An Agile ISM Band Frequency Synthesizer with Built-In GMSK Data Modulation

Norman M. Filiol, Thomas A. D. Riley, *Member, IEEE*,
Calvin Plett, *Member, IEEE*, and Miles A. Copeland, *Fellow, IEEE*

***Abstract*— In this paper, a high-resolution fractional-N RF frequency synthesizer is presented which is controlled by a fourth-order digital sigma–delta modulator. The high resolution allows the synthesizer to be digitally modulated directly at RF. A simplified digital filter which makes use of sigma–delta quantized tap coefficients is included which provides built-in GMSK pulse shaping for data transmission. Quantization of the tap coefficients to single-bit values not only simplifies the filter architecture, but the fourth-order digital sigma–delta modulator as well.**

The synthesizer makes extensive use of custom VLSI, with only a simple off-chip loop filter and VCO required. The synthesizer operates from a single 3-V supply, and has low power consumption. Phase noise levels are less than −90 dBc/Hz at frequency offsets within the loop bandwidth. Spurious components are less than −90 dBc/Hz over a 19.6-MHz tuning range.

***Index Terms*— Frequency synthesizer, GMSK modulation, sigma–delta modulation, sigma–delta quantized tap coefficients, radio transmitters.**

I. Introduction

FREQUENCY synthesizers for use in mobile radio applications must be well suited for VLSI implementation and have low power consumption. In addition to providing a fixed local oscillator (LO) in a receiver, a frequency synthesizer may be called on to be agile in its frequency control, for instance, in spread-spectrum applications. Synthesizers used in such applications must have fast settling times and good spurious performance. Direct modulation capability is also an asset.

It has been shown that sigma–delta controlled fractional-N frequency synthesis can meet these requirements [1]–[3]. Unlike direct digital synthesis (DDS), no up-conversion is required, and the modulated signal can be synthesized directly at the RF frequency with comparable performance. The high resolution achieved allows accurate continuous-phase constant-amplitude modulation of the carrier at RF frequencies. Such a synthesizer is shown in Fig. 1.

In order to transmit digital data efficiently, some form of baseband pulse shaping is required to control the RF bandwidth [4]. This pulse shaping can be achieved using a digital filter. A method has been presented whereby the tap coefficients for such a filter could be quantized to single-bit values, using a software sigma–delta modulator, prior to storage in a read-only memory (ROM) [5] with the potential for greatly reducing the required storage space and simplifying the multiplications required to single-bit operations. Although the work presented in [5] showed the feasibility of this method, the demonstration was partly done at the software level, and a hardware modulator was never combined with an integrated sigma–delta controlled synthesizer. A recent synthesizer presented in [3] allows high-speed modulation outside the loop bandwidth through the use of a preemphasis filter, but the data filter was not integrated along with the synthesizer. In this paper, it will be shown, using experimental results, that through careful design and consideration, a low-power transmitter can be built by combining the above two techniques. Also, an error in the work presented in [5] has been corrected. The filter response stored in the ROM should have been that of a Gaussian pulse convolved with a square pulse. Instead, the nonconvolved pulse was used. The design makes use of custom VLSI with only a few external components required. Many of the tradeoffs and issues that were addressed in order to ensure good performance and low-power consumption are explained in detail. It will also be shown that careful design using ECL/CML logic can result in low power consumption at high operating frequencies.

Before presenting the synthesizer design in detail, it is useful to review sigma–delta controlled fractional-N frequency synthesis and MASH architecture sigma–delta modulators.

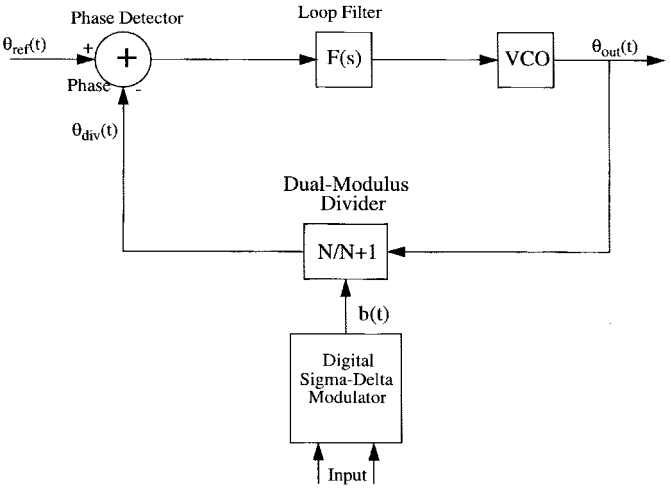

Fig. 1. Sigma–delta modulator controlled phase-locked loop.

Manuscript received November 10, 1997; revised February 5, 1998. This work was supported in part by the Telecommunications Research Institute of Ontario, Micronet (a Federal Centre of Excellence), and by the Canadian Microelectronics Corporation.

N. M. Filiol, C. Plett, and M. A. Copeland are with the Department of Electronics, Carleton University, Ottawa, Ont., K1S 5B6 Canada.

T. A. D. Riley was with the Department of Electronics, Carleton University, Ottawa, Ont., K1S 5B6 Canada. He is now with the University of Oulu, Oulu, Finland.

Publisher Item Identifier S 0018-9200(98)03431-3.

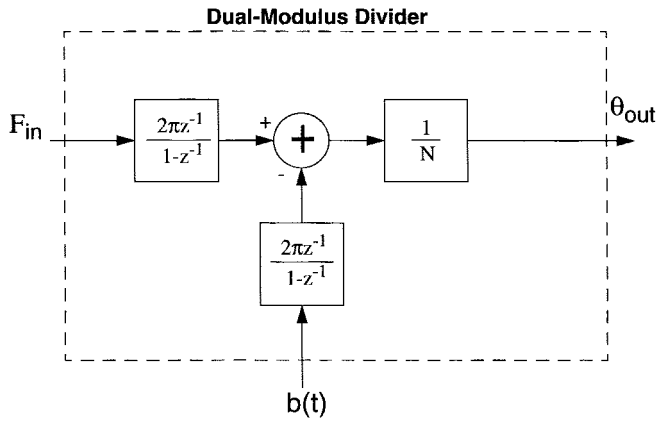

Fig. 2. Model of dual-modulus divider.

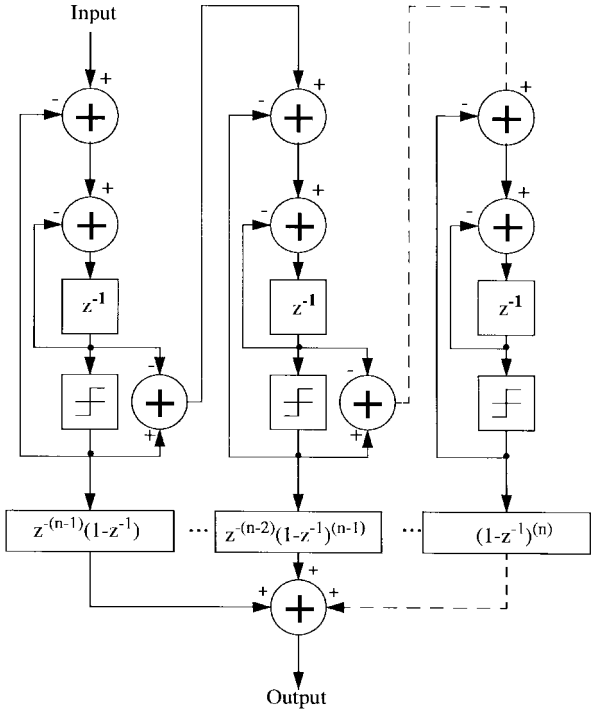

Fig. 3. MASH architecture sigma–delta modulator.

II. Sigma–Delta Controlled Fractional-N Frequency Synthesis

The synthesizer shown in Fig. 1 operates as follows. The desired output frequency is fed into a digital sigma–delta modulator as a digital word. The resulting output bitstream $b(t)$ is used to control the dual-modulus divider in the PLL. The average value of the bitstream is the division ratio required to synthesize the desired output frequency. The noise shaping and oversampling of the modulator push the quantization noise to higher frequencies. This high-frequency noise is then filtered out by the PLL.

To understand the operation of the dual-modulus divider in Fig. 1, consider a phase-locked loop (PLL) with a fixed divide by N in the feedback loop. We could postulate a digital-to-phase (D/P) converter which would allow a fixed amount of phase to be added or subtracted from the divider output every reference cycle, providing digital control of the output frequency. The amount of phase added or subtracted would depend on the digital input to the D/P. The output of this D/P could be seen to quantize phase in the same way that a D/A converter quantizes voltage. The resolution achieved would depend on the number of phase quantization levels and their linearity. Building an accurate multibit phase quantizer is a difficult design problem. However, single-bit phase quantizers are readily available in the form of dual-modulus dividers. The input to the divider can be seen as a single control bit which allows 0 or 2π rad of phase (0 or 1 period of the VCO in Fig. 1) to be subtracted once every reference cycle from the input signal. This subtraction is followed by a fixed divide by N. This divider model is shown in Fig. 2. The sampled data integrators shown represent conversion from frequency to phase. Using oversampling and noise-shaping techniques, high resolution can be achieved using this single-bit phase quantizer in the same way that high resolution is obtained from single-bit amplitude quantizers [6], [7]. Thus, the large knowledge base available on the design of sigma–delta modulators is also applicable to the design of high-resolution frequency synthesizers.

III. MASH Architecture Sigma–Delta Modulators

Sigma–delta modulators achieve high resolution from a single-bit quantizer through the use of noise-shaping and oversampling techniques. Higher order modulators have fewer limit cycle tones, and higher in-band signal-to-noise ratios. When designing higher order sigma–delta modulators, stability becomes a concern due to high-order feedback around the loop, and the design can be quite challenging. An alternative to this approach is to use a MASH architecture [8].

A MASH architecture sigma–delta modulator is shown in Fig. 3. In this case, it consists of a cascade of first-order sigma–delta modulators. The quantization error of each stage is fed forward to the next stage, whose output bitstream is a sigma–delta quantized estimate of the error from the previous stage. The outputs are then combined in a noise-shaping block which cancels the noise from the first $n-1$ stages, producing a multibit output which has nth-order noise shaping given as

$$N(z) = (1 - z^{-1})^n. \tag{1}$$

The advantage of this modulator architecture is that it is unconditionally stable since no nth-order feedback is present and the first-order stages have been proven to be stable [8]. The main disadvantage is that the output is n bits wide, making a multibit decimator necessary in data converter applications and a multimodulus divider necessary in PLL applications, as here.

Analog implementations of MASH architectures suffer from gain mismatches between stages which result in incomplete cancellation of the quantization noise generated in the first $n-1$ stages. Digital implementations of MASH architectures do not suffer from these mismatches. Thus, digital MASH architectures are capable of achieving complete cancellation of the quantization noise produced in the first $n-1$ stages of the modulator.

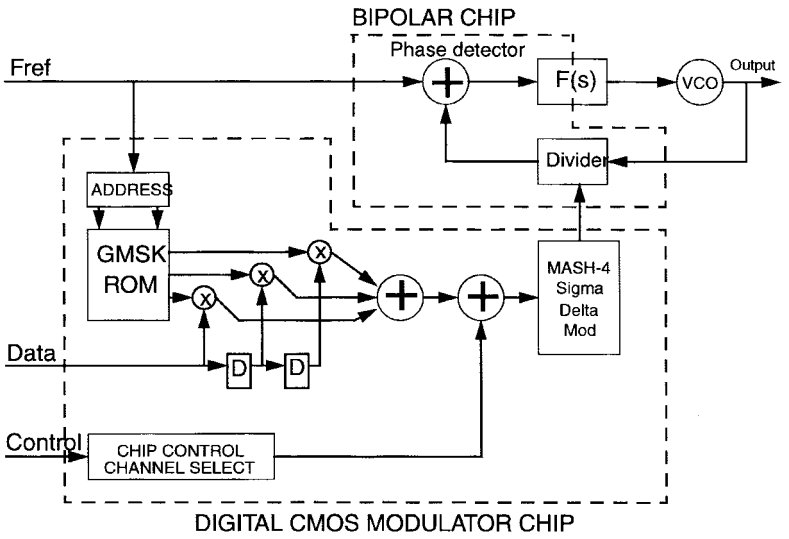

Fig. 4. Block diagram of synthesizer architecture.

It can be shown [1] that a digital accumulator is equivalent to a first-order sigma–delta modulator, as in one of the stages of Fig. 3, where the parallel input to the accumulator is the input to the modulator and the carry output is the sigma–delta modulated bitstream. The parallel output of the accumulator is a measure of the quantization noise. In this paper, a simple and low-power fourth-order digital MASH architecture formed from four digital accumulators is used to control an ECL/CML bipolar multimodulus divider circuit.

IV. Synthesizer Design

A block diagram of the frequency synthesizer is shown in Fig. 4. While the ultimate goal is to integrate the entire synthesizer onto a single BiCMOS chip, the design presented here was partitioned such that the high-frequency components of the PLL are implemented in a high-quality bipolar process and the digital sections are implemented in a high-quality CMOS process. This partitioning allows the best features of both technologies to be exploited while avoiding the problems that arise when mixing noisy digital circuits with sensitive analog circuits. The two custom chips were mounted on a four-layer printed circuit board along with an external loop filter and a commercial VCO. This approach avoided the difficulties of integrating a VCO onto a chip along with other noisy circuits. The designs of the custom CMOS chip and the custom bipolar chip are covered in subsequent sections.

The synthesizer was designed to be part of a frequency-hopping transmitter that would operate in part of the lower instrumentation scientific and medical (ISM) band from 902 to 928 MHz. The channel spacing was chosen to be 78.125 kHz, with a frequency settling time of less than 50 μs. The data transmission rate is 62.5 kbits/s with GMSK pulse shaping with a BT of 0.5, giving a frequency deviation of ± 15.625 kHz. The GMSK transmitter architecture used eliminates the need for the I and Q channels, D/A converters, summers, and up-conversion normally used in GMSK data transmission [5].

The digital data to be transmitted are input serially to the data input of the digital CMOS chip in Fig. 4, and filtered by the GMSK data filter to be described later, to shape the transmitted frequency spectrum. The synthesizer transmit frequency is controlled via a 3-bit serial interface (the control input in Fig. 4). This control interface reduces the number of pads required for the digital CMOS chip, and allows easy microprocessor interfacing for synthesizer control.

In order to achieve good spurious performance and low-phase-noise levels in the synthesizer, the interactions among the loop bandwidth, settling time, data rate, reference frequency, quantization noise, and VCO phase noise were considered carefully in the design, as follows.

The settling time of a PLL can be approximated as being equal to four times the inverse of the closed-loop bandwidth [9]. Therefore, in order to achieve a settling time of less than 50 μs, the loop bandwidth must be at least 80 kHz wide. Also, since the loop used in the synthesizer is being directly modulated with data, the bandwidth must be at least as wide as the modulation rate in order to avoid unwanted filtering of the transmitted data. The loop bandwidth was chosen to be 100 kHz. This is wide enough to achieve fast settling times and accommodate the transmitted data. If faster settling times were needed, the PLL could be modified such that the loop bandwidth could be increased temporarily during frequency changes [10].

The PLL (bipolar chip) suppresses quantization noise from the digital CMOS chip outside its loop bandwidth. Inside the PLL loop bandwidth, the quantization noise suppression is achieved by the noise-shaping properties of the digital sigma–delta modulator. Therefore, the modulator oversampling ratio and order both have a great influence on the noise performance of the synthesizer. Not only is the frequency reference ($F_{\rm ref}$ in Fig. 4) used to lock the PLL, but it serves as the clock for the sigma–delta modulator and the oversampled data filter. This means that the reference frequency sets the oversampling ratio of both the data filter and the sigma–delta modulator. It is important to note that as the reference frequency is increased, the power consumption of the CMOS chip also increases. Therefore, the tradeoff for

additional bandwidth improvement is increased power consumption. Similarly, higher order modulators to reduce in-band phase noise and increase the bandwidth of the noise shaping require more circuitry, also increasing power consumption.

Inside the loop bandwidth, the dominant sources of phase noise are the reference frequency, divider, and phase detector. Outside the loop bandwidth, the dominant sources of phase noise are the filtered sigma–delta quantization noise, filtered divider phase noise, and unfiltered VCO phase noise. Therefore, the loop filter must be designed in such a way as to achieve low in-band and out-of-band phase noise levels. The loop filter used in this design consists of an integrator with phase lead correction forming a second-order PLL. Two additional real poles were added to the loop in order to reduce reference feedthrough and out-of-band quantization noise. These poles were placed in such a way that they would not affect loop stability. The reference frequency was chosen to be 20 MHz, and the sigma–delta modulator order 4.

The multimodulus divider makes it simple to add an integer offset (a dc component of the digital control input) to the modulator output, allowing different fractional bands to be selected for data transmission. In this case, the bands are 20 MHz wide. For widely spaced bands, the loop filter must be modified in order to take into account the change in loop gain brought about by the change in division ratio that is required. A different VCO might be required as well. The loop presented in this paper was designed such that it would be usable over the entire band from 902 to 928 MHz without any modifications.

A. Digital Modulator Chip

Fig. 4 includes a block diagram of the CMOS digital modulator chip, which was captured in the Verilog hardware description language (HDL) and implemented in a 3-V CMOS process using a low-power standard cell library. Preliminary testing of the design was accomplished using a standard Xilinx field-programmable gate array (FPGA). The use of Verilog allowed the modulator design to be technology independent, facilitating the move from FPGA to standard cell CMOS. The main sections of the digital modulator chip are the GMSK data filter and the MASH-4 sigma–delta modulator. Each of these two blocks is described in greater detail in the following sections.

1) GMSK Data Filter: Fig. 4 includes a block diagram for the GMSK data filter designed using techniques from [5]. Three data symbols, past, present, and future, are multiplied by the filter tap coefficients and combined to produce Gaussian pulse shaping with a BT of 0.5. It is important to note that in a traditional GMSK data filter, a data bit consists of a rectangular pulse which is then filtered by a Gaussian pulse. Thus, the output is a Gaussian pulse convolved with a rectangular pulse. The implementation used here directly looks up this convolved response in a ROM, at the oversampled rate with one filter coefficient for each sample. To reduce the ROM size to 1 bit per sample, the samples are quantized to single-bit values, prior to storage in the ROM, using a third-order software sigma–delta modulator. The average value of the tap coefficients is equal to the desired Gaussian pulse, and the short-term changes represent the high-pass filtered quantization noise which is subsequently filtered out by the PLL. In this way, high resolution is obtained from single-bit tap coefficients. The multipliers, shown in Fig. 4, are now reduced to XNOR gates followed by the summing action shown. The output of the filter is fed to the MASH-4 sigma–delta modulator such that the total area under the Gaussian pulse results in a carrier phase shift of one quarter cycle, or $\pi/2$ rad, giving minimum shift keying. The oversampling ratio is 320, giving 320 single-bit coefficients and a data rate of 62.5 kbits/s for a reference frequency of 20 MHz. The total memory storage space required to store the filter tap coefficients is 960 bits, significantly lower than what would be required if a standard filtering approach were used.

The GMSK filter response is plotted in Fig. 5. Note that Fig. 5 shows the response of the nonconvolved pulse in order to facilitate comparison to other GMSK filters. The high-pass filtered quantization error of the tap coefficients can clearly be seen on the plot. The stopband frequency attenuation is 38 dB. The in-band characteristics of the filter are as expected for a Gaussian low-pass filter, and the out-of-band characteristics depend on the MASH-4 sigma–delta noise shaping and on how well the PLL filters out the quantization error shown in Fig. 5. The flat portion of the characteristic from 180 to 900 kHz is a result of end effects, specifically, caused by the truncation of the sigma–delta quantized tap coefficients.

Since the GMSK data filter makes use of sigma–delta quantized tap coefficients, the synthesizer which follows the filter can have lower frequency resolution than what would normally be required. This is true because the filter output is quantized to a fixed number of levels which are oversampled. This means that frequency values between the levels can be achieved. For instance, if the output of the filter is quantized to two levels, proportional to 0 and 20 kHz, values between these levels can be synthesized due to the oversampling and noise-shaping nature of their quantization. For example, a resolution of 20 kHz in the synthesizer will allow filtering which would normally require higher resolution. Therefore, the filter architecture used, in addition to reducing the ROM storage space and hardware complexity required for GMSK data filtering, also reduces the complexity required in the synthesizer sigma–delta modulator, lowering the power consumption of the control chip. The resolution required in the synthesizer in order to accommodate the GMSK data filtering was found to be 12 bits. This gives a frequency resolution of 4.882 kHz for a 20-MHz clock frequency. If a more conventional approach were used for the data filtering, and the tap coefficients were quantized to 8-bit values, the resolution required in the synthesizer would be 122 Hz. This would require an accumulator width of 18 bits in the sigma–delta modulator. Through the use of sigma–delta quantized tap coefficients, the resolution required in the sigma–delta modulator has been reduced by a factor of 64, significantly lowering power consumption.

An extra input is included in the modulator chip which allows the GMSK modulator to be disabled when not in use. This allows improved power budgeting in the transmitter.

2) MASH-4 Digital Sigma–Delta Modulator: As mentioned in Section III, the fourth-order digital sigma–delta

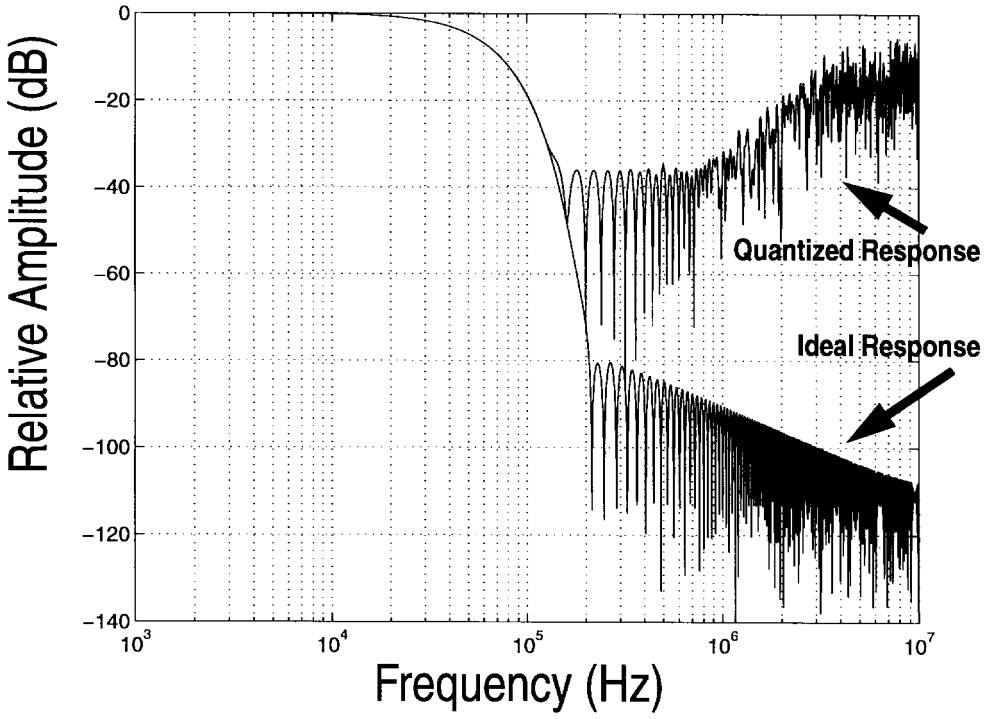

Fig. 5. Gaussian filter frequency response.

modulator required for the synthesizer is formed using a MASH architecture that makes use of four digital accumulators. Although the accumulators allow a simple implementation of the fourth-order sigma–delta modulator, they are still responsible for a large portion of the power consumption of the modulator chip. In order to achieve the lowest power consumption possible, the number of bits in the accumulator must be minimized. The use of sigma–delta quantized tap coefficients, as explained in the previous section, allows this.

One way to prevent limit cycles from occurring in the modulator, due to certain dc inputs, is to "plant a seed" in the accumulator each time the circuit is reset. This is done by toggling the least significant bit in the first accumulator on for one reference cycle and then turning it off. For "busy" inputs, this step is not required, but is included in order to ensure proper operation of the synthesizer for all inputs.

The output of the sigma–delta modulator is 4 bits wide, making it necessary to have a minimum of 16 available modulae in the multimodulus divider. Fifty-five division ratios were provided in the implementation of the multimodulus divider to provide greater flexibility in use. In this way, the synthesizer can be used with different VCO's and reference frequencies. An adder is included after the sigma–delta modulator, allowing an integer offset to be added to the output bits, selecting different fractional bands for data transmission. A translation circuit is included to convert the output to the format required by the divider control input.

B. Bipolar Synthesizer Chip

A block diagram of the bipolar chip is shown in Fig. 6. The chip consists of a multimodulus divider and a phase/frequency detector with an output charge pump. The divider is capable of dividing by any integer value between 9 and 64. CML logic is used wherever possible in the design, and the tail currents were set to the lowest value possible while maintaining functionality. Although the lower current leads to increased phase noise in the divider, and hence in the synthesizer noise floor, it was felt that the reduction in power consumption was worthwhile for this application. The noise floor achieved in the synthesizer chip was between -90 and -95 dBc/Hz. A more stringent phase noise floor specification would require higher power in the bipolar chip.

1) Multimodulus Divider: Looking at the multimodulus divider part of the block diagram shown in Fig. 6, the first section is a dual-modulus divider (DMD), which divides by 3 or 4, depending on the polarity of the control line from the A counter. When the A counter output is high, the DMD divides by 3, and when the A counter output is low the DMD divides by 4. The A counter is reloaded each time the M counter reaches its terminal count. The A counter then counts down to 0, meanwhile outputting a 1. When the count has reached 0, the counter remains at 0 until reloaded, and puts out a 0. For example, if the A counter is loaded with 3, it will count down to 0, while putting out three 1's to the DMD, and stay there. This will produce three "short" DMD cycles, which are three VCO cycles long as opposed to four cycles long. If a 0 is loaded, then no "short" cycles are produced. In this way, the number of "short" cycles is controlled by the A counter.

The M counter determines the number of DMD cycles that occur before both counters reload. Therefore, between every reload, there are $M+1$ DMD cycles, A of which are "short." The number of VCO cycles between reloads is four times the number of "short" cycles plus three times the number of "long"

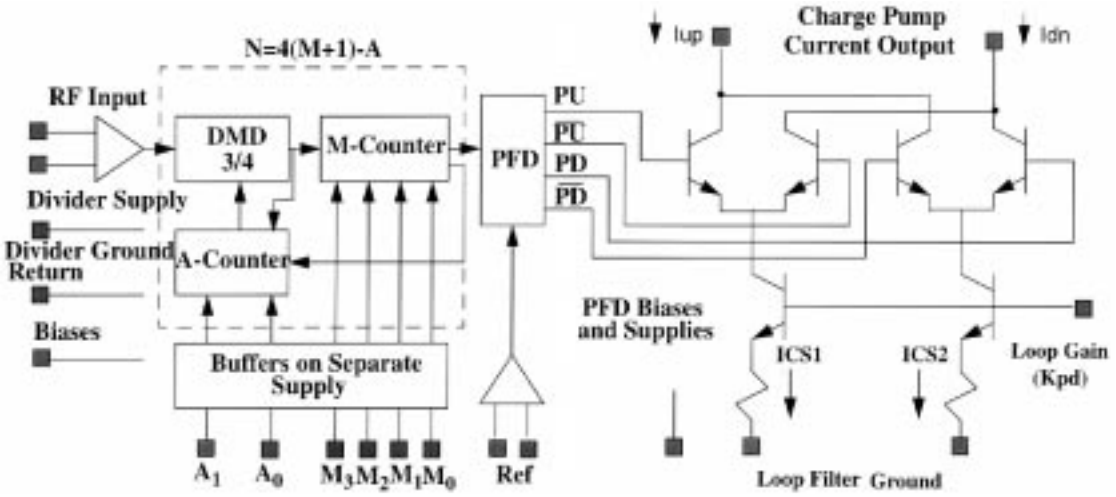

Fig. 6. Block diagram of bipolar chip.

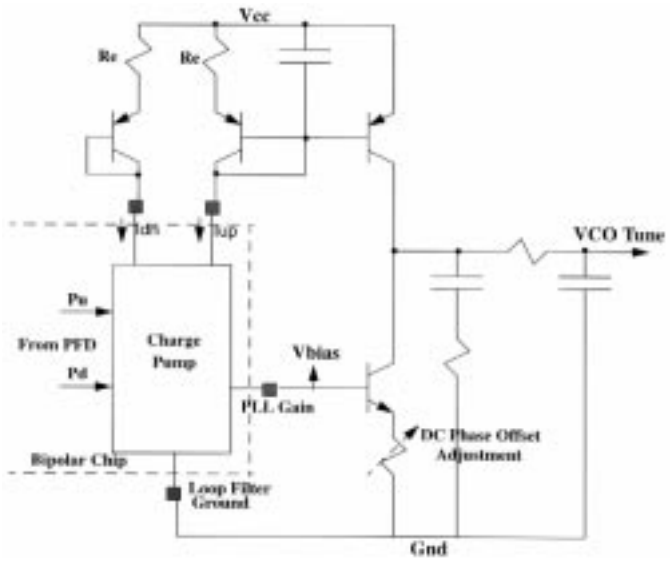

Fig. 7. Loop filter.

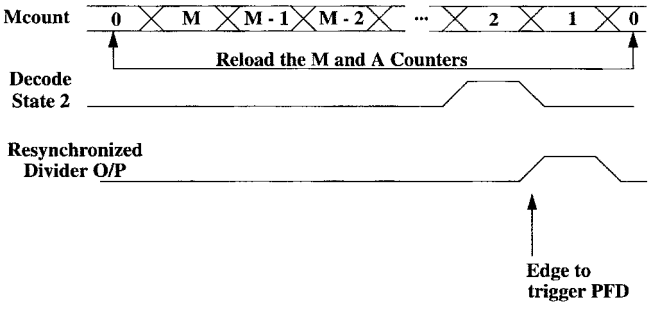

Fig. 8. Timing diagram: decoding the divider output.

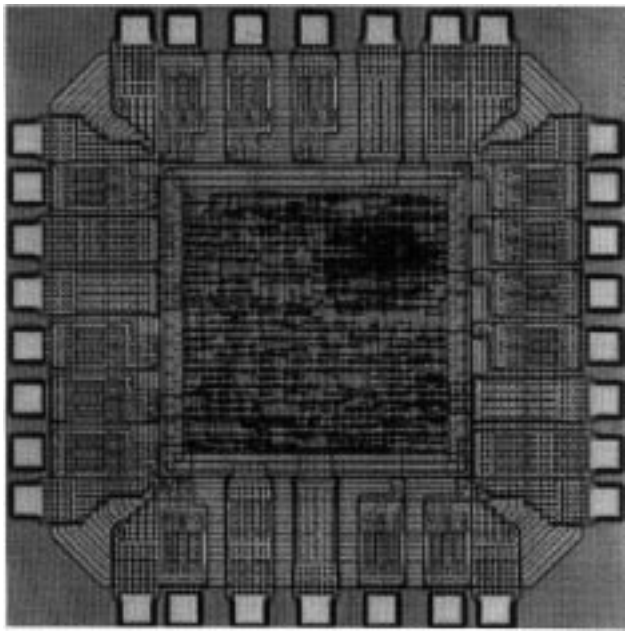

Fig. 9. Photomicrograph of digital CMOS modulator chip.

cycles. That is, $N = 4((M+1) - A) + 3A$. This simplifies to $N = 4(M+1) - A$.

In this design, the M counter is 4 bits wide and the A counter is 2 bits wide. Since the minimum value of M is 2, the range of available division ratios is 9–64.

a) Digital versus semianalog: Although the function of the divider is simply to count up to some numerical value and then provide an output pulse, the edge timing of the output pulse is an analog quantity. Thus, a digital function can have an analog error in the timing. If these errors are larger than other sources of error such as VCO phase noise or quantization error, then the phase noise of the overall synthesizer can be limited by noisy (i.e., random) timing errors of the divider. To minimize these errors, several strategies are employed in the divider chip, and are described in the following paragraphs.

The path from the off-chip VCO to the on-chip phase/frequency detector, PFD in Fig. 6, is fully differential. To use fully differential latched NOR gates with a 3-V supply requires setting the DMD clock at about 1.5 V_{be} drops below the positive supply. Initially, the VCO signal is converted to a differential signal with an off-chip center-tapped transformer. It is anticipated that future versions at higher frequencies will incorporate this transformer on chip [11] or the VCO itself [12] on chip. The M counter uses single-ended logic internally, but its output is retimed (i.e., resynchronized) by the DMD output clock with a fully differential flip-flop.

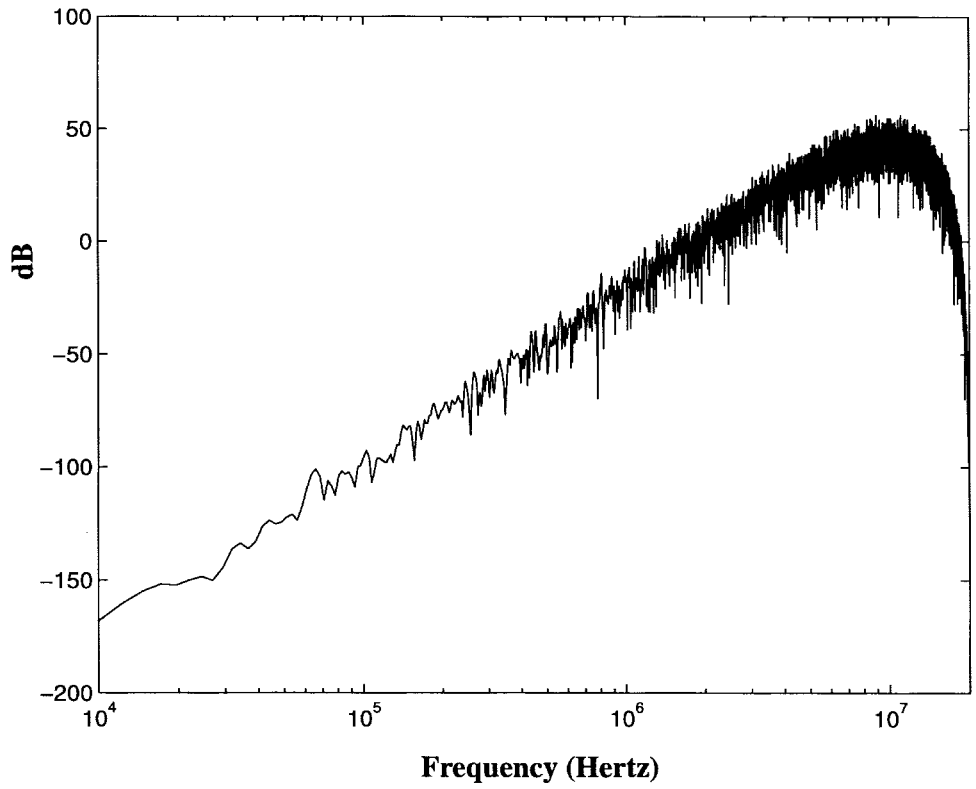

Fig. 10. Frequency spectrum of sigma–delta modulator output bits.

Previous work has presented techniques for time domain isolation of the divider output signal from the reference signal, and also from power supply interference generated by the MASH-4 [2]. We call this interference "digital feedthrough." The motivation for this is to force the rising edge of the divider output to occur at a different time from the rising edge of the reference output. By having the two events occur at different times, interactions between the two signals can be reduced. An older and simpler technique, which is used here, is to introduce an offset into the loop filter integrator which in turn forces a phase offset into the locked loop. This is accomplished with the loop filter shown in Fig. 7. An adjustable phase offset is not required, but is convenient for debugging purposes. This phase offset technique increases the reference feedthrough slightly, but it can be removed with subsequent filtering [9], [13] when high reference frequencies are used. In addition to time domain isolation, this technique also provides wider pulses for controlling the loop filter. These wide pulses are intended to provide a reliable phase detector gain even when the loop is in lock [9]. Without the offset, narrow spikes result in dead zones in the phase detector which can give less reliable performance.

Another simple technique to obtain time domain isolation is to decode the divider output in such a way that it occurs at a time when there is minimal digital activity in the divider. As illustrated in the timing diagram of Fig. 8, the divider loads a new value of N during its terminal count of 0. By decoding state 2 as the divider output pulse, the PFD can be triggered before the divider reloads. This allows almost an entire reference cycle for any parasitic effects from divider loading to settle before the next divider output pulse is applied to the PFD.

ECL/CML logic can be thought of as a series of switched current sources pulling various resistors to some voltage below the positive supply voltage. For this reason, the positive supply is taken as a local common voltage in each major functional block in the chip. By arranging the layout of the cells for each functional block around a central positive supply, IR voltage drops along the length of the positive supply are minimized within each functional block.

By using differential signals between functional blocks, differences between the power supplies of the functional blocks are converted to common mode and attenuated by the CMRR of the receiving block. In contrast, CMOS complementary logic would be more susceptible to power supply noise, and would generate larger current spikes on the power supplies.

Without exhaustive testing, it is difficult to isolate which of these approaches was most critical in its effects on phase noise or linearity. It is hoped that this discussion will stimulate further research and reporting on such approaches.

2) Phase Detector and Charge Pump: The phase/frequency detector (PFD) indicated in Fig. 6 is a standard architecture similar to the MC4044 or 12040. It is built in ECL rather than CML to allow easy design of a reference voltage for single ended logic gates.

The output of the charge pump after the PFD is open collector as shown in Fig. 6. The operation is as follows. Normally, neither the PU nor PD signals are active, and the two current sources ICS1 and ICS2 are shunted one to each transistor in the off-chip load, not shown. If the VCO is leading, PD will be active more often and I_{dn} will have twice as much current pulled through it, while I_{up} will be pulled up by the off-chip load. If the VCO is lagging, the opposite is true.

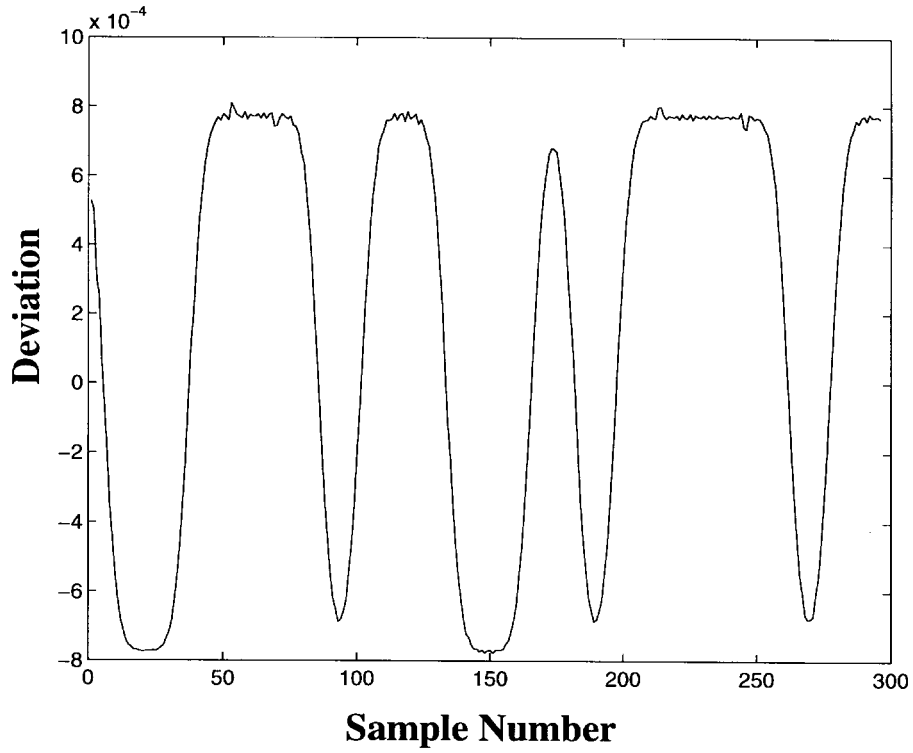

Fig. 11. GMSK data bits.

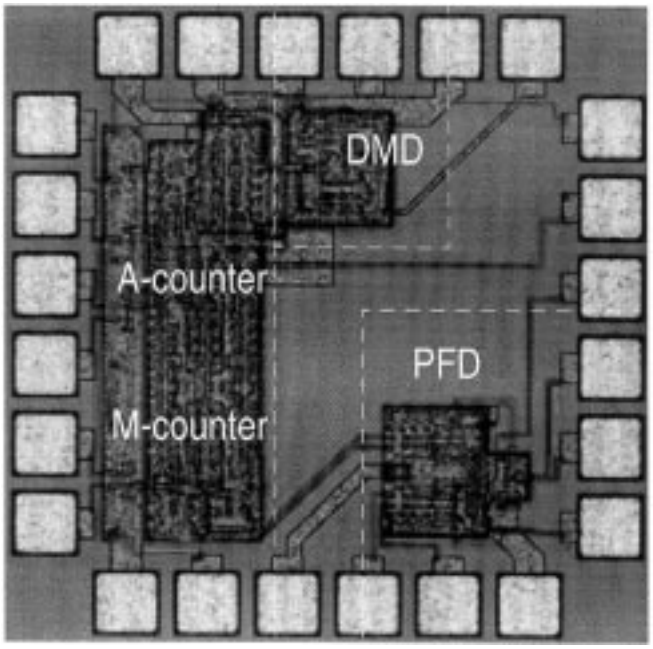

Fig. 12. Photomicrograph of bipolar chip.

V. Measurement Results

In this section, experimental results for the CMOS modulator chip, the bipolar chip, and the complete frequency synthesizer are presented.

A. CMOS Synthesizer Modulator Chip

A photomicrograph of the CMOS modulator chip is shown in Fig. 9. The circuit was tested by capturing a large number of the divider control words using a logic analyzer and processing them in MATLAB. The frequency spectrum of this bitstream for a fixed input is shown in Fig. 10. The fourth-order high-pass noise shaping can clearly be seen from this plot.

In order to test the GMSK data filtering, a pseudorandom (PN) code generator was used to feed random data bits into the data input of the digital modulator chip. Once again, the words from the modulator chip were captured using a logic analyzer. Fig. 11 shows a plot of the output bits for a particular input bit sequence. This plot was obtained in MATLAB using the decimation function to process the sigma–delta bitstream. This function reduces the sampling rate and removes the noise-shaped quantization error. The data symbols shown in the plot represent the "1's" and "0's" generated by the PN code generator and filtered by the Gaussian filter in the digital modulator chip. The horizontal axis is representative of time, and the vertical axis is representative of frequency deviation. However, the scales used have no particular significance. The small "dimples" that can be seen in the data bits are due to the truncation to a 1 or -1 that occurs at the end of the impulse response when it is quantized and then truncated.

The maximum clock frequency was found to be 50.0 MHz, allowing GMSK data rates of up to 156.25 kbits/s. The power consumption, measured with the data filter and sigma–delta modulator active, was found to be 14 mW with a 20-MHz clock.

B. Bipolar Synthesizer Chip

A photomicrograph of the bipolar synthesizer chip is shown in Fig. 12. The bipolar chip was tested for functionality over a wide range of division ratios. The maximum frequency for the RF input was found to be 1.2 GHz. The maximum usable

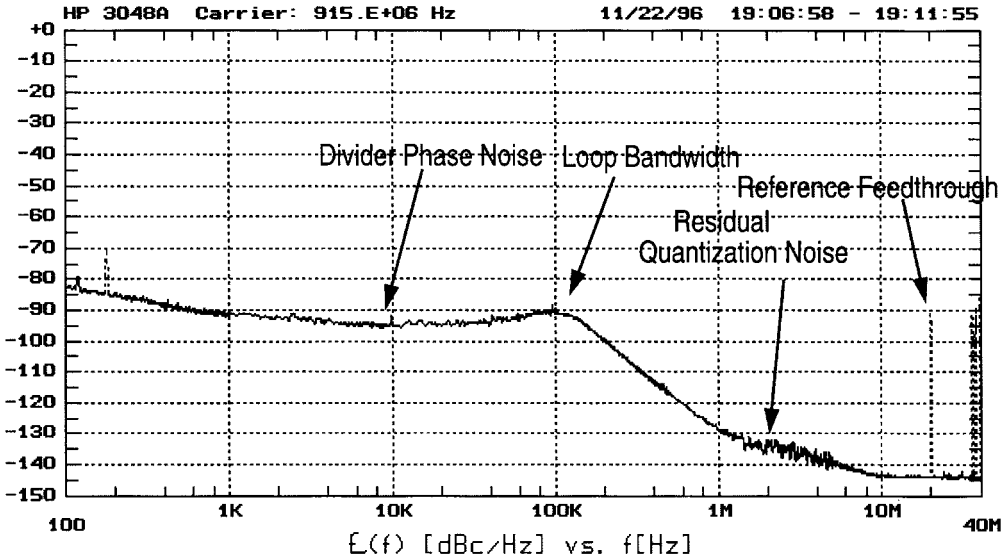

Fig. 13. Single-sideband phase noise plot of synthesizer for 915-MHz output.

reference frequency was found to be 300 MHz. The chip consumes 18.1 mW from a 3-V supply.

C. Overall Synthesizer Results

A single-sideband (SSB) phase noise plot for the overall synthesizer in fractional-N operation at a frequency of 915 MHz is shown in Fig. 13. It can be seen from the plot that the only spurs present in the output spectrum are the 20-MHz reference frequency feedthrough at -90 dBc and a harmonic of 60-Hz power line noise at -70 dBc. The power supply harmonic could be eliminated through improved shielding or by running the synthesizer from a battery. It is important to note that the in-band phase noise at the synthesizer output is the same for both integer operation and fractional operation. This means that the sigma–delta modulator does not affect the in-band phase noise characteristics of the synthesizer. By varying the dc phase offset adjustment of the synthesizer, in-band noise could be increased or decreased. It was found that the in-band noise was lowest when the divider edges and reference edges were separated as widely as possible in time while maintaining lock. This setting also provided the largest reference feedthrough. This result is consistent with the theory that "digital feedthrough" does, in fact, perturb the divider edges. The phase noise of the open-loop VCO was measured separately from the synthesizer, and was found to be less than -110 dBc/Hz at offsets greater than 10 kHz from the carrier. This means that the in-band phase noise floor of the synthesizer, at -90 to -95 dBc/Hz, is due to the bipolar synthesizer chip. This is due to the low current densities that were used in the logic gates. This lowered the transconductance of the devices, as well making high valued pull-up resistors necessary to maintain signal swings.

In Fig. 13, the phase noise floor from 1 to 9 MHz is due to residual sigma–delta quantization noise. The sigma–delta quantization noise is not truly white, but rather shows periodicity. The level of this noise, however, is at the predicted value. This periodicity occurs because the sigma–delta modulator used is a digital state machine. In fact, because the noise from the first three sigma–delta stages is cancelled, the number of possible quantization noise states is determined by the number of possible states in the final accumulator. If less periodicity is desired, increasing the number of bits in the accumulator or applying a dither signal to the first accumulator would be desirable.

It was found that as the frequency being synthesized approaches the fractional band edges (integer multiples of the reference frequency), to within one loop bandwidth, in-band spurs appear at a level of -40 dBc. The spurs are present on both sides of the carrier at a frequency offset exactly equal to the carrier offset from the band edge. For example, if the frequency being synthesized is $F_{\text{ref}} * N$ plus 78.125 kHz, the carrier will be centered at this frequency, and spurs will appear on both sides of it at a frequency offset of 78.125 kHz. These spurs do not appear in the sigma–delta bitstream when it is captured and processed in MATLAB. The spurs arise due to nonlinear effects in the synthesizer. The spur levels show a strong dependence on the input level of the loop reference frequency, indicating a possible nonlinearity in the phase detector. The exact source and elimination of this nonlinear effect is the subject of ongoing research. However, if the band edges are avoided, a spur-free range of 19.6 MHz is available for data transmission.

The approximate settling time of the synthesizer, to 80% of the final frequency, was measured for a frequency step of 8.4 MHz, and was found to range from 30 to 60 μs. This range is due to the fact that for this size of step, the synthesizer loses lock and must reacquire lock before settling. For smaller steps, the settling time is lower than 50 μs, as desired.

Random data bits were applied to the data input of the digital modulator chip at a rate of 62.5 kbits/s, resulting in a GMSK modulated frequency spectrum centered at 915 MHz. The modulated carrier was down-converted to 1.0 MHz in

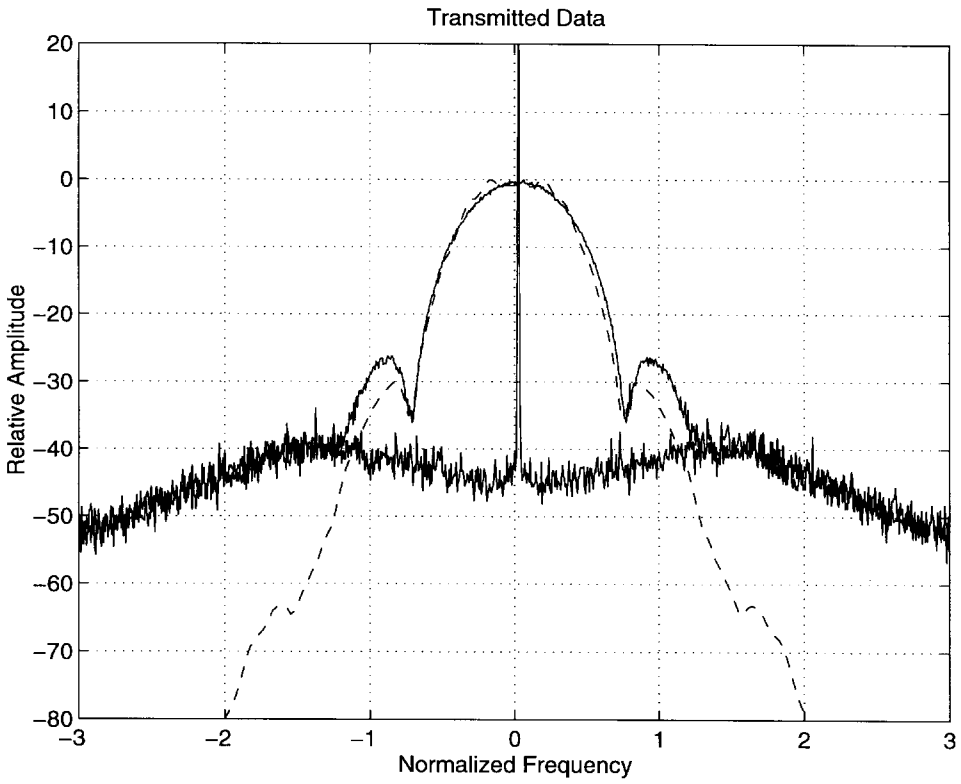

Fig. 14. GMSK modulated carrier (measured versus theoretical).

order to take advantage of the narrow-resolution bandwidth of the 40-MHz HP spectrum analyzer. A plot of the down-converted spectrum is shown in Fig. 14, along with a plot of the theoretical spectrum for GMSK with a BT of 0.5. The horizontal frequency axis is normalized by the data rate. A plot of the down-converted and unmodulated carrier is included as well. From this plot, it can be seen that the theoretical and measured spectra agree closely. The level of the sidelobes for the measured spectrum is slightly higher than theory due to the limited stopband attenuation of the Gaussian filter. By comparing the spectra of the modulated and unmodulated carriers, it can be seen that the noise floor of the modulated carrier is due to the in-band and out-of-band phase noise of the PLL, and is not a nonideal effect in the GMSK filtering.

The power consumption of the two chips was 32.1 mW. The synthesizer presented in [3] had a power consumption of 27 mW, but did not include on-chip data filtering. The overall power consumption of the synthesizer was found to be 491 mW. Most of the power consumption is due to the two "drop-in" amplifiers used in the synthesizer. These amplifiers are rated at 85 mA from a 5-V supply. Another 30 mW of the power consumption is due to the VCO. If the VCO were integrated on the chip as well, the power consumption of the synthesizer could be reduced even further [12].

VI. Conclusion

It has been shown, through the design and results presented in this paper, that by incorporating a digital sigma–delta modulator, a fractional-N frequency synthesizer can be designed which has good spurious performance, low phase noise, and fast settling times. The synthesizer can be controlled with high accuracy, allowing direct digital modulation with GMSK data, within the loop bandwidth.

The data filter and the sigma–delta modulator were simplified through the quantization of the tap coefficients using a software sigma–delta modulator prior to storing them in an on-chip ROM. This reduced the required multipliers to XNOR gates.

The design was partitioned into a PLL section and a digital modulator section. A bipolar process was used to implement the PLL section, incorporating a multimodulus divider, phase/frequency detector, and charge pump, using ECL/CML logic. The bipolar chip required 18.1 mW at 3.0 V. A CMOS FPGA was used to prototype the digital modulator section, incorporating a fourth-order sigma–delta modulator, a GMSK data filter, and serial interface. The final digital modulator was implemented in a CMOS process and required 14 mW for 20 MHz operation at 3 V. The power consumption of the digital chip could be reduced significantly further by moving from a standard cell approach to a full custom digital design. In order to simplify the design, an external loop filter and commercial VCO were used, but in future versions, the VCO will be included on chip where phase noise requirements allow.

Acknowledgment

The authors thank Nortel Department 5C60 for supporting the initial synthesizer work, including the bipolar design, and Department 5S12 for supporting the CMOS digital modulator work.

REFERENCES

[1] T. A. D. Riley, M. A. Copeland, and T. A. Kwasniewski, "Delta–sigma modulation in fractional-N frequency synthesis," *IEEE J. Solid-State Circuits*, vol. 28, pp. 553–559, May 1993.

[2] B. Miller and B. Conley, "A multiple modulator fractional divider," in *Proc. IEEE 44th Annu. Symp. Frequency Contr.*, 1990, pp. 559–567.

[3] M. H. Perrott, T. L. Tewksbury, and C. G. Sodini, "A 27 mW fractional-N synthesizer/modulator IC," in *Proc. IEEE Solid-State Circuits Conf.*, San Francisco, CA, Feb. 1997, pp. 366–367.

[4] K. Murota and K. Hirade, "GMSK modulation for digital mobile radio telephony," *IEEE Trans. Commun.*, vol. COM-29, pp. 1044–1050, July 1981.

[5] T. A. D. Riley and M. A. Copeland, "A simplified continuous phase modulator technique," *IEEE Trans. Circuits Syst.*, vol. 41, pp. 321–328, May 1994.

[6] J. C. Candy and G. C. Temes, "Oversampling methods for A/D and D/A conversion," in *Oversampling Delta–Sigma Converters*. New York: IEEE Press, 1991.

[7] P. M. Aziz, H. V. Sorensen, and J. Van Der Spiegel, "An overview of sigma–delta converters," *IEEE Signal Processing Mag.*, pp. 61–84, Jan. 1996.

[8] K. Uchimura, T. Hayashi, T. Kimura, and A. Iwata, "Oversampling A-to-D and D-to-A converters with multistage noise shaping modulators," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 259–258, Dec. 1988.

[9] J. A. Crawford, *Frequency Synthesizer Design Handbook*. Norwood, MA: Artech House, 1994.

[10] J. Browne, "BiCMOS PLL's conserve current in wireless designs," *Microwaves RF*, pp. 136–140, Oct. 1996.

[11] J. R. Long and M. A. Copeland, "Modeling of monolithic inductors and transformers for silicon RFIC design," in *Proc. IEEE MTT-S Int. Symp. Technol. for Wireless Appl.*, Vancouver, B.C., Canada, Feb. 1995, pp. 129–134.

[12] L. Daupinee, M. Copeland, and P. Schvan, "A balanced 1.5 GHz voltage controlled oscillator with an integrated LC resonator," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 1997, pp. 390–391.

[13] U. L. Rhodes, *Digital PLL Frequency Synthesizers*. Englewood Cliffs, NJ: Prentice-Hall, 1983.

Norman M. Filiol received the B.Eng. degree in electrical engineering in 1992 and the M.Eng. degree in electrical engineering in 1994, both from Carleton University, Ottawa, Ont., Canada. He is currently a Ph.D. candidate at Carleton University in a cooperative research agreement with Nortel Semiconductor.

From 1989 to 1992, he spent three consecutive summer terms at Transport Canada working on navigational aids. His research interests are in the area of mixed-signal VLSI design for mobile radio. Ongoing research is making use of CMOS, bipolar, and BiCMOS technologies to develop low-power, highly integrated building blocks for frequency synthesis and frequency discrimination. He is funded, partially, by Nortel Semiconductor and by Communications and Information Technology Ontario (CITO), a provincially funded university-based research organization.

Thomas A. D. Riley (S'81–M'82) received the B.E.Sc. degree in electrical engineering from the University of Western Ontario, London, Ont., Canada, in 1982.

From 1982 to 1986, he worked in the Power Products Division of Phillips Cables Ltd., Brockville, Ont. His M.Eng. research was performed at Carleton University, Ottawa, Ont., and Mitel Semiconductor on modeling sigma–delta modulator architectures for primary rate ISDN. In 1988, he joined the VLSI in Communications Group at Carleton University as a Research Engineer. He is currently pursuing research in IF A/D converters at the University of Oulu, Finland.

Calvin Plett (S'82–M'82) received the B.A.Sc. degree in electrical engineering from the University of Waterloo, Waterloo, Ont., Canada, in 1982, and the M.Eng. and Ph.D. degrees from Carleton University, Ottawa, Ont., in 1986 and 1991, respectively.

From 1982 to 1984, he worked with Bell-Northern Research. In 1989, he joined the Department of Electronics, Carleton University, where he is now an Assistant Professor. Since 1995, he has done consulting work for Nortel Semiconductor in the area of RF integrated circuit design. His research interests are in the area of analog integrated circuit design, including filters, radio-frequency front-end components, and communications applications.

Miles A. Copeland (M'65–SM'85–F'89) received the B.Sc. degree in electrical engineering from the University of Manitoba, Winnipeg, Man., Canada, in 1957, and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Toronto, Toronto, Ont., Canada, in 1962 and 1965, respectively.

He has been a Professor in the Department of Electronics, Carleton University, Ottawa, Ont., for 31 years, and a Consultant with Nortel Technology Ltd. He is now officially retired, but maintaining an office and the honorary title of Distinguished Research Professor. He is active in research and graduate student supervision, and is presently a coinvestigator at Carleton University in the VLSI for Communications Group. As well, some of his research is supported through cooperative graduate student thesis projects with industry.