Digitally Controlled Oscillator (DCO)-Based Architecture for RF Frequency Synthesis in a Deep-Submicrometer CMOS Process

Robert Bogdan Staszewski, *Member, IEEE*, Dirk Leipold, Khurram Muhammad, and Poras T. Balsara, *Senior Member, IEEE*

***Abstract*—A novel digitally controlled oscillator (DCO)-based architecture for frequency synthesis in wireless RF applications is proposed and demonstrated. It deliberately avoids any use of an analog tuning voltage control line. Fine frequency resolution is achieved through high-speed $\Sigma\Delta$ dithering. Other imperfections of analog circuits are compensated through digital means. The presented ideas enable the employment of fully-digital frequency synthesizers using sophisticated signal processing algorithms, realized in the most advanced deep-submicrometer digital CMOS processes which allow almost no analog extensions. They also promote cost-effective integration with the digital back-end onto a single silicon die. The demonstrator test chip has been fabricated in a digital 0.13-μm CMOS process together with a DSP, which acts as a digital baseband processor with a large number of digital gates in order to investigate noise coupling. The phase noise is −112 dBc/Hz at 500-kHz offset. The close-in spurious tones are below −62 dBc, while the far-out spurs are below −80 dBc. The presented ideas have been incorporated in a commercial Bluetooth transceiver.**

***Index Terms*—Deep-submicrometer CMOS, digital compensation, digital control, digitally controlled oscillator (DCO), frequency synthesizer.**

I. Introduction

TRADITIONAL designs of commercial frequency synthesizers for multigigahertz mobile RF wireless applications have almost exclusively employed the use of a charge-pump *phase-locked loop* (PLL), which acts as a *local oscillator* (LO) for both a transmitter and a receiver. Unfortunately, the design flow and circuit techniques required are analog intensive and utilize process technologies that are incompatible with a *digital baseband* (DBB). Nowadays, the DBB design constantly migrates to the most advanced deep-submicrometer digital CMOS process available, which usually does not offer any analog extensions and has very limited voltage headroom. Consequently, the aggressive cost and power reductions of high-volume mobile wireless solutions can only be realistically achieved by the highest level of integration, and this favors digitally-intensive approach in the most aggressive deep-submicrometer process.

Deep-submicrometer CMOS processes present new integration opportunities on one hand, but make it extremely difficult to implement traditional analog circuits, on the other. For example, frequency tuning of a low-voltage deep-submicrometer CMOS oscillator is an extremely challenging task due to its highly nonlinear frequency versus voltage characteristics [1] and low-voltage headroom making it susceptible to the power/ground supply and substrate noise. In such low supply voltage case, not only the dynamic range of the signal suffers but also the noise floor rises, thus causing even more severe degradation of the signal-to-noise ratio. At times, it is possible to find a specific solution, such as utilizing a voltage doubler [2]. Unfortunately, with each CMOS feature size reduction, the supply voltage needs also to be scaled down, which is inevitable in order to avoid breakdown and reliability issues. Circuits designed to ensure proper operation of RF amplifiers, filters, mixers, and oscillators depend on circuit techniques that operate best with long-channel, thick-oxide devices with supply voltage of 2.5 V or higher. The process utilized in this paper is optimized for short-channel, thin-oxide devices operating as digital switches at only 1.5 V. In order to address the various deep-submicrometer RF integration issues for frequency synthesis, digitally-intensive techniques need to be developed such that analog imperfections are compensated for by using advanced signal processing techniques.

So far there have not been any reports in literature (except recently in [1]) on the *fully* digital control of oscillators for RF applications. Lack of the fully digital control is a severe impediment for the total integration in a deep-submicrometer CMOS process, without which signal processing algorithms cannot be used to control or drive the analog circuits. There have been several disclosures on ring-oscillator-based DCOs for clock recovery and clock generation applications [3] [4]. However, the frequency resolution is low and spurious tone level is high for these DCOs, which seem to become an effective deterrent against digitally-intensive RF synthesizers for wireless communications. DCO in reference [1] deliberately avoids any analog tuning voltage controls. This allows for its loop control circuitry, including loop filter, to be implemented in a fully digital manner. That DCO, however, provides only a raw and bare minimum of functionality from a signal processing standpoint. This paper introduces a circuitry built around it for the purpose of adding a hierarchical layer of arithmetic abstraction that makes it easier to operate the DCO from higher logical levels in a wide variety of digitally-intensive synthesizer architectures.

The organization of this paper is as follows. Section II gives an overview of the digitally controlled oscillator, whose time-

Manuscript received May 11, 2003; revised July 2003. This paper was recommended by Guest Editor M. Perrott.

R. B. Staszewski, D. Leipold, and K. Muhammad are with the Texas Instruments Inc., Dallas, TX 75243 USA (e-mail: b-staszewski@ti.com).

P. T. Balsara is with the Center for Circuits and Systems, University of Texas at Dallas, Richardson, TX 75083 USA.

Digital Object Identifier 10.1109/TCSII.2003.819128

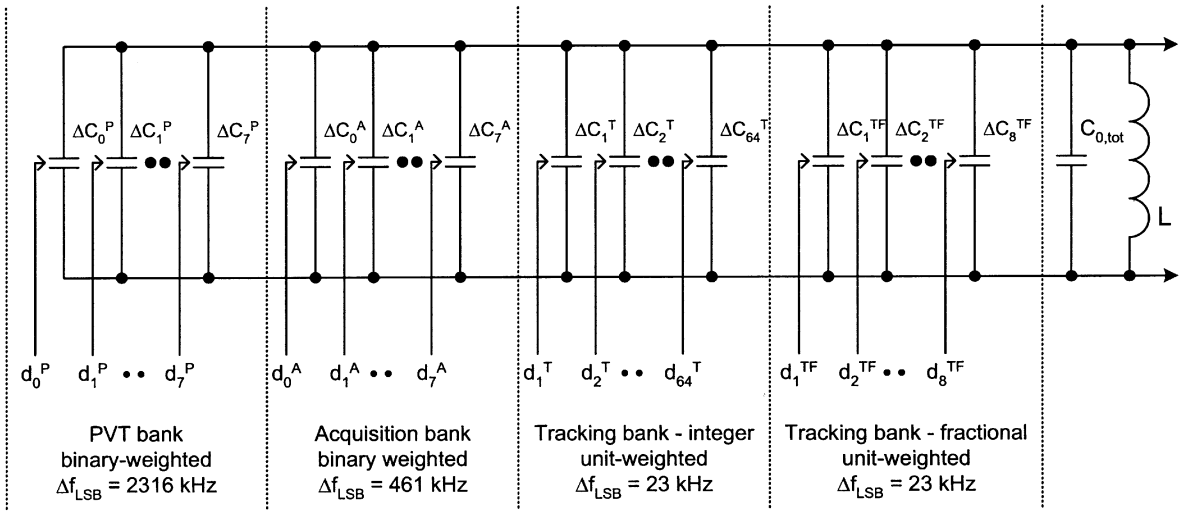

Fig. 1. *LC* tank with dedicated discrete capacitor banks for each of the three operational modes.

domain mathematical model is derived in Section III. Section IV describes the normalized DCO and its time-domain model. The objective there is to suggest possibility of a fully-digital PLL implementation. Section V describes synchronously-timed adjustments of a tuning word to the DCO input. PVT and acquisition interfaces are presented in Section VI, whereas the tracking bit interface is shown in Section VII. An example of a digital PLL architecture possible with the DCO is given in Section VIII. The implementation and measured results are presented in Section IX.

II. Digitally Controlled Oscillator (DCO)

Advanced CMOS process lithography allows nowadays to create extremely small size but well-controlled varactors. The switchable capacitance of the finest differential varactor in this 0.13-μm CMOS process is 38 attofarads, which corresponds to frequency step size of 23 kHz at 2.4 GHz. Frequency tuning of the presented DCO is accomplished by means of a quantized capacitance (with no analog tuning voltage control) of the *LC* tank (Fig. 1) based oscillator [1].

There are three operational modes of the DCO as implemented in the presented IC chip:

- *Process/voltage/temperature (PVT)—calibration mode*: Active during cold power-up and on as-needed basis. Places the nominal center frequency of the DCO in the middle of the Bluetooth band. Possible to use this mode on a regular basis as an ultra-fast acquisition before the regular acquisition mode.
- *Acquisition mode*: Active during channel select.
- *Tracking mode*: Active during the actual transmit and receive. The fractional bits undergo high-speed dithering to increase frequency resolution.

The desired oscillating frequency is acquired step-by-step by traversing through the three capacitor banks with progressively finer resolution. At the end of the PVT and acquisition modes, the terminating-mode capacitor state is frozen and it now constitutes a new center frequency from which the frequency offsets, during the following modes, are referenced.

Fig. 2 demonstrates numerical example of the frequency transversal for the implemented DCO. The acquisition mode starts from the mid-point reset value of 128. The desired center frequency lies two channels or 2-MHz lower from the center channel of the Bluetooth band. This translates to between four and five acquisition steps. As a result, the loop will first quickly move several steps lower. After reaching the point about 2 MHz away, it will dither between the codes of 123 and 124, not being able to resolve any finer. In this example, the transition to the tracking mode happens when the acquisition code is 123. This code stays frozen for the duration of the packet. The tracking mode always starts from the mid-point value of 31. It happens that the desired center frequency is located about 230-kHz higher from that point—this corresponds to 10 tracking steps.

The presented oscillator is built as an ASIC cell (Fig. 3) with truly digital inputs and outputs (even at the RF frequency of 2.4 GHz with rise and fall times specified to be 50 ps) operating in the discrete-time domain, even though the underlying functionality is mainly continuous-time and continuous-amplitude in nature. The RF signal digitizer is a differential-to-single-ended converter that transforms the analog oscillator waveform into the zero-crossing digital waveform with a high degree of common-mode rejection. The digitizer stops the analog nature from propagating up in the hierarchy, right at the interface level. The analog design, modeling and simulation constraints of the system are thus vastly simplified. An analogy could be drawn here to a flip-flop and its fundamental role in the sequential digital circuits, even though its underlying nature and internals are analog. The digital nature of the DCO allows the frequency synthesizer that controls it to be implemented in a fully-digital manner, thus extending scalability, insensitivity to process and environmental variations and general ease of design.

III. Time-Domain Mathematical Model of the DCO

Due to the fact that the conventional RF synthesizers are based on the frequency-domain model, whereas the proposed discrete-time architecture is rooted in time domain, this section introduces basic time-domain equations and modeling concepts that are used in the proposed architecture. It should be noted

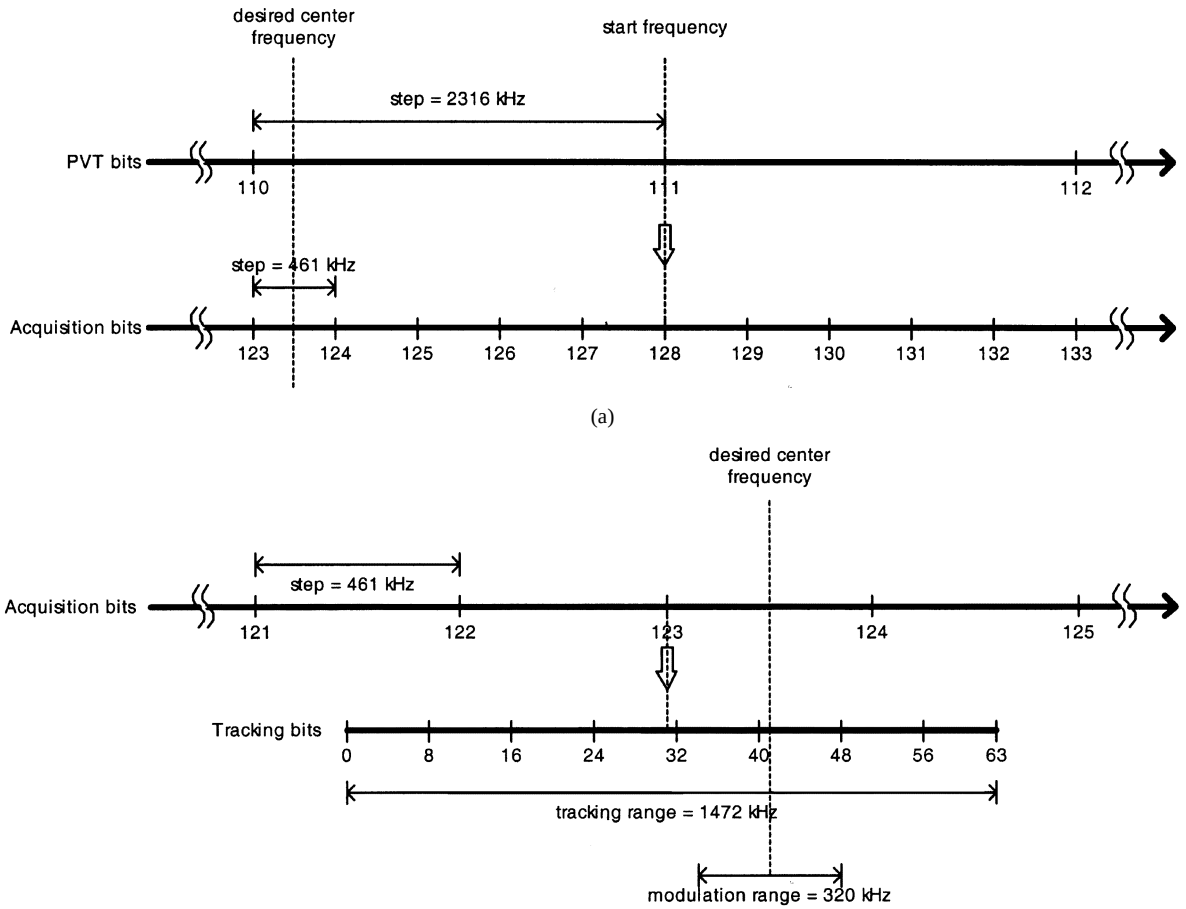

Fig. 2. Frequency transversal example for the implemented DCO. (a) PVT to acquisition. (b) Acquisition to tracking. PVT is calibrated to the middle of the Bluetooth band with code 111.

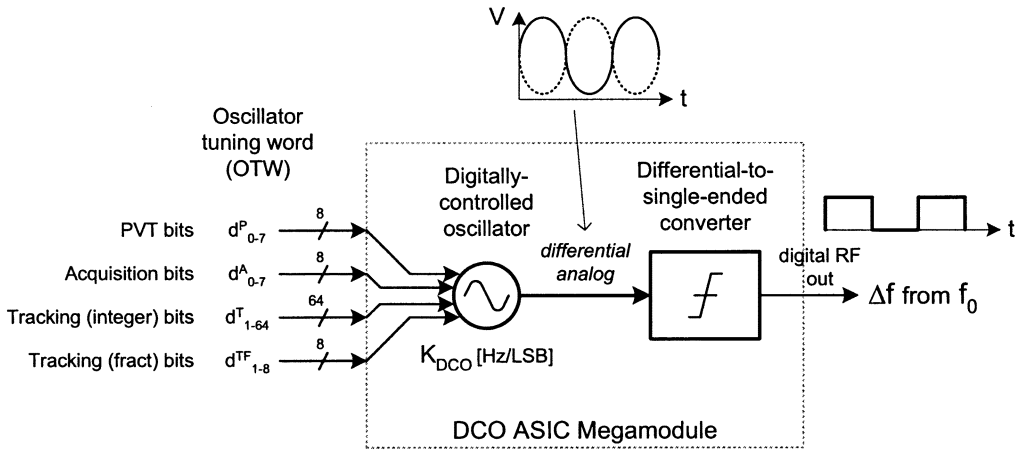

Fig. 3. DCO as an ASIC cell with digital I/Os.

here that recently there have been other attempts to model RF frequency synthesizers in the time domain, such as in [5] for $\Sigma\Delta$ fractional-*N* PLLs.

Let the nominal frequency of oscillation be f_0. It is related to the nominal clock period T_0 by its inverse $f_0 = 1/T_0$. If the clock period is shortened by ΔT, the new clock period will be $T = T_0 - \Delta T$. This will result in a higher frequency of oscillation of $f = f_0 + \Delta f$. Let's determine the relationship between Δf and ΔT. Expanding $f = 1/T$ results in

$$f_0 + \Delta f = \frac{1}{T_0 - \Delta T} = \frac{f_0}{1 - \frac{\Delta T}{T_0}}. \quad (1)$$

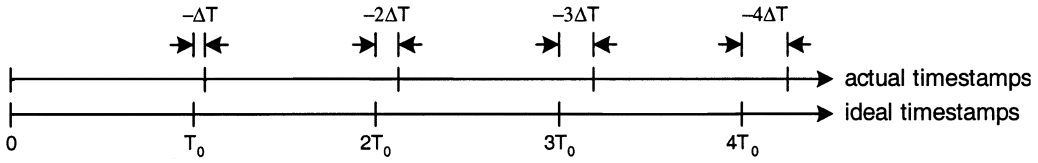

Fig. 4. Development of an accumulative timing deviation (TDEV).

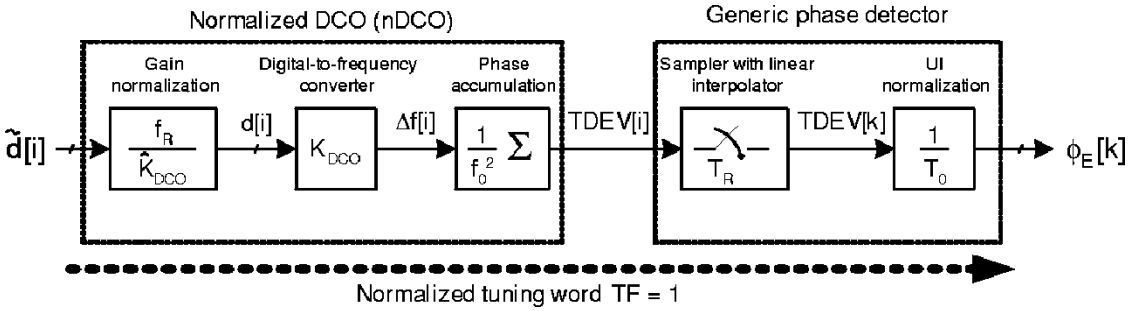

Fig. 5. nDCO time-domain model a with generic phase detection. The complete transfer function is unity. The generic phase detector suggests possibility of a fully-digital PLL implementation.

For $\Delta T/T_0 \ll 1$, using the approximation formula $(1/1-\varepsilon) \approx 1+\varepsilon$, (1) simplifies to

$$\Delta f \approx f_0 \frac{\Delta T}{T_0} = f_0^2 \Delta T = \frac{\Delta T}{T_0^2} \tag{2}$$

(2) was used extensively in this design as a conversion formula for system analysis and simulation. The simulation environment is based on VHDL which, being an event-driven digital simulator, is foreign to the concept of frequency and exclusively operates in the native time domain. As an example of (2), 1 fs of a period deviation will cause 5953-Hz frequency deviation in the middle of the Bluetooth band at 2440 MHz. It is obvious that a fine timing resolution is required at RF frequencies for time-domain simulation tools. In fact, it was necessary to resort to the finest timing resolution of 1 fs that the VHDL standard provides. From a physical viewpoint, a femtosecond time deviation is quite meaningless for a single observation, and only an averaged value could make sense.

For a time-invariant oscillator with a fixed frequency excursion, ΔT period deviation from f_0 will result in ΔT deviation from ideal timing instances within one oscillator clock period, $2\Delta T$ within two clock periods, etc., as shown in Fig. 4. Within i oscillator clock cycles, the accumulated *timing deviation* (TDEV) will reach

$$\text{TDEV}[i] = i \cdot \Delta T = i \cdot \frac{\Delta f}{f_0^2}. \tag{3}$$

For varying ΔT values, (3) could be rewritten as

$$\text{TDEV}[i] = \sum_{l=1}^{i} \Delta T[l] = \sum_{l=1}^{i} \frac{\Delta f[l]}{f_0^2}. \tag{4}$$

Equation (4) states that the TDEV defined as the difference between actual and ideal timing instances is an integral of the oscillator frequency deviation. The ΔT direction is selected toward shortening the period such that ΔT and Δf signs agree.

The *oscillator tuning word* (OTW) at the DCO input (active d^X bits in Fig. 3, where X is P, A or T and TF) will change its operating frequency by $\Delta f[i] = d[i] \cdot K_{\text{DCO}}$. On every rising DCO edge event, the DCO event output Δf multiplied by a "constant" $1/f_0^2$ will be accumulated. At the end of i cycles, it will accumulate the TDEV timing deviation according to (4). The accumulated timing deviation is only defined at the end of the DCO clock cycle with each rising clock edge. The timing deviation is a measure of "badness" and signifies the departure from the desired timing instances that has to be corrected by a feedback loop mechanism.

IV. DCO Gain Normalization and Its Time-Domain Model

At a higher level of abstraction, the DCO oscillator, together with the DCO gain normalization $f_R/\widehat{K}_{\text{DCO}}$ multiplier, comprise the *normalized DCO*. The DCO gain normalization decouples the phase and frequency information throughout the system from the process, voltage and temperature variations that normally affect the DCO gain K_{DCO}, whose estimate is denoted as $\widehat{K}_{\text{DCO}}$. The frequency information is normalized to the value of the external *reference frequency* f_R. The digital input to the nDCO is a fixed-point *normalized tuning word* (NTW), whose integer part LSB bit corresponds to f_R. For additional discussion on the DCO gain normalization issue, see [6].

The time-domain nDCO model is presented in Fig. 5. Provided the DCO gain is estimated correctly, the *normalized tuning word* (NTW) (denoted as $\tilde{d}[i]$) at the nDCO input will change its operating frequency by $\Delta f = \tilde{d}[i] \cdot f_R$. On every rising DCO edge event, Δf multiplied by $1/f_0^2$ will be accumulated. The accumulation interval is established by the inverse of the reference frequency $T_R = 1/f_R$. It is related to the nominal oscillation frequency f_0 by $T_R = N \cdot T_0$, where N is the traditionally-defined (possibly fractional) frequency division ratio. During the T_R interval, the $\tilde{d}[i]$ frequency tuning input is assumed constant. It is justified by the fact that between the reference events there are not any updates to the tuning

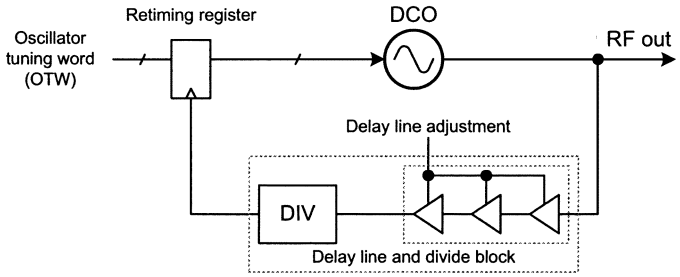

Fig. 6. Synchronously-optimal sampling and timing adjustment of the DCO input.

word. At the end of N cycles, the accumulated timing deviation TDEV will be sampled with value

$$\begin{aligned}\text{TDEV}[k]-\text{TDEV}[k-1] &= N\cdot\Delta T\\ &= N\cdot\frac{\Delta f}{f_0^2}\\ &= \tilde{d}[i]\cdot\frac{N\cdot f_R}{f_0^2}\\ &= \tilde{d}[i]\cdot T_0 \end{aligned} \tag{5}$$

where $k=\lfloor i/N\rfloor$. The samples out of the nDCO are at the DCO rate while the generic phase detector operates at the reference clock. The change of data rate is accomplished using a sampler which employs a first-order interpolation [7].

The DCO phase accumulation is not tied to any hardware. It simply reflects the workings of a progression of time. However, the sampling mechanism requires an explicit hardware that would take periodic snapshots of the evolving TDEV. Time-domain model of the entire normalized DCO is also given in Fig. 5. A generic hardware circuitry is added to the nDCO that would detect TDEV and convert it to the digital bit format. At the same time, it deals with the troublesome unit of time by performing normalization to the clock period of the DCO oscillation, defined as *unit interval* (UI). The diagram does not suggest any particular mechanism, it merely indicates and mathematically describes a timing deviation detector that would determine any frequency and phase deviations of the oscillator which would then be fed back as loop corrections. The transfer function from the normalized tuning word input to the detector output is 1 [bits/bits] within one frequency reference clock cycle. An appropriate digital scaler/filter between ϕ_E and $\tilde{d}$ will give rise to the phase-locked loop (PLL). An example is shown in Section VIII.

V. Synchronously Optimal Oscillator Tuning Word Retiming

Fig. 6 shows a principle of the synchronously-optimal DCO input tuning word retiming method. This idea is based on the observation that changing the tuning control input of an oscillator, in order to adjust its phase/frequency in a normal PLL operation, is quite a disturbing event that reveals itself as jitter or phase noise [8]–[10]. This is especially noticeable in case of a sample-mode oscillator, such as the DCO, where its oscillating frequency is commanded to change at discrete times. Since the oscillating frequency of an *LC* tank is controlled by a voltage-to-capacitance conversion device (e.g., varactor), the instances when the oscillating energy is fully stored in a capacitor are the worst moments to change the capacitance. The total charge must be preserved, so changing the capacitance at those moments causes the electrical potential to exhibit the largest change ($\Delta V = Q/\Delta C$), as shown plot in Fig. 7(a). These perturbations are then AM-to-PM translated by the oscillator circuit into timing jitter. Changing the varactor capacitance at times when it is fully discharged will hardly affect its voltage and thus hardly contribute to the oscillating jitter [Fig. 7(b)].

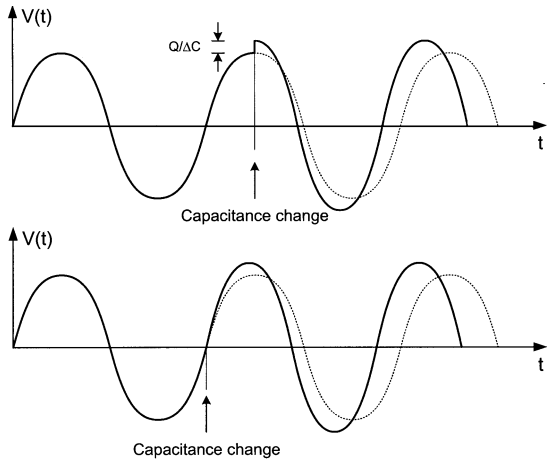

Fig. 7. Waveforms for capacitance change of an *LC*-tank oscillator.

The proposed solution is to control the timing moments when the varactor capacitance change is allowed to occur, thus minimizing jitter due to the tuning word update. This is implemented by feeding the delayed oscillator edge transitions back as the clock input to the *synchronous* register retiming stage, as shown in Fig. 6. The retiming stage ensures that the input control data, as seen by the oscillator, is allowed to change at precise and *optimal* time after the oscillator zero crossings. The feedback loop delay is set algorithmically to minimize the oscillator jitter. The algorithm to control the delay line value for the optimal timing adjustment takes advantage of the fact that the phase error, which is related to the DCO jitter metric to be minimized, is now in digital form and readily available for processing. Various statistics of this signal, such as mean squared error, could be thus optimized by utilizing the digital signal processing hardware. The tightly-integrated companion DSP engine is capable of transferring the digital phase error samples to its own memory at the 13-MHz reference frequency and then to postprocess them.

The actual delay could be accomplished by a voltage-controlled delay line. In the testchip, we chose a long string of inverters while externally controlling their V_{dd} supply voltage. The allowed range of the delay line V_{dd} voltage is quite limited (from 1 to about 1.8 V) so the delay change contribution per inverter is not very significant. However, the delay multiplied by the total number of inverters can exceed the DCO clock cycle, thus guaranteeing the full 360° coverage.

VI. PVT and Acquisition DCO Banks Interface to Digital Logic

Fig. 8 illustrates the interface control structure for the 8-bit binary-weighted acquisition bits (the structure for PVT bits is

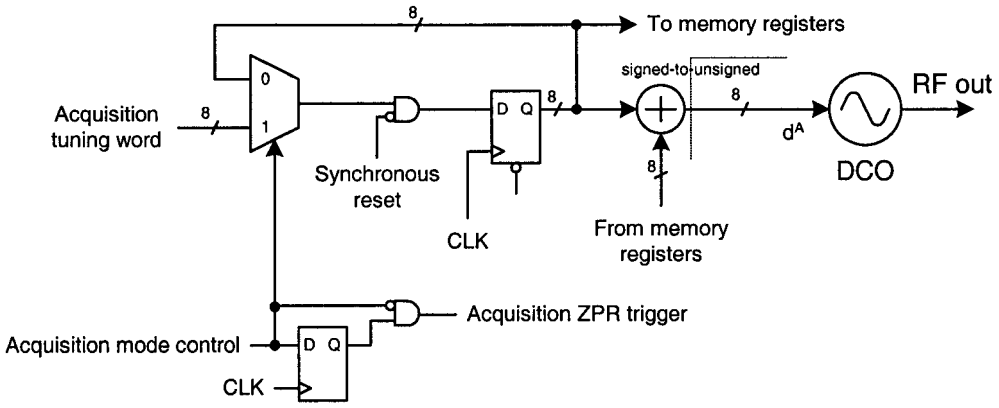

Fig. 8. Control circuit of the oscillator acquisition bits.

almost identical). Their direct digital control is arithmetically expressed as *unsigned* number arguments into the effective resonating capacitance of the *LC* tank. The digital PLL loop control, however, natively operates at an offset frequency with respect to a certain center or "natural" frequency. A conversion mechanism, by simply inverting the MSB bit, is thus only required at the interface point.

The two sets of interface words from memory registers could be the last frequency estimate from the controller's lookup table in order to speed up the loop operation. The two sets of frequency offset status words are reported back to the controller through the memory registers. At reset, the DCO is placed at the center of the operational frequency range through asynchronous clear of the driving registers. This mechanism prevents the oscillator from failing to oscillate if the random power-up values of tuning word registers set it above the oscillating range, which might happen at the slow process corner.

During the active mode of operation, the new tuning word is latched by the register with every clock cycle. Upon the DCO operational bank mode change-over, the last stored value of the tuning word is maintained by the register. Consequently, during the regular operation, only one interface path can be active at a given time, whereas the previously executed modes maintain their final DCO control states. A *zero phase restart* (ZPR) mechanism is used to zero out the phase detector output to avoid any discontinuities in the oscillator tuning word during the mode switchover. A short explanation of the ZPR principle is as follows: At the mode switchover, the tuning word of the last mode corresponds to a certain value of the phase error. This tuning word is now frozen, so the phase error value that maintains it is not longer needed. However, the new mode is always referenced to the new center frequency established by the last mode. Consequently, it operates on the *excess* phase error rather than absolute. Therefore, the old value of the phase error that corresponds to the frozen tuning word of the last mode would have to be constantly subtracted from the new phase error. A better solution is to use the proposed method of zero phase restarting. In this way, a hitless progression through the three DCO operational modes is accomplished.

VII. Tracking DCO Bank Interface to Digital Logic

The PVT and acquisition bits described above are used in the preparatory steps to quickly establish center of the operating frequency. They are inactive during the subsequent normal operation when the settled synthesized frequency is used. The tracking bits of the DCO oscillator, on the other hand, need a much greater care and attention to detail since any phase noise or spurious tone contribution of the tracking bits will degrade the synthesizer performance.

Fig. 9 shows the idea to increase frequency resolution of the DCO. Tracking part of the *oscillator tuning word* (OTW) is split into 6 integer bits and 5 fractional bits. The LSB of the integer part corresponds to the basic frequency resolution of the DCO oscillator. The integer part is thermometer encoded to control the same-size DCO varactors of the *LC*-based tank oscillator. In this scheme, all the varactors are unit weighted but their switching order is predetermined. This guarantees monotonicity and helps to achieve an excellent linearity, especially if their switching order agrees with the physical layout. The transients are minimized since the number of switching varactors is no greater than the code change. This compares very favorably with the binary-weighted control, where an incremental change can cause all the varactors to toggle. In addition, due to equal load throughout for all bits, the switching time is equalized in response to code changes. In this implementation, a slightly more general unit-weighted capacitance control is used to add some extra coding redundancy which lends itself to various algorithmic improvements of the system operation, as described below.

The fractional part, on the other hand, employs a time-averaged dithering echanism to further increase the frequency resolution. The dithering is performed by a digital $\Sigma\Delta$ modulator that produces a high-rate integer stream whose average value equals the lower-rate fractional input. The digital $\Sigma\Delta$ Modulator (SDM) is considered an essential part of the proposed DCO solution for wireless applications. $\Sigma\Delta$ techniques have been used successfully for over two decades in the field of analog data converters. This has developed a rich body of knowledge for other applications to draw upon [11]. A simple first-order SDM pattern [12] is not random at all and is likely to create spurious tones. If the LSB varactor has a frequency resolution of Δf and is dithered at a rate of f_m then it will produce two spurs f_m away on both sides of the oscillating frequency with the power level of $20\log(\beta/2)$ relative to the carrier, where $\beta = (2/\pi)(\Delta f/f_m)$. In this implementation, the step size of the LSB varactors is $\Delta f = 23$ kHz and if the dithering clock is 600 MHz such

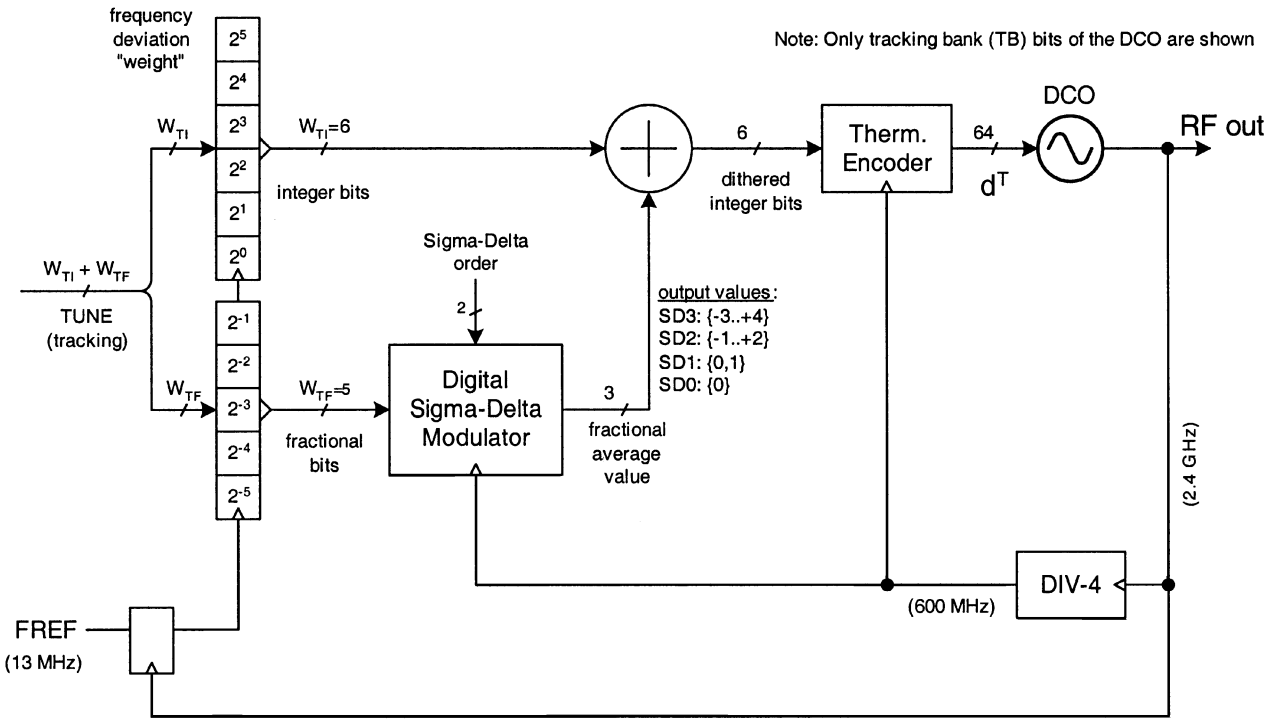

Fig. 9. Improving frequency resolution with $\Sigma\Delta$ dither of DCO tracking bits.

that the maximum dithering rate $f_m = 300$ MHz (for the fractional input of 0.5), then the generated spur level is only -92 dBc. It could be shown through Fourier series decomposition of the modulating wave that for the slowest nonzero dithering rate of $f_m = 18.75$ MHz (corresponding to the fractional input of 1/32) and beyond, the generated spur rises only 4 dB to -88.3 dBc. Even though this level is sufficiently low for most wireless applications, it is far from representing the worst case. First of all, the DCO input, and consequently the f_m dithering rate, is not constant but is subject to a continual change during the normal closed-loop PLL operation thus widely spreading the spur energy. Second, the $\Sigma\Delta$ modulator is normally selected to operate in the second or third order to further randomize the dithering pattern. With such insignificant amount of $\Sigma\Delta$ quantization energy, no phase noise degradation could be observed in our system during normal operation.

The integer part of the tuning word is then added to the integer-valued high-rate-dithered fractional part. The resulting binary signal is thermometer encoded to drive the sixty-four tracking bank varactors. In this simplest embodiment, the high-rate fractional part is arithmetically added to the low-rate integer part thus making its output, as well as the entire signal path terminating at the varactors inside the DCO, high rate. The separate DCO fractional bits d^{TF} of Fig. 1 are not used here. A preferred solution to implement this approach is presented below. It should be noted here that we purposefully left uncorrected a small delay mismatch between the integer and fractional parts due to the SDM group delay. The precise alignment was not necessary since it was determined that the resulting degradation was insignificant due to the high-rate dithering and small amount of quantization energy.

The dithering method trades the sampling rate for the frequency granularity. Here, the frequency resolution of the 2.4-GHz DCO is $\Delta f^T = 23$ kHz with a 13-MHz update rate. Consequently, the effective time-averaged frequency resolution, within one reference cycle, after the 600 MHz $\Sigma\Delta$ dither with five sub-LSB bits would be $\Delta f^{T-\Sigma\Delta} = 23\,\text{kHz}/2^5 = 718$ Hz. The frequency resolution improvement achieved here is $2^5 = 32$. This roughly corresponds to the sampling rate speedup of $600\,\text{MHz}/13\,\text{MHz} = 46.2$.

The structure of the digital $\Sigma\Delta$ modulator is depicted in Fig. 10. It is implemented as a third-order MASH-type architecture [13] that could be conveniently and efficiently scaled down to a lower order. It is clocked by the 600 MHz divided-by-4 oscillator clock. Its topology is based on [12]. The original structure is not the best choice for high-speed designs because the critical path spans through all the three accumulator stages and the carry sum adders. A critical path retiming transformation needed to be performed in order to shorten the longest timing path to only one accumulator so that the 600-MHz clock operation could be reached. Since the structure is highly modular, the lower-order modulation characteristics could be set by disabling the tail accumulators through gating off the clock, which is a preferred method from power saving standpoint.

The combiner circuit merges the three single-bit carry-out streams such that the resulting multi-bit output satisfies the third-order $\Sigma\Delta$ spectral property. The $\Sigma\Delta$ stream equation below is a result of register retiming of the architecture originally described in [12]

$$\text{out}_{\Sigma\Delta} = C_1 \cdot D^3 + C_2 \cdot (D^2 - D^3) + C_3 \cdot (D - 2D^2 + D^3) \quad (6)$$

where $D \equiv z^{-1}$ is the delay element operation. This equation is easily scaled down to the second- or first-order $\Sigma\Delta$ by disregarding the third or third and second terms, respectively.

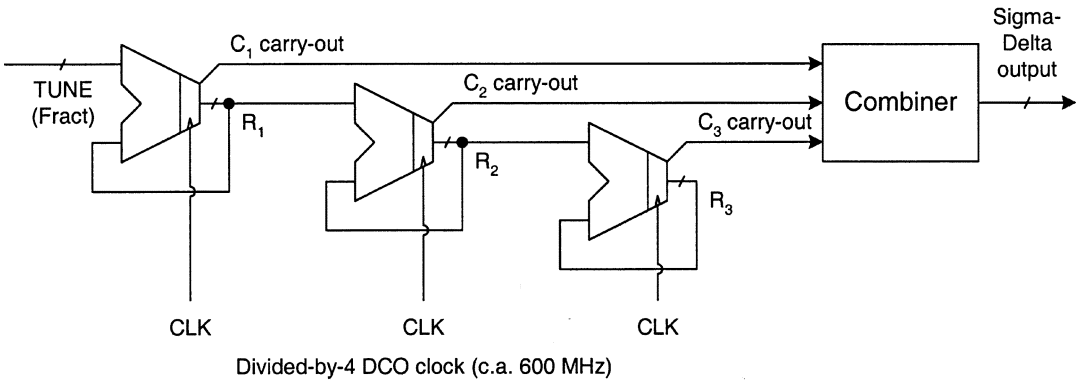

Fig. 10. MASH-3 $\Sigma\Delta$ digital modulator structure.

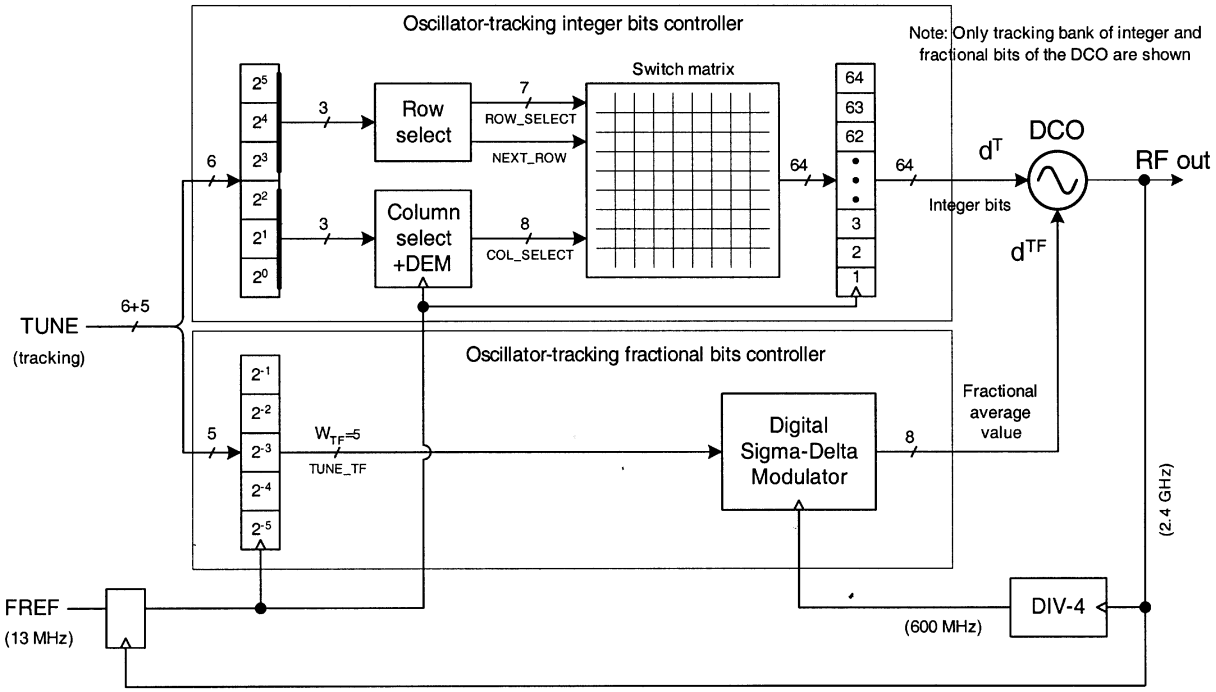

Fig. 11. Implementation block diagram of the DCO tracking bits with DEM of the integer part and $\Sigma\Delta$ dithering of the fractional part. Critical high-speed arithmetic operations are performed in "analog" domain through capacitative additions inside the DCO.

Fig. 11 reveals the preferred method of implementing the integer and fractional oscillator tracking control from a lower power standpoint. The fractional path of the DCO tracking bits, which undergoes high-rate dithering, is entirely separated from the lower-rate integer part. It even has a dedicated DCO input just to avoid "contaminating" the rest of the tracking bits with frequent transitions. The switch matrix, together with the row and column select logic, operates as a binary-to-unit-weight encoder in response to the integer part of the tracking tuning word. The $\Sigma\Delta$ modulator is responsive to only the fractional part of the tracking tuning word. The actual merging of both parts is performed inside the oscillator through time-averaged capacitance summation at the *LC* tank.

Another important benefit of the chosen approach is that the high-speed arithmetic operation of the (6) combiner is now trivial. Fig. 12 shows the proposed implementation. All that is required are flip-flop registers (for the delay operation) with complementary outputs (for the negation). The arithmetic addition is performed inside the oscillator through capacitance summation.

Fig. 13 illustrates the second-order MASH-type $\Sigma\Delta$ modulation of the fixed-point tracking DCO tuning control word with five-fractional bits. The fixed-point tuning word (upper plot) consists of six-integer bits and five fractional bits and is clocked at the 13-MHz reference frequency. The $\Sigma\Delta$ modulates the five-bit fractional part at 600-MHz clock rate and outputs the integer stream that controls the DCO frequency. The lower plot shows the $\Sigma\Delta$ output stream "merged" with the 6-bit integer part stream. For the purposes of visualization only, the integer stream is mathematically decoded into an unsigned number representation and added to the mathematically decoded signed fractional stream.

A. Dynamic Element Matching of the Varactors

Ideally, each of the unit-weighted capacitors of the tracking bank has the exact same capacitative value. Using real-world fabrication process, however, the capacitative value of each capacitor will vary slightly from the ideal. As capacitors are turned

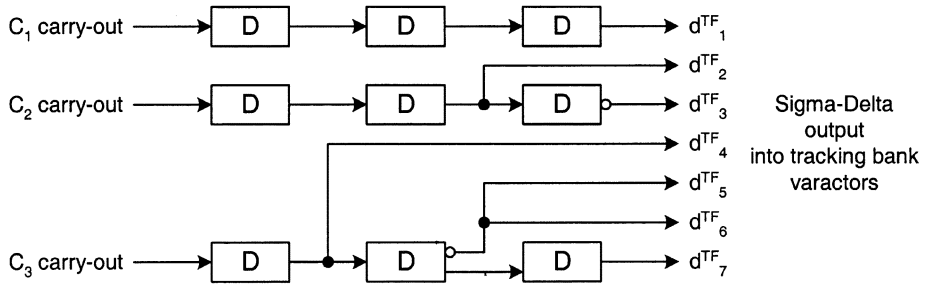

Fig. 12. $\Sigma\Delta$ modulator carry-out combiner structure.

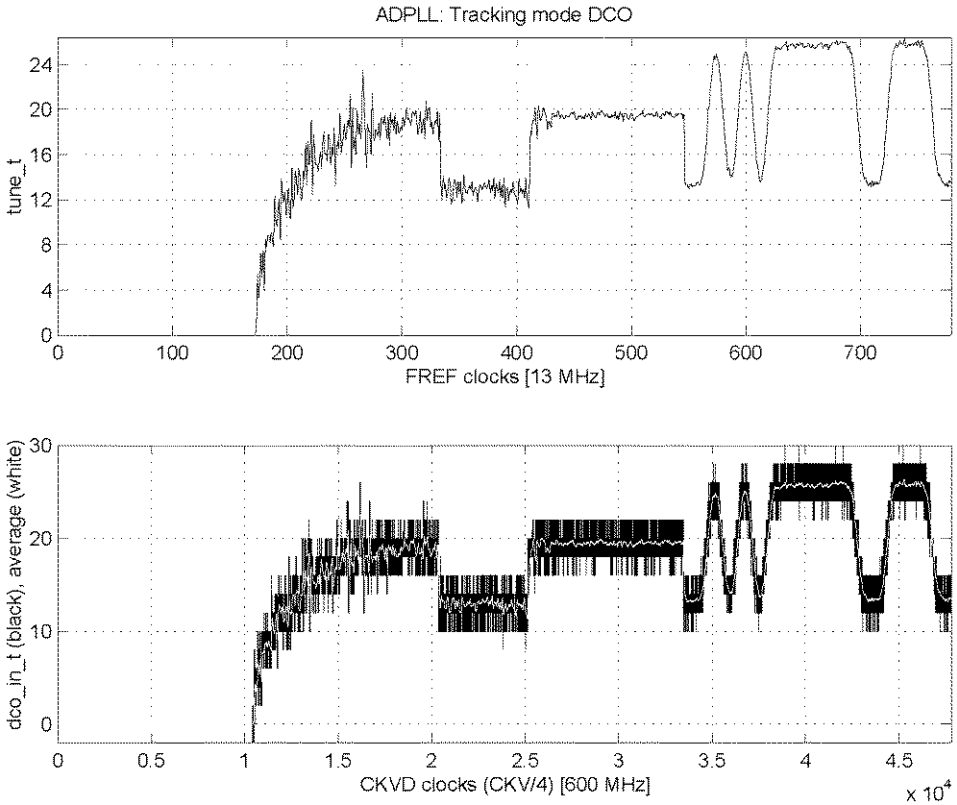

Fig. 13. Simulation plot using the $\Sigma\Delta$ modulation of the fractional part of the tracking tuning word. Top: fixed-point tuning word at 13-MHz frequency reference. Bottom: decoded merged DCO integer input at 600 MHz with running average.

on and off by the integer tracking oscillator controller, nonlinearities will be evident in the output due to variations in capacitative values, as shown in Fig. 14.

A method to improve the digital-to-frequency conversion linearity is also revealed in Fig. 11. It cyclically shifts the unit-weighted varactors using the *dynamic element matching* (DEM) method recently being employed in digital-to-analog converters [14]. The integer part of the tuning word is split into upper and lower bits. The upper bits are encoded and control the row selection of the switch matrix. The lower bits are also encoded and select the next column of the switch matrix. The cyclic shift of unit-weight varactors is performed within the row (see Fig. 15) but could also extend to other rows. However, the number of active switches does not change for the same control input.

In Fig. 15, the capacitors associated with an unfilled row ("next row" signal of Fig. 11) of the switch matrix are rotated on each clock cycle. Initially, the first three columns of row three are enabled. On the next clock cycle, columns two through four, rather than columns one through three are enabled. On the next clock cycle, columns three through five are enabled, etc. Accordingly, on each clock cycle, the set of capacitors used in the

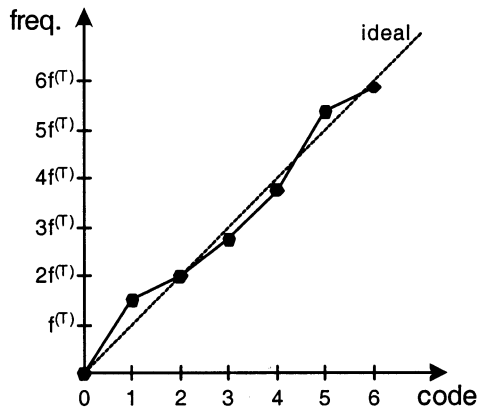

Fig. 14. Cumulative nonlinearity of the DCO tracking bits.

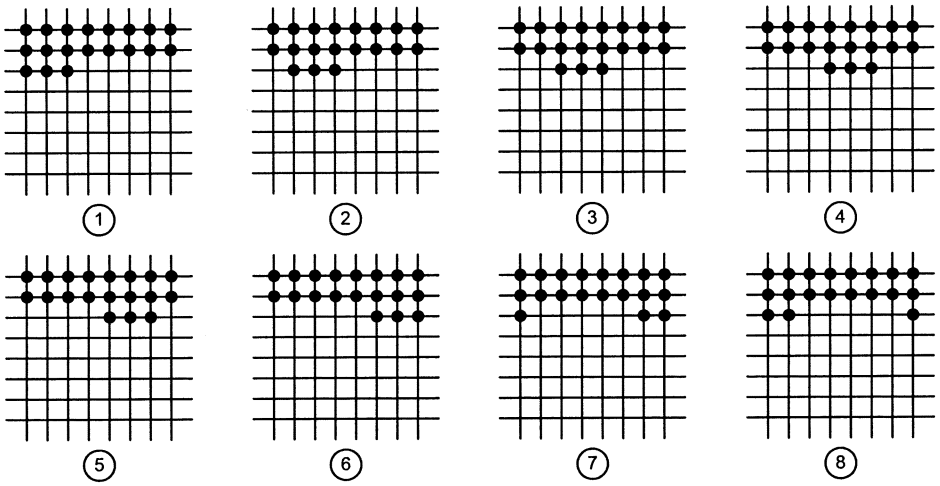

Fig. 15. Dynamic element matching through cyclic shift within a matrix row.

64-element array changes slightly. Over time, the nonlinearities shown in Fig. 14 average out, thereby producing a much more accurate output.

With this DEM scheme, the enabled switches for a single row are rotated. This is accomplished by modulo incrementing the starting column of the enabled switches on each clock cycle. This method could be varied slightly by including two (or more) rows in the rotation. As a result, a larger frequency range would be subject to the beneficial time averaging, by including a greater number of capacitors in the rotation, but at a cost of a longer repetition cycle. An alternative method of increasing the DEM frequency span would be to lengthen the number of columns per row, thus creating a nonsquare matrix.

The output bits of the switch matrix are individually coupled to the bank of sixty-four resampling drivers, which are implemented as flip-flop registers. Each driver controls a single unit-weighted varactor of the *LC* tank. Using resampling by the clock eliminates delay mismatches due to path differences, such that the timing points of varactor transitions coincide. This helps with the spurious noise control. It should be noted that while the switch matrix is shown (from an algorithmic standpoint) in a row/column configuration, the actual implementation is not a precise grid. In fact, a group of rows could be physically combined into a single line.

The principal difference in DFC versus DAC specification requirements is that the full dynamic range is not required for the available number of controlled units. In the DFC application, the frequency headroom is required because it is not expected that the oscillator operates at the precisely specified frequency before entering the tracking mode.

B. DCO Varactor Rearrangement

As illustrated in Fig. 16, the sixty-four integer tracking-bit varactors of the *LC* tank have a physical layout of two long columns and the fractional tracking-bit varactors are arranged separately. However, the controlling circuitry is located only on one side. This creates an unbalanced structure in which routing to one varactor column is significantly shorter with easier access than to the other and, therefore, their transient response is different. Moreover, the spatially separated devices are likely to be more mismatched due to the process gradient. If, during the course of operation, the varactor selection transitions through the column boundary, it is likely to create larger switching perturbations. It is proposed that before entering the finest tracking mode, a rearrangement of DCO varactors to be performed such that "lower quality" capacitors be filled in order to maximize the frequency dynamic range of the most preferred capacitor section. In this architecture, this could be done upon switchover from fast tracking to tracking.

It should be noted that, while in other designs, the tracking capacitors could be arranged differently depending upon various layout issues, a certain set of capacitors will always be favored based on the proximity metric to the control logic.

Fig. 17 illustrates a method of improving the quality of the DFC conversion. In the initial state, half of the capacitors in each column are turned on (as designated by a "+") and half are turned off (as designated by a "−"). During fast tracking, the capacitors of the less desirable right-hand column are enabled or disabled in order to fine tune the oscillator to the selected channel, to the extent possible. If additional capacitors need to be enabled or disabled, the capacitors from the left-hand column may be used, preferably those capacitors at the edges of the column. After channel tuning, the capacitors in the left-hand column are used for modulation and drift control. In this way, the most desirable capacitors are used for maintaining lock and for generating the signal once data is being transmitted.

VIII. Example of a Digital PLL Architecture

The presented DCO-based system architecture allows for the frequency synthesizer to be implemented in a fully-digital manner and places little restrictions on the specific architecture of the PLL loop. For this reason, the phase detector mechanism depicted in Fig. 4 was described in generic terms. Since the focus of this paper is on the DCO and its surrounding circuitry, a brief description of a specific PLL structure used in our work is provided below as a means of further motivation.

The PLL architecture shown in Fig. 18 belongs to a class of PLL frequency synthesizers operating in the phase-domain, which was proposed in [15]. It is a type-I structure that operates in a digital fixed-point phase domain. The variable phase $R_V[i]$

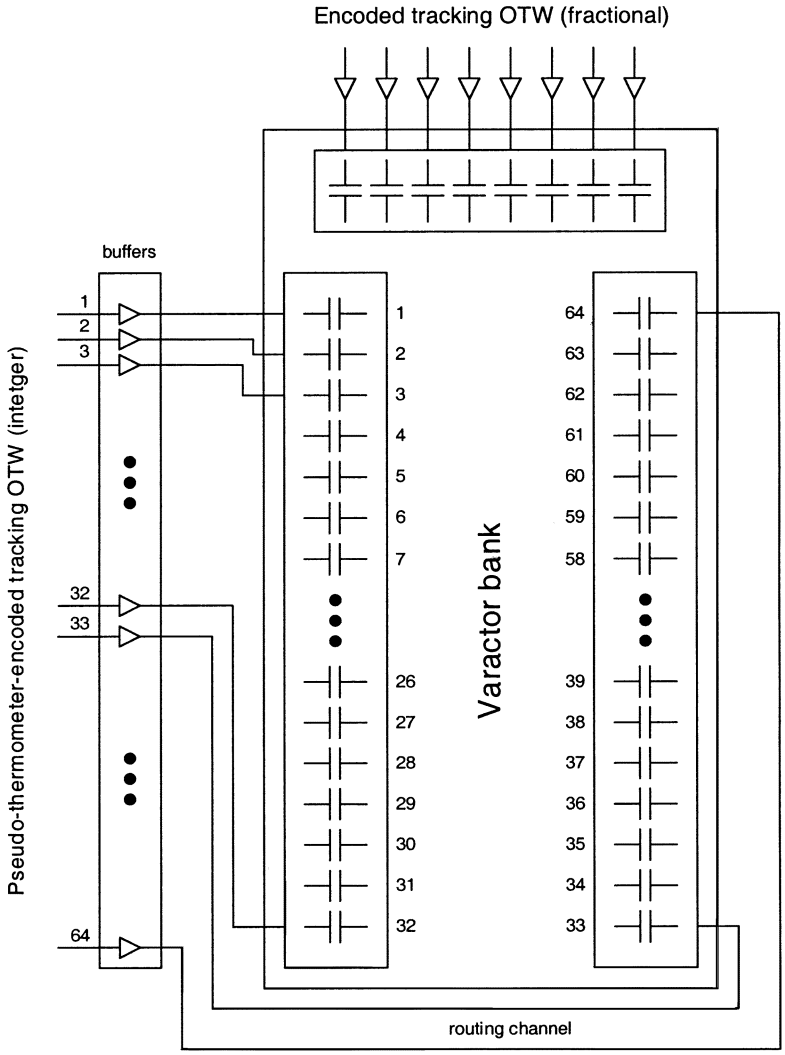

Fig. 16. Layout diagram of the tracking capacitors.

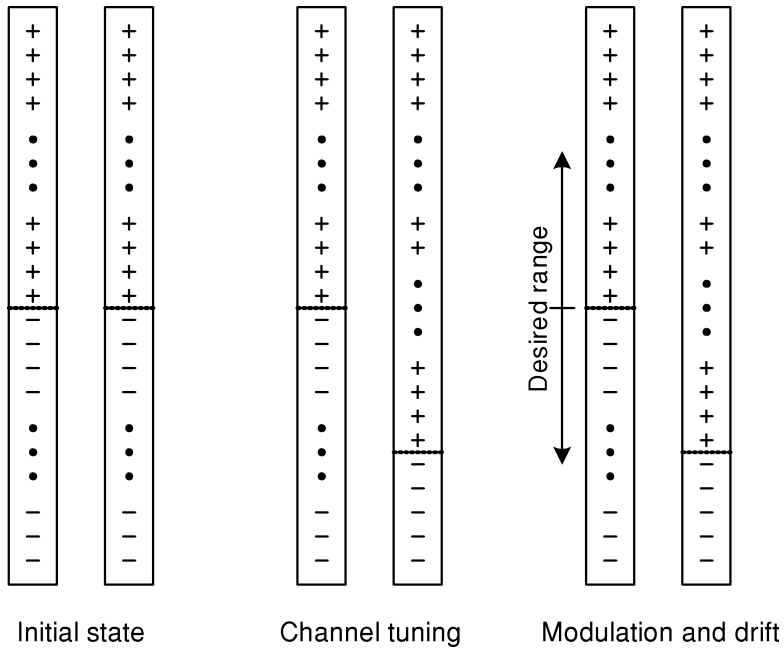

Fig. 17. Tracking capacitors rearrangement after initial settling.

is determined by counting the number of rising-edge clock transitions of the *digitally controlled oscillator* (DCO) clock. It is sampled by the *frequency reference* (FREF) clock and adjusted through linear interpolation [7] for the fractional time difference

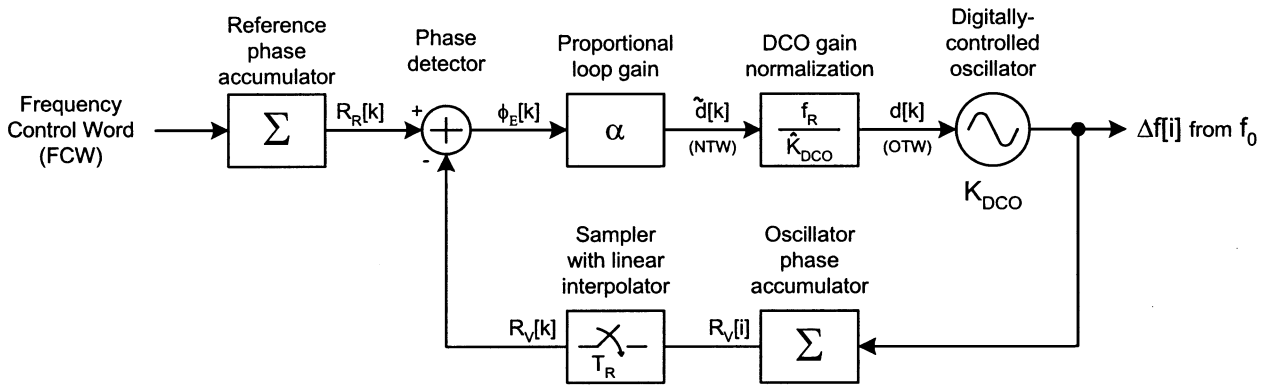

Fig. 18. Digital PLL architecture.

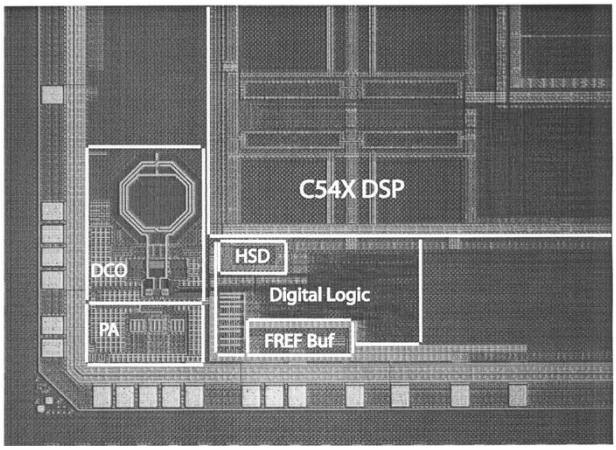

Fig. 19. Micrograph of the RF area residing in the lower-left corner of the chip is shown.

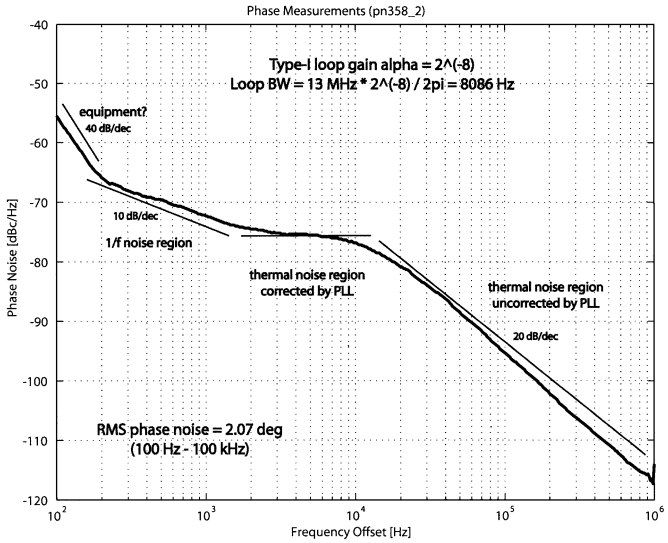

Fig. 20. Measured synthesizer phase noise with the presented DCO: wide loop bandwidth of 8 kHz. Used HP8563E spectrum analyzer with HP85671A phase noise measurement utility.

between the DCO and FREF clock edges. The reference phase $R_R[k]$ is obtained by accumulating the FCW with every rising edge of the FREF clock. The sampled variable phase $R_V[k]$ is subtracted from the reference phase $R_R[k]$ in a digital-arithmetic phase detector. The digital phase error is then multiplied by a proportional loop gain constant α and then normalized by the DCO gain.

The chief advantage of keeping the phase information in fixed-point digital numbers is that, after the conversion, it cannot be further corrupted by noise. Consequently, the phase detector could be simply realized as an arithmetic subtractor that performs an exact digital operation. Therefore, the number of conversion places is kept at minimum.

IX. Experimental Results

Fig. 19 is a die micrograph of the RF frequency synthesizer area. It is located in the lower-left corner and occupies 0.54 mm^2. The *LC*-tank inductor itself occupies a 270 μm $\times$ 270 μm square. High-speed Digital (HSD) running at 600 MHz (divide-by-4 DCO clock) performs the $\Sigma\Delta$ dithering of the DCO varactors. The companion TMS320C54X DSP (used in cellular phones) digital baseband occupies 6 mm^2. This photo dramatically illustrates the high cost (in terms of digital gates) of conventional RF components in high-density modern CMOS processes (this process features 150 k gates per mm^2). Consequently, in order to optimize cost, the number of classical RF components shall be minimized with proper architectural and circuit design choices. The synthesized RF output is buffered to the external pins through a class-E power amplifier (PA), which was chosen due to its digital-friendly characteristics. The DCO core consumes 2.3 mA from a 1.5 V supply and has a very large tuning range of 500 MHz. Its frequency pushing is 600 kHz/V.

Measured phase noise of an all-digital frequency synthesizer with the presented DCO is shown in Fig. 20. The PLL loop forms a type-I first-order structure with the 3-dB loop bandwidth of 8 kHz. Fig. 21 reveals spurious tones emitted during locked operation in the wide span of 1.5 GHz when the DSP is turned on. The close-in spurs lie below -62 dBc and are due to suboptimal layout of power and ground lines causing excessive FREF coupling. The spur level could be even lower if the layout is improved. The far-out spurs are vanishingly small and well below the floor of -80 dBc. Key measured performance parameters are summarized in Table I. The overall performance easily meets the Bluetooth spec.

By scanning the DCO feedback delay line adjustment voltage (Fig. 6), a 1–2 dB variation in phase noise was observed. Since

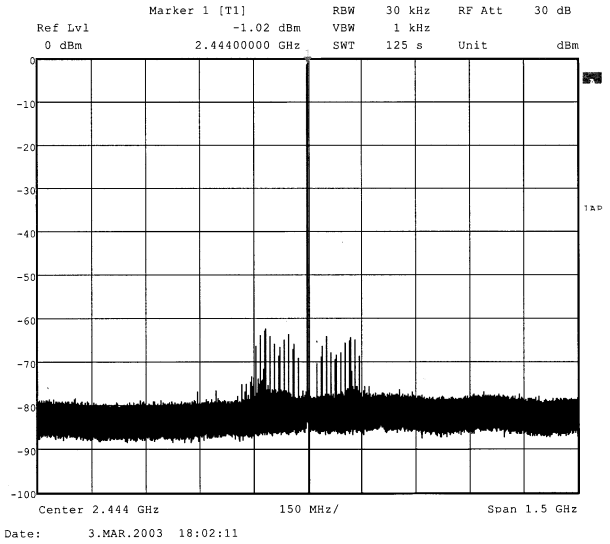

Fig. 21. Spurious tones generated by a synthesizer with the presented DCO using $\Sigma\Delta$ varactor randomization. Measured with Rohde&Schwarz (R&S) FSIQ-7 signal analyzer.

TABLE I
MEASURED KEY SYNTHESIZER PERFORMANCE

Phase noise	$\leq$-112 dBc/Hz @500 kHz offset
RMS phase error	2.1 deg
Close-in spurious tones	$\leq$-62 dBc
Far-out spurious tones	$\leq$-80 dBc
Settling time	$\leq$50 μs

the oscillator performance significantly exceeds the requirements of the targeted Bluetooth specification, the inaccuracy in the DCO timing adjustments was considered inconsequential to warrant added software complexity. The presented results therefore use a nominal value of the adjustment voltage. This tradeoff should be revisited for more challenging standards, such as GSM.

X. CONCLUSION

A novel *digitally controlled oscillator* (DCO)-based system architecture for wireless RF applications has been proposed and demonstrated. This enables to employ fully-digital frequency synthesizers in the most advanced deep-submicrometer CMOS processes with almost no analog extensions. It allows cost-effective integration with the digital back-end onto a single silicon die. The proposed combination of various circuit and architectural techniques has brought to fruition a fully digital solution that has a fine-frequency resolution with low-spurious content and low phase noise. A 2.4-GHz DCO and its peripheral circuitry have been fabricated in a digital 0.13 μm CMOS process and integrated with a DSP in order to investigate noise coupling. This paper demonstrates feasibility and attractiveness of the digitally controlled oscillator in a digital synthesizer architecture for RF multigigahertz applications.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewer for constructive comments that helped to improve quality of the paper.

REFERENCES

[1] R. B. Staszewski, D. Leipold, C.-M. Hung, and P. T. Balsara, "A first digitally controlled oscillator in a deep-submicron CMOS process for multi-GHz wireless applications," in *Proc. 2003 IEEE Radio Frequency Integrated Circuits (RFIC) Symp.*, sec. MO4B-2, June 2003, pp. 81–84.

[2] G. K. Dehng, C. Y. Yang, and J. M. Hsu *et al.*, "A 900-MHz 1-V CMOS frequency synthesizer," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1211–1214, Aug. 2000.

[3] J. Dunning, G. Garcia, J. Lundberg, and E. Nuckolls, "An all-digital phase-locked loop with 50-cycle lock time suitable for high performance microprocessors," *J. Solid-State Circuits*, vol. 30, pp. 412–422, Apr. 1995.

[4] R. E. Best, *Phase Locked Loops: Design, Simulation and Applications*, 3rd ed. New York: McGraw-Hill, 1997.

[5] M. H. Perrott, M. D. Trott, and C. G. Sodini, "A modeling approach for Σ-Δ fractional-*N* frequency synthesizers allowing straightforward noise analysis," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1028–1038, Aug. 2002.

[6] R. B. Staszewski, D. Leipold, and P. T. Balsara, "Just-in-time gain estimation of an RF digitally controlled oscillator for digital direct frequency modulation," *IEEE Trans. Circuits Syst. II*, vol. 50, pp. 887–892, Nov. 2003.

[7] F. M. Gardner, "Interpolation in digital modems—part I: fundamentals," *IEEE Trans. Commun.*, vol. 41, pp. 501–507, Mar. 1993.

[8] T. H. Lee and A. Hajimiri, "Oscillator phase noise: a tutorial," *IEEE J. Solid-State Circuits*, vol. 35, pp. 326–336, Mar. 2000.

[9] A. Hajimiri and T. H. Lee, "A general theory of phase noise in electrical oscillators," *IEEE J. Solid-State Circuits*, vol. 35, pp. 326–336, Feb. 1998.

[10] ——, *The Design of Low Noise Oscillators*. Norwell, MA: Kluwer, 1999.

[11] J. C. Candy and G. C. Temes, "Oversampling methods for A/D and D/A conversion," in *Oversampling Delta-Sigma Data Converters*. Piscataway, NJ: IEEE Press, 1991.

[12] B. Miller and R. J. Conley, "A multiple modulator fractional divider," *IEEE Trans. Instrum. Meas.*, vol. 40, pp. 578–583, June 1991.

[13] Y. Matsua *et al.*, "A 16-bit oversampling A/D conversion technology using triple integration noise shaping," *IEEE J. Solid-State Circuits*, vol. SC-22, pp. 921–929, Dec. 1987.

[14] R. E. Radke, A. Eshraghi, and T. S. Fiez, "A 14-bit current-mode Σ-Δ DAC based upon rotated data weighted averaging," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1074–1084, Aug. 2000.

[15] A. Kajiwara and M. Nakagawa, "A new PLL frequency synthesizer with high switching speed," *IEEE Trans. Veh. Technol.*, vol. 41, pp. 407–413, Nov. 1992.

Robert Bogdan Staszewski (M'94) received the BSEE (summa cum laude), the M.S.E.E., and Ph.D. degrees from the University of Texas, Dallas in 1991, 1992, and 2002, respectively.

From 1991 to 1995, he was with Alcatel Network Systems, Richardson, TX, where he worked on Sonnet cross-connect systems. He joined Texas Instruments, Dallas, TX, in 1995 where he is currently a Senior Member of the Technical Staff. At Texas Instruments, he has been engaged in advanced CMOS read channel development and since 1999, he has been involved in RF CMOS synthesizer design for short-distance wireless and cellular phones. His research interests include RF transceivers, frequency synthesizers, high-speed and low-power digital circuits and system implementation in a deep-submicrometer CMOS process.

Dirk Leipold received the Diploma and Ph.D. degrees in physics from the University of Konstanz, Germany, in 1991 and 1995, respectively.

From 1991 to 1995, he was with the Paul Scherrer Institute, Zurich, Switzerland, where he worked on Smart pixel optoelectronics. In 1995, he joined Texas Instruments, Germany, where he was involved in RF process integration, device characterization, and modeling, and in the development of RF-CMOS technologies on high-resistivity substrates. From 1998 to 1999, he represented Texas Instruments in the ETSI Hiperlan2 Committee, where he was Editor for the PHY layer technical specification. In 1999, he moved to Texas Instruments, Dallas, where he is currently Manager of the Digital Radio Processor Architecture Group. His research interests include advanced RF architectures, nanometer scale CMOS process, and quantum electronics.

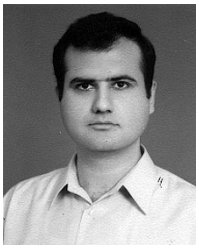

Khurram Muhammad received the B.Sc. degree from the University of Engineering and Technology, Lahore, Pakistan, in 1990, the M. Eng. SC. degree from the University of Melbourne, Parkville, Victoria, Australia, in 1993, and the Ph.D. degree from Purdue University, West Lafayette, IN in 1999.

From 1990 to 1991 and 1993 to 1994, he worked with Carrier Telephone Industries, Islamabad, Pakistan, in Research and Development where he was involved in the development of board level designs for rural and urban telecommunication projects. From 1994 to 1995, he was a Research Associate, Department of Electronic Engineering, Ghulam Ishaq Khan Institute of Engineering Science and Technology. In 1995, he was with Hong Kong University of Science and Technology, Kowloon, where he developed fast simulation techniques for DS/CDMA systems in multipath fading. Since 1999, he has been working at Texas Instruments Inc., Dallas.

Poras T. Balsara (S'83–M'85–SM'95) received the L.E.E. diploma in electronics from The Victoria Jubilee Technical Institute, Bombay, India in 1980, the B.E. degree from the University of Bombay, India, in 1983, and the M.S. and Ph.D. degrees in computer science from Pennsylvania State University, University Park, in 1985 and 1989, respectively.

Currently, he is a Professor of Electrical Engineering with the Erik Jonsson School of Engineering and Computer Science, The University of Texas at Dallas, Richardson. His research interests include, VLSI design, design of energy efficient digital circuits and systems, computer arithmetic, application-specific architecture design, and reconfigurable computing. He has published several journal and conference papers in these areas.