Temperature Compensation of Crystal References in NB-IoT Modems

Tuomas Haapala, *Student Member, IEEE*, Antonio Liscidini, *Senior Member, IEEE*, and Kari A. I. Halonen, *Member, IEEE*

***Abstract*— The low-cost nature of NB-IoT modems encourages them to deploy uncompensated crystal oscillators (XOs) as frequency references. The frequency offset of an uncompensated XO, however, makes network acquisition inefficient under low network coverage. In the worst case, the bulk of an NB-IoT modem's power is consumed in network acquisition. This work demonstrates a discrete frequency synthesizer prototype that employs its phase-locked loop to compensate for the frequency offset of its reference XO. We propose an accurate crystal model and a compensation logic that are suitable for 32-bit microprocessors, commonly available in NB-IoT modems. Alternatively, when synthesized in a commercial 65-nm process in 24-bit precision, the proposed compensation logic is simulated to require a total area of 0.029 mm^2 and power of 290 nW at a refresh rate of 1.4 kHz. Moreover, the XO model requires only 0.16 kB of RAM. The prototype achieves a compensation error level of down to 27 ppb (3σ) over the temperature range from −40 to 85 °C. The dominant error sources of the prototype are discussed in detail. In addition, we show that an NB-IoT modem can use the information acquired from sequential network acquisitions to compose and update its XO model in the field. This support for field calibration removes the need for XO characterization in production and ensures the validity of the model over the NB-IoT modem's lifespan. The prototype achieves a compensation error level of down to 50 ppb (3σ) in an emulated in-field calibration test, sufficient to ensure energy-efficient low-latency network acquisition under low network coverage.**

***Index Terms*— Low-power wide-area networks, phase-locked loop, temperature compensation, field calibration, crystal oscillator, narrowband Internet of Things, NB-IoT, LPWAN.**

I. Introduction

THE number of smart devices connected to the internet is predicted to grow exponentially in the near future. To support this trend, low-power wide area network (LPWAN) standards, such as NB-IoT, LoRa and SigFox, have emerged with the goal of providing a wireless range in the order of several kilometers, together with a low level of power consumption. A common target for the maximum operation time of an LPWAN modem is more than 10 years on a single coin battery [1]. This operation time requirement is equal to an average power consumption level of approximately several tens of microwatts.

In practice, the required microwatt-level power budget is achieved by aggressive duty cycling. The NB-IoT standard employes a mode named extended Discontinuous Reception for this purpose. In this mode, a user connects to the network in cycles ranging from 20.48 seconds to almost 3 hours [2]. The radio transceiver is kept in a deep-sleep state in the middle of a cycle, during which it leaks a minimal amount of power.

LPWAN standards utilize ultra-narrow communication bandwidths to support a long wireless range. As an example, the minimum bandwidths for LTE-M, NB-IoT and SigFox are specified as 180 kHz, 3.75 kHz and 100 Hz, respectively [1]. Due to the narrowness of the specified bandwidth, an NB-IoT modem allows for a frequency error of only ±200 ppb and ±100 ppb for sub-GHz and over-one-GHz transmissions, respectively [3]. Meanwhile, the frequency tolerance of a low-cost crystal oscillator (XO) is approximately ±50 ppm over the industrial temperature range from −40 to 85 °C.

An NB-IoT modem compensates for the excess frequency offset of its low-cost reference XO based on known network syncronization signals. Nevertheless, the duration of network acquisition is greatly affected by the initial frequency offset of the XO under low network coverage. At worst, the energy consumed for network acquisition may become a dominant contributor to the total power consumption of the modem [4].

The frequency offset problem could be remedied by replacing the XO with a discrete temperature-compensated XO (TCXO), but TCXOs are expensive [5]. Alternatively, the XO can be compensated on-chip by means of a digital-to-analog converter (DAC) [6] or a switched-capacitor array [5], [7]. MEMS oscillators are commonly compensated by adding a highly tunable sequential phase-locked loop (PLL) [8], [9]. However, the above solutions have drawbacks in terms of area, power consumption and noise level. They also require an XO model that is calibrated in production, which increases costs. Furthermore, the calibration becomes obsolete over time due to long-term drifts in XO frequency.

Due to its ultra-narrow bandwidth, an NB-IoT modem requires a high-resolution fractional-N PLL (FPLL) in its frequency synthesizer. Using this existing FPLL, the modem

Manuscript received October 4, 2019; revised January 18, 2020; accepted February 4, 2020. Date of publication February 24, 2020; date of current version July 1, 2020. This work was supported in part by the Naked Approach (40336/14, 3246/31/2014) and Towards Digital Paradise (2727/31/2016) projects granted by Business Finland and in part by the LightningSense project (319043) granted by the Academy of Finland. The work of Tuomas Haapala was supported in part by the Foundation for Aalto University Science and Technology under Travel Grant and in part by the Aalto ELEC Doctoral School Grant. This article was recommended by Associate Editor G. Jovanovic Dolecek. *(Corresponding author: Tuomas Haapala.)*

Tuomas Haapala and Kari A. I. Halonen are with the Department of Electronics and Nanoengineering, Aalto University, 02150 Espoo, Finland (e-mail: tuomas.haapala@aalto.fi; kari.halonen@aalto.fi).

Antonio Liscidini is with the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: antonio.liscidini@utoronto.ca).

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

Digital Object Identifier 10.1109/TCSI.2020.2973908

can compensate for its reference offset with little power, area or noise overhead. In addition, the support for field calibration would remove the need for calibration in production and allow the modem to adapt to the long-term frequency drift of its XO. Radio-frequency transceivers utilizing their existing PLL in XO compensation have been previously described, for instance, in [10], and the deployment of wireless satellite or terrestrial signals in the composition of an XO model in the field has been previously suggested, for instance, in [11]. In this paper, however, we provide a detailed analysis of these techniques applied specifically to NB-IoT modems:

1) We propose a temperature-compensation logic for the FPLL that remains accurate enough for commmon XOs with a 24-bit floating-point implementation. Current commercial NB-IoT modems are commonly based on 32-bit microprocessors [12], [13].
2) We propose a simple but accurate field calibration algorithm that uses averaging to compensate for the inaccuracy of an NB-IoT modem's network detector.
3) We show that a high-resolution look-up-table (LUT) paired with a quadrature interpolator allows an excellent XO modeling accuracy and simple maintenance in the field compared to popular polynomial models.
4) We demonstrate the functionality of the compensation logic, field calibration algorithm and XO model in a discrete prototype and analyze its major error sources. We provide measured data about the static and dynamic behavior of low-cost XOs and discuss their impact on the performance of the prototype. We provide estimates of the area, power and memory consumption of the prototype in an integrated implementation scenario.

The discrete temperature-compensated PLL (TCPLL) prototype achieves an accuracy of down to 27 ppb (3σ) over the industrial temperature range using an up-to-date XO model. In an emulated in-field calibration test, the prototype achieves an accuracy of 50 ppb (3σ), which is adequate for ensuring energy-efficient low-latency network acquisition and data transfer under low network coverage. When synthesized in a commercial 65-nm CMOS process as a hardware accelerator, the model and logic measure approximately 0.029 mm^2 and are simulated to consume 290 nW of power at a refresh rate of 1.4 kHz. In addition, the proposed XO model requires only 0.16 kB of RAM while the NB-IoT modem in [12], for instance, has 256 kB of RAM reserved for user applications.

While this paper is focused on NB-IoT, its analysis and results are applicable in any radio transceiver equipped with a means of sensing temperature, a PLL and a possibility for the detection of an accurate reference frequency in the field.

This paper is organized as follows. Section II gives a brief description of the acquisition of an NB-IoT network and gives an estimate of the TCPLL's effect on the power consumption of a strongly duty-cycled NB-IoT modem. Section III gives a general introduction to the TCPLL concept. Section IV details our implementation of the TCPLL's control logic. Section V presents our XO model implementation, discusses its main error sources and describes a simple field-calibration algorithm. Section VI presents the measurement setup and Section VII details the design of the TCPLL prototype. The performance of the TCPLL prototype is evaluated in Section VIII and compared with conventional XO compensation methods in Section IX.

II. Acquisition of an NB-IoT Network

After waking up from hibernation, an NB-IoT modem has to acquire its frequency and timing offset from an NB-IoT base station and detect the base station's cell identity. For this purpose, an NB-IoT base station transmits two recurring synchronization signals known as the Narrowband Primary Synchronization Signal (NPSS) and the Narrowband Secondary Synchronization Signal (NSSS). An NB-IoT modem uses the NPSS to acquire coarse esimates for its frequency and timing offset, and the NSSS to acquire its fine timing offset (frame boundary) and the identity of the base station [14]. While the NSSS can also be used to refine the frequency offset estimate [15], many works base their frequency estimation solely on the NPSS [5], [16]. A common accuracy target for the frequency estimation is $\pm$50 Hz ($\pm$50 ppb at 1 GHz) [5], [15], [16] to meet the $\pm$100-ppb accuracy requirement established for NB-IoT user equipment [3].

An NB-IoT modem detects the NPSS by means of, for instance, matched and differential filtering [4] or a maximum likelihood detector [17]. However, the duration of the network acquisition is greatly affected by the modem's initial frequency offset. Under low network coverage, an accurate frequency reference allows the acquisition within a couple of tens of milliseconds, while a high initial frequency offset may extend the acquisition time up to more than a second [15]–[17]. Since the actual data exchange between a modem and a base station may contain a few hundreds of bytes, lasting for only several tens of milliseconds, the energy consumed for network acquisition may become a significant contributor to the total energy budget. It is estimated in [4] that fast NPSS detection can improve the efficiency of the modem by 24 %. Consequently, a considerable effort has recently been put on developing network searchers for NB-IoT modems that are tolerant to large initial frequency offsets [4], [17], [19]. While these network searchers relieve the detection problem slightly, they also bring about an increased level of complexity and power consumption. Alternative network synchronizations signals have also been proposed that provide an improved detection performance [18].

After an NB-IoT modem has detected and compensated its frequency offset, the residual frequency offset can be detected and compensated at the NB-IoT base station. A small residual frequency offset level allows an NB-IoT modem to save further energy under low network coverage because the base station requires a smaller number of repetitions to successfully decode the transmitted data. It is estimated in [20] that an NB-IoT modem under low network coverage with a 0.1-% duty-cycle can achieve a 22.5-% efficiency gain by decreasing its residual frequency offset from $\pm$100 Hz ($\pm$100 ppb at 1 GHz) to less than $\pm$10 Hz ($\pm$10 ppb at 1 GHz).

III. Temperature-Compensated Fractional-N PLL

The temperature-dependent frequency offset of a low-cost AT-cut XO is shown in Fig. 1. The frequency offset of an XO

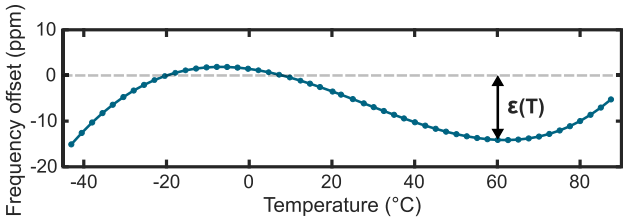

Fig. 1. Frequency offset of an uncompensated AT-cut XO.

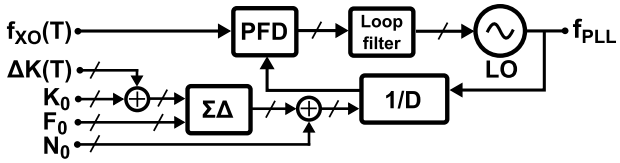

Fig. 2. Block diagram of a fractional-N phase-locked loop with a temperature-compensation factor ΔK. The loop consist of a phase-frequency detector (PDF), a loop filter, a radio-frequency local oscillator (LO) and a frequency divider (1/D) modulated by a sigma-delta modulator ($\Sigma\Delta$).

at temperature T is defined as

$$\epsilon(T) = \frac{f_{XO}(T)}{f_{XO,0}} - 1, \tag{1}$$

where $f_{XO,0}$ is the specified output frequency of the XO and f_{XO} is the actual output frequency of the XO.

An NB-IoT modem requires an FPLL for translating its low-frequency reference signal f_{XO} into a radio-frequency signal f_{PLL} that is utilizable by its radio transceiver front-end. The analysis of this paper is applicable to both digital and analog FPLLs. A temperature-compensated FPLL is depicted in Fig. 2. The output frequency of the FPLL is given by

$$f_{PLL}(T) = D(T) f_{XO}(T), \tag{2}$$

where $D(T)$ is the loop divider. The loop divider is composed of its integer part N_0, fractional denominator F_0 and fractional numerator $K(T)$ as

$$D(T) = \left(N_0 + \frac{K(T)}{F_0}\right) = \left(N_0 + \frac{K_0 + \Delta K(T)}{F_0}\right), \tag{3}$$

where the fractional numerator is divided into its fixed part K_0 and a temperature-dependent offset $\Delta K(T)$. In practice, the loop divider block of an FPLL can only produce integer division ratios, and the fractionality is produced by sigma-delta modulating the fractional part of (3) with an integer-output quantizer. The function of PLLs is detailed, for instance, in [21]. $D(T)$ can also be written as

$$D(T) = \frac{D_0}{1 + \kappa(T)}, \tag{4}$$

where D_0 is the value of the loop divider when $\Delta K = 0$ and κ is a temperature compensation variable.

Substituting terms $D(T)$ and $f_{XO}(T)$ in (2) for the same terms from (1) and (4) results in

$$f_{PLL}(T) = \frac{1 + \epsilon(T)}{1 + \kappa(T)} D_0 f_{XO,0}. \tag{5}$$

This expression shows that the output frequency becomes independent of temperature when the compensation variable $\kappa(T)$ is equal to the frequency offset $\epsilon(T)$ of the XO. In practice,

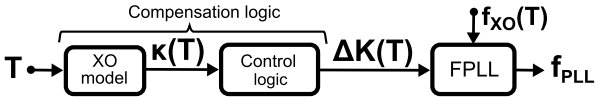

Fig. 3. Outline of a TCPLL.

$\kappa(T)$ is determined by an XO model. The numerator offset ΔK of a temperature-compensated PLL can now be set based on a modeled temperature offset κ of an XO. This numerator offset is found by combining (3) and (4), giving

$$\Delta K(T) = \left(\frac{D_0}{1 + \kappa(T)} - N_0\right) F_0 - K_0. \tag{6}$$

The basic operation of a TCPLL is depicted in Fig. 3. The compensation logic of the TCPLL is comprised of an XO model and a PLL control logic. The XO model provides a prediction κ of the frequency offset of a reference XO at temperature T. The control logic translates κ further into the corresponding temperature-dependent FPLL compensation offset ΔK. In effect, ΔK acts as a frequency demodulator opposing the modulation effect of a changing reference frequency.

IV. Control Logic

Since the magnitude of $\kappa(T)$ ranges from 0 to the order of several tens of ppm over the industrial temperature range, the term $(1+\kappa(T))^{-1}$ in (6) is not properly interpretable with 32-bit floating point precision of up to 7.2 decimal digits. However, the baseband processors of many NB-IoT modems are based on a 32-bit microprocessor architecture [12], [13]. To decrease the precision requirement, $\Delta K(T)$ can be approximated with

$$\Delta K(T) \approx \left(-\kappa(T) + \kappa(T)^2\right) D_0 F_0, \tag{7}$$

which is derived by substituting the term $(1 + \kappa(T))^{-1}$ in (6) for its second-order Taylor series approximation about the point $\kappa = 0$. This derivation is presented in Appendix A. Eq. (7) is not vulnerable to round-off errors.

A custom 24-bit floating point implementation of (7) was written in Verilog and simulated in ModelSim by Mentor Graphics. The used exponent and mantissa lengths are 7 and 16, respectively, with the remaining bit reserved for the sign. The error of the implemented control logic is evaluated from

$$E_c(\kappa) = \frac{f_{PLL}(\Delta K_c(\kappa))}{f_{PLL}(\Delta K_{ref}(\kappa))} - 1, \tag{8}$$

where $\Delta K_c(\kappa)$ is computed by the implemented logic and $\Delta K_{ref}(\kappa)$ by a 64-bit reference implementation of (6) in MATLAB by MathWorks. We examine here the error of the control logic alone assuming an ideal XO model, i.e. that κ is equal to ϵ.

The simulated error of the custom 24-bit control logic is shown in Fig. 4 along with the errors of 32-bit implementations of (6) and (7) that were simulated in MATLAB. In the simulation, F_0 was set to $2^{24} - 1$ to represent a 24-bit FPLL, and N_0 was set to 40 to represent a frequency translation from 25 MHz to about 1 GHz. K_0 is given a random value between 0 and F_0 from trial to trial. At low XO frequency offsets below

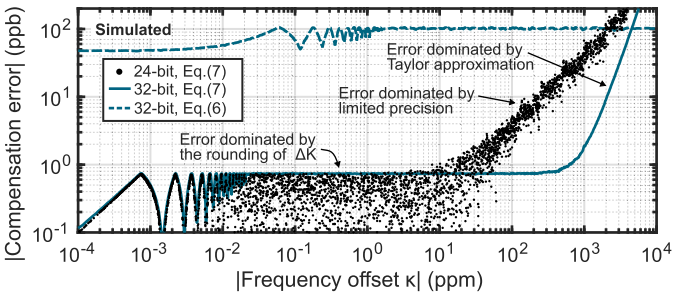

Fig. 4. Compensation error of the proposed PLL control logic. The figure shows the errors of a 24-bit and a 32-bit floating point implementation of (7) and a 32-bit implementation of (6). The error of the 24-bit implementation is shown as raw data. The errors of the 32-bit implementations are shown as envelopes for clarity.

±10 ppm, the error of the control logic is dominated by the rounding of ΔK that is forwarded to an FPLL as an integer. The error floor level is slightly below 1 ppb for a 24-bit FPLL and every 1-bit decrease in the FPLL order elevates the floor level by two. At high XO frequency offsets above ±10 ppm, the error of the 24-bit implementation of (7) is dominated by its limited floating point precision. However, the result shows that, despite the limited precision, the implementation can be used with reference oscillators of up to a ±100-ppm frequency offset with a reasonable one-decade margin to the ±100-ppb accuracy requirement of NB-IoT modems. Eq. (7) is employable in 32-bit precision up to a reference frequency offset of ±2000-ppm despite the Taylor approximation. Due to its large range, the proposed compensation logic can also be used with MEMS oscillators.

V. XO Model

This section discusses the modeling of XOs. The discussion is focused on AT-cut XOs, commonly employed in low-cost applications. The performance of polynomial and LUT-based XO models are compared based on measured data. We discuss the frequency drift and dynamic behavior of the XOs and analyze their effect on modeling accuracy. A field calibration algorithm for an LUT-based XO model is presented at the end of this section.

The frequency offsets of 16 low-cost AT-cut XOs from two different manufacturers, A and B, were characterized over the industrial temperature range from −40 to 85 °C with a temperature step size of 2.5 °C. The measured frequency offsets of the XOs are presented in Fig. 5(a) and their corresponding temperature sensitivities (f-T slopes) are shown in Fig. 5(b). The XOs have a specified output frequency of 25 MHz and their specified tolerances over the target temperature range are ±25 ppm and ±50 ppm for manufacturers A and B, respectively. Herein the XOs are named commonly as device-under-test (DUT) XOs or separately as XOs A0–A15 and B0–B15. Three of XOs A are from a different manufacturing lot than the rest. The measurement setup is detailed in Section VI. The characterization of Fig. 5 was performed with chamber LabEvent-LC.

Typically, an activity dip results from an overlap of the fundamental and a high-order mode of a crystal at a narrow

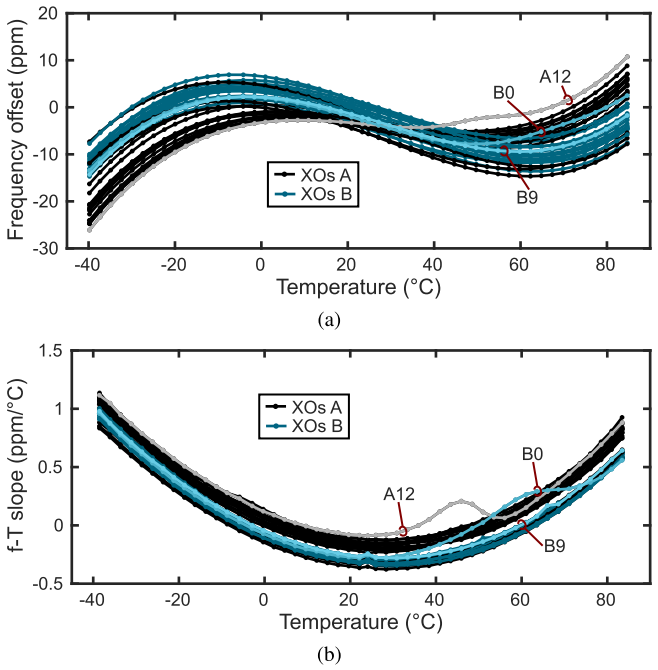

Fig. 5. Frequency offsets of the DUT XOs. DUT XOs A12, B0 and B9 are highlighted due to their activity dips. (a) Frequency offsets. (b) Corresponding f-T slopes.

temperature range, which increases the motional resistance of the crystal. The frequency pull caused by an activity dip ranges usually from 2 to 20 ppm [22]. These dips cause an extra modeling error since they create an additional local non-linearity in the f-T characteristics of an XO. We refer to an activity dip that causes major or minor local non-linearities as a sharp or blunt dip, respectively. Two of the DUT XOs, A12 and B0, showed a blunt activity dip at around 40 °C. In addition, XO B9 showed a sharp activity dip at around 65 °C. These three XOs are highlighted in Fig. 5. The small peaks shown at around 25 °C in Fig. 5(b) are caused by the non-linearity of the chamber's integrated temperature sensor.

A. Comparison of XO Models

The frequency offset of an XO or MEMS oscillator over temperature is often modeled by means of a polynomial [7]–[9], or a look-up table (LUT) paired with a linear [6] or quadrature [23] interpolator. The n-th degree polynomial model of an XO is given as

$$\kappa(T) = \sum_{i=0}^{n} a_i T^i, \tag{9}$$

where $\kappa(T)$ is the modeled temperature offset of the XO at temperature T, and a_i are the model coefficients. An LUT coupled with a linear interpolation would result in a systematic error of about ±100 ppb in the high and low end of the target temperature range, as shown in Appendix B. To avoid this systematic compensation error, we employ instead a quadratic interpolator based on the Newtons Divided

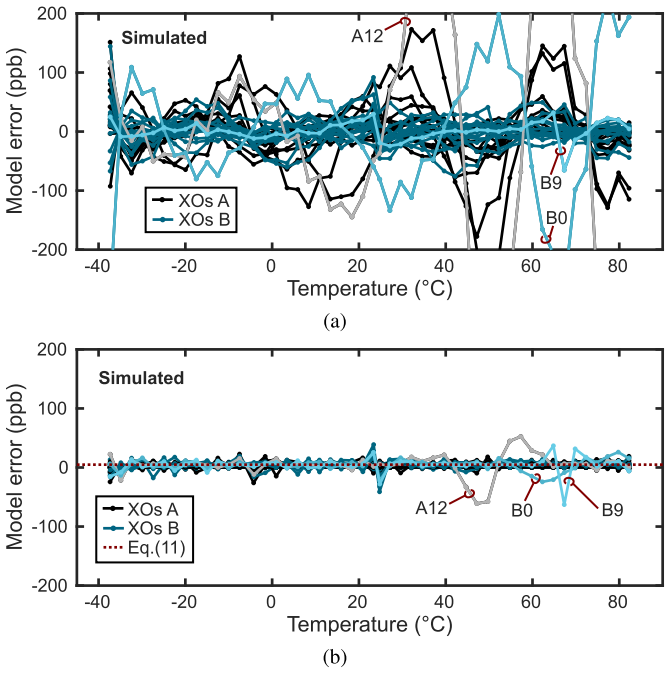

Fig. 6. Errors of a polynomial and an LUT-based XO model estimated by the cross-validation of the data in Fig 5(a). (a) 7th order polynomial model. (b) 5-°C-step LUT with a quadratic interpolator. DUT XOs A12, B0 and B9 are highlighted due to their activity dips.

Difference Polynomial Method [24]

$$\kappa(T) = \epsilon_0 + \frac{\epsilon_1 - \epsilon_0}{T_1 - T_0}(T - T_0) + \frac{\frac{\epsilon_2 - \epsilon_1}{T_2 - T_1} - \frac{\epsilon_1 - \epsilon_0}{T_1 - T_0}}{T_2 - T_0}(T - T_0)(T - T_1), \quad (10)$$

where $\{\epsilon_0, T_0\}$, $\{\epsilon_1, T_1\}$ and $\{\epsilon_2, T_2\}$ define the tabled ϵ-T pairs and $T \in I_T$. I_T is the smallest interval containing the temperatures T_0–T_2. The quadratic interpolator is composed of a linear interpolator accompanied by a quadratic term that can handle minor local non-linearities. The maximum systematic error of this interpolator can be estimated from

$$E_2 = \frac{2\sqrt{3}}{9}\Delta T^3 \alpha, \quad (11)$$

where ΔT is the resolution of the LUT and α is a constant. This error is not a function of temperature and it assumes that $T_0 \leq T \leq T_1$. The value of coefficient α is about 0.1 ppb/$(°C)^3$ for the DUT XOs. Eq. (11) and this value for α are derived in Appendix B.

The computational errors of a 7th-degree polynomial model and a 5-°C-step LUT with a quadrature interpolator were estimated in MATLAB by applying leave-one-out cross validation to the data in Fig. 5(a). For a given XO, one data sample was left out and the models were build upon the remaining data. The process was repeated over the measured temperature range. The estimated model errors are shown in Fig. 6(a) for the polynomial model and Fig. 6(b) for the LUT model. The polynomial model gives a significant general error level of 160 ppb (3σ). In particular, the activity dip of XO A12 gives a maximum modeling error of −610 ppb. In contrast, the quadrature interpolator gives a general compensation error level of only 24 ppb (3σ) and a maximum error of −61 ppb for the same activity dip. The estimated error of the quadrature interpolator has a mean of 5 ppb which is the same as the prediction given by (11). The small increases in the model error level at around −5 and 25 °C are caused by the non-linearity of the integrated temperature sensor of the LabEvent-LC that affects the LUT's local resolution.

The realization of a 7th-degree polynomial model of (9) requires 7 multiplications and 7 summations whereas the quadrature interpolation of (10) requires 11 summations, 4 divisions and 3 multiplications. Due to the divisions, the quadrature interpolator requires a more complex hardware. However, an LUT is maintained in the field by simply updating the tabled ϵ-T pairs. In contrast, updating a polynomial model requires the re-computation of the polynomial coefficients by, for instance, applying ordinary least squares estimation, which requires solving a system of linear equations. Consequently, the polynomial model would be significantly more laborious to maintain in the field.

Assuming a 24-bit implementation of the compensation logic, a 7th-degree polynomial model requires 72 B of RAM to save its 8 model coefficients and at least 8 f-T pairs for computing the coefficients in the field. A 5-°C-step LUT requires 26 f-T pairs, equal to 156 B of RAM, to cover the frequency range from −40 to 85 °C. While the LUT requires a double amount of memory, the total memory usage remains insignificant compared to the commonly available memory size in NB-IoT modems. The NB-IoT modem in [12], for instance, has 256 kB of RAM reserved for user applications.

In this work, we employ a 5-°C-step look-up table with an quadrature interpolator due to its accuracy and simple maintenance in the field.

B. Error Sources of a Static XO model

The frequency offset of an XO drifts over time for reasons such as aging, thermal hysteresis and micro-jumps. Due to these sources of drift, an XO model may become invalid in the time span of just a few weeks.

The two dominat causes of XO aging are contamination transfer in the crystal and stress relief in its mounting structure [25]. An uncompensated low-cost XO and a TCXO commonly have a specified aging rate of ±5 ppm and ±1 ppm per year, respectively. These rates tend to decelerate with XO age. An accelerated aging characterization of the DUT XOs is shown in Fig. 7, measured in chamber VTM7004. The XOs were incubated at 105 °C over two successive one-week periods for 14 days in total. This total aging time corresponds to roughly two years of aging at 25 °C [26]. The aging rate of the XOs decelerated over time. One of the measured XOs proved unstable at 105 °C and its data is not shown.

Thermal hysteresis occurs when an XO experiences thermal gradients. Some suggested causes of thermal hysteresis include contamination redistribution, strain changes in a mounting structure and changes in the quartz or oscillation circuitry [27]. Additionally, when an XO and its temperature sensor are separated, their temperature difference and different thermal lag cause so-called apparent hysteresis [27].

The thermal hysteresis of the DUT XOs is characterized in Fig. 8. The measurement was performed at around 60 °C

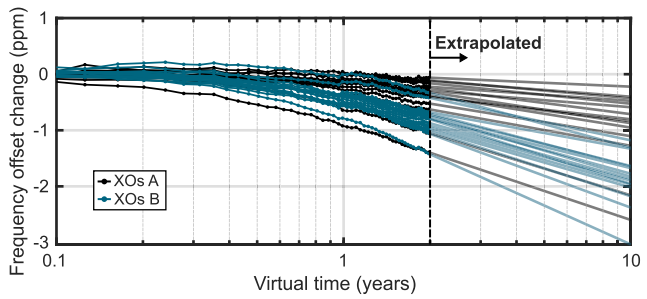

Fig. 7. Accelerated aging of the DUT XOs. Equivalent aging rate at 25 °C, scaled according to [26]. Original aging data was measured at 105 °C.

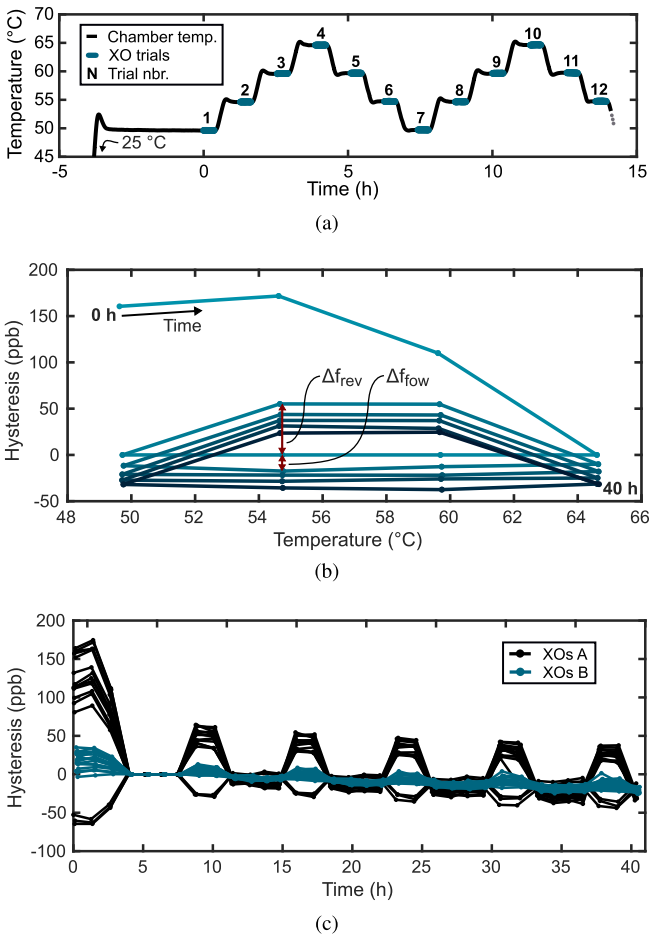

Fig. 8. Thermal hysteresis of the DUT XOs during successive temperature cycles. (a) Measurement transient. (b) Hysteresis of XO A12. (c) Hysteresis of the DUT XOs.

where the f-T slopes of the DUT XOs are small, minimizing the effect of small temperature changes on the result. Fig. 8(a) shows the transient of a cycled thermal hysteresis measurement. After a 3-h period of incubation at 50 °C, the DUT XOs were characterized for a total of 40 h in successive temperature cycles that had a magnitude of 15 °C. The characterization of a DUT XO after a temperature step constitutes a trial. All DUT XOs are trialled after each temperature step. The temperature steps of the measurement transient show slight overshoots or undershoots, which presumably have only a minor effect on the measurement result due to their low magnitude and duration compared to the main temperature cycle. The small temperature difference between trials taken at the same temperature level (such as trials 3, 5, 9, etc.) are compensated numerically based on the f-T slopes of the DUT XOs estimated from trials 4–7. Fig. 8(b) shows the thermal hysteresis of XO A12. The shown data is referenced to trials 4–7. The magnitude of measured thermal hysteresis is large when a trial is compared to a trial taken during the opposite direction of the thermal gradient (reverse hysteresis, Δf_{rev}). However, there is also a small frequency shift remaining after each full temperature cycle (forward hysteresis, Δf_{fow}). The thermal hysteresis of the DUT XOs is shown in Fig. 8(c), also referenced to trials 4–7. The periods of reverse hysteresis are easily distinguishable due to their large magnitude. XOs A show a significantly higher level of hysteresis than XOs B. Furthermore, the direction of observed hysteresis is different for the three XOs A from a different manufacturing lot. The mean forward hysteresis level of the DUT XOs has shifted to -23 ppb at around 40h. One frequency reading gave a distinct value that differed from the common trend by about 40 ppb, which is the cause of the divergent line at about 32 h in Fig. 8(c). The reason for the exceptional reading is unknown.

Micro-jumps are sudden step changes in the frequency of an XO. Their typical magnitude is in the order of 10 ppb and peak magnitude as high as 300 ppb [28]. Some suggested contributors include the surface preparation of a crystal [29] and crystal drive level [30], but apparently there is currently no consensus regarding the likely causes of the jumps [28], [31]. A micro-jump may be intermittent, discrete or abrupt [32].

Since aging affects the frequency offset of an XO over the time frame of months, it can be accounted for by the update of the XO's model in the field. However, the effect of thermal hysteresis is immediate and thus cannot be accounted for by a static crystal model with a long refresh period. Modeling and compensating the effect of thermal hysteresis accurately would instead require a dynamic model that takes into account the recent thermal history of an XO. While such models have been developed [33], the implementation of one would add greatly to the complexity of the compensation logic and its memory usage. In addition to changing over temperature, the thermal hysteresis of an XO may also change with XO age and duty-cycle. As shown Fig. 8(c), the error level of a static model after a rather large temperature transient is, however, less than ± 100 ppb even for XOs that show a high level of thermal hysteresis. Depending on its magnitude, a micro-jump may have a significant sudden effect on the validity of an XO model. However, the effect of an abrupt micro-jump will be compensated during successive network acquisitions over the time frame of a few days. Another dynamic effect not taken into account by static models is XO start-up, which is discussed in more detail in Section VI.

Other environmental factors such as acceleration may have a significant effect on the frequency of an XO, but their analysis is beyond the scope of this paper. Further information on the different environmental sensitivities of XOs can be found, for instance, in [22].

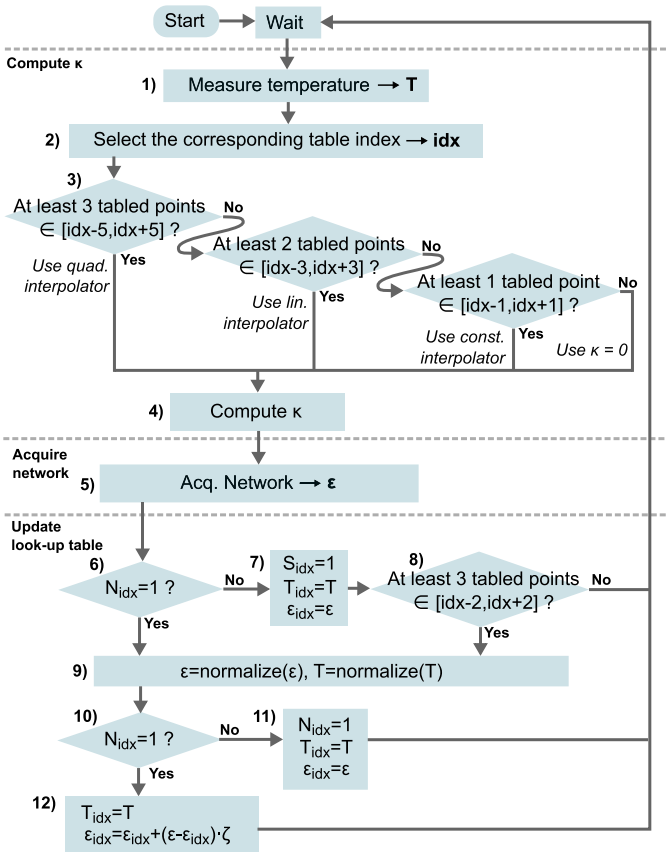

Fig. 9. Field calibration algorithm for the XO model of an NB-IoT modem.

C. Field Calibration of an XO Model

In network acquisition, an NB-IoT modem may sweep the ΔK of its FPLL based on (2) until it detects the NB-IoT network synchronization signals. After acquiring a satisfying value for ΔK, the modem can compute the corresponding XO frequency offset based on the expression

$$\epsilon(\Delta K) = -\frac{\Delta K}{N_0 F_0 + K_0 + \Delta K}, \tag{12}$$

which is derived by solving κ from (6) and marking it as equal to ϵ. This expression can be evaluated with sufficient accuracy using 32-bit precision because its dominating term ΔK in the numerator does not suffer from round-off errors.

The acquired estimate of ϵ can be used to update an LUT-based XO model based on a simple algorithm. Such an algorithm is depicted in Fig. 9 and its functionality is demonstrated in Section VIII. In addition to an ϵ-T pair, each LUT index contains status bits S and N. Status bit S tells whether its corresponding ϵ-T pair contains any meaningful data and is utilized in algorithm phases 3) and 8). Status bit N tells whether the data has additionally been normalized. Normalized data is referenced to a normalized integer temperature point, such as 25 °C instead of 26.235 °C. This normalization ensures a predictable maximum interpolation error level according to (11) and allows averaging ϵ-T pairs over multiple network acquisitions, virtually improving the precision of an NB-IoT modem's frequency detector. It would also allow the tracking of an XO's frequency drift statistically to, for instance, detect micro-jumps in a more sophisticated LUT update algorithm.

The status of all ϵ-T pairs of the LUT is initially invalid ($S_{\text{idx}} = 0\ \forall$ idx) and no ϵ-T pairs are normalized ($N_{\text{idx}} = 0\ \forall$ idx). When triggered, the proposed algorithm 1) measures the temperature T of an XO and 2) selects the corresponding LUT index idx. The algorithm then 3) chooses an LUT interpolation method based on the available valid (previously acquired) ϵ-T pairs near the idx. Multiple tabled ϵ-T pairs near the idx allow the use of the quadrature interpolator but a smaller amount of tabled ϵ-T pairs may still allow for utilizing the linear or constant part of (10). If all found tabled ϵ-T pairs locate above or below the idx, the selected interpolator acts in practice as an extrapolator. The found tabled ϵ-T pairs and the selected interpolator are then used 4) to compute an estimate κ for the XO's frequency offset that can be used as a basis for network acquisition. The network acquisition results in 5) a new estimate for the frequency offset ϵ of the XO. The temperature T of the XO may have to be remeasured at this point due to the possible long duration of the network acquisition. 6) If the ϵ-T pair at the idx, ϵ_{idx}-T_{idx}, has not previously been normalized, 7) the acquired ϵ-T pair is used to update the LUT as is. However, 6) if ϵ_{idx}-T_{idx} has previously been normalized or 8) there is a sufficient amount of nearby tabled ϵ-T pairs available, the algorithm moves to a normalization phase. In the normalization phase, the acquired ϵ-T pair is 9) re-referenced to the normalized integer temperature point corresponding to the idx by employing the quadrature interpolator of the XO model. Depending on 10) whether the data at the idx has been previously normalized, the normalized ϵ-T pair is 11) used to update the table as is, or 12) with averaging where the constant $\zeta \in [0, 1]$. A high value for ζ results in a highly dynamic but less averaging LUT whereas a low value results in a less dynamic but highly averaging LUT.

VI. Measurement Setup

The measurements were performed with two temperature chambers, Vötsch Technik models VTM7004 and LabEvent-LC, based on the temperature stability requirement of a measurement and chamber availability. The VTM7004 has a broader temperature range but also a high temperature drift level at the low end of the targeted temperature range. The LabEvent-LC offers a better stability and was used for measurements that required a high level of accuracy, although it has a non-linear integrated temperature sensor.

To emulate the operation of the XOs in an NB-IoT modem, the XOs were duty-cycled in all measurements during idle periods with a sleep time of about 40 s, followed by an active time of about 5 s.

The DUT XOs are characterized with the setup depicted in Fig. 10(a). The output of the XO board f_{XO} is filtered by a 30-MHz low-pass filter and evaluated with a signal source analyzer that uses a rubidium frequency standard as its reference f_{REF}. The specified frequency tolerance of the 10-MHz output of the rubidium standard is $\pm$0.05 ppb. The TCPLL output f_{PLL} is evaluated after mixing it down to an intermediate frequency f_{IF} of approximately 25 MHz,

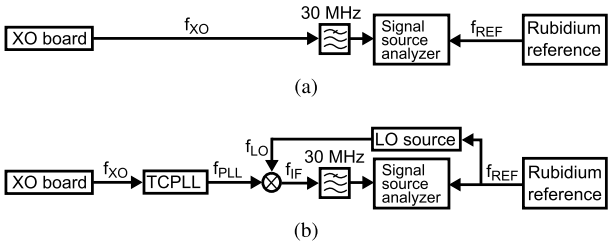

Fig. 10. Measurement setup for (a) XO characterizations and (b) TCPLL prototype characterizations.

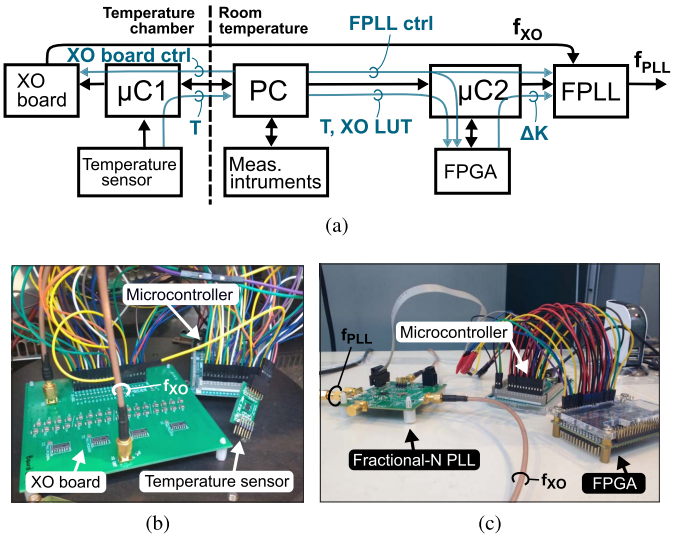

Fig. 11. TCPLL prototype. (a) Block diagram and the flow of the core signals. (b) Prototype components in a chamber. (c) Prototype components at room temperature.

as shown in Fig. 10(b). The mixing is performed because the signal source analyzer offers a better frequency resolution for low-frequency input signals. The employed local oscillator (LO) for the mixer uses the rubidium standard as its reference as well.

VII. TCPLL Prototype

The TCPLL prototype is depicted in Fig.11(a). The prototype employs an LMX2571 FPLL from Texas Instruments, a TSYS01 temperature sensor from TE Connectivity and a De0-Nano FPGA board from Terasic. The FPLL uses its maximum order of 24 bits and a loop divider value of about 38.65, corresponding to an output frequency of about 966.3 MHz. The components are interfaced with a PC and two Aruduio Due boards that are based on a 32-bit ARM core microcontroller. The DUT XOs are soldered to a custom XO board. The XO board, temperature sensor and one of the microcontrollers, shown in Fig. 11(b), are placed in a temperature chamber while the rest of the components, shown in Fig. 11(c), are kept at room temperature.

The FPGA board is programmed with the control logic presented in Section IV and the quadrature interpolator presented in Section V, both employing the custom 24-bit floating-point format. In addition to forwarding ΔK to the FPLL, the FPGA can forward the output of its interpolator to the PC to support the normalization phase of the field calibration algorithm. The interpolator can be programmed to output only the constant or linear part from (10) to provide the selection of different interpolators required by the algorithm. The full compensation logic has a slightly higher precision-dominated error level compared to the corresponding error of the control logic alone in Fig. 4. However, its simulated maximum error level is still smaller than $\pm$10 ppb for reference oscillators of a $\pm$50 ppm frequency offset. The error of a 32-bit implementation of the compensation logic is similar to the error of the control logic alone since its dominant error sources remain the same.

The compensation logic was synthesized in a commercial 65-nm CMOS process to analyze its performance as a hardware accelerator. The design reuses its floating point arithmetic units in sequential clock cycles to save area. The computation of (7) and (10) is controlled by a finite-state machine. The core of the synthesized 24-bit logic measures approximately 170 μm x 170 μm and was simulated to consume 250 nW of active power at a clock rate of 32.7 kHz and 40 nW of leakage power (typical process corner at 25 °C) from a 1.1-V supply. The computation of both the quadrature XO model and the corresponding ΔK requires 22 clock cycles, resulting in a refresh rate of 1.4 kHz. The computation of ΔK alone according to a new set of FPLL settings only requires 5 clock cycles, corresponding to a refresh rate of 6.6 kHz (0.15 ms). The frame structure of the NB-IoT standard is divided to time slots of 0.5 ms [3].

The apparent start-up time t_a of the DUT XOs is defined by their quality factor Q and frequency according to

$$t_a \approx \frac{Q}{\pi f_{XO}}. \tag{13}$$

The apparent start-up takes several milliseconds for a commercial tens-of-MHz XO. However, on top of the apparent start-up, the settling time of an XO is defined by self-heating, the effect of which is a function of the XO's f-T slope according to

$$\Delta f_T \approx \Delta T \cdot \frac{\mathrm{d}f}{\mathrm{d}T}, \tag{14}$$

where $\mathrm{d}f/\mathrm{d}T$ is the f-T slope of the XO and ΔT is the change in XO temperature due to self-heating.

The settling time of the DUT XOs is measured in Fig. 12 in chamber LabEvent-LC with 10-°C steps over the temperature range from -35 to 85 °C. The f-T slopes of the DUT XOs are estimated from the data of Fig. 5(b). The settling of XO A14 is shown in Fig. 12(a) as an example. All XOs A behaved in a similar fashion. Eight of XOs B, however, showed a distinct start-up behavior, as demonstrated by XO B5 in Fig. 12(b). An initial rise in their output frequency is followed by a sudden drop at the low end of the target temperature range. The delay of the frequency drop increases with decreasing temperature, taking as long as almost 3 minutes to occur. Re-soldering some of the distinct XOs had no significant effect on their start-up behavior. The measured changes in the frequency offset of the DUT XOs during start-up are plotted as a function of their f-T slopes in Fig. 12(c). The distinct XOs are omitted from the shown data. The figure shows that the sign and magnitude

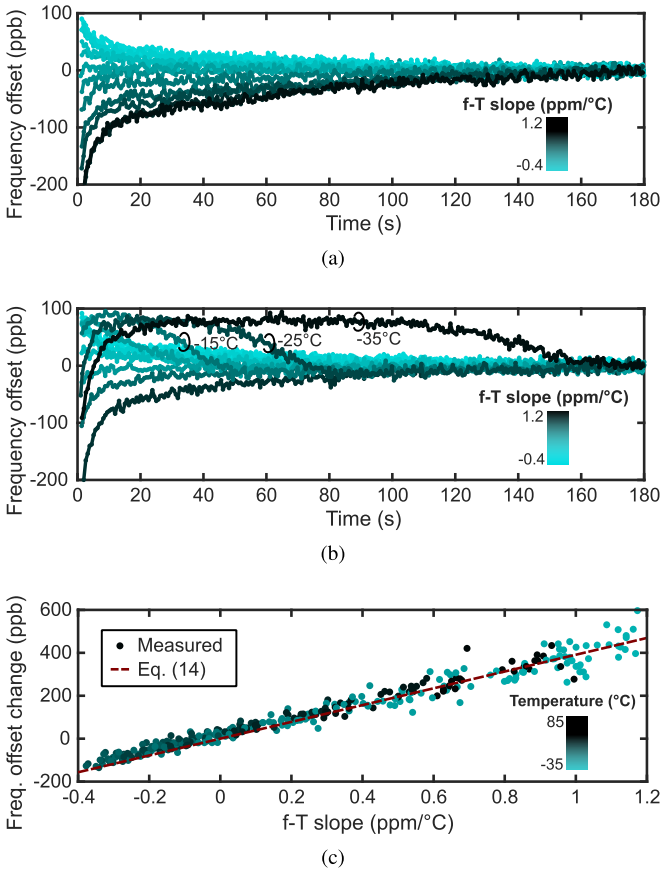

Fig. 12. Self-heating of the DUT XOs. (a) Self-heating of XO A14. (b) Self-heating of XO B5. (c) Change in the frequency offsets of the DUT XOs during start-up. The f-T slopes of the DUT XOs are estimated from the data of Fig. 5(b).

of the frequency change is an approximately linear function of the XOs' f-T slopes, as descibed by (14), rather than a function of XO temperature. The shown linear fit uses (14) with $\Delta T = 0.4$ °C, confirming that the DUT XOs suffer from a significant level of self-heating.

An XO can drain a significant amount of energy during its long warm-up time. Since XOs have fixed f-T slopes, the effect of self-heating can be diminished only by decreasing the magnitude of self-heating ΔT according to (14). This could be achieved in an NB-IoT modem by employing an ultra-low power XO, such as the one in [34], and by isolating a crystal resonator thermally from its oscillator circuitry. Additionally, we observed in our measurements that an active XO would also allow nearby XOs to heat-up faster. This suggests that a DUT XO does not only heat up itself but also its surroundings on our XO board design. A better thermal isolation between a DUT XO and the rest of the PCB by means of, for instance, air gaps would allow the XO to reach a thermal equilibrium faster. In this work, however, we allow the DUT XOs a significant warm-up time. Since we use the same DUT XO warm-up during LUT generation and TCPLL employment, the minor residual frequency offsets are included in the XO models and are compensated by the TCPLL.

The temperature offset of the prototype's temperature sensor is measured in Fig. 13 using the integrated sensors of the

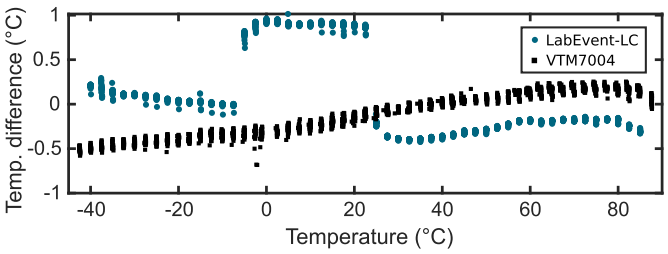

Fig. 13. Characterization of the TCPLL prototype's TSYS01 temperature sensor. Temperature offset with chamber LabEvent-LC and VTM7004 as a reference.

two temperature chambers as references. The chambers have a resolution of 0.1 °C, which causes the low measurement precision. The accuracy of the discrete sensor is similar to integrated ultra-low power temperature sensors that support 1-point temperature-based [35] or voltage-based [36] calibration. 1-point calibration decreases sensor manufacturing costs since the single calibration point can be taken at room temperature. The minor non-linearities of the prototype's temperature sensor are compensated out by the TCPLL because its XO model is constructed based on readings from the same sensor. The measurement also shows that the integrated sensor of the LabEvent-LC experiences strong non-linearities at around -5 and 25 °C. While these non-linearities do not affect the function of the TCPLL per se, they do affect the spacing of the LUT in some measurements and, consequently, the resulting systematic error. This effect is clearly visible in the cross-validation in Fig. 6 and the measurement in Fig. 15.

The prototype's temperature sensor has a self-heating magnitude of about 80 m°C during a time span of 2 minutes. The self-heating magnitude was measured at a 20-Hz sample rate at room temperature. The effect of this error source can be estimated from (14). However, the effect is diminished greatly in the prototype due to its low sample rate and systematic warm-up time during LUT generation and TCPLL employment. As with the reference XOs, the self-heating of temperature sensors encourages careful thermal design and the use of low-power sensor types.

The physical distance between a temperature sensor and an XO leads to a temperature difference ΔT between them that changes over time. This causes another compensation error whose magnitude can be estimated from (14). Due to the high temperature homogeneity of chamber LabEvent-LC, however, this error source did not seem to have a significant effect on the accuracy of the prototype. The compensation errors caused by the self-heating of an XO and a temperature sensor, and the physical separation between the two could be eliminated altogether by employing a dual-mode XO. A dual-mode XO can act as its own temperature sensor, as detailed in [23].

Due to its low integration level, the TCPLL prototype has a rather slow refresh rate of about 0.6 s that includes taking a temperature sample, sending it along with an LUT to the compensation logic, applying the computed ΔK to the FPLL and reading the frequency of the FPLL output from the signal source analyzer. The slow refresh rate causes a compensation error if the frequency of an XO is changing rapidly. The consequent frequency error Δf_t at a refresh rate of Δt can

TABLE I
MAIN TCPLL PROTOTYPE ERROR SOURCES

Error source	Approx. effect	Comp. method	Approx. residual effect
XO freq. offset	30 ppm[5(a)]	XO model	30 ppb[6(c)]
XO activity dips	3 ppm[5(a)]	XO model	50 ppb[6(c),⋆]
XO aging	1 ppm/year[7(b)]	Field calib.	negligible[†]
XO micro-jumps	10 ppb[⋆]	Field calib.	negligible[‡]
XO thermal hyst.	50 ppb[8(c)]	-	50 ppb
XO self-heat	500 ppb[12(c)]	Warm-up XO model	10 ppb
T. sens. accuracy	500 ppb[13,×]	XO model	negligible
T. sens. self-heat	100 ppb[×]	1.7-Hz RR[◇] XO model	negligible
T. gradients	1 ppm/°C[5(b)]	1.7-Hz RR[◇]	negligible[14]
T. uncertainty	<20 ppb	-	<20 ppb
Freq. meas.	10 ppb[14]	Averaging	negligible

A superscripted number refers to the corresponding figure.
[⋆] From [28].
[◇] System refresh rate.
[⋆] Sharp activity dips result in a significant compensation error.
[†] Slow phenomenon compared to field calibration.
[‡] Managed by field calibration with a latency.
[×] Assuming an f-T slope of 1 ppm/°C.

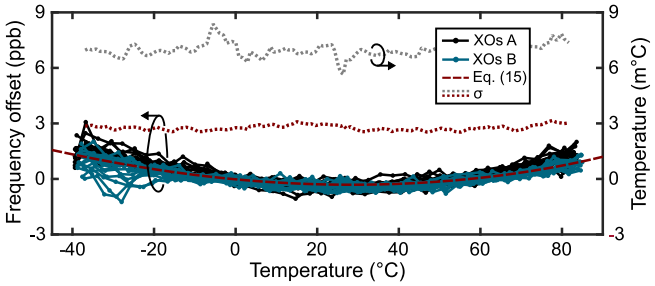

Fig. 14. TCPLL performance with an ideal XO model.

be approximated from

$$\Delta f_t \approx \Delta t \cdot \frac{\mathrm{d}f}{\mathrm{d}t} = \Delta t \cdot \frac{\mathrm{d}f}{\mathrm{d}T} \cdot \frac{\mathrm{d}T}{\mathrm{d}t}, \tag{15}$$

where $\mathrm{d}T/\mathrm{d}t$ is the temperature drift of the XO. This error is rather small but detectable in the measurement of Fig. 14. In an NB-IoT modem, this error is further reduced by a significantly faster refresh rate enabled by a high integration level.

The discussed compensation error sources and how they are managed by the TCPLL prototype are summarized in Table I. A major residual error source originates from activity dips that, however, affect only over a limited temperature range. In a more sophisticated compensation algorithm, the local resolution of the LUT could be increased around the detected activity dips to decrease their error level. Thermal hysteresis is not modeled by the prototype due to its complexity and, consequently, it constitutes another major source of error.

VIII. TCPLL PROTOTYPE EVALUATION

The performance of the control logic in Fig. 3 is evaluated in Fig. 14. The temperature of the LabEvent-LC was swept from −40 to 85 °C with a ramp rate of approximately 7 °C/h. Each DUT XO was compensated successively over an 18-sample window. Instead of using an XO model, the XOs were compensated based on the measured frequency offset samples. Fig. 14 reports the mean of the frequency offset samples over each window. Additionally, the figure reports the standard deviations of the frequency offset and temperature samples as a moving average over 32 successive measured windows. The sample-to-sample precision of the measurement setup is about 3 ppb (1σ) and the chamber has a uniform stability over the target temperature range. The measured accuracy of the compensation logic is in good agreement with the simulated result in Fig. 4. The small systematic compensation error is explained well by the constant temperature ramp of the chamber and the slow refresh rate of the prototype, as predicted by (15). Some XOs B give a slightly differing mean compensation error at around −40 to −20 °C due to their distinct start-up behavior: a rapid drop in the frequency of an XO during a measurement window translates into a negative systematic compensation offset, as also predicted by (15).

The performance of the TCPLL with an up-to-date LUT is evaluated with the LabEvent-LC in Fig. 15. The transient of the measurement is shown in Fig. 15(a). The spacing of the LUTs is composed based on the integrated sensor of the chamber. The momentarily increase in chamber temperature before the construction of the LUTs is to ensure that the DUT XOs are biased in forward hysteresis during the measurement. The compensated frequency offset of the TCPLL is presented in Fig.s 15(b)-(c) for XOs A and B, respectively. The moving 3σ compensation errors averaged over 32 successive trials are represented by the dashed lines. XOs A12, B0 and B9 are highlighted due to their activity dips. The local error caused by the dips is clearly visible in the measured data. XO B11 causes an additional error peak at around 29 °C that might be caused by a very narrow activity dip not visible in Fig. 5. The increase in the compensation error levels for all DUT XOs at around 25 °C is most likely caused by the strong non-linearity of the chamber's temperature sensor, as predicted in Fig. 6(c). The general error level of XOs A is significantly higher than that of XOs B due to their large thermal hysteresis. The distributions of the compensation errors for XOs A and B are shown in Fig.s 15(d)-(e), respectively. The TCPLL compensates XOs B with a very small error level of 27 ppb (3σ). The error level for XOs A is also acceptable despite their thermal hysteresis.

The field calibration algorithm proposed in Section V-C is evaluated in Fig. 16. The LabEvent-LC was set to follow the weather data of the observation station in the Tapiola area of Espoo, Finland, from 1 to 12 August 2018 [37] at a fourfold speed. Each DUT XO was measured and its LUT was updated once in 0.6 h (real time), corresponding to an aggressively duty-cycled NB-IoT modem with a 2.4-h network connection cycle (virtual time). The measured XO frequency offsets were given an additional random normally distributed 50-ppb (3σ) error before being updated to the LUTs to simulate the imprecision of the frequency detector of an NB-IoT modem [5], [15], [16]. The employed temperature data and the corresponding chamber temperature readings are shown in Fig. 16(a). The figure also shows the approximate time points when the TCPLL moves between different interpolators

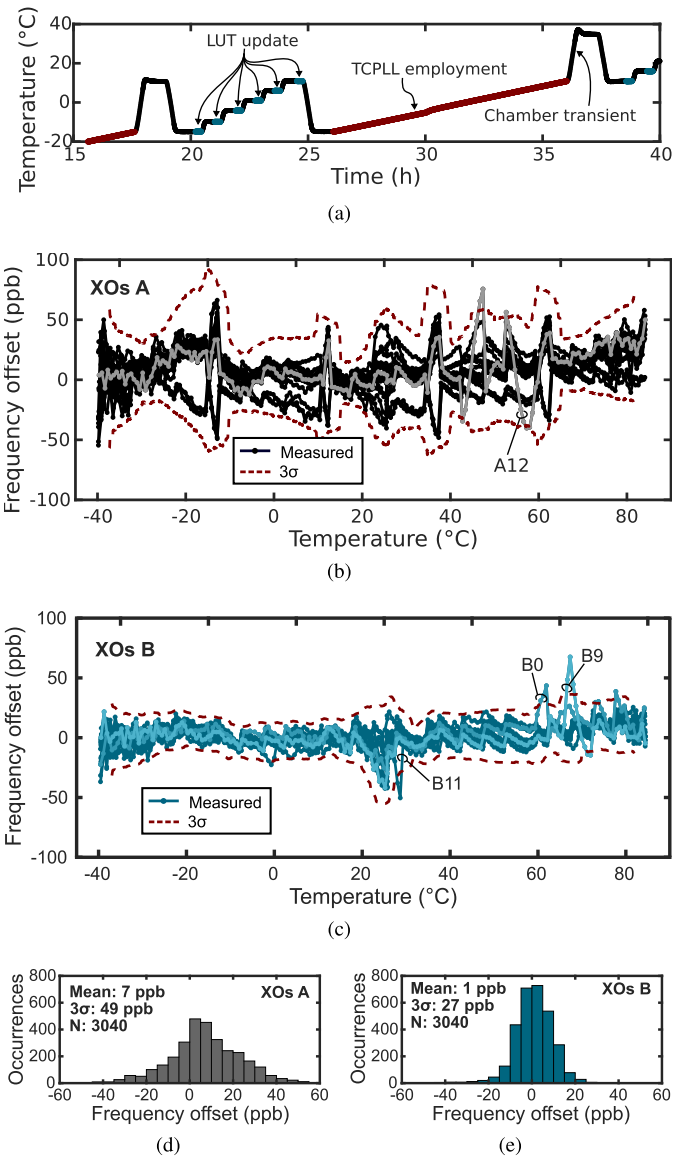

Fig. 15. TCPLL performance with an up-to-date LUT. (a) Measurement transient. (b) Compensated TCPLL frequency offset for XOs A and (c) XOs B. (d) Corresponding distributions for XOs A and (e) XOs B.

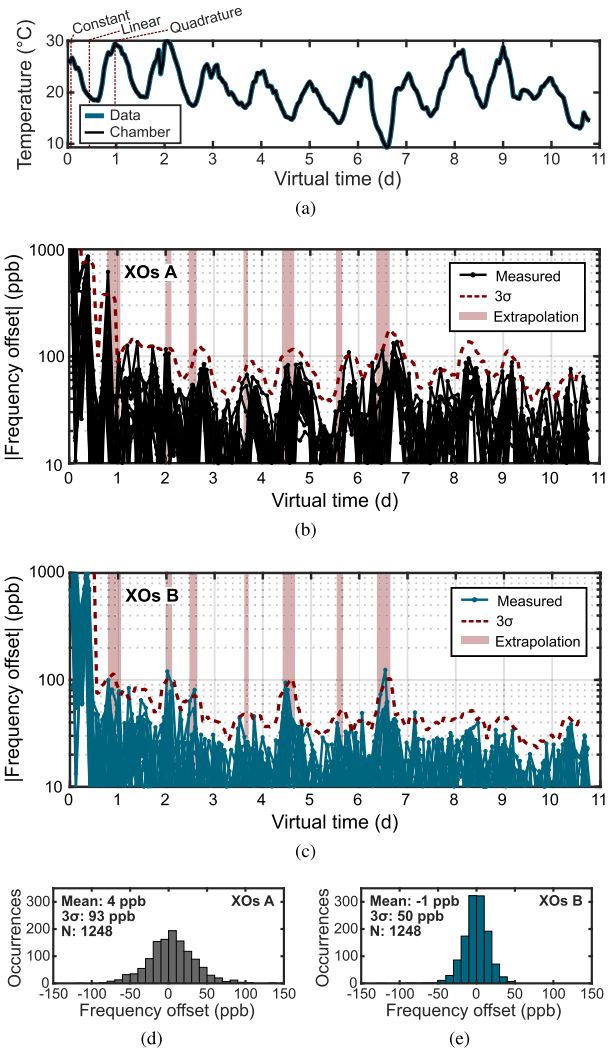

Fig. 16. TCPLL performance with its LUT calibrated in the field. The temperature chamber follows real-life temperature data at a fourfold speed: four virtual days correspond to one measurement day. (a) Measurement transient and approximate time points when the TCPLL switches to using a more accurate interpolator. (b) Compensated TCPLL frequency offset for XOs A and (c) XOs B. (d) Corresponding distributions from 3 d onward for XOs A and (e) XOs B.

at phase 3) of its calibration algorithm. The compensated frequency offset of the TCPLL is presented in Fig.s 16(b)-(c) for XOs A and B, respectively. The time points where κ is computed based on LUT extrapolation instead of interpolation are shadowed in red, and the moving 3σ compensation error averaged over 32 successive trials are represented by the dashed lines. The measurement is started with the TCPLL in a non-calibrated state. Consequently, the compensation error remains high until the LUTs are updated with the first valid ϵ-T pairs. The compensation error drops significantly at around 0.5 d where the TCPLL starts linear interpolation. The error of the linear interpolator is small around room temperature as shown in Appendix B. Gathering more ϵ-T points allows the TCPLL to begin quadrature interpolation at around 1 d. The compensation error level for XOs B settles to around 40 ppb (3σ) in the time span of the next 2 days, which is even smaller than the error level of the frequency detector. This is allowed by LUT normalization and the averaging in algorithm phase 12) with $\zeta = 0.33$. The error level is slightly elevated during extrapolations. As expected, XOs A give a significantly larger error level due to their thermal hysteresis. The distributions of the TCPLL's compensation error starting from the beginning of day 3 for XOs A and B are shown in Fig.s 16(d)-(e), respectively. Despite occasional extrapolation, the TCPLL achieves an accuracy of 93 ppb (3σ) for XOs A and 5 ppb (3σ) for XOs B. In an otherwise identical measurement with $\zeta = 1$ (no LUT averaging), the respective achieved accuracy levels were 120 ppb (3σ) and 72 ppb (3σ).

IX. Comparison With Conventional Methods

The commercial NB-IoT modem in [12] has support for an external TCXO. TCXOs can have an initial accuracy of less than ± 200 ppb but they suffer from an aging rate of up to ± 1 ppm/year [38] and a high cost. Consequently, the NB-IoT modems in [5] and [39] feature instead an

TABLE II
COMPARISON OF XO COMPENSATION METHODS FOR NB-IOT MODEMS

	This work	[8] Perrott JSSC 2013	[7] Tran Sensors J. 2017	[6] Huang TCAS-II 2015
Comp. target	FPLL	MEMS osc.	XO	XO
Comp. method	TCPLL	Sequent. TCPLL	DCXO	VCTCXO
CMOS process (nm)	65	180	180	-
Area overhead‡ (mm^2)	**0.029**†	>0.63	0.35	-
Current overhead‡	**260nA**†	>16mA	-	>380μA×
Tuning range (ppm)	100 / **4000**†	>4000	41	12
Noise penalty	**Low**	High	Low	Low
Osc. Model	LUT+ quad. interp.	5th-deg. polyn.	3rd-deg. polyn.	LUT+ lin. interp.
Cal. points	26	12	-	10
Ram usage (B)	156	54⋆	36⋆	60⋆
Accuracy (ppb)	**27**▷	±200◁	±200◁	±450◁
Temp. range (°C)	−40–85	−40–85	−40–85	−30–60
Meas. samples	16	101	1	1
Field cal.	**Yes**	No	No	No

‡ Temperature sensor omitted.
† Simulated.
× From the specifications [40] of the DAC used in [6].
⋆ Not reported. Estimated according to Section V-A assuming 24-bit precision.
▷ 3σ error.
◁ Reported maximum error.

integrated digitally-controlled crystal oscillator (DCXO) based on a switched-capacitor array. A TCXO can also be founded on a varactor array, named as a voltage-controlled TCXO (VCTCXO). A VCTCXO can be made programmable by using a DAC. A sequential extra FPLL is commonly deployed with MEMS oscillators.

The prototype of this paper is compared with conventional reference compensation methods in Table II. The compared state-of-the-art works also include an oscillator model. They are designed for generic use and do not presume existing circuitry. The works based on an integrated capacitance array [6], [7] suffer from a trade-off between achievable tuning range, tuning resolution and the area of the array. The DAC of the VCTCXO in [6] also consumes both area and power. Using an extra sequential FPLL circumvents the tuning range issue but requires a great amount of power and area. Because of deploying existing circuitry, however, the TCPLL method brings about little power and area overhead. The tuning range of the TCPLL prototype is ultimately limited by the precision of its 24-bit or 32-bit PLL control logic, as shown in Fig. 4. In 32-bit precision, the control logic covers a similar range to [8], but the integrated area and simulated power consumption of the compensation logic increase to about 0.044 mm^2 and 380 nA. All the compensation methods inflict a minor low-frequency noise component that originates from the compensation activity. However, the extra FPLL in [8] comes with a large additional noise penalty, notably induced by the FPLL's phase detector and oscillator [21]. The performance of [7] and [8] would benefit from a more modern process but the resulting power and area overheads would still remain significant.

The TCPLL prototype achieves a high compensation accuracy by combining a high-resolution LUT with the quadrature interpolator. The same accuracy level has been achieved by the MEMS oscillator in the TCXO mode in [9] but with similar sacrifices as in [8]. While the LUT also requires a high number of calibration points, manufacturing costs are not increased since the points can be acquired in the field. Due to its support for field calibration, the TCPLL prototype can compensate for long-term XO frequency drift, which is not taken into account by the conventional methods.

X. CONCLUSION

In this paper, we showed that an NB-IoT modem can deploy its existing FPLL to compensate for the frequency offset of its reference XO. We proposed a temperature compensation logic for the FPLL that remains accurate with down to 24-bit floating-point precision and a simple field calibration algorithm. By utilizing the algorithm, the NB-IoT modem can compose an accurate XO model in the field based on a high-resolution LUT and a quadrature interpolator. We demonstrated a discrete TCPLL prototype that achieved an accuracy of down to 50 ppb (3σ) in an emulated in-field calibration test. The high accuracy of the proposed system allows energy savings in the NB-IoT modem due to low-latency network acquisition and a small number of data transmission repetitions.

APPENDIX A
DERIVATION OF AN EXPRESSION FOR $\Delta K(T)$

Eq. (6) suffers from a high sensitivity to roundoff errors because its critical term $\kappa(T)$ is summed with a significantly larger term 1. The N-th order Taylor series for the term $\frac{1}{1+\kappa}$ about the point $\kappa = 0$ is

$$\frac{1}{1+\kappa(T)} \approx \sum_{n=0}^{N} (-\kappa(T))^n = 1 + \sum_{n=1}^{N} (-\kappa(T))^n . \quad (16)$$

Pairing (6) and (16) yields

$$\begin{aligned}\Delta K(T) &\approx \left[D_0 \left(1 + \sum_{n=1}^{N} (-\kappa(T))^n \right) - N_0 \right] F_0 - K_0 \\ &= \left(\sum_{n=1}^{N} (-\kappa(T))^n \right) D_0 F_0. \end{aligned} \quad (17)$$

This expression is not sensitive to roundoff errors since $\kappa(T)$ is not summed with significantly larger terms. We use the second-order Taylor approximation for $\Delta K(T)$ according to (17), resulting in

$$\Delta K(T) \approx \left(-\kappa(T) + \kappa(T)^2 \right) D_0 F_0. \quad (18)$$

A custom 24-bit floating point implementation of this approximation achieves a frequency error of less than 10 ppb for XOs with a maximum frequency offset of ±100 ppm, as demonstrated in Fig. 4.

APPENDIX B
ERRORS OF A LINEAR AND QUADRATURE INTEPOLATOR

The error of a polynomial interpolator at point x for a real-valued function f that is continuous up to its $(n+1)$-th derivative inside the smallest interval I_x containing the nodes x_0, x_1,..., x_n and x is given by [24]

$$E_n(x) = \frac{\omega_{n+1}(x)}{(n+1)!} f^{(n+1)}(\xi), \tag{19}$$

where $\xi \in I_x$ and $\omega_{n+1}(x)$ is the nodal polynomial of degree $n+1$, defined as

$$\omega_{n+1}(x) = \prod_{i=0}^{n} (x - x_i). \tag{20}$$

We assume that $\epsilon(T)$ follows approximately a third-order polynomial according to

$$\epsilon(T) = \alpha T^3 + \beta T^2 + \gamma T + \delta, \tag{21}$$

where $\alpha - \delta$ are the polynomial coefficients. Applying polynomial fitting to the data of Fig. 5(a) in MATLAB gives coefficients α and β average values of 0.1 ppb/$(°C)^3$ and -8 ppb/$(°C)^2$, respectively. A linear interpolator is defined as

$$\kappa(T) = \epsilon_0 + \frac{\epsilon_1 - \epsilon_0}{T_1 - T_0}(T - T_0), \tag{22}$$

where ϵ_0 and ϵ_1 are the tabled temperature offsets of the XO at temperatures T_0 and T_1, respectively, and $T \in [T_0, T_1]$. Consequently, the error of the linear interpolator ($n = 1$) is

$$E_1(T) = \frac{(T - T_1)(T - T_0)}{2!} \epsilon''(\xi). \tag{23}$$

where $\xi \in [T_0, T_1]$. By solving $\frac{\mathrm{d}E_1(T)}{\mathrm{d}T} = 0$, we find that the maximum absolute error occurs at $T = \frac{T_1+T_0}{2}$. At this point, for an LUT step size of ΔT, T_1 and T_0 can be reformulated as $T_1 = T + \frac{\Delta T}{2}$ and $T_0 = T - \frac{\Delta T}{2}$. The consequent error becomes

$$\begin{aligned} \max_{T_0 \le T \le T_1} |E_1(T)| &= \left| -\frac{\Delta T^2}{4}(3\alpha\xi + \beta) \right| \\ &\approx \left| -\frac{\Delta T^2}{4}(3\alpha T + \beta) \right|, \end{aligned} \tag{24}$$

which is a linear function of temperature. The magnitude of this error is approximately ±100 ppb for a 5-°C LUT around the high and low end of the target temperature range. The magnitude of this error is zero at about 10 °C.

To avoid the significant systematic error of a linear interpolator, we apply instead a quadratic interpolator ($n = 2$) based on Newton's Divided Difference Interpolation in (10), whose error is given by

$$E_2(T) = \frac{(T - T_2)(T - T_1)(T - T_0)}{3!} \epsilon'''(\xi). \tag{25}$$

For simplicity, we assume that the applied LUT is evenly spaced and, consequently, $T_1 = T_0 + \Delta T$ and $T_2 = T_0 + 2\Delta T$. By solving $\frac{\mathrm{d}E_n(T)}{\mathrm{d}T} = 0$, we find that the maximum errors occur at $T = T_0 + \Delta T \left(1 \pm \frac{\sqrt{3}}{3}\right)$ where

$$\max_{T_0 \le T \le T_1} |E_2| = \left| \frac{2\sqrt{3}}{9} \Delta T^3 \alpha \right|, \tag{26}$$

$$\max_{T_1 \le T \le T_2} |E_2| = \left| -\frac{2\sqrt{3}}{9} \Delta T^3 \alpha \right|. \tag{27}$$

The errors are functions of the LUT step size ΔT and polynomial coefficient α but are not functions of temperature. The magnitude of the maximum errors is approximately 5 ppb, assuming a 5-°C LUT resolution.

REFERENCES

[1] W. Yang *et al.*, "Narrowband wireless access for low-power massive Internet of Things: A bandwidth perspective," *IEEE Wireless Commun.*, vol. 24, no. 3, pp. 138–145, Jun. 2017.

[2] *Architecture Enhancements to Facilitate Communications With Packet Data Networks and Applications*, 3GPP document TS 23.682, version 15.4.0, 3rd Generation Partnership Project, Mar. 2018.

[3] *User Equipment (UE) Radio Transmission and Reception*, 3GPP document TS 36.101, version 15.1.0, 3rd Generation Partnership Project, Dec. 2017.

[4] W. Yang *et al.*, "Enhanced system acquisition for NB-IoT," *IEEE Access*, vol. 5, pp. 13179–13191, 2017.

[5] M. Korb *et al.*, "A dual-mode NB-IoT and EC-GSM RF-SoC achieving -128-dBm extended-coverage and supporting OTDOA and A-GPS positioning," in *Proc. IEEE 44th Eur. Solid-State Circuits Conf.*, Sep. 2018, pp. 286-289.

[6] X. Huang, D. Liu, Y. Wang, P. Chen, and W. Fu, "100-MHz low-phase-noise microprocessor temperature-compensated crystal oscillator," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 62, no. 7, pp. 636–640, Jul. 2015.

[7] T. H. Tran, H. W. Peng, P. C. P. Chao, and J. W. Hsieh, "A low-ppm digitally controlled crystal oscillator compensated by a new 0.19-mm^2 time-domain temperature sensor," *IEEE Sensors J.*, vol. 17, no. 1, pp. 51–62, Jan. 2017.

[8] M. H. Perrott *et al.*, "A temperature-to-digital converter for a MEMS-based programmable oscillator with <±0.5-ppm frequency stability and <1-ps integrated jitter," *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 276–291, Jan. 2013.

[9] M. H. Roshan *et al.*, "A MEMS-assisted temperature sensor with 20-μK resolution, conversion rate of 200 S/s, and FOM of 0.04 pJK^2," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 185–197, Jan. 2017.

[10] "Improving the accuracy of a crystal oscillator," Appl. Note, AN1200.07 Rev. 1, Semtech, Jan. 2009.

[11] D. G. Farmer, "Crystal oscillator calibration," U.S. Patent Appl. 20 140 004 887 A1, Jan. 2, 2014.

[12] *Product Specifications*, document nRF9160 v1.0, Nordic Semiconductor, May. 2019.

[13] Ericsson. *Examples of Partners in the IoT Accelerator Ecosystem.* Aug. 30, 2019. [Online]. Available: htpps//:www.ericsson.com/en/Internet-of-things/iot-platform/iot-ecosystem/partners

[14] Y.-P.-E. Wang *et al.*, "A primer on 3GPP narrowband Internet of Things," *IEEE Commun. Mag.*, vol. 55, no. 3, pp. 117–123, Mar. 2017.

[15] A. Adhikary, X. Lin, and Y.-P.-E. Wang, "Performance evaluation of NB-IoT coverage," in *Proc. IEEE 84th Veh. Technol. Conf. (VTC-Fall)*, Sep. 2016, pp. 1–5.

[16] *NB-PSS and NB-SSS Design (Revised)*, document R1-161981, Agenda Item 2.2.5, Qualcomm, 3GPP TSG RAN WG1 NB-IoT Ad-Hoc Meeting, Sophia Antipolis, France, Mar. 2016.

[17] H. Kröll, M. Korb, B. Weber, S. Willi, and Q. Huang, "Maximum-likelihood detection for energy-efficient timing acquisition in NB-IoT," in *Proc. IEEE Wireless Commun. Netw. Conf. Workshops (WCNCW)*, Mar. 2017, pp. 1–5.

[18] J. Zhang, M. M. Wang, and T. Xia, "Practical synchronization waveform for massive machine-type communications," *IEEE Trans. Commun.*, vol. 67, no. 2, pp. 1467–1479, Feb. 2019.

[19] A. Ali and W. Hamouda, "On the cell search and initial synchronization for NB-IoT LTE systems," *IEEE Commun. Lett.*, vol. 21, no. 8, pp. 1843–1846, Aug. 2017.

[20] N. Mysore Balasubramanya, L. Lampe, G. Vos, and S. Bennett, "Low SNR uplink CFO estimation for energy efficient IoT using LTE," *IEEE Access*, vol. 4, pp. 3936–3950, 2016.

[21] C. Quemada, G. Bistue, and I. Adin, *Design Methodology for RF CMOS Phase Locked Loops*, Norwood, MA, USA: Artech House, 2008.

[22] F. L. Walls and J.-J. Gagnepain, "Environmental sensitivities of quartz oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 39, no. 2, pp. 241–249, Mar. 1992.

[23] J. A. Kusters, M. C. Fischer, and J. G. Leach, "Dual mode operation of temperature and stress compensated crystals," in *Proc. 32nd Annu. Symp. Freq. Control*, 1978, pp. 389–397.

[24] A. Quarteroni, R. Sacco, and F. Saleri, *Numerical Mathematics*. New York, NY, USA: Springer-Verlag, 2000.

[25] J. R. Vig and T. R. Meeker, "The aging of bulk acoustic wave resonators, filters and oscillators," in *Proc. 45th Annu. Symp. Freq. Control*, May 1991, pp. 77–101.

[26] *Performance Specification Crystal Units, Quartz General Specification For*, document MIL-PRF-3098H, 27 Aug. 1997.

[27] J. A. Kusters and J. R. Vig, "Hysteresis in quartz resonators–A review," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 38, no. 3, pp. 281–290, May 1991.

[28] L. D. Vittorini, "Micro-jump screening station for GPS user equipment," in *Proc. Int. Freq. Control Symp.*, May 1997, pp. 373–381.

[29] M. E. Frerking, B. W. Paisley, and W. C. Thomas, "Micro-frequency jump design of experiments investigation," in *Proc. IEEE Int. Freq. Control Symp.*, May 1998, pp. 116–120.

[30] M. Koyama, Y. Watanabe, H. Sekimoto, and Y. Oomura, "An experimental study of frequency jumps during the aging of quartz oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 43, no. 5, pp. 907–910, Sep. 1996.

[31] *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology–Random Instabilities*, IEEE Standard 1139, IEEE Standards Coordinating Committee 27, 2008.

[32] K. K. Tuladhar and G. Jenni, "Frequency jumps on BVA and other precision quartz crystal resonators and burst-noise on overtone mode high-frequency quartz crystal resonators," in *Proc. IET 10th Eur. Freq. Time Forum (EFTF)*, 1996, pp. 111–115.

[33] A. V. Kosykh and B. P. Ionov, "Dynamic temperature model and dynamic temperature compensation of crystal oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 41, no. 3, pp. 370–374, May 1994.

[34] S. Iguchi, T. Sakurai, and M. Takamiya, "A low-power CMOS crystal oscillator using a stacked-amplifier architecture," *IEEE J. Solid-State Circuits*, vol. 52, no. 11, pp. 3006–3017, Nov. 2017.

[35] K. Souri, Y. Chae, F. Thus, and K. Makinwa, "12.7 A 0.85V 600nW all-CMOS temperature sensor with an inaccuracy of ± 0.4°C (3σ) from -40 to 125°C," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 222–223.

[36] K. Souri, Y. Chae, and K. A. A. Makinwa, "A CMOS temperature sensor with a voltage-calibrated inaccuracy of ± 0.15°C (3σ) from -55°C to 125°C," *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 292–301, Jan. 2013.

[37] *Espoo Tapiola observation station meteorological data. Distributed by the Finnish Meteorological Institute. Licensed under CC BY 4.0*. Accessed: Aug. 30, 2019. [Online]. Available: https://ilmatieteenlaitos.fi/havaintojen-lataus#!/

[38] *TCXO Specification–Models TL602 TVL602*, document Tx414, Rev. 2, Connor-Winfield, Mar. 2015.

[39] Z. Song, X. Liu, X. Zhao, Q. Liu, Z. Jin, and B. Chi, "A low-power NB-IoT transceiver with digital-polar transmitter in 180-nm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 64, no. 9, pp. 2569–2581, Sep. 2017.

[40] *STM32F103xC Datasheet-Production Data*, document DS5792, Rev. 13, ST Microelectronics, Jul. 2018.

Tuomas Haapala (Student Member, IEEE) received the M.Sc. degree in biotronics from Aalto University, Espoo, Finland, in 2015, where he is currently pursuing the Ph.D. degree under the supervision of Prof. Kari Halonen.

He joined the Department of Electronics and Nanoengineering, Aalto University, in 2011. In 2018, he was a visiting Ph.D. student at the University of Toronto, Toronto, ON, Canada, under the supervision of Prof. Antonio Liscidini. His research interests include low-power analog, digital and RF circuits for the massive IoT.

Antonio Liscidini (Senior Member, IEEE) was born in Tirano, Italy, in 1977. He received the Laurea *(summa cum laude)* and Ph.D. degrees in electrical engineering from the University of Pavia, Pavia, Italy, in 2002 and 2006, respectively.

He was a summer Intern with National Semiconductors, Santa Clara, CA, USA, in 2003, studying poly phase filters and CMOS low-noise amplifiers. From 2008 to 2012, he was an Assistant Professor with the University of Pavia and also a Consultant with Marvell Semiconductors, Pavia, in the area of integrated circuit design. In 2012, he moved to the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada, where he is currently an Associate professor. In 2019, he became a Consultant for Huawei Technology Group in the area of RFIC for optical communication. His research interests are focused on analog mixed signal interfaces, with a particular emphasis on the implementations of transceivers and frequency synthesizers for wireless-wireline communication and ultra-low power applications.

Dr. Liscidini has been a member of the ISSCC TPC (2012–2017), the ESSCIRC TPC (2010–2018), and a member of the CICC TPC (since 2019). He was a recipient of the Best Student Paper Award at the IEEE 2005 Symposium on VLSI Circuits, a co-recipient of the Best Invited Paper Award at the 2011 IEEE CICC, and a Best Student Paper Award at the 2018 IEEE ESSCIRC. He has served as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS (2008–2011) (2017–2018), a Guest Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS in 2013 and 2016, respectively, and a Guest Editor for the IEEE RFIC VIRTUAL JOURNAL (2018). From 2016 to 2018, he has been a Distinguished Lecturer of the IEEE Solid-State Circuits Society.

Kari A. I. Halonen (Member, IEEE) received the M.Sc. degree in electrical engineering from Aalto University (formerly the Helsinki University of Technology), Espoo, Finland, in 1982, and the Ph.D. degree in electrical engineering from Katholieke Universiteit Leuven, Leuven, Belgium, in 1987.

He became the Head of the Electronic Circuit Design Laboratory in 1998, and was appointed as the Head of the Department of Micro- and Nanosciences, Aalto University, from 2007 to 2013. Since 1988, he has been with the Electronic Circuit Design Laboratory, Helsinki University of Technology (since 2011 Aalto University). Since 1993, he has been an Associate Professor with the Faculty of Electrical Engineering and Telecommunications, where he has also been a Full Professor since 1997. He specializes in CMOS and BiCMOS analog and RF integrated circuits, particularly for telecommunication and sensor applications. He has authored or coauthored over 450 international and national conference and journal publications on analog and RF integrated circuits. He was a TPC member of ESSCIRC and ISSCC. He was a recipient of the Beatrice Winner Award of the ISSCC Conference 2002. He has been an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I: REGULAR PAPERS, a Guest Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, and the Technical Program Committee Chairman of the European Solid-State Circuits Conference in 2000 and 2011.