth of 2 MHz.

The 802.15.4 standard additionally specifies performance limits on compliant systems. The hardware must have an absolute clock accuracy of  $\pm 40$  ppm, ensuring tight bounds on the transmitter carrier frequency and allowing the receiver to use a narrow channel-select filter to minimize out-of-band noise power. The standard defines sensitivity to be measured at 1% packet error rate (PER) with 26 byte packets, corresponding to a symbol error rate (SER) of  $1.9 \times 10^{-4}$  [3]. In practice, this translates to a minimum acceptable signal to noise ratio (SNR) at the receiver to fulfill the 802.15.4 specifications.

However, [3] also notes that in real-world WSNs, the vast majority of links have significantly better than the minimum SNR required to hit the acceptable PER, and so other requirements, in particular the clock accuracy, may be relaxed while still maintaining network performance.

### B. Oscillator performance

In order to achieve the 40 ppm frequency accuracy required by the 802.15.4 specification, a high quality quartz crystal (with typical accuracy on the order of ones of ppm) is required as a timing reference. An oscillator fabricated in the silicon would perform far worse: most on-chip oscillators only guarantee accuracy on the order of one percent. It is possible to design oscillators with slightly better performance: a digitally controlled relaxation oscillator was designed in a  $0.18 \mu\text{m}$  CMOS process; the measured frequency accuracy over time of the actual hardware is displayed in figure 1. <sup>1</sup> With focused design, then, an on-chip frequency reference can be made to maintain  $\pm 1000$  ppm accuracy.

<sup>1</sup>Image and data courtesy Mark Lemkin, used with permission.

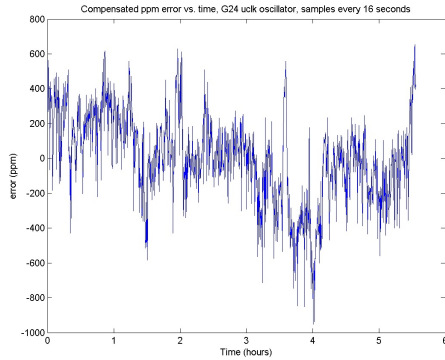


Fig. 1. The frequency of an on-chip relaxation oscillator drifts mostly between  $\pm 600$  ppm.

### III. CHANNEL SELECT FILTERING

A standard receive chain involves mixing the received RF signal down to a specified intermediate frequency (IF) (which can be at DC for a direct-conversion receiver). The down-converted signal is then passed through a channel-select filter then demodulated. Typically, an accurate frequency reference drives the local oscillator (LO) used in the mixer. This takes a received signal from a given RF carrier frequency down to the expected IF. However, if the LO is not tuned to the same frequency as the transmitter, the downconverted signal will be offset in frequency from the expected IF by the same amount as the difference between the LO and RF carrier frequencies.

The receiver includes the channel-select filter to improve sensitivity by attenuating out-of-band noise power. The filter needs to be wide enough to contain the signal bandwidth with sufficient margin to accommodate clock offset. For an 802.15.4 compliant system, the filter bandwidth (at baseband) needs to be at least  $(2 \text{ MHz})/2 + (40 \text{ ppm} \times 2.48 \text{ GHz}) \approx 1.1 \text{ MHz}$ . However, to accommodate clock offsets of up to 1000 ppm, the filter bandwidth needs to be expanded considerably lest the signal power itself be filtered out. This serves to increase the noise bandwidth and therefore the out-of-band noise power input to the demodulator, thus increasing the required in-band SNR for a given performance.

Simulations were run where a generated 802.15.4 transmission at RF was input to a receiver consisting of a mixer down to baseband (BB, zero IF), then a low pass channel-select filter, followed by a digital demodulator to recover the data. The LO of the receiver was allowed to vary up to 1000 ppm offset from the actual carrier frequency, and a range of noise power (in-band SNR values) was simulated.

Figure 2 compares the effect of varying the channel-select filter bandwidth. As expected, at zero clock offset, the SER drops with increasing SNR, becoming negligible above some threshold. That SNR threshold rises if the filter bandwidth is increased, as more out-of-band noise power is admitted to the demodulator. However, as seen in figure 3, at 1000 ppm clock offset between receiver and transmitter, no amount of SNR returns any usable signal with a 1.5 MHz filter: the




Fig. 2. The effect of increasing the filter bandwidth is to increase the out-of-band noise power admitted to the receiver, increasing the required in-band SNR to maintain performance.

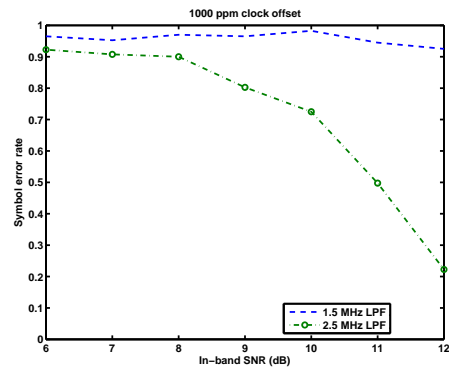


Fig. 3. Before increasing the filter bandwidth, the input signal is completely filtered out, leading to no output signal regardless of the SNR.

entire signal has been filtered out and only noise is captured. Expanding the filter to 2.5 MHz bandwidth allows some of the signal through, though there is enough noise power that the SER never really drops to zero.

The full simulation results with a 2.5 MHz filter, shown in figure 4, demonstrate the ability of a receiver to trade SNR for frequency offset tolerance, successfully decoding 802.15.4 transmissions under conditions that can be expected from replacing a crystal with an on-chip oscillator. As [3] points out, the additional SNR required to compensate for the clock mismatch is often readily available.

### IV. FREQUENCY OFFSET COMPENSATION

The aforementioned noise performance can be significantly improved with a digitally controllable oscillator. Many on-chip oscillator designs include controllability, typically through the use of switched resistors or capacitors. With sufficiently high resolution, such an oscillator can be commanded to get within a few ppm of any frequency, and thus, the LO can be driven to match the RF carrier frequency if it can be estimated.

The downconverted signal from the mixer is offset from the expected IF by the difference between the LO and RF carrier frequency. There have been algorithms devised to estimate this offset directly [4], [5], thereby recovering the clock error. However, in a standard digital OQPSK-HSS demodulator, a




Fig. 4. The noise performance of an 802.15.4 compatible receiver at varying clock offsets.

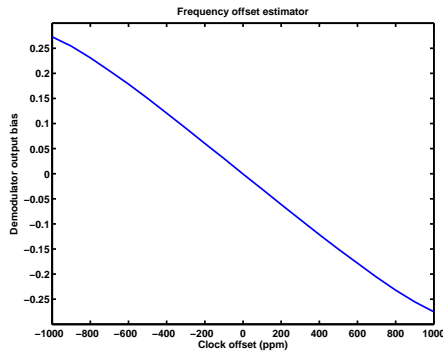


Fig. 5. The DC bias of the demodulator output (in arbitrary units) is directly proportional to the frequency offset between the receiver LO and the transmitter's carrier frequency. By applying a feedback loop on this bias value to the LO generator, the receiver clock can be driven to match the transmitter clock.

frequency shift in the input signal (at IF) translates to a constant DC bias in the demodulated output. Thresholding the demodulator output at this bias point returns the original data chips; the bias value itself can be used to extract the frequency offset between the LO and RF carrier frequencies.

The demodulator used in the above simulations extracted a DC bias as described; that value is plotted in figure 5, and its linear relationship to the frequency offset can be clearly seen. By using this measured bias as an error value, a feedback controller can be wrapped around a variable on-chip clock as outlined in figure 6. At the extremes of frequency offset, the bias / offset relationship deviates from linear; nonetheless the clear monotonicity remains present, and a feedback loop can drive the error signal to zero, aligning the LO with the RF carrier frequency.

The 802.15.4 standard includes a known 5 octet long PHY synchronization header on which this feedback loop can operate. Successive refinements can allow the feedback to be robust to an unknown slope, centering the LO on the RF carrier frequency for the PHY payload. At this point, the output of the mixer is at the expected IF, obviating the need for a wide bandwidth channel-select filter. By allowing the channel-select filter to be of a variable bandwidth, or by including a second stage filter (potentially in the digital domain), the excess noise




Fig. 6. A block diagram depicting how the additional frequency offset estimate demodulator output can be used to adjust the LO frequency until it matches the RF carrier frequency

bandwidth can be filtered out, restoring the original noise performance of a crystal driven 802.15.4 network.

## V. CONCLUSION

It is often simply assumed that a crystal frequency reference is necessary for RF communication, and as such any wireless sensor network solution will require some board level design. However, using data from actual hardware as a base, this work has presented simulations that demonstrate robust, narrow-band, standards compliant RF communication is indeed possible without an external quartz oscillator. By including the RF oscillator on chip, the path is paved to a single-chip, fully integrated wireless sensor mote.

However, there are still issues that need to be addressed. Research is ongoing into expanding the solution presented in this work into a protocol capable of multi-channel mesh networking. The current work synchronizes a receiver to a transmitter. Further study needs to address the dynamics that may arise across a network of multiple nodes, all synchronizing to each other. Additionally, at the limits of the possible frequency offsets, adjacent channels may overlap, thus necessitating more a more sophisticated approach to medium access.

Nonetheless, this work takes the first step, by demonstrating the feasibility of a point-to-point crystal-free wireless link.

## ACKNOWLEDGMENT

The authors would like Mark Lemkin for his work on the hardware used as a basis for this research.

## REFERENCES

- [1] S. Drago, D. Leenaerts, F. Sebastiano, L. Breems, K. Makinwa, and B. Nauta, "A 2.4ghz 830pj/bit duty-cycled wake-up receiver with -82dbm sensitivity for crystal-less wireless sensor nodes," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, feb. 2010, pp. 224 –225.
- [2] "IEEE std 802.15.4-2006: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (WPANs)," <http://ieeexplore.ieee.org/servlet/opac?punumber=11161>.
- [3] S. Lanzisera, A. Mehta, and K. Pister, "Reducing average power in wireless sensor networks through data rate adaptation," in *Communications, 2009. ICC '09. IEEE International Conference on*, Jun. 2009, pp. 1 –6.
- [4] M. Morelli and U. Mengali, "Joint frequency and timing recovery for msk-type modulation," *Communications, IEEE Transactions on*, vol. 47, no. 6, pp. 938 –946, jun 1999.
- [5] R. Mehlan, Y.-E. Chen, and H. Meyr, "A fully digital feedforward msk demodulator with joint frequency offset and symbol timing estimation for burst mode mobile radio," *Vehicular Technology, IEEE Transactions on*, vol. 42, no. 4, pp. 434 –443, nov 1993.