A Low Power Direct Digital Frequency Synthesizer with 60 dBc Spectral Purity

J.M.P. Langlois
Department of Electrical and Computer Engineering
Royal Military College of Canada
Kingston, Ontario, Canada
+1-613-541-6000 x6860
pierre.langlois@rmc.ca

D. Al-Khalili
Department of Electrical and Computer Engineering
Royal Military College of Canada
Kingston, Ontario, Canada
+1-613-541-6000 x6401
alkhalili-d@rmc.ca

ABSTRACT

We present a low-power sine-output Direct Digital Frequency Synthesizer (DDFS) realized in 0.18 µm CMOS that achieves 60 dBc spectral purity from DC to the Nyquist frequency. No ROM or multipliers are used, but an external DAC is required if an analog output is desired. Power consumption is 10 mW for a 100 MHz clock, which is significantly less than figures reported previously. System complexity is greatly reduced by using an efficient linear interpolation scheme to approximate a sinusoid function. This has resulted in silicon area utilization of 0.011 mm^2. The design would be suitable as an IP core in a low power digital RF transceiver ASIC.

Categories and Subject Descriptors

B.7.1 [**Integrated Circuits**]: Types and Design Styles – *algorithms implemented in hardware, VLSI.*

General Terms

Algorithms, Design, Performance.

Keywords

DDFS, DDS, phase to sine amplitude conversion, low power

1. INTRODUCTION

Modern spread spectrum frequency hopping communication systems require fast and precise tuning over a wide bandwidth. Direct Digital Frequency Synthesizers (DDFS) offer interesting advantages over analogue approaches such as phase-locked loops. These include unparalleled settling times and tuning bandwidths, phase continuity of the output, and the ability to include phase and amplitude modulation with almost trivial modifications. However, DDF Synthesizers are often considered power-hungry and can occupy a large area.

GLSVLSI'02, April 18-19, 2002, New York, New York, USA.

The so-called Table Look-Up Direct Digital Frequency Synthesizer (DDFS) architecture was introduced by Tierney, Rader and Gold in 1971 [17]. As shown in Figure 1, this synthesizer is based on three blocks: a phase accumulator, a phase-to-sine-amplitude converter, and a Digital-to-Analog Converter (DAC) with a Low Pass Filter (LPF).

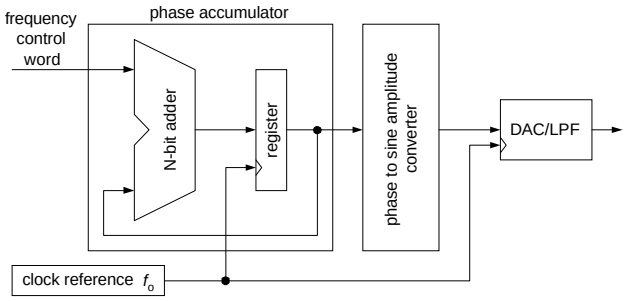

Figure 1. Basic Direct Digital Frequency Synthesizer

The system has two inputs: a clock reference and a frequency control word (FCW). The phase accumulator integrates the value of the FCW on every clock cycle, producing a ramp whose slope is directly proportional to the FCW. Assuming an N bit wide accumulator, the frequency of this ramp is given by:

$$f_{out} = f_0 \times \frac{FCW}{2^N}. \qquad (1)$$

The phase accumulator output is converted to an approximated sine amplitude by the phase-to-sine amplitude converter. In the simplest case, this converter is a Read Only Memory (ROM). A DAC and a LPF are used to convert the sequence of sinusoid samples to an analog wave.

In this paper, we present a low power DDFS ASIC realized in 0.18 µm CMOS. It uses a hardware-efficient linear interpolation scheme for the phase-to-sine amplitude converter block [7][9]. The paper is organized as follows. In section II, we review DDF Synthesizers recently reported in the literature. We describe our phase-to-sine amplitude converter block in section III, and present our system's architecture in section IV. Implementation details are given in section V. Section VI contains performance comparisons and design metrics between our design and others reported in the literature. Conclusions are presented in section VII.

2. REVIEW OF EXISTING WORK

Sine-output DDF Synthesizers based on the architecture shown in Figure 1 differ mostly in their implementation of the phase-to-sine amplitude converter block. Since this block is normally the slowest and most power hungry, it has been the target of much research effort in the past 30 years.

The simplest approach for the realization of the phase-to-sine amplitude converter is to use a ROM. However, its size grows exponentially with the width of the phase accumulator, *N*. Since one normally desires a large *N* to achieve fine frequency tuning, several techniques have been devised to limit the ROM size while maintaining adequate performance.

One method consists of exploiting the quarter wave symmetry of the sine function to reduce by 4 the number of angles for which a sine amplitude is required [5]. Truncating the phase accumulator output is another common method, although it introduces spurious harmonics [6]. Still, this approach is almost universally adopted because of fine frequency resolution requirements requiring a very large value for *N*.

Various angular decomposition methods have been proposed [5] to further reduce memory requirements with LUT-based methods. They normally consist of splitting the ROM into a number of smaller units, each addressed by a portion of the truncated phase accumulator output. Data retrieved from each sub-ROM is added to yield a sinusoid approximation.

In order to reduce the ROM size further, researchers have proposed making an initial approximation of the sine amplitude from the value of the phase angle, and to use the ROM or a combination of ROMs to store correction values [12][20][14][8]. These are often referred to as amplitude compression techniques. Used in combination with ROM segmentation, an effective ROM size compression ratio of 37:1 has been attained [12].

The ROM size can be further reduced or be eliminated altogether by increasing the amount of calculations performed to obtain a sine approximation. Several researchers have proposed approaches using piecewise continuous polynomials to approximate the first quadrant of the sine function. These have been based either on trigonometric approximations [2], on a Taylor-series expansion [19][1], on a simplified 4^{th} degree polynomial [15], on 4^{th} degree Chebyshev polynomials [13] or on linear segments of unequal lengths [10].

All these approaches are based on the following premise: that the amplitude error on the sinusoid approximation for any phase angle must be reduced to what would be achievable with a ROM Look-Up Table. For example, for a ROM width of *L* bits, the maximum amplitude error should be less than 2^{-L}. In order to meet this goal, all architectures require an increased complexity in the approximation circuit. Thus, approaches presented in [2][19][1][15][13] include multiplication and/or squaring circuits. The one based on unequal linear segments [10] involves complex control circuitry to identify the segments and several additions for the calculation of each of these segments.

3. PHASE-TO-SINE AMPLITUDE CONVERTER BLOCK

Our first goal in designing the phase-to-sine amplitude converter block is to minimize its complexity. We have therefore proposed [9] approximating the first quadrant of the sine function with eight equal-length piecewise continuous linear segments s_i of the form:

$$s_i(x) \cong m_i(x - x_i) + y_i,\ x_i \le x < x_{i+1} \qquad \textbf{(2)}$$

where m_i is a segment's slope, x_i defines segment bounds, and y_i is the initial approximation for each segment. From this equation, it is obvious that our approach requires no squaring circuit. However, equation (2) includes one multiplication.

The first important feature of our architecture, when compared to previous work, is that we constrain the quantization of the segment slopes such that they are represented with at most two non-zero binary digits. We exploit the well known principle that multiplication by a factor of two can be accomplished with a trivial bit shift, and that multiplication by a factor equal to a sum of two powers of two can be accomplished with at most two trivial bit shifts and one addition. Consequently, implementing the multiplication in equation (2) requires at most one addition.

The second important feature of our architecture is that we limit the dynamic range of each slope m_i so that each can be expressed with four bits. This implies 16 possibilities, however we use only a subset of 12. We discard those slopes with more than two non-zero digits but accept -1 as a valid digit. We scale first quadrant angles from the interval [0, π/2] to the interval [0, 1], in order to represent them as a binary fraction. Hence, the first derivative of the sine function in the first quadrant is scaled from the range [0, 1] to [0, π/2] ≈ [0, 1.57]. Consequently, we select segment slopes in the following set:

{1.5, 1.25, 1.125, 1, .875, .75, .625, .5, .375, .25, .125, 0}.

As mentioned previously, common wisdom in designing a DDF Synthesizer says that one should minimize the amplitude error on sinusoid amplitudes calculated for any phase angle. While this may be an important performance parameter for a sine function block, it is not necessarily so for a DDF Synthesizer. Spectral purity, which is defined as the ratio of the power in the desired frequency to the power in the greatest harmonic across the synthesizer's tuning bandwidth, is much more important. Spectral purity is an essential design parameter for synthesizers in communications systems, ensuring that undesired in-band signals remain below a given threshold and are not detected.

In order to achieve a desired spectral purity, we evaluate different sets of eight pairs of m_i and y_i coefficients, and select the best one meeting our requirements. We solve this optimization problem with a Genetic Algorithm [3], with the fitness function equal to the spectral purity. All calculations are done taking finite bit-width effects into account.

Equation (3) below gives the slopes and *y* approximations that we have used for this architecture. They meet the requirement of 60 dBc spectral purity. Figure 2 shows the corresponding output for angles in the first quadrant.

$$\sin(x\times\frac{\pi}{2})\cong\begin{Bmatrix}(1+\frac{1}{2})(x-\frac{0}{8})+\frac{2}{1024},\frac{0}{8}\le x<\frac{1}{8}\\(1+\frac{1}{2})(x-\frac{1}{8})+\frac{191}{1024},\frac{1}{8}\le x<\frac{2}{8}\\(1+\frac{1}{4})(x-\frac{2}{8})+\frac{384}{1024},\frac{2}{8}\le x<\frac{3}{8}\\(1+\frac{1}{8})(x-\frac{3}{8})+\frac{552}{1024},\frac{3}{8}\le x<\frac{4}{8}\\(1+0)(x-\frac{4}{8})+\frac{697}{1024},\frac{4}{8}\le x<\frac{5}{8}\\(\frac{1}{2}+\frac{1}{4})(x-\frac{5}{8})+\frac{819}{1024},\frac{5}{8}\le x<\frac{6}{8}\\(\frac{1}{2}+0)(x-\frac{6}{8})+\frac{909}{1024},\frac{6}{8}\le x<\frac{7}{8}\\(\frac{1}{8}+0)(x-\frac{7}{8})+\frac{971}{1024},\frac{7}{8}\le x<\frac{8}{8}\end{Bmatrix} \quad \textbf{(3)}$$

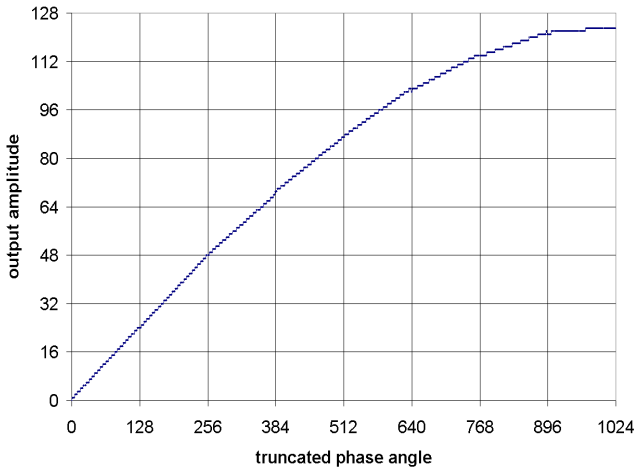

Figure 2. First quadrant sine approximation

In Figure 2, the 8 segments are noticeable, as are the amplitude quantization effects for each angle. Discontinuities at quadrant transitions may also be observed. The maximum amplitude is equal to 123/128 or 0.9609375. Taking the Discrete Fourier Transform of a full period of data reveals that the amplitude of the fundamental is approximately 123.1/128. This reduction from a maximum of 127/128 in full-scale output is inconsequential from a system perspective.

Figure 3 shows the amplitude error for each angle in the first quadrant. It is seen that the maximum error amplitude on any sample is between the approximate bounds -1.3 and +0.75 LSB. This is somewhat greater than results achieved by many architectures described in section 2, but inconsequential as spectral purity remains acceptable.

Figure 4 shows the output spectrum for a Frequency Control Word of 10001, or an output frequency approximately equal to $0.1526 \times f_0$. This spectrum was obtained by taking a Discrete Fourier Transform of one grand repetition period of system output data (here 2^{16} samples). The scale of the plot has been adjusted to bring the fundamental frequency to 0 dB. It can be seen that all spurs are at least 60 dB below the fundamental.

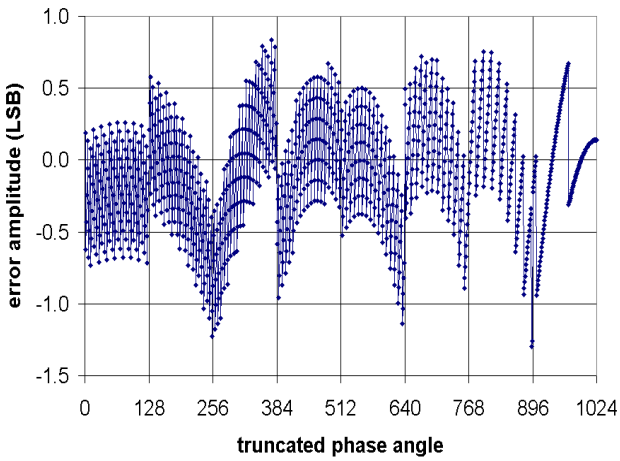

Figure 3. Sine amplitude error for first quadrant angles

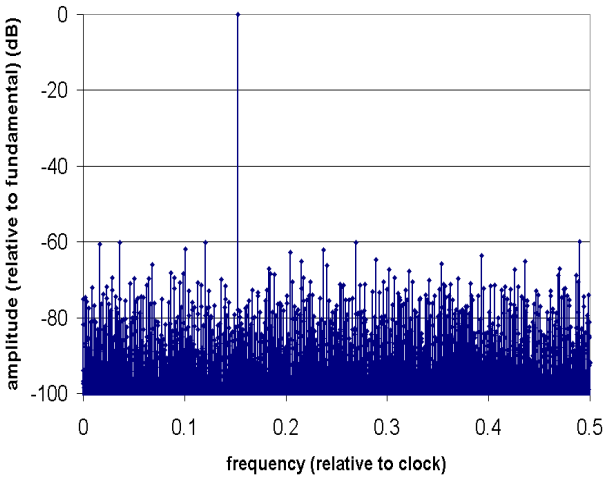

Figure 4. Calculated output spectrum, FCW = 10001

4. SYSTEM ARCHITECTURE

The system architecture is shown in Figure 5.

The phase accumulator's 16 bits are truncated to 12. This limits spurs due to phase truncation to approximately -72 dBc [6]. The two MSBs are used for quadrant symmetry [5]. The first MSB determines the sign of the output data. It controls a format converter block which modifies the sign and magnitude format to the two's complement format required by the DAC. The second MSB controls a 1's complement block, which inverts the remaining phase accumulator bits for angles in quadrants 2 and 4. The consequence is that the ramp output from the phase accumulator is converted to a triangular wave of equal frequency and twice the amplitude.

The next three MSBs identify one of eight linear segments, and thus they control the multiplexers that implement equation (3), which is defined in 8 parts. The remaining 7 bits identify different sub-angles, or positions along any of the 8 segments. In equations (2) and (3), these 7 bits are equal to the quantity (x - x_i), so this operation does not require any processing.

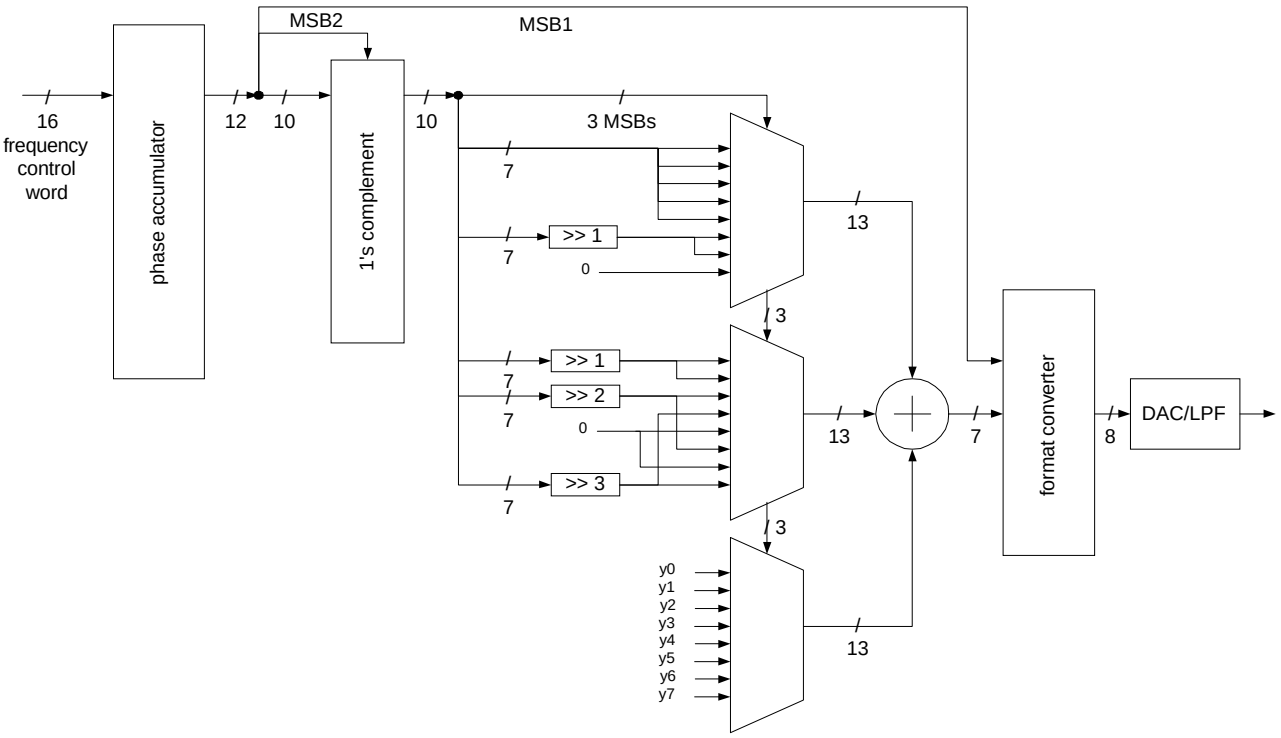

Figure 5. 60 dBc purity DDFS architecture using 8 linear segments

The two upper multiplexers select shifted versions of the 7 least significant phase bits, passing them to the three-operand adder according to the corresponding segment. In the figure, the notation $\{>>n\}$ signifies a right shift by n bits, or division by 2^n. The addition of two shifted versions of an angle x realizes the multiplication operation of an angle x by a slope m_i in equation (3). The bottom multiplexer selects one of eight initial approximations and also passes it to the three operand adder.

The output from the multiplexers is shown to be 13 bits wide, in order to properly align the three terms to be added. In actual fact, the first three bits of the two upper multiplexers are 0, as are the last three bits of the lower multiplexer.

The three-operand adder adds the multiplexer outputs together and rounds the result to 7 bits. The rounding operation is accomplished by adding the 8^{th} bit to the truncated 7-bit sum.

This architecture is significantly less complex than all those listed in section 2 for a similar output spectral purity performance. It does not include a ROM. No multipliers nor squaring circuits are required. Equal-length segments are used to simplify control circuitry. Only 3 integers need to be added, and the multiplexers shown in Figure 5 can be optimized by combining similar inputs, and be implemented with combinational logic.

5. IMPLEMENTATION DETAILS

The system was described in VHDL with less than 200 lines of code. Synthesis was performed using a 0.18 μm CMOS library of standard cells. A total of 155 core cells are used, and they occupy a total area of 5773 μm^2. With clock and reset signals, and power and ground pins, a total of 36 pins are necessary. The total chip area, including I/O cells and I/O pads, is approximately 1.5 mm^2, making this a severely I/O bound chip. The actual core area used is approximately 11250 μm^2.

During placement and routing with automated tools, a clock constraint of 125 MHz was easily met without having to add pipelining registers. Pipelining would increase this maximum clock rate, but at the expense of a longer latency when changing the synthesizer's output frequency. Power consumption is estimated at under 10 mW for a 100 MHz clock, or 0.1 mW/MHz.

The Frequency Control Word is 16 bits wide, yielding a frequency resolution of approximately 1526 Hz for a 100 MHz reference clock. The 8-bit wide output data is in two's complement format, compatible with most commercial DACs.

As stated above, this design is severely IO bound. This is a direct consequence of the tremendous reduction in complexity when compared to other previously reported designs for similar spectral purity. It is also a consequence of using an advanced 0.18 μm CMOS technology. Due to limited allocation of silicon area, it was decided not to increase the phase accumulator width to 32 bits, as is common. This would have added 16 pins to the chip and approximately 300 μm to each side of the die. The phase control word input could also have been serialized, but that would have increased the tuning latency.

If system frequency resolution requirements called for a 32 bit wide accumulator and the same 125 MHz clock rate, a modest increase in system complexity would follow. This is because several pipelining registers would be required. Alternatively, a more efficient adder configuration would have to be used with a corresponding increase in the number of cells. In any case, the present core is very small, which makes it an ideal building block in a System On a Chip digital receiver.

Table 1. Comparison of recent DDFS designs

ref./ year	process	complexity	core area (mm^2)	power (mW/MHz)	purity pre-DAC (dBc)	purity post-DAC (dBc)	notes
[16]/ 1995	0.8 µm	52000 transistors	15.9	10.0	84.3		ROM segmentation quadrature output + modulation circuitry
[18]/ 1998	0.8 µm	19100 transistors	3.9	4.0		52	ROM segmentation on-chip DAC
[4]/ 1999	0.5 µm	17803 transistors	9.0	3.3	72.2	47	ROM segmentation quadrature output + modulation circuitry
[11]/ 1999	1.0 µm	58000 transistors	12.0	14.0	100.0		CORDIC quadrature output
[1]/ 2000	0.8 µm		0.9	0.3	60.0		Taylor series expansion
[13]/ 2000	0.35 µm	5900 gates			70.0		Chebyshev polynomials
[10]/ 2001	0.6 µm	11721 transistors	1.1	1.0	67.6		linear segments, unequal lengths
our design	0.18 µm	155 cells	0.011	0.1	60.0		linear segments, equal lengths, restricted slopes

The chip is being fabricated, and should be tested and characterized in March 2002.

6. COMPARISON WITH OTHER DESIGNS

Table 1 gives a comparison between our design and other designs in terms of complexity, power consumption and output spectral purity. Empty cells represent data not included in the reference. The notes column gives the technique used to realize the phase to sine amplitude conversion block, and specifies whether the design has quadrature output, modulation capability, and an on-chip DAC. These features obviously entail greater design complexity.

Complexity units are shown as given by the authors. Since we used a black-box cell library, we do not have an exact transistor count for the 155 cells in our design. However, from inspection of the synthesized netlist, we estimate an average not greater than 15 transistors per cell, for a total under 2500 transistors. This is less than 25% of the least complex design listed in the table, and one order of magnitude smaller than most others.

Two spectral purity columns are given, one for pre-DAC purity and one post DAC. Most designs do not report post-DAC spectral purity. Achieved post-DAC purity is rarely above 50 dBc throughout the literature on DDFS. However, several researchers, as shown here, have produced designs with pre-DAC purity that greatly exceed any results that could be achieved or required in practice.

In the table, power consumption information was normalized to the clock rate. Power consumption is also affected by the design's complexity, technology and power supply voltage level. All of these factors contribute to our design achieving very low power consumption.

7. CONCLUSIONS

We have presented a low-power sine-output Direct Digital Frequency Synthesizer (DDFS) realized in 0.18 µm CMOS that achieves 60 dBc spectral purity from DC to the Nyquist frequency. It includes no ROM and no multipliers but requires an external DAC if an analog output is desired. Power consumption is 10 mW for a 100 MHz clock, which is significantly less than figures reported previously. System complexity is greatly reduced by using an efficient linear interpolation scheme to approximate a sinusoid function. Only 155 cells from a standard library are required, for a total core area of 0.011 mm^2. We have compared our design to recent ones published in the literature and showed that it is significantly less complex than previous work, for adequate output performance in many communications systems. The synthesizer would be suitable as an IP core in a low power digital RF transceiver ASIC.

8. ACKNOWLEDGMENTS

Chip fabrication is being made possible by a grant from the Canadian Microelectronics Corporation. This work was supported in part by Canada's National Science and Engineering Research Council.

9. REFERENCES

[1] A. Bellaouar, M.S. O'brecht, A.M. Fahim, and M.I. Elmasry, "Low-power direct digital frequency synthesis for wireless communications," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 3, March 2000, pp. 385-390.

[2] R.A. Freeman, "Digital sine conversion circuit for use in direct digital synthesizers," U.S. Patent 4,809,205, 28th February 1989.

[3] D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, 1989.

[4] M. Kosunen, J. Vankka, M. Waltari, L. Sumanen, K. Koli and K. Halonen, "A CMOS quadrature baseband frequency synthesizer/modulator," in *Analog Integrated Circuits and Signal Processing*, vol. 18, no. 1, January 1999, pp. 55-67.

[5] V.F. Kroupa, Ed., *Direct Digital Frequency Synthesizers*, IEEE Press, 1999.

[6] V.F. Kroupa, V. Cizek, J. Stursa, and H. Svandova, "Spurious signals in direct digital frequency synthesizers due to the phase truncation," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 47, no. 5, September 2000, pp. 1166-1172.

[7] J.M.P. Langlois and D. Al-Khalili, "Multiplier-less direct digital frequency synthesizer", U.S. patent provisional application, August 2001.

[8] J.M.P. Langlois and D. Al-Khalili, "ROM size reduction with low processing cost for direct digital frequency synthesis," *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, August 2001, pp. 287-290.

[9] J.M.P. Langlois and D. Al-Khalili, "Hardware optimized direct digital frequency synthesizer architecture with 60 dBc spectral purity", accepted for presentation at the *2002 IEEE International Symposium on Circuits and Systems (ISCAS 2002)*, May 2002.

[10] S.-I. Liu, T.-B. Yu and H.-W. Tsao, "Pipeline direct digital frequency synthesiser using decomposition method," *IEE Proceedings on Circuits, Devices and Systems*, vol. 148, no. 3, June 2001, pp. 141-144.

[11] A. Madisetti, A. Y. Kwentus, and A. N. Willson, Jr., "A 100-MHz, 16-b, direct digital frequency synthesizer with a 100-dBc spurious-free dynamic range," in *IEEE Journal of Solid-State Circuits*, vol. 34, no. 8, August 1999, pp. 1034-1043.

[12] H.T. Nicholas III, H. Samueli and B. Kim, "The optimization of direct digital frequency synthesizer performance in the presence of finite word length effects," *Proceedings of the 42nd Annual Frequency Control Symposium*, 1988, pp. 357-363.

[13] K.I. Palomaki and J. Niitylahti, "Direct digital frequency synthesizer architecture based on Chebyshev approximation," *Proceedings of the 34th Asilomar Conference on Signals, Systems and Computers*, Oct. 29th – Nov. 1st., 2000, pp. 1639-1643.

[14] A. M. Sodagar and G. R. Lahiji, "Mapping from phase to sine-amplitude in direct digital frequency synthesizers using parabolic approximation," in *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 47, December 2000, pp. 1452-1457.

[15] A.M. Sodagar and G.R. Lahiji, "A novel architecture for ROM-less sine-output direct digital frequency synthesizers by using the 2nd-order parabolic approximation," *Proceedings of the 2000 IEEE/IEA International Frequency Control Symposium and Exhibition*, 7-9 June 2000, pp. 284-289.

[16] L. K. Tan and H. Samueli, "A 200 MHz quadrature digital synthesizer/mixer in 0.8 μm CMOS," in *IEEE Journal of Solid-State Circuits*, vol. 30, no. 3, March 1995, pp. 193-200.

[17] J. Tierney, C.M. Rader and B. Gold, "A digital frequency synthesizer," *IEEE Transactions on Audio and Electroacoustics*, vol. AU-19, 1971, pp. 48-57.

[18] J. Vankka, M. Waltari, M. Kosunen and K. A. I. Halonen, "A direct digital synthesizer with an on-chip D/A converter," in *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, February 1998, pp. 218-227.

[19] L.A. Weaver Jr. and R.J. Kerr, "High resolution phase to sine amplitude conversion," U.S. patent 4,905,177, February 27th, 1990.

[20] A. Yamagishi, M. Ishikawa, T. Tsukahara, and S. Date, "A 2-V, 2-GHz low-power direct digital frequency synthesizer chip-set for wireless communication," in *IEEE Journal of Solid-State Circuits*, vol. 33, February 1998, pp. 210-217.